


Semiconductors in Focus: Trends Shaping the Next Wave of Innovation

David Tsoi, CFA, CAIA, FRM, CESGA, CAMS, *Head of Index Research, APAC*

Artificial intelligence (AI) continues to be the most transformative technology of our era, with semiconductor companies leading the charge and powering groundbreaking advancements. Following a robust recovery in 2024 driven by demand for logic and memory chips, the global semiconductor market is forecast to grow by 18% this year, reaching total revenue of US\$800 billion¹, with the Americas and Asia Pacific expected to lead the growth. The expansion of data centers continues to drive significant growth, especially for companies specializing in AI and semiconductor innovations.



Source: Semiconductor Industry Association. As of October 3, 2025.

AI growth remains intact

As AI-driven monetization opportunities begin to take shape, hyperscaler capital spending remains on the rise despite mounting tariff and economic headwinds. Global data center capex soared by 43% year-over-year in Q2 2025.² Microsoft, Amazon and Google reported that demand persistently exceeds available infrastructure capacity for AI workloads, with projections indicating that additional capacity will continue to expand throughout the year. OpenAI is committing about US\$400 billion to build five new US data centers in collaboration with Oracle and SoftBank. These sites are designed to deliver 7 gigawatts of power capacity, dramatically expanding computing resources to scale OpenAI's services. Meta's capital spending could increase further in 2026, as it is building multiple multi-gigawatt data center clusters to fuel its AI ambitions, with the first facility slated to go live next year. The company has made AI central to its advertising strategy and plans to enable brands to fully design and target campaigns using AI tools by the end of next year. Based on the client's budget, these new tools would generate the entire advertisement, including images, videos and text, and deliver it to the targeted audience.³

¹ <https://my.idc.com/getdoc.jsp?containerId=prUS53791725/>

² <https://www.delloro.com/news/hyperscaler-ai-deployments-lift-data-center-capex-to-record-highs-in-2q-2025/>

³ <https://www.wsj.com/tech/ai/meta-aims-to-fully-automate-ad-creation-using-ai-7d82e249/>

In the past, AI demand has primarily focused on training workloads, particularly for frontier models. While leading tech companies continue pouring resources into building ever-larger AI models, they are also reallocating more investment toward inference. Inference is the stage where trained AI models process new data to generate insights, make predictions or support decision-making. While training a model is essentially a one-time expense, prompting a model (inference) produces tokens, each of which carries a cost. During the Google I/O 2025 keynote, Alphabet CEO Sundar Pichai shared that the firm processed 480 trillion tokens across its products and APIs in April 2025, 50 times more than the same month a year earlier.⁴ The rapid surge in token volume reflects growing usage and adoption of AI models, signaling a greater need for computing power and driving higher demand for chips.

The age of AI reasoning

The shift in investment toward inference has also gained momentum with the launch of new reasoning models. While traditional AI models respond swiftly and excel at pattern recognition, they often fail to understand broader contexts and struggle with complex reasoning. Reasoning models are built to deconstruct complex problems into smaller, manageable steps and solve them through explicit logical reasoning. They are specifically trained to show their work and follow a more structured thought process, which results in longer computation times for user queries. These models demand significantly more compute during inference to reason through intricate problems. This evolution from basic pattern recognition to structured reasoning is pivotal to AI, unlocking its potential to tackle complex real-world challenges effectively. As AI adoption rapidly expands, demand for inference will correspondingly intensify.

The ascending wave of AI agents

AI agents are positioned to revolutionize how organizations function, delivering breakthroughs in productivity and operational efficiency. They are intelligent systems designed to execute tasks independently by comprehending objectives, formulating decisions and taking actions to achieve predetermined goals. While humans define the desired outcomes, AI agents autonomously select optimal actions required to accomplish those goals. These agents boast a broad spectrum of uses, from supporting academic research and streamlining online purchases to planning leisurely vacations. Customer service, sales and marketing, and IT and cybersecurity are the three business functions where AI agents are most frequently deployed or planned for implementation in the next six months.⁵ As enterprises progressively integrate AI agents across diverse operational applications, demand for computational infrastructure is escalating dramatically.

