BERNSTEIN
SOCIETE GENERALE GROUP

21 September 2026

U.S. Semiconductors and Semiconductor Capital Equipment

U.S. Semiconductors and Semicap Equipment: Notes from the road - takeaways from management meetings in Silicon Valley



Stacy A. Rasgon, Ph.D.
+1 213 559 5917
stacy.rasgon@bernsteinsg.com



Alrick Shaw
+1 917 344 8454
alrick.shaw@bernsteinsg.com



Arpad von Nemes
+1 917 344 8461
arpad.vonnemes@bernsteinsg.com



Eva Zhang
+1 212 845 7839
eva.zhang@bernsteinsg.com

Last week we hosted a semiconductor investor bus tour in Silicon Valley, meeting with 11 companies and management teams over a 2-day period (5 covered, and 6 non-covered). Specifically we met with:

- Broadcom (Hock Tan, President and CEO, Amie Thuener, CFO and Ji Yoo, Head of IR) (*rated OP, $575 TP*)
- NVIDIA (Colette Kress, CFO, and Toshiya Hari, VP IR) *(rated OP, $400 TP)*
- Intel (John Pitzer, Corporate VP, Treasurer and Investor Relations) *(rated MP, $110 TP)*
- NXP Semiconductors (Jeff Palmer, SVP of Investor relations) *(rated MP, $290 TP)*
- Applied Materials (Zhebo Chen, Head of product management for ALD & MDP, Bala Haran, VP Integrated Materials Solutions, and Marty Parker, Senior Director Investor Relations) (*rated OP, $700 TP)*
- Monolithic Power Systems (Tony Balow, VP Finance) (*not covered*)
- Astera Labs (Jitendra Mohan, CEO, Des Lynch, CFO, and Nicholas Aberle, SVP Finance and Investor Relations) *(not covered)*
- Lattice Semiconductor (Rick Muscha, Vice President, Investor Relations) (*not covered*)
- Global Foundries (Michael Hogan, Chief Business Officer) *(not covered)*
- Ambarella (Fermi Wang, CEO, John Young, CFO, Louis Gerhardy, VP Corporate Development and IR) *(not covered)*
- Lightmatter (Nick Harris, Founder and CEO, Simona Jankowski, CFO) (*private*)

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.

First Published: 21 Sep 2026 04:14 UTC Completion Date: 21 Sep 2026 04:14 UTC

BERNSTEIN TICKER TABLE

			18 Sep 2026 Closing		TTM	Adjusted EPS				Adjusted P/E (x)		
Ticker	Rating	Cur	Price	Price Target	Rel. Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
AMAT (Applied Materials)	O	USD	444.57	700.00	118.9%	USD	9.42	12.76	18.30	47.2	34.8	24.3
AVGO (Broadcom)	O	USD	357.61	575.00	(11.8)%	USD	6.82	11.64	18.63	52.4	30.7	19.2
INTC (Intel)	M	USD	108.60	110.00	239.9%	USD	0.43	1.55	1.81	254.7	70.2	59.9
NVDA (NVIDIA)	O	USD	222.27	400.00	10.8%	USD	4.77	9.20	15.80	46.6	24.2	14.1
NXPI (NXP Semiconductors)	M	USD	227.99	290.00	(14.7)%	USD	11.81	15.02	16.85	19.3	15.2	13.5
SPX			7,650.50									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

NVDA base year is 2026;

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

AMAT (OP, $700): Exposure to key inflections is strong & valuation vs peers is attractive.

AVGO (OP, $575PT): A strong 2026 AI trajectory seems set to markedly accelerate into 2027 and 2028.

INTC (MP, $110): Server strength is helping the company get back on their feet, and narrative/headlines may fuel the vibe for now.

NVDA (OP, $400): The datacenter opportunity remains enormous.

NXPI (MP, $290): The pace and makeup of recovery remains open for debate, though company execution is solid and valuation is increasingly inexpensive.

DETAILS

BROADCOM (AVGO, RATED OP, TP $575) - MEETING WITH HOCK TAN (PRESIDENT AND CEO), AMIE THUENER (CFO) AND JI YOO (HEAD OF IR)

Management emphasized that the custom AI accelerator opportunity is broadening well beyond Google and that the demand attributed to Anthropic is separate from Google's internal TPU demand. Specifically, Broadcom characterized the 5GW Anthropic opportunity as Anthropic-specific demand (they are writing the checks), regardless of where the infrastructure is ultimately hosted. Management also clarified that its AI revenue outlook and customer GW roadmaps are related but not directly equivalent: the GW figures represent customer demand, while revenue forecasts incorporate judgment around how fully that demand can actually be deployed. If customers ultimately deploy the full GW amounts, revenue would be higher than the current outlook of $115B/$230B in 2027/28, respectively. Google remains a stable-to-growing business, with Broadcom participating in Sunfish in 2027 and management is expecting to remain involved in subsequent generations. More importantly, management believes OpenAI and Anthropic could each become at least if not larger than Google, making the custom accelerator business considerably more diversified by 2028.

Near-term demand also appears stronger than anticipated. Management expects sequential growth through 2027, with OpenAI looking to pull in its roadmap and accelerate deployment of Jalapeño beyond what is embedded in Broadcom's current outlook. Meta is similarly attempting to pull in its roadmap, although Meta remains primarily a 2028 opportunity with potential initial revenue in late 2027. The gating factors are increasingly physical infrastructure and semiconductor supply rather than demand, particularly power, wafers and memory. Broadcom has detailed visibility into the physical sites underlying its 2027-28 forecasts, noting that 2027 sites are largely complete, and 2028 sites are already sufficiently advanced to provide confidence in the outlook.

Management spent considerable time addressing the competitive debate around Google's future TPU generations. Broadcom reiterated that its agreement with Google covers long-term TPU development and supply, and while declining to validate speculation around specific future program naming or share allocations management expects to participate in the generations following Sunfish as well as arguing that technical challenges are becoming substantially harder with every generation, particularly around high-speed mixed-signal design and SerDes. More broadly, management's message was that investors are too focused on Google share risk relative to the size of the incremental opportunities at OpenAI, Anthropic and Meta. Google already has a large installed compute base, while the frontier labs are building capacity from much lower starting points, creating significantly greater incremental silicon requirements. Management also pushed back on the view that NVIDIA-based infrastructure will dominate, including Elon Musk's comments favoring an all-NVIDIA approach. Broadcom argued that frontier labs with sufficiently large, predictable workloads have a strong economic and performance incentive to develop custom silicon because they understand the models and workloads they intend to run. In addition, software is becoming less of a barrier as well, noting that its custom XPUs can increasingly support broader PyTorch-based models in addition to the workloads for which they were specifically optimized, albeit with the best performance/watt still coming from purpose-built hardware/software co-design. Management similarly argued that closed frontier models remain structurally better positioned to fund increasingly expensive compute infrastructure than open-weight models, given their ability to monetize proprietary capabilities.

