



FRANKLIN
TEMPLETON

AI in Emerging Asia: A Dynamic Ecosystem

Not FDIC Insured | No Bank Guarantee | May Lose Value



Emerging Asia has taken center stage in the artificial intelligence (AI) revolution. From hyperscalers to chip producers and infrastructure enablers, companies across emerging Asia anchor a vibrant ecosystem that investors should consider as they seek exposure to the AI theme.


To identify the real long-term winners in this AI ecosystem, investors need to appreciate the AI innovations and growth drivers across emerging Asia. They need to understand, for instance, developments in China, where the “DeepSeek moment” in January 2025 heralded the rapid emergence of home-grown AI models that compete well on cost efficiency versus the incumbent models in the United States.

But the AI story in emerging Asia goes far beyond China. It is also about the semiconductor powerhouses in Taiwan and South Korea; the crucial OSAT (outsourced semiconductor assembly and testing) players in Southeast Asia; and the integration of AI into software services in India. At every layer of the so-called “AI stack,” the end-to-end complex of hardware, software and technologies that drive the development, training, scaling and deployment of AI solutions, we can find emerging Asia companies playing the critical role of value chain enabler.

At Templeton Global Investments (TGI), we are actively exploring the diverse opportunity set across the entire AI value chain in emerging Asia. This paper aims to share some of our perspectives on this dynamic ecosystem. We will adopt a structured approach, looking at four distinct layers that combine to form the entire AI stack, each with its own technological nuances, industry characteristics and growth driver:

As we go through the four layers, we will explore industries and value chains that each occupy a share of the entire emerging Asia AI investible universe.*

Click on the chart components to directly access their corresponding content.



*The investible universe comprises listed emerging Asia companies selected based on TGI's research view and their significance within the regional AI ecosystem. This chart aims to illustrate the potential market size and breadth of the opportunity set across individual AI-related sectors and industries. It does not represent a benchmark and it does not constitute an investment recommendation.



Layer One Hyperscalers AI Model Developers

Hyperscalers provide the cloud services and data center infrastructure critical for the computing, training, inference and deployment of AI solutions. AI model developers write and train the foundational models that have now become synonymous to AI. Together, they constitute the cerebral cortex of the AI stack.

While US technology mega-caps continue to lead the way at this layer, emerging Asia has its share of cutting-edge companies, with those in China taking major strides forward despite geopolitical constraints on technology access.

Hyperscalers

China is today the world's second largest AI cloud market after the United States, with its market size expected to see a compound annual growth rate (CAGR) of 72% between 2024 and 2029, hitting RMB218 billion (around US\$32 billion) over the period.¹ In tandem, token demand growth is also surging, with one forecast projecting a 330% CAGR for China's domestic inference token demand through 2030.²

Chinese hyperscalers are both the drivers and beneficiaries of this exponential growth. This is the world of China's internet and e-commerce mega-caps that have spent years adopting cloud-based architecture, building data centers while amassing a large inventory of memory chips and graphics processing units (GPUs). These companies are now fully capable of deploying their massive cloud computing infrastructures to commercialize AI.

Among them, China's e-commerce and cloud leader was the country's first company to migrate all operations onto the cloud in 2021.³ The Cloud + AI segment has been a core part of its growth engine since, and the company is now targeting some US\$100 billion of AI cloud revenue over the next five years, implying a close to 40% CAGR.⁴

Other Chinese hyperscalers are similarly benefitting from cloud and AI-driven growth. China's largest internet search engine operator entered 2026 with the cloud infrastructure segment as its major growth driver across the whole business portfolio. Separately, a Chinese technology conglomerate straddling the internet, social media and gaming sectors anticipates rising demand for its platform-as-a-service (PaaS) solutions, having expanded its global footprint to 65 availability zones in early 2026.⁵

Importantly, the sector's growth is anchored by a potentially improving profitability outlook, as AI cloud margins gradually improve on the back of price hikes. We saw this in March 2026, when at least two of China's leading hyperscalers announced a price increase of up to 30% and 34%, respectively.⁶

1. Source: "China's AI Path: Monetizing Surging Token Use via AI Cloud." Morgan Stanley Research. As of March 16, 2026.

2. Source: China Equity Strategy, J.P. Morgan Global Markets Strategy. As of June 18, 2026.

3. Source: "The Asia AI Ecosystem." Goldman Sachs Global Investment Research. As of October 23, 2023.

4. Source: "Alibaba Targets \$100 Billion of AI Revenue in Five Years." Bloomberg. As of 19 March 2026.


5. Source: [Tencent 2026 first quarter results](#). As of May 13, 2026. An availability zone refers to one or a cluster of data centers that are independent of the other zones in terms of infrastructure, power and cooling.

6. Source: "Chinese cloud providers hike up prices for cloud and AI services." Datacenterdynamics.com. March 19, 2026.

Investors should calibrate the growth narrative with attention to the capital expenditure (capex) intensity of AI infrastructure build-up. Since late 2025, market concerns surrounding the return on investment (ROI) of aggressive AI capex have intensified, impacting hyperscaler stocks, particularly those in the United States.


