AI OPEN
DATA CENTER

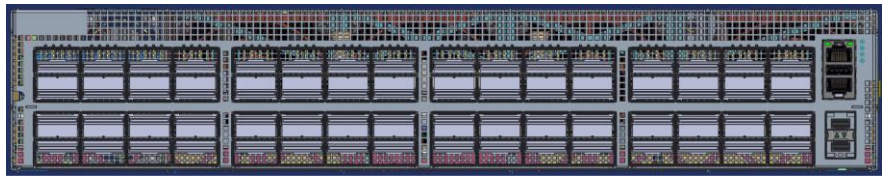
Co-Designing High-Speed Copper and Optical Interconnects: A Passive Interconnect Approach for Next-Gen AI Data Centers

Brian Kirk

Chief Technology Officer (CTO), Amphenol

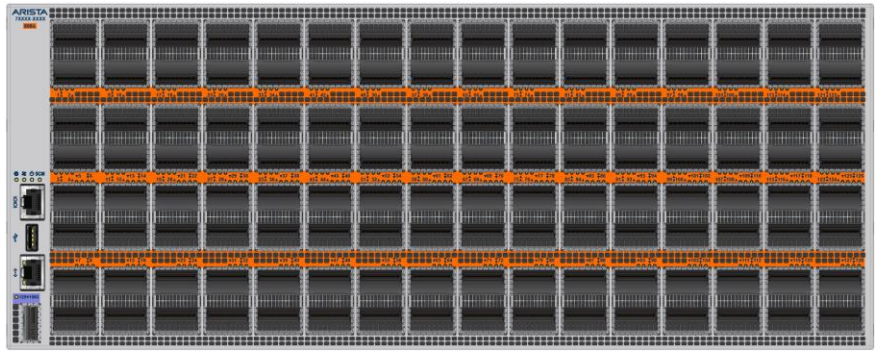
OCP SILVER Member

Beyond 102.4 Tb/s Systems



102.4 Tb/s

64 × 1.6T OSFP



204.8 Tb/s

128 × 1.6T OSFP



204.8 Tb/s

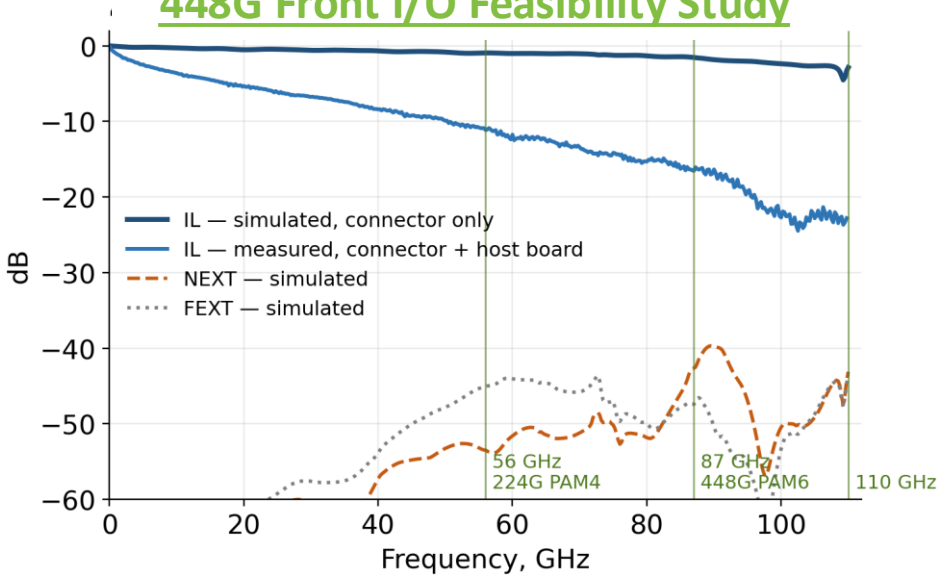
16 × 12.8T XPO

IO Density at the faceplate now gate scaling more than the switch ASIC itself

Beyond 224 Gb/s per Lane Channels







Leave the PCB as early as possible, new form factor(s) at the Port