A key differentiator in custom silicon is Broadcom's ability to architect substantially more than the accelerator itself. With less mature silicon organizations such as OpenAI and Anthropic, Broadcom is helping define the XPU, memory subsystem, networking architecture and rack design, which should create more content per deployment than with an established customer such as Google. Management was dismissive of the idea that competitors currently provide equivalent capabilities: MediaTek participates in pieces of the TPU opportunity but does not provide the same end-to-end architecture, while management characterized Marvell's current exposure as XPU-adjacent rather than a meaningful custom accelerator business.

Networking remains an important component of the company's competitive advantage with management suggesting that Broadcom's 400G-class SerDes allows customers to retain copper within AI racks for longer, which they prefer because copper remains cheaper, lower-latency and lower-power than optics. In management's view, competitors discussing NPO/CPO more aggressively partly reflects weaker capabilities, whereas Broadcom can defer the optical transition until it is economically necessary. In that context, they noted that roughly 80% of connectivity within Google's scale-up topology remains copper, while optical connections are already standard for longer-distance scale-out. As per-XPU bandwidth rises toward 200-400Gbps per lane and electrical reach becomes insufficient, optics will progressively move closer to compute, but

management does not expect CPO to become a major scale-up technology until roughly 2028/29.

Supply is increasingly central to Broadcom's ability to convert demand into revenue. TSMC remains the primary foundry partner, and management sees near-term alternatives as limited because these products require not only leading-edge digital processes but also mature mixed-signal PDKs developed through years of collaboration. As such, Samsung discussions are currently primarily memory-related, while management does not currently view Intel as having a production-ready process for Broadcom's requirements. Memory is similarly being secured directly by Broadcom through LTAs, with costs passed through to customers; and management believes it has sufficient memory supply through 2029. Broadcom is also building proprietary advanced-packaging capacity in Malaysia because future architectures become increasingly difficult to manufacture (e.g. 2 dies per package today, to 4 next year, and potentially 16 the following year) and the company noted that this capacity is intended for Broadcom's own custom-chip customers rather than as a merchant packaging business and will supplement, rather than replace, CoWoS.

NVDIA (NVDA, RATED OP, TP $400) - MEETING WITH COLETTE KRESS (CFO) AND TOSHIYA HARI (VP IR)

Management pushed back strongly on the view that custom ASICs will structurally erode NVIDIA's position, arguing that the value of its platform comes from performance, fungibility and the broader CUDA ecosystem rather than the GPU alone. Based on NVIDIA's hyperscaler revenue relative to publicly disclosed capex, management estimates its share of hyperscaler capex has risen from roughly 5% pre-ChatGPT to ~20% and more recently the mid-20% range, despite continued TPU and custom ASIC deployments. Management cited Elon Musk's decision to standardize Colossus around NVIDIA as an example of the value of fungible infrastructure: GPUs can support changing workloads or potentially be rented to third parties, whereas an ASIC is optimized for a narrower workload. NVIDIA also disagrees that its CUDA moat is weaker in inference, arguing that CUDA, CUDA-X libraries and the developer ecosystem together continue to offer better performance and ease of deployment than porting workloads to alternative hardware. Anthropic is one area where NVIDIA starts from relatively low share, making increased engagement with Anthropic, Microsoft and Lambda an incremental opportunity despite Anthropic's efforts to diversify across AMD and custom silicon.

NVIDIA also believes its content opportunity is expanding materially beyond the GPU. Management discussed the progression in data center content toward roughly $40B/GW, reflecting the addition of CPUs, LPUs, networking and other system-level components rather than simply higher GPU pricing. The roughly $20B of Vera Rubin revenue discussed for the current quarter refers primarily to the standard Vera Rubin platform and mostly excludes separate CPU rack and LPU opportunities, though with LPUs beginning to ramp this quarter. This broader system strategy is central to NVIDIA's differentiation versus custom silicon vendors, which management views as primarily participating at the chip level. Longer term, NVIDIA prefers a highly co-designed but general-purpose platform that can accommodate changing workloads rather than narrowly customized processors, arguing that flexibility becomes more valuable as AI models and workloads continue to evolve.

Demand remains substantially ahead of supply. NVIDIA previously indicated approximately 70% supply-constrained growth in FY27, with the constraints spanning wafers, CoWoS, memory and, increasingly, optics/transceivers and available land, power and shell capacity. Management believes substrates are manageable but emphasized that virtually every part of the supply chain is being pushed aggressively. NVIDIA is using its balance sheet primarily to secure and expand supply, providing commitments where necessary to enable suppliers to invest. The company would welcome additional leading-edge foundry alternatives (Intel?) but stressed that qualifying a new foundry would require more than 2-3 years given the technical work and depth of the manufacturing relationship required. The discussion also addressed how quickly the ecosystem can absorb the capacity NVIDIA is enabling. Management sees land, power and shell availability becoming an increasingly important constraint alongside semiconductor supply, although customers are working aggressively to bring infrastructure online. NVIDIA appreciates the SpaceX's commitment to its platform, but also emphasized that a potential 5-10GW deployment next year remains subject to site readiness, financing and NVIDIA's broader allocation process. NVIDIA is not prepared to divert supply simply on the basis of announced GW targets. More broadly, management's allocation decisions increasingly incorporate whether the underlying data center infrastructure is sufficiently advanced to absorb hardware as it becomes available.

Cash return sounds like it is set to inflect. Once supply and infrastructure commitments are addressed, management intends to return essentially all remaining FCF to shareholders through buybacks and potentially a growing dividend.

NVIDIA's compute-rental backstops with neoclouds were another focus of investor questions. Management framed these arrangements as a mechanism to unlock capacity for enterprises that need compute but are unwilling to commit to a five-year contract. NVIDIA can provide a baseload commitment that enables the neocloud to stand up infrastructure, after which shorter-duration enterprise customers consume that capacity at higher rental rates. Management emphasized that no GPUs are currently sitting idle and expects these arrangements to reach a few $B of revenue next year, still relatively small versus NVIDIA's overall business. NVIDIA receives the upfront economics from selling the hardware and can participate in subsequent usage revenue; only if the compute went unused would NVIDIA potentially consume it internally and recognize the associated expense in R&D. Management therefore does not view these arrangements as financing weak demand, but as helping address a mismatch between long-duration infrastructure commitments and customers seeking shorter-term compute.

