

KASASV

KOREAN AMERICAN SEMICONDUCTOR ASSOCIATION IN SILICON VALLEY

March 2026

최고의 선수 vs 최고의 팀

2004년 아테네 올림픽에서 미국의 드림팀은 금메달을 따는데 실패했다. 어떻게 이런일이 있을 수 있지? 미국 농구팀은 전원이 NBA 리그에 속해 있었다. 그런데 어떻게 다른 나라가 미국을 이길 수 있었을까? 답은 간단하다. 미국은 **최고의 선수들을** 출전시켰고 다른 나라는 **최고의 팀을** 출전시켰기 때문이다. 혼자서는 큰일을 할 수 없다.

농구에서 골을 넣으려면 열개의 손가락이 필요하다. 어느 조직이나 중요한 역할을 하는 사람이 있기 마련이다. 하지만, 실제로는 각 개인의 팀전체의 성공에 이바지하고 있어야 한다. 무엇보다도 중요한 것은 자신의 성과가 어떻게 팀의 이익과 생존에 영향을 끼치는지 이해하고 있는 것이다.

개인의 위대함은 그 사람이 맡은 역할의 크기가 아니라 그 **역할에 쏟아부은 노력의 크기**에 따라 결정된다. 내가 다른 사람보다 더 나은지 걱정하지 마라. 대신에 네가 될 수 있는 최고의 사람이 되기 위해 끊임없이 노력해라

- '88 연승의 비밀' 에서, 존우드

INTRO

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

3월에는 2026년 올해 첫번째 골프토너먼트가 있습니다. 3월 28일(토) 예정입니다. 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 있습니다. 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

Emerging Markets And The AI Surge

Feb 24, 2026. Mutual Fund Commentaries

Summary

- Seven of the ten largest contributors to the MSCI EM Index return in 2025 were AI-related and accounted for more than 40% of the index's 34% return.

- Significant portion of the AI physical infrastructure is sourced from EM companies, especially Asia-based global enterprises such as the semiconductor manufacturing giant TSMC in Taiwan.
- The AI boom is engendering structural changes in the market for memory products that should support a period of higher and more consistent profitability across the memory industry.
- HBM production is highly capacity intensive; as a growing share of fabrication capacity is allocated to these chips, the industry's available capacity for conventional DRAM is tightening.

The following segment was excerpted from the [Harding Loevner Emerging Markets Equity Composite Q4 2025 Commentary](#).

Of the 10 largest contributors to the MSCI EM Index return in 2025, seven were AI-related and accounted for more than 40% of the index's 34% return. The surge in AI-related stocks reflects the sharply accelerating capital investment into the physical infrastructure that enables AI development. Hyperscalers, such as Amazon (AMZN) and Meta (META), and other Cloud Service Providers (CSPs), which are leading the large-scale build-out of AI data centers, have repeatedly increased capital expenditure plans. Alliance Bernstein (AB) estimates that CSP capex could grow at a 36% annual growth rate from 2024 to 2027, reaching approximately US\$630 billion.

EMs have been a standout beneficiary of this investment wave because a significant portion of the AI physical infrastructure is sourced from EM companies, especially Asia-based global enterprises such as the semiconductor manufacturing giant TSMC (TSM) in Taiwan. These businesses sit at the technological frontier across several industries that are critical to AI development, including leading-edge foundry services, advanced chip packaging, high-bandwidth memory manufacturing, server and rack integration, and power and cooling system design and production. The current AI computing buildout would not be happening without EM companies.

The surge in AI-related investment and the resulting uplift in valuations across the supply chain do not come without risk. While real-world AI applications are expanding rapidly, it remains uncertain whether the profits generated from AI will eventually be large enough, and expand quickly enough, to justify the scale of capital being deployed. Concerns around an AI bubble have intensified. We therefore continue to focus on identifying where AI investment is leading to genuine structural improvement in companies' competitive positions within their industries and thereby sustainably enhancing their business prospects. One area where we project such an improved outlook is the memory chip industry, led by South Korea's Samsung Electronics (SSNLF) and SK hynix (HXSCL) ('HXSCL').

The industry's core product is DRAM (Dynamic Random Access Memory) chips, which are essential for the operation of computer systems, smartphones, and servers. Manufacturing this product is highly capital intensive, with profitability historically reliant upon scale, cost leadership, and sustained investment in R&D. The industry is also notoriously cyclical, with participants experiencing violent swings in pricing and profit margins as demand and supply have, historically, failed to remain in balance for any significant period of time. Higher demand and pricing have precipitated bouts of over-investment leading to gluts of mostly commoditized product flooding markets as demand inevitably fades. As a result, over the last couple of decades the industry has become increasingly concentrated; Samsung and SK hynix each enjoy over 30% market share. Many smaller competitors with weaker scale economies and hence lower cash generation to fund R&D have been unable to keep pace with technological advances and have been gradually forced to exit the industry.

Increasing industry concentration has underpinned our investment case for Samsung Electronics, a core holding in our portfolio for more than two decades. We believed the company would benefit, together with other industry leaders, from more rational coordination of capital investment leading to healthier, less-volatile profit cycles. However, the past decade has provided only limited evidence that greater industry concentration has actually dampened cyclicality. Nevertheless, the AI boom is engendering structural changes in the market for memory products that should support a period of higher and more consistent profitability across the memory industry while strengthening the competitive advantages of Samsung and SK hynix, its leading players.

There are three key developments that are changing the industry's dynamics:

Growing Demand for Customized, High-Value Memory Products

One reason memory manufacturers have struggled to sustain pricing power and returns on capital is that for many applications DRAM has become a standardized product. The JEDEC Solid State Technology Association establishes industry-wide standards that have historically allowed for interchangeability in memory products across suppliers, which kept switching costs low for buyers. But today, while JEDEC standards remain critical, the demands of AI workloads and the physical limits of traditional chip design are placing a premium on customized, higher-value solutions. Traditional planar scaling is increasingly giving way to more complex 3D architectures, including vertically stacked and tightly packaged memory solutions. The greater complexity of 3D architecture and customers' demands for chip customization raises barriers to entry for memory manufacturing. These factors also increase the pricing power of the leading manufacturers capable of consistently delivering high-performance, customized memory products at scale.

The Need to Surmount the Memory Wall

The gap between compute throughput and memory bandwidth—the “memory wall,” in industry parlance—has widened. This limits system-level performance despite rapid advances in computing chips, both graphics processing units (GPUS) and application-specific integrated circuits (ASICs), which are specialized semiconductors designed to execute specific tasks. Training and inference workloads now require not only large memory capacity, but also sustained, high-bandwidth memory access well beyond what conventional DRAM can deliver. HBM (High Bandwidth Memory), which vertically stacks multiple DRAM dies on a chip and is tightly integrated with GPUs and ASICs through advanced packaging technologies, is therefore critical for AI training and operations. Importantly, HBM development cycles are now closely linked to GPU and hyperscaler design cycles, increasing co-development requirements and compressing memory-product qualification timelines. This further strengthens the competitive positions of leading memory manufacturers able to meet customers' more onerous production requirements.

Increasing Constraints on Memory Manufacturing Capacity

HBM production is highly capacity intensive; as a growing share of fabrication capacity is allocated to these chips, the industry's available capacity for conventional DRAM is tightening. In late 2024, this emerging structural constraint was obscured by cyclical factors. Customers overbuilt their inventories of conventional memory products earlier in the year. The process of digesting this inventory was prolonged by weak smartphone and PC demand, temporarily placing pressure on DRAM spot prices. Chinese producers also began to have a disruptive impact on DRAM prices. CXMT, China's leading DRAM manufacturer, has gained share, particularly within China's smartphone market; some earlier estimates suggested it could reach around 8% of global DRAM supply in 2025.

However, through the course of 2025, DRAM supply tightened, first modestly then acutely, as AI-related demand expanded beyond just HBM products and across other high-performance DRAM products used in AI servers. The expected glut of legacy DRAM product from Chinese suppliers also did not materialize as CXMT and other Chinese manufacturers instead prioritized the development of their own HBM products in a long-term bid to reduce China's reliance on global suppliers. Customers scrambled to gain preemptive DRAM supply commitments, resulting in sharp price increases toward year end.

The AI boom is engendering structural changes in the market for memory products that should support a period of higher and more consistent profitability across the memory industry while strengthening the competitive advantages of its leading players.

These three changes in the memory market caused largely by AI industry demand not only support the current memory upcycle but should also underpin structurally higher and less-volatile profitability for leading memory manufacturers over the long term. Based on this assessment, we added to our long-held holding in Samsung Electronics and initiated a new one in SK hynix.

Historically, we have viewed SK hynix as a structurally weaker player in the memory industry. Unlike Samsung's more diversified business, SK hynix's sole focus on memory amplified its earnings volatility and led to weak cash generation during downturns, constraining its strategic flexibility. While SK hynix was a top-tier producer by scale, its weaker balance sheet and subordinate position relative to Samsung as a developer of leading-edge technology further discouraged us from investing in the company. As recently as 2023, SK hynix had substantial operating losses that raised concerns about its rising debt burden (approaching a 30% net debt/equity ratio) and onerous debt maturity schedule. But the company's quality and growth status has since improved significantly. The company made an early strategic commitment to HBM, investing ahead of the demand curve in both the technology and the capabilities needed to partner with customers on memory chip development. Combined with execution setbacks at Samsung, this enabled SK hynix to emerge as the global leader in HBM and the primary supplier to NVIDIA (NVDA). SK hynix's deep collaboration across the semiconductor supply chain, including close integration with TSMC on advanced process technologies, has translated into superior manufacturing yields, faster qualification cycles, and consistently higher-performing chips.

SK hynix's shift toward HBM has materially improved its financial profile. Customized, system-critical memory products command significantly higher average selling prices and margins, supporting stronger free cash flow that has allowed SK hynix to repair its balance sheet. The company achieved a net-cash position in its third-quarter earnings report. While memory remains cyclical, SK hynix's repositioning as a critical enabler of AI should result in structurally higher returns on capital. We estimate through-cycle operating margins of over 25%, above historical levels, giving us a positive view of long-term return potential despite the near 300% rise in the company's share price in 2025.

Improved industry dynamics also benefit Samsung's memory business, which is expected to generate record profits in 2026 as its efforts to regain HBM market share bear fruit and the company benefits from the tight supply conditions in conventional memory.

Forecasting the precise duration of the current memory upcycle is difficult, and risks remain. Most notably, if AI infrastructure spending proves front-loaded rather than durable and hyperscalers decrease the pace of their capex growth, then the industry could rapidly transition from supply scarcity to oversupply. The other relevant

risk is whether the expected cadence of system development and reliance upon ever-more-powerful memory components continues. If it slows, there is a risk that design and manufacturing complexity also decelerates and facilitates renewed standardization and a re-commoditization of high-performance DRAM. This would likely erode the current strong pricing power of memory players and crimp profit margins. Despite these risks, we think the favorable business environment for Samsung and SK hynix is likely to extend into late 2026 and potentially beyond. More importantly, we believe the structural forces behind the current "super" cycle will persist across multiple cycles, supporting higher and more consistent returns for the memory industry.

