

Tech Conf. Takeaways: AI broadening, supported by order/supply visibility

08 June 2026

Industry Overview

Takeaways from 2026 BofA Global Tech Conference

We recap takeaways from our Tech Conference, where 37 global semiconductor companies attended including in compute/AI (NVDA, AMD, INTC, ARM), connectivity (MRVL, CRDO, ALAB, MTSI, COHR, LITE), analog (ADI, TXN, MCHP, ALGM), EDA (CDNS), semicap/complexity (LRCX, KLAC, AMAT, NVMI, CAMT, ACLS, MKSI, AEIS), and consumer (QCOM, AMBA, AMBQ) markets. See inside for summary of each company discussion.

AI Compute: CPUs additive to XPU, n-t supplies secured

Overall tone was very positive across both CPUs and XPU, with every major CPU vendor confident in growing strongly on agentic AI deployments. CPU TAM estimates vary between \$100-200bn for ~CY30, but we suspect the differences come in ASP differences (varies \$2-2.5k to \$4-5k), rather than units (likely ~20-25% CAGR). Supplies remain tight across the board, but each CPU/XPU vendor has secured sufficient capacity (memory, wafers, substrates, packaging, etc.) for its forecasted growth, based on strong order visibility and customer commitments (i.e. NVDA \$124bn) for the next 1-2 years.

Connectivity: heterogeneous architecture of copper/optics

Outlook for pluggable optics in scale-out remains robust, with the current 1.6T gen ramping with high DSP, TIA, driver, etc. content. For scale -up, AVGO sees continued use of copper through 200G/400G SerDes generation, while others are generally on track to move to optics-based NPO in 2027-28 before transitioning to CPO in 2028+. Meanwhile, the current EML supply-demand imbalance is expected to remain over the next few years despite capacity ramp ups, suppliers have order visibility through LTAs (2-5 yrs).

Semicap/EDA: cleanroom constraints extending WFE cycle

Semicap tone was very positive, with 2026 WFE tracking potentially above ~\$140bn+ and another strong year likely in 2027 as cleanroom constraints push demand forward. AI-led logic, DRAM/HBM and advanced packaging are driving mix toward higher etch/dep, process control and metrology intensity, while NAND remains later-cycle upside. China/trailing-edge looks steady rather than accelerating, and margin upside is increasingly tied to pricing/value capture, mix and new product ramps. In EDA, CDNS sees agentic AI tools as TAM expanders amid semi R&D reacceleration.

Analog: recovery broadening, AI upside emerging

Analog tone was more constructive, with recovery broadening into industrial/mass-market demand while auto remains more content-led than SAAR-led. Data center is becoming a more measurable growth vector across power, sensing, thermal, PCIe/CXL and 800V, with exposures ranging from L-MSD to ~20% of sales today and content growing to over \$100k per rack. Pricing has turned more supportive to offset inflation, while margin upside should come from utilization, mix and higher-value new products.

Consumer: edge AI accelerating while handsets still mixed

AI continues to move to the edge with purpose-built silicon and ultra-low-power edge AI MCUs. Key constraints remain power, software, and rising memory prices. QCOM sees mixed handset outlook, vs. improving AI product sentiment ahead of 6/24 Investor Day.

AI Compute

NVIDIA (NVDA)

Following 'GTC Taipei' CEO keynote on Monday (see takeaway [here](#)), we hosted NVDA CFO Colette Kress for a well-attended keynote on Thursday, and Sr. Director, IR Stewart Stecker for a bus tour investor meeting.

Rising content, increasing opp'ty, solid supply visibility

Key takeaways: 1) Vera CPU opp'ty is roughly split 50/50 between head node (compute rack) and agentic AI (standalone CPU rack), 2) NVDA's \$200bn CPU TAM only includes agentic AI, RL-type workloads, not general purpose traditional CPU workloads, 3) NVDA content to rise meaningfully every system gen, at ~\$40bn/GW for Blackwell, ~\$60 - 80bn/GW for Rubin/Rubin Ultra, and potentially ~\$100bn/GW for Feynman, as NVDA consistently adds new chips/racks to the system; 4) NVDA's full-stack architecture and ecosystem incumbency provides advantages in both hyperscale and ACIE; 5) GPU useful life may be longer than depreciable life, as 3-4 year old Hoppers are still fully utilized; and 6) NVDA's \$124bn purchase commitment suggests robust demand/supply visibility ahead for AI infra. Maintain Buy, \$350 PO, and a top AI pick, given attractive 16x CY27 PE or 0.4x PEG vs. Mag-7 avg of 1.9x.

Vera CPU: head node vs. agentic AI likely 50/50 opp'ty

NVDA clarified the ~\$20bn in Vera CPU sales expected in F2H27 is roughly split evenly between head node sockets (traditional GPU support role in NVL racks) and standalone agentic AI sockets (reinforcement learning, orchestration roles in CPU racks). NVDA has shipped 2.5mn Grace CPUs to date, and the follow-on Vera shipment could already reach ~4-5mn units in just the first two quarters of launch, at an estimated ASP of \$4-5k per CPU. NVDA also noted that the Vera builds on Grace's existing system-level software and optimization on NVDA tech stack, a key differentiation vs. first-gen Grace. Overall, the total CPU unit demand could potentially be higher than GPU over time, reinforcing the CPU-GPU ratio argument of 1-to-1 across both racks (vs. 1-to-2 and 1-to-4 prior).

NVDA content \$/GW rising materially gen-over-gen

NVDA believes every new generation of its systems can increase the company's addressable market materially, from the current ~\$40bn/GW for Blackwell Ultra, toward \$60-80bn/GW for Vera Rubin and Rubin Ultra, and toward ~\$100bn/GW for Feynman maybe. While exact figures may vary, NVDA continues to address more parts of the AI system each generation, with the current Vera Rubin solution comprised of 7 different chips across 5 different racks (vs. Blackwell just 1 main compute rack), and Rubin Ultra (2027) expected to add another rack on top. Beginning the Feynman generation (2028), NVDA also plans to offer optional optical scale-up (via CPO) for customers wanting a larger NVLink domain (vs. regular copper up to 576-die or 144 GPU packages).

"King of diversity": ACIE as important as hyperscale/cloud

NVDA's sales mix between ACIE (AI Clouds, Industrial, Enterprise) and hyperscale is ~50/50 today, with ACIE likely the faster growing segment over time. NVDA is the "king of diversity", as it supplies to every hyperscale, every major model maker, and has significant incumbency and moat in ACIE-type workloads (we view custom XPU's have limited opportunities here). NVDA's continuous upgrades to system software (i.e. Hopper/Blackwell) and new model rollouts are also key competitive advantages.

Advanced Micro Devices (AMD)

Agentic AI driving CPU upside and TAM expansion

We hosted CFO Jean Hu and Head of IR Matt Ramsay, who highlighted the inflection of agentic AI as the key change, with workloads shifting toward orchestration-heavy, multi-step inference requiring incrementally higher CPU performance, driving demand across CPU, GPU, and accelerators.

AMD reiterated a three-part CPU TAM (enterprise, head node, agentic), with agentic AI the fastest-growing and supporting a TAM of >\$120bn as adoption accelerates.

Agentic racks largest opportunity; mix shift favorable

Mgmt sees agentic AI server racks (>50% of TAM over time) as the primary growth driver, with increasing need for high core/thread count CPUs and rising CPU:GPU ratios supporting head node growth.

Positioning remains performance/TCO-driven (not x86 vs. ARM), with strength in both high-core agentic workloads and high-performance head nodes.

Supply tight; growth volume-led with ASP upside over time

Server CPU share ~46%, with demand exceeding supply given tight leading-edge capacity; Q2 growth remains primarily unit-driven (~2/3), though ASP upside is expected as mix shifts toward more complex, higher-performance SKUs.

MI450 ramp on track; hyperscale demand tracking above plan

AMD is on track for MI450 ramp (Q3 start, Q4/Q1 acceleration), with early system deployments underway.

Hyperscale demand (incl. OpenAI, Meta) is tracking above prior 2027 expectations, with additional multi-GW opportunities emerging. Future engagements do not need to give warrants out as part of deployment deals.

Intel (INTC)

Execution focus remains central as CPU/AI tailwinds build

We hosted CFO Dave Zinsner and IR John Pitzer, who emphasized that while CPU demand is inflecting sharply on AI (inference → agentic → RL), Intel's primary issue remains execution, with recent efforts focused on cultural change, flattening org structure, and improving operational discipline.

Mgmt views CPU demand as structurally strong, with both total capex and CPU mix increasing at hyperscalers, supporting multi-year growth if supply can be delivered.

CPU demand strong but supply-driven share dynamics near-term

Near-term, mgmt characterized the CPU market as supply-constrained ("if you can make it, you can sell it"), with share largely dictated by capacity ramp rather than product.

Intel expects meaningful supply growth across Intel 3 / 18A over the next ~12 quarters, while also increasing wafer starts near -term (incl. Intel 7), positioning for improved competitiveness as supply normalizes.

On pricing, ASP upside is driven more by mix (core count growth) than shortages, with increasing use of LTAs to lock in pricing/volumes vs. relying on spot dynamics.

