KASASV

KOREAN AMERICAN SEMICONDUCTOR ASSOCIATION IN SILICON VALLEY

May 2026

낮술

이러면,
안되는데…

– 시인 김상배
(단 두줄이지만 공감되고 강력한 시…)

INTRO

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

NGL 2026 Conference 에서는 ***SPIE Advanced Lithography + Patterning 2026 (SPIE ALP 2026)*** **학회의 주요 기술 동향과 인사이트를 정리한 Conference Summary 세션**을 마련하였습니다.

- 본 행사는 **온라인으로** 진행되며, 사전 등록을 완료하신 분들에 한하여 **Zoom 접속 링크를** 개별 안내드릴 예정입니다.

일시: **2026 년 5 월 7 일(목) 13:00 ~ 17:00**

- 장소: 온라인 회의 (Zoom Live)
- 참가비: 학술대회 회원 및 후원사 무료
- 사전 등록 기간: 2026 년 4 월 28 일 ~ 5 월 4 일
- 사전 등록 방법: **구글폼 신청 (QR 코드)**
- **QR code 및 기타 자세한 내용은 웹사이트(announcement)를 참조하시기 바랍니다.**

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

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

– 사이먼 리 드림/협회 상근임원

www.kasainsv.com 문의 info@kasainsv.com

감사합니다

TSMC Tech Symposium 2026, By The Numbers

April 27, 2026 Barry Pangrle

TSMC announced three new advanced process technologies at its North America Technology Symposium last week A13, A12 and N2U as part of its latest advanced technology roadmap.

Compared to the roadmap presented at last year's tech symposium, the new roadmap shows the three new process technologies in gold text

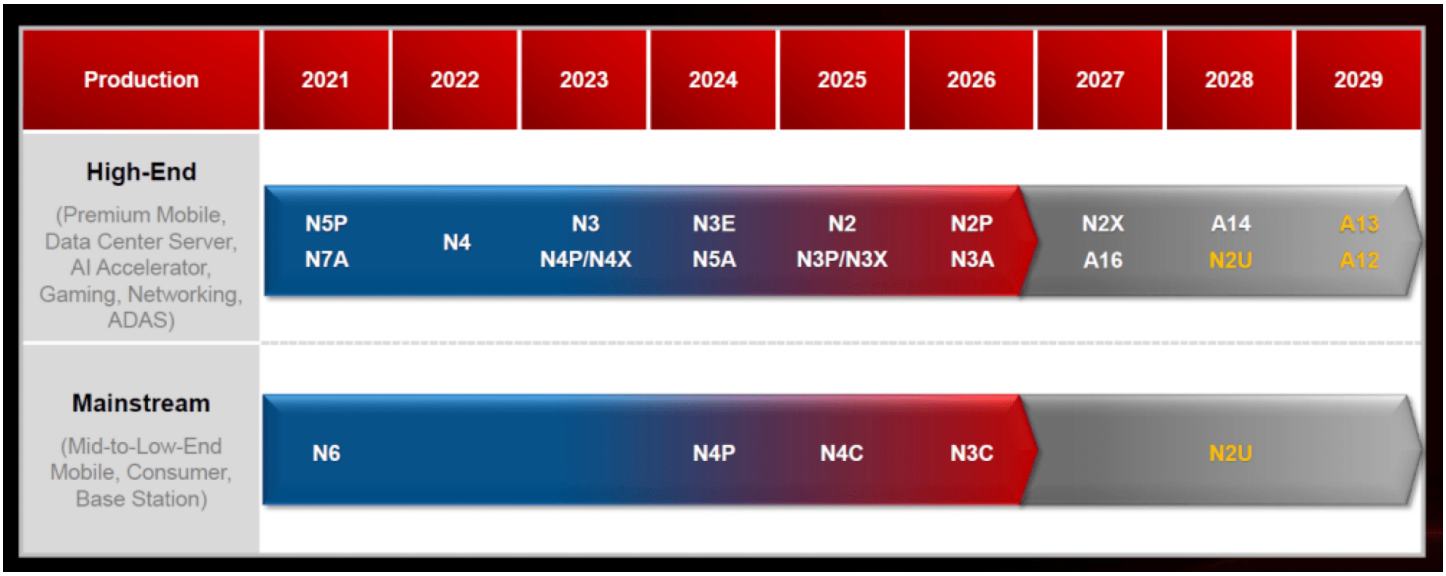

Figure 1: TSMC's Advanced Technology Roadmap. Source: TSMC

To put these three into context, N2 is TSMC's first node to use NanoFlex, and A16 will be TSMC's first node to use Super Power Rail (SPR) backside power delivery technology. N2 is in production now, and TSMC claims that process has the strongest ever customer adoption, with more than 20 customer tape-outs received so far and over 70 in the pipeline. We expect to see N2 revenue showing up soon.

TSMC's A14 uses its NanoFlex Pro technology and is targeted for 2028. A13 is a direct shrink of A14. It provides 6% area savings, and the design rules are fully backward compatible with A14. A13 is targeted for 2029, along with A12, which is an A14 platform enhancement with SPR. Figure 2, below, shows N2U's incremental improvements over N2P and it is scheduled for production in 2028.

N2U PPA (vs. N2P)		
Speed Improvement at Same Power	Power Reduction at Same Speed	Logic Density
3~4%	8~10%	1.02~1.03X

Figure 2: N2U Enhancement Advantages vs. N2P Source: TSMC

N2P PPA (vs. N3E)			
Speed Improvement at Same Power	Power Reduction at Same Speed	Logic Density	Chip Density
~18%	~36%	~1.2X	>= 1.15X

TSMC A16™ PPA*			
Technology	Speed Improvement at Same Vdd	Power Reduction at Same Speed	Chip Density
A16 (Ref: N2P)	8~10%	15~20%	1.07~1.10X

*Datacenter AI products

A14 PPA (vs. N2)			
Speed Improvement at Same Power	Power Reduction at Same Speed	Logic Density	Chip Density
10~15%	25~30%	~1.23X	~1.2X

Figure 3: PPA Comparisons from 2025 Tech Symposium. Surce: TSMC

Figure 3 shows PPA comparisons from last year's TSMC Tech Symposium. The A14 numbers were presented again this year without any changes. A16 had been advertised earlier for production in late 2026, but Figure 1 now shows it in 2027. The announcement of A12 for 2029 is an aggressive speedup in delivering a new technology node. TSMC also claimed that N2 is ramping better in terms of defect density reduction than N3. This is quite impressive, given that N3 is the last finFET node and N2 is the first NanoFlex node.

The demand for advanced packaging technologies is very high and largely driven by AI compute scaling requirements. TSMC is manufacturing the world's largest 5.5-reticle size CoWoS, with > 98% yield in 2026. Figure 4 shows the roadmap for CoWoS and SoW support out to 2029, with an SoW-X projected to accommodate 64 HBM stacks. If we expect 64GB HBM stacks for 2029, that will imply a total of 4TB of HBM on an SoW-X.

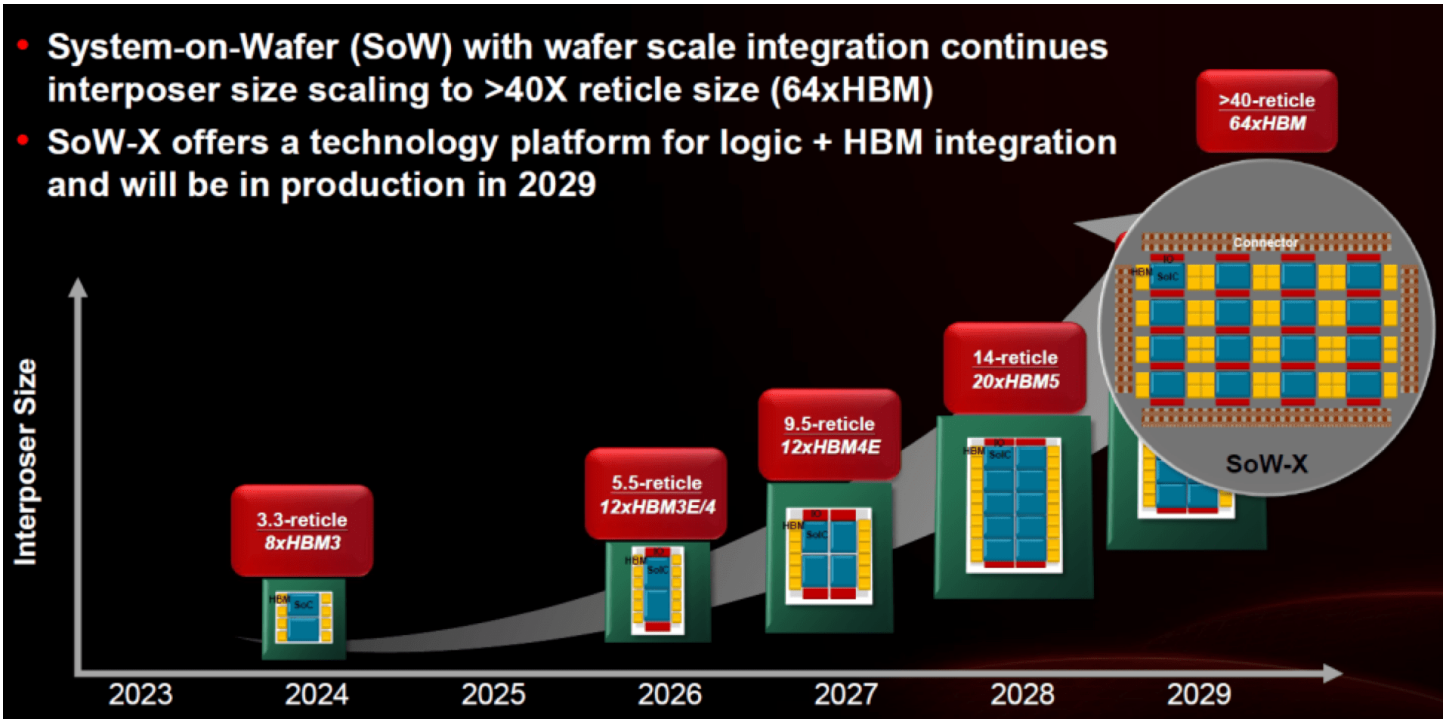

Figure 4: TSMC CoWos and SOW Integration Roadmap. Source: TSMC

Another important technology to enable system scaling for AI is co-packaged optics. TSMC's Compact Universal Photonics Engine (COUPE) is bringing optical signaling closer to the chip, and along with it come improved energy efficiency and latency reduction. Figure 5, below, shows the roadmap progression for CPO with COUPE.

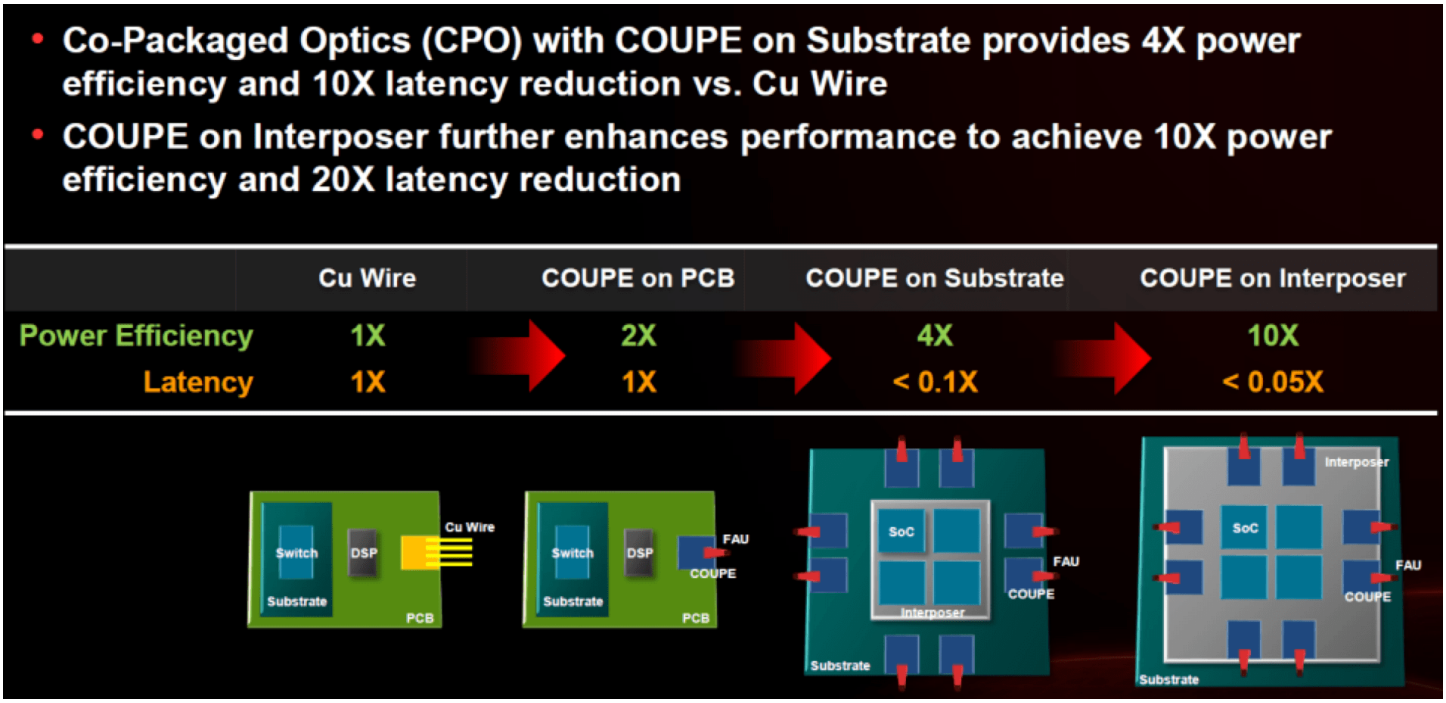

Figure 5: Co-Packaged Optics Roadmap using COUP. Source: TSMC

TSMC also had its first quarter of 2026 earnings call on April 16. Figures 6 through 10 below show the dramatic impact that AI/HPC has had on revenue, and how it has displaced smartphones as the top revenue generator

for TSMC. Figure 10 also shows how consistent the revenue percentages for N7, N5 and N3 have been for the last 6 quarters. It will be interesting to see how N2 will still start to impact these percentages moving forward as more finished wafers are delivered.