Onto Innovation projects over 30% advanced packaging growth in 2026 amid record backlog and AI demand

Feb 19, 2026 AI-Generated Earnings Calls Insights

Earnings Call Insights: Onto Innovation Inc. (ONTO) Q4 2025

Management View

- CEO Michael Plisinski highlighted that Onto Innovation ended 2025 "on a high note with orders from 2.5D packaging for AI devices more than doubling in the quarter, contributing to a record revenue of \$267 million." He emphasized the positive impact of AI investments on semiconductor capital equipment spending, referencing forecasts for global AI infrastructure growth and increased CapEx by industry leaders.
- Plisinski announced a volume purchase agreement from an HBM customer covering Dragonfly 2D and 3D bump metrology demand through 2027, valued at over \$240 million, including over \$60 million in systems for 3D bump metrology. He stated, "our expanding portfolio of technology is putting us in a position to increase the value we deliver to our customers, serving the seemingly insatiable demand for AI."
- The advanced packaging business grew over 25% sequentially, driven by demand for Dragonfly inspection and Iris films metrology, with multiple customer evaluations underway for next-generation inspection systems. Plisinski noted, "preliminary feedback on system performance has been positive with customers acknowledging significant improvement in optical performance and higher throughput."
- The company received additional orders in 3Di metrology, including from an OEM for panel level process development, and sees growth in panel-level packaging as a trend for enterprise server and AI device designers.
- CFO Brian Roberts reported "record revenue of \$267 million, representing a 22% increase from Q3," and "a record level of \$95 million of cash in the quarter for a cash conversion of approximately 150% of non-GAAP net income." Roberts added, "we delivered more than 50% of our tools in Q4 from our extended factories, completed the acquisition of Semilab in mid-November and implemented a more robust forecasting and spending control process."

Outlook

- Plisinski stated, "we estimate advanced packaging revenue to grow over 30% in 2026, resulting in a new revenue record for this market." He also projected revenue for the first quarter to be in the range of

\$275 million to \$285 million and expects "demand to continue to increase in the second quarter with revenue exceeding \$300 million."

- Roberts forecasted operating margins for Q1 to be "approximately 25.5% to 26.5%" and earnings per share for the quarter in the range of \$1.26 to \$1.36 per share, with an estimated tax rate of approximately 16% and about 49.9 million shares outstanding.
- Management expects power semiconductor revenue to decline around 10% in 2026 and similar decreases for Semilab, reflecting weakening EV demand and infrastructure spending.

Financial Results

- Advanced packaging and specialty devices in Q4 represented approximately \$145 million in revenue, with sales from the 2.5D packaging business doubling compared to Q3.
- Advanced nodes revenue grew by over 30% sequentially to \$72 million in Q4, with annual advanced nodes revenue more than doubling to \$308 million in 2025.
- Gross margin improved by about 50 basis points to 54.6% from Q3, and operating margins improved to 25.2% in Q4, an increase of 410 basis points.
- The backlog nearly doubled over the last three months to a record level of about two quarters.
- Upon closing the Semilab acquisition, \$445 million was paid in cash and 641,771 shares of common stock were issued.

Q&A

- Ezra Weener, Jefferies: Questioned market outlook and capacity. Plisinski responded, "our capacity, we're set up to be able to serve a \$2 billion run rate...I don't see capacity being a big issue for us. It's more on the supply chain side."
- Craig Ellis, B. Riley: Asked about advanced packaging growth contour and notable wins. Plisinski said, "advanced packaging revenue to be relatively stable throughout...the adoption of G5...could certainly add even stronger upside for the second half."
- Edward Yang, Oppenheimer: Inquired about the \$240 million VPA for HBM. Plisinski explained, "it's a 2-year, so it's extending into 2 years. It was more 2027 weighted, so maybe 2/3, 1/3."
- Yu Shi, Needham: Sought clarity on AI packaging and litho business. Plisinski indicated, "the vast majority of that 30% growth is all tied to supporting the strong demand in AI," and noted growing engagement and resumed orders in the litho business.
- Vedvati Shrotre, Evercore: Asked about Semilab contributions. Roberts stated, "they did about \$9 million of revenue contribution in Q4 after the close in mid-November. For 2026...somewhere low \$100 million to \$110 million was kind of the revenue."
- Brian Chin, Stifel: Questioned new product contributions. Plisinski said, "these are early penetrations. So you're talking maybe 10% of the business, maybe a little less, but certainly growing into 2027."

Sentiment Analysis

- Analysts pressed for specifics on growth drivers, timing of revenues, and customer adoption, reflecting a slightly positive and constructive tone, with curiosity about the implications of large VPAs and product adoption.
- Management maintained a confident and positive tone in prepared remarks, citing "great momentum" and "meaningful improvement" in key metrics, while in Q&A, they acknowledged some uncertainties in timing and supply chain, but reiterated strong demand and capacity.

- Compared to the previous quarter, analyst tone shifted from cautious optimism to greater focus on the scale and timing of new agreements and growth, while management's confidence increased, supported by record backlog and new long-term agreements.
- Quarter-over-Quarter Comparison
- Guidance language became more specific, with clear projections for Q1 and Q2 2026 and targeted growth percentages for advanced packaging and nodes, versus broader directional statements in Q3.
- Strategic focus shifted from initial product launches and integration to scaling production, managing supply chain constraints, and executing on large VPAs.
- Management tone grew more confident, emphasizing execution and expanded visibility, while analysts focused on the sustainability of growth and the impact of new agreements.
- Key metric changes include a record \$267 million in Q4 revenue, doubling of 2.5D packaging orders, and the announcement of a \$240 million VPA.

Risks and Concerns

- Management identified supply chain constraints, particularly in precision optics, as a key challenge, stating, "the rapid increase in orders...is putting a strain on some of our suppliers."
- Power semiconductor and Semilab revenue are expected to decline due to weakening EV demand and infrastructure spending.
- Analysts queried the timing and magnitude of new customer agreements and the pace of product adoption, suggesting some uncertainty in the trajectory of advanced nodes and packaging revenues.

Final Takeaway

- Onto Innovation closed 2025 with record revenue and cash generation, driven by accelerating demand in AI and advanced packaging, highlighted by a significant \$240 million volume purchase agreement. Management projects advanced packaging revenue to grow over 30% in 2026, with robust visibility supported by a record backlog and expanding customer engagements. While monitoring supply chain capacity and anticipating some headwinds in power semiconductors, the company expects continued margin expansion and is confident in its positioning for sustained growth in the rapidly evolving semiconductor market.

Camtek Ltd. (CAMT) Q4 2025 Earnings Call

Feb 18, 2026 AI-Generated Earnings Calls Insights

Management View

- Record Q4 revenue: \$128M
- FY2025 revenue: \$496.9M (+16% YoY) — reached the \$0.5B milestone
- AI-related applications now ~50% of revenue
- Strong conviction that 2026 will be another double-digit growth year
- Visibility materially improved vs. prior quarter
- Management repeatedly emphasized:
 - “Significantly increased order flow”

- “Much better visibility”
- “Very high confidence in second half”
- Key strategic themes:
 - HBM4 transition is a major growth catalyst
 - Hawk + Eagle Gen5 adoption accelerating
 - Market share expected to increase, not just maintain
 - Capacity expanded to >\$700M revenue capability
 - Early signals of 2027 demand emerging

Outlook

- Q1 2026 Guidance
 - Revenue: ~\$120M (down sequentially)
 - Slow start due to customer order timing
- 2026 Full-Year Outlook
 - Double-digit growth year
 - Stronger growth in:
 - H2 2026
 - Especially memory/HBM capacity builds
 - Gross margin:
 - 1H: ~50.5–51.5%
 - 2H: Improvement expected
 - OpEx:
 - R&D investments increasing in 1H
- Visibility
 - Strong H2 2026 backlog
 - Early discussions already beginning for 2027 shipments

Financial Results (Q4 & FY2025)

<u>Metric</u>	<u>Q4 2025</u>	<u>YoY</u>
Revenue	\$128.1M	+9%
Gross Margin	51.1%	Slight expansion
Operating Margin	28.6%	Slight decline QoQ
Net Income	\$40.7M	
EPS	\$0.81	

FY2025:

- Revenue: \$496.9M (+16%)
- Gross margin: 51.6%
- Operating margin: 30%

Balance Sheet

- Cash: \$851M
- Strong FCF: \$61M in Q4
- Inventory normalized after Hawk/G5 ramp
- DSO improved (81 → 65 days)

Financial profile = clean, strong, self-funded growth

Quarter-over-Quarter Comparison

<u>Area</u>	<u>Q3 → Q4 Trend</u>
Revenue	Slight increase
Gross Margin	Stable
Operating Margin	Slight compression
Cash	Increased
Inventory	Reduced significantly

Q1 sequential dip is purely timing-related, not demand weakness.

Key Business Drivers

1. HBM4 Transition

- More inspection-intensive
- Higher density structures
- Greater metrology requirements
- Positive mix impact
- Hawk adoption likely to increase

2. Hawk vs Eagle Gen 5

- 2025 mix: ~30% of revenue
- 2026 expected: ≥50%
- Hawk:
 - High throughput

- Hybrid bonding capable
- High-accuracy
- G5:
 - OSAT stronghold
 - Flexible, strong ROI
- Future capacity builds skewing more toward Hawk.

3. OSAT Strength

- ~50% of business
- CoWoS-like technologies moving into OSATs
- Real purchase orders in hand

Competitive & Market Position

- Competitor references (implicitly vs. KLA and others):
 - No share loss
 - Expect share gain
 - Expanding into additional process steps
 - 3D + 2D combined capability is key differentiation
- Management message:
 - “We are the tool of reference for 3D metrology.”
- They believe they are structurally positioned better than broader WFE due to AI/HPC concentration.

Risks & Concerns

1. Timing Risk

- Heavy second-half weighting
- Execution risk if customer CapEx shifts

2. Customer Concentration

- AI/HBM memory suppliers driving growth
- If memory CapEx moderates → sensitivity

3. Competitive Intensity

- KLA reporting strong packaging growth
- Camtek insists markets not apples-to-apples

4. Sequential Decline vs Test/Probe Peers

- Divergence may concern investors
- Management attributes to order timing only

Sentiment Analysis

Indicators:

- Repeated use of “very confident”
- Capacity expansion underway
- Early visibility into 2027
- Explicit statement: “We did not lose market share”
- Expect meaningful share increase

Confidence level vs prior call: Materially higher

Management energy level: 8.5/10

Final Takeaway (Executive Summary)

Camtek exits 2025 in strong shape with:

- Record revenue
- 51%+ gross margins
- \$851M cash
- AI exposure now 50% of revenue
- Hawk/G5 ramp accelerating
- HBM4 transition highly favorable
- Share gain expected
- 2026 double-digit growth with strong H2 weighting
- Early visibility into 2027

Camtek is leveraged to the most inspection-intensive segments of AI-driven advanced packaging and believes it will grow faster than broader WFE again in 2026.

Axcelis anticipates flat 2026 revenue as memory growth offsets declines in power and mature markets

Feb 17, 2026 AI-Generated Earnings Calls Insights

Earnings Call Insights: Axcelis Technologies (ACLS) Q4 2025

Management View

- CEO Russell Low highlighted that "we generated solid results in the fourth quarter with revenue of \$238 million and non-GAAP earnings per diluted share of \$1.49, both exceeding our outlook." He attributed the outperformance to stronger CS&I aftermarket revenue, which also benefited gross

margins. CEO Low noted significant progress toward the Veeco merger, stating, "Shareholders of both companies voted in favor of the transaction at our special meetings on February 6." Regulatory clearance has been received in several jurisdictions, with final approval from China pending and closing expected in the second half of 2026.

- CEO Low explained that mature-node applications, particularly in power and general mature segments, accounted for the majority of system shipments, with bookings improving significantly on a sequential basis. He also introduced the Purion H6 high current ion implanter, emphasizing its advancements in dose repeatability, purity, precision, and productivity.
- In the power segment, CEO Low observed a slight moderation in silicon carbide shipments and ongoing customer focus on technology transitions. He reported growing interest in Axcelis's high-energy channeling capabilities, particularly for superjunction development.
- CEO Low stated, "We expect strong long-term demand through the cycles, driven by clear secular trends such as the growing penetration rate of silicon carbide into electric vehicles...and the growing adoption of silicon carbide outside of the automotive sector."
- In memory, CEO Low said demand improved sequentially for DRAM and HBM, and momentum is expected to extend into 2026. He highlighted a key customer win in North America, broadening Axcelis's presence beyond Korea.
- CFO James Coogan reported, "Fourth quarter revenue was \$238 million, with systems revenue at \$156 million and CS&I revenue reaching another quarterly record of \$82 million, both above our expectations for the quarter."

