KASASV

KOREAN AMERICAN SEMICONDUCTOR ASSOCIATION IN SILICON VALLEY

Jan. 2026

어떻게 쓰이고 싶은가?

지금 하고 있는 일은 처음부터 원했던 게 아니었지만 **계속 해도 괜찮다는 답이 나오면 그렇게 해나가면 됩니다.**

그러나 **'나는 회사에 이렇게 쓰이고 싶진 않다'는 생각이 명확하다면 다르게 쓰일 길을 찾아야죠.**

중요한 것은, 어떻게 쓰이고 싶은가에 대한 질문에 자신만의 답을 찾아가는 겁니다.

우리는 누군가에게 도움이 될 때도 즐거움을 느끼는 존재들입니다.

당장 이익이 생기지 않는다 해도 내가 맡아서 한 업무를 통해 누군가 도움을 받고 기뻐하는 걸 보면 기분이 좋아지고 심지어 행복해질 때까지 있죠. '나에게 이런 면이 있었나?' 하고 돌아보게 되기도 하고요. 즉 어떻게 쓰이고 싶은지,

내가 아는 나의 재능과 취향, 선호를 어떻게 썼을 때 자신의 성장과 더불어 내가 속한 곳에 대한 기여도 커질 수 있을지에 대해선 계속 생각해 봐야 합니다.

이 질문은 평생 가까이 해야 합니다. 인생의 어느 시점에서 답을 찾았다 하더라도 시간이 흐르고 경험이 쌓이며 상황이 변하면 답은 또 달라질 수 있으니까요.

-내가 가진 것을 세상이 원하게 하라, 최인아

INTRO

E-mail 주소는 blind 로 나가기때문에 공유가 되지는 않습니다. 혹시 개인 메일주소를 원하시면 저에게 알려주시면 update 하도록 하겠습니다. 그리고 Newsletter distribution 을 원치않으시면, 알려주시기 바랍니다.

12 월 27 일(토)에는 올해 4 번째 Golf Tournament 가 있었습니다. 참석해주신 모든 분들께 감사를 드립니다.

자세한 내용은 웹사이트(announcement)를 참조하시기 바랍니다.

Linkedin 에 KASASV 를 officially launch 했습니다. 많은 관심 부탁드립니다.

저희는 미국 Internal Revenue Code 501(c)(3)에 의거하여 Federal Income tax exemption 을 받은 비영리단체로 donation/협찬으로 운영이 되고 있습 니다. 관심이 있으시면 저희 웹사이트 Partners tab 을 보시면 좀 더 자세히 보실수 있습니다.

- 사이먼 리 드림/협회 상근임원

www.kasainsv.com 문의 info@kasainsv.com

감사합니다

3D-IC Market Outlook: Technology Roadmaps, Readiness, And Design Implications

Advanced packaging technologies are reshaping how compute platforms are conceived, optimized, and manufactured.

Dec 23, 2025 Reela Samuel

The 3D-IC market outlook is entering a decisive phase as the semiconductor industry transitions beyond the limits of traditional Moore's Law scaling. As performance, power efficiency, and system complexity outpace what planar integration can deliver economically, vertical integration and heterogeneous system design are no longer experimental; they are becoming foundational. Advanced packaging technologies such as stacked dies, hybrid bonding, and chiplet-based architectures are reshaping how future compute platforms are conceived, optimized, and manufactured.

This shift is not incremental. It reflects a broader architectural inflection point where performance gains increasingly come from proximity—bringing compute, memory, and specialized accelerators closer together in three dimensions. Industry analysts project that the global 3D-IC packaging market will reach approximately USD 32.9 billion by 2030, growing at a rate of roughly 15% CAGR—fueled by AI acceleration, high-performance computing (HPC), memory stacking, and aggressive packaging capacity expansion across Asia, North America, and Europe. In this context, 3D-IC technology has moved from a niche enabler to a mainstream design imperative.

This blog explores the key 3D-IC market trends, adoption drivers by vertical, the 3D-IC technology roadmap for packaging and interconnect innovation, ecosystem standardization progress, and the resulting design implications for tools and methodologies. It concludes with a practical roadmap checklist to help engineering organizations prepare for the next generation of heterogeneous integration.

Market drivers and investment signals by vertical

Growth in the 3D-IC and heterogeneous integration market is rooted in fundamental changes in compute demand and semiconductor economics. As traditional transistor scaling delivers diminishing returns in performance per watt and cost per function, system architects are shifting toward vertical integration and chiplet-based partitioning to sustain generational improvements

AI and HPC packaging demand

AI and HPC workloads, particularly large language models (LLMs), recommendation engines, and real-time analytics, demand extreme memory bandwidth, low latency, and tightly coupled compute-memory architectures. 3D-IC designs that integrate logic dies with multiple HBM stacks address these requirements directly, enabling bandwidth scaling that would be impractical with off-package memory. Hyperscalers and accelerator vendors are increasingly signaling roadmaps that prioritize memory stacking and die-to-die interconnect density as primary performance levers.

The design implication here is that memory hierarchy, floorplanning, and power delivery must now be optimized at the package level, not just at the die level.

Telecom and 5G infrastructure

The rollout of 5G infrastructure and edge compute nodes continues to drive demand for compact, high-throughput modules. Advanced OFDM and massive MIMO radios benefit from 2.5D and 3D packaging approaches that support dense interconnects, minimize signal loss, and facilitate the tighter integration of RF, baseband, and processing elements.

Automotive and IoT adoption

Automotive electronics, including ADAS and autonomous platforms, represent one of the fastest-growing adoption vectors for heterogeneous integration. These systems must balance high compute performance with stringent thermal, reliability, and lifecycle requirements. Similarly, consumer and internet of things (IoT) devices leverage stacked architectures to achieve form-factor reduction and power efficiency without sacrificing functionality.

Geopolitical and foundry investment trends

On the supply side, advanced packaging capacity is expanding rapidly. The Asia-Pacific region remains the manufacturing center of gravity, while North America and Europe are investing heavily in sovereign semiconductor initiatives to localize advanced packaging and chiplet assembly capabilities. These investments reinforce the long-term viability and scalability of the 3D-IC ecosystem.

Together, these drivers indicate sustained, multi-vertical momentum for 3D-IC adoption, anchoring advanced packaging firmly within future semiconductor roadmaps.

Packaging and interconnect innovations on the horizon

The advanced packaging innovation roadmap is accelerating, with multiple technologies redefining interconnect density, bandwidth scaling, and thermal constraints.

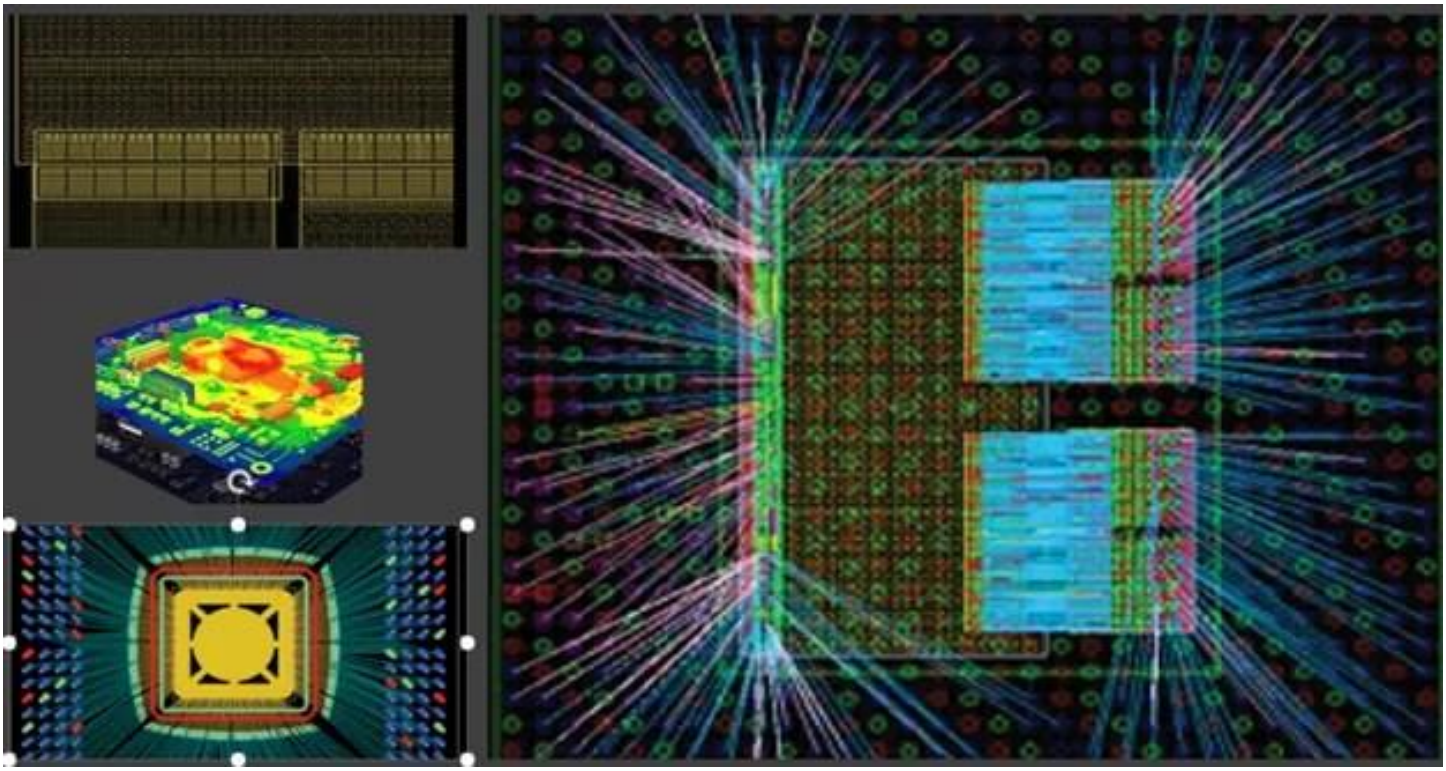

Hybrid bonding advancements

Hybrid bonding, enabling direct copper-to-copper connections at pitches below 10 micrometers, is emerging as a cornerstone of next-generation 3D integration. Compared to traditional micro-bump approaches, hybrid bonding dramatically increases I/O density, reduces parasitics, and lowers z-height, making it essential for high-bandwidth AI and HPC systems.

The architectural impact is that finer-pitch bonding enables more granular die partitioning and tighter coupling between compute and memory, reshaping how architects allocate functionality across dies.

TSVs and next-generation interposers

While hybrid bonding adoption accelerates, TSVs remain foundational for vertical connectivity. In parallel, innovation in silicon, glass, and advanced organic interposers aims to strike a balance between electrical performance, cost, and manufacturability. These developments extend the scalability of redistribution layers (RDL) and substrate technologies.

Thermal and material innovations

As stacked designs push power densities beyond 1 W/mm², thermal management becomes a first-order design constraint. Research and early commercialization efforts are exploring microfluidic cooling, enhanced thermal via networks, and advanced die-attach materials to mitigate hotspots and improve long-term reliability.

Wafer-to-wafer and chiplet stacking

The commercial deployment of wafer-to-wafer hybrid bonding and multi-die HBM3E stacks, which exceed 1TB/s aggregate bandwidth, demonstrates that true 3D stacking is now viable for production-scale systems. These advances significantly expand the performance envelope of 3D-IC technologies.

Standardization and the chiplet ecosystem

Mainstream adoption of 3D-ICs depends on a robust, interoperable chiplet ecosystem supported by industry-wide standards.

UCIe and interoperability

The Universal Chiplet Interconnect Express (UCIe) standard defines physical, electrical, and protocol layers for die-to-die communication. Its adoption reduces integration risk, accelerates chiplet reuse, and enables multi-vendor modular systems—critical for ecosystem scalability.

DFT and 3D test standards

Standards such as IEEE 1838 address test access architectures for stacked ICs, enabling known good die (KGD) strategies and improving yield visibility across complex assemblies. Testability is increasingly a differentiator as stacks grow deeper and more heterogeneous.

Consortium-driven collaboration

Ecosystem progress relies on close collaboration among foundries, OSATs, EDA vendors, IP providers, and system companies. Consortium-driven models help align standards, manufacturing readiness, and tool qualifications, reducing friction across the value chain.

Implications for tool flows and methodologies

The transition to 3D-ICs fundamentally changes how systems are designed, analyzed, and signed off.

Unified multi-die EDA platforms

Design teams require integrated platforms that span chip, package, and system domains. Unified multi-die environments enable early co-design of partitioning, interconnect planning, and power/thermal tradeoffs—reducing late-stage surprises.

