Hardware & Networking

OCP Insights: Takeaways from Meetings with 650 Group, KEYS, TEL, CRDO and JBL

IT Hardware/ Telecom & Networking Equipment

Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

Manmohanpreet Singh [AC]
(1-212) 622-4527
manmohanpreet.singh@jpmchase.com

Marc Vitenzon
(1-212) 622-3342
marc.vitenzon@jpmchase.com
J.P. Morgan Securities LLC

At OCP 2026 in San Jose, California, we hosted meetings with 650 Group, Keysight, TE Connectivity, Credo, and Jabil that reinforced the view that AI infrastructure is shifting from single-building scale-out towards multi-building campuses and more heterogeneous scale-up architectures. This shift is driving significant incremental demand across technologies, such as coherent-lite, OCS, new pluggable optics, CPO/NPO, copper interconnect, and memory disaggregation. Importantly, the faster technology cadence in combination with the complexity and broadening aperture of technologies being used across more distinct deployments are also expanding the test, validation, and systems-integration opportunity. Across the coverage, companies with broad portfolios and execution depth appear best positioned as the AI data center buildout continues to evolve heading into the latter part of the decade. *Takeaways from individual meetings include*:

650 Group

We hosted an investor meeting with Alan Weckel and Chris DePuy, co-founders of the 650 Group. Key takeaways are below:

- **Multi-building AI DC campuses are driving demand for scale-across, including coherent lite products.** As AI clusters extend beyond a single building, hyperscalers are building campuses with 1–10 km links between buildings, e.g., Meta's keynote showed such a campus with trenches carrying ~1,000 fibers per conduit, most of which are expected to connect via coherent lite. Such deployments support a coherent lite forecast of ~20 mn units in the out-years, with millions of units expected in 2027 and a volume ramp in 2028, and is expected to benefit existing ZR suppliers. On the networking side, Arista is benefiting from scale-across as the DCI spine for some data center deployments.
- **OCS is expected to broaden in terms of customers and use cases by the end of the decade.** While current OCS deployments are limited primarily to Google, more than half of OCS revenues are expected to come from non-Google customers by 2030. Expected new use cases include: (1) scale-across - dynamically allocating bandwidth between buildings on a campus, which is a near-term application with immediate payback; and (2) scale-up (sub-50m) - provisioning additional racks of GPU, CPU and memory into a cluster on demand, e.g., for large-context frontier workloads. The OCS market is expected to reach ~$12 bn by 2030, with Lumentum and Coherent viewed as the primary beneficiaries.
- **Robust ramp expected in global laser shipments, with different laser types scaling across different use cases.** Laser shipments are expected to reach ~220–240 mn units in 2026 across transceivers, ELSFPs and CPO, and then expand to multiple billions by 2030 and eventually >10 bn. At that scale, a single laser type cannot serve every application, so the industry is likely to move toward purpose-built lasers matched to distance and use case (e.g., VCSELs for scale-up, microLEDs for slow-and-wide memory pooling applications). In the near to medium term, high-power CW lasers are expected to win in CPO/ELSFP, given limited implementation risk when supporting a

See page 7 for analyst certification and important disclosures.

$5-10 mn rack.

