

# 400 Gbps per lane PHYs and Interfaces

Call for Interest (CFI) Consensus Building

Kent Lusted, Synopsys

Mark Nowell, Cisco

# Today's Panel

- Speakers
  - Kent Lusted, Synopsys
  - Adam Healey, Broadcom
  - Mark Nowell, Cisco
- Additional Panelists
  - John D'Ambrosia, Futurewei, US Subsidiary of Huawei
  - Mark Kimber, Semtech
  - Rob Stone, Meta

# Contributors

- John D'Ambrosia, Futurewei, US Subsidiary of Huawei
- Adam Healey, Broadcom
- Halil Cirit, Meta
- Rob Stone, Meta
- Matt Brown, Qualcomm
- Adee Ran, Cisco
- Alan Weckel, 650 Group
- David Ofelt, HPE
- Jose M Castro, Panduit
- Baron Fung, Dell'Oro Group
- Sameh Boujelbene, Dell'Oro Group
- Xiaojie Fan, Huawei
- Ram Huggahalli, Microsoft
- Nic Viljoen, Meta
- Kapil Shrikhande, Marvell
- Tony Chan Carosone, Qualcomm
- Nathan Tracy, TE Connectivity
- Ashika Pandankeril Shaji, TE Connectivity
- Andrew Josephson, Samtec
- Vivek Shah, Molex
- Sam Kocsis, Amphenol
- Xiang He, Huawei
- Peter Graumann, Microchip
- John Johnson, Broadcom
- John Calvin, Keysight Technologies
- Naim Ben-Hamida, Ciena
- Mike Li, AMD

# Two Efforts Forming In the Pipeline

## Moving Forward in IEEE 802.3

NEA (E4AI Assessment) – Consensus / Incubation - Ongoing stakeholder input

Focus of this CFI =  
400 G/lane  
signaling

Effort #1

SG

Baseline  
Adoption

TF Draft

Effort #2

SG

Baseline  
Adoption

TF Draft

4 – 8 mos

>= 3.2 Tb Ethernet  
and PHYs not in  
Effort #1

\* Each new project will potentially build on the previous project(s).  
802.3 has the "Ethernet for AI" NEA effort which will sustain  
incubation across many of these distinct Task Force projects to  
encourage a fast-follow cadence.

Additional  
Efforts

18

[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0819/dambrosia\\_e4ai\\_01\\_250819.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0819/dambrosia_e4ai_01_250819.pdf)

IEEE 802.3 400G per Lane Signaling CFI, March 2026

# Objective

- Build consensus on starting a study group investigating a “400 Gbps per lane PHYs and Interfaces” project
- We do not need to:
  - Fully explore the problem
  - Debate strengths and weaknesses of solutions
  - Choose a solution
  - Create a PAR or 5 Criteria or the SG objectives
  - Create a standard
- Anyone in the meeting may vote or speak

# Supporters (1/2)

- Osorio, Luz - Nokia
- Ghiasi, Ali - Ghiasi Quantum, Marvell
- Ofelt, Dave - HPE
- Choudhury, Mabud - Genuine Optics
- El-Chayeb, Ahmad - Keysight Technologies
- Nicholl, Gary - Cisco
- Zimmerman, George - CME Consulting/Cisco, Eliyan
- Mascitto, Marco - Nokia
- Galan, Jose - Maxlinear
- Stephens, Jeremy - AMD
- Weaver, Jim - Arista
- Goergen, Joel - Cisco
- D'Ambrosia, John - FutureWei, US subsidiary of Huawei
- Calvin, John - Keysight Technologies
- Moorwood, Andy - Keysight Technologies
- Rabinovich, Rick - Keysight Technologies
- Sun, Junqing (Phil) - Credo
- Ramesh, Sridhar - Maxlinear
- Estes, Dave - Viavi
- Brooks, Paul - Viavi
- Nakamoto, Ed - Viavi
- Lewis, Jon - Dell Technologies
- Stone, Rob - Meta
- Zhang, Geoff - AMD
- Li, Tobey - MediaTek
- Wu, Mau-Lin - MediaTek
- Nicholl, Shawn - AMD
- Klempa, Michael - Qualcomm
- Brown, Matt - Qualcomm
- Vidal, Or - Qualcomm
- Bowman, Kurtis - AMD/UALink
- Akinwale, Femi - Intel
- Johnson, John - Broadcom
- Parthasarathy, Vasudevan - Broadcom
- Muth, Karl - Broadcom
- Liu, Cathy - Broadcom
- Mellitz, Rich - Samtec
- Gore, Brandon - Samtec
- Cirit, Halil - Meta
- Opsasnick, Eugene - Broadcom
- Simms, Bill - Nvidia
- Tooyserkani, Pirooz - Cisco
- Kabra, Lokesh - Synopsys
- Bhatt, Vipul - Coherent
- Cole, Chris - Coherent
- Rodes, Roberto - Coherent
- Biederman, Dan - Eridu AI
- Patra, Lenin - Marvell
- Kochuparambil, Beth - Cisco
- Tang, Yi - Cisco
- Noujeim, Leesa - Google
- Kareti, Upen Reddy - Cisco
- Ben-Artzi, Liav - Marvell
- Williams, Tom - Cisco
- Welch, Brian - Cisco
- Dudek, Mike - Marvell
- Kimber, Mark - Semtech
- Ran, Adeel - Cisco
- Kutscher, Noam - Marvell
- He, Xiang - Huawei
- Mi, Guangcan - Huawei
- Sommers, Scott - Molex

# Supporters (2/2)

- Marshall, John - AMD
- Li, Mike - AMD
- Issenhuth, Tom - Huawei
- Chen, David CC - AOI
- Tracy, Nathan - TE
- Heck, Howard - TE
- Huggahalli, Ram - Microsoft
- Maniloff, Eric - Ciena
- Wingrove, Mike - Ciena
- Yu, Rangchen Ryan - Terahop
- Huber, Tom - Nokia
- Bruckman, Leon - nvidia
- Marques, Flavio - Lightera
- Palkert, Tom - Samtec
- Wang, Sharon - TE Connectivity
- Parsons, Earl - CommScope Amphenol
- He, Michael - Terahop
- Pepper, Jerry - Keysight
- Bernier, Eric - Huawei
- Slavick, Jeff - Broadcom
- Kocsis, Sam - Amphenol
- D'Silva, Hansel - Amphenol
- Shakiba, Hossein - Huawei
- Jackson, Kenneth - Sumitomo
- Zivny, Pavel - Multilane
- Muhigana, Ernest - Macom
- Landry, Gary - Texas Instruments
- Tartaglia, Antonio - Ericsson
- Parkholm, Ulf - Ericsson
- Lessard, Stephane - Ericsson
- Chan, Jason - Arista
- Yu, Xu Helen - Huawei
- Healey, Adam - Broadcom
- Tran, Viet - Keysight Technologies
- Donahue, Curtis - Rohde & Schwartz
- Ninomiya, Tiger - Accelink
- Maki, Jeff - Hewlett Packard Enterprise
- Sakai, Toshiaki - Socionext
- Theodoras, Jim - SCINTIL Photonics
- Jones, Chad - Cisco
- Jones, Peter - Cisco
- Sekel, Steve - Wilder Technologies
- DeKoos, Andras - Microchip
- Mammenga, Kevin - Wilder Technologies
- Chang, Frank - Source Photonics
- Shrikhande, Kapil - Upscale AI
- Potterf, Jason - Cisco

# Overview

To keep pace with the technology demands necessary to support the projected growth in AI Infrastructure, 400 Gbps/lane data rate PHYs need to be defined in IEEE 802.3

To explain this better and propose the formation of a new Study group we plan to walk through the following material:

- Market Drivers – Mark Nowell
- Technical Feasibility – Adam Healey
- Why Now & Summary – Kent Lusted
- Q&A (With panel)

# CFI Request

With the rapid adoption of Artificial Intelligence (AI) tools and technology, there is a desire to use Ethernet in high-bandwidth, high-growth AI cluster networks, as well as traditional front-end networks for server and switch attachment. The Ethernet community must quickly respond to support the AI network time-to-market requirements and use cases. A number of technology improvements should be investigated to increase the per-lane signaling rate of electrical and optical Ethernet interconnects for the radix-optimized connectivity essential for communication between the high-performance compute devices and intermediate network switch devices.

**This call for interest is to initiate a Study Group to develop a Project Authorization Request (PAR) and Criteria for Standards Development (CSD) for Physical Layer specifications supporting existing Ethernet rates using 400 Gbps-per-lane signaling for electrical interconnects as well as optical interconnects with reaches of up to 500 meters.**

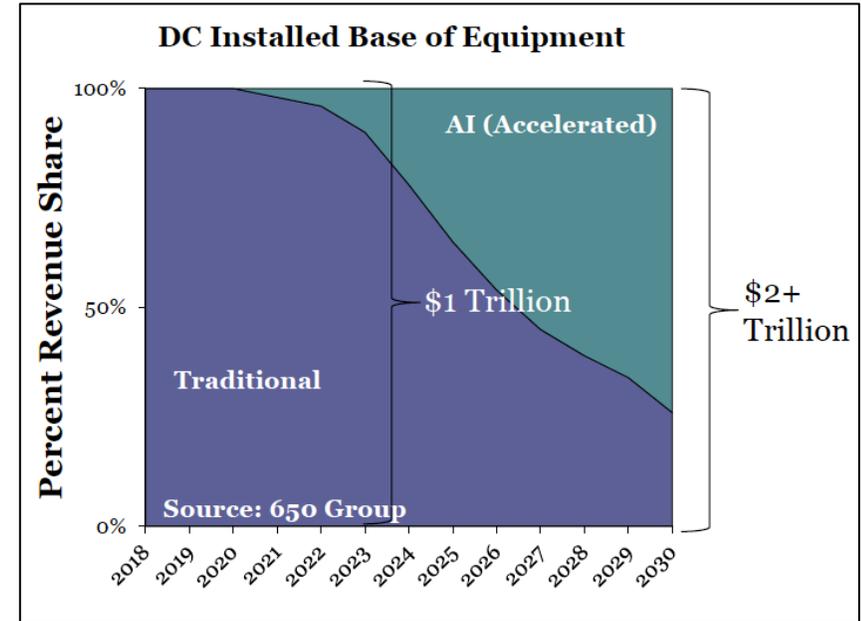
[https://ieee802.org/3/cfi/request\\_0326\\_1.html](https://ieee802.org/3/cfi/request_0326_1.html)

