28 Sep 2026 00:58:52 ET | 26 pages

US Communications Equipment

OCS an $11B Opportunity; Raising LITE TP to $1,400

CITI'S TAKE

Following Citi TMT Conference and ECOC where we met management teams of COHR and LITE, it is clear to us that OCS is turning into a multi-year trend in AI Optics. Putting it all together, Citi debuts its OCS TAM Model that indicates OCS could be an ~$11B opportunity for OCS merchant vendors such as COHR and LITE (see detailed assumptions on section 6). As we revise our prior LITE OCS estimates (see section 7), we raise our LITE sales estimates in FY2027E and FY2028E by +5% and +3%. This coupled with management comments on OCS and UHP margins being in the 60s% and 70s% respectively put our FY2027E and FY2028E EPS estimates at $25 ($23 prior) and ~$41 ($37 prior). We subsequently raise our LITE TP to $1,400 on consistent 35x PE while maintaining COHR's at $420.

Papa Sylla [AC]
+1-212-816-9476
papa.sylla@citi.com

Atif Malik
+1-415-951-1892
atif.malik@citi.com

TAM Bottom Up — Our model estimates a $11B Merchant OCS TAM by 2030. Our bottom focused on XPU-per-port model indicates a ~$5B of merchant OCS TAM by FY2030 from the scale-up layer alone, plus ~$2B from a newly rebuilt, port-driven Spine & Other Related Applications bucket, for a ~$7B FY2030 Combined Total (excluding non-Google OCS reconfiguration). We fielded a Lumentum management comment characterizing a reconfiguration opportunity — distinct from the Google-anchored buckets above — as potentially ~10x the size of Google's, on roughly a 1:1 ratio with AI servers, materializing as a "late 2028/2029" story. Citi believes this to be likely tied to NVDA. Thus, adding a third bucket based on NVDA's GPU and NVL server shipments that we estimate can add another ~$4B to the TAM to ultimately reach $11B.

Key Winners — We expect LITE and COHR to remain the key drivers of OCS and expect the two to accumulate 80%+ of the merchant OCS wallet share.

															Current Fiscal Year		Next Fiscal Year	
						Rating		Short-Term	Target Price						EPS		EPS	
Company	Ticker	Ccy	Price	Mkt Cap (M)	Date & Time	Old	New	View	Old	New	ESPR (%)	Div Yld (%)	ETR (%)	Last Rpt Yr	Old	New	Old	New
Coherent Corp	COHR	US$	295.83	57,933	25 Sep 16:00	1	nc	N/A	420.00	nc	42.0	0.0	42.0	Jun-26	9.57	nc	16.11	nc
Lumentum Hldg	LITE	US$	941.65	84,466	25 Sep 16:00	1	nc	N/A	1,200.00	1,400.00	48.7	0.0	48.7	Jun-26	23.03	25.14	37.21	40.90

1 = Buy, 2 = Neutral, 3 = Sell, H = High Risk
ESPR = Expected Share Price Return, ETR = Expected Total Return, nc = no change
Source: Citi Research
^Catalyst Watch

See Appendix A-1 for Analyst Certification, Important Disclosures and Research Analyst Affiliations.

Earnings Estimates

				Last Reported Year					Current Fiscal Year					Next Fiscal Year				
Company Name	Ticker	Last Rpt Year	Currency	1Q	2Q	3Q	4Q	FY0	1Q	2Q	3Q	4Q	FY1	1Q	2Q	3Q	4Q	FY2
Coherent Corp	COHR	Jun-26	US$	1.16	1.29	1.41	1.74	5.59	1.99	2.21	2.55	2.81	9.57	3.41	3.83	4.22	4.64	16.11
Lumentum Hldg	LITE	Jun-26	US$	1.10	1.67	2.37	3.23	8.68	4.18	5.10	6.54	9.24	25.14	10.44	10.97	11.55	8.06	40.90
Old		Jun-26	US$	1.10	1.67	2.37	3.23	8.68	4.18	4.97	6.21	7.60	23.03	8.62	9.06	9.58	10.02	37.21

Source: Citi Research

1. Investment Thesis: Why OCS, Why Now

Optical circuit switching (OCS) routes light directly from an input fiber to an output fiber by physically or optically redirecting a beam, with no optical-electrical-optical (O-E-O) conversion and no packet processing in the data path. That makes it fundamentally different from the electrical packet. An OCS switch is reconfigurable but not packet-aware, so it trades microsecond-scale, per-packet flexibility for near-zero latency, at a fraction of the power draw.

At ECOC in September 2026, OCS moved from a single-customer curiosity (Google's internal Palomar/Jupiter deployments, which date to the TPU v4 generation) to a multi-customer, multi-layer commercial category. Three data points anchor that shift, each covered in detail later in this note: (1) Lumentum disclosed a ~$400M OCS order backlog from three customers, guided to $400M of OCS revenue in 2H CY26 and point OCS to the main reason the company is comfortable setting a $40 FY2028 (Jun-28) EPS floor; (2) Coherent has now disclosed a specific, standalone "more than $4B" addressable OCS opportunity across DCI, scale-out, and scale-up applications, doubled from the $2B figure it gave at OFC in March 2026; and (3) a second manufacturer, HUBER+SUHNER/Polatis, opened a dedicated new OCS production line in September 2026 specifically to keep up with demand. All three recent developments point to OCS being a multi-year, multi-company trend.

For our coverage, the read-through is direct: LITE and COHR are the two established, merchant-scale OCS suppliers, and this note's model focuses on OCS as an incremental, high-growth product line layered on top of both companies' core transceiver/CPO franchises.

This note updates and formalizes work we have been building over the past several weeks following positive comments on OCS by both LITE and COHR at the companies' earnings, Citi Technology Conference color and ECOC commentary. Notably at Citi TMT Conference, LITE provided an XPU-per-port bottom-up model that started as a scale-up-only build for the key OCS adapter Google (covered by Ron Josey). This was subsequently extended with a second demand bucket for spine and other non-scale-up applications. Section 6 walks through the different buckets' current mechanics in full; an accompanying Excel workbook is available on request.