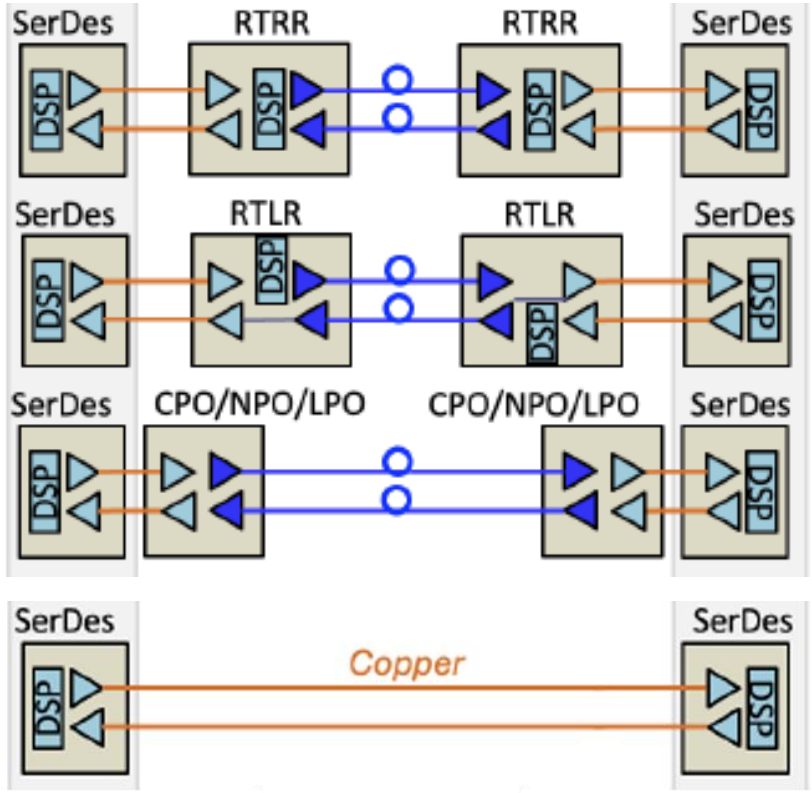
Maximizing the Port Capabilities

Retimed Tx
Retimed Rx
(fully retimed)

Retimed Tx
Non-retimed Rx
(RTLR, LRO, TRO)

Non-Retimed Tx
Non-retimed Rx
(LINEAR, LPO)

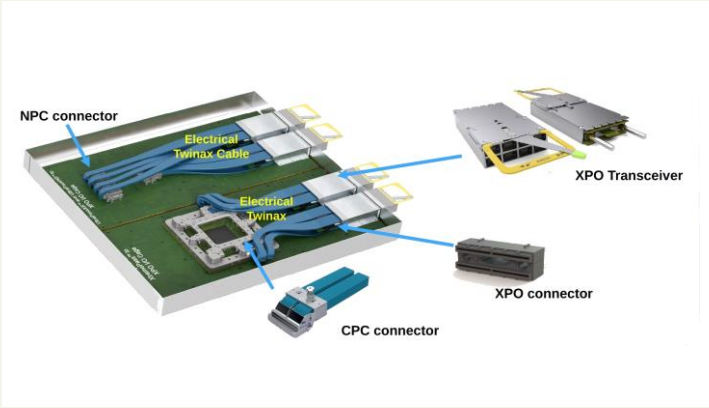
Passive Cu Cable



Electrical Channel	Key Capability
Compliant	Interoperability, plug-and-play
Better	Overall compromise on channel/power/latency
Best	Lowest Optical Power, Latency
Best	Lowest Cost, Power, Latency

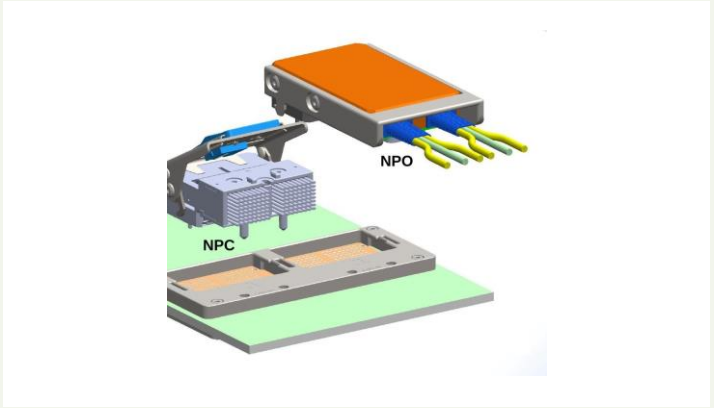
The Linear (Copper or Optical) Ports are defined by the Electrical Channel performance