We believe this is less of a concern for Chinese hyperscalers, whose AI capex relative to operating cash flow remains far lower than their US peers, at 58% vs. 89% for 2026⁷ (see Exhibit 1). The aggregate capex of major Chinese hyperscalers is forecast to be US\$74 billion in 2026, dwarfed by the US\$715 billion projected by their US peers⁸ (see Exhibit 2).

Exhibit 1: Chinese Hyperscalers Have Lower AI Capex-to-Operating-Cash-Flow Ratios



Source: Goldman Sachs Global Investment Research. April 2, 2026.

Exhibit 2: Chinese Hyperscaler Capex vs. US Peers



Source: Goldman Sachs Global Investment Research. April 2, 2026.

Spending projections coupled with domestic demand growth imply that Chinese hyperscalers should have the means to fund research and expand capacity to stay competitive without sacrificing capital discipline and shareholder value. This will mark a clear contrast against US hyperscalers, which face tangible concerns around capital discipline.

AI Model Developers

AI models put the “intelligence” into artificial intelligence. They are the core program behind the deep learning, data processing and content generation of AI. From the first-generation generative AI (GenAI) models now widely used to create or enhance text, visual and audio content, to agentic AI models that can increasingly act as autonomous problem-solving systems, the innovations of model developers are the forefront catalysts of AI ecosystem growth. This is also where China has solidified its position as one of the global leaders.

China is today recognized as both a trailblazer and growth engine of AI model development. A vibrant cluster of AI model startups has taken root, following the meteoric rise of DeepSeek in 2025 as a more cost-efficient alternative to the incumbent US offerings.



DeepSeek’s emergence is a story of necessity driving invention. Given US restrictions on technology exports to China since 2019, DeepSeek is among the Chinese technology companies grappling with a lack of access to top-tier US chips and GPUs. Through architectural innovations and an open-source approach, DeepSeek effectively set the template for Chinese AI model startups to compete with advanced US models on superior cost efficiency.


In developing its titular models, DeepSeek adopted and fine-tuned the mixture-of-experts (MoE) model architecture. A MoE architecture uses conditional routing for parameter activation. This means that when processing prompt tokens and generating responses, it only needs to access a subset of the model’s “brain” that is specialized in said prompt, instead of its entire computing capacity. The result is meaningfully lower cost per token, without significantly impacting output quality. Additionally, DeepSeek’s open-source approach further eliminates the costs associated with licensing fees and vendor markups.

7. Source: “Navigating China Internet.” Goldman Sachs Global Investment Research. April 2, 2026. There is no assurance that any estimate, forecast or projection will be realized.

8. Ibid.

Since the DeepSeek moment, Chinese AI models have firmly established themselves as cheaper but still competent solutions. While various Chinese models have recently begun to raise token prices, they are still far more affordable versus premium US models. Depending on tasks, the pricing gap can sometimes be as wide as 80%–90%. For instance, DeepSeek’s latest flagship model V4 charges roughly US\$5.22 for one million input and output tokens. In the same comparison, GPT-5.5 charges around US\$35.⁹ Notably, their cost competitiveness is achieved with limited impact on performance, based on the latest GDPval-AA ranking¹⁰ (see Exhibit 3).

Exhibit 3: Chinese Models on the GDPval-AA Leaderboard



Source: Artificial Analysis. As of July 2026. This chart is a partial representation of the entire leaderboard.

Among these Chinese models are offerings from high-profile Chinese AI startups that are following DeepSeek’s path to push the boundaries of China’s AI ecosystem. They are often headlined by the so-called Six AI Tigers: Zhipu AI, Moonshot AI, MiniMax, Baichuan AI, 01.AI and StepFun.

In the financial world, two of the Six AI Tigers have successfully launched their Hong Kong initial public offerings (IPOs) in January 2026; another company filed its listing application in June 2026. In the realm of AI technologies, these companies are advancing model architectures and expanding end-market applications across worker productivity, entertainment, visual and audio generation, achieving strong user traction in the process. This dovetails with the global advent of agentic AI and multi-modal AI applications.

For instance, one of these AI developers is notable for its foundation models that have long context processing capacity, capable of understanding, generating and integrating a wide range of modalities, including text, video and audio. These models power a suite of AI products that collectively have accumulated over 200 million users as of September 2025.¹¹ They include two sophisticated AI chatboxes functioning as “companions” that can interact and role-play via text and audio exchange.

These AI products are forecast to position the company for accelerating revenue growth through to at least 2030.¹² Taking advantage of the company’s multi-modal capabilities, the AI products can provide companionship, stimulate creativity and offer a level of personalized experience that has the potential of tapping the emotional and psychological needs of AI natives. This has positioned said company as an emerging leader in the field of “AI as entertainment.”

9. Source: “DeepSeek-V4 arrives with near state-of-the-art intelligence at 1/6th the cost of Opus 4.7, GPT-5.5.” VentureBeat. April 24, 2026.

10. The GDPval-AA framework tests AI models on real-world tasks across 44 occupations and nine major industries. Source: [Artificial Analysis](#).

11. Source: MiniMax IPO prospectus. December 2025.

12. Source: “MiniMax revenue seen rising to \$219M in 2026, reaching \$5.8B by 2030.” Visible Alpha, S&P Global Market Intelligence. As of April 22, 2026.



