AMERICAS TECHNOLOGY: HARDWARE & NETWORKING

Communacopia + Technology 2026 — Key Takeaways

Michael Ng, CFA
+1(212)902-8618 | michael.ng@gs.com
Goldman Sachs & Co. LLC

Katherine Murphy
+1(212)902-1151 | katherine.a.murphy@gs.com
Goldman Sachs & Co. LLC

Yash Goenka, CFA
+1(212)934-6312 | yash.goenka@gs.com
Goldman Sachs India SPL

Lindsey Shema
+1(212)902-6506 | lindsey.shema@gs.com
Goldman Sachs & Co. LLC



We hosted 13 companies from the Hardware and Networking sector at the Goldman Sachs Communacopia + Technology Conference in San Francisco on September 8-11. In this note, we discuss key themes and takeaways emerging from the conference, where we hosted hardware and commtech companies across our coverage. Company sections are listed in order of appearance.

Hosted Companies: Arista Networks, Inc. (ANET), **Axon Enterprise, Inc.** (AXON), **NetApp, Inc** (NTAP), **Cisco Systems, Inc** (CSCO), **HP, Inc** (HPQ), **Celestica, Inc** (CLS), **Ingram Micro** (INGM), **F5, Inc** (FFIV), **Dell Technologies, Inc** (DELL), **Hewlett Packard Enterprise Co** (HPE), **Super Micro Computer, Inc** (SMCI), **Penguin Solutions, Inc** (PENG), and **Cricut, Inc** (CRCT).

Key Takeaways: Hardware & CommTech

AI-driven networking demand remains robust. Companies across our coverage are experiencing robust, multi-year demand tailwinds catalyzed by AI workloads. CSCO highlighted that the shift toward agentic AI is driving a 5x increase in network bandwidth consumption, underpinning a durable networking supercycle across hyperscale, neocloud, and enterprise customers. The company reiterated its fiscal 2027 outlook, which includes 15% revenue growth at the midpoint with $7.5 bn in hyperscale AI revenue supported by $9.3 bn in fiscal 2026 hyperscaler AI orders. Similarly, ANET is observing accelerating demand across complex back-end and scale-across deployments, as power, cooling, and space constraints force customers to expand clusters across multiple locations. Furthermore, ANET views scale-up Ethernet as a significant medium-term growth vector as the market transitions toward open standards, anticipating material volume by early 2028. Additionally, CLS noted that to support this massive AI-driven compute and networking demand, it is actively expanding capacity. CLS noted that its networking portfolio is rapidly transitioning to higher speeds, with 1.6T programs coming online in 4Q26 and

Goldman Sachs does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. For Reg AC certification and other important disclosures, see the Disclosure Appendix, or go to www.gs.com/research/hedge.html. Analysts employed by non-US affiliates are not registered/qualified as research analysts with FINRA in the U.S.

expected to surpass 800G revenue in 2027. For FFIV, AI is driving both indirect demand through increased bandwidth requirements for traditional ADCs and direct use cases. The company is actively monetizing these direct AI opportunities, which drove roughly $50 million of bookings in the first half of the year across AI data delivery, AI runtime security, and AI factory load balancing.

Margin resilience and operating leverage despite lingering supply chain constraints. Despite ongoing supply chain challenges and inflationary pressures, networking companies are successfully defending margins through operating leverage, pricing actions, and strategic business models. ANET maintained its 62% to 64% gross margin guidance for 2026, utilizing targeted price increases and supply chain efficiencies to offset higher component costs, while leveraging $9.7 bn in multiyear purchase commitments to secure long-lead components. CLS similarly committed operating margins to expand into the mid-to-high 8% range over the next 1-2 years, driven by operating leverage and productivity gains, even as supply constraints persist for memory and Tomahawk silicon. For CSCO, while the company is navigating slight near-term gross margin compression as it leans into hardware-heavy hyperscale and enterprise AI deployments, management emphasized record operating efficiency, noting that a recent decline in gross margins was more than offset by a nearly 4 point reduction in OpEx margins as a percentage of revenue. Similarly, FFIV reiterated its fiscal 2027 gross margin guidance of 80% to 82%, expecting long-term stabilization as component availability improves and strategic inventory builds mitigate mixed memory pricing signals. Additionally, FFIV anticipates continued operating margin expansion by leveraging automation and AI to scale expenses efficiently.

Embedded security emerges as a critical differentiator. As AI deployments and hybrid multi-cloud architectures introduce significant complexity, vendors are increasingly differentiating through deeply integrated security solutions. CSCO is fusing security directly into the network fabric by integrating Data Processing Units (DPUs) and Neural Processing Units (NPUs) onto top-of-rack smart switches, significantly reducing latency and embedding security to monitor behavioral drift and secure agentic AI at runtime. FFIV is similarly expanding its security portfolio to address the growing threat landscape at the front door of mission-critical applications. FFIV stated that supported by its recent acquisition of SurePath AI, it is adding guardrails, red teaming, and discovery capabilities to its AI-powered WAF, driving momentum within highly regulated sectors such as financial services and government agencies.

Enterprise campus refresh driven by AI demand and infrastructure modernization. We observe a healthy enterprise demand environment characterized by: AI-driven capacity expansion and traditional infrastructure refresh. CSCO noted that agentic AI is driving a 5x increase in network bandwidth consumption, accelerating campus networking growth from historical 3% to 4% levels to double-digits. This AI momentum is further underpinned by a broader >$100 billion end-of-life infrastructure refresh opportunity over the next few years. ANET is successfully capitalizing on this enterprise modernization trend by extending its data center networking expertise into the campus market. Utilizing its unified EOS software stack, ANET is executing a dual-pronged go-to-market motion: cross-selling into existing data-center customers and securing "campus-first" wins against legacy incumbents. Concurrently, FFIV is experiencing a "refresh plus expansion" dynamic within its Systems segment. Customers are replacing legacy infrastructure and upgrading to higher-end appliances to support AI-driven

workload growth, further accelerated by the upcoming end of software support for the iSeries in F2Q27.

Enterprise AI hardware demand remains robust and is broadening beyond initial deployments. (1) NTAP highlighted that demand in its most recent quarter was stronger than expected and expanding beyond dedicated AI storage into broader database, analytics, streaming, and data infrastructure upgrades, with customers increasingly moving from pilot projects to production AI deployments, supporting its raised FY27 outlook. **(2) HPQ**'s PC portfolio is skewing more toward higher-end commercial devices, with customers increasingly favoring hybrid AI architectures and local AI-capable PCs rather than cloud-only computing because of more favorable token costs & security/privacy. **(3) INGM** is seeing robust GPU and AI infrastructure project demand, as well as emerging cloud infrastructure-as-a-service demand from SMBs. **(4) DELL** characterized AI adoption as still being in the early innings, with most enterprises only beginning to move beyond productivity applications toward rearchitecting full business workflows, supported by more than 6,500 AI Factory customers. **(5) HPE** pointed to accelerating enterprise demand for inference-oriented computing driven by token-cost optimization, data sovereignty, and datacenter refresh cycles. **(6) PENG** highlighted growing AI Factory deployments for enterprise customers, particularly around coding assistants and other inference applications. **(7) SMCI** likewise emphasized growing momentum in enterprise and sovereign markets as customers increasingly seek fully integrated AI solutions rather than standalone hardware.

