28 May 2025 11:15:37 ET | 64 pages

Global Semiconductors

The AI Hardware Shift in IT Devices: The Era of Personal AI Server Begins

Super - Sector Analysis

CITI'S TAKE

With the expansion of AI distillation technology, we anticipate a significant reduction in AI model sizes which should accelerate architectural changes in portable edge AI devices where AI models can be directly integrated. We expect portable AI servers to emerge as key applications and to significantly drive on-device AI semiconductor demand. With this Super-Sector report, we introduce new architectures for edge AI devices for the first time in the industry, along with advanced AI products and key beneficiaries.

Portable AI era to come: distilled AI models to accelerate on-device AI demand — We project that AI servers will become so efficient and extremely compact that they can be integrated into various small-sized form factors. Recent AI distillation technology pioneered by DeepSeek has enabled minimization of LLMs in a much smaller parameter scale. We believe this advance will support portable AI servers, accelerating on-device AI demand as well as AI semiconductor content growth.

New architectures for on-device AI: miniaturized AI servers — We foresee three major architectural shifts in on-device AI products: 1) AI kits will be added to conventional Von Neumann architectures, integrating AI modules in PCIe; 2) use of near memory or LPDDDR6 with increased I/Os & bandwidth near NPU & TPU; and 3) implementation of LPW/LLW DRAM near NPU/TPU, which can be seen as mini Nvidia's AI server architecture.

Implications: SoIC, LPDDR6/LPW DRAM, and hybrid bonding — Portable edge AI servers must be cost- and energy-efficient to fully operate AI models in edge devices. In order to realize such characteristics, we anticipate broader adoption of heterogeneous integration from a hardware perspective, with TSMC's SoIC playing a central role in enabling this. We further expect mobile DRAM to be increasingly adopted for edge devices. Specifically, we foresee LPDDR6 adoption for flagship devices to start in 2026E and LPW DRAM to become mainstream by 2028E. Lastly, die-to-die integration via hybrid bonding should become more widely adopted.

Edge AI smartphones & PCs to drive AI chip growth — Among various form factors, we see AI smartphone/AI PC/robotics markets as promising given their progress of gen-AI use cases. IDC forecasts AI phone shipments to register +78% CAGR (2023-28E) and AI PC shipments to see +28% CAGR (2024-29E). Along with the explosive growth of edge AI devices, we expect overall AI DRAM demand to grow at +75% CAGR to 331bn pcs (1Gb eq.) in 2028E from 35bn pcs (1Gb eq.) in 2024, driven by both shipment and content growth of next-gen memory solutions. We forecast on-device AI DRAM demand for AI smartphones/PCs/robots with new computing architecture to register +198%/+104%/+239% CAGR (2024- 28E) respectively.

Key beneficiaries — We believe that device architectures adopted across portable IT devices will ultimately evolve into AI server-like architectures where DRAMs are placed adjacent to NPU/TPU to maximize AI operation. We identify potential beneficiaries of this shift: Hynix, Samsung, TSMC, Micron, Nvidia, Mediatek, Dell, Besi, ASMPT, Advantest, Disco, KLA, Lam Research, AMAT, Techwing, Mitsubishi Electric, Panasonic , Ibiden, Asus, MSI, Lenovo, Tokyo Ohka, SUMCO.

Peter Lee AC
+82-2-3705-0720
peter.sc.lee@citi.com

Laura (Chia Yi) Chen AC
laura.cy.chen@citi.com

Kota Ezawa AC
kota.ezawa@citi.com

Andrew M. Gardiner, CFA AC
andrew.gardiner@citi.com

Christopher Danely AC
christopher.danely@citi.com

Masahiro Shibano AC
masahiro.shibano@citi.com

Atif Malik AC
atif.malik@citi.com

Asiya Merchant, CFA AC
asiya.merchant@citi.com

Kyna Wong AC
kyna.wong@citi.com

Carrie Liu AC
carrie.liu@citi.com

Takayuki Naito AC
takayuki.naito@citi.com

Yuta Nishiyama AC
yuta.nishiyama@citi.com

Josh Yang AC
josh.yang@citi.com

Kevin Chen AC
kevin.y.chen@citi.com

Papa Sylla

Michael Cadiz

Jayden Oh

Elizabeth Sun

See Appendix A-1 for Analyst Certification, Important Disclosures and Research Analyst Affiliations.

Contents

Our Thesis in Charts

Figure 1. Computing Trend: Centralized → Distributed

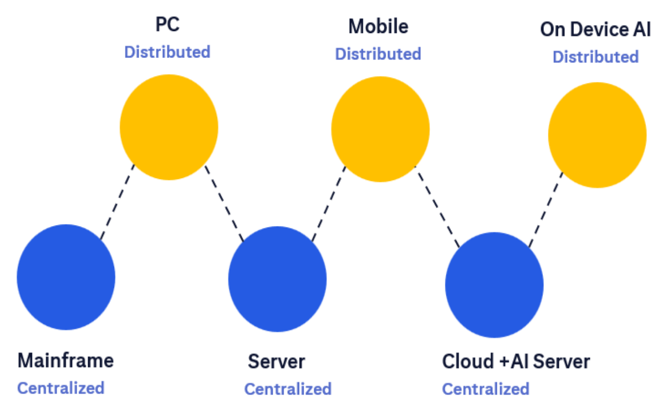

Source: Citi Research

Figure 2. Mega Trends of AI → Real-Time Processing

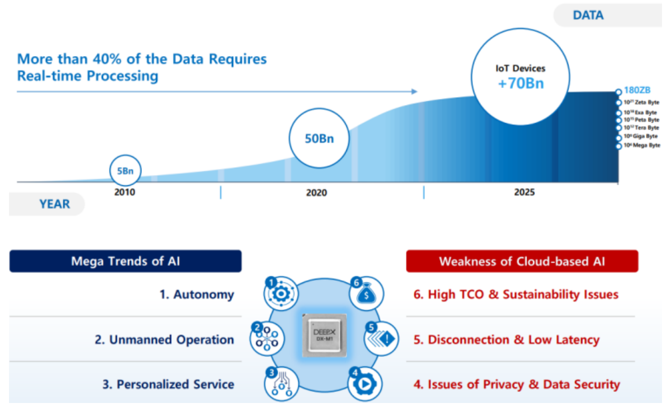

Source: DeepX. Citi Research

Figure 3. Smartphone → Mini AI Server

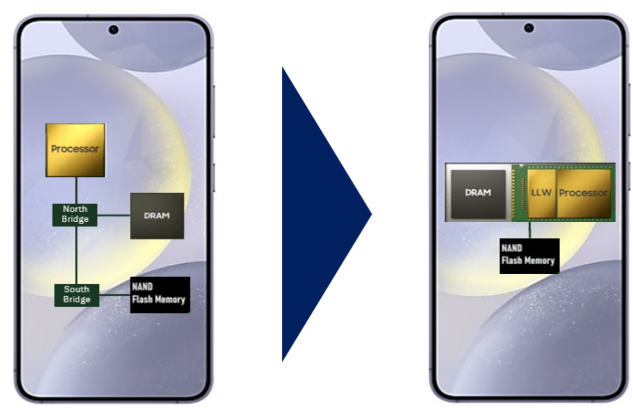

Source: Samsung, DeepX, Citi Research

Figure 4. Portable AI Technology Direction

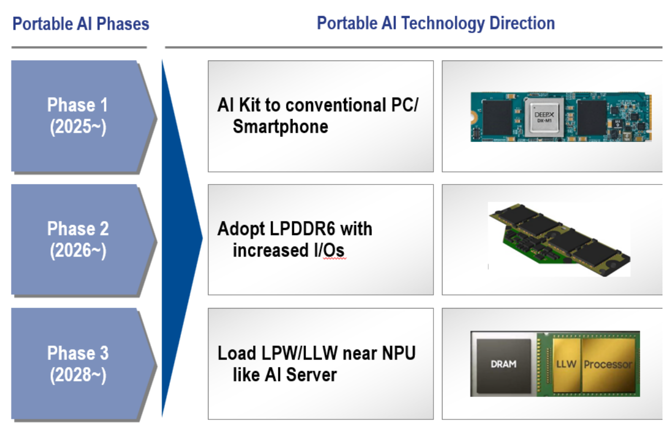

Source: Samsung, DeepX, Citi Research

Figure 5. SoIC + LPDDR/LPW + Hybrid Bonding

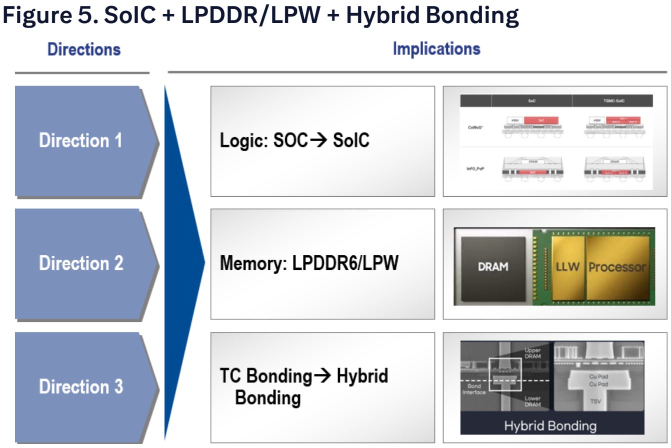

Source: TSMC, SK Hynix, Citi Research

Figure 6. Direction 1: System on Integrated Circuit (SoIC)

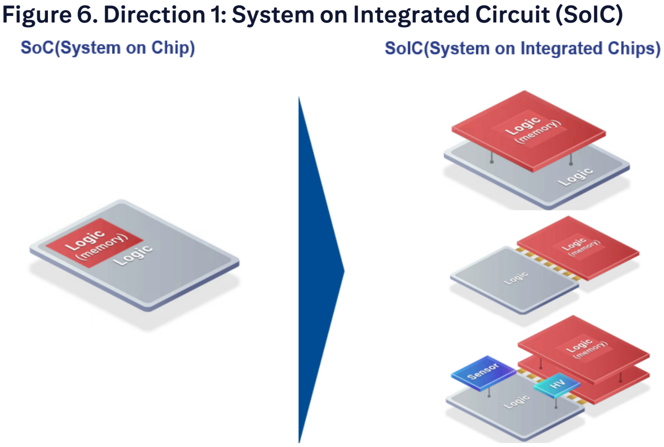

Source: TSMC, Citi Research

Figure 7. InFO PoP with SoIC

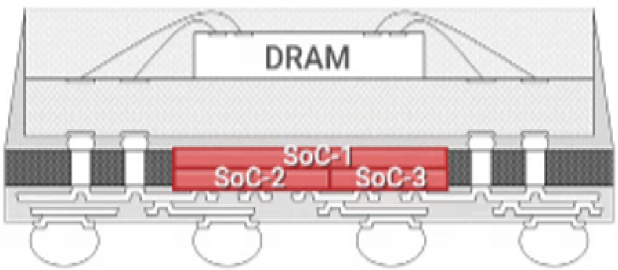

Source: TSMC, Citi Research

Figure 8. Direction 2-1 LPDDR6

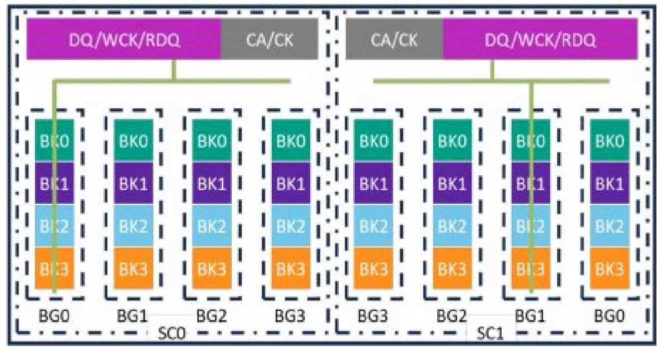

Source: JEDEC, Citi Research

Figure 9. Direction 2-2 LPW(Low Power Wide IO)

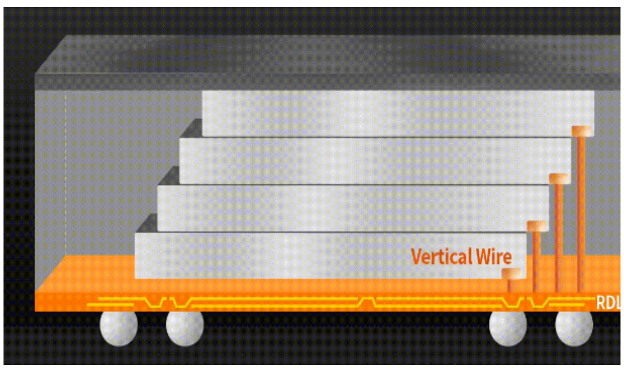

Source: SK Hynix

Figure 10.AI DRAM Demand Portion to Increase

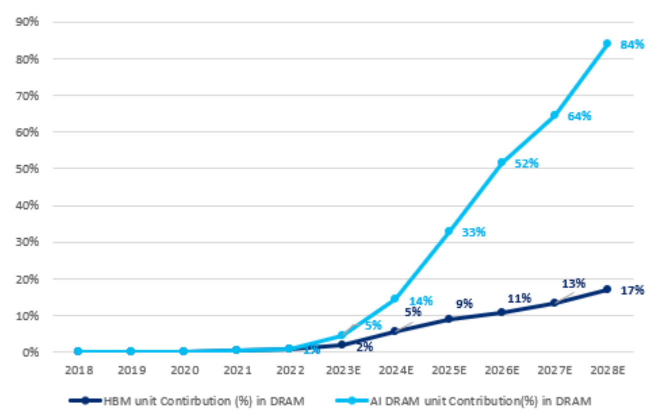

Source: Citi Research estimates

Figure 11. Hybrid Bonding

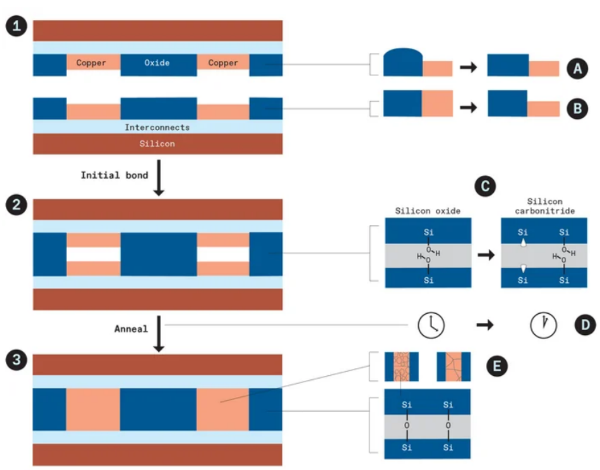

Source: Company data, Citi Research

Figure 12. On-Device AI Stocks

Category	Company Name	Product	BBG Ticker
GlobalSemiconductor	TSMC	Foundry	2330 TT EQUITY
	SK Hynix	HBM/LLW/eSSD	000660 KS EQUITY
	Micron	HBM/eSSD	MU US EQUITY
	Samsung Electronics	LLW/eSSD	005930 KS EQUITY
SPE - Frontend	Applied Materials	Etch/Deposition	AMAT US EQUITY
	TEL	Etch	8035 JP EQUITY
	LAM Research	Etch	LRCX US EQUITY
	Eugene Tech	Depostion	084370 KS EQUITY
	KLA	Metrology	KLA US EQUITY
SPE - Backend	BESI	Hybrid Bonding	BESI NA EQUITY
	ASMPT	Bonding	0522 HK EQUITY
	Sony Group	Hybrid Bonding	6758 JP EQUITY
	Disco	Dicing	6146 JP EQUITY
	Techwing	HBM Test Handler	089030 KS EQUITY
	Nextin	HBM Inspection	348210 KS EQUITY
	Parksystems	HBM Inspection	140860 KS EQUITY
	Advantest	Testing	6857 JP EQUITY
	Teradyne	Testing	TER US EQUITY
IT Hardware	Nvidia	AI GPU	NVDA US EQUITY
	Mediatek	AI ASIC	2454 TT EQUITY
	Dell Technologies	AI PC	DELL US EQUITY
	ASUS	AI PC	2357 TT EQUITY
	Lenovo	AI PC	2377 TT EQUITY
Components/Materials/Others	Ibiden	AI Substrate (ABF)	4062 JP EQUITY
	Panasonic Holdings	Advanced Package Material	6752 JP EQUITY
	Tokyo Ohka	Advanced Package Material	4186 JP EQUITY
	ISC	Socket	095340 KS EQUITY
	Leeno Industrial	Socket	058470 KS EQUITY

Source: Citi Research

Executive Summary

AI era could see most consumers owning a portable personal AI server

We project AI servers to become extremely efficient and so compact that one day people will carry them in their hands as easily as they do a smartphone. We expect these personal AI servers in the future to become portable in various form factors such as smartphones, laptops, smart glasses, and robots.

Looking at the history of computing, early computers were confined to laboratories and utilized for limited purposes such as military or national defense. However, with the emergence of PC in 1980s, computing entered everyday life. In the 2000s, laptops and smartphones were invented and mass-produced, enabling people to own and carry their own devices. Similarly, we anticipate a transition in AI computing from centralized server-based infrastructure to personal on-device AI servers. Contrary to the current AI operation concentrated in heavy server-level infrastructure facilities, as personalized AI devices emerge, we expect an era in which most consumers own a portable personal AI server.

DeepSeek impact: Distilled AI models should accelerate on-device AI demand

While AI-trained models have traditionally required heavy hardware investment, the emergence of distilled AI models such as DeepSeek has sent a disruptive signal across the industry given its 1) significantly lower training and development cost, 2) benchmark performance comparable to cutting-edge AI models, and 3) open-source accessibility.

DeepSeek has implemented highly efficient algorithmic ways to maximize model efficiency. In order to boost cost efficiency, three major algorithm-based technologies synergistically work in DeepSeek's model: 1) Mixture-of-Experts (MOE) architecture allows for efficient scaling of model capacity, 2) Group Relative Policy Optimization (GRPO) enables cost-effective reinforcement learnings, and 3) Knowledge Distillation transfers these capabilities to smaller, more deployable models.

All in all, DeepSeek's comprehensive approach has significantly advanced the field by showing how strategic combinations of reinforcement learnings and knowledge distillation can dramatically improve reasoning capabilities across various model sizes. We believe such advances will accelerate the advancement of AI models that can be minimized in a small parameter scale to operate in edge AI devices. As a result, we project the emergence of on-device AI demand, which should drive the change of computing structure and content growth of semiconductors.

We believe that Von Neumann computing structure, widely adopted across IT devices including PCs, smartphones, and TVs, will ultimately evolve into AI server-like architectures where LPW (or LLW) DRAMs are placed adjacent to NPU/TPU to maximize AI processing efficiency.

Three directions of future computing architectures for on-device AI

We foresee the emergence of portable AI servers triggering three main architectural shifts in on-device AI products.

- The first architecture direction is adding an AI kit to conventional Von Neumann architecture. While this architecture retains the conventional structure of PC and smartphone, it adds an AI module via PCIe to boost the capability. This is the most convenient way to upgrade without changing the architecture of conventional IT devices. However, we believe this does not fundamentally solve the limitation of the existing architecture and lacks the optimization needed for AI computing.

- The second architecture direction involves the use of near memory or LPDDDR6 with increased I/Os & bandwidth near NPU & TPU. This is a more advanced and promising architecture than the first one, offering a more AI-friendly data path that improves computing efficiency.

- The third architecture direction places LPW/LLW DRAM directly adjacent to NPU/TPU, which is similar to Nvidia's AI server architecture. The third architecture offers the most outstanding performance, but it comes with a highest cost.

Our checks suggest that PC suppliers are considering the first architecture, and we believe key industry participants are also exploring the second and third.

With the expansion of AI distillation technology such as that pioneered by DeepSeek, we expect a significant reduction in AI model sizes, accelerating architectural changes in portable devices where AI models can be embedded directly into edge devices. As a result, we believe mini AI servers will emerge as key applications, and expect smartphones – which are the most responsive to rapidly changing use cases – will take the lead in terms of the architecture change.

Edge devices require low power consumption and better thermal efficiency to become a viable portable AI server. With limited network throughput in edge AI environments, unlike datacenter-use AI servers, portable edge device AI servers must optimize cost efficiency and energy efficiency to fully operate AI models.

Technology directions of on-device AI computing architectures

Thus, future semiconductors for mini-Ai servers are likely to be designed with mobile-device-like characteristics such as low power consumption and cost efficiency. To achieve such efficiency, we highlight three major hardware architecture schemes below:

- We anticipate broader adoption of heterogeneous integration, where high-end mobile APs combine processors, GPUs, SRAM, and other components. TSMC's SoIC are expected to play a central role in enabling this integration within edge devices.

- We expect the adoption of mobile AI DRAM will continue to expand. Specifically, LPDDR6 and next-gen mobile DRAMs are projected to enter mass production as early as 2027E, with some products likely emerging in 2H25E. LPDDR6 is expected to be prioritized for flagship devices starting in 2026E, while by 2028E, LPW DRAM is likely to become the mainstream choice, effectively increasing bandwidth and serving as an alternative to HBM.

- We expect die-to-die integration to widely adopted, which should result in growing importance of hybrid bonding technology. Hybrid bonding is essential for connecting between two dies and is the most efficient technology for that purpose. Although it has the drawback of higher cost compared to current TC bonding technology, hybrid bonding would likely be adopted to realize mini AI server architecture.

The Era of Portable AI Server

On-device AI tech & new computing architecture support portable personalized AI servers

We project AI servers to become extremely efficient and so compact that one day people will carry them in their hands as easily as they do a smartphone. These personal AI servers are expected to become portable in various form factors such as smartphones, laptops, smart glasses, and robots in the future.

We anticipate a transition in AI computing from centralized server-based infrastructure to personal on-device AI servers. Contrary to the current AI operation concentrated in heavy server-level infrastructure facilities, as personalized AI devices emerge, we expect an era to come where most consumers own a portable personal AI server.

Computing trend: centralized to distributed to centralized to distributed

Looking at the history of computing, early computers were confined to laboratories and utilized for limited purposes such as military or national defense. However, with the emergence of PCs in the 1980s, computing entered everyday life. In the 2000s-2010s, PCs and smartphones were invented and mass-produced, enabling people to own and carry their own devices. The tech industry has recently shifted from 'distributed' to 'centralized' with the advent of new technologies such as server and cloud/AI computing. However, similar to what we have seen historically, as personalized AI devices emerge, we expect tech mega trends to shift back to "distributed" in the future.

Figure 13. Technology Advancement from AI

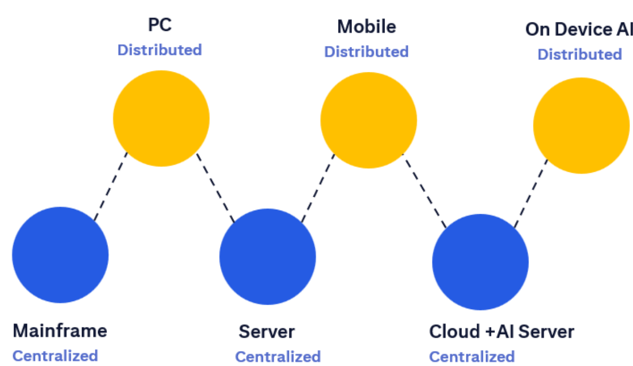

Source: Citi Research

Currently, tech companies that aim to capitalize on AI computation have accelerated the conversion of conventional servers into AI servers at data centers where most of AI computing work is done in heavy server-scale facilities. However, as an overwhelming amount of data is accumulated at edge devices, the need for imminent training and reasoning within the edge device is expected to significantly increase to effectively address the data, in our view.