Outlook

- CFO Coogan provided guidance for Q1 2026, stating, "We expect revenue in the first quarter of approximately \$195 million." Non-GAAP gross margins are projected at approximately 41%, with non-GAAP operating expenses of around \$59 million and adjusted EBITDA of about \$26 million. Non-GAAP diluted earnings per share is expected to be approximately \$0.71.
- For full year 2026, CFO Coogan said, "We anticipate overall revenue to be relatively flat compared to 2025 levels." He expects memory market growth to offset declines in the power and general mature markets, with revenue second half weighted.
- Non-GAAP gross margins for 2026 are estimated in the low to mid-40% range, with a modest tariff impact.

Financial Results

- Axcelis reported Q4 revenue of \$238 million, systems revenue of \$156 million, and CS&I revenue of \$82 million.
- Non-GAAP gross margin was 47.3%, above the outlook of 43%. GAAP operating margin was 15.2%, and non-GAAP operating margin was 21.1%. Adjusted EBITDA was \$55 million, with an adjusted EBITDA margin of 22.9%.
- GAAP diluted earnings per share was \$1.10; non-GAAP diluted earnings per share was \$1.49.
- Free cash flow in Q4 was negative \$9 million, affected by the timing of sales and \$5 million in transaction expenses related to the Veeco merger. Full year free cash flow was \$107 million.
- Axcelis repurchased \$25 million in shares during Q4 and exited with \$557 million in cash, cash equivalents, and marketable securities.

Q&A

- Jonathan Dorsheimer, William Blair: Asked about memory market capacity expansion and DRAM demand. CEO Low said growth is limited by cleanroom space but expects "DRAM significantly up in 2026 compared to 2025, although somewhat of a lower base in '25."
- Dorsheimer followed up on CS&I silicon carbide upgrades. CEO Low stated, "Only a very small part of our installed base has gone up to 200-millimeter...ample opportunity to continue to provide these really high-value upgrades."
- Denis Pyatchanin, Needham: Asked about segments driving bookings. CEO Low said bookings matched prior segments, mainly general mature and power, with a bump in memory.
- Pyatchanin questioned the outlook for power and general mature. CEO Low indicated both segments are "slightly down," with memory picking up the slack as cleanroom space becomes available.
- Christian Schwab, Craig-Hallum: Inquired about memory CapEx cycles and ion implant intensity per 100,000 wafer starts. CEO Low confirmed about \$150 million to \$200 million in capital for implants in DRAM per 100,000 wafer starts.
- Craig Ellis, B. Riley: Asked about mature foundry utilization rates and China demand intensity. CEO Low highlighted high utilization rates in China and ongoing demand for equipment to meet self-sufficiency goals. CFO Coogan said, "We see it to be relatively flat to down slightly" for China in 2026.
- Jack Egan, Charter Equity: Asked about power and general mature guidance. CEO Low cited careful capacity additions in China and rising utilization rates outside China. Egan also queried memory order cycle dynamics, with CFO Coogan clarifying that customer forecasts are typically multi-quarter, supporting confidence in memory recovery.
- Duksan Jang, BofA: Asked about Q1 guidance components and 2026 second-half revenue weighting. CFO Coogan explained the Q1 guide reflects both pull-in and pushout effects, and memory is expected to drive the second-half weighting.
- David Duley, Steelhead: Asked about memory growth percentages and historical memory system revenue. CFO Coogan did not offer a specific 2026 percentage but expects healthy growth from a low 2025 base. CEO Low and CFO Coogan anticipate acceleration in 2027.

Sentiment Analysis

- Analyst tone was predominantly constructive yet probing, with focus on memory momentum, power and general mature weakness, and China demand. Skepticism emerged around the sustainability of CS&I upgrades and timing of a broader recovery.
- Management maintained a confident and optimistic tone in prepared remarks, emphasizing strategic wins and long-term opportunities, but responses in Q&A were more measured, especially regarding power and general mature segments, and included cautious language such as "slightly down" and "relatively flat."
- Compared to the previous quarter, analyst sentiment shifted to a more positive outlook on memory, while management's tone became more confident about memory but remained cautious on power and general mature recovery.

Quarter-over-Quarter Comparison

- Management in Q4 guided for flat overall 2026 revenue, compared to Q3's outlook of flattish Q4 and Q1 2026 revenues, now specifying memory as the growth driver offsetting declines in power and general mature markets.

- In Q4, management reported record CS&I revenue and highlighted the launch of Purion H6, whereas Q3 focused on new product announcements and ongoing strategic initiatives.
- Q4's Q&A included more detailed discussion of memory's role and cleanroom constraints, while Q3 analysts focused on China demand and power segments. Management's confidence in memory has increased, while its stance on power and general mature recovery remains cautious.
- Risks and Concerns
- Management noted muted near-term demand for silicon carbide and power markets as customers absorb existing capacity, with memory growth constrained by cleanroom space.
- Tariff impacts are expected to modestly affect margins in 2026.
- Q1 revenue and gross margin are projected to decline sequentially due to less favorable mix, lower CS&I revenue, and tariff impacts.
- Analysts raised concerns about the sustainability of CS&I growth and the timing of recovery in power and general mature segments.

Final Takeaway


- Axcelis Technologies delivered strong Q4 results, driven by record CS&I aftermarket revenue and strategic wins in the memory segment. Management anticipates overall 2026 revenue will be flat compared to 2025, with growth in memory, particularly DRAM, offsetting softness in power and general mature markets. The company remains focused on disciplined execution, the rollout of new products like the Purion H6, and integration planning for the pending Veeco merger, while maintaining a solid balance sheet and strong cash position to support future opportunities.

The Capex Cycle And Semiconductors

Feb 10, 2026 Semiconductor Analyst

Summary

- Big Tech's 2026 capex guidance — led by AMZN, GOOGL, META — signals another aggressive investment cycle, driving bullish prospects for semiconductors.
- Memory remains in a super cycle: HBM demand is surging, DRAM supply is tight, and sellers like MU, SK Hynix, and Samsung should benefit through 2026.
- TSMC's capex hike to ~\$54B and 63–65% gross margin guidance reinforce strength in AI, advanced packaging, and foundry leadership.
- Semicap equipment makers (LRCX, KLAC, ASML, AMAT) are well-positioned, while consumer electronics face margin pressure from rising DRAM prices.
- NVIDIA (NVDA) and Broadcom (AVGO) should be direct winners from the 2026 capex step-up, with NVDA's platform cadence and AVGO's Google-led TPU ramp positioning both to capture outsized AI infrastructure spend.



J Studios/DigitalVision via Getty Images

Big Tech capex has turned into the most important semiconductor demand signal again.

Big Tech earnings season has been marked by the shocking spending plans released by Amazon (AMZN), Google (GOOGL), and Meta (META), which disclosed figures far above prior market expectations. As you can see in the table below, consensus failed to anticipate another leg up in capex, similar to what we saw from 2024 to 2025.

Company	Today (USD)	Above 1 month ago
AMZN	-\$200,000.00	46%
MSFT	-\$117,787.57	7%
META	-\$125,000.00	15%
GOOGL	-\$180,000.00	54%
ORCL	-\$56,382.41	0%
Total	-\$679,169.98	28%

Capex levels for 2026 (AMZN, GOOGL, and META reflect company-provided midpoint guidance, while MSFT and ORCL are S&P Capital IQ consensus averages)

Capex across these five companies rose 72% in 2025, and the figures above imply another ~65% increase versus 2024. Given that Microsoft did not provide full-year guidance and Oracle's last earnings update was earlier, I think the eventual 2026 total could end up even higher.

While these increases have reignited the debate over whether the trend is economically sustainable and whether these investments will actually pay off, which has helped drive some valuation compression across Big

Tech, that's a debate for another forum. From a semiconductor perspective, I see this as a pretty bullish setup for another strong year.

Much of this capex will go into logic chips: GPUs, ASICs, or TPUs. Microsoft, for instance, said on its latest earnings call that two-thirds of capex was on short-lived assets, primarily GPUs and CPUs.

So... which companies should benefit?

The memory super cycle

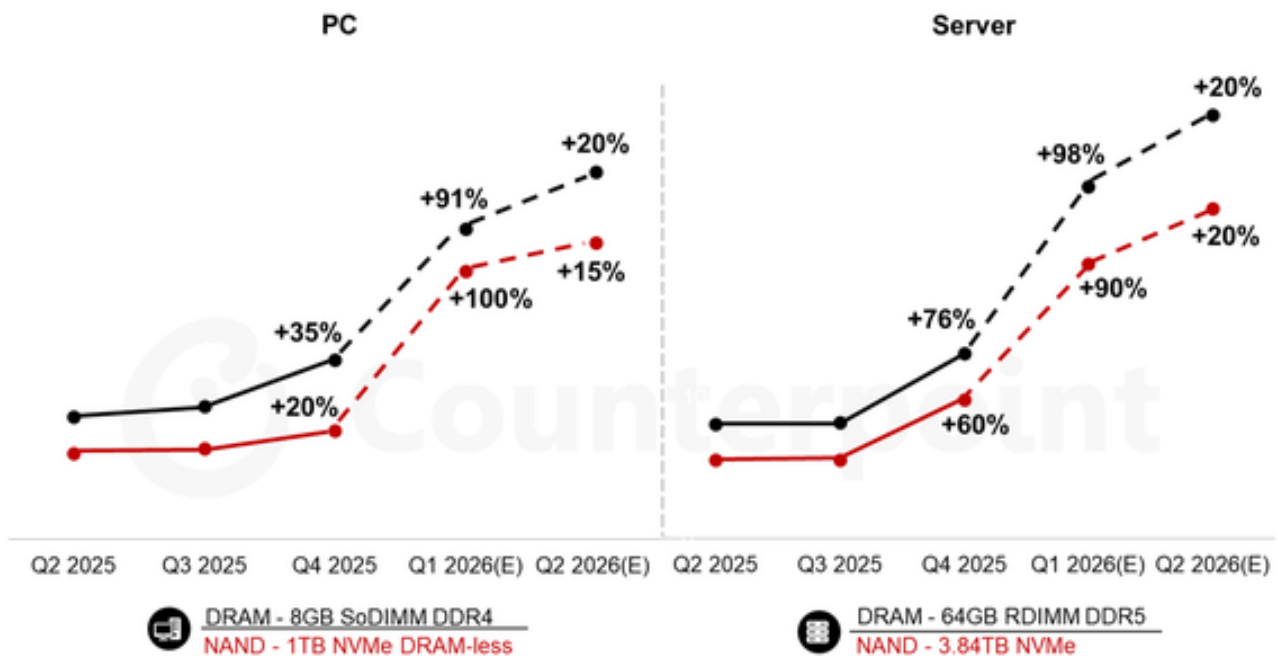
I'd like to start with memory and its implications because it has been the hottest topic in semis lately, and it has spillovers into a lot of segments. I believe these capex estimates are yet another indication that today's memory cycle still has room to run, and possibly evidence that we are not at peak yet.

DRAM scaled faster than logic through most of semiconductor history and, with much less differentiation among players, it ultimately became a commoditized segment. That also led to a consolidation of the market into three main players. Cycles are typically driven by new technologies, and this time is no different.

HBM has consolidated itself as a key differentiator for AI products when combined with logic, and its content keeps rising. For example, NVIDIA is raising HBM intensity again with Blackwell Ultra and Rubin, and Rubin Ultra looks like a step-change in system memory (roughly ~288GB to ~1TB). The same trend is showing up in custom silicon, with Google's TPU and Amazon's Trainium moving from HBM3E 8-Hi to 12-Hi, and AMD also lifting memory from ~288GB on MI350 to ~432GB on MI400.