The customer mix should also continue broadening. Hyperscalers are expected to remain close to 50% of the business longer term, but management believes the other half, including neoclouds, industrial, enterprise, sovereign and AI labs, can ultimately grow faster. Neoclouds are currently the largest component of that non-hyperscale bucket, followed by areas such as industrial and enterprise, while sovereign deployments frequently flow through cloud providers rather than directly. **Frontier AI labs represent roughly 20% of NVIDIA revenue**, including purchases flowing through

clouds and neoclouds, meaning OpenAI and Anthropic are important but far from representing the entirety of the demand base. Management is also indifferent as to whether frontier labs ultimately self-build or use CSPs: several initially favored first-party infrastructure but discovered both the operational complexity and financing advantages of using third parties, suggesting the two models will coexist.

Management pushed back against the argument that open models undermine AI economics and NVIDIA expects roughly 80% of models ultimately to be open noting that open models does not make the underlying compute free. Instead, management views broader open-model adoption as positive for accelerated computing because it expands the universe of enterprises and developers capable of deploying AI. This also relates to management's response to custom silicon: NVIDIA believes specialized ASICs can work well for predictable workloads, but a broadly programmable system is more valuable when workloads and models are changing rapidly. The company's objective is therefore not to customize silicon for every workload, but to use extreme hardware/software co-design to maximize performance while retaining platform fungibility.

Finally, management continues to see significant opportunity to improve efficiency within that full-stack architecture. It expects memory consumed per token to improve over time through continued hardware and software optimization and remains willing to reassess component specifications as future systems evolve, although management indicated it is not currently pursuing memory de-specification for existing architectures (including Vera Rubin). At the same time, rapidly increasing performance requirements mean absolute memory and infrastructure demand can still grow substantially even as efficiency improves. Gross margins remain influenced by rapidly changing product mix and elevated input costs, particularly memory, but management indicated selective pricing actions have been effective and does not expect supply tightness to disappear quickly.

INTEL (INTC, RATED MP, TP $110) - MEETING WITH JOHN PITZER (CORPORATE VP INVESTOR RELATIONS)

Management attributed the recent gross margin upside to better product mix, volume and pricing, but emphasized that the broader margin recovery remains a multi-year cost and scale story. Intel 18A and Intel 3 are currently below corporate-average gross margins as they ramp, while mature Intel 7/10 is well above average, and the mix shift over time from the latter to the former will be a headwind. Management noted that node economics are driven by utilization, yield and scale, with scale currently the missing piece on 18A. Intel expects to more than double 18A wafer capacity next year, with Nova Lake becoming the primary volume driver. At the same time, rapidly growing ASIC and advanced packaging businesses will be gross-margin dilutive near term, partially offsetting improvements in the core CPU business.

Longer term Lip-Bu is not optimizing for gross margin, but rather revenue growth with a "rule-of-45" framework (revenue growth plus operating margin). Management noted that even during Intel's historical period of 60%+ gross margins, revenue growth of only ~3% and a heavy operating expense structure meant the company struggled to achieve even a Rule of 30. By contrast, management estimates the TAM Intel is addressing today is approximately 15x larger than its historical opportunity set, and as a result, the company is willing to accept some near-term gross margin dilution if the incremental businesses generate attractive returns and provide attractive revenue growth to capture the large TAM opportunity.

Perhaps surprisingly, the company was direct that the current lack of cost competitiveness is driven more by the Products than the Foundry business, reflecting that Intel Products uses substantially more silicon than peers, as historically design teams optimized primarily for performance and treated silicon area as a "free resource." The company provided an illustrative example, noting that a synthetic, integrated AMD/TSMC model might incur only ~$15-20 of wafer cost per $100 of revenue, versus roughly $60 at Intel today, with only part of that difference stemming from the Foundry. This as an issue with Granite and Diamond Rapids and the company said that Lip-Bu Tan is putting a greater emphasis on die size, tile reuse and overall cost competitiveness with Panther Lake and Nova Lake representing incremental progress, and Coral Rapids potentially improving server economics. However, reaching competitive cost levels will remain a multi-year process with Foundry further along in its profitability improvement than Intel Products.

Server demand remains exceptionally strong, with management expecting to sell essentially everything it can manufacture in 2027 and tightness potentially extending into 1H28. Near-term constraints are not limited to wafers: substrates were the bigger server constraint in the most recent quarter, and management noted that customers may be willing to take CPUs even if data center power availability delays actual deployment. On the competitive front, Granite Rapids remains competitive with AMD Turin in several workloads, but management acknowledged that Diamond Rapids could be a relative step back versus the subsequent AMD generation given Intel's lack of multithreading, although agentic workloads' emphasis on single-thread performance could mitigate this. Consequently, management suggested that Intel's 2027 server share may depend as much on available wafer capacity, particularly from Ireland, as on underlying demand.

Advanced packaging represents another major revenue opportunity, although profitability will take time to mature. Intel continues to target approximately 40% gross margins for advanced packaging over time, but cautioned against assuming EMIB-T reaches that level by late 2027. The principal challenge is obtaining acceptable large-format substrates and improving substrate supplier yields rather than Intel's internal packaging yield once substrates arrive. Recent capital raising was partly intended to secure substrate capacity, and Intel is working directly with suppliers to facilitate investment after many were burned by the post-COVID inventory correction. Management noted that Intel remains the world's largest substrate buyer, with NVIDIA now approaching similar scale, but does not assume its EMIB-T position will remain exclusive as competing technologies emerge.

The ASIC business could become an increasingly important growth leg and was one area where Lip-Bu appears to have materially changed Intel's strategy of investing behind, instead of divesting it. ASIC revenue has increased from only a few hundred million dollars roughly a year ago to a ~$2B annualized run-rate, with management expecting it to reach ~$4B run rate relatively quickly. Today, essentially all revenue is networking ASIC-related, but Intel sees opportunities across networking, compute and accelerator ASICs. Hyperscalers have expressed interest in custom x86 silicon and while Intel does not intend to license x86 directly, it intends to monetize its x86 IP through custom ASIC engagements, with management noting that a compute ASIC win in 2027 would be an important KPI. On accelerators, Intel acknowledged that like others, its SerDes technology currently trails Broadcom's, but sees differentiation from offering design services, wafers and advanced packaging as an integrated, one-stop-shop solution. Intel also pushed back on the notion that Arm has an inherent power-efficiency advantage, arguing that Panther Lake demonstrates competitive power efficiency in client, with the next challenge being to prove the same in servers. Hyperscalers historically developed Arm CPUs partly because custom x86 was not available to them

at the time, creating an opportunity for Intel's ASIC model to offer customization while preserving the x86 ecosystem.