We believe current devices have leveraged on "Von Neumann" architecture, which cannot fully enable diverse AI functions. We think current architecture has limitations on the full-scale operation of "real" AI functions.

Figure 14. Mainframe → Portable Computers → AI Server → Portable AI (On-device AI)

Source: AMAT, SIA, Citi Research

Meanwhile, the recent development of the DeepSeek model suggested a possibility that AI models could deliver a satisfactory level of functionality even at edge level when the AI model could be minimized efficiently. We believe this advance implies edge AI has strong potential to be aggressively adopted in various areas such as smartphones, 5xPC, robots, and automotive, along with new computing architecture.

Figure 15. Semiconductor Market Outlook

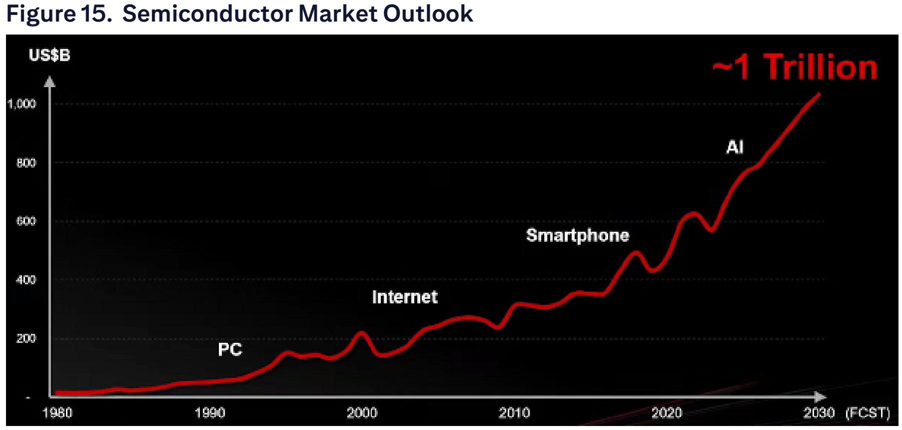

Source: TSMC

Figure 16. Technology Advancement from AI

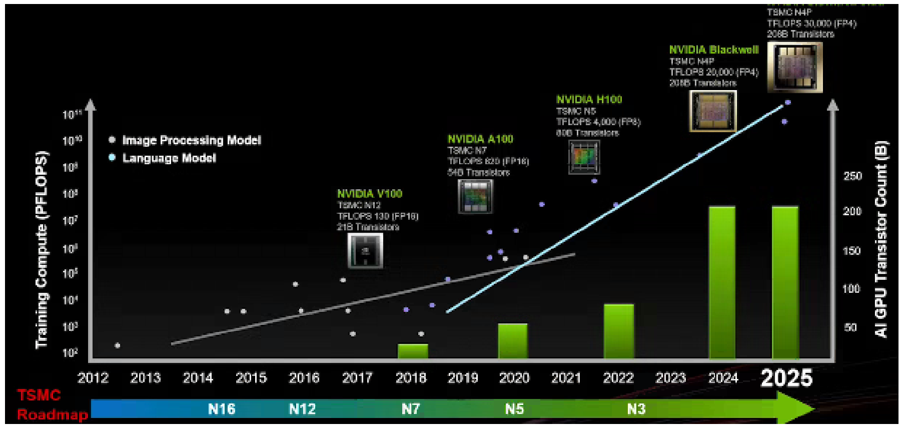

Source: TSMC

Figure 17. Technology Platform for HPC/AI

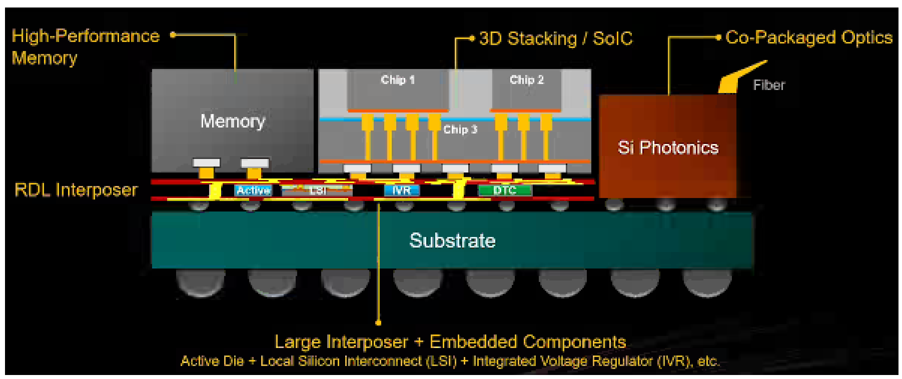

Source: TSMC

Figure 18. Global Semiconductor Industry Direction

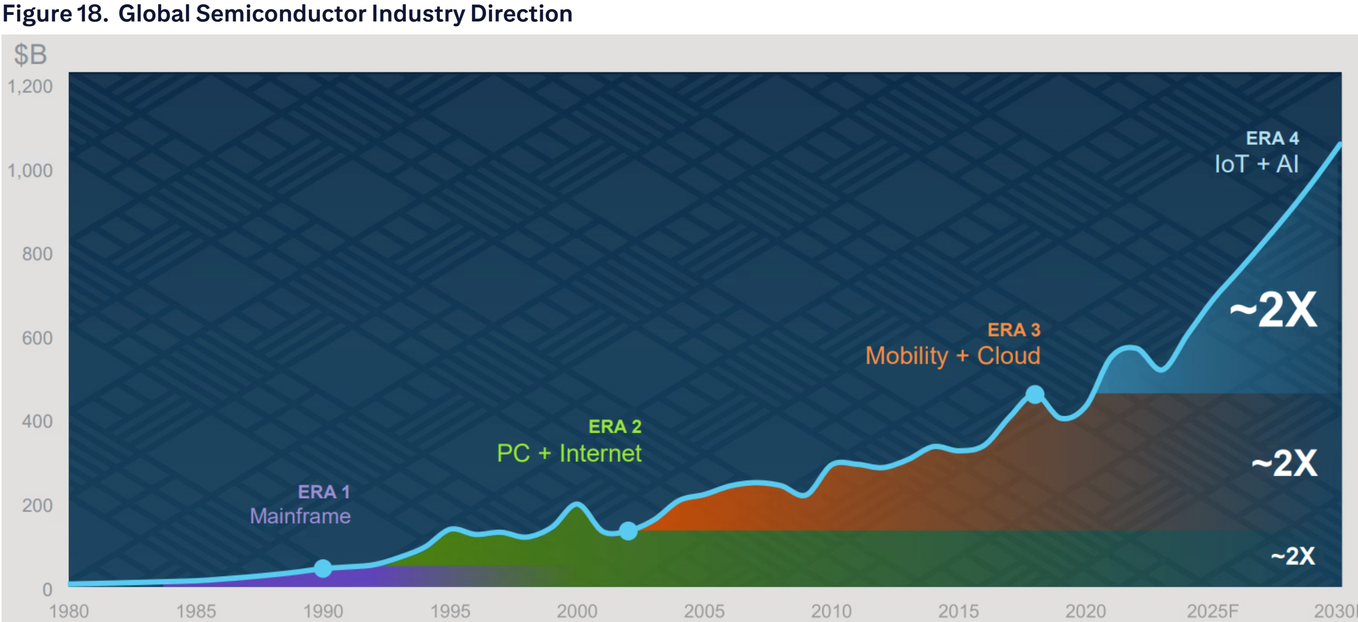

Source: AMAT, SIA, Citi Research

AI Model Efficiency Unchained

DeepSeek opens door to efficiency-optimized on-device AI models

While AI-trained models have traditionally required heavy hardware investment, the emergence of distilled AI models like DeepSeek has sent a disruptive signal across the industry given its 1) significantly lower training and development cost, 2) benchmark performance comparable to cutting-edge AI models, and 3) open-source accessibility.

Figure 19. Benchmark Performance of DeepSeek-R1

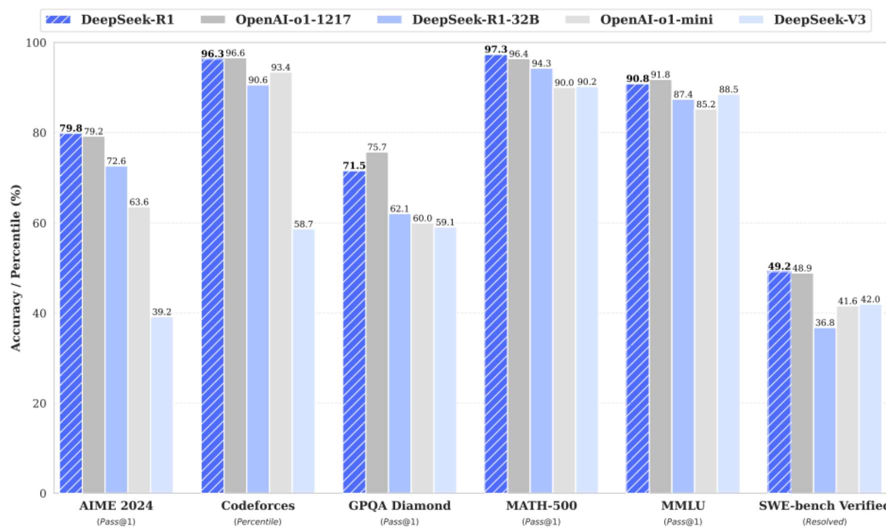

Source: DeepSeek, Citi Research

According to DeepSeek, despite having leveraged on fewer GPUs and lower versions of GPUs, its distilled DeepSeek-R1 model demonstrated comparable capabilities (or outperformed in several benchmarks) even compared with the latest OpenAI-o1-mini and Open-o1 models (Figure 1). This came as a significant surprise to the industry, as the training and development cost of DeepSeek's models was claimed as significantly lower, under US$6mn, while using Nvidia's less-advanced H800 chips (Reuters, 27-Jan-2025).

DeepSeek's algorithmic innovations maximize efficiency of AI computations

DeepSeek has adopted pure algorithmic ways to maximize model efficiency. Such efficiency originates from DeepSeek-V3-Base, which utilizes Mixture-of-Experts (MOE) architecture featuring 671bn total parameters, but only 37bn activated during inference. The architecture allows each input token to be routed to only a subset of the available neural networks so that the model can pursue efficient balance between knowledge capacity and computational requirements, compared to other LLMs. As 5.5% of the total parameters for any given input gets activated, this dramatically lower its GPU dependency.

With the foundation knowledge generated by DeepSeek-V3-Base model, DeepSeek enhance sand refines data through parallel paths: DeepSeek-R1-Zero and DeepSeek-R1, then democratizes the outcomes in smaller architectures through its Distilled DeepSeek R1 model. We highlight key technology of each model below.

DeepSeek-R1-Zero applies reinforcement learnings directly to the base model without intermediate fine-tuning. This approach utilized Group Relative Policy Optimization (GRPO), a reinforced algorithm eliminating the need for a separate critic model, significantly reducing computational costs. Through GRPO, the model targets for accuracy, proper formatting, and sophisticated reasoning behaviors without explicit programming.

DeepSeek-R1, which incorporated lessons from R1-Zero, employs four-stage training pipelines: 1) Cold Start with Supervised Data, Reasoning-Oriented Reinforcement Training, 3) Rejection Samling and Mixed Fine-Tuning, and 4) Comprehensive Reinforcement Learnings.

DeepSeek-R1-Distill model is the most fine-tuned model, which transfers the reasoning capabilities from DeepSeek-R1 to smaller, more efficient architectures through Knowledge Distillation. Distillation refers to leveraging few excellent samples (around 800K) carefully curated during DeepSeek R1's development and then allowing small model mimics the reasoning patterns and problem-solving approaches of the larger model, R1. DeepSeek presented this approach proved significantly more effective than applying reinforced learnings directly to smaller models.

Figure 20. Evolution of DeepSeek Models

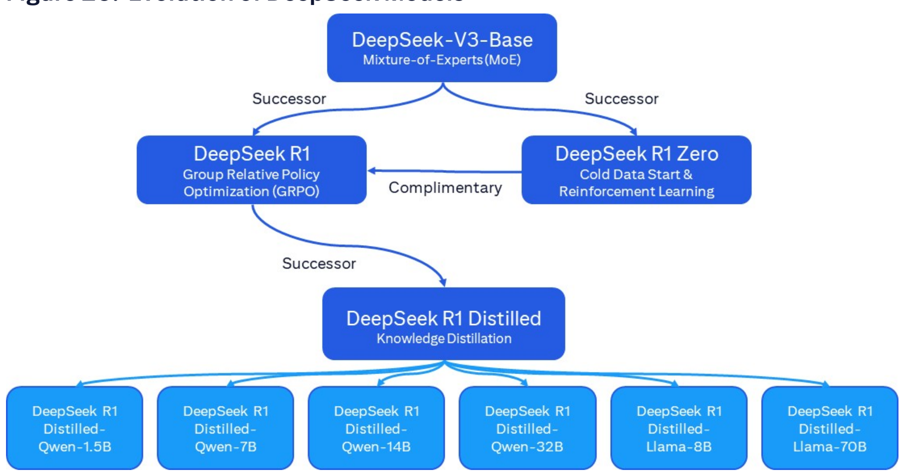

Source: DeepSeek, Citi Research

As a result, three major technologies synergistically work as following: 1) Mixture-of-Experts (MOE) architecture allows for efficient scaling of model capacity, 2) Group Relative Policy Optimization (GRPO) enables cost-effective reinforcement learnings, and 3) Knowledge Distillation transfers these capabilities to smaller, more deployable models.

Hybrid AI model supports edge AI boom

High performance, lower-power processors or dedicated accelerators of GPU and TPUs are commonly used on a cloud AI. Though LLMs are cloud-based models and mostly share a common thread in the broader landscape of AI hardware, with DeepSeek's launch, we now see increasing capability of models with fewer parameters processed on the edge, enabling more context-aware and responsive

Laura Chen
+886-2-8726-9090
laura.cy.chen@citi.com

devices. A hybrid method combining both cloud and edge processing is a realistic way, in our view, to strike a balance between local responsiveness and the power efficiency of cloud-based resources. To fit LLM into edge devices with limited hardware resources, model compression techniques are crucial, including quantization, pruning, knowledge distillation and model-size reduction, while preserving essential information at the same time. More importantly, some edge devices with sufficient computing power can perform local training, especially for smaller models, which would become more common when the model needs to adapt to specific user preference or local conditions. With support from on-device training, these AI service providers could adapt and improve their language models based on interactions with users without relying on constant updates from the cloud.

For more complex or resource-intensive processing tasks, edge devices may offload certain computation to the cloud. Edge devices may dynamically load a specific portion of the language model based on the task at hand. This adaptive loading helps to optimize resource usage and responsiveness, with the process involving sending data to a cloud server, where a larger, more powerful LLM can process the information and send back the results. And cloud-based services can facilitate continuous learning by updating and fine-tuning the language model based on aggregated data from various edge devices. This will enable the model to evolve and improve over time.

In addition, processing data locally on the device enhances data privacy by reducing the need to send sensitive information to the cloud, providing better security and privacy protection. This is especially important in applications where user privacy is a primary concern. Even when data needs to be sent to the cloud, secured communications protocols can be implemented to protect sensitive information during transit.

The evolution of LLM, as exemplified by DeepSeek, demonstrates the move to cost-efficient innovation and ability to reproduce. We view the semiconductor supply chain, in particular TSMC, as playing a critical role in enabling AI advancement. We also expect RTX-based PC to be adopted for AI-inference and edge applications, making AI PC a bridge between large-scale cloud deployment and local computing. We note that MediaTek recently announced its partnership with Nvidia on GB10-chip for Project Digits that aligns it with the rising trend of edge AI and personal computing. This collaboration allows MediaTek to diversify beyond its consumer electronics focus, positioning it to benefit from more AI popularity. We expect to see MediaTek's computing business outgrow its smartphone business in the following two years.

Deepseek impact: distilled AI models accelerate on-device AI demand

All in all, DeepSeek's comprehensive approach has significantly advanced the field by showing how strategic combinations of reinforcement learnings and knowledge distillation can dramatically improve reasoning capabilities across various model sizes. We believe this advance will accelerate the development of AI models in a way of being minimized in a small parameter scale that could operate in edge AI devices. As a result, we project on-device AI demand to emerge in 2H25E, which should drive the change of computing structure and content growth of semiconductors.

IT Device Architecture Direction

IT device architecture to transform like mini AI server

We believe that Von Neumann computing structure adopted across IT devices including PC, smartphone, and TV will ultimately evolve into AI server-like architecture where LPW (or LLW) DRAM locates right next to NPU/TPU to maximize AI functionality.

Figure 21. Von Numan Architecture → Mini AI Server Architecture

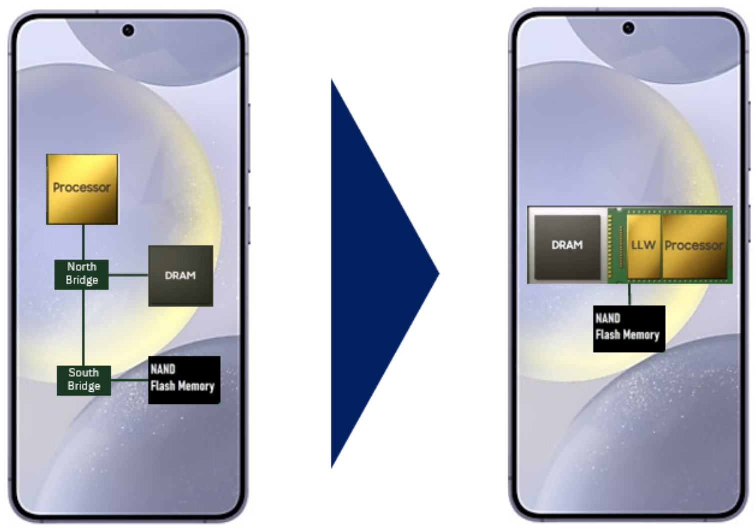

Source: Samsung, Citi Research

Three directions of computing architectures for on-device AI

We foresee the emergence of portable AI servers would entail three types of architectural changes in on-device AI products (see Figure 22).

Figure 22. Von Numan Architecture → Mini AI Server Architecture

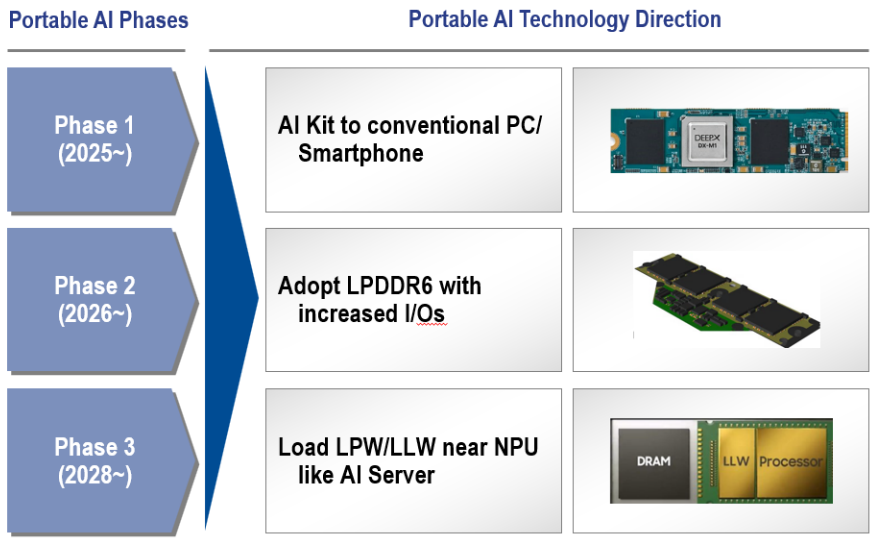

Source: DeepX, Samsung, Citi Research

The first architecture direction is adding an AI kit to conventional Von Neumann architecture. While this architecture maintains the conventional structure of PC and smartphone, it adds an AI module in PCIe slot to boost the capability.

This is the most convenient way to upgrade without changing the architecture of conventional IT devices. However, we believe this is not a fundamental solution to fully enable AI computing, which lacks the capability to optimize AI computing.

Figure 23. Phase 1: AI Kit to Conventional PC/Smartphone

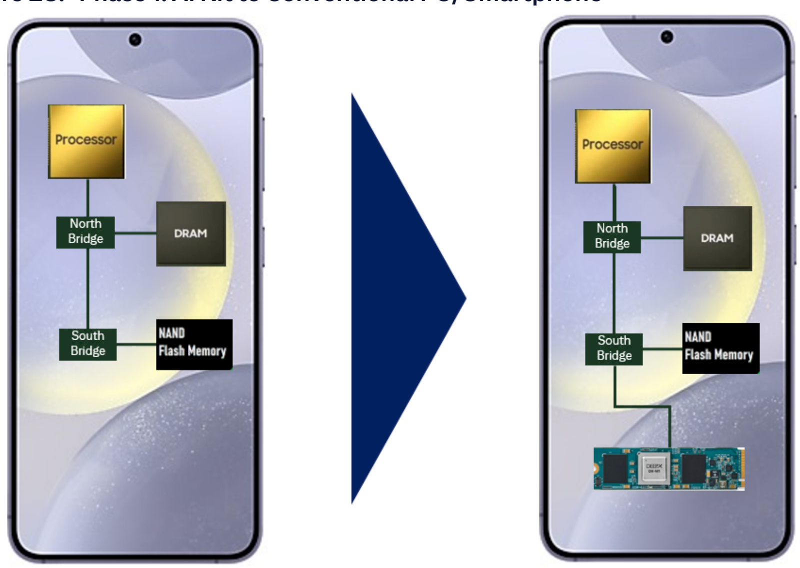

Source: Samsung, DeepX, Citi Research

The second architecture scheme enhanced from the first one is utilizing near memory or LPDDDR6 with increased I/Os & bandwidth near NPU and TPU. The second scheme is a more powerful technology that could enable AI computing more effectively compared to the first one.

Figure 24. Phase 2: Adopt LPDDR6 with Increased I/Os

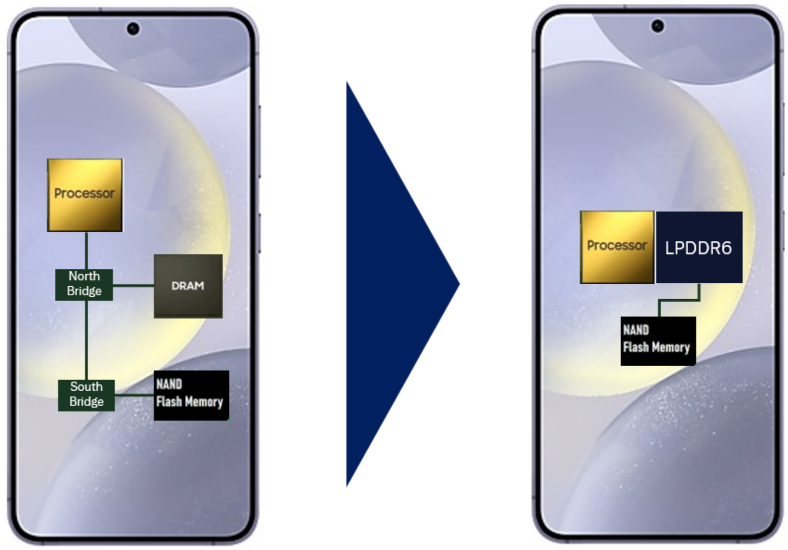

Source: Samsung, Citi Research

The third architecture direction is loading LPW/LLW DRAM near NPU/TPU, which is similar to Nvidia's AI server architecture. The third architecture would offer the most outstanding performance, but it would be the least favorable option in terms of the cost.