1.1 A brief history: from Google's Palomar to a $multi-billion-dollar commercial category

OCS is not a new invention — MEMS-based optical switching has existed in telecom labs and test equipment for over a decade (Polatis's roots trace back well before the current AI cycle). What changed is the buyer. The timeline below traces the shift from a single hyperscaler's internal infrastructure project to a broad-based, multi-vendor commercial category with company-disclosed, multi-billion-dollar TAM figures.

Figure 1. OCS Milestones

Approx. period	Milestone
2015 2019	Google begins deploying MEMS-based OCS internally (Palomar switches) to reconfigure its Jupiter data center fabric; largely undisclosed publicly at the time.
2021-2022	Google publishes technical detail on OCS-based Jupiter spine replacement and TPU v4's 128-port MEMS switch architecture (arXiv:2304.01433), the first detailed public disclosure of hyperscale OCS in production.
2023-2024	Lumentum and Coherent both begin shipping merchant OCS products; Coherent announces a dedicated optical circuit switch product line (Mar-2024).
Early 2025	Cignal AI's first OCS market sizing places the market's prior (Jan-2025) forecast meaningfully below where it would revise to later in the year.
Dec-2025	Cignal AI revises its OCS market forecast to "exceed $2.5B in 2029," ~40% above its Jan-2025 forecast, citing broader multi-vendor adoption.
Mar-2026	OFC 2026: Coherent first discloses a standalone OCS addressable opportunity figure (~$2B) as part of a broader $23B company growth framework; NVIDIA's $2B investments each in Lumentum and Coherent are in place by this point.
Apr 2026	Open Compute Project publishes its first dedicated OCS white paper, co-authored by iPronics, Lumentum, Ciena, Lumotive, and Carnegie Mellon University — the clearest signal yet of industry-wide standardization interest.
Sep-2026	NVIDIA invests $125M in iPronics; ECOC 2026 (Malaga) sees Coherent unveil PhotonLink and HUBER+SUHNER/Polatis open dedicated new OCS manufacturing capacity in Poland.

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research, Company Reports, Cignal AI, OCP, Thomson Reuters

2. OCS Technology Primer

2.1 How OCS differs from electrical packet switching

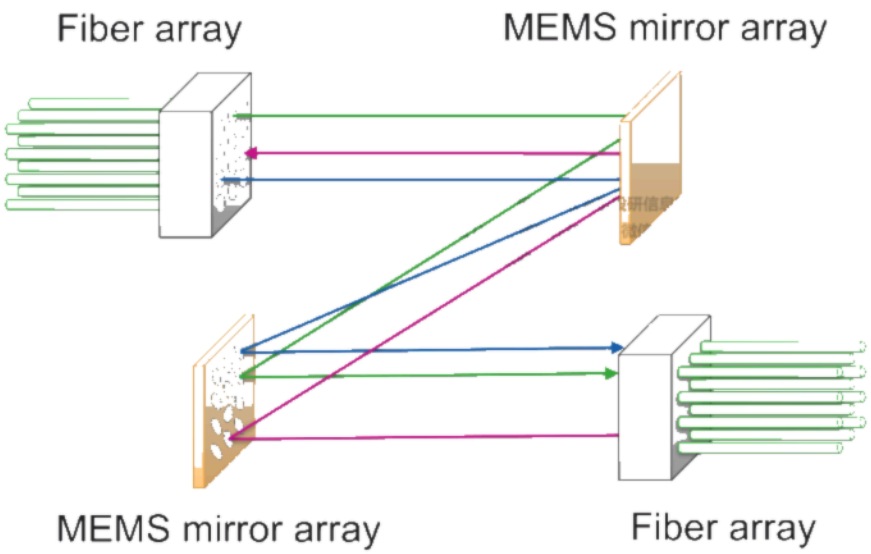
A traditional data center spine or leaf switch receives an optical signal, converts it to an electrical signal, buffers and inspects each packet, makes a forwarding decision, and re-converts the outgoing signal back to light. That O-E-O conversion and packet-level intelligence is what lets Ethernet fabrics do sophisticated things — congestion control, load balancing, telemetry, multi-tenant isolation — but it also consumes power, adds latency, and ties the switch's maximum throughput to the specific ASIC generation inside it (Broadcom's Tomahawk 6, for example, tops out at 102.4 Tbps per chip).

OCS instead reconfigures a physical or optical path between fixed sets of fibers and then gets out of the way: once set, the connection is a piece of glass, not a piece of silicon, so it is effectively speed-transparent. The Open Compute Project's April 2026 OCS white paper (co-authored by iPronics, Lumentum, Ciena, Lumotive, and Carnegie Mellon University) frames the core trade-off cleanly: OCS provides "ultra-low latency and congestion-free operation enabled by fully optical signal paths," but reconfiguration happens on a millisecond-to-second timescale, not a per-packet basis — so it augments, rather than replaces, packet switching for most of the network.

2.2 Switching technologies

- MEMS beam-steering (dominant, in-production): Microscopic mirrors physically tilt to redirect a free-space beam between fiber ports. This is the technology behind Google's Palomar switches, Lumentum's R300/R64 switches. Millisecond-scale reconfiguration, high port count (commercially up to ~300-384 ports today), low insertion loss, and the most mature manufacturing base.

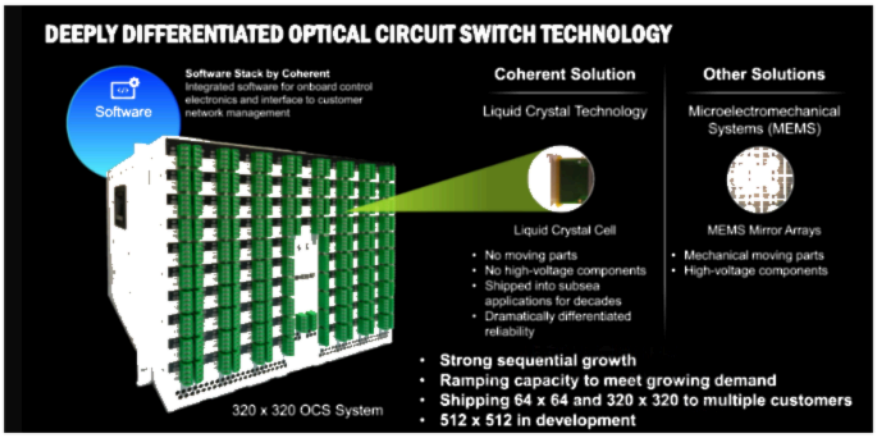
Figure 2. MEMS Technology



Source: Lumentum

- Digital Liquid Crystal (DLC): Uses liquid-crystal cells to steer polarized light. Offers solid-state reliability (no moving parts) but is polarization-dependent and historically has had a narrower port-count range than MEMS though COHR at ECOC showcased its SONiC-Enabled Optical Switching that can reach 520x520 ports.