Challenges and Progress



With hyperscalers taking the lead, along with AI model developers and semiconductor companies (which we will subsequently touch on), China's AI stack is rapidly growing into one of the world's most innovative and dynamic technology ecosystems. Headwinds to further growth remain as geopolitical and trade risks may yet disadvantage China relative to the dominant US stack, but domestic developments are helping to narrow that gap, in our view.

US restrictions on technology exports will remain a blockage that could tangibly hold back AI development, particularly in the area of AI inference and training, which heavily relies on advanced chips and GPUs. As mentioned, local innovations have yielded impressive results, but that will not completely remove the hurdle. At the same time, the United States and Europe are limiting Chinese models' access to their own markets and data bases, often citing security concerns.

We are encouraged to see that China is responding well to the challenges by actively pushing for technology localization. Progress on this front is notable. Domestically produced AI chips and GPUs made up nearly 41% of China's AI accelerator server market in 2025, with Huawei Technologies as the dominant driver and beneficiary of this paradigm shift.¹³ In March 2026, the company launched its next-generation Ascend 950PR chips, which offer near-compatibility with NVIDIA's CUDA programming architecture, potentially shrinking NVIDIA's technology moat.

The AI ecosystem in China is also supported by a bevy of favorable domestic drivers. An abundant access to rare earth minerals is an advantage few countries in the world enjoy. Cheaper energy, a large pool of engineering talents and accommodative government policies are all reasons to be confident of the future of China's AI stack, in our view.

13. Source: "Chinese chipmakers claim nearly half of local market as Nvidia's lead shrinks." Reuters. As of April 1, 2026.



Layer Two Chip Producers IC Design and OSAT Companies


If hyperscalers and models form the cerebral cortex of AI, semiconductors will be the fundamental building blocks of its brain matter and nervous system. It is estimated that a single server rack in an AI data center contains over 4,500 chips and a leading AI data center can contain over 45 million chips.¹⁴ In this foundational part of the AI value chain, emerging Asia companies play a dominant role.

A top cadre of world-class companies in South Korea and Taiwan firmly occupy global leadership positions across logic and memory chip production, as well as integrated circuit (IC) design. Elsewhere, Southeast Asia's prominent role in semiconductor backend assembly and testing services also bears watching.

Semiconductor Producers

By some estimates, the global rush for AI infrastructure buildout may push spending to US\$3 trillion–US\$4 trillion by 2030.¹⁵ The overall AI data center semiconductor revenue should grow in tandem, potentially from US\$670 billion in 2026 to US\$1.2 trillion in 2028¹⁶ (see Exhibit 4). Critical semiconductor producers in South Korea and Taiwan should remain the prime beneficiaries throughout this period of exponential growth.

Exhibit 4: AI Data Centre Buildouts Fuelling Semiconductor Revenue Growth



Source: "Powering AI: The Semiconductor Ecosystem at the Foundation of Data Centres." Semiconductor Industry Association. As of June 2026. There is no assurance that any estimate, forecast or projection will be realized.

14. Source: "Powering AI: The Semiconductor Ecosystem at the Foundation of Data Centres." Semiconductor Industry Association. As of June 2026.


15. Source: "Nvidia's profit triples as Jensen Huang predicts further revenue growth." Reuters. As of May 21, 2026.

16. Source: "Powering AI: The Semiconductor Ecosystem at the Foundation of Data Centres." Semiconductor Industry Association. As of June 2026.

As evidence, the Taiwan-based global leader in advanced chip fabrication and packaging has reported consecutive quarters of record earnings and a robust 2026 revenue growth guidance.¹⁸ AI accelerator chips remain a structural growth driver. While the company previously projected a compound annual growth rate of mid-to-high 50% for its AI-related revenue, the growth trajectory has strengthened further in 2026,¹⁹ supporting our strong conviction on the company.

The robust outlook of this Taiwanese chip foundry giant is firmly anchored by its relentless push for process technology leadership at the cutting edge of advanced node architectures. In broad terms, advanced nodes are at or below 7 nanometres (nm). The ‘smaller’ the nodes are, the higher the transistor density on a chip, which means better performance and energy efficiency. Already dominant in the 3nm process market—which accounts for the bulk of the current-generation AI accelerator chips—the Taiwanese company is now rapidly ramping up its production capacity for the 2nm process. The new architecture marks the company’s transition from the currently standard FinFET transistor design to a whole new “gate-all-around” (GAA) nanosheet transistor design, which boosts performance by 10%–15% over the 3nm process at the same power (see Exhibit 5).

Exhibit 5: Inside the GAA Architecture




A transistor typically comprises the channel (green) through which a current flows and the gate (red) that regulates the flow. The transition from FinFET to GAA transistors means turning the channel into stacked nanosheets that are wrapped on all four sides by the gate. This reduces leakage and increases drive current, leading to superior electrical signals pass through and between the transistors, thereby improving chip performance.¹⁷

While the company did not pioneer the GAA technology, it stands out with peerless manufacturing yield and the largest, most trusted production line in the industry. Efforts to expand the 2nm production capacity is a priority in 2026, even as the company gears up for the next-generation 1.6nm and 1.4nm processes. Set for mass production in 2028, the 1.4nm process may offer another 10%–15% performance boost over the 2nm process at the same power.

