Scaling Up Networking for AI: Unlocking the Full Potential of Optical Interconnects

OCP Educational Webinar Series

October 7, 2025

Dell'Oro Group

Lightmatter

Open Compute Project®

Community-driven hyperscale innovation for all

Scaling Up Networking for AI: Unlocking the Full Potential of Optical Interconnects

Cliff Grossner, Ph.D.
Chief Innovation Officer
OCP Foundation

Sameh Boujelbene
VP Market Research
Dell'Oro Group

Ritesh Jain
SVP Engineering & Operations
Lightmatter

Steve Klinger
VP Product
Lightmatter

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

AI Keeps Changing and Evolving at a Rapid Pace

- New age of AI reasoning driving new scaling laws: from pre-training scaling to post-training and test-time scaling
- Jensen Huang at GTC25: “Demand for compute is 100X what we predicted a year ago”
- AI clusters exploding in size

	Pre-2023	2023	2024	2026+
Cluster Size (XPUs)	1 K	25 K	100 K	1 M
	↓ 2X	↓ 3X	↓ 5X	↓ 10X
Interconnects	2K	75 K	500 K	10 M

DELL'ORO GROUP

Scale-up vs. Scale-out Network Design Options for AI Clusters

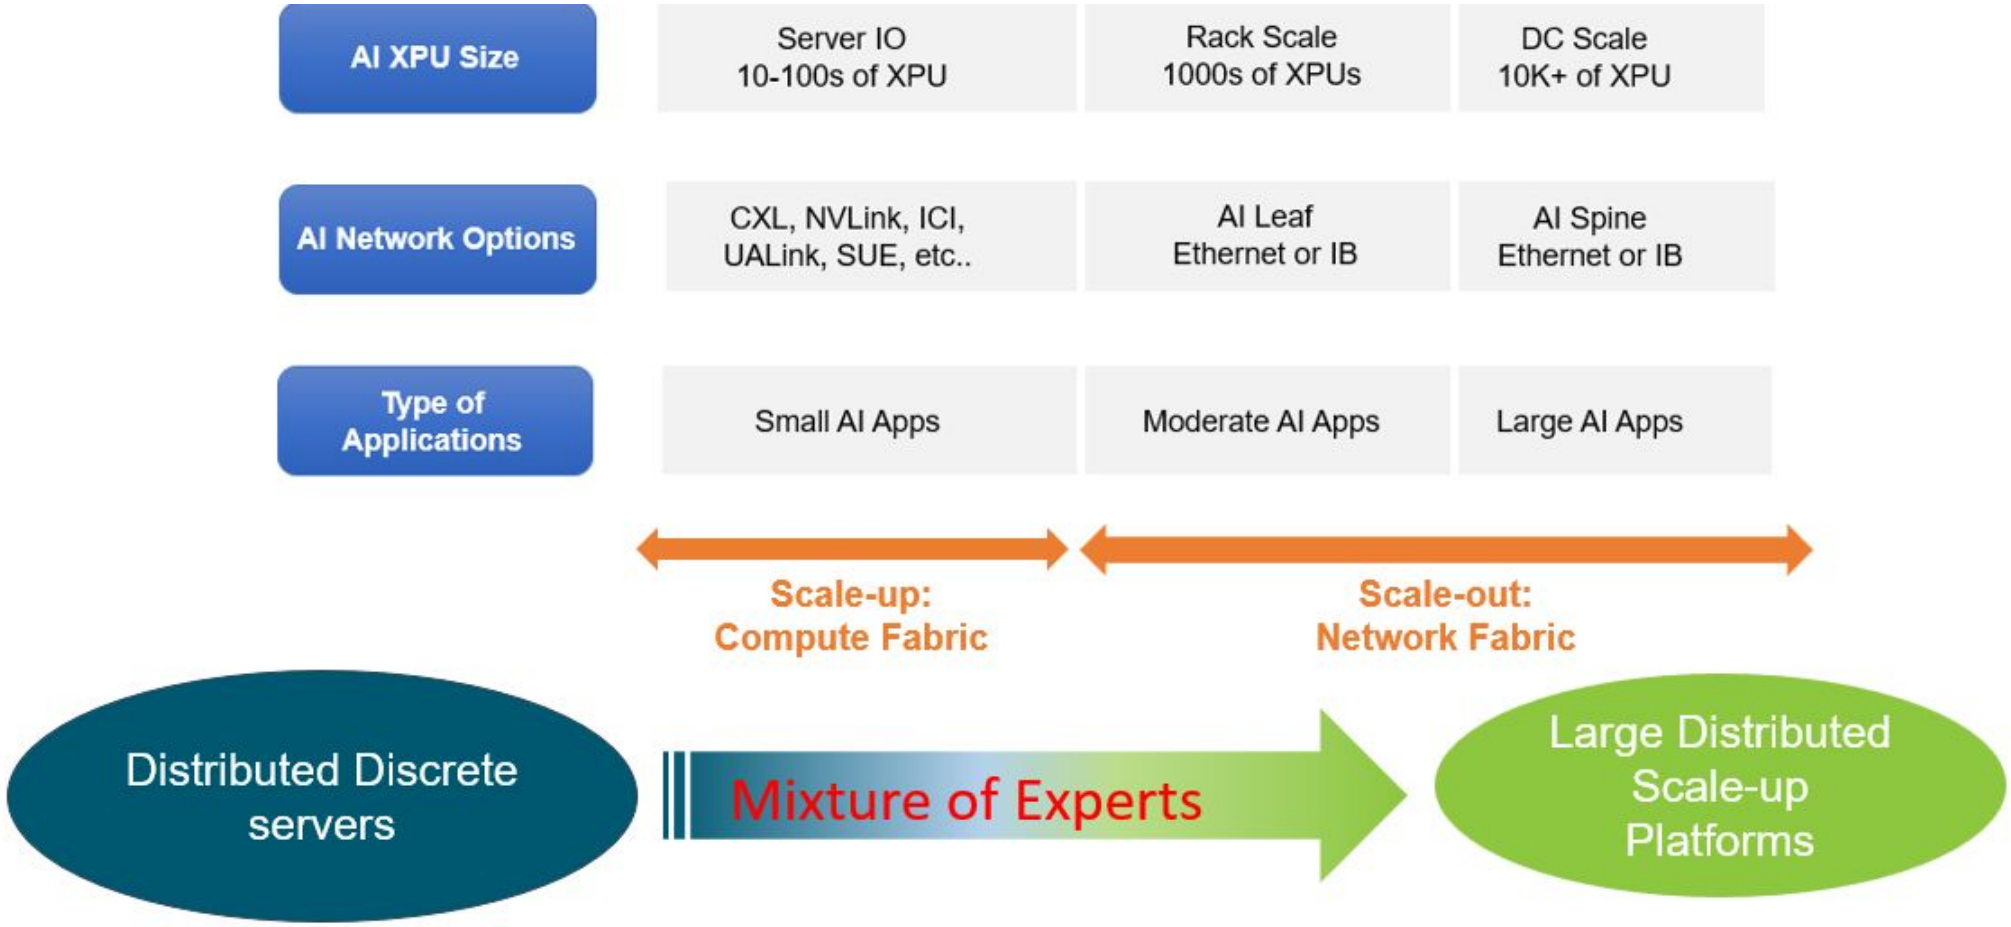

XPU could be GPU, TPU or any other type of accelerators

Scale-up Compute Fabric Advantages and Attributes

Advantages	Attributes
• Enables multiple XPUs to appear as one • Increase compute capacity • Increase main memory capacity	• Extremely high bandwidth • Extremely low latency • Deterministic performance • Lossless

Dell'Oro Group

Scale-up vs Scale-out Cumulative Switch Ports/Links (2024-2029)