Figure 25. Phase 3: Mini AI Server Architecture with LPW/LLW

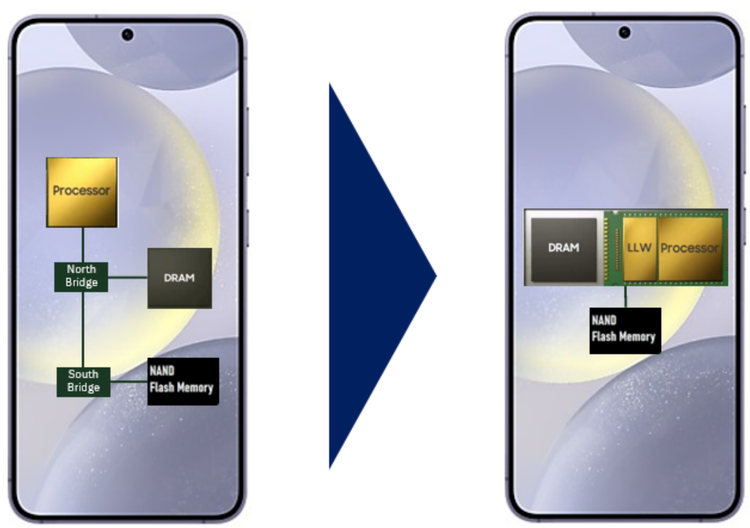

Source: Samsung, Citi Research

Our supply-chain checks suggest that PC suppliers are considering the first architecture for on-device AI version, and we believe key industry participants are also exploring the second and third.

On-Device AI Development Status

On-Device AI Trend

Atif Malik
+1 415-951-1892
atif.malik@citi.com

Asiya Merchant
+1 415-951-1752
asiya.merchant@citi.com

Papa Sylla

Michael Cadiz

Elizabeth Sun

AI on PC

While underlying demand drivers for Windows 10 EOL remain intact, our expectation for PC demand in 2025 is more tempered given the potential impact from tariff-induced price inflation and macro uncertainty. 1Q demand has been positively influenced by tariff pull-forward, and we expect some momentum to sustain in 2Q (given tariff pause). However, uncertainty surrounding US tariffs and associated inflationary pressure and global recessionary risks should negatively impact demand for PCs in the following quarters in 2025. Given Covid PC hardware installations are running long in the tooth, pushing five years vs. typical replacement of 3-4 years, we would expect demand shift into CY2026 and model PC unit growth of +4.7% yy.

Within this total PC backdrop, 2024 AI PC shipments (defined by IDC as containing either NPUs or discrete GPUs of 100TOPs+ (trillions of operations per second)) reached 76.7M, or ~30% share, and are projected to grow at a 28.4% CAGR (2024-2029E) to 267M units and or ~98% share, with PC revenue TAM expected to grow at 6-7% CAGR as AI PCs proliferate and raise ASPs from current levels, and largely consistent with Intel's recent PC TAM growth estimate of 3-5% growth. We do note, however, that vs. six months ago, IDC's five-year CAGR (then 2023-2028E) has moderated from 42.1% to the current 28.4% (2024-2029E), directionally in line with Intel's recent earnings commentary that lower-than-expected 1Q AI PC growth should be also then be followed with lower growth. For now, AI PC demand remains somewhat muted as enterprises continue to face a lack of killer apps and attractive use cases and are waiting for Intel's Panther Lake, other incremental silicon advancements, as well as improved scale and volume which would then help drive system ASPs below $1,000 before mass adoption of AI PC systems, where more attractive price points are needed given the current tariff-induced macro.

Meanwhile from a product perspective, vs six months ago, many new SKUs have also come to market, including Intel Core Ultra Series 2, Qualcomm's Snapdragon X Plus (leading the drive for sub$1,000 price point), while Apple announced M4 family Pro and Max variants (38TOPS), plus NVIDIA Series 50 GPUs, including RTX 5090 (3,000TOPS). The industry is expecting improved demand in 2026 as price points improve and more compelling use cases emerge, including efforts by NVIDIA, Intel, AMD, Qualcomm, and Microsoft working with independent software vendors (ISVs) to build new AI features that can better leverage local silicon, which should accelerate traction through the coming years, including agentic AI which IDC expects to serve as a significant inflection point.

Why AI on PCs?

While running AI software in the cloud (whether public or private) is typical, the practice can become cumbersome and prohibitively expensive, and we believe we will see AI directly on client devices primarily on:

- Improved performance by eliminating interfacing with the cloud and over network.

- Enhanced privacy/security, by keeping the data at the source – where there is increased risk of private corporate data comingling with public sources.
- Lowered costs by limiting the (costly) access needed to the cloud.

Currently, new AI PCs will be powered by specific system-on-a-chip (SoC) capabilities, including standard CPUs, GPUs, NPUs, and AI CPUs (GB10 developed by Nvidia and MediaTek). We would expect additional announcements by Nvidia and MediaTek to advance AI on PC silicon developments.

AI PC Leaders

Numerous companies, including Intel, AMD, Apple, and Qualcomm (with its Snapdragon X chips) have developed silicon with purpose-built AI accelerators residing on-chip alongside standard CPU and GPU cores.

Microsoft's and Intel's co-developed definition defines that an AI PC will come with a Neural Processing Unit (NPU), CPU, and GPU that support Microsoft's Copilot and comes with a physical Copilot key directly on the keyboard, replacing the second Windows key on the right side of the keyboard. Copilot is an AI chatbot powered by an LLM currently being rolled into newer versions of Windows 11.

Currently, Copilot computation occurs in the cloud, but executing the workload locally will provide improved latency, performance, and privacy benefits as well as free up cloud servers for more compute-intense workloads. The goal is to create systems that are quicker at executing AI tasks in a more energy-efficient manner, and eliminate the need to send data – especially sensitive data – to cloud-based AI servers for processing. This approach ensures that systems can operate independently of an internet connection and enhance security by retaining the data locally.

From a hardware perspective, AI-enabled PC hardware is typically defined by its minimum CPU computational requirement of approximately 40TOPS. Early AI PC silicon offerings were only capable of 10-16 TOPS (typically INT4), while Apple (on ARM architecture) announced its Pro and Max variants of the M4 family coming in at 38TOPS, Nvidia/MediaTek currently maintain a roadmap of 180-200 TOPS compute performance, and Nvidia's DGX Spark personal supercomputer is rated at 1,000TOPS (at FP4 precision) and is powered by its GB10 processor, while the company's new Series 50 GPUs can reach 3,352TOPS. All in, IDC notes that while silicon builds out installed base and software compatibility, OEMs strive to get systems into trials and pilots, while IT departments then try to identify appropriate use cases.

Additional hardware requirements and leading players

While CPU & GPU play critical roles in system performance, memory also plays a critical role. So from a memory perspective, the amount required for optimal performance depends on the complexity of the task, size of datasets and complexity of algorithms. Larger datasets and complex algorithms require more memory to process optimally, and AI training models usually require more memory vs. inference. For AI-enabled PCs to function optimally, parameters require preinstallation, with the computer requiring at least 1.8-2.0GB of memory for each 1B parameters, according to dialogues with our experts. So for example, running GPT-3.5 (with approximately 17B parameters) would require 17-20 gigabyte of memory. Larger datasets and more complex algorithms often require 32GB or more for smooth processing and memory bottleneck avoidance, with vast datasets requiring memory in the range of 64GB to as high as 128GB for optimal performance.

Performance has improved over the course of a year, where most currently, Nvidia DGX Spark (GB10 processor, 1,000TOPS, starting at US$3,999) features 128GB of coherent unified system memory, while DGX Station (GB300, Blackwell Ultra GPU, starting at US$99,000) can be configured between 160GB-784GB of coherent memory.

Similarly, Dell's partnership with Nvidia includes laptops and desktops with Blackwell and Blackwell Ultra GPUs, designed for pre-deployment AI model testing. Pro Max GB300 (Blackwell Ultra GPU + Grace CPU) includes 784GB of unified system memory (288GB HBM3e memory + 496GB LPDDR5X memory) for up to 460 billion parameters, while the Pro Max GB10 (Blackwell GPU + Grace CPU) can handle AI models with up to 200 billion parameters.

Education drives use cases and, ultimately, adoption

We believe continued education and increased familiarity with the technology will build positive momentum surrounding AI-driven use cases that spur demand for hardware upgrades (newer/faster CPUs/GPUs) to locally process large datasets more securely, cost effectively, and with lower latency (vs. accessing cloud). And while cloud currently accommodates large models (1.76tn parameters as GPT-4), experts expect PCs can process models locally (10-20bln parameters) in the near future, necessitating newer/faster hardware at the user level.

Enterprise users are currently more keenly aware of the benefits of AI on servers and in the cloud, while users at the local level are only beginning to learn of the benefits of AI. From the user perspective at the outset, AI is being marketed initially as a true enhancement to personal productivity (meeting notes, transcripts, customer service) and computing security – reducing/eliminating time-consuming tasks or complementing one's creative work adaptively, and predictively eliminating security threats. Longer-term, AI on PCs should prove to be more cost-effective vs. continuously interfacing with the cloud, in both a time (latency) and cost (transaction and access costs) perspective, in addition to being more secure locally, reducing computations relayed between the cloud and edge.

Figure 26. AI PC Definition and AI PC Shipment Forecasts

IDC's definition of AI PCs include the following categories:

- **NPU only** are PCs that have integrated NPUs but do not have discrete GPUs.
- **GPU only** are PCs that have 100TOPS+ dGPUs but lack NPUs.
- **NPU + GPU** are PCs that have both integrated NPUs and 100TOPS+ dGPUs.

Excluded are PCs that lack both an NPU and a 100TOPS+ dGPU.

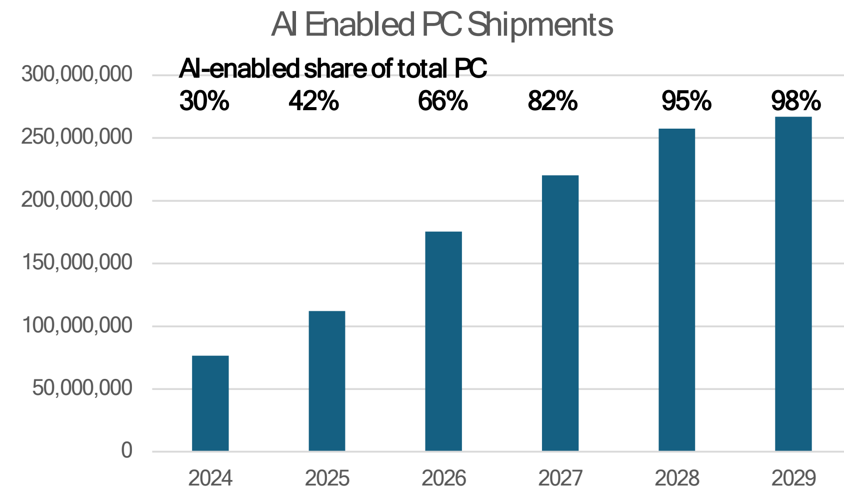

Source: IDC

Figure 27. IDC: Why an AI PC?

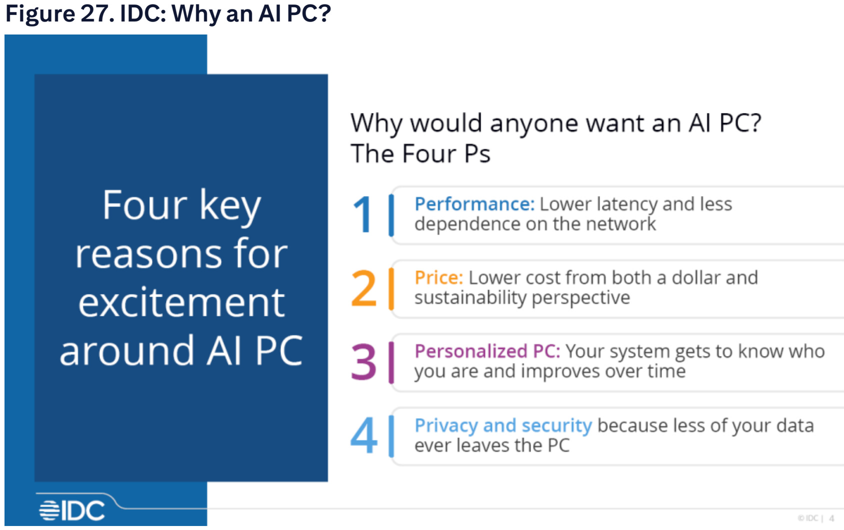

Source: IDC

AI on Smartphones

Overview

Smartphones capable of running on-device AI have been developed for nearly a decade; think, natural language processing (NLP) and computational photography. The majority of smartphones shipped today have silicon that provides some integrated AI capabilities. However, the industry is rapidly embracing the next-generation chips that will provide new and exciting features and interaction modes. IDC defines two types of AI smartphones.

Figure 28. IDC Identifies Two Types of NPU-enabled AI Smartphones

Types of NPU-Enables AI Smartphones	NPU Performance by TOPS* using int-8 data type	AI Features	Vendors of smartphone SoCs
Hardware-Enabled	<30	These smartphones utilize accelerators, or specialized processors apart from the main application processors, to operate on-device AI at lower power. In recent times, this includes a shift to the use of neural processing unit (NPU) cores. These smartphones have been on the market for almost a decade, (think, natural language processing (NLP) and computational photography).	
Next-Generation	>30	These smartphones utilize SoCs that can operate on-device GenAI models more quickly and efficiently. Examples of on-device GenAI include numerous large language models (LLMs). This category of smartphones was first introduced in the second half of 2023.	Apple MediaTek Qualcomm

*tera opeations per second

Source: IDC

Why is AI phone so important?

IDC forecasts Gen-AI smartphones to grow at +78.4% CAGR from 2023 to 2028E, compared with 2.3% for total smartphone market, which means more than 70% of the smartphones in 2028E will have gen-AI capabilities. IDC defines next-gen AI smartphone as using SoCs capable of running on-device GenAI models efficiently and an NPU (neural processing unit) with at least 30 Tera operations per second (TOPS). Apple's A17 Pro currently used on iPhone 15 Pro models, as well as MediaTek Dimensity 9300 and Qualcomm Snapdragon 8 Gen 3, fall into the category. On the memory side, 16GB RAM is considered to be a minimum requirement for next-gen AI phones. We note that Samsung Galaxy S25 Ultra (launched in January 2025) has 12GB/16GB RAM, Xiaomi's 15 Ultra has up to 16GB RAM (launched in March 2025), and OnePlus 13 (launched in January 2025) has as high as 24GB RAM. Even though the iPhone 16 Pro offers only 8GB RAM, iPhone performance outpaces competition like Android phones with less RAM because of its seamless integration of hardware and software.

AI phone use cases

Even though Gen-AI capable phones are expected to take ~20% total smartphone share this year, it's still early in the adoption cycle in terms of how AI phones can actually help with consumers' daily life. Some of the most demonstrated use cases include summarizing emails, screenshots, phone calls and meetings for both daily life and work environment, interactive learning experience for education, and also personalized emoji/image/video generation for communications. AI can also be helpful with tracking and analyzing vital body signals for wellbeing, create travel plans and enhance gaming experience to a new level. It will take time for consumers to test all these new capabilities and really see the impact of how it can improve daily life, before the adoption of AI phones go into the mass consumer market, and we view Apple as in the best position to make it possible given its leading position in the premium smartphone market and seamless integration of software and hardware.

Is the AI on smartphones race already on?

Apple has been staggering the rollout of its Apple Intelligence features promised in 2024 WWDC over the past year. The expectation has been that the big update of Siri with on-screen awareness, personal context and deep app integration will be released sometime in April/May, but Apple confirmed on March 7th that the company sees delay in these features and now expects to roll them out in the coming year. That said, in addition to AI Intelligence, we expect the company to take a hybrid approach, with perhaps some of the computation being done on cloud and others with higher latency and privacy requirements to run directly on the device. Apple potentially taking the AI on edge approach would not be new. Qualcomm has been pushing for edge AI for some time. Qualcomm previously announced working with Meta to enable on-device AI applications using the latter's LLM model, Llama 3. In fact, Qualcomm allows the export of 8B version of Llama 3 to any Qualcomm AI Engine enabled device. Additionally, the firm claims on-device AI implementation helps to increase user privacy, address security preferences, enhance applications reliability and enable personalization – at a significantly lower cost for developers compared to the sole use of cloud-based AI implementation and services. Coming back to Apple, we foresee the company's upcoming devices to further build on their predecessors' AI computing processing capabilities. The A18 pro chips powering its iPhone 16 Pro/Pro Max models contain each a 16-core Neural Engine and a 6-core GPU that is 40% faster than iPhone 16's A16 chips with 35% less power.

The memory wall challenge

As smartphones remain the consumers' primary device of choice, a widespread consumer adoption of AI applications will likely require smartphone makers to overcome couple of challenges. One sticks out, the memory wall. LLMs require parameters to be constantly going back and forth from the RAM to the processor to predict the next token. This subsequently requires parameters to be stored as close as possible to the processor (in the device most likely). For mobile devices, LLMs with parameters numbered in the 100B+ would take more memory space than desired or even possible. One of the potential solutions being explored is the reduction of model sizes. While model accuracy is a trade-off, solutions such as quantization and sparsification can shrink the 1B-parameters model with limited accuracy trade-off.

On-device AI by Companies

Key global tech companies have been making various efforts to adopt and support on-device AI features, including launch of specific chips, such as Nvidia DGX Spark's designed to support on-device AI and Qualcomm's Snapdragon X Elite, as well as adoption of edge AI features by key smartphone makers such as Samsung and Huawei to run AI locally on their devices. We discuss efforts by key tech companies in adopting on-device AI below.

Nvidia

Nvidia offers two product line-ups tailored to meet on-device AI demands: DGX Spark (formerly known as Project DIGITS), and the Jetson platform.

DGX Spark (Project DIGITS) – Project DIGITS is Nvidia's compact AI supercomputer designed to bring high-performance computing to individual users. DGX Spark will feature NVIDIA GB10 Grace Blackwell Superchip, which combines Blackwell GPU with 5th-gen Tensor Cores & FP4 support and Grace CPU, delivering 1 petaflop of AI performance. In terms of memory, Spark will be equipped with 128GB DRAM and 1TB or 4TB of NVMe storage to support efficient handling of large AI models. Spark will be small enough to fit on a desk, making advanced AI computing accessible to individuals. Spark is expected to be available by summer of 2025 with pricing starting at US$3,999.

Figure 29. NVIDIA DGX Spark

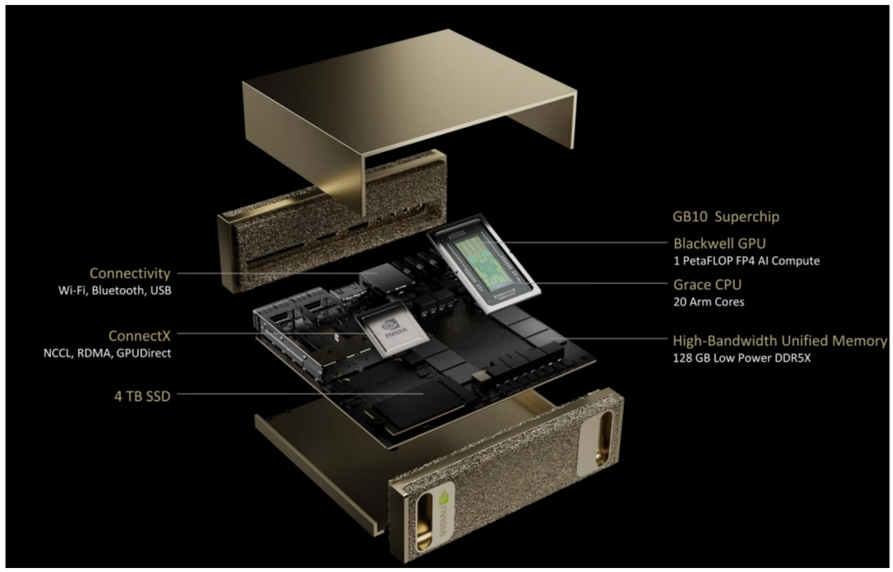

Source: Nvidia

Jetson Platform – The Jetson platform is designed to support edge AI applications, aiming to offer a range of modules to suit performance and power requirements for different edge AI applications. Jetson platform includes Jetson Orin series, Jetson Xavier series, and Jetson Nano series.

- **Jetson Orin series:** Jetson Orin series includes modules like Jetson AGX Orin, delivering up to 275 TOPS of AI performance. Jetson Orin series features 1792/2048-core NVIDIA Ampere GPUs and 8/12-core Arm Cortex CPUs, and 32-64GB LPDDR5 DRAM and 64GB eMMC NAND. Jetson Orin series is designed to be suitable for next-gen advanced robotics and autonomous machine applications.

- **Jetson Xavier series:** Jetson Xavier series includes Jetson AGX Xavier and Jetson Xavier NX, providing up to 32 TOPS of AI performance. Specifically, Jetson AGX Xavier features 512-core Volta architecture GPU and Nvidia Carmel Arm CPU, and 32-64GB LPDDR4x, while Jetson Xavier NX series features 384-core NVIDIA Volta architecture and 6-core Carmel Arm CPU, and 8-16GB LPDDR4x. Xavier series is designed for applications requiring high compute density and energy efficiency, including robotics and edge AI applications.

- **Jetson Nano series: Jetson** Jetson Nano series is an entry-level module delivering 472 GFLOPS of AI performance. Jetson Nano series is powered by Nvidia Maxwell GPU and Quad-core ARM Coretix-A57. Jetson Nano series features 4GB of 64-bit LPDDR4 and 16GB eMMC. Nano series is designed for developing AI in cost-effective and power-efficient edge devices.

Apple

Apple is advancing edge AI capabilities on iPhones through hardware innovation and software integration, including introduction of neural engine, upgraded memory capacity, and Apple intelligence.

In terms of hardware innovation, Apple has introduced neural engine on its iPhone with the adoption of A11 Bionic chip in 2017 that is dedicated to improve AI functionality on iPhones. The latest A18 and A18 Pro chips on iPhone 16 feature 16-core neural engine capable of performing up to 35 trillion TOPS to facilitate real-time AI tasks such as facial recognition and AR features. Apple is also improving memory capability on its iPhones to support AI features, and is expected to increase DRAM capacity on the upcoming iPhone 17 Pro that will be launched in 2H25E to 12GB from current 8GB, in our view. In addition, Apple is reportedly working with key memory vendors to develop a mobile version of HBM in an effort to enhance edge AI capabilities in future devices. The initiative, which seeks to integrate HBM's high-speed data processing to mobile devices, could be featured from iPhone 18 series that will be launched in 2H26E.