Figure 3. DLC Technology



Source: Coherent

- Silicon photonics / Mach-Zehnder interferometer (MZI) mesh: Fully solid-state, CMOS-compatible switching fabric etched onto a silicon photonic chip. Fastest reconfiguration speed (microseconds, approaching packet-switching timescales) but historically the most insertion-loss-constrained at higher radix, which limits scalability. iPronics's newly launched switch (with NVIDIA, Maverick Silicon, and Light Street Capital backing, announced Sep-2026) targets exactly this constraint — a second-generation silicon-photonics OCS promising sub-millisecond reconfiguration at higher port density than first-generation MEMS.

- Metasurfaces / 2D photonic beamforming: An emerging approach (Lumotive, backed by a $45M Series B led by Gates Frontier and MetaVC) that uses programmable optical metamaterials to steer beams across a wide angular range (±80°) in a compact form factor. Still pre-commercial-scale for OCS specifically.

- Robotic / free-space fiber-patch systems: The highest-radix, slowest-reconfiguration approach (physically moving fiber patch cords), generally used for static or rarely-changed topology management rather than dynamic AI cluster reconfiguration.

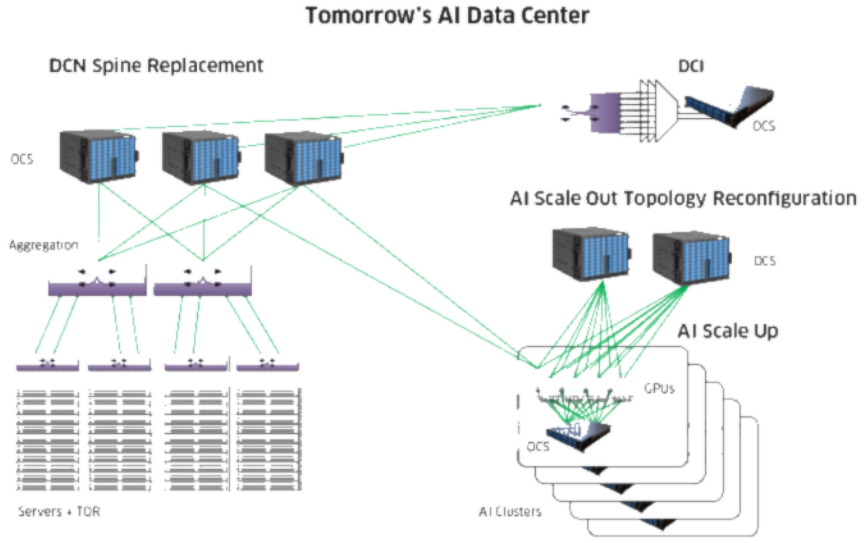
2.3 Port radix and where the market is today

Port count (radix) has scaled quickly: Google's original TPU v4-era Palomar switch was 136×136 (128 usable ports plus 8 spares); TPU v7/Ironwood-era switches are commonly cited at 300+ ports; and the commercial vendors have both followed and in some cases exceeded that. Lumentum's R300 is a 300×300 platform (its highest-volume product to date, per management commentary), though the company also has a smaller R64 (64×64, expandable to 72×72) for lower-port-count applications; Coherent's flagship is a 320×320 platform; and HUBER+SUHNER/Polatis's line spans from 8×8 up to 384×384. In our model, we use a blended, rising port-radix assumption (300 ports in FY2025-27, stepping up to 600 by FY2030) as the OCS switch port radix driver — see Section 6.

3. Applications & Use Cases

We group commercial and near-commercial OCS use cases into five layers.

Figure 4. OCS Use Cases



Source: Lumentum

3.1 Reconfiguration: the core mechanism

Before walking through the five network layers, it's worth separating the layer (where two endpoints sit in the fabric) from the capability that makes OCS valuable there in the first place: reconfiguration. Because AI training traffic is large, sustained, and relatively predictable — heavy, east-west, all-to-all

communication between XPU groups — a network built on OCS can be physically re-wired in software to match each job's communication pattern, rather than running one fixed, general-purpose topology for every workload. We see two distinct flavors of this in the literature and in Google's own disclosures, plus a third, closely related resiliency benefit:

- Static-to-dynamic, per-job reconfiguration (in production today). Google's TPU Superpods use OCS between servers so the cluster can be re-shaped into job-specific topologies — such as a torus — based on that workload's communication pattern, rather than accepting a single fixed fabric for every job. Google has reported this delivered a 3.3x performance improvement and a 9% power reduction on LLM training versus a static electronic-switch fabric of the same scale (OCP white paper, citing Google's published TPU v4 results).

- In-workload, dynamic reconfiguration More recent academic work pushes reconfiguration further: adjusting the optical topology while a job is still running, to match shifting traffic patterns inside a single training run rather than fixing topology once at job start. MixNet (arXiv:2501.03905), a runtime-reconfigurable optical-electrical fabric for Mixture-of-Experts training, reports outperforming TopoOpt's one-shot (job-start-only) reconfiguration approach by up to 2.5x, assuming millisecond-scale MEMS switching latency (10-25ms, benchmarked against commodity Polatis switches). ATRO (arXiv:2507.13717), a newer, solver-free topology-and-routing optimization algorithm for reconfigurable datacenter networks, targets the same problem — jointly optimizing topology and routing rather than treating them separately — and reports outperforming prior approaches on runtime and maximum link utilization, though without a single headline multiplier comparable to MixNet's. We have also seen a claim, from a source we could not independently verify, of a ~2.3x in-workload improvement at sub-100-microsecond switching latency; we flag it as unconfirmed rather than include it as a sourced figure.

