TEF 2025
Ethernet for AI

December 2-3, 2025
Hyatt Centric Mountain View, CA, USA

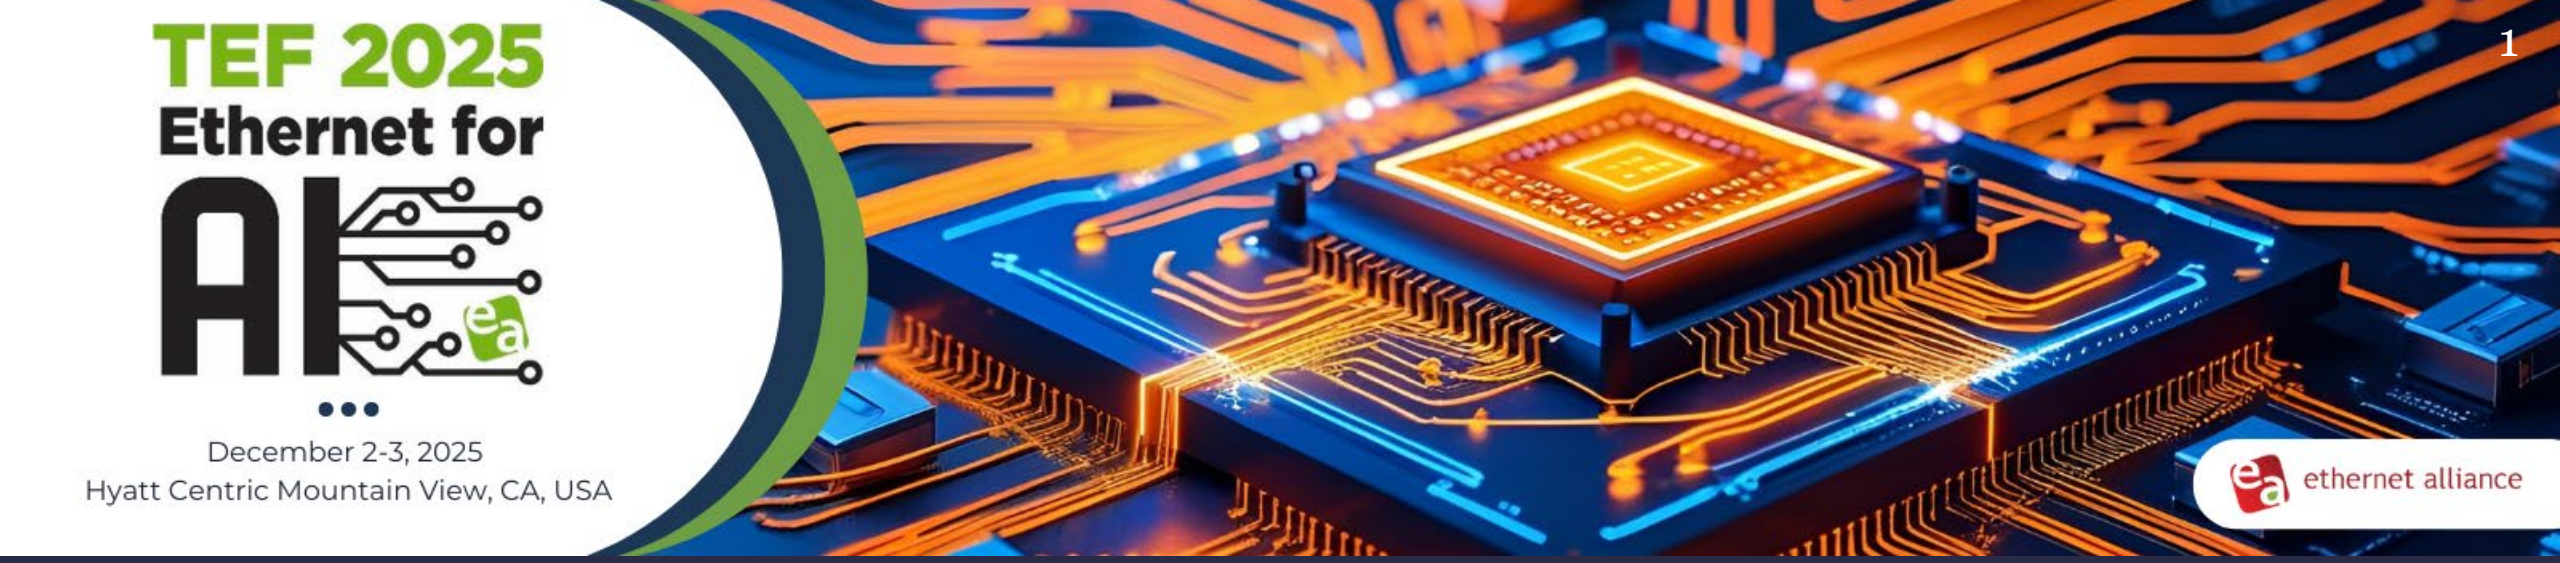

ethernet alliance

Copper Interconnect for 400 Gb/s Signaling for AI Networks

December 2-3, 2025

This presentation has been developed within the Ethernet Alliance and is intended to educate and promote the exchange of information. Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance

Market Update and Panel Introduction

Lisa Huff, Chief Analyst, DC Tech Analysis

Datacom Transceiver Roadmap

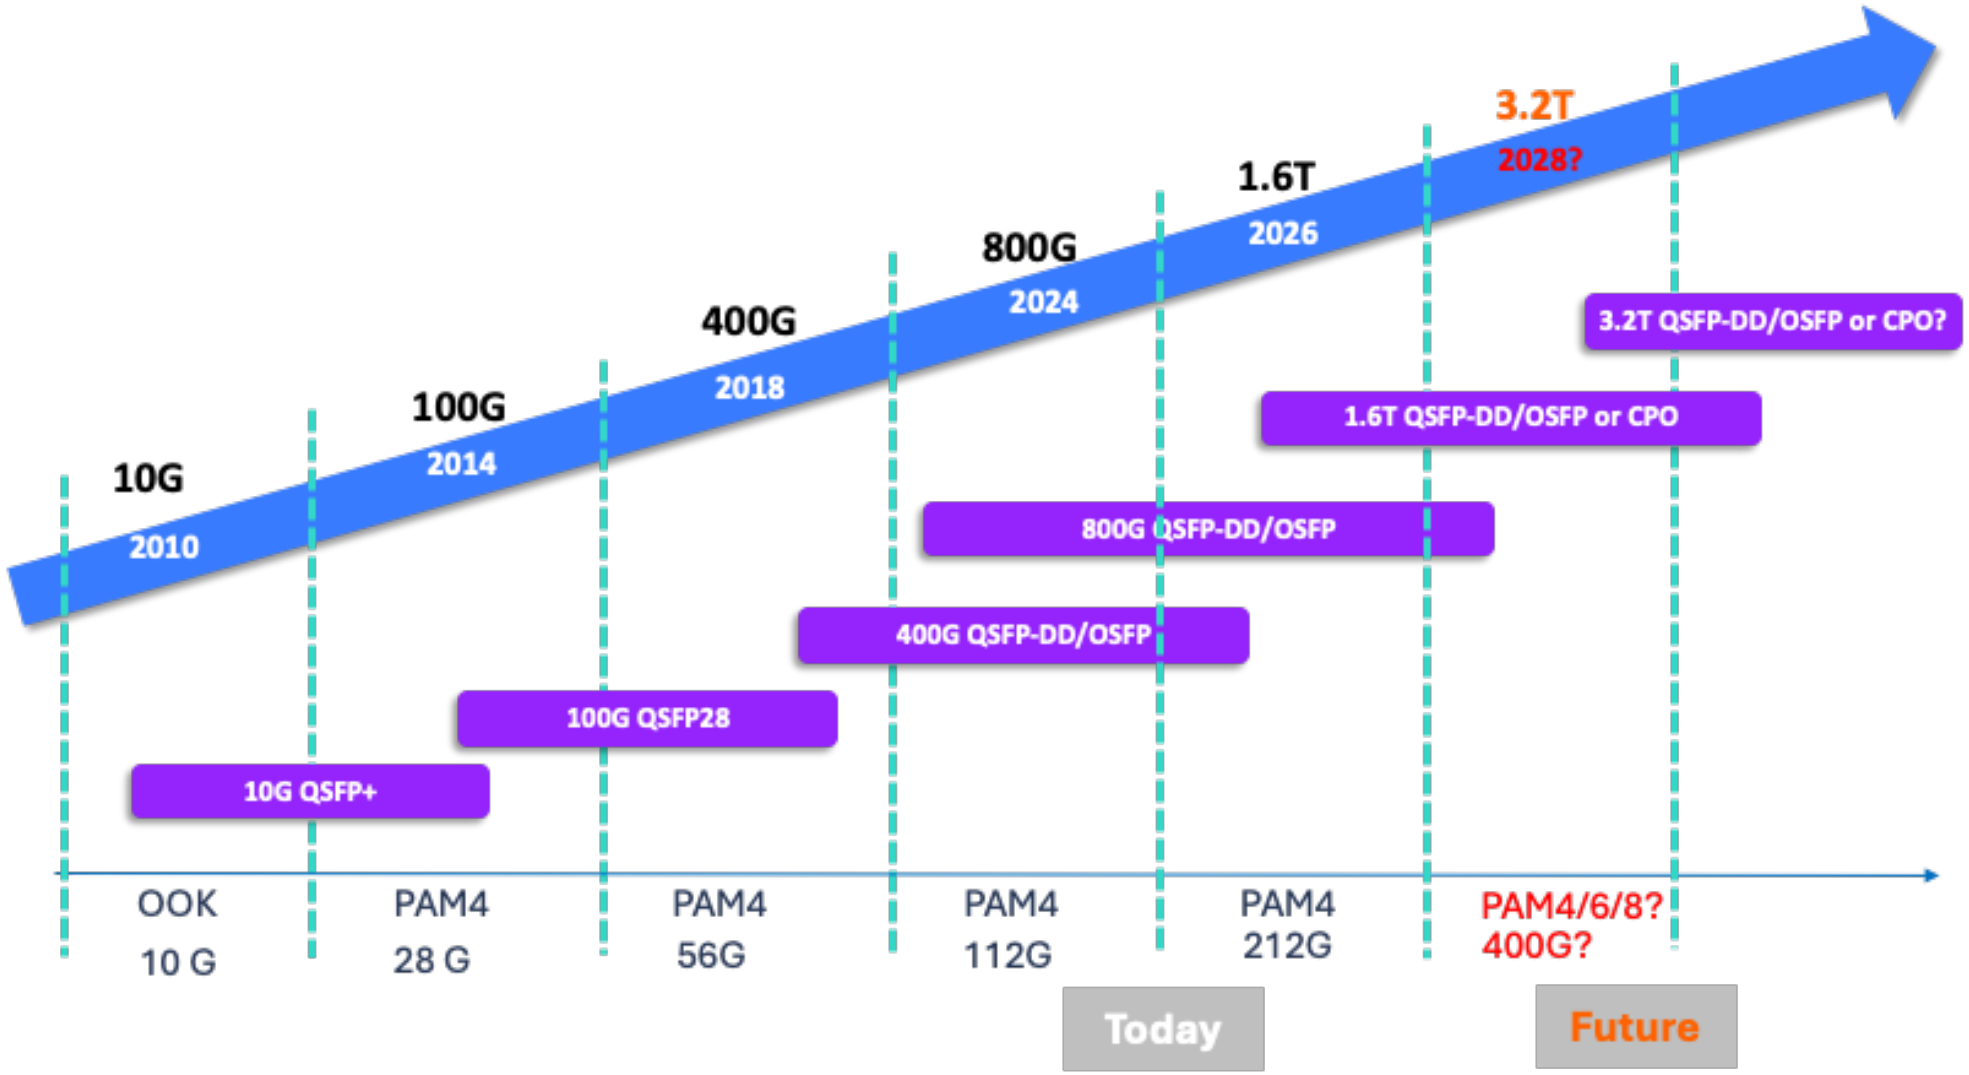

Four major limiting factors for AI data centers

- **Power: To the server rack and to the building:** New server nodes are pushing the AI server power envelope to new heights. With xPU rack density skyrocketing, the power envelope per rack is increasing beyond 200kW. Place rows of these racks in a data center and it becomes difficult to get enough power delivered from the power grid.
- **Cooling in the rack:** Liquid cooling is needed for the higher density AI data centers. While many cloud service providers (cloud SPs) have designs for this, it is a new technology and takes longer to deploy. Either it must be retrofit into existing data centers, or entirely new data centers must be built to support it.
- **AEC/DACs, AOCs, and optical transceiver manufacturing capacity:** While there is already high demand for AECs/DACs, AOCs, and optical transceivers to support new AI deployments, it will accelerate even more this year and beyond. Whether suppliers can keep up with it remains to be seen.
- **Mounting community and political pressure:** Communities do not want to see their electric bills increase and when data centers are built within them, electric bills rise and the availability of electric decreases. Cloud SPs have several tactics to counter this resistance:
  - Their green initiatives
  - Alternatives to connecting to existing power grids.
    - Constellation Energy to restart the Unit 1 reactor on Three Mile Island in Pennsylvania to support AI data center expansion.
    - Adding their own “portable” nuclear energy sources to their new data center cites.
    - New “scale across” architectures to split AI workloads between two adjacent power grids

Global Port Forecasts for AI Networks (both SU & SO)

