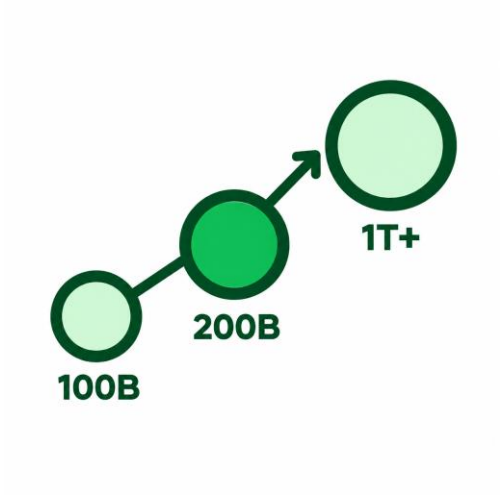
Light on the Horizon: Architecting Scale-Up Fabrics with NPO & CPO for the Next-Gen Data Center

Chris Blackburn *Director of Platform Solutions Architecture, Astera Labs*

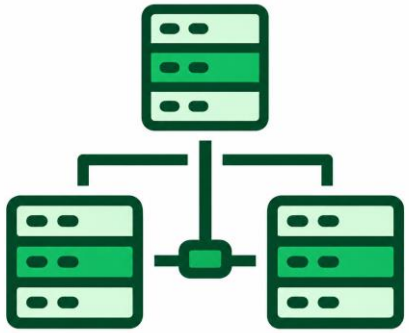
Jignesh Shah *Sr. Director of Product Marketing, Signal Connectivity Group, Astera Labs*

OCP PLATINUM Member

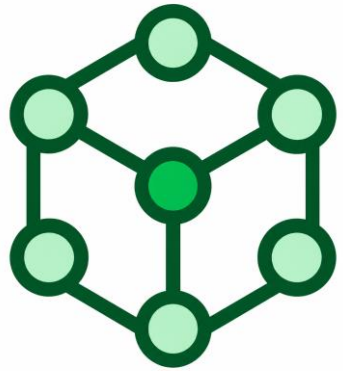
AI is Rewriting the Rules of the Data Center



Model sizes doubling every 12–18 months with parameter counts moving from hundreds of billions to **trillions**



GPU/accelerator clusters scaling from dozens to **thousands** of endpoints



Scale-up fabric is now a **first-order design constraint**, not an afterthought

The **interconnect** is the bottleneck — compute is only as fast as the fabric allows

Three Forces Reshaping Scale-Up Fabrics

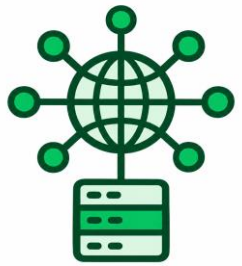
Trend	Today	Tomorrow
Scale-Up Pod Size	8–72 XPUs / pod	144–1000+ XPUs / pod
Topology	Scale-Up then Scale-Out	Increased Scale-Up Domain
Bandwidth	PCIe 5/6, 200G/lane	PCIe 7 → 8, 400G+/lane

- Scale-up size increasing to keep pace with AI workloads (ex. Agentic inference, Recursive Self-improvement)
- Single-hop topologies reduce latency and increase XPU utilization — critical for all-reduce collective operations
- PCIe 6→7→8, UALink & ESUN 200G/lane → 400G/lane: bandwidth doubling, increased frequencies, shrinking electrical channel distance

System **architectures are evolving** to align with these key trends

The Five Pillars of Customer Decision Making

Performance



- Signal integrity
- Maximum interconnect distance at Nyquist frequency
- Bundle size and density

Power



- Watts per XPU-XPU bit
- Total rack / pod power envelope
- Power usage effectiveness (PUE)

Cost



- CAPEX (NRE & BOM)
- OPEX (power, spares, maintenance, uptime)

Speed to Deployment



- Supply chain readiness
- Cabling complexity
- Production & installation timelines