Over in the memory segment, a memory supercycle is also unfolding, as linear supply expansion struggles to keep up with exponential demand growth (see Exhibit 6). This is creating major growth tailwinds for the world’s top two memory chip producers, both in South Korea. The duopoly accounts for around 67% of the global DRAM market share by revenue. In the high bandwidth memory (HBM) market, which primarily feeds the AI data center memory needs, their combined market share is 79%.²⁰ One of these companies, the clear global leader in the HBM market, has seen a significant profit surge amid a favorable pricing environment. Market research indicates the company may see its operating profit soar 533% to KRW299 trillion in 2026 and KRW448 trillion in 2028.²¹

Exhibit 6: A Widening Gap Between Exponential Memory Demand vs. Linear Supply



This chart is for concept illustration only and is not based on data. KV cache refers to short-term AI memory that optimizes token processing and reduces computing strain.

Not resting on its laurels, these companies are also putting technological leadership and production expertise at the core of their competitive strength. The HBM leader was the world’s first to launch the product in 2013, and it has been staying ahead across all generations of HBM since, including the flagship HBM3E memory chips adopted for NVIDIA’s Blackwell Ultra AI accelerators, as well as proprietary application-specific integrated circuits (ASICs) developed by Alphabet and Amazon.com. The company was also the first to start mass production of HBM4 memory in early 2026. Compared to HBM3, HBM4 benefits from an architectural redesign that enables doubling of interface width and a 40% power efficiency improvement. When applied, HBM4 could boost AI service performance by up to 69%.²²

17. Source: “What is a gate-all-around transistor?” ASML.

18. Source: TSMC financial results – 2026Q2. July 16, 2026.

19. Source: Goldman Sachs Global Investment Research. July 16, 2026. There is no assurance that any estimate, forecast or projection will be realized.


20. Source: “Global DRAM and HBM Market Share: Quarterly,” Counterpoint. June 9, 2026.

21. Source: “SK Hynix”. Mirae Asset Securities. June 9, 2026.

22. Source: “SK Hynix Completes World’s First HBM4 Development and Readies Mass Production.” SK Hynix. As of September 12, 2025.

The competition is far from over, however, with the HBM leader's main rival across town surprising the market with its shipment of the next-generation HBM4E memory in May 2026. Not to be left behind, the HBM leader announced its own HBM4E shipments in June, suggesting an even match as the next leg of the HBM race heats up. Compared to the baseline architecture, HBM4E may deliver an over 20% increase in stable pin speed.²³ As AI models become increasingly sophisticated, their computing needs will demand both companies to keep pushing the envelope in the memory space.

Overall, we remain confident of the outlook for the semiconductor heavyweights in Taiwan and South Korea. These bellwether companies command strong or improving technological moats, production expertise and client relationships. Supply constraints also grant them the pricing power to boost margins and profitability. As the AI boom continues, we expect the semiconductor producers in emerging Asia to sustain robust earnings growth.



China is steadily emerging as another key supplier of memory chips in emerging Asia. As part of the country's active policy push to onshore technology value chains and reduce reliance on Western imports, a small group of domestic memory producers have grown in stature and output.

Leading the pack is a Hefei-based dynamic random access memory (DRAM) maker. The company is currently the world's fourth-largest DRAM supplier, with a global DRAM market share of around 7.7%.²⁴ Similarly buoyed by the AI-driven memory supercycle and price hikes, the company expects a major revenue bump in 2026.²⁵ China's domestic memory market growth shows a similar momentum, as the market size may expand 262.9% to reach US\$449.6 billion in 2026.²⁶

The Hefei company's emergence reflects the aspirations and challenges of the broader memory value chain in China. In 2019, just as the US government ramped up the technology access restrictions on China, the company laid its technological foundation with licensed patents from a bankrupt German DRAM maker. It acquired the buried word line (BWL) + stacked capacitor cell architecture, which is applied to the 10nm class that currently entails the bulk of its products. As the BWL + stacked capacitor architecture is the memory industry's shared platform with proven physics, the company has essentially minimized its research, development and know-how gap versus the incumbent manufacturers without having to self-invent alternatives. As well, it lends clarity on the incumbents' forward roadmaps, further reducing technology risks.

Despite their growth, we continue to see a tangible competitive gap between Chinese memory producers versus the South Korean and US leaders. For the Hefei company, restrictions on access to cutting-edge wafer fabrication equipment (WFE), especially extreme ultraviolet (EUV) lithography systems, will remain a critical issue. EUV lithography—which uses very sharp ultraviolet beams to print ultra-fine circuit patterns onto silicon wafers—is essential for producing advanced DRAM memory chips with cost efficiency and a high yield. Without that, the company must rely on deep ultraviolet (DUV) machines, which have a much larger wavelength and hence require multiple passthroughs to print advanced nodes. The result is not only a roughly 30% increase in cost per bit, but also a major yield loss for HBM chips. The company only has a roughly 35% yield for its HBM3 chips, compared to the 85%–90% industry standard.