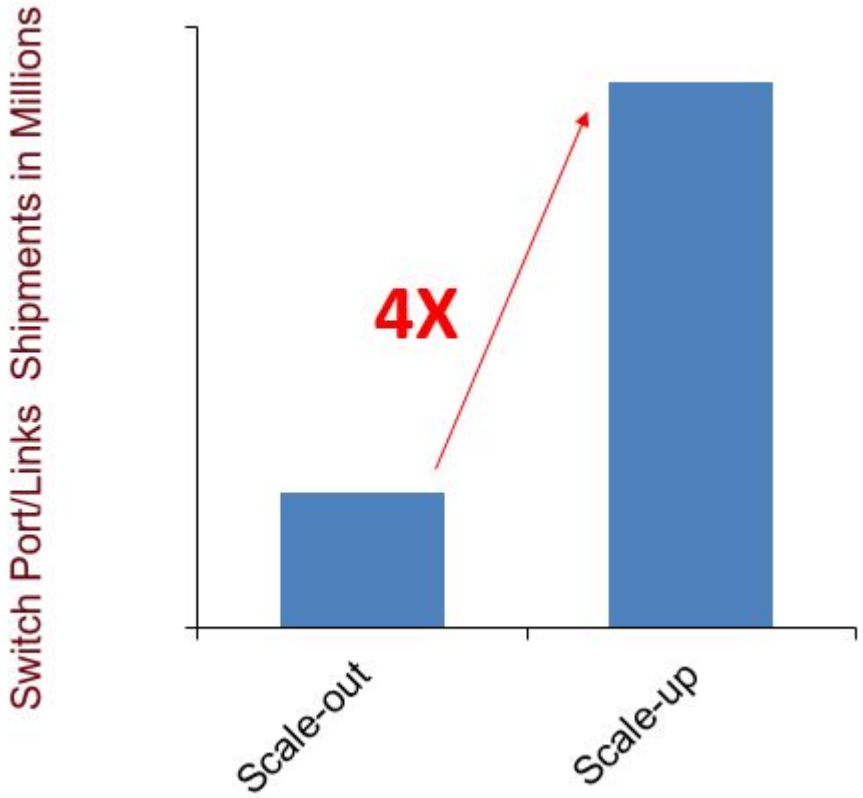

- Number of ports in Scale-up: 4X scale out
- Bandwidth per GPU in Scale-up: 10X scale-out

Source: Dell'Oro AI Back-end Networks Report

Scaling-up the Scale-up Domain

< 10 XPU	< 100 XPU	> 100 XPU
A Tray	A Rack	A Row
Inches Reach	1 to 3m Reach	10 m Reach
PCB traces	Copper	Fiber
Shorter Reach, Lower Power, Lower Cost	Limited Reach, Lower Power, Lower Cost	Longer Reach, Higher Power, Higher Cost

Source: Dell'Oro AI Back-end Networks Report

Scaling-up the Scale-up Domain

Problem Statement:
copper reach is one of the biggest physical challenges to scaling up the scale-up domain

Solution:
High-bandwidth, cost-efficient, low-power optics

Dell'Oro Group

CPO Switch Port Shipment Forecast in AI Back-end Networks

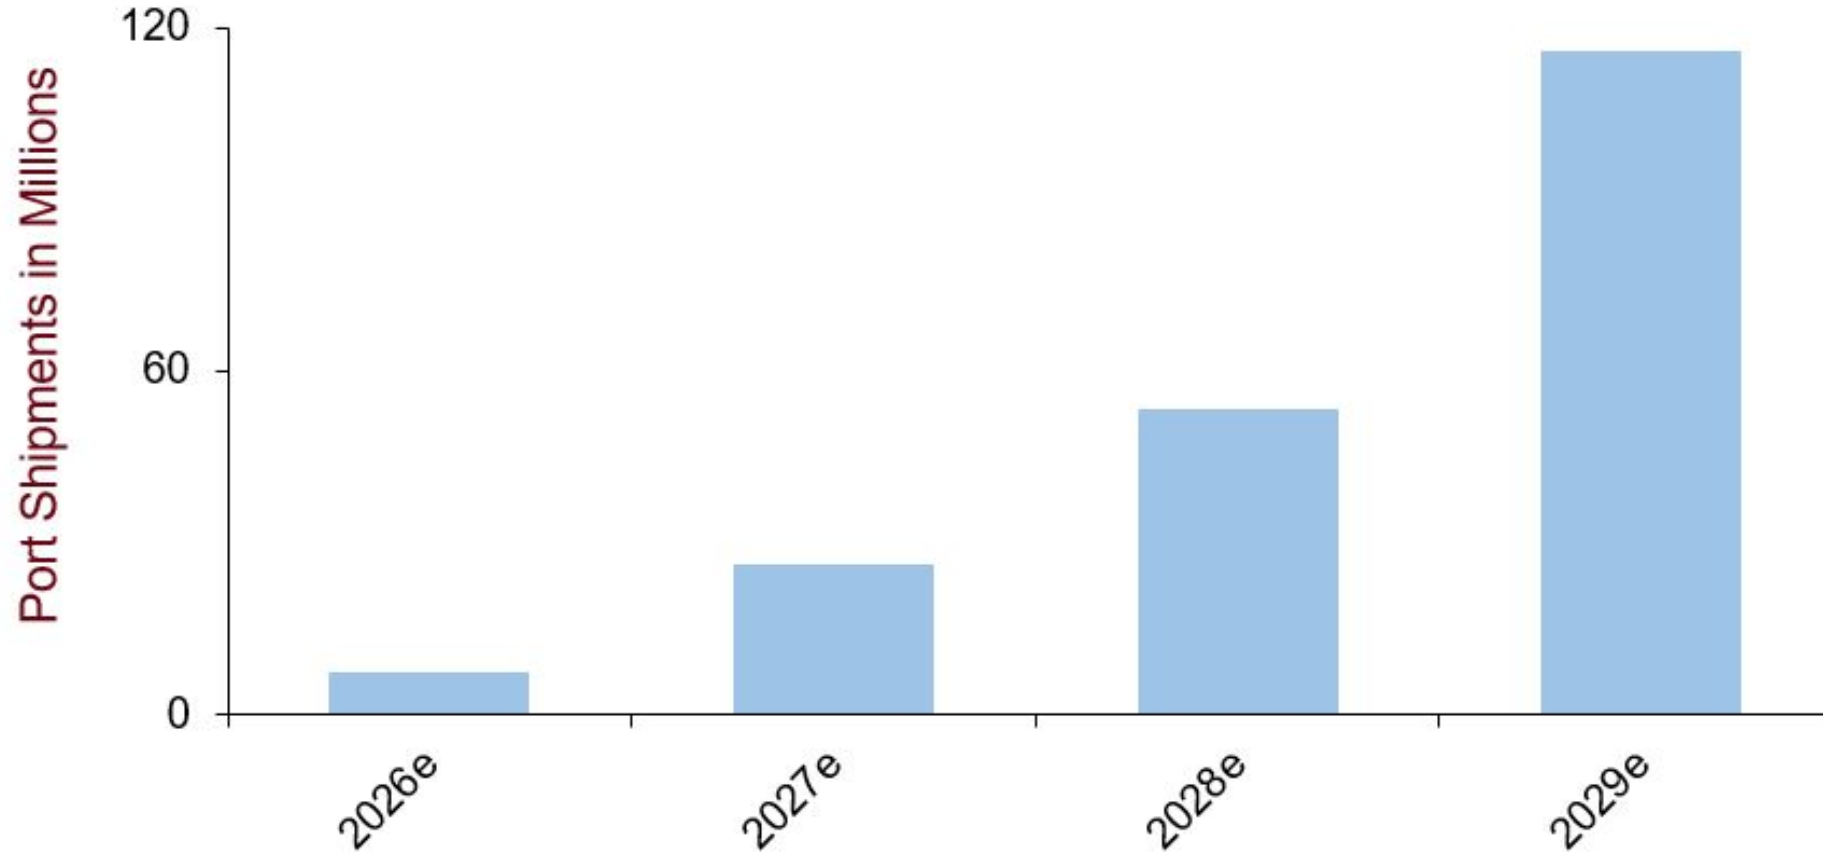

Source: Dell'Oro AI Back-end Networks Report
Includes Scale-out and Scale-up Switches