- Failure resiliency (a reconfiguration side benefit, in production today). The same software-reconfigurable property lets OCS bypass a failed link or node by re-routing the optical path instead of relying on electronic re-routing or a job restart.

3.2 Scale-up (intra-pod)

OCS connects XPUs (or XPU groups) within a single training pod — Google's canonical use case. TPU v4 "Superpods" used 128-port MEMS switches to connect compute cubes, yielding a disclosed 3.3x performance improvement and 9% power reduction for large-model training versus a fixed electrical topology (OCP white paper, citing Google's own published results). TPU v7/Ironwood scaled this to 9,216-chip pods built from 4×4×4-cube blocks interconnected via 300-port MEMS switches. Google's newest architectures (TPU 8t's 3D torus and TPU 8i's hierarchical "Boardfly" topology, both disclosed in 2026) continue to use OCS at this layer, even as copper handles the tightest, lowest-radix connections within a board or ring.

3.3 Scale-out (pod-to-pod)

A related but distinct layer: connecting multiple pods together, at a materially higher chips-per-switch ratio than the intra-pod layer (industry estimates for Google's Ironwood-generation scale-out layer run in the 500-600 chips-per-OCS-switch range — see Section 7.2 for how we use this to anchor our Spine & Other build). This is the layer where non-Google adoption is least mature today.

since AWS Trainium, Meta MTIA, and Microsoft Maia have not disclosed a confirmed OCS-based scale-out fabric as of this note.

3.4 Spine-layer replacement

The most-cited efficiency case in the industry: replacing electrical packet switches at the data center spine layer with large-radix OCS. Google has publicly disclosed a 41% reduction in power consumption and a 30% reduction in capex from OCS spine replacement in its Jupiter network fabric, alongside a 5x increase in network throughput (OCP white paper, citing Google's published figures). This is also the layer most discussed as a long-term structural risk to traditional EPS spine vendors (see Section 4.6).

3.5 Pooled / disaggregated resources

A more nascent use case: using OCS to dynamically wire compute to pooled memory, storage, or accelerator resources on demand, rather than fixing that topology at design time. This is closely related to hardware failure resiliency — OCS can route around a failed link or node automatically, at the physical layer, without a job restart or measurable performance hit.

3.6 Campus data center interconnect (DCI)

OCS used to interconnect buildings within a single hyperscale campus (as distinct from long-haul, inter-city DCI, which remains a coherent-optics/DWDM application, not an OCS one). Cignal AI names this, alongside spine replacement and pooled resources, as part of its OCS total addressable market.

3.7 Emerging / research-stage

The in-workload reconfiguration research introduced in Section 3.1 (MixNet, ATRO) sits here as well, alongside one additional item:

- Dynamic in-workload reconfiguration ("ACTINA"): Columbia University / NVIDIA research on an OCS-BCube topology with wavelength-reconfigurable transceivers, showing 1.84x faster iteration times than a static 3D torus in simulation (OCP white paper) — a smaller claimed improvement than MixNet's 2.5x vs. one-shot reconfiguration (3.1), though the two papers compare against different baselines and workloads, so the figures are not directly comparable.

- Ultrafast OCS via soliton microcombs: Microsoft Research has published academic work on this approach ("Ultrafast optical circuit switching for data centers using integrated soliton microcombs"), evidence of internal Microsoft research interest in OCS — but we are not aware of any Microsoft production-network disclosure that confirms OCS deployment in Fairwater or any other live Azure AI infrastructure. We treat this as R&D-stage interest, not a modeled revenue driver, and flag it as a reason Microsoft's OCS attach-rate assumptions in our model remain the most conservative of the four platforms (Section 7).

3.8 Use-case maturity summary

"Confirmed deployer" means a company has publicly disclosed the specific use case in question; it does not mean other hyperscalers are not doing the same thing without disclosure. This distinction is exactly why our AWS/Meta/Microsoft attach-rate assumptions in Section 7 are calibrated estimates rather than sourced figures.

4. Competitive Landscape & Key Players

The table below summarizes the vendor landscape as we understand it as of this note. "Disclosed OCS TAM/backlog" figures are company-specific disclosures, not our model output (see Section 7 for our estimates).

4.1 Lumentum (LITE)

Lumentum is, alongside Coherent, one of the two established merchant OCS suppliers and the vendor with the longest track record as Google's primary external OCS partner. Its product line spans the R300 (300×300 MEMS, management's highest-demand platform) and the newer R64 (64×64, expandable to 72×72; >100 Tbps optical capacity; <150W power draw, an ~80% reduction versus packet-based switches; sampling Q4 2025, GA H2 2026). Management disclosed a ~$400mm OCS order backlog from three customers and, on the F4Q26 earnings call (Aug-2026), reiterated a $400mm OCS revenue target for 2H CY26 and guided to OCS's first triple-digit-million revenue quarter in F1Q27 — management's own words: "demand is expanding faster than several capacity ramps." In our model, we hold Lumentum's share of merchant OCS revenue at a flat 45% throughout the forecast — see Section 7.4 for the allocation methodology and its limitations.

4.2 Coherent (COHR)

Coherent is the second established merchant OCS supplier, built around a 320×320 MEMS platform and, as of ECOC 2026 (Sep-2026), a newly unveiled "PhotonLink" integrated-optics platform spanning the full optical signal chain (light generation, beam shaping, transmission, detection, and O-E conversion for the xPU or switch chip). Most notably for this note, Coherent has now put a specific number on its OCS opportunity: management disclosed "more than $4 billion of addressable opportunity across DCI, scale-out, and scale-up networks" on the F4Q26 call (Aug-2026), doubled from the $2bn figure given at OFC in March 2026, with management characterizing even $4bn as potentially conservative. This sits inside Coherent's broader, company-wide $23bn multi-year growth-opportunity framework (CPO $15bn, OCS $4bn, thermal management $2bn, multi-rail transport $2bn) disclosed around OFC 2026. Coherent did not break out OCS revenue as a separate line item on the F4Q26 call beyond noting it "increased sequentially," and demand is broadening beyond scale-out into scale-across and scale-up applications, with production expanding across two manufacturing locations. We hold Coherent's share of merchant OCS revenue at a flat 40% in our model.