Modularity starts at the Chip – NPC, CPC, NPO, CPO



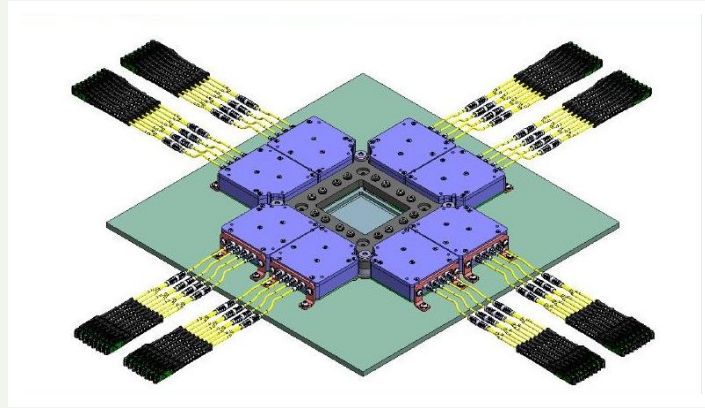
NPC/CPC → XPO

Cabled-Host,
Front-panel Pluggable



NPC → NPO

Near-Package module,
One socket electrical or optical

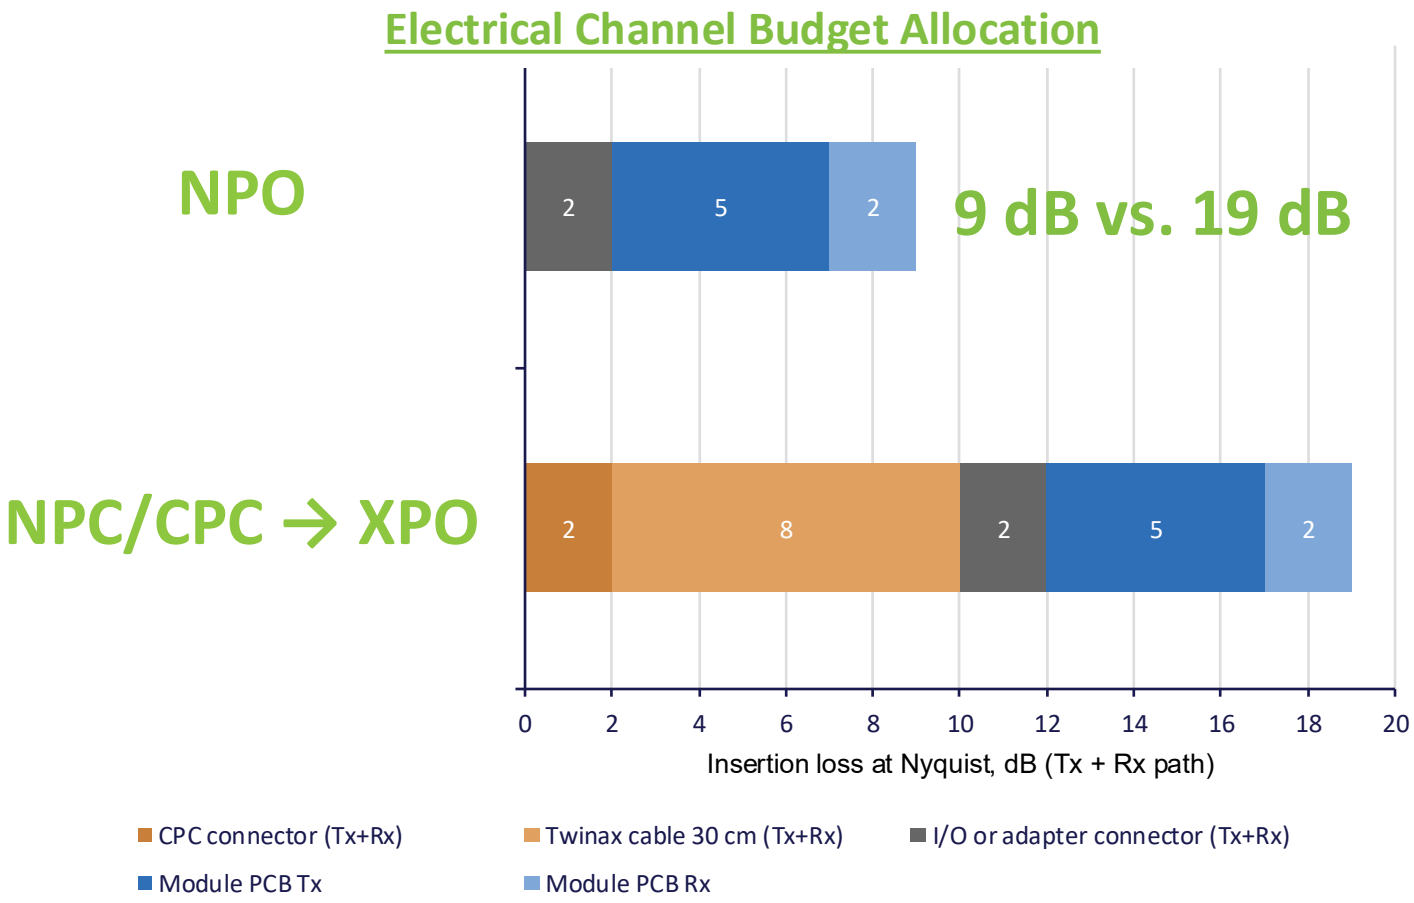


CPO

Co-Packaged optics

(1) Host SerDes, (3) Host Architecture Endpoints

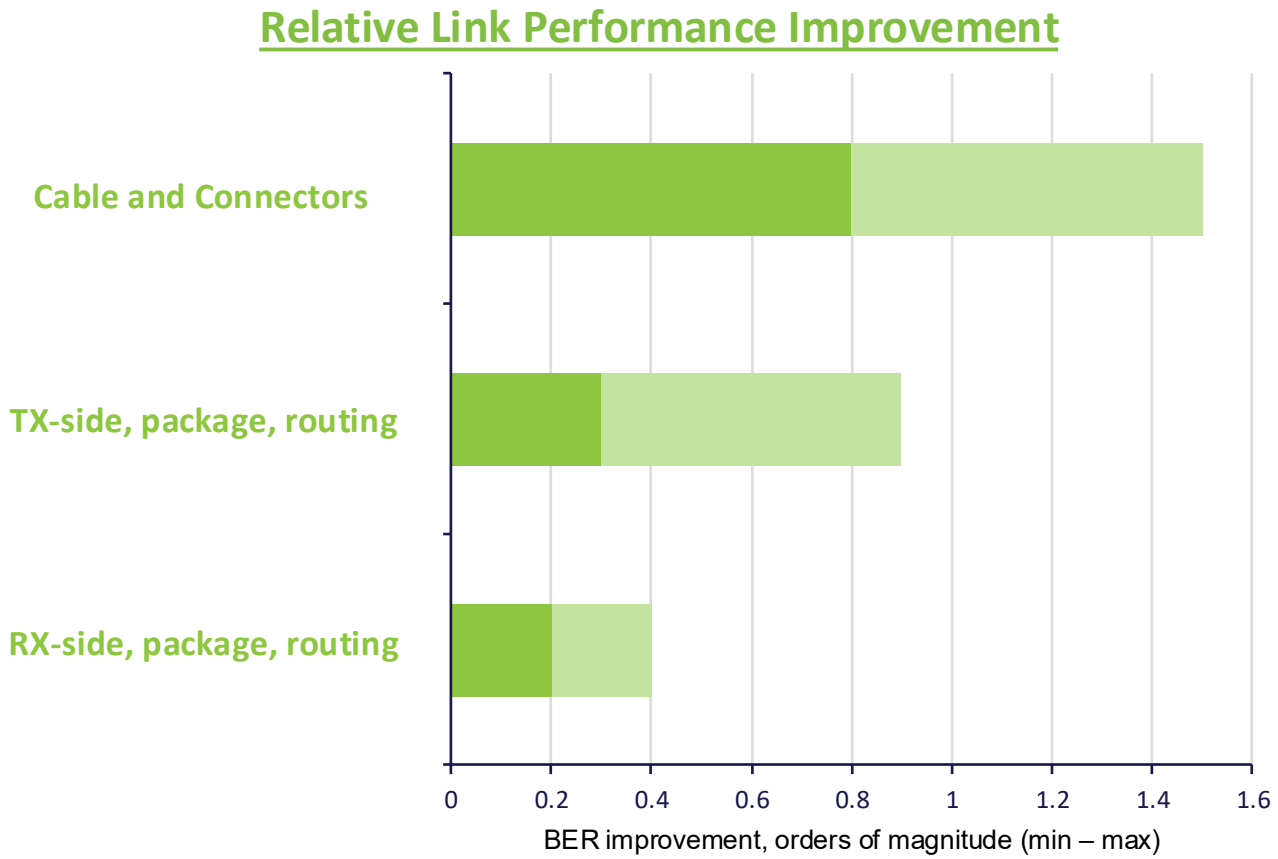
Modularity Trade-offs