Dell'Oro Group

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

Deeper Challenge: Escape Bandwidth

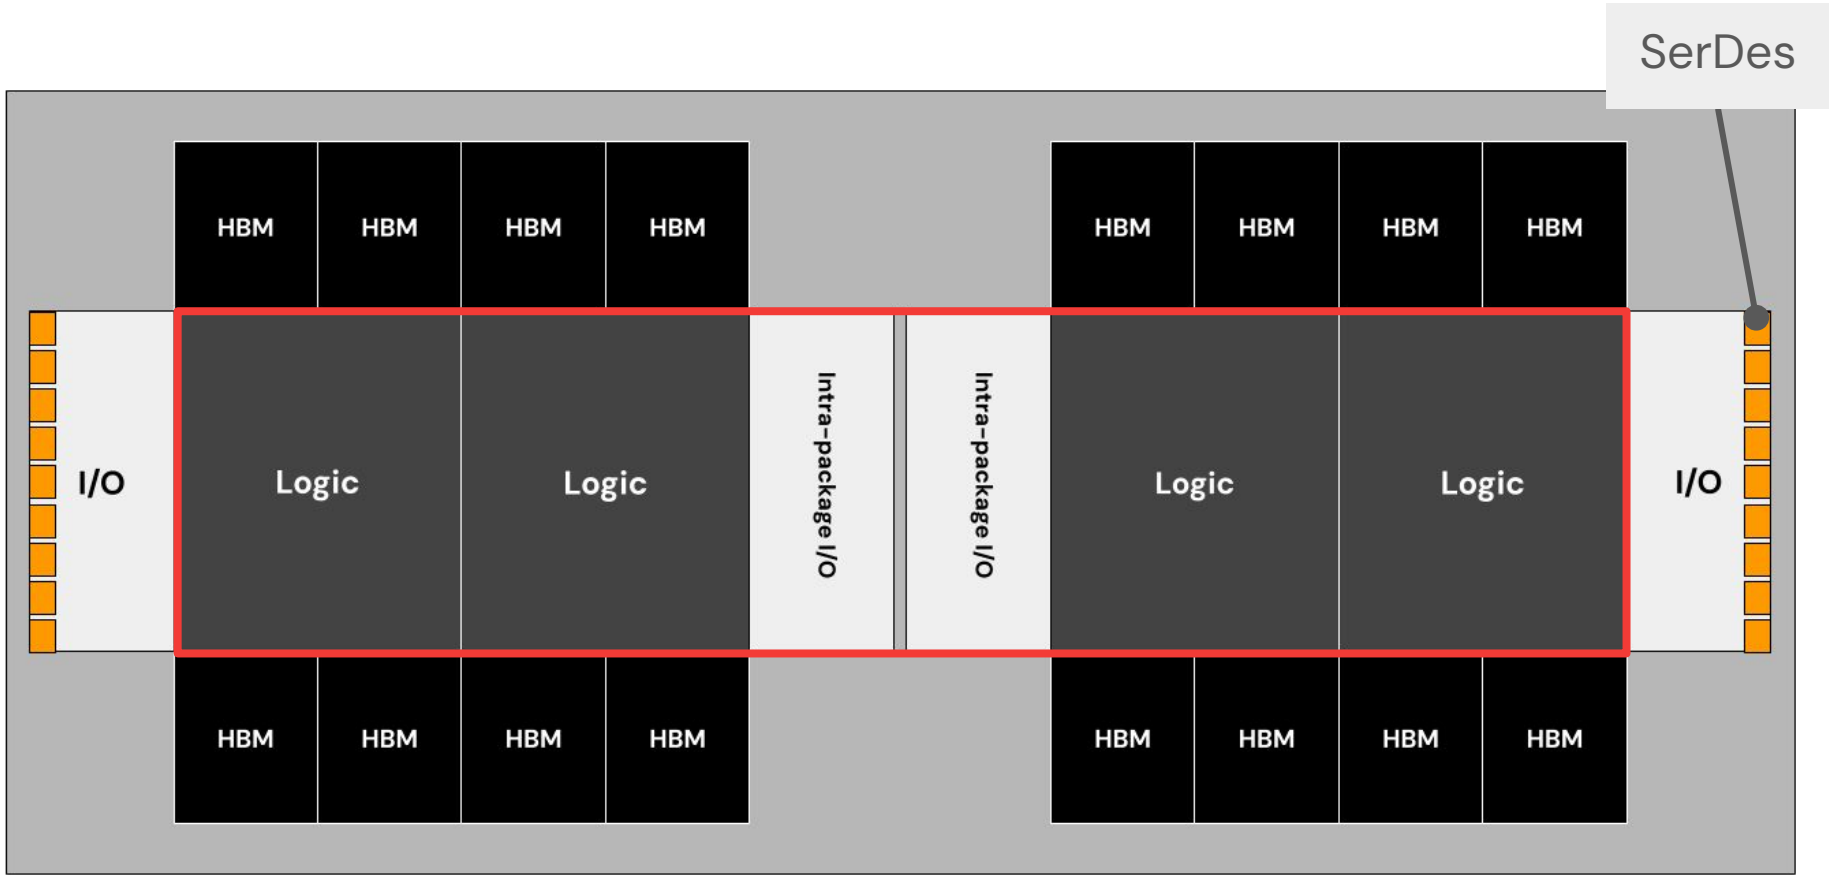

Communication happens on the chip perimeter

There is not enough shoreline, resulting in bandwidth constraints

LIGHTMATTER

Ritesh Jain

Evolving Interconnect

Bandwidth

100x —

10x —

1x —

3D Photonic Integration

- ASICs 3D stacked on PIC
- Advanced package, OCS
- ~2.5 pJ/bit

3D Co-Packaged Optics (CPO)

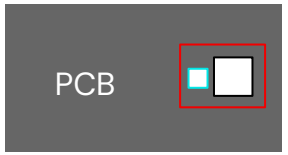

- 3D stacked PIC/EIC chiplet
- Standard pkg, UCIe
- ~3 pJ/bit

2D Co-Packaged Optics (CPO)

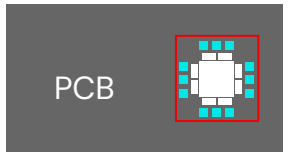

- I/O chiplet (PIC+EIC)
- In-package, XSR SerDes
- ~6-10 pJ/bit

Near Package Optics (NPO)

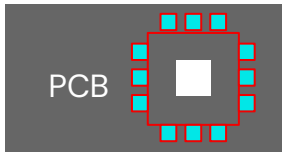

- Pre-packaged OE (PIC+EIC)
- LR SerDes, Re-timer
- ~8-15 pJ/bit

Power Efficiency

Highest

Lowest

Photonic Integrated Circuit: PIC
Electronic Integrated Circuit: EIC
Optical Engine (OE) = PIC+EIC