Reliability & Serviceability

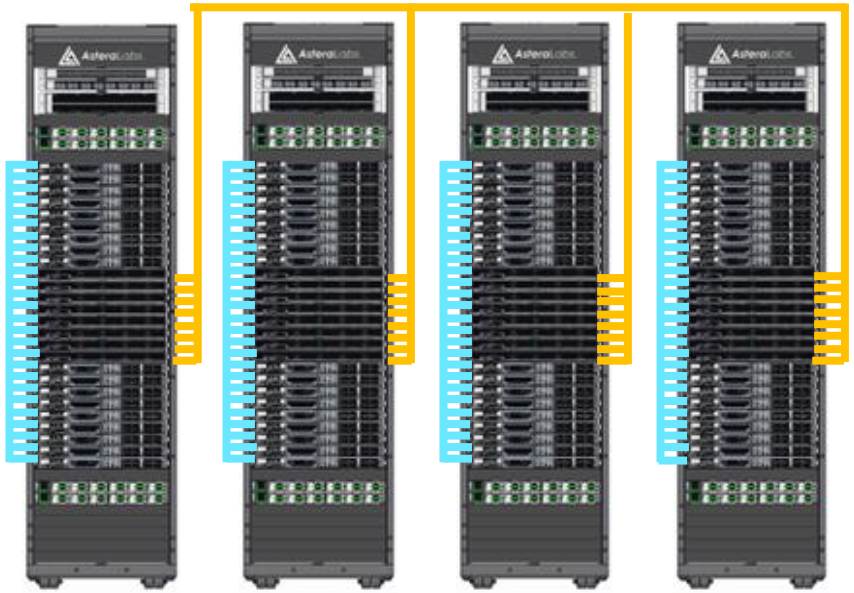


- MTBF
- Redundancy architectures
- Fabric healing
- Hot-swap
- FRU granularity
- Diagnostics

No single technology wins on all five dimensions. The right answer **depends on the application**.

One Size Does NOT Fit All

Training Cluster



- Long sustained all-reduce collectives
- High bisection bandwidth is critical
- Large physical distances between nodes

Inference Cluster



- Short burst, latency sensitive traffic
- TCO sensitivity, prioritizes bandwidth per dollar
- Shorter physical differences between nodes (but growing!)

Copper

Optics

Technology choice is not binary – **copper and optics will coexist**

Today’s Fabrics: Proven, Copper-based, and Evolving





- NVLink, PCIe, & UALink scale-up protocols viable over copper for this generation and next generation
- Copper cable backplanes are ubiquitous scale-up interconnect architecture today
- DACs, ACCs, AECs handle inter-rack scale-up connections