Key Attribute	Modularity Preference
Reliability	Pluggable
Serviceability	Pluggable
Electrical Channel	NPO
Bandwidth Density	NPO
Availability	-

RAS favors Pluggables, but the gap is shrinking

Return on Investment – Optimizing the Link



Current Development Example:

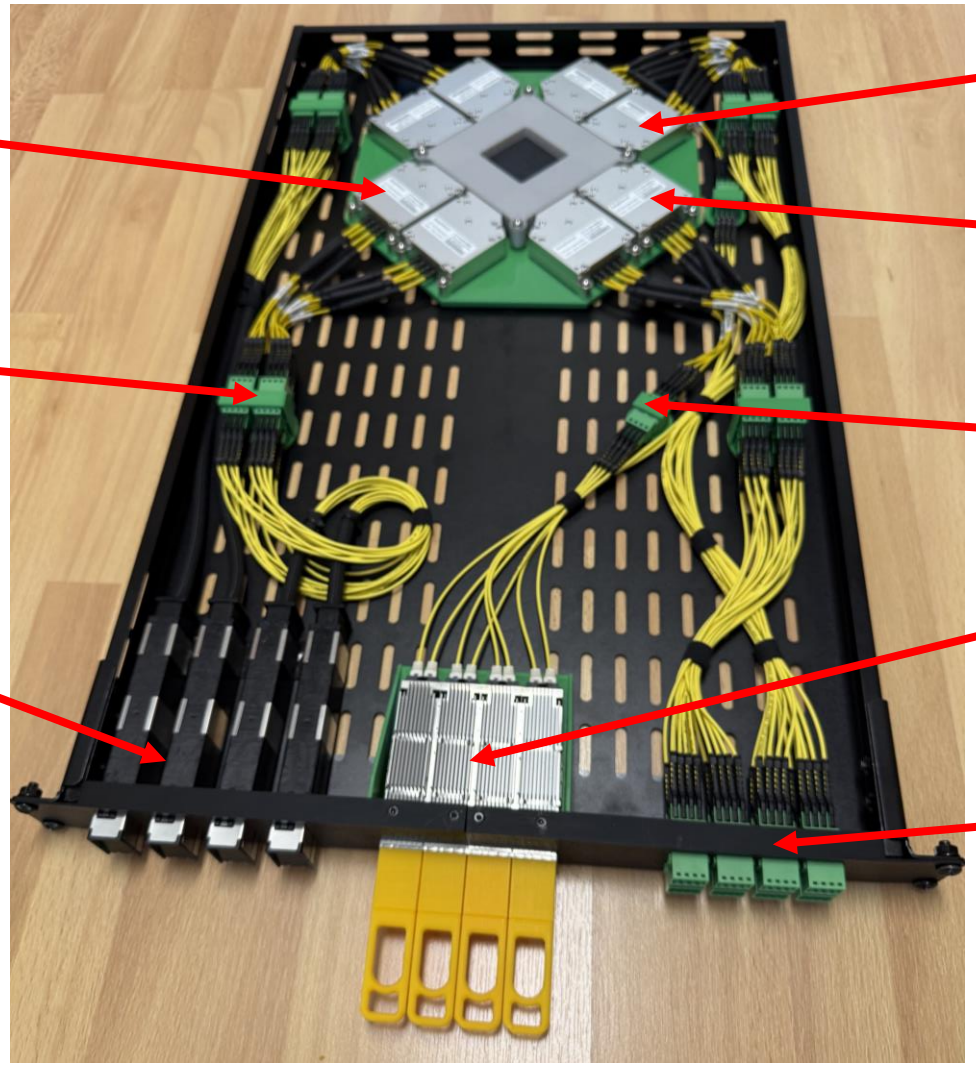
- 64-lane XPO-LPO
- 3nm switch Host
- 15dB Host + Cable IL
- All lanes active