For memory players, HBM became an opportunity to differentiate in a commoditized market and take share, mainly because there is no true standard or default yet. It also offered a path to grow in a more predictable way and with better ASPs. That is why capex has been heavily directed to HBM and less to the commoditized DRAM side.

This capex allocation has ultimately created a supply-demand imbalance in commodity DRAM, pushing DRAM prices much higher.



DRAM Prices (Counterpoint)

These price surges typically lead customers to build inventory, which can add another layer of inflation. For memory players, it creates a weird dichotomy where commoditized DRAM can sometimes carry better margins than its “chic brother” HBM. Does that mean more capex will move back to commodity DRAM? Possibly. But with these rising capex plans, it is less obvious, because no one wants to lose share on the HBM front. I believe the imbalance should persist across both DRAM markets through 2026.

Therefore, 2026 should still be a great year for memory sellers like Micron (MU), SK Hynix, and Samsung.

On the other hand, for products that depend heavily on commodity DRAM, this is not a great setup. That is true for cellphones in particular, but also for PCs and laptops. Higher DRAM prices hit margins and can pressure demand through higher end prices.

Qualcomm (QCOM) has already disappointed the market in its latest earnings, and Arm (ARM) and MediaTek (MDTK.F) have noted similar issues. Apple (AAPL) has more pricing power over its content, so it is more protected, but even they have signaled that if this trend persists, 2H can show more impact of memory prices on gross margins. Both Intel (INTC) and AMD (AMD) also flagged concerns in their last earnings calls. As AMD has been taking share, I think this could be especially painful for Intel if price increases persist.

Memory players are starting to react more aggressively and plan capacity expansions, but increasing capacity is not easy, so this imbalance should not be resolved this year. I’ve commented on that in my latest ASML article.

Company	6M capex revision			3M capex revision		
	CY2026	CY2027	CY2028	CY2026	CY2027	CY2028
Micron	36%	76%	62%	18%	23%	28%
Intel	-11%	-2%	27%	-3%	-1%	-6%
TSMC	14%	16%	9%	18%	15%	18%
SK Hynix	30%	38%	56%	18%	26%	0%
Samsung	4%	3%	7%	6%	7%	0%
Total	13%	19%	20%	12%	13%	10%

CAPEX estimates evolution (S&P CIQ)

These capex revisions from memory players also support the thesis for semicap equipment, especially for Lam Research, which has historically higher memory exposure, but also for KLA (KLAC), ASML (ASML), and Applied Materials (AMAT).

Logic

Still, much of this capex will be invested into logic: GPUs, ASICs, or TPUs. Microsoft, for instance, said on their latest earnings that two-thirds of capex was on short-lived assets, primarily GPUs and CPUs. For other Big Techs, I estimate that share in the 70% to 80% range.

I believe this capex surprise is positive for the obvious names in logic. Memory, however, can still affect margins, as the share of memory content keeps increasing, as discussed. The more premium the product, the less impact. That's something to monitor.

While the fight for share has been heating up, with Broadcom, AMD, and MRVL (MRVL) all trying to steal NVDA's share, one company is still the producer of them all: Taiwan Semiconductor. At the end of the day, it is the factory behind these chips, regardless of whose logo is on the box.

Taiwan Semiconductor (TSM)

The company has been raising its capex from ~\$30B in 2024 to ~\$41B in 2025, and is now guiding ~\$54B at the midpoint for 2026. They tend to position themselves conservatively when it comes to capex, so I see that as a very bullish sign. AI has better margins too, and CoWoS is another growth vector and a competitive differentiator tied to it.

The capex guidance from TSMC further strengthens the semicap call I've already laid out on the memory front.

The midpoint gross margin guide of ~63% to 65% for 2026 is the real tell here. It suggests any dilution from overseas fabs is being more than offset by productivity gains, plus the pricing power they can command on advanced nodes, especially the AI accelerator mix.

On their last earnings call, they guided for ~30% growth YoY, and these capex raises should benefit the AI part of the business (high-teens %). On the other hand, cellphones (~32%), laptops, and PCs still represent a big part of the mix too.

Overall, I think the setup is positive, with the company raising the long-term gross margin floor by 3pp.


However it is worth monitoring the electronics market to see if it worsens. Apple is a meaningful part of mobile and is somewhat protected, as we discussed. On the lower end, we still need to see if weakness leads to canceled wafer bookings. For now, there are no signs of that in fab utilization.

Broadcom (AVGO)

Google is still the firm's lead AI customer, and it has historically relied on AVGO to manufacture its TPUs. The nuance now is that Google is actively building redundancy, with MediaTek increasingly discussed as a second-source for TPU variants. That does not mean Broadcom gets replaced in 2026, but it does matter for the long-term pricing and share narrative.

Broadcom has flagged the risk of hyperscalers doing more of the design work in house. In theory, they would still need a semiconductor partner for back end integration, packaging coordination, and managing the foundry relationship, so I see this as a tail risk rather than the base case for now. Still, it does give Big Tech more leverage.

Among the companies shown in the first table, Google is increasing capex the most in percentage terms versus 2025. It's also the one where analysts have revised estimates up the most versus a month ago, and Broadcom's stock didn't really react to that.



Increase on Capex from 2025 (AMZN, GOOGL, and META reflect company-provided midpoint guidance, while MSFT and ORCL are S&P Capital IQconsensus averages)

Price action in Broadcom has surprised me negatively lately, and I think 2026 growth is still backed by Google's capex plans. Also, the TPU is no longer a Google-only product, and we now expect meaningful orders from Anthropic and Meta for these chips and systems.

NVIDIA (NVDA)

NVDA stock hasn't really moved much over the last six months. Their sales are still heavily dependent on Big Tech spending, with these large customers accounting for more than 50% of data center revenue, and data center representing roughly ~90% of total revenue.

They've also been caught up in the OpenAI investment discussion which is clearly weighing on price action.

Still, the capex plans from their major customers are, of course, bullish for the year. It's worth noting that this trend could actually make Nvidia even more dependent on Big Tech over time.

There's always the risk of customers moving away, but that's a very hard thing to execute. At CES 2026, they marketed Vera Rubin as delivering roughly 5x the performance of current-generation Blackwell for certain inference workloads. And it's not just the GPU. Nvidia is also layering in a host of storage, optical, and networking upgrades at the rack level. As long as Nvidia keeps shipping meaningful system-level improvements, they should continue to capture the biggest share of that capex.

Conclusion

Big Tech capex continues to accelerate, and the mix is exactly what semis want: more compute, more HBM, and more advanced packaging. That should keep demand tight for memory and leading-edge foundry capacity, which also strengthens the case for semi-cap equipment. On the other hand, consumer electronics should feel the impact of an AI-driven memory cycle as DRAM supply and capex get pulled toward servers, tightening availability and crowding out spend elsewhere.

Editor's Note: This article discusses one or more securities that do not trade on a major U.S. exchange. Please be aware of the risks associated with these stocks.

'Impossible': Taiwan pushes back against Washington's 40% chip supply relocation goal

Feb 9, 2026 Dylan Butts

Key Points

- In an interview, Taiwan's top trade negotiator described U.S. calls to relocate 40% of the island's semiconductor production capacity stateside as "impossible."
- The statements push back on recent onshoring goals laid out by the Commerce Secretary Howard Lutnick earlier this year.



Sopa Images | Lightrocket | Getty Images

Taiwan has told Washington that its proposal to move 40% of the island's semiconductor supply chain to the U.S. was "impossible," Taipei's top tariff trade negotiator said in an interview.

Speaking on a local television broadcast Sunday, Vice Premier Cheng Li-chiun said she had made it clear to Washington that the country's semiconductor ecosystem, built over decades, could not simply be relocated.

Taiwan's international expansion, including its investments in the U.S., is predicated on the notion that the industry remains' rooted in Taiwan and continues to expand domestic investments, she said in Mandarin, translated by CNBC.

The comments push back against onshoring targets outlined by U.S. Commerce Secretary Howard Lutnick in a CNBC interview in January, shortly after the latest U.S.-Taiwan trade agreement. Lutnick said he wanted 40% of Taiwan's entire chip supply chain to shift to the U.S. within President Donald Trump's ongoing term.

Under the deal, the Taiwanese government promised \$250 billion in direct investments by its tech companies, with an additional \$250 billion in credit provided for them to expand their production capacity in the U.S.

Washington on its part lowered levies on most goods from Taiwan to 15% from 20%, waived tariffs on generic drugs and ingredients, aircraft components and natural resources unavailable domestically, and promised higher quotas for tariff-free exports of Taiwanese chips to the U.S.

Taiwan Semiconductor Manufacturing Co, the world's leading contract chipmaker and producer of the most advanced semiconductors, has already been working to better align with U.S. policy interests.

The company has committed more than \$65 billion to U.S. manufacturing in recent years, with plans to expand that to \$165 billion, as it produces chips for American clients Apple and Nvidia. The investments have also leveraged grants under the U.S. CHIPS and Science Act.

But according to Lutnick, Washington is also looking for hundreds of other smaller companies in the chip supply chain to come to the U.S.

"We're going to build giant semiconductor industrial parks in America ... This is a \$500 billion down payment on let's bring those semiconductors home," he said in January, adding that Taiwan-based chip companies that don't build in the U.S are likely to face a 100% tariff Trump has threatened against the sector.

However, semiconductor analysts broadly agree with Cheng's assessment that Washington's most ambitious onshoring plans are unfeasible, citing the difficulties of relocating such an advanced supply chain.

Analysts and industry officials point to Taiwan's deeply integrated semiconductor ecosystem, U.S. labor shortages and elevated costs as some of the key obstacles.

Geopolitical analysts have also pointed to the so-called "Silicon Shield" theory, which posits that the island's pivotal role in global chip supply makes safeguarding its autonomy a U.S. strategic imperative, deterring a potential Chinese aggression. Beijing claims sovereignty over the democratically governed island.

This Silicon Shield could further discourage Taiwan from shifting its supply chains abroad.

Taiwanese authorities have already implemented a policy requiring TSMC's overseas plants to operate using technologies at least two generations behind the cutting-edge ones being deployed in Taiwan, a policy often referred to as the N-2 rule.

The U.S. Commerce Department did not immediately respond to a request for comment on Cheng's statement.

TSMC shares were trading up 2.75% in Taiwan on Tuesday.

— CNBC's Matthew Chin contributed to this report.

Amkor projects over 20% computing growth for 2026 while advancing Arizona expansion and HDFO ramp

Feb 9, 2026 AI generated call summary

Management View

- Kevin Engel, in his first remarks as CEO, emphasized "how honored I am to speak with you for the first time as the CEO of Amkor" and outlined a leadership approach grounded in transparency, disciplined execution, and a strong customer focus. Engel highlighted the company's "record advanced and computing revenue, driven by deep customer engagements across AI and HPC," and noted the successful ramp of Amkor's first High-Density Fan-Out (HDFO) programs into high-volume production, positioning for strong 2026 tailwinds.
- Engel stated Amkor's three strategic pillars: elevate technology leadership, expand geographic footprint, and enhance strategic partnerships in focus markets. He announced operational milestones, including Vietnam reaching breakeven in Q4 and the start of construction on the Arizona campus, with Phase 1 underway.

- Investment will continue in advanced packaging, including HDFO and test, with two additional HDFO programs in final qualification for AI data centers and the majority of 2026 equipment investments focused on HDFO and test. "As AI and HPC demand continues, we are well positioned to enable the next wave of advanced products," Engel said.
- Engel detailed customer commitments supporting capacity investment, with an HDFO CPU device ramping in 2026, and stressed margin improvement driven by operational excellence, optimization in Japan, efficiencies in Vietnam, and a shift toward high-value advanced packaging.
- CFO Megan Faust reported, "Fourth quarter revenue was \$1.89 billion, down 5% sequentially and up 16% year-on-year." Faust highlighted strong year-on-year revenue growth in communications, computing, and automotive/industrial, as well as breakeven achieved in Vietnam. She noted, "Gross profit for the quarter was \$315 million, which included a benefit of approximately \$30 million from asset sales." Faust also confirmed capital expenditures for 2025 were \$905 million, below guidance due to timing of Arizona facility payments, with investments shifting into 2026.