- **Celestica and Arista are ceding share at Meta, but TAM growth leaves revenue growth intact.** Meta is now among the largest networking buyers, and it is deliberately multi-sourcing at both the silicon and system level. While Celestica and Arista have historically accounted for 90%+ of Meta's switching volume, their share is being diluted as other vendors, including Cisco, Nvidia (primarily via Accton), and others gain volume. However, with Ethernet switching growing 40%+ annually, share loss no longer means lost business: incumbents are being augmented rather than displaced, and share loss in % terms does not imply any growth concerns for Celestica or Arista.
- **DSP supplier base is widening at 1.6T; Credo positioned to gain despite early ZeroFlap ramp delays.** With demand running well ahead of forecasts, 1.6T is set to have more DSP suppliers than any prior generation. Share shifts are expected to come through two channels: (1) supply shortages, where hyperscalers that under-ordered go to whoever can deliver first; and (2) NeoCloud and Oracle, which lack hyperscaler-grade telemetry and stand to benefit most from Credo's ZF optics; positioning Credo well on both fronts. The ZF ramp delays are attributed to typically poor early DSP yields and less experience in this market, neither of which is seen as a structural concern.
- **Memory disaggregation is opening up a large, underappreciated networking market.** As context windows grow, every hyperscaler is expected to deploy its own custom memory appliance, essentially a KV-cache system optimized for its workloads, as a source of cost or performance advantage. Pairing a custom memory controller with a custom network switch lets hyperscalers optimize the full memory stack, potentially delivering equivalent performance with cheaper memory, e.g., DDR4 instead of DDR5, lower-end HBM stacks, or flash in place of DRAM. Memory networking is seen as at least a $10 bn market by the end of the decade. Beyond Credo (Weaver Gearbox), there are few established vendors openly addressing the opportunity, while most activity is coming from startups.
- **Hyperscalers will run a diverse fleet of accelerators; Maia is less likely to reach meaningful volumes for Microsoft.** Over time, each hyperscaler is expected to operate a dozen or more accelerators, each optimized for a specific algorithm, energy profile, workload or sovereign requirement. On that basis, a custom, energy-efficient Copilot processor makes strategic sense for Microsoft, however Maia 300 is not expected to reach any meaningful volumes.

Keysight

We hosted an investor meeting with Kailash Narayanan (President, Communications Solutions Group) and Ram Periakaruppan (VP and GM, Wireline Network and Security). Key takeaways are below:

- **AI infrastructure is driving broad-based testing demand.** Keysight highlighted several structural drivers supporting its AI infrastructure testing business, including accelerating network-speed transitions, proliferation of networking protocols and architectures, and increasing validation requirements across components, systems, racks, and data centers. Importantly, management noted that components and systems are increasingly being tested multiple times throughout the development cycle, with hyperscalers independently validating equipment performance, interoperability, and workload characteristics even after suppliers complete their own qualifications. This creates incremental hardware and software opportunities beyond traditional component-level testing, with demand driven by the number of development programs, engineering teams, and technologies under evaluation.
- **Testing and early engagements are occuring across 800G, 1.6T and 3.2T in tandem.** Management highlighted accelerating network-speed transitions and increasingly overlapping development cycles, with 800G in volume manufacturing, 1.6T entering

initial manufacturing alongside ongoing R&D, and 3.2T already generating early customer engagements. Keysight claimed to be capturing the lion's share of the early 1.6T testing business, while also noting that initial 3.2T testing solutions could emerge in 2H27 or early 2028, depending on the availability of supporting SerDes and silicon technologies. Importantly, Keysight is already participating in early physical-layer characterization ahead of complete protocol-testing solutions.