The rise of custom AI chips

Hyperscalers are increasingly focused on ASIC (application-specific integrated circuits) infrastructure to meet surging AI demand. ASICs are custom-built for specific workloads and can execute those tasks far more efficiently and at a substantially lower cost than high-performance GPUs. Although the initial investment to develop ASIC infrastructure is considerable, the long-term cost of running GenAI workloads on these chips is expected to be lower once the upfront expense is absorbed. For example, in April 2025, Google unveiled Ironwood, its seventh-generation Tensor Processing Unit (TPU), specifically designed for inference workloads.⁶ While Google's in-house TPUs were once limited to internal use, the company is expanding external access to drive faster growth of its cloud business. Marvell Technology projects the custom computing device market will surge to US\$55.4 billion by 2028, more than eight times its size in 2023.⁷

⁴ <https://blog.google/technology/ai/io-2025-keynote/>

⁵ Source: PwC's AI Agent Survey (May 2025)

⁶ <https://blog.google/products/google-cloud/ironwood-tpu-age-of-inference/>


⁷ <https://www.marvell.com/content/dam/marvell/en/company/assets/marvell-custom-ai-investor-event-2025.pdf>

AI drives robust demand for high-bandwidth memory (HBM) technology

HBM represents a cutting-edge memory technology engineered to deliver faster data access while reducing energy consumption, which is critical for the performance of AI processing. HBM's market share in the dynamic random access memory (DRAM) segment is forecast to leap from 18% in 2024 to more than 50% by 2030.⁸ Starting with the next-generation HBM4, the base die will be produced using logic processes, enabling lower power consumption and customizable features tailored to client requirements. Driven by escalating computational demands from AI training and inference workloads, HBM's market outlook remains strong. As the primary HBM supplier for Nvidia, holding a 62% share of global HBM shipments in Q2 2025⁹, SK Hynix projects the global HBM market to expand by 30% annually through 2030.¹⁰

SOX™ – the leading index for the semiconductor industry

Covering the 30 largest US-listed stocks and ADRs of companies primarily involved in the design, distribution, manufacture and sale of semiconductors, Nasdaq's PHLX Semiconductor™ Index (SOX) posted a total return of 185% over the past three years, outperforming the NYSE Semiconductor Index by 20 percentage points and delivering more than 1.6 times the gain of the S&P Semiconductors Select Industry Index.



Source: Nasdaq Global Indexes, Bloomberg. As of September 30, 2025.


SOX is a modified market capitalization-weighted index, with the top three constituents by market capitalization capped at 12%, 10% and 8%, respectively, and the rest capped at 4% during quarterly rebalancing. For the full index methodology, please visit our [website](#).

As of the end of September 2025, the 10 largest constituents accounted for 60.9% of the index weight. 77.4% of the index weight is in the Semiconductor Subsector, with the rest in the Production Technology Equipment Subsector, according to the Industry Classification Benchmark (ICB) classification system.

⁸ <https://www.yolegroup.com/strategy-insights/memory-industry-at-a-crossroads-why-2025-marks-a-defining-year/>


⁹ <https://www.counterpointresearch.com/en/insights/samsungs-q2-2025-memory-performance-disappoints-but-signals-h2-recovery/>

¹⁰ <https://www.reuters.com/world/asia-pacific/sk-hynix-expects-ai-memory-market-grow-30-year-2030-2025-08-11/>



Source: Nasdaq Global Indexes, FactSet. As of September 30, 2025.

Nine of the ten largest holdings posted positive total returns over the past 12 months. On average, the top 10 firms achieved a one-year total return of 39%. Within the same industry, stock performances varied significantly. The difference in one-year total returns between the best- and worst-performing stocks among the top 10 constituents (Broadcom and AMD) was a staggering 95 percentage points. This demonstrates the importance of diversification, even when investing in a single sector or theme.



Source: Nasdaq Global Indexes, FactSet. As of September 30, 2025.

Nvidia (weight: 11.3%)

As the largest constituent in SOX, Nvidia was the fourth-best performer among the ten largest constituents, gaining 54% over the past 12 months. It became the first publicly traded company in history to achieve a US\$4 trillion market valuation in July 2025, after joining the trillion-dollar club in May 2023.¹¹ Nvidia's technological edge remains pronounced, as Blackwell shipments accelerate, driven by soaring AI reasoning demand and realized economies of scale. Beyond AI, the company identifies robotics as its most substantial addressable growth opportunity, with autonomous vehicles representing the first major commercial deployment. The firm recently unveiled new NVIDIA Omniverse libraries and NVIDIA Cosmos world foundation models that empower developers to build next-generation robots and autonomous vehicles by integrating AI reasoning with scalable, physically accurate simulations.¹² Nvidia plans to invest up to US\$100 billion in OpenAI to accelerate the build-out of data centers with at least 10 gigawatts of capacity, and US\$5 billion in Intel to co-develop next-generation data center and personal computing products.