Outlook

- For Q1 2026, management expects revenue between \$1.6 billion and \$1.7 billion, marking a 25% year-on-year increase at the midpoint. Gross margin is projected to be between 12.5% and 13.5%, with operating expenses around \$135 million and a full-year effective tax rate of about 20%. Net income is forecasted between \$45 million and \$70 million, resulting in EPS between \$0.18 and \$0.28.
- For the full year 2026, Engel stated, "we expect revenue growth to be driven by continued acceleration in computing with expected to grow over 20% and continued strong growth in advanced automotive. The remainder of our business is expected to grow in the single digits."
- CapEx for 2026 is expected to be \$2.5 billion to \$3 billion, with 65%-70% allocated to facility expansion, including the Arizona campus, and 30%-35% to HDFO, test, and advanced packaging capacity.

Financial Results

- Q4 revenue was \$1.89 billion and EPS was \$0.69. Gross margin was 16.7%. Operating income was \$185 million, operating income margin was 9.8%, and net income was \$172 million. EBITDA was \$369 million, with an EBITDA margin of 19.5%.

- For the full year 2025, revenue reached \$6.7 billion, gross profit was \$939 million, and gross margin was 14%. Net income was \$374 million, with EPS of \$1.50. Free cash flow was \$308 million. Year-end cash and short-term investments totaled \$2 billion, with total liquidity of \$3 billion, and total debt of \$1.4 billion.

Q&A

- Yu Shi, Needham & Company: Questioned the front-loaded CapEx for Arizona and customer commitments. CEO Engel explained, "between this year and the first half of next year, you get an idea of that capital spending for the Phase 1 build-out," and noted strong demand for advanced packaging as the reason for the increased equipment investment. Engel added, "that just highlights the strong demand that we're seeing in this advanced packaging area."
- Shi followed up on CapEx and government incentives. CFO Faust stated, "those government incentives as well as the investment tax credits as well as the grants, those are going to come in on a lag," and clarified that 2026 guidance includes minimal offsets from subsidies.
- Craig Ellis, B. Riley Securities: Asked about computing growth and HDFO data center program ramps. Engel responded, "when we look at the 2.5D and HDFO platforms, we're expecting that to nearly triple over the course of this year." One device will be in high volume, with the other also ramping.
- Steve Barger, KeyBanc: Asked about recovering gross margin headwinds from Vietnam. Faust stated, "I want to reiterate Kevin's remarks that we did have, I would say, breakeven in Q4, which was a significant milestone for us in Vietnam, and we see that continuing in Q1."

Sentiment Analysis

- Analyst questions focused on CapEx timing, funding sources, and the sustainability of margin improvements, with a generally neutral tone but some skepticism on capital allocation and margin progression.
- Management maintained a confident and detailed tone in prepared remarks and in responses, frequently referencing customer commitments and operational milestones. Engel used phrases such as "we are well positioned" and "I'm confident in our ability to execute."
- Compared to the previous quarter, management's tone shifted from transition and continuity under outgoing CEO Rutten to forward-looking confidence and operational detail from new CEO Engel.

Analysts remained probing on capital intensity and margin drivers but were less focused on strategic risk.

Quarter-over-Quarter Comparison

- The most notable change is the CEO transition from Giel Rutten to Kevin Engel, marking a new leadership phase. The strategic focus remains on advanced packaging and geographic expansion, but management now provides more explicit forward guidance for 2026, including computing growth over 20% and CapEx plans.
- CapEx guidance increased substantially for 2026, with clear allocation toward Arizona and HDFO expansion. Margin discussion shifted from managing costs in Japan and Vietnam ramp-up to margin improvement from operational leverage, with Vietnam breakeven and the Arizona facility underway.
- Analysts in both quarters focused on CapEx, margin drivers, and the impact of strategic investments, but the current quarter's Q&A had heavier emphasis on the timing and structure of government incentives and customer commitments.

Risks and Concerns

- Management cited continued monitoring of export control and trade policies as well as substrate, advanced silicon, and memory supply dynamics as factors considered in Q1 guidance.
- Faust discussed the lag in receiving government subsidy benefits, which may impact near-term cash flows despite front-loaded CapEx.
- Engel acknowledged space and labor as potential constraints, specifically on R&D and facility expansion, and outlined mitigation by shifting SiP production to Vietnam and increasing cleanroom space in Korea.

Final Takeaway

- Amkor enters 2026 with strong momentum supported by new leadership, a record year in advanced packaging and computing, and aggressive investment in capacity and geographic expansion. The company expects over 20% growth in computing, continued strong automotive growth, and significant CapEx to support these initiatives. Strategic customer partnerships, margin improvement from

operational efficiency, and targeted capital allocation underpin Amkor's outlook as it positions for the next wave of AI and high-performance computing demand.

ASE Technology targets \$3.2B LEAP revenue in 2026 driven by AI demand and advanced packaging

Feb 5, 2026 AI generated Q4, 2025 call summary

Management View

- Group COO Tien Wu outlined a multiyear outlook, emphasizing that "the AI server cycle continues, primarily led by hyperscaler and the data center development." Wu reported that mainstream business recovery is underway and stated, "We believe the mainstream, namely the IoT, the automotive, the general sector, the mainstream business will recover better this year comparing to last year."
- Wu highlighted manufacturing leadership in Taiwan, noting that "the Taiwan cluster has demonstrated the best efficiency and speed in terms of manufacturing ramp-up." He said ASE is expanding outside Taiwan, especially in Penang for automotive and robotics, and building footprints in Korea and the Philippines.
- Wu recapped 2025 with "consolidated revenue grew 12% at a holdco level with ATM revenue up 23% led by Leading Edge Advanced Packaging services and Testing business." LEAP services reached \$1.6B, accounting for 13% of ATM revenue, and testing business grew 36% year-on-year. Machinery CapEx totaled \$3.4B and building/facilities CapEx was \$2.1B, largely for LEAP services and testing.
- Wu said, "We expect revenue uptrend to continue 2026 and beyond, driven by leading-edge solutions and broad-based semiconductor demand related to AI proliferation and general market recovery. ATM business, leading-edge assembly packaging service to double from USD 1.6 billion to USD 3.2 billion."
- Head of Investor Relations Ken Hsiang stated, "Our ATM factories in Taiwan ran at or near full capacity with LEAP and traditional advanced packaging utilization rates outpacing that of wirebond."
- Group CFO Joseph Tung provided guidance: "First quarter revenue should decline only by 5% to 7% quarter-over-quarter. First quarter gross margin should decline by 50 basis points to 100 basis points quarter-over-quarter."

Outlook

- Joseph Tung stated, "For ATM, as Tien mentioned, we expect 2026 leading-edge revenue to at least double compared with last year, while demand continues to significantly exceed supply."
- Tung added, "We expect ATM gross margins to stay within our structural margin range throughout the year and to improve every quarter, while second half gross margin to reach the upper end of the range."
- CapEx for 2026 will remain aggressive, with "another USD 1.5 billion in machinery on top of last year's USD 3.4 billion, of which about 2/3 will be for leading-edge services."

Financial Results

- For Q4 2025, ASE reported fully diluted EPS of \$3.24 and basic EPS of \$3.37. Consolidated net revenues were \$177.9B, up 6% sequentially and 10% year-over-year. Gross profit was \$34.7B with a

gross margin of 19.5%, improving 2.4 percentage points sequentially and 3.1 percentage points year-over-year.

- Operating profit was \$17.7B, up \$4.5B sequentially and \$6.5B year-over-year. Operating margin was 9.9%, up 2.1 percentage points sequentially. Net income for the quarter was \$14.7B.
- ATM business posted record revenues of \$109.7B, with test revenues growing 13% sequentially and 33% annually. ATM gross profit margin reached 26.3%, up 3.7 percentage points sequentially. EMS business revenues were flat at \$69B, with gross margin declining 0.2 percentage points to 9%.
- Full year 2025, consolidated net revenues improved 8% versus 2024, ATM business improved by 20%, and EMS business declined by 5%. Net income for the year increased by 25% to \$40.7B.

Q&A

- Sunny Lin, UBS: Asked about breakdown and growth potential of the LEAP business. Joseph Tung responded, "We expect to triple our full process revenue this year to reach about 10% of the overall LEAP service revenue."
- Lin also inquired about margin expansion and long-term LEAP gross margin potential. Tien Wu replied, "We would like to take 1 year at a time. So we'll talk about this next year."
- Haas Liu, BofA: Asked about mainstream segment growth and LEAP business upside. Wu explained, "The pricing environment, I would say, friendly... our best view is we believe we can achieve \$3.2 billion comfortably."
- Charlie Chan, Morgan Stanley: Asked about the impact of the EugenLight acquisition on the CPO business. Wu stated, "The optical business is an important direction... USI is trying to piece together early deployment in terms of the future optical road map, and this is part of it."
- Rick Hsu, Daiwa: Asked if EMS is being downsized. Wu clarified, "No, we're not trying to downsize, but the market is shifting... we just have to change accordingly."
- Gokul Hariharan, JPMorgan: Asked about CapEx allocation and full process packaging. Wu confirmed, "We're looking at about 10% full process revenue by the end of this year."

Sentiment Analysis

- Analysts pressed for specifics on LEAP revenue projections, margin expansion, and CapEx discipline, displaying a slightly positive but probing tone.
- Management showed confidence in growth and technology investments but remained hesitant to provide multi-year margin targets, frequently emphasizing a one-year-at-a-time approach. Wu said, "We would like to take 1 year at a time."
- Compared to the previous quarter, analyst tone shifted from broad recovery questions to more granular focus on advanced packaging, CapEx, and margin structure. Management's tone stayed confident but was more explicit about capacity constraints and CapEx pressures.

Quarter-over-Quarter Comparison

- Guidance for LEAP-related revenue was raised from \$1.6B to \$3.2B for the next year, compared to the previous quarter's expectation of \$2.6B.
- Gross margin and operating margin improved both sequentially and year-over-year, with ATM business driving the margin gains.

- CapEx intensity and facility expansion plans are more aggressive, with increased focus on Penang and other non-Taiwan sites. Strategic focus shifted to executing on global customer requirements and accelerating advanced packaging.
- Analyst questions became more focused on the sustainability and scalability of leading-edge business and capital allocation.

Risks and Concerns

- Management highlighted supply constraints and capacity ramp-up challenges, particularly in LEAP and advanced packaging.
- Analysts voiced concerns about EMS business stagnation, future CapEx discipline, and the pace of transition in EMS toward AI and industrial applications.
- Wu acknowledged, "Processes, equipment lead time and all of the management and also the engineering training, they're not easy."

Final Takeaway

- ASE Technology Holding Co., Ltd. signaled strong momentum in leading-edge packaging and testing, with LEAP revenue expected to double to \$3.2B in 2026 and continuous margin improvement anticipated. Management sees ongoing demand outpacing supply, aggressive CapEx plans, and a strategic focus on expanding both capacity and geographic footprint to capture AI-driven growth. The company remains committed to operational discipline while balancing resource constraints and customer needs in a rapidly evolving semiconductor environment.

AI Memory Boom Squeezes Legacy DRAM Supply, Pushing Prices Higher

Feb 3, 2026 Markit

Summary

- High-bandwidth memory, or HBM, has become one of the most lucrative niches in semiconductors, benefiting from explosive demand linked to AI model training and inference.
- As capacity tilts toward HBM, traditional DRAM is experiencing a supply shortage, driving up prices.
- Visible Alpha consensus estimates point to a steep acceleration in average selling prices for conventional DRAM in 2026.