Copper, Pluggables

LIGHTMATTER

Ritesh Jain

Silicon Photonics Fundamentals

Integrate the Communication Technology onto the IC Package

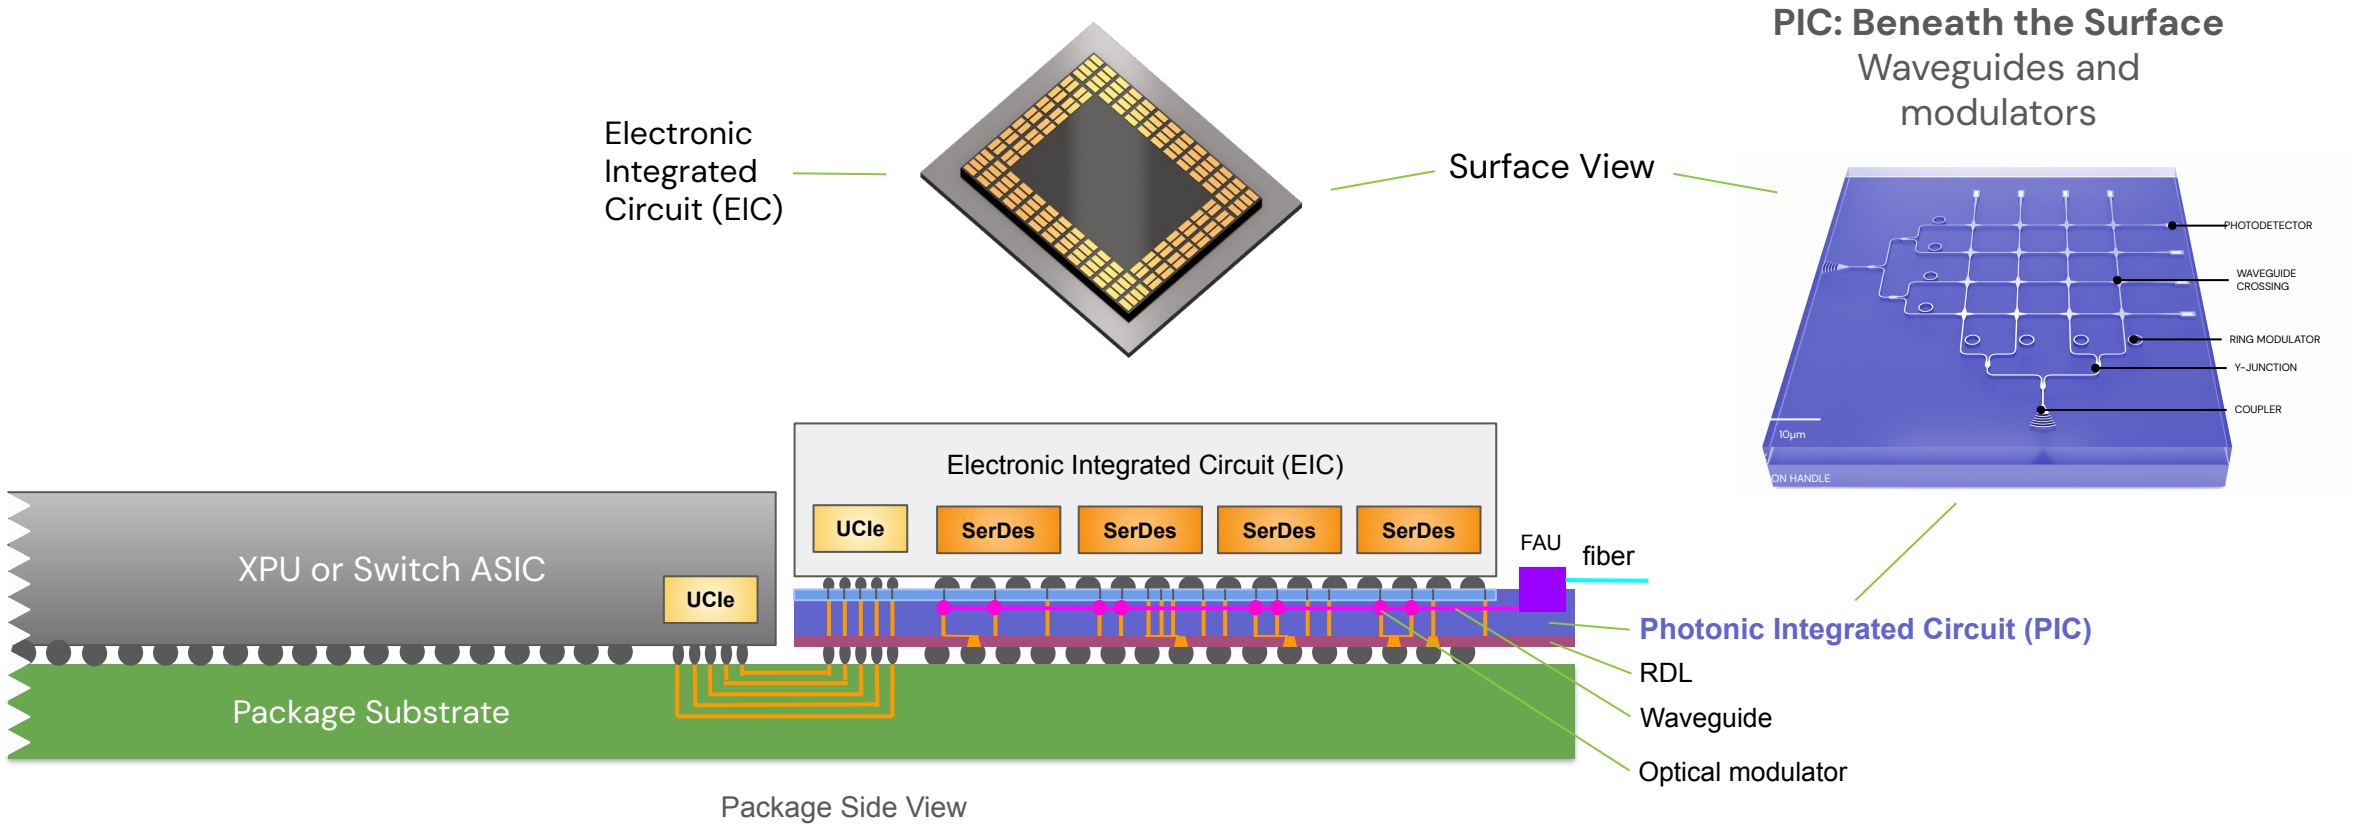

LIGHTMATTER
Ritesh Jain

Silicon Photonics Key Technologies

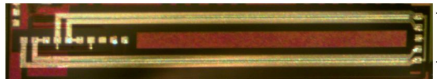

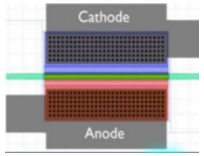

Metric	Mach Zehnder Modulator (MZM) [1]	SiGe/Ge Electro-Absorption Modulator (EAM) [2, 3]	Microring Modulator (MRM) [4]
Size	>1000µm	50–100µm	~15µm in diameter
Requires a multiplexer?	Yes	Yes	No, already a multiplexer
Thermal stability	Stable, thermal gradient must match between both arms.	Stable with feedback loop. Modulation is temperature sensitive.	Stable with feedback loop. Transmission depends on ring resonance.
Power consumption [5]	~50 mW	~10 mW	~1 mW
Transmitter loss	<5 dB	<10 dB	<5 dB
Optical bands	O-band & C-band	C-band only	O-band & C-band

[1] M. Streshinsky, et al. Opt. Express 21, 30350-30357 (2013).
[2] J. Fujikata, et al., Opt. Express 31, 10732-10743 (2023).
[3] Y. Liu, et al., PRJ 8, 1474 (2023).
[4] E. Timurdogan, et al. Nature Communications 5, 4008 (2014).
[5] DAB Miller, et al. Opt. Express 20, A293-A308 (2012).