(1) lane alone: 1e-10 to 1e-12.

All (64) lanes on: ~1e-7 The gap is crosstalk

The Link Performance is defined by the Electrical Channel

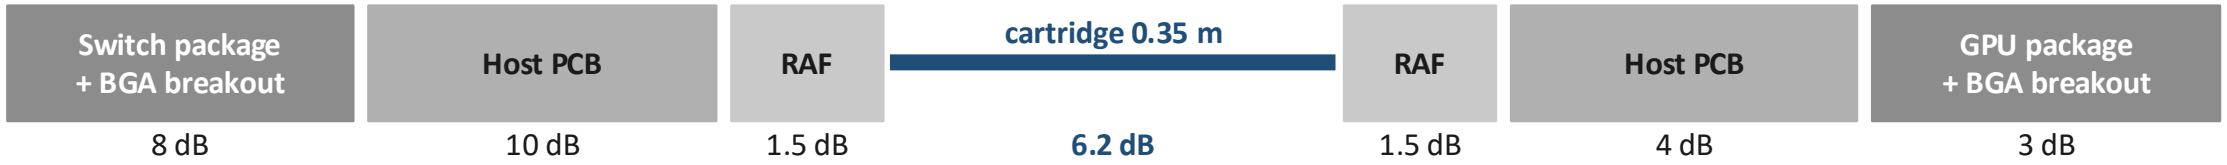
Considerations towards CPO Deployment



The End of Copper is not the Cartridge

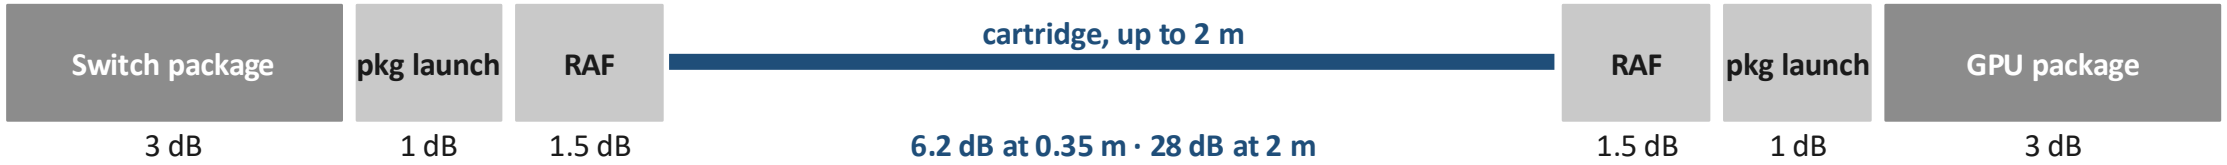
Today: ASIC → PCB → RAF → cartridge → RAF → PCB → ASIC

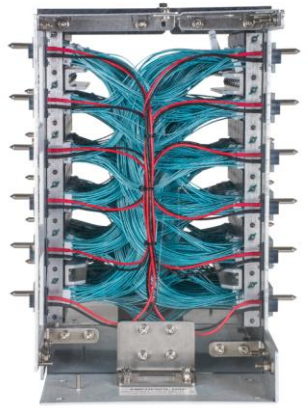
53 GHz Nyquist, 224G PAM4, twinax 14 dB/m



Launch at the package: ASIC → cartridge → ASIC

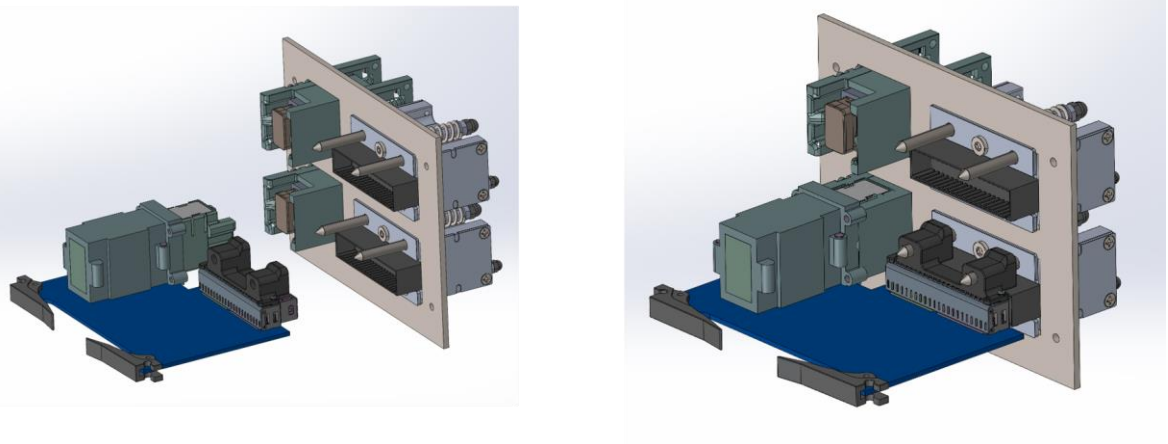
same budget, same cartridge, no host PCB