Foundry and capital intensity remain critical pieces of the debate. Intel continues to target 40% gross margin / 30% operating margin for Foundry over time, with mature capacity potentially becoming an underappreciated asset. Management highlighted that Intel 7/10 currently generates approximately $6-8 of annual FCF for every $1 of net book value, reinforcing the value of extending older nodes into applications such as base dies rather than simply retiring them as leading-edge products migrate forward.

Management would not comment on speculation around Hynix, but characterized base dies as an attractive business and noted Intel 3 could have a long life in such applications. Further addressing recent Sk Hynix/Intel rumors, the company noted that they have no desire to re-enter the memory business as a discrete player, but emphasized its deep memory IP portfolio and willingness to consider opportunities that address the industry's memory bottleneck if they create shareholder value.

Finally, client remains an important source of both near-term risk and cash generation. Higher memory pricing is beginning to weigh on PCs, although management characterized current weakness as demand deferral rather than demand destruction and continues to see PCs as important endpoints for agentic and edge AI longer term. The inventory charge taken in Q2 was in fact client work-in-progress that was scrapped to make room for server. More broadly, management views the largest 24-month risks as PC memory inflation and distorted demand signals created by industry-wide supply tightness, rather than an imminent collapse in AI spending. Lastly, management believes the AI environment remains tight into 2027 and potentially into 1H28, but acknowledged that the more important test comes in late 2027/2028, when current capacity investments begin arriving and the industry gets a clearer view of normalized demand. Regarding their own capex, which is expected to increase meaningfully next year, the company noted that the recent equity raise was mainly driven to lock down tool orders (in addition to secure substrate supply) and that the raise is probably the last time Intel needs to access the public equity markets, while leaving open the possibility of further strategic customer investments.

NXP SEMICONDUCTORS (NXPI, RATED MP, TP $290) - MEETING WITH JEFF PALMER (SVP OF INVESTOR RELATIONS)

NXP's tone was broadly constructive, particularly around the automotive and industrial outlook. Auto inventory normalization appears largely complete, with channel inventory returning to ~11 weeks, the first time it has returned to pre-COVID levels. Importantly, management does not see evidence of broad customer restocking, as inventory remains uneven across the customer base. Auto demand appears to have troughed in 3Q24, and management continues to view the industry as being in the early stages of the software-defined vehicle (SDV) transition, which remains the primary driver of long-term content growth.

In automotive, the core investment thesis remains centered on content expansion rather than vehicle unit growth. NXP reiterated its long-term algorithm of 8-12% automotive growth, assuming relatively flat industry production, with roughly half of the business tied to higher-growth areas such as SDVs, processing, connectivity, and vehicle electronics. Pipeline activity around SDV programs continues to increase, and management noted that zonal architectures typically require meaningfully more microcontroller content than traditional vehicle designs. While Tesla has accelerated development cycles across the industry, management believes Chinese OEMs have been far more successful than Western manufacturers at adopting that operating model.

The discussion around China was balanced. NXP noted that only 17% of revenue is tied to China under its current reporting methodology and emphasized that the company is far more exposed to electronic content growth than domestic vehicle unit trends. While Chinese domestic demand remains mixed, export-oriented automotive production remains strong. Competitive pressure from Chinese semiconductor suppliers is most evident in power discretes, but management continues to see meaningful capability gaps in microcontrollers and broader platform-level solutions, where local competitors remain less developed.

Industrial remains an area of particular confidence. NXP noted that the business is currently tracking roughly a year ahead of internal expectations, supported by strong end customer backlog visibility (via tight integration with distributors systems) and growing demand for higher-performance edge processing. Historically, much of the industrial market was served through catalog products, but customers increasingly require processors capable of running AI models locally. The recently acquired Kinara platform appears to be gaining traction, with management citing a roughly $1.5 billion pipeline. While much of that pipeline remains at proof-of-concept stage, management described it as the fastest adoption curve they have seen for a new product family and is building software and model optimization capabilities alongside the hardware offering.

While investors continue to focus on AI, the company emphasized that NXP's data center business remains relatively small at roughly 5% of revenue, albeit growing at over 20% annually. The company participates primarily in networking, top-of-rack switching, cooling systems, and power management applications rather than accelerator silicon itself. Management argued that this business is likely more durable than investors appreciate given its infrastructure-oriented nature and sticky customer relationships.

Finally, management spent considerable time discussing margin expansion and capital allocation. The company remains confident in its ability to drive gross margins toward 60% by the end of 2028 through a combination of manufacturing restructuring, strategic foundry partnerships, and the ramp of ventures such as VSMC. Management pushed back on the notion that future growth will be driven by pricing, noting that pricing has contributed little recently and that the long-term story remains about content growth and operating leverage. With major equity investments largely behind the company, management also suggested that greater capital returns could become a larger component of the investment story over the next several years.



APPLIED MATERIALS (AMAT, RATED OP, TP $700) - MEETING WITH ZHEBO CHEN (HEAD OF PRODUCT MANAGEMENT FOR ALD & MDP), BALA HARAN (VP INTEGRATED MATERIALS SOLUTIONS), AND MARTY PARKER (IR)

Management's conviction in the WFE outlook has strengthened meaningfully through the year, with Semi Systems growth expectations progressing from >20% in February to >30% in May and ~40% by August. Management continues to expect very strong growth in 2027, with roughly 80% of incremental WFE growth coming from leading-edge logic, DRAM and advanced packaging. Customer visibility has improved, and AMAT is preparing to double quarterly output capacity by 2028, with discussions with some customers extending through 2030. While management is not explicitly forecasting a $300B WFE market, its planned output capacity would support a market of that magnitude based on current share, versus Street expectations of roughly $250-260B WFE in 2028. Management stressed that it is not planning around 2028 as a peak and views AI infrastructure as a more durable demand driver than prior PC- and smartphone-led cycles.

Near-term capacity appears constrained more by customer cleanroom availability than underlying demand. AMAT is tracking >100 fabs globally, with significant DRAM capacity under construction, and indicated that revenue could be growing materially faster without cleanroom constraints. Customers continue to find ways to squeeze incremental capacity out of existing facilities while accelerating new construction. For AMAT itself, much of the physical infrastructure required to double output is already in place: manufacturing floor space has already doubled, with the remaining work focused on hiring, adding shifts, outfitting bays and optimizing flow. With approximately 80% of the cost structure variable, management believes it can prepare for a much larger WFE environment without taking much fixed-cost risk.