# Market Need

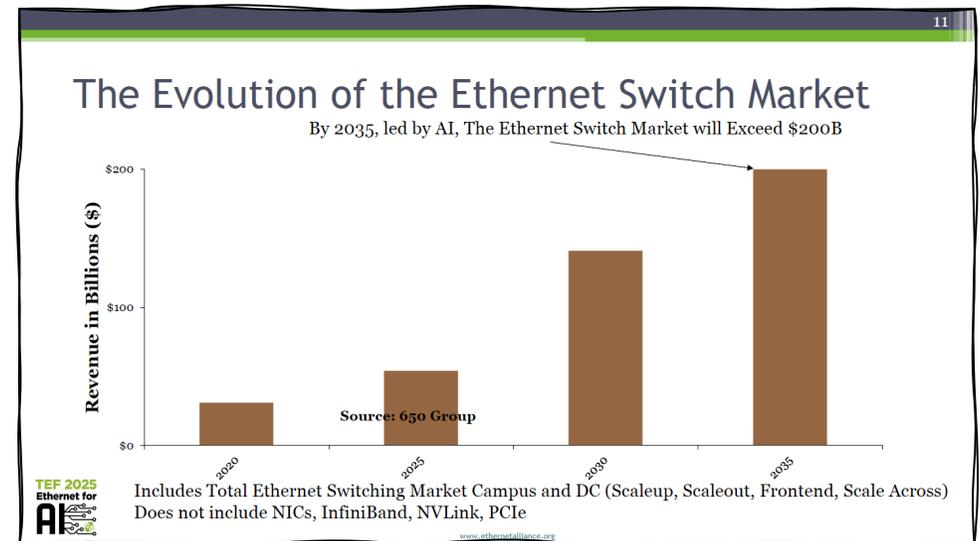
Mark Nowell

# AI Explosion

- The enormous growth of AI/ML is placing unprecedented **demand on system interconnects and on scaling the bandwidth**, specifically xPU-to-xPU (via switches)
- Ethernet is uniquely positioned to be the building block of choice for AI/ML system interconnects
- The scale of the deployments requires a robust ecosystem to support the volumes.
- In networking, Ethernet wins due to its robust ecosystem and supply chain



Source: Alan Weckel, EA's TEF 2025, "Setting the stage for networking in an AI world"

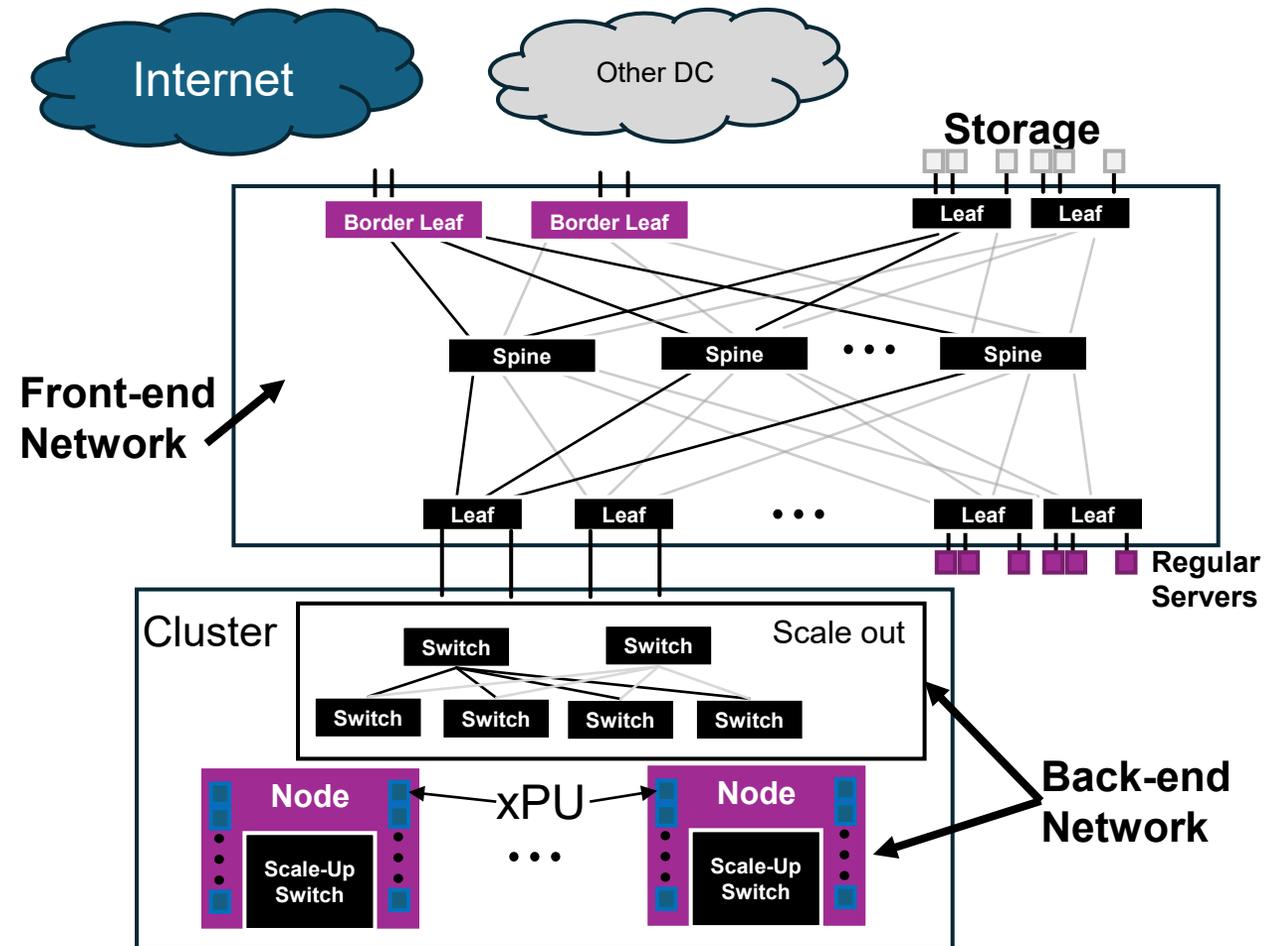


Source: Alan Weckel, EA's TEF 2025, "Setting the stage for networking in an AI world"

# Network Evolution

The network is evolving into distinct networks with both common yet unique needs

- Front-end is the classic data center switching network connecting servers, storage, internet, etc.
- One of the “servers” can be a large AI cluster
- An AI cluster is a set of compute resources (CPUs, GPUs, accelerators) tied together by a “back-end” network.
  - The entire cluster can be viewed as a single server



Source: David Ofelt, HPE

# Looking Closer at the Backend Network

Front-end

Today's Ethernet

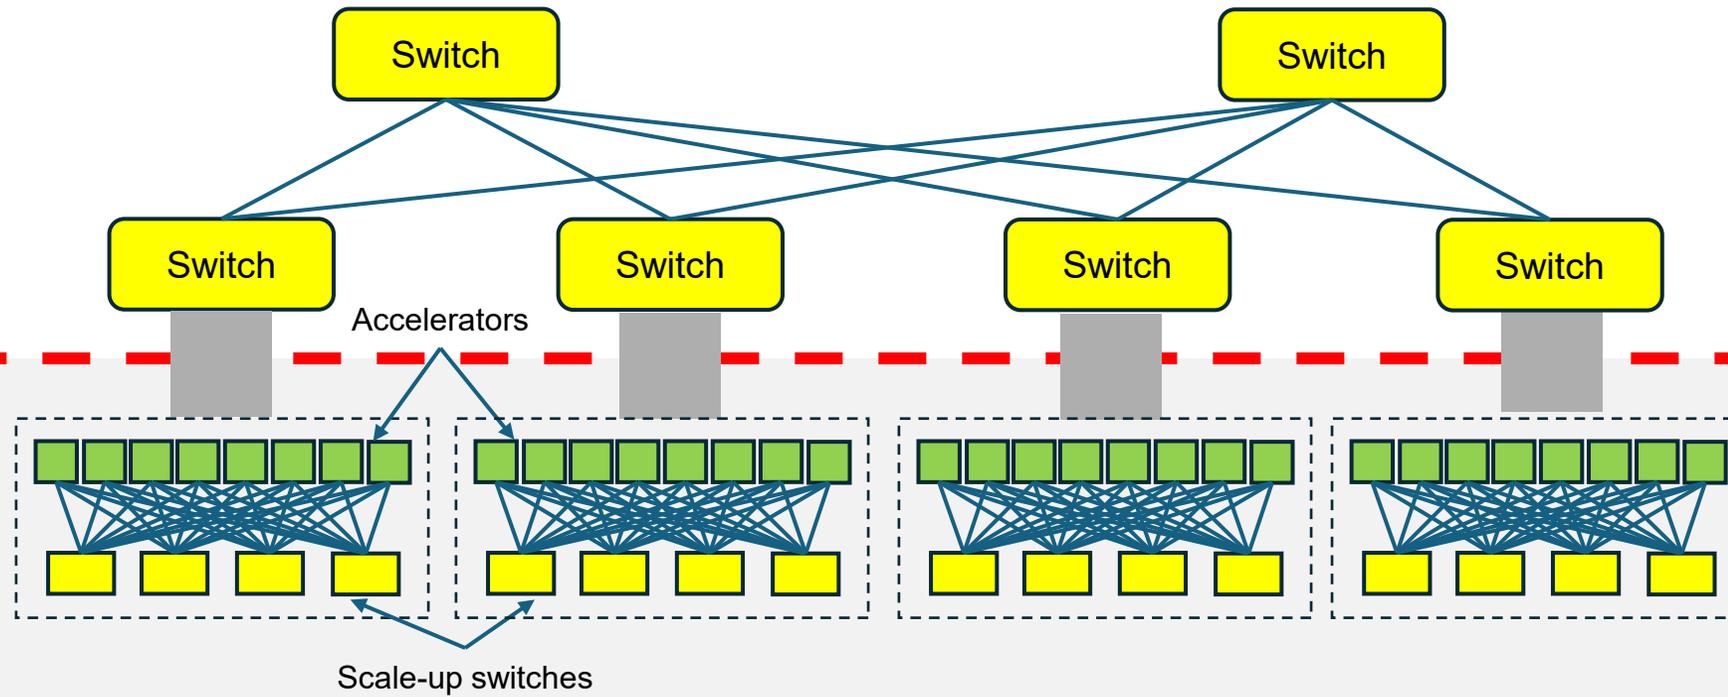
Higher MAC rates needed

Back-end Scale-Out

Radix & increased port bandwidth

Back-end Scale-Up

High radix & higher data rate needed



Based on AI Datacenters and their Diverse Network Requirements, Ram Huggahalli (Microsoft), Ethernet Alliance TEF 2024, Oct 2024.