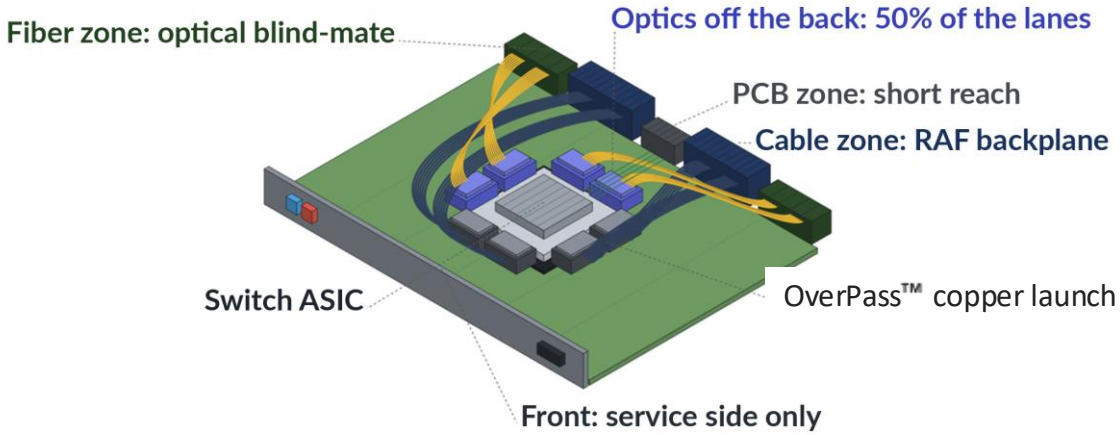


The cartridge is the smallest allocation in the budget

Copper and Optics can Co-Exist



RAF twin-ax and a floating optical blind-mate on one edge



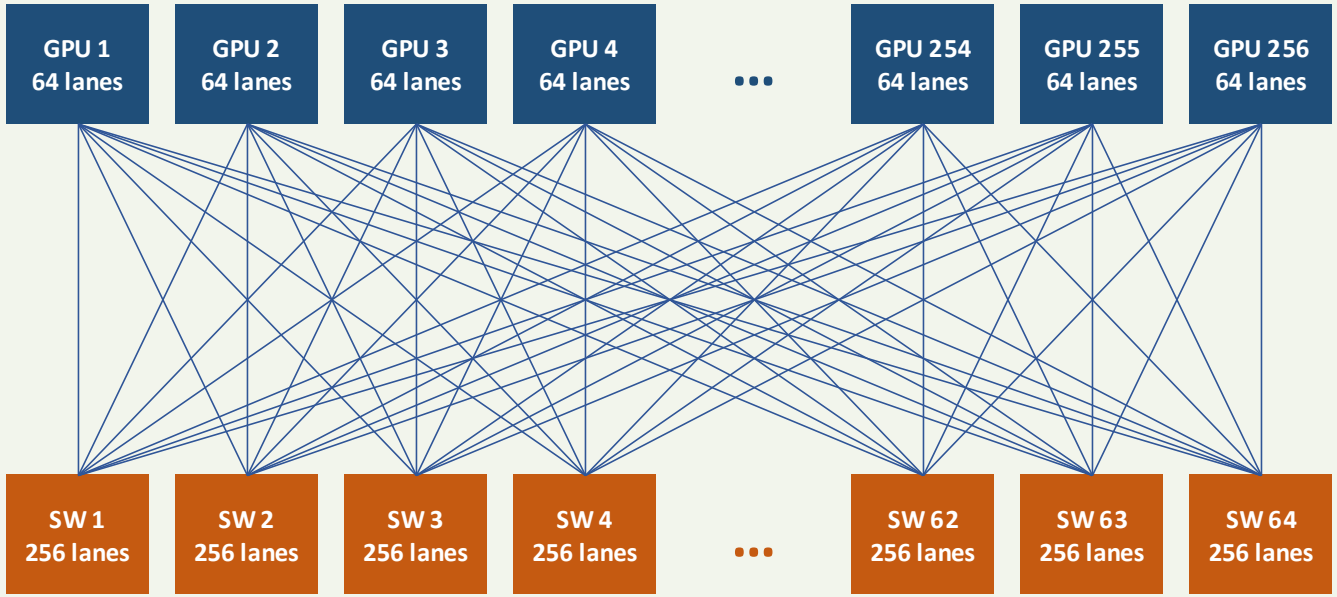
Conceptual Implementation

Blind-Mating Copper and Optics

Modularity at the Chip, is Modularity at the System

Network Architecture Example

256 GPUs, 64 switch ASICs, one hop



Every line = 1 lane (TX/RX). Each GPU spends its 64 lanes as 1 lane to each of the 64 switches. Each switch spends its 256 lanes as 1 lane to each of the 256 GPUs.