- **Networking scale-across and scale-out opportunities on display.** Keysight demonstrated 800G traffic emulation alongside Cisco's P200 data center solution, simulating DCI workloads to evaluate congestion, latency, throughput, and overall performance. Additionally, Keysight demonstrated eight AresONE systems emulating 64 endpoints at 800G each to validate a Cisco Nexus switch platform against MoE workload patterns. Management emphasized that similar capabilities are required across networking OEMs as well as hyperscalers evaluating vendor performance, published specifications, and interoperability. These systems are primarily deployed in R&D environments, with separate testing equipment typically used during manufacturing.
- **Higher switch radix drives demand for denser 1.6T traffic emulation.** Keysight introduced a new 1.6T traffic emulator supporting eight 1.6T ports, doubling the four-port density of its previous generation. Management emphasized that increasing switch radix, including transitions from 512 to 1,024 ports, requires proportionally greater testing capacity, with customers ideally validating every port under simultaneous traffic conditions. Higher-density systems allow customers to scale test environments more efficiently, while potentially increasing value per platform. Management noted that new technology generations generally command higher ASPs given increasing component costs and technical complexity.
- **Ethernet scale-up is starting to gain traction.** Management highlighted the emergence of competing scale-up architectures, including Ethernet-based solutions, UALink, and Nvidia's proprietary NVLink Fusion, creating incremental testing requirements across protocols and hardware implementations. Keysight demonstrated Ethernet scale-up testing using Broadcom Tomahawk 6 and Tomahawk Ultra switching platforms, including a system associated with Alibaba. The company currently supports UALink testing but does not support NVLink Fusion protocol emulation, while management expects the broader UALink ecosystem to expand as switch, endpoint, and IP suppliers introduce new products. Separately, Keysight demonstrated multipath networking with HPE, simulating congestion, packet timing, and failover across multiple switching planes at speeds ranging from 400G to 1.6T.
- **CPO testing opportunities build ahead of manufacturing maturity.** Management highlighted ongoing CPO and silicon photonics R&D engagements with leading industry participants, including a major GPU supplier, while suggesting that CPO manufacturing remains three to five years away from achieving greater maturity given challenges around production stability, quality, and cost. Keysight expects copper and optical interconnect to coexist for an extended period, citing copper's established supply chain, cost advantages, and technological maturity. Management referenced an illustrative 80% copper / 20% optical mix over time. Importantly, Keysight expects to benefit from both technologies, with CPO already generating R&D testing opportunities ahead of broader manufacturing adoption.
- **AI inference workload emulation targets time to first token, latency, and cost per token.** Keysight demonstrated inference workload testing, allowing customers to evaluate networking and infrastructure configurations before physically deploying data centers. The solution also supports physical inference clusters utilizing various accelerator architectures, measuring performance metrics including time to first token, latency, and cost per token to assess infrastructure efficiency and return on investment. Importantly, Keysight can ingest Chakra workload traces to replicate customer-specific communication and compute patterns without accessing proprietary underlying data.

While primarily targeted at R&D and infrastructure design today, management identified potential production applications involving troubleshooting and root-cause analysis.

- **AI infrastructure complexity expands test content opportunity**. Management emphasized that Keysight's content opportunity is increasingly determined by protocol proliferation, port density, network speed, rack architecture, interconnect technology, and the expansion of scale-up beyond an individual rack. Hyperscaler-specific system architectures and proprietary workload characteristics introduce additional software and hardware testing requirements, while newer speed generations generally support higher initial ASPs before moderating as technologies mature. Importantly, management cautioned against evaluating Keysight's equipment pricing using conventional bandwidth metrics, given its focus on specialized R&D applications rather than high-volume networking deployments.

TE Connectivity

We hosted a booth tour and investor meeting with Vishwas Rao, SVP & GM of Digital Data Networks (DDN). Key takeaways are below:

- **Collaboration on USFP optics important in terms of showcasing content in optical as well as partnership with Nvidia.** The company is providing the critical high-speed connectivity for NVIDIA and Cisco's new Ultra-Dense Small Form Factor Pluggable (USFP) architecture (see press release here), including MegaLane (connector) technology to support 400G per lane, with a path to 12.8T of aggregate capacity in the OSFP sized footprint, along with cages and modules. The new USFP optics will deliver 4x bandwidth density and ~30% better power efficiency, relative to 200G per lane optics. *Note: Press release also names Amphenol as a partner.*
- **Showcased a 448G CPC demo; Copper expected to co-exist, with notable preference in CPU based architectures.** Copper and optics are expected to co-exist, with CPU-heavy agentic and inference architectures being even more copper-friendly. As a continued pace of innovation on copper interconnects, TE demonstrated a full 448G co-packaged copper (CPC) to backplane channel of about 1m, with the channel staying within the ~40dB loss budget that customers typically design to. The backplane connector is dense enough that some customers are evaluating it as a near-chip connector instead.
- **ICA sockets help deliver serviceability in Near Packaged solutions.** Signal path starts at the ICA socket, which sits on the substrate (next to the main silicon) in a small footprint and fans out to a replaceable outer board carrying NPO/NPC ports. Designs of the ICA socket helps support both NPO and NPC solutions, by varying the ports/connectors that sit on the fanned out edge, and also help deliver serviceability of those copper/optical ports. A failed optical or copper channel can be fixed by swapping a low-cost board instead of scrapping the chip or tray, and copper and optical ports can be mixed on the same board.
- **Breadth of offerings helps support signal paths from the chip to the backplane.** As the optical signal exits the optical engines (*that are mounted on ports of the ICA socket*), TE can help support the passive optical length of the signal via key connection points like - one on each side of the fiber shuffle, and then the optical backplane. On the chip end, the RAM Photonics acquisition adds a fiber array unit (FAU) platform that actively aligns each fiber into a lens array and laser-welds it to hold sub-micron precision, that can help deliver connections like fiber attach to CPO/NPO chips, or link backplane cartridges.
- **Showcased liquid-cooled cable solutions for higher power racks; and smart busbars with built-in telemetry.** As racks shift to 800V and dedicated power sidecars, TE is developing liquid-cooled power cable targeting ~1,500A per cable. At that rating, fewer cables could support a 1MW rack, driving significant reduction in cables relative to air-cooled variants. Alongside this, the company showcased smart busbar that embeds