Broadcom (weight: 9.8%)

As the second-largest constituent in SOX, Broadcom emerged as the best performer among the ten largest index components, generating a one-year total return of 93% through September 2025. The company continues to dominate the AI ASIC and AI networking semiconductor markets. The chipmaker is engaged with the top seven hyperscalers for custom silicon products, including accelerators for Google, Meta and ByteDance. OpenAI has also committed to purchasing custom chips and networking components from Broadcom, adding 10 gigawatts of AI data center capacity slated to begin in the second half of 2026 and finish by the end of 2029.¹³ Moreover, Broadcom has significant exposure to enterprise software following its acquisition of VMware in 2023. Its infrastructure software division posted a 77% operating margin last quarter, up from 67% a year ago.¹⁴

Micron Technology (weight: 5.0%)

With 22% of the market share in DRAM revenue in Q2 2025¹⁵, Micron Technology stands as the fourth-largest SOX constituent, delivering a 62% total return over the past 12 months. The company expects its November-quarter gross margin to reach 51.5%¹⁶, marking its first time surpassing the 50% threshold since 2019. Driven by the ramp-up of HBM3E products, Micron reached an annualized HBM revenue run-rate of nearly US\$8 billion. The firm continues to build on this momentum into next year as it ramps HBM4 and expands its customer base to six. In smartphones, a rising share of AI-ready devices remains a key driver of DRAM content growth in mobile devices.

TSMC (weight: 4.3%)

As the world's largest contract semiconductor manufacturer, TSMC delivered a 63% total return over the trailing twelve months.¹⁷ With net revenue from its high-performance computing segment expanding to 60% from 52% a year ago, AI-driven demand remains the primary growth catalyst for TSMC, while its market dominance provides it with pricing power. Advanced chips with sizes 7nm or smaller accounted for 74% of the company's total wafer revenue in Q2 2025. Despite the US president's recent announcement of tariffs on semiconductor imports, Taiwan confirmed that TSMC secured an exemption from the latest levies given its substantial American manufacturing investments.¹⁸

¹¹ <https://www.reuters.com/technology/nvidia-sets-eye-1-trillion-market-value-2023-05-30/>

¹² <https://nvidianews.nvidia.com/news/nvidia-opens-portals-to-world-of-robotics-with-new-omniverse-libraries-cosmos-physical-ai-models-and-ai-computing-infrastructure/>

¹³ <https://openai.com/index/openai-and-broadcom-announce-strategic-collaboration/>

¹⁴ <https://investors.broadcom.com/static-files/24714868-15da-4d0b-87b6-2ad418441b27>

¹⁵ <https://www.trendforce.com/presscenter/news/20250902-12694.html>

¹⁶ <https://investors.micron.com/static-files/5fb98d73-2134-4446-8d1b-0f90285f6c02>

¹⁷ Source: FactSet. Total return for TSMC (US listing).

¹⁸ <https://www.bloomberg.com/news/articles/2025-08-07/taiwan-chip-giant-surges-on-exemption-from-tough-new-trump-tariffs-on-chips>

Conclusion

Beyond the escalating demand for advanced training capabilities to support increasingly complex AI model architectures, AI inference has surfaced as a pivotal growth driver. The rapid increase in token volume signifies expanding usage and adoption of AI models. The proliferation of AI agents is poised to transform various industries and substantially increase compute demand. This dynamic landscape highlights the pivotal role of the semiconductor sector in driving the next wave of AI innovation.

Nasdaq's PHLX Semiconductor Index (SOX) delivered a total return of 185% over the three years through September 2025. Funds tracking SOX include the Invesco PHLX Semiconductor ETF (Nasdaq: SOXQ), the Mirae Asset TIGER US PHLX Semiconductor Sector Nasdaq ETF (South Korea: 381180), the Cathay PHLX Semiconductor ETF (Taiwan: 00830), the Global X Semiconductor ETF (Japan: 2243) and the Yurie PHLX Semiconductor Index Fund (South Korea: 7D01596). The Mirae Asset TIGER Synth-US PHLX Semiconductor Sector Leverage ETF (South Korea: 423920) tracks SOX with two times leverage.

Disclaimer:

Nasdaq®, PHLX Semiconductor™, SOX™ and Nasdaq-100® are registered trademarks of Nasdaq, Inc. The information contained above is provided for informational and educational purposes only, and nothing contained herein should be construed as investment advice, either on behalf of a particular security or an overall investment strategy. Neither Nasdaq, Inc. nor any of its affiliates makes any recommendation to buy or sell any security or any representation about the financial condition of any company. Statements regarding Nasdaq-listed companies or Nasdaq proprietary indexes are not guarantees of future performance. Actual results may differ materially from those expressed or implied. Past performance is not indicative of future results. Investors should undertake their own due diligence and carefully evaluate companies before investing. **ADVICE FROM A SECURITIES PROFESSIONAL IS STRONGLY ADVISED.**

© 2025. Nasdaq, Inc. All Rights Reserved.