4.3 HUBER+SUHNER / Polatis

Polatis (a HUBER+SUHNER brand) is the longest-established all-optical MEMS switch vendor, with a product line spanning 8×8 up to 384×384 ports and a long history in telecom/lab test-and-measurement applications before AI data centers. At ECOC 2026, HUBER+SUHNER announced a new dedicated Polatis OCS production site in Pisary, Poland, citing a "huge surge in demand for OCS solutions" from hyperscale data center operators and AI compute cluster providers.

4.4 Calient Technologies

Calient is the fourth established OCS pioneer (its S-Series switches predate the current AI cycle by more than a decade). In 2026, Suzhou Chunxing Precision Mechanical, a Chinese precision-manufacturing company, acquired a controlling stake in Calient, pairing its manufacturing scale with Calient's optical switching IP; Calient continues to operate under its own brand from Goleta, California. We flag

this as a geopolitical/supply-chain consideration for US hyperscaler customers evaluating Calient specifically, and as a data point supporting the broader observation that OCS technology and manufacturing capacity is attracting capital and consolidation interest well beyond the two US-listed merchant suppliers (COHR, LITE) in our coverage.

4.5 iPronics, Lumotive, and next-generation entrants

iPronics, a silicon-photonics OCS start-up, received a $125M investment in September 2026 from NVIDIA (alongside Maverick Silicon and Light Street Capital), earmarked for a second-generation switch promising sub-millisecond reconfiguration and higher port density than first-generation MEMS. iPronics is also a named co-author of the OCP's April 2026 OCS white paper, alongside Lumentum, Ciena, Lumotive, and Carnegie Mellon University — evidence of an industry-wide standardization push rather than single-vendor lock-in. Lumotive (metasurface/2D photonic beamforming, $45M Series B led by Gates Frontier and MetaVC) is the most credible next-generation entrant we are tracking outside the MEMS incumbents.

4.6 Hyperscalers, NVIDIA, and the EPS incumbents

Google remains the OCS demand anchor and the only platform with fully confirmed, multi-generation production deployment (Palomar/Jupiter spine, TPU v4 through TPU 8t/8i scale-up and scale-out). Microsoft, NVIDIA, and Meta have all joined the Open Compute Project's OCS initiative, and NVIDIA separately backs iPronics directly and co-authored the Columbia "ACTINA" scale-up research — but NVIDIA's own production optical roadmap remains CPO-centric (Quantum-X/Spectrum-X today, Vera Rubin Ultra's CPO-based scale-up next), not OCS-based.

5. Public Company Commentary

All commentary below is drawn from public disclosures — SEC filings, earnings call transcripts, investor conference remarks, and company press releases — each dated and sourced. We have no non-public channel checks or expert-network commentary to include in this note.

5.1 Lumentum (LITE) — F4Q26 earnings call, August 2026

- Total company revenue of $1.01bn for the quarter, ahead of the $988.6mm consensus estimate.
- Management guided to OCS's first triple-digit-million-dollar revenue quarter in F1Q27 ("meaningfully above $100 million").
- Reiterated the prior $400M OCS revenue target for 2H CY26 — management specifically noted the company is on track but not ahead of that target, reflecting earlier supply-chain constraints that have since "returned to plan."
- Management characterized underlying demand as "expanding faster than several capacity ramps," and identified OCS scaling as a standing management priority alongside 1.6T transceiver growth.
- No specific customer names or unit economics were disclosed on this call; a separately reported ~$400M order backlog is attributed to three customers, consistent with management's broader messaging that OCS customer breadth

is "increasingly widening" beyond Google, with some non-Google customers "much more substantial than Google in our early rounds of shipments."

5.2 Coherent (COHR) — F4Q26 earnings call, August 2026

- OCS revenue "increased sequentially" in the quarter (no absolute dollar figure disclosed).
- Management disclosed "more than $4B of addressable opportunity across DCI, scale-out, and scale-up networks" for OCS — doubled from the $2B figure given at OFC in March 2026, with management noting even the higher figure may be conservative.
- Demand is described as broadening beyond scale-out deployments into scale-across and scale-up applications; production is expanding across two manufacturing locations, with the 320×320 platform anchoring the current product lineup.

6. Our TAM Model & Assumptions

Our model estimates a $11B Merchant OCS TAM by 2030. Our bottom-up, XPU-per-port model (focused on Google TPU) reaches ~$5B of merchant OCS TAM by FY2030 from the scale-up layer alone, plus ~$2B from a newly rebuilt, port-driven Spine & Other Related Applications bucket, for a ~$7B FY2030 Combined Total (excluding non-Google OCS reconfiguration). Our team fielded a Lumentum management comment characterizing a reconfiguration opportunity — distinct from the Google-anchored buckets above — as potentially ~10x the size of Google's, on roughly a 1:1 ratio with AI servers, materializing as a "late 2028/2029" story. Citi believes this to be likely tied to NVDA. Thus, adding a third bucket based on NVDA's GPU and NVL server shipments that we estimate can add another ~$4B to the TAM to ultimately reach $11B.

Figure 5. Citi Estimate of OCS TAM

COMBINED TOTAL INCLUDING BUCKET 3 — Merchant GPU Rack-Scale Reconfiguration						
($mm unless noted)	CY2025	CY2026	CY2027	CY2028	CY2029	CY2030
Combined Total OCS TAM (Buckets 1+2 base case)	$283	$853	$3,115	$4,271	$5,888	$7,138
Bucket 3 OCS Reconfiguration (Beyond Google)				$734	$1,808	$3,579
Combined Total OCS TAM incl. Bucket 3	**$283**	**$853**	**$3,115**	**$5,005**	**$7,696**	**$10,717**

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research

6.1 Bucket 1: XPU Scale-Up Build

For each platform and year, the chain is: XTPU unit shipments × optical ports per XPU × OCS attach rate ÷ OCS switch port radix × blended external port ASP = merchant OCS TAM.