System-level multiphysics analysis

Stacked architectures amplify interactions between electrical, thermal, and mechanical effects. Integrated multiphysics analysis for SI, PI, EM, thermal stress, and warpage is essential for achieving predictable signoff and manufacturing yield.

3D design kits and IP reuse

Pre-validated 3D design kits (3DKs) and reusable chiplet IP reduce integration risk and compress development timelines, particularly in multi-vendor environments.

Silicon-to-system convergence

Preserving connectivity from silicon through package to system enables manufacturing-aware design decisions, improves yield predictability, and supports system-level optimization in heterogeneous architectures.

Roadmap checklist: Readiness for upcoming nodes

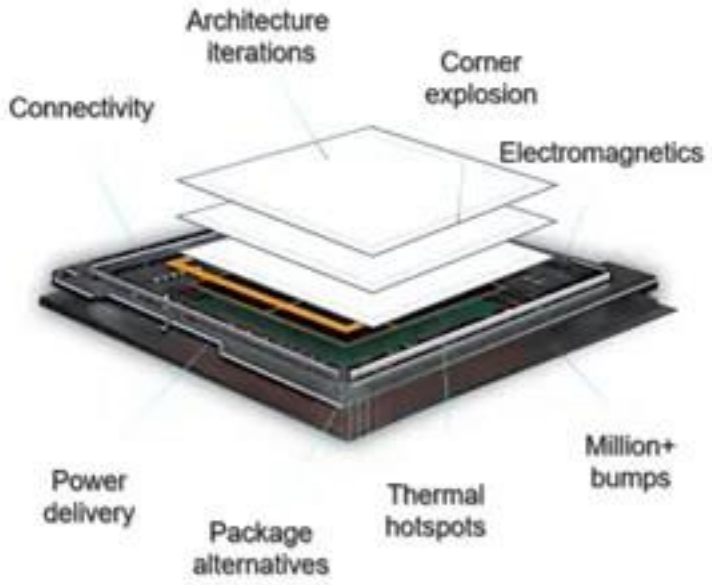

To prepare for future 3D-IC nodes, design teams should focus on the following:

- Adopt a Heterogeneous Integration Strategy: Architect systems as collections of optimized chiplets rather than monolithic SoCs.
- Invest in Advanced EDA Toolchains: Ensure support for unified visualization, automated floorplanning, and large-scale multiphysics analysis.
- Implement Industry Standards: Embrace UCIe and IEEE 1838 to improve interoperability, testability, and yield.
- Engage with the Ecosystem: Participate in consortia and closely track foundry and OSAT roadmaps.
- Plan Thermal and Power Early: Address thermal constraints during architecture definition, not after layout.
- Align with Manufacturing Partners: Understand hybrid bonding readiness, wafer-level processes, and supply chain constraints early.

A transformational era for 3D-IC technology

The future of 3D-IC technology reflects a structural transformation across the semiconductor landscape. From AI and HPC to automotive and communications, vertical integration and heterogeneous innovation are redefining performance scaling, manufacturing strategies, and design methodologies. Advanced interconnects, maturing standards, and expanding packaging capacity are accelerating the shift from experimental adoption to mainstream deployment.

Cadence's Integrity 3D-IC Platform—combined with system-aware, multiphysics analysis workflows—supports this transition by enabling predictable, manufacturing-aware design closure across complex 3D systems. As the industry rethinks how compute platforms are architected and delivered, 3D-IC technology stands at the center of a new era—one defined by modularity, proximity, and system-level optimization rather than transistor scaling alone.

Inside Intel's new Arizona fab, where the chipmaker's fate hangs in the balance

Dec 19, 2025 Katie Tarasov

KEY POINTS

- Intel's advanced chip node, 18A, is now in high-volume production at its new Arizona fab, but no major outside customers have emerged
- The fab is meant to help Intel catch back up to TSMC after years of missteps, with executives saying it's "turned the corner."
- Lifelines include $8.9 billion from the U.S. government and $5 billion from Nvidia, although the AI chip leader didn't commit to making chips at Intel Foundry

Intel was once the world's largest semiconductor company, but its market cap plummeted in recent years as the chipmaker fell behind Taiwan Semiconductor Manufacturing Co. and spent billions of dollars trying to catch up.

Now, Intel has entered high-volume production of 18A, the new chip node it says will turn things around.

The biggest problem? Convincing a big chipmaker to trust Intel with manufacturing on the new node. For now, Intel's only major customer is itself. The company's long-awaited Core Ultra series 3 PC processor, code-named Panther Lake, will come to PCs in January as the first major product made on 18A.

"It's become an internal node for now," said Daniel Newman, CEO of Futurum Group. "So many companies have made such massive investments into TSMC to ensure yield, to ensure capacity wafers that they just will not make the switch just yet."

Intel is pinning its hopes of attracting customers on a new chip fabrication plant, Fab52, in Chandler, Arizona, where CNBC got an exclusive on-camera tour in November. Some 50 miles north, in Phoenix, TSMC also has a new fab, where it's making chips with 4 nanometer technology. Its most advanced 2nm tech is currently only made in Taiwan.

Intel's 18A is generally on par with TSMC's 2nm node in some metrics such as transistor density. But as Intel works out the kinks after years of delays on previous nodes, some 18A wafers have had defects, making for a lower number of usable chips per wafer, typically referred to as yield.

"Yields are always an issue at the advanced node. This is not an uncommon problem," said Harvard Business School professor David Yoffie, who served on Intel's board from 1989 to 2018. He pointed to early yield issues with Nvidia

's Blackwell GPUs at TSMC that were quickly fixed.

Intel's renewed focus on foundry — manufacturing chips for outside clients — came when Pat Gelsinger took the helm as CEO in 2021. Gelsinger was pushed out last December and replaced by Lip-Bu Tan in March.

"Over the past several years, the company invested too much, too soon – without adequate demand," Tan wrote in a memo in July.

Intel's campus in Chandler, Arizona, now includes five chip fabrication plants, with Fab52 being the newest addition, shown here on Nov 17, 2025. Tony Puryol

With Intel awaiting a big outside customer, the U.S. government stepped up in August, taking a 10% stake in the company with an $8.9 billion investment, primarily coming from grants promised under the CHIPS Act signed by President Joe Biden in 2022.

Days earlier, SoftBank invested $2 billion in Intel, followed by a $5 billion investment in September from Nvidia, which agreed to use some Intel technology but didn't commit to using its foundry.

Here's a look behind the curtain of Intel's new chip factory where it hopes to find major foundry customers and, with them, redemption.

Fall of a giant

Founded in 1968 by Silicon Valley chip pioneers Robert Noyce and Gordon Moore and legendary investor Arthur Rock, Intel brought the world's first commercially available microprocessor to market just three years later.

From the late 1970s through the early 2000s, Intel pumped out increasingly advanced process nodes at a rapid pace, leading to the term "Moore's Law" — the doubling of components on a chip every couple years.

"The 1990s was a period of wonder and excitement at Intel," Yoffie said. "We were the world's largest semiconductor company, the world's most profitable."

But Intel largely missed the mobile revolution, famously turning down a deal to make Apple's processors for the original iPhone. Then came a whiff in AI.

In 2024, Intel saw its worst year ever, losing about 60% of its value. The plunge came after two of its previous chip nodes, 10nm and 7nm, were delayed by several years. Analysts say the delays may have been triggered by an earlier choice to hold off on using ASML's costly Extreme Ultraviolet Lithography machines.

"I think we lost the discipline of cycle time," said client computing head Jim Johnson, who joined Intel more than 30 years ago. "Cycle time requires you to commit and deliver, and we started talking ourselves into, hey, we can have longer cycle times and try and lift more or do more."

As it hustles to get back on track, Intel told CNBC there will be at least 15 EUV machines in Fab52.

Intel 18A production manager Lea Tensuan shows CNBC's Katie Tarasov the EUV machines inside Fab52 in Chandler, Arizona, on Nov 17, 2025

By 2021, TSMC had become the node leader, and Intel began to outsource some leading-edge chip production to the Taiwanese giant. Around the same time, Apple

began replacing Intel chips in Mac computers with its own M-series chips, also manufactured primarily at TSMC.

In his earlier stint at Intel, more than a decade before rejoining as CEO, Gelsinger "was given the responsibility to build a GPU to compete with Nvidia," Yoffie said. "Unfortunately, that project failed and that ultimately meant we ended up not playing a significant role in the AI revolution."

Intel may now be considering a deal to buy custom AI chip design startup SambaNova for $1.6 billion, though the company declined to comment on the matter.

'Changing our culture'

The trademark of Gelsinger's tenure as CEO was Intel's focus on chip manufacturing. His ambitious roadmap had Intel catching back up to TSMC by releasing five nodes in four years.

Now, Tan is CEO and Naga Chandrasekaran is in charge of foundry.

"We are making yield improvements, defect density improvements, month-over-month and hitting our goals," Chandrasekaran told CNBC in an interview in November. "So I believe we have turned the corner."

Chandrasekaran joined Intel last year after more than two decades at leading memory maker Micron

. He said his top goal is finding foundry customers.

"I have to become part of their team and convince them that they can trust Intel Foundry to execute," Chandrasekaran said. "That's number one. And to do that, we are changing our culture. We are bringing a huge execution focus internally into Intel Foundry."

Chandrasekaran told CNBC that Fab52 is capable of more than 10,000 18A wafer starts per week. There's more than a million square feet of manufacturing cleanroom space in Arizona, with five fabs all connected by 30 miles of overhead track moving wafers between them. A sixth fab, Fab62, is expected to be ready around 2028.

18A also uses RibbonFet, Intel's gate-all-around architecture that improves power control by fully surrounding the transistor, unlike previous designs that only contact the top and sides. Chandrasekaran said 18A offers "more than 15% performance per watt improvement" over Intel 3.

Perhaps the biggest way Intel stands out is in advanced packaging, the assembly and connections of chips onto the final systems where they appear in real-world applications.

Intel engineer Shripad Gokhale shows its next Xeon data center chip in intel's advanced packaging lab in Cahndler, Arizona. Tony Puyol

CNBC went to Intel's advanced packaging lab in Chandler to see several steps in the process, such as protecting chips with a polymer-based seal, and exposing them to a liquid that detects any defects. Yoffie said Intel's advanced packaging "can help mitigate some of the power consumption problems."

"One of the biggest problems today for everybody making chips for data centers is the power that it consumes," Yoffie said.

Chandrasekaran said the Arizona fab is on almost 100% renewable energy. As for water, Intel's Arizona facilities used more than 3 billion gallons in 2024 and returned 2.4 billion gallons to the local supply through a water recycling plant it has on site.

'No blank checks'

Tan's message to employees when it comes to spending on future foundry nodes is clear: "No more blank checks." The company needs customers.

Intel's big new Ohio chip fab is delayed until at least 2030, and Tan has made major cost cuts by slashing 15% of the workforce in July and axing projects in Germany and Poland.

"That's what the company needed," said Newman of Futurum. "It needed to be faster. It needed to be leaner. It needed to be more focused. It needed someone that would be a little bit more shrewd."

Tan is waiting to see how demand shapes up before giving solid details about Intel's next node, 14A. Chandrasekaran told CNBC it will first be developed in Oregon, with a goal of volume production in 2028.

Finding customers for 18A won't be easy. Unlike TSMC which only makes chips for outside clients, Intel also makes devices powered by its chips, positioning it as a competitor to some of the customers it hopes to land.

"If I'm an Nvidia or AMD or Qualcomm, or Broadcom, do you really want to put your secret sauce into a manufacturing operation where you're giving Intel access to that secret sauce?" Yoffie said.

He suggests breaking out foundry into a different company.

"If you actually separated the two, I think you'd give Intel a much better shot at being successful," Yoffie said. "And you'd also give the United States a much stronger position for being the home of a major semiconductor manufacturing organization."

Intel client computing head Jim Johnson gives an early look at its Panther Lake CPU in Santa Clara, Ca. Marc Ganley

For now, Intel hopes Panther Lake will be a big proof point when it debuts in PCs from major companies like Samsung, Dell, HP, Lenovo, Asus and Acer in January. Intel's next data center chip, Xeon 6+, is also made on 18A.

"If you're a major company that wants to bet on a process node, you're going to feel a lot more comfortable if you see Intel ramping the heart of their client product line to high volume on that process node," Johnson said.

Microsoft and Amazon signed early deals last year committing to use Intel's foundry for some of their in-house custom chips.

"It's a good sign, but of course their volumes are very small relative to Nvidia and the other major chip companies," Yoffie said.

Recent reports suggest AMD is considering manufacturing at Intel, and one analyst predicts Apple may once again make some Mac chips at Intel by 2027.