Tackling this equipment shortfall will be a key obstacle to overcome if the Hefei company and its local peers are to consistently scale up their advanced memory platforms. Success on this front will be the key to sustain growth beyond the current AI-driven market upcycle or the next node transition. This will in turn depend on the growth of China's domestic WFE industry, a structural trend that we believe is also in motion as part of China's technology localization push.

23. Source: "Samsung Electronics Begins Shipment of Industry-First HBM4E Samples." Samsung Electronics. As of May 29, 2026. Stable pin speed refers to the rate of data transfer through a circuit connection without signal error. It measures the maximum performance of a memory chip where data transfer remains stable.

24. Source: CXMT IPO prospectus. As of May 17, 2026.

25. Ibid.

26. Source: Omdia forecast. As of July 2026.



IC Design and Outsourced Semiconductor Assembly and Testing (OSAT)

In addition to its semiconductor heavyweights, emerging Asia is home to critical enablers across the semiconductor value chain, including IC design houses in the upstream and OSAT services providers in the downstream.

IC design companies refer to the likes of NVIDIA, Qualcomm and AMD. They are 'fabless' companies that do not handle semiconductor manufacturing or fabrication, instead focusing on designing chips such as GPUs, CPUs and ASICs. As innovators that push the boundary of AI's capabilities, IC design companies are situated at the top end of the semiconductor value chain.

While the US is firmly leading on this front, a number of emerging and established names stand out in emerging Asia. In Taiwan, for instance, the leading IC design house is also the world's fifth-largest, having carved up major market share in the mobile and Internet of Things (IoT) segments. The company is now actively growing its ASIC market share. It is working on two projects related to Alphabet's Tensor Processing Unit (TPU). We expect revenue contribution to start in the second half of 2026, with potentially strong upside in the coming years.²⁷

Developments in China are also noteworthy. The country's IC market is expected to grow at a compound annual rate of 9.8% over 2023–2032, outpacing the global IC market growth of 8.5% over the same period. Amid said growth, China's policy push for technology localization is yielding tangible results, with domestic self-sufficiency ratio likely improving from 18% in 2022 to 27% in 2028.²⁸ A group of rising stars are driving this wave of domestic growth, including a Shanghai-based chip designer that scored a successful Hong Kong IPO in January 2026.

OSAT providers ensure chips work as intended. Due to their complex design and delicate nature, AI chips require rigorous post-production testing to ensure their functionality. Once defects are filtered out, the functional chips are then assembled onto a substrate through packaging, before another round of final testing is conducted. This critical and sophisticated process is the speciality of OSAT companies.

In emerging Asia, Taiwan is also a leader in the OSAT industry, with one major company ranked as the world's top pure-play OSAT services provider. Elsewhere, Southeast Asia continues to grow as a key OSAT hub, showcasing an increasingly diverse technology landscape across emerging Asia.

Among the Southeast Asian countries, Malaysia has secured an entrenched position in the backend of the semiconductor value chain, holding some 13% share of the global assembly, testing and packaging activities.²⁹ The leading OSAT company in Malaysia is a close partner of global semiconductor majors, and it has near-monopoly over Broadcom's orders.

In a shift that bodes well for its outlook, this Malaysian company is aiming to venture into more advanced packaging platforms, such as flip-chip chip-scale packages, flip-chip ball grid arrays and 2.5D/3D stacked packaging. This should further align the company with the proliferation of advanced AI technologies.

With a cluster of world-class and emerging chip producers, IC design companies and OSAT providers, the emerging Asia semiconductor complex is firmly positioned as the backbone of the global technology value chain. Asian technology companies in this layer of the AI stack have been heavily favored by investors over the past two years. Attention to their valuation levels is warranted and portfolio allocations must be calibrated by risk/reward considerations, but we believe their intrinsic worth and long-term growth fundamentals remain compelling.

27. Source: Templeton Global Investments research. May 2026.

28. Source: "China IC Design Sector." BOCOM International. As of October 2024. There is no assurance that any estimate, forecast or projection will be realized.

29. Source: "Malaysia's Position in the Global E&E Value Chain and Prospects." Bank Negara Malaysia. As of end-2024.



Layer Three


Server ODMs PCB and Component Providers Energy Infrastructure Enablers

AI models and semiconductors cannot function in isolation. They exist in an environment wrought by a multitude of specialized producers and providers. In this third layer of the AI stack, cloud server original equipment and design manufacturers (OEMs and ODMs), printed circuit board (PCB) makers as well as secondary component companies come together to create the server infrastructure that hyperscalers require to deploy their cloud AI solutions at scale. Also essential are adjacent enablers such as energy and cooling solutions providers, without which data centers will struggle to operate.

To understand the infrastructure layer of the AI ecosystem, let us first consider the rough composition of a server. Once chips and memory are packaged and tested, as described in Layer Two of this paper, they are assembled onto PCBs together with an array of active and passive components. This is usually completed by server manufacturers, which manage the last mile of server manufacturing, creating the server racks that form the basic units of a data center (see Exhibit 7 for illustration).