DRAM-related technology pull is something management has "never seen" during its career as AI drives greater customization and DRAM increasingly adopts logic-like transistor technology. Roughly 60% of DRAM is the memory array and 40% the transistor/periphery, with much of the innovation in custom DRAM occurring in the latter. DRAM is migrating toward FinFET-based architectures, while future transitions include $4F^2$ and ultimately 3D DRAM, which require additional materials engineering and high-aspect-ratio conductor etch. The company expects a strong 2H26 DRAM environment extending into 2027 and believes AMAT's existing leading-edge logic capabilities can increasingly be leveraged into memory as DRAM architectures become more complex.

Advanced packaging is also increasingly converging with front-end wafer manufacturing. Management highlighted backside power as an early example and expects it to proliferate across high-performance compute, followed over time by denser nanosheet architectures and eventually stacked transistors/CFET. The key metric across AI infrastructure is increasingly performance per watt, driving customers to reduce resistance, capacitance and transistor-level variability while improving power efficiency. This is expanding materials intensity and creating new opportunities across AMAT's portfolio. Management also noted that AI has changed the historical semiconductor development model: rather than mobile leading a node followed later by HPC, multiple GPU, custom ASIC, DRAM and packaging architectures are now being developed concurrently, increasing both process complexity and customer-specific requirements.

Another major theme was the increasing importance of integrated systems as transistor and interconnect complexity rises. Integrated solutions have grown from roughly 20% of Semi Systems revenue several years ago to ~30% today, with management expecting the mix to trend higher. At the most advanced interconnect levels, processes that historically required four steps can now require 6-7 sequential deposition, clean and materials steps performed in the same system under tightly controlled conditions. Similar integration is occurring in gate-all-around transistors, where multiple films only 1-2nm thick must be deposited in confined structures. AMAT believes this favors its broad portfolio across PVD, CVD, ALD, etch, e-beam and other technologies, particularly as customers increasingly want co-optimized solutions rather than individual process tools. The upcoming EPIC center should deepen those engagements by allowing customers and AMAT engineers to collaborate earlier on future architectures. Management argued that this integrated approach represents an underappreciated competitive advantage. The value is not simply AMAT's breadth of individual process tools, but its ability to combine them into a solution and co-optimize the overall process. As technologies previously deployed at leading-edge logic customers migrate into custom DRAM and other applications, AMAT can leverage its existing process and tool expertise without transferring customer IP. Management believes this creates both greater customer stickiness and better economics, while increasingly complex integrated systems should also pull through additional services revenue over their lifetime.

The shift toward AI is also changing both the technology roadmap and AMAT's customer engagement model. Historically, mobile would lead a new node, followed by HPC and eventually automotive; today, GPU, custom ASIC, DRAM and packaging architectures are advancing simultaneously, with performance/watt becoming the most

important metric. Customers increasingly need AMAT to help lower resistance and capacitance, tighten transistor-level variability, and improve power efficiency. Advanced packaging is also moving into the 300mm fab, with backside power an early example and denser nanosheet architectures and eventually stacked transistors/CFET creating additional materials-engineering challenges. AMAT's new EPIC center is intended to deepen engagement earlier in the development cycle and allow multiple process steps to be optimized together.

On margins, management remains confident in the longer-term trajectory/ Semi Systems gross margin is already >55%, and company gross margin has improved roughly 300bps over the past three years. More than 1,500 recently hired engineers and customer-support employees are currently in training and not yet supporting revenue, creating an approximately 30bps near-term headwind vs what they would otherwise be delivering, but margins remain over 50% and are seen ticking up quarter over quarter from here. Over time, AMAT sees upside from higher-value integrated systems and value-based pricing. Management indicated that essentially the entire go-to-market strategy has shifted toward pricing for customer value rather than simply lowering tool cost, with differentiated technologies carrying the greatest pricing opportunity. Integrated systems should also be accretive to the services business because their greater complexity increases customers' reliance on AMAT over the equipment lifecycle.

MONOLOTHIC POWER SYSTEMS (MPWR, NOT COVERED) – MEETING WITH TONY BALOW (VP FINANCE)

Management was notably constructive on both near-term demand and the longer-term AI infrastructure opportunity. Despite significant debate around AI spending, LLM monetization, and data center investment sustainability, the company indicated that nothing has changed in the underlying demand environment. Bookings were very strong in 2Q, particularly within enterprise data center, resulting in a stronger backlog and book-to-bill ratio that supported the recent outlook increase. Management's view is that the continued proliferation of AI models ultimately requires more compute and, by extension, more power, reinforcing confidence in current demand trends. Looking ahead, the company sees strong demand signals extending into 2027, supported by both unit growth and increasing power content across AI systems.

The discussion highlighted broad-based strength across MPWR's major AI-facing businesses. Enterprise data center remains the fastest-growing segment, while communications is emerging as the second-largest growth contributor, driven by power solutions for switches, DPUs, and optical networking infrastructure. Management was particularly positive on optical demand and believes there remains additional share gain opportunity. Storage has also become a larger contributor than expected, benefiting from DDR5 adoption and greater power management content, with management noting that storage is now a larger business than compute/storage originally appeared earlier in the year. By contrast, notebooks remain the weakest end market due to limited TAM growth, although management believes share gains and eventual replacement demand could create upside when the market recovers. Automotive remains a favorable long-term growth opportunity, with expectations for a stronger second-half ramp and continued growth into 2026.

A significant portion of the meeting focused on the evolution of AI power architectures and MPWR's positioning within future platforms. Management expressed confidence in its internal roadmap, including integrated voltage regulator (IVR) technologies and modular power solutions, arguing that customer requirements are increasingly moving toward higher efficiency and denser power delivery. The company believes modular solutions offer meaningful advantages versus traditional discrete approaches through improved efficiency, reduced footprint, and support for vertical power architectures. While there is substantial industry discussion around various approaches to power delivery, management appeared confident that MPWR can address these opportunities organically without requiring major acquisitions. The company is already engaged with customers several generations ahead, including platforms beyond current AI accelerator roadmaps.

The conversation also addressed competitive positioning and share sustainability in AI power systems. Management acknowledged that the largest AI platforms are unlikely to remain single-sourced over time, particularly as power franchises grow into multi-billion-dollar opportunities. However, the company believes technical differentiation, customer engagement, time-to-market, supply assurance, and cost competitiveness continue to support share gains. Importantly, management noted that design wins are only the first step; procurement decisions often determine final share allocations, creating some variability in market share even after qualification. Nevertheless, the company indicated it feels good about current share trends and emphasized that its customer base has become significantly more diversified over the past year.