IEEE 802.3 400G per Lane Signaling CFI, March 2026

Accelerators (e.g. GPU, xPU)  
Switches

# Looking Forward

- AI back-end “scale-up” networks are driving higher signaling rates first
  - Followed by scale-out
  - Then front-end networks
- Scale-up = high radix
  - Flat network, lots of one-lane and two-lane links
- Copper excels due to low cost & power and high reliability, but the reach is limited and shrinking
- Optics have a growing footprint in scale-up (and, of course, scale-out) to meet reach and cluster size requirements
  - Power optimized solutions, such as Co-packaged optics (CPO) implementations, are increasingly important

## Projected relative growth of different networks (relative to front-end)

Front-end	1x
Back-end network (scale-out)	10x
Back-end network (scale-up)	100x

[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0327/weckel\\_e4ai\\_01\\_250327.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0327/weckel_e4ai_01_250327.pdf)

# Why Higher Data-Rate Physical Layers?

- Compute devices and Switch devices are being driven to higher capacities
- Beachfront density limits the number of lanes; thus, a higher rate per lane is required
- Consistent with historical progression of Ethernet lane speeds for data center applications

## Beachfront Bottleneck of Scale in Systems

- AI system performance is dependent on 3 major pillars which need to improve all together
- Memory and Compute takes full benefit (Linear scaling) of advanced nodes while SerDes design can't
- Beachfront is the major limitation on scaling the Network IO BW

Performance Dependencies

MEMORY BW COMPUTE TOP NETWORK IO BW

IO Chiplets HBMs XPU Chiplet HBMs HBMs XPU Chiplet HBMs

Backplane Cable Connector

CPC, Connector Route & Cable pitch restrictions

CPC = Co-Packaged Copper

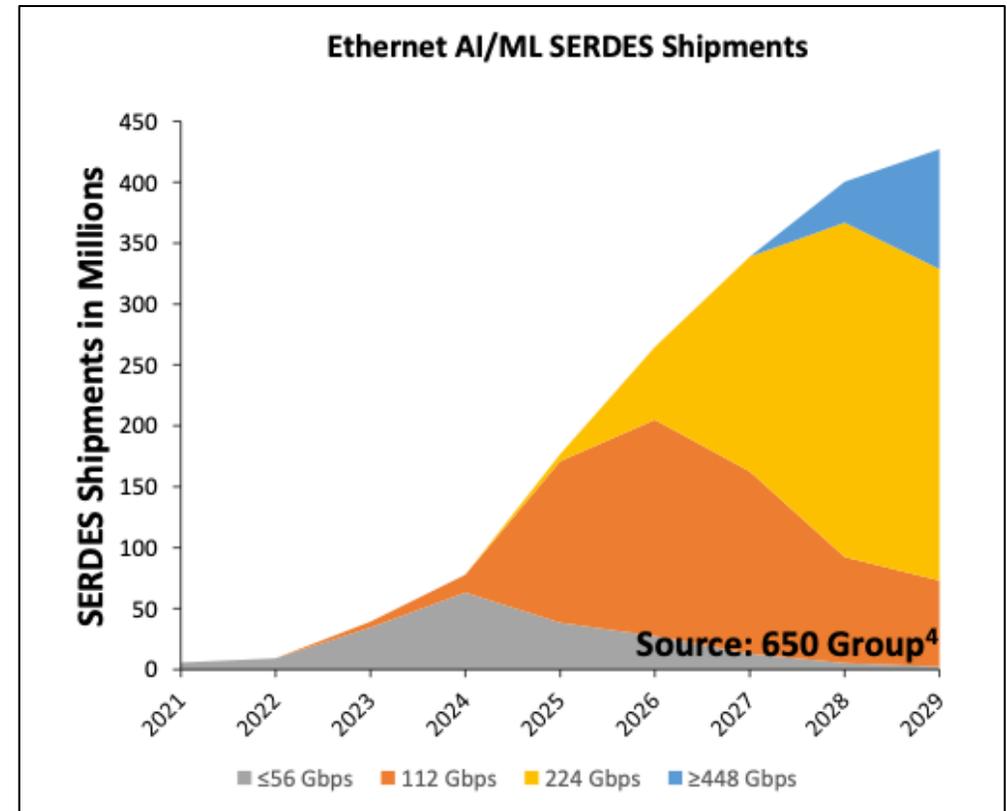
TEF 2025 Ethernet for AI

xN Computing  
xN Memory BW  
x1 Network IO BW

Source: Halil Cirit, EA's TEF 2025, "Transitioning to 400G SerDes: Key Drivers and System Design Implications for Future AI Workloads"

# Market Projections for 400 G/lane

- Market analysis is predicting the market need for 400 Gb/s signaling starting in 2027
- Driven by AI networks getting bigger and constraints at the Switch/GPU, the transition to 400 G/lane technology is natural

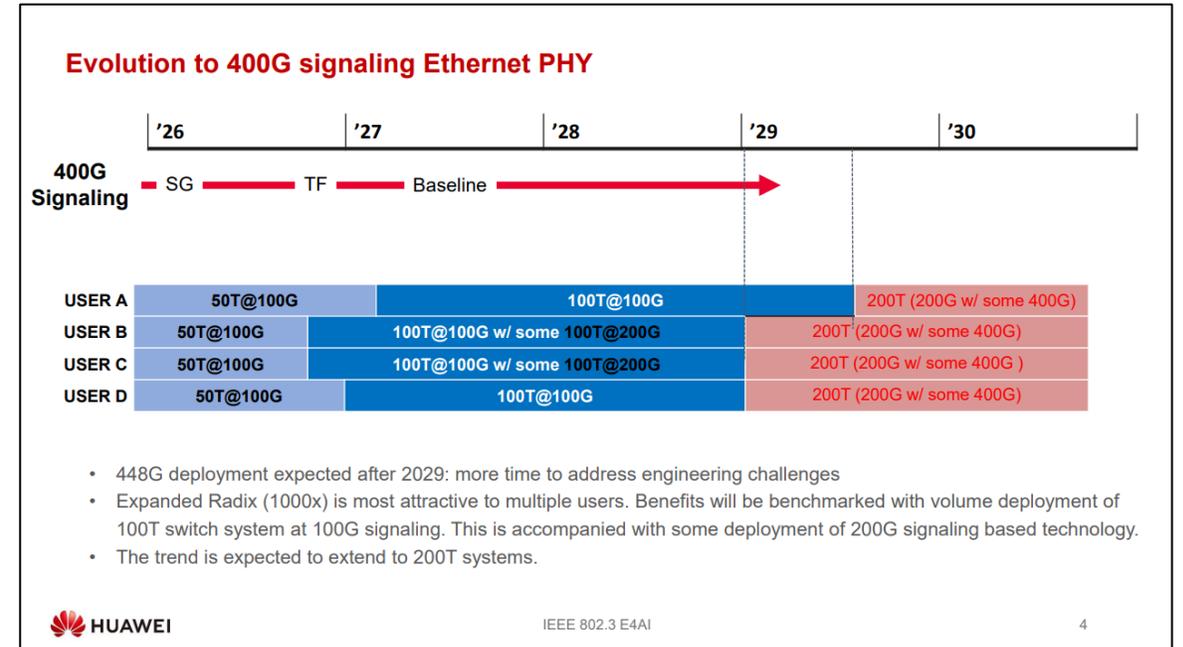


[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0327/weckel\\_e4ai\\_01\\_250327.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0327/weckel_e4ai_01_250327.pdf)

# Global Market Need

- Data-center operators are rapidly scaling AI workloads and expect to adopt 400 Gb/s per-lane Ethernet PHY, driven by the need for higher bandwidth and expanded optical interconnect use

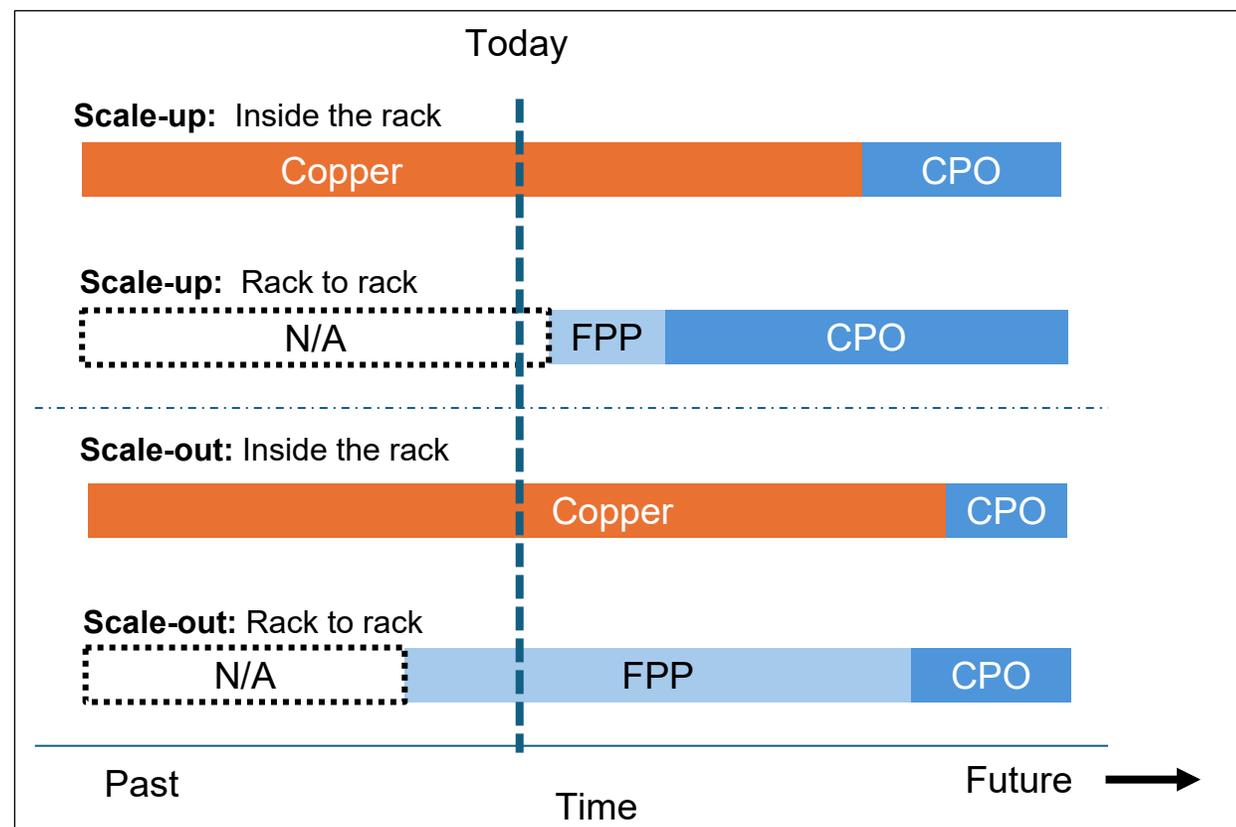
Input from Cloud and AI providers based in China



[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/26\\_0219/mi\\_e4ai\\_01\\_260219.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/26_0219/mi_e4ai_01_260219.pdf)

# I/O Transitions for AI

- Current and future transitions involve Copper, front panel pluggable (FPP) optics, and CPO/NPO
- Copper dominates inside-the-rack connections today but shifts toward optics over time
- Rack-to-rack connections transition from pluggable to CPO
- Market requires both copper and optical PMDs @ 400G



Based on AI Networking: What do scaleup and scaleout really mean for networking demand, Alan Weckel (650 Group), [https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0327/weckel\\_e4ai\\_01\\_250327.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0327/weckel_e4ai_01_250327.pdf)

## Requirements for BOTH electrical and optical PMDs

# Many Consortia are Relying on Ethernet

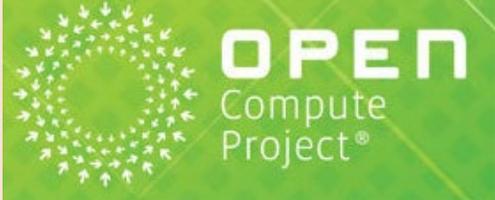
To support the breadth of AI applications, multiple organizations exist:



Deliver an Ethernet based open, interoperable, high performance, full-communications stack architecture



An open, high-bandwidth and low latency interconnect for connecting AI accelerators (such as GPUs) and switches



SUE and ESUN build on Ethernet by introducing evolutionary enhancements that improve throughput, latency, and reliability to meet the demands of Scale-Up GPU connectivity.