Operational expertise, deployment speed, and services driving durability within the Neocloud opportunity. (1) DELL highlighted its ability to accelerate time-to-first-token for neocloud customers and simplify deployment through services and financing. **(2) HPE** cited services, installation, and lifecycle management as key differentiators in recent enterprise wins, including its gigawatt scale Oracle deal. **(3) SMCI** emphasized that reducing deployment times can create meaningful economic value for customers by accelerating revenue generation. **(4) PENG**'s design-build-deploy-manage model and 3-5 year services contracts are increasingly helping it win AI Factory projects.

Hardware vendors are helping customers navigate component cost increases through the breadth of their portfolios. (1) NTAP is helping customers navigate flash prices increases through sophisticated tiered-storage architectures where customers can optimize data storage across flash and disk storage. **(2) HPQ**'s Workforce Experience Platform (WXP) gives customers operational insights that helps them configure their equipment purchases specifically to the needs of its employees. **(3) INGM** noted that memory-driven component inflation continues to lift ASPs across the industry, but the company's scale, sourcing capabilities, and broad vendor portfolio allow it to secure supply, offer alternative products, and pass higher costs through to customers rather than absorb them. **(4) DELL** highlighted advantages from its scale, supply-chain relationships, financing capabilities, and portfolio spanning PCs, servers, storage, networking, and services, arguing this breadth helps customers navigate both component shortages and increasingly complex AI deployments. **(5) HPE** similarly emphasized its ability to combine networking, compute, storage, services, and financing into a single solution. **(6) PENG** highlighted its AI Factory platform, which combines ClusterWareAI, MemoryAI, ComputeAI, and managed services to optimize infrastructure performance and economics, particularly for inference workloads. **(7) SMCI**'s DCBBS strategy bundles power, cooling, cabling, software, and services to improve

time-to-value.

Arista Networks, Inc. (ANET)

Presenters: ANET's President & CTO, Ken Duda, and CFO, Chantelle Breithaupt, presented at the Communacopia + Technology Conference.

Bottom Line: ANET reiterated its 2026 revenue and margin guidance, supported by improved supply chain visibility and demand across its two strategic focus areas: AI infrastructure and enterprise campus deployments. First, Arista is leveraging its best-in-breed data center networking capabilities to address demand across front-end and back-end including scale-out and scale-across deployments. Second, the company is successfully extending its data center expertise into the enterprise campus market, utilizing its unified EOS software stack to capture its fair share within campus. Third, price increases and supply chain efficiencies could help Arista achieve similar gross margins in 2027E, despite lapping tariff refunds (30 bps in 2026E).

Key Takeaways

Reiterated revenue guidance and stable margin trajectory supported by supply chain visibility. Management noted that ANET raised its 2026E revenue guidance in 2Q26 to $12.6 billion, aided by easing supply chain constraints and improved demand visibility. Management maintained its gross margin guidance of 62% to 64% despite inflation and mix shifts, helped partly by ~30 bps of tariff refunds this year and targeted price increases intended to offset higher component costs. Furthermore, large multiyear purchase commitments of $9.7 billion in 2Q26 reflect robust demand rather than just component inflation, and Arista is willing to lean into long-lead components because its portfolio flexibility and customer roadmap visibility limit obsolescence risk.

ANET differentiated by EOS software quality and single-stack architecture. First, ANET's EOS operating system is unique in the industry as it runs as a single image across ANET's entire portfolio of products (v. different OS for different products at competitors). Second, EOS delivers high reliability and consistency across hyperscaler routing, neoclouds, and enterprise environments. For neoclouds and enterprises, the value proposition centers on the manageability and quality of a single software stack across both campus and cloud use cases. Third, operating at scale mandates best-in-breed solutions, allowing ANET to win against turnkey alternatives that are typically relegated to low-volume customers.

AI networking demand accelerating across complex back-end and scale-across deployments. ANET is seeing strong demand across its AI networking domains, particularly in back-end networks that connect GPU accelerators. Scale-across deployments are seeing significant demand as customers face power, cooling, and data center space constraints, requiring them to expand clusters across multiple locations. Additionally, management noted that deferred revenue tied to newer AI deployments may continue to convert unevenly; recognition depends on customer acceptance and performance validation, and deployment timing can be delayed by data-center readiness issues such as construction, power, cooling, and cabling. Finally, while neocloud customers are an attractive growth segment that values open, best-of-breed

AI networking, management noted that Arista is monitoring customer credit risk.

Scale-up Ethernet represents a new growth vector for the medium term. While scale-up Ethernet is still an early opportunity, Arista expects the market transition away from proprietary interconnects toward open standards to create a significant new growth vector. The company anticipates this market to begin ramping in the later part of 2027, with more material volume expected in early 2028.

Opportunity to extend ANET's best-in-breed solutions into the enterprise campus. ANET is successfully extending its data center networking expertise into the enterprise campus market, representing a significant growth opportunity. The value proposition remains consistent with its data center offerings, providing quality and the same EOS software and cloud vision management stack built by the same engineering team. This growth is supported by a dual-pronged go-to-market motion: successful cross-selling into existing data-center customers and a growing number of "campus-first" wins driven by customers who are dissatisfied with legacy incumbents.

Valuation and key risks

Valuation. We are Buy-rated on ANET with a 12-month target price of $225 based on 36X our NTM+1Y EPS.

Key downside risks include slower cloud capex spending; customer concentration, with META and MSFT representing 16% and 26% of total revenue in 2025 respectively; competition, including from Cisco, major Chinese providers Huawei & H3C and whitebox switching solutions; margin degradation from investment in enterprise sales force, elevated costs associated with supply chain headwinds; broader pricing pressure from commoditization of networking hardware; inability to execute on enterprise/campus strategy & development of new SD-WAN platform.

Investment thesis summary. As the leading branded provider of switches to US hyperscalers, ANET is well positioned to capitalize on the ongoing growth in data, the continued digital transformation driving workloads from on-premise to public and hybrid-cloud, and the growing demand for higher bandwidth, faster speed, and lower latency. ANET has the most revenue exposure to cloud spending among our coverage universe with 48% of its 2025 revenue from its Cloud Titans segment (META was 16%, MSFT was 26%) and 20% from tier 2 cloud and service providers. In addition, ANET is leveraging its dominant position in data center switches to other networking solutions, particularly in enterprise (e.g., campus switching, wireless, routing, telemetry).

ANET should deliver strong double-digit revenue and EPS growth over the several years reflecting strong growth in data center and expansion into enterprise, as well as continued R&D and S&M investments particularly as ANET builds out its enterprise go-to-market capabilities.

Axon Enterprise, Inc. (AXON)

Presenters: AXON's Chief Financial Officer & Chief Operating Officer, Brittany Bagley, presented at the Communacopia + Technology Conference.

Bottom Line: Axon is demonstrating strong business momentum across its interconnected ecosystem, reiterating full-year 2026 revenue guidance of 32-34%. AXON reiterated that 2H26 revenue growth will be weighted toward 4Q, similar to last year. The company remains underpenetrated in its ~$160 billion TAM, capturing only 1-5% of police budgets today, providing a long runway for growth as it increases share of wallet through continuous innovation and value-add pricing.