Copper Delivers the Backbone for Next-Gen Scale-Up

Performance
Continues to support next-gen data rates

Power
Lowest power per bit

Cost
Lowest bandwidth per dollar

Speed to Deployment
Proven, ready supply chain

Reliability & Serviceability
Decades of maturity

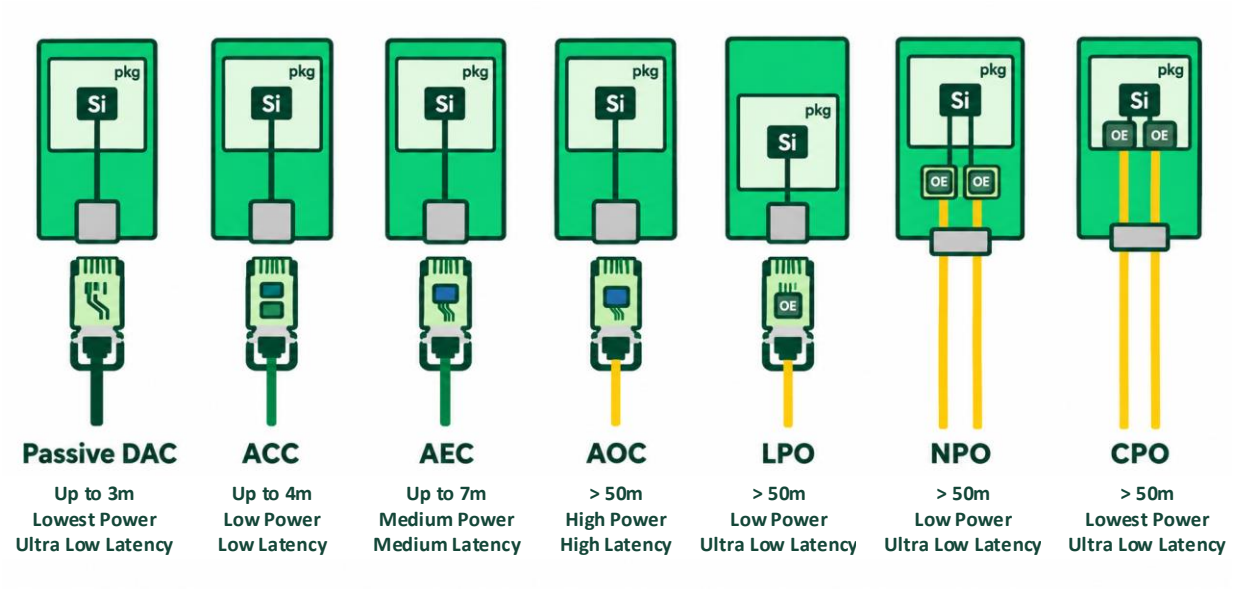


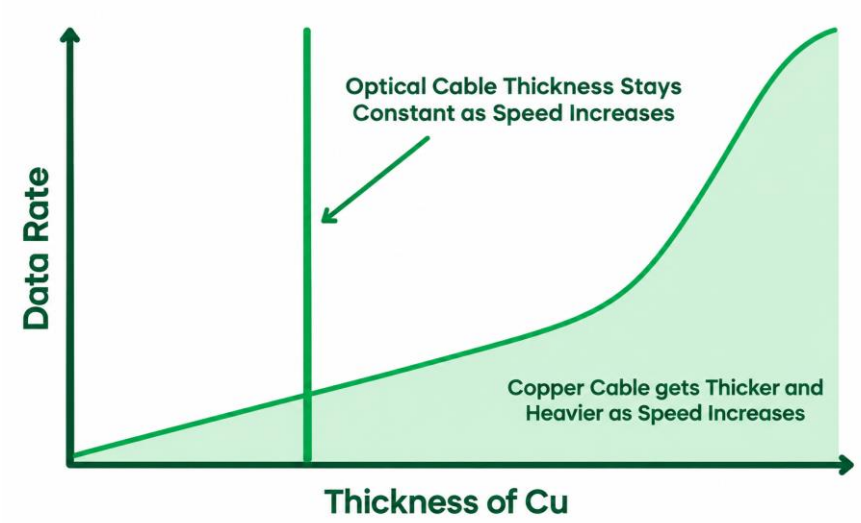
Copper will not disappear. The question is where the boundary sits — and that boundary is evolving.

Why Optics is Coming to Scale-Up

- Pod sizes beyond 144 GPUs require physical distances that need the reach of Optics
- High Radix Switches and XPUs require higher density connectivity solutions
- Increasing data-rates brings optics closer to the Silicon
- Lower system power, lower link latency with optics