Alexander Sikov/iStock via Getty Images

Traditional DRAM vs. HBM ASP comparison among leading memory manufacturers

Traditional DRAM ASPs (\$)

HBM DRAM ASPs (\$)



Data compiled Jan. 28, 2026.

DRAM = dynamic random access memory; HBM = high-bandwidth memory; ASPs = average selling prices. Projections for 2026 and beyond are projected estimates.

Source: Visible Alpha, a part of S&P Global Market Intelligence. © 2026 S&P Global.

As leading memory semiconductor manufacturers, Samsung Electronics Co. Ltd. (SSNLF), SK hynix Inc. (HXSC.F), and Micron Technology Inc. (MU), divert production capacity toward high-bandwidth memory, a premium form of DRAM used in AI data centers, supplies of conventional DRAM are tightening, driving prices higher.

High-bandwidth memory, or HBM, has become one of the most lucrative niches in semiconductors, benefiting from explosive demand linked to AI model training and inference.

The problem for enterprise IT buyers is that the same production lines also make standard DRAM, which supports servers, PCs, and a wide range of consumer electronics. As capacity tilts toward HBM, traditional DRAM is experiencing a supply shortage, driving up prices.

Visible Alpha consensus estimates point to a steep acceleration in average selling prices for conventional DRAM in 2026. Samsung's revenue per bit from traditional DRAM is forecast to rise 116% year-on-year to \$0.79 in 2026, up from \$0.36 last year. SK Hynix is expected to see a 78% increase to \$0.70 from \$0.40, while Micron's traditional DRAM ASP is projected to climb 54% to \$1.06.

By contrast, pricing gains in HBM look far more restrained, reflecting both rapid capacity expansion and the already elevated starting point. Analysts expect HBM ASPs to rise by about 8% at Samsung, 1% at SK Hynix, and 22% at Micron in 2026.

The imbalance reflects clear economic incentives. HBM carries meaningfully higher margins than conventional DRAM, encouraging all manufacturers to prioritize AI-linked products over legacy enterprise and consumer memory. That strategy supports profitability, but risks pushing up costs for corporate IT budgets at a time when data center investment is already under pressure.

Micron has moved to ease near-term constraints. The US chipmaker recently struck a strategic partnership with Taiwan's Powerchip Semiconductor Manufacturing to access additional DRAM capacity, while also stepping up investment in new fabrication plants in the US. Those facilities, however, will not come on stream for several years. In the meantime, Micron is relying on existing clean-room infrastructure in Taiwan to lift output.

AMD Vs. Intel: AMD Takes The Lead In 2026

Feb 2, 2026 Khaveen Investment

Summary

- Intel's CPU performance in desktop and notebook computers outperforms AMD, but server performance is where AMD's huge advantage really shows.
- AMD's data center expansion, robust product roadmap, and fabless model drive superior revenue growth and margin profile versus INTC.

- AMD's diversified portfolio and market share gains contrast with INTC's concentrated, capital-intensive strategy and ongoing restructuring.



JHVEPhoto/iStock Editorial via

Getty Images

Since our previous comparison of both AMD (AMD) and Intel (INTC) back a year ago, we believed both companies were undervalued, with AMD having significant strengths over Intel, but at the same time, our valuations indicated Intel was oversold. Both companies' stock prices have since surged, meeting our price targets and continuing to rise above them. Intel's stock is up by a whopping 132% over the past 1 year, while AMD is up 99%.

In this analysis, we compare the two companies again to determine whether these companies are still attractive and investable at current levels given the strong run-up of their stock prices. We analyze both companies with a direct comparison in terms of their revenue performance and growth outlook, in terms of specific drivers, which focused on AMD's continued data center growth and CPU strengths, and Intel's strategy revolving around its foundry expansion and process roadmap execution. We then provide a contrast of both companies' strategy execution and direction, as well as a comparison of the two companies' financials, highlighting the strength of AMD's fabless model over Intel's capital-intensive approach, especially for its foundry ambitions. We conclude with a comparison of both companies with updated valuations based on DCF.

Revenue Growth Performance and Outlook Comparison

In this section, we will compare both companies' revenue segments and performance to see which company is better and has better growth prospects.

Performance Comparison

Intel Segment Growth %	Intel (FY2025)	Intel (5-Year Average)	AMD Segment Growth %	AMD (Q3 2025 YTD)	AMD (5-Year Average)
CCG	6.4%	-3.9%	Data Center	29.1%	71.0%
DCAI	32.0%	-4.6%	Client	59.1%	12.4%
Mobileye	11.8%	16.6%	Gaming	50.9%	14.7%
Intel Foundry (External)	107.4%	-7.2%	Embedded	-4.9%	451.5%
Intel Others	-20.8%	67.3%			
Intel Total Growth %	-0.1%	-7.0%	AMD Total Growth %	34.4%	30.4%

Khaveen Investments

We compiled the revenue segment breakdown summary above, for AMD. Based on its Q3 2025 YTD results, its AMD revenue growth has been impressive (34.4% YoY) and in line with its 5-year average. Its strong growth is mainly driven by its new Ryzen AI CPUs, while its gaming GPUs also recovered strongly from a weak prior year. Data Center revenue growth was strong too but weaker than its 5-year average, while Embedded revenue was particularly weak, likely attributable to weak industrial demand.

Intel's average growth performance, on the other hand, has been much weaker than AMD's in almost all accounts. In 2025, Intel's revenue was flattish (-0.1% YoY), and it had reorganized its NEX segment revenue, which was previously broken down separately but has now been included as part of CCG and DCAI segments from 2025 onwards, thus the growth rates seen in the table reflect the NEX segment inclusion. Excluding the NEX reorganization effect, its DCAI revenue growth was only 4.9%, supported by the Granite Rapids ramp, but still much weaker than AMD's data center growth. The CCG segment's actual growth without the NEX reorganization effect was actually negative (-3.4%) despite the new launch of Lunar Lake PC CPUs, due to "wafer fabrication supply constraint" according to its annual report, and also highlights the competition from AMD which had positive growth. Mobileye was really the only noteworthy segment with stronger performance in 2025, a recovery from the previous year, as well as its IFS foundry revenue with a strong 107% YoY growth, but it is only minimal (0.58% of total revenue). In terms of overall growth performance, AMD has a clearly much stronger and more consistent growth track record than Intel in the past 5 years.


Revenue Exposure Comparison

Category (\$ mln)	AMD	% of AMD Revenue	Our 3-yr Growth Forecasts	Intel	% of Intel Revenue	Our 3-yr Growth Forecasts
PC CPUs	7,054	27.36%	16.09%	29,453	55.73%	2.86%
PC GPUs	1,044	4.05%	11.72%	837	1.58%	2.86%
Data Center CPUs	7,479	29.01%	35.79%	12,317	23.30%	8.00%
Data Center GPUs	5,000	19.39%	72.27%	500	0.95%	8.00%
Data Center DPUs	100	0.39%	18.90%	0	0.00%	
Gaming Console SoCs	1,551	6.02%	5.30%	0	0.00%	
Foundry Services	0	0.00%		300	0.57%	172.40%
All Others	3,557	13.79%	10.50%	9,446	17.87%	13.52%
Total Revenue	25,785	100%	34.8%	52,853	100%	5.3%


Khaveen Investments

AMD vs Intel Revenue Breakdown

AMD Revenue Breakdown



Intel Revenue Breakdown



■ PC CPUs ■ PC GPUs ■ Data Center CPUs ■ Data Center GPUs ■ Data Center DPUs ■ Gaming Console SoCs ■ Foundry Services ■ All Others

www.khaveen.com
Copyright 2026. All Rights Reserved


Source: Company Data, Khaveen Investments

Khaveen Investments

We compiled above the revenue breakdown that shows the revenue concentration of the companies based on their main product categories that we calculated using available company segment data and estimates. As seen, AMD is overall more diversified than Intel in terms of its revenue mix. Intel's revenues are largely concentrated in the CPU market, with PC CPUs as its largest end market (55.47% of revenue), followed by servers (23.20% of revenue), and this is negative because Intel continues to face competitive struggles against AMD and also ARM-based alternatives, while its faster-growth segments like Mobileye are only a small portion of its revenues, and its foundry business, which it is mainly focusing on to drive growth, is minuscule (0.58% of revenue).


Even though we have factored in new developments such as Intel's secured contract wins with Amazon, Microsoft, and the U.S. military (172.40% 3-year CAGR), it is too small to give a meaningful contribution to Intel's overall growth in the next three years. Whereas AMD is not only more diversified with CPUs, GPUs, and other SoCs and FPGAs, but its growth performance has also trounced Intel, and our three-year forward forecasts indicate strong continued growth, supported by data centers, where AMD highlighted it expects roughly 80% CAGR in data center accelerator revenue, consistent with our forecast of 72.27% CAGR for data center GPUs.

Market Share Comparison




Mercury Research, Khaveen investments

From 2018 to Q3 2025, AMD’s x86 CPU market share grew impressively, whereas Intel's worsened considerably. We previously expected AMD’s gain could happen due to AMD’s stronger competitive position, earlier product launches, and better execution. For example, AMD’s Ryzen AI 300 series was released much quicker than Intel’s, while also offering better performance-per-dollar. AMD’s advanced execution using TSMC’s leading nodes has also supported AMD’s market share.



JPR, IDC Khaveen Investments

In the desktop discrete GPU market, NVIDIA's (NVDA) share had dominantly improved to 94%, whereas AMD's share declined unexpectedly to 6% but was still better than Intel, which has fallen out of the market share ranking. NVIDIA's strong growth is attributable to high demand for its RTX 50-series GPUs, whereas AMD's RX 9000 witnessed lower demand. NVIDIA has also partnered with and invested \$5 billion in Intel to build custom CPUs and an x86 SoC, which might benefit Intel if they developed integrated chips as planned. AMD, despite being down, is not giving up in the GPU space, however, as it released its new Radeon series last year based on its next-gen RDNA 4 architecture.



Khaveen Investments

For the FPGA market, in Q3 2025, AMD steadily improved its lead in the FPGA market while Intel lost its market share. Intel then fully exited the market when it divested its Altera business recently. For other competitors, microchip share was relatively stable, whereas Lattice saw a slight improvement. We see that AMD has been outperforming Intel, as its CPU share has been gaining against Intel across all segments. We also found that AMD's GPUs still hold a notable market position compared to Intel's. This is even with NVIDIA dominating. We also see that AMD has steadily widened its lead over Intel.

Revenue Growth Performance and Outlook	Winner
Growth Performance	AMD
Revenue Exposure	AMD
Market Share	AMD
Overall	AMD


Khaveen Investments

Based on these areas that we compared, we see AMD trouncing over Intel with a superior revenue growth track record, and we expect it to continue being strong based on our 3-year forward revenue forecast averaging 34.8% vs. just 8.1% for Intel, due to its tailwinds as it focuses on its data center expansion strategies, which is an area that Intel has struggled with. We also see AMD continuing to carve out market share gains across the PC and data center CPU markets, which Intel is much more reliant on based on its revenue mix.

Management Strategy and Execution

In this section, we compare the strategies of AMD and Intel, particularly for CPU product launches. This is because our earlier analysis showed that these companies' CPU performance has been steadily improving. For this analysis, we assess how effective their execution is and whether launches remained on track. We then analyze their data center expansion strategies and the reason for AMD's larger major cloud customer base. We would then evaluate their manufacturing strategies, as Intel has pursued IDM 2.0 for its internal advanced nodes, while AMD maintained a fabless model built on TSMC's nodes. We would then examine the possible challenges faced by Intel's internal foundry and what it could mean for its future competitiveness.