Key Takeaways

Robust bookings momentum and expanding share of wallet: Management noted that AXON introduced a new five-year normalized bookings metric in 2Q26, normalized for contract duration, which grew over 30% in 2Q26. Gross bookings grew 20% yoy in 2Q26, comping a 50% growth rate in the prior year. Management expects normalized bookings to grow 30%-plus for the full year, with future contracted bookings growing above that rate. AXON currently captures roughly 1% to 5% of state and local police budgets (averaging ~3%). The company sees runway to increase its share of wallet by solving mission-critical problems and driving departmental ROI.

Revenue and gross margin dynamics: Axon reiterated its full-year 2026 revenue guidance to 32-34% yoy, supported by strong future contracted bookings and a fourth quarter that is expected to be the highest growth period of the year. On margins management noted that Software-only margins remain steady at over 80%, with any fluctuations in the broader segment driven by the mix and deployment costs of services across its various products. Furthermore, within Connected Devices, management noted that margins are influenced by three primary drivers: First, TASER margins remain stable with no material changes. Second, camera margins are facing headwinds from increased memory costs. Third, the rapid scaling of the Dedrone requires upfront hardware investments that act as a near-term drag on connected device margins. However, management emphasized this is not a permanent headwind as they optimize scaling efficiencies over time.

Core products remain strong alongside rapid AI software adoption: AXON's core TASER business remains highly resilient, highlighted by the largest individual TASER order in company history in 2Q26 - an upgrade to TASER 10 occurring three-plus years into the refresh cycle. Meanwhile, software adoption is accelerating, with non-evidence software such as 911, Fusus, and the AI Era Plan growing. The AI Era Plan is a standout, that includes high-ROI tools like Draft One, which saves officers up to 4 hours per week, as well as Form One and Translation. AXON views the subscription bundle as future-proofing AI for departments, allowing them to seamlessly adopt new features as they are released.

Opportunity in nascent markets including Enterprise, International, and Counter-Drone: AXON is seeing significant traction across adjacent, more nascent markets. Enterprise bookings are being driven by trials and phased rollouts of the Axon Body Workforce and the newly launched ABW Mini, alongside the Axon 911 ecosystem - leveraging the Carbyne and Prepared acquisitions to deliver sticky, cloud-native AI overlays that seamlessly integrate with legacy CAD providers. International bookings also tripled year-over-year in 2Q26, supported by TASER 10 adoption and large AI Era Plan deals (three of the top five AI Era Plan deals last quarter were international). Furthermore, management noted that the Counter-Drone business (Dedrone) is emerging as a massive opportunity, as drone mitigation becomes a standard

requirement for major events and federal agencies.

Valuation and key risks

Valuation. We are Buy-rated on AXON with a 12-month target price of $715 based on 17X our NTM+1Y EV/Gross Profit.

Key risks: patent expiry, anti-trust lawsuits, new product innovation, slower-than-expected international expansion, changes in TASER regulation, changes in public perception, government budget volatility, and near-term supply chain constraints.

Investment thesis summary. AXON is the manufacturer of the TASER, the leading conducted energy weapon in the United States and globally. AXON's market leadership with TASER has resulted in deep customer relationships with state and local law enforcement (e.g., 17k of 18k US law enforcement agencies) and a highly valuable installed base of law enforcement personnel. In addition to opportunities for TASER upgrades and installed base expansion, AXON is well positioned to cross-sell sensors, such as body cameras and fleet cameras, and related software and cloud storage and analytics services. Growing demand for body cameras, fleet cameras, and less-lethal weapons against the backdrop of initiatives to improve accountability and effectiveness in public safety and alongside AXON's 10+ year future contracted revenue provides visibility into revenue growth. In addition, AXON should benefit from the expansion into new geographies, new customer segments (e.g., federal, consumer), and adjacent products and services.

NetApp, Inc (NTAP)

Presenters: NTAP's Chief Executive Officer, George Kurian, presented at the Communacopia + Technology Conference.

Bottom Line: NTAP reiterated its F2027 guidance (+17% revenue, +22% EPS growth at midpoint) driven by strong, broad-based demand, including both AI-specific configurations and non-mission critical upgrades. First, Product segment gross margin outlook is improved (v. initial F2027 guidance) as the company increases visibility into supply & takes pricing actions, though strong growth in Product may weigh on overall corporate gross margins. Second, NTAP is bullish on Public Cloud growth remaining strong in the high teens to low 20s% growth rate. Third, NTAP will introduce new products addressing the needs of high end neoclouds at its INSIGHT event later this month, which could expand the company's addressable market between enterprise storage buyers.

Key Takeaways

Product margin outlook better than 90 days ago, though mix (i.e. strong Product growth) may weigh on corporate gross margins. NTAP's raised its Product gross margin expectations on earnings relative to its initial FY27 guidance reflecting better cost visibility and more confidence in recovering component inflation through pricing

actions. The company is helping customers navigate prices increases by (1) optimizing mix of flash/disk storage; and (2) optimizing compute usage across cloud and on-prem. NTAP works with suppliers across a broad range of commercial vehicles including buying in the open market as well as through long-term agreements to assure access to available supply at competitive cost structures.

Public Cloud momentum should remain robust. The company remains constructive on sustaining high teens-low 20s% growth in the Public Cloud segment from (1) secular growth in hyperscalers; (2) broader product offerings; and (3) go-to-market improvements. The company views AI use cases as an incremental driver of demand as enterprises look to utilize data across hybrid environments.

Neocloud use cases could expand NTAP's TAM. NetApp expects to make further product announcements at its annual INSIGHT event later this month around new products addressing the needs of higher-end neoclouds and frontier-model training environments. The company attributed its right-to-win with neocloud buyers to: (1) track record of hyperscale level data management infrastructure; (2) go-to-market models that supports collaboration between vendors & sales; and (3) hybrid/sovereign capabilities.

Valuation and key risks

Valuation. We are Buy rated on NTAP with a 12-month target price of $210 reflecting 19x our NTM+1Y EPS.

Key risks: (-) Key supply shortages or pricing increases that weigh on product gross margins; (-) structural declines in demand for on-premise enterprise storage appliances; (-) increased competitiveness of storage market from legacy players & new-entrants, particular in growth areas like parallel file systems; (-) evolving demand for as-a-service offerings like Keystone in a changing interest rate environment; (-) impact of global macro weakness on demand for storage

Cisco Systems, Inc (CSCO)

Presenters: CSCO's CFO, Mark Patterson, and President & Chief Product Officer, Jeetu Patel, presented at the Communacopia + Technology Conference.

Bottom Line: CSCO is well positioned to capitalize on a multi-year, AI-driven networking supercycle, with demand across all customer types and its entire product portfolio. AI demand is driving tailwinds across hyperscale, neocloud, and enterprise, as well as data center switching, routing, security, observability, and campus networking. CSCO reiterated its F2027 revenue outlook, which includes 15% revenue growth at the midpoint with $7.5 bn of hyperscale AI revenue.

Key Takeaways

Agentic AI driving a multi-year networking supercycle: Management highlighted a secular shift driven by AI agents. Agentic AI is driving a 5x increase in network bandwidth consumption, catalyzing a durable networking demand cycle across all of

CSCO's customer segments including hyperscale, neocloud, and enterprise, as well as its product segments. For instance, this includes back-end and front-end networking, security & observability, as well as campus networking, where growth has accelerated from historical 3-4% levels to double-digits. Physical AI likely will drive an additional vector of growth.