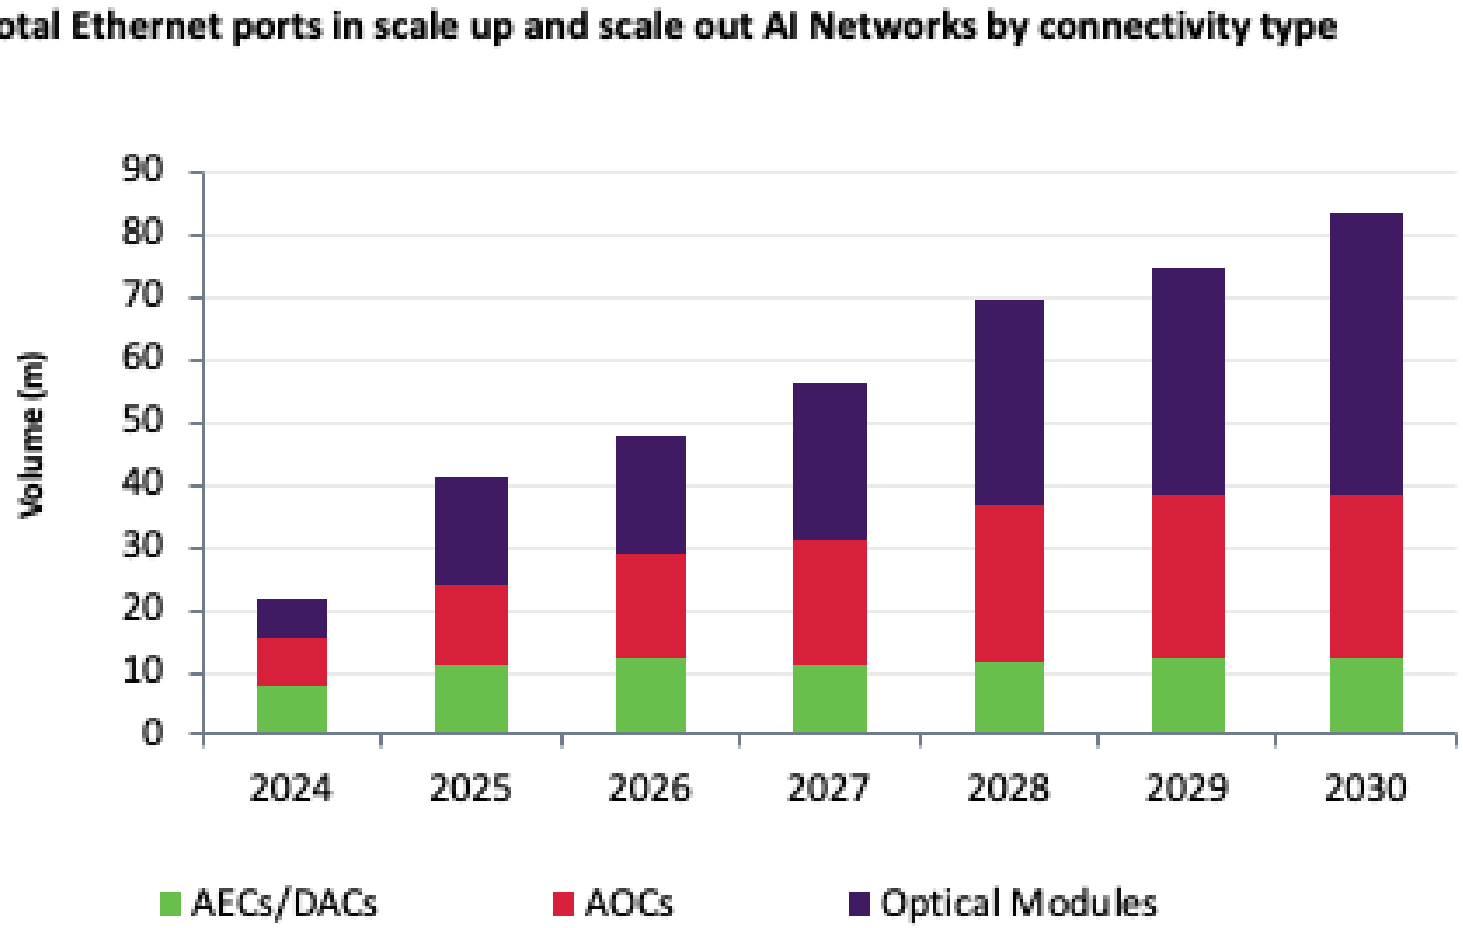

Source: DC Tech Analysis

Global Port Forecasts for AI Networks (both SU & SO)

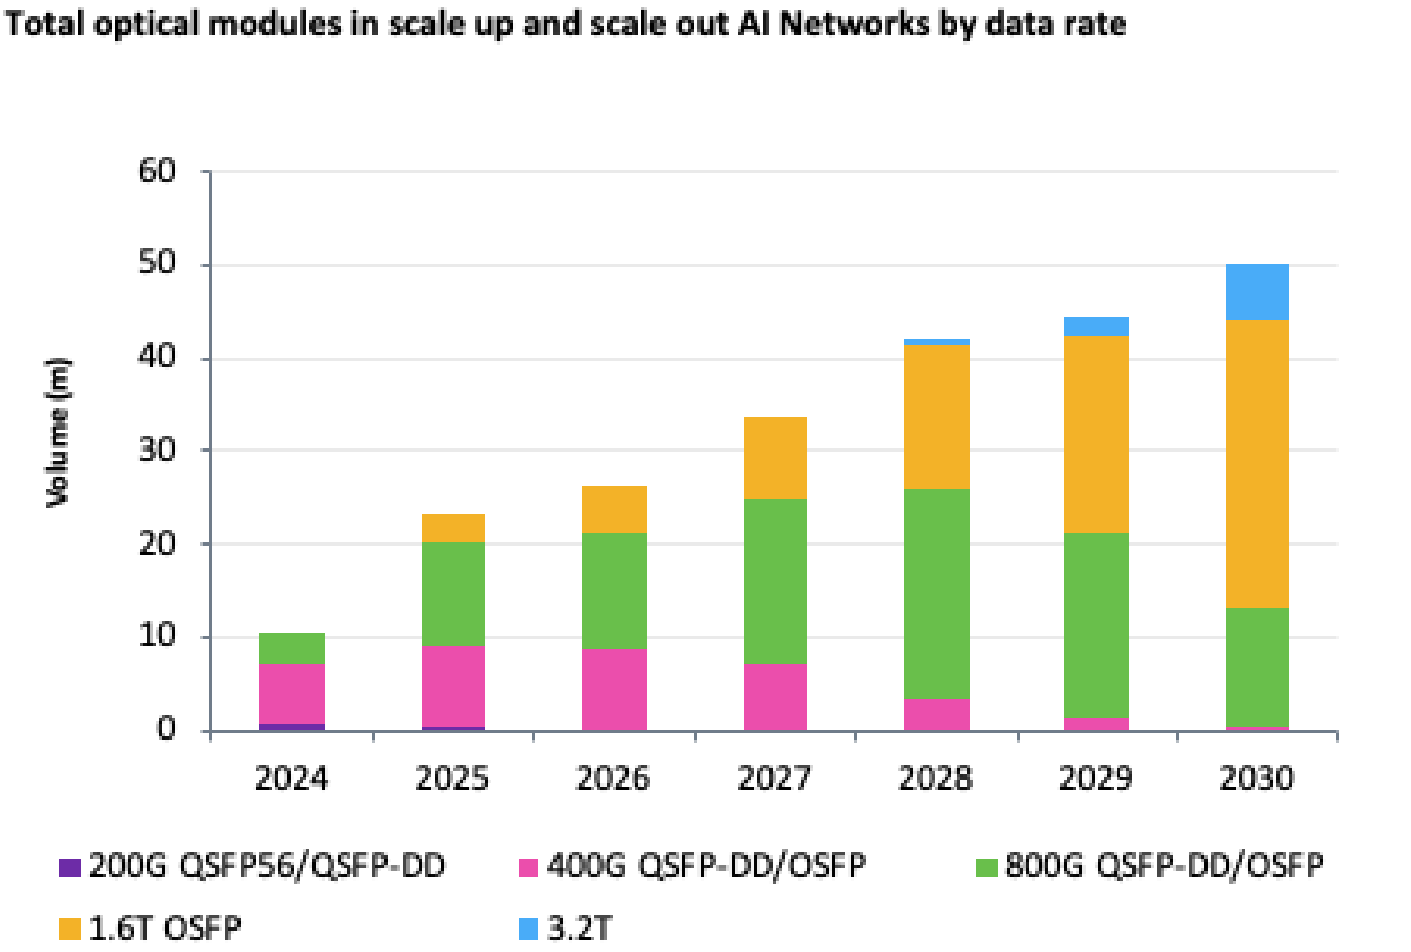

Source: DC Tech Analysis

Global Port Forecasts for AI Networks (both SU & SO)

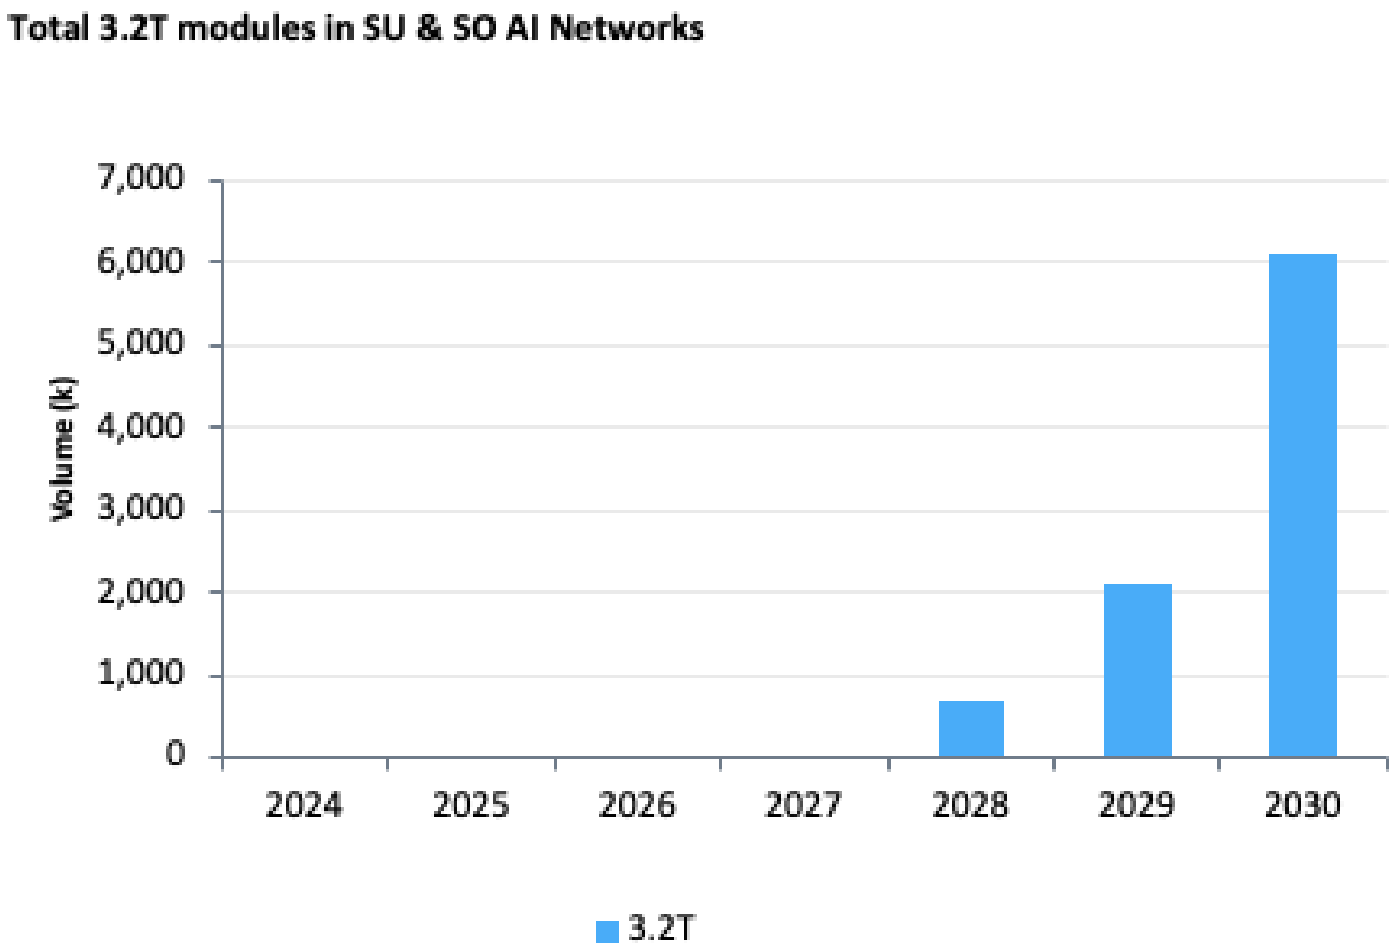

Source: DC Tech Analysis

QUESTIONS?

