Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

IDEA

October 16, 2026 04:28 AM GMT

Telecom & Networking Equipment | North America

OCP 2026: Still Plenty of Hard Problems to Solve, But Problems Largely Known

OCP was less incremental this year, as challenges hyperscalers face in scaling data centers have been well explored during GTC, ECOC, OFC, Computex. In selloff, favor LITE/COHR given NPO catalyst, but OCP did not change opinion.

Key Takeaways

- OCP looks to bring together hyperscalers and the AI ecosystem to find common solutions to problems.
- In past years, this has been incremental with the Ultra Ethernet Consortium or ESUN pointing to ways to move towards common networking standards.
- The focus this year was on grid to data center power connectivity, CPU diversity and continued NPO / CPO debate.
- With NPO / CPO debate well explored recently at ECOC, less incremental than we thought.
- Ethernet scale-up opportunity remains most incremental in near term, a positive for CSCO / ANET; NPO/CPO most incremental for LITE / COHR / KEYS.

We attended the Open Compute Project Global Summit in San Jose, CA from October12th - 15th.

15K people cramped into the San Jose Convention Center, largely learning what they already knew. OCP, as with most AI ecosystem conferences, saw a tremendous increase in attendance this year, very much outgrowing the SJCC, even as the event had to turn away 3K addition attendees. While there were 4 days of detailed sessions, all diving into the debates on how to enable the most compute at the lowest power, after so many well attended industry conferences this year, we did not walk away thinking that OCP moved any of those debates forward significantly (particularly in light of the way it has in the past, e.g., regarding Ethernet challenging Infiniband or ESUM challenging NVLink). Overall, share performance aside, the hyperscalers made it very clear this week that building more compute is the challenge they are solving for, and they need more power, GPUs/CPUs, connectivity and optics to do that (aka demand is good). In terms of beneficiaries, we would note Ethernet continues to be the favored connectivity standard, a positive for ANET / CSCO; optics in scale-up is an inevitability, a positive for LITE / COHR; KEYS will be more necessary given variety of architectures being trialed and scale of optics needing to be tested.

Getting more from the infrastructure is becoming as important as adding

MORGAN STANLEY & CO. LLC

Meta A Marshall

Equity Analyst

Meta.Marshall@morganstanley.com [redacted]

Antonio Jaramillo

Research Associate

Antonio.Jaramillo@morganstanley.com [redacted]

TELECOM & NETWORKING EQUIPMENT

North America

Industry View **In-Line**

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

capacity. Networking is taking on a larger role in AI systems, both to connect more accelerators and to get more work from installed compute. Faster links remain necessary, but fewer interruptions and better utilization are increasingly part of the same investment decision. We see a growing opportunity across our networking names rather than a single product cycle. However, customers are also working to reduce the equipment and power required to support larger clusters. A larger TAM remains supportive, but bandwidth growth alone will not determine supplier revenue. Customer deployments, equipment content and operating performance should become more important differentiators as architectures mature.

Open systems are moving toward implementation, with supplier choice still a priority. Common switch platforms and more portable software are intended to reduce repeated engineering work and make new hardware easier to introduce. Optical designs are following a similar path, with open specifications moving toward products and qualification. Customers want lower power and higher density without giving up multiple sources or making repairs more disruptive. We think the implication is a longer period of coexistence rather than a quick transition to one architecture. ANET and CSCO will need to differentiate through integration and operations, while for LITE and COHR, broader optical adoption should be considered separately from the content retained in each design.

Ethernet is moving beyond the initial scale-up standards discussion. Shipping Ethernet scale-up switches and further development of common fabric requirements make the opportunity more tangible, although adoption of newer specifications remains ahead. We see a broader addressable market for ANET and CSCO as customers connect larger pools of compute using a common networking foundation. However, greater openness does not necessarily mean less competition from NVDA (covered by Joseph Moore), which supports multiple network operating systems and is partnering across the networking stack. Winning the architecture debate and winning the system business remain separate questions, with production platforms and supplier selections the next milestones.

Scale-across is becoming part of the cluster design rather than a separate network upgrade. Connecting compute across sites brings latency, congestion and recovery directly into the economics of AI workloads. Longer distances support the optical transport opportunity, which we view as particularly relevant for CIEN. Routing and switching opportunities also expand, although approaches differ on buffering and where network intelligence resides. We would not assume every distributed cluster requires the same packet architecture. Capacity deployed between sites should matter alongside the mix of optical and routing equipment customers select.