Figure 30. A18 PRO with 16-core Neural Engine

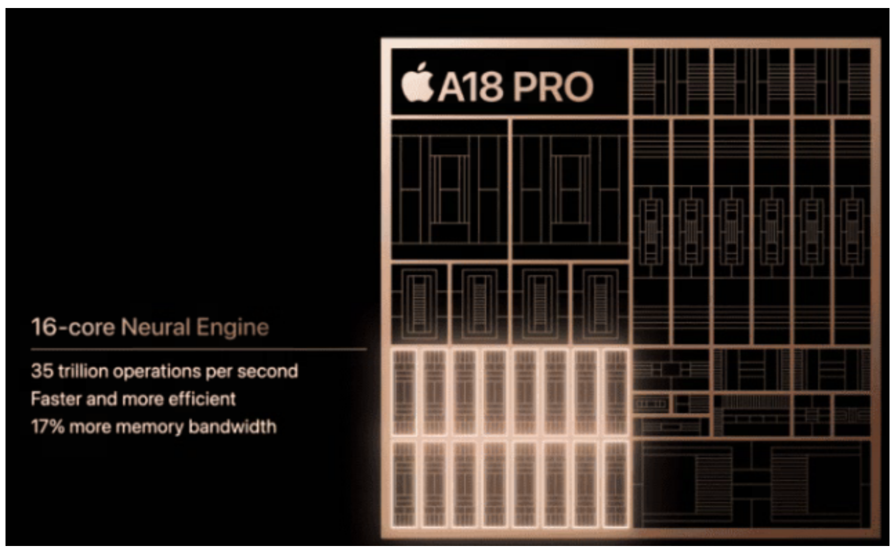

Source: Apple

In terms of software innovation, Apple has launched its own AI system Apple Intelligence in October 2024. Apple Intelligence emphasizes on-device processing, which enhances user privacy by minimizing data transmission to external servers, as well as quicker processing time vs. centralized AI models. Key features of Apple Intelligence include enhanced writing tools, on-device image generation, enhanced Siri, and Photo App improvements to improve user experience. In addition, Apple has incorporated ChatGPT on Siri to provide more comprehensive responses for certain requests.

Samsung

Samsung has also been actively adopting edge-AI features on its devices, including in Galaxy series to enhance user experience, privacy, and functionality. Samsung first introduced Galaxy AI on Galaxy S24 series that was launched in January 2024, which featured various AI applications such as real-time translation, circle-to-search, call transcript, AI-driven photo editing, and generative search tools, which significantly boosted the company flagship Galaxy S series shipment in 2024, according to the company.

To support various AI features, Samsung's Galaxy S24 series integrated Snapdragon 8 Gen 3, which featured enhanced NPU, to support new AI features such as Galaxy AI. The latest flagship, Galaxy S25, series is powered by Snapdragon 8 Elite for Galaxy, which delivers 40% improvement in NPU performance compared to its predecessor, to support more complex AI tasks, including Generative Edit on the edge-level.

Figure 31. Galaxy AI: Call Transcript

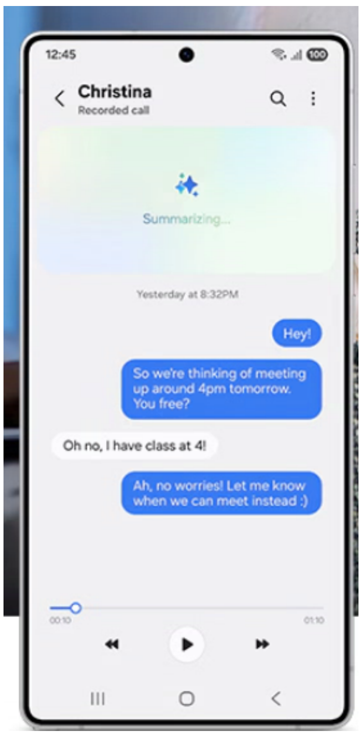

Source: Samsung Electronics

Figure 32. Galaxy AI: Circle to Search

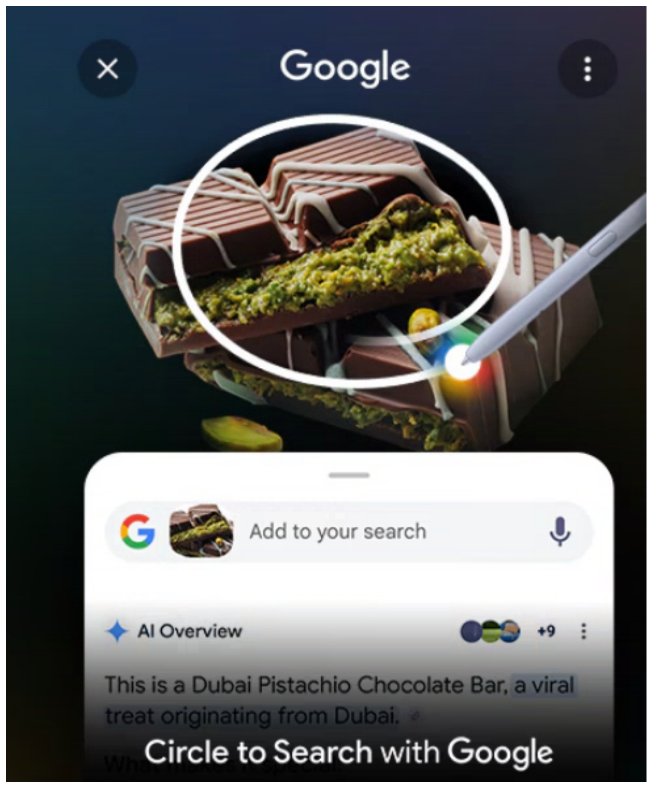

Source: Samsung Electronics

Qualcomm

Qualcomm aims to address edge-AI market with its Snapdragon X series lineup, including Snapdragon X Elite, Snapdragon X Plus, and Snapdragon X, which are ARM-based processors designed to support AI functionality on laptops and PCs. Snapdragon X platform incorporates a Hexagon NPU to support advanced AI functionalities, including Microsoft's Copilot+ while its Oryon CPU aims to provide high yet energy efficient performance. Qualcomm has also launched Snapdragon 8 Elite designed to support on-device AI functionalities on smartphones.

Snapdragon X Elite – Snapdragon X Elite 12 high-performance Oryon CPU cores with Adreno GPU delivering up to 4.6TFLOPS and Hexagon NPU capable of 45 TOPS. Snapdragon X Elite platform is featured in flagship AI laptops including Microsoft's Surface Laptop (7th gen) & Surface Pro (11th gen), Lenovo's Yoga Slim 7x, as well as Dell, HP, ASUS, and Samsung's laptops.

Snapdragon X Plus – A slightly scaled-down version of the X Elite, the X Plus offers 10 Oryon CPU cores with Adreno GPU and Hexagon NPU. Microsoft's Surface Laptop (7th gen) Surface Pro (11th gen).

Snapdragon X – Snapdragon X is Qualcomm's entry-level processor in the Snapdragon X series aimed at mainstream and budget laptops, and features eight Oryon CPU cores and 45 TOPS Hexagon NPU, supporting AI functionalities. Laptops featuring Snapdragon X include Azus Zenbook A14, Asus Vivobook 14,and Asus Vivobook S14 &S16.

Figure 33. Qualcomm Snapdragon X Platform

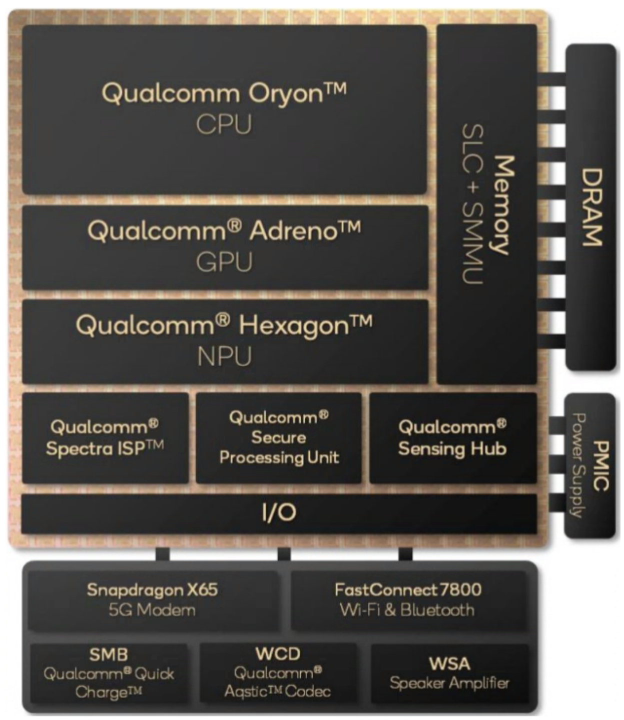

Source: Qualcomm

Snapdragon 8 Elite – Qualcomm's Snapdragon 8 Elite features 8-core Oryon architecture CPU and Adreno 830 GPU and Hexagon NPU, offering 45% improved CPU performance, 40% faster graphic rendering, and 45% AI performance improvement compared to its predecessor, to support edge AI functionalities on smartphone devices. Snapdragon 8 Elite powers several high-end smartphones, including Samsung Galaxy S25, OnePlus 13, Xiaomi 15, ASUS ROG Phone 9, and Vivo iQOO 13. Qualcomm also plans to launch the next version of the chip, Snapdragon 8 Elite Gen 2, in 4Q25E, which is anticipated to feature 2nd gen Oryon Cores and Adreno 840 GPU.

Qualcomm Edge AI Box – Qualcomm Edge AI Box is an advanced edge computing device designed to perform real-time data processing directly on the device, rather than cloud, making it particularly suitable for applications in smart retail, security, and surveillance. Edge AI Box offers: [1] low-power yet high-performance AI processing, [2] versatile connectivity, and [3] compact design suitable to be installed on legacy devices.

Key China smartphone makers

Key Chinese smartphone makers such as Xiaomi, HONOR, and Vivo have been actively integrating edge AI capabilities into their smartphone devices, employing advanced techniques such as quantization and pruning to reduce AI model sizes.

In terms of optimizing AI models for devices, leading Chinese makers such as Xiaomi, HONOR and Vivo have made meaningful strides in refining LLMs for efficient on-device performance. For example, Xiaomi's MiLM2 model reduced parameters from 6bn to 4bn, while HONOR and Vivo reduced their parameters from 7bn to 3nm to support on-device AI capabilities. Through the optimizations, they were able to lower computation costs and improve inference speeds to enable AI functionalities on their flagship smartphones more seamless and efficient.

Chinese smartphone makers have also collaborated with leading AI engine providers to integrate advanced AI functionalities. For example, Chinese smartphone makers such as Xiaomi, HONOR, and Oppo showcased AI-driven smartphones featuring Google's Gemini for AI functionalities, such as photo editing and restaurant reservations, in MWC 2025.

On the hardware side, Chinese smartphone makers have been adopting advanced chipsets, not only from Qualcomm for also from internally-developed Kirin series, to support on-device AI functionalities. Qualcomm's Snapdragon 8 Elite chipset, designed to run generative AI models locally on edge devices, has been adopted by leading Chinese smartphone makers such as Xiaomi, OnePlus, Realme, and HONOR. In addition, there have been efforts to adopt home-grown AI chips as well. For example, HiSilicon has developed the Kirin 9000 series chipset, which incorporates Huawei's Da Vinci architecture 2.0, to enhance AI processing. Kirin 9000 chip powers Huawei Mate 40 series, Huawei Mate X2, and Huawei P50 Pro.

Looking at Chinese smartphone makers' efforts in adopting on-device AI:

HONOR: HONOR unveiled its AI UI Agent, a mobile AI capable of understanding and interacting with graphical user interfaces on screens, at MWC 2025. The agent, which leverages in-house execution models as well as Google's Gemini 2, can perform tasks such as booking a restaurant without relying on external APIs.

Huawei: Huawei's HarmonyOS NEXT integrates native AI capabilities throughout its PAnGu-LLM and MindSpore framework. The system supports on-device AI processing, enhancing features such as document analysis and voice-command recognition, to offer an intelligent user experience across Huawei devices.

Figure 34. Chinese Smartphone Makers' AI Adoption

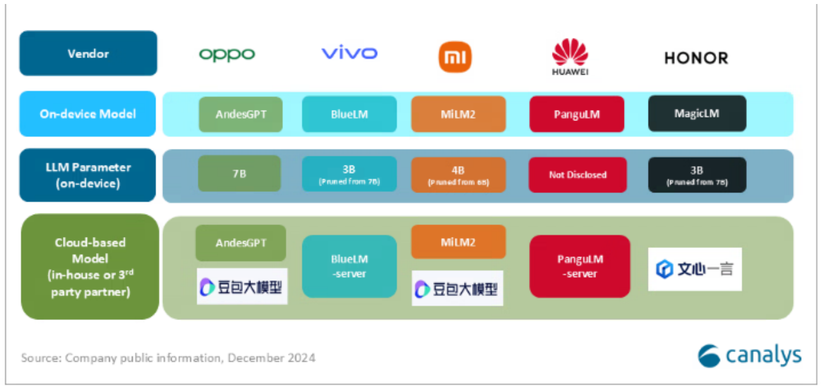

Source: Canalys

Technology Implications

Overview

With the expansion of AI Distillation technology, we expect a significant reduction in AI model sizes, accelerating architectural changes in portable edge devices where AI models can be embedded directly into the devices. As a result, we believe mini AI servers will emerge as key applications, and expect smartphones – which are the most responsive to rapidly changing use cases – will take the lead in terms of the architecture change.

We expect an edge device inevitably requires low-power configuration and better thermal efficiency to support AI operations within highly integrated circuits. With the limited network throughput of edge devices, unlike datacenter-use AI servers, we anticipate the AI functionality of portable edge device would depend on cost efficiency and the energy efficiency.

Figure 35. Mini AI Server Architecture for Edge Device (Smartphone/PC/Robot)

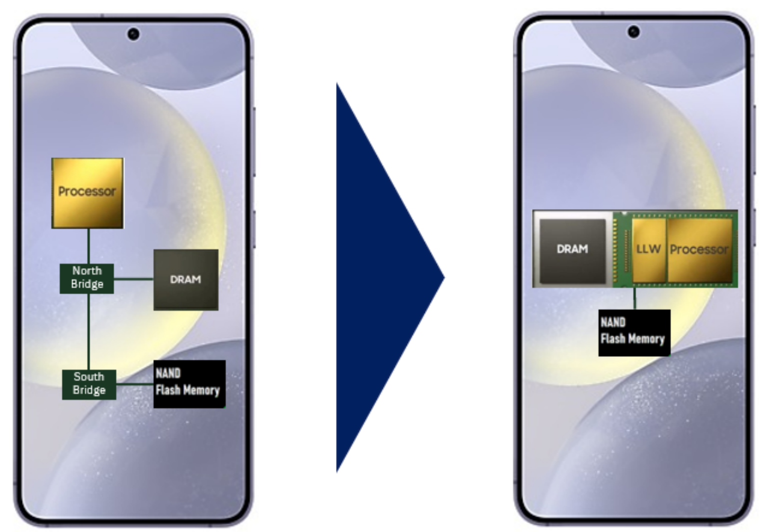

Source: Samsung, Citi Research

As future semiconductors for mini edge AI servers are likely to be designed with mobile-device-like characteristics such as low power consumption and cost efficiency, we analyze upcoming AI hardware architectural shifts to three major schemes below:

- First of all, we anticipate a broader adoption of heterogeneous integration, where high-end mobile APs combine processors, GPUs, SRAM, and other components are near memory. TSMC's SoIC are expected to play a central role in enabling this integration within edge devices.
- Secondly, we expect the adoption of next-gen mobile DRAM will continue to expand. Specifically, LPDDR6 and LPW/LLW DRAMs are projected to enter mass production as early as 2027E, with some products likely emerging in 2H25E. We expect LPDDR6 to get adopted in flagship devices starting in 2026E, while by

2028E, LPW/LLW DRAM is likely to become the mainstream choice, effectively increasing bandwidth and serving as an alternative to HBM.

- Lastly, we expect die-to-die integration to get widely adopted, resulting in growing importance of hybrid bonding technology. Given the scale of semiconductor miniaturization and closed gap between dies, hybrid bonding would become more essential for connecting between two dies. Although it has the drawback of higher cost compared to current TC bonding technology, hybrid bonding will likely be adopted to realize on-device AI computation.

Figure 36. Three Directions from On-Device AI

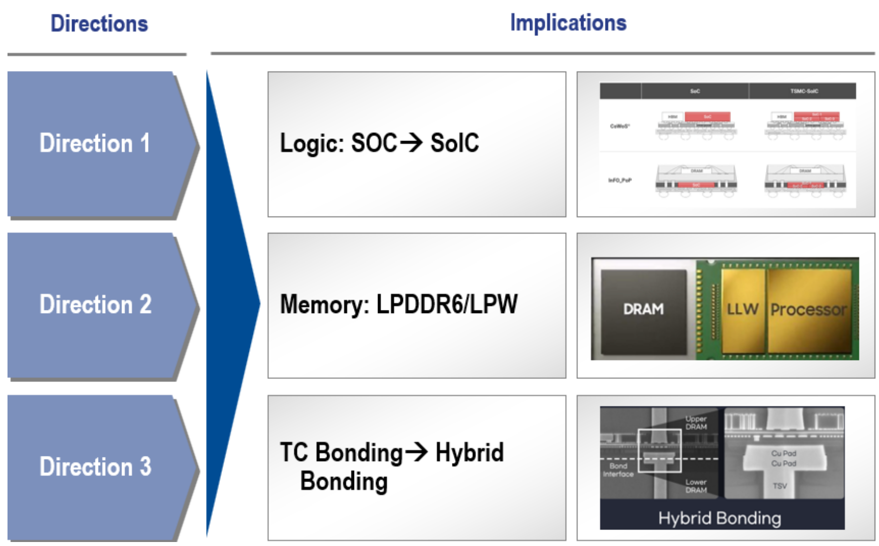

Source: TSMC, SK Hynix, Citi Research

Figure 37. SoC→ SoIC

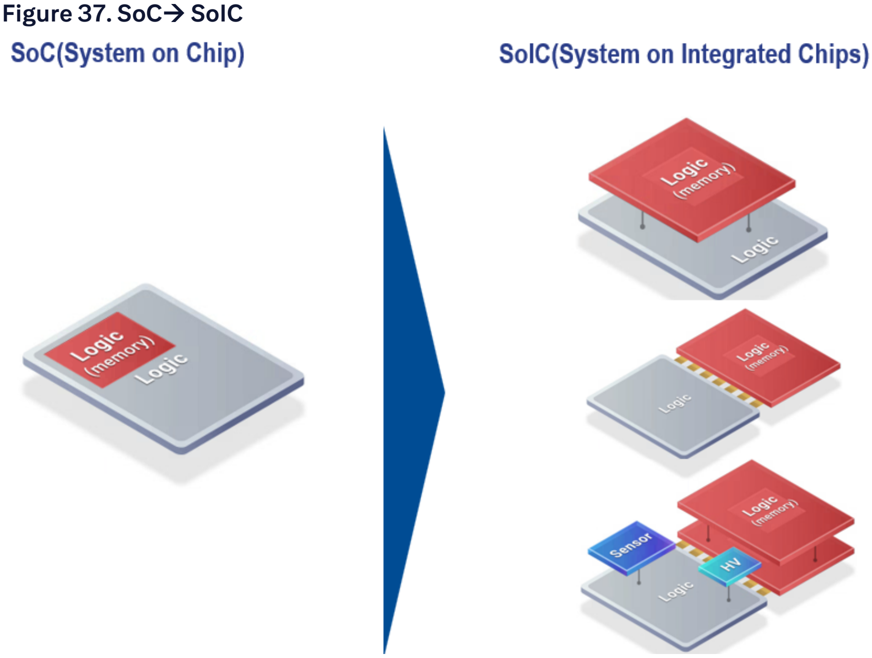

Source: TSMC, Citi Research

Laura Chen
+886-2-8726-9090
laura.cy.chen@citi.com

Direction 1: Chip integration with SoIC

AI chip integration with SoIC

As the leading semiconductor foundry, TSMC remains at the forefront of innovation when it comes to advanced packaging technologies. At the latest 2025 TSMC Technology Symposium, advanced packaging technologies, particularly CoWoS, took center stage, reflecting its strong capabilities to address the growing demands of AI, HPC and next-generation networking applications. TSMC's advanced packaging technologies covers CoWoS, InFO, SOIC and SOW. These technologies are critical for enabling heterogenous integration, where multiple chiplet are combined to optimize performance, power, and cost for AI and HPC workloads.

Figure 38. TSMC SoIC

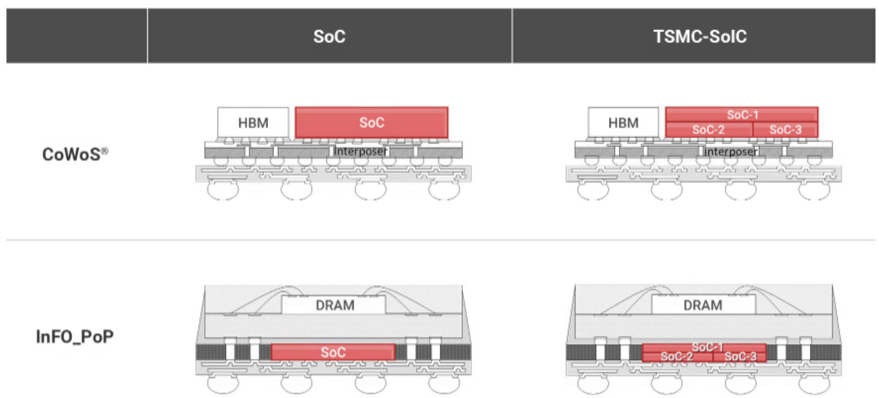

Source: TSMC, Citi Research

CoWoS – Transition from CoWoS-S to CoWoS-L

TSMC's CoWoS technology has been instrumental in enabling HBM integration with logic dies, facilitating efficient data processing. The traditional CoWoS-S utilized a silicon interposer to connect multiple dies. However, to address scalability and cost effectiveness, TSMC introduced CoWoS-L, which incorporates a local silicon interconnect (LSI) bridge embedded in an organic substrate, forming a reconstituted interposer. This design maintain the benefits of CoWoS-S while mitigating yield issues associated with large silicon interposers.

Figure 39. TSMC 3DFabric Technology Portfolio

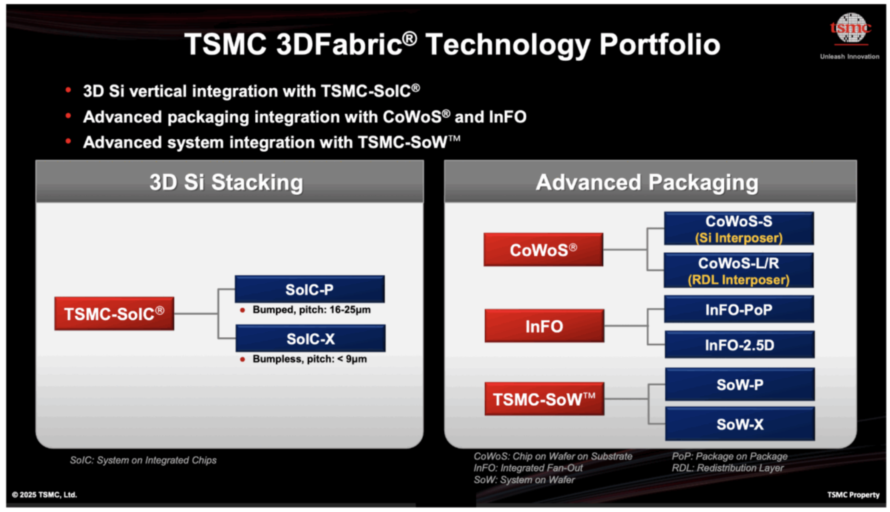

Source: TSMC, Citi Research

We note that TSMC's roadmap includes expanding CoWoS capabilities to support large interposer sizes, accommodating more complex chiplet configurations. The integration of CoWoS with other packaging technologies, such as SoIC, enables the creation of sophisticated SiP solutions, combining 2.5D and 3D integration techniques.