current and thermal sensors in each connector, delivering telemetry into busbars, to catch overcurrent or thermal events before they cause damage.

Credo

We hosted an investor meeting with Bill Brennan (CEO). Key takeaways are below:

- **Broadening the pluggable opportunity beyond AECs.** Management emphasized its strategy of addressing the broader pluggable connectivity market, spanning 1m to 2km, through a portfolio of AECs, ZF Optics, and ALCs, rather than viewing each as a separate opportunity. While AECs offer the lowest cost and shortest reach, ZF Optics provides the longest-distance at the highest cost, and ALCs extend reach to 50m at above AEC cost (*similar to conventional optics*), with all targeting comparable reliability. Importantly, management expects ALCs to displace conventional optics at distances below 50m and potentially cannibalize AECs, with the power consumption arguably tracking lower than even passive copper (*when including high-power SerDes*); although, it views that substitution favorably given Credo's ability to capture a broader portion of spending.
- **Highlights ALC engagements with Microsoft and Oracle.** Management highlighted Microsoft as a participant in its ALC initiative and indicated Oracle intends to join the MSA, with the latter particularly interested in addressing reliability and maintenance challenges. Initial ALC applications are expected to include NIC-to-TOR and potentially switch-to-switch, with shipments anticipated in 2027.
- **Micro-VSCEL simplifies ALC manufacturing and provides path to 3.2T.** The transition from microLED to micro-VCSEL was driven by increasing requirements for 50m reach, compared with the previous 10m target. The shift increases optical lane speeds from 2Gb/s to 26Gb/s, reducing the number of fibers required for a 1.6T ALC from 2.5k to 200, including 15% allocated to redundancy. Importantly, increasing optical lanes to 50Gb/s was suggested to provide a relatively straightforward path toward 3.2T.
- **Scale-up represents another significant growth opportunity.** Scale-up was identified as another significant growth opportunity, with connectivity volumes tracking 8-10x higher than the pluggable market. Management highlighted NPO, CPO, and optical backplane opportunities, with micro-VCSELs targeting distances below 50m and traditional laser-based optics targeting longer distances. On that note, Credo specifically highlighted an optical backplane incorporating 1mn fibers and supporting up to 30Pb/s of aggregate bandwidth, or 30x Nvidia's current capability.
- **OmniConnect offers potentially signficant incremental content per accelerator opportuntiy.** With Positron production expected to begin in roughly one year and customer plans contemplating a ramp toward 100k accelerators per quarter, Credo highlighted OmniConnect as a sizable opportunity. Specifically, OmniConnect content is expected to track to $3k per accelerator (vs. ~$1k suggested prior), which, combined with $1k from two associated AECs per accelerator, translates into an annual revenue opportunity of $1.6bn. Beyond Positron, management disclosed multiple signed OmniConnect customers; although, fewer than five, with additional advanced engagements underway and plans to announce further customers once it reaches five signed contracts. Credo also announced engagements with switch manufacturers.
- **Nvidia qualification removes key barrier to ZF Optics adoption.** Management disclosed that ZF Optics passed Nvidia qualification, potentially opening a broader opportunity with Nvidia-based neocloud customers. Previously, some customers had been reluctant to deploy ZF Optics because non-qualified components could jeopardize support arrangements, creating a meaningful barrier to adoption despite customer interest. Qualification now establishes a clearer path toward deployments within the Nvidia ecosystem.