Reiterated fiscal 2027 guidance driven by core refresh and AI momentum: CSCO's fiscal 2027 guidance implies 15% total revenue growth at the midpoint, supported by an estimated double-digit ex-AI baseline growth. This acceleration is underpinned by a >$100 billion end-of-life/end-of-support infrastructure refresh opportunity over the next few years.

Hyperscaler AI orders scale rapidly alongside vertical integration strategy: CSCO secured $9.3 billion in fiscal 2026 hyperscaler AI orders, an acceleration from nearly zero just two years prior. The company is leveraging a vertically integrated, co-designed full stack, deploying scale-out systems powered by the G300 Silicon One chip and scale-across P-Series chips. Management noted that by 2029, CSCO will be completely independent of merchant silicon providers.

Navigating near-term margin compression offset by operating leverage while maintaining robust shareholder returns: CSCO's F1Q27E gross margins guidance of 65%-66% reflects a slight headwind as the company leans into hardware-heavy hyperscale and enterprise AI deployments where revenue is recognized on day one, while attached software and services are recognized ratably. However, management emphasized record operating efficiency, noting that in F4Q26, gross margins declined ~2 points year-over-year, but OpEx margins declined by almost 4 points as a percentage of revenue, driving record operating margins.

Fusing security into the network fabric to capture machine-scale defense: CSCO reiterated F2027 security revenue growth to accelerate to high single-digits yoy, driven by the lapping of Splunk's cloud transition, >30% firewall growth over the past two quarters, and new product customers momentum. The company is leveraging its unique position to integrate Data Processing Units (DPUs) and Neural Processing Units (NPUs) directly onto top-of-rack smart switches, significantly reducing latency and embedding security into the network architecture to monitor behavioral drift and secure agentic AI at runtime.

Strategic partnerships expanding the Secure AI Factory: CSCO is simplifying the deployment of on-premises AI workloads through its expanded Secure AI Factory partnership with NVIDIA and Super Micro. By integrating Super Micro's rack-scale compute capabilities, CSCO can deliver full-stack solutions (server, networking, security, observability, and data) with favorable net margin treatment, avoiding the historical margin headwinds of the compute business while driving significant cross-sell potential across its proprietary technologies.

Valuation and key risks

Valuation. We are Neutral-rated on CSCO with a 12-month target price of $125 based on 25x our NTM+1Y (Q5-Q8) EPS.

Key upside risks include secular tailwinds including hybrid work, multi-cloud network architecture adoption, broader roll-out of WiFi 6/6E and 5G, increasing edge compute use cases; New disaggregated consumption models to target previously underserved cloud providers, near-term revenue visibility from elevated backlog; **downside risks** include competition, including from major Chinese providers Huawei & H3C and White Box solutions; margin degradation from mix shift into more cloud customers, elevated costs associated with supply chain headwinds; broader pricing pressure from commoditization of networking hardware; dilutive acquisitions.

Investment thesis summary. CSCO is the global market share leader of networking equipment and solutions with a comprehensive platform of offerings across networking, security, collaboration, applications, and cloud. As a leading #1 or #2 incumbent in nearly every product type and category it participates in, CSCO has been subject to market share losses over the last decade, particularly due to lower cost competitors (e.g., White Box) and smaller, more nimble companies with specific vertical solutions (e.g., ANET in data center switching, FFIV in ADC). That said, CSCO has the widest offering of networking solutions among any of its key competitors, cementing its role as a trusted networking solutions provider, as well as a key procurement partner for its suppliers. Product backlog, RPO, and ARR provide ample revenue visibility even if order growth slows in the near term as customers absorb previously placed orders & more closely scrutinize IT budgets in a challenging macro environment.

HP, Inc (HPQ)

Presenters: HPQ's Chief Financial Officer, Karen Parkhill, presented at the Communacopia + Technology Conference.

Bottom Line: HPQ re-affirmed its expectations for Personal Systems margins to trough in F4Q before recovering sequentially back into the company's long-term range (5-7%) in F2027. HPQ continues to execute against its 4-pillar strategy to mitigate the impacts of input costs on margins by (1) securing supply; (2) portfolio mix shift (i.e., more premium); (3) operational levers like repricing LT enterprise contracts and product redesign. The company is also prioritizing growth in Print in key areas including Big Tank (low-to-mid 20s% market share in big tank v. 30s% share in overall print) and industrial.

Key Takeaways

PC portfolio mix-shifting towards premium categories. HPQ's ~40% year-over-year increase in ASPs last quarter was largely driven by mix-shift towards more premium PC categories like AI PCs and workstations. Demand for AI PCs is no longer just driven by future-proofing, but by enterprises looking for an alternative to conducting AI entirely in the cloud because of the (1) significant risk in token costs; (2) higher security/privacy from running AI locally. HPQ is also seeing strength in its peripherals & services offering including its Workforce Experience Platform (WXP) where customers can monitor and optimize their equipment.

Personal Systems margins to recover from 4Q trough based on (1) repricing long-term enterprise contracts (subset of HPQ's ~70% commercial business), (2) redesigning products to lower memory usage (i.e., de-speccing), and (3) continued mix towards

premium PC categories and peripherals, which should benefit ASPs. HPQ has secured the supply required to reach F2026 guidance and is working towards securing supply for F2027, including through qualifying new suppliers like CXMT, which can address HPQ's PC portfolio in China and some other countries.

Prioritizing high-value & growth areas in print including Big Tank, industrial. First, HPQ gained ~4 points of share in Big Tank in F3Q as the company more aggressively goes after the market. Big Tank printers represented ~1/3 of the 17 mn global printers sold last quarter. HPQ is currently under indexed in Big Tank relative to the rest of its portfolio, with low-to-mid 20s% market share in Big Tank as compares to low-to-mid 30s% market share overall. Second, HPQ's core competencies in home & office print should position the company to go after the industrial print opportunity as (1) the market transitions from analog to digital; and (2) AI-enabled workflow management tools like HP Neo add value to large enterprise customers like Shutterfly.

Most remaining tariff benefits should be recognized in F4Q26, thought some benefit may fall into F1Q27.

Valuation and key risks

Valuation. We are Sell rated on HPQ with a 12-month target price of $22 based on 6.5x our NTM+1Y EPS.

Key risks: (+) Better-than-expected industry PC demand; (+) Better-than-feared impact of higher input costs on margins; (+) Market share gains in PC, especially in higher-margin categories like gaming PCs and peripherals; (+) Recovery in office and consumer industry demand; (+) Print portfolio shift towards growth opportunities, such as Big Tank (profit up front model), consumer subscriptions (e.g., Instant Ink, Instant Paper, HP+), and industrial print.

Celestica, Inc (CLS)

Presenters: CLS's CFO, Mandeep Chawla, presented at the Communacopia + Technology Conference.