Investors also spent considerable time discussing the industry's migration toward higher-voltage data center power architectures. Management characterized 800V power distribution as a meaningful future opportunity but one that is likely to be pushed toward late 2027 or 2028 rather than become a material contributor next year. The company expects the first phase of adoption to focus on converting 800V power down to intermediate voltages before broader architectural changes occur across data centers. While debate continues around technologies such as GaN and SiC, management's focus remains on overall system performance rather than material-specific solutions. Importantly, MPWR believes the opportunity spans multiple stages of the power conversion stack and can leverage both its power board expertise and modular power IP as deployments expand.

On operations and profitability, management maintained a confident tone regarding both supply chain readiness and long-term margin performance. Visibility has improved meaningfully relative to last year, with customers providing better demand forecasts and the overall supply environment becoming more orderly. The company's multi-foundry strategy remains a differentiator, with manufacturing spread across numerous partners and fabs, allowing capacity to be added incrementally while maintaining flexibility and supply assurance. Management also reiterated confidence in sustaining gross margins within its historical framework, noting that input cost inflation has remained manageable and largely offset through pricing and product mix. As growth accelerates, the company expects operating leverage to become increasingly evident given that revenue can scale materially faster than headcount-driven operating expenses.

ASTERA LABS (ALAB, NOT COVERED) - MEETING WITH JITENDRA MOHAN (CEO), DES LYNCH (CFO), AND NICHOLAS ABERLE (IR)

Management's core message was that Astera remains in the early stages of a multi-year AI infrastructure buildout, with its opportunity expanding well beyond its original PCIe connectivity business. The company's latest Leo 2X CXL memory expansion product targets growing AI memory requirements, particularly around KV-cache and agentic AI workloads where context lengths are increasing rapidly. Management believes memory expansion solutions become increasingly valuable as customers look for alternatives once scarce and expensive HBM capacity is exhausted.

The discussion reinforced how much larger the AI opportunity is relative to the traditional server market. While conventional servers remain a meaningful business, AI infrastructure carries substantially higher content opportunities (order of magnitude) due to the complexity of scale-up architectures. Management expects content per accelerator to continue increasing across future generations and believes the largest long-term opportunity comes from participating deeper in the connectivity stack, particularly within AI clusters rather than traditional enterprise servers.

A key theme was Astera's effort to become a comprehensive connectivity provider across multiple AI ecosystems. The company has expanded into PCIe, retimers, switches, CXL memory, UALink, and NVLink Fusion-related solutions, with management emphasizing a strategy of delivering complete platform solutions rather than isolated point products. Astera currently has two active NVLink Fusion engagements and views itself as a neutral enabler for hyperscalers and ASIC providers seeking alternatives to fully integrated proprietary ecosystems.

Management was also constructive on the scale-up roadmap. UALink remains a 2027 deployment story with more meaningful revenue expected in 2028, while the broader objective is to participate in scale-up architectures across all major hyperscalers. The company believes its opportunity set increases materially when involved in scale-up networks, where content per accelerator is significantly higher than in traditional scale-out deployments. Interest from hyperscalers remains healthy, and management's long-term goal is to establish a presence across multiple customer architectures rather than relying on any single ecosystem.

Optical networking emerged as another important long-term growth vector. Management expects copper and optical interconnects to coexist for several years but believes growing cluster sizes will ultimately drive broader optical adoption. Near-packaged optics (NPO) could begin appearing around 2027, while co-packaged optics (CPO) is likely further out given cost and supply chain constraints. The company is investing in silicon photonics and optical switching technologies and views optics as both a content expansion and share gain opportunity as AI clusters continue to scale.

Finally, management expressed confidence in customer demand, product roadmaps, and supply readiness heading into 2027. While wafer supply remains concentrated at TSMC, the company believes capacity is sufficient and currently views substrates and backend manufacturing as larger areas of focus. Overall, management argued that investors may underestimate both the breadth of Astera's portfolio and the duration of the opportunity, pointing to CXL memory expansion, UALink, NVLink Fusion, optical networking, and increasing content per accelerator as drivers of sustained growth over the next several years.

LATTICE SEMICONDUCTOR (LSCC, NOT COVERED) - MEETING WITH RICK MUSCHA (IR)

Lattice was constructive on the AI infrastructure opportunity and believes FPGA content growth is still in the early innings. The company's server attach rate has increased materially over the past several years, rising from roughly 0.25x per traditional server in 2018-19 when Lattice primarily provided power sequencing and I/O functions, to over 2x today as security and infrastructure management content increased. AI servers are driving another step-function higher and with their CPU/GPU agnostic approach, they noted that customer deployments have evolved from roughly 50 FPGAs per rack in 2024 to some designs utilizing 100+ FPGAs, while AI racks spanning compute, networking, power, and cooling can require several hundred FPGAs. As a result, blended server attach rates increased from approximately 1.5x in 2024 to 2.5x in 2025, with expectations for 3.5-3.6x in 2026.

The company remains particularly bullish on the data center growth outlook. Server TAM is expected to expand from approximately 16.5 million servers in 2025 to roughly 20 million in 2026, but they believe FPGA growth will significantly outpace server unit growth. FPGA server units are expected to grow more than 100% this year, while ASPs are increasing ~30%, driving server revenue growth of ~150%. Management expects server revenue to exceed $300M this year, with the business now comprising ~50% traditional servers and 50% AI infrastructure. Importantly, Lattice emphasized its position as a neutral supplier across CPUs, GPUs, TPUs, and custom XPUs, allowing it to benefit regardless of which compute architectures ultimately gain share.

Competition was also addressed following Altera's separation from Intel. According to the company, it remains differentiated through its low-power architecture, which targets a lower-end segment of the FPGA market than AMD/Xilinx and Altera. Whereas AMD/Xilinx and Altera have historically focused on larger, performance-oriented devices built around 6-LUT architectures, Lattice's 4-LUT architecture is optimized for power efficiency and smaller form factors. Management acknowledged reports that Altera intends to move further into lower-end FPGAs following the majority sale to Silver Lake but also noted that Altera is going through a headcount reduction and that such design efforts would take years as competitors have not meaningfully refreshed their roadmaps in this segment. More broadly, they emphasized that increasing server complexity, security requirements, and infrastructure management needs continue to expand the addressable opportunity faster than competitive pressures are emerging.