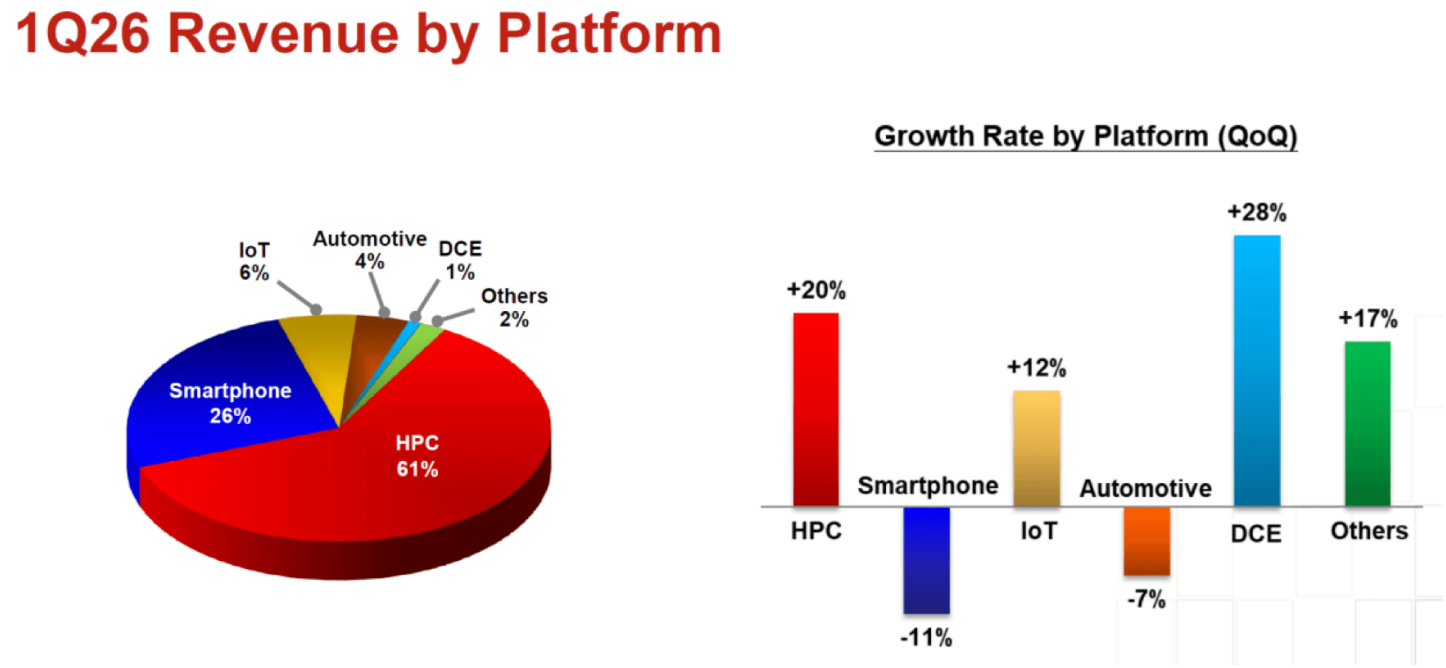

Figure 6: Revenue by Platform. Source: TSMC

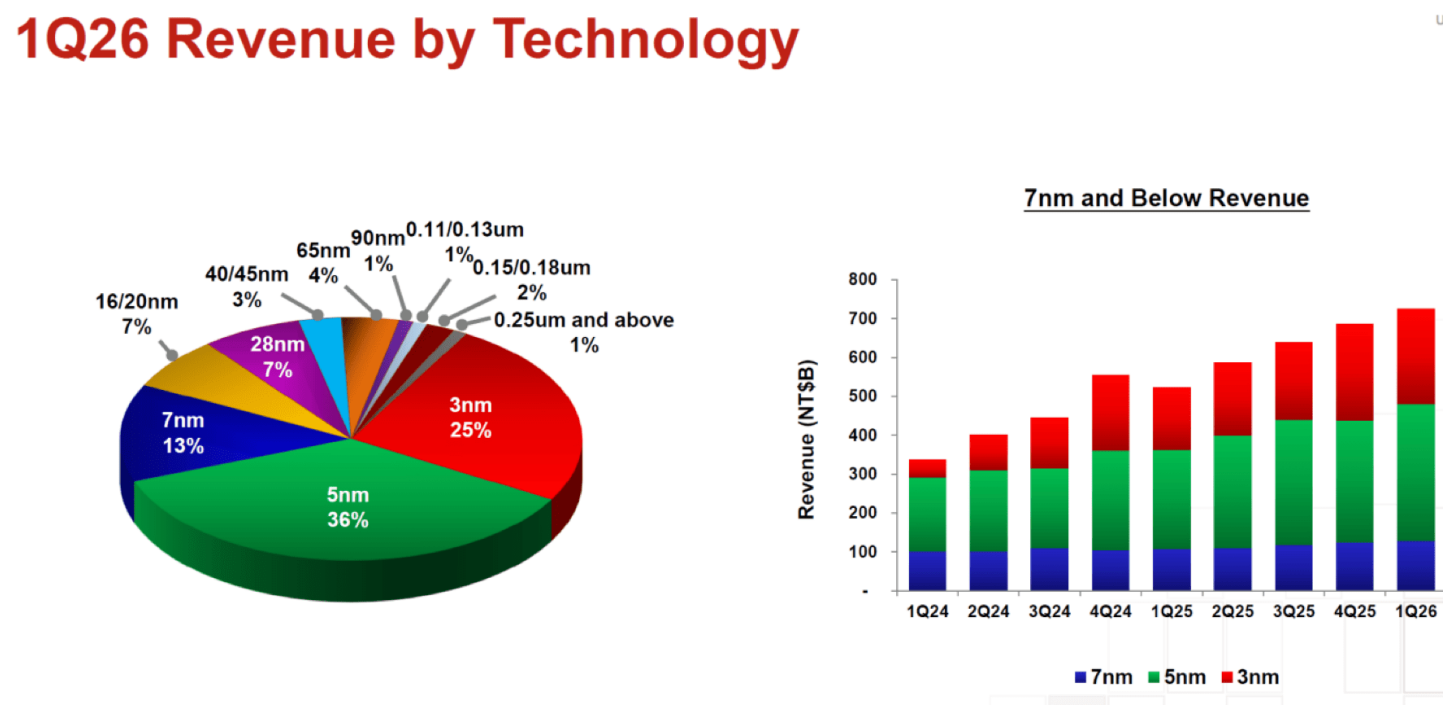

Figure 7: Revenue by Technology. Source: TSMC

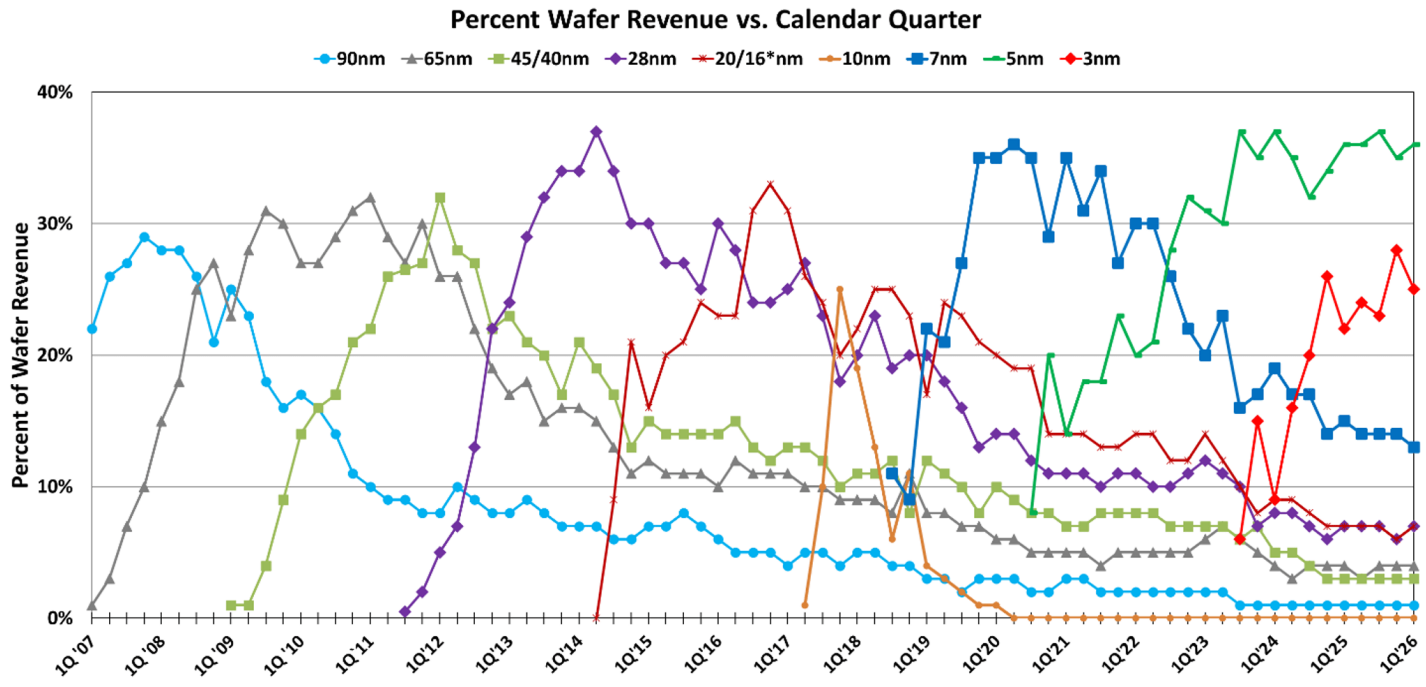

Figure 8: Historical Percent Wafer Revenue by Node Per Calendar Quarter (19 years). Source: TSMC

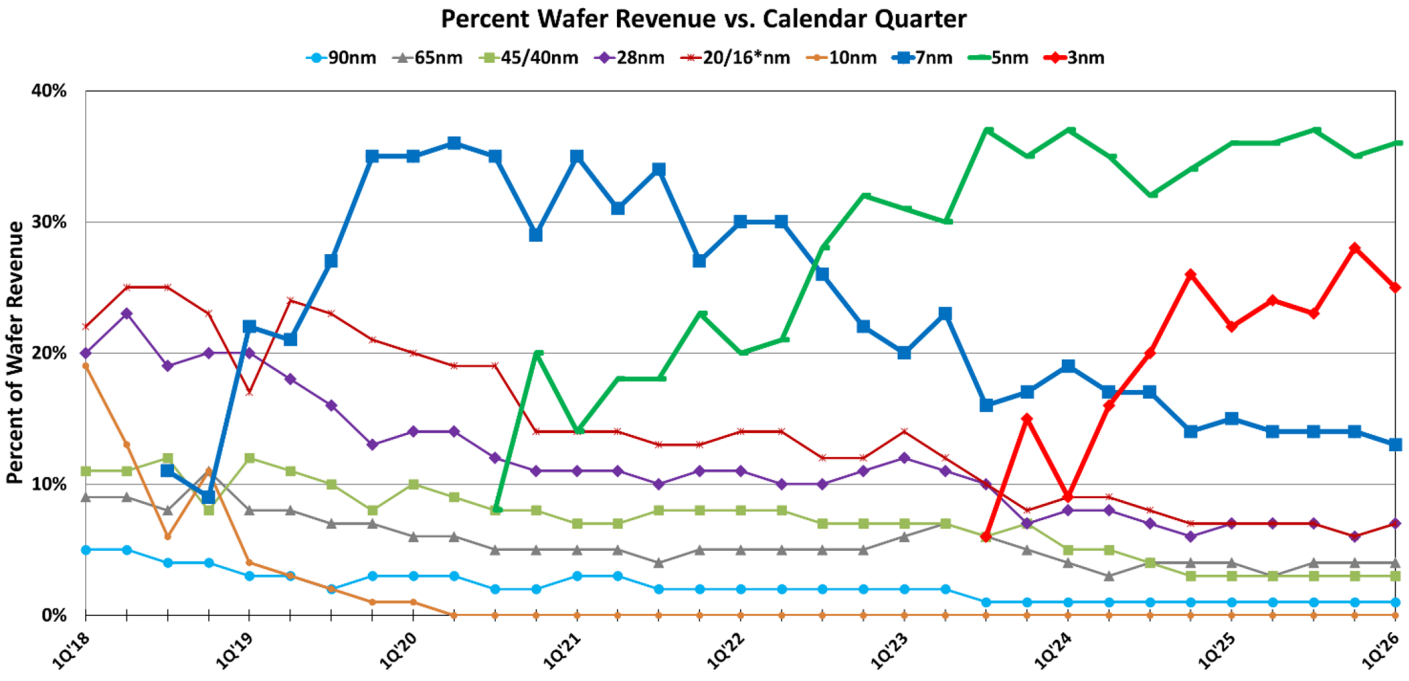

Figure 9: Percent Wafer Revenue by Node Per Calendar Quarter (8 years, N7 and below Era). Source: TSMC

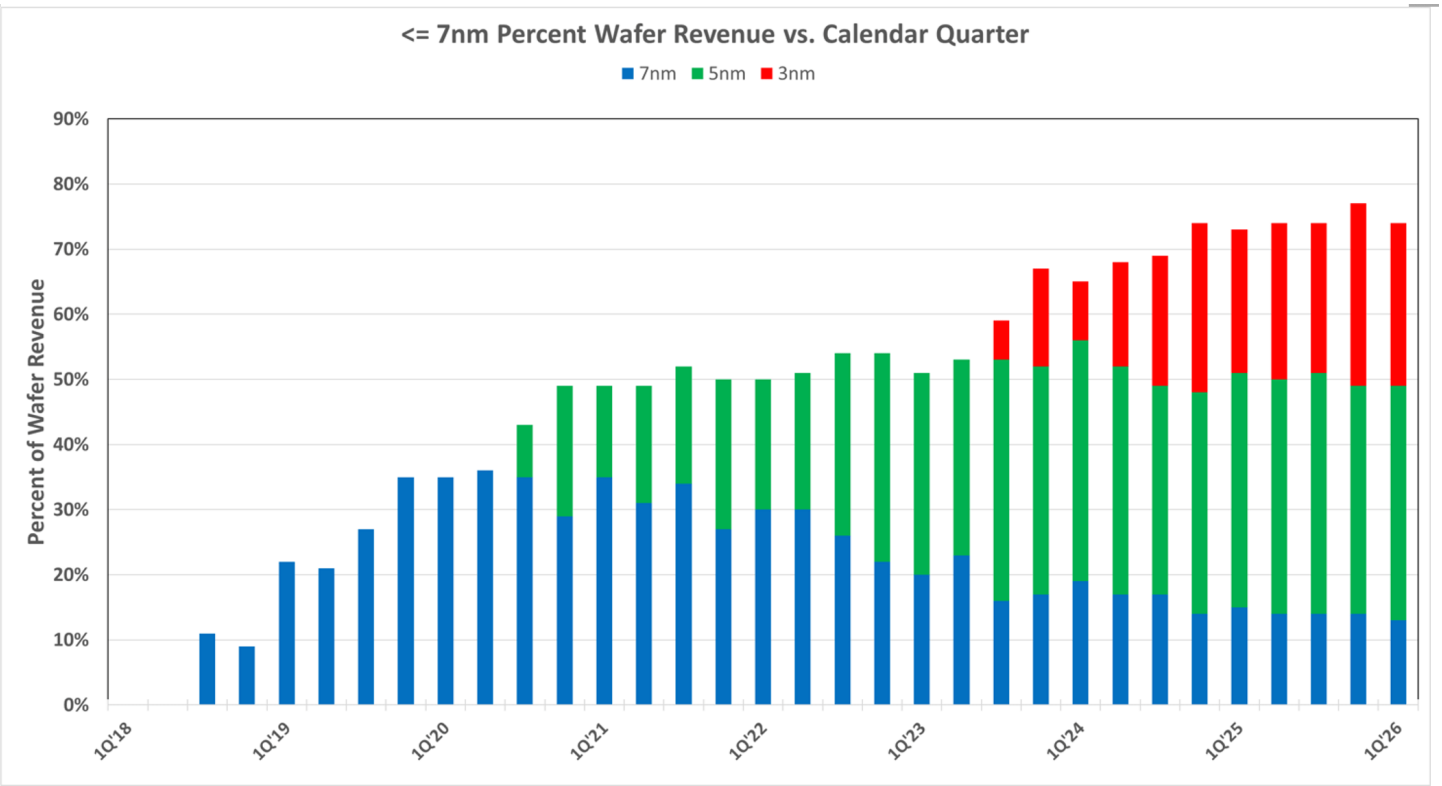

Figure 10: Percent wafer revenue for </= 7nm by Calendar Quarter. Source: TSMC.

TSMC Launches A13 Process and Advanced Technologies

April 27, 2026 Article By TSMC

A13 maintains full backward compatibility with A14 design rules, enabling faster customer migration to the company's nanosheet transistor architecture.

Taiwan Semiconductor Manufacturing Co. Ltd (TSMC) introduced its A13 process technology at its 2026 North America Technology Symposium, positioning it as a direct shrink of the previously announced A14 node to support increasing computational demands across artificial intelligence (AI), high-performance computing (HPC), and mobile applications.

A13 delivers 6% area reduction compared with A14 and maintains full backward compatibility with A14 design rules, enabling faster customer migration to the company's nanosheet transistor architecture. The process also provides improved power efficiency and performance through design-technology co-optimization and is scheduled for production in 2029, one year after A14.

"At TSMC, we understand our customers are always looking ahead to their next innovation and they come to us for a reliable stream of new silicon technologies, like A13, meticulously engineered to be ready for high-volume production right when their visionary new designs demand them," said TSMC Chairman and CEO Dr. C.C. Wei. "TSMC's advanced process technologies lead the industry in density, performance and power efficiency, and we continually strive to make them even better for our customers' future products, ensuring customers' success as their most reliable technological partner."

Advanced logic and 2nm platform developments

TSMC also previewed A12, an enhancement of the A14 platform featuring Super Power Rail technology for backside power delivery in AI and HPC applications, with production targeted for 2029.

The company further advanced its 2nm roadmap with N2U, which delivers 3–4% speed improvement or 8–10% power reduction and a 1.02–1.03× logic density increase over N2P. N2U is scheduled for production in 2028.

3DFabric and packaging innovations

To address AI-driven demand for higher compute density, TSMC expanded its CoWoS® advanced packaging technology. The company is producing 5.5-reticle CoWoS and developing a 14-reticle version capable of integrating approximately 10 compute dies and 20 high-bandwidth memory stacks, targeted for production in 2028, with further scaling planned in 2029. These developments complement the 40-reticle SoW-X System-on-Wafer platform expected in 2029.

TSMC also introduced updates to its TSMC-SoIC® 3D stacking technology, with A14-to-A14 SoIC planned for 2029, delivering 1.8× higher die-to-die I/O density than N2-on-N2 SoIC.

In photonics, the company's TSMC-COUPE™ technology will enable co-packaged optics production starting in 2026. By integrating optical engines within the package, it achieves 2× power efficiency and 10× latency reduction compared with pluggable solutions, demonstrated in a 200Gbps micro-ring modulator.

Automotive, robotics, and specialty technologies

For automotive and physical AI applications, TSMC announced N2A, its first automotive-grade nanosheet transistor process. N2A offers 15–20% speed gains at the same power versus N3A and is expected to complete AEC-Q100 qualification in 2028. The company is also enabling earlier automotive design through "Auto-Use" kits within its N2P process design kit.

N3A is set to enter production in 2026, with more than 10 automotive products planned under the N3 "Auto Early" program.

In specialty technologies, TSMC introduced its N16HV high-voltage FinFET process for display driver applications. Compared with N28HV, N16HV increases gate density by 41% and reduces power consumption by 35% for smartphone displays, while enabling over 20% power reduction and 40% die size shrink for near-eye display applications.

The announcement was made at the symposium held in Santa Clara, which carries the theme “Expanding AI with Leadership Silicon” and serves as TSMC’s primary annual customer event.

Sandisk: A High Beta Play For NAND Memory As AI Demand Surges

April 21, 2026 Robert Glanville

Summary

- Sandisk has delivered a 2,424% return since its 2025 divestiture, driven by surging NAND memory demand from AI hyperscalers.
- SNDK benefits from a joint venture with Kioxia, high average selling prices, and compressed cash conversion cycles, signaling robust operational efficiency.
- Cloud/datacenter sales grew 123% from Q3 2025 to Q2 2026, reflecting SNDK's hyperscaler foothold, but cyclicality and supply discipline remain critical to future margins.
- With no valuation premium and rising supply risks, I see more downside than upside near-term and am holding pending key SID drivers on the April 30th earnings call.

LisalnGlasses/iStock via Getty Images

Sandisk Corporation (SNDK), which used to be controlled by Western Digital Corporation, has accumulated a whopping 2,424% return since it was divested in February of 2025. The company’s current Seeking Alpha ratings are all in the green ahead, as it rose another 15% in the last week. But what do they sell? Sandisk produces memory storage, which in recent months has been in high demand for AI hyperscalers. To run large language models, High Bandwidth Memory (HBM) or Dynamic Random Access Memory (DRAM) is required while the model is in use. But this is not the Sandisk pure play. Sandisk sells NAND memory, the inactive ingredient used when the model isn’t in use. NAND stores data that is used to train LLMs, and this specific product is where we are seeing incredible returns. Quarterly revenues have only risen roughly 59% in the same time, while the stock price has returned over 2,000%. What's the reason for the boom? Will these conditions

continue? And what should we look out for on the next earnings call? I'll touch on all of that below, but to summarize, I am holding until April 30th on the next earnings call.

Supply

The story of the NAND boom cycle starts in 2023-2024, when the memory market was oversupplied. Economic theory tells us this played a role in driving prices and margins down. Suppliers would then slow production, so big-name competitors like Micron Technology, Inc. (MU), Kioxia Holdings Corporation (KXIAY), and Sandisk invested their money elsewhere. Now, with a new and substantially lower equilibrium inventory level, the AI boom comes into play. Capex investment in AI data centers caused a rise in demand for NAND memory. Short supply with high demand creates high prices and a high-margin environment. For Sandisk, annual revenues in 2023 were $6 billion relative to today's TTM of almost $9 billion. There are 2 key factors to look at here:

First, Sandisk has a hidden advantage. They entered into a joint venture with Kioxia Holdings. Instead of building their own manufacturing facility, they both split the cost of production out of one jointly owned facility. They are basically roommates in an expensive house.

Second, margins are propped up by high average selling prices (ASP). As stated in the most recent 10-Q filing, "Net revenue increased 61% in the three months ended January 2, 2026, from the comparable period in the prior year, primarily due to a 36% increase in average selling prices ('ASP') per gigabyte and a 22% increase in exabytes sold." While the joint venture keeps costs low, higher prices go straight to the margin.

The result is that costs of revenues are almost identical comparing 2023 to the present, resulting in a $3 billion boost in gross margin. Low overhead and high prices are every company's dream. Additionally, supply is a de facto oligopoly:

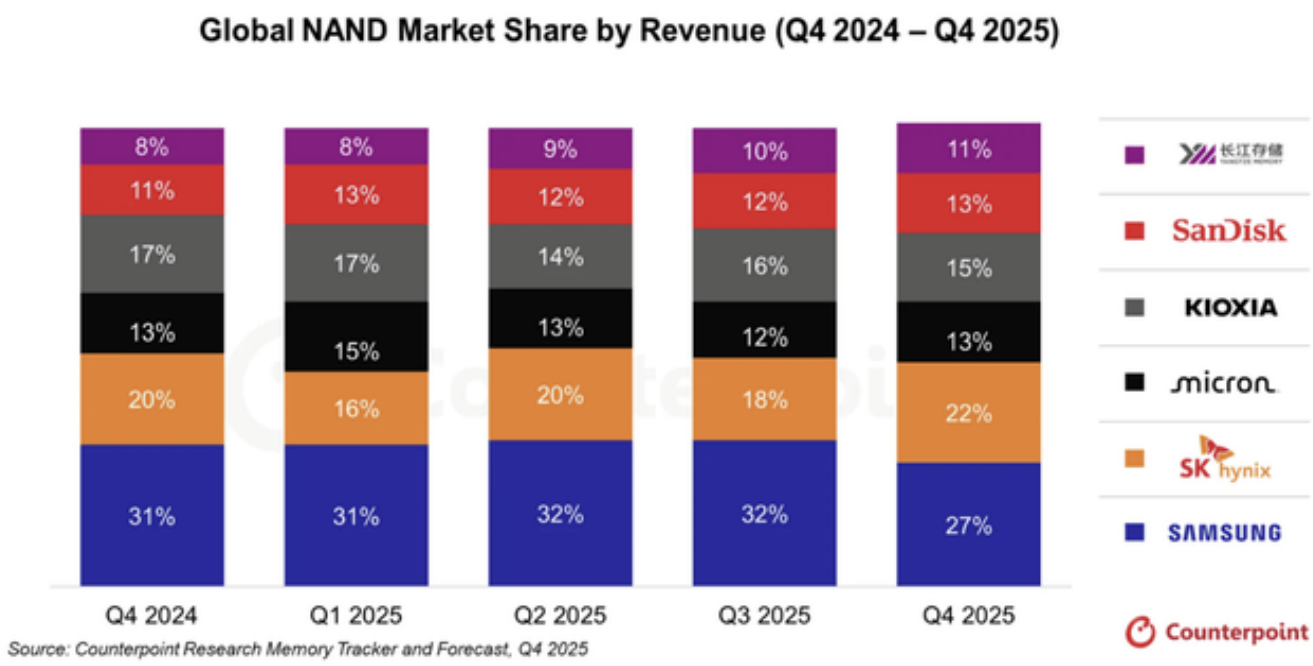

Counterpoint

I see a market this concentrated and think there would be some pricing power going on here. But Sandisk is the only pure play, whereas other competitors don't benefit nearly as much from high prices. So they have less of an incentive to restrict supply to keep prices high. My current assessment is that we will likely see

overproduction here soon, bringing prices back down to normal levels. But at the current moment, we still may be seeing increases in the short term. Here is a note from Micron's investor presentation:

% ASP Change	FQ1-26 Q/Q
DRAM	Increased approximately 20%
NAND	Increased in the mid-teens percentage range

% ASP Change (Micron Investor Presentation)

So, the key to keep in mind is capex from the broader market moving forward. If we see supply ramping up, we are likely to see another cyclical downturn as prices level.

Inventory

Dissecting inventory turnover, I like to look at the overall cash conversion cycle (CCC). This gives us a better understanding of whether they are recklessly riding a demand boom or if they are displaying operational efficiency. When looking at 2025-2026, we see the CCC has compressed from 150 to 104 days. Shocking, I know… with higher demand comes more cash. But what's good for Sandisk is that a falling DIO with unchanged DPO means Sandisk is turning over inventory faster without shortening its timeline to pay suppliers. These are signs of health.

Days' inventory outstanding is also my key metric moving forward. We talked earlier about prices being a major driver for margins, but demand for cloud services is the other side of the coin. If DIO falls, it means demand is still increasing. But if it rises… the opposite is true. In which case, I would recommend selling.