LIGHTMATTER
Ritesh Jain

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

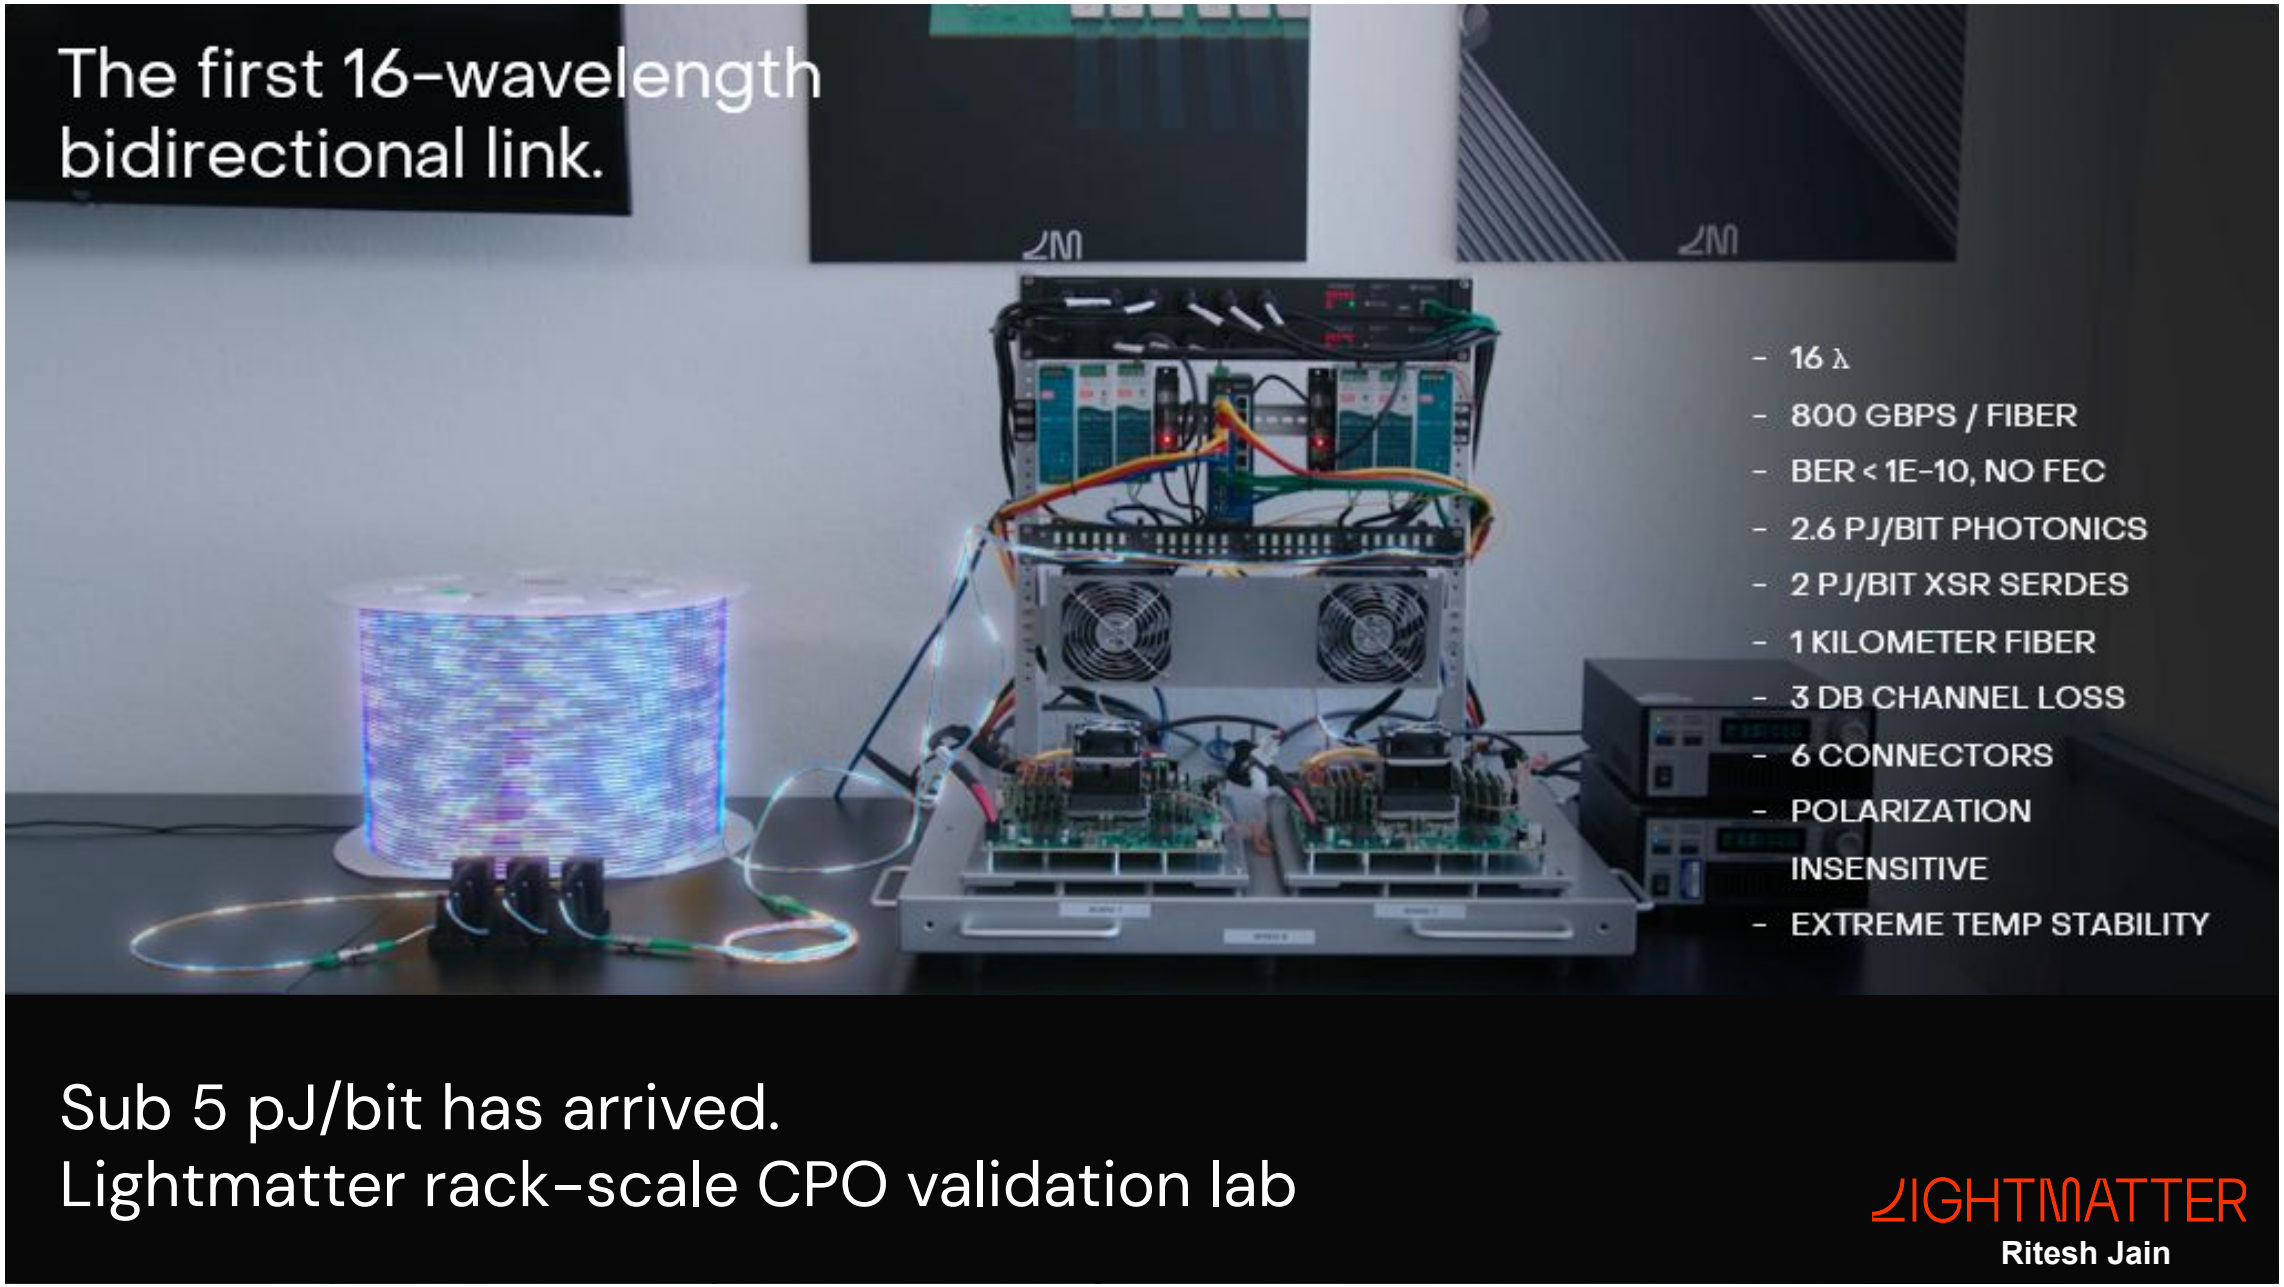

https://youtu.be/Gjee92kYmwg

Operating Temperature / Thermal Cycling

Long Term BER over Temp Cycling

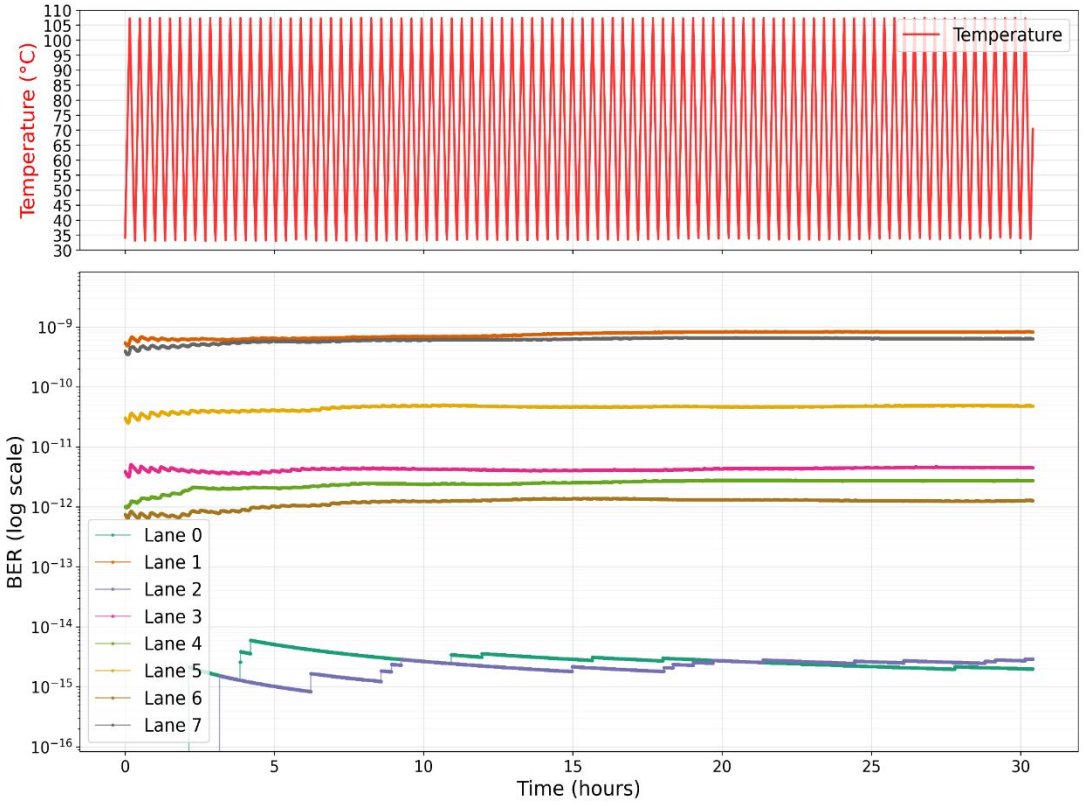

Measuring robustness against 25-105C thermal aggressors

BER over Rapid (800C/s) Temp Cycling

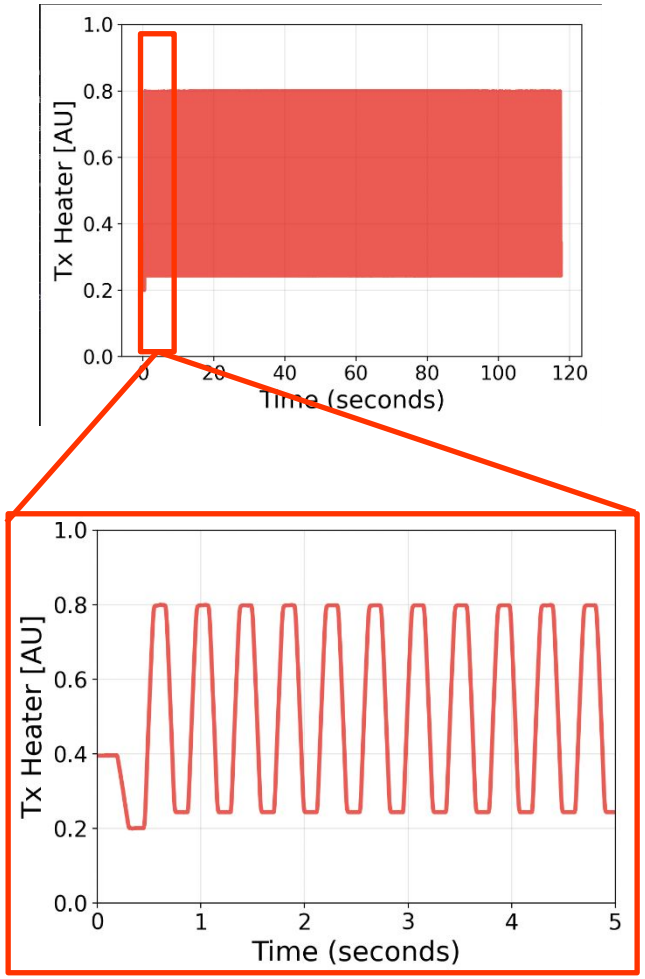

- A swept tunable laser is used to emulate a fast temperature aggressor
- A 50 nm/s ramp rate is used to emulate a 800℃/s temperature aggressor