Security was another focus of the discussion, with the company highlighting what it believes is a meaningful competitive advantage in post-quantum cryptography (PQC) and secure root-of-trust solutions. They indicated that it is currently among the few semiconductor vendors compliant with emerging U.S. and European security standards and expects PQC-related revenue to begin contributing around mid-to-late 2027. While management declined to quantify the precise ASP uplift, it characterized it as meaningful. The combination of reprogrammable hardware and evolving security standards is viewed as a particularly attractive use case for FPGAs, as customers can update cryptographic implementations without redesigning hardware.

The recently announced AMI acquisition for ~$1.5B was described as highly strategic and significantly expands Lattice's role within server infrastructure. AMI is expected to generate roughly $200M of revenue in 2025, growing approximately 19% this year and 25% next year. Management believes the acquisition transforms Lattice from a component supplier into a broader infrastructure management platform provider through capabilities spanning boot firmware, secure boot, and system manageability software. The company cited opportunities to pair a roughly $5 security FPGA with a software and firmware solution worth five to six times the hardware value. Beyond standalone growth, management sees substantial revenue synergies around secure infrastructure, power and cooling management, and heterogeneous AI systems. The transaction expands Lattice's TAM from approximately $6B today to $12B by 2030.

Lattice was confident in the longer-term aspirational outlook of $3B by 2030. Near-term, while consensus sits near $1.3B of revenue in 2027, the company suggested a range of $1.3-1.5B appears achievable, including approximately $250-300M of AMI revenue. This outlook assumes recovery in industrial and embedded markets, which they expect to return to growth in 2026, alongside continued strength in compute and communications. The newer Avant FPGA family remains small today, generating only low-single-digit millions of revenue this year, but management pointed to strong design-win momentum and expects a more meaningful contribution beginning in 2027.

Looking beyond AI data centers, the company highlighted emerging opportunities in physical AI, robotics, and edge inference. Lattice already has exposure to warehouse automation, industrial robotics, machine vision, and autonomous systems, but believes humanoid robotics could become a meaningful future market. The company sees

opportunities in vision processing, sensor aggregation, and motor control, suggesting future humanoid platforms could contain as many as 75 FPGAs at roughly $5 ASPs, creating content opportunities comparable to some AI server deployments. While these opportunities are likely a 2027-28 and beyond story, they view them as another example of how increasing system complexity continues to expand the role of low-power programmable devices.

GLOBALFOUNDRIES (GFS, NOT COVERED) - MEETING WITH MICHAEL HOGAN (CHIEF BUSINESS OFFICER)

Management's tone was positive, with improving conditions across most end markets and particular strength tied to AI infrastructure. The company continues to benefit from AI-related demand for power delivery, silicon photonics (SiPh), and silicon germanium (SiGe), while inventory corrections in broader infrastructure and industrial markets appear to be easing. Utilization reached the high-80% range last quarter, and management believes it can ultimately move into the 90-95% range, supporting both revenue growth and margin expansion. The company also highlighted over $675M of government awards across photonics ($300M) and quantum-related programs ($375M), reinforcing its strategic position within the U.S. semiconductor ecosystem. GFS also emphasized its broader tech portfolio in addition to its traditional foundry business and the company believes that the market underestimates its technology portfolio, pointing to areas such as photonics, RF, power, GaN, advanced packaging, and specialty nodes where innovation is not necessarily dependent on transistor scaling. The acquisitions of MIPS and processor IP assets from Synopsys have also expanded the company's capabilities, making GFS the largest provider of RISC-V IP and allowing it to engage customers earlier in the design cycle. Management noted that discussions increasingly begin around architecture and IP before customers ultimately select a manufacturing process, creating a tighter link between its IP and foundry businesses.

Silicon photonics was the most discussed growth opportunity. Management expects photonics revenue to double in 2025 and double again in 2026, reaching a $1B annualized run-rate by 2028, while noting that the adjacent SiGe business is already larger and growing on a similar trajectory. Unlike many competitors which may possess some but not all required building blocks to compete effectively, GFS believes it possesses the combination of technology, customer relationships, and 300mm manufacturing scale required to participate meaningfully in the market. The company has approximately 300k-350k wafers of photonics-capable capacity in its New York facility that can be converted as demand grows, and stressed that capacity, not customer demand, is the primary bottleneck today. Importantly, management noted that its $1B target is largely driven by existing pluggable optics demand and does not require a rapid transition to NPO or CPO architectures.

Capacity was a major focus of the discussion. GFS currently has approximately 2.8M wafers/year of tooled capacity, with existing cleanroom infrastructure allowing it to expand to roughly 3.2M wafers/year, which management estimates could support at least $9-10B of annual revenue. Importantly, the revenue opportunity could be higher as capacity is remixed toward higher-value wafers: the corporate average wafer ASP is roughly $2,500, while SiPh wafers can generate multiples of that with fewer process layers and shorter cycle times. In its New York fab alone, GFS has roughly 300K-350K wafers of "dry powder" that could be converted to photonics through incremental tooling or by reallocating lower-value capacity. Management emphasized that it already has the customers and tapeouts, making capacity expansion the primary constraint rather than demand generation. In addition to pure capacity, geographic diversification is becoming another meaningful driver of customer engagement with significant customer interest in non-China/non-Taiwan supply, including customers that specifically want U.S. demand served from U.S. fabs. GFS's largest fab remains in Singapore, which management views as a geopolitically attractive location, while its U.S. manufacturing footprint provides additional differentiation.

This existing capacity also supports the margin upside. GFS reiterated its target for 40% gross margin exiting 2028 and 45% longer term, with every 500bps increase in utilization worth roughly 200bps of gross margin. More importantly, management expects 500-600bps of improvement from mix as it reallocates capacity away from lower-margin businesses toward constrained technologies such as SiPh, SiGe and 22nm. Customers are increasingly helping fund these expansions through LTAs, take-or-pay agreements and other capacity commitments. Management acknowledged that LTAs proved less durable than expected following COVID but argued that current agreements are increasingly tied to sole-sourced technologies and more durable markets, making the economics and utilization profile structurally better.

Smartphones which historically represented roughly 50% of GFS revenue and were still 36% last quarter, are no longer expected to be a meaningful future growth driver with management's longer-term share expectation of ~25%. In the short term, the company expects the market to remain pressured due to memory, although GFS's mix is somewhat insulated because ~2/3 of its handset exposure is tied to the premium segment. Management's strategy is increasingly to harvest the smartphone business while reallocating investment toward infrastructure, auto and other more durable markets, while leveraging technologies developed for smartphones in areas such as wearables and physical AI.