Demand—The Jevons Paradox

To add to the falling demand concerns, innovations, like TurboQuant, are threatening to make the AI process and memory market more efficient. Running LLMs requiring less memory demand would be unfortunate for memory producers, right? This is a modern-day Jevons Paradox. Just like in 1865, when steam engines became more efficient, demand for coal increased rather than decreased. Greater efficiency actually encouraged wider use. I see the same story here with AI memory demand. Sure, in the short term, we may see reduced storage requirements, but in the long term, demand will drastically expand to meet broader AI applications. When I see headlines of memory compression or efficiency, I think of seeded growth.

From a financial perspective, Sandisk saw a major shift in its product segment sales mix. From Q3 2025 to Q2 2026, Cloud/Datacenter sales grew by 123%. In the same period, sales for Edge/Client and Consumer rose 81% and 59%, respectively. What this tells me is that they have locked in growth hyperscalers (cloud/datacenter), who don’t switch suppliers easily.

From a demand perspective, I see more efficient memory requirements, paired with a strong hyperscaler foothold, to be 2 strong indicators for future sales.

Key SID Drivers

The next earnings call is scheduled for 30 April 2026. I've synthesized the key SID drivers that should be in focus.

- **Supply Discipline and Manufacturing Utilization**—High average sales prices have been the tailwind for growth. Increases in industry-wide capex would signal a future supply boom, decreasing ASP, and drastically reducing margins.
- **Inventory Digestion**—Sandisk's cash conversion cycle has fallen by 46 days, which will be a strong benchmark moving forward. A growing inventory (DIO), which currently sits at 121 days, is a clear signal of oversupply, which would have downward pressure on prices and margins.
- **Sales Mix/Demand—**Growth in revenues is driven by cloud/datacenter expansion. Share price performance will be largely dependent on those growth expectations.

Financials

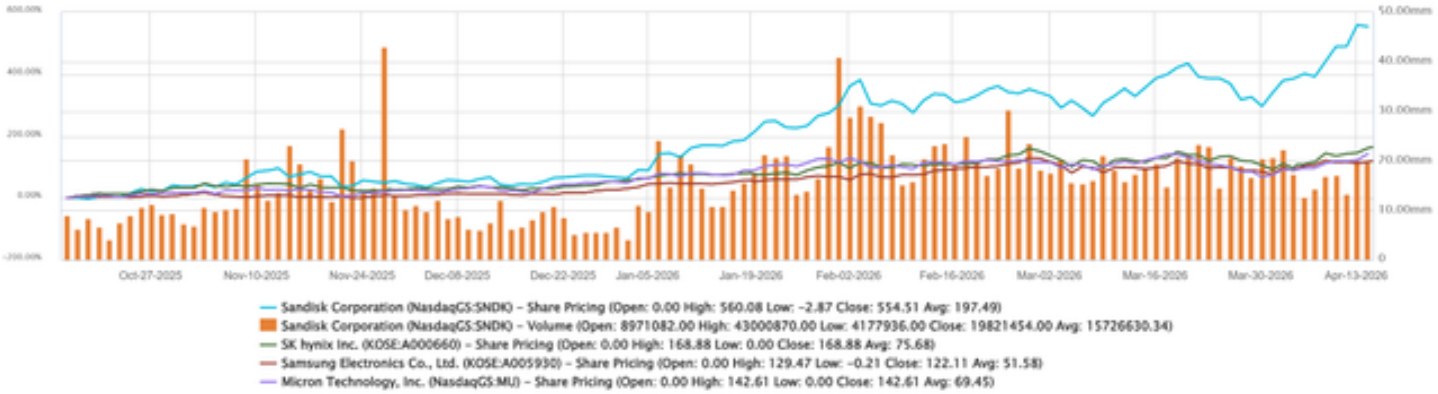

Price Performance (Capital IQ)

To illustrate the aggressive performance relative to memory-producing competitors, the above chart from Capital IQ includes SK Hynix and Samsung Electronics Co., Ltd. (SSNLF), and Micron. First, Sandisk is the only pure play in the mix here, achieving the concentrated gains from high prices and demand. Although subtle, you can see a very similar pattern, showing the cyclicality of the returns and the sensitivity to NAND price increases. Being the pure play option of the market, Sandisk has minimal downside protection, and we could see these returns drastically reverse should we enter a low-price environment.

Additionally, from a forward P/E perspective, we do not see a clear premium relative to competitors. Earnings have now been priced into the equation and have converged the market to lower multiples. This could mean we are in the late stages of the memory cycle, which would indicate a peak in earnings is among us. However, it would be unwise to assume this cycle will be like the last. With AI capex investment, there is no clear indication that memory supply will outpace demand anytime soon.

Risks

Among the risks I mention in this article, I think there are 2 key areas of my investment thesis worth reiterating. First are the AI efficiency headwinds. As I talked about the Jevons Paradox, there could be a short-term downside caused by memory innovation. How long exactly is short-term? That's left to be unsaid. With inflation high, uncertainty in future rate cuts, and global conflict, the macroeconomic situation makes future AI infrastructure demand a hard nail to pin down.

Second, I mentioned earlier the joint venture with Kioxia. This partnership is invaluable because it allows Sandisk to flatten costs. However, as we've seen with the current administration, tariffs are never off the table. With manufacturing occurring overseas, another high-tariff environment could offset the synergies from the partnership and make the memory market more complicated to price.

Conclusion

Sandisk is a boat on the tide of NAND pricing. When the tide rises, it might as well be a cruise ship, but when it recedes, you might find yourself sitting in the marsh. Market expectations have lowered for Sandisk, and although we see great tailwinds for demand, the lack of a valuation premium, a severely low PEG ratio, and rising supply risk indicate to me that there is more downside than upside potential. On April 30th, I will be listening for those 3 key SID drivers to forecast the next quarter. Until then, I am holding.

ASML Q1 Earnings Review: Sales Growth Slows, Memory Dominates Compute

April 16, 2026 Sandeep G. Rao

Summary

- ASML Holding N.V. exceeded Q1 expectations with €8.8B net sales and €2.76B net profit, while sentiment remains muted amid AI cycle concerns.
- Growth is slowing: FY2026 net sales may be half of FY2025, with memory chip equipment now dominating net system sales.
- Geopolitical constraints, notably reduced China sales, and rising grassroots opposition to datacenters in the U.S. pose structural risks to ASML's forward demand.
- Cyclical headwinds, competitive pressures in DUV, and sustainability concerns for AI-driven demand are currently dominant points of consideration against ASML.

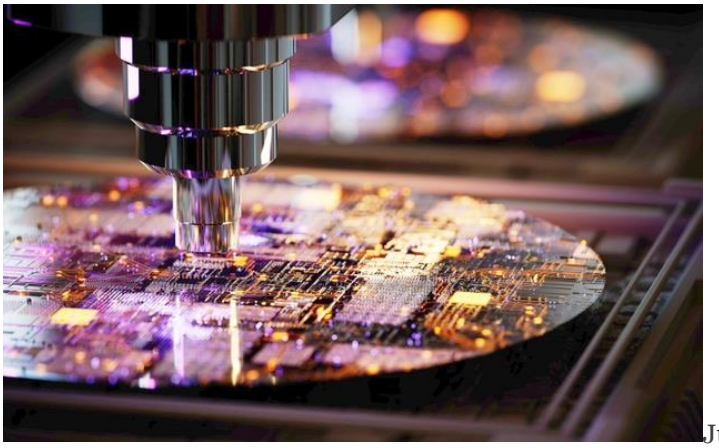

Just_Super/iStock via Getty Images

ASML Overview

Dutch semiconductor manufacturing equipment provider ASML Holding N.V.'s (ASML) stock price didn't generate a substantial investor sentiment despite its Q1 earnings release on the 15th of April indicating that it had exceeded expectations with €8.8 billion in net sales and €2.76 billion in net profit.

The factors behind this are centered on the overall expectations of the AI Hype cycle, to which ASML has been linked to.

Trend Drilldown

Early trends indicate that net sales (or revenue) could close Fiscal Year (FY) 2026 at roughly half the growth exhibited in FY 2025.

	Q1 2026, % of 2025	2025 vs 2024	2024 vs 2023	2023 vs 2022	2022 vs 2021	2021 vs 2020
		YoY Trends				
Net Sales	27%	16%	3%	30%	14%	33%
System Sales	26%	12%	-1%	42%	13%	32%
Service and Field Option Sales	30%	26%	16%	-2%	16%	35%
Cost of Sales	27%	12%	3%	28%	19%	23%
Research & Development	25%	9%	8%	22%	28%	16%
Selling, General & Administrative	24%	8%	5%	18%	30%	33%
Operating Income	28%	25%	0%	39%	-4%	67%
Net Income	29%	27%	-3%	39%	-4%	66%

Created by Author using ASML's Financial Statements

The bulk of the slowdown can be attributed to the sale of lithography machines used in the manufacturing of advanced chips. However, as past data reveals, this is a highly cyclical business: the machines sold are expensive and long-lived and go on to generate income for the company via service fees, et al., which are captured in the "Service and Field Option Sales" line item and which are trending at 20% growth in FY2026.

Expenses such as R&D expenditure and selling costs are running at par with the previous FY, while the relatively higher ticket price of machines sold is aiding in boosting net income to trend at 16% growth, which is, again, lower than that seen through most of the 2020s.

A very noticeable change, at least in this quarter, is the composition of the purpose for which machines are being acquired by clients: memory manufacturing dominates over the sale of machines for the manufacturing of GPUs ("Graphical Computation Units") that are the core of the "brain" employed for running AI.

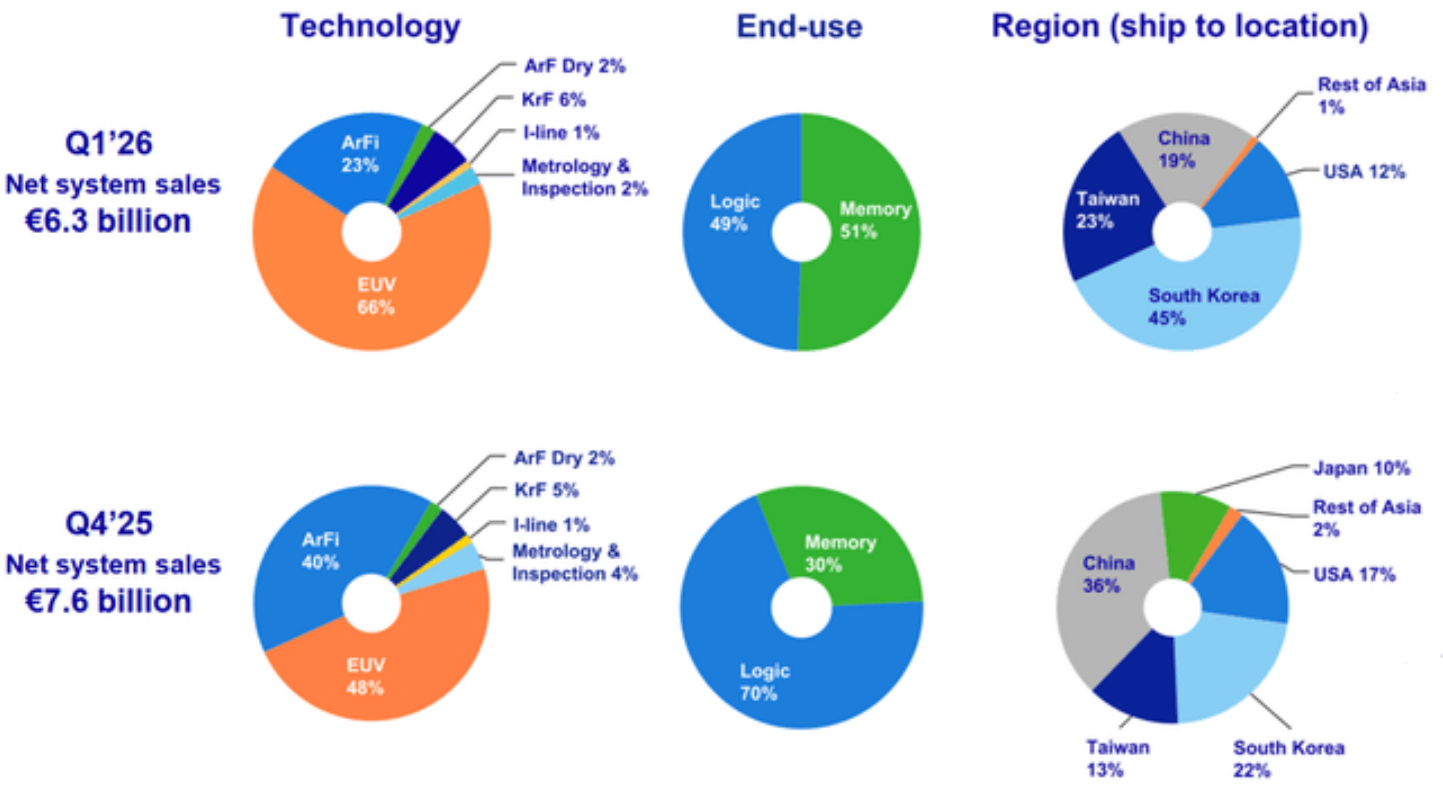

ASML Q1 2026 Presentation

Within a single quarter, machines for the manufacturing of memory chips have gone from 30% of all system sales to 51%. This is borne out by the massive increase in sales share in South Korea from 22% to 45%, where Samsung Electronics (SSNLF) and SK Hynix (HXSCL) have been ramping up production of memory chips. Furthermore, of relevance is the fact pattern that EUV ("extreme ultraviolet") lithography machines comprise 66% of all sales now (up from 48%). ASML is the world's sole supplier of EUV machines employed in the manufacturing of advanced chips. While companies such as Japan's Nikon and Canon supply lithography tools as well, they are generally deemed not sufficiently capable enough for the thin-wafer technology being employed in AI-relevant chips.

Now, while EUV machines find particular favor for the manufacturing of AI-relevant chips, the dynamics are different for memory chips. Memory manufacturers typically use EUV for critical layers and rely heavily on DUV ("Deep Ultraviolet") lithography machines for other layers. During the earnings call, CFO Roger Dassen highlighted that DUV would be an area of interest for the company. While the company leads in high-end DUV ("immersion") and EUV, it faces intense competition in the DUV space: Nikon, for instance, is a particularly staunch competitor, while Canon has been actively advocating for cost-effective alternative technologies to the type of high-end DUV that both ASML and Nikon exist in. In the "non-EUV" business, Dassen continues, the company expects to see an increase in demand in 2026 as opposed to previous expectations of running flat relative to the previous FY.

As a percentage of sales, however, it bears noting the company shows quite strong symmetry across FYs:

	Q1 2026	FY 2025	FY 2024	FY 2023	FY 2022	FY 2021	FY 2020
Key Line Items		as a % of Net Sales					
System Sales	72%	75%	77%	80%	73%	73%	74%
Service and Field Option Sales	28%	25%	23%	20%	27%	27%	26%
Cost of Sales	47%	47%	49%	49%	49%	47%	51%
Research & Development	14%	14%	15%	14%	15%	14%	16%
Selling, General & Administrative	3%	4%	4%	4%	4%	4%	4%
Operating Income	36%	35%	32%	33%	31%	36%	29%
Net Income	31%	29%	27%	28%	27%	32%	25%

Created by Author using ASML's Financial Statements

While net income has improved as a result of higher-margin EUV sales, they do not project massive impacts in passthrough efficiency. The company essentially prices to perfection and collects on installed capacity like one would collect on an annuity.

Consequences Of Control And Concentration

ASML also indicated that China has undergone a near-halving of revenue share in Q1 2206 relative to Q4 2025 as the company seemingly limits sales to previously-inked contracts (as indicated in the previous earnings call) under the weight of technology export restrictions. This represents a long-term bottleneck for the company: its ability to turn out high-quality products effectively turns it into a pawn of global geopolitics, potentially leaving a significant portion of demand untapped.

This untapped demand, in turn, becomes the potential nucleus of permanent substitution. Given that compute and memory are the cornerstones of technology employed in modern economies, national demands must be met. In countries with high friction, sovereign policymaking trends, and large talent pools, this will inevitably promote calls for self-reliance. Restrictions on China or even a "soft ban" on (for example) Indian enterprises acquiring the technology to meet domestic needs in the future under the banner of geopolitical wrangling will inevitably lead to the creation of alternatives.

Proponents of ASML who mention "moats" might like to bear in mind that moats surrounding castles were regularly breached in the Middle Ages. While companies think in quarters and fiscal years, countries necessarily think in decades at least.

Another aspect that investors are likely considering is the sustainability of demand. Purchases of ASML's machines are predicated on sustained demand for AI-relevant chips in datacenters, which is currently under fire in arguably one of the world's biggest markets for AI: the U.S. The state of Maine has passed a "temporary" ban on the construction of datacenters. This was merely the legislative affirmation of a larger trend: between May 2024 and early 2026, local opposition in the U.S. has significantly impacted the data center industry, with $64 billion to $162 billion in projects being blocked or delayed.

The grassroots of the movement against datacenters have a strong bipartisan flavor: around 55% of the local officials opposed to datacenters are Republicans (i.e., ostensibly the same party in power in the White House). The factors for this opposition have been well-documented in recent times: datacenters consume massive

volumes of water and electricity that often have the communities they're situated in - usually rural and relatively impoverished - reeling from the effects of and footing the bill for.

In contrast, the likes of China and India might have strategic intents behind building datacenters (as opposed to the mostly commercial in the U.S.). However, given their economic size and talent availability, owning a stake in the production of assets is a point of contention they would seek to negotiate, but ASML is unable to comply with, given the laws it must follow. Whether ASML would ever be able to is currently deemed to be somewhere between uncertain and unlikely.

The uncertainty behind widespread adoption given the high costs, the grassroots opposition, and the geopolitical constraints are combining to create an impact on the forward valuation of the likes of ASML and potentially other AI-adjacent stocks in this current earnings season.

TSMC: The AI Silicon Shortage Is About To Get Worse

April 13, 2026 Doug Collins

Summary

- TSMC is the critical bottleneck in the AI semiconductor supply chain, with 2nm and 3nm capacity fully sold out into 2027 and beyond.
- TSM's advanced node and CoWoS packaging capacity are running at or above 100% utilization, with demand outstripping supply despite aggressive expansion.
- TSM trades at a significant valuation discount (10x FY27E EPS) versus peers, despite superior growth, margins, and multi-year visibility on AI-driven demand.

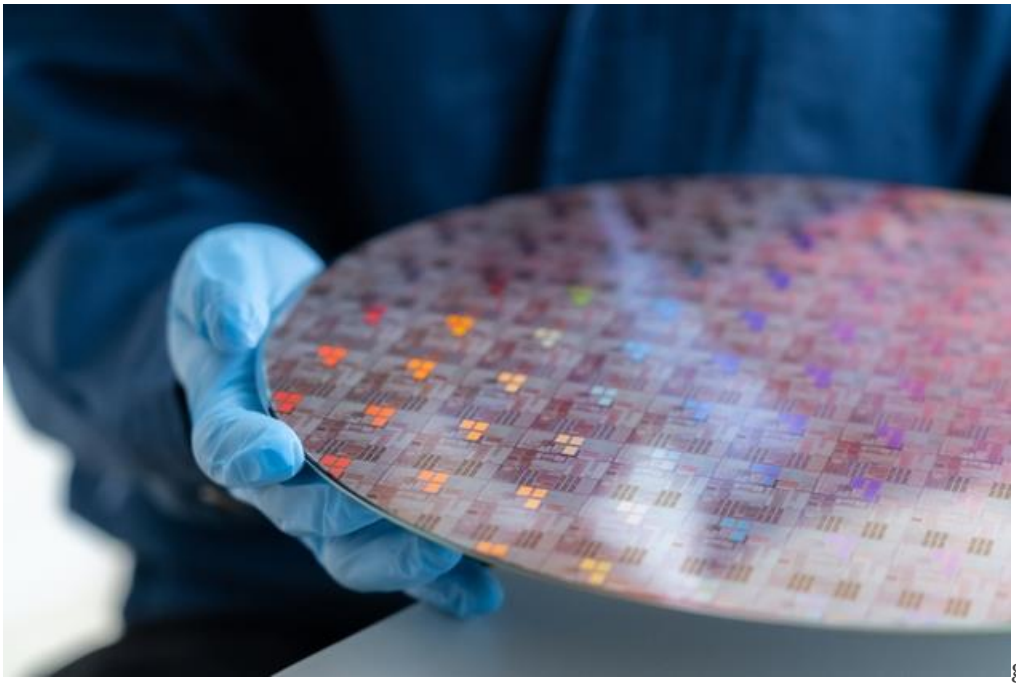

gentlelight/iStock via Getty Images

I have learned a lot since I started researching the AI supply chain and sharing that work here, and a lot of the credit for that goes to Seeking Alpha and their editors, who have helped me along the way with every piece I have published. I wouldn't have learned this much if I weren't publishing here. This platform is basically peer-reviewed journaling for me, and every article has pushed the research a layer deeper.