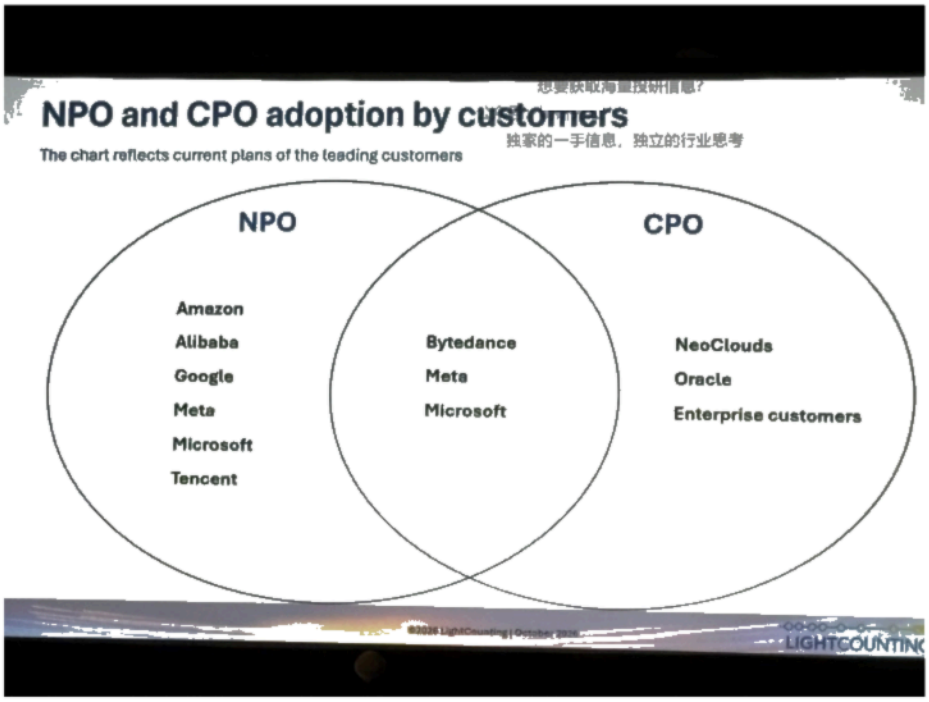
Quick hits:

- **CPO reliability is improving, but remains implementation specific.** META (covered by Brian Nowak) reported a 6x improvement in interruption rates versus pluggables in its more mature lab systems, with no non-serviceable failures over 1.35 million system hours. Encouraging results strengthen the deployment case, although subsequent platforms still require qualification before extending the conclusion across vendors.

- **Pluggables still have a roadmap, but laser supply matters.** NVDA's OCP Q&A pointed to USFP production next year, implying 2027, with CSCO identified separately as a development partner. For EML-based designs, we think existing 200G EML supply tightness could complicate 400G EML procurement as suppliers move beyond demonstrations and controlled availability toward a broader commercial ramp.
- **External lasers are not a settled choice.** LITE said an ELSFP module is in production. Faceplate-pluggable laser sources preserve field serviceability, while embedded designs aim to reduce power and assembly complexity. We would track platform selections and laser content alongside the overall CPO ramp rather than assume a uniform opportunity across designs.
- **Flatter fabrics change the equipment mix.** META's multiplane architecture moves from three switch tiers to two, while NVDA cited hardware-based failure detection and rebalancing below 3 milliseconds. Larger clusters support demand, but fewer tiers and more intelligence at the endpoints make network content per accelerator more important than headline bandwidth alone.
- **Inference broadens the networking discussion beyond the GPU back end.** NVDA described a context-scale storage network optimized for KV cache, alongside scale-in networking connecting users, models and data while extending security across the AI factory. We think the broader architecture creates additional areas to watch for ANET and CSCO beyond training fabrics, with supplier participation still dependent on customer implementations