Competitive positioning mixed near-term, roadmap improving

Mgmt acknowledged mixed competitiveness in data center CPUs (stronger in single-thread, weaker in multi-thread today), with Granite improving vs. prior gen, Diamond stepping back, and Coral Rapids expected to re-accelerate competitiveness.

Competition remains TCO-driven across x86, ARM, and ASIC-based solutions, with hyperscalers primarily optimizing for performance/cost rather than architecture.

AI ecosystem strategy: partnership + ASIC model + platform breadth

Intel framed its role in AI as broader than merchant CPUs, including partnership with NVDA (client + server), internal accelerators, and increasing focus on ASIC-style engagements (co-design + manufacture) similar to industry models.

Over time, mgmt expects CPU demand expansion beyond servers into client/edge/physical AI, supporting a broader CPU TAM tied to agentic and edge deployment.

Foundry progress improving; financial model still investment-heavy

On foundry, mgmt highlighted improving 18A yields (ahead of plan) and a clearer path toward margin expansion and breakeven exiting 2027, with 14A tracking ahead of 18A at similar stages.

Financially, Intel reiterated a multi-year path to improved growth and profitability (“Rule of 45”), though near-term remains cash-intensive given foundry build-out, with FCF improving but contingent on continued execution.

Arm (ARM)

Near-term royalty growth muted; mix headwinds despite strong DC momentum

We hosted IR Alexis Waadt, who indicated FY27 royalty growth (~20% YoY) is impacted by mix, with smartphone weakness (low-end units, Samsung insourcing) and more limited incremental v9 uplift near-term.

In data center, growth remains strong, though early hyperscale ramps (AWS, NVDA, AAPL) carry lower royalty rates, creating near-term mix headwinds despite strong volume.

v9/CSS driving long-term uplift; data center content scaling meaningfully

Arm highlighted continued v9 royalty expansion (mid-4% → 5%+) and CSS progression (8–9% → 10%+), supporting structural uplift over time.

Content per server CPU continues to scale, with ~\$0.50/core (v9), ~\$1/core (CSS), and \$1.50-\$2 longer-term (CSSv2), with rising core counts across hyperscalers further amplifying royalty growth.

AGI/CPU platform expansion broadening beyond hyperscalers

Mgmt sees growing traction in Arm-based infrastructure CPUs (AGI), with demand expanding beyond initial hyperscale customers (OpenAI, AWS ecosystem) to new cloud/enterprise adopters (e.g., ByteDance, Oracle, neoclouds).

Early AGI deployments are margin-dilutive (~high-30% GM), but roadmap targets ~50% GM over time as Arm increases in-house IP and reduces partner mix.

Adoption broadening across end markets, though uneven by segment

v9 adoption is now well-advanced across smartphones and data center, with CSS fully v9-based, while auto remains slower but still growing (low-/mid-20% YoY content growth).

Across markets, growth continues to be driven by higher content per device (core counts, complexity) rather than unit growth alone.

Connectivity (Optical/Networking)

Marvell (MRVL)

Connectivity-led AI infrastructure story with accelerating growth

We hosted CEO Matt Murphy and IR Ashish Saran, who framed MRVL as a connectivity-first AI infrastructure provider (vs. compute), with ~70% of revenue tied to connectivity/IP that is complementary to hyperscaler compute ecosystems.

Mgmt highlighted strong momentum driven by data center expansion (now ~75% of revenue), following a decade-long transformation into a data infrastructure leader.

AI cycle viewed as structural, not cyclical; constraints support durability

Mgmt emphasized the current AI cycle as a multi-year, structural infrastructure build (industrial revolution scale) rather than a typical semiconductor cycle.

Despite strong demand, supply constraints (logic wafers, memory, power) remain limiting factors, with bookings/backlog strong and upside dependent on AI capex trajectories.

Diversification improving materially; multiple \$1bn+ franchises emerging

MRVL expects increasing revenue diversification, with multiple new \$1bn+ businesses (cloud switching, optical/DC interconnect, analog) layering in ahead of scale-up ramps, and 15–20 products in production by 2028.

Mgmt emphasized the model is not dependent on any single socket, customer, or technology, with broad exposure across XPU, XPU attach, and connectivity layers.

Custom silicon gaining scale, but remain a secondary driver vs connectivity

Custom compute continues to scale (>2x growth expected), with multiple XPU and XPU-attached sockets and improving visibility toward \$8-10bn potential by 2028, though mgmt remains disciplined on near-term growth guidance.

Importantly, MRVL views connectivity as the larger controllable TAM, with custom compute more customer-directed but still benefiting from rising die size, complexity, and ASP trends.

Scale-up opportunity (CPO, switching) largely incremental and under-penetrated

Mgmt highlighted scale-up as incremental TAM beyond current scale-out driven growth, with:

- ~\$300m scale-up optics revenue next year (from ~\$0 today) -
- Greenfield opportunity in scale-up switching (no incumbent)
- Platform advantage via integration across optics, interconnect, switching, and XPU attach

Overall, MRVL’s positioning remains centered on owning the connectivity layer across scale-out, scale-up, and scale-across, with significant optionality as AI architectures evolve.

Lumentum (LITE)

Execution bottleneck remains supply (fab, test/assembly, supply chain)

We hosted Wupen Yuen (President, Global Business Units), who emphasized that the key challenge is execution vs. demand, with focus on fab scale, assembly/test ramps, and supply chain management.

Mgmt highlighted that AI-related components remain broadly supply constrained, with internal efforts focused on capacity expansion and moving legacy products to CMs to prioritize high-growth optical segments.

Supply still tight into scale-out; scale-up creates a second demand wave

Lumentum continues to see meaningful EML supply-demand imbalance (~30%), driven by strong scale-out demand, with no near-term resolution despite capacity ramps.

Importantly, mgmt sees this as a two-phase demand cycle, with scale-out absorbing capacity over the next few years followed by optical scale-up in ~2028 driving 3x+ increases in optical intensity, implying a second, larger demand wave before supply

Differentiation and pricing power improving on yield/complexity

Mgmt highlighted increasing technical differentiation (yield, complexity, design) as lasers move from scale-out to scale-up/CPO, with higher-performance designs harder to manufacture and more critical at system level.

As a result, LITE is seeing improved pricing dynamics vs. historical deflation, driven by quality/yield advantages and spot opportunities amid tight supply.

CPO remains a 2028 scale-up inflection; NPO interim step

Mgmt reiterated that CPO is a scale-up (not scale-out) story, driven by inference-driven bandwidth/power constraints, with adoption on schedule and increasingly urgent across customers.

Near-term architectures remain mixed (pluggable, NPO, CPO), though mgmt expects NPO as a temporary bridge, with convergence toward CPO post-2028 as scale-up deployments ramp.

Good visibility via LTAs; capacity remains the gating factor

LITE indicated strong visibility through multi-year take-or-pay LTAs (2–5 years) with hyperscalers, supporting aggressive capacity expansion.

However, supply (not demand) remains the key constraint, with mgmt focused on not fully locking capacity to retain flexibility as AI optical demand continues to scale.

Coherent (COHR)

AI optics order book extending into 2028

We met Paul Silverstein, Senior Vice President, Investor Relations and Corporate Communications, who highlighted broad data center strength across transceivers and emerging product categories. Mgmt noted the order book is greater than 4x revenue with orders extending into 2028, while overlapping ramps in OCS, CPO, multi-rail, lasers, and related optical content should support growth and margins.

InP capacity is the key supply constraint

Mgmt framed indium phosphide capacity as COHR’s largest near-term operational challenge. COHR expects to double capacity one quarter ahead of schedule and more than double it again next year, helped by existing facility space in Texas and Sweden, though tool lead times and qualifications remain the gating factor.

Component breadth limits architecture risk

Mgmt pushed back on the idea that CPO, NPO, XPO, or other architectures materially change COHR’s opportunity, given the company can sell transceivers, lasers, ELS, CW lasers, passive optics, and other components across multiple approaches. For COHR, the key point is optical content growth, not which specific architecture hyperscalers ultimately choose.

CPO and scale-up remain large future opportunities

Mgmt expects scale-out CPO revenue in 2H26 and scale-up in 2H27, with scale-up potentially driving an order-of-magnitude increase in unit demand. COHR does not view the UHP laser opportunity as limited to a smaller allocation and sees the broader CPO opportunity as materially larger than \$2bn-\$3bn over time.

Isolators and garnet are underappreciated bottlenecks

COHR highlighted its position in isolators, which are required for CW and EML lasers and depend on garnet supply. Mgmt said COHR has roughly 70% share in this area and does not face the same garnet constraint that could affect other suppliers.

GM path still has self-help ahead

Mgmt remains confident in reaching the 42%+ GM model over the next 3-4 years, but emphasized 42% is not the ultimate destination. Self-help remains ongoing, with additional yield, volume, mix, and ASP optimization still ahead after meaningful improvement since 2023.

Industrial still a recovery option

Mgmt noted industrial has been treading water for roughly two years, but semicap and precision manufacturing trends are improving. ASP optimization, services tied to industrial lasers, and future opportunities such as microLED remain potential incremental contributors outside the core data center optics story.