All are looking for IEEE 802.3 to define the Ethernet Physical Layer Specifications

# Summary: “Market need”

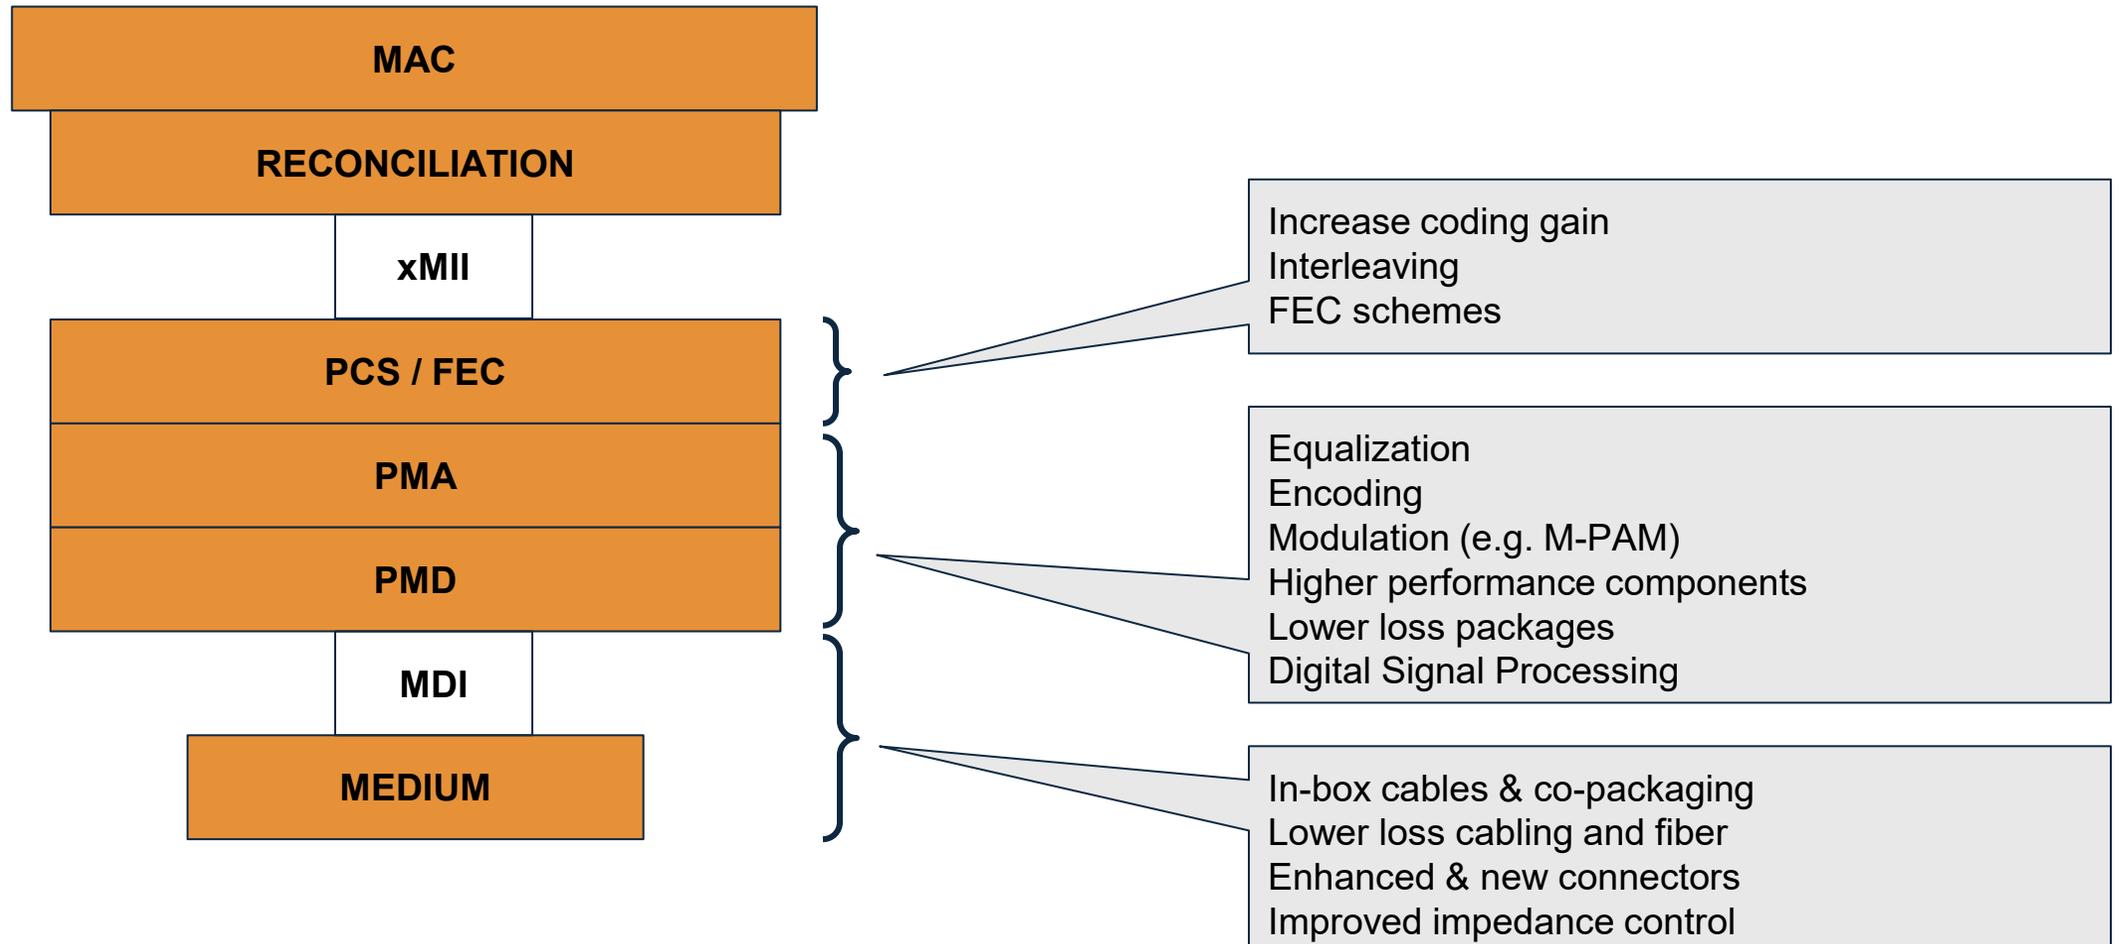
- AI scale-up networks are the fastest growing segment and require 400 Gb/s lane rates
- AI networks needs rapid development of the next physical layer data rate
- Ecosystem dynamics are all pointing to Ethernet (and IEEE 802.3 specifically) to provide the industry with these key building blocks
- Initial market is for high radix solutions supporting both electrical and optical interfaces

# Technical Feasibility

We have done this before...

1G → 10G → 25G → 50G → 100G → 200G

# Opening the Toolbox Again



# SERDES Capability is a Key Enabler

<b>Data rate per lane, Gb/s</b>	10	25	50	100	200	<b>400</b>
<b>Modulation</b>	PAM-2	PAM-2	PAM-4	PAM-4	PAM-4	<b>?</b>
<b>Nominal cable reach, m</b>	7	5	3	2	1~2 <sup>1</sup>	<b>?</b>
<b>Technology added</b>	DFE	RS FEC	Stronger RS FEC	Floating-tap DFE	MLSD	<b>?</b>
<b>Year<sup>2</sup></b>	2007	2014	2018	2022	2026 (est.)	<b>?</b>