Google's OCS attach rate is held at 100% (OCS is in production across Google's TPU fleet); its merchant capture rate ramps from 15% to 65% as Lumentum (potentially and potentially Coherent) take share from Google's historically more captive/co-designed sourcing.

6.2 Bucket 2: Spine & Other Applications Build (rebuilt on a port basis, Sep-2026)

This bucket captures OCS demand outside the XPU scale-up relationship: spine-layer replacement, pooled/disaggregated resources, and campus DCI. We rebuilt it

using a port-focused method, structured identically to the scale-up build above: total XPU unit shipments (all platforms) × an incremental scale-out/spine optical-ports-per-chip assumption × an OCS attach rate × a merchant capture rate × a blended port ASP.

The ports-per-chip anchor is Google's disclosed TPU v7/Ironwood pod-to-pod interconnect ratio: 9,216-chip pods connected via 300-port OCS switches at an estimated ~500-600 chips-per-switch ratio at the scale-out layer, which implies roughly 0.5-0.6 incremental ports/chip for Google alone which we blend down initially due to the likely lower ratio by non-Google deployments but ramp it up over the forecast as adoption broadens, consistent with Lumentum management's of OCS customer breadth as "increasingly widening."

The merchant capture rate for this bucket is set above the XPU scale-up bucket's Google-specific rate because the spine/pooled-resource layer is less captive to in-house/co-designed silicon than intra-pod scale-up — consistent with Lumentum management's comment that some non-Google OCS customers are "much more substantial than Google in our early rounds of shipments" at this layer.

Figure 6. OCS TAM (excluding non-Google Reconfiguration)

COMBINED TOTAL OCS TAM (XPU SCALE-UP + SPINE & OTHER APPLICATIONS)							
	($mm unless noted)	CY2025	CY2026	CY2027	CY2028	CY2029	CY2030
XPU scale-up merchant OCS TAM		$263	$740	$2,496	$3,062	$4,199	$4,811
Spine & Other applications OCS TAM		$20	$86	$322	$803	$1,128	$1,648
DCI and Others		$6	$18	$297	$407	$561	$680
Combined Total OCS TAM		**$283**	**$853**	**$3,115**	**$4,271**	**$5,888**	**$7,138**
YoY growth			201.8%	265.0%	37.1%	37.8%	21.2%

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research

6.3 Bucket 3: Merchant-GPU rack-scale reconfiguration

During this note's preparation, our team fielded a Lumentum management comment characterizing a reconfiguration opportunity — distinct from the Google-anchored buckets above — as potentially ~10x the size of Google's, on roughly a 1:1 ratio with AI servers, materializing as a "late 2028/2029" story. Citi is publishing its analysis of what it would imply.

Working through the unit economics, we think "1:1 with AI servers" most plausibly means 1:1 with an NVL-class rack (NVIDIA's unit of topology — the boundary at which reconfiguration matters, analogous to how Google attaches OCS at the TPU pod/cube level rather than per-chip. As Citi believes this to be likely tied to NVDA and built an estimate based on GPU shipments, NVL server mix, OCS reconfiguration mix of NVL deployment, external merchant capture rate and ASP assumptions.

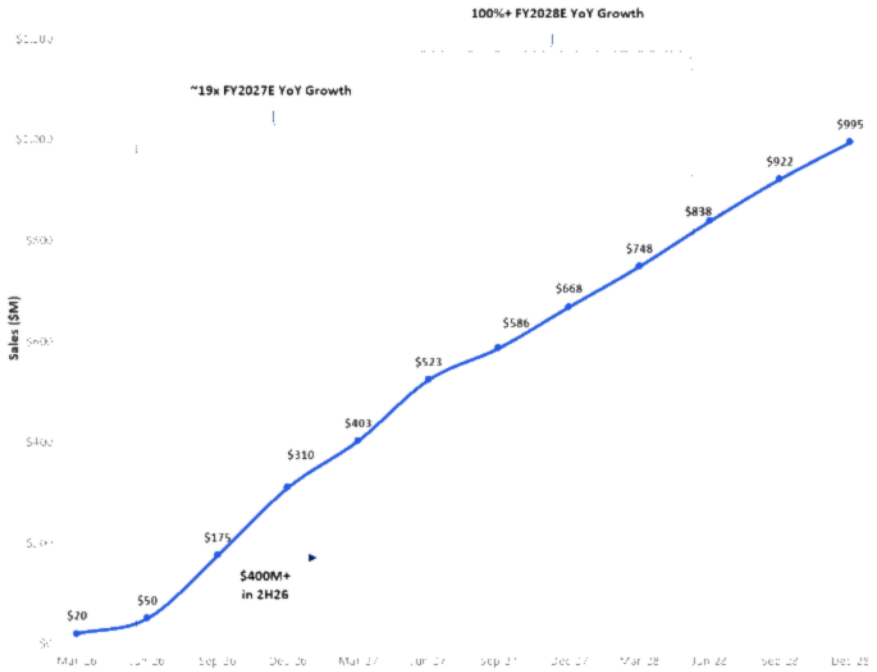
6.3 Output: TAM by platform, year, and combined total

We expect LITE and COHR to be leaders of the merchant market with 80%+ of the shares by 2030 and expect OCS to be a key driver for growth.

7. Model implies Upside to our Prior LITE forecast

Our model estimates LITE to deliver $55M, $2.2B, $3.5B, $5B and ~$7B in OCS sales in C2026, C2027, C2028, C2029 and C2030 sales. This puts us at $485M OCS sales in 2H2026 (20% above the $400M 2H26 prior OCS guide and 15% above prior estimates. As we look at FY2027E and FY2028E we now model $1.4B (vs $1.1B prior) and $2.8B (vs $2.5B prior).

Figure 7. LITE OCS Quarterly Sales Estimates



© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research

As we revise our prior LITE OCS estimates (see section 7), we raise our LITE sales estimates in FY2027E and FY2028E by +5% and +3%. This coupled with management comments on OCS and UHP margins being in the 60s% and 70s% respectively put now our FY2027E and FY2028E EPS estimates at $25 ($23 prior) and ~$41 ($37 prior). We subsequently raise our LITE TP to $1,400 on consistent 35x PE while maintaining COHR's at $420.