In the meantime, Intel got a lifeline with the U.S. government's 10% stake.

"It shows the confidence that the U.S. government has in Intel and the belief that we need to have leading edge R&D and manufacturing on U.S. soil," Chandrasekaran said.

The government investment came days after President Donald Trump called for Tan to resign, then reversed course.

"I worry sometimes about the scope creep here and how the U.S. could decide to take stakes in all kinds of things," Newman said. "But you have industries that we have let leave the U.S. to an extent that put us into indefensible risk, and we need to bring them back."

Some 92% of the world's most advanced chips are made in Taiwan, following a decades-long decline in the percentage of chips made in the U.S.

"The stakes are incredibly high for Intel, for the U.S. and for the world," Yoffie said. "The whole idea that the world's most advanced products are dependent on a single location in an island a few miles off the Chinese coast is a terrible situation for the whole world to have to deal with."

Chandrasekaran, for his part, is committed to turning Intel into a manufacturer of advanced chips.

"As a semiconductor community, we have to enable this solution for the world to move forward with AI," he said. "There's no other option than to be successful."

Broadcom: Q4 Confirms The Company's Position At The Core Of AI Infrastructure

Dec 17, 2025 Research Wise

Summary

- Broadcom Inc. reported robust double-digit top-line growth driven by hyperscaler demand for custom accelerators and switches, with semiconductor segment AI revenue up 74%.
- Robust momentum in the AI semiconductor business and elevated Backlog expected to support sustained double-digit growth through FY26.
- With fundamentals intact, strong demand visibility, and valuation reset after the recent pullback, AVGO offers an attractive long-term opportunity, supporting a Buy rating.

JHVEPhoto/iStock Editorial via Getty Images

Introduction And The Thesis

Broadcom Inc. (AVGO) is a leading global semiconductor and infrastructure software company. In the past few years, the company has pivoted from a diversified semiconductor company into a leading supplier for AI racks to hyperscalers, supplying them with custom accelerators (XPUs) and high-bandwidth switches. This transition has materially boosted the company's top-line growth over the past years.

Continuing the strong momentum further, AVGO reported another solid quarter exiting FY25, with a double beat in its Q4 earnings, primarily driven by AI-related growth in the semiconductor segment. While the demand environment for the company's offering across the segment remains robust with strong hyperscaler spending, the record order backlog provides solid visibility for consolidated top-line growth through FY26, despite expected softness in the non-AI semiconductor business.

For margins, while there might be near-term pressure from an unfavorable mix of lower-margin AI-related revenue, I think operating leverage from strong revenue growth should continue to support overall margins in FY26 and beyond. The AVGO stock has corrected approximately 18% by the time of writing this article due to concerns around margin softness in the near term. The pullback has reset the valuation to a more reasonable multiple. While the growth prospects appear strong with a combination of scale, visibility, and AI-driven growth, I believe that AVGO stock is a compelling buy for the long term.

AVGO's Q4 2025 Highlights

Although AVGO continued to deliver steady double-digit growth throughout 2025, it was slightly modest versus the prior year due to tougher comparisons from the integration of VMware in 2024, which significantly boosted the company's revenue for the year. Last week, the company reported its Q4 results for 2025 with a solid 28.2% year-over-year growth across its consolidated top-line, reaching a record $18 billion, as both of the company's semiconductor solutions and infrastructure software segments continued to show strength.

However, this growth was primarily driven by the semiconductor segment, which grew a solid 35% during the quarter with robust momentum, particularly on the AI semiconductor side. Meanwhile, the Infrastructure software segment was up about 19%, as the overall demand environment remains healthy with continued adoption of VMware Cloud Foundation (VCF) during the quarter. For the full year 2024, the company's consolidated top-line grew approximately 23.9% to $63.9 billion.

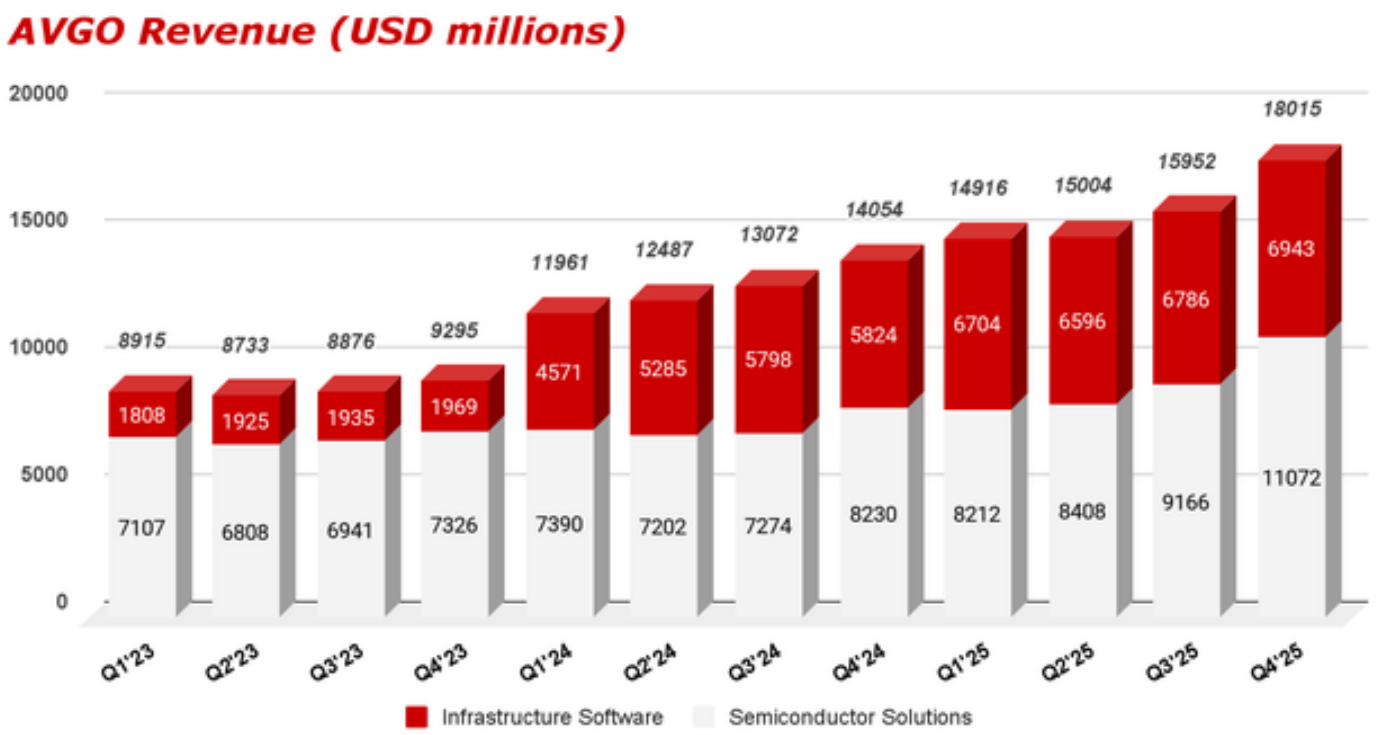

AVGP's quarterly revenue (Research Wise)

The company's profitability also remained strong during the quarter, as its adjusted EBITDA margin expanded 260 bps year-on-year to 67.3% in the fourth quarter of 2025. Although R&D expenses increased notably during the quarter, the margins primarily benefited from operational leverage of strong sales. Increased EBITDA also benefited the company's bottom-line performance, as its adjusted EPS jumped to $1.95 from just $1.42 a year ago, surpassing the consensus estimate by $0.08 during Q4 2025.

Latest Quarter's Earnings

Announce Date	12/11/2025
EPS Normalized Actual	$1.95 (Beat by $0.08)
EPS GAAP Actual	$1.74 (Beat by $0.62)
Revenue Actual	$18.02B
Revenue Surprise	Beat by $555.89M

AVGO's Q4 performance (SA)

Comment On Outlook

After a robust FY24 with the addition of VMware, AVGO closed FY25 on a strong note with solid double-digit growth representing clear momentum across the consolidated top-line with an accelerating AI demand cycle. Going ahead, I expect Broadcom to sustain the momentum entering into FY26, as hyperscaler spending remains strong and adoption of custom accelerators continues to expand.

In the past quarter, the company's semiconductor solutions segment saw a sharp increase of nearly 74% in its AI revenue, primarily driven by strong adoption of AVGO's custom accelerators (XPUs). In the second half of FY25, the company secured approximately $21 billion worth of XPUs orders, with $11 billion in Q4 alone underscoring the deepening scale of deployments.

Additionally, the company also added a new hyperscaler XPU customer with an order of $1 billion, which further expanded the company's footprint in the custom silicon space. In my view, while most hyperscalers are now increasingly favoring tailored silicon solutions to optimize their performance and power efficiency, and spending also continues to scale significantly, AVGO should continue to see strong order momentum as it progresses through 2026.

On the other hand, demand on the AI networking side also remains quite strong, as the company reported another quarter of record bookings for its 102 terabit per second Tomahawk 6 switches, with order backlog for AI switches now exceeding $10 billion. Customers are continuing to build data center infrastructure ahead of deploying AI accelerators. AI system architectures are also becoming more and more network-intensive. With customers scaling new-generation AI clusters that require ultra-high bandwidth and low-latency networking, I think demand for AVGO's switching and connectivity solutions should stay strong and remain a key revenue driver.

However, the company's non-AI Semiconductor business, which accounts for roughly 40% of the semiconductor segment revenue (and ~25% of consolidated revenue), continues to face headwinds from muted enterprise spending. In my view, with almost no signs of recovery in customer spending in this space, I expect the non-AI revenue growth to remain almost flat to low single digit in the short term, which could weigh on the segment growth through FY26. But, with robust momentum on the AI side and a solid $73 billion in AI backlog, the company's key semiconductor solutions segment should sustain decent double-digit growth in the quarters ahead.

Turning to the company's Infrastructure Software segment, demand remains healthy and highly visible, with a solid $10.4 billion in new contracts in Q4, pushing the total infrastructure software backlog to approximately $73 billion. While the adoption of VMware Cloud Foundation continued, I believe support from the elevated pipeline should continue to benefit the segment's top-line growth in FY26 and beyond. Overall, despite flat growth expectations in the company's non-AI business, a robust demand environment in the AI semiconductors, along with a record order backlog and steady growth on the industrial software side, should help the company sustain its double-digit growth through FY26.

AVGO's outlook beyond FY26 also appears compelling with its dominant positioning across various levels of the AI infrastructure stack, which should continue to create durable demand for the company's custom accelerators, switching and networking solutions, and other related offerings. Apart from this, the company's total backlog currently stands at a whopping $162 billion, which keeps growing quarter on quarter with strong order momentum from existing customers and new customer additions, further providing multi-year demand visibility.

To sum up, while AVGO's growth over the coming quarters is set to be driven by AI semiconductors, its long-term story appears equally compelling due to the company's scale, diversified exposure, and deep customer ties, which position it well to sustain its strong top-line growth over the coming years.

Valuation Is Reasonable

Over the past year, AVGO stock has been up roughly 36%, outperforming the S&P 500 as it continues to deliver steady growth across its top and bottom line. Following a strong double-digit surge after the Q3 earnings earlier in September, AVGO shares extended the rally and touched a new high around $414 just ahead of Q4 FY25 results.

AVGO stock performance in last one year (SA)

Although Broadcom once again reported a strong double beat in Q4, the company's stock pulled back strongly after the earnings release. This reaction was mainly driven by concerns around potential margin weakness in FY26.

As a result, the stock has stayed nearly flat since the levels post Q3 earnings. However, the Q4 dip, along with upward EPS revisions following a strong upbeat earnings result, has significantly cooled down the valuation multiple. Currently, the company's stock is trading at a forward non-GAAP P/E ratio of 33.77, based on FY26 (ending Oct'26) EPS estimate of $10.06.

When compared to its close peers like NVIDIA Corporation (NVDA), Advanced Micro Devices, Inc. (AMD), and Marvell Technology, Inc. (MRVL) across the GPU and networking space, having a forward P/E multiple of around 37.4x, 52.3x, and 29.7x, respectively, the AVGO stock appears fairly valued. And considering its strong double-digit forward growth expectations, notably higher than the historical average, as can be seen below, the stock valuation looks really attractive.