Credo (CRDO)

Connectivity exposure expanding across copper + optics

We hosted CEO Bill Brennan and CFO Dan Fleming, who emphasized the end of the copper vs optics debate, with a heterogeneous connectivity stack emerging across data center architectures.

Mgmt framed CRDO as a full-spectrum connectivity provider (AECs + optics + DSP + die-to-die), with reliability a core differentiator given lack of redundancy in GPU-to-ToR links, where failures can disrupt entire clusters.

AEC remains core growth driver; optical scaling faster from smaller base

CRDO continues to see strong momentum in AECs, with no visible peak in demand, benefiting from rapid growth in short-reach connectivity alongside pluggable optics.

At the same time, ZF Optics + optical DSP/Silicon Photonics are scaling faster, addressing the full pluggable TAM with feature-rich, reliability-enhanced solutions, implying additive (not substitutive) growth vs copper.

Positioned across AI clusters (training, inference, agentic)

Mgmt highlighted strong exposure to AI connectivity across training, inference, and emerging agentic workloads, with AECs a direct beneficiary as every server/CPU node must connect to ToR switches.

More broadly, CRDO sees the AI connectivity opportunity as expanding structurally with cluster size, density, and workload diversity.

Custom optical + system-level approach driving share gains

CRDO differentiated its model as system-level (cable/module ownership) vs. chip-only, with deep integration across SerDes, DSP, firmware, and qualification, enabling higher reliability and faster customer ramps.

This approach is supporting share gains (>\$1bn revenue trajectory), with optical expected to scale to ~\$600m+ near-term and grow faster than AECs, driven by demand for reliability and performance.

Long runway in AEC + 1.6T transition; supply largely secured

CRDO sees continued tailwinds from higher speeds (200G/lane, 400G roadmap) and increasing content per link, with AECs remaining viable across future generations (incl. “wide & slow” optical alternatives).

On supply, mgmt indicated visibility through FY27 with manageable exposure to leading-edge constraints, given smaller die sizes vs compute customers and strategic positioning in cluster deployments.

Astera Labs (ALAB)

CXL reaccelerating; inferencing/KV cache emerging as next leg

We hosted CFO Desmond Lynch and IR Nick Aberle, who highlighted renewed interest in CXL over the past ~18 months, with Leo positioned as a multi-year opportunity.

Mgmt sees initial CXL revenue in 2H26, becoming more material in 2027, with design wins tied to AI inferencing / KV cache alongside more general-purpose memory expansion use cases.

Scorpio-X broadens scale-up switch opportunity beyond initial sockets

ALAB continues to expand Scorpio from initial P-Series applications into broader P/X-Series scale-up switching, with 10+ unique engagements across PCIe and UALink -based architectures.

Mgmt sees the non-NVLink scale-up switch opportunity as large, with customers likely standardizing around one scale-up fabric over time; current engagements often start with PCIe scale-up and migrate toward UALink.

Optical roadmap building toward NPO/CPO; architecture-agnostic approach

Optical strategy is developing around NPO starting in 2H27, supported by connector/glass technology from aiXscale and internal analog/mixed-signal + SiPho/PIC work.

Mgmt views NPO as the intermediate step before CPO, with ALAB aiming to stay agnostic across customer-owned SiPho/PIC choices while anchoring the roadmap around Scorpio-X.

PCIe retimer lead still meaningful, but competition expected

ALAB remains early in PCIe 6 retimer/fabric ramps, with strong design win momentum and broad interoperability across CPUs, GPUs, networking, storage, and memory.

Mgmt does not expect 95% share to persist indefinitely, but believes its time-to-market, interoperability, and deployed base support an outsized position into PCIe 7.

GM profile intact; supply mostly manageable

CXL and custom products remain silicon-based and should track near the company's ~70% long-term GM target.

On supply, mgmt feels good on front-end capacity, though back-end constraints remain a near-term watch item.

MACOM (MTSI)

Defense strength broadening; GaN remains key differentiator

We hosted CFO Jack Kober and SVP Corp Dev/IR Stephen Ferranti, who highlighted continued strength in A&D, with defense growing double-digits despite historically LSD/GDP-type market expectations.

Defense is now ~70% of I&D mix (vs. ~35% three years ago), supported by demand across drones, missiles, radar, and other RF/GaN applications, where supplier availability remains limited.

LEO/Satcom opportunity intact, though timing somewhat lumpy

Mgmt reiterated LEO/Satcom remains a key growth vector, led by broader LEO customers, though some customer efforts to insource remain a watch item.

A previously announced \$50mn+ design win is now expected to begin revenue in C4Q26, spread across roughly CY27–28, with some launch-schedule-driven lumpiness but no change in overall opportunity.

Data center optics opportunity architecture-agnostic

In data center, MTSI remains positioned as a merchant supplier across LRO /LPO/NPO/CPO architectures, with the ability to monetize higher-value requirements through DSPs, TIAs, drivers, PDs, and lasers.

Mgmt sees accelerator vendors more focused on CPO, while hyperscalers appear more interested in NPO given desire to pair custom ASICs/processors with nearby optical subassemblies.

Capacity/absorption improving; GM path intact

Mgmt noted improving industrial demand, PD ramps, and better fab absorption are helping offset prior underutilization pressure.

Most data center products use third-party foundries, while PDs and I&D are internal, and mgmt remains focused on yield improvement and more die per wafer.

Gross margin should continue to improve, with mgmt targeting 60%+ exiting FY26.

Analog

Texas Instruments (TXN)

Cycle recovery broadening across industrial

We hosted Mike Beckman, Head of Investor Relations, who highlighted a strong 1H with broader regional strength and growth across every industrial sector. Long-tail customers that were on pause last year resumed activity in 1Q, while A&D continues to make new highs for secular reasons and data center remains a strong incremental driver.

Industrial still below prior peak ex-A&D

A&D is the only industrial sector making new highs every quarter and is now roughly a \$1bn run-rate business. The rest of industrial remains below prior peak, though energy infrastructure and test/measurement are growing faster given data center exposure and higher content versus the last cycle.

Data center emerging as third growth leg

Mgmt framed data center as a third leg of growth alongside industrial and auto, with DC up ~90% YoY in 1Q and at roughly a \$2bn run rate off 1Q levels. TXN sizes the 2025 data center TAM at ~\$7.5bn growing to ~\$12.5bn, with data center compute ~50% of the TAM today and broader opportunity across rack power, thermal management, pumps, liquid cooling, chillers, networking, and compute.

Auto steady at elevated levels

Auto remains steady at elevated levels, with TXN well exposed to China EVs where content is higher. Mgmt did not point to major regional exposure differences, instead framing growth around content opportunity and sustained elevated demand rather than a sharp unit-driven inflection.

Pricing is constructive but not a lever to gain share

Mgmt noted pricing typically declines LSD long term, but stabilized in 1Q and could become more constructive in 2H. TXN does not view price as the primary share-gain mechanism in analog, given the effort required to redesign components; the better share lever is having the right product and available supply when competitors do not.

SLAB deal clears high strategic and financial bar

Mgmt said the SLAB acquisition fit TXN's strict M&A criteria, combining SLAB's broad industrial portfolio and installed base with TXN's manufacturing technology and channels. Financially, the deal was helped by TXN's 28nm process development for Lehi, which can drive meaningful COGS synergies.

Capex tied increasingly to demand-driven tools

TXN expects 2026 capex of \$2bn-\$3bn, with any acceleration tied to nearer-term demand rather than brick-and-mortar construction. Mgmt framed maintenance capex at roughly 4% of revenue, with long-term capex intensity around 1.2x growth as TXN prepares for 7%-10% revenue growth while preserving the strategy of growing

FCF/share.

Analog Devices (ADI)

Cycle recovery broadening beyond idiosyncratic strength

We hosted CFO Rich Puccio, who highlighted that ADI is now seeing classic upcycle signals, with significant growth in the mass-market portion of the business over multiple quarters. Mgmt had already seen strength in ATE, Aerospace and Defense, and data center communications, but the incremental point is that cyclical recovery is now adding to those growth vectors.

Industrial demand still has room to recover

Mgmt pushed back on double-ordering concerns, noting industrial is still ~20% below

~6-7 weeks. Infrastructure, Aerospace and Defense, and data center-related industrial demand are easier to trace to end-market activity, while the rest of industrial remains below the long-term consumption line.

Data center now close to 20% of sales

Data center exposure across Power, Optics, and ATE is close to 20% of ADI sales, with power and optical roughly equal within the data center portfolio. Mgmt sees power as the larger incremental growth opportunity, particularly as higher voltage architectures and vertical power adoption drive demand for IVR and last-millimeter power delivery.

Auto strength driven by content, not units

Mgmt reiterated that ADI expected auto to grow YoY despite a record 2025, with demand accelerating sooner than expected late in the last quarter. Growth is not dependent on SAAR recovery, but on content and share gains, particularly China EVs, L2/L3 ADAS adoption, and BMS, which returned to YoY growth for the first time in two years.

Pricing is cost recovery, not shortage-driven

ADI's first broad price increase in 2026 was framed as recovery of accumulated input cost inflation, not a response to shortages. Lead times remain largely within normal ranges, and mgmt emphasized that the more important pricing lever remains design-in value capture, where newer products carry roughly double the ASP of products 10 years and older.