Figure 8. LITE FY P&L Changes

Key Financial Data -- Income Statement ($M), except per share	FY25 Jun-25	FY26 Jun-26	Old FY27E Jun-27	New FY27E Jun-27	Old FY28E Jun-28	New FY28E Jun-28
Revenue ($M)	1,645	3,014	6,828	7,155	10,532	10,953
YY	21.0%	83.2% 0.0%	126.5%	137.4%	54.3%	53.1%
Gross Margin - GAAP	28.0%	41.7%	50.5%	52.4%	53.1%	55.8%
Gross Margin (ex-1, ex-amort, inc-SBC)	34.7%	46.0%	52.4%	54.3%	54.3%	57.0%
RD % sales	15.6%	10.1%	6.4%	6.4%	5.0%	5.0%
SG&A % sales	9.4%	6.1%	4.5%	4.5%	4.0%	4.0%
OpEx non GAAP % Sales	25.0%	16.2%	10.9%	10.9%	9.0%	9.0%
R&D non GAAP	256	305	435	454	527	548
SG&A non GAAP	155	185	307	322	421	438
OpEx (GAAP)	675	731	1,120	1,170	1,443	1,499
OpEx (non GAAP)	411	490	742	776	948	986
Operating Income non GAAP	160.1	897.0	2,837.8	3,105.9	4,771.3	5,253.1
Operating Margin non-GAAP	9.7%	29.8%	41.6%	43.4%	45.3%	48.0%
Interest Expense	(22.2)	(21.8)	(25.2)	(25.2)	(25.2)	(25.2)
Other Income	30.2	80.8	90.4	90.4	90.4	90.4
Pretax Income non GAAP	175.3	936.9	2925.8	3193.9	4859.3	5341.1
Tax Rate non GAAP	16.5%	16.5%	18.0%	18.0%	17.0%	17.0%
Share Count	71.2	90.2	104.2	104.2	108.4	108.4
EPS (ex-1, inc-SBC, ex-amort)	$2.06	$8.68	$23.03	$25.14	$37.21	$40.90
Consensus Sales				$6,241		$9,597
Consensus EPS non GAAP				$17.97		$32.14

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

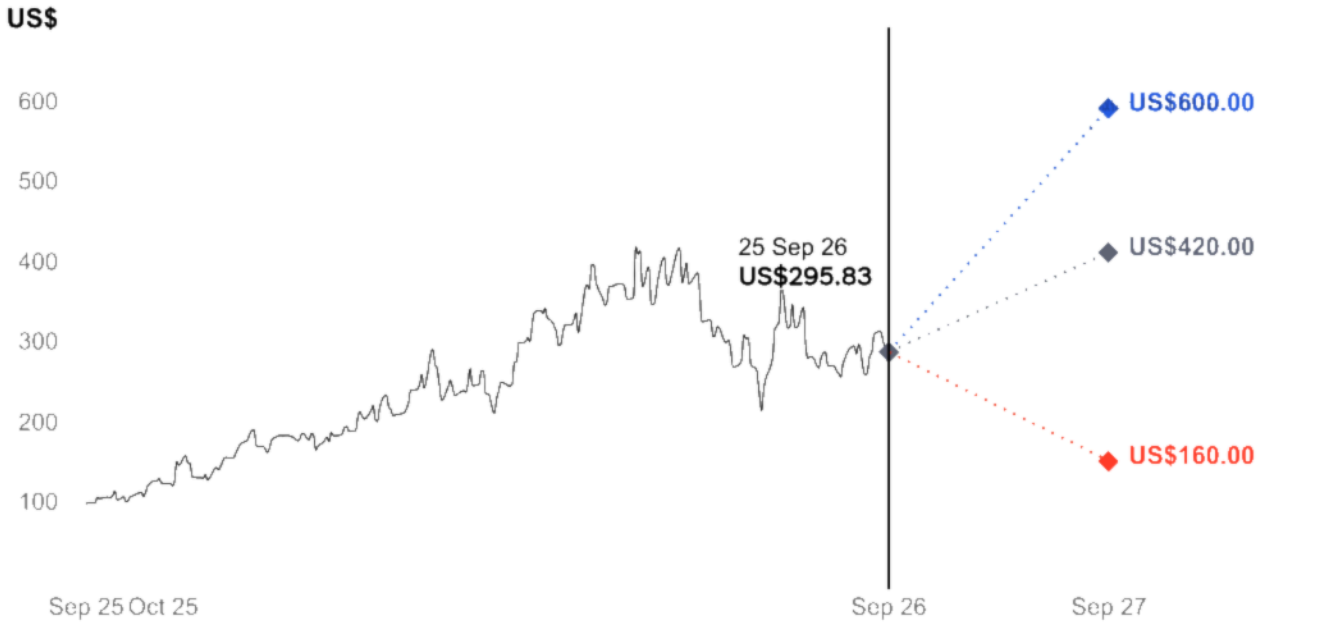
Source: Citi Research, Company Reports, FactSet

Bull/Bear: Coherent Corp (COHR.K)

US$ 600.00
▲103% Upside

US$ 420.00
▲42% Upside

US$ 160.00
▼46% Downside



Spread 149pp
Current Price and expected returns (upside/downside) as of 25 Sep 2026

BULL Assumptions

- Sustained margin improvement drives EPS accretion backed by robust AI/SiC ramp.
- Faster-than-expected ESP growth amid new management's effort to drive margins.
- 35x PE

BASE Assumptions

- Halo effect on AI helps expand multiple.
- Strong execution by new management to drive margins and EPS expansion.