AVGO Growth Grade A-

	Sector Relative Grade	AVGO	Sector Median	% Diff. to Sector	AVGO 5Y Avg.	% Diff. to 5Y Avg.
Revenue Growth (YoY)	A-	23.87%	8.90%	168.34%	19.96%	19.59%
Revenue Growth (FWD)	A	35.30%	8.03%	339.66%	16.39%	115.36%
EBITDA Growth (YoY)	B+	33.34%	11.41%	192.22%	24.83%	34.31%
EBITDA Growth (FWD)	A-	39.26%	12.31%	218.99%	18.33%	114.26%
EBIT Growth (YoY)	A-	62.44%	12.36%	405.37%	43.20%	44.53%
EBIT Growth (FWD)	A	40.58%	12.76%	218.04%	18.97%	113.89%
EPS Diluted Growth (YoY)	A+	269.77%	14.02%	1,824.74%	53.06%	408.38%
EPS Diluted Growth (FWD)	A	41.51%	12.87%	222.55%	18.25%	127.47%

AVGO Growth table (SA)

I expect the company's top-line to continue to expand in double digits through FY26 as demand remains robust, particularly on the AI semiconductor side, and a robust order backlog continues to provide strong demand visibility.

The overall profitability, on the other hand, might face a slight headwind with a higher mix of AI revenue, which carries comparatively lower margins. As the CFO noted during the call:

We expect Q1 consolidated gross margin to be down approximately 100 basis points sequentially, primarily reflecting a higher mix of AI revenue. As a reminder, consolidated gross margins through the year will be impacted by the revenue mix of Infrastructure Software and Semiconductors, and also product mix within Semiconductors.

However, I don't expect this anticipated impact of the unfavorable mix to impact the company's EBITDA margins meaningfully and expect them to remain close to current levels as the company benefits from operating leverage on anticipated higher revenue growth through FY26. With margin strength, the company's bottom line should also continue to expand steadily and enhance the stock's valuation further in the future.

Risks

Although Broadcom's overall outlook appears good, its non-AI semiconductor business remains under pressure due to weak enterprise spending. If this softness persists longer than expected, it could weigh on overall top-line growth, given its meaningful share of about 25% in total revenue.

Additionally, AVGO's AI growth is mostly driven by a concentrated set of hyperscale customers. If there's any slowdown in customer spending or shifts in custom silicon strategies, the company is likely to face near-term revenue volatility.

Lastly, while the recent drop has cooled down the stock valuation, it continues to trade at a decent premium. Therefore, any weakness in performance in the upcoming earnings could negatively impact the company's stock performance in the short term, despite strong long-term fundamentals.

STMicroelectronics (STM) - Barclays 23rd Annual Global Technology Conference Summary

Date: December 11, 2025. AI generated summary SL

Primary Speaker: Lorenzo Grandi – President & CFO

Moderator: Simon Coles – Barclays

Executive Overview

Lorenzo Grandi (CFO) stated that 2025 marked the end of major inventory corrections in automotive and industrial markets. Sequential growth resumed from Q2 onward, and visibility into 2026 is meaningfully better than last year, with normalized seasonality expected starting Q1 2026.

STM expects revenue recovery, margin improvement, and structurally healthier demand as inventory normalizes, manufacturing capacity is reshaped, and new growth drivers (Silicon Photonics, AI power, MEMS, satellites) begin contributing.

📈 Market & Demand Trends

Automotive & Industrial Recovery

Speaker: Lorenzo Grandi (CFO)

Automotive: Inventory correction largely completed in H1 2025; sequential growth resumed from Q2 and expected to continue at mid-single-digit Q/Q.

Industrial: Distribution inventory normalized; bookings improving with book-to-bill at or above parity.

2026 visibility: Backlog higher than 2024 entry point, though not yet at pre-COVID levels.

💰 Gross Margin Outlook

Speaker: Lorenzo Grandi (CFO)

2025 gross margin ~33.8%, heavily impacted by >$400M under-utilization charges (~300+ bps headwind).

2026 margin expected to improve, driven by:

Lower under-utilization as revenue recovers

Manufacturing footprint reshaping (SiC 150mm reduction, tuned 200mm/300mm mix)

Improved product mix (Industrial, MCU, Analog recovery)

Headwinds in 2026:

Capacity reservation fees declining (loss of ~70–80 bps benefit)

FX pressure (EUR/USD moving from ~1.11 → ~1.15–1.16)

Bottom line: Margin recovery expected, stronger in 2H 2026 than 1H.

📉 Operating Expenses (OpEx)

Speaker: Lorenzo Grandi (CFO)

Cost-reduction plan on track: $300–360M savings vs 2024 baseline

~$110M achieved in 2025; similar or slightly higher in 2026

OpEx to rise slightly due to:

Lower grant contribution

Start-up costs for 200mm SiC facilities

FX headwinds (Euro-denominated expenses)

🚀 Strategic Growth Drivers

Silicon Photonics

Speaker: Lorenzo Grandi (CFO)

Strong customer traction; qualification underway

Meaningful revenue starting 2026, accelerating in 2027

Long-term potential: ~$0.5B revenue within 3–4 years

Already profit accretive

AI Power & Data Center Infrastructure

Speaker: Lorenzo Grandi (CFO)

Portfolio spans SiC, GaN, power MOSFETs, MCUs, analog

Cooperation with NVIDIA targeting 800V AI server architectures

Revenue contribution expected from 2027 onward

Long-term opportunity similar to silicon photonics: ~$0.5B+ by late decade

Satellite Communications

Speaker: Lorenzo Grandi (CFO)

Long-standing leadership (incl. Starlink)

Expanding customer base (US, Europe, China constellations)

Growth expected in 2026–2027

Strong differentiation via front-end + back-end integration

Silicon Carbide (SiC)

Speaker: Lorenzo Grandi (CFO)

2025 = transition year, revenues down due to customer destocking and China pricing pressure

2026: return to growth via new sockets in Europe & China

2027: potential return close to 2024 revenue levels

Long-term demand intact (EV, industrial, AI servers)

Humanoid Robotics

Speaker: Lorenzo Grandi (CFO)

STM well positioned (MCUs, sensors, analog, motion control)

Content per robot: $200–300, potentially higher

Market still early; already shipping to early customers

Applications beyond consumer: industrial & defense

🤝 M&A: NXP MEMS Acquisition

Speaker: Lorenzo Grandi (CFO)

~$1B all-cash acquisition

Strategically rebalances MEMS exposure toward automotive & industrial

Positions STM as top-2 global MEMS supplier

MEMS = one of STM's most profitable product lines

Balance sheet remains strong post-acquisition

📌 Key Takeaways

Inventory correction largely complete; 2026 visibility improving

Gross margins recovering as utilization improves and mix normalizes

Silicon Photonics & AI power emerging as multi-year growth engines

Satellite & MEMS provide profitable, differentiated expansion

STM positioning for profitable growth beyond the cyclical recovery

10 Risks For The Global Economy In 2026

Dec 04, 2025 ING Economic and Financial Analysis

Summary

- From AI bubbles to government spending splurges, property crashes to oil price spikes, these are the 10 ways our economic outlook for 2026 could go wrong – or right.
- US tech companies fail to monetise AI, questioning the logic of immense investment in hardware/software and related industries.
- At 15%, the eurozone savings rate has consistently been three percentage points above its pre-Covid average.
- The key upside risk to oil prices remains Russian supply, due to both US sanctions and persistent Ukrainian attacks on Russian energy infrastructure.


Ismagilov/iStock via Getty Images

The AI bubble bursts

US tech companies fail to monetise AI, questioning the logic of immense investment in hardware/software and related industries. Tech stocks crash, hitting the top 20% of American earners who own the lion's share of US equities held domestically.

Having enabled consumer spending to grow over the past couple of years, even as the bottom 60% have struggled, lower household wealth causes a fall in consumption in 2026.

AI investment falls abruptly, weighing on the construction and investment that has likely contributed around one percentage point to US growth in 2025 (though less once imported equipment is netted off). This is enough to push the US jobs market into a full-blown recession.

Impact: The US falls into recession, while Europe is less affected. The Fed cuts rates more aggressively.

Congress approves 'tariff rebates' ahead of mid-terms

Fiscal policy is a major upside risk to growth and inflation in 2026. President Trump is pushing Congress to hand out $2,000 'tariff rebate' checks to 150 million Americans, resurfacing memories of the Covid-era stimulus that helped turbocharge inflation. Though the maths doesn't totally add up - and tariffs already helped justify the One Big Beautiful Bill - pressure could build into November's mid-term elections.

This would help the bottom 60% of American consumers struggling under the weight of the cost of living, though much of it may simply be used to pay off debt, and the impact on growth may be more muted than in 2020/21.

Impact: US growth is higher and inflation rises. The Fed becomes more hawkish, depending on political influence over the committee.

Inflation resurgence on AI supply bottlenecks

Many economists - not least the Fed's doves - expect AI to be a massive positive for productivity, which pushes down inflation. But what if that's wrong? In the short term, massive investment in AI infrastructure could crowd out other forms of economic activity. Data centres are expected to account for 10% of US power demand by 2030. Electricity grids globally will be under increasing strain, risking blackouts and higher prices. Rising investment needs also risk fresh supply shortages, at a time of tighter immigration rules in the US and Europe. Wage growth risks turning higher again.

Impact: Global inflation rises. Central bank rate hikes draw nearer.

President Trump slashes tariffs as negative impact grows

There are two ways the US average tariff - currently around 16% - could fall. First, the US administration opts to lower tariffs ahead of the elections, just as it has done with certain food products recently. The resulting fall in revenue would complicate efforts to convince Congress to approve 'tariff rebate' checks, but it's possible that once this is done, the president will begin to roll back trade barriers in a bid to lower consumer bills.

Alternatively, the Supreme Court rules that tariffs imposed under emergency powers - most country-level levies - are illegal. The president uses other means, such as Section 122, which allows 15% tariffs for 150 days, to rebuild trade barriers - but the result is messier. He could also widen the scope of sector-specific tariffs, though this would take time. The result may well be a lower average tariff level.

Impact: Growth rises, inflation eases, but the former is judged as the dominant factor by the Fed. US rate cuts are curtailed.

European consumers start splashing the cash

At 15%, the eurozone savings rate has consistently been three percentage points above its pre-Covid average. Savings intentions remain ultra-high. But having had time to rebuild savings after the 2022 energy crisis - and after a period of stable 2% inflation - it's possible consumers will begin to splash the cash more willingly in 2026. At least if governments can take away policy uncertainty regarding pensions.

Impact: The eurozone grows above trend (upwards of 1.5% annual growth). The ECB hikes rates in late 2026.

US-China relations sour, hitting supply of rare earths

The US-China friction has been alleviated after a face-to-face meeting between Presidents Trump and Xi led to a 12-month truce, which theoretically would leave tariffs and export controls unchanged for much of 2026. However, the truce remains fragile, and any miscalculations along the way could derail the deal. If cooler minds do not prevail, non-tariff barriers such as rare earth controls could be enacted.

Impact: Direct impact on semiconductor, auto, and defence sectors, potentially resulting in shortages and price surges on affected products, contributing to higher inflation.

Oil prices spike on renewed geopolitical tensions

The key upside risk to oil prices remains Russian supply, due to both US sanctions and persistent Ukrainian attacks on Russian energy infrastructure. The widely held view is that Russian oil will find ways to circumvent sanctions. However, if sanctions prove to be more effective than thought, this potentially reduces the scale of the oil surplus expected in 2026, and is an upside risk to our view that Brent will average $57/bbl next year.

The recent escalation between the US and Venezuela also leaves uncertainty over Venezuelan supply, while the fragility of the Israel/Gaza ceasefire means that supply risks from the Middle East could re-emerge.

Impact: Weaker global growth and higher inflation. Central banks are more likely to hike rates/curtail rate cuts to lean against inflation risk.

Budget crises loom as bond investors lose confidence

Investors have been surprisingly immune to concerns about the trajectory of the US fiscal deficit this year, helped perhaps by lingering concerns about the US macro outlook and resumption of Fed rate cuts. But America's public finances are precarious; the deficit is expected to remain at 6-7% for some time. There's a risk investors begin to baulk at the volume of debt issuance, potentially in the aftermath of our second scenario - fiscal profligacy combined with a perception of too loose monetary policy and hints of inflation.

Europe remains vulnerable too; the situation in France could become more widespread as spending pressures mount - not least from defence. Bond yields would spike, and the economic impact would depend greatly on how central banks react. Do we see a return to QE - or in Europe's case, the first use of its 'Transmission Protection Instrument'? If not, governments may be forced to embrace austerity.

Impact: Governments - particularly in Europe - forced into painful spending cuts in a bid to halt the bond sell-off. Growth slumps.

China enters downturn on deeper property price correction

After stabilising at the start of 2025, the downturn in property prices began to steepen once more, starting in mid-2025. Prices have fallen, inventories remain high, and property investment remains a major drag on growth. Default concerns have re-emerged after property developer Vanke asked for a 1-year extension on a bond payment.