<sup>1</sup> Nominal cable reach is a function of host class

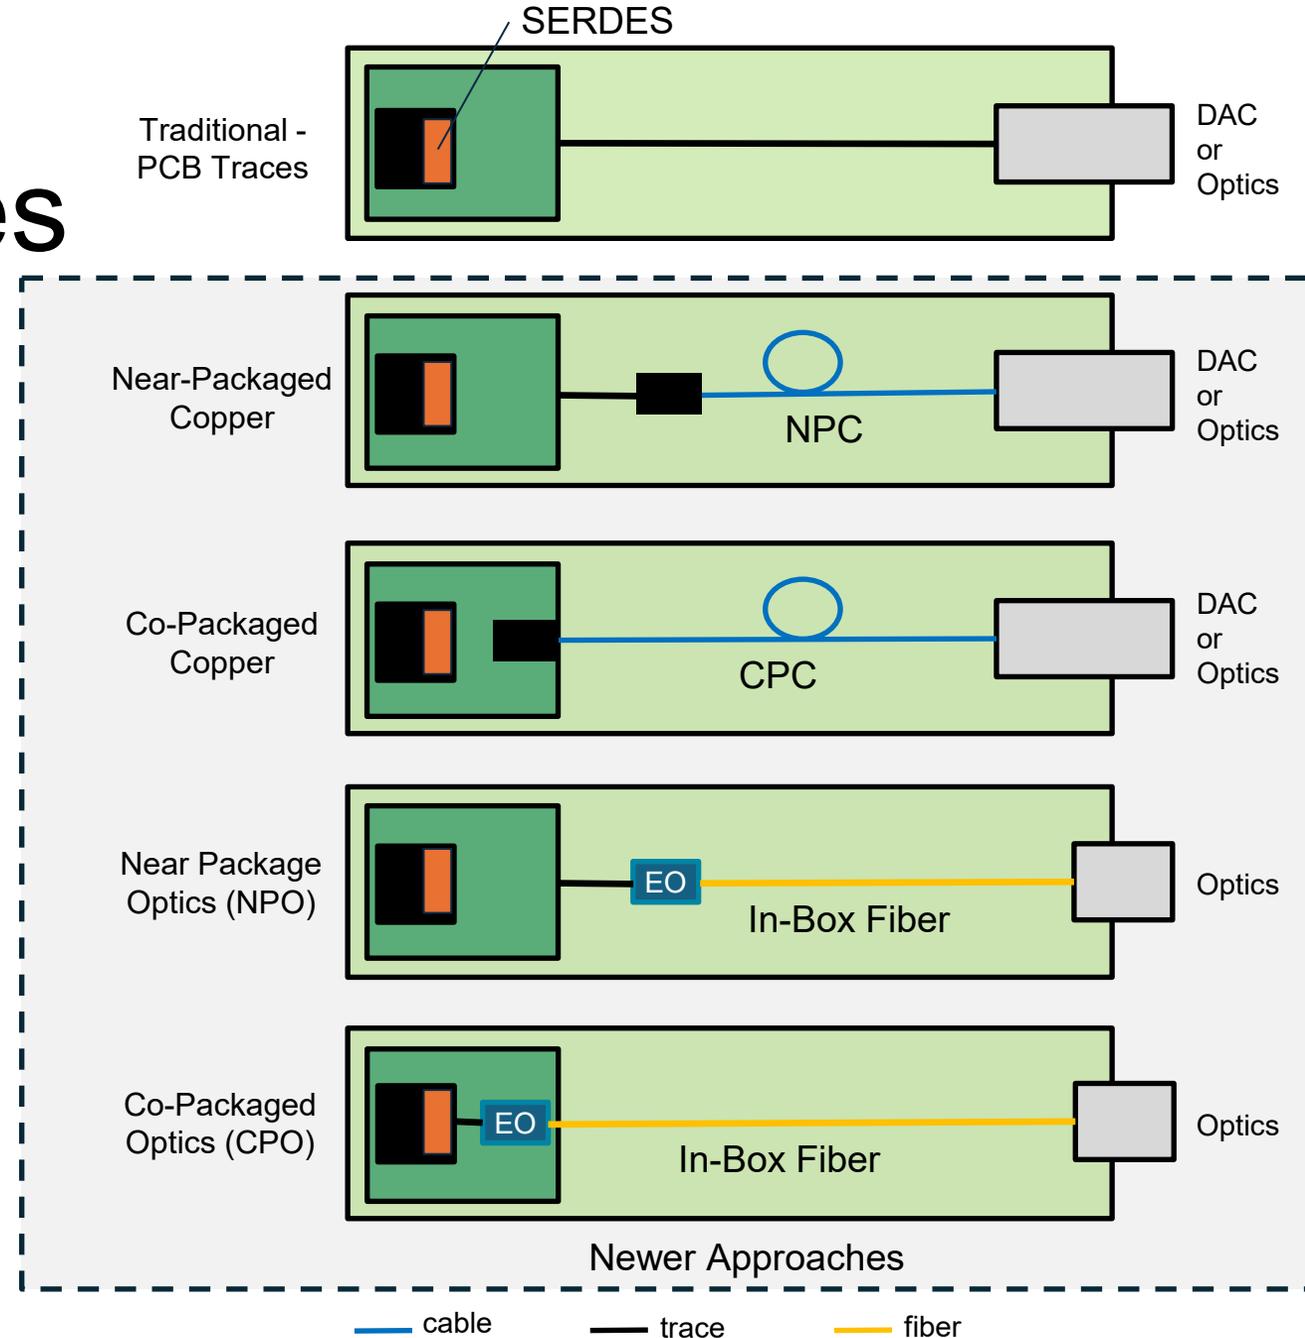
<sup>2</sup> Year of publication of the IEEE 802.3 Ethernet standard.

DFE = decision feedback equalizer  
 FEC = forward error correction  
 MLSD = maximum-likelihood sequence detector  
 RS = Reed-Solomon

- New standard for doubling the data rate every 4 years
- New capabilities added for each generation to address challenges presented by doubling the data rate
- What will be required for the next doubling of data rate? Numerous options have already been explored...

# System Architectures

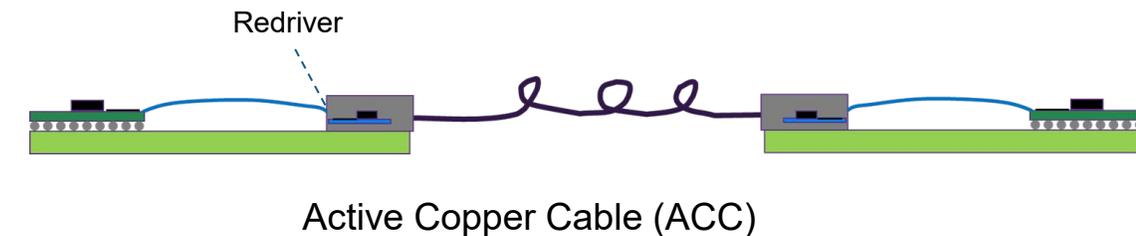
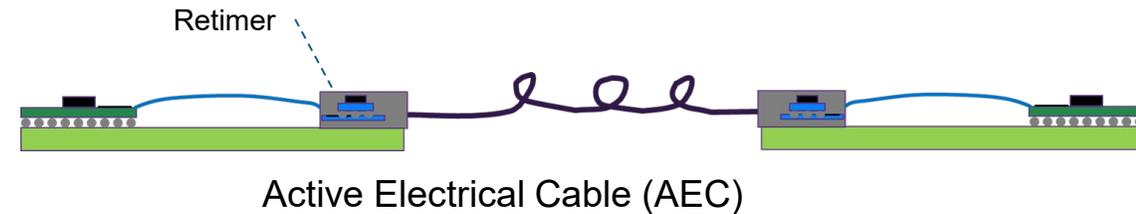
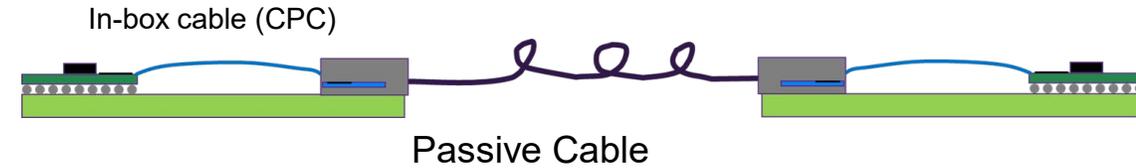
- Broader mix of system architectures than before
  - PCB traces to front panel pluggable (FPP)
  - Near-packaged and Co-packaged copper to FPP
  - Near-Package Optics (NPO)
  - Co-Packaged Optics (CPO)
- A variety of approaches depending on system requirements & limitations
- All need SERDES



# Copper Interconnect

- Numerous copper solutions can be used to address the needs of the market
  - Passive:
    - Well known and well-established methodologies
  - AEC:
    - Comparable to a pluggable optical transceiver
  - ACC:
    - Not well specified across industry

Type	IC	Reach	Power
Passive	None		
Active Electrical Cable (AEC)	Retimer		
Active Copper Cable (ACC)	Redriver		



# Optics and Fiber

Optics provide reach beyond copper to extend scale-up networks (e.g., beyond a single rack)

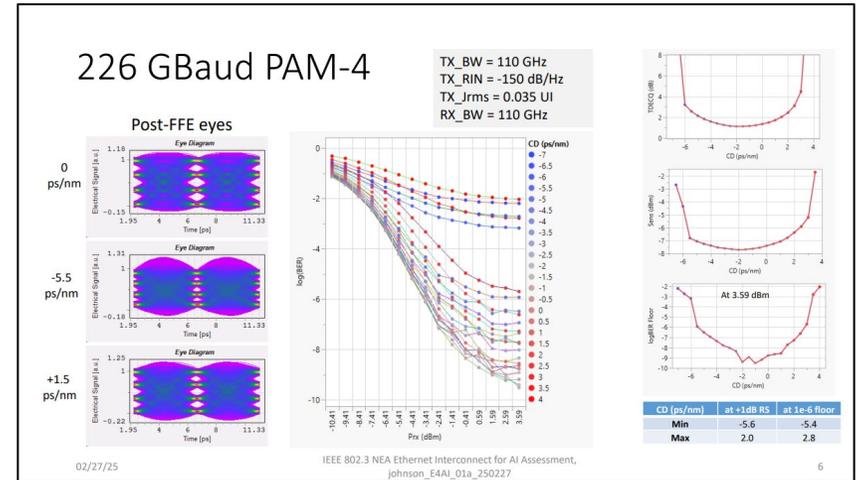
Single-wavelength optics and parallel single-mode fiber cables provide the best support for high-radix scale-up (and some scale-out) networks

- High bandwidth per fiber enables dense connectivity
- Low dispersion from single-wavelength operation

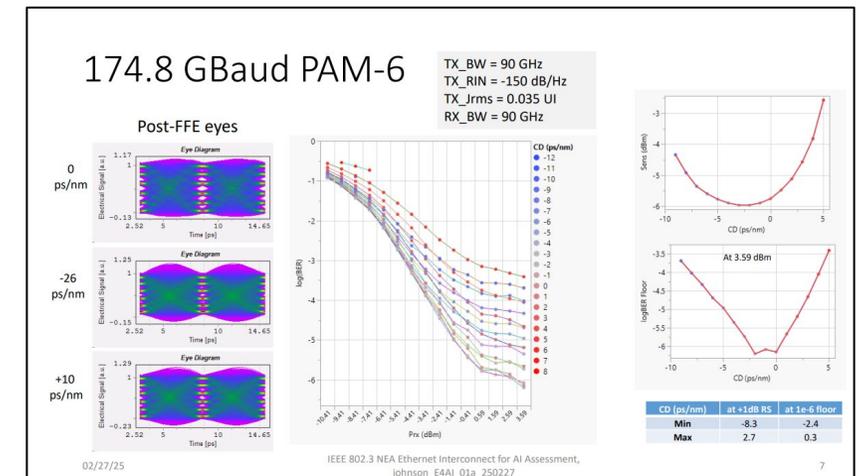
Both PAM-4 and PAM-6 modulation are feasible, each having advantages and disadvantages

Single-wavelength optics are aligned with industry directions (e.g., break-out, pluggables, CPO/NPO)

On-going IEEE 802.3 NEA work to explore the future of fiber (e.g., multi-core fiber, hollow-core fiber)