Bottom Line: CLS has successfully transitioned from a traditional EMS provider to a higher-value ODM Plus model, positioning itself as a strategic co-development partner to capitalize on AI-driven networking and compute demand. Management reiterated confidence in accelerating revenue growth through 2027 driven by core networking demand and major program wins with OpenAI and AMD, and expects operating margins to expand into 2027E and beyond, supported by operating leverage, recent capacity expansions, and a healthy balance sheet despite ongoing supply chain constraints.

Key Takeaways

Accelerating revenue growth supported by robust demand and new program wins. Management reiterated its 2026 revenue growth guidance of 65% yoy, with expectations for further acceleration in 2027 driven by strength across both networking and compute. To support this demand, CLS is actively bringing new capacity online

across Texas, Thailand, and Malaysia, with growth driven by: (1) strong networking demand from three large hyperscalers ; (2) steady 400G scale-up demand alongside scaling 800G programs; (3) accelerating 1.6T ramps expected to surpass 800G revenue; (4) mass production for OpenAI Jalapeno accelerator in 2027; (5) the ramp of CLS's scale-up networking switch for AMD's Helios rack; and (6) major next-generation program wins on the compute side.

Networking portfolio transitioning rapidly to 800G and 1.6T. CLS is experiencing volume growth in high-speed networking driven by massive AI workloads. While 400G demand remains robust and is expected to grow yoy in 2026 (and remain flat in 2027), 800G revenue is expanding and will continue growing in 2027. Furthermore, the transition to 1.6T is accelerating, with 10 programs won and starting to come online in 4Q26; management expects 1.6T revenue to ramp at a rate that will surpass 800G revenue in 2027. Regarding co-packaged optics (CPO) switch adoption, management anticipates mass adoption alongside 3.2T switches, though it is not expected to be a meaningful contributor until 2029-2030.

Strategic capital allocation and equity raise to fund multi-year growth. To support demand visibility extending into 2029, CLS executed an equity raise last month. Management reiterated its capital expenditures guidance of $1 billion for the current year and expects capex to increase to the high-end of $1.5 billion to $2 billion in 2027.

OpenAI consignment model progressing with a path to strong long-term margins. CLS is currently the sole partner for its current OpenAI custom ASIC rack-scale program, operating under a consignment model where silicon is provided. While management noted that compute margins are currently dilutive compared to accretive networking margins, management remains comfortable with the profitability profile. As a multi-generation partner, CLS expects additional opportunities over time as it works on next-generation designs and adds services revenue.

Operating margin expansion expected despite ongoing supply chain constraints. Management sees room for operating margins to expand from the low 8% range in 2Q26 to the mid-to-high 8% range in the next 1-2 years. This expansion is driven by operating leverage, productivity gains on R&D spend, and margin opportunities within the Advanced Technology Solutions (ATS) segment. Supply constraints persist, particularly for memory (which is sold out for next year) and Tomahawk silicon (which currently has lead times extending well over a year).

Valuation and key risks

Valuation. We are Buy rated on CLS with a 12-month target price of $505, reflecting 18X our NTM+1Y EPS.

Key downside risks include: (-) continued competition from branded data center equipment OEMs and ODMs, which could pressure market share; (-) AI server market competitive intensity, which can pressure AI server margins; (-) high customer concentration could cause revenue volatility; (-) potential declines in data center investment demand; (-) reductions in CLS's ability to win higher margin design-deals, which could limit CLS' ability to grow EBIAT margins over time; and (-) volatility in end-market demand within CLS' ATS segment.

Celestica supplies, designs, and manufactures IT hardware and industrial tech solutions for hyperscale and OEM customers. CLS delivers solutions under two segments: (1) Connectivity and Cloud Solutions (CCS, ~75% of 2025 revenue), which comprises both Communications (Networking, ~58% of total revenue), and Enterprise (Servers and Storage, 17% of 2025 revenue); and (2) Advanced Technology Solutions (ATS, 26% of 2025 revenue), which delivers various industrial solutions across its subverticals including (a) Industrial & Smart Energy; (b) Aerospace & Defense; (3) Capital Equipment; and (4) Health Tech. We view Celestica as a winner amongst growing AI data center build deployments due to their competitive advantage against other ODMs and EMS players in engineering, as well as the company's ability to quickly and reliably scale production. As such, Celestica should benefit from (a) ongoing growth amongst hyperscale and large tier-2 cloud data center spend, where the company has demonstrated its ability to win customer programs; and (b) as Celestica gains traction within higher-margin ODM server business, relative to its traditional position as an EMS server player.

Ingram Micro (INGM)

Presenters: INGM's CFO, Michael Zilis, presented at the Communacopia + Technology Conference.

Bottom Line: (1) AI-related demand driven both by GPU and AI infrastructure deals (doubled YoY) and emerging investment from SMB in infrastructure-as-a-service; (2) INGM's proprietary xVantage platform is helping the company to grow above market with margin expansion opportunities in the future; (3) INGM announced that its primary sponsor Platinum Equity is to sell 13.1 mn shares, its third follow-on offering of 2026; and (4) at its analyst day next week, INGM will provide a financial update with long term expectations.

Key Takeaways

AI-related demand driven both by (1) GPU and AI infrastructure deals; and (2) emerging investment from SMBs, largely on the infrastructure-as-a-service side. First on GPU and AI infrastructure deals, INGM is well positioned to address this net-new vector of demand from large enterprise/neoclouds because of its global scale and reach. While these fulfillment-type deals are gross margin rate dilutive, they carry a very low cost to serve and are accretive to operating profit margins & ROIC. Second, SMBs are starting to ramp their investments in AI-ready capabilities primarily through IaaS offerings, though it is still early days (driving 44% growth in Cloud segment last quarter). Cloud as <1% of revenue but high-teens share of profit dollars in the 1H26.

Proprietary xVantage platform is helping INGM grow above market. Through xVantage, INGM has been able to remove ~$200 mn in annualized opex from its business (late 2023-2025) and grow above market (particularly in regions like APAC) driven by differentiated capabilities in automation. Over time, INGM expects xVa
gross margin accretion through dynamic pricing models, demand optim
learning.

Further refresh opportunity in the PC market. INGM sees further run
refresh opportunity, with an estimated ~30% of the global installed bas



running aged equipment. INGM expects that growing adoption of AI PCs (~30% of PC sales) could smooth some of the typical seasonality in the market.

Continuing to execute against capital allocation strategy, which prioritizes investment into the business (both organic and inorganic), opportunistically paying down debt, and shareholder returns (dividends, eventually buybacks). INGM announced that its primary sponsor Platinum Equity is to repurchase 13.1 mn shares, its third follow-on offering of 2026.

Analyst day planned for September 15. At its analyst day next week, INGM expects to provide a financial update with long term expectations.

Valuation and key risks

Valuation. We are Neutral rated on INGM with a 12-month target price of $33 based on 9x our NTM+1Y EPS.

Key risks: (-) OEMs may shift more of their sales to direct-to-end user or 1-tier distribution models, altogether bypassing the distributor layer in which Ingram Micro operates; (-) INGM's gross margins are low, magnifying the impact of supply chain disruption and price competition on operating income; (-) Ability to execute on capital allocation strategy, where INGM's capital allocation strategy is dependent on tuck-in M&A, maintaining a favorable net leverage ratio, and dividends, and INGM may be challenged in identifying M&A targets in line with their targets (e.g., cybersecurity, cloud, deeper geographic presence in market like Japan); (-) INGM's ability to delever may be at risk if the company is unable to generate sufficient free cash flow (e.g., lower levels of IT spend, market share loss, margin compression from business mix shift); (+) Upgraded outlook from credit rating agencies could drive down interest expense, ultimately benefiting profitability & shareholder returns; (+) Better-than-expected IT spending environment driven by ongoing refresh on traditional data center and campus equipment; (+) INGM's broader participation in AI-related demand via the sale of GPU servers.