After rolling out myriad policies to help stabilise the markets in 2024, the momentum has slowed in 2025, with more voices advocating for just letting the cycle fully play out naturally over the coming years instead of trying to stabilise the market - a move which could have serious implications. The spillover effects could be significant if the downward momentum is not contained.

Impact: Household wealth destruction, deteriorating bank asset quality, and entrenched pessimism would all be potential consequences if the property downturn continues. It would hamper efforts to transition toward domestic demand-driven growth, as well as reduce the growth outlook in the near term.

Ukraine war ends with full and enduring peace agreement

If peace negotiations are successful, the wider economic impact will likely depend on the extent to which trickier topics - such as territorial recognition - are addressed, and how enduring any ceasefire is perceived to be. In a more optimistic scenario - where a credible, long-term agreement is reached and investors feel confident about redeploying money in Ukraine - reconstruction efforts would likely have wider ripple effects on activity and, more importantly, sentiment in Eastern Europe.

Lower energy prices, depending on the extent of sanction removal, could also have a stimulative effect on global consumers. However, our energy team notes that the Russian oil supply hasn't materially fallen in recent years, so the impact on the global supply balance may not be significant. Although, admittedly, it would reduce a large amount of supply risk hanging over the oil market. The impact on the gas market would be more significant, but this would require Europe to start resuming its purchases of Russian natural gas.

Impact: Lower energy prices boost global growth. Some central banks (e.g., the Bank of England) may counterintuitively react dovishly, having recently reacted hawkishly to price spikes on fears of supply-driven inflation.

Micron: Why An Exit From Consumer Isn't Scary

Dec. 04, 2025 Tech Stock Pros

Summary

- Micron Technology announced that it'll exit its consumer business by February 2026 to prioritize AI-driven memory demand and capitalize on industry-wide shortages.
- The market was confused about how to understand the news; our take is that this won't negatively impact MU, considering the consumer is not what's driving the company's outperformance.
- Instead, that would be DRAM, and HBM in specific, where the company holds the second-largest market cap after SK Hynix at 21%.
- We expect MU to continue to benefit from the price uplift amid the shortage, and the consumer business exit speaks to a larger memory hyperscale that has.

JHVEPhoto/iStock Editorial via Getty Images

What Micron's Exit From the Consumer Business Means

Micron Technology, Inc. (MU) issued a press release on Wednesday stating its intention to exit its consumer business and refocus capacity towards AI-led hyper demand for memory, as the AI infrastructure build-out shows no signs of slowing, and memory requirements continue to balloon. AI is memory-bound, and so it's eating up quite literally all the capacity within the supply chain.

The announcement reads as follows:

December 3, 2025 at 11:40 AM EST

Micron Announces Exit from Crucial Consumer Business

BOISE, Idaho, Dec. 03, 2025 (GLOBE NEWSWIRE) -- Micron Technology, Inc. (Nasdaq: MU), a leader in innovative memory and storage solutions, today announced its decision to exit the Crucial consumer business, including the sale of Crucial consumer-branded products at key retailers, e-tailers and distributors worldwide.

Micron will continue Crucial consumer product shipments through the consumer channel until the end of fiscal Q2 (February 2026). The company will work closely with partners and customers through this transition and will provide continued warranty service and support for Crucial products. Micron will continue to support the sale of Micron-branded enterprise products to commercial channel customers globally.

"The AI-driven growth in the data center has led to a surge in demand for memory and storage. Micron has made the difficult decision to exit the Crucial consumer business in order to improve supply and support for our larger, strategic customers in faster-growing segments," said Sumit Sadana, EVP and Chief Business Officer at Micron Technology. "Thanks to a passionate community of consumers, the Crucial brand has become synonymous with technical leadership, quality and reliability of leading-edge memory and storage products. We would like to thank our millions of customers, hundreds of partners and all of the Micron team members who have supported the Crucial journey for the last 29 years."

This decision reflects Micron's commitment to its ongoing portfolio transformation and the resulting alignment of its business to secular, profitable growth vectors in memory and storage. By concentrating on core enterprise and commercial segments, Micron aims to improve long-term business performance and create value for strategic customers as well as stakeholders.

Micron intends to reduce impact on team members due to this business decision through redeployment opportunities into existing open positions within the company.

Micron

Basically, MU will exit its consumer business starting fiscal Q2, or February 2026. The stock was down initially on the news by about 2.6%. Now, in pre-market, it's back to being flat, and that is because this move points to how unprecedented the memory shortage is, more than anything else, or any financial problems for MU.

The memory shortage is an industry-wide phenomenon, and everyone is feeling the heat. Things are so bad that a major smartphone manufacturer had 30-40% of its orders cut;

Now, it's not about how much demand you have as a mobile phone manufacturer, but about how much inventory Samsung and SK Hynix [or Micron] can allocate to you."

Xiaomi warned investors in late November that while the company secured 2026 memory pricing, shortages "could push smartphone prices higher." Even Samsung is looking at potentially rising prices for its upcoming Galaxy S26, which is particularly notable as Samsung has kept prices relatively flat for a few years.

For MU, this is a strategic move to prioritize advanced memory chips for AI data center use, which is smart considering the higher ASP and overflowing demand on that front. It also speaks to the fact that this memory price surge is likely to stick around longer than we're used to from typical memory cyclicality. This cycle, demand continues to outpace supply, and there are limited ways to boost industry supply. So, MU's exit from the consumer business is an attempt to reallocate supply and consolidate more capacity towards DRAM, and specifically HBM, which AI has driven into shortage. The shortage is also reflected in spot prices, which, for those who aren't familiar, are the real-time current market rates for memory chips.

The chart below showcases the DRAM average spot price as of November 19th. DRAM spot prices have been going crazy; DDR4 1Gx8 chips' price is up 158% since September, while DDR5 2Gx8 chips' spot price increased 307%. So, shortage is seriously pushing up pricing, and while this price surge won't continue forever, we think MU is smart in trying to maximize its profit out of it, particularly being an HBM industry go-to.

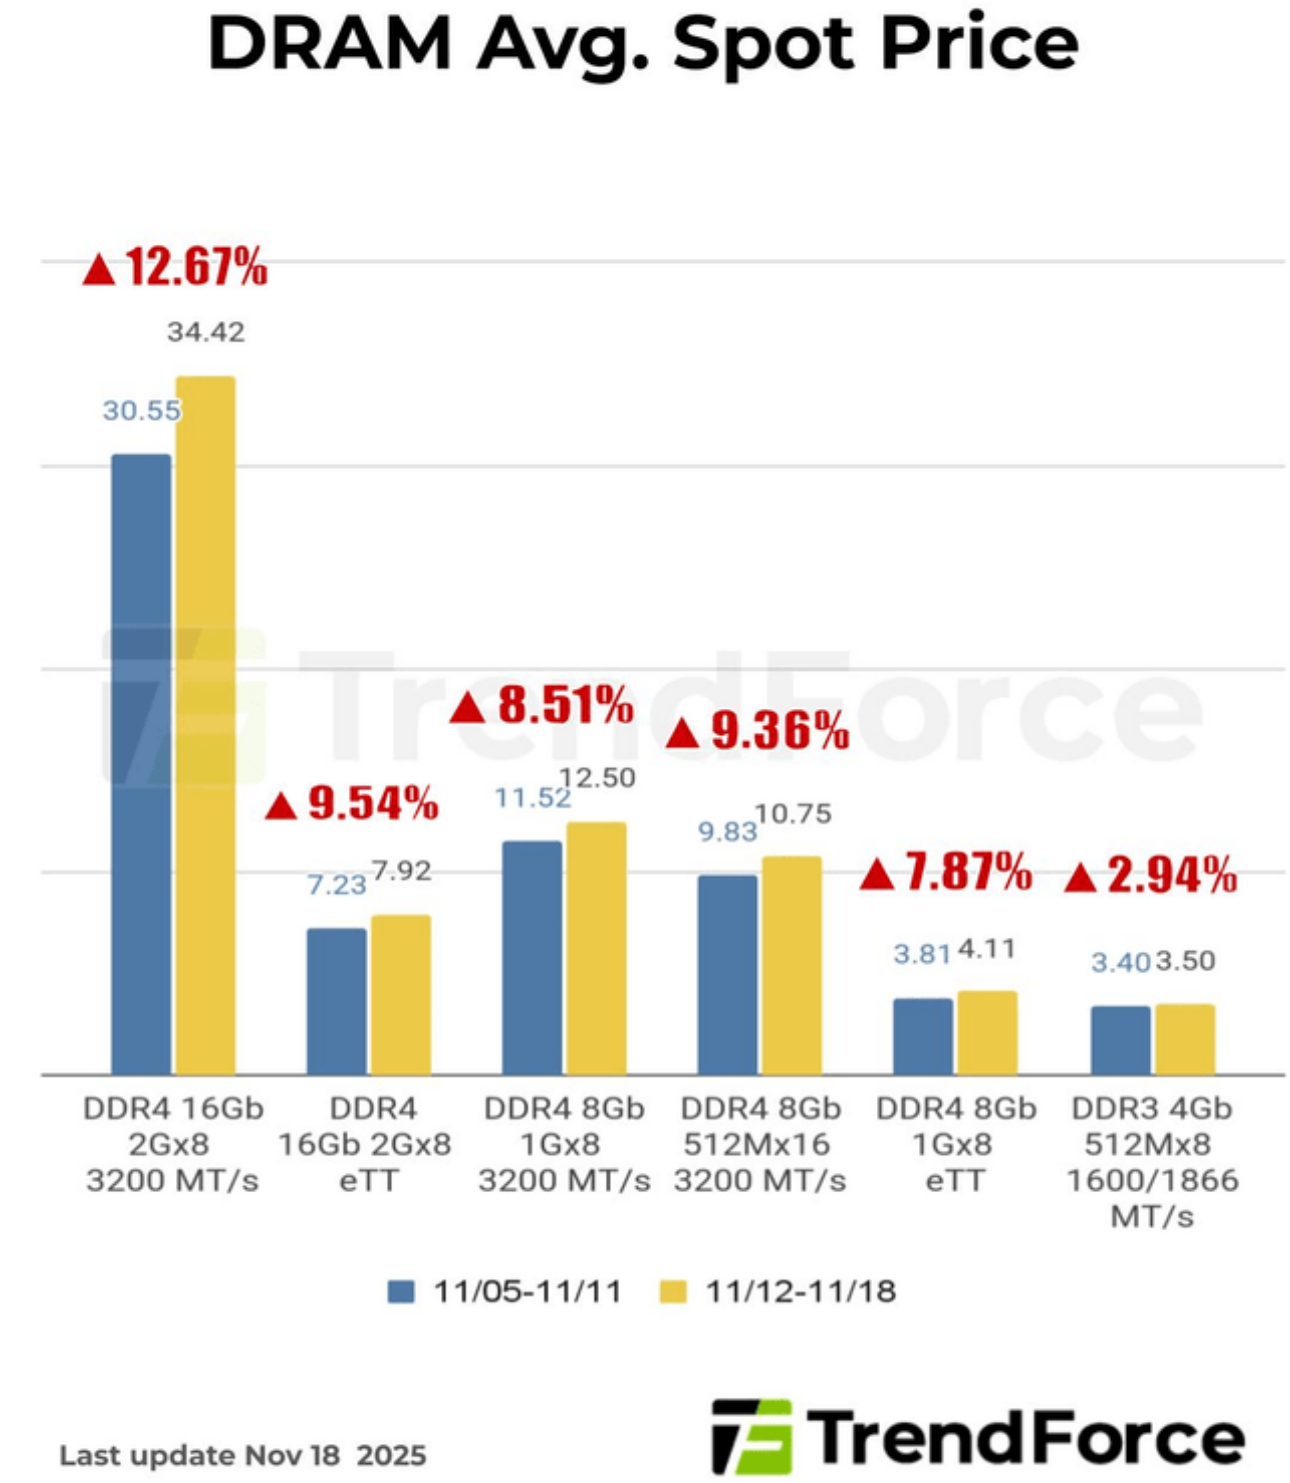

TrendForece

For the first time ever, MU has overtaken Samsung (OTCPK:SSNLF) in market share when it comes to HBM. MU captured 21% of HBM shipments during Q2 CY25, second to SK hynix (OTCPK:HXSC.F), which holds a comfortable lead with a 62% share, ahead of Samsung's 17%. We think MU is focusing on AI-led memory demand to grow its share as the entire HBM market opportunity grows. In fact, the HBM market is poised for a 26.10% CAGR by 2034.