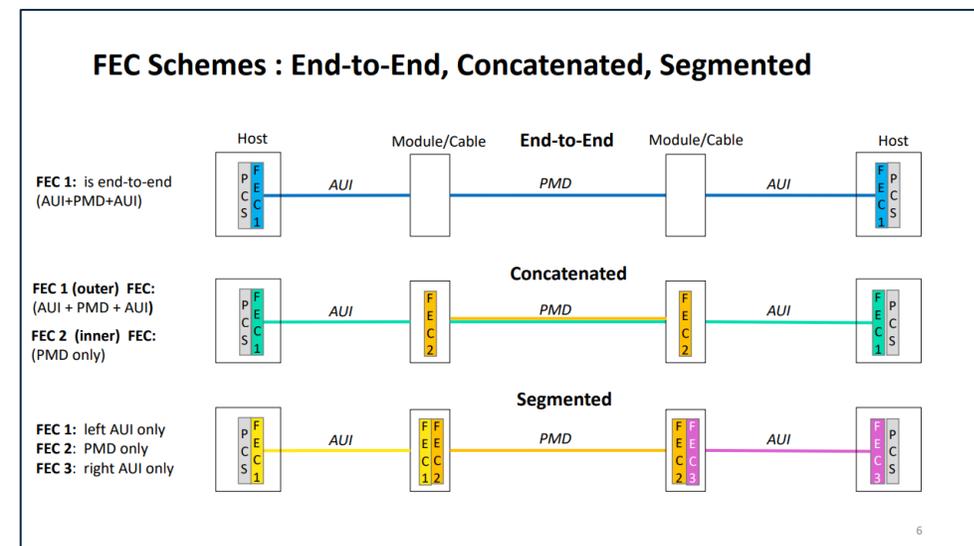
[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0227/johnson\\_e4ai\\_01a\\_250227.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0227/johnson_e4ai_01a_250227.pdf)



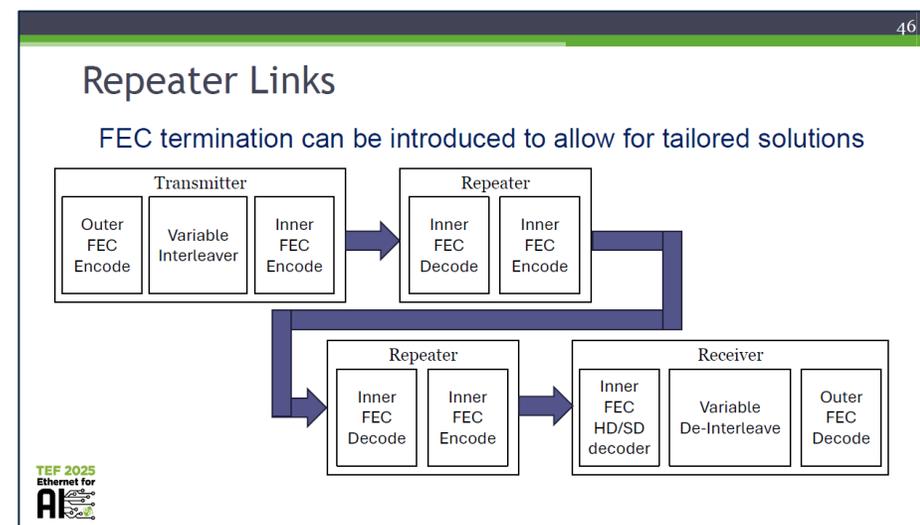
[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0227/johnson\\_e4ai\\_01a\\_250227.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0227/johnson_e4ai_01a_250227.pdf)

# FEC Schemes

- Multiple FEC schemes are available for 400 G/lane
- Each type has distinct latency, coding-gain, and implementation-complexity tradeoffs that require detailed study
- The diversity of options enables selecting the FEC scheme and associated architecture that best aligns with system-level performance, power, and interoperability requirements



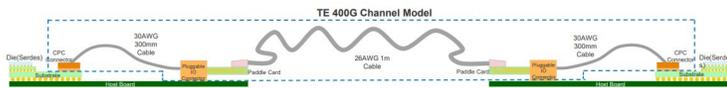
[https://www.ieee802.org/3/df/public/22\\_02/shrikhande\\_3df\\_01a\\_220203.pdf](https://www.ieee802.org/3/df/public/22_02/shrikhande_3df_01a_220203.pdf)



Source: Tony Chan Carusone, EA's TEF 2025, "Channel-Aware Modulation and FEC Selection for 400G+ Ethernet"

# Work is Already Underway

## Direct Attach Copper Cable Channel (tracy\_efai\_03\_250430)



### CPC Connector

- Includes connection to substrate and internal cable termination.
- Does not include substrate footprint and breakout.

### Pluggable IO Connector

- Includes internal cable termination.

### Paddle Card

- Includes external cable termination and paddle card routing and bulk cable

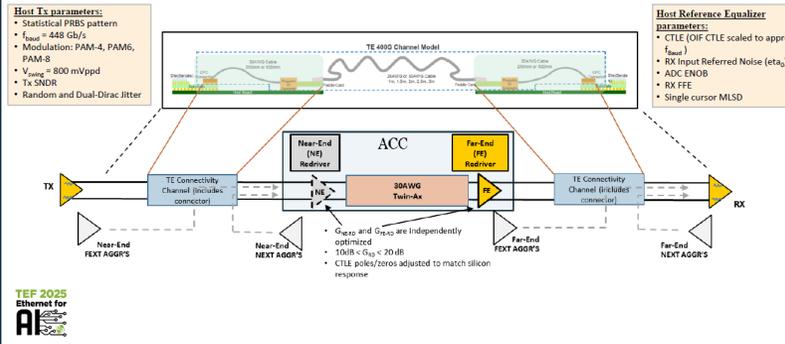
Channel File Name: tracy\_efai\_03\_250430.

Models are preliminary and under development.

[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0430/tracy\\_e4ai\\_01\\_250430.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0430/tracy_e4ai_01_250430.pdf)

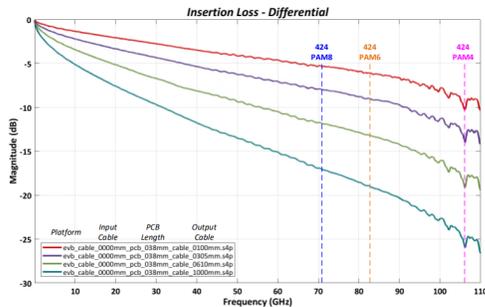
- Cable, Backplane and C2M, C2C channels
- Active Cables
- Connectors

## ACC Detailed System Block Diagram



Source: Ashika Pandankeril Shaji, EA's TEF 2025, "448G/lane Interconnects for AI: Channel Feasibility, Reach Extension, and Early Measurement Results"

## EVb Measured Insertion Loss



- 4.2dB to 14dB 212Gbps Channels
- No notch filter response to 110GHz
- -1dB ILD BW ~ 95GHz
- -15dB RL BW ~ 95GHz
- P/N Skew < 0.5ps
- 1.65dB/25mm PCB Striplines @53GHz
- Independently measured coaxial 2-ports provided as low skew 4 ports.

Contents in : josephson\_e4ai\_02\_250430.zip

IEEE 802.3 New Ethernet Applications Ad Hoc

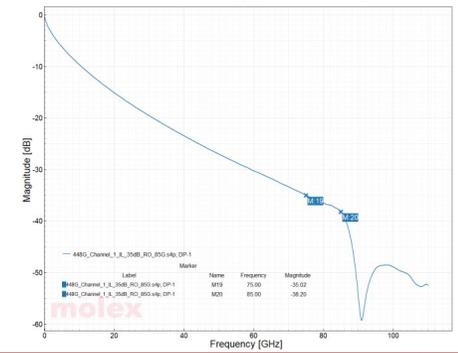
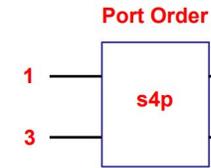
[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0430/josephson\\_e4ai\\_01a\\_250430.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0430/josephson_e4ai_01a_250430.pdf)

The TEF 2025 presentations reprinted with permission from the Ethernet Alliance and the respective copywriters as listed in "TEF 2025: Ethernet for AI" © 2025

IEEE 802.3 400G per Lane Signaling CFI, March 2026

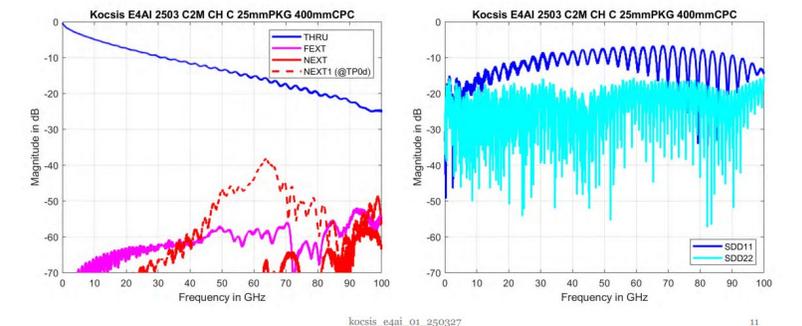
## IL Channels

Filename: 448G\_Channel\_1\_IL\_35dB\_RO\_85G.s4p



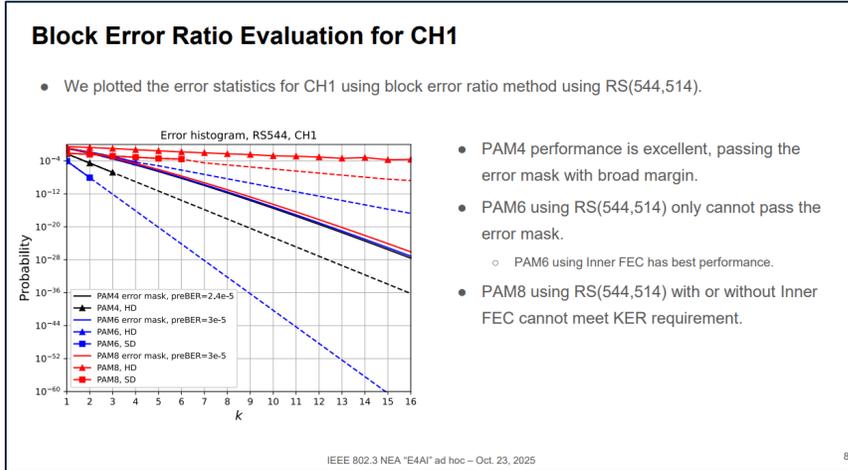
[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0430/shah\\_e4ai\\_01a\\_250430.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0430/shah_e4ai_01a_250430.pdf)