Copper Interconnect for 400 Gb/s Signaling for AI Networks Panel

Panelists

- **Kasthuri Damodharan, Amphenol**: “Solving 400Gb/s at the Speed of Interconnect”
- **Ashika Pandankeril Shaji, TE Connectivity:** “400G/lane Interconnects for AI: Channel Feasibility, Reach Extension, and Early Measurement Results”
- **Augusto Panella, Molex**: “Copper Backplane Considerations for 400G Signaling”

Solving 400Gb/s at the Speed of Interconnect

Sam Kocsis, Director of Standards at Amphenol
Kasthuri Sankar Damodharan, Solution Architect at Amphenol

Background

- Interconnect plays a critical role in shaping key decision for next gen 400 per lane of Ethernet
- Focus on high-speed copper links, co-packaged topologies
- Amphenol was the first interconnect company to contribute models (revisit) for the IEEE E4AI project, in March 2025
- Observed a significant interest on defining copper implementations in the anticipated new IEEE project
  - Chip-to-Module / Chip-to-Chip
  - Pluggable IO / Cabled-Backplane
- New form factors take time to develop consensus
  - Standardization – interoperability and broad deployment at scale
- Received feedback on the *Call to Action* items, from March 2025

Evolution of Copper Links

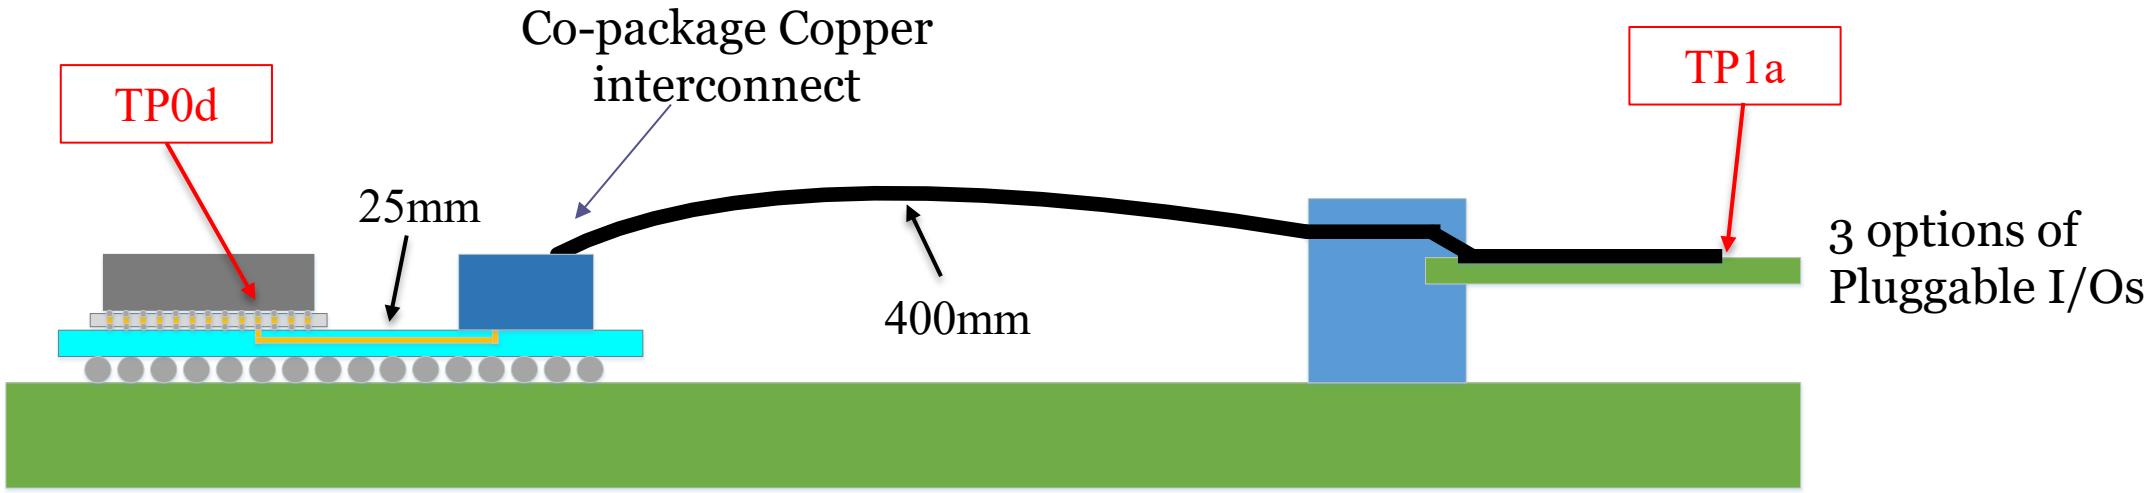

Call to Action:

1. Feedback on topology relevance
2. Thoughts on reference/compliance points
3. Feedback on model bandwidth
4. Opinions on data rate and modulation

Source: [kocsis_e4ai_01_250327]

Topology

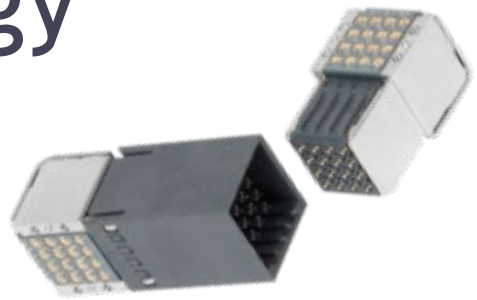

Backplane Core Technology

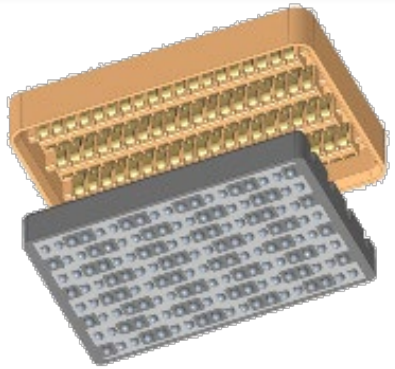

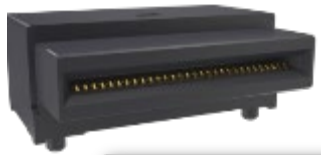

IO Connector

Cabled-Backplane

Twin-ax Cable

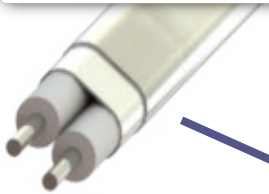

Pluggable Modules

Near-chip / On-chip

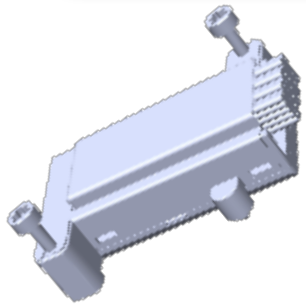

Cabled-Host

- **All** copper interconnect options are on the table for 448G
- Current gen are centered around twin-ax cabling, Twin-ax performance is critical
- Cable standard requirements - Cabled-backplane “KR” strongly resembles “CR” implementations

Compliance Points

- Compliance points at mechanical interfaces, can be problematic

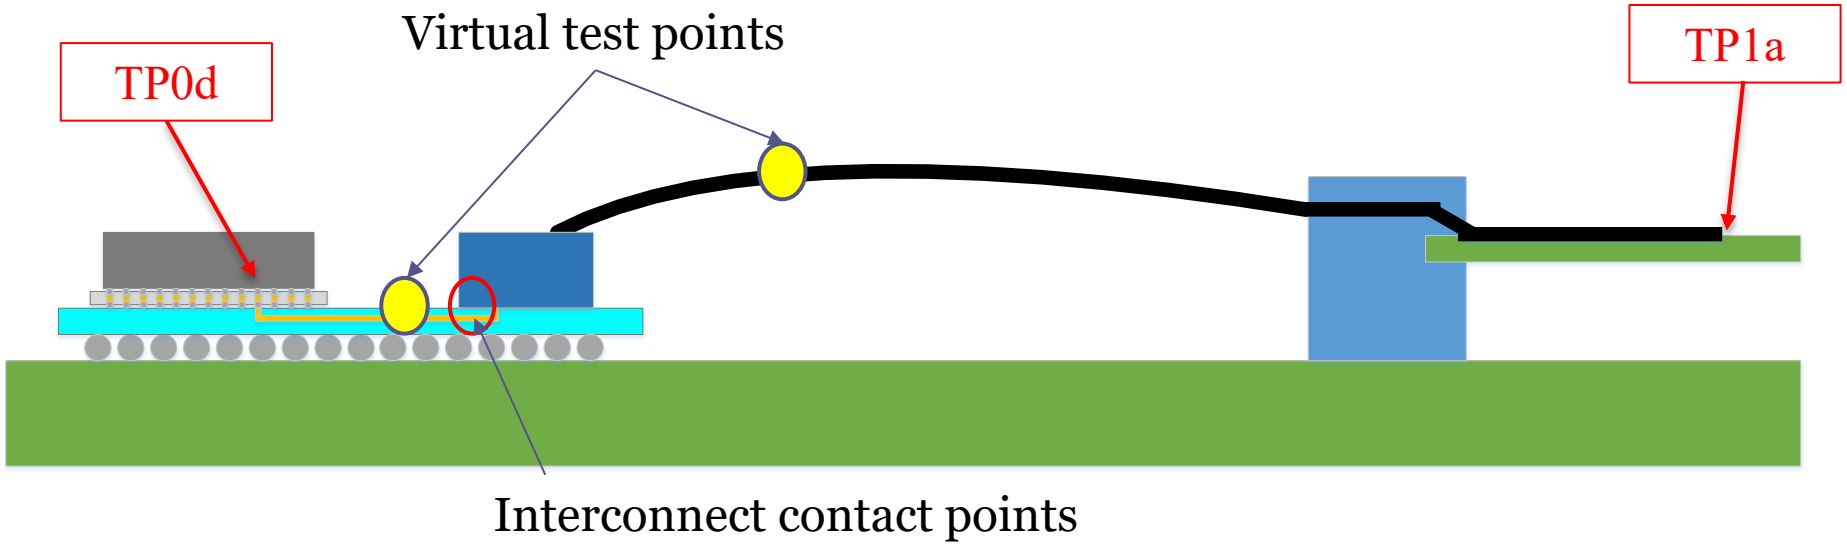

- Cabled-Host topologies introduce new targets for test points
- Virtual compliance points may provide multi-source flexibility
  - New test fixtures and measurement capability are necessary

Model Bandwidth

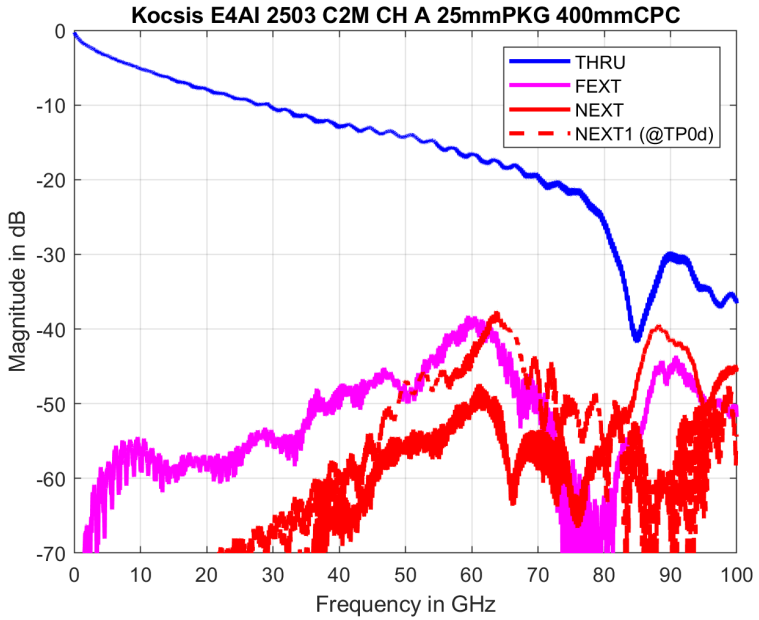

"Limit of today's Pluggable"

Channel A

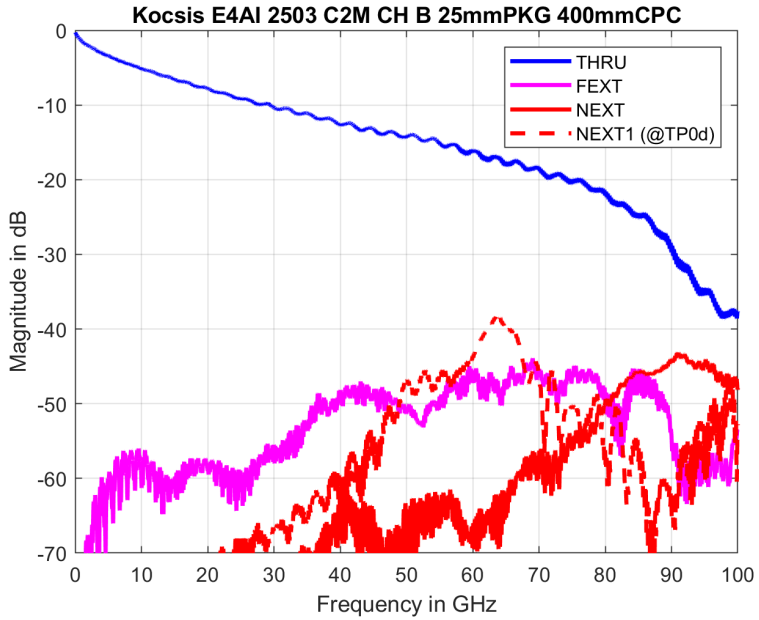

"New Pluggable, Familiar Feel"

Channel B

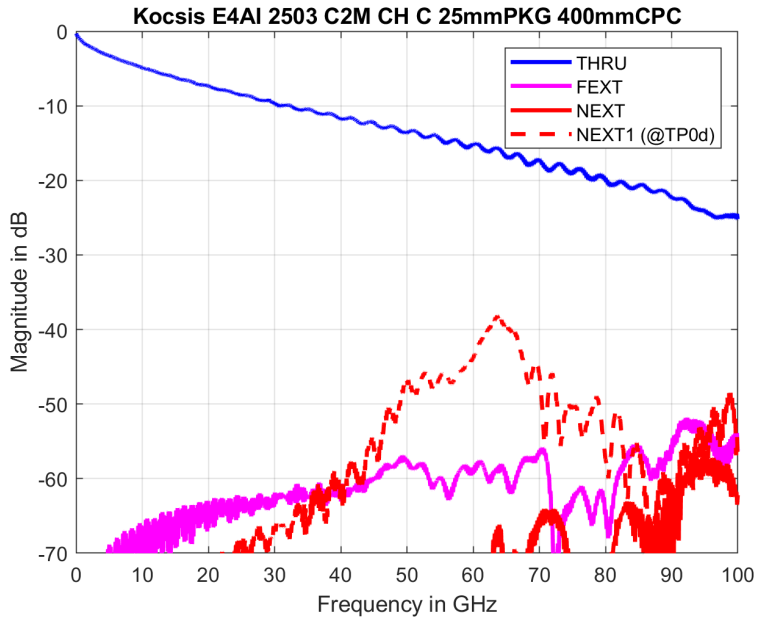

"New Pluggable, New Paradigm"

Channel C

- 100 GHz is insufficient bandwidth to make key decisions, but…

Functional Bandwidth

- *Functional bandwidth* refers the frequency range used to characterize the interconnect for quality and compliance

IEEE Project	Nyquist Frequency	Functional Bandwidth	
"3bj"	12.89 GHz (NRZ)	19 GHz	+47%
"3cd"	13.28 GHz(NRZ)	25 GHz	+88%
"3ck"	26.56 GHz (NRZ)	40 GHz	+50%
"3dj"	53.125 GHz(PAM4)	67 GHz	+26%
TBD	106.25 GHz	130 GHz	+22%