16,384

lanes in the fabric

256 GPUs × 64 lanes

64

switch ASICs

16,384 ÷ 256 lanes per switch

1 hop

between any two GPUs

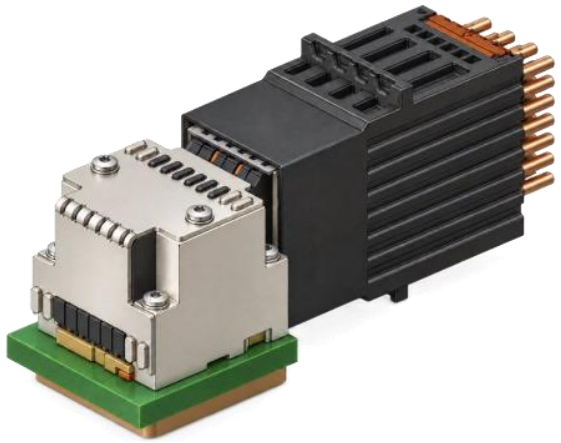
no switch-to-switch links, nothing blocks

Every GPU reaches every switch directly. Any two GPUs share all 64 switches, so GPU-to-GPU bandwidth is the full 64 lanes, with no second tier and no switch-to-switch hop.

A Copper Lane – Package-to-Package



GPU package: twinax flyovers launched at the substrate edge (CPC)



RAF backplane connector: cabled, blind-mate, SI held through 1 mm of de-mate

GPU pkg	launch	RAF	twinax 0.8 m	RAF	launch	switch pkg
3 dB	1 dB	1.5 dB	11 dB	1.5 dB	1 dB	3 dB

No host PCB in the path: 22 dB of a 40 dB budget at 53 GHz, 18 dB left

22 dB

used, full-rack run

0.8 m twinax at 14 dB/m, both ends launched at the package

18 dB

headroom

versus 5.8 dB when the ASIC launches through host PCB

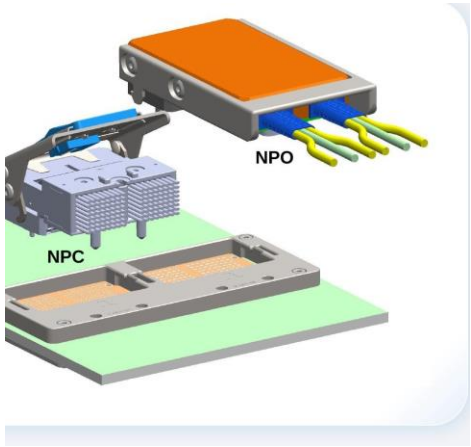
3.8 m

reach at 6 dB/m

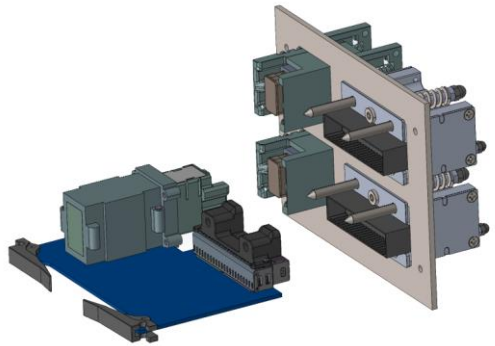
next-gen dielectric: two racks could stay copper

Launch from the package and the budget goes to the cable

A Fiber Lane – NPO-to-NPO



NPO on the shared NPC/NPO footprint: ~9 dB electrical, the shortest channel on the tray