So anyway, over the last couple of weeks, my research has been converging on one point. The sheer scale of AI demand, and the fact that all of it is arriving at once, has created a bottleneck at the very base of the supply chain. We have a power shortage because the grid physically cannot support the load growth coming from AI data centers. We have a huge HBM shortage, which I covered in detail in my Micron piece. And if you read my latest Intel article, we have a CPU shortage, too.

So, shortages everywhere across the supply chain.

And every single one of these components, except electricity, is silicon. It is all semiconductor wafers. GPUs, HBM dies, CPUs, networking ASICs—it doesn't matter. They all funnel down into the same physical process. And basically, one company has a monopoly on that process.

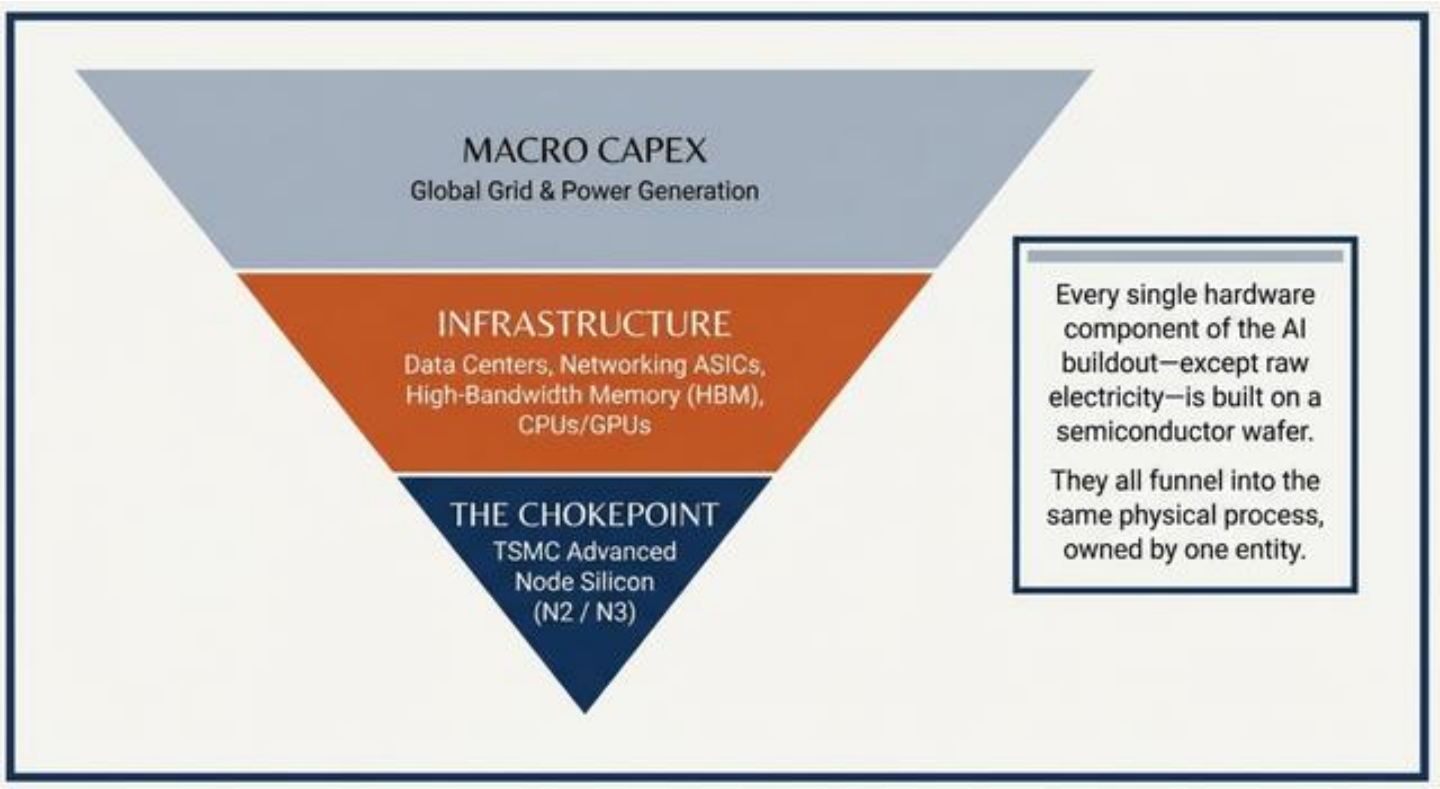

Supply Chokepoint (Author)

Which brings us to Taiwan Semiconductor Manufacturing Company Limited (TSM). They report Q1 2026 earnings on April 16th, a couple of days from now, and I will be watching this print very closely. I believe we are about to have a huge shortage of 2nm and 3nm advanced chips. A multi-year waitlist kind of shortage.

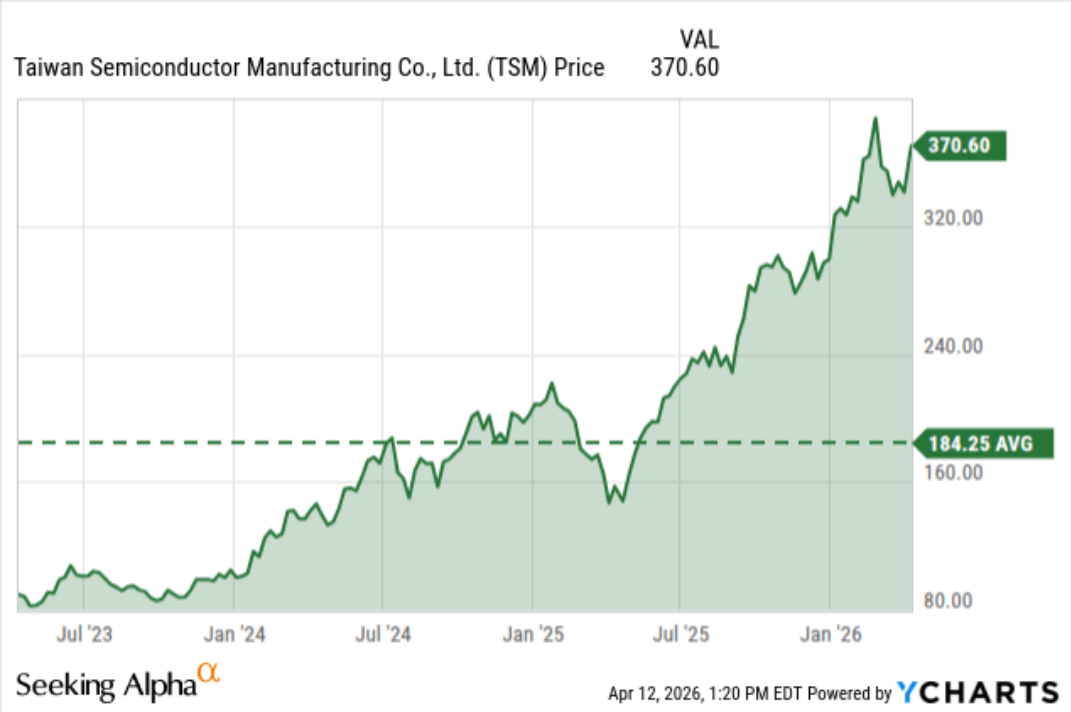

YCharts

Because Everything Is Sold Out

If I had to pick a single factor that is going to drive TSMC's stock over the next three years, it would be that TSMC has become a monopoly on the only production process that can build every single chip that matters for AI. And everything else in this piece is me trying to convince you that the market has not yet fully repriced the stock to reflect what that actually means.

Because the demand trajectory is compounding every year at a very high pace, and the supply to resolve it has a lead time of four to five years. Their 3nm nodes are going to run above 100% effective utilization in the second half of this year. You can go above 100% of theoretical capacity by pushing existing lines harder, rerouting process layers, and pulling every last wafer out of tools that were never designed to run this hot. They are doing all of that right now. And it is still not enough.

At the same time, N2 entered mass production in Q4 of last year. Fab 20 in Hsinchu and Fab 22 in Kaohsiung started at the same time, which is the first time TSMC has ever ramped a new node at two sites simultaneously. That's why both are fully sold out for 2026 and nearly sold out for 2027.

At the end of 2025, they had roughly 40k to 50k wafers per month of 2nm capacity and about 150k WPM of 3nm. By the end of 2026, they are trying to ramp 2nm to 80k to 130k WPM, and 3nm is targeted at 180k to 220k WPM. It sounds like a lot of wafers. Until you look at the demand. Apple has locked up more than 50% of the 2nm capacity for the A20, M6, and R2. That's before NVIDIA Corporation's (NVDA) migration of their chips to 2 nm, which starts in 2027, and not to mention AMD's demand for MI400 on N2.

Broadcom Inc. (AVGO) has pre-reserved a huge slice of 2nm for Alphabet Inc.'s (GOOG) (GOOGL) TPU and OpenAI. I am actually writing a separate piece right now on Broadcom and the position that they have built for themselves because of the capacity they secured from TSMC years in advance. So make sure to follow and stay tuned for that one.

An important thing to note here is that I also wrote in my Micron Technology, Inc. (MU) piece. HBM4 base dies alone would add hundreds of millions of units to TSMC's N3/N5 load. Because the HBM4 base die is made

on advanced FinFET on N3 or N5. So every HBM4 stack going into a Blackwell Ultra or a Rubin is basically eating TSMC's advanced node wafers twice.

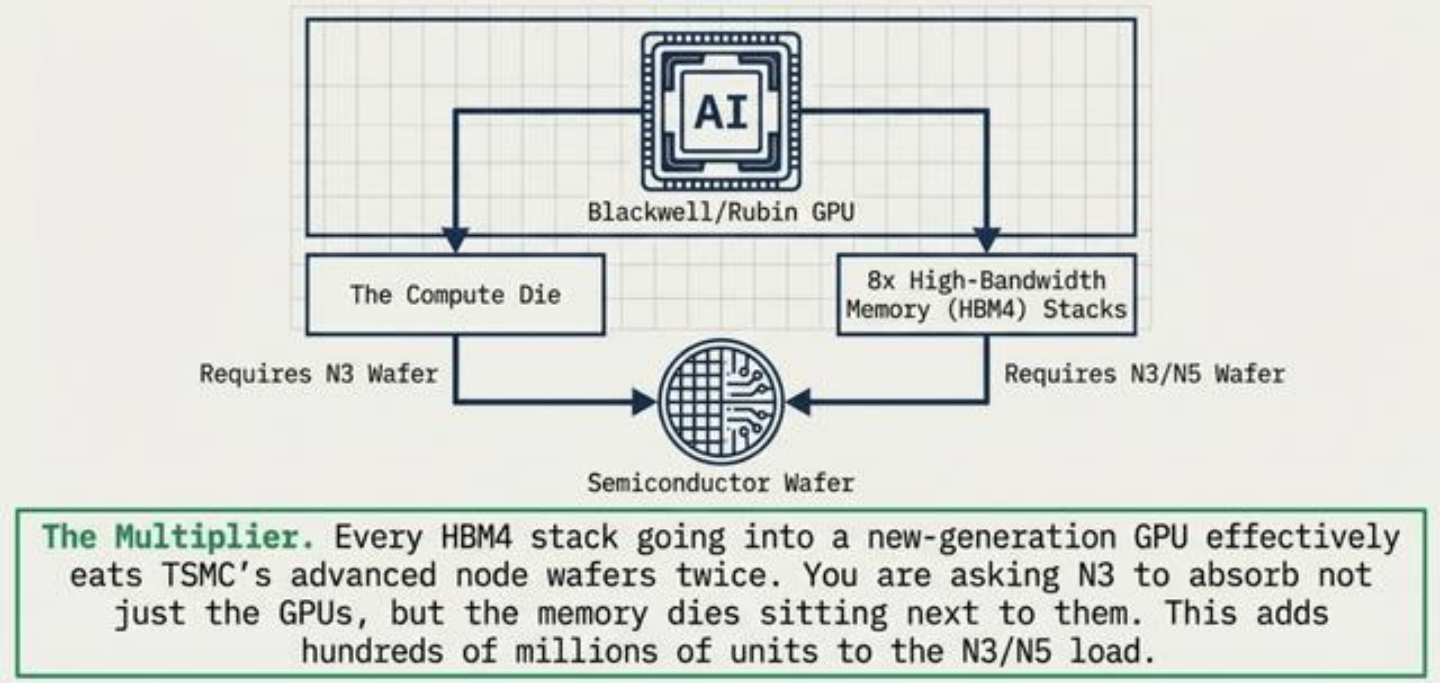

Multiplier Effect of Next-gen Chips (Author)

Once for the GPU compute die, once for the HBM base die on it. And every single Rubin GPU needs 8 or more of them stacked on it. You are basically asking N3 to absorb not just the GPUs but also the memory dies that go next to the GPUs.

CoWoS Is Also Overcrowded

And even once you get the wafer, you still have to package it. CoWoS advanced packaging is the pipe that connects the GPU die to the HBM stacks. TSM has ramped up CoWoS from roughly 35k wafers per month in late 2024 to 75 to 80k by the end of last year and is targeting 120 to 130k by the end of 2026. And it is still not enough.

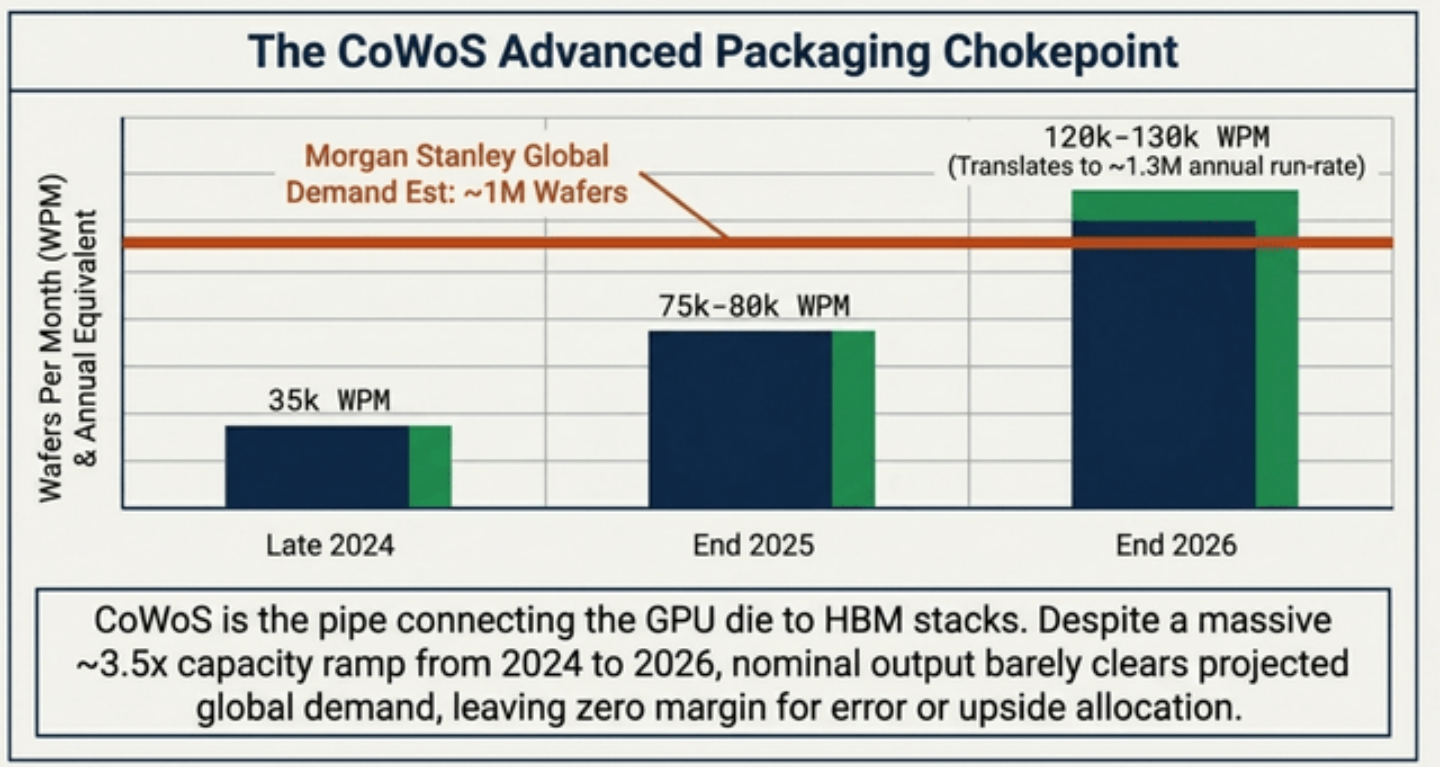

CoWos Ramp (Data Complied By Author)

Morgan Stanley has global CoWoS demand at roughly 1 million wafers in 2026. TSMC will ship, let's call it 1.3 million at the 120k WPM run rate by the end of the year, so nominally the numbers look close.

But if you look at the bookings:

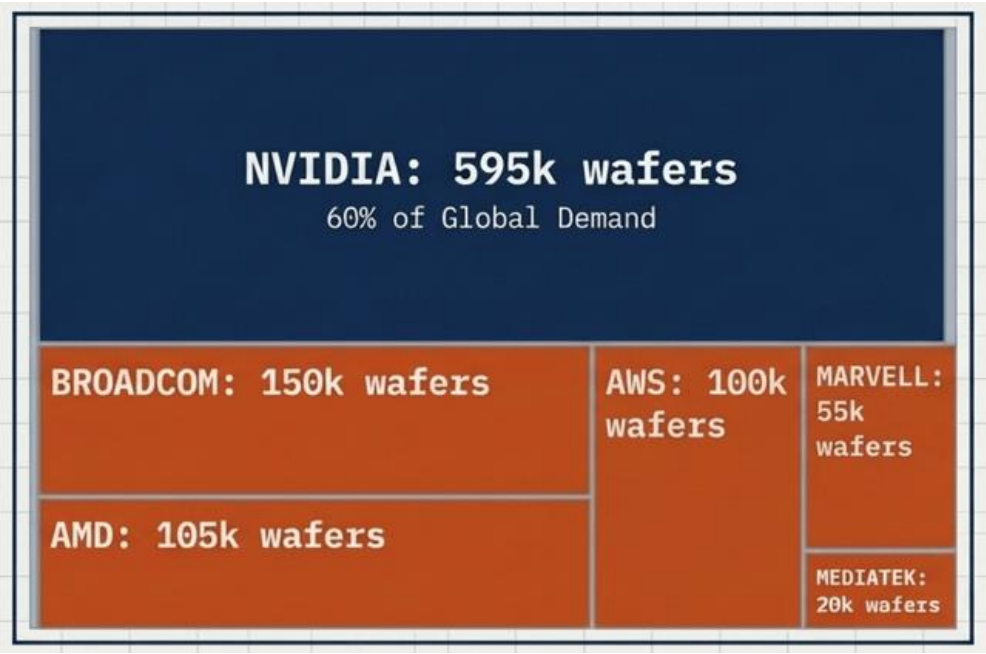

TSM CoWos Bookings (Author)

The math tells you the top 3 alone consume 85% of the output. Google already had to cut its 2026 TPU production target from 4 million units to 3 million because it could not get enough CoWoS slots.

What I'm Actually Watching On April 16

Okay, so with all of that as context, let's get into what I will be looking for in the upcoming print. TSMC has already guided Q1 to $34.6 billion to $35.8 billion in revenue, or about 38% YoY at the midpoint, with a gross margin of 63 to 65%. Consensus sits at roughly 35 billion. I think they will most likely beat it. But here is what I am most interested in:

First, the revenue guide for Q2 and any update to the full year. They guided 2026 to grow “close to 30%” in US dollar terms back in January. Since then, we've had Jensen's GTC keynote, where he casually dropped $1 trillion in orders for Blackwell and Vera Rubin through 2027. We have had Nvidia taking over as their number one customer, from Apple, by revenue. And we have had Broadcom's AI revenue running at a $33 billion annualized rate with a $73 billion backlog, all of which has to be built somewhere, and that somewhere is TSMC. If the 2026 guide does not get revised upward, I will be genuinely surprised. If it gets revised upward by a meaningful amount, that is the bullish signal.

Second would be the gross margins, but specifically the Q2 and full-year guide. The Q1 guide of 63 to 65% already bakes in the overseas ramp dilution and the early N2 drag. What I care about is the next quarter. If TSMC guides Q2 flat or up, it means the pricing power on N3P is offsetting the Arizona drag and that Apple has most likely swallowed the 2nm wafer price hike.