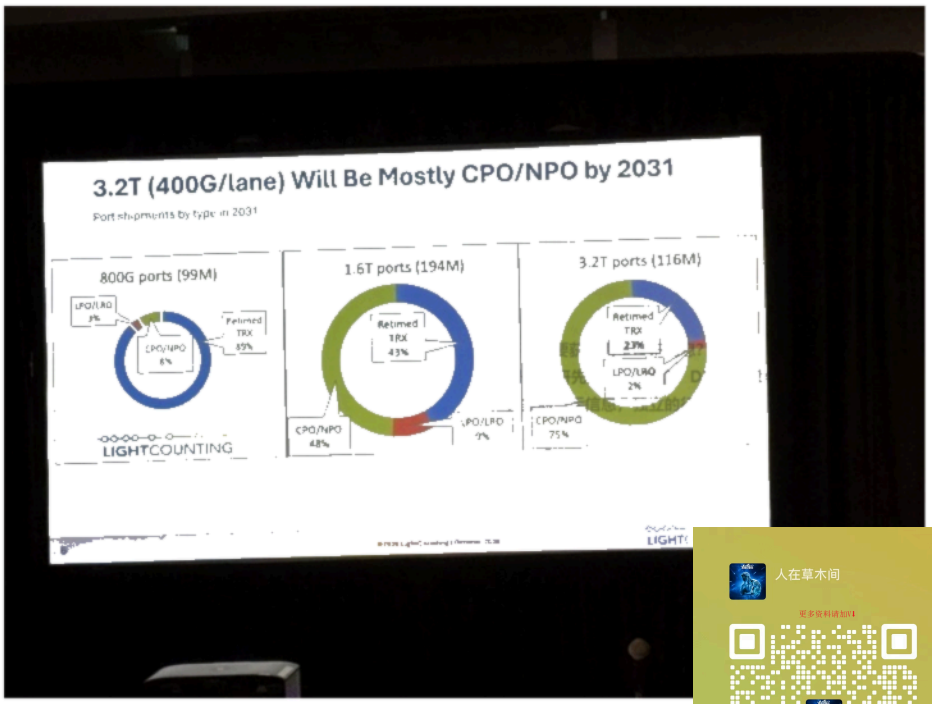
Exhibits

Exhibit 1: Ecosystem More Beholden to NVDA to Be Early CPO Adopters, Most Others Exploring NPO



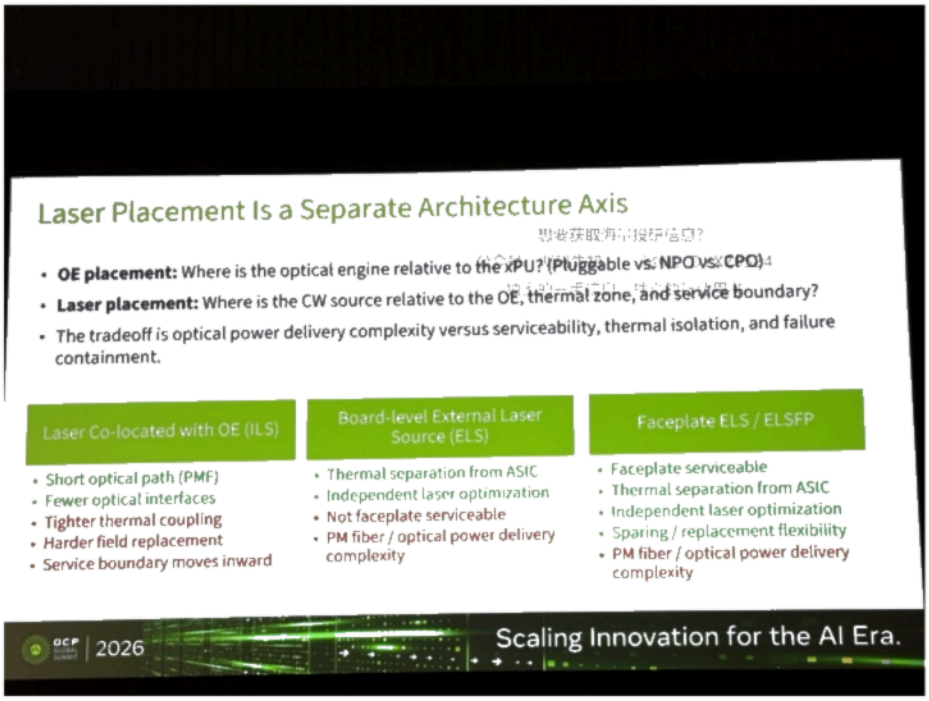
Source : LightCounting

Exhibit 2: According to Light Counting, CPO to be More Mainstream By 2031



Source : LightCounting

Exhibit 3: Potential impact of laser placement remains highly debated



Source: Lumentum OCP Presentation

Valuation Methodology and Risks

Coherent Corp (COHR.N)

~28x Base FY28e EPS (~$13.3). Multiple trades at premium to historical trading range given AI exposure, trading more in-line with AI comps.

Risks to Upside

- Pricing increases help gross margins
- 1.6T AI opportunity ramps faster than expected, benefitting margins
- Recovery in non-datacomm portions of business
- CPO / NPO ramps quicker than expected

Risks to Downside

- CPO / NPO delayed
- AI build out pause
- Competitive pressures in datacomm

Lumentum Holdings Inc (LITE.O)

25x CY28 P/E of ~$40. ~25x is a premium to historical trading range given AI growth outlook

Risks to Upside

- Datacom ramps and revenue mix more attractive vs. expectations
- OCS ramps quickly, bringing additional revenue / profitability source
- AI demand remains strong, causing EML business to outperform and drive profitability
- CPO / NPO ramp faster than expected

Risks to Downside

- EML pricing falls given oversupply on reduced AI demand
- CW takes share from EML
- Cloud Light hurts ability to improve margins
- Pushout in CPO timing