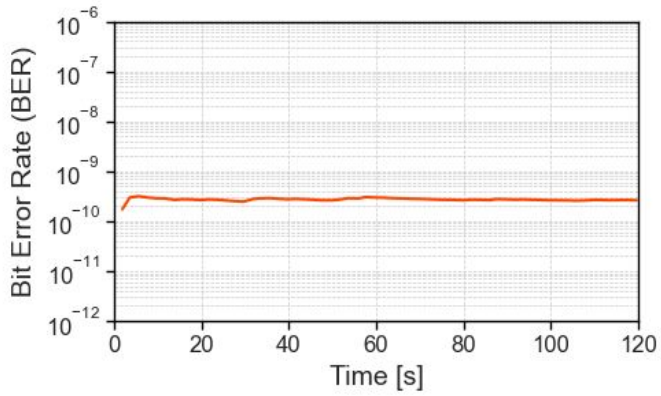

LIGHTMATTER

Ritesh Jain

Robust HVM: 3D CPO Assembly & Test Flow

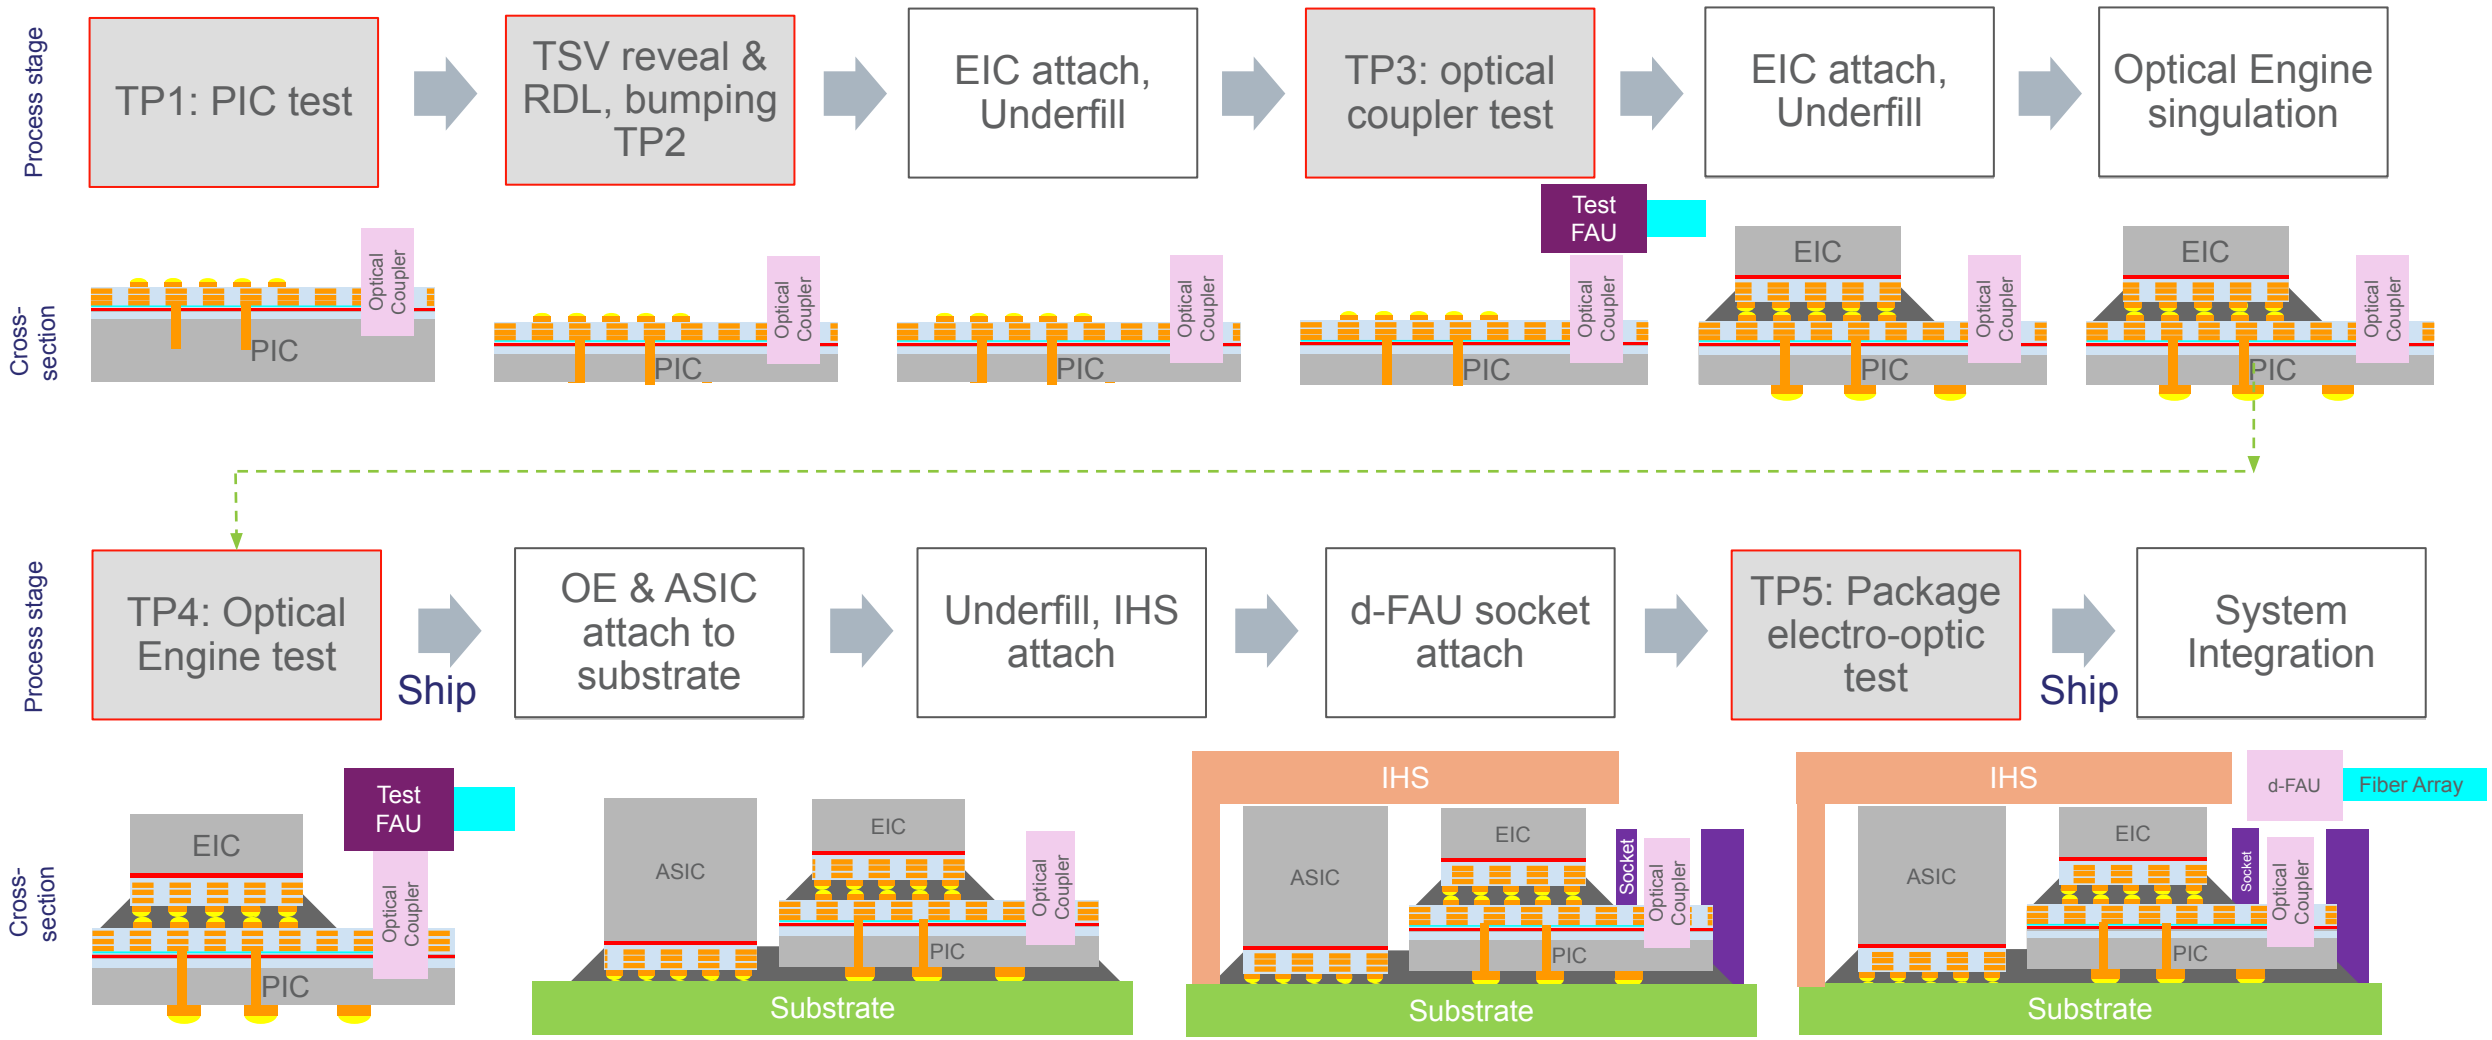

LIGHTMATTER

Ritesh Jain

Audience Poll Question

Which statement best reflects your organization's position on CPO adoption in the next 2-3 years?

1. We plan to deploy CPO, starting with GPUs
2. We plan to deploy CPO, starting with switches before GPUs.
4. We plan to deploy CPO, on both GPUs and switches.
5. We do not plan to deploy CPO within the next 4–5 years.

Audience Poll Results

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

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

Example 3D Photonic Implementations

3D CPO

XPU application

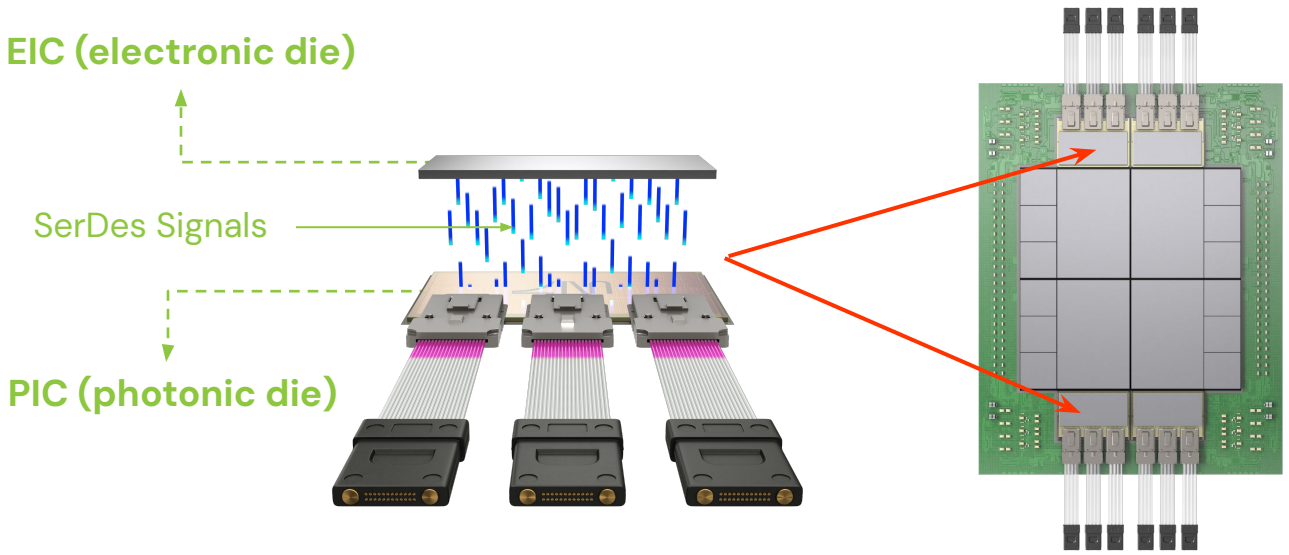

Passage™ L-Series

3D Photonic Interposer