GM upside shifts from utilization to mix and revenue growth

ADI is holding ~72.5% GM, with trough-to-peak improvement helped by higher utilization. From here, mgmt sees incremental GM upside coming more from mix, particularly industrial returning toward prior peak mix, and revenue growth with flow-through, rather than further utilization gains.

Cash use focused on synergies and targeted tuck-ins

Mgmt reiterated ADI is on track for \$1bn of revenue synergies by 2027, with hundreds of millions delivered in 2025 and more expected in 2026. Large consolidation is tougher given regulatory constraints, so MS&A focus is on tuck-ins in software, digital, AI, and power that accelerate time-to-market or solve specific customer problems.

Microchip (MCHP)

Data center disclosure clarifies upside

We hosted Chairman and CEO Steve Sanghi and CFO Eric Bjornholt, who clarified that the dedicated data center solutions business generated \$302.7mn in CY25 and is expected to grow to ~\$500mn in CY26. Mgmt emphasized this is not total data center exposure, as MCHP also sells MCUs, analog, timing, security, memory, and other catalog products into data center, within a broader data center/compute end market that was ~\$787mn in CY25.

Pricing now shifting to cost recovery

Mgmt announced plans to raise prices after absorbing higher costs across labor, chemicals, gases, transportation, energy, gold, copper, foundry, and OSAT. The price action is not shortage-driven and is intended to recover input cost inflation with a modestly accretive bias, though MCHP does not plan to disclose the average increase.

Customer relationships improving despite pricing

Mgmt framed MCHP's delayed pricing action as a customer relationship positive, after prior cycle friction from PSP and non-cancelable, non-returnable terms. Competitors have already taken multiple price increases, while MCHP held off, which mgmt believes has helped drive design wins and customer re-engagement.

Gen6 PCIe switch is key data center swing factor

MCHP’s data center solutions unit includes storage controllers/accelerators/expanders, PCIe/CXL memory controllers, and PCIe switches, with a fourth product line coming. Mgmt acknowledged missing Gen5, but said Gen6 is on time, built on 3nm, offers 30%-40% lower power, and has already won eight switch designs plus one retimer design before production.

Analog recovery shifting from inventory to innovation

Mgmt said the first phase of growth was customers and distributors moving purchases back toward consumption after inventory depletion. The next phase is driven by design wins and innovation across data center, auto, A&D, and industrial, with inventory-depletion growth expected to largely end in the next quarter or so.

Margins moving back toward target

MCHP guided June quarter GM to 62.75%, with a path back toward the 65% target as underutilization charges decline. March quarter underutilization charges were \$46.6mn, and adding those back would put GM essentially at 65%; OpEx is ~29% of revenue versus a 25% long-term target, with leverage expected as revenue grows.

Prior peak less useful as a benchmark

Mgmt pushed back on using the prior revenue peak to frame recovery, noting the last cycle included significant shipments into customer inventory, broad price increases, and mix shifts that distort both revenue and unit comparisons. The focus from here is on forward opportunity rather than a return-to-peak framework.

Onsemi (ON)

Q1 marked bottom, but auto not yet in cyclical recovery

We hosted CEO Hassane El-Khoury and CFO Thad Trent, who reiterated Q1 as an inflection point based on broader visibility, improving backlog layering, modest lead-time extension, and PMI moving above 50. Mgmt was careful to frame auto as moving from under-shipping to shipping to demand, not yet a cyclical recovery.

2H supported by program ramps and backlog visibility

Mgmt expects 2H to be better than 1H, driven by full-quarter benefits from program ramps that started in Q1/Q2 and stronger booking visibility into Q3/Q4 versus prior years. AI data center revenue was 2x better than expected in Q1, while renewable energy, microgrid, and ESS are expected to resume growth in 2H.

Auto content growth beyond SiC

ON sees auto growth driven by EV content, zonal architecture, 10BASE-T1S Ethernet, ultrasonic sensing, and inductive position sensing. Mgmt emphasized that content has outgrown SAAR by high-single-digits historically, with new products adding content that was not present two years ago.

SiC differentiation is device, not substrate

Mgmt pushed back on China SiC substrate concerns, framing substrate as the least differentiated part of the stack. ON differentiates through device performance and supply assurance, with mgmt citing >50% share in China auto and strength with exporters, while noting the module-to-device rebalancing is largely behind the company.

AI data center content step-up tied to 800V

ON expects AI data center revenue to double this year, based on existing wins already ramping. Mgmt framed the larger opportunity around the move to 800V racks, where ON content could rise from ~\$15k per rack today to ~\$115k, with high-voltage content moving from roughly ~\$5k to ~\$55k.



Vertical GaN and SiC JFET key differentiators

Mgmt highlighted vertical GaN and SiC JFET as core technologies for higher-power data center architectures. ON believes its vertical GaN enables 800V conversion in a single chip, supporting higher power density and smaller system footprint versus competing approaches.

GM leverage remains the key model debate

Mgmt expects GM expansion through the rest of the year and over a multi-year period, with utilization at 77% versus 68% last quarter. Each utilization point contributes 25-30bps of GM two quarters later, while Fab Right, mix, internalization of divested fab output, and pricing each provide additional levers toward a GM above +50%.

Allegro MicroSystems (ALGM)

Data center content expanding beyond GPU fans

We met Michael Doogue, President and CEO, and Jalene Hoover, Vice President of Investor Relations and Corporate Communications, who highlighted rapid data center growth across fan driver ICs, current sensors, and emerging gate drivers. Mgmt noted ~80% of current DC revenue is fan driver ICs, increasingly used in power supplies as well as GPU cooling, which should limit the impact from liquid cooling.

800V drives larger sensor and power opportunity

Mgmt framed 400V/800V data center architectures as a major tailwind for ALGM's magnetic current sensors, motor drivers, and isolated gate drivers. Rack content has expanded from ~\$150 a year ago to ~\$425, with more than two-thirds tied to motor drivers and current sensors, and current sensors the largest piece.

Gate drivers add longer-term upside

ALGM is investing in isolated gate drivers used to turn GaN and SiC devices on/off in high-voltage applications. Mgmt sees the integrated gate driver market at ~\$3bn by 2028, with ALGM targeting at least 10% share over time, though meaningful revenue is still several years out.

Auto content growth remains above SAAR

Mgmt expects auto to grow at least ~10% above SAAR, supported by electric braking, steering, safety-critical magnetic sensing, and ADAS-adjacent content. Electric braking and steering transitions are still early, with China ramping near term and North America also contributing.

China design-win momentum strong

Mgmt noted China auto dollar content is much stronger than five years ago, with strong design-win velocity and sticky safety-critical content. China is ~90% auto for ALGM, and mgmt emphasized that even if local suppliers copy a few components, the company's broader content position remains resilient.

Robotics could become a meaningful adjacency

Mgmt highlighted meaningful existing exposure to robotics through magnetic sensors and motor control in robotic joints. Some designs can require 75 to 90 chips per robot, with ASPs well above \$15 per robot, and mgmt expects to revise its robotics opportunity higher as adoption broadens.

Pricing moving higher

ALGM recently raised prices by ~10%, with mgmt suggesting a more gradual pricing approach over the next few years and a focus on holding price once implemented. Auto contracts remain a key mechanism, with additional increases possible where customers do not convert certain materials such as gold to copper.

Semi Capital Equipment

Applied Materials (AMT)

Orders drive Semi Systems step-up

We hosted CFO Brice Hill, who highlighted a step-up in Semi Systems growth to >30%, driven by incremental orders received over the past quarter. Mgmt framed the upside as broad-based across leading-edge logic, DRAM, and advanced packaging, with AI data center the primary demand driver and advanced logic/DRAM growing faster than the overall business.

Cleanroom remains the bottleneck and not AMAT capacity

Mgmt reiterated that near-term growth is gated by cleanroom availability rather than AMAT's ability to build tools, noting the company has ample floor space, better supply chain visibility, and a more diversified supplier footprint than during COVID. Customers are addressing the constraint by optimizing existing fab floor space and accelerating new fab construction, but new fab timelines remain the limiting factor.

Mix shifting toward leading-edge; ICAPS digesting but stable

WFE and revenue mix are skewing back toward leading-edge logic and memory, while ICAPS is digesting after several years of significant growth when it became >50% of WFE/foundry revenue. Mgmt expects ICAPS to be flat to slightly up this year globally and in China, with demand supported by data center related power semis, analog, and photonics, though not framed as a high-growth market.

China still important, but concentrated in ICAPS

China remains the largest ICAPS region and is now effectively all trailing-edge for AMAT given restrictions on leading-edge logic and memory. Mgmt noted AMAT is doing well where it can compete, particularly around 28nm, and characterized share as strong in the accessible China market.

DRAM share gains remain a key confidence point

Mgmt reiterated confidence in DRAM after gaining ~10pts of share over the last decade, supported by technologies that track the logic roadmap where AMAT has historically been strong. The company highlighted its broad conductor etch, deposition, modification, and process control portfolio as synergistic, with 4F2 and future 3D DRAM expected to remain favorable inflections.