If Q2 guidance implies sequential margin compression, then I am looking for a valid explanation from them. If it's because of the overseas ramp, then that’s a good sign it will start to normalize in the next couple of quarters. But I do not expect this to happen, because when N3 pricing is going up 5 to 10% and N2 is booked out to 2028, the mix effect alone should outrun the dilution. It will be the single most important line in the release.

Third, and this is where you have to read between the lines on the call, I want to hear management talk about utilization and the booking curve. Last quarter, the CEO used the language that “focused more on short-term output” for N2. This quarter, I expect that language to get stronger, not weaker.

Specifically, I am listening for

Any update on how much of 2027 N2 is already allocated.

How they talk about N3X and whether the premium variants are getting pulled forward or not.

If management signals that 2027 N2 is sold out, that would be a huge forward indicator because it means every new AI chip design decision being made right now for the 2028 product is either getting denied or being pushed to N3 with a price hike.

What This Demand Would Look Like Financially

I have walked you through the supply/demand math. Now let's talk about what all of this is actually worth. At the current price, TSM trades at roughly 25x forward P/E and 14x forward EV/EBITDA, with a PEG ratio of 0.83 on a forward non-GAAP basis.

TSM Valuation Grade D-

	Sector Relative Grade	TSM	Sector Median	% Diff. to Sector	TSM 5Y Avg.	% Diff. to 5Y Avg.
P/E Non-GAAP (TTM)	D+	34.80	21.62	60.95%	24.70	40.89%
P/E Non-GAAP (FWD)	C	25.04	21.78	15.01%	22.11	13.29%
P/E GAAP (TTM)	B-	29.83	31.72	-5.98%	-	-
P/E GAAP (FWD)	B-	25.52	29.87	-14.57%	22.25	14.68%
PEG GAAP (TTM)	B-	0.64	0.92	-30.37%	-	-
PEG Non-GAAP (FWD)	B+	0.83	1.34	-37.97%	-	NM
EV / Sales (TTM)	D	12.94	3.14	311.78%	8.48	52.71%
EV / Sales (FWD)	D	9.95	2.81	254.53%	7.55	31.79%
EV / EBITDA (TTM)	C+	18.86	17.33	8.85%	12.47	51.23%
EV / EBITDA (FWD)	C	14.05	12.87	9.10%	11.08	26.79%
EV / EBIT (TTM)	B-	25.46	25.78	-1.25%	18.92	34.53%
EV / EBIT (FWD)	C+	18.19	17.41	4.52%	16.93	7.47%

SA

I think you can safely say that TSMC is in the range of fairly valued to slightly undervalued. Because they are growing twice as fast as the sector, at margins five to seven times higher, with a forward P/E that is only 15% above the sector median. On a PEG basis, it is trading at a 38% discount to the sector median.

TSMC did $122.4 billion in 2025 revenue, up 35.9% YoY in USD terms. They guided 2026 to grow "close to 30%" in US dollar terms despite FX headwinds, and I think that guide may get revised upward on April 16th for reasons I walked through earlier.

So, here is my revenue forecast through 2030:

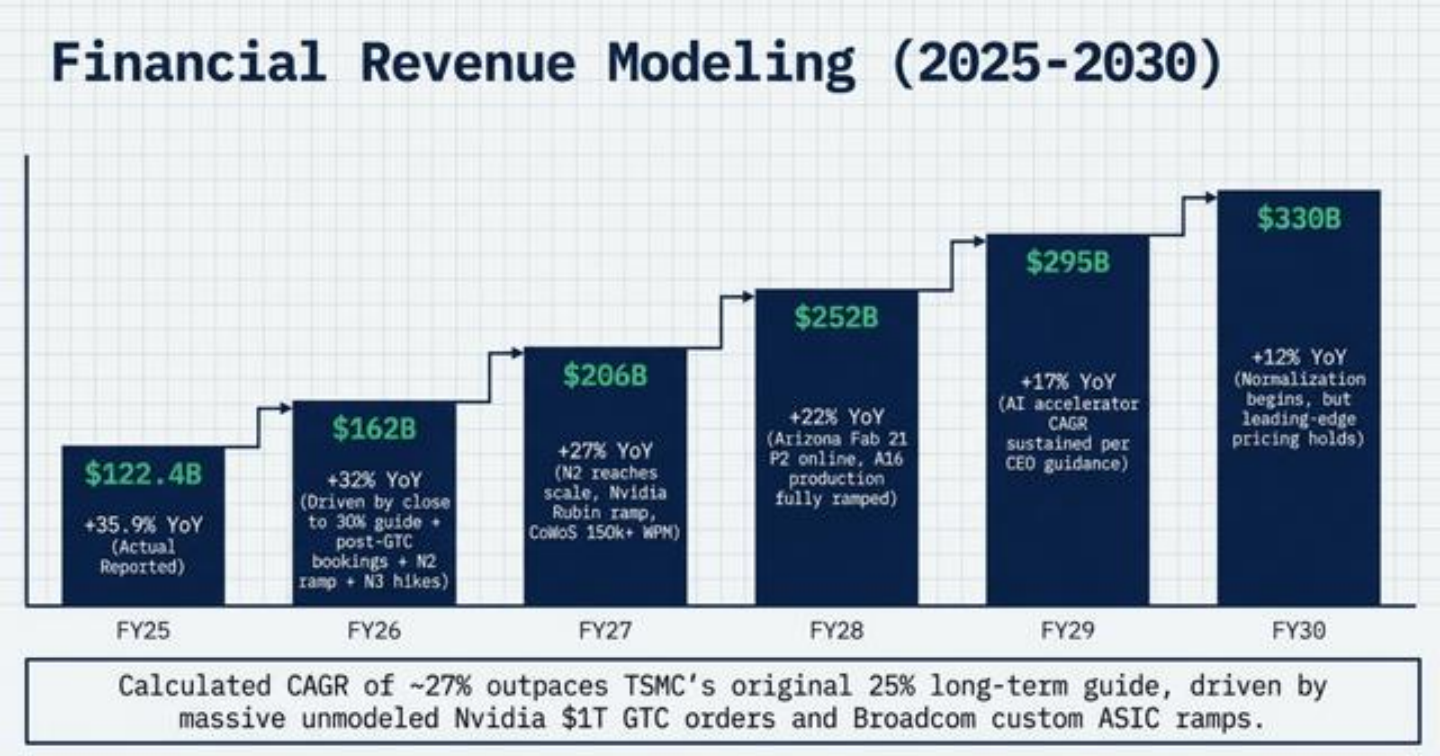

My Revenue Forecast for TSM (Author)

A few things to note here. Their own guide was 25% CAGR in USD terms for the 5-year period starting in 2024. And AI accelerators specifically compounding in the mid-50% range for 2024 to 2029.

My revenue forecast works out to a CAGR of roughly 27%, which is slightly above their number. The reason for that is that Nvidia's GTC orders of $1 trillion through 2027 were not in their long-term guide when they first issued it. Moreover, Broadcom's custom ASIC ramp for OpenAI and Anthropic is added to what was modeled, and there are 5 to 10% price hikes on N3 and N5 already in motion.

Now, if you apply a normalized 50% net margin on that FY27 revenue number of $206 billion, you get net income of roughly $103 billion. Divide that by 25.9 billion diluted shares (about 5.2 billion ADR equivalents), and you are looking at FY27 EPS of around $19.8 per ADR.

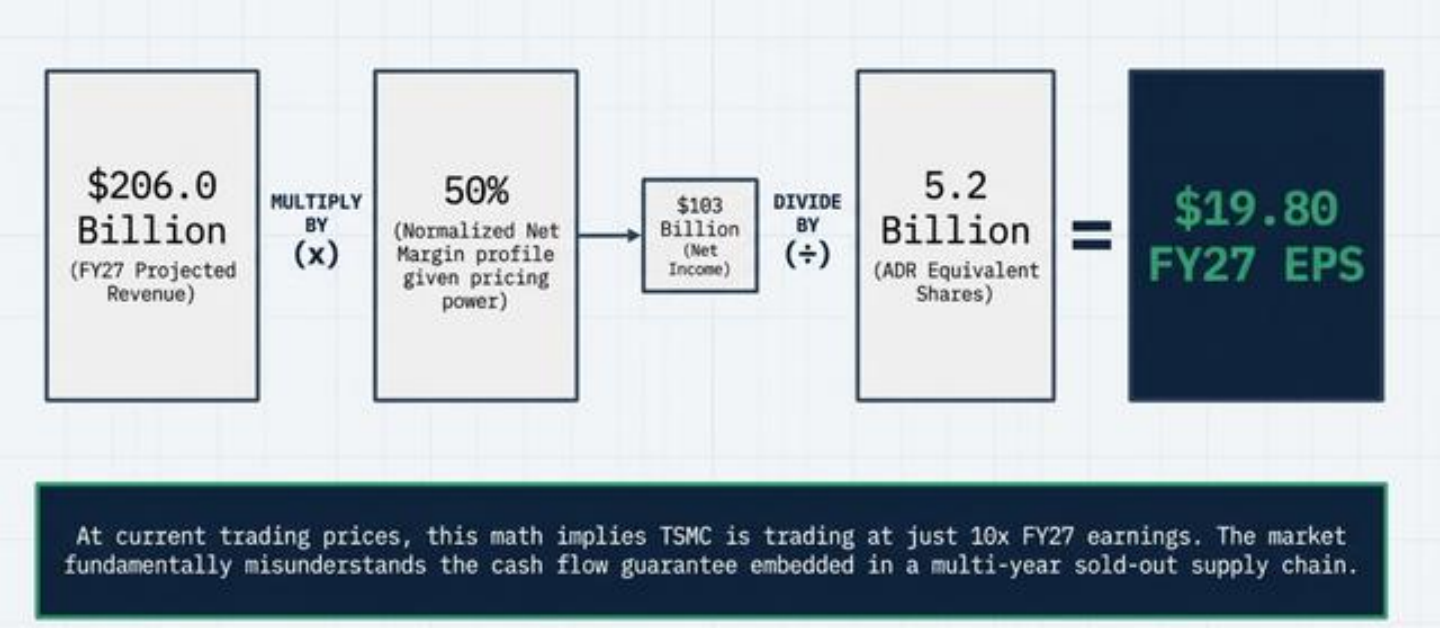

Valuation Calculation (Author)

The stock right now is trading at just 10x FY27 earnings. And keep in mind that Nvidia trades at 28x forward earnings. Broadcom trades at 31x.

What Could Go Wrong

And I think the reason for this valuation discount is a Taiwan-specific risk. A Taiwan Strait event is a very low probability but high risk and unmodelable. This risk is what keeps them from reaching Nvidia or Broadcom levels and, honestly, probably always will. But in my opinion, the valuation still has enough room to run because of the sheer scale of demand.

The risk to them is not the competition. Samsung's advanced node yields are still not competitive. Intel 18A is a credible process for Intel's own products, but not a threat to its external foundry business.

The biggest risk apart from a Taiwan Strait event is a what-if question. What if the AI capex cycle breaks and hyperscalers are forced to a capex cut because of a macro recession or financial crisis? Although the probability of that is very low, there is still a risk associated with many of these AI buildout investment theses.

Final Thoughts

On April 16, I will be watching the revenue guide, the Q2 margin guide, the language around N2 bookings, and the CoWoS ramp update. If management does what I expect them to do, which is raise the full-year revenue outlook, hold margin guidance flat, and reiterate that the capacity is booked out beyond 2028, then the stock is materially undervalued at the current valuation.

Micron: The Peak Everyone Sees Isn't There

April 13, 2026 Future Stack Investment

Summary

- Micron Technology, Inc. remains a Buy as AI-driven memory demand, HBM strength, and margin expansion underpin a robust structural setup.
- Q2 FY25 saw revenue surge over 3x YOY to $23.86B and non-GAAP gross margin jump to 74.9%, reflecting accelerating AI adoption.
- Supply constraints, complex HBM production, and longer-term agreements support pricing power and visibility, while capex timing delays near-term supply risk.

JHVEPhoto/iStock Editorial via Getty Images

Elevator Thesis

Micron Technology, Inc. (MU) remains hot, as it is at the heart of the AI infrastructure development. Interestingly, change in the AI infrastructure continues to redefine the cycle and is gradually pushing Micron towards a more systematic outlook.

I emphasized the AI memory demand, high bandwidth memory ("HBM") strength, and margin expansion as the core pillars in my previous coverage. Indeed, these drivers have held firm, and conviction has enhanced with the course of time. Every earnings print, in turn, supports the same structural setup with increased visibility.

Additionally, Micron had an outstanding Q2 FY 25. Its revenues and gross margins ramped up. The quality of such growth was of equal importance as its scale. The AI-related memory did the heavy lifting and tilted the mix more to higher-value segments.

Despite this, the stock has not been on a smooth path since my last writing. After slipping into the $360-$370 range, it has found its footing again around $420 (as of writing). In fact, the market is striking a balance between good execution and the renewed cycle anxiety.

The price action simply means that the investors are resetting positions around their expectations.

Moreover, tools such as Google's (GOOG, GOOGL) TurboQuant also bring up a concern that smart memory usage may alleviate the demand growth in the long run.

Nonetheless, the bigger picture looks pretty different.

The AI workloads continue to grow. Instead of decreasing the total demand, they consume efficiency improvements and convert them into even higher usage.

Memory supply will ultimately react at some point, and the cycles will come under pressure. Until that time, demand continues to outrun supply.

Overall, the current setup keeps supporting a Buy rating.

Demand is Structural Across the AI Stack

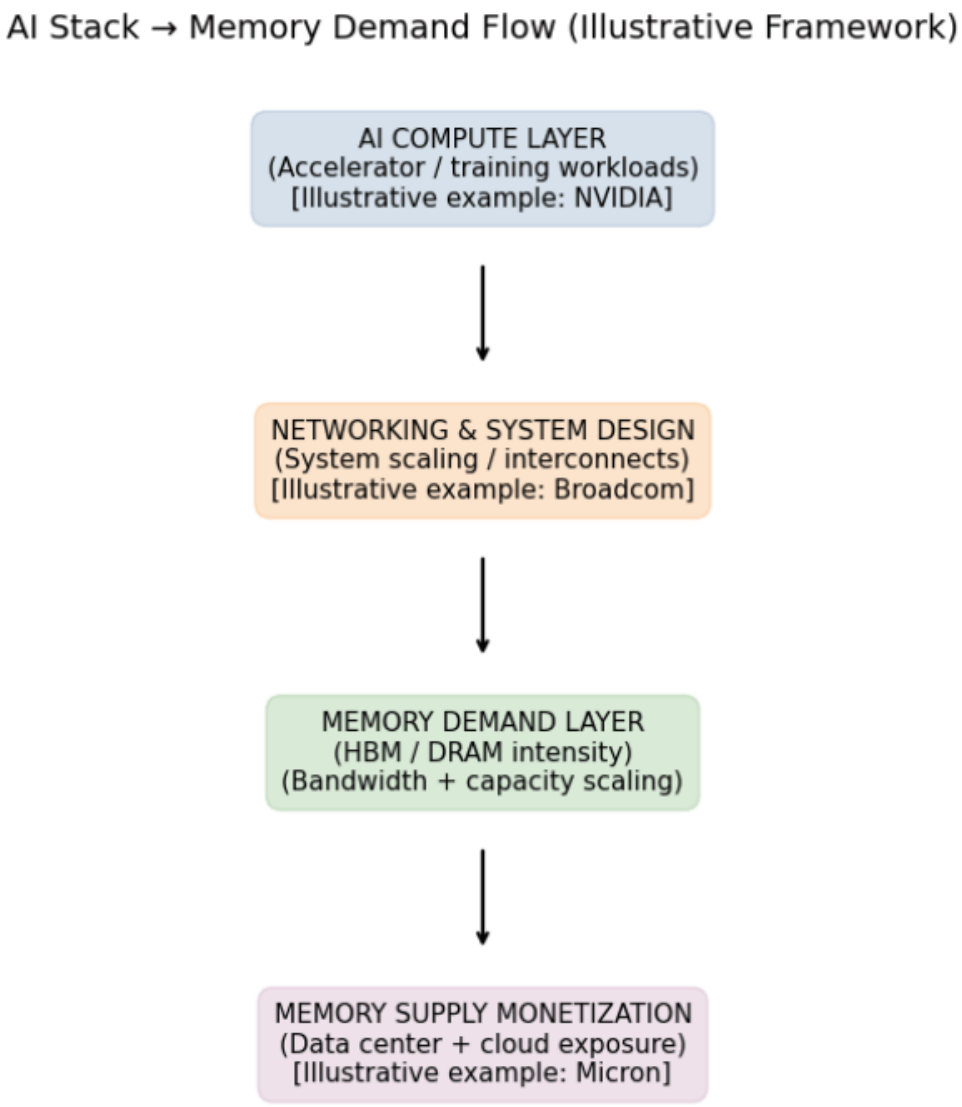

Illustrative AI stack (Author)

The idea of a peak cycle often comes from looking at Micron in isolation. However, the concept gets more meaningful when we begin with the AI stack and work backwards.

It starts with AI compute, transitions through networking and system design, and finally appears in DRAM and HBM demand.

For some color, let's see what's happening higher up the stack.

At the compute layer, Nvidia (NVDA) posted Q4 FY 26 revenue of $68.1 billion. It includes the data center revenue of $62.3 billion, up 75% YOY. At such a pace of deployment expansion, HBM and high-density DRAM scale with it. The reason behind this is pretty simple, as they are at the core of the AI memory bandwidth needs.

There's the same momentum in custom silicon and networking, too. For perspective, Broadcom (AVGO) generated AI revenue of $8.4 billion in Q1 FY 26 at the platform level. It's a whopping 106% upswing from Q1 FY 25. A boom in custom silicon and networking generally raises the amount of memory content per deployed rack.

Well, all of this flows back into Micron. Its core data center revenue rose over two-fold to $5.69 billion (sequentially), and operating margins clocked in at 67%.

Additionally, its Cloud segment brought in $7.75 billion in revenue with 74% gross margins, as compared to 66% in Q1.

Furthermore, Micron's Q2 .25 made it clear that this demand is no longer a forward-looking story. The revenue surged more than 3x YOY to $23.86 billion. Alongside, its non-GAAP gross margin crushed the guidance and rose to 74.9% from 56.8% in Q1.

This kind of margin expansion generally doesn't happen towards the end of a phase. If anything, it appears more that demand is yet to catch up with supply.

The peak cycles usually exhibit cracks in demand visibility, but it's not the case right now.

Currently, the revenue is growing faster, the margins are increasing, and the mix is becoming more AI-driven.

Metric	Value
KV-cache memory reduction (TurboQuant)	Up to 6x decrease
Inference speedup (TurboQuant)	Up to 8x increase
Context window size (Gemini 3 Pro / Llama 4 Scout)	~10 million tokens
SK hynix HBM shipment growth	>2x YoY
Samsung semiconductor operating profit (Q1 2026)	57.2 trillion won
Micron FY26 Q3 revenue guidance	$33.5 billion
Samsung semiconductor contribution to profit surge	~8x increase

Google's TurboQuant poses a valid question for the memory cycle. It aims at the key value cache usage in long-context inference loads.

It can slash the memory needs by up to 6x and can deliver a speedup inference of up to 8x. On paper, this appears to be a huge headwind for demand.

When all requests become six times more efficient, it might be assumed that the pressure on memory vendors would eventually decrease.

Yet, the real system-level data points otherwise. The main driving force here is a mix shift, and companies like SK hynix (HXSCL) and Samsung (SSNLF) are already proving this.

Let's put this into perspective.

SK hynix experienced tremendous growth in shipments of HBM, and the figures have increased by over 2x annually.

In addition, Samsung is anticipating an unprecedented Q1 FY 26 operating profit of 57.2 trillion won. It's roughly 8-9x higher than last year and is largely due to its semiconductor business.

This means that memory per unit of compute is on the upward trajectory while software is becoming more intelligent.

Moreover, the Jevons Paradox is usually triggered by the efficiency gains in AI. When the price of a commodity such as memory footprint decreases by 6x, the usage doesn't drop proportionally. The users would rather intensify their ambition. I can easily see this in context windows.

At the beginning of 2026, Gemini 3 Pro and Llama 4 Scout had already reached the 10 million token limit. These ultra-long contexts necessitate huge resident memory in the GPU, even when compressed.

This change is supported by the financials. Samsung has confirmed an upswing in its DRAM price in Q1 FY 26. Any space saved by decreasing the use of KV cache is consumed by these larger context windows and denser request batches. The system basically fills all the available headroom that is created.

After all, TurboQuant is a marginal adjustment within a system that continues to scale at a breakneck pace. I find the most obvious sign of it at Micron.

Recently, Micron forecasted a Q3 FY 26 revenue of $33.5 billion. This is higher than the annual revenue that it earned in nearly any year before 2025. Efficiency cannot outrun the sheer scale of the AI buildout, with its HBM production already sold out through 2026's end.