In this third layer of the AI stack, emerging Asia companies are the global leaders, playing critical roles in every step of server construction. Among them are server ODMs in Taiwan, as well as optical module and power equipment providers in China, all with major or rapidly growing market share in the global AI value chain. Collectively, they provide investors with another cluster of AI growth proxies alternative to the first-order beneficiaries at the hyperscaler and semiconductor levels.

Exhibit 7: The Birth of a Server




Source: Semiconductor Industry Association. June 2026.

Server Manufacturers

Server OEMs are companies that provide their own branded servers—think Dell or Hewlett Packard in the developed market world. ODMs instead design and build “white-box,” non-branded servers specific to a client’s need.

Taiwanese companies command the lion’s share of the global server ODM market and are thus strongly positioned for the rising capex on AI server investments. The global server total addressable market size is expected to cross US\$1 trillion in 2028, with AI servers accounting for 78% of the value (see Exhibit 8).

Exhibit 8: AI Servers Dominate Global Server Market Growth



Source: Goldman Sachs Global Investment Research. As of June 24, 2026. % in chart indicates AI server revenues mix. There is no assurance that any estimate, forecast or projection will be realized.

Riding this sustained growth momentum, a Taiwan-based contract electronics manufacturer with a globally leading position in the server ODM market was able to book record high profits in 2025, and it expects AI servers to remain a major growth engine in 2026.³⁰


To maintain their competitive position, Taiwanese server ODMs do not stop at downstream manufacturing. They are actively pushing the cutting edge of server architecture in a technological pursuit similar to that of model developers and semiconductor companies. That is why, as of early 2026, major Taiwanese ODMs have all kicked off production of NVIDIA GB300 NVL72, a liquid-cooled, rack-scale supercomputer built for the kind of massive-scale inference required for agentic AI. We believe these sustained efforts to innovate and collaborate will be a driver for significant margin enhancements and greater pricing power. This will be a key differentiating factor in a congested downstream market where margins are typically thinner.

PCBs and Components

Among the foundational components for server production, PCBs is an area where Chinese companies dominate. Between China and other key markets in emerging Asia, such as Taiwan, Asia ex-Japan commands some 85% of the global PCB market (see Exhibit 9).

Additionally, Chinese manufacturers are the world’s top component producers. They provide the secondary components—capacitors, resistors, transformers, diodes, to name just a few—critical for the power and signal regulation that ensures a server functions in optimal conditions.

Exhibit 9: Asia Dominates the Global PCB Market



Source: CMB International Global Markets Equity Research. November 2025.

Optical modules, for instance, form the critical data bridge between trays. With Chinese optical module producers claiming seven of the global top 10 spots and a combined market share of over 60%,³¹ we expect AI infrastructure buildouts to also drive growth across select companies in the industry. This was reflected by the success of the world’s top optical module provider. Looking ahead, this Shandong-based company expects a stable growth outlook, benefitting from its early-mover advantage in silicon photonic module capabilities, which offer the top-tier speed and power efficiency coveted by advanced data centers.³²

30. Source: [Hon Hai Technology Group Q1 2026 financial results](#). May 14, 2026

31. Source: “[From optical modules to chips – China’s tech supply chains sustain global AI growth](#)”. Xinhua. June 7, 2026.

32. Source: Templeton Global Investments research. As of September 2025.


Energy Enablers

Data centers are enormously hungry for energy. The International Energy Agency estimates that data centers globally consumed around 485 terawatt-hour (TWh) of electricity in 2025,³³ enough to power close to 50 million US households for an entire year.³⁴ That hunger will only grow in the coming years, potentially hitting 950 TWh in 2030.³⁵

To sustain the projected pace of AI infrastructure buildout requires overcoming the bottlenecks in the existing power grids. That means demand for solutions that can stabilize or scale up power network efficiency and energy storage systems that minimize outage disruptions, among others.

These opportunities should benefit a range of industrial and energy companies in emerging Asia. At TGI, we have a constructive view on several established or emerging energy enablers that are helping data centers globally to bypass strained power grids and secure more stable energy sources.


In China, for instance, a diesel engine manufacturer is notable for its pivot towards the data center power equipment business. These include conventional diesel-powered data center back-up systems, a segment where the company targets a roughly 30% global market share and unit growth of at least 250% between 2025 and 2027.³⁶ Progress in deploying solid oxide fuel cells (SOFC) marks another growth catalyst that bodes well for investors, in our view. The company reports strong SOFC demand in North America, Europe and Southeast Asia, and is targeting capacity expansion from 30 megawatts (MW) in 2026 to 200 MW in 2027.³⁷



What Are Solid Oxide Fuel Cells (SOFCs)?

SOFCs are batteries that generate power through electrochemical reactions between their fuels—typically hydrogen or hydrocarbons—and oxygen, without combustion. This ensures the kind of highly efficient and environmentally friendly on-site power generation that data centers require for uninterrupted and sustainable operations. SOFC demand is set to reach 22GW by 2030, driven by applications in multiple fields, but data centers are set to be the main consumers (see Exhibit 10).

Exhibit 10: End-Market Breakdown of SOFC Demand by 2030



Source: Ceres Power Holdings.