CPU Product Performance Strategy



Khaveen Investments

Both Intel and AMD have been highly competitive across the 3 segments. In 2024, Intel leads in desktop and laptop CPU scores due to strong performance from its Arrow Lake and Raptor Lake Refresh. Intel has also planned to launch its Core Ultra Series 3 CPU and its Arrow Lake Refresh later this year. Despite Intel's stronger performance, AMD performs better than Intel in servers. This is because its EPYC Genoa and Bergamo CPUs have achieved higher benchmarks. AMD is also planning to launch its latest EPYC Venice server CPUs soon to further solidify its position.

CPU Product Launch Roadmap Comparison

 AMD vs Intel CPU Roadmap Comparison				
Company	Metric	2024	2025	2026
AMD (Desktop)	Product Name	Ryzen 9000 (Granite Ridge)	Ryzen 9000 X3D / Threadripper 9000	Ryzen 10000 (Medusa Ridge, Zen 6)
	Launch Period	Aug 2024	Q1 2025 (X3D) / H2 2025 (TR)	Late 2026 (Paper Launch) / 2027 (Vol)
	Process Node	TSMC 4nm	TSMC 4nm	TSMC 3nm / 2nm
Intel (Desktop)	Product Name	Arrow Lake-S (Core Ultra 200S)	No Major Release	Arrow Lake Refresh (Early) / Nova Lake (Late)
	Launch Period	Oct 2024	-	H1 2026 (Refresh) / Q4 2026 (Nova)
	Process Node	TSMC N3B + 20A	-	Intel 18A (Nova) / TSMC N3 (Refresh)
AMD (Laptop)	Product Name	Ryzen AI 300 (Strix Point)	Fire Range / Ryzen AI Max (Strix Halo)	Krackan Point (Entry) / Gorgon Point
	Launch Period	Jul 2024	Early 2025	Early 2026
	Process Node	TSMC N4P	TSMC N4P / N3E	TSMC N3
Intel (Laptop)	Product Name	Lunar Lake (Core Ultra 200V)	Arrow Lake-H / HX / U	Panther Lake (Core Ultra 300)
	Launch Period	Sep 2024	Q1 2025	Jan 2026 (CES Launch)
	Process Node	TSMC N3B	TSMC N3B	Intel 18A (Compute Tile)
AMD (Server CPU)	Product Name	EPYC 9005 (Turin, Zen 5)	EPYC 9005 (Turin Dense)	EPYC Venice (Zen 6)
	Launch Period	H2 2024	H1 2025	H2 2026
	Process Node	TSMC N4P	TSMC N3E	TSMC N2
Intel (Server CPU)	Product Name	Granite Rapids-AP	Granite Rapids-SP (Mainstream)	Clearwater Forest / Diamond Rapids
	Launch Period	Q3 2024	H1 2025	H1 2026 (Clearwater) / Late 2026 (Diamond)
	Process Node	Intel 3	Intel 3	Intel 18A (Both)

www.khaveen.com
Copyright 2026. All Rights Reserved

Khaveen Investments

In desktops, AMD moved to 2 nm in 2026 from 4nm in 2024. We also believe its reliance on TSMC's mature nodes has likely improved its overall execution, as its desktop performance has improved. While AMD's manufacturing has been steady, Intel's internal manufacturing has faced disappointing setbacks. Intel had not launched any major desktop CPUs in 2025, as it had delayed its Arrow Lake-S Refresh until 2026.

For laptops, both AMD's and Intel's releases have been steady. AMD launched its Strix Point and Strix Halo for 2024 and 2025. Later this year, AMD is also expected to release Gorgon Point on schedule. Like AMD, Intel has also introduced products such as Lunar Lake and Arrow Lake on time and should likely launch its Panther Lake promptly this year.

For servers, AMD has impressively improved from EPYC Turin (4nm) in 2024 to EPYC Venice (2nm) in 2026, while Intel faced continuous launching issues. Intel had initially planned to release its Diamond Rapids and Clearwater Forest in 2025, but has now pushed it to 2026. Intel's product delays are not surprising to us. This is because of Intel's long history of delays, which we analyzed previously. We reasonably believed these setbacks had led to average revenue losses of about 0.7% from market share declines and around 1% from weaker pricing. Intel also struggles with internal process node ramps, while AMD relies fully on TSMC's proven 4 nm, 3 nm, and 2 nm nodes. If Intel can't bring 18A into stable volume production by 2026, AMD could gain a lasting edge in power efficiency and competitiveness across all segments. We believe that AMD should be better than Intel, due to AMD's more consistent product roadmap and TSMC's nodes.

Data Center Expansion Strategy

Intel still holds a lead in data center and server CPUs, but that lead has continued to weaken in 2025. This is due to AMD's better CPU performance. We previously explained that AMD CPUs had benefited from more

advanced technologies, while Intel remained behind. To combat this, Intel had initially planned an ambitious roadmap with upcoming products based on its 18A node, but had recently delayed its 18A. This makes us believe that AMD should maintain its advantage.

NVIDIA continues to dominate with over 90% market share in data center GPUs, while AMD and Intel have been slowly improving their GPUs. AMD offers the Instinct GPU lineup, while Intel competes with its Gaudi series. We believe that our analysis shows that AMD is likely ahead of Intel. This is because AMD's Instinct has more than a few times the revenue of Gaudi in 2024. To further our data center analysis, we will compare the roadmaps of AMD and Intel below.

KHAWEEN Investments		AI Accelerator Comparison			
Company	Metric	2024	2025	2026	
AMD	Product Name	Instinct MI325X	Instinct MI350X (CDNA 4)	Instinct MI400 (CDNA Next)	
	Launch Period	Q4 2024	Q2 2025	2026	
	Process Node	TSMC 5 nm	TSMC 3nm	TSMC 2 nm	
	Pricing	\$15,000	\$30,000	-	
	FP8 Performance	2.6 PFlops	4.6 PFLOPS	20.0 PFlops	
Intel	Product Name	Gaudi 3	Falcon Shores	Jaguar Shores	
	Launch Period	Q3 2024	2025 (Cancelled)	H2 2026	
	Process Node	TSMC N5 (5nm)	-	Intel 18A (1.8 nm)	
	Pricing	\$15,625*	-	-	
	FP8 Performance	1.8 PLops	-	-	

www.khaveen.com
Copyright 2026. All Rights Reserved

Khaveen Investments. * Estimate

AMD has been quickly moving from the MI325X (TSMC, 5 nm) in late 2024 to the MI400 (2 nm) by 2026. Its MI325x FP8 performance was also much higher than Intel's Gaudi 3 FP8 performance, despite both companies having the same process nodes. AMD also had no delays, whereas Intel cancelled its Falcon Shores in 2025.

We believe that AMD likely targets the higher-end segments. This is because the MI350X is priced around twice as much as Intel's Gaudi 3, while offering better F8 performance than Intel's. FP8 is critical for AI workloads, as higher F8 helps speed up inference, along with lowering power consumption. AMD's positioning is more focused towards premium customers and higher performance, whereas Intel appears targeted at lower-cost options with modest performance.

Cloud Provider Partnerships

Customer / Cloud Provider	AMD Instinct (MI300X/MI350X)	Intel Gaudi (Gaudi 3)
AWS	No	Yes
Microsoft	Yes (announced partner)	Yes
Google	Yes (announced partner)	No
Meta	Yes (announced partner)	No
Oracle Cloud Infrastructure	Yes	No
IBM Cloud	Yes	Yes

Khaveen Investments

AMD has better engagement than Intel with leading hyperscalers. This is because AMD’s AI accelerators are used by Microsoft, Google, Meta (META), Oracle (ORCL), IBM (IBM), whereas Intel only partnered with IBM Cloud, Microsoft, and AWS.

Due to AMD’s clearer roadmap, better performance chips, and higher cloud adoption, we reasonably believe that AMD is better than Intel for AI data center accelerators.

Manufacturing Strategy Comparison (Intel Foundry Strategy IDM 2.0)

Intel launched its IDM 2.0 strategy in 2021, but its foundry has not delivered any 18A chips. Its chips also have yields (10%) far below TSMC’s yields (70%). Intel’s new CEO has now refocused and plans to release its 14A node in 2027. However, without a “significant external customer” for 14A and the foundry’s high costs, Intel’s CEO indicated he may “reconsider the viability” of continuing IDM. We believe it could be too early to tell if Intel’s foundry could be successful, due to its slow foundry progress and low revenues.

While Intel focused on internal manufacturing, AMD has remained committed to its fabless model, focusing on chip design by using TSMC’s industry-leading foundry capabilities. We believe this strategy has served AMD well, as AMD has executed consistently across its CPU and AI accelerator roadmaps.

Cost Management and Leadership Stability

Since 2024, Intel has been restructuring, laying off around 15,000 employees (15% of the workforce). After Lip-Bu Tan becomes CEO in 2025, Intel intends to trim its 75,000 core employees. This is to reduce its operating expenses by \$500 million and \$1 billion for 2025 and 2026, respectively. We also believe this strategy is unique to Intel because there aren’t any similar AMD restructurings. AMD’s management has instead been focused on expanding its CPU and AI accelerator products and maintaining its fabless model.

Intel's COGS % of revenue climbed steadily from 45% TTM as weaker revenues reduced economies of scale, and depreciation increased. Key expenses such as R&D, SG&A, and R&D also rose. Comparing Intel's and AMD's expenses, Intel's R&D % is now higher than AMD's despite AMD's R&D % increase. AMD's COGS % halved from 2015 to TTM, likely because its higher-margin data center business has grown more quickly than its other segments.

In terms of expansion, Intel had cancelled its proposed fab in Germany and stopped expansions in Poland and Ohio. Intel has also consolidated its assembly and testing sites to Malaysia and Vietnam to cut capex by around \$2 billion and \$18 billion for 2025. AMD does not need to expand compared to Intel because it is fables and already running with much lower capex at only \$636 million for 2024.

As part of its cost rationalization strategy, Intel sold its 51% Altera stake to Silver Lake for \$4.5 billion, valuing Altera at \$8.8 billion. This is well below the \$16.7 billion Intel had originally paid. Intel also exited its automotive chip and camera businesses. These exits are intended to simplify its business model towards its core CPU, data center, and AI chips, while freeing up resources. For AMD, we did not come across any restructuring. Instead, AMD has been acquiring companies such as ZT Systems to strengthen its data center business.

In August 2025, the US government invested \$8.9 billion into Intel to help fund Intel's domestic supply chain expansion. This is because Intel is planning to build semiconductor factories across 4 states to make its leading-edge 18A chips, which are expected to cost over \$100 billion. The investment is also to make sure that the U.S. can manufacture chips domestically. In Intel's Q4 FY2025 earnings transcript, the company also announced that its external foundry revenues for the quarter were "driven by projects with the U.S. government."

Strategy and Execution	Winner
CPU Product Launch and Performance	AMD
Data Center Expansion	AMD
Manufacturing and Expansions	Intel
Cost Management and Leadership Stability	AMD
Overall	AMD



Khaveen Investmetns

In terms of the comparison of management strategies and execution, we believe AMD is undoubtedly ahead of Intel. While we believe Intel's efforts are positive for the long-term survival of the company, and some signs indicate Intel has started to get some momentum in its foundry business with the ramp-up of its Intel 18A node and expectations of Amazon, Microsoft, and the US military contracting its foundry business and finally providing a jumpstart to its external foundry revenue, challenges persist, and whether it will encounter any setbacks to its expansions remains too early to tell. On the other hand, AMD sits comfortably, benefiting from long-term strategies that it has long committed to by enhancing its product development, and has a better track record with its roadmap, leveraging TSMC as a partner.

Financials and Margins Comparison

In this section, we assess which company has the stronger financial profile by analyzing and comparing profitability metrics, including gross, operating, net, and free cash flow margins.

