

External Laser Source Architectures for Deployable CPO-Based AI Networks

Christian Urricariet

Senior Director of Product & Strategic Marketing

Lumentum



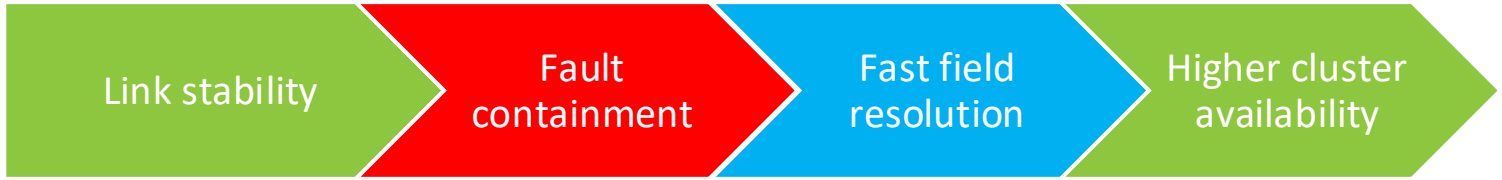
CPO: From Energy Efficiency to Deployment Readiness

- CPO adoption was first justified by energy efficiency and bandwidth density.
- At AI scale, reliability and service model expected to become equally important.
- Optical failures can affect realized cluster utilization, not just component yield.
- The question shifts from “Can CPO work?” to “Can CPO be operated?”

Power Savings: Pluggables → CPO

- ~30W → ~9W per port(*)
- ~3.5x improvement(*)
- Lower electrical loss from shorter ASIC-to-optics traces
- Optimized Serdes – No DSP

(*) Lumentum illustrative estimates; results may vary by configuration.

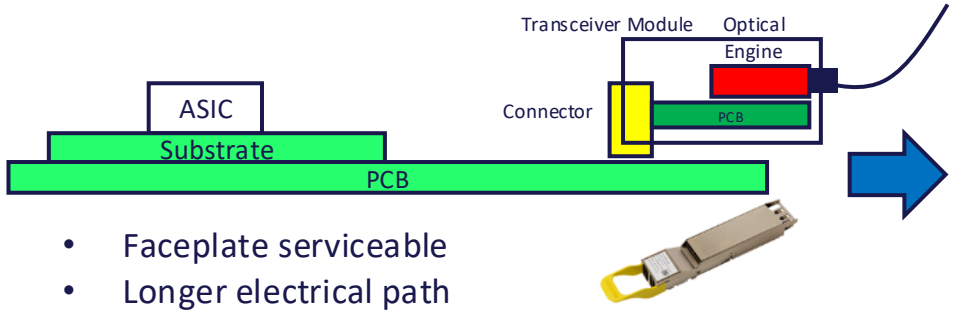


CPO adoption may depend on delivering power savings without compromising reliability, serviceability, or operational control.

Moving Optics Closer to Silicon Changes System Boundary

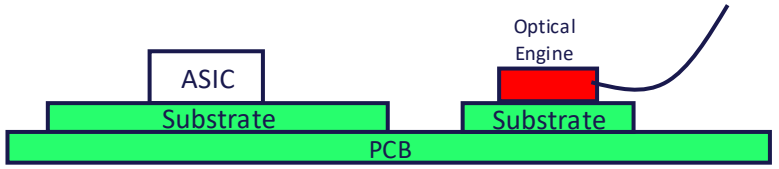
- Shorter electrical paths improve power efficiency and bandwidth density.
- The optical engine (OE) moves closer to high-value silicon.
- Service access, thermal coupling, and failure-domain impact become first-order design considerations. The optical architecture must be treated as part of the system design, not just as a module-placement choice.

Pluggable Optics



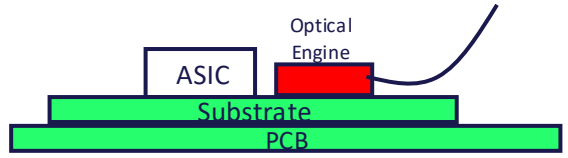
- Faceplate serviceable
- Longer electrical path

Near-Packaged Optics (NPO)



- Optics adjacent to the ASIC package (socketed or board-mounted)
- Reduced electrical loss

Co-Packaged Optics (CPO)



- Optical engine integrated on ASIC package
- Minimum electrical I/O

Laser Placement Is a Separate Architecture Axis

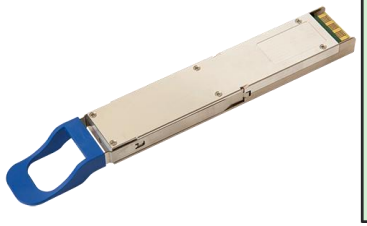
- **OE placement:** Where is the optical engine relative to the xPU? (Pluggable vs. NPO vs. CPO)
- **Laser placement:** Where is the CW source relative to the OE, thermal zone, and service boundary?
- The tradeoff is optical power delivery complexity versus serviceability, thermal isolation, and failure containment.