In South Korea, a lithium-ion battery producer also anticipates healthy data center demand for its energy storage systems (ESS). It targets a 50% growth in ESS sales in 2026 and expects to triple its US capacity to 30 gigawatt-hour (GWh) in 2027. The sustained growth and robust profitability on this front will be critical for the company to mitigate the impact of its loss-making electric vehicle (EV) battery business, as it looks to achieve operating breakeven by the end of 2026.³⁸

33. Source: "Key Questions on Energy and AI." International Energy Agency. As of April 2026.

34. This illustrative comparison is based on average US household electricity consumption per year, which is around 10,500 kilowatt-hour. Source: US Energy Information Administration.

35. Ibid.

36. Source: Goldman Sachs Research. As of May 19, 2026. There is no assurance that any estimate, forecast or projection will be realized.

37. Source: Templeton Global Investments research. As of March 2026. There is no assurance that any estimate, forecast or projection will be realized.

38. Source: Templeton Global Investments research. As of May 2026.



Layer Four

End-Market Applications

At the very end of the AI value chain sit customer-facing applications, software, services and devices. As AI continues to fundamentally disrupt the way we work, live and play, effective integration of AI into products and services will be decisive for the next phase of business growth across multiple industries.

In this fourth layer of the AI stack, emerging Asia companies have demonstrated a high degree of agility and innovations in executing their AI strategies. From internet search, e-commerce, gaming, food delivery, online travel agencies to robotaxi services, a large array of consumer-facing platforms and solutions in the region are now integrated with AI. In the corporate world, enterprise software and fintech solutions are also looking to draw synergies from agentic AI.

To illustrate this paradigm shift, we can examine a few notable examples across emerging Asia:



China's e-commerce and cloud hyperscale leader

continues to showcase the growth and evolution of its comprehensive AI stack. In tandem with its heavy investments in hyperscale infrastructure, discussed in Layer One, the company has progressed far in its AI “super app” transformation.

In January 2026, the company has completed the full integration of its core services—including e-commerce, navigation, payments and travel—into its flagship AI assistant, driven by its proprietary model. This is reportedly China's first AI assistant capable of executing large-scale complex tasks in one stop.³⁹

Given its full-stack AI capabilities, we consider the company among the world's foremost in adopting an AI ecosystem approach to reinforce its market leadership.⁴⁰

39. Source: [Alibaba December Quarter 2025 results](#), March 19, 2026.

40. Source: Templeton Global Investments research, May 2026.



In India, a major merchant commerce platform provider is actively pursuing its AI transition strategy, against the backdrop of AI disruption to the incumbent business models of software and services.

In June 2026, the company rolled out the country's first AI agent payment protocol that can complete a Unified Payments Interface (UPI) payment without human authentication. The agent potentially allows autonomous transactions—browsing and paying—following a one-time authorization versus a merchant.⁴¹

Fresh off its first full financial year of profitability, the company is among the emerging Asia fintech innovators that are now pursuing an AI-centric strategy for business growth and efficiency enhancement.



For the operator of India's largest food delivery and quick commerce platform, AI is the key to enhancing both user and partner experience. Demand prediction, catalogue quality, route optimization, supply-chain management and partner support are just some aspects of the business benefitting from AI integration.

An early AI mover, the company had already launched an in-app generative AI chatbot in 2023 that functions as a personalized foodie companion. More recently, the company is reportedly piloting AI-powered voiced ordering, underlining its multi-modal AI strategy.

Amid concerns that AI will replace commerce services, the company believes AI can instead broaden its competitive moat by reducing friction between consumers, partners and businesses in the physical world.

41. Source: [Pine Labs Launches P3P](#), June 11, 2026.

In other areas of applications, emerging Asia companies have brought AI into our homes, tools, devices and vehicles, aiming to stand out in a crowded consumer market through differentiated “smart” products and experiences.



China's leading EV company has developed an AI ecosystem that encompasses automated manufacturing, in-house semiconductor development and intelligent vehicles. In May 2026, these AI initiatives culminated in the launch of reportedly China's first 4nm automotive driving chip to power its drive assistance system.⁴² This breakthrough marks a major step in the company's progress towards fully autonomous self-driving vehicles. Securing market leadership on this frontier may be a key factor for the company to stand out from the intense competition in China's EV market.



AI is integral to the growth strategies and product developments of Asia's Internet of Things (IoT) leader. Among other offerings, the Beijing-headquartered company is notable for its range of AI-powered vacuum cleaners. The machines reportedly benefit from onboard AI that powers computer vision and object recognition, which then enables intelligent route planning.

More broadly, the company has also ventured into multimodal AI model development, with its flagship model touted as the centerpiece of its smart ecosystem that connects users, homes and cars.




42. Source: [BYD](#). As of May 28, 2026.

Conclusion

Across all four key layers of an AI stack, we see evidence of emerging Asia companies playing the role of critical manufacturers, developers and service providers. In several aspects, these companies are not merely emerging—they are leading the rest of the world in the global AI value chain.

As a global asset manager, TGI invests across the developed and emerging markets. Our comprehensive research and firm conviction in the emerging Asia AI ecosystem complement our understanding of the global technology complex.