Profitability Margins




Margins	Intel (TTM)	Intel (5-year Average)	AMD (TTM)	AMD (5-year Average)
Gross Margin	36.56%	43.21%	51.50%	33.40%
EBITDA Margin	23.90%	26.04%	19.30%	49.40%
EBIT Margin	1.75%	6.66%	9.50%	20.80%
EBT Margin	2.95%	4.60%	9.30%	10.30%
Net Margin	-0.51%	1.03%	10.30%	10.10%
Free Cash Flow Margin (Total Investing Cash Flow)	-9.57%	-7.64%	3.60%	12.10%
Free Cash Flow Margin (Capex Only)	-9.24%	-13.98%	16.80%	10.10%

Khaveen Investments

The table above compares Intel and AMD across gross, EBITDA, EBIT, EBT, net, and free cash flow margins on both a TTM and 5-year average basis. AMD outperforms Intel in nearly every metric. On a TTM basis, AMD's gross margin is 51.5% compared to Intel's 36.6%, and AMD also posts stronger EBIT, EBT, and net margins compared to Intel's EBIT margin, which is 1.8%, and net margin, which is -0.5%.

Free cash flow margins also favor AMD, with 16.8% (capex basis) compared to Intel's -9.2%. Even over the 5-year average, AMD leads in most categories, despite Intel's historically higher gross margin from its vertically integrated model. Intel's recent margin weakness includes \$2.2 billion in restructuring charges, \$505 million in acquisition-related costs, \$10.3 billion of operating losses in its foundry operations, and execution issues in process technology ramp-ups like 18A. Its FCF margins are also weighed down by high capex, as its 2025 capex was \$14.6 billion. In comparison, AMD's 2024 capex was much lower at \$636 million, highlighting its lean, fabless model. In our previous analysis, we also highlighted that Intel's profitability has been under pressure from product delays and underutilization in manufacturing, while AMD's fabless model supports its lower capital intensity and stronger profitability. In Intel's latest earnings transcript, management continued to point at operating losses of its foundry business. We thus believe AMD is better due to stronger margins, lower capital intensity, and consistent execution compared to Intel's margin pressure and high capex.

Financial Position (Cash-to-Debt)



Cash-to-Debt	Intel (TTM)	Intel (5-year Average)	AMD (TTM)	AMD (5-year Average)
Debt	47,109	45,501	3,870	1,876
Cash & Cash Equivalents	37,416	28,421	7,243	4,532
Net Debt (Net Cash)	9,693	17,081	-3,373	-2,656
Cash to Debt Ratio	0.8x	0.6x	1.9x	3.2x

Khaveen Investments

Intel's cash-to-debt ratio is weaker than AMD's at 0.8x on a TTM basis compared to AMD's 1.9x. This reflects Intel's higher total debt of \$47.1 billion relative to its \$37.4 billion in cash and equivalents. Over the 5-year average, Intel's ratio is 0.6x but still trails AMD's much stronger 3.2x, supported by AMD's lower debt levels (\$3.9 billion TTM) and higher relative cash balance (\$7.2 billion TTM). We believe this overall highlights AMD's stronger liquidity position and lower leverage risk compared to Intel.

Valuation Ratios

Valuation Comparison	EV/EBITDA (5Y Avg)	P/S (5Y Avg)	P/E (5Y Avg)	P/FCF (5Y Avg)
US Chipmaker Average	31.51	9.47	42.06	43.77
INTC	12.52	2.41	12.36	6.58
AMD	46.94	8.67	44.64	89.95

Khaveen Investments

Except for P/S, AMD's valuation multiples are far higher than Intel's and above the US chipmaker industry average. AMD's EV/EBITDA and P/S are over 3 times higher than Intel's. AMD's P/E also trades much higher than Intel's and is slightly above the industry's average. P/FCF shows the widest gap, as AMD is 15 times larger

than Intel and 2 times larger than the industry. We believe Intel's lower reflections likely reflect better value and lower growth expectations.

Valuation Comparison	EV/EBITDA (5Y Avg)	P/S (5Y Avg)	P/E (5Y Avg)	P/FCF (5Y Avg)
US Chipmaker Average	31.51	9.47	42.06	43.77
INTC	12.52	2.41	12.36	6.58
AMD	46.94	8.67	44.64	89.95

Khaveen Investments

AMD is better than Intel in terms of both profitability and financial position. This is due to AMD's stronger margins as well as a lean fabless model with lower capital intensity. AMD also maintains a stronger liquidity profile, with a higher cash-to-debt ratio supported by lower debt and higher relative cash balances. In contrast, Intel's margins have been pressured by high restructuring and impairment charges. Losses in its foundry operations, product delays, and heavy capex requirements have also lowered profitability. Intel's valuation multiples are lower than both AMD's and the US chipmaker industry average. If Intel could address its strategy, its lower valuation could be appealing to investors seeking value.

Risk: End Market Impact Due to Rising Memory Prices

One of the key issues that could affect the end market outlook growth for both AMD and Intel in the PC and possibly the data center market is the rising memory prices, as DRAM and NAND chip prices have been surging due to shortages. According to the IDC, they project up to a 9% decline in the PC market this year due to the higher prices. This could affect both companies indirectly as customers put on hold CPU and GPU purchases. Whereas in the data center market, companies have to reduce purchases of data center chips to offset the rising memory costs unless they raise capex. AMD and Intel may also have to cut production of GPUs in the PC market, as these chips feature RAM, due to higher costs, which may impact sales.

Valuation

AMD Valuation

Discounted Cash Flow Model												
Calendar Year	Entry	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	Terminal
Financial Year		31/12/2025	31/12/2026	31/12/2027	31/12/2028	31/12/2029	31/12/2030	31/12/2031	31/12/2032	31/12/2033	31/12/2034	31/12/2034
Revenue Growth		34.80%	34.66%	34.62%	35.23%	28.19%	22.53%	18.04%	14.43%	11.55%	9.24%	
Revenue		34,757	46,805	63,011	85,212	109,231	133,862	158,010	180,814	201,689	220,318	
EBIT		3,439	5,837	9,181	13,694	18,555	23,529	28,400	33,002	37,219	40,989	
Less: Cash Taxes		731	1,240	1,943	2,890	3,911	4,956	5,979	6,945	7,830	8,622	
Plus: D&A		3,121	3,343	3,177	3,019	2,868	2,726	2,590	2,461	2,339	2,222	
Less: Capex		(5,206)	(742)	(705)	(670)	(637)	(605)	(575)	(546)	(519)	(493)	
Less: Changes in NWC		(648)	2,350	4,082	5,593	6,050	6,205	6,083	5,744	5,259	4,693	
Plus: Other Cash Transactions												
		1,606	2,162	2,911	3,937	5,047	6,185	7,300	8,354	9,318	10,179	
Unlevered FCF		2,877	7,010	8,539	11,497	15,872	20,674	25,654	30,581	35,267	39,582	
Year Fractions		-	1	1	1	1	1	1	1	1	1	
Net FCF		-	7,010	8,539	11,497	15,872	20,674	25,654	30,581	35,267	39,582	1,263,485
												EV/EBITDA 29.24
WACC	Cost	Value (\$M)	Weight	Intrinsic Value								
Risk Free Rate	4.45%			Enterprise Value	446,631	Equity Value/Share	\$ 277.36					
Equity Risk Premium	5.4%			Plus: Cash	7,243	Current Price	\$ 251.31					
Beta	1.95			Less: Debt	2,321	Upside/Downside	10.37%					
Equity	15.0%	409,143	99.4%	Less: Minority Interest	-	Rating	Buy					
Yield on Debt	2.9%			Equity Value	451,553	Average Revenue Growth %						
Tax Rate	21.0%			Assumptions		Past 5-years	33.4%					
Debt	2.3%	2,321	0.6%	Depreciation (% of PPE)	6.4%	Forward 5-years	31.1%					
Discount Rate	14.9%	411,464	100.0%	Capex (% of Fixed Assets)	3.2%							

Khaveen Investments

For our valuation, we updated our DCF models for both companies, taking into account our updated revenue and margin projections. For AMD, we continue to see its strong growth due to its strong position in the data center market, with a total 5-year revenue forecast averaging 31.1%. Based on a discount rate of 14.9% (the company's WACC) and a terminal value based on the top US chipmakers' 5-year average EV/EBITDA, our model shows an upside of 19.69%.

Discounted Cash Flow Model												
Calendar Year	Entry	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034 Terminal	
Financial Year		27/1/2026	31/12/2026	31/12/2027	31/12/2028	31/12/2029	31/12/2030	31/12/2031	31/12/2032	31/12/2033	31/12/2034	31/12/2035
Revenue Growth			7.43%	6.19%	4.97%	4.72%	4.48%	4.26%	4.03%	3.84%	3.65%	2.80%
Revenue			56,782	60,299	63,294	66,281	69,252	72,202	75,123	78,010	80,859	83,123
EBIT			2,633	6,959	9,707	12,733	19,399	20,285	21,128	21,928	22,686	23,150
Less: Cash Taxes			535	1,507	2,084	2,720	4,120	4,306	4,483	4,651	4,810	4,907
Plus: D&A			11,706	11,545	12,037	12,468	12,927	13,418	13,939	14,486	15,058	15,651
Less: Capex			(18,000)	(17,059)	(17,906)	(18,751)	(19,592)	(20,426)	(21,253)	(22,069)	(22,875)	(23,516)
Less: Changes in NWC			(904)	1,962	120	100	78	295	292	289	285	227
Plus: Other Cash Transactions												
			2,784	2,957	3,104	3,250	3,396	3,541	3,684	3,825	3,965	4,076
Unlevered FCF			(507)	933	4,737	6,881	11,932	12,216	12,723	13,230	13,739	14,227
Year Fractions			0.93	1	1	1	1	1	1	1	1	1
Net FCF			(470)	933	4,737	6,881	11,932	12,216	12,723	13,230	13,739	14,227
												490,049
												12.63
EV/EBITDA												
WACC	Cost	Value (\$M)	Weight	Intrinsic Value				Enterprise Value		Equity Value/Share	\$	43.60
Risk Free Rate	4.45%			Enterprise Value				239,509		Current Price	\$	42.49
Equity Risk Premium	5.4%			Plus: Cash				37,416		Less: Debt		
Beta	1.34			Less: Minority Interest				12,079		Rating	Hold	
Equity	11.7%	212,195	81.8%	Equity Value				217,737		Average Revenue Growth %		
Yield on Debt	2.3%			Assumptions						Past 5-years		-7.1%
Tax Rate	21.0%			Depreciation (% of PPE)				7.8%		Forward 5-years		5.5%
Debt	1.8%	47,109	18.2%	Capex (% of Revenue)				29.0%				
Discount Rate	9.9%	259,304	100.0%									

www.khaveen.com
Copyright 2026. All Rights Reserved

Khaveen Investments

Whereas for Intel, our DCF model shows limited upside from its current price despite factoring in revenue to pick up for its foundry business and a stronger but gradual recovery in margins. Based on a discount rate of 9.9% (the company's WACC) and a terminal value based on its 5-year historical average of 12.83x, which is considerably lower than the industry average, as we believe this reflects Intel's lower and uncertain growth outlook, our model shows an upside of 7.49%.

Verdict

Category	Winner
Growth	AMD
Strategy and Execution	AMD
Financials	AMD
Valuation	AMD
Overall	AMD

Khaveen Investments

All in all, our comparison summary shows AMD outshines Intel in all categories. We believe this reasonably reflects AMD's strengths due to its long-term commitment to its data center expansion strategy with a diverse product portfolio, though it shines in the CPU and GPU space, where it continues to be the top-choice CPU supplier and its progress in securing GPU customers has been very positive, which guides our strong revenue forecasts for the company, on top of its superior margin profile as a fabless company. On the other hand, while Intel has shown positive signs under new leadership with the recent ramp of its Intel 18A, and developments

with Amazon, Microsoft, and the US military could finally see Intel growing its foundry business as a start, this remains too little to give a meaningful contribution to the total company while still facing strong competitive headwinds from none other than AMD in the CPU market, and its long-term foundry ambitions still hinges on important milestones yet to be seen.