InFO (Integrated Fan-Out) – more for mobile device

TSMC's InFO technology provides a substrate-less packaging alternative that route I/Os through high density RDLs. It supports thinner form factors, making it ideal for mobile and cost-sensitive applications. Apple has been working with TSMC on InFO process since 2016. Its InFO-PoP is the industry first high-density 2.5D/3D fan-out wafer-level packaging. According to TSMC, InFO-PoP leverages high-density redistribution layers (RDL) and fine pitch through InFO via (TIV) technologies to integrate mobile application processors (AP) into a compact 3D system with cost and performance benefits. It features a thinner profile, thicker silicon for better thermal performance, and improved power and signal integrity (PI/SI) while eliminating a organic substrate and C4 bumps. InFO-PoP is ideal for 5G and artificial intelligence applications.

We estimate Apple will continue to leverage InFo technology for its new product in 2025 and 2026. Unlike CoWoS with a silicon substrate, InFO utilizes an organic interposer with a fan-out design. This organic interposer is typically made of polymer-based material with embedded routing traces. The fan-out design allows for a higher density of redistribution layers compared to transitional 2D packaging. This enables more connections between the logic die and passive components (e.g. capacitors, resistors) embedded within the package. With the Fan-out design, it can also bring the benefit of shorter interconnects and lower power consumption.

System on Wafer (SoW), a paradigm shift

TSMC's SoW technology represents a significant advancement in semiconductor packaging. By placing a large array of dies on 300mm wafer, SoW offers increased compute power while occupying less space. The SoW based on CoWoS is scheduled for 2027 production.

Figure 40. TSMC SoW

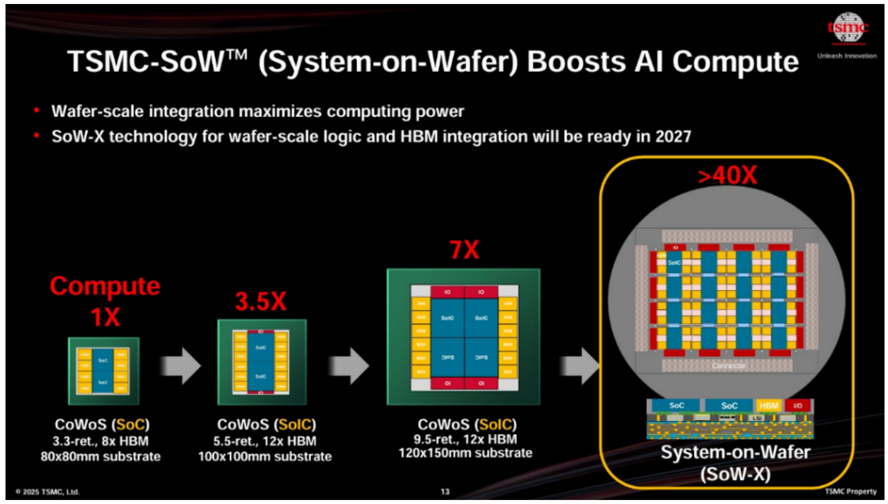

Source: TSMC, Citi Research

Silicon Photonics and Co-Packaged Optics

To address the growing demand for high-speed data transmission, TSMC introduced the Compact Universal Photonic Engine (COUPE) technology. COUPE employs SoIC-X chip stacking to integrate electrical and photonics dies, reducing impedance and enhancing energy efficiency. TSMC plans to qualify COUPE for small form factor pluggable in 2025, followed by integration into CoWoS packaging as co-packaged optics (CPO) in 2026/2027.

The CPO technology would replaces transitional electrical signal transmission with optical communications, significantly increasing bandwidth and reducing power consumption. This advancement is crucial for AI and HPC applications, where data throughput and energy efficiency are paramount.

Figure 41. TSMC Technology Platform for HPC/AI

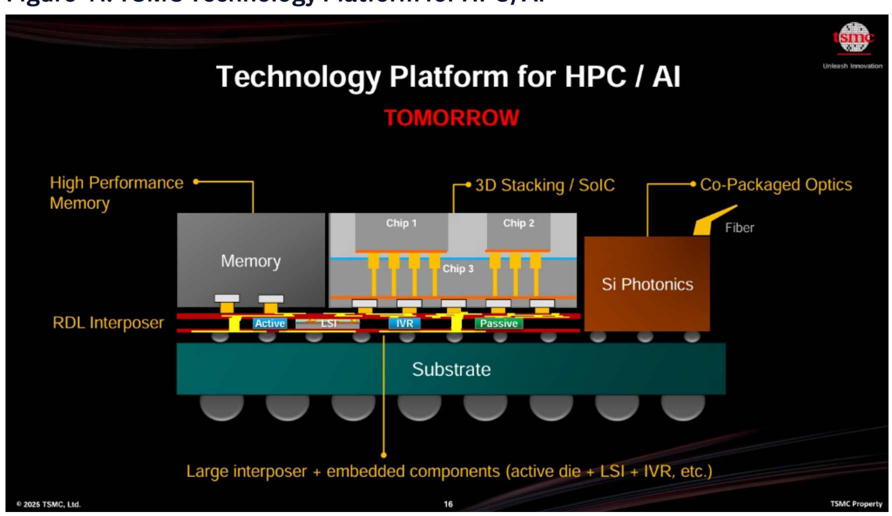

Source: TSMC, Citi Research

Figure 42. TSMC SoIC Technology Roadmap

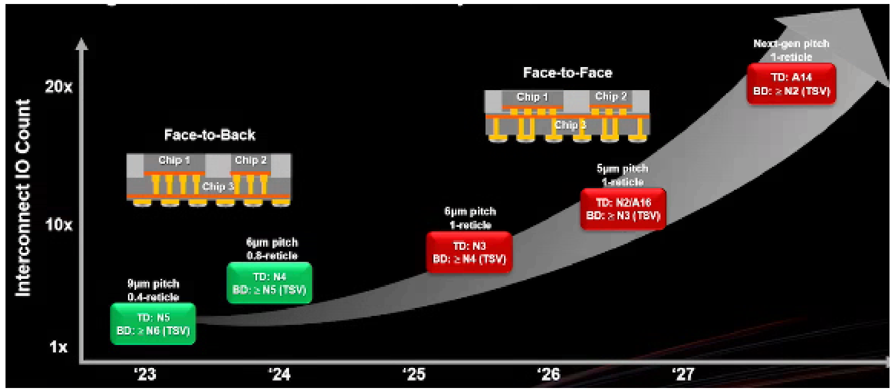

Source: TSMC, Citi Research

Looking forward when moving to 3DIC and SOIC, it would combine most of the key advantages of tighter chip stacking and shorter signal paths and potentially reduces power consumption. Nevertheless, it also requires extremely clean and smooth surface to achieve high yield rate. Particle control would also be a concern and requires advanced cleanroom environment.

We believe TSMC might involve initial SOIC solution focusing on integration of logical and passive component, followed by more advanced integration that incorporate memory and other functionalities. The applications would be primarily used for high performance computing. With advancement in die to wafer bonding (sorting good dies from the top wafer and singulation them for individual placement), it could further increase cost-effectiveness and enable more heterogeneous integration options.

Direction 2: LPDDR6 and LPW for on-device AI

Memory: LPDDR6 and LPW

We expect that the SRAM contribution in logic chips for future edge AI computing will either decline or become limited, while the use of LPDDR6 and Low Latency Wide I/O (LPW) is likely to expand. In particular, for AI phones, LPDDR6 is expected to be primarily used first starting in 2026E, but after 2028E, LPW offering a better bandwidth is anticipated to become mainstream, in our view.

LPDDR6 is the next-generation low-power double data rate memory standard, designed for mobile and other applications requiring both high performance and low power consumption. It builds upon the existing LPDDR5 standard and offers significant improvements in speed, bandwidth, and efficiency.

LPDDR6 is expected to have data rates ranging from 10.667 to 14.4 Gbps. LPDDR6 utilizes a 24-bit wide data bus, a significant increase from LPDDR5's 16-bit bus. This wider bus and higher data rates result in a significant increase in memory bandwidth, reaching up to 38.4 GB/s, according to TrendForce. LPDDR6 introduces a new 24-bit burst length, up from LPDDR5's 15-bit. We see LPDDR6 is being developed with a focus on supporting on-device AI processing in mobile devices and other applications.

Figure 43. LPDDR6 – Channel View

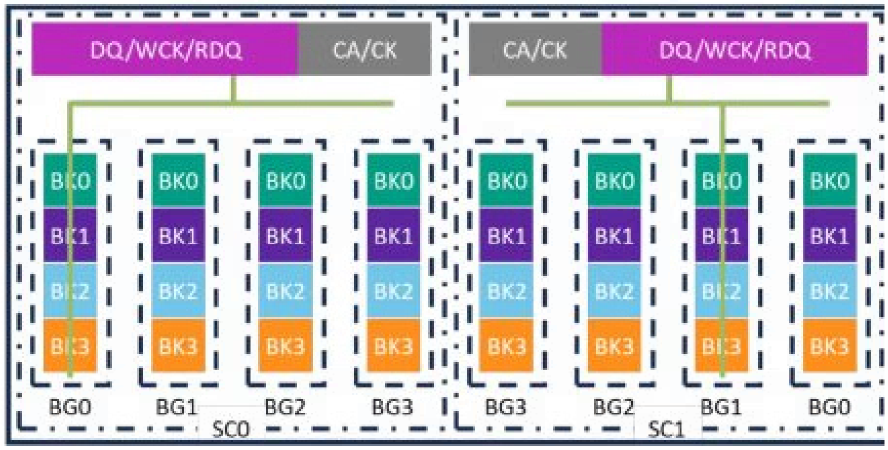

Source: JEDEC

Figure 44. LPDDR Specs by Generation

	LPDDR4	LPDDR4X	LPDDR5	LPDDR5X	LPDDR5T	LPDDR6
Max Density	64 Gbit		32 Gbit		32 Gbit (expected)	Unknown
Max Data Rate	4.3 Gbps		6.4 Gbps	8.5 Gbps	9.6 Gbps	14.4 Gbps
Channels	2		1		4	2
Width	x32 (2x x16)		x16		x64	x24 (2x 12)
Banks (Per Channel)	8		8-16		16	16
Bank Grouping	No		Yes		Yes	Yes
Prefetch	16n		16n		16n	Unknown
Voltage	1.1v		Variable Nominal 1.05v Max: 1.1v		Variable Nominal 1.05v Max: 1.12v	Unknown

Source: JEDEC, Synopsis

Figure 45. LPDDR6 CAMM2

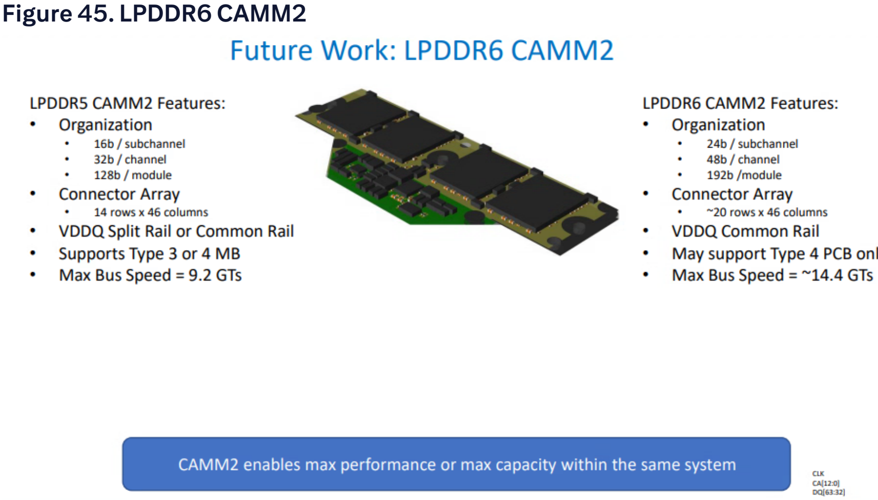

Source: JEDEC, DELL

LPW(Low Power Wide IO)

We anticipate LPW(Low Power Wide IO) DRAM, which could be seen as mini HBM for mobile, is set to emerge as a high bandwidth solution in 2028E. LPW DRAM is fabricated based on Vertical wire Fan-Out (VFO) technology based on the idea that data transfer gets quicker and shorter to travel when data transfer is conducted through vertical wires than curved wires. We anticipate increasing adoption of LLW DRAM as [1] a high bandwidth solution for mobile applications and [2] an alternative to SRAMs in processors.

LPW is a technology that increases capacity (bandwidth) by increasing the number of I/O counts compared to the conventional LPDDR products. When LPW DRAM is placed close to processors, power efficiency could be improved by about 70% compared to the general DRAMs. LPW aims to increase I/O to deliver performance like HBM in a way to suit mobile customers and LPW would be rapidly adopted starting from 2028E.

We anticipate LPW DRAM to be fabricated based on Vertical wire fan-out (VFO) technology. VFO is based on the idea that data transfer speed gets much quicker and shorter to travel a straight line than a curved one. VFO minimizes space and reduces power consumption by connecting wires vertically instead of curving them. It has also revolutionized conventional fan-out wafer-level package (WLP), a packaging technology which connects I/O terminals with wires from the outside of chips.

Recently, SK Hynix completed the development of VFO technology and started verification process. SK Hynix revealed that the wire length was reduced 4.6 times and power efficiency improved by 4.9%, compared to products that use conventional wiring. While heat dissipation also increased by 1.4%, the most notable advancement was that package thickness was reduced by 27%.

Samsung Electronics also introduced VCS technology, aiming to develop LPW (or (LLW) DRAM, which stacks four LPDDR5 DRAM dies like HBM with vertical wires in the shape of stairs, resembling SK Hynix's VFO technology.

Figure 46. Vertical Wire Fan-Out (VFO) Technology

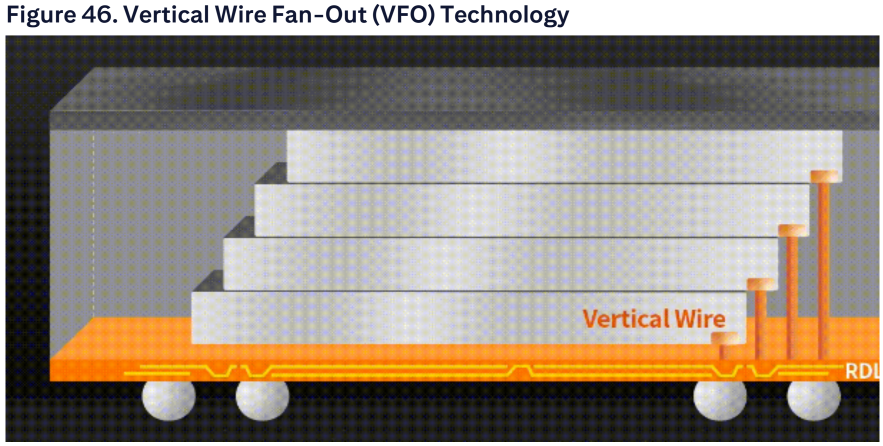

Source: SK Hynix, Citi Research

Figure 47. Number of I/Os: LPDDR5/LPDDR5X/LPDDR6

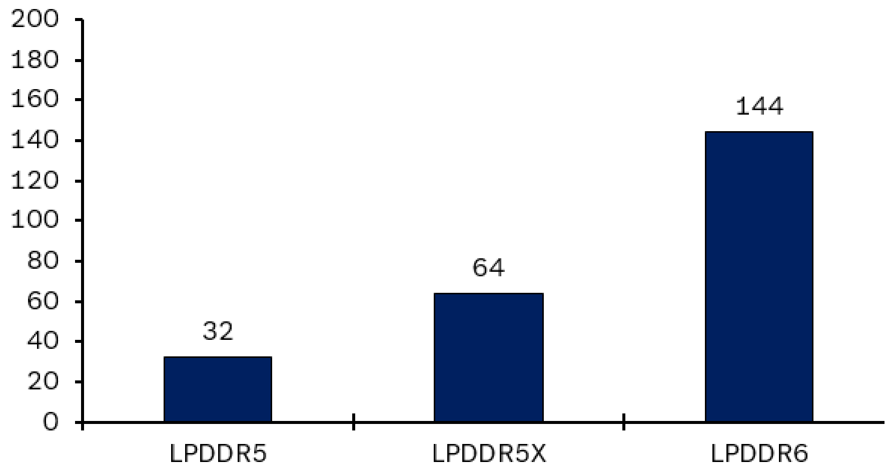

Source: Citi Research

Memory content to increase

We expect memory content per box to increase significantly driven by the emergence of efficient AI models, which would bring forward potential launches of portable AI devices. Given the nature of AI computation and large-enough parameters, even if AI models are distilled or scaled down effectively, we expect edge devices to require high-density DRAM in order to operate AI computations.

Moreover, we believe SRAM in AP or CPU would become insufficient to operate AI computation, thus computing architecture would evolve into utilizing higher-density mobile DRAMs. So far, SRAM has been used as level-1 cache memory, but has limitation in terms of capacity expansion. As a result, manufacturers would minimize and replace SRAM in AP and CPU with high-density mobile DRAMs, in our view.

All in all, we expect overall AI DRAM demand to grow at +75% CAGR to 331bn pcs (1Gb eq.) in 2028E from 35bn pcs (1Gb eq.) in 2024. We project on-device AI DRAM demand for AI smartphones with new computing architecture to reach 53.0bn pcs (1Gb eq.) by 2028E from 0.7bn pcs (1Gb eq.), representing +198% CAGR from 2024 to 2028E. For AI PC, we project AI PC DRAM demand to see +104% CAGR to 10.8bn pcs (1Gb eq.) by 2028E from 0.6bn pcs (1Gb eq.) in 2024. For AI robotics, we believe robots will meaningfully adopt on-device AI DRAM starting from 2026E, forecasting on-device AI DRAM demand for robotics to see +239% CAGR (2024-2028E).

Figure 48. AI Memory Demand Forecast

		2018	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E
Total DRAM Demand(Mpcs)		107,176	125,847	153,845	191,673	193,221	188,683	245,768	290,561	341,370	393,969	393,969
	HBM portion(%)	**0.0%**	**0.1%**	**0.2%**	**0.3%**	**0.9%**	**1.8%**	**5.5%**	**8.9%**	**10.8%**	**13.3%**	**16.9%**
HBM DRAM (Mpcs, 1Gb eq)		45	115	260	659	1,709	3,403	13,491	25,877	36,714	52,420	66,619
	HBM1(1Gb eq.)	45	81	71	27	10	-	-	-	-	-	-
	HBM2(1Gb eq.)	-	35	142	211	202	82	248	217	-	-	-
	HBM2E(1Gb eq.)	-	-	47	421	1,440	1,632	1,488	1,012	356	-	-
	HBM3(1Gb eq.)	-	-	-	-	58	1,689	3,571	2,536	1,245	649	435
	HBM3E(1Gb eq)	-	-	-	-	-	-	5,952	5,853	5,334	1,679	1,349
	HBM4(1Gb eq.)	-	-	-	-	-	-	-	1,626	10,801	20,148	13,492
HBM Price	HBM1(1Gb eq.)	1.4	1.3	1.2	1.2	1.1	1.0	1.0	1.0	0.9	0.9	0.8
	HBM2(1Gb eq.)	-	1.4	1.3	1.2	1.2	1.1	1.1	1.0	1.0	0.9	0.9
	HBM2E(1Gb eq.)	-	-	-	1.3	1.3	1.3	1.2	1.1	1.0	1.0	0.9
	HBM3(1Gb eq.)	-	-	-	-	1.6	1.6	1.6	1.6	1.5	1.4	1.4
	HBM3E(1Gb eq)	-	-	-	-	-	-	1.8	1.8	1.7	1.6	1.5
	HBM4(1Gb eq.)	-	-	-	-	-	-	-	2.6	2.6	2.5	2.4
HBM Rev	HBM1(1Gb eq.)	61	105	88	32	11	-	-	-	-	-	-
	HBM2(1Gb eq.)	-	48	186	263	239	92	282	217	-	-	-
	HBM2E(1Gb eq.)	-	-	-	553	1,800	2,040	1,786	1,113	372	-	-
	HBM3(1Gb eq.)	-	-	-	-	-	2,415	4,604	3,106	1,448	718	456
	HBM3E(1Gb eq)	-	-	-	-	-	-	9,300	9,365	8,107	2,425	1,851
	HBM4(1Gb eq.)	-	-	-	-	-	-	-	3,849	25,922	45,939	29,223
	Total HBM Revenue	61	153	274	847	2,050	4,547	19,885	43,309	81,725	154,607	269,244
AI DRAM by Applicaton	HBM	45	115	260	659	1,709	3,403	13,491	25,877	36,714	52,420	66,619
	AI Server DDR5						5,104	20,237	64,693	110,142	157,259	199,856
	On Device AI_Smartphone						-	674	3,950	24,846	36,715	52,975
	On Device AI_PC						-	623	1,298	4,914	7,124	10,793
	On Device AI_Robot							4	17	86	333	590
AI DRAM Demand		45	115	260	659	1,709	8,507	35,029	95,835	176,703	253,851	330,833
AI Unit Demand Portion(%)		**0.0%**	**0.1%**	**0.2%**	**0.3%**	**0.9%**	**4.5%**	**14.3%**	**33.0%**	**51.8%**	**64.4%**	**84.0%**

Source: Company Reports and Citi Research Estimates

Figure 49. Detailed Assumptions for AI Memory Demand Forecast

			2018	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E
Smartphone	Shipment(Mset)		1,887	1,795	1,550	1,668	1,461	1,393	1,404	1,543	1,643	1,639	1,655
	AI Phone Portion(%)								3%	8%	27%	35%	50%
	Additional DRAM(GB) per AI Phone								2.0	4.0	7.0	8.0	8.0
	On Device AI DRAM Demand from Smartphone								674	3,950	24,846	36,715	52,975
			2018	**2019**	**2020**	**2021**	**2022**	**2023**	**2024**	**2025E**	**2026E**	**2027E**	**2028E**
PC	Shipment(Mset)		255	262	297	335	284	248	260	270	273	262	265
	AI PC Portion(%)								3%	5%	15%	20%	30%
	Additional DRAM(GB) per AI PC								10.0	12.0	15.0	17.0	17.0
	On Device AI DRAM Demand from PC								623	1,298	4,914	7,124	10,793
			2018	**2019**	**2020**	**2021**	**2022**	**2023**	**2024**	**2025E**	**2026E**	**2027E**	**2028E**
Robots	Shipment(Mset)	Indutrial	0.42	0.37	0.38	0.52	0.55	0.54	0.56	0.58	0.60	0.62	0.64
		Service	0.27	0.33	0.39	0.47	0.56	0.67	0.80	0.96	1.15	1.38	1.65
		Home	16.30	19.00	22.00	26.00	31.00	37.00	44.00	52.00	61.00	71.00	82.00
		Total	17	20	23	27	32	38	45	54	63	73	84
	AI Robot Portion(%)								10%	30%	50%	70%	90%
	Additional DRAM(GB) per AI Robots								10.0	12.0	36.0	96.0	128.0
	On Device AI DRAM Demand from Robots								4	17	86	333	590

Source: Company Reports and Citi Research estimates; Note: *AI smartphone & PC refer to those with three new on-device architectures mentioned in the report.

Direction 3: Hybrid bonding technology for on-device AI

Hybrid bonding

We believe heterogeneous integration to become a prerequisite technology for on-device AI in order to achieve sufficient cost efficiency and low power characteristics. To physically realize the new integration, we expect hybrid bonding technology will become much more critical, while the increasing number of I/Os would expand the testing and socket demand in the future.