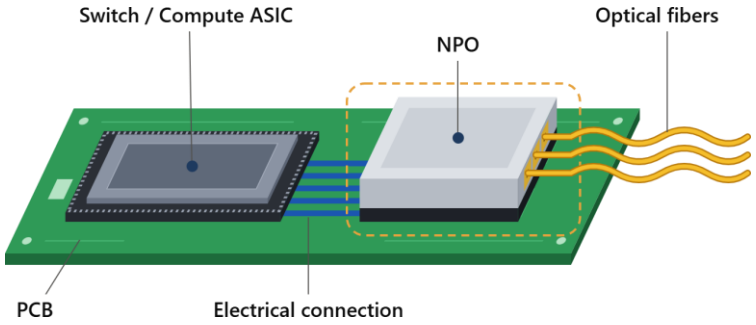




Copper Cable Limitiations, www.silicon-line.com

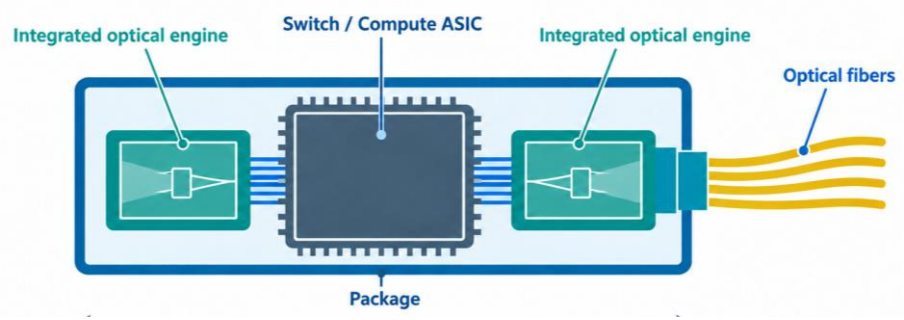
NPO: Near-Packaged Optics

- **Density:** Higher beachfront density of NPO than front panel pluggable modules
- **Power:** Significantly lower than retimed pluggable
- **Latency:** Ultra-low with linear optics
- **Optical connectivity:** Requires robust on-board fiber routing
- **Serviceability:** Electrically pluggable, fiber-pigtailed
- **Ecosystem maturity:** Uses existing front panel module manufacturing and test ecosystem
- **Yield / Cost:** Increased lane count as compared to pluggables

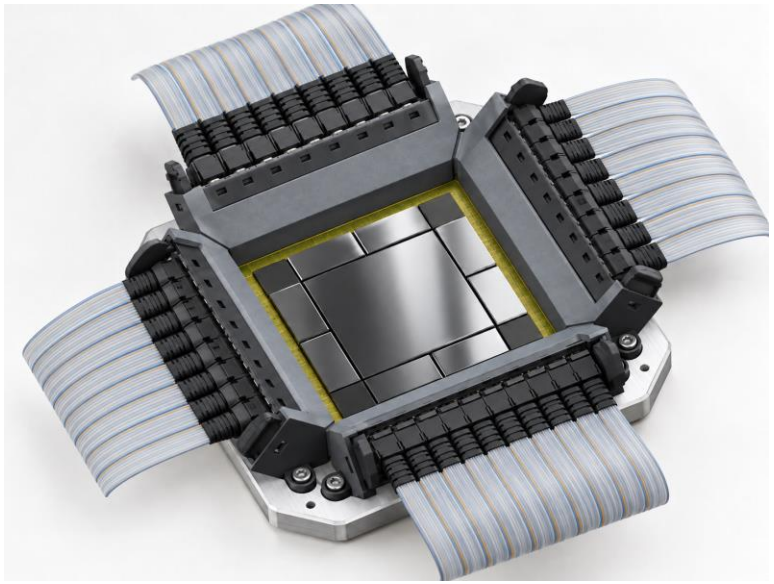




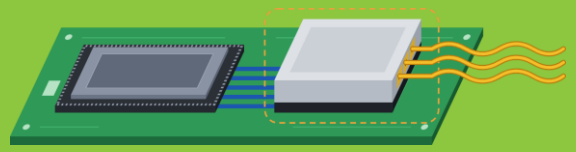
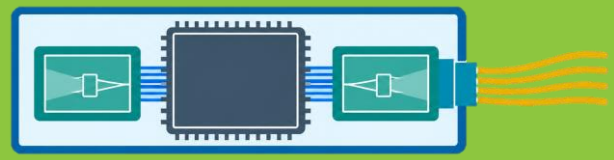
CPO: Co-Packaged Optics



- **Density:** High Radix scalability, Highest beachfront utilization
- **Power:** Best-in-class W/Tbps
- **Latency:** Ultra-low — no additional signal conditioning in end-to-end link
- **Optical connectivity:** Detachable Fiber-to-package interface
- **Serviceability:** Not easy as Co-packaged ‘Optics + Switch’
- **Ecosystem readiness:** WIP for Supply chain, Advanced packaging, Testing and Standards
- **Yield / cost:** Increased package complexity affects overall yield



NPO & CPO Comparison

	NPO 	CPO 
Beachfront Density	Enough for near-term applications	Scalable for future high radix applications
Power	Significantly lower than pluggables	Lowest of all optical architectures
Serviceability	Good, Electrically detachable like pluggables	Difficult, if optical engine integrated with Switch/XPU
Form Factor Commonality	Cu and Optics	Optics only
Manufacturing Complexity	Low, leverages existing ecosystem	High, new ecosystem of partners needed
Most Likely Use Case	Near-term multi-rack AI scale up	Next-generation switch and XPU platforms