- 130 GHz is the practical limit for the next IEEE project

Data Rate and Modulation

	Channel A	Channel B	Channel C	Comments
Backwards Compatible	++	-	-	Backwards compatibility not required
Interface Bandwidth	85 GHz	95 GHz	100 GHz+	*Approximately, Slide 9*
Crosstalk	-	-	-	Package must improve
Supported Modulation	PAM-8	PAM-8/6	PAM-8/6/4	Channel C preferred for all signaling rates
HVM Readiness	++	+	-	Time to market is critical
Time-to-Market	2Q	3Q	6Q	*As a "standard" solution*

	"224G"				"448G"		
Number of signal levels (M)	8	6	4		8	6	4
Bits per symbol	3	2.5	2		3	2.5	2
Signaling rate*, GBd	75	90	112		150	180	224
Nyquist frequency, GHz	37.5	45	56	...	75	90	112

* Signaling rate rounded for simplicity

System-level Architectural Goals

	Scale-Up	Scale-Out	Scale-Across
Expansion Architecture	Vertical *Intra-rack*	Horizontal *Inter-rack*	Vertical + Horizontal *Inter-cluster*
Application	Compute Fabrics + Storage/Memory	Data Center Fabrics	Global Fabrics
Transmission Distance	<10m	10s-100s of m	2km+
Transmission Medium	Copper/Optics	Optics	Optics
Transmission Technology	Ethernet	Ethernet	Ethernet

- *Density* is key on backplane, *Scalability* is key on front-panel
- Front-panel support for both copper and optics, in the same port
- Flexibility for module cooling, not the focus of this contribution

Emerging Form Factor(s)

- Based on *OSFP-like* envelop
- Copper "fit" in OSFP-XD is not the same as OSFP
- 64-lane form factor advantageous for optics
- 8-lane form factor preferred for copper
- A *new* 8/16-lane form factor preserves the copper/optics ecosystem

Faceplate Size

Lanes per module	1RU	2RU	4RU	
8	256	512	1024	16x across 1RU
16	512	1024	2048	
32	512	1024	2048	8x across 1RU
64	1024	2048	4096	

Aggregate IO Lane Density versus Faceplate Size

Twin-ax Copper Size

Lanes per module	26AWG	32AWG
8	32%	14%
16	64%	29%
32	36%	17%
64	72%	34%

Copper Density Utilization

Summary

- Progress over the last 6months shows that an ecosystem of interconnect solutions with >100 GHz bandwidth is possible
- Channel C is the winner!
  - New form factor, standardized by 2027
  - 8-lane, extensible to 16-lane, form factor is preferred
  - Functional bandwidth of 130GHz is the practical target
- Copper, until the limits give out
- PAM-4, as far as it will go
- Strong engagement between Silicon, Packaging and Interconnect is necessary to develop the supply chain for 400Gb/s

QUESTIONS?

448G/lane Interconnects for AI: Channel Feasibility, Reach Extension, and Early Measurement Results

Ashika Pandankeril Shaji, TE Connectivity

Megha Shanbhag, TE Connectivity

Nathan Tracy, TE Connectivity

Jason Ellison, TE Connectivity

Ed Frlan, Semtech

Sameh Elnaggar, Semtech

448Gbps – PAM 6 vs PAM 4

- Modulation scheme is being debated
- PAM 4 loss is high for passive channels
  - **Chip to module short reach with CPC feasible**
  - Longer channels will be very lossy
  - Lower power **active copper channels ACC/AEC** over optics
- PCB host routing not viable for 448Gbps
  - Need **Co-packaged Copper(CPC)** cabled solution
  - Density is key – need thinner AWG cables
- Traditional IO pluggable connectors have performance limitations
  - Need **new IO form factors** to achieve 448Gbps

** Early Estimates*

Component	PAM 4 @~112GHz	PAM 6 @~85GHz
Bulk Cable	~17dB/m	~14dB/m
IO Connector*	~4dB	~2.5dB
Co-Packaged Connector	~1 dB	~0.7 dB
Substrate Loss	~10-15dB	~5-8dB
Total Loss for Typical Channel (TP0-TP5)	>42dB	~36dB
Chip to Module Channel	~28dB	**~18dB**