Figure 50. Hybrid Bonding

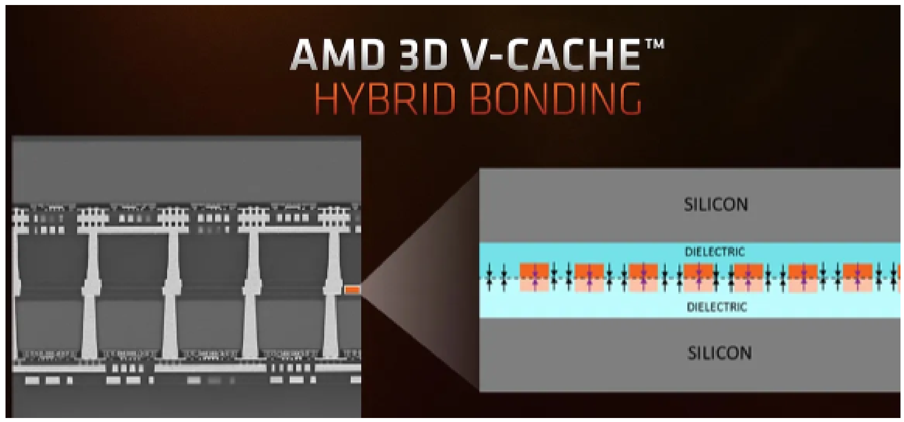

Source: AMD, Citi Research

Hybrid bonding is form of a bonding that uses hybrid bond between a dielectric bond (SiO2) and copper metal bond. Hybrid bonding extends fusion bonding with embedded metal pads in the bond interface, allowing face-to-face connection of the wafers. Hybrid bonding is an advanced bonding technique that bonds a SiO2-based dielectric surface with another dielectric surface while bonding copper-based metal surface with another metal surface.

We believe demand for hybrid bonding will be increasing due to three reasons: [1] reduced bump pitch to 5um from 100um, [2] improved signal transfer efficiency, and [3] reduced chip thickness.

To achieve this, the bump pitch must be reduced. Given that hybrid bonding doesn't use bumps to connect, bump pitch must be reduced from 20-50um in the conventional bonding to 2um level. Therefore, hybrid bonding makes it easier to form larger number of interconnects in a limited space.

The second reason for using hybrid bonding is reduced signal loss. When using solder balls or bumps, resistance occurs, resulting in signal loss and heat generation. On the other hand, with hybrid bonding, there is almost no signal loss and heat generation.

The third advantage of the hybrid bonding is that hybrid bonding can reduce chip thickness, which in turn would allow stacking of increased number of chips without significantly increasing the overall chip stack height. As hybrid bonding doesn't require intermediaries, such as solder ball or bumps, stacked chip's thickness can be significantly reduced.

Hybrid bonding process

The Hybrid Bonding process can be categorized into three stages. The first process is the preparation process, which involves die surface preparation. The second process is the bonding process, which involves die to wafer bonding. The third process is the annealing process, which involves copper to copper bonding.

Figure 51. Hybrid Bonding Process

Source: Company Data, Citi Research

Figure 52. Hybrid Bonding Process

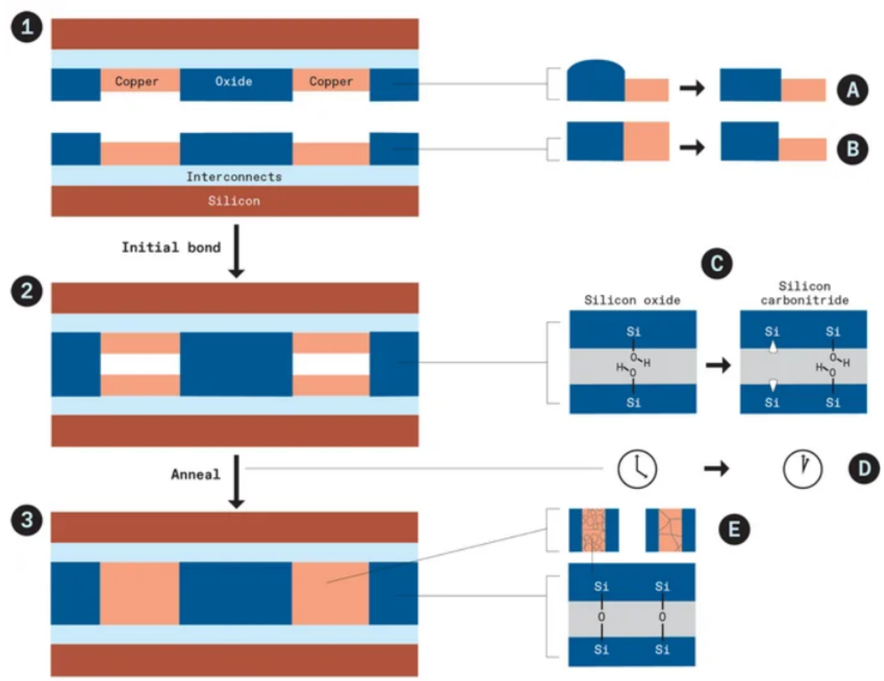

Source: IMEC, Citi Research

The dielectric portion is flattened out while empty space is intentionally created through dishing on the copper portion due to the difference in thermal expansion coefficients between dielectrics and metals. Afterwards, the surface is activated through plasma treatment to create a hydrophilic condition. Afterwards, preparations for preliminary bonding are made. The most important process in the first stage is the CMP process.

The second step is temporary dielectric bonding. The via holes are positioned to face each other and then pressurized with a low pressure of 15 N, which creates hydrogen bonding between dielectrics due to Van der Waals forces. This is not a strong bond as it is a molecule-to-molecule bond, and only dielectrics are connected to each other while the metal is not bonded yet.

The third step is annealing for metal bonding. First, by annealing at 150-degree, the covalent bond between dielectrics is created, and leads to complete bonding of dielectric portion. Afterwards, by annealing at 350-degree, copper portion expands and fills the empty space, and forms a metallic bond, which leads to complete bonding of copper portion.

As explained earlier, because metal has a higher coefficient of thermal expansion than dielectric, empty space is intentionally created with dishing in the first stage. In the third stage, the metal expands further and copper to copper bond is created.

Hybrid bonding suppliers

Andrew Gardiner
+44-20-7986-4206
andrew.gardiner@citi.com

Hybrid bonding is in the early stages of adoption, with leading-edge logic the primary driver of orders and revenue at present. In supporting this market, Besi is the leading supplier of hybrid bonding systems, having an installed base of around 100 units, with around half in Taiwan, most of the remainder in the U.S., and testing tools elsewhere. The installed base has roughly doubled in size over the past year or so.

In 1Q25, Besi received hybrid bonding orders from two leading memory producers for HBM4 applications, plus follow-on orders from an Asian foundry for logic (TSMC). Intel already took significant volume of hybrid bonders in late 2024 / early 2025. While advanced logic chips from Intel and AMD (via TSMC) are in production using hybrid bonding, the HBM DRAM market is still in evaluation mode. We are encouraged to see two leading makers order hybrid bonding tools in 1Q25, but per management's comments, these remain evaluation tools rather than destined immediately for volume production. Equipment makers are still uncertainty whether hybrid bonding will be adopted for HBM4E or 5HBM.

Besi is the primary supplier for these applications and benefits from longstanding relationships with TSMC and Intel and a partnership with Applied Materials. Besi currently offers 100nm machines (more advanced than peers) and will ship the first 50nm machines by the end of 3Q/early 4Q25. These latest tools will become production ready during 2026, with mainstream production possible in 2027.

We also see EV Group (Austria, private), ASMPT, Hanmi Semiconductor and Hanwha vision are also developing hybrid bonding tools.

Other implications: ABF substrates

AI architecture evolution contributing to demand growth for package substrates

Kota Ezawa
+81-3-6776-4640
kota.ezawa@citi.com

Takayuki Naito
+81-3-6776-4644
takayuki.naito@citi.com

We expect growth in AI server GPUs and edge AI devices to drive growth in ABF substrate demand. We have a Buy rating on the Japanese electronic component maker Ibiden as a leading supplier.

ABF substrates are plate-shaped components that play a role in, among other things, 1) protecting semiconductor chips from the external environment, 2) increasing the mounting surface by virtue of their high density, and 3) ensuring electrical connectivity with PCBs. Chiplet architecture is increasingly being adopted as chip performance advances.

Chiplet architecture increases package substrate size and connects semiconductors (logic chips) and memory on an interposer. The bandwidth obtained with HBM chiplets enables fewer components and shorter circuitry, enhancing electrical signal transmission characteristics and holding down power consumption.

After expanding supply capacity for the recovery in PC demand and growth in general-purpose server demand during the pandemic, ABF substrate players were confronted with protracted inventory adjustments due to a general-purpose server correction and fallback in PC demand. However, demand for chiplet architecture is now growing on the back of AI server GPU demand, and ABF substrates are seeing renewed growth centered on the high end.

The adoption of chiplet designs with interposers is increasing ABF substrate size as noted above, and is resulting in increased value added as well as higher operating rates via larger production loads. The standard ABF substrate size for PCs was 37.5mm*37.5mm with eight layers, whereas the chiplet product standard is 90mm*90mm with fourteen layers. Size and number of layers is also increasing in GPUs for AI servers. For example, given the production load of a single AI GPU's ABF substrate as 1.0 for the H100, the level roughly doubles for Blackwell architecture and rises 3.6x for Rubin. We therefore expect value added to rise along with production load.

At present, Ibiden and Unimicron supply ABF substrate for AI GPUs, and we understand Ibiden's share is around 90%. We expect Ibiden to maintain its high share as a beneficiary of AI device expansion thanks to its strengths in cutting-edge products given that improvements in product yield are an issue for large multilayer products with over 10,000 circuit vias as well as the need to counter heat-related substrate warpage.

Other implications: Semiconductor materials

Could be broad benefits from uptake of on-device AI

Yuta Nishiyama
+81-3-6776-4643
yuta.nishiyama@citi.com

We think there will be broad benefits for semiconductor materials suppliers if the emergence of low-cost AI models such as DeepSeek drive accelerated uptake of on-device AI. Over the past few years, there has been a decoupling in the semiconductor market between AI-related applications and other chips. Semiconductor material demand is fundamentally correlated with semiconductor volumes, and as AI-related applications have only a low weighting on a volume basis they have lacked momentum relative to the semiconductor market itself (Shin-Etsu Chemical estimates AI-related applications account for only around 10% of the overall market on a wafer-area basis). We think that as uptake of on-device AI increases semiconductor materials will increasingly benefit from a broadening in the AI-related user base (growth in volumes).

The use of silicon interposers more than doubles the amount of wafers used. Increased layering in high bandwidth memory (HBM) also increases wafer area. Increased uptake of hybrid bonding should help reduce chip thickness, so if this also increases laying, there should be benefits for wafer manufacturers. In addition, hybrid bonding could increase chemical-mechanical planarization (CMP) processes by 10–20%, and this could drive growth in the volumes of metallization CMP slurry that are used. Tokyo Ohka has large market shares in micro-bump resists and interposer redistribution layer (RDL) resists. Nissan Chemical supplies temporary adhesives used in through-silicon via processes and targets sales of more than ¥1.0bn in 2027E and more than ¥3.0bn in 2030E.

Risks to Our Thesis

While we believe the long-term outlook for on-device AI remains compelling, we analyze a number of key risks which may decelerate the shift into on-device AI architectures versus our expectations discussed above.

Demand risks

Potential lack of compelling use cases could pose a demand risk to edge AI device demand. If there is no killer use-case for edge AI, consumers may remain unconvinced as to the desirability of such devices, which may hinder a potential shift from cloud-based AI to on-device AI.

Tariff-related risks

Trump administration tariff policies concentrated on IT devices may also delay progress of on-device AI projects of device makers, as tariffs could create inflationary pressures on goods and weaken consumer demand.

Hardware risks

We identify rising difficulty of hardware manufacturing as another risk factor for slower-than-expected integration of on device AI. Currently, only few major companies such as Qualcomm and Apple are developing and securing edge AI devices. As the new AI-on-device architectures remain immature, the majority of semiconductor OEMs and ODMs may take more time to secure relevant development capability or tool chains to support full operation of edge AI devices. Moreover, if next-gen advanced packaging technology is further delayed, this could lead to slower-than-expected adoption of new edge AI architectures.

Regulatory risks

Given that AI-powered services can draw in extensive regulatory scrutiny – with current examples including regulations such as the EU AI Act and US AI safety guidelines – we project edge AI may lead to additional governmental frameworks to ensure compliance and liability with safety rules, which could result in slower-than-expected adoption. However, we believe one potentially mitigating factor is that edge AI operation can reinforce the security of individual information compared to cloud-based AI operation.

Companies

Acer (2353.TW)

Carrie Liu
+886-2-8726-9086
Carrie.liu@citi.com

We rate Acer shares as Sell as we expect its target PC market to shrink, in addition to market-share loss amid PC industry consolidation. We believe the uncertain macro and FX environment in both DM and EM would weigh on consumer PC demand and negatively impact Acer. Hence, we project earnings to decline. Further, the weakening macro in EM would negatively impact Acer's cash collection/inventory turnover and would weigh on its FCF.

Advantest (6857.T)

Masahiro Shibano
+81-3-6776-4641
masahiro.shibano@citi.com

We see Buy-rated Advantest as a likely AI beneficiary via rising demand for SOC testers for high-performance GPU and DRAM testers for HBM. In SOC testers, substantial installations through 2022 combined with waning demand for smartphone APU testing in 2023 resulted in surplus capacity in the installed base that had 2023's growing need for GPU and other HPC testing failing to drive any increase in new tester demand. This surplus capacity has now been consumed. Demand for new SOC testers began to ramp up in July-September 2024. Advantest subsequently expanded its supply capacity, and as a result, FY3/25 sales reached ¥400bn (+2.7x YoY). As of May, it appears likely that demand will remain strong through CY25 H1. Company guidance for H2 onward (announced April 25) is conservative, but Advantest appears optimistic. Advantest expects demand for ASIC testers for AI applications to ramp up in earnest in CY25 H2, and we believe this will create a business opportunity for the company. In DRAM testers, the company says concrete inquiries have been materializing since around August 2023. Sales of DRAM testers (including for HBM) topped ¥77bn in FY3/24, and reached ¥15obn in FY3/25. Advantest expects to be the tester of choice for the time being, as both customers and vendors are still going through a trial-and-error process to establish specifications for HBM and lack the engineering resources to evaluate multiple testers in parallel. Tester demand is also seeing a tailwind from increased stack counts, which increase test difficulty.

Applied Material (AMAT.O)

Atif Malik
+1 415-951-1892
atif.malik@citi.com

We are Buy-rated on AMAT as we see it well exposed to secular megatrends like GAA, advanced packaging, and ICAPS in the next 2-3 years. We expect the migration to 3D devices will generate significant changes and opportunities for materials and equipment suppliers, and especially for etch and deposition with increasing steps. We view LRCX/AMAT, with highest exposures to etch and deposition, as best positioned among the US semiconductor equipment group to benefit from the rise in etch/deposition spending intensity from the trend of hybrid bonding/3D stacking, and also the broader megatrend of GAA, backside power and advanced packaging including HBM DRAM.

ASMPT (0522.HK)

Kevin Chen
+852-2501-2125
kevin.y.chen@citi.com

We believe Buy-rated ASMPT is set to benefit from increasing AI-driven demand for its advanced packaging solutions, notably for its TCB (thermo-compression bonder). ASMPT is a leading supplier of TCB with the largest tool install base worldwide (300+ by 2023). The company estimate the global TCB market to grow from US$300mn in 2024 to US$1bn by 2027, at 49% CAGR, driven by AI GPU and HBM. TCB is currently the preferred choice for stacking HBM, given the better cost/performance supporting up to 20H. ASMPT is also developing hybrid bonders, with the next-generation products to come out in 2H25 aiming for opportunities in 2026-27.

We expect robust advanced packaging revenue growth to continue through 2025 and beyond. For logic chips, ASMPT remains the sole supplier of C2S (chip-to-substrate) TCB at the leading foundry, with repeat orders in 2H25. ASMPT has developed TCB with active oxide removal capability, making its TCB highly competitive for C2W (chip-to-wafer) adoption in 2H25 or 2026. On the memory side, TCB is theoretically capable of memory stacking up to 20H at the current 775µm height requirement for HBM4. ASMPT has already achieved 16H stacking using TCB, and believes higher stacking to be possible if JEDEC relaxes the height requirement further.

ASUS (2357.TW)

Carrie Liu
+886-2-8726-9086
Carrie.liu@citi.com

We rate ASUS at Buy. We believe Asus is well-positioned to capture the growth opportunities brought about by rising AI PC penetration, proliferation of edge AI, as well as robust growth in gaming PC/VGA and AI servers. The company's strong balance sheet and solid cash yield would also help cushion on the downside in case of market volatility. ASUS expects AI PCs to reach 25-30% of its shipment share by 2025, with continued ramp into 2026. Management emphasizes the need for local inference capability, particularly for creator-oriented and gaming applications, and has committed to supporting all major chip platforms.

BE Semiconductor Industries (BESI.AS)

Andrew Gardiner
+44-20-7986-4206
andrew.gardiner@citi.com

Neutral-rated Besi is a leading supplier of assembly equipment for the semiconductor industry, specializing in leading-edge assembly processes and equipment for leadframe, substrate and wafer level packaging applications. It generates c80% of revenue from the die attach market, selling both to logic and memory markets as well as IDMs, foundries and OSATs. Besi has a disproportionate exposure to advanced bonding and packaging, the assembly of semiconductor devices with geometries <28nm and placement accuracy <10um. It claims 74% market share of the advanced die placement market.

Besi is the leading supplier of hybrid bonding tools, with its 8800 CHAMEO Ultra Plus tool. Today it has limited competition, although ASM Pacific is beginning to ship a small volume of tools. Besi has supported the initial ramp of chiplets in logic, but more recently has seen its customer base expand from logic and into memory. As discussed above, hybrid bonding was initially expected to be adopted in HBM4 around 2H25/1H26, but this is now more likely to be for HBM4E or HBM5. Besi has acknowledged such a delay, but is still seeing strong interest in its hybrid bonding tools. Besi's management's view is that it is a question of economics and timing, with TCB more viable for longer as a known quantity that is economical; hybrid bonding is the better solution, but is more expensive (primarily due to the need to have hybrid bonders located in the clean room, whereas lesser packaging technologies do not require full front-end chip manufacturing levels of cleanliness) and not yet at the sweet spot of adoption for memory.

While 1Q25 orders in total were below market expectations, Besi did receive multiple customer orders for hybrid bonding systems, across both logic and memory. We expect HBM orders to improve further later this year, but really building in volume during 2026, with memory production ramping in 2027. The further ramp in logic orders and the pending adoption of hybrid bonding for HBM gives us confidence in our 23% CAGR in revenue from 2024-28E, which makes Besi one of the fastest growing stocks in European Technology, if not the industry more globally.

Dell Technologies (DELL.N)

Asiya Merchant
+1 415-951-1752
asiya.merchant@citi.com

We rate Dell shares Buy. Within the overall enterprise hardware space, we believe Dell's broad exposure, size, and scale should enable it to outperform broader IT spending trends. We believe a cyclical downturn tech hardware should see recovery in CY24/FY25 (following a downturn in CY23), and we believe Dell's leading share in commercial PCs as well as across several of its product categories (servers, storage infrastructure) should drive top-line growth at least in line with the core markets they participate, with AI and AI PC offerings to further Dell's moat in these core markets. Additionally, we expect FCF margins to return to normalized levels at 100% FCF/NI conversions, which should aid in increasing its level of shareholder returns.

Disco (6146.T)

Masahiro Shibano
+81-3-6776-4641
masahiro.shibano@citi.com

We see Buy-rated Disco as positioned to benefit from AI via the TSV process for HBM and the CoWoS process for logic, with equipment shipment value expanding to ¥60n in 2024. On memory, TSV is used in DRAM stacking for HBM, and we expect Disco systems to be used in edge trimming to minimize wafer edge chipping and dicing that peels off the support substrate and also offer greater cleanliness in response to wafer thinning in the front-end process. With the migration from HBM3 to HBM4 increasing the number of layers in DRAM chips, we expect increased demand for grinders for DRAM die thinning. The use of hybrid bonding will also require greater cleanliness (= higher added value) in the stacking process. In logic, we expect increased business opportunities via the use of laser grooving, which reduces thermal damage when processing heat-sensitive low-K films and increases the number of grinding passes in backside PDN. Although Disco is a back-end manufacturer, we expect its systems to take on increasing added value (as cleanliness requirements become more stringent), which bodes well for margins. As of May 2025 Disco explained that HBM-related inquiries have eased compared with a year ago and AI-package related investment is advancing in line with plan.

Hitachi (6501.T)

Kota Ezawa
+81-3-6776-4640
kota.ezawa@citi.com

With industrial electronic equipment expected to proactively embrace edge AI we think the industrial device and foundry automation businesses at Hitachi and MELCO will adopt edge AI with the aim of growing (not so much immediately but over the longer term). We look for industry reorganization among SMEs that produce and sell machinery and among companies with specific machinery technology, and believe large companies that can provide solutions that incorporate edge AI will see growth. We rate Hitachi shares as Buy.

Ibiden (4062.T)

Takayuki Naito
+81-3-6776-4644
takayuki.naito@citi.com

We have a Buy rating on Ibiden shares. The company is the top supplier of ABF substrates for AI GPUs. In AI GPUs Ibiden has high shares in ABF substrates for Nvidia and AMD. The Nvidia share is particularly high, at over 90%. We believe operating profit is on track for renewed growth with bottoming out in demand from a major CPU customer (a recent drag on earnings) and continued expansion in demand for AI servers. Ibiden intends to secure capacity for AI GPUs with the operation of a newly completed plant and conversion of existing facilities. In a change of stance, it says it will not implement additional large-scale capex. With FY24 results Ibiden disclosed a medium-term outlook targeting FY27 sales of ¥600bn and an OPM of 15.0%, followed by ¥750bn and 20.0% in FY30. FY24 actuals were ¥369.4bn in sales and a 12.9% OPM. We have a strong recommendation on this Japan electronic components name as a major beneficiary of growth in AI demand.

KLA Corp. (KLAC.O)

Atif Malik
+1 415-951-1892
atif.malik@citi.com

We are Buy-rated on KLAC as a good play on growing logic WFE spend mix in 2025E/26E. Citi models 2025E/26E WFE -5%/+2% or $94Bn/$96Bn on tariffs and global uncertainty. We expect the investments in leading edge AI, HBM and NAND upgrade to be more than offset by weak mature logic industrial/auto demand, weaker consumer demand, and lower China spending. We see the semiconductor equipment group stuck in Phase 1 of our stock-picking framework where we are Buy-rated on the whole group. We see KLAC continue to benefit the most from 2nm GAA and advanced packaging secular growth drivers.

Lam Research (LRCX.O)

Atif Malik
+1 415-951-1892
atif.malik@citi.com

We are Buy-rated on LRCX as a good stock play on 3D devices megatrend in 2025/26. We expect the migration to 3D devices will generate significant changes and opportunities for materials and equipment suppliers, and especially for etch and deposition with increasing steps. With the upgrade opportunities and technology inflections across NAND (especially on the upgrades), DRAM and Foundry/logic, Lam expects SAM to grow from low 30% WFE to high 30%, and targets to win >50% of the incremental SAM. Despite some headwind in Reliant in 2025, CSBG is expected to grow at LDD CAGR in longer term and higher gross margin from the Asia manufacturing footprint also drive EPS growth.

Lenovo (0992.HK)

Carrie Liu
+886-2-8726-9086
Carrie.liu@citi.com

Lenovo, the world's largest PC brand by volume, takes an aggressive view on AI PC outlook. It targets 25% of its PC shipments in 2025E to be AI-enabled, growing to 80% by 2027E. Lenovo's product strategy spans commercial and consumer lines, and the company benefits from vertical integration and early Copliot + PC deployment in China and Europe market. We rate Lenovo at Buy as we expect ISG to turn around at the PTI level in FY26E/27E; this, coupled with the company's strong performance in PC segments, including commercial and gaming, as well as the smartphone business, should help drive sales and earnings growth.