The cycle continues to be strongly propelled by the growth of computing power, with optimization being in the back seat.

Supply Risk is Real but Structurally Constrained

Following the Q2, investors paid a lot of attention to the decision to advance the fiscal capex by more than $25 billion.

This was despite the fact that the company guided towards another level of profitability for Q3.

That response takes the immediate fear that it is another peak on the rise.

Q2 also shows silently how this cycle is different. Margins remain robust as Micron moves to higher-value AI memory. That is important, since supply expansion in this cycle is all about creating more complexity-intensive products.

This is where the commodity system begins to fail, and HBM is at its center. It is based on 3D stacking and sophisticated packaging processes, which add actual yield and integration bottlenecks. That is, the output is not linear to capex, as in the case with the older memory cycles.

Such limitations go beyond Micron.

Another common choke point in AI supply chains is advanced packaging. The speed with which the ecosystem can grow is limited by capacity constraints at partners such as Taiwan Semiconductor (TSM).

The market is also changing in terms of its behavior as procurement is stretching out.

According to industry statistics, memory agreements are shifting towards 3–5-year arrangements. Micron is taking the lead in this change with its first 5-year Strategic Customer Agreement. This directly hedges demand visibility as opposed to exposing it to spot volatility.

The last and most crucial constraint is timing. The additions to supply are intentionally back-ended by Micron. Idaho output is projected in mid-2027, and the only year that Singapore's NAND capacity ramps is 2028. The near-term incremental supply is therefore pushed out.

The $1.8 billion acquisition of Powerchip fab is only predicted to contribute any significant DRAM wafer output from the second half of 2027. The lag continues to tighten the market until this year's end.

In simple words, the cycle has become longer, more irregular, and essentially more predictable than any other era.

Valuation Discussion

Micron will no longer enter earnings with low expectations.

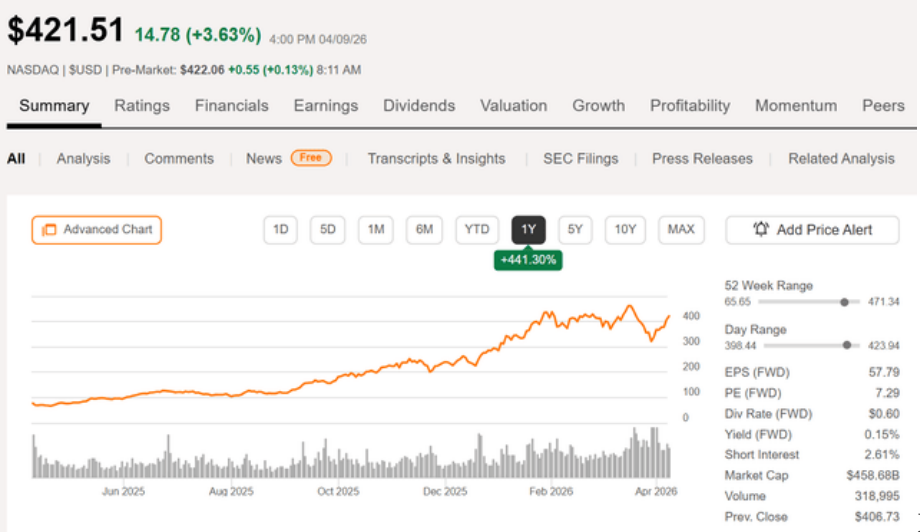

MU stock price (SA)

MU stock plummeted to roughly $377 a few days back, but then it climbed back again to about $420. The 12% move is more about positioning as opposed to fundamentals.

Basically, the market initially traded on uncertainty, then immediately retreated as the wider AI demand story held. This rebound narrows the margin for error.

Annual EPS Estimate

Fiscal Period Ending	EPS Estimate	Forward PE	Low	High	# of Analysts
Aug 2026	57.79	7.29	28.42	64.37	29
Aug 2027	97.94	4.30	44.83	142.48	30

Annual Revenue Estimate

Fiscal Period Ending	Revenue Estimate	Forward P/S	Low	High	# of Analysts
Aug 2026	108.50B	4.38	72.22B	116.93B	33
Aug 2027	164.47B	2.89	91.35B	240.36B	37

NU consensus estimates (SA)

Most of the heavy lifting is carried out by consensus estimates. The earnings are projected to rise aggressively. The FY 26 and 27 EPS estimates are trending towards the high $50s and $100, respectively.

The same is the case with revenue. This means a structural step-change and not a regular cyclical recovery.

Earnings roll over into 2027 and then start to normalize, which is why forward multiples will compress within a short period of time. Micron is currently trading at approximately 7x FY 26 earnings and a bit closer to 4-5x FY 27. This continues to add some level of apprehension regarding durability.

At about $420 today, the stock has already priced in strong earnings momentum, but it has not yet been fully re-rated.

My base case sits in the $425 to $430 range. This remains the most realistic near-term result. It presupposes a robust quarter and optimistic commentary on DRAM pricing and HBM demand. In such a case, Micron will not require a significant revaluation. It just has to hold its existing narrative and allow estimates to catch up.

That said, the bull case between the $440 and $460 range needs more than that. It's required to have greater belief in the fact that supply constraints go further into 2026.

The last drop of $377 means that the buyers are ready to intervene quickly. However, it is also an indicator that much of the near-term optimism must have been pulled forward.

Micron is in a great spot fundamentally, but the stock is moving from a quick recovery into a phase where it needs to prove itself.

This is why I am remaining realistic with my base case at the moment. The bull case will begin to seem more likely as it will become apparent that earnings will only get better.

Risks to Thesis and Takeaway

The greatest risk in this case is the same at its core. Memory is still a cyclical business. The current cycle has been enhanced by AI, but it has not eliminated that nature.

History is a good guide. Micron has experienced steep earnings growth in previous cycles. Interestingly, it has an equal experience of sharp declines afterward when supply returns to equilibrium.

This trend does not really point to a collapsing demand. It just shows supplies returning to normal. So, in such a kind of situation, timing is more crucial than the direction.

Yes, the structure is robust today, but it remains vulnerable to changes in cycles. MU's forward valuation already prices in much strength. The stock is trading at around 7x forward earnings. It is also close to 4.3x EV/Sales, as stability is already in the expectations.

The second risk is supply, though it is postponed. In Samsung and SK hynix, capex in the industry is increasing once again. This doesn't affect the short-run pricing. That said, it would matter more for 2027 and beyond, when the new capacity will be introduced.

Furthermore, gains such as TurboQuant are currently absorbed by AI workloads with increased usage. However, efficiency ceases to be a tailwind and instead becomes a constraint to memory intensity when spending on hyperscalers decelerates. Such a change would directly affect the demand per workload.

Regardless of these risks, we have to count on positioning. Micron has a good underlying demand. This can be seen as it corrected sharply and then rebounded again to around $420.

Right now, Micron is no longer a secret success story, as the market already sees it. Now, the question is to what extent the AI memory cycle can run before supply catches up.

On the whole, MU's valuation remains conservative, and earnings continue to go up. The opportunity lies in the gap.

The growth of earnings is continuing, and a complete re-rating is yet to come.

It's sensible to stay with the trend if you are already holding. Meanwhile, if you are not entirely positioned, it appears like a decent entry point.

BE Semiconductor: The AI Packaging Bottleneck Play Investors Overlook

April 7, 2026. Dhierin Bechai

Summary

- BE Semiconductor remains a Buy, with 35% upside and a $301.48 price target despite premium valuation and near-term hybrid bonding uncertainty.
- Q4 revenue rose 25.4% sequentially and 8.5% YoY, driven by strong 2.5D packaging and photonics demand, while orders surged 43.3% sequentially.
- BESIY's book-to-bill ratio consistently above 1 signals robust 2026 demand, with hybrid bonding positioning the company as a technology front-runner.
- Risks include weak ex-AI markets and potential hybrid bonding delays, but BESIY's innovation, margin expansion, and strong cash generation support long-term value.

BE Semiconductors, also known as Besi (BESIY, BESVF), has gained 11.5% since my last report, putting its strength as a die-attach and packaging expert for chips on display, outperforming the S&P 500's 5.2% loss. The increase in the stock price comes despite lower year-on-year sales and doubts about hybrid bonding techniques. In this report, I discuss the company's financial results and the color provided for 2026, and I update my stock price target with a continued buy rating as its strong expertise and supply chain constraints position Besi strongly to capitalize.

Innovation And Mainstream Market Weakness Masks High AI Packaging Demand

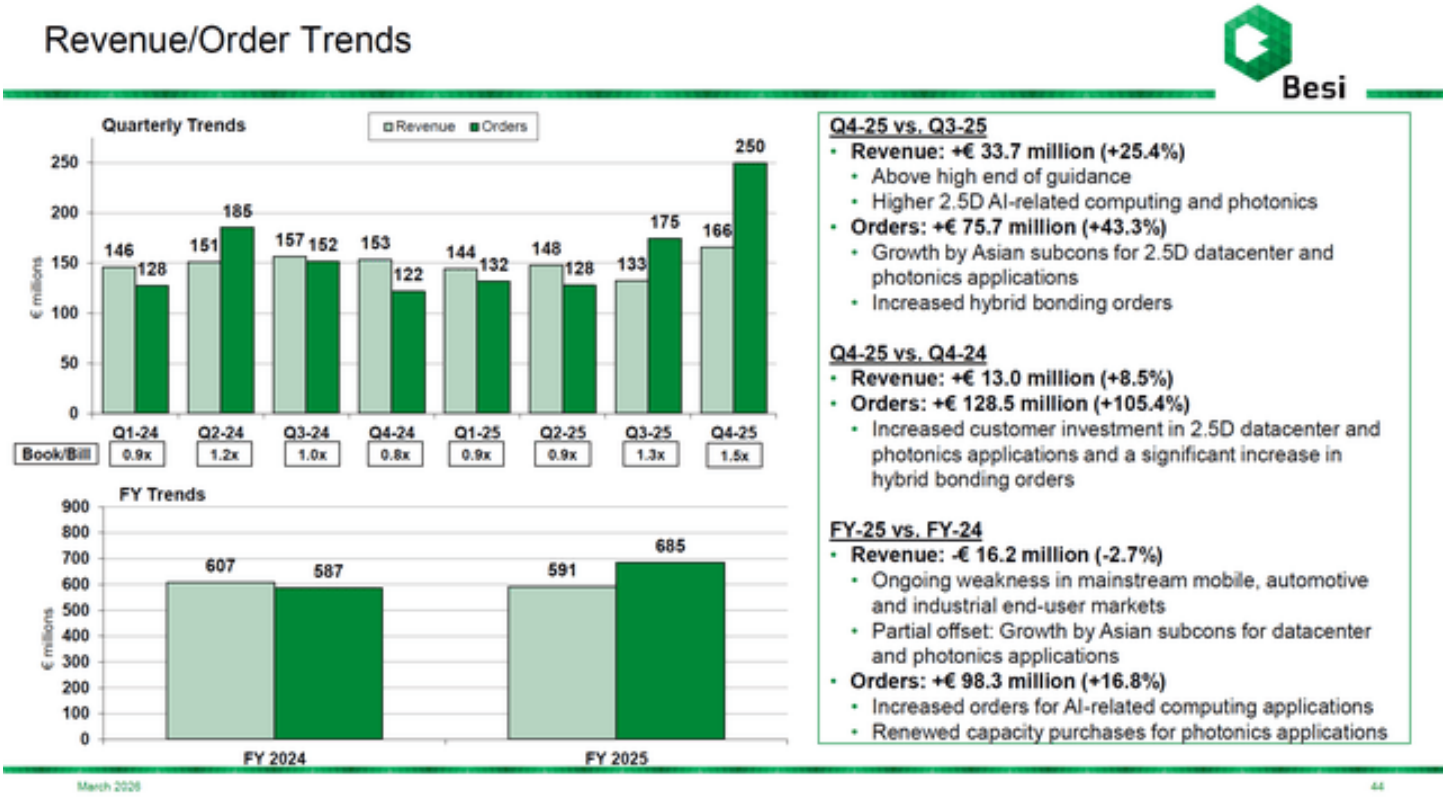

Besi

The revenue and order trends, in my view, are encouraging. Fourth quarter sales rose 25.4% to €166 million and 8.5% year-on-year. The growth exceeded the company's expectations, driven by 2.5D packaging demand as well as photonics. 2.5D is a chip packaging approach where multiple chips are placed side by side rather than stacking, which is truly 3D packaging. For the full year, revenues declined 2.7% to €591 million, driven by strong AI growth but weakness in other end markets.

Orders for the fourth quarter increased 43.3% sequentially and more than doubled year-over-year to €250 million, indicating a book-to-bill ratio of 1.5x. Since Q3 2025, we have seen the book-to-bill ratio consistently above one. That is an indication of customer planning translating to orders in support of higher production in 2026, with bottlenecks such as packaging being one of the pain points that companies are investing in. For the full year, the book-to-bill ratio was 1.16x, showing that the strong book-to-bill was not something from the last two quarters only, but the book-to-bill ratio was strong throughout the year.

Year-on-year, revenues were down 3% to €591.3 million, and that shows the cyclicality of the business that AI is not yet fully offsetting. However, gross margins declined from 65% to 63%. This is not driven by pricing pressure but by the volume weakness in the mainstream markets, such as mobile, auto, and industrial.

Furthermore, hybrid-bonding, which should be the company's future growth and margin driver, is still in its early stages. Operating income declined 11.5% to €173.1 million, driven by the same items that drove the gross margins lower and almost fully reflecting the lower gross profits translated to the bottom line. EBITDA declined 7.8% to €206.8 million, with margins declining from 37% to 35%.

Revenue Growth Set To Continue

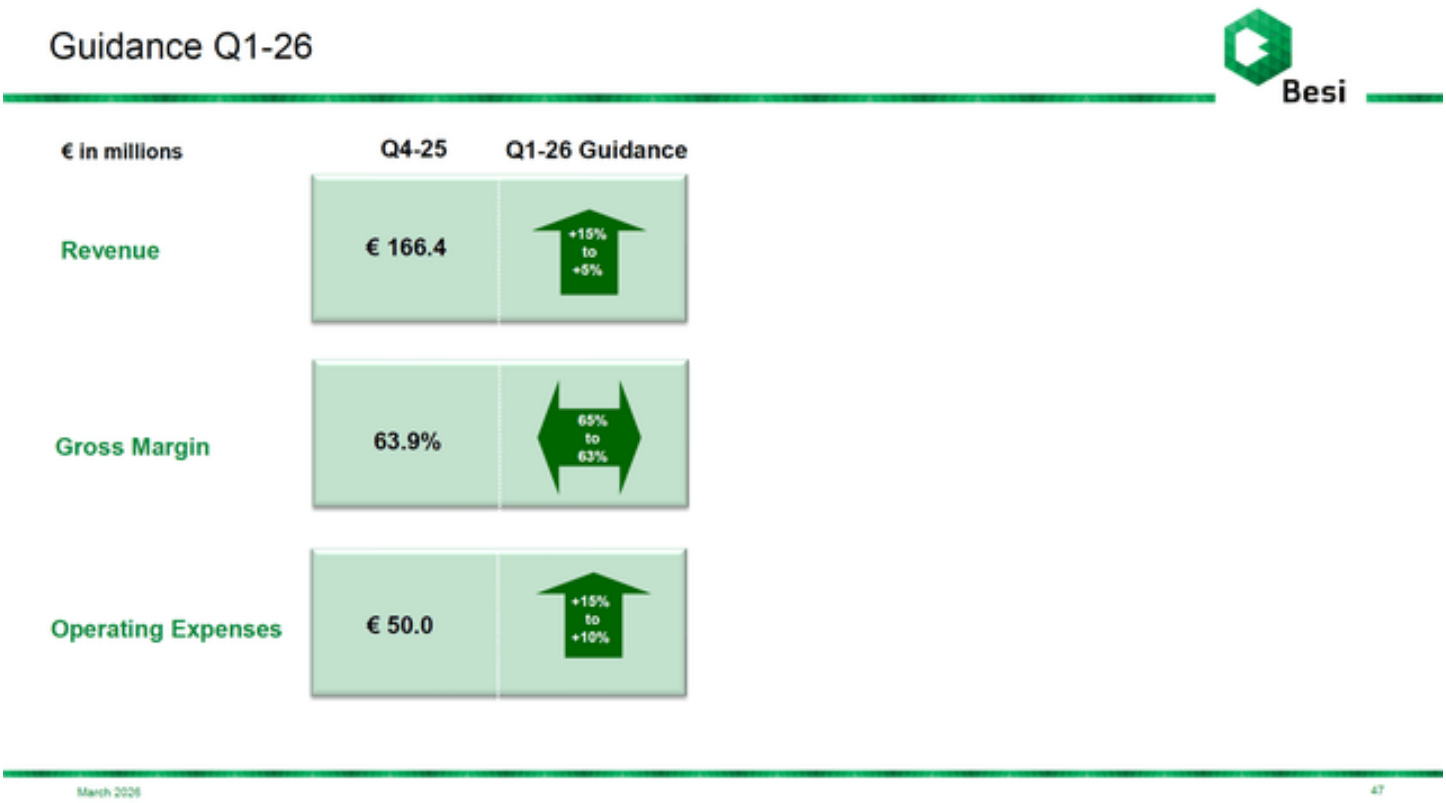

Besi

For the first quarter of 2026, the company expects 5%-15% revenue growth underpinned by strong order intake in the prior quarters. Gross margins are expected to be in the 63%-65% range, with operating expenses rising 10%-15%. At the midpoint, this would suggest EBIT margins to come down further but with 8% EBIT growth. It likely reflects the same hybrid bonding and weak ex-AI market environment as seen in the prior quarters. The sequential revenue growth is somewhat limited against the strong Q4 order intake, and that is because many of the orders were placed towards the end of the quarter. Including lead times, this makes it impossible to turn these orders into sales in the subsequent quarter.

Hybrid Bonding Concerns And Take Over Rumors Contradict Each Other

Early in March, Besi stock plunged 15% as artificial intelligence was said to consider a slowdown on the adoption of hybrid bonding and instead ease requirements on the thickness of memory chips. The news comes at a time when potential customers are starting qualification of the hybrid bonding process in the second quarter to make a decision by year-end on whether this process technology will be used for HMB 4E or for future generations of HBM chips.

It seems more that AI chip manufacturers are contemplating a slowdown to give other companies the chance to catch up on hybrid bonding techniques. For Besi, as a frontrunner, it would not be good news in the near term,

but it also does show how far ahead the company is with its hybrid bonding technology. So, a delayed adoption would be disappointing for the company, but it seems that the question regarding hybrid bonding is not whether it will be used but rather when it will be used.

Somewhat interesting is that not much later reports emerged that there are parties interested in acquiring BE Semiconductor, and I believe that really shows how far the company is ahead with its hybrid bonding and established on the current thermos-compression bonding technique.