Tray rear edge: RAF twinax and a floating optical blind-mate side by side

2.5 dB

optical budget, end to end

OCI MSA link model, 500 m SMF

≈0.3 dB

per separable connection

six to eight of them between the two NPOs

0.1 / 0.35 / 0.7

dB: V-groove / MPO / EBO

the connector choice decides whether 12,288 fiber lanes close

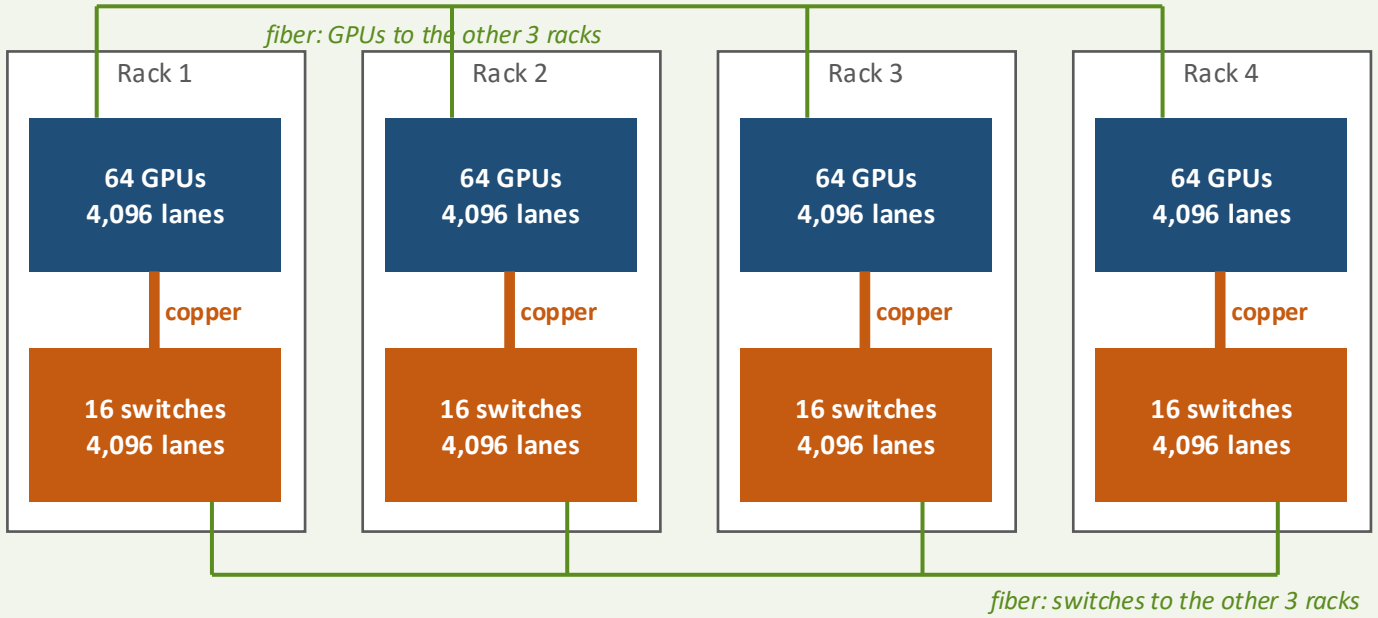
GPU pkg	NPO	sep.	blind-mate	shuffle / trunk	blind-mate	sep.	NPO	switch pkg
9 dB el.	Tx	0.3	0.3	0.3 + 0.3	0.3	0.3	Rx	9 dB el.

Six separable optical connections ≈ 1.8 dB of a 2.5 dB optical budget; 0.25 dB for fiber

Copper or Fiber, a Lane is a chain of component budgets

256 GPUs – Copper in-rack, Fiber-between

Four racks, each 64 GPUs + 16 switch ASICs



16 + 48

lanes per GPU

16 copper to its rack, 48 fiber to 3 racks

64 + 192

lanes per switch

64 copper from its rack, 192 fiber from 3 racks

25 / 75

copper / fiber, %

4,096 copper lanes, 12,288 fiber lanes

Inside a rack 64 GPUs × 16 switches × 1 lane = 1,024 copper lanes.
Between any two racks 64 GPUs × 16 remote switches, both directions = 2,048 fiber lanes.

Copper carries every link in the rack, Fiber carries every link that leaves the rack

Questions for the Open Community

1. At 448G can we define a link budget for a broadly adopted implementation for a Linear link across the host, channel, and module?

2. As architectures evolve are we limited to fixed test-point limits or can we deploy a more dynamic negotiated loss-and-crosstalk budget?

3. Which point of modularity (chip, module, system) will best achieve both the RAS goals and the TCO requirements for will meet industry requirements for next-gen AI data centers?

4. For Blind-mate optical backplanes: which interface (MPO, MMC, self-cleaning) and which float and cleanliness spec should will meet industry requirements for scale-up racks?

Value of Industry Consensus (Standardization)

ELECTRICAL			
IEEE 802.3dv 400G/lane PMDs	**OIF CEI-448G** 448G Physical Layer	**IEEE 802.3 E4AI** Technology Pathfinding	**OIF CEI-224G-Linear** Linear Pluggables
OPTICAL			
OCI MSA Optical Scale Up	**400G Optical MSA** 400G Pluggable Optics	**OCP: SROI** Optical Compute	**OIF 12.8T NPO Module** NPO @200Gb/s per Lane
SYSTEM/FORM FACTOR			
OCP: ESUN, Open Rack Open AI infrastructure	**XPO MSA** Optics Innovation	**OpenCPX MSA** Standardized NPO/CPO	**More to come!!**

Industry Goal: Identify consensus implementations to keep pace with AI demand

Call to Action

Problem to solve

- Linear optics and 448G copper both push the link budget into the passive interconnect, which has no open, reusable definition from package to faceplate to backplane

Timeline

- Q4 2026: scope and reference channels
- 1H 2027: first contributed measurements and test-point proposal
- 1Q 2028: specification consensus and first deployment

Where to find more

- IEEE 802.3 | OIF Forum | XPO MSA | OpenCPX MSA | OCP SROI | OCP ESUN | OCP Server | OCP Optics

Thank You!


Amphenol
Booth Number: C4
The Full Spectrum of Data Center Connectivity
End-to-End Optical and Copper Interconnects, Power Delivery, and Thermal Management for AI-scale Infrastructure
Optical
Copper
Power
Thermal
OPEN Compute Project®
Next Gen
448G Performance
Visit our booth
Watch live demos and meet our experts