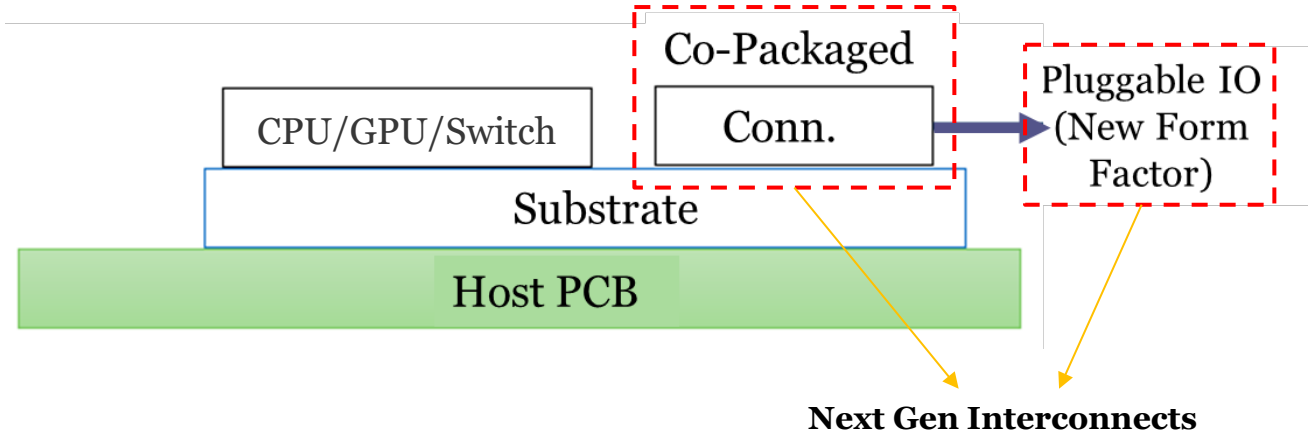

448G New I/O Form Factor – Connector SI Results

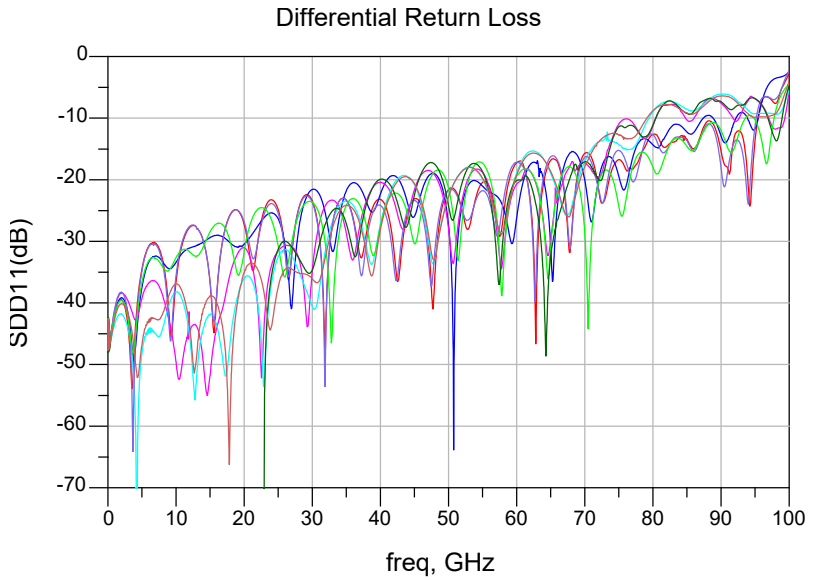

Model Details OTB Version

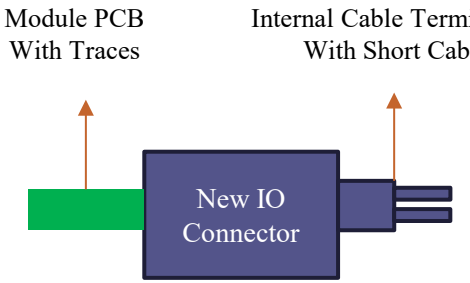

Modeled data

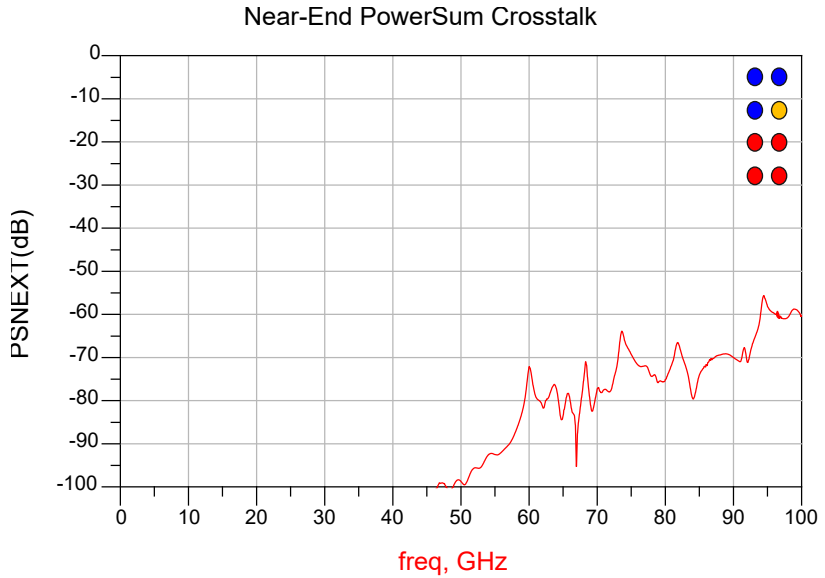

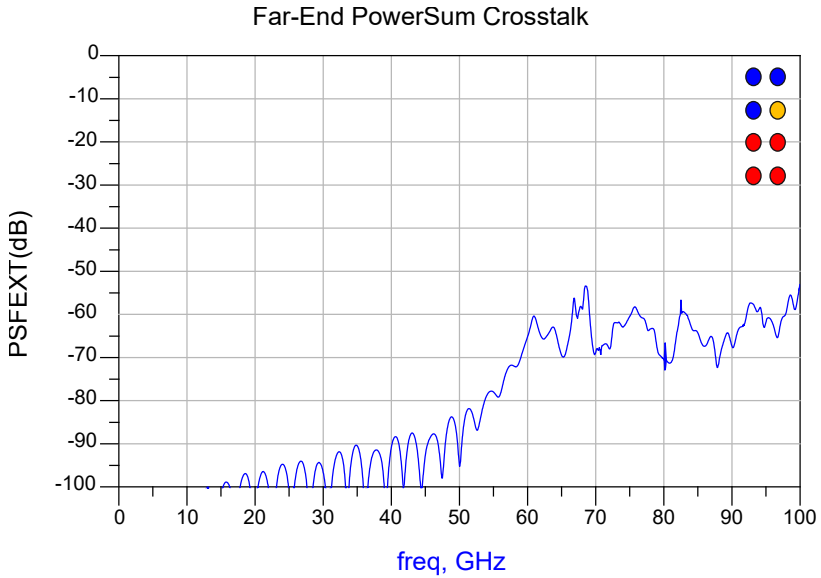

- Victim pair
- Near-end XT aggressor pairs
- Far-end XT aggressor pairs

448G CPC Connector– SI Results

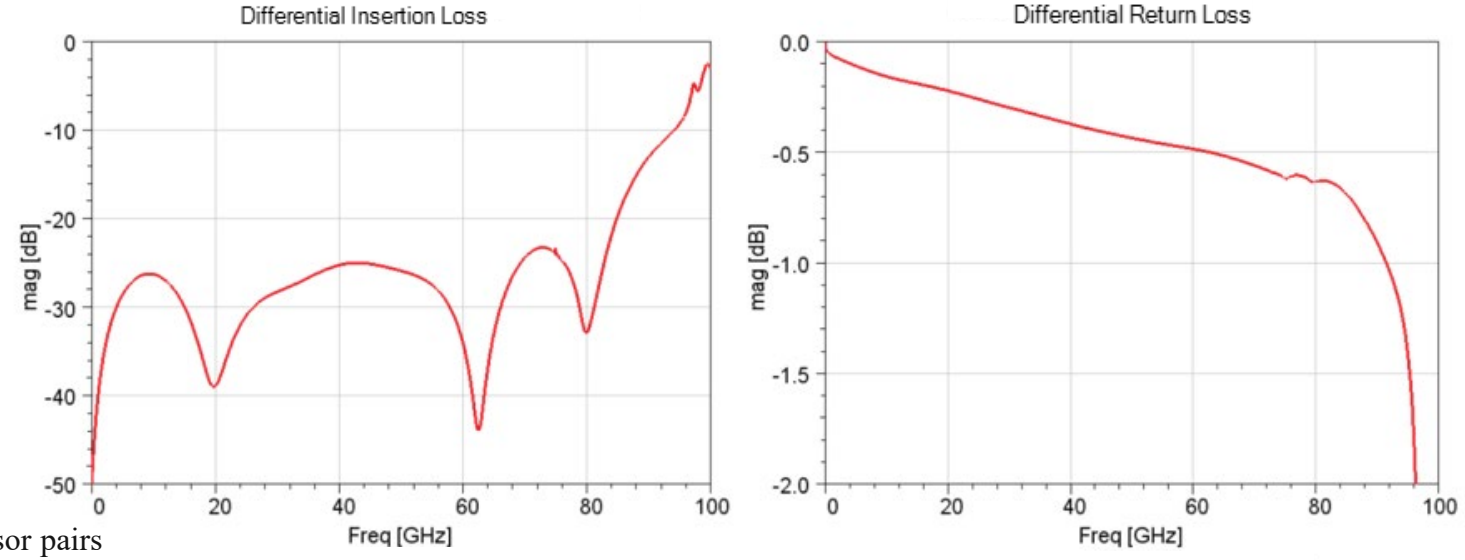

- Victim pair
- Near-end XT aggressor pairs
- Far-end XT aggressor pairs

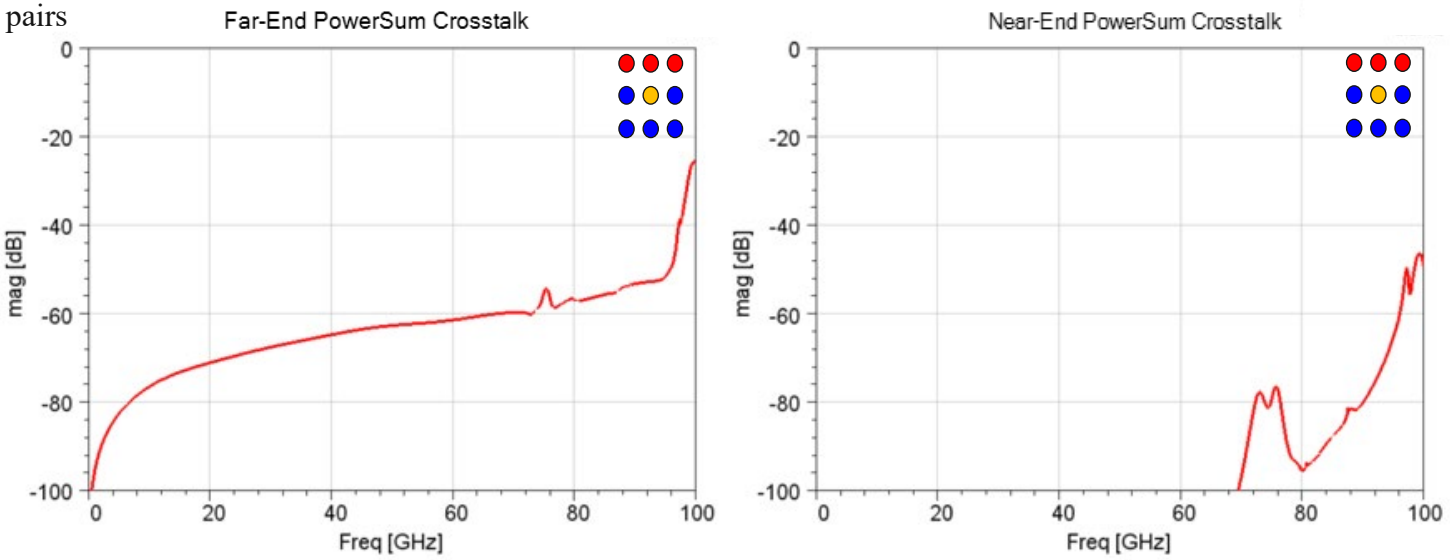

Model Corresponding to SI Results

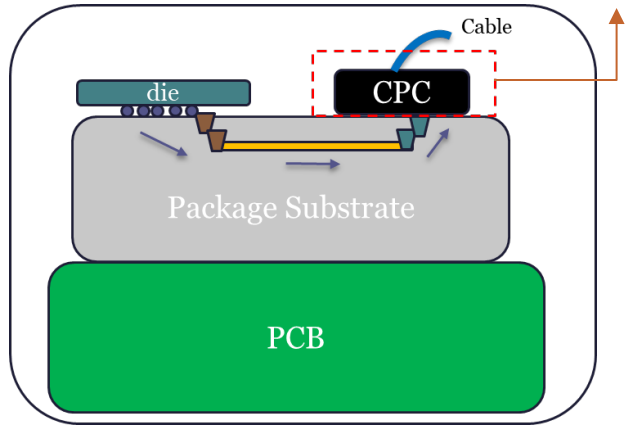

CPC Implementation

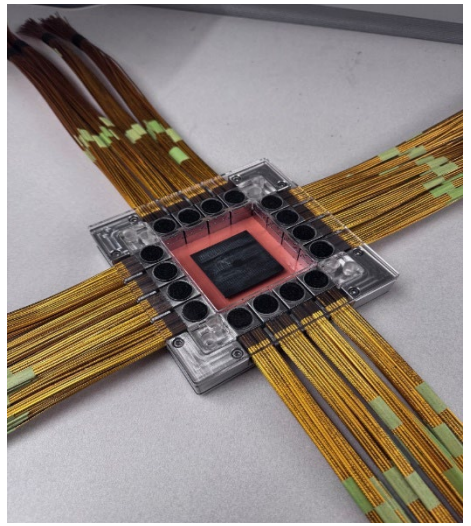

TEF 2025 Ethernet for AI

Channel 1: Chip to Module 30AWG 250mm(CPC to IO)