We believe a similarly holistic view should benefit investors as they seek diversified exposure to the AI growth theme. The ability to appreciate the full dynamics of emerging Asia's AI ecosystem is the key to navigating this once-in-a-generation technological leap.



WHAT ARE THE RISKS?

All investments involve risks, including possible loss of principal.

Equity securities are subject to price fluctuation and possible loss of principal.

To the extent a portfolio invests in a concentration of certain securities, regions or industries, it is subject to increased volatility.

Large-capitalization companies may fall out of favor with investors based on market and economic conditions. Small- and mid-cap stocks involve greater risks and volatility than large-cap stocks.

International investments are subject to special risks, including currency fluctuations and social, economic and political uncertainties, which could increase volatility. These risks are magnified in emerging markets.

The government's participation in the economy is still high and, therefore, investments in **China** will be subject to larger regulatory risk levels compared to many other countries.

There are special risks associated with investments in **China, Hong Kong and Taiwan**, including less liquidity, expropriation, confiscatory taxation, international trade tensions, nationalization and exchange control regulations and rapid inflation, all of which can negatively impact the fund. Investments in Hong Kong and Taiwan could be adversely affected by its political and economic relationship with China.

Any companies and/or case studies referenced herein are used solely for illustrative purposes; any investment may or may not be currently held by any portfolio advised by Franklin Templeton. The information provided is not a recommendation or individual investment advice for any particular security, strategy or investment product and is not an indication of the trading intent of any Franklin Templeton managed portfolio.

IMPORTANT LEGAL INFORMATION

This material is provided for general informational purposes only and should not be considered individualized investment advice, a recommendation or a solicitation to adopt any investment strategy. It does not constitute legal or tax advice. Franklin Templeton accepts no liability for losses arising from use of this material.

The views expressed are those of the investment manager as of the publication date and may change without notice. These opinions and analyses are based on certain assumptions, including market conditions that may change. They may differ from those of other portfolio managers or from the firm as a whole.

This material is not intended to provide a complete analysis of all material facts regarding any country, region or market. No assurance can be given that any forecast, projection or prediction regarding economies or financial markets will be realized. References to specific securities are for illustrative purposes only and should not be interpreted as recommendations or a solicitation to buy, sell or hold any security.

Any research or analysis in this material has been prepared by Franklin Templeton for its own purposes and is provided incidentally. While the information included is believed to be reliable, its accuracy and completeness cannot be guaranteed, and it is subject to change without notice.

Past performance does not guarantee future results or any profit or gain. All investments involve risks, including possible loss of principal.

Franklin Templeton offers environmental, social and governance (ESG) capabilities, though not all strategies or products incorporate ESG as part of the investment process.

Investment strategies and services may not be available in all jurisdictions. Please consult your financial professional or Franklin Templeton contact for further information.

Brazil: Issued by Franklin Templeton Brasil Ltda. **Canada:** Issued by Franklin Templeton Investments Corp. **Offshore Americas:** In the United States, this publication is made available by Franklin Templeton. **United States:** Issued by Franklin Templeton. **Investments are not FDIC insured; may lose value; and are not bank guaranteed.**

Europe: Issued by Franklin Templeton International Services S.à r.l., 8A, rue Albert Borschette, L-1246 Luxembourg. **Poland:** Issued by Templeton Asset Management (Poland) TFI S.A.; Rondo ONZ 1; 00-124 Warsaw. **Saudi Arabia:** Issued by Franklin Templeton Financial Company, 13512 Riyadh, Saudi Arabia. Regulated by CMA. License no. 23265-22. **South Africa:** Issued by Franklin Templeton Investments SA (PTY) Limited, which is authorised by the FSCA as a Financial Service Provider (No.44475). **Switzerland:** Issued by Franklin Templeton Switzerland Ltd, Talstrasse 41, CH-8001 Zurich. **Middle East & Africa (ex South Africa):** Issued by Franklin Templeton Investments (ME) Limited, which is regulated by the Dubai Financial Services Authority. Address: Franklin Templeton, The Gate, East Wing, Level 2, Dubai International Financial Centre, P.O. Box 506613, Dubai, U.A.E. Tel: +971(04) 428 4100. **United Kingdom:** Issued by Franklin Templeton Investment Management Limited (FTIML), registered office: Cannon Place, 78 Cannon Street, London EC4N 6HL.

Australia: Issued by Franklin Templeton Australia Limited (ABN 76 004 835 849) (Australian Financial Services License Holder No. 240827), Level 47, 120 Collins Street, Melbourne, Victoria 3000. **Hong Kong:** Issued by Franklin Templeton Investments (Asia) Limited. **Japan:** Issued by Franklin Templeton Japan Co., Ltd. **South Korea:** Issued by Franklin Templeton Investment Advisors Korea Co., Ltd. **Malaysia:** Issued by Franklin Templeton Asset Management (Malaysia) Sdn. Bhd. & Franklin Templeton GSC Asset Management Sdn. Bhd. **Singapore:** Issued by Templeton Asset Management Ltd. (UEN) 199205211E.

Access your local website at www.franklinresources.com/all-sites.