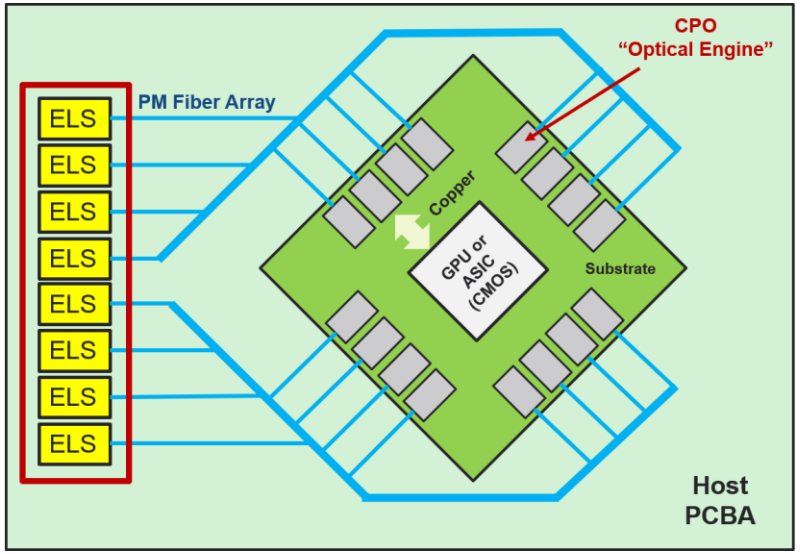
Laser Co-located with OE (ILS)	Board-level External Laser Source (ELS)	Faceplate ELS / ELSFP
• Short optical path (PMF) • Fewer optical interfaces • Tighter thermal coupling • Harder field replacement • Service boundary moves inward	• Thermal separation from ASIC • Independent laser optimization • Not faceplate serviceable • PM fiber / optical power delivery complexity	• Faceplate serviceable • Thermal separation from ASIC • Independent laser optimization • Sparing / replacement flexibility • PM fiber / optical power delivery complexity

External Laser Source Architecture: What Changes

- ELS/ELSFP provides CW/UHP optical power to silicon photonics optical engines from outside the CPO package environment.
- The laser source becomes independently serviceable, monitorable, and optimizable.
- Laser thermal design is decoupled from the ASIC/package thermal zone.
- The system designer gains new options for sparing, redundancy, optical power distribution, and failure-domain containment.

- Service boundary at the faceplate
- Thermal boundary away from the ASIC hot zone
- Laser optimization boundary independent from ASIC
- Failure recovery domain defined at system level





Deployment Metrics Beyond Power per Bit

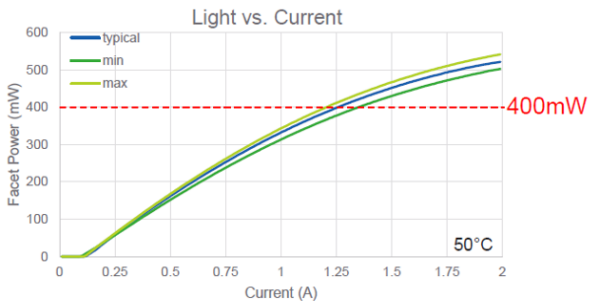
- The following system-level metrics are proposed as important considerations for evaluating CPO deployment readiness:

Optical power delivery	Power / wavelength, loss budget, aging margin
Thermal envelope	Laser temp, ASIC isolation, cooling assumptions
Reliability target	Laser FIT, module FIT, system qualification
Serviceability	FRU boundary, faceplate access, replacement workflow
Telemetry / RAS	Optical power, temperature, alarms, fault isolation
Recovery domain	Sparing, redundancy, MTTR, workload impact

- As these requirements evolve, laser placement becomes an architectural lever and not just a component choice.