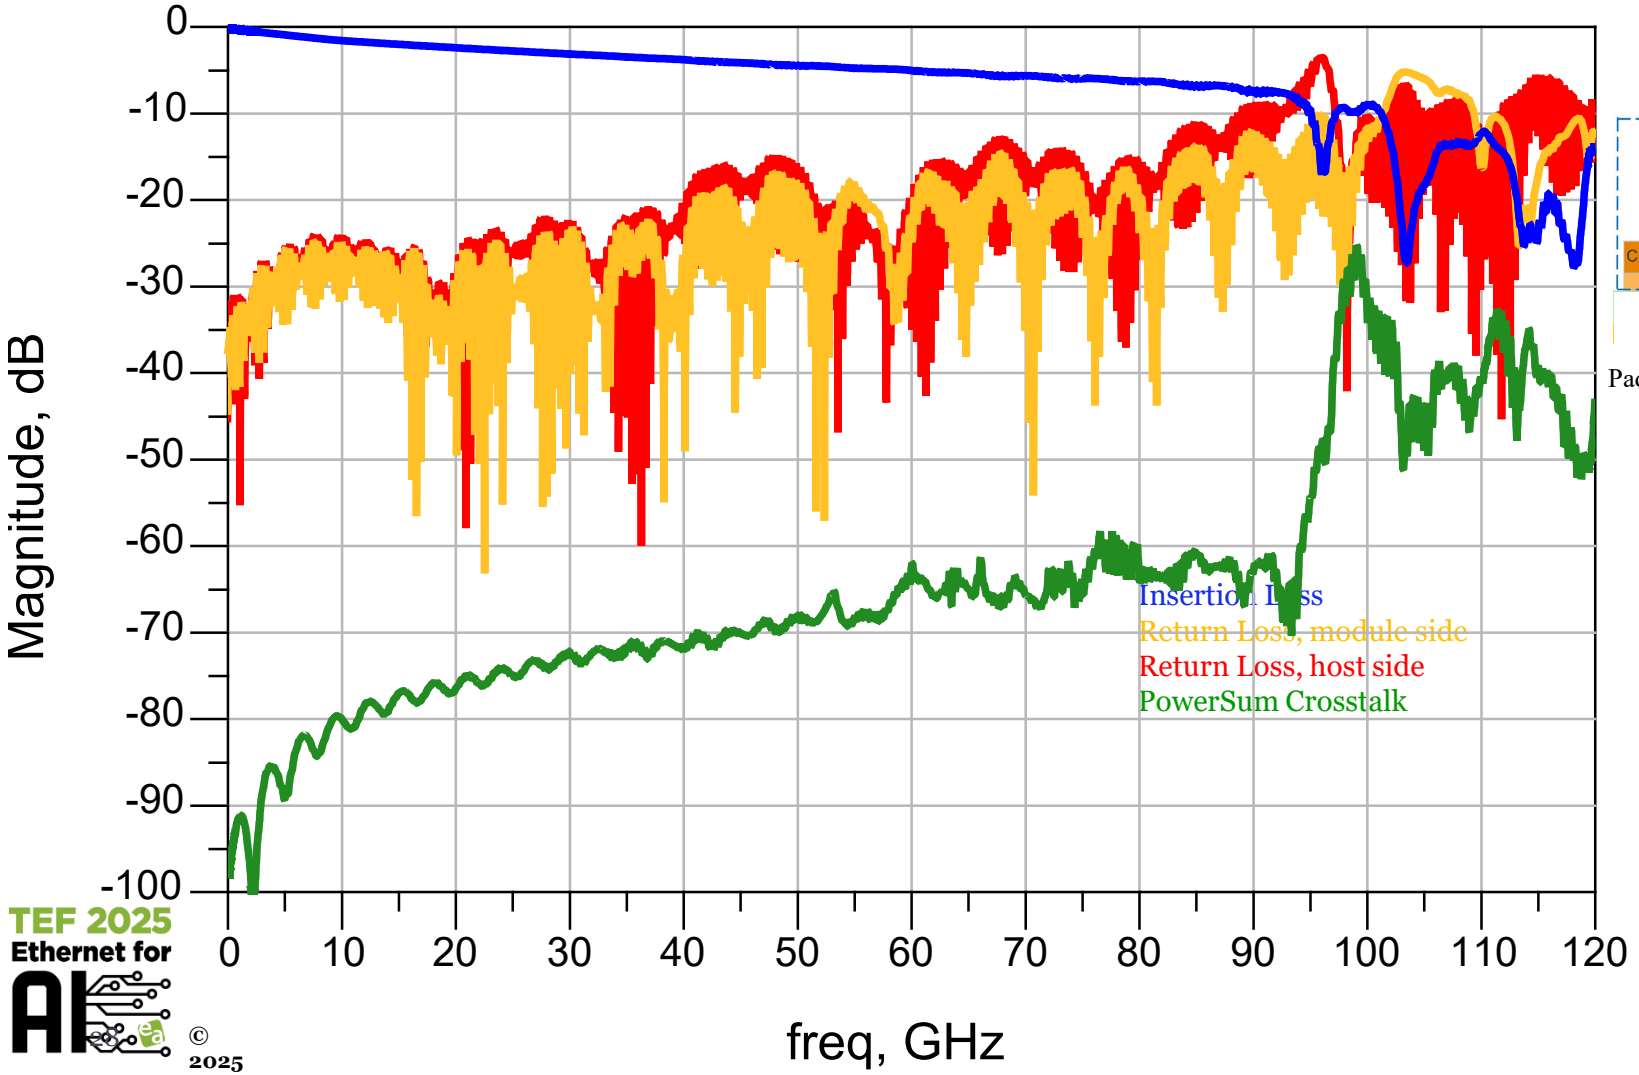

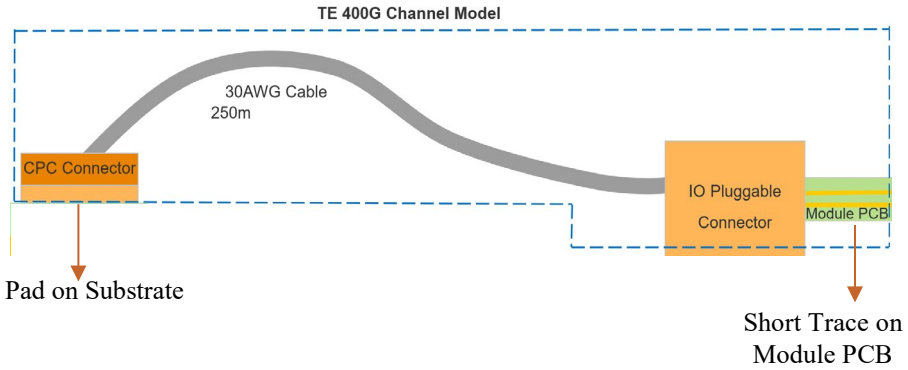

Crosstalk includes 5 aggressors
3 FEXT and 2 NEXT, module side

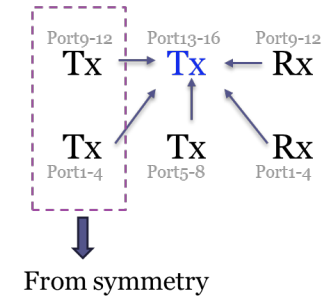

Channel 2: CPC/DAC/CPC 1M 26AWG DAC & 30AWG 250mm Internal Host

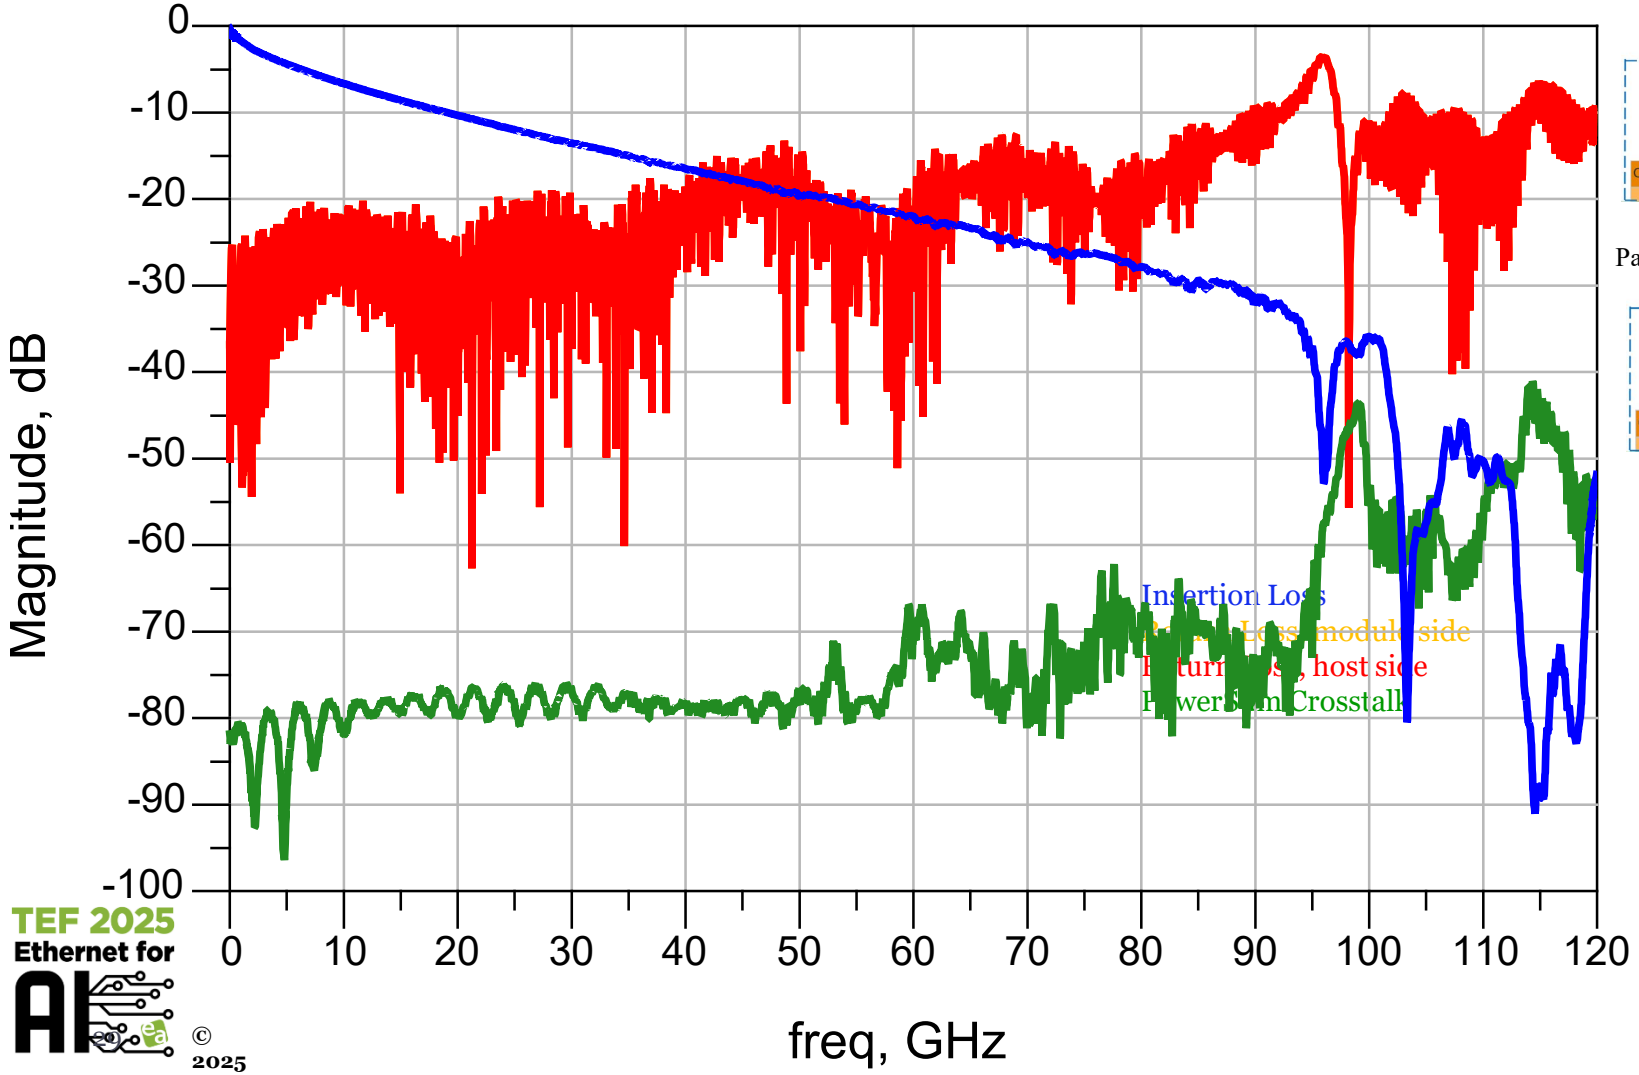

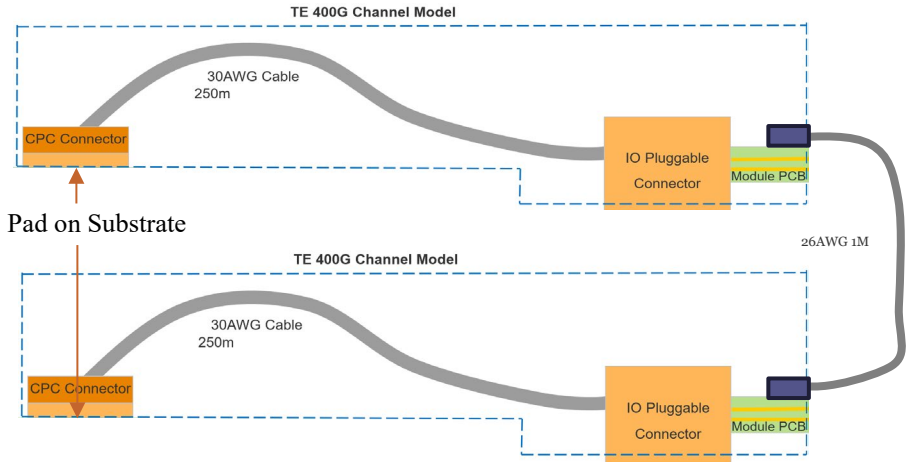

Crosstalk includes 5 aggressors
3 FEXT and 2 NEXT, module side

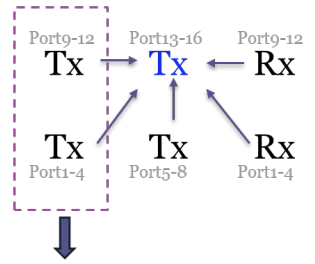

110GHz Current Testing Challenges

- Test lead length is limited to 6-inches
  - 64 port multiport VNAs are needed
- De-embedding up to 110 GHz is challenging
  - High attenuation and test point tuning difficulties
- 1mm multipin test points reduce loss at the cost of crosstalk degradation
- Mechanical calibration standards are limited to 250 mating cycles
  - Require 118 measurements to calibrate a 110 GHz 16 port VNA.
- 1mm connectors are easily damaged due to pin and socket misalignment.
  - Over torquing grooved ring or coupling nut
  - PCB mounting hole tolerance to the signal pad and PCB layer-to-layer

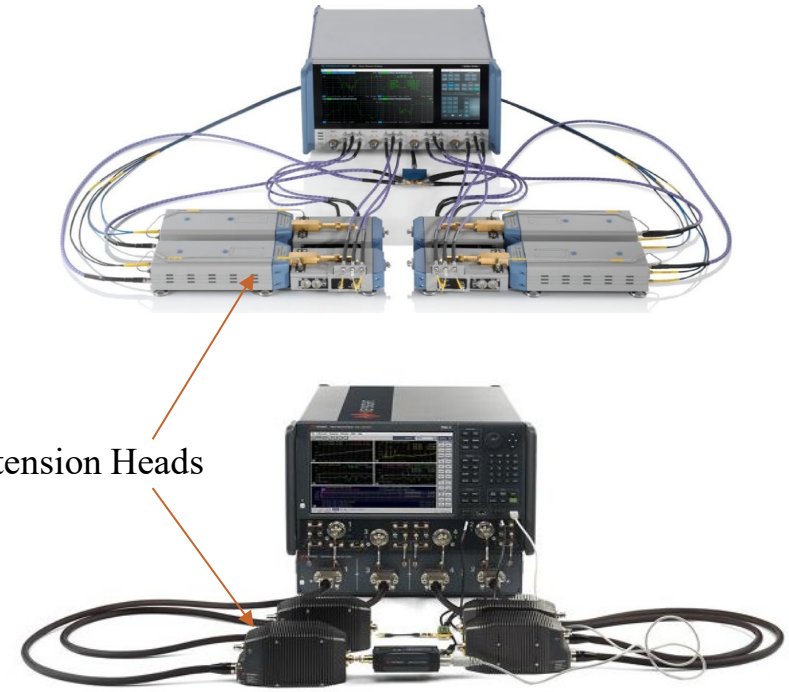

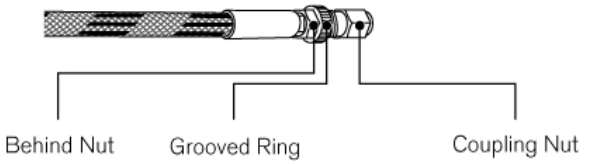

1mm Test Point

448Gbps Measured Data

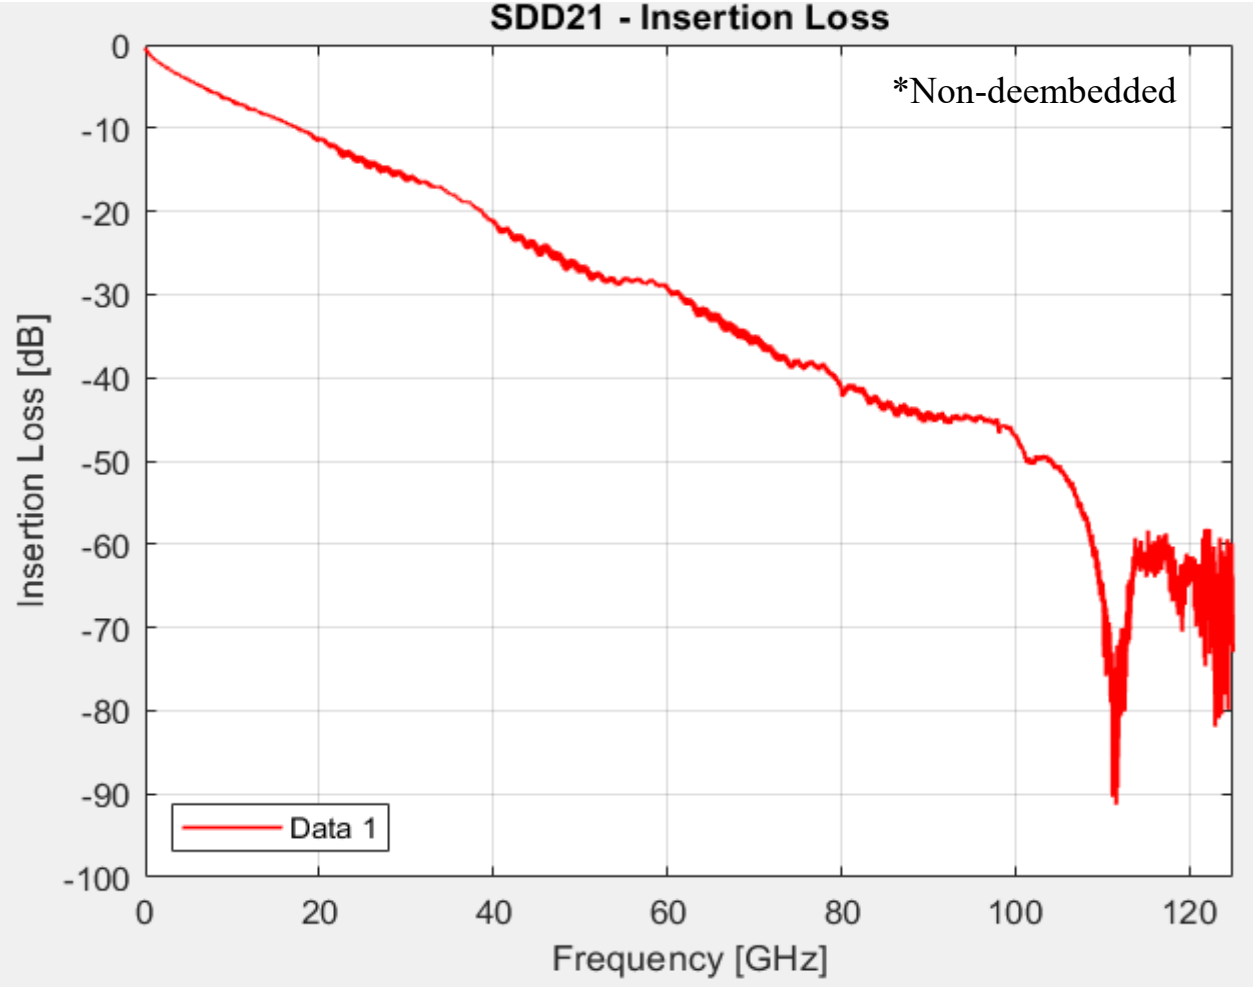

448Gpbs Live Demo at OCP 2025

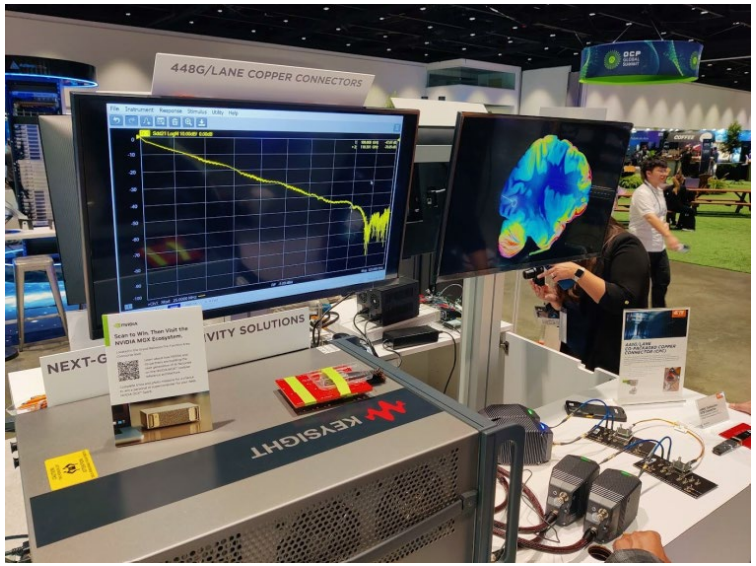

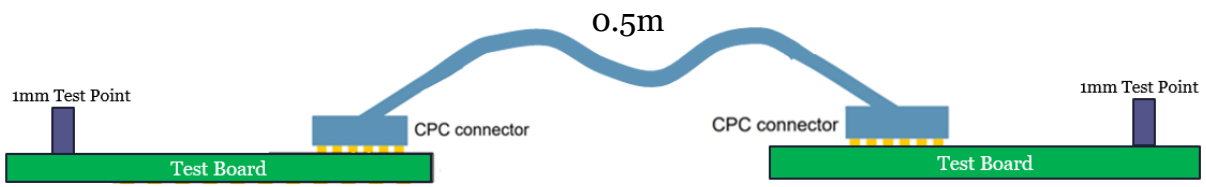

- 0.5m 30AWG “448G” cable assembly
- 448G test boards with 1mm test points
  - Insertion loss of 2x thru ~30dB
- 125GHz VNA
- Resonance free till 105GHz

System BER Simulation

- Three external (front-panel pluggable) cable configurations were simulated at 448Gb/s per lane with co-packaged copper host channel:
  - Traditional passive direct attach copper (DAC)
  - Far end equalized active copper cable (FEE ACC)
  - Dual end equalized active copper cable (DEE ACC)
- Initial results show DAC is feasible with ACCs extending reach and margin

ACC Detailed System Block Diagram

Host Tx parameters:
- Statistical PRBS pattern
- f_{baud} = 448 Gb/s
- Modulation: PAM-4, PAM6, PAM-8
- V_{swing} = 800 mVppd
- Tx SNDR
- Random and Dual-Dirac Jitter