Services guide raised; margin progress gradual

AMAT raised Services growth expectations to mid-teens from low-double-digit, helped by installed base growth and the roughly one-year lag from systems shipment to service contract revenue. On margins, Semi Systems GM is now disclosed at ~55%, with quarterly movement largely tool-mix driven; pricing is the bigger lever from here, while cost and supply chain optimization have already captured much of the low-hanging fruit.

Lam Research (LRCX)

WFE upside broad-based, cleanroom remains gating factor

We hosted EVP and CFO Doug Bettinger, who highlighted WFE moving from LRCX's initial ~\$135bn view to ~\$140bn with a bias to upside, driven by incremental cleanroom availability and project pull-ins across segments. Mgmt framed the industry as undersupplied, with cleanroom availability the key constraint rather than demand, supporting a favorable setup into 2027 as additional projects come online.

AI demand driving advanced equipment intensity

Mgmt reiterated AI as the primary demand driver, showing up across accelerator chips, HBM/DRAM, NAND, and advanced packaging. Higher etch/deposition intensity is the critical driver for LRCX as more architectures move into 3D structures, including GAA,

SAM expansion already visible; foundry/logic outgrowth underappreciated

LRCX reiterated its path from low-30% WFE SAM toward high-30%, noting it is already in the mid-30s as architecture inflections increase etch/deposition intensity. Mgmt noted its diversifying mix as foundry represented 54% of systems revenue last quarter and 59% in the prior quarter, supported by investments made several years ago.

NAND conversions pulled forward, but DRAM getting cleanroom priority

Mgmt said the previously discussed ~\$40bn NAND conversion opportunity over several years is now expected to be largely complete by the end of next year. LRCX remains relatively agnostic between conversions and greenfield capacity, though current cleanroom allocation is skewing toward DRAM given stronger profitability; NAND should accelerate over time as constraints persist and profitability improves.

CSBG durable, but utilization tailwind moderates

GM held near 50%. CSBG was \$2.1bn last quarter and remains a highly recurring, profitable part of the model, driven by installed base, spares, service, upgrades, and Reliant tools. Mgmt noted industry utilization is effectively at 100%, limiting further utilization upside, while GM is expected to remain around the recent ~50.5% guide as pricing and efficiencies offset inflation and customer mix headwinds

KLA Corp (KLAC)

WFE upside with strong 2027 visibility

We hosted CFO Bren Higgins, who highlighted 2026 WFE up high-teens to ~20%, with the market likely above the prior ~\$140bn+ view. Mgmt also noted unusual visibility into 2027, with expectations for another strong year as new fabs and expansions come online across logic/foundry, DRAM, and some recovery in flash.

Process control still gaining share of wallet

Mgmt reiterated that process control has gained share of WFE alongside lithography over the last five to six years, driven by design complexity, larger die, tighter performance specs, and higher value die. KLA targets moving from ~7.5% of WFE to ~9% by 2030, supported by continued share gains in semiconductor process control.

Advanced packaging becoming material

Packaging has moved from a low-relevance market for KLA to a meaningful growth driver as requirements become more front-end like. Mgmt expects packaging revenue to reach ~\$1bn this year, up from ~\$635mn last year and ~\$300mn the year before, with share rising from ~1% a few years ago to mid-7% or better this year.

Memory intensity improving towards logic levels

Mgmt pushed back on the view that a more memory-heavy cycle is less favorable for KLA. HBM is process-control intensive in certain product lines, while more EUV in DRAM should lift memory intensity. KLA still assumes logic is ~60% of WFE longer-term, but expects memory to support continued outperformance.

Supply chain being positioned for growth

Mgmt is focused on ensuring KLA is not a bottleneck as industry growth accelerates. 1H growth was constrained partly by component availability, particularly optics, but capacity improves in 2H and KLA is giving suppliers multi-year visibility and investing where needed to secure dedicated capacity.

Margins best-in-class with modest upside

KLA is operating near 62% GM today and sees a path to 63%-64% over the next several years. Mgmt emphasized value pricing tied to new product cadence rather than broad pass-throughs, with service dilutive to GM but accretive to operating margin.

Services remains unusually durable

KLA raised service growth expectations to 13%-15%, supported by longer tool lives, high uptime requirements, and ~80% contract revenue. Tool lifetime has expanded from ~10 years to over 20 years, while contract value as a percent of ASP is 1-2pts higher than a decade ago.

Teradyne (TER)

ATE intensity rising with AI/advanced packaging; TAM path expanding

We hosted CEO Greg Smith, who highlighted an inflection in test equipment buy rates, with ATE now rising despite AI semi growth being heavily ASP/margin-driven.

Mgmt sees test intensity increasing with wafer output, advanced packaging, HBM stacking, and accelerator complexity, where downstream failure costs justify more upstream test. In a larger WFE environment, TER sees ATE becoming a larger slice of semicap spend, with a \$250bn WFE market potentially implying ~\$20bn ATE TAM.

Evergreen model implies share gains, not just TAM growth

TER's long-term model assumes a \$12-14bn semi ATE TAM, under which Teradyne would be a ~\$6bn revenue company, driven by both market growth and share gains.

Mgmt sees ATE share moving from ~30% in 2025 to ~35-38%, helped by improving position in compute, DRAM/HBM, and recovery in mobile/power/industrial/flash.

Compute opportunity broadening; CPU helpful but accelerators remain larger

Mgmt framed CPU as an incremental opportunity, but still smaller than accelerator test: accelerator test intensity is ~4x a data center CPU, and current accelerator:CPU ratios imply much larger accelerator-driven TAM.

That said, TER sees meaningful share opportunity in ARM-based CPUs, where it is qualified or qualifying across major programs, while x86 opportunity is more back loaded, with one internal-test customer potentially opening from 2028+. -end

CPO test becoming a meaningful new TAM

TER sees CPO test as early but material: roughly \$100mn of CPO test equipment in 2026, rising to \$300-700mn by 2028, with potential to exceed \$1bn+ as scale-up CPO follows scale-out.

TER is better positioned in scale-out/networking CPO, while Advantest is stronger in accelerator/scale-up exposure, setting up a competitive duopoly as the market scales.

Large compute customer share ramp remains gradual

At a major merchant compute customer, TER has moved from qualification into a "fast follower" phase, with initial production orders and follow-on engagements.

Mgmt sees a natural ceiling of ~30% share while following the incumbent, with a mature dual-vendor strategy potentially emerging over the next 3-4 years based on platform capability, throughput, and reliability rather than incumbency.

Camtek (CAMT)

Advanced packaging is the core business

We hosted CFO Moshe Eisenberg, who highlighted that CAMT's long-standing focus on advanced packaging is paying off as AI demand increasingly depends on CoWoS, EMIB, HBM, and related 2.5D architectures. Advanced packaging is now ~70% of revenue, with roughly half of that tied to AI applications.

2H ramp driven by HBM and CoWoS-like demand

CAMT reported \$120mn of revenue in Q1 and guided Q2 to \$130mn, with a significant 2H26 ramp expected. Mgmt expects 2H revenue to grow ~30% HoH, driven primarily by HBM and CoWoS-like 2.5D applications, with Q4 revenue expected to be ~40% above Q1.

2027 visibility improving on new orders

Mgmt expects growth to accelerate further in 2027, supported by a \$105mn order announced earlier in the week from two customers, split roughly 50/50 between a tier-1 OSAT and an HBM vendor. These orders are expected to be delivered within 2027, with additional orders expected for both 2026 and 2027, particularly from HBM.

HBM transition increases inspection intensity

Mgmt sees a major HBM inflection from both more wafer volume and technology transitions from HBM3e to higher stacks and HBM4. More complex wafers, tighter scaling, and longer process times require new inspection and metrology tools, including CAMT's fifth-generation Eagle platform.

OSATs becoming more strategic in AI supply chain

Mgmt highlighted growing OSAT importance as TSMC capacity spills over to Amkor, ASE, and Chinese OSATs, while some OSATs are also receiving business directly from hyperscalers and custom silicon vendors. CAMT is positioned across memory, TSMC CoWoS, and OSATs serving both TSMC and hyperscaler ecosystems.

GM and capacity support leverage

Mgmt sees a clear path to GM improving toward 53%-53.5% in the next few quarters, with a \$750mn revenue model supporting 54%+ GM over time. CAMT is investing in Germany to add capacity and does not view capacity as a limiting factor, with supply chain initiatives also expected to support margin improvement.

Nova Measuring Instruments (NVMI)

GAA and DRAM driving WFE outperformance

We hosted CEO Gaby Waisman and CFO Guy Kizner, who highlighted continued WFE outperformance driven by Nova's differentiated dimensional, material, and chemical metrology portfolio. Mgmt sees momentum in advanced logic, especially GAA, and noted Nova is positioned across all four GAA players.

\$500mn GAA opportunity on track

Mgmt reiterated that Nova is on track for its cumulative ~\$500mn GAA revenue opportunity across 2024-2026. GAA increased metrology intensity by ~30%, with additional upside possible as the node broadens beyond the initial leading customer and more players ramp production.

Advanced packaging now >20% of sales

Advanced packaging has grown from essentially no business five years ago to >20% of revenue today, with roughly two-thirds logic and one-third HBM. Mgmt sees continued double-digit growth as hybrid bonding, tighter specs, and front-end-like process steps drive higher metrology intensity.