Finally, the company only expects minor impact from Apple's move away from Qualcomm's cellular modem. GFS's exposure to the affected Qualcomm content was limited primarily to small mmWave and envelope-tracking

opportunities, while its more material smartphone exposure remains in the RF front end, power management, audio and analog. Management therefore does not view Qualcomm's loss of the Apple modem socket as a meaningful read-through for GFS and believes it can offset the limited lost content elsewhere.

AMBARELLA (AMBA, NOT COVERED) - MEETING WITH FERMI WANG (CEO), JOHN YOUNG (CFO) AND LOUIS GERHARDY (VP CORPORATE DEVELOPMENT AND IR)

Management remains constructive on the long-term edge/physical AI opportunity, recently increasing AMBA's SAM to $23B by CY31/FY32, driven primarily by higher-ASP edge infrastructure products. The company believes it can grow faster than the underlying market on a CAGR basis from 2026-31, although management cautioned against expecting smooth annual growth. As context, its prior SAM implied an ~18% CAGR, while AMBA revenue grew 26% two years ago and 37% last year before decelerating this year. Management views itself as an established physical AI supplier rather than a new entrant: roughly 90% of current revenue fits its definition of physical AI, and the company has shipped ~50M SoCs across a portfolio spanning 1 to 1,000 TOPS, supported by a common SDK.

Physical AI is expected to develop through a broad collection of use cases rather than a single "killer app", leading AMBA to pursue a portfolio approach, casting a wide net of solutions. AMBA already participates in enterprise cameras, robotics and other sensor-based applications, and its role is evolving from perception only toward systems that also perform planning and ultimately take action. Management sees humanoid robotics as potentially significant but believes widespread adoption remains relatively far out, with more incremental robotics applications likely to scale first. The principal challenge is optimizing models sufficiently to run locally within edge power and cost constraints; a very large model provides little benefit if it cannot be distilled effectively for an endpoint device. AMBA believes its architecture is differentiated, with a proprietary inference engine and >200 model architectures already ported to its platform, although management expects cloud services to remain part of the overall solution.

Regarding competition with Qualcomm, management noted that the two companies have competed for ~15 years, but believes AMBA's architecture is better optimized for edge AI performance/watt and cost given that Qualcomm's architecture originated as a smartphone application processor, whereas AMBA designs specifically around workloads executed locally, reducing reliance on cloud processing and associated latency, security and connectivity requirements. The newer CV72 and CV75 5nm products are currently ramping rapidly and support both CNN and transformer-based workloads. Partnerships with system integrators such as Capgemini should broaden distribution, particularly among customers finding conventional GPUs too power-intensive, while other partners can economically address fragmented opportunities where individual customers may require only ~2k chips/month.

The largest near-term uncertainty is memory. AMBA does not procure DRAM directly, but its memory controller means it must work with customers whenever they change suppliers, migrate between DDR generations or de-spec memory. Management is seeing customers pursue multiple strategies to address rising prices and constrained availability, but visibility remains poor: some customers still lack firm allocations/pricing for November deliveries, with even less clarity for December. AMBA remains comfortable with its 3Q guidance, but views memory availability as a potential constraint into next year. Chinese DRAM suppliers could indirectly alleviate pressure by serving domestic Chinese customers and freeing capacity elsewhere, although management is not seeing U.S. customers broadly adopt Chinese memory.

Supply inflation extends beyond memory, with AMBA seeing price pressure across much of its supply chain. In response, the company is implementing broad-based price increases for the first time in its history, although these should take roughly six months to flow through to customers. Management still expects to maintain its long-term ~59-60% gross margin framework and guided 3Q to approximately 59.5% at the midpoint. Inventory had also risen as AMBA built protection against a potential Samsung strike and supported the steep CV72/CV75 ramp, but management indicated that inventories have begun to decline and remain concentrated in high-volume products.

LIGHTMATTER (PRIVATE) - MEETING WITH NICK HARRIS (FOUNDER AND CEO), AND SIMONA JANKOWSKI, (CFO)

The technology discussion focused heavily on Lightmatter's differentiation at the photonic interconnect level. Management believes one of the biggest challenges as scale-up domains expand is the number of physical fibers required to connect accelerators, particularly given that conventional links use separate fibers for transmit and receive. Lightmatter has demonstrated 16 wavelengths on a single bidirectional fiber, allowing one fiber to perform both functions while materially increasing bandwidth density. Its link provides 800Gbps per fiber, 400Gbps in each direction, and management believes this type of bidirectional architecture becomes increasingly important as scale-up domains grow. Management noted that it has already demonstrated scale-up domains incorporating roughly 131k GPUs, while longer term expecting architectures to move away from pure all-to-all connectivity because the number of required connections grows exponentially becomes increasingly difficult to scale.

Lightmatter is building a broader photonics stack. Their Guide laser platform uses a 300mm wafer-based approach to integrate large numbers of lasers, rather than relying exclusively on traditional discrete laser architectures. Management discussed configurations integrating 128 lasers per chip and 16 wavelengths of light, which should improve both density and manufacturability. The laser products can be sold independently but also power Lightmatter's own Passage interconnect platform, giving the company control over both the light source and photonic interconnect. The strategy is particularly relevant as increasing optical bandwidth requires greater laser density rather than simply increasing laser power.

Another key area of differentiation is Lightmatter's ability to combine its interposer, laser and modulation technologies. Management highlighted micro-ring modulators and its ability to separate and control multiple wavelengths while maintaining low insertion loss, with bidirectional transmission particularly difficult to execute. The company's architecture enables multiple wavelengths to share the same fiber rather than simply adding more fibers as bandwidth grows. This is intended to address both the bandwidth and physical I/O limitations that emerge as accelerator clusters scale, and management believes it is materially ahead in bidirectional wavelength management.

Near-packaged optics (NPO) has emerged as the most important near-term opportunity, with customer trials expected in 2027 and a much larger volume ramp in 2028. The advantage of NPO is that it delivers much of the bandwidth and power benefit of CPO while leveraging the existing pluggable optics manufacturing and test infrastructure, significantly reducing implementation risk.

On ecosystem positioning, management indicated Lightmatter is working with multiple compute platforms, including Qualcomm, and has joined NVIDIA's NVLink Fusion ecosystem, reinforcing the company's strategy of remaining compute-architecture agnostic. The objective is to provide the optical connectivity layer as scale-up networks expand, irrespective of the underlying XPU. Management expects initial product sampling in 1Q followed by shipments during 2027, with a much larger opportunity developing in 2028 as NPO adoption broadens.