Host Reference Equalizer parameters:
- CTLE (OIF CTLE scaled to approx. f_{Baud})
- RX Input Referred Noise (eta$_0$)
- ADC ENOB
- RX FFE
- Single cursor MLSD

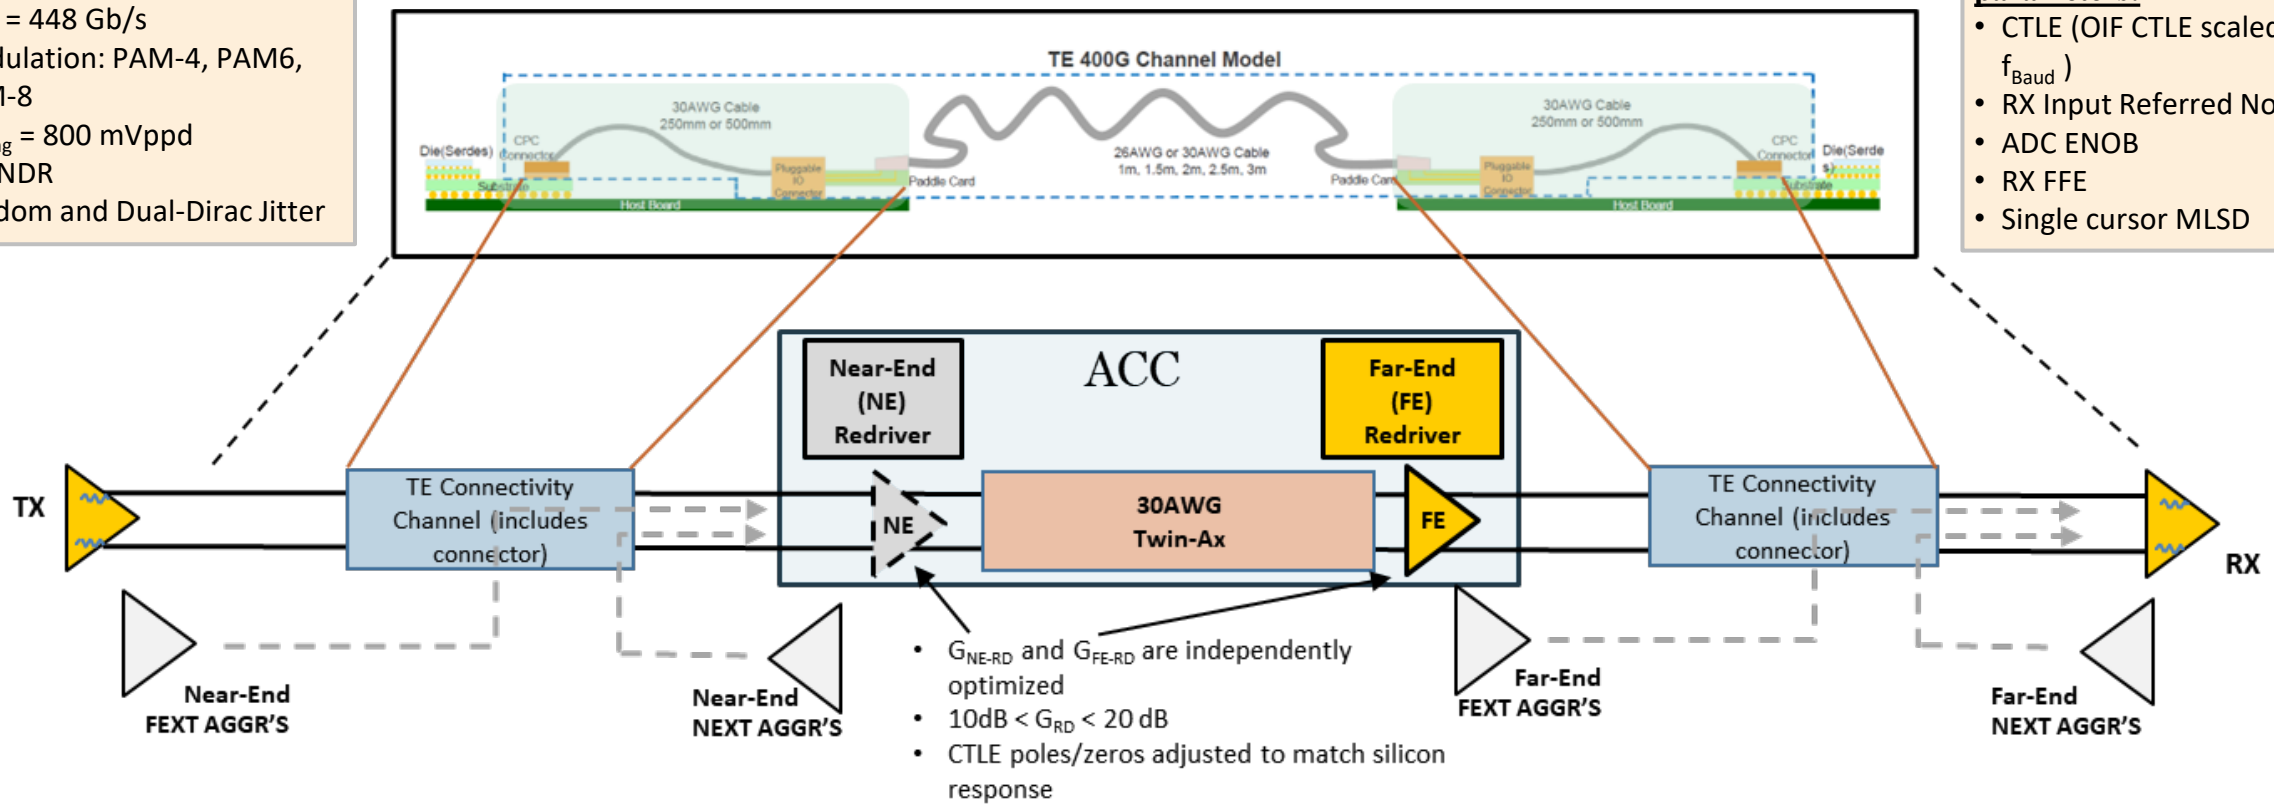

BER vs Modulation – 30AWG 1M Cable Length

PAM4

Configuration	BER
DAC	1.40×10^{-1}
ACC, Far-End	8.28×10^{-2}
ACC, Dual-End	7.26×10^{-2}

PAM6

Configuration	BER
DAC	2.00×10^{-5}
ACC, Far-End	$\sim 1.00 \times 10^{-7}$ *
ACC, Dual-End	$\sim 1.00 \times 10^{-7}$ *

PAM8

Configuration	BER
DAC	3.00×10^{-5}
ACC, Far-End	6.17×10^{-6}
ACC, Dual-End	3.7×10^{-6}

- CPC and improved connectors enable DAC cables up to 1m
- ACC improves BER and allows longer reach up to 2m

ACC BER vs Cable Length

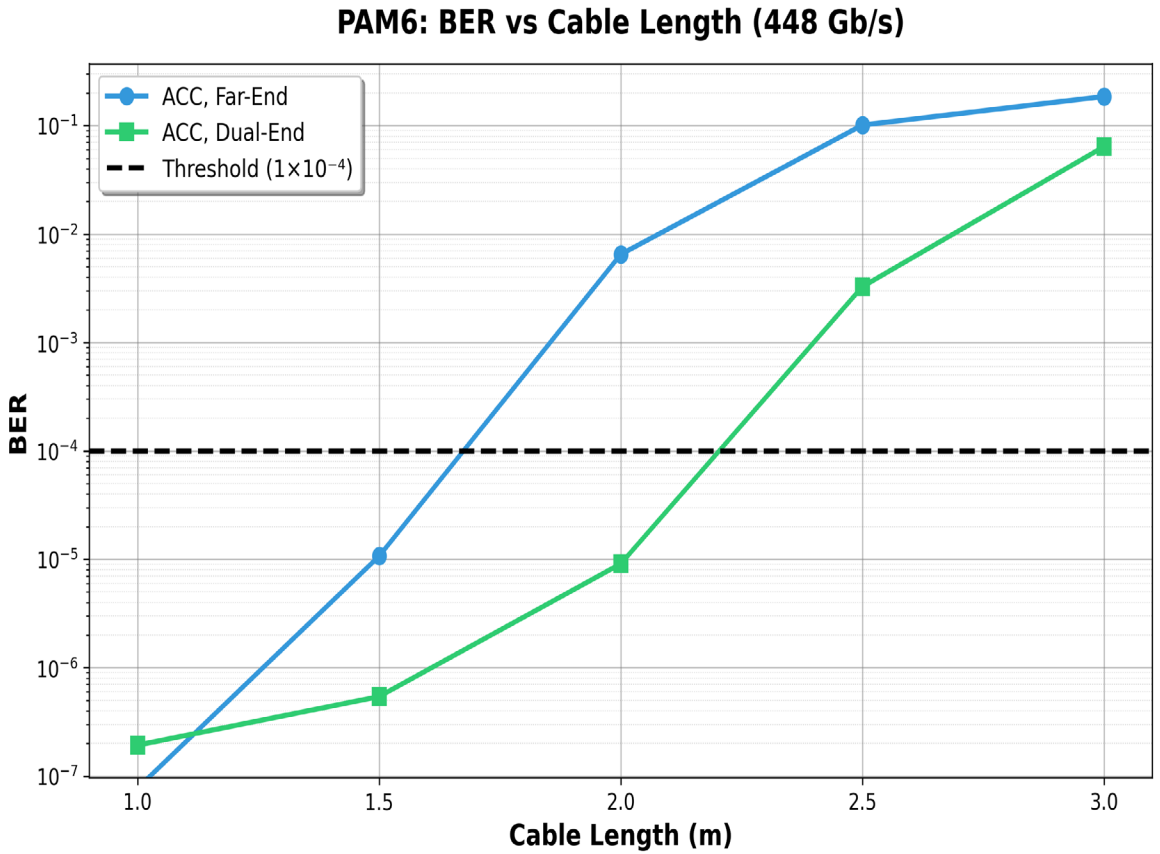

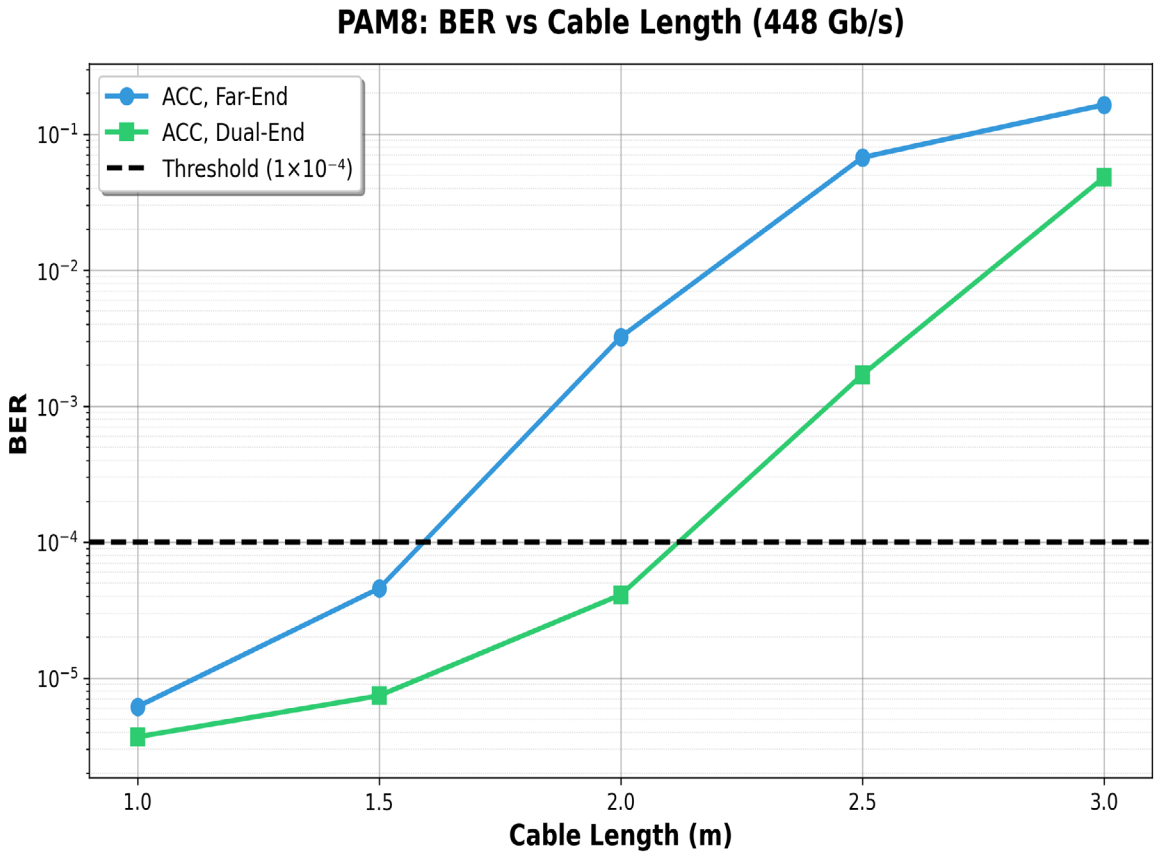

Summary

- 448G electrical links are feasible
- Path to 448-PAM4 with passive copper links might be challenging
- Passive copper cable assemblies could be feasible with PAM6/PAM8
- Alternate form factor allows clean slate for robust interconnect design for 448G
- Active components like redrivers can be used to extend reach for electrical links

QUESTIONS?

Copper Interconnect for 400 Gb/s Signaling for AI Networks

Gus Panella; Director, Interconnect Technology; Molex, LLC]

Abstract

- The advancement of 448G signaling architecture is determined by the cumulative budget of all components within a signaling channel. Both AI and HPC implementations, whether utilizing optical or copper solutions, will encounter industry-specific challenges.

- The mechanics of copper interconnects play a crucial role in defining the performance range required to understand impact on large-scale implementations. This study focuses on an idealized copper cable backplane channel, examining key mechanical impacts on signaling objective.