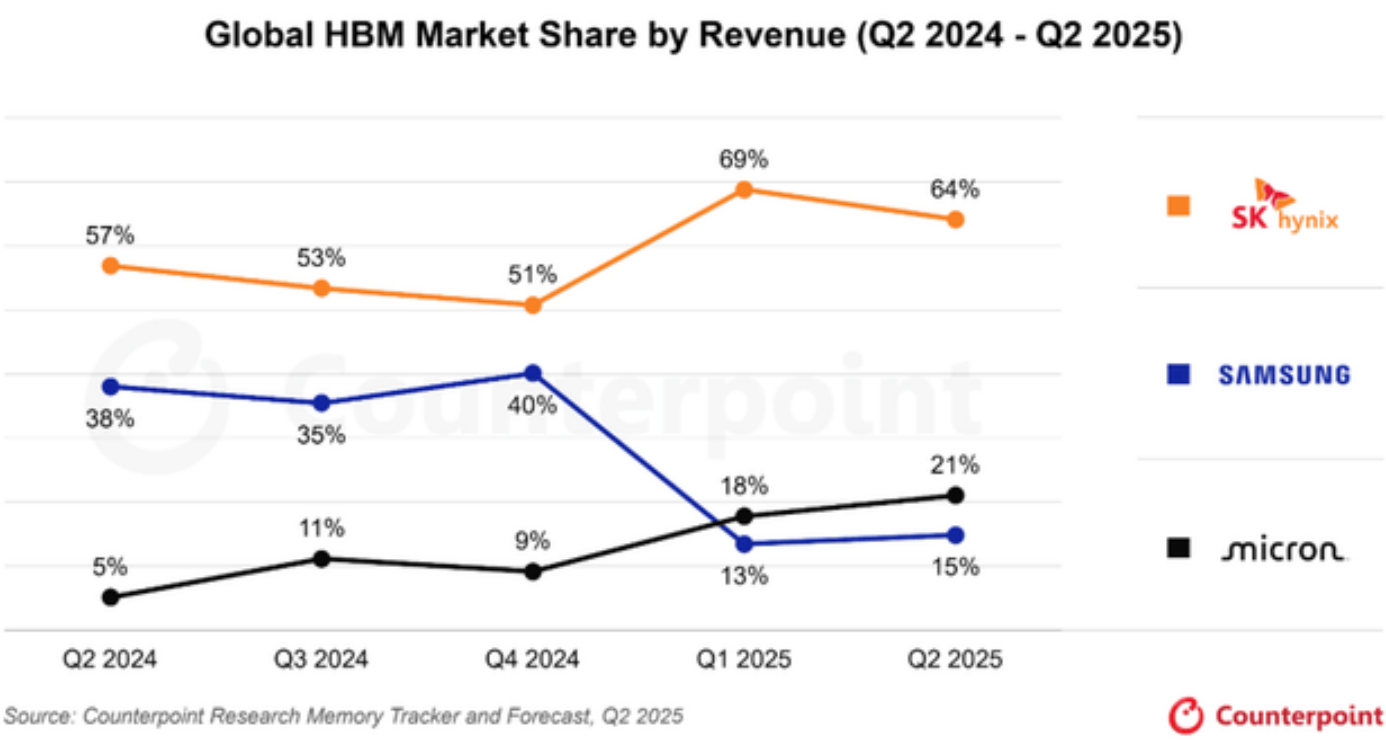

Counterpoint

So, looking at MU's financials, the company derives the bulk of its revenue from DRAM, up from 69% of total sales in a year-ago quarter to 76% last quarter (Q3 FY25) and 79% in 4QFY25. DRAM will likely make up an even greater chunk of total sales next year, with our estimate being for north of 80%, given that HBM3E supply and pricing agreements for 2026 are already in place. As of Q4, HBM revenue grew to $2B, indicating an annualized run rate of $8B, and we're expecting this number to grow substantially as MU benefits from the new product cycle coming up.

HBM4 pricing is expected to be at least 50% higher than HBM3E, within the mid-$500s range per unit, versus HBM3E's pricing range within the mid-$300s. SK hynix is rumored to be pricing HBM4 up to 70% higher. HBM4 samples are already at customers, and the HBM4E base logic die is being manufactured in partnership with TSMC (TSM). So, we should be looking at a continued price uplift for MU, particularly as AI servers' DRAM consumption is several times higher than standard cloud machines.

MU is also investing in its Japan production capabilities to expand supply in hopes of meeting the growing demand for advanced memory. This, however, will take time. It's not an overnight kind of fix. So, we're seeing MU think ahead and position itself strategically towards AI with Wednesday's news.

What to Do With the Stock?

MU has been our top memory pick; it was one of our highest conviction ideas on the April crash when it hit the high $60s range. The stock is up around 45% since our last buy call in late September, shown below.

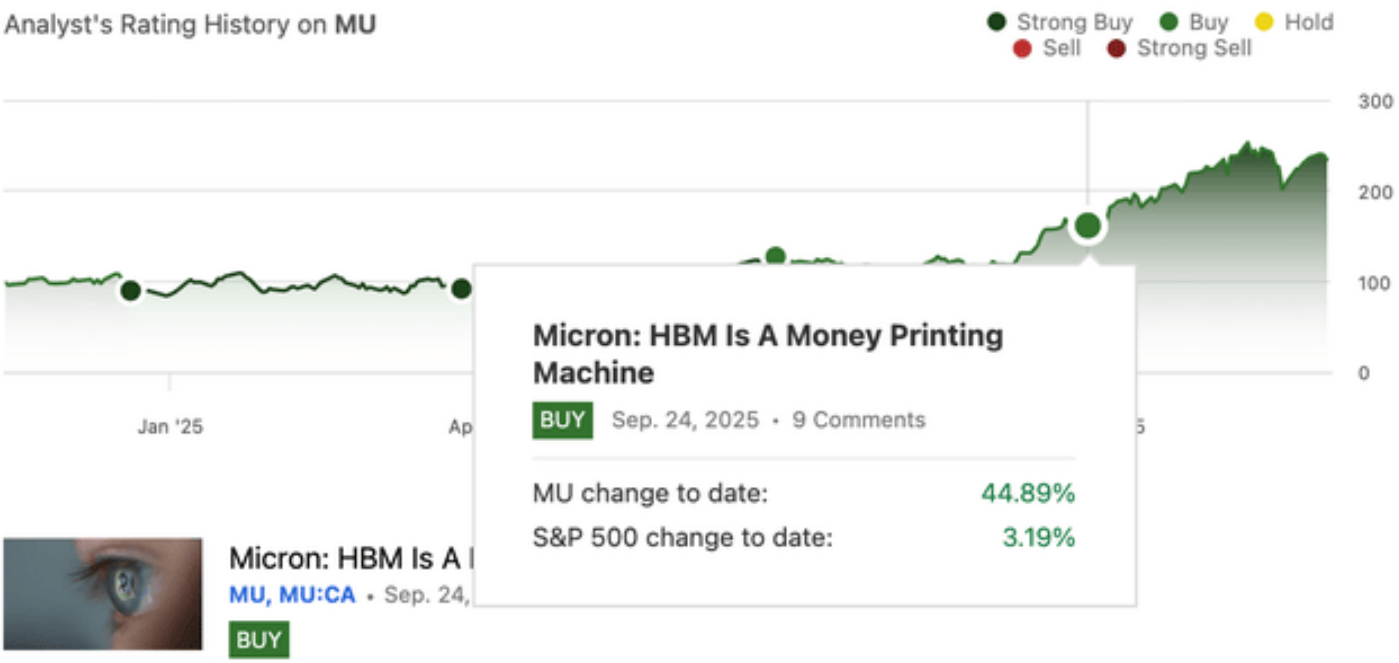

SA

MU stock, along with our memory/storage-focused companies, has been skyrocketing. The YTD chart below showcases performance from MU, Seagate (STX), and Western Digital (WDC) against the S&P 500.

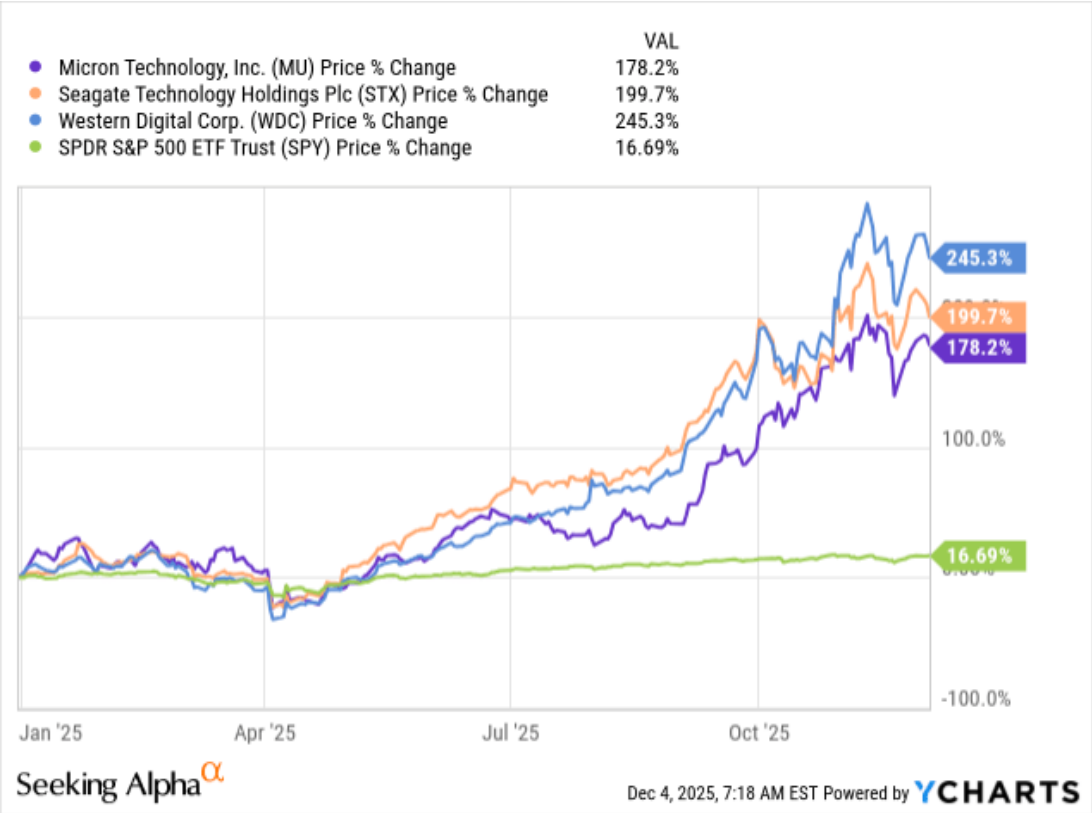

YCharts

The surge in AI demand for memory/storage has also pushed valuations up, big time, too.

Will Molybdenum Make Lam Research The Next ASML?

Dec 02, 2025 Jia Ming Eow

Summary

- Lam Research plays a critical role in the AI supply chain, providing advanced deposition machines for next-generation semiconductor manufacturing.
- LRCX's ALTUS Halo enables high-volume molybdenum deposition, supporting surging demand in 3D NAND and advancing DRAM applications.
- Despite a clarified growth trajectory for 2026, the company faces near-term caution due to significant revenue exposure to China.

AlexLMX/iStock via Getty Image

Lam Research (LRCX) remains an important part of the AI supply chain, supplying deposition machines like their ALTUS Halo, which enables atomic layer deposition (ALD) machines to volume produce molybdenum, moving on from tungsten. The use case of molybdenum is described in my last Lam article. AI manufacturing firms previously raised caution with possibilities of no growth YoY in 2026, which has recently been clarified with a clearer trajectory for 2026. However, companies like Lam still have a huge proportion of their revenue from China, so despite having a unique moat and a differentiated product, caution still remains in the near term. With the surging demand in both 3D NAND and DRAM, molybdenum is increasingly being used for 3D NAND, with advancing application research and development in DRAM. With valuations relatively high now as manufacturing stocks have recently had a run-up, I'm maintaining Lam at a 'Hold'.

Molybdenum, DRAM, and NAND

As the AI boom isn't slowing down, with top hyperscalers continuing to spend hundreds of billions on AI infrastructure, there has been a surge in demand for DRAM and NAND, which are some of the key hardware within the AI ecosystem to enable training and inference of AI, as well as future use cases.

DRAM, or dynamic random access memory, is the fast, temporary memory that a CPU or GPU uses to process data and only stores data while powered. It essentially feeds data to CPUs/GPUs at a very high bandwidth, which is required in AI training since it requires extremely high memory bandwidth for reading data and making conclusions/predictions.