F5, Inc (FFIV)

Presenters: FFIV's CFO, Cooper Werner, and COO, Tom Fountain presented at the Communacopia + Technology Conference.

Bottom Line: FFIV is capitalizing on the growing complexities of hybrid multi-cloud architectures and AI-driven demand with its comprehensive application delivery and security platform. The company is experiencing strong momentum in Systems refresh and capacity expansion, while software growth is poised to accelerate in F2027 driven by favorable renewal cohorts and increasing platform adoption. Looking ahead, management remains excited about four key priorities: capitalizing on the secular shift to hybrid multi-cloud, capturing growing ADC demand, driving platform adoption across the installed base, and seizing new AI opportunities.

Key Takeaways

FFIV is well positioned to capture demand from hybrid multi-cloud & AI trends. Management noted that 90% of organizations today have a hybrid multi-cloud architecture, and are operating across an average of 22 different environments, creating significant complexity that FFIV is uniquely equipped to manage. Furthermore, AI is driving both indirect demand (increased bandwidth & capacity needs for traditional ADCs) and direct use cases. FFIV is actively monetizing direct AI opportunities which drove roughly $50 million of bookings in the first half of the year particularly in AI data delivery, AI runtime security, and AI factory load balancing (via a partnership with NVIDIA on DPUs).

Systems revenue supported by accelerated refresh cycles, capacity expansion, & competitive takeouts. FFIV is seeing a healthy demand environment characterized by a "refresh plus expansion" dynamic, where customers are not only replacing legacy infrastructure but also adding capacity and moving up to higher-end appliances to support workload growth. The upcoming end of software support for the iSeries in F2Q27 is driving a more orderly and accelerated refresh motion. Additionally, FFIV continues to gain market share through competitive displacements, benefiting from a pricing umbrella and competitors' underinvestment in systems innovation.

Software growth expected to accelerate in F2027 alongside strong Distributed Cloud Services adoption. While F2026 software growth faces a headwind from a flattish renewal cohort, management expects double-digit software revenue growth in F2027 as this dynamic becomes a tailwind. Growth is primarily driven by expansion within existing subscribers, higher consumption rates, and broader portfolio adoption. The SaaS transition headwind is largely complete, and DCS continues to see strong traction, with 33% of the top 1,000 customers adopting the platform (up from 24% a year ago), enabling consistent policy management from the data center to the edge.

Expanding security portfolio supported by SurePath AI acquisition. FFIV sits at the front door of mission-critical applications, serving 80% of the Fortune 500, which positions it well to address the expanding threat landscape. The company is seeing strong momentum with its AI-powered WAF and recently acquired SurePath AI to add guardrails, red teaming, and discovery capabilities to its AI security offering, which is resonating particularly well within financial services & government agencies.

Reiterated 2027 gross margins guidance amidst supply chain dynamics & consistent capital allocation. Management reiterated their F2027 gross margins guide of 80-82%, while noted that it has seen recent stability in memory pricing, though they remain cautious due to mixed signals from providers and will provide an updated outlook in October. Longer-term, gross margins are expected to stabilize as the availability of scarce components improves, supported by FFIV's strategic inventory build of long-lead-time components to assure supply. Additionally, management highlighted room for continued operating margin improvement as the company leverages automation and AI to drive sales productivity and scale opex efficiently. Capital allocation remains consistent, with a commitment to return at least 50% of free cash flow to shareholders via share repurchases, supplemented by strategic tuck-in acquisitions and increased capex to build out DCS points of presence.

Valuation and key risks

Valuation. We are Neutral rated on FFIV with a 12-month target price of $400 based on 20X NTM+1Y EPS forecast.

Key risks to our Neutral rating include: **(+)** Multi-cloud network architecture adoption should drive demand for a unified application delivery/security platform across both monolithic & cloud based apps; near-term revenue visibility from elevated backlog and hardware redesign; **(-)** Slower enterprise IT spending; margin degradation from investment elevated costs associated with supply chain headwinds; execution of cloud-based, SaaS-distribution platform.

Investment thesis summary. F5 is an industry leading application delivery and security company that provides customers with a unified experience across on-premise, public/private cloud, multi-cloud, and edge environments. Secular trends should have a mixed impact on demand for FFIV's legacy ADC product, which should be partially offset by the diversification into security & observability features. In our view, broad adoption of hybrid multi-cloud network architectures will require enterprises to continue to invest in their on-premise data centers, though we expect a majority of eligible workloads to transition into the cloud to the extent it makes sense for business operations. As such, we expect that enterprises will continue to need ADC hardware and software in their data centers to better manage load balancing, app acceleration, rate shaping, etc. FFIV has made meaningful progress in diversifying its business away from the legacy ADC market to also offer application security & observability features, which in our view are natural adjacencies to ADC and where FFIV has the right-to-win given its dominant market share position in ADCs.

Dell Technologies, Inc (DELL)

Presenters: DELL's CEO, Michael Dell, presented at the Communacopia + Technology Conference.

Bottom Line: DELL believes the market is in the early stages of a durable, multi-year AI infrastructure buildout, with demand broadening beyond neoclouds into enterprise and sovereign AI, while replacement cycles across PCs and traditional servers provide an additional layer of durability. The company is uniquely positioned to capture that demand through its scale, supplier relationships, and end-to-end portfolio spanning endpoints, compute, storage, networking, services, and financing. Management views execution and margin expansion as key differentiators for DELL that are not fully reflected in the current debate around AI infrastructure demand.

Key Takeaways

AI adoption curve is still looks early. DELL framed AI as a multi-year infrastructure S-curve, with enterprises only beginning to move beyond productivity tools (where organizations may see 10-20% improvements), toward rearchitecting end-to-end workflows. Dell's 6,500 AI Factory customers include a growing number of enterprise deployments, which reinforce management's view that adoption remains in the early stages. Regarding neoclouds, Dell looks through the neoclouds to the underlying end

customers, and management said it is only interested in serving neocloud customers that are backed by contracts from hyperscalers and frontier model providers such as OpenAI, Anthropic, Microsoft, Google and Meta. The company views itself as well positioned to continue addressing neocloud demand because of their ability to deliver clusters faster and accelerate time-to-first-token, which creates meaningful economic value for customers.

Enterprise and sovereign AI are becoming increasingly important demand vectors. While neocloud remains strong, Dell highlighted accelerating adoption among enterprise customers and sovereign AI projects globally. Management noted many regions are 12-18 months behind the U.S., but indicated the long-term sovereign opportunity is sizable and could represent ~1/2 of the opportunities outside of the U.S.

Demand remains durable and, if anything, deferred. Dell pointed to roughly 400 million PCs over four years old and 1.3 million servers running on 13G or older platforms as evidence of significant replacement demand still ahead. The company continues to see tight component availability, potentially worse in 2027 than 2026, while emphasizing that its direct customer relationships provide strong demand visibility with no signs of meaningful double-ordering or underutilization.