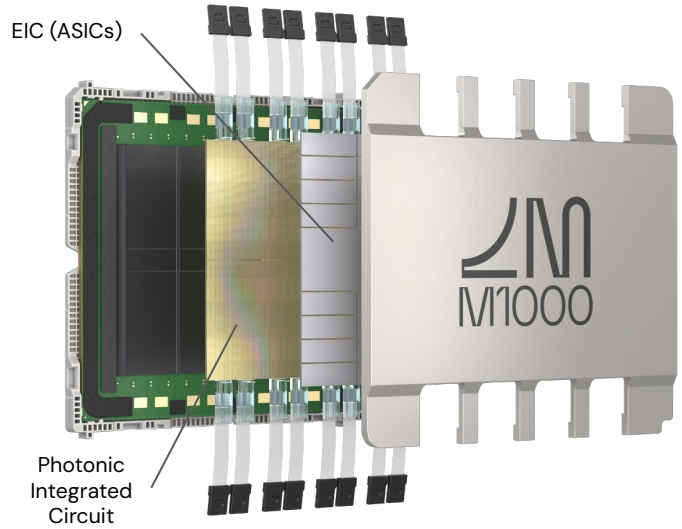

Passage™ M-Series

LIGHTMATTER

Steve Klinger

Support for Multiple Link Format Options

Bandwidth	Modulation	Number of Wavelengths per Fiber	Transmission Type	WDM Type CWDM/DWDM
56 Gbps	NRZ	16	Bi-directional	DWDM
56 Gbps	NRZ/PAM4	16	Uni-directional	DWDM
112 Gbps	PAM4	16	Uni-directional	DWDM
224 Gbps	PAM4	4	Bi-directional	CWDM
224 Gbps	PAM4	4	Uni-directional	CWDM

Enabling unidirectional bandwidth per XPU of 32T and above, up to 1 km

LIGHTMATTER

Steve Klinger

Scale Up Requires an Optimized Solution

Requirement: deliver increasing scale-up bandwidth with high radix to a growing number of XPUs, over longer reaches, in nanoseconds

Scale Up

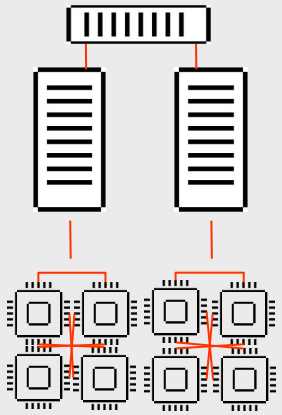

→1024 GPUs
(needs to extend beyond rack)