Jabil

Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

J.P.Morgan

We hosted an investor meeting with Matt Crowley (EVP, Intelligent Infrastructure), Ed Bailey (CTO/SVP, Intelligent Infrastructure) and Adam Berry (IR). Key takeaways are below:

- **Joint venture with Adani brings together complementary capabilities and creates value.** In relation to the JV, Adani brings to the table access to ports, local government relationships, and substantial power (up to ~8GW). Jabil, meanwhile, brings expertise in manufacturing complex AI racks at scale, compute, storage and networking manufacturing, and data center infrastructure content including chillers and low- and medium-voltage switchgear. Jabil also brings customer relationships, and would likely bring the JV's anchor customer, possibly one of its existing hyperscalers. Management sees a potential IPO of the JV in India as one of the end goals, given favorable local market conditions.
- **Positioned for strong growth in Trainium ecosystem.** Jabil's East Coast retrofits were planned for longer-term demand around next-gen liquid-cooled Trainium racks, and continue to expect robust demand from that business. Strong current position is due to the engineering and architecture team built for this customer seven years ago, which now supports 40–50 k racks a year. Management is not concerned about any share shift from current PCBA manufacturers, Accton and Fabrinet, potentially looking to build racks.
- **Boasts broad range of capabilities across optical; business likely to attract capex near term.** Intel transceiver acquisition was to expand capability across co-packaged optics (CPO), near-packaged optics (NPO), XPO and co-packaged copper, not just selling pluggables. That flexibility lets Jabil pair optical engines with different switch silicon inside the rack. Management sees NPO as closer to real scale than CPO, noting there are many more customer conversations on NPO and only one for CPO. Importantly, optics is a part of the business that might attract more capex in the near future, however, corporate capex is still expected to track within 1.5-2.0% of revenues.
- **Networking capabilities include rack level solutions as well.** Networking is expected to ramp from ~$2.0 bn in FY26 to ~$2.9 bn in FY27. It is shifting away from the top-of-rack box toward a more complex full rack-level system that integrates fiber shuffle boxes, fiber, liquid manifolds, cold plates, firmware and systems management. Jabil was among the first to work with switch OEMs on these designs, given incumbent relationship around building switches and experience shipping ~50 k racks a year. Networking relationships are not just limited to Ethernet OEM and span the silicon providers as well.
- **Importance of execution increases as the narrative shifts from "what" gets built to "how" it gets built and "when" it gets deployed.** Demand continues to be robust but the timing of how systems come together is now the main variable, which favors Jabil, given its ability to engineer, pre-manufacture and pre-fabricate in the factory environment to speed actual deployments. Additionally, Jabil's reliable execution is also important in the environment as the number of companies building AI infrastructure has grown much faster than the pool of engineers and subject matter experts, which is translating to hyperscale customers witnessing "B-team" experiences from some large ODM competitors.
- **Data center demand is not shifting away from US; International builds driven by funding unlocks and in-region inferencing.** India JV was strategic and not driven by U.S. power shortages. Customers are adding capacity outside the U.S. because they need capacity everywhere, and the overall demand pie is continuing to grow with no dramatic mix shift away from the U.S. International markets are on a conceptually similar trajectory to the U.S. but move more slowly given tougher regulation. Funding is now being unlocked in certain countries, and those builds are typically smaller than in the U.S. but still larger than traditional. Management expects the trend to continue as more inference is deployed and used in-country.