Recommendations

- In some case, 0.25mm demate was determined to be adequate
- Flush mating for 224G recommendation will continue to 448G
- For 448G, we expect the mechanics to drive line-to-line mated interfaces

Configuration

- 26AWG backplane connection
- Wire-to-wire
- Swept de-mate
  1.00mm, 0.75mm, 0.50mm, .25mm, Flush

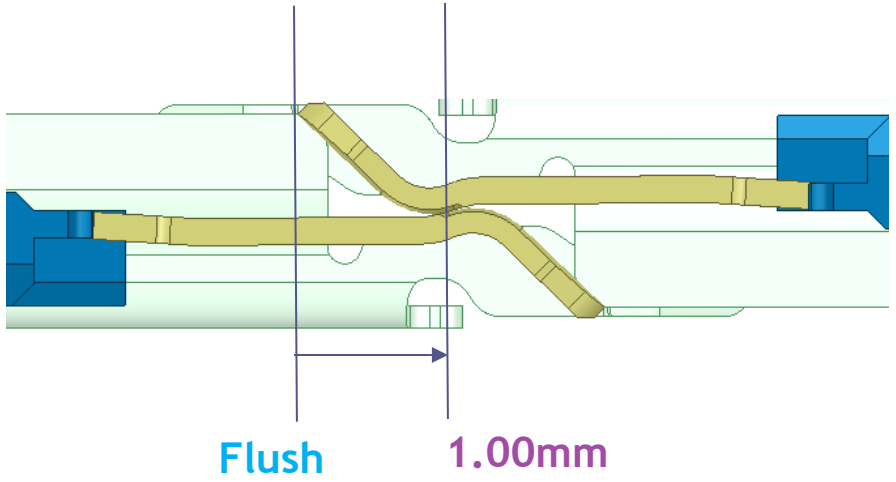

1.00mm
0.75mm
0.50mm
0.25mm
Flush

SDD21

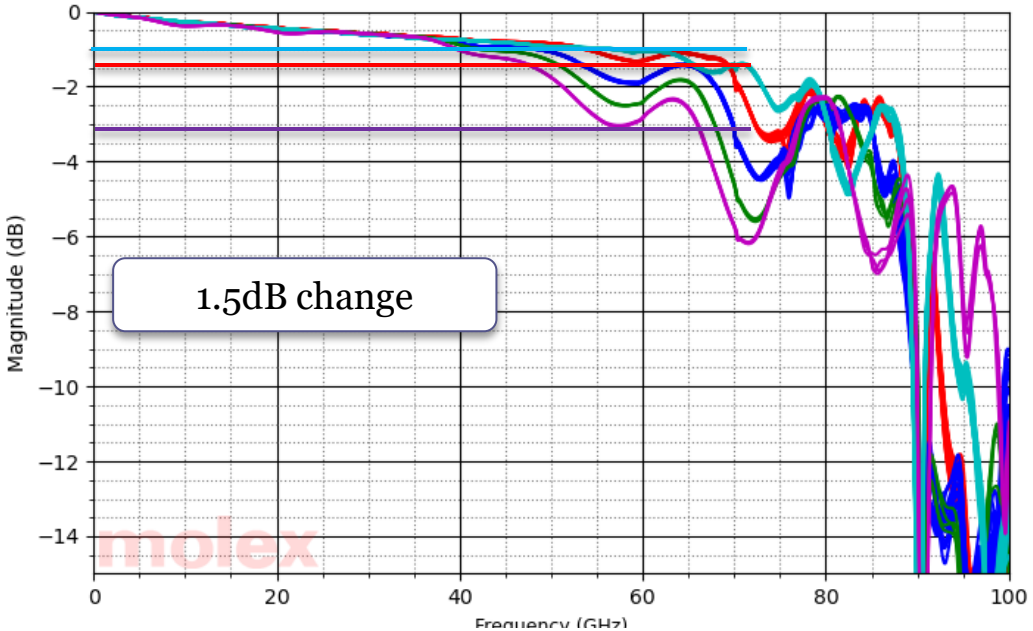

SDD11

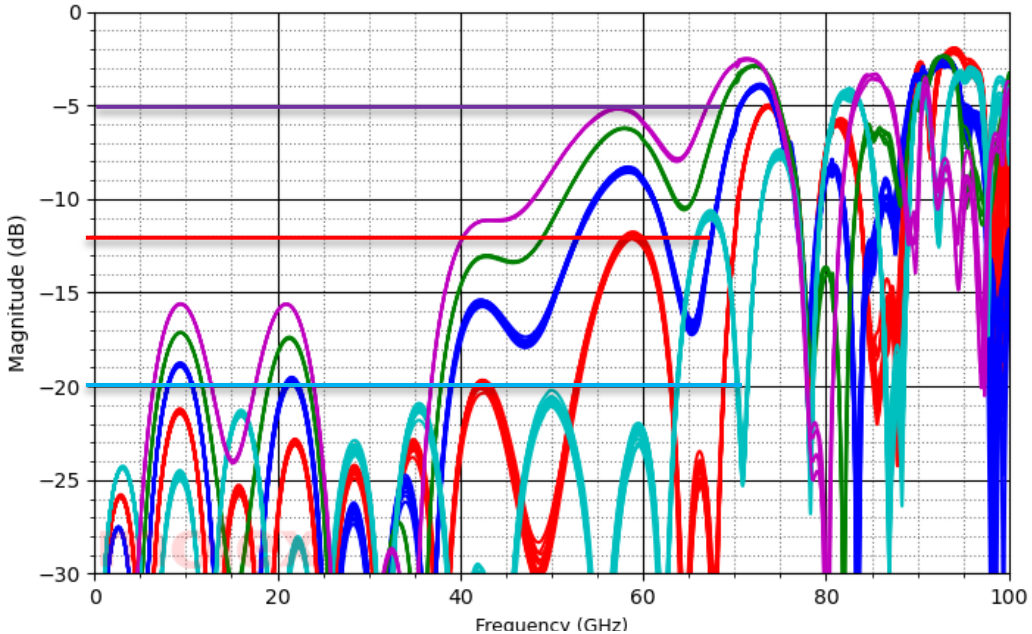

SCD21

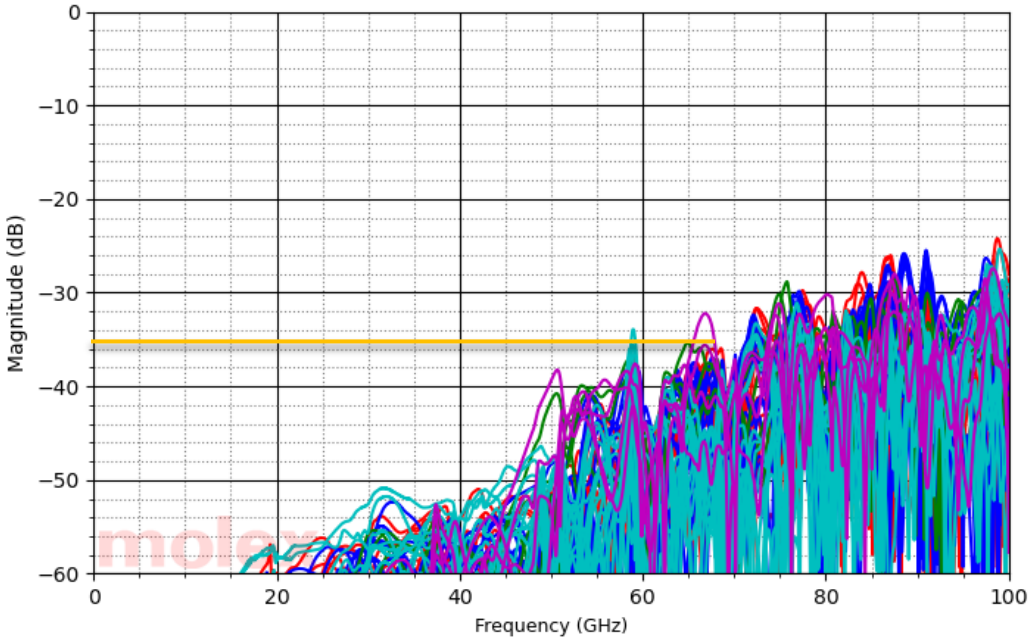

TDD11

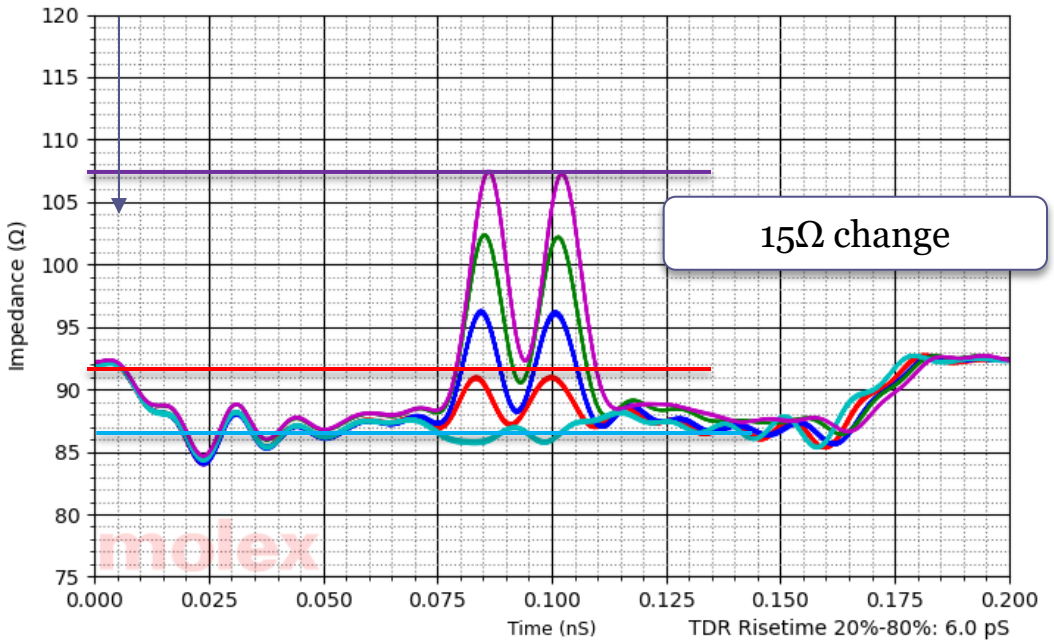

TEF 2025
Ethernet for

1.00mm
0.75mm
0.50mm
0.25mm
Flush

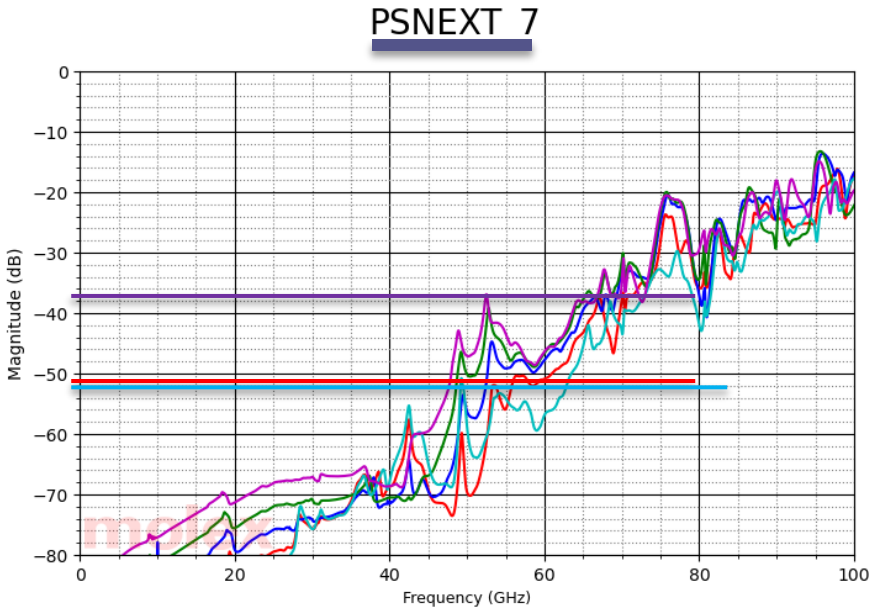

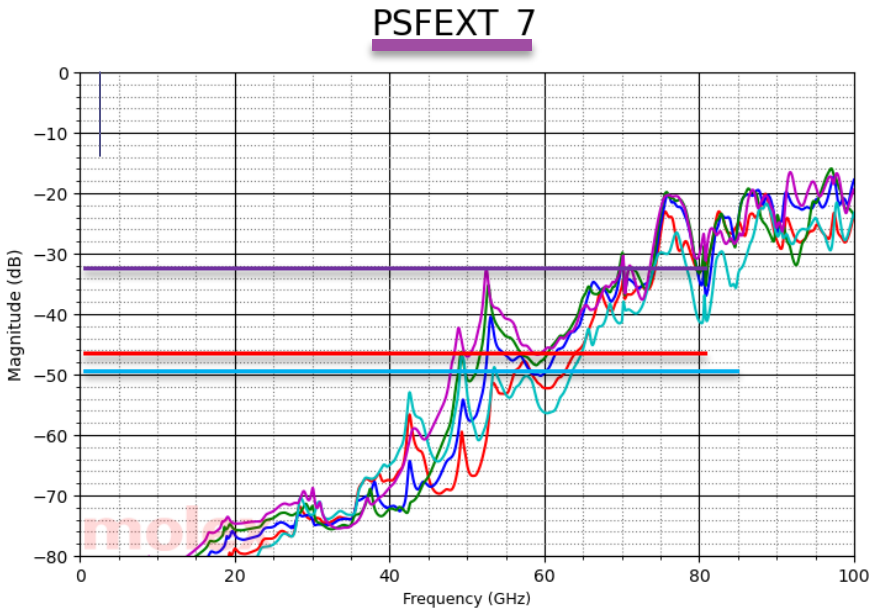

12dB to 14dB change

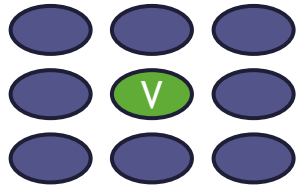

Victim7

TEF 2025
Ethernet for
AI

QUESTIONS?