## Channel "C" "New Pluggable, New Paradigm"



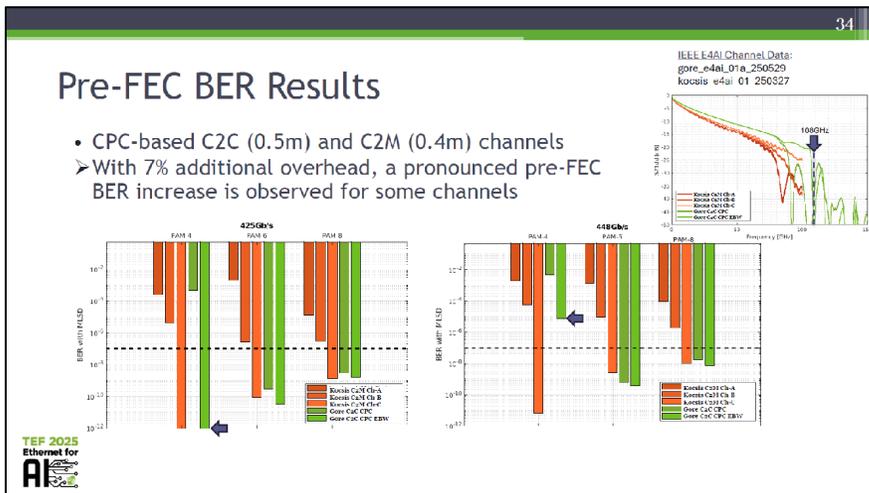
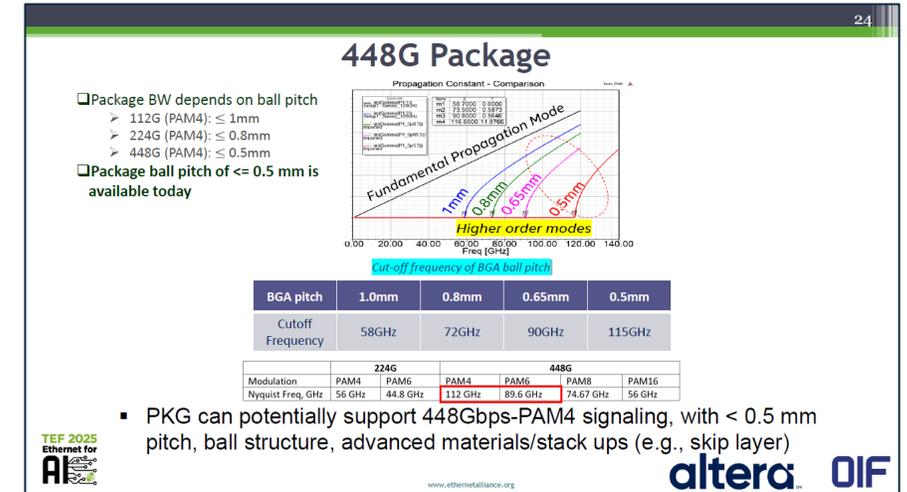
[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0327/kocsis\\_e4ai\\_01\\_250327.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0327/kocsis_e4ai_01_250327.pdf)

# Underway... (2/3)

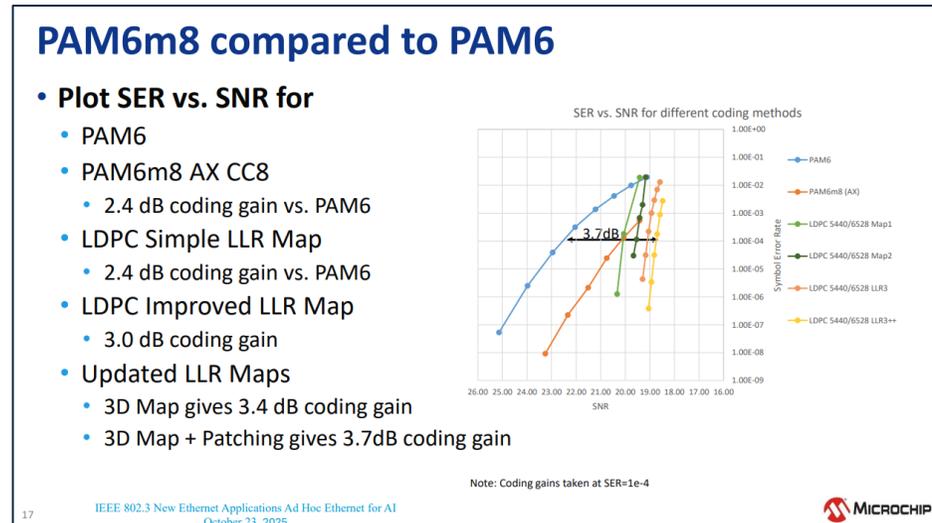


[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_1023/he\\_e4ai\\_01\\_251023.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_1023/he_e4ai_01_251023.pdf)

- FEC
- Package
- Encoding & Modulation
- COM analysis

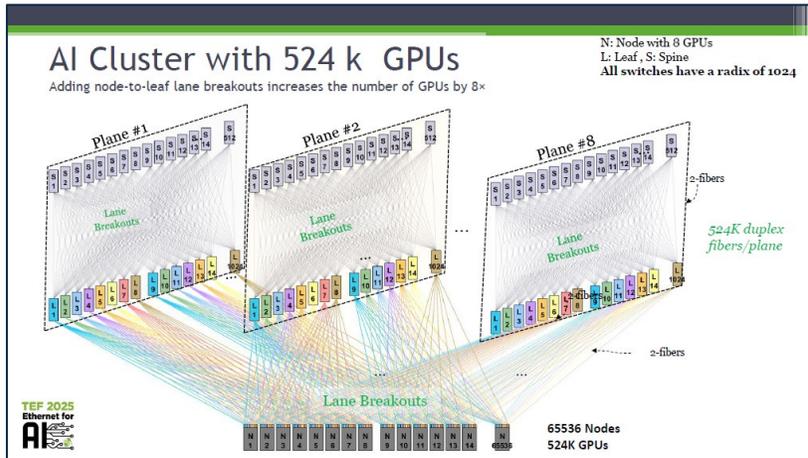


Source: Tony Chan Carusone, EA's TEF 2025, "Channel-Aware Modulation and FEC Selection for 400G+ Ethernet"



[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_1023/gramann\\_e4ai\\_01\\_251023.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_1023/gramann_e4ai_01_251023.pdf)

# Underway... (3/3)



- Optics Analysis
- Test & Measurement
- Demos

### 425Gb/ Electrical – Equalizer Investigation

**Important note**

- No input referred noise ( $\eta_0=0$ )
- SNR does not degrade after the channel

**Matlab post processing**

- PR-MSLD
- SER

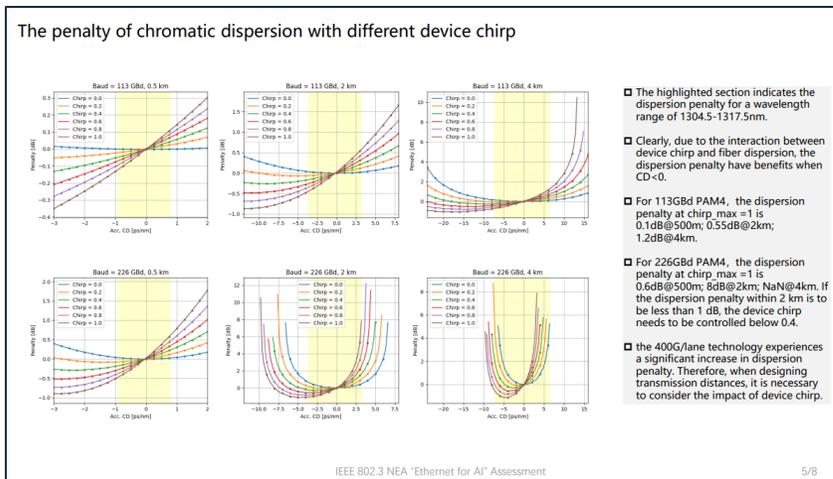
**PR-MSLD**

- 15-Tap pre-shaping filter (in addition to 22-Tap FFE)
- 1-transition memory
- Trellis length: 10 symbols

FFE & DFE are co-optimized (pulse Response & MMSE)

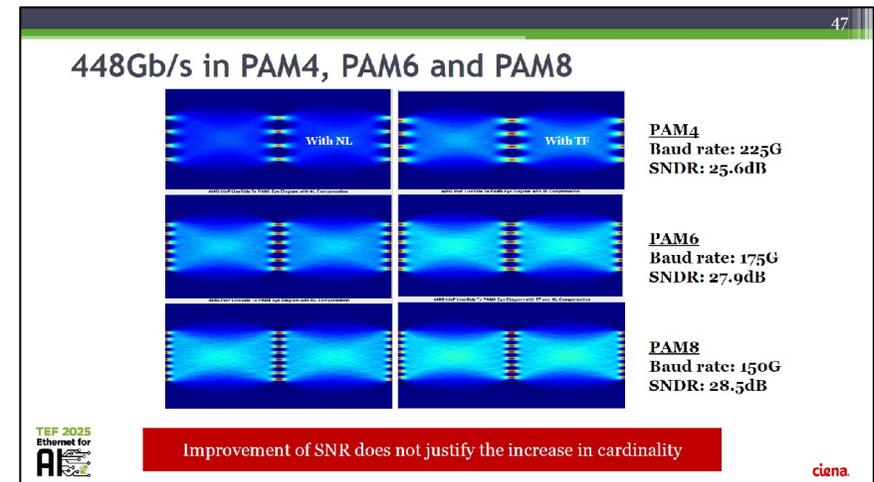
[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0529/calvin\\_e4ai\\_01a\\_250529.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0529/calvin_e4ai_01a_250529.pdf)

Source: Jose M Castro, EA's TEF 2025, "Enabling Massive Scale-Out AI Networks with Ethernet and Optical Lane Breakouts"



[https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0819/fan\\_e4ai\\_01\\_250819.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0819/fan_e4ai_01_250819.pdf)

IEEE 802.3 400G per Lane Signaling CFI, March 2026



Source: Naim Ben-Hamida, EA's TEF 2025, "Ethernet for AI: 448G technology for next gen networking"

# Summary: Technical Feasibility

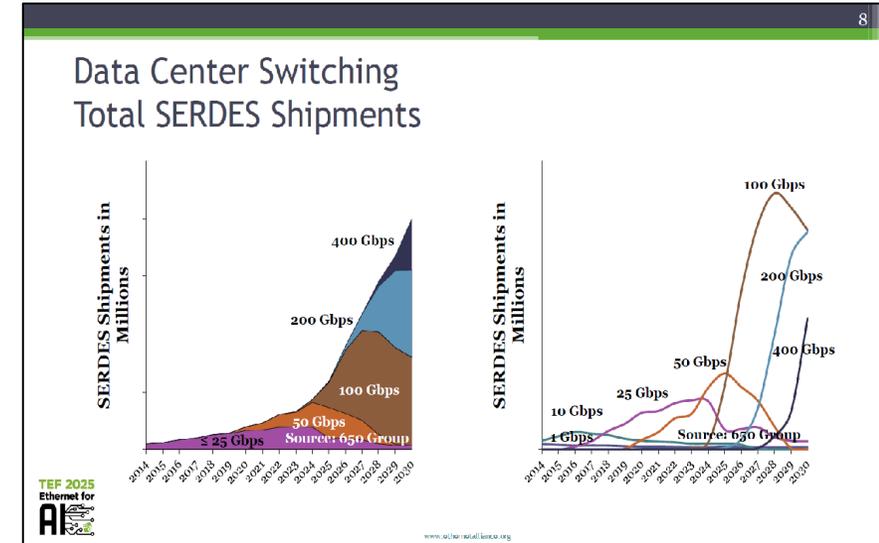
- Technical feasibility is no longer the question — industry is already investing across SERDES, channels, connectors, modulation, FEC, packaging, etc.
- IEEE 802.3 must define the interoperable standard
- The work is already underway across industry
  - IEEE 802.3 E4AI, EA TEF, SNIA, UEC, UALink, OCP ESUN, OIF ([OIF 448G Framework Document](#))