Hybrid bonding qualifications approaching key decision point

Nova is already in qualifications to become tool of record for hybrid bonding, with some selections made and other evaluations in advanced stages. Mgmt expects to have a clearer view of share by early next year, with the opportunity tied to planarity, TSV, RDL topography, and warpage requirements.

Lab-to-Fab still expanding

Mgmt highlighted Lab-to-Fab as a core growth framework, moving technologies like XPS, XRF, inline SIMS, and Raman from the lab into high-volume fabs. XPS/XRF adoption has moved from tool-per-customer toward tool-per-fab, while inline SIMS is still earlier in penetration but saw a record quarter in Q1.

Capacity and M&A support next leg

Nova doubled capacity in 2024-2025 and is adding a production center in Asia to support 2027 and beyond. On M&A, mgmt noted >\$1.7bn of cash, with most allocated toward inorganic growth, but remains disciplined around accretive deals with clear top-

Share gains remain broad-based

Mgmt cited Gartner data showing share in CD and FinFET metrology rising from ~20.3% in 2023 to almost 28.8% in 2025. Future gains may not come at the same pace, but mgmt sees opportunities across the portfolio as new products and applications scale.

MKS Instruments (MKSI)

WFE exposure should outperform on the way up

We hosted Ram Mayampurath, Executive Vice President and Chief Financial Officer, Pareto Misra, Vice President of Investor Relations, and Keith Loop, Vice President of Strategic Marketing and Global Service. Mgmt framed MKSI as typically outperforming OEM builds in an upcycle as customers build inventory ahead of tool shipments. Mgmt is planning around strong customer demand signals and sees the current setup as a long-cycle opportunity, with NAND still largely absent and therefore potential upside when greenfield activity returns.

Dielectric etch position remains strong

Mgmt was confident in MKSI's position in high-aspect-ratio dielectric etch, where the company plays through critical power and plasma solutions. MKSI also sees opportunity in pulsed DC for conductor etch applications, where it has been investing and could gain share as new technology inflections ramp.

Greenfield more favorable than upgrades

Mgmt noted that upgrades are primarily a power opportunity, but greenfield fab expansion carries significantly more MKSI content because all new tools require its subsystems and components. Lead times remain relatively short, and the new Malaysia factory opens this month, though mgmt said it is not needed to meet this year's demand.

Capacity planning shifting toward 2028

MKSI does not need incremental capacity beyond Malaysia for 2026, but is already discussing additional capacity needs for the 2028 timeframe. Capex should remain around 4%-5% of revenue, with the company able to scale near-term demand largely through variable costs and existing cash flow rather than customer prepayments.

Chemistry benefiting from AI-driven PCB complexity

Mgmt highlighted unexpectedly strong PCB demand as AI drives higher layer counts, tighter lines, and more complex interconnect across package substrates, HDI boards, and multilayer boards. PCB chemistry is growing double-digits despite historically tracking closer to GDP-plus, with AI exposure rising as servers become a larger demand vector.

All three chemistry segments improving

MKSI's chemistry business is roughly split across multilayer boards, HDI, and package substrates. Mgmt sees growth across all three, with package substrates tied directly to AI, HDI benefiting from high-end smartphones and server applications, and multilayer boards supported by conventional server growth.

Margins supported by technology value, not opportunistic pricing

Mgmt said it does not use tight demand as a reason to raise price, but does monitor inflation and price for the value delivered. In higher-end interconnect chemistry, MKSI believes its margins are meaningfully above competitors given stronger technology, R&D depth, and exposure to more complex applications.

2026 an investment year, leverage longer term

Tariffs and metals inflation are embedded in guidance, with mgmt noting GM would have been 30-40bps higher absent tariff headwinds. Underutilization benefits are not expected in 2026, but Malaysia ramp and supply chain initiatives should support future

Capital allocation focused on deleveraging

Mgmt does not see a need for additional platform M&A after Atotech and is focused on organic growth, debt reduction, and liquidity. Target leverage remains 2.0x-2.5x, with opportunistic buybacks possible once leverage reaches the desired range.

Advanced Energy (AEIS)

Semis and data center driving broad strength

We hosted CEO Steve Kelley and CFO Paul Oldham, who highlighted strong demand across semis, data center, and industrial/medical with no major areas of weakness. Semi revenue is expected to grow >20% YoY in Q2, with 2H sales +30% above 2H25. Q2 Data Center moderation is temporary (customer supply constraint) and the segment should return to strong growth in the 2H, alongside improving I&M throughout the year.

New semi products support share gains

Mgmt reiterated confidence in gaining share from eVerest, eVoS, and NavX, which have been in development and customer qualification over the past three years. These products replace platforms that have been in place as long as 10 years. Revenue from these products was \$10-\$20mn in 2025 and should at least double this year, with more meaningful ramps in 2027 for conductor etch and 2028 for dielectric.

Design wins create multi-year revenue annuities

AEIS noted leading-edge semi wins as largely sole-source once qualified, with wins typically lasting several years and in some cases supporting nodes for 5-10 years. New products carry higher pricing and margins than mature platforms, while some customers may also back-convert the technology to more mature processes.

Data center demand strong, supply the gating factor

Data center revenue grew 107% in 2025 and is forecast to grow at least 35% in 2026, with upside dependent on downstream supply constraints around memory and logic. Mgmt said most customers are providing at least 12 months of visibility, “second wave” neocloud/enterprise customers expected to begin ramping in 2027-28. AEIS continues to focus on the high-end, high-density, high-efficiency board/rack-level power conversion vs. main competitor Delta offering system/-infra-level portfolio.

800V architecture is a 2027/2028 opportunity

AEIS is focused on rack-level 800V power conversion, with current modular products already sampling and positioned around high power density and efficiency. Mgmt expects limited 800V revenue until 2H27, with 2028 the more meaningful ramp, and sees higher content and better margins as products shift toward modules and boards rather than boxes. They did not specify exact content/ASP/unit opportunities associated with the 800V architecture change/ramp.

Thailand pull-in supports semis and data center

AEIS is opening its Thailand factory a quarter early, with production starting in Q4 for both semiconductor and data center products. Mgmt sees nameplate capacity of ~\$3.5bn, with ability to stretch to ~\$4.5bn-\$5bn if needed, suggesting capacity should not be a limiting factor near term.

Gross margin milestone reached, more to go

AEIS crossed 40% GM in Q1 and continues to expect exiting the year at ~41%, with a longer-term goal of 43%+. New products should add 200-300bps to margin as they scale, while volume, manufacturing efficiency, and OpEx growing at roughly half the rate of revenue support 35%-45% incremental operating margin flow-through.

Axcelis (ACLS)

Veeco deal integration planning underway

We hosted David Ryzhik, SVP of Investor Relations and Corporate Strategy and interim CFO, who highlighted that ACLS and Veeco are still operating as separate companies while integration planning continues. The deal is on track to close in 2H26, with ACLS planning around cost and revenue synergies but not yet implementing integration actions.

Cost synergies straightforward; revenue synergies more strategic

Mgmt reiterated the \$45mn run-rate cost synergy target within two years, driven by normal public company and corporate cost overlap. Strategically, implant and laser annealing sit adjacent in the fab, while Veeco brings ion beam, MOCVD, and advanced logic exposure that complements ACLS's strength in ion implant, SiC, and Korean memory.

Power semi remains muted, but SiC can grow again

Mgmt expects ACLS revenue to be roughly flat in 2026, with memory stronger but power and general mature down YoY. SiC has been treading water as EV growth decelerated, but penetration remains low and data center power infrastructure is an incremental driver, supporting the view that SiC could return to growth next year if current trends continue.

Memory strength concentrated in DRAM

Mgmt said memory strength today is primarily, if not entirely, DRAM. NAND implant remains muted because customers are focused on vertical scaling, where added layers drive more etch/dep than implant; once that scaling path is exhausted, new NAND capacity could become more relevant for ACLS.

High current implant a focus area

ACLS is the number two player globally in high current implant and has been investing in the segment for several years, including the recently introduced H6 platform. Mgmt noted progress with memory customers, particularly in Korea, and recently announced a win with a North American memory customer that was attributed to technology rather than commercial pricing.

GaN and next-gen SiC add optionality

Veeco is expected to bring a competitive position in 300mm GaN MOCVD, while ACLS remains more leveraged to SiC where implant intensity is higher. Mgmt also highlighted superjunction SiC and JFET as potential next-generation opportunities, where high energy implant could be particularly relevant.

China competition still limited

Mgmt acknowledged local Chinese implant competitors are trying to gain share but still see very small share in China. Customers remain focused on yield, contamination, and aftermarket support, where mgmt views ACLS and AMAT as the only meaningful options today.

Cadence (CDNS)

EDA/IP setup strongest in years; AI a net positive

We hosted CEO Anirudh Devgan, who highlighted Cadence's ~15% revenue CAGR over the last five years and Rule of 60 financial profile, with competitive positioning now the strongest it has been.

Mgmt sees the value of EDA/IP materially higher as AI and semiconductor complexity accelerate, with EDA and IP both having potential to grow meaningfully over the next five years while SDA remains additive.