Scale and portfolio breadth are increasingly differentiating DELL from its competitors. Management highlighted advantages from supplier relationships, supply chain execution, and a portfolio spanning PCs, servers, storage, networking, services, and financing. Storage was a particular focus, with Dell arguing AI is creating a virtuous cycle where more intelligence generates more data, driving both storage demand and customer efforts to reorganize and architect data for AI workloads.

Operating leverage is underappreciated, per management. Dell reduced OpEx to roughly 8% of sales through a multi-year AI-driven modernization effort that decoupled revenue growth from cost growth. Management argued investors remain focused on AI infrastructure demand but may be underestimating the earnings power created by Dell's operating leverage and execution advantages.

Valuation and key risks

Valuation. We are Buy rated on DELL with a 12-month target price of $570 reflecting 18X our NTM+1Y EPS.

Key risks: (-) Weaker than expected consumer and commercial PC market demand; (-) Worse-than-expected enterprise IT spending; (-) headwinds from hybrid work could result in reduced spending in enterprise IT; (-) Macroeconomic weakness could result in lowered consumer demand; (-) Pricing pressure could occur as competitors discount due to excess channel inventory and reduced demand, particularly in PCs; (-) Margin pressure from higher input costs; (-) Competition from white box manufacturers could occur, particularly in servers and storage; (-) Structurally lower AI server demand from neocloud cohort.

Hewlett Packard Enterprise Co (HPE)

Presenters: HPE's SVP, Chief Strategy Officer, Investor Relations, Shannon Cross, presented at the Communacopia + Technology Conference.

Bottom Line: HPE emphasized its constructive outlook for both AI networking and inference-driven compute demand. Management highlighted momentum across scale-up, scale-out, and scale-across networking opportunities, including AMD Helios, where HPE will participate through both full racks and white-label trays, and provided additional detail on its gigawatt-scale Oracle win, which was driven by HPE's networking portfolio, scale, and financing capabilities. Regarding the raised F27 outlook (on earnings), the company emphasized the benefit of new wins, strong core demand (across networking, storage, and servers), and improving supply visibility ($30 bn in purchase commitments). HPE argued that enterprise inference demand is proving durable, driven by infrastructure refreshes, token-cost optimization, and data sovereignty concerns, citing its recently announced $3.5 billion enterprise inference win as evidence that customers are increasingly willing to deploy AI compute on-prem with traditional server-like economics.

Key Takeaways

Data center networking momentum in scale up, scale across, and scale out. First on AMD Helios, HPE plans to participate in two ways: (1) providing the full rack, combining HPE server and networking IP, and (2) supplying white-label trays to other server vendors. Management has not yet quantified the revenue opportunity but views Helios as an attractive way to expand into hyperscale and neocloud AI deployments where HPE has historically been more selective (i.e., higher margins because networking IP). Second, HPE is seeing strong demand for routing and switching from hyperscalers and neoclouds, with growing enterprise interest as customers modernize infrastructure for AI. Third, HPE also provided additional color on its recent gigawatt-scale Oracle win announced on earnings, emphasizing that it is a net-new back-end AI networking deployment, with HPE's end-to-end portfolio (including the QFX direct liquid-cooled Tomahawk 6 platform), scale, and financing capabilities helping secure the deal.

Raised FY27 revenue outlook on (1) new deals; (2) momentum in core demand; and (3) improved supply availability. Management emphasized that demand remains robust across networking, storage, and traditional servers, with networking orders up 36%, storage orders growing roughly twice the rate of revenue, and traditional server orders increasing 75% year-over-year. Improved supply visibility also supported the raise, with purchase commitments increasing to roughly $30 billion from $6 billion last quarter, including a doubling of networking-related commitments year-over-year. Management reiterated its focus on converting backlog into revenue while maintaining its free cash flow framework of at least $3.75 billion in FY26 and $5 billion in FY27, with a commitment to return at least 75% of FCF to shareholders.

HPE sees durable, accelerating demand for inferencing based compute. HPE remains increasingly constructive on the durability of enterprise inference demand on (1) customers refreshing legacy infrastructure to free up data center capacity; (2) efforts to contain token costs through on-prem AI deployments; and (3) growing concerns around data privacy and sovereignty. The company highlighted a recently announced $3.5

billion enterprise win with a large cloud customer for internal AI inference workloads, noting the opportunity carries traditional server-like margins rather than typical AI server economics. Management believes HPE's services, installation, maintenance, and financing capabilities were key differentiators and sees a growing pipeline of similar opportunities across both enterprise and sovereign customers.

Valuation and key risks

Valuation. We are Buy rated on HPE with a 12-month target price of $75 reflecting 14x our NTM+1Y EPS.

Key risks: (-) Lower-than-expected corporate IT spending and data center capex could result in worse-than-expected hardware sales; (-) Competition from white box manufacturers could result in cannibalization in HPE's server market share; (-) Market share loss from customer churn during Juniper integration; (-) Component costs could be higher than expected, resulting in worse gross margins; (-) Corporate actions from strategic investors; (-) Failure to deliver market share gains in storage driven by a strategy shift towards selling more owned-IP storage products that carry higher gross margins; (-) Worse-than-expected demand for enterprise and sovereign AI data center infrastructure.

Super Micro Computer, Inc (SMCI)

Presenters: SMCI's Senior Vice President, Corporate Development, Michael Staiger, presented at the Communacopia + Technology Conference.

Bottom Line: (1) FY27 outlook underpinned by order visibility, customer diversification, and expanding share opportunities; (2) Expectations to return to stable double-digit gross margins driven by a richer solutions mix; (3) Neocloud AI server opportunity is durable, driven by complexity of solutions & SMCI's faster time-to-value proposition.

Key Takeaways

FY27 outlook underpinned by order visibility, customer diversification, and expanding share opportunities. Management pointed to its roughly $60 billion order book as the primary foundation for FY27 guidance while highlighting a broadening customer base and a multi-year AI infrastructure buildout cycle. SMCI believes it is increasingly positioned to serve the enterprise and sovereign markets as customers shift from buying components to buying complete solutions. Driving DCBBS adoption, which bundles power, cooling, cabling, storage, software, and services into integrated deployments, is core to SMCI's enterprise strategy. SMCI designs most of these components internally while leveraging partners where appropriate. While working capital intensity remains elevated given the pace of growth, management is focused on improving customer terms and delivery timing to reduce capital requirements over time and ultimately position the business to return capital to shareholders.

Expectations to return to stable double-digit gross margins driven by a richer solutions mix. Management reiterated that its path back to sustainable double-digit

gross margins is being driven by increasing content per deployment rather than a change in the overall demand environment. While the market previously experienced intense price competition among suppliers focused on a handful of large AI buyers, customers are increasingly prioritizing integrated solutions that reduce deployment times and improve reliability. Management noted that FY4Q benefited from favorable mix, including more traditional server content, while FY1Q will see some pressure from delayed AI-server shipments. SMCI remains committed to investing in engineering, software, services, and go-to-market capabilities to support its expanding enterprise and sovereign customer base.