MediaTek (2454.TW)

Laura Chen
+886-2-8726-9090
laura.cy.chen@citi.com

In the AI ASIC chip supply chain, MediaTek's view and growth outlook on AI ASIC was the focus during its earnings call. The management reiterated its positive view on the AI accelerator long-term upside. Aside from Mediatek's advanced node design capability and Serdes capability, the company has also engaged with Nvidia to leverage its NVLink IP for future AI chip development. MediaTek indicates that its AI ASIC project would be able to contribute in late 2026E at the worst case, which we view as conservative. To note that the company previously guided about US$1bn revenue in 2026E and we see more upside into 2027E.

In the edge AI, We believe Buy-rated MediaTek is on the right track gaining market share in the high-end flagship smartphone as we note that leading Android bands including Xiaomi, OPPO, and Vivo are increasing their flagship/high-end models with MediaTek's D9400 series product. We believe there is good chance that MediaTek can get 20%+ market share in the Android flagship smartphone market thanks to its comprehensive AI functions and NPU support. MediaTek has also had more affordable AI solutions in D8000 series; we believe the increasing AI function will allow MediaTek to maintain its ASP thanks to the increasing silicon content.

Micro-Star (MSI) (2377.TW)

Carrie Liu
+886-2-8726-9086
Carrie.liu@citi.com

We rate MSI shares a Buy. MSI is leveraging its strong presence in the gaming and creator segments to introduce AI PC models with integrated AI engines, targeting low-latency content creation. Apart from publishers, competition organizers, platforms and players, hardware vendors also play a key role in the ecosystem and would therefore benefit. We believe MSI is well positioned to capture the growth opportunities, backed by its strong design innovation, fast time-to-market, and a comprehensive product portfolio that caters to both hardcore and casual gamers, as well as professional consumers.

Micron Technology Inc (MU.O)

Christopher Danely
+1-415-951-1885
christopher.danely@citi.com

Micron's HBM3E design utilizes its 1-beta process technology and advanced through-silicon via (TSV). In F2Q25, Micron's HBM sales increased by more than 50% QoQ to over $1.0 billion. Micron is currently shipping HBM3E, and we expect $8.3 billion (up 690% YoY) in HBM sales in C25E and $13.2 billion (up 59% YoY) in C26E. Micron's 12H HBM3E is designed into the GB300, and Micron expects that product will make up the majority of its HBM shipments in 2H25E. We believe HBM has a 6X pricing increase versus normal DRAM with higher margins as well. We have a Buy rating on MU due to the DRAM market upturn and its AI exposure.

Mitsubishi Electric (6503.T)

Kota Ezawa
+81-3-6776-4640
kota.ezawa@citi.com

With industrial electronic equipment expected to proactively embrace edge AI, we think the industrial device and foundry automation businesses at Hitachi and MELCO will adopt edge AI with the aim of growing (not so much immediately but over the longer term). We look for industry reorganization among SMEs that produce and sell machinery and among companies with specific machinery technology, and believe large companies that can provide solutions that incorporate edge AI will see growth.

We rate MELCO shares as Buy and believe its factory automation (FA) business will post growth above industry cyclicality. We forecast FA business sales will grow 12% CAGR in the next three years with OP growing 40%. FA accounted for 12% consolidated OP in FY24 and we expect the weighting to be 23% in FY27. MELCO has had an edge AI-equipped FA product line-up since 2020 and aims to appeal based on its capacity in deep learning and big data analysis. Deep learning initializes algorithms and big-data analytics enhance historical data exploitation, enabling cameras, sensors, and satellites in space to have better identification and recognition capabilities, production machinery itself to detect root causes of problems, and to generally implement predictive diagnosis. MELCO focuses on customer criteria such as plants, offices, and residential buildings as well as autonomous driving. Based on the higher-level business collaboration concept "Serendie" announced in 2024, multiple business groups integrate big data and connect systems to drive better edge AI capability. The major focus for the moment is the FA business working together with the electric power equipment division and the elevator business working together with air conditioner division.

NVIDIA (NVDA.O)

Atif Malik
+1 415-951-1892
atif.malik@citi.com

The AI Infrastructure Corner Stone — We continue to expect NVDA representing ~95% of the AI GPUs market in C25 and C26 backed by an accelerated roadmap and a technology leadership going from the chip level to the software level. In terms of units, we are modeling 6M (+74%)/6.4M (+20%) in C25E/C26E. we expect Blackwell chips to represent respectively 66%/51% in C25E/C26E of units as Blackwell can deliver up to 4x faster training and 30x faster inference than the Hopper at a system level. Boosted by rising units and weighted average ASPs of $28K/$34K in C25E and C26E, we expect NVDA AI GPU sales to grow 60% and 29% to $168B/$217B in C25E/C26E. Maintain Buy.

Panasonic (6752.T)

Kota Ezawa
+81-3-6776-4640
kota.ezawa@citi.com

We rate Panasonic shares as Buy. We believe Panasonic will conduct corporate-level restructuring in 2025 with multiple businesses divestitures and spin-offs. In our view, the printed circuit board material Megtron is one Panasonic's best products and believe the related businesses category will post 25% sales growth CAGR, driving the company' s next growth story. We believe divestiture or spin-offs would create better value recognition in the market if Panasonic decides to do this as part of ongoing restructuring. Until the final decisions are made, we believe Megtron will see growth in sales for AI chip package substrates and server motherboards thanks to its dominant market share driven by top-level performance. Currently the business is at an early stage in its growth phase as next-generation AI chips penetrate the AI server space. We believe growth will be led by high-end AI-related products as well as ordinary computing and mobile products. In the AI server area Panasonic provides hybrid capacitors for AI server motherboards and Li-Ion-based ESS for AI data centers, which we think will also be recognized as growth drivers.

Samsung Electronics (005930.KS)

Peter Lee
+82 2 3705 0720
peter.sc.lee@citi.com

Despite initial delay in HBM3e solution, Samsung expects the recovery of conventional memory market, and its overhaul of front-end process is well progressing. The company is preparing next-gen memory products such as SOCAMM/LPDDR6/LPW DRAM to address edge AI demand. Further, Samsung expects DRAM S/D to improve throughout 2025E given limited supply growth falling short of overall market demand. Samsung expects it would fully digest its accumulated inventory to meet this year's demand and aims to develop and mass produce next- gen products such as LPW DRAM by 2028E. We rate the shares as Buy.

SK Hynix (000660.KS)

Peter Lee
+82 2 3705 0720
peter.sc.lee@citi.com

We expect SK Hynix to benefit from increasing on-device AI demand, with its LPDDR5T demand expected to increase in 2H25E according to customer's smartphone launch. Hynix anticipates robust AI DRAM demand growth to continue throughout 2026E/2027E driven by Nvidia's Rubin launch, introduction of HBM4, and increasing demand for next-gen on device AI DRAM such as LLW DRAM. Looking toward the future, the company expects strong memory content growth, driven by increasing number of HBM layers, die density increasing to 32Gb from 24Gb, and strong HBM demand from ASIC customers. With ASIC HBM revenue contribution expected to exceed 50% in 2027E from the current 30%, SK Hynix believes custom HBM solutions to gain momentum from late 2027E or 2028E. We rate the shares as Buy.

Sony Group (6758.T)

Kota Ezawa
+81-3-6776-4640
kota.ezawa@citi.com

We rate Sony shares as Neutral but believe its semiconductor technology has considerable potential in hybrid bonding. Sony has developed proprietary Cu-Cu bonding technology used in CMOS image sensors (CIS) and appears to have begun mass production in 2023. CIS seems to require relatively long steps as part of its back-end processes, mainly in chip bonding and micro lens processing. At the moment Sony does not outsource those processes (even to close partner TSMC) and maintains all CIS back-end production in-house. Although Sony does not have an official plan to deploy Cu-Cu technology to products other than CIS or develop it externally, we see a chance of including edge AI logic cores in the CIS graphic processing logic chip.

SUMCO (3436.T)

Yuta Nishiyama
+81-3-6776-4643
yuta.nishiyama@citi.com

Shin-Etsu Chemical estimates AI-related applications (HBM and GPUs) account for only c10% of the wafer market on an area basis, but the company also thinks that as the adoption of on-device AI increases the weighting of higher priced AI-related wafers will rise as well. Wafer consumption volumes can also be expected to grow thanks to an increase in interposer usage and HBM layering. SUMCO's share price has sagged due to the stagnation of OP because of an increase in wafer inventories and growth in depreciation, and this appears to have priced in the deterioration in fundamentals to some degree. We think the focus could shift to potential upside if the weighting of AI-related wafers starts to grow. We rate SUMCO shares as Buy.

Techwing (089030.KQ)

Josh Yang
+82-2-3705-0723
josh.yang@citi.com

We expect wider AI adoption by edge devices is likely to increase testing intensity for memory chips, and thus benefit Buy / High Risk rated Techwing, the only supplier of HBM-specialized handler. We foresee edge devices will adopt specialized memory chips, such as LLW, to power AI functions on device. Given increased value of each memory chip, we expect memory makers will increase testing intensity to ensure full functionality of each chip. As the global leader in memory handler with >70% global m/s and the only supplier of HBM-specialized handler, which we think can also be used to support testing of mini-HBMs that will be used on edge devices, Techwing is projected to benefit from increased testing intensity. Of note, Techwing is likely to begin supply of Cube Prober to all three HBM makers from 3Q25E after passing Samsung's qual in 4Q24, Hynix's qual in 2Q25, and Micron's qual in 3Q25, which should position Techwing well for future growth in AI-specialized memory chip testing demand.

Tokyo Ohka (4186.T)

Yuta Nishiyama
+81-3-6776-4643
yuta.nishiyama@citi.com

We think that as the adoption of on-device AI increases there will be growth in volumes for AI-related semiconductors, and this will have broad benefits for semiconductors materials, with the resists supplied by Tokyo Ohka being a prime example. Buy-rated Tokyo Ohka also supplies materials for rear-end processes such as RDL resists for chip on wafer substrate (CoWoS) and micro-bump resists for HBM. We estimate that sales of HBM micro-bump resists grew by around 2.5x between 2022 and 2023, and then again roughly doubled between 2023 and 2024. January-March OP of ¥9.8bn was higher than expected, but we expect ongoing modest growth in resist demand from April-June as semiconductor manufacturers' capacity utilization improves. We therefore think the OP consensus for FY12/25 will likely rise from ¥38b to around ¥40bn.

TSMC (2330.TW)

Laura Chen
+886-2-8726-9090
laura.cy.chen@citi.com

We rate TSMC shares as Buy. We believe TSMC would continue outgrowing the industry thanks to its solid industry position particularly in advanced nodes and promising AI demand. TSMC expects AI GPU/accelerator and CPU revenue contribution to grow more than double YoY in 2025 and its long term mid-40% CAGR during 2024-2029E intact. Staring from 2026, we believe AI chip's transition from N5/4 to N3 in AI GPUs/accelerators brings some architecture changes; one of the most importance is the separation of compute die (top die) and SerDes I/O die. It will be implemented using chiplet and advanced packaging like CoWoS. That is mainly due to the different technology optimization requirement.

For compute logic, it would benefit from cutting-edge node like 3nm for higher performance and energy efficiency. At 3nm, higher transistor density allows more cores, which pushes the compute portion close to or even beyond reticle limits if I/O is still integrated on the same die. Nevertheless, for Serdes I/O circuit, it won't be scaled well with smaller nodes as they are mix-signal heavy. Separating the dies allows each block to be fabricated on the best-fit process node with better performance and avoids overheat issue.

Even with macro uncertainties ahead, we note that TSMC has been outperforming global semiconductor market and global GDP growth in the past decades even during the downcycle (see our early-April report TSMC (2330.TW): Not Immune from Global Trade Tensions; Maintain Buy, but Lower Estimates). We also believe management's firm and prudent comment on news of JV or technology transfer clears the air of the Intel engagement and will relieve some of the inventor's concern on TSMC's long-term competitive advantage.

Valuation comps summary

Figure 53. Valuation Comps Summary (pricing as of 27-May-2025 closes)

2025-05-27							2025E							2026E						
(LCYbn)	Ticker	Mkt Cap (LCYbn)	Rating	Price (LCY)	TP (LCY)	ETR (%)	Sales	OP	OPM (%)	P/E (x)	P/B (x)	ROE (%)	Div Yld (%)	Sales	OP	OPM (%)	P/E (x)	P/B (x)	ROE (%)	Div Yld (%)
Acer	2353.TW	106	Sell	35	27	-18.9	289	6	2.0	17.0	1.4	8.2	4.8	307	6	2.0	16.4	1.4	8.5	0.0
Advantest	6857.T	5,239	Buy	7.1k	11.0k	+49.4	756	232	30.7	27.3	8.1	33.6	0.6	812	258	31.8	25.0	6.3	28.4	0.6
Alphabet	GOOGL.O	1,154	Buy	173	200	+16.2	380	129	33.9	18.0	5.0	31.3	0.0	415	147	35.4	17.1	3.9	25.5	0.0
Apple	AAPL.O	2,990	Buy	200	240	+20.4	406	127	31.2	27.6	47.3	179.3	0.5	433	135	31.2	25.0	38.4	167.2	0.6
Applied Materials	AMAT.O	130	Buy	162	190	+18.4	28.4	8.5	29.9	17.1	5.6	32.5	1.1	29.3	8.7	29.7	16.5	4.6	30.4	1.1
ASMPT	0522.HK	22	Buy	52	65	+27.5	14.4	1.1	7.6	25.5	1.4	5.5	3.4	16.9	2.0	11.6	13.6	1.3	10.1	6.5
ASUS	2357.TW	452	Buy	608	730	+23.4	656	39	5.9	11.2	1.6	14.8	5.8	711	38	5.4	11.2	1.5	14.1	6.1
BESI	BESI.AS	9	Neutral	109	105	-3.6	0.6	0.2	31.4	54.4	19.2	33.1	1.7	0.9	0.3	39.2	29.4	15.0	57.1	3.0
DELL	DELL.N	78	Buy	114	128	+14.2	96	6	6.5	13.0	396.1	nm	1.8	101	7	7.0	11.2	28.4	378.0	2.0
Disco	6146.T	3,670	Buy	33.9k	43.0k	+26.7	393	167	42.4	28.6	6.7	25.3	1.2	420	174	41.4	26.9	5.7	22.9	1.3
Eugene Tech	084370.KQ	733	Buy	32.0k	47.0k	+45.5	401	63	15.7	12.9	1.6	13.5	0.3	413	80	19.3	10.7	1.4	14.1	0.3
Hanmi Semi	042700.KS	7,951	Buy	82.3k	133.0k	+57.3	734	353	48.0	26.2	10.1	46.3	0.8	843	411	48.8	23.2	7.5	37.0	0.8
Hitachi	6501.T	17,991	Buy	3.9k	4.7k	+20.5	9,662	967	10.0	23.5	2.9	12.6	1.3	10,447	1,141	10.9	19.9	2.7	13.9	1.5
Ibiden	4062.T	777	Buy	5.6k	6.0k	+4.7	366	43	11.8	23.5	1.6	6.7	0.7	417	49	11.8	18.3	1.5	8.3	0.7
Intel	INTC.O	90	Neutral	21	21	+2.2	48.7	-0.5	-0.9	-94.0	0.8	-0.4	0.0	50.0	4.9	9.7	19.8	0.7	3.8	0.4
KLA	KLAC.N	105	Buy	789	835	+6.8	12.1	5.0	41.5	25.1	19.5	90.3	0.9	12.4	4.9	39.7	23.4	11.9	61.4	1.0
Lam Research	LRCX.O	107	Buy	84	87	+4.7	18.1	5.7	31.8	22.3	9.1	45.3	0.0	17.7	5.3	29.7	21.6	7.1	37.0	0.6
Lenovo	0992.HK	116	Buy	9	13	+42.6	69.1	2.2	3.5	10.9	2.7	26.2	4.1	74.0	3.2	3.5	10.1	2.4	25.1	4.1
MediaTek	2454.TW	2,050	Buy	1.3k	1.6k	+28.7	617	113	18.4	18.3	4.9	27.2	4.4	791	115	14.6	17.6	4.6	26.5	4.6
Meta	META.O	1,615	Buy	642	690	+7.8	187	74	39.7	20.2	6.9	31.8	0.0	211	88	41.6	18.3	5.2	27.7	0.0
Micron	MU.O	108	Buy	96	110	+14.6	35	8	22.2	14.0	1.8	14.6	0.5	40	13	31.9	9.2	1.5	18.2	0.5
Mitsubishi Elec.	6503.T	6,361	Buy	3.1k	3.5k	+16.1	5,475	401	7.3	17.9	1.5	8.7	1.9	5,661	456	8.1	15.1	1.5	9.7	2.1
MSI	2377.TW	116	Buy	137	180	+33.6	226	10	4.4	13.3	1.9	15.2	4.7	247	12	4.9	11.3	1.7	15.9	5.5
Nextin	348210.KQ	556	Buy	53.1k	82.1k	+50.4	144	58	40.6	10.4	2.5	27.5	0.0	165	74	44.7	8.2	1.9	26.4	0.0
Nvidia	NVDA.O	3,305	Buy	136	150	+10.7	130	81	62.4	31.0	23.8	90.8	0.1	206	120	58.3	25.1	14.5	68.8	0.1
Panasonic	6752.T	3,685	Buy	1.6k	2.3k	+47.1	8,320	375	4.5	11.1	0.7	6.7	2.8	7,736	420	5.4	9.3	0.7	7.6	3.1
Qualcomm	QCOM.O	162	Neutral	149	145	-0.2	43	12	28.0	15.9	5.8	37.1	2.4	41	11	27.2	16.0	5.4	33.6	2.4
Samsung Elec.	005930.KS	319,068	Buy	53.9k	83.0k	+52.6	354,774	30,313	8.5	14.8	0.9	6.4	2.6	385,908	31,977	8.3	16.0	0.9	5.7	2.6
Shin-Etsu Chemical	4063.T	8,611	Buy	4.6k	5.1k	+13.2	2,561	742	29.0	16.7	2.0	12.0	2.5	2,575	750	29.1	16.4	2.0	12.2	2.5
SK Hynix	000660.KS	147,420	Buy	202.5k	350.0k	+67.6	88,938	38,787	43.6	5.0	1.5	34.9	1.2	101,434	42,600	42.0	4.7	1.1	27.8	1.4
Sony Group	6758.T	22,894	Neutral	3.8k	4.1k	+6.8	13,326	1,389	10.4	20.0	2.6	14.3	0.8	13,140	1,508	11.5	19.2	2.4	13.6	0.6

SUMCO	3436.T	335	Buy	1.0k	1.2k	+27.0	391	1	0.3	-79.5	0.6	-0.7	2.1	417	-4	-1.0	-26.2	0.6	-2.2	2.1
Techwing	089030.KQ	1,277	Buy	34.2k	90.0k	+157.4	556	191	34.3	11.9	4.5	45.3	0.4	847	313	37.0	5.6	2.5	57.2	0.4
Teradyne	TER.O	13	Buy	82	94	+15.6	3	1	21.0	23.8	5.5	18.2	0.0	4	1	25.3	16.6	4.3	24.4	0.0
Tokyo Electron (TEL)	8035.T	10,626	Buy	23.2k	35.0k	+51.8	2,418	695	28.7	20.0	5.1	27.1	2.5	2,470	710	28.7	17.3	4.4	27.2	2.9
Tokyo Ohka	4186.T	425	Buy	3.5k	4.4k	+24.5	226	38	16.8	18.0	2.0	11.5	1.9	250	41	16.4	16.8	1.8	11.4	2.1
TSMC	2330.TW	25,025	Buy	1.0k	1.1k	+13.3	3,654	1,708	46.7	16.8	4.7	31.2	1.4	4,045	1,763	43.6	15.7	3.8	26.8	1.4
Unimicron	3037.TW	169	Buy	111	100	-9.4	127	7	5.8	24.9	1.7	7.2	1.6	132	9	7.0	19.2	1.6	8.8	2.1
Xiaomi	1810.HK	1,338	Buy	52	74	+42.6	501	43	8.6	25.9	4.3	20.5	0.0	591	43	8.6	20.3	3.5	19.9	0.0

Source: dataCentral, Citi Research

If you are visually impaired and would like to speak to a Citi representative regarding the details of the graphics in this document, please call USA 1-888-800-5008 (TTY: 711), from outside the US +1-210-677-3788

Appendix A-1

ANALYST CERTIFICATION

The research analysts primarily responsible for the preparation and content of this research report are either (i) designated by "AC" in the author block or (ii) listed in bold alongside content which is attributable to that analyst. If multiple AC analysts are designated in the author block, each analyst is certifying with respect to the entire research report other than (a) content attributable to another AC certifying analyst listed in bold alongside the content and (b) views expressed solely with respect to a specific issuer which are attributable to another AC certifying analyst identified in the price charts or rating history tables for that issuer shown below. Each of these analysts certify, with respect to the sections of the report for which they are responsible: (1) that the views expressed therein accurately reflect their personal views about each issuer and security referenced and were prepared in an independent manner, including with respect to Citigroup Global Markets Inc. and its affiliates; and (2) no part of the research analyst's compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views expressed by that research analyst in this report.