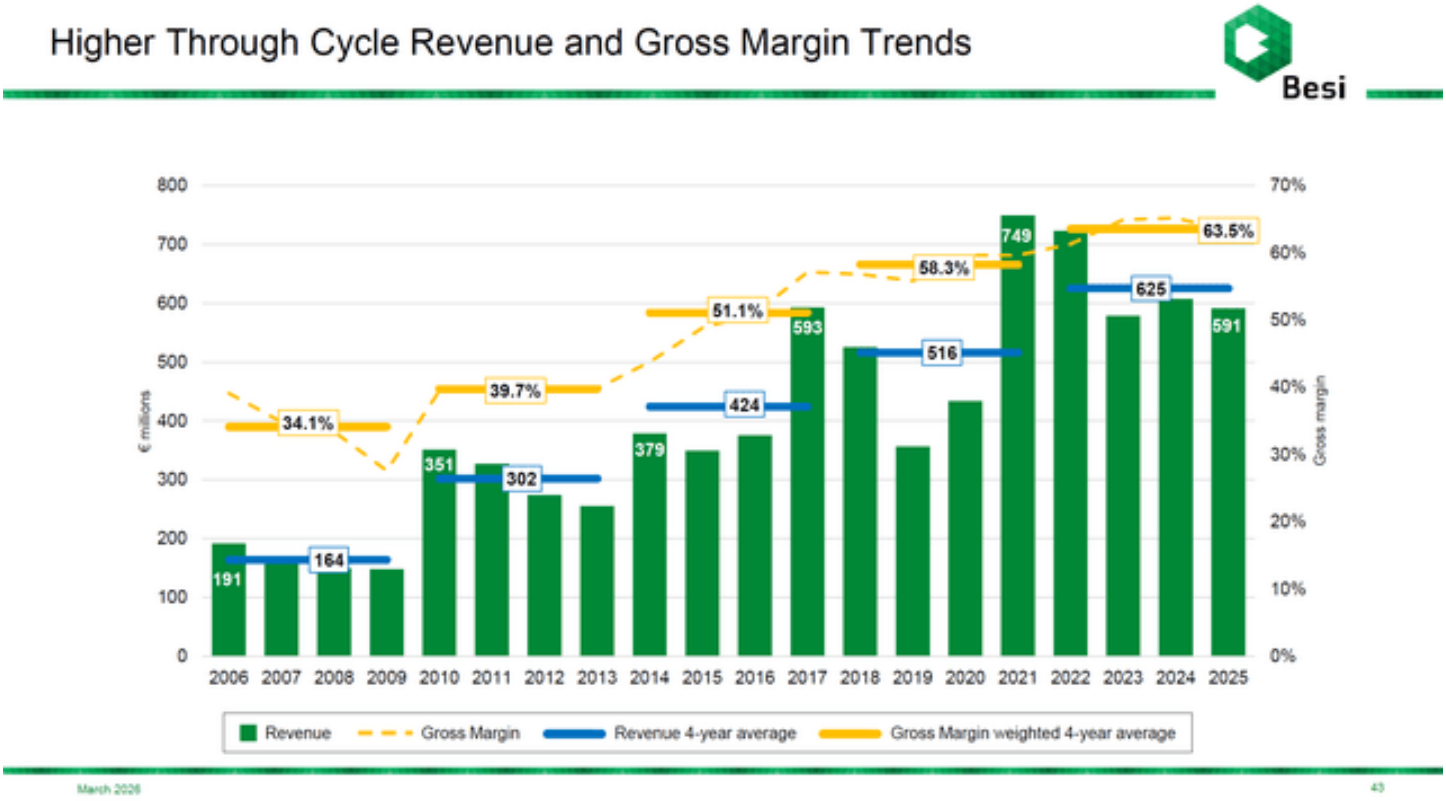

Besi

One of the best slides that I have ever seen is the one above. It shows that regardless of technology injection points and the cyclical nature of the business, Besi has been executing a consistent path of growing its average revenues throughout the cycles along with substantial step-ups in the margins, and I believe that for investors, that is all you need to know. Throughout cycles, the company is capitalizing on demand through innovative solutions.

BE Semiconductor Stock Remains Attractive Despite Premium Pricing

To determine multi-year price targets, The Aerospace Forum has developed a stock screener that uses a combination of analyst consensus on EBITDA, cash flows, and the most recent balance sheet data. Each quarter, we revisit those assumptions and update the stock price targets accordingly. In a separate blog, I have detailed our analysis methodology.

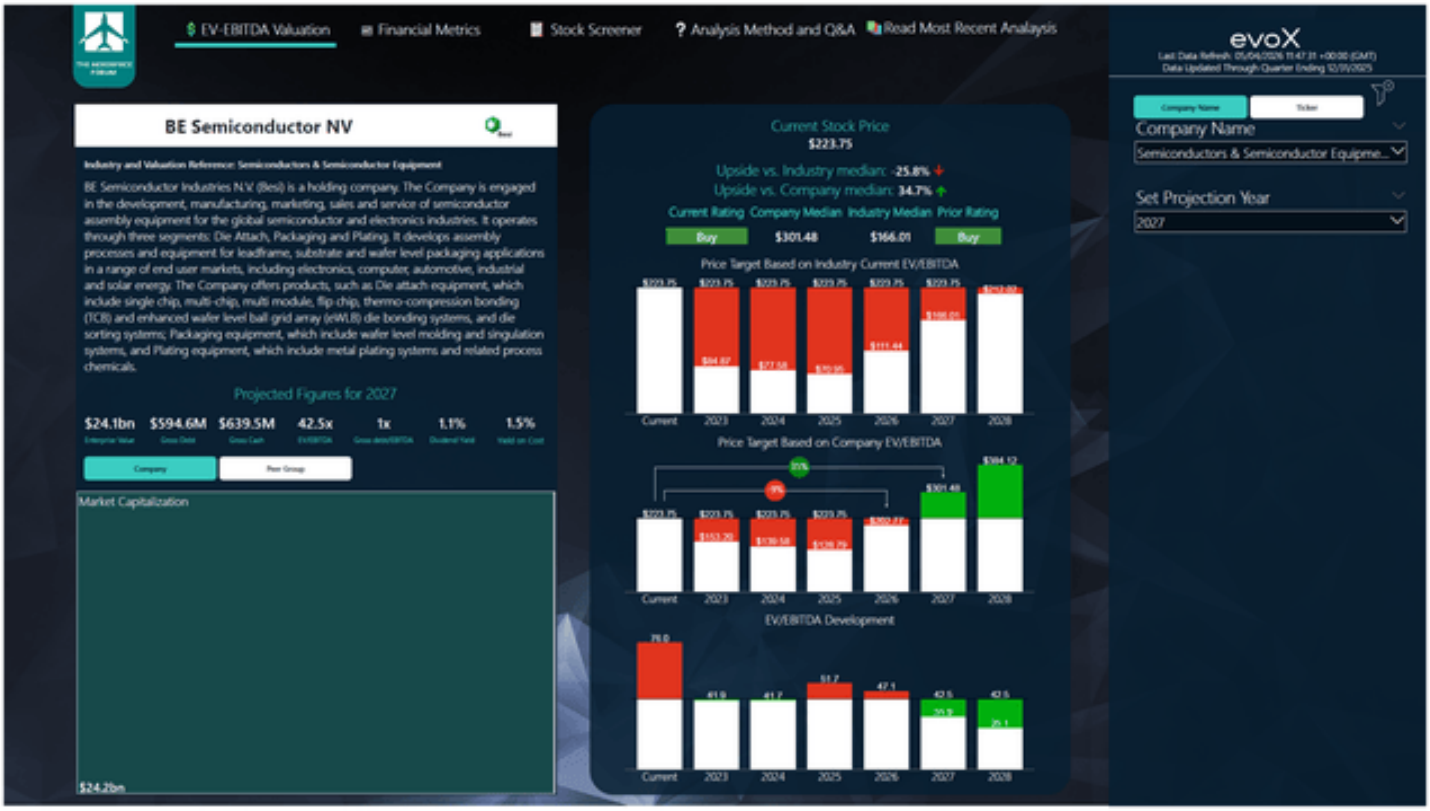

The Aerospace Forum

BE Semiconductor stock is definitely not cheap, but I believe the stock is just slightly overvalued against the expected earnings for this year and offers 35% upside for next year with a $301.48 price target. EBITDA has been revised down by 6.6% for this year, likely driven by the weak ex-AI market, and by 3.5% for next year. Free cash flow estimates declined by 4.8% for this year but were revised up by 3.5% for next year. Sales growth is expected to be 33% annually through 2028, with 44.3% growth in EBITDA and 56.2% growth in free cash flow. With these growth figures, any acquisition offer that Besi may receive needs to be exceptionally good. There are some risks, and those are the ex-AI market and possible delays in hybrid bonding adoption, but that would be more driven by Besi's competitors being behind on developing solutions that could compete with Besi rather than Besi falling behind.

The Aerospace Forum

Even without some margin pressure from hybrid bonding and the weak mainstream market, the EBITDA margins should improve substantially with extremely strong cash flow generation with lower capital intensity. Even with the modeled dividend payments and share repurchases, the cash position is expected to substantially increase in 2027 and 2028, providing an opportunity for Besi to increase its share repurchase activity beyond levels currently modeled.

Conclusion: BE Semiconductor Remains Attractive

While the weak market conditions and potential delay of hybrid bonding adoption pose a risk to BE Semiconductor, I believe that the stock remains attractive. Through each cycle, the company has managed to improve its revenues and margins, and the company is seen as a frontrunner in hybrid bonding. The fact that companies in the semiconductor equipment industry are considering purchasing the company speaks volumes to the key positioning of BE Semiconductor, and whether the company ends up being acquired or operates as a stand-alone business, I believe there is substantial value to be generated as the company uses innovation as a cornerstone to deliver shareholder value.

Moor Insights & Strategy

Sandisk: Strong Fundamentals With Visible Growth Path

March 31, 2026 ZK Research

Summary

- Sandisk is exceptionally well-positioned in the NAND Flash memory market following its spin-off from Western Digital.
- Competitors are reallocating capacity to high-margin HBM, tightening NAND supply and enabling SNDK to expand both margins and market share.
- The Kioxia joint venture secures fixed costs and stable wafer access, supporting margin expansion during surging AI-driven storage demand.
- Recent concerns over Google's TurboQuant are misplaced; the technology does not threaten NAND demand and may ultimately boost storage needs.

Getty Images

Introduction

Sandisk (SNDK) has soared since its spinoff from Western Digital last year, climbing as much as 1,851% by March 24th. After a run like that, investors naturally get cautious and start jumping at every piece of concerning news. Throw in a macro crisis and a sensational blog post hyping a shiny new technology, and you've got the perfect recipe for an eye-popping 11% single-day drawdown.

Sure, you could say we shouldn't read too much into one day's price action - and you'd be right. Still, after wading through the investor reactions, I felt it was worth stepping back to clarify a few things. Without a solid grasp of the fundamentals and the bigger picture in the memory market, it's easy to draw the wrong conclusions.

I rate Sandisk a Strong Buy. The company is exceptionally well-positioned in the NAND flash market at a time when the other dominant NAND Flash producers are deliberately constraining NAND wafer output to prioritize higher-margin HBM capacity. Combined with Sandisk's locked-in, low-cost wafer supply through its extended joint venture with Kioxia, this dynamic should drive both margin expansion and market share gains. Meanwhile, Google's TurboQuant announcement poses zero threat to NAND demand; on the contrary, more efficient AI models are likely to accelerate data generation and storage needs. At current levels, the risk/reward strongly favors the upside.

In this article, I'll walk you through the memory market, Sandisk's place in it, and what Google's TurboQuant algorithm might actually mean for memory stocks.

Main Product

We can divide computer memory into two main segments: primary and secondary memory.

Primary memory is used for immediate data storage. While the logic units (GPUs, CPUs, and newly LPUs) are busy doing the math, it stores their temporary data. Not for too long, just for a few milliseconds. Because of the

need for immediate access, their response capabilities have to be lightning-fast. The currently used high-end HBM3e memories have a bandwidth of about 1200 GB/s. This kind of extreme speed has a price: their capacity is significantly lower compared to secondary memory (24-32 GB per HBM3e stack).

Secondary memory is used for long-term data storage, and this is where Sandisk operates. A typical enterprise-level NAND SSD can store data for years and has a capacity of a few thousand GB. So far so good, but they have a downside, of course. Their bandwidth is about 10-14 GB/s.

Primary memory (HBM3e)	Secondary memory (NAND Flash)	
Volatility	Volatile, loses data without power	Non-volatile, keeps data without power
Storage time	Milliseconds	Years
Bandwidth (speed)	1200 GB/s	10-14 GB/s
Capacity	24-32 GB/stack	up to 30.000 GB/drive
Used for	Help the calculations	Store data permanently

Market and Business Model

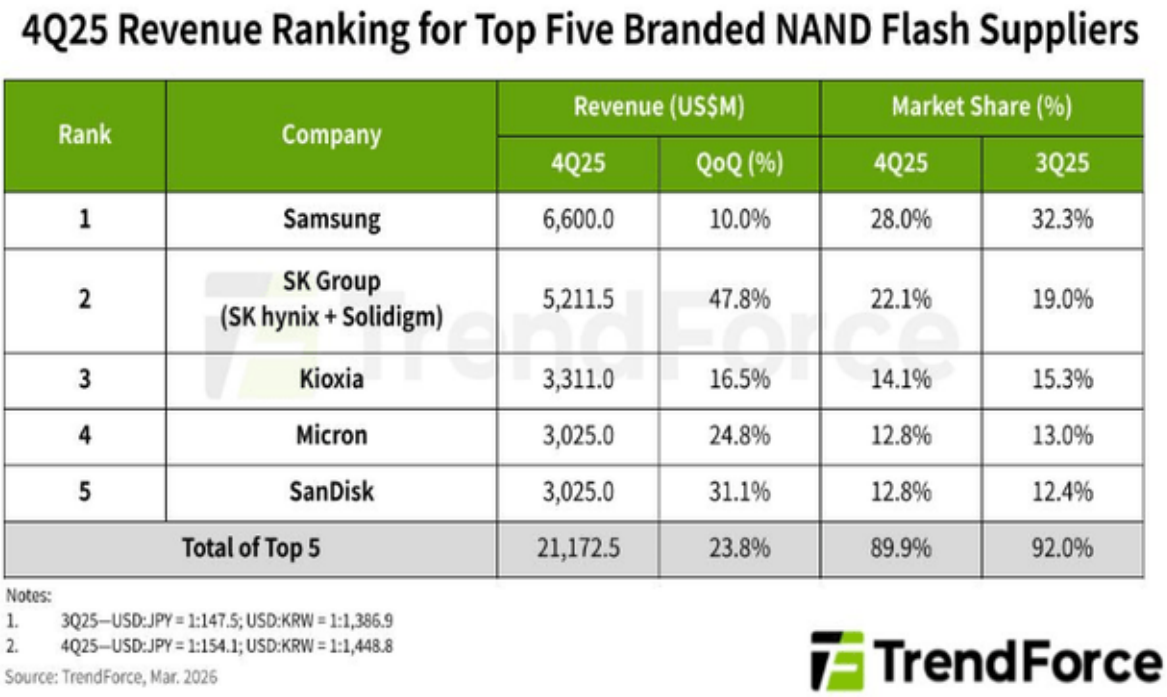

4Q25 Revenue Ranking for Top Five Branded NAND Flash Suppliers

Rank	Company	Revenue (US$M)		Market Share (%)	
		4Q25	QoQ (%)	4Q25	3Q25
1	Samsung	6,600.0	10.0%	28.0%	32.3%
2	SK Group (SK hynix + Solidigm)	5,211.5	47.8%	22.1%	19.0%
3	Kioxia	3,311.0	16.5%	14.1%	15.3%
4	Micron	3,025.0	24.8%	12.8%	13.0%
5	SanDisk	3,025.0	31.1%	12.8%	12.4%
Total of Top 5		21,172.5	23.8%	89.9%	92.0%

Notes:
1. 3Q25—USD:JPY = 1:147.5; USD:KRW = 1:1,386.9
2. 4Q25—USD:JPY = 1:154.1; USD:KRW = 1:1,448.8

Source: TrendForce, Mar. 2026

TrendForce

NAND Flash market

This is a highly concentrated oligopoly with only five major global producers. (And a very capable Chinese one, but they won't ship anything to US data centers anytime soon.) Despite the current Korean dominance, there are two main reasons why Sandisk has a promising future.

SK Hynix, Samsung, and Micron run the entire HBM market; they produce over 97% of those chips. And they can't produce enough-this chip is the current chokehold of the AI buildout. Today, they can sell their HBM chips with an +80% margin, while NAND sits at about +50%. **They are reducing or constraining NAND wafer output to reallocate fab capacity and wafers toward HBM output.**

The increased demand from the customer side and decreased supply from competitors give Sandisk and Kioxia a rare opportunity to raise both margins and market share at the same time.

Speaking of Kioxia, this Japanese producer is my second main reason why Sandisk has a positive outlook.

Sandisk and Kioxia have a joint-venture partnership; they operate two main fabs in Japan. Kioxia owns and operates the factories, while Sandisk invests capital. They run the R&D together, and both get the wafers with minimal markup from the JV factory. The partnership is so balanced that they recently extended their agreements through 2034.

This gives Sandisk predictable and fixed costs on the buy side, while they can widen their margins on the sell side. An extremely well-positioned company during a demand boom!

Recent News, Google TurboQuant "Threat"

Now that we have a quick glimpse into the fundamentals of Sandisk's products and market environment, this solid base helps us estimate the impact of this recent announcement. (Spoiler: the negative impact is zero.)

Google researchers published a blog post on March 24th about a new type of data compression technique. It can shrink AI inference memory usage by 6x with almost zero accuracy loss.

Instead of looking at a memory vector using standard coordinates (i.e., X, Y, Z) that indicate the distance along each axis, PolarQuant converts the vector into polar coordinates using a Cartesian coordinate system.

To be honest, with sentences like that, I feel my initial fears are somehow validated. :) It'd be a worrisome situation if the people who develop algorithms like this were coming for my investment's market share. Fortunately, this is not the case.

Their original publication is almost a year old, and it is linked in the blog post, too. The industry has been well aware of this specific technique for a year, and still, we haven't seen any diminished memory demand.

This is about KV cache memory optimization, which happens on the GPU side and not on the storage drives! It has absolutely nothing to do with NAND Flash drives, as they are for storage, not for data processing!

(Yes, you can offload some of the KV cache to NAND, but it is a suboptimal process and will never be the primary use case for them.)

This academic paper is an important step toward more capable and efficient models. With the help of more efficient models, AI labs can produce more tokens and more compute, and they can serve more customers. More AI data means more demand for fast storage, not less.

While I don't see any significant and immediate effect on the memory business, the direction itself clearly aligns with the interests of storage producers.

Fundamental Analysis

The past week provided an opportunity to enter or re-enter the storage sector at a discounted price.

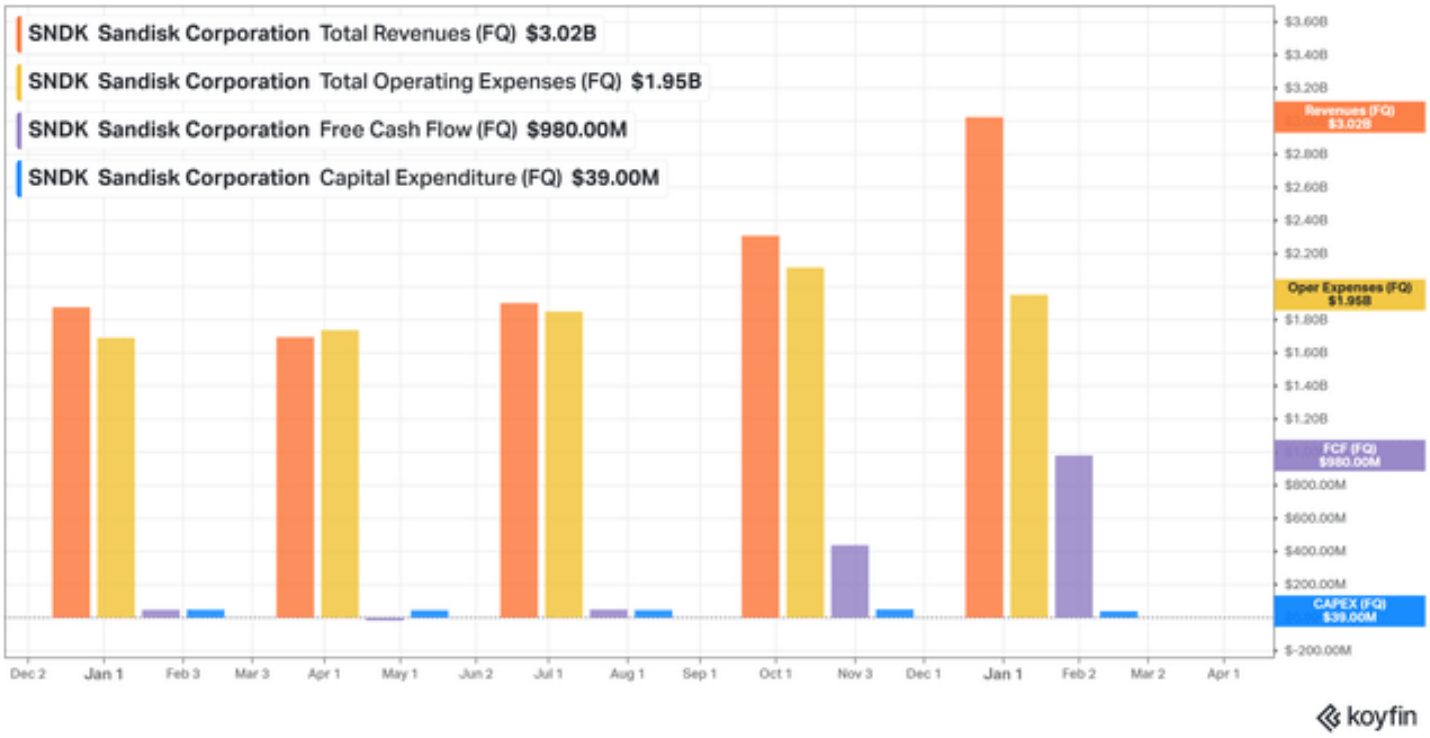

Chart1

Chart 1 shows the rapid acceleration in Sandisk's financials over the past five quarters. Revenue has climbed from roughly $1.7 billion in early 2025 to $3.02 billion in the January 2026 quarter. Operating expenses have remained disciplined at $1.95 billion in the latest print. Most telling, the Capital Expenditure, which has stayed extremely low, was just $39 million in the recent quarter. This is the direct result of the Flash Ventures JV with Kioxia: Sandisk does not carry the full burden of owning and operating the fabs. The combination of surging revenue, controlled operating expenses, and minimal capex converts directly into $980 million of free cash flow - a dramatic inflection from the near-zero FCF levels seen throughout 2025. In my view, this chart is the clearest proof that Sandisk is now in a true cash-generation phase, thanks to the capital-light structure of the JV.