Neocloud AI server opportunity is durable. Management argued that as AI infrastructure grows more complex, customers increasingly value fast deployment, operational readiness, and reliability over component-level cost optimization. SMCI highlighted an example where an 8 rack deployment taking 30 days to become operational could represent roughly $3 million of lost value, underscoring the importance of time-to-revenue. To support future demand, the company is investing aggressively in capacity beyond its current 6,000 racks-per-month capability, including new facilities, clean rooms, manufacturing capacity, and expanded go-to-market resources, with additional details expected in the future.

Valuation and key risks

Valuation. We are Sell rated on SMCI with a 12-month target price of $34 reflecting 7.5x NTM+1 EPS.

Key risks: (+) Stronger-than-expected demand for AI servers; (+) Market share gains; (+) Core operating margin improvement; (+) Customer diversification.

Penguin Solutions, Inc (PENG)

Presenters: PENG's CEO, Kash Shaikh, presented at the Communacopia + Technology Conference.

Bottom Line: PENG's differentiated position at the intersection of AI infrastructure and memory, combined with ClusterWareAI and recurring 3-5 year services contracts, is helping the company win larger projects in both the neocloud & enterprise. Neocloud demand should be durable, while enterprise interest will continue to be driven by customers seeking greater data control and better economics from on-prem AI deployments. Longer term, management sees additional upside from inference-focused memory solutions, including CXL-based products today and photonic memory architectures beginning in 2H27.

Key Takeaways

Focus on serving both neocloud and enterprise customers with end-to-end AI infrastructure solutions. Management highlighted accelerating demand from neocloud providers, which it views as a durable opportunity given continued demand (hyperscaler outsourcing, enterprise AI adoption) and the need for third-party operators to design,

build, and manage AI factories. PENG's enterprise value proposition is built around four offerings: (1) AI Factory solution, including the proprietary ClusterWareAI operating system; (2) MemoryAI products; (3) ComputeAI offerings (GPU servers); and (4) end-to-end services. The company's land-and-expand model often starts with hardware or software and expands into broader infrastructure deployments, with 3-5 year management contracts creating recurring and increasingly sticky revenue streams.

Enterprise AI factory wins increasingly driven by economics and operational simplicity. PENG highlighted a recent Tier 1 financial services win where the customer elected to build an on-prem AI factory to support new AI applications and reduce the long-term economics associated with consuming AI infrastructure from external providers. PENG believes it won the engagement because of its (1) comprehensive design-build-deploy-manage offering; (2) its hardware-agnostic ClusterWareAI platform; and (3) its MemoryAI appliances. The company noted that coding assistants are emerging as one of the most common enterprise AI use cases, driven by both productivity gains and cloud cost considerations. Internally, PENG is also leveraging on-prem AI infrastructure and agentic AI tools to accelerate software development for ClusterWareAI and drive operating efficiencies.

Integrated Memory business increasingly aligned with AI inference workloads. PENG's integrated memory business is built around its ability to engineer, validate, manufacture, and supply highly reliable memory modules for major networking and compute OEMs. Management is increasingly prioritizing data center customers and AI-driven opportunities, where demand appears more durable than historical memory cycles. A key area of focus is CXL-based memory expansion cards, which helps address growing inference requirements by expanding memory capacity and improving GPU utilization. Beyond CXL memory cards, the company has developed its MemoryAI KV Cache Appliance (part of the Accelerated Compute segment) to reduce recomputation and improve LLM response times. Looking further out, photonic memory solutions, developed through its relationship with Celestial AI, are expected to become available in 2H27 and could enable high-speed shared memory architectures for future AI infrastructure deployments.

Valuation and key risks

Valuation. We are Buy rated on PENG with a 12-month target price of $75 based on a blended fundamental/M&A valuation methodology. Our fundamental value of $72.00 (85% weight) is based on 18 NTM+1Y P/E. 18x is largely in line with PENG's enterprise AI infrastructure peers DELL, HPE, and NTAP. Our theoretical M&A value of $90.00 (15% weight) is based on 2.0X NTM+1Y EV/Sales and is largely in line with precedent transactions in the advanced computing and integrated memory space.

Key risks: (-) Memory market cyclicality; (-) Competition from OEMs for Advanced Computing; (-) China revenue exposure; (-) Supply chain disruptions; (-) Enterprise demand for AI

Cricut, Inc (CRCT)

Presenters: CRCT's CFO, Kimball Shill, presented at the Communacopia + Technology Conference.

Bottom Line: CRCT looks to return to sustainable, profitable growth through investment in better user experiences, rather than aggressive monetization, as the company remains focused on making its offering affordable for users. Management believes re-engineered hardware, simplified guided workflows, AI-enabled creation tools, and a more compelling out-of-box experience (via bundles) should drive stronger engagement and customer retention over time.

Key Takeaways

Investing in hardware design and platform improvements to enhance the customer experience and strengthen the flywheel. CRCT highlighted the recent launch of the re-engineered Maker 5, the first major redesign of CRCT's flagship Maker platform since 2014, which reduces component count while improving performance at an attractive price point. Cutting machines remain the primary entry point into the ecosystem, driving downstream monetization through subscriptions and accessories/materials. At the same time, CRCT continues to invest in platform improvements, including guided flows, AI-powered project creation, and simplified onboarding experiences that make the most popular use cases easier for both new and existing users. The company is also making smaller, targeted monetization investments through Cricut Labs offerings such as patterns, paint-by-number, and coloring pages, primarily aimed at driving engagement among subscribers.

Expectations for sustainable platform growth remain underpinned by improving engagement metrics and monetization opportunities. Management emphasized that subscription growth has remained consistent, with platform revenue continuing to expand and account for roughly 80% of company profitability. Recent engagement metrics showed encouraging signs, including year-over-year growth in active users and stabilization in 90-day engaged users. CRCT is also seeing early success from selective pricing actions, including premium pricing within the iOS channel and the roll out of a higher-tier subscription plan that includes additional AI credits and benefits. Importantly, management has not yet implemented a broad-based price increase on its core $9.99 plan, noting that stronger engagement trends would likely need to be sustained before taking that step. Alongside monetization efforts, the company continues to invest in user acquisition through the "Think It. Make It. Cricut." campaign and international expansion, with Cricut.com traffic serving as an important measure of marketing effectiveness and purchase intent.

Capital allocation remains focused on investing in growth while maintaining meaningful shareholder returns. First, the company will prioritize investment in hardware innovation, platform development, and marketing to drive awareness and user acquisition, particularly in international markets where management ultimately believes the opportunity could rival the size of North America. Second, CRCT remains open to strategic M&A opportunities that accelerate existing priorities, though management reiterated it has no interest in pursuing bolt-on acquisitions simply for scale. Third, the company intends to continue returning excess capital through share repurchases, a

recurring semiannual dividend, and special dividends when appropriate.

Valuation and key risks

Valuation. We are Sell rated on CRCT with a 12-month target price of $4.50 reflecting a blended fundamental/M&A valuation methodology. Our fundamental value of $4.00 (85% weight) is based on 18x NTM+1Y P/E. Our theoretical M&A value of $6.50 (15% weight) is based on 1.5X NTM+1Y EV/Sales.

Key risks: (+) Higher than expected connected machine sales; (+) Expanding TAM; (+) Higher levels of engagement; (+) Long-term operating margins expansion