IMPORTANT DISCLOSURES

TechWing (089030.KQ)

Ratings and Target Price History
Fundamental Research

Analyst: Josh Yang

#	Date	Rating	Target Price	Closing Price
1	11-Mar-24 17:00:45	*1H	*41,000.00	23,900.00
2	14-Apr-24 18:50:06	1H	*49,000.00	38,600.00
3	14-Jul-24 16:00:12	1H	*90,000.00	65,700.00

*Indicates Change

Rating/target price changes above reflect Eastern Time

SK Hynix (000660.KS)

Ratings and Target Price History
Fundamental Research

Analyst: Peter Lee

#	Date	Rating	Target Price	Closing Price
1	27-May-22 07:23:49	1	*170,000.00	106,000.00
2	21-Jun-22 02:35:43	1	*163,000.00	95,200.00
3	27-Jul-22 09:23:09	1	*158,000.00	100,000.00
4	27-Sep-22 13:21:16	1	*143,000.00	82,000.00
5	11-Oct-22 12:50:34	1	*135,000.00	90,200.00
6	27-Oct-22 11:20:33	1	*130,000.00	90,000.00
7	18-Nov-22 05:26:56	1	*123,000.00	88,400.00
8	30-May-23 08:36:29	1	*160,000.00	110,300.00
9	30-Jun-23 05:08:18	1	*170,000.00	115,200.00
10	30-Aug-23 11:34:39	1	*180,000.00	119,400.00
11	04-Oct-23 06:14:16	1	*185,000.00	115,400.00
12	13-Nov-23 09:52:50	1	*190,000.00	131,800.00
13	01-Jan-24 18:20:58	1	*230,000.00	141,500.00
14	18-Mar-24 10:47:11	1	*234,000.00	164,300.00
15	01-Apr-24 10:34:31	1	*238,000.00	185,500.00
16	12-Apr-24 08:17:40	1	*310,000.00	187,400.00
17	25-Jun-24 06:27:18	1	*350,000.00	225,000.00
18	25-Jul-24 03:44:27	1	*337,000.00	190,000.00
19	09-Sep-24 02:31:32	1	*310,000.00	157,000.00
20	24-Oct-24 06:29:30	1	*330,000.00	198,200.00
21	12-Nov-24 10:31:16	1	*350,000.00	185,800.00
22	31-Dec-24 06:45:21	1	*340,000.00	173,900.00
23	03-Mar-25 15:00:11	1	*350,000.00	190,200.00

*Indicates Change

Rating/target price changes above reflect Eastern Time

Samsung Electronics (005930.KS)

Ratings and Target Price History
Fundamental Research

Analyst: Peter Lee

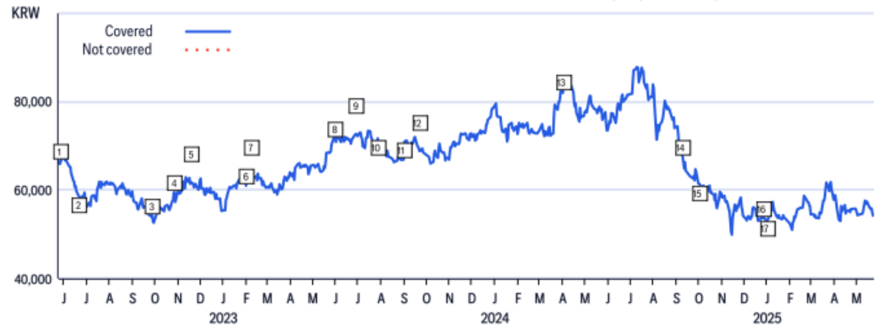

#	Date	Rating	Target Price	Closing Price
1	27-May-22 10:28:13	1	*110,000.00	66,500.00
2	21-Jun-22 02:34:29	1	*105,000.00	58,500.00
3	27-Sep-22 13:24:35	1	*95,000.00	54,200.00
4	27-Oct-22 10:52:58	1	*100,000.00	59,500.00
5	18-Nov-22 04:29:36	1	*93,000.00	61,800.00
6	31-Jan-23 06:05:23	1	*90,000.00	61,000.00
7	08-Feb-23 03:46:21	1	*87,000.00	63,100.00
8	31-May-23 03:36:58	1	*100,000.00	71,400.00
9	28-Jun-23 08:19:44	1	*105,000.00	72,700.00
10	27-Jul-23 05:53:47	1	*110,000.00	71,700.00
11	31-Aug-23 06:54:46	1	*120,000.00	66,900.00
12	22-Sep-23 09:53:22	1	*110,000.00	68,800.00
13	01-Apr-24 05:23:20	1	*120,000.00	82,000.00
14	09-Sep-24 02:31:35	1	*110,000.00	67,500.00
15	02-Oct-24 05:16:51	1	*97,000.00	61,300.00
16	26-Dec-24 07:42:50	1	*87,000.00	53,600.00
17	31-Dec-24 06:21:22	1	*83,000.00	53,200.00

*Indicates Change

Rating/target price changes above reflect Eastern Time

Samsung Electronics (005930.KS)

Short-Term View/Catalyst Watch Research

Analyst: Peter Lee

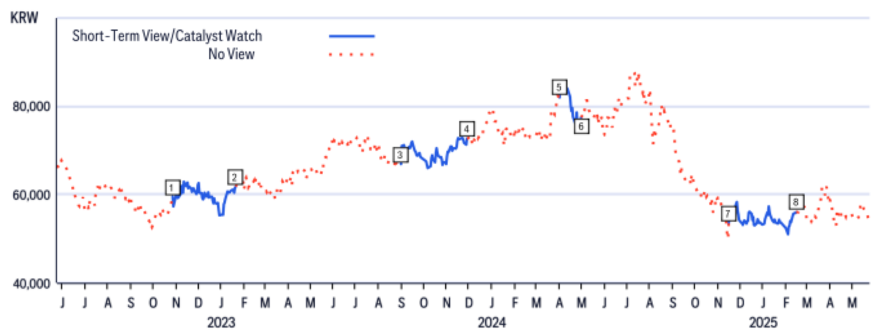

#	Date	Action	Expected Direction	Duration	Closing Price
1	27-Oct-22 06:52:58	Add CW	Upside	90 Days	59,500.00
2	24-Jan-23 21:49:55	Remove CW	Upside	90 Days	61,800.00
3	31-Aug-23 02:54:46	Add CW	Upside	90 Days	66,900.00
4	28-Nov-23 21:19:48	Remove CW	Upside	90 Days	72,700.00
5	01-Apr-24 01:23:20	Add CW	Upside	30 Days	82,000.00
6	30-Apr-24 22:52:33	Remove CW	Upside	30 Days	77,500.00
7	17-Nov-24 11:09:02	Add CW	Upside	90 Days	53,500.00
8	14-Feb-25 12:17:17	Remove CW	Upside	90 Days	56,000.00

CW - Catalyst Watch , STV - Short-Term View

Rating/target price changes above reflect Eastern Time

SK Hynix (000660.KS)

Short-Term View/Catalyst Watch Research

Analyst: Peter Lee

#	Date	Action	Expected Direction	Duration	Closing Price
1	27-Oct-22 07:20:33	Add CW	Downside	90 Days	90,000.00
2	25-Jan-23 21:30:26	Remove CW	Downside	90 Days	91,400.00
3	30-May-23 04:36:29	Add CW	Upside	90 Days	110,300.00
4	27-Aug-23 22:58:59	Remove CW	Upside	90 Days	116,500.00
5	04-Oct-23 02:14:16	Add CW	Upside	90 Days	115,400.00
6	01-Jan-24 21:19:41	Remove CW	Upside	90 Days	141,500.00
7	03-Mar-25 10:00:11	Add CW	Upside	90 Days	190,200.00

CW - Catalyst Watch , STV - Short-Term View

Rating/target price changes above reflect Eastern Time

The Firm has made a market in the publicly traded equity securities of Lenovo Group Ltd on at least one occasion since 1 Jan 2024.

Citi is advisor to Silver Lake Technology Management LLC on a pending M&A transaction involving Intel Corporation.

The Firm has made a market in the publicly traded equity securities of ASMPT Ltd on at least one occasion since 1 Jan 2024.

The Firm has made a market in the publicly traded equity securities of Xiaomi Corp on at least one occasion since 1 Jan 2024.

Citigroup Global Markets Inc. or its affiliates has a net long position of 0.5% or more of any class of common equity securities of Teradyne Inc.

Within the past 12 months, Citigroup Global Markets Inc. or its affiliates has acted as manager or co-manager of an offering of securities of Alphabet Inc,Apple, Inc.,Applied Materials Inc,Meta Platforms Inc,Micron Technology Inc,Qualcomm Inc.

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Intel Corp,KLA Corp,Lam Research Corp,Meta Platforms Inc,Micron Technology Inc,NVIDIA Corp,Qualcomm Inc,Samsung Electronics,Sony Group,TSMC.

Citigroup Global Markets Inc. or its affiliates has received compensation for investment banking services provided within the past 12 months from ASUSTeK Computer,Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Hitachi,Intel Corp,KLA Corp,Lam Research Corp,Lenovo Group,Meta Platforms Inc,Micron Technology Inc,NVIDIA Corp,Qualcomm Inc,SK Hynix,Sony Group,Xiaomi.

Citigroup Global Markets Inc. or its affiliates expects to receive or intends to seek, within the next three months, compensation for investment banking services from ASUSTeK Computer.

Citigroup Global Markets Inc. or its affiliates received compensation for products and services other than investment banking services from ASMPT,ASUSTeK Computer,Acer,Advantest,Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Disco,Hitachi,Ibiden,Intel Corp,KLA Corp,Lam Research Corp,Lenovo Group,Mediatek,Meta Platforms Inc,Micro-Star,Micron Technology Inc,Mitsubishi Electric,NVIDIA Corp,Nextin,Panasonic Holdings,Qualcomm Inc,SK Hynix,Samsung Electronics,Shin-Etsu Chemical,Sony Group,Sumco,TSMC,TechWing,Teradyne Inc,Tokyo Electron,Xiaomi in the past 12 months.

Citigroup Global Markets Inc. or its affiliates currently has, or had within the past 12 months, the following as investment banking client(s): ASUSTeK Computer,Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Hitachi,Intel Corp,KLA Corp,Lam Research Corp,Lenovo Group,Meta Platforms Inc,Micron Technology Inc,NVIDIA Corp,Qualcomm Inc,SK Hynix,Sony Group,Xiaomi.

Citigroup Global Markets Inc. or its affiliates currently has, or had within the past 12 months, the following as clients, and the services provided were non-investment-banking, securities-related: ASMPT,ASUSTeK Computer,Acer,Advantest,Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Hitachi,Intel Corp,KLA Corp,Lam Research Corp,Lenovo Group,Mediatek,Meta Platforms Inc,Micron Technology Inc,Mitsubishi Electric,NVIDIA Corp,Panasonic Holdings,Qualcomm Inc,SK Hynix,Samsung Electronics,Shin-Etsu Chemical,Sony Group,TSMC,Tokyo Electron,Xiaomi.

Citigroup Global Markets Inc. or its affiliates currently has, or had within the past 12 months, the following as clients, and the services provided were non-investment-banking, non-securities-related: ASMPT,ASUSTeK Computer,Acer,Advantest,Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Disco,Hitachi,Ibiden,Intel Corp,KLA Corp,Lam Research Corp,Lenovo Group,Mediatek,Meta Platforms Inc,Micro-Star,Micron Technology Inc,Mitsubishi Electric,NVIDIA Corp,Nextin,Panasonic Holdings,Qualcomm Inc,SK Hynix,Samsung Electronics,Shin-Etsu Chemical,Sony Group,Sumco,TSMC,TechWing,Teradyne Inc,Tokyo Electron,Xiaomi.

Citigroup Global Markets Inc. and/or its affiliates has a significant financial interest in relation to ASUSTeK Computer,Alphabet Inc,Apple, Inc.,Applied Materials Inc,Dell Technologies,Hitachi,Intel Corp,KLA Corp,Lam Research Corp,Lenovo Group,Meta Platforms Inc,Micron Technology Inc,NVIDIA Corp,Qualcomm Inc,SK Hynix,Samsung Electronics,Sony Group,TSMC,Tokyo Electron,Xiaomi. (For an explanation of the determination of significant financial interest, please refer to the policy for managing conflicts of interest which can be found at www.citiVelocity.com.)

Analysts' compensation is determined by Citi Research management and Citigroup's senior management and is based upon activities and services intended to benefit the investor clients of Citigroup Global Markets Inc. and its affiliates (the "Firm"). Compensation is not linked to specific transactions or recommendations. Like all Firm employees, analysts receive compensation that is impacted by overall Firm profitability which includes investment banking, sales and trading, and principal trading revenues. One factor in equity research analyst compensation is arranging corporate access events between institutional clients and the management teams of covered companies. Typically, company management is more likely to participate when the analyst has a positive view of the company.

For financial instruments recommended in the Product in which the Firm is not a market maker, the Firm is a liquidity provider in such financial instruments (and any underlying instruments) and may act as principal in connection with transactions in such instruments. The Firm is a regular issuer of traded financial instruments linked to securities that may have been recommended in the Product. The Firm regularly trades in the securities of the issuer(s) discussed in the Product. The Firm may engage in securities transactions in a manner inconsistent with the Product and, with respect to securities covered by the Product, will buy or sell from customers on a principal basis.

The Firm is a market maker in the publicly traded equity securities of ASMPT,Lenovo Group,Xiaomi.

Unless stated otherwise neither the Research Analyst nor any member of their team has viewed the material operations of the Companies for which an investment view has been provided within the past 12 months.

For important disclosures (including copies of historical disclosures) regarding the companies that are the subject of this Citi Research product ("the Product"), please contact Citi Research, 388 Greenwich Street, 6th Floor, New York, NY, 10013, Attention: Legal/Compliance [E6WYB6412478]. In addition, the same important disclosures, with the exception of the Valuation and Risk assessments and historical disclosures, are contained on the Firm's disclosure website at https://www.citivelocity.com/cvr/eppublic/citi_research_disclosures. Valuation and Risk assessments can be found in the text of the most recent research note/report regarding the subject company. Pursuant to the Market Abuse Regulation a history of all Citi Research recommendations published during the preceding 12-month period can be accessed via Citi Velocity (https://www.citivelocity.com/cv2) or your standard distribution portal. Historical disclosures (for up to the past three years) will be provided upon request.

Citi Research Equity Ratings Distribution

	12 Month Rating			Catalyst Watch		
Data current as of 01 Apr 2025	**Buy**	**Hold**	**Sell**	**Buy**	**Hold**	**Sell**
Citi Research Global Fundamental Coverage (Neutral=Hold)	58%	33%	9%	36%	48%	17%
% of companies in each rating category that are investment banking clients	37%	39%	29%	39%	36%	36%

Guide to Citi Research Fundamental Research Investment Ratings:
Citi Research stock recommendations include an investment rating and an optional risk rating to highlight high risk stocks. **Risk rating** takes into account both price volatility and fundamental criteria. Stocks will either have no risk rating or a High risk rating assigned.
Investment Ratings: Citi Research investment ratings are Buy, Neutral and Sell. Our ratings are a function of analyst expectations of expected total return ("ETR") and risk. ETR is the sum of the forecast price appreciation (or depreciation) plus the dividend yield for a stock within the next 12 months. The target price is based on a 12 month time horizon. The Investment rating definitions are: Buy (1) ETR of 15% or more or 25% or more for High risk stocks; and Sell (3) for negative ETR. Any covered stock not assigned a Buy or a Sell is a Neutral (2). For stocks rated Neutral (2), if an analyst believes that there are insufficient valuation drivers and/or investment catalysts to derive a positive or negative investment view, they may elect with the approval of Citi Research management not to assign a target price and, thus, not derive an ETR. Citi Research may suspend its rating and target price and assign “Rating Suspended” status for regulatory and/or internal policy reasons. Citi Research may also suspend its rating and target price and assign "Under Review" status for other exceptional circumstances (e.g. lack of information critical to the analyst's thesis, trading suspension) affecting the company and/or trading in the company's securities. In both such situations, the rating and target price will show as “—“ and “-“ respectively in the rating history price chart. Prior to 11 April 2022 Citi Research assigned “Under Review” status to both situations and prior to 11 Nov 2020 only in exceptional circumstances. As soon as practically possible, the analyst will publish a note re-establishing a rating and investment thesis. Investment ratings are determined by the ranges described above at the time of initiation of coverage, a change in investment and/or risk rating, or a change in target price (subject to limited management discretion). At times, the expected total returns may fall outside of these ranges because of market price movements and/or other short-term volatility or trading patterns. Such interim deviations will be permitted but will become subject to review by Research Management. Your decision to buy or sell a security should be based upon your personal investment objectives and should be made only after evaluating the stock's expected performance and risk.

Catalyst Watch/Short Term Views (“STV”) Ratings Disclosure:
Catalyst Watch and STV Upside/Downside calls: Citi Research may also include a Catalyst Watch or STV Upside or Downside call to indicate the analyst expects the share price to rise (fall) in absolute terms over a specified period of 30 or 90 days in reaction to one or more specific near-term catalysts or events impacting the company or the market. A Catalyst Watch will be published when Analyst confidence is high that an impact to share price will occur; it will be a STV when confidence level is moderate. A Catalyst Watch or STV Upside/Downside call will automatically expire at the end of the specified 30/90 day period (though only a Catalyst Watch expiration will generate a published research note). The Catalyst Watch will also be automatically removed with a published research note if share price performance (calculated at market close) exceeds 15% against the direction of the call (unless over-ridden by the analyst). The analyst may also remove a Catalyst Watch or STV call prior to the end of the specified period in a published research note. A Catalyst Watch/STV Upside or Downside call may be different from and does not affect a stock’s fundamental equity rating, which reflects a longer-term total absolute return expectation. For purposes of FINRA ratings-distribution-disclosure rules, a Catalyst Watch/STV Upside call corresponds to a buy recommendation and a Catalyst Watch/STV Downside call corresponds to a sell recommendation. Any stock not assigned to a Catalyst Watch Upside, Catalyst Watch Downside, STV Upside, or STV Downside call is considered Catalyst Watch/STV No View. For purposes of FINRA ratings distribution-disclosure rules, we correspond Catalyst Watch/STV No View to Hold in our ratings distribution table for our Catalyst Watch/STV Upside/Downside rating system. However, we reiterate that we do not consider No View to be a recommendation. For all Catalyst Watch/STV Upside/Downside calls, risk exists that the catalyst(s) and associated share-price movement will not materialize as expected.

RESEARCH ANALYST AFFILIATIONS / NON-US RESEARCH ANALYST DISCLOSURES
The legal entities employing the authors of this report are listed below (and their regulators are listed further herein). Non-US research analysts who have prepared this report (i.e., all research analysts listed below other than those identified as employed by Citigroup Global Markets Inc.) are not registered/qualified as research analysts with FINRA. Such research analysts may not be associated persons of the member organization (but are employed by an affiliate of the member organization) and therefore may not be subject to the FINRA Rule 2241 restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

Citigroup Global Markets Inc.	Papa Sylla; Elizabeth Sun; Michael Cadiz; Atif Malik; Christopher Danely
Citigroup Global Markets Korea Securities Limited	Jayden Oh; Peter Lee; Josh Yang
Citigroup Global Markets Taiwan Securities Company Limited	Carrie Liu; Laura (Chia Yi) Chen
Citigroup Global Markets Asia Limited	Kyna Wong; Kevin Chen

Citigroup Global Markets Canada Inc.	Asiya Merchant, CFA
Citigroup Global Markets Japan Inc.	Kota Ezawa; Yuta Nishiyama; Takayuki Naito; Masahiro Shibano
Citigroup Global Markets Limited	Andrew M. Gardiner, CFA

OTHER DISCLOSURES

Any price(s) of instruments mentioned in recommendations are as of the prior day's market close on the primary market for the instrument, unless otherwise stated.

The completion and first dissemination of any recommendations made within this research report are as of the Eastern date-time displayed at the top of the Product. If the Product references views of other analysts then please refer to the price chart or rating history table for the date/time of completion and first dissemination with respect to that view.

Regulations in various jurisdictions require that where a recommendation differs from any of the author's previous recommendations concerning the same financial instrument or issuer that has been published during the preceding 12-month period that the change(s) and the date of that previous recommendation are indicated. For fundamental coverage please refer to the price chart or rating change history within this disclosure appendix or the issuer disclosure summary at https://www.citivelocity.com/cvr/eppublic/citi_research_disclosures.

Citi Research has implemented policies for identifying, considering and managing potential conflicts of interest arising as a result of publication or distribution of investment research. A description of these policies can be found at https://www.citivelocity.com/cvr/eppublic/citi_research_disclosures.

The proportion of all Citi Research research recommendations that were the equivalent to "Buy","Hold","Sell" at the end of each quarter over the prior 12 months (with the % of these that had received investment firm services from Citi in the prior 12 months shown in brackets) is as follows: Q1 2025 Buy 33% (63%), Hold 44% (51%), Sell 23% (49%), RV 0.5% (87%); Q4 2024 Buy 32% (64%), Hold 44% (52%), Sell 23% (44%), RV 0.4% (92%); Q3 2024 Buy 33% (61%), Hold 44% (52%), Sell 23% (49%), RV 0.4% (80%); Q2 2024 Buy 33% (61%), Hold 44% (52%), Sell 23% (48%), RV 0.5% (90%). For the purposes of disclosing recommendations other than for equity (whose definitions can be found in the corresponding disclosure sections), "Buy" means a positive directional trade idea; "Sell" means a negative directional trade idea; and "Relative Value" means any trade idea which does not have a clear direction to the investment strategy.

European regulations require a 5 year price history when past performance of a security is referenced. CitiVelocity's Charting Tool (https://www.citivelocity.com/cv2/#go/CHARTING_3_Equities) provides the facility to create customisable price charts including a five year option. This tool can be found in the Data & Analytics section under any of the asset class menus in CitiVelocity (https://www.citivelocity.com/). For further information contact CitiVelocity support (https://www.citivelocity.com/cv2/go/CLIENT_SUPPORT). The source for all referenced prices, unless otherwise stated, is DataCentral, which sources price information from Thomson Reuters. Past performance is not a guarantee or reliable indicator of future results. Forecasts are not a guarantee or reliable indicator of future performance.

Investors should always consider the investment objectives, risks, and charges and expenses of an ETF carefully before investing. The applicable prospectus and key investor information document (as applicable) for an ETF should contain this and other information about such ETF. It is important to read carefully any such prospectus before investing. Clients may obtain prospectuses and key investor information documents for ETFs from the applicable distributor or authorized participant, the exchange upon which an ETF is listed and/or from the applicable website of the applicable ETF issuer. The value of the investments and any accruing income may fall or rise. Any past performance, prediction or forecast is not indicative of future or likely performance. Any information on ETFs contained herein is provided strictly for illustrative purposes and should not be deemed an offer to sell or a solicitation of an offer to purchase units of any ETF either explicitly or implicitly. The opinions expressed are those of the authors and do not necessarily reflect the views of ETF issuers, any of their agents or their affiliates.

Citigroup Global Markets India Private Limited and/or its affiliates may have, from time to time, actual or beneficial ownership of 1% or more in the debt securities of the subject issuer.

Please be advised that pursuant to Executive Order 13959 as amended (the "Order"), U.S. persons are prohibited from investing in securities of any company determined by the United States Government to be the subject of the Order. This research is not intended to be used or relied upon in any way that could result in a violation of the Order. Investors are encouraged to rely upon their own legal counsel for advice on compliance with the Order and other economic sanctions programs administered and enforced by the Office of Foreign Assets Control of the U.S. Treasury Department.

This communication is directed at persons who are "Eligible Clients" as such term is defined in the Israeli Regulation of Investment Advice, Investment Marketing and Investment Portfolio Management law, 1995 (the "Advisory Law"). Within Israel, this communication is not intended for retail clients and Citi will not make such products or transactions available to retail clients or to non-Eligible Clients. The presenter is not licensed as investment advisor or investment marketer by the Israeli Securities Authority ("ISA") and this communication does not constitute investment or marketing advice. The information contained herein may relate to matters that are not regulated by the ISA. Any securities which are the subject of this communication may not be offered or sold to any Israeli person except pursuant to a security offering exemption according to the Israeli Securities Law, 1968and the public offering rules provided thereunder.

Citi Research broadly and simultaneously disseminates its research content to the Firm's institutional and retail clients via the Firm's proprietary electronic distribution platforms (e.g., Citi Velocity and various Global Wealth platforms). As a convenience, certain, but not all, research content may be distributed through third party aggregators. Clients may receive published research reports by email, on a discretionary basis, and only after such research content has been broadly disseminated. Certain research is made available only to institutional investors to satisfy regulatory requirements. The level and types of services provided by Citi Research analysts to clients may vary depending on various factors such as the client's individual preferences as to the frequency and manner of receiving communications from analysts, the client's risk profile and investment focus and perspective (e.g. market-wide, sector specific, long term, short-term etc.), the size and scope of the overall client relationship with the Firm and legal and regulatory constraints.

Pursuant to Comissão de Valores Mobiliários Resolução 20 and ASIC Regulatory Guide 264, Citi is required to disclose whether a Citi related company or business has a commercial relationship with the subject company. Considering that Citi operates multiple businesses in more than 100 countries around the world, it is likely that Citi has a commercial relationship with the subject company.

Disclosure for investors in the Republic of Turkey: Under Capital Markets Law of Turkey (Law No: 6362), the investment information, comments and recommendations stated here, are not within the scope of investment advisory activity. Investment advisory service is provided in accordance with a contract of engagement on investment advisory concluded between brokerage houses, portfolio management companies, non-deposit banks and clients. Comments and recommendations stated here rely on the individual opinions of the ones providing these comments and recommendations. These opinions may not fit to your financial status, risk and return preferences. For this reason, to make an investment decision by relying solely to this information stated here may not bring about outcomes that fit your expectations. Furthermore, Citi Research is a division of Citigroup Global Markets Inc. (the "Firm"), which does