BEAR Assumptions

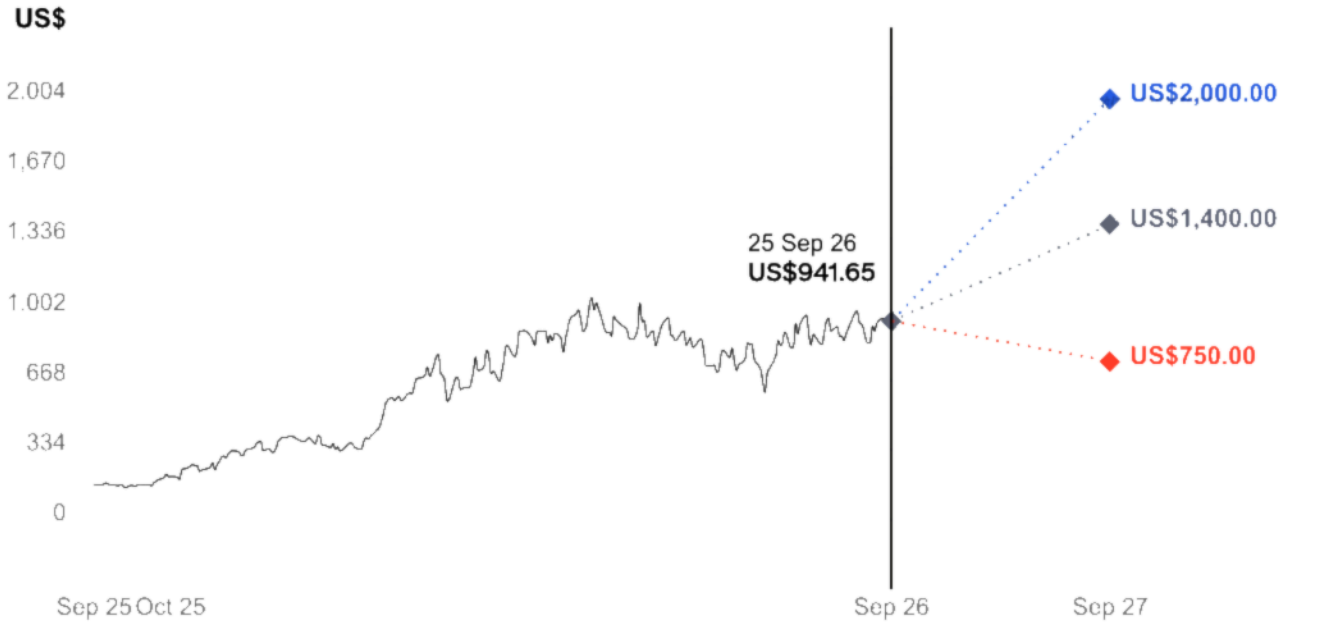
- Inability to achieve margin targets drives EPS lower.
- Frontloaded AI demand.
- 25x PE

Bull/Bear: Lumentum Holdings Inc (LITE.O)

US$2,000.00
▲112% Upside

US$1,400.00
▲49% Upside

US$750.00
▼20% Downside



Spread 133pp
Current Price and expected returns (upside/downside) as of 25 Sep 2026

BULL Assumptions

- Sustained AI sales drives sales and EPS faster-than-previously expected.
- Faster-than-expected ESP growth amid new management's effort to drive margins.
- 40x PE with $50 FY2028 EPS.

BASE Assumptions

- Halo effect on AI helps expand multiple.
- Strong execution by new management to drive margins and EPS expansion.

BEAR Assumptions

- Inability to achieve margin targets drives EPS lower.
- Frontloaded AI demand.
- 30x PE with ~$25 FY2028E EPS.

人在草木间

Coherent Corp

Company description

COHR is a vertically integrated manufacturing company that develops engineered materials and opto-electronic components for industrial materials processing, communications, aerospace and defense, consumer electronics, semiconductor capital equipment, life sciences and automotive end markets. COHR produces a wide variety of application-specific photonic and electronic materials and components, and deploys them in various forms, including integration with advanced software. COHR's customer base includes original equipment manufacturers, laser end users, system integrators of high-power lasers, manufacturers of equipment and devices for industrial, optical communications, consumer electronics, security and monitoring applications, U.S. government prime contractors, and various U.S. government agencies. COHR is a headquartered in Pennsylvania.

Investment strategy

We are Buy-rated on COHR. We believe that the stock's multiple will benefit from the AI appeal while its fundamentals should expand on new management's strong execution and focus on profitability.

Valuation

We value COHR at $420 or 30x forward P/E applied to FY2028 EPS of ~$16 discounted back. Our COHR valuation is nevertheless at a 10% discount to the average forward P/E of AI optics peers to reflect a lower AI sales mix and 3-year EPS CAGR vs peers such as LITE.

Risks

Downside risks to the achievement of our target price include: 1) increased competition on AI and SiC leads to limited market share growth and limited revenue upside, 2) inability to achieve synergies and margin targets established post acquisition of Finisar, and 3) prolonged economic recession leads to slower 5G build-out and lower consumer demand driving revenue of COHR down.

Upside risks to the achievement of our target price include: 1) gaining share in AI and SiC market through new product offerings and first to market status on larger wafer sizes, 2) quicker and stronger adoption of AI workloads, electric and autonomous vehicles, and 3) quicker return of margins to pre-Finisar acquisition levels.

Lumentum Holdings Inc

Company description

Lumentum Holdings Inc. is a leader in optical and photonic market which delivers essential products to various end-markets including cloud, AI/ML, ns, consumer, and industrial. LITE specifically focuses in

two end-market focused reportable segments, Cloud & Networking and Industrial Tech.

Investment strategy

We are Buy-rated on LITE. We believe that the stock's multiple will benefit from the AI appeal while its fundamentals should expand on new management's strong execution and focus on profitability.

Valuation

Our 35x P/E multiple is a ~10% premium vs the average of AI optics peers' valuation. We model LITE to grow its EPS between 2025-2028E 2-3x faster. Therefore, on an earnings growth adjusted basis, we think a higher P/E multiple vs peers is justified. Our TP on LITE is $1400.

Risks

Downside risks to the achievement of our target price include: 1) increased competition on AI leads to limited market share growth and limited revenue upside; 2) Frontloaded or slower-than-expected AI ramp affecting EML ramp; and 3) Slower than expected adoption of LITE CPO and OCS new products.

Upside risks to the achievement of our target price include: 1) gaining share in AI market through new products on the AI transceivers, CPO and OCS fronts; 2) quicker and stronger adoption of AI workloads; and 3) faster-than-expected margin expansion.