AI expands TAM rather than displacing EDA usage

Mgmt pushed back on AI disruption risk, arguing Cadence benefits from both building AI chips and applying AI to design workflows.

Agentic tools like ChipStack/ViraStack/InnoStack create a new TAM by automating manual design work, while also increasing usage of base tools such as Jasper/Xcelium.

The key point: AI helps blunt an otherwise unrealistic engineering headcount curve as transistor/system complexity rises exponentially.

Design activity improving across hyperscalers, system companies, analog

Design starts remain strong despite rising leading-edge costs, driven by hyperscaler verticalization, AI custom silicon, physical AI, and improving analog/mixed-signal activity.

Mgmt cited broader engagement from hyperscalers and system companies, with physical AI/edge AI beginning to drive early design activity across autos, robotics, and China system OEMs.

IP growth accelerating from focused "star IP" strategy

IP has become a faster growth area as Cadence focuses on advanced-node DDR, PCIe, UCIe, HBM, and SerDes IP, where AI disaggregation and advanced packaging are increasing value.

Mgmt noted IP performance has improved as R&D quality strengthened, with share gains supported by better PPA and deeper engagement with leading-edge customers.

Margin leverage aided by internal AI productivity

Cadence is applying agentic workflows internally, especially in IP, targeting 2-3x productivity improvement through lower headcount per project and shorter schedules.

Mgmt expects these tools to support both customer value capture and internal leverage, with recent incremental margins already around 60%.

Consumer

Qualcomm (QCOM)

Data center AI strategy broadening

QCOM IR previewed a broader AI/data center strategy ahead of their AI-focused Investor Day, with FY27 custom silicon revenue from one hyperscaler expected to be “material” and represent the lion’s share of initial data center revenue.

QCOM also plans to address the market with merchant AI accelerators (AI200/AI250), where AI250 introduces a differentiated memory architecture aimed at memory bandwidth constraints, plus future CPU engagements and Alphawave-enabled connectivity IP.

Right to win centered on power efficiency, CPU IP, and flexible custom model

Mgmt framed QCOM’s data center edge around high-performance / low-power compute, track record in custom CPU design, and willingness to combine customer IP with QCOM’s execution engine.

On CPU architecture, QCOM wants flexibility across ARM and RISC-V, with RISC-V viewed as customer-led, particularly where hyperscalers or China customers prefer a license-free/open ISA model.

Handsets pressured by Apple roll-off and China Android memory headwinds

Mgmt reiterated expectations for ~20% iPhone 18 modem share, implying roughly \$2bn Apple revenue including new and legacy volume.

China Android weakness appears inventory-driven, with OEMs reducing inventories by roughly one month in both March and June quarters as memory prices rose; mgmt sees September/October launches as supporting a bottom, though handset headwinds could persist through 2027 from memory cost pressure.

6G broader than prior telecom cycles

Mgmt sees 6G as underappreciated, with a symbolic start around 2028 and more material ramps in 2029–2030.

The opportunity could be broader than a normal wireless cycle, with carriers potentially evolving toward edge/cloud-like infrastructure as uplink-heavy AI endpoints, drone tracking, digital twins, and distributed processing requirements grow.

Mgmt does not see a near-term step-function opex need, but expects incremental spend as data center revenue scales; FCF returns remain the framework, though inventory/working capital is intentionally running higher to support future ramps.

Ambarella (AMBA)

Edge AI positioning centered on power-efficient vision + mature software

We hosted CEO Fermi Wang, who emphasized AMBA’s focus on edge AI, with differentiation from an extensive silicon roadmap, strong video quality/power efficiency, and a mature software stack/SDK that helps customers build applications on top of AMBA processors.

Mgmt framed AMBA’s advantage as application-first silicon design vs. more general-purpose architectures, enabling better die size, power, and performance optimization for embedded AI workloads.

Auto solid; IoT remains bigger edge AI growth vector

Auto should reach record revenue this year, with growth beyond autonomous driving from fleet/telematics, DMS, and other automotive applications.

IoT remains the larger long-term growth engine, with near-term opportunities in drones, robots, and edge infrastructure appliances that aggregate sensors and enable GenAI on existing infrastructure using low-power silicon.

Hanwha LTA adds visibility and broader edge AI optionality

Mgmt highlighted the Hanwha agreement as strategically important, with Hanwha looking to standardize on a unified edge AI platform across cameras, retail, drones, robots, and other applications.

Hanwha Vision is already a MSD% revenue customer, with upside from higher AMBA share, higher ASPs, and multi-generation adoption across broader Hanwha use cases; near-term upside is not fully reflected in current expectations.

Growth constrained more by end-market/memory than AMBA supply

Mgmt reiterated FY growth of 10-15%, with no AMBA supply limitation, but memory cost inflation remains a constraint for customers and could affect demand/pricing.

Next year, CV7 (4nm) and CV8 (2nm, taped out in January, expected back shortly) should support ASP/mix upside, though ramp timing depends on customer adoption and memory availability.

Software/Cooper platform key to stickiness and scaling

AMBA's Cooper SDK is used by nearly all customers, supports ~200 model architectures, and enables easier software portability across AMBA silicon.

Mgmt views mature software as critical to scaling through an indirect sales/channel model, especially as robotics and edge AI opportunities broaden.

Ambiq Micro (AMBQ)

Ultra-low power edge AI leader; SPOT is core differentiation

We hosted CEO and CTO/Founder, who framed AMBQ as an ultra-low-power edge AI semiconductor provider, enabling meaningful battery life for AI-enabled endpoint devices.

Mgmt highlighted SPOT (Subthreshold Power Optimized Technology) as the key platform advantage, enabling 2-5x lower power and positioning AMBQ well as AI moves out of the data center and energy becomes the key constraint for edge compute.

Wearables remain core; AI expands use cases and ASP opportunity

Mgmt sees wearables evolving from basic step counters into AI-enabled health/fitness agents, with customers using AMBQ's power efficiency to enable longer battery life, more functionality, and smaller batteries.

Customer demand appears broad rather than duplicative, with different end markets across watches, bands, rings, and health-focused devices; mgmt sees strong backlog and little/no channel inventory.

Edge AI TAM growing faster than expected; demand exceeds planning

Mgmt indicated edge AI demand is growing faster than internal expectations, with 2Q/2H demand above plan and customers paying for expedite.

Unlike 2023, mgmt sees demand as end-market driven rather than inventory build, with 22-week lead times and hard backlog supporting growth through the rest of the year and into 2027.

Expansion beyond consumer: medical, industrial, smart home, AR/VR
Roughly a quarter of business is already outside consumer, with growth opportunities in medical, industrial, agriculture, smart home/building, and AR glasses.

Medical is viewed as a significant future opportunity as wearable/medical boundaries blur, while AR glasses represent another natural fit given battery-life constraints.

Licensing/semi-custom could unlock larger markets longer term
Mgmt is evaluating IP licensing and semi-custom models for markets AMBQ cannot serve directly near-term, including auto and potentially data center, with longer-term timing around 2028–2030.

Nearer term, focus remains on expanding product cadence, moving from ~1 product/year historically toward multiple launches, while building software/middleware, AI kernels, runtimes, and reference designs to increase stickiness and monetization.

Glossary: 800

V: 800 Volt

10BASE-T1S: Single-pair Ethernet standard used in automotive zonal/vehicle networks 2

.5D: 2.5-dimensional advanced packaging architecture using an interposer or substrate-level integration

3D DRAM: Three-dimensional DRAM architecture

4F2: Four-F-squared DRAM cell architecture shorthand

14A: Intel 14 Angstrom process node

18A: Intel 18 Angstrom process

node A&D: Aerospace and Defense

A/P: Advanced Packaging

AAPL: Apple

AC: Alternating Current

ACC: Active Copper Cable

ACIE: AI Clouds, Industrial, Enterprise

ACLS: Axcelis Technologies

ADAS: Advanced Driver-Assistance Systems

ADI: Analog Devices

AEIS: Advanced Energy Industries

AEC: Active Electrical Cable

AGI: Artificial General Intelligence; in Arm context, Arm-based AI infrastructure CPU opportunity

AI: Artificial Intelligence

AICAPEX: AI Capital Expenditures

ALAB: Astera Labs

AMAT: Applied Materials

AMBA: Ambarella

AMBQ: Ambiq Micro

AMD: Advanced Micro Devices

AMRS: Americas Research

AMT: Applied Materials; use AMAT as the preferred ticker unless AMT is intentional

AMX: Application Matrix Switch

ARM: Arm Holdings; also Arm processor architecture

ARR: Annualized Revenue Run Rate

ASIC: Application-Specific Integrated Circuit

ASML: ASML Holding

ASP: Average Selling Price

ATE: Automated Test Equipment

AWS: Amazon Web Services

AVGO: Broadcom

BMS: Battery Management System

BOM: Bill of Materials

bps: Basis Points

CAMT: Camtek

CAGR: Compound Annual Growth Rate

Capex: Capital Expenditures

CD: Critical Dimension

CDNS: Cadence Design Systems

CIS: CMOS Image Sensor

CM: Contract Manufacturer

CMOS: Complementary Metal-Oxide-Semiconductor

COGS: Cost of Goods Sold

COHR: Coherent

Cooper: Ambarella software platform / SDK

CoWoS: Chip-on-Wafer-on-Substrate; TSMC advanced packaging technology

COVID: Coronavirus Disease; used as shorthand for COVID-era supply-chain conditions