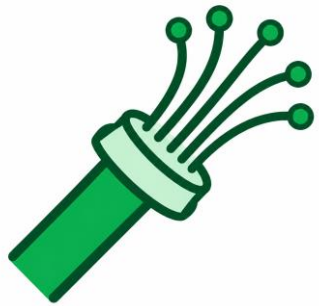
The Sweet Spot: Hybrid Copper & Optics

Copper



- Utilize copper for links within chassis, within rack, and for rack to rack
- Use retimers and redrivers to extend the reach of copper where necessary
- Optimize for cost and power

Optics



- Utilize optics for links from rack to rack and across the pod
- Use linear optics, NPO, & CPO technologies to reduce pJ/bit
- Optimize for physical reach and density

Flexibility in Connectivity enables Optimized AI Infrastructure

Copper & Optics Form Factor Commonality

Pluggable Form Factors

DAC / ACC / AEC



AOC



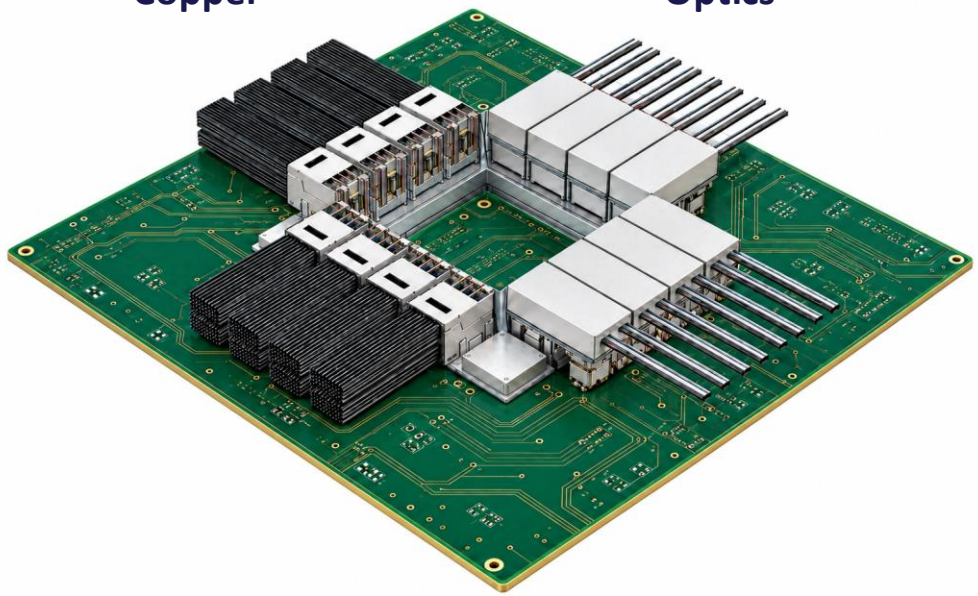
Transceiver



Near-Packaged / Co-Packaged

Copper

Optics



Hybrid (Copper & Optics) Pod Architecture

Hybrid Copper / Optical Pod



- Continue to use copper cable backplanes to scale up within the rack
- Use NPO to scale up rack to rack
  - Direct to GPU or Switch-Switch
- Targeted towards scale up sizes between 72 and 576 accelerators

Full Optical Pod



- Dedicated compute and switch racks
- Use NPO / CPO to scale-up whole pod
- Targeted towards scale-up sizes >576 accelerators

Copper
Optics

The Path is Clear. Let’s Build it Together.

The next generation of AI infrastructure will be defined by the interconnect, not the compute.

Copper built the foundation. Optics will build the ceiling.

The architectures that win will use both — intelligently, in the right places, for the right reasons.

NPO and CPO are not a threat to existing investments.

They are the **natural evolution** of a fabric that has to scale by orders of magnitude in the next five years.



Copper continues to thrive as Optics comes alive

Call to Action

- Engage across OCP & other industry standards to drive ecosystem alignment
- Visit the **Astera Labs Booth B34** to learn more about the switches that enable these architectures
- Visit **AsteraLabs.com** for more information
- Come chat with us at the show or drop us a line:

  Jignesh.Shah@AsteraLabs.com & Christopher.Blackburn@AsteraLabs.com

Thank You!