3D NAND is a flash memory that stacks vertically, unlike traditional 2D, which is horizontal, to increase storage density. NAND is used to store data when the power is off, which is the main differentiation from DRAM. 3D NAND is in SSDs to reduce the cost of SSDs since they are more expensive than HDDs. From my Western Digital article, HDD prices are as much as 6.6x that of SSD, as SSD has lower latency and shorter read/write times. Since storage is so important to the AI boom, the demand for SSDs and HDDs has been booming, as the demand for 3D NAND has also been higher than it has ever been before.

Lam's ALD machines essentially enable the volume production of chip-grade molybdenum. Molybdenum has advantages over tungsten, such as electromigration resistance, low resistivity in scaled features, good adhesion, high stability, and reduced process steps. It is finding a use in NAND, DRAM, and logic chips.

There has been a huge surge in demand for DRAM and NAND, with South Korean technology companies like SK Hynix and Samsung raising prices by up to 30% for Q4, as there has been a supply shortage in both memory and storage. In my combined Seagate and Western Digital article navigating the HDD market, Seagate's supply is way smaller than their demand, with orders all the way until the end of 2026, as Seagate has the best technology right now, and Western Digital customers are willing to wait for volume shipping in 2027 (qualification starts at 1H of 2026). The point is that demand for memory and storage has seen a huge surge in demand across the board, and some of the key products are transitioning to using molybdenum rather than tungsten. Lam is essentially positioned to fulfill the manufacturing equipment for molybdenum, and with essentially a shortage in memory and storage products across the board, Lam's demand will likely continue to surge until supply > demand, which doesn't seem likely to happen at least in the near term.

Lam's Similarity to ASML

From the earnings call:

We're also seeing solid demand for our atomic layer deposition or ALD products, including a recent key win at a major NAND manufacturer for a critical high aspect ratio dielectric deposition application. Lam's differentiated conformal fill capability using a higher temperature process was fundamental in securing this win.On the metal side, Lam's Halo moly ALD tool has been selected as the tool of record for 3 consecutive nodes at a leading customer, including for devices with more than 500 layers. This further reinforces our leadership in the 3D NAND wordline application, a step that is fundamental to building the higher performance devices required for ESSDs.

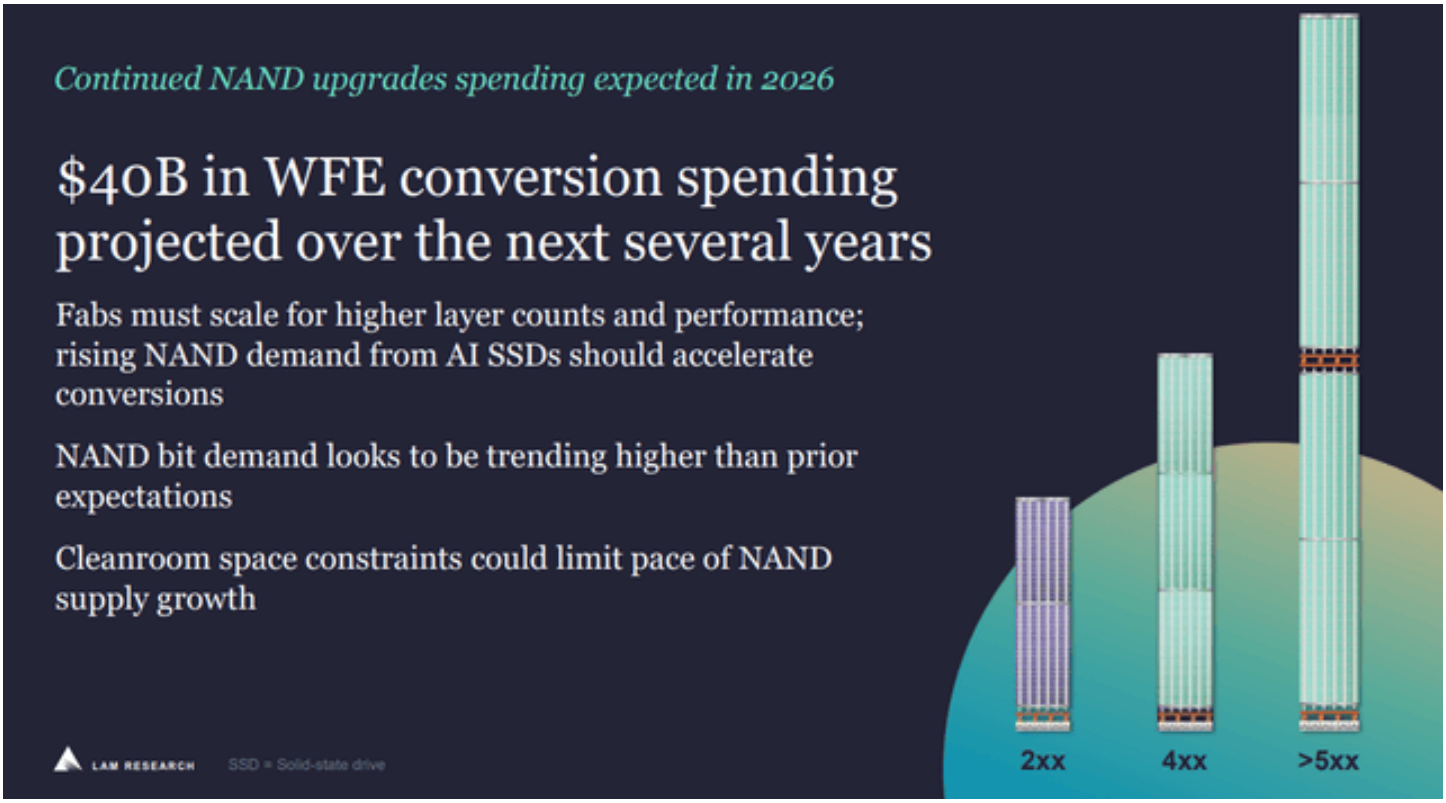

WFE Conversion Spending (LAM Investor Relations)

Lam's moat with their ALD machines has close similarities to ASML's EUV machines, being essentially the only machine in the market that is highly demanded in the market, as Lam is one of the key semiconductor manufacturing equipment suppliers to continue to fuel the AI boom. There are clear differences, as ASML's moat is way stronger, considering there is no other viable solution to EUV (DUV is extremely costly and hard to scale), while semiconductors are still transitioning from the use of tungsten to molybdenum, so molybdenum is not the norm yet. This essentially means that ASML doesn't have competitors, since the switching cost is so high (because there is no other option), while Lam also has a high switching cost, but there is another option. With continued NAND upgrades in 2026, as well as Lam's conversion spending from wafer fabrication equipment (or WFE) continuing to grow fast, I doubt that Lam's demand will fall off a cliff anytime soon.

Lam obviously doesn't only supply machines that make molybdenum precursors that make it into wafer-thin chip-grade molybdenum, as a majority of their revenue still stems from systems and machines in foundries. However, for specifically their ALTUS HALO machine, I see it as a similar but weaker moat compared to ASML Holding's (ASML) EUV machines, essentially being the only machine in the market, as the other options are simply high cost and not as effective as other options (can use tungsten but shifting to molybdenum). This essentially gives Lam an edge over some of its peers like Applied Materials (AMAT), since it is one of the only machines in the market and an important one for the advancement of AI.

Q1 FY26

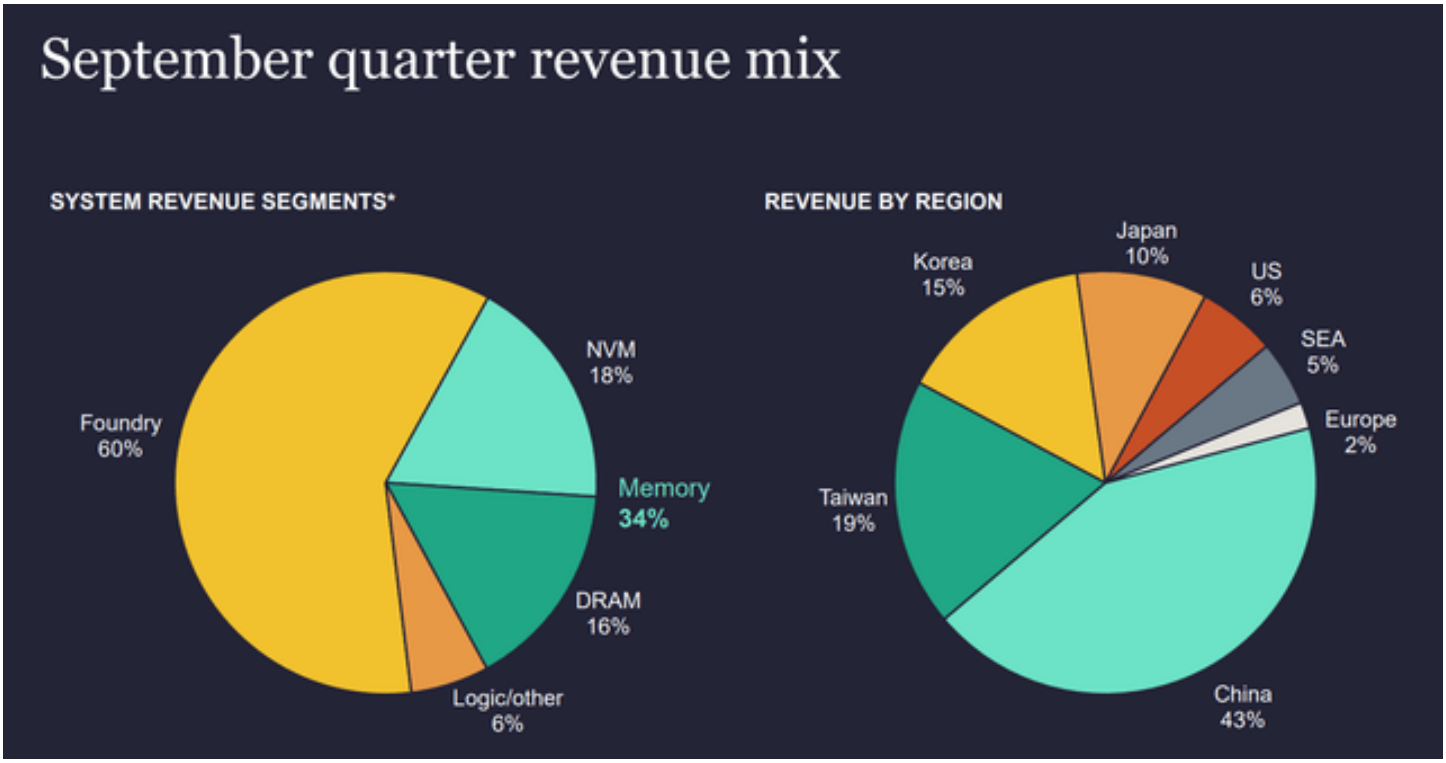

Revenue by System and Geography (LAM Investor Relations)

Lam continues to beat both top and bottom lines in Q1 FY26, but the main focus and takeaway from this earnings quarter was how they are adapting to the Chinese restrictions. China in this quarter stood at 43% of revenue, which is up from Q4 FY25, with only 35%.