# Why now?

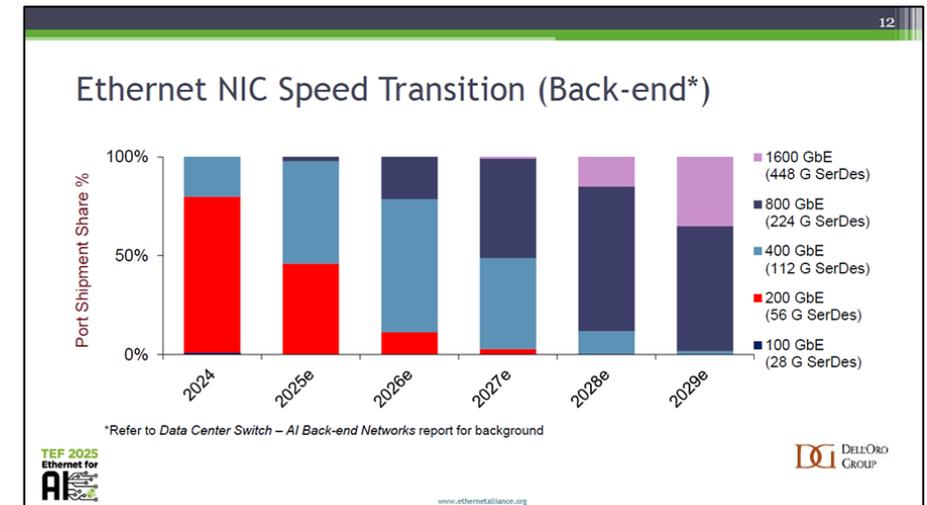
Kent Lusted

# Growth & Bandwidth Needs

- AI scale-up networks need 400G/lane for xPU-xPU and flattest network
  - Scale-out and front-end networks will leverage this work
- xPU and Switch designs need higher speed interconnect to remain feasible
- Copper interfaces specs are needed to define SERDES capabilities and for low power and low cost
- Optics interface specs are needed to enable reach beyond copper



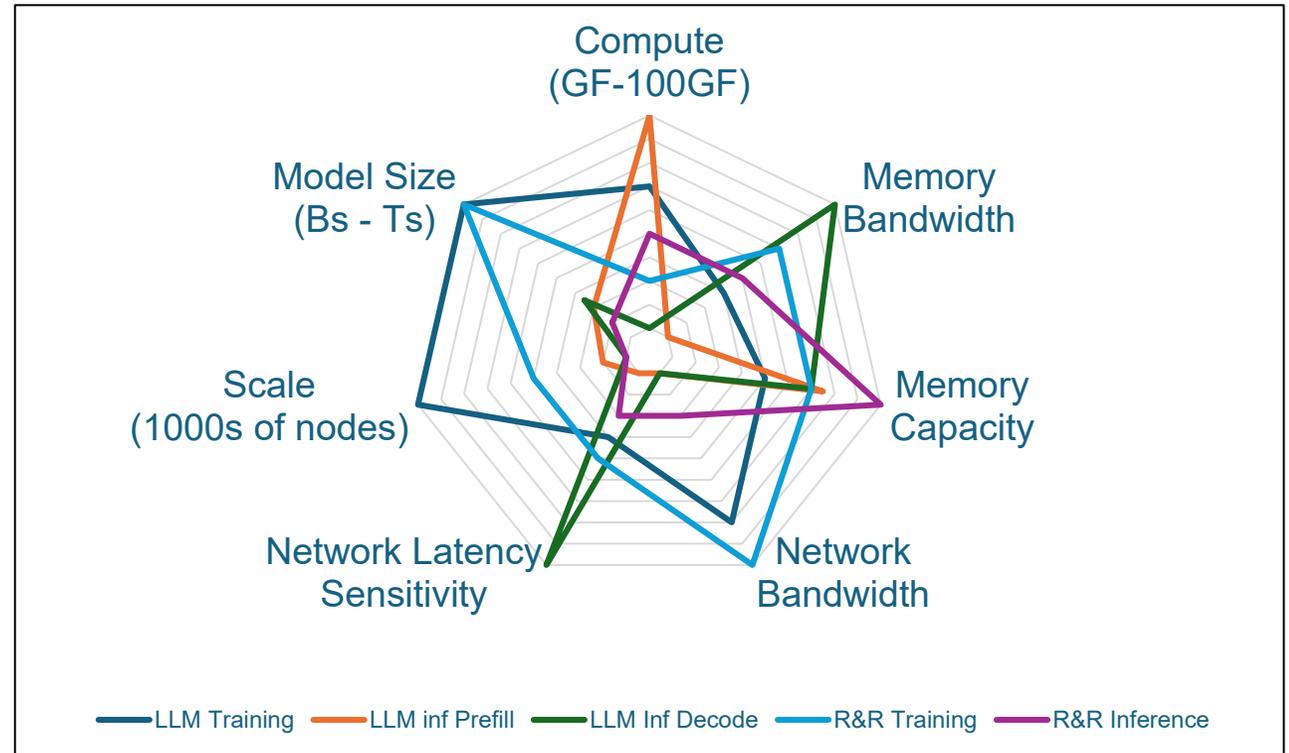
Source: Alan Weckel, EA's TEF 2025, "Setting the stage for networking in an AI world"



Source: Baron Fung, EA's TEF 2025, "XPU's and the Future of AI Connectivity"

# AI Application Requirements

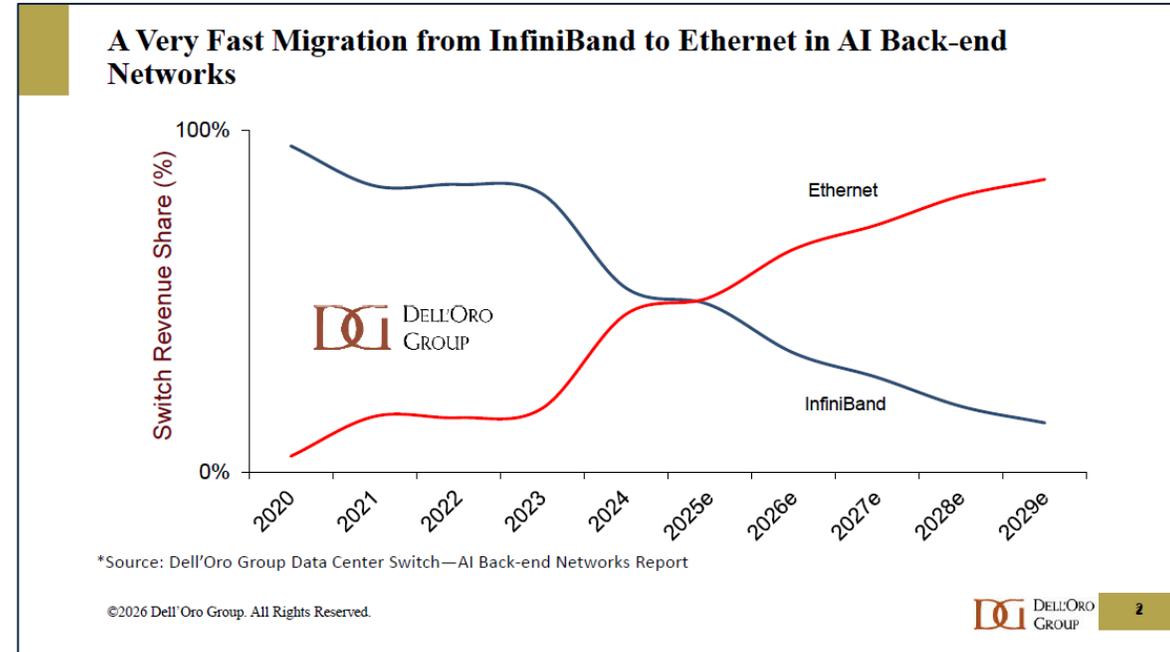
- AI workloads differ widely: training, prefill, and decode each impose different bandwidth, memory, and scale needs
- Latency sensitivity depends on the specific operation in the model
  - Decode paths can be extremely latency-sensitive, while training can tolerate more latency in exchange for resilience, reach, or power efficiency
- Each of these workloads are big enough to justify their own network



Source: John D'Ambrosia. Based on Ethernet Alliance TEF 2024 Keynote – "Ethernet – The foundation of AI @ Meta", Nic Viljoen

# Ethernet is the Choice

- Ethernet standards are the physical layer building blocks for AI networks
- Industry partners want to leverage Ethernet physical layer specs for broadly interoperable AI applications (scale-up and scale-out)



Source: Sameh Boujelbene, Dell'Oro Group

“Ethernet is winning the race in AI back-end networks and was able to eclipse InfiniBand in 2025.”, Sameh Boujelbene, Dell'Oro Group

# Why Now?

- 400 Gbps/lane development is already underway, but no IEEE 802.3 standard exists, creating urgency to align industry investment and avoid fragmentation
- AI network architectures require 400 Gbps/lane to enable high-radix systems – demand is immediate and growing
- Standardization ensures multi-vendor interoperability and leverages Ethernet's broad ecosystem for faster, lower-cost deployment of next-generation technologies

# Summary

- AI infrastructure demands capacity growth
- Ethernet physical layer is the foundation for emerging networks (UALink, UEC, OCP, ESUN, etc)
- Scale-up is the fastest growing and largest segment, followed by scale-out and front end
- 400 Gbps/lane is required for AI network architectures to enable high-radix systems
  - Evolution of known techniques plus advancing technology provide paths forward
- Meeting time-to-market needs requires a focused Ethernet PHY project now
  - Additional PHY developments can follow in future efforts
- Let's form a Study Group!!

# Thank You!

# Next Steps

- A consensus building presentation will be given during the IEEE 802.3 March 2026 Plenary meeting:
  - Tuesday, 10 March 2026 18:30-19:45 PDT (UTC-7)
  - Details [here](#)
- The vote to determine whether a Study Group will be formed will take place at the IEEE 802.3 WG Closing Plenary on the afternoon of Thursday, 12 March 2026
- I would be happy to speak with anybody (in person or remote; please contact me at [kent.lusted@synopsys.com](mailto:kent.lusted@synopsys.com) for remote) to address questions or concerns
- I would be happy to add anybody who wishes to be listed as a supporter
  - You must ask to be a supporter before the motion to form the Study Group is considered

# Straw Polls

Should a study group be formed to develop a PAR, CSD responses, and objectives for “400 Gbps per lane Signaling PHYs and Interfaces”?

Y: 122, N: 0, A: 1

If formed, will you participate in this Study Group?

Y (Individuals): 111

Unique affiliations for those who indicated they'd participate (post-processed): 54