CPO: Co-Packaged Optics

CPU: Central Processing Unit

CPU-GPU: Central Processing Unit-to-Graphics Processing Unit

CRDO: Credo Technology

CSBG: Customer Support Business Group; Lam Research services/spares/upgrades business

CSS: Compute Subsystems; in Arm context, Arm’s higher-value subsystem IP offering

CSSv2: Second-generation Compute Subsystems

CXL: Compute Express Link

CW: Continuous Wave DAC:

Direct Attach Copper

DC: Data Center; also Direct Current in “800V DC” context

DC-DC: Direct Current-to-Direct Current power conversion

DDR: Double Data Rate memory

DDR4: Fourth-generation Double Data Rate memory

DDR5: Fifth-generation Double Data Rate memory

Dep: Deposition

DMS: Driver Monitoring System

DoE: Degree of Engagement

DPU: Data Processing Unit

DRAM: Dynamic Random-Access Memory

DSP: Digital Signal Processor

EBIT: Earnings Before Interest and Taxes

EBITDA: Earnings Before Interest, Taxes, Depreciation and Amortization

EDA: Electronic Design Automation

EDM: Electric Discharge Machining

ELS: External Laser Source

EMIB: Embedded Multi-die Interconnect Bridge

EML: Electro-absorption Modulated Laser

EPS: Earnings Per Share

ESS: Energy Storage System

EUV: Extreme Ultraviolet lithography

EV: Electric Vehicle

Fab: Semiconductor fabrication facility

FCF: Free Cash Flow

FinFET: Fin Field-Effect Transistor

FLOPs: Floating Point Operations per Second

FPO: Final Purchase Order

FPGA: Field-Programmable Gate Array

GAA: Gate-All-Around transistor architecture

GaN: Gallium Nitride

GB: Gigabyte

GB200: NVIDIA Grace Blackwell 200 system platform

GDP: Gross Domestic Product

GEM: Gross Equipment Margin

GenAI: Generative Artificial Intelligence

GFS: GlobalFoundries

GM: Gross Margin

GPU: Graphics Processing Unit

GTC: GPU Technology Conference / NVIDIA GTC

GW: Gigawatt

H6: Axcelis high-current implant platform

HBM: High Bandwidth Memory

HBM3e: Enhanced third-generation High Bandwidth Memory

HBM4: Fourth-generation High Bandwidth Memory

HDI: High-Density Interconnect

HoH: Half-on-Half

IC: Integrated Circuit

ICAPS: IoT, Communications, Automotive, Power and Sensors

IBC: Intermediate Bus Conversion / Integrated Bus Conversion, depending on power architecture context

I&D: Industrial and Defense

I&M: Industrial and Medical

IFX: Infineon Technologies

InP: Indium Phosphide

INTC: Intel

IoT: Internet of Things

IP: Intellectual Property; in semiconductor context, reusable chip design blocks or licensed technology

IR: Investor Relations

ISA: Instruction Set Architecture

IT: Information Technology

IVR: Integrated Voltage Regulator

JFET: Junction Field-Effect Transistor

KLAC: KLA Corp.

KV cache: Key-Value cache used in transformer inference workloads

L2: Level 2; typically automotive automation / ADAS context

L3: Level 3; typically automotive automation / ADAS context

Lab-to-Fab: Moving metrology/inspection techniques from lab environments into high-volume manufacturing fabs

LCA: Lifecycle Assessment

LEO: Low-Earth Orbit

LITE: Lumentum

LLM: Large Language Model

LM: Linear Modulator

L-MSD: Low- to Mid-Single Digit

LPO: Linear Pluggable Optics

LPDDR: Low-Power Double Data Rate memory

LPDDR5X: Enhanced fifth-generation Low-Power Double Data Rate memory

LPU: Language Processing Unit

LRCX: Lam Research

LRO: Linear Receive Optics / Linear Receiver Optics

LSCC: Lattice Semiconductor

LSD: Low Single Digit

LT: Long Term

LTA: Long-Term Agreement

LTAs: Long-Term Agreements

LV: Light Vehicle

M&A: Mergers and Acquisitions

MACOM: MACOM Technology Solutions

Mag-7: Magnificent Seven mega-cap technology companies

MBE: Molecular Beam Epitaxy

MCHP: Microchip Technology

MCU: Microcontroller Unit

MI450: AMD Instinct MI450 AI accelerator

MKSI: MKS Instruments

mn: Million

MOCVD: Metal-Organic Chemical Vapor Deposition

MPWR: Monolithic Power Systems

MRVL: Marvell Technology

MS&A: Mergers, Strategic investments and Acquisitions / M&A-style strategic activity

MSD: Mid-Single Digit

MSD%: Mid-Single-Digit Percent

MTSI: MACOM Technology Solutions

MU: Micron Technology

NAND: NAND flash memory

NAVI: Nova Measuring Instruments; likely should be NVMI in this report context

NavX: Advanced Energy semi power product platform

NB: Narrowband

NPO: Near-Packaged Optics

NVDA: NVIDIA

NVFP4: NVIDIA 4-bit floating-point format

NVL: NVIDIA NVLink rack/system family shorthand

NVL72: NVIDIA 72-GPU NVLink rack system

NVLink: NVIDIA high-speed GPU/CPU interconnect

NVMI: Nova Measuring Instruments / Nova

NXPI: NXP Semiconductors

OCS: Optical Circuit Switch

OEM: Original Equipment

Manufacturer ON: onsemi

OpEx: Operating Expenses

Opto: Optoelectronics

OSAT: Outsourced Semiconductor Assembly and Test

p.a.: Per Annum

P/B: Price-to-Book

PCB: Printed Circuit Board

PCIe: Peripheral Component Interconnect Express

PD: Photodiode

PF: Photonic Fabric

PIC: Photonic Integrated Circuit

PMI: Purchasing Managers' Index

PMIC: Power Management Integrated Circuit

PO: Price Objective

PPA: Power, Performance and Area

PSP: Preferred Supplier

Program QCOM: Qualcomm

QoQ: Quarter-over-Quarter

QTD: Quarter-to-Date

R&D: Research and Development

RDL: Redistribution Layer

RF: Radio Frequency

RISC-V: Open-standard Reduced Instruction Set Computer architecture

RL: Reinforcement Learning

ROI: Return on Investment

RPO: Remaining Performance Obligation

SAAR: Seasonally Adjusted Annual Rate

SaaS: Software as a Service

SAM: Serviceable Addressable Market

SATCOM: Satellite Communications

SDK: Software Development Kit

SD: Secure Digital; can also mean Single Digit depending on context

SDA: System Design and Analysis; used in Cadence context

SDRAM: Synchronous Dynamic Random-Access

Memory SerDes: Serializer/Deserializer

SiC: Silicon Carbide

SiPho: Silicon Photonics

SIMS: Secondary Ion Mass Spectrometry

SKU: Stock Keeping Unit

SLAB: Silicon Labs

SNPS: Synopsys SoC

: System-on-Chip

SPOT: Subthreshold Power Optimized Technology; Ambiq ’s ultra-low-power platform

SRAM: Static Random-Access Memory

SSD: Solid-State Drive

STM: STMicroelectronics

STMicro: STMicroelectronics

SWKS: Skyworks Solutions

TAM: Total Addressable

Market

TB: Terabyte

TCO: Total Cost of Ownership

TER: Teradyne

TIA: Transimpedance Amplifier

Tn: Trillion

ToR: Top-of-Rack switch

TSMC: Taiwan Semiconductor Manufacturing Company

TSV: Through-Silicon Via

TXN: Texas Instruments

UALink: Ultra Accelerator Link / Universal Accelerator Link; open accelerator interconnect standard

UCIe: Universal Chiplet Interconnect Express

UHP: Ultra-High Power / Ultra-High Performance

v9: Armv9 architecture generation

vGaN: Vertical Gallium Nitride

WFE: Wafer Fab Equipment

WOLF: Wolfspeed

Xcelium: Cadence simulation/verification platform

x86: Intel/AMD processor instruction set architecture

XPO: Generic shorthand for optical packaging approaches such as CPO/NPO and related architectures

XPS: X-ray Photoelectron Spectroscopy

XPU: Generic accelerator/processor category, often encompassing CPU, GPU, ASIC, FPGA and other custom compute

XRF: X-ray Fluorescence

YMTC: Yangtze Memory Technologies

YoY: Year-over-Year

YTD: Year-to-Date

ZF Optics: Credo ZF Optics platform; used as a branded optical connectivity platform

Price objective basis & risk

NVIDIA Corporation (NVDA)

Our \$350 PO is based on 26x CY27E PE ex cash, within NVDA’s historical 25x-56x forward year PE range, which we believe is justified by NVDA's leading share in fast-growing AI compute/networking markets, offset by lumpiness in global AI projects, cyclical gaming market, and concerns around access to power.

Downside risks are: 1) weakness in consumer driven gaming market, 2) Competition with major public firms, internal cloud projects and other private companies in AI and