Network Type	XPUs	Latency	Bandwidth Per XPU	Energy
Scale-up	→ 1024	100-250 **ns**	> 12.8 Tbps	< 5 pJ/bit
Scale-out	> 100k	multi-hop 2-10 us	→ 1.6 Tbps	16 pJ/bit

Scale Out

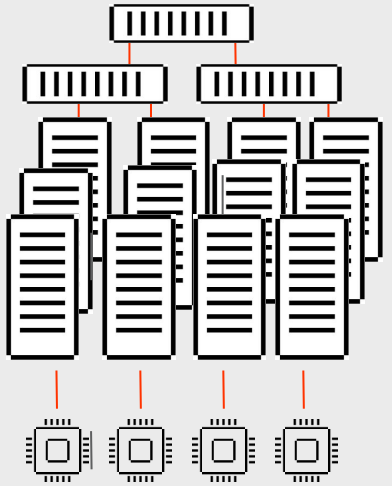

100k+ GPUs over 100's of racks

Scale-up pods to 1,024 XPUs and across multiple racks

LIGHTMATTER

Steve Klinger

Impact of Pod Size + Bandwidth +Radix

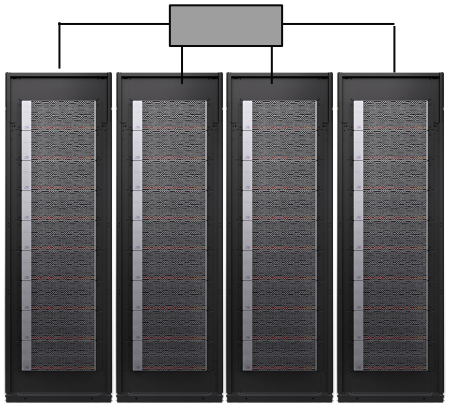

32T 3D CPO SCALE-UP POD
576-Pod Passage 3D CPO
512 active GPUs / Pod (2,048 dies)
Optical 32T per GPU
Scale-out 1.6T per GPU

Versus

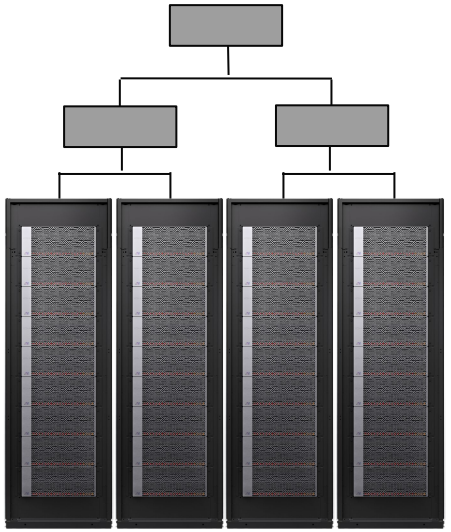

14.4T ELECTRICAL SCALE-UP+1.6T SCALE-OUT
4 x 144-Pods
512 active GPUs/ 4 Pods (2,048 dies)
Electrical 14.4T per GPU
Scale-out 1.6T per GPU

3D CPO 3–7X Training Speedup

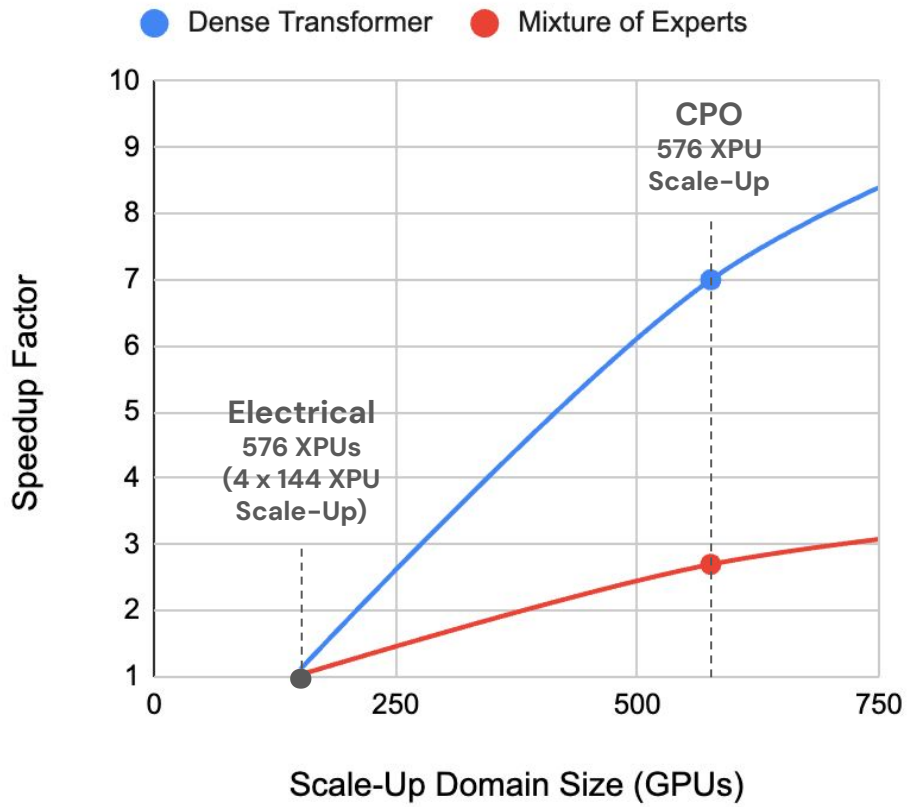

LIGHTMATTER

Steve Klinger

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

Lightmatter 3D CPO Product and Ecosystem

Customer ASIC with UCIe IP	**Multiple fabs**
Passage™ EIC (partners)	tsmc
LIGHTMATTER Passage™ PIC	GF tsmc
Detachable Fiber Coupler	**Multiple Suppliers**

Guide™ Light Engine

ASE

Amkor Technology®

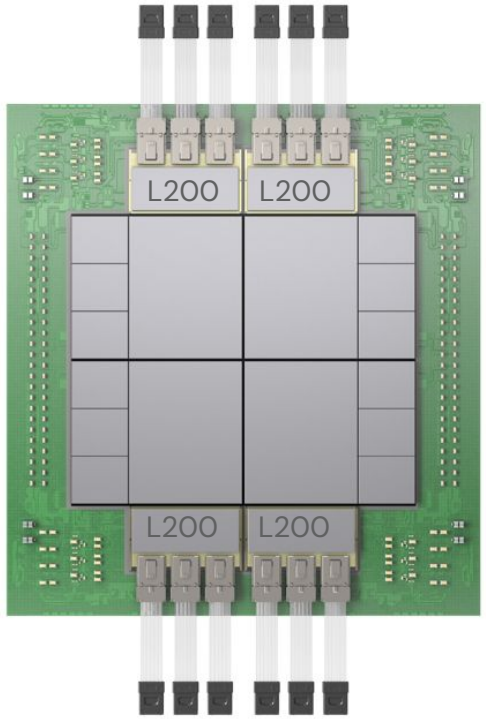

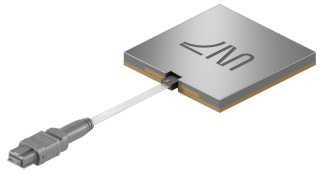

- Global ecosystem
- High-volume readiness
- Quality and reliability

LIGHTMATTER

Steve Klinger

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

Take Aways

- AI keeps changing and evolving at a very rapid pace
- The rise of Mixture of Experts (MoE) is driving the need for much larger single scale-up GPU domains
- The scale-up domain is the unit of high-performance compute
- Copper reach is one of the biggest physical challenges to scaling-up the scale-up domain
- Breakthrough innovations are needed to achieve high-bandwidth, cost-efficient, low-power optical interconnect

DELL'ORO GROUP

Agenda

- CPO for Scale-up, When and How Much?
- Moving to Optical Interconnects & CPO
- Practical Realities Delivering Optical Interposers
- AI Cluster Interconnect Challenges
- Sponsor Product & Ecosystem
- Conclusion
- Live Q&A

Live Q&A

Cliff Grossner, Ph.D.
Chief Innovation Officer
OCP Foundation

Sameh Boujelbene
VP Market Research
Dell'Oro Group

Ritesh Jain
SVP Engineering & Operations
Lightmatter

Steve Klinger
VP Product
Lightmatter

Thank you for attending

Scaling Up Networking for AI: Unlocking the Full Potential of Optical Interconnects

OCP Educational Webinar Series

October 7, 2025

OPEN Compute Project®

Dell'Oro Group

LIGHTMATTER

Community-driven hyperscale innovation for all