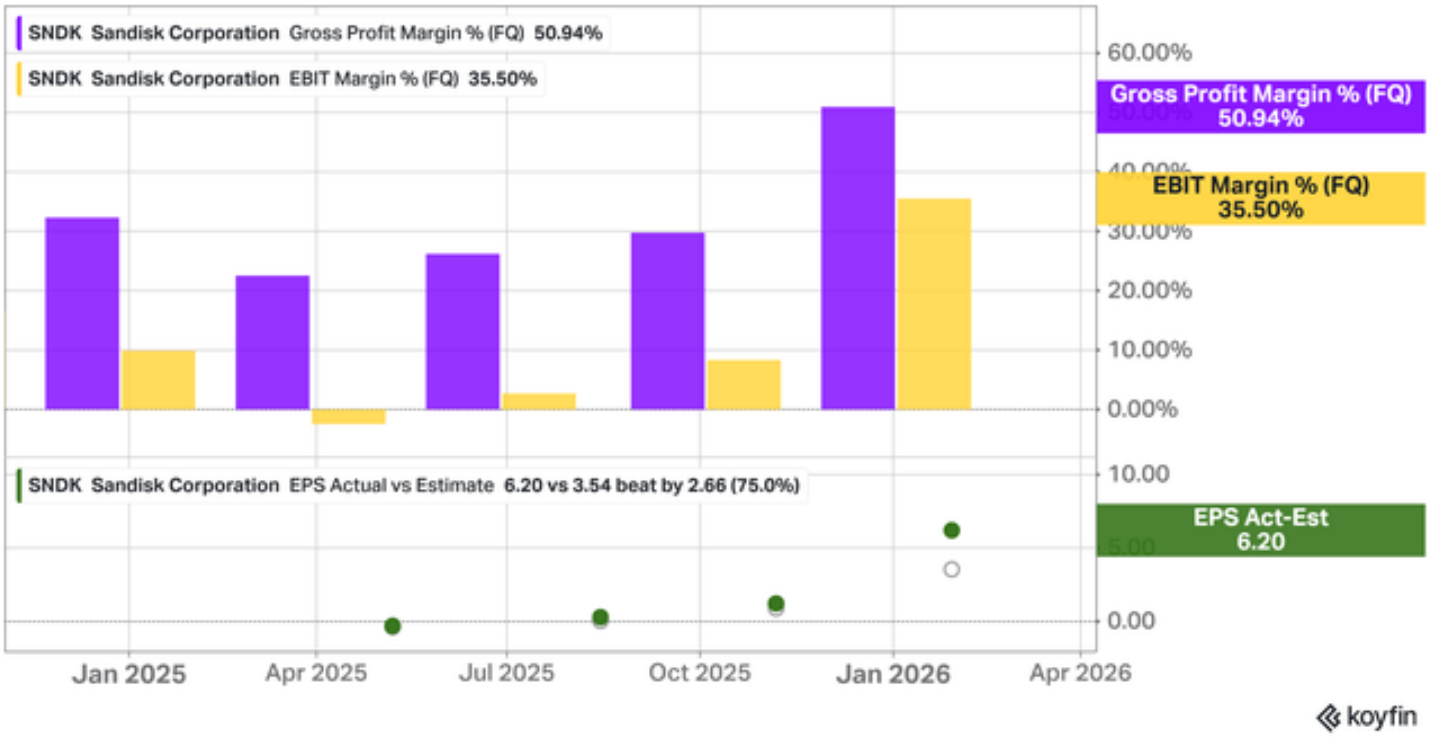

Chart2

Chart 2 highlights the margin expansion and execution quality that accompany that growth. Gross profit margin reached 50.94%, and EBIT margin hit 35.50% in the most recent quarter. Equally impressive is the 75% EPS beat ($6.20 actual versus the $3.54 consensus estimate). These results reflect both favorable NAND pricing and the operating leverage inherent in the JV model: fixed wafer costs on the supply side paired with rising sell-side prices.

Chart3

Chart 3 puts Sandisk's performance in context against Nasdaq peers. The scatter plot of EPS estimate changes versus growth expectations shows SNDK as a clear positive outlier (red arrow). While names cluster around more modest beats or flat estimates, Sandisk is delivering outsized earnings surprises on the back of the AI-driven storage tailwind. This reinforces my belief that the growth story is structural, not purely cyclical.

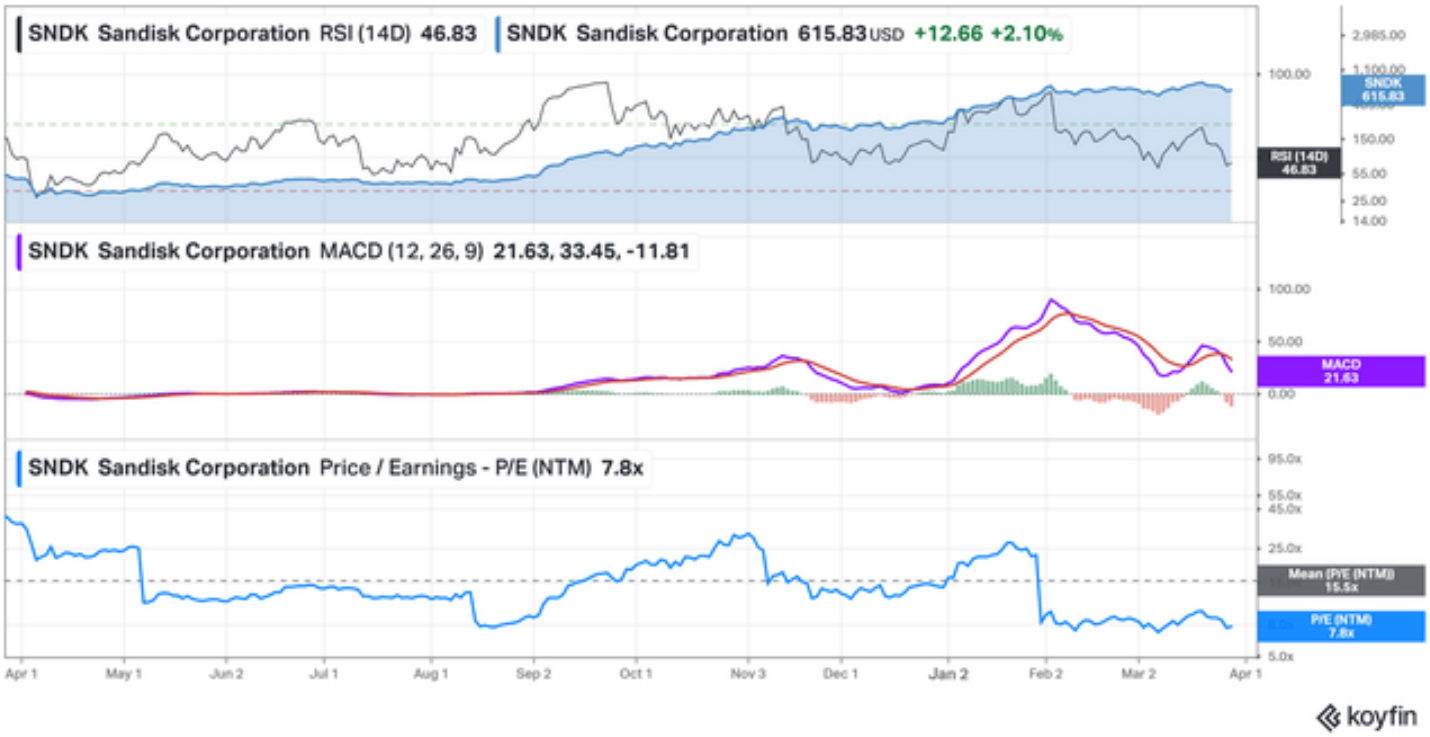

Chart4

Chart 4 provides a clear overlay that reinforces why the current levels are an attractive entry point. The stock is trading at $615.83, yet the forward P/E (NTM) stands at only 7.8x, roughly half the 15.5x long-term mean shown on the chart. This valuation compression has occurred even as the company has delivered accelerating revenue, expanding margins, and massive earnings beats. The 14-day RSI sits at a neutral 46.83, well away from overbought territory. The MACD histogram is mildly negative at -11.81, indicating a short-term pause in momentum rather than a reversal. The overall price trend has remained firmly upward over the past year. In my view, fundamentally strong results paired with a depressed forward multiple and non-extended technicals make the current price an excellent risk/reward setup for new or add-on positions.

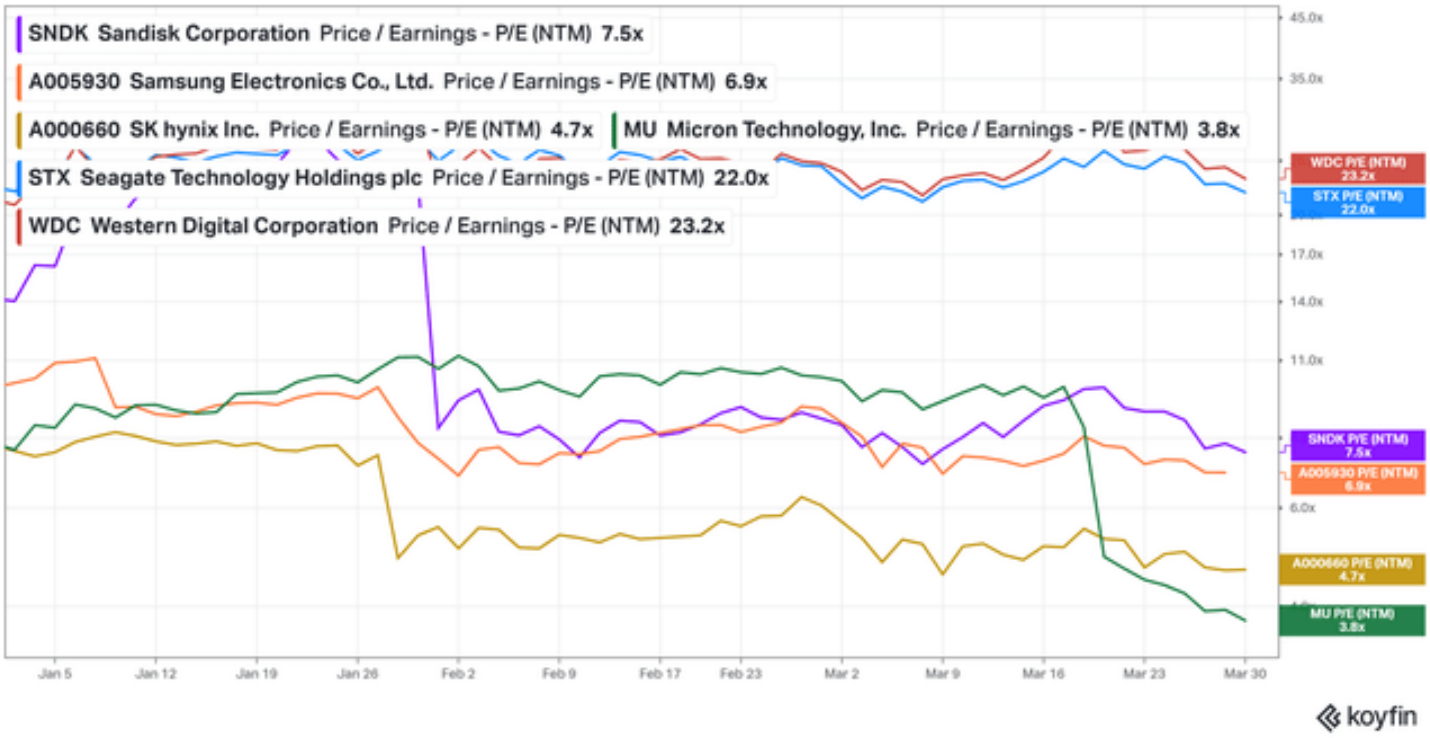

Chart5

Chart 5 shows Sandisk's forward P/E valuation relative to key peers. At 7.5x, SNDK trades at a modest premium to the HBM-heavy players (Samsung 6.9x, SK Hynix 4.7x, Micron 3.8x) but at a significant discount to the legacy HDD/storage names (Seagate 22.0x, Western Digital 23.2x). This positioning is attractive: combined with the strong earnings momentum shown in Charts 1-3, the 7.5x multiple underscores that the stock remains undervalued on a relative and absolute basis.

Note: Numbers before February 2025 are from Western Digital's Flash segment results.

Target Price

Reverse DCF

Inputs	
Current stock price	600,00
Shares outstanding (million)	147,6
Perpetuity growth rate	3,0%
Desired return (discount rate)	10,00%
FCF year 1 (million)	4910,00
Growth rate year 1-10	6,62%

Adjusted FCF	0
FCF	
Stock based compensation (million)	
CapEx (million)	
Depreciation & Amortization (million)	

Inputs
Formula

Cash flow estimates	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10
FCF	4 910,00	5 235,22	5 581,98	5 951,71	6 345,92	6 766,25	7 214,42	7 692,28	8 201,78	8 745,04
Growth		6,62%	6,62%	6,62%	6,62%	6,62%	6,62%	6,62%	6,62%	6,62%

Discounted FCFs	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10
FCF	4 910,00	5 235,22	5 581,98	5 951,71	6 345,92	6 766,25	7 214,42	7 692,28	8 201,78	8 745,04
Present value	4 463,64	4 326,63	4 193,82	4 065,10	3 940,32	3 819,37	3 702,14	3 588,50	3 478,36	3 371,59
Sum discounted CFs	38 949,46									

Terminal value	
TV	128 676,95
Present Value TV	49 610,54

Total equity value	
Total equity value	88 560,00

Value per share	
Intrinsic Value	600,00
Upside/downside	0,00%

Reverse DCF

To quantify the market's embedded expectations, I ran a reverse discounted cash flow model. At the price of $600 and $4.91 billion of free cash flow (annualized from last FCF) in Year 1, the market expects 6.62% compound annual growth for the next 10 years. So...there is 6.62% compound annual growth priced in. We have to check whether it is realistic or not. Using Sandisk's Q2 results and consensus estimates, I modeled the company's valuation under different growth rates. The base facts and forward estimates are shown in the table below. Even under the conservative 10% 10-year CAGR assumption, the DCF yields a target price of $775. The bull-case scenario of 22% 10-year CAGR (aligned with the reported 22% YoY bit-volume growth and the aggressive revenue ramp expected in FY2026-2027) implies a target of $1,693. This demonstrates that the price of $600 already embeds only modest long-term growth, while the structural tailwinds in NAND supply and AI-driven demand support significantly higher outcomes.

Item	Value	Source/Notes
Q2 FY2026 Adjusted FCF	$843 million	Company earnings release
Q2 FY2026 FCF Margin	27.9%	Company earnings release
Q3 FY2026 Revenue Guidance	$4.4 - $4.8 billion	Company guidance (low end shown)
FY2026 Consensus Revenue	$15.58 billion	Wall Street consensus
FY2027 Consensus Revenue	$26.84 billion	Wall Street consensus
FY2026 Consensus EPS	$40.27	Wall Street consensus
FY2027 Consensus EPS	$87.73	Wall Street consensus

Growth Assumption	Target Price
22% growth rate	**$1,693**
15% growth rate	**$1,074**
10% growth rate	**$775**
5% growth rate	**$560**

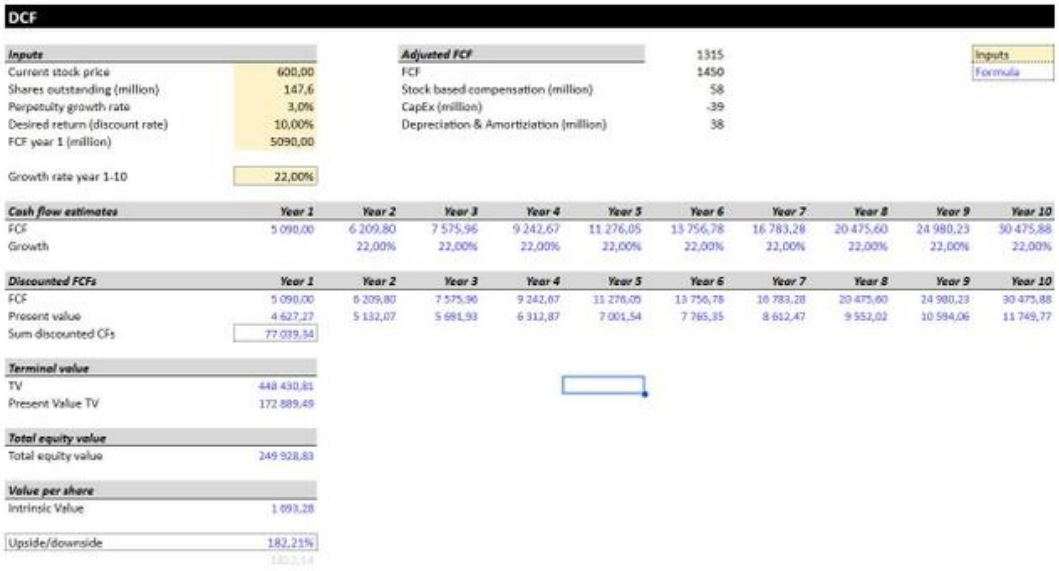

DCF

DCF										
Inputs		**Adjusted FCF**								
Current stock price	600,00	FCF	1450							
Shares outstanding (million)	147,6	Stock based compensation (million)	58							
Perpetuity growth rate	3,0%	CapEx (million)	-39							
Desired return (discount rate)	10,00%	Depreciation & Amortization (million)	38							
FCF year 1 (million)	5090,00									
Growth rate year 1-10	22,00%									
Cash flow estimates	**Year 1**	**Year 2**	**Year 3**	**Year 4**	**Year 5**	**Year 6**	**Year 7**	**Year 8**	**Year 9**	**Year 10**
FCF	5 090,00	6 209,80	7 575,96	9 242,67	11 276,05	13 756,78	16 783,28	20 475,60	24 980,23	30 475,88
Growth		22,00%	22,00%	22,00%	22,00%	22,00%	22,00%	22,00%	22,00%	22,00%
Discounted FCFs	**Year 1**	**Year 2**	**Year 3**	**Year 4**	**Year 5**	**Year 6**	**Year 7**	**Year 8**	**Year 9**	**Year 10**
FCF	5 090,00	6 209,80	7 575,96	9 242,67	11 276,05	13 756,78	16 783,28	20 475,60	24 980,23	30 475,88
Present value	4 627,27	5 132,07	5 691,93	6 312,87	7 001,54	7 765,35	8 612,47	9 552,02	10 594,06	11 749,77
Sum discounted CFs	77 039,34									
Terminal value										
TV	448 430,81									
Present Value TV	172 889,49									
Total equity value										
Total equity value	249 928,83									
Value per share										
Intrinsic Value	1 693,28									
Upside/downside	182,21%									

Under the 22% growth scenario, free cash flow reaches $30.4 billion by Year 10. At the current 28% FCF margin, this implies Year 10 revenue of approximately $108.6 billion. This remains a relatively modest

estimate, as consensus already projects FY2027 revenue at $27 billion, meaning the model assumes only moderate further acceleration beyond the near-term outlook.

All things considered, I view $1,693 as a realistic and achievable target for the next 12 months.

Risks to the Thesis

Geopolitical escalation in the Middle East could disrupt sovereign wealth funding for hyperscalers, slowing AI capex and NAND demand. Additionally, any broader slowdown in AI spending or macro weakness would compress NAND pricing and margins. On the supply side, unforeseen disruptions in the Kioxia joint venture could pressure Sandisk's cost advantage. Finally, following its massive post-spinoff rally, the stock is priced for perfection, risking a sharp de-rating on any missteps. However, Sandisk's capital-light model, locked-in supply, and the AI storage boom outweigh these risks.

Conclusion

Sandisk has strong fundamentals and a clearly visible growth path. Still, the stock price is volatile and heavily undervalued.

As far as I can see, the main reason for this is not found in the fundamentals or the balance sheet. The TurboQuant moment highlighted the main reason, I think. There is negative market sentiment surrounding AI stocks, specifically memory stocks. Because of past boom-bust cycles, investors are concerned about these companies' sustainable growth. The current outlook for these companies seems almost too good to be true. If you invest in such stocks, you can expect similar short, volatile periods and negative overreactions. Still, in the long run, you can expect growing revenues and healthy balance sheets.