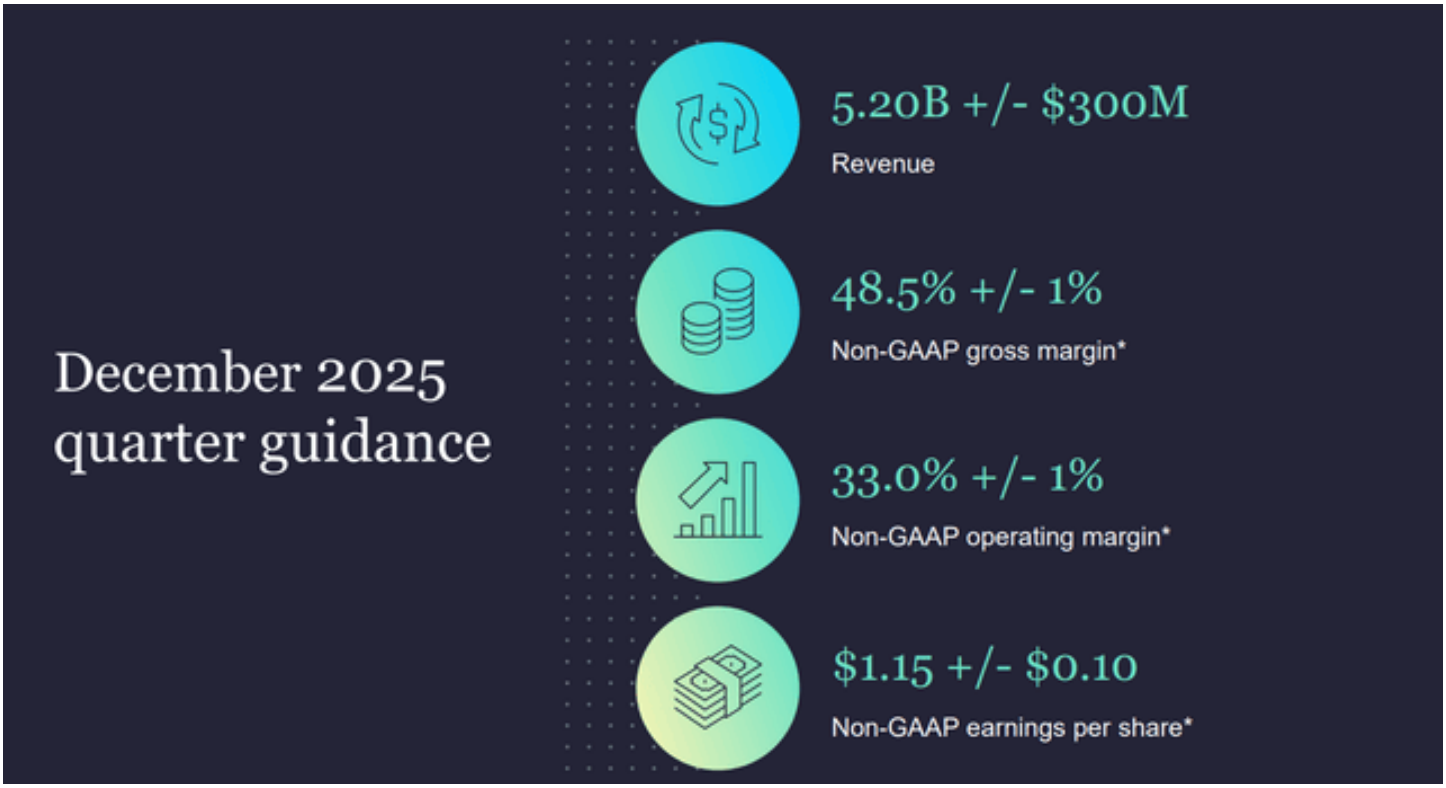

Q2FY26 Guidance (LAM Investor Relations)

Lam has been quite flat since post-earnings as valuations are getting higher and higher; however, Lam's next quarter is expecting negative sequential growth QoQ and growth rates significantly decelerating (18.7% YoY in Q2 FY26 to 27.74% YoY in Q1 FY26), and margins are also slightly contracting, so from the surface without context, Lam seems like it will continue to underperform, but these results are actually amazing from Lam.

The Affiliate Rule

Lam is one of many semiconductor manufacturing equipment suppliers to get affected by the 50% affiliate rule, which creates huge restrictions on exports towards entities that are on the entity list, essentially disrupting Lam's ability to ship manufacturing equipment (at least the most advanced ones) to China, which has also already happened to Nvidia, being restricted from shipping their most advanced GPUs to China and needing to make a custom one just for the China market. This essentially rippled now to Lam, as peer ASML has long been restricted from shipping its most advanced EUV lithography machines to China. This directly eats into Lam's revenue, especially considering the fact that China is currently the largest market that they serve. Considering there was already negative sentiment before the rule, such as ASML and AMAT in the prior quarter, mentioning a potential no growth in 2026 with poor guidance. Due to these factors, I believe there was a narrative problem for Lam, and this quarter essentially clarifies to investors and analysts that the effects won't be a huge toll on Lam, as the prior negative sentiment has been removed, as the future trajectory is clearer, with the guidance for CY26 announced in the January earnings report as always.

Impact of the affiliate rule from the most recent earnings call:

Our December quarter guidance does contemplate roughly a $200 million revenue impact from the recently announced 50% affiliate rule, restricting shipments to certain domestic China customers. Currently, we expect this rule to impact our calendar year 2026 revenues by approximately $600 million. This impact, together with the strong growth anticipated in worldwide fabrication equipment or WFE spending, leads us to expect the China region to represent less than 30% of our overall revenues in calendar year 2026.

The next quarter has been slightly affected with a $200M impact, which means the midpoint for next quarter would be $5.4B without the affiliate rule. This suggests a 23.3% YoY growth without the impact. Lam isn't the only one hit by this affiliate rule, with AMAT expected to get hit $110M for the remainder of the year and $600M in 2026. Despite this, I believe that this has shown investors that, despite the China uncertainty, Lam continues to execute well and expects China revenue to fall below 30% of their overall revenue, which is very bullish, essentially moving on from the uncertainty. This is also led by their EPS continuing to beat Wall Street estimates, as Lam remains resilient despite major headwinds in the near term.

The Narrative Change: What to Look Out For Next Quarter

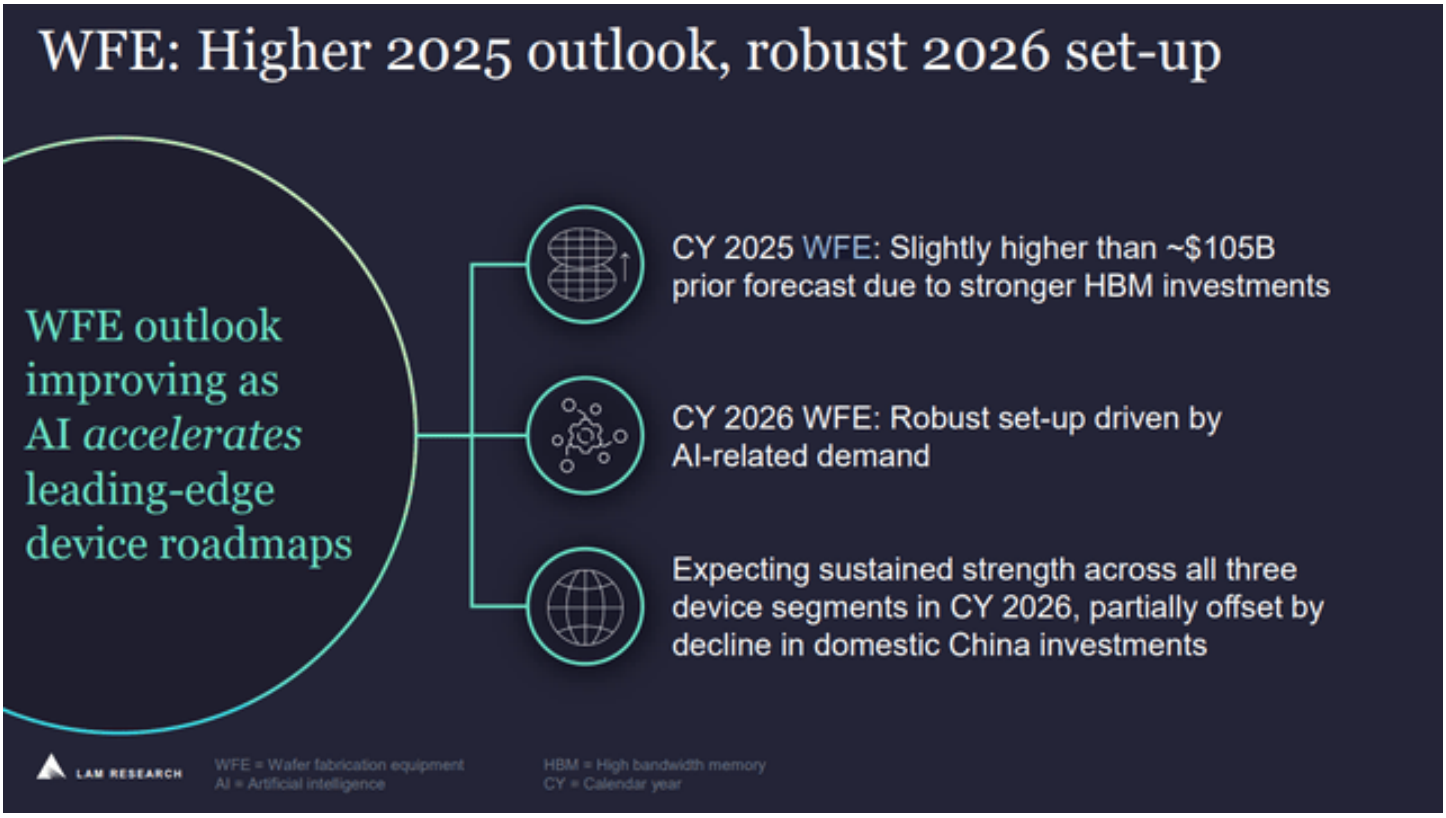

2025 WFE Spending Outlook (LAM Investor Relations)

The narrative has changed in the near term, and the next quarter is very important for investors since the January call and presentation will outline their full outlook for 2026, as well as WFE spending for 2026. The demand for WFE has been slightly higher this year despite the huge tariff fears back in April, as the huge capex and demand for hardware in the AI ecosystem will highly likely continue to drive up demand for WFE, as supply right now remains short of demand across the board.

From the earnings call:

While we are not yet quantifying the level of growth in WFE, I will tell you that calendar '26 looks somewhat second half weighted as we sit here today. The newly restricted China entities would have been weighted to the first half of next year. Customer mix will be a headwind to gross margin a bit next year as the China mix normalizes.

As we look ahead, we see a robust setup for equipment spending in calendar year 2026. AI-related demand should support sustained strength in leading-edge foundry logic and DRAM as well as continued NAND upgrade spending.

Lam's markets, etch and deposition, will outgrow WFE because of nearly every trend that's taking place technology-wise in semiconductor manufacturing, whether that's 3D devices in foundry/logic, in NAND, whether it's smaller, higher aspect ratio devices in stacking using advanced packaging and DRAM, whether it's advanced packaging itself, they're all deposition and etch-intensive products. And therefore, I think over the longer term, confidence to outperform WFE is very high.

And honestly, as we sit here right now and look at the stack up of everybody's plans in China, it's going to be less. That's the best I can tell you. And yes, you're absolutely right. A year ago, when we were sitting here, we would have seen the same thing and then it strengthened through the year. I just right now don't see where that's going to come from next year, Tim.

The reason why the next quarter's outlook is so important is that uncertainty remains, as growth rates have clearly been stunted next year. As the first statement mentions, the toll from lower China revenue will affect 1H of 2026, as well as affect gross margins. The growth of CY26 is also more weighted towards 2H of 2026, so the following 2 quarters are also important to understand how much impact on revenue the China loss would have for 2H 2026, as well as their growth rate, considering it will likely decelerate YoY. WFE spending from China is also likely to fall from statement 4, so the next few quarters are ever more important to monitor the extent of the impact China would have, as a 35% in the prior quarter and 43% this quarter to a target below 30% is quite significant, as the near future is rather uncertain despite huge confidence by investors with their high valuations.

The WFE spending growth is also important, as that will likely create a base case scenario for Lam's growth in the near future. From statement 2, they believe WFE spending YoY will be robust, but we haven't gotten a figure yet, so we are still in the dark. From statement 3, it was mentioned that their growth rate to outperform the WFE spending rate is high, so the projected growth rate on WFE spending could perhaps be used as a base case for Lam's growth in 2026.

With TSMC showing higher capex in the most recent earnings call:

Now let me make some comments on our 2025 CapEx. As the structural AI-related demand continues to be very strong, we continue to invest to support our customers' growth. We are narrowing the range of our 2025 CapEx to be between USD 40 billion and USD 42 billion as compared to USD 38 billion to USD 42 billion previously.

As well as Micron, a major player in DRAM and NAND's most recent presentation (slide 20):

Micron invested $13.8 billion in capex in fiscal 2025. As we continue to make 1γ DRAM and HBM-related investments, we expect fiscal 2026 capex to be higher than fiscal 2025 levels.

All signs remain bullish for 2026, as the trend continues to seem like higher WFE spending next year, as Lam is set to capitalise, especially with the transition from tungsten to molybdenum as well as NAND upgrades.

Final Thoughts

Out of all the WFE names out there, my top preference would probably be Lam, as I do believe that there is slight underappreciation with Lam and their moat compared to the likes of ASML, which has been widely recognised by the market for their indispensable moat.

A lot of bullish comments have been mentioned this quarter, as the sentiment has certainly changed from the prior quarter, as WFE manufacturers were facing restrictions every other week as well as fear from customers. The restrictions have now cooled down significantly since then as the future is way clearer, but uncertainty remains. The narrative has changed, but Lam needs to live up to the expectations, as investors have put a high premium on Lam, so all eyes are on the next quarter for outlook and projected WFE spending for 2026.