Evidence Point: Reliability Designed into the Light Source Platform

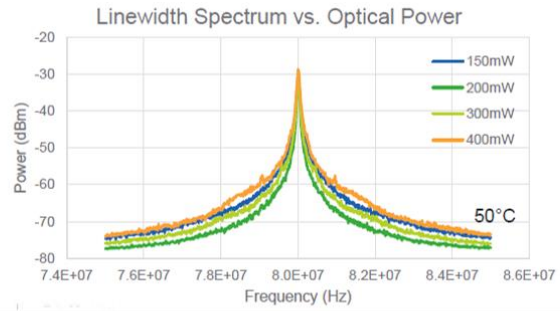
- Deployable CPO requires more than an ELSFP form-factor module.
- Laser-source readiness depends on technology heritage, thermal design, qualification, burn-in, and manufacturing control and readiness.
- High-power operation must be demonstrated with stability and reliability at data center conditions.
- Lumentum's DWDM UHP/ELSFP platform is one example of applying proven laser-platform heritage to CPO external light sources.



>350 mW at 50 °C – reliable headroom for SiPh-based CPO in datacenter thermal environments



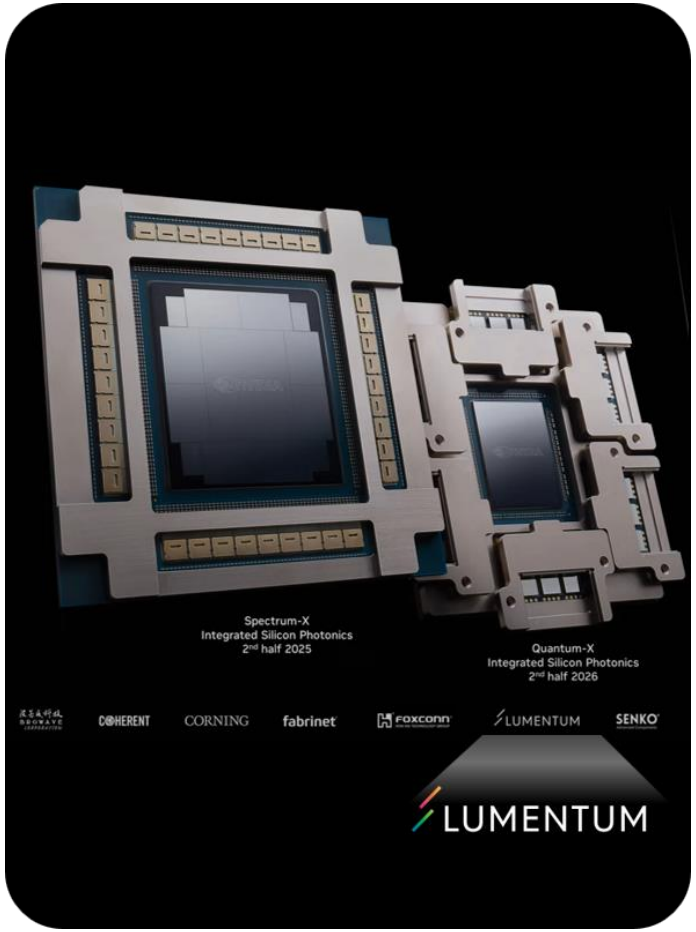
>20% PCE at 50 °C – lowers energy use and cooling costs



<500 kHz linewidth (150–400 mW) – clean spectra for high-speed modulation

From Switch CPO to Future Scale-Up CPO Architectures

- CPO is expected to be a near-term deployment path for certain scale-out architectures, e.g., NVIDIA announced CPO-based Spectrum-X / Quantum-X switches using ELS.
- As OEs move closer to accelerators/scale-up, laser-source placement, service boundary, telemetry, and recovery domains remain system-level decisions, based on how the readiness metrics evolve.
- Emerging architecture discussions are treating photonics as platform infrastructure rather than isolated optical components.
- This talk does not endorse one implementation; it argues that laser-source placement should be evaluated as part of deployability, reliability and serviceability for every application and/or system generation.



CTA: Define Deployable CPO Laser-Source Requirements

- **Problem to solve:** CPO adoption needs system-level requirements for external laser source deployment, not only component specifications.
- **How to get involved:** Engage in relevant OCP Photonics / AI systems / platform-integration discussions on ELS serviceability, thermal envelopes, telemetry, and recovery domains.
- **Where to find more information:**
  - OIF ELSFP implementation agreement.
  - OCI MSA and Open-CPX MSA.
  - Lumentum CPO / ELSFP technical materials.
  - OCP project pages / contribution database.

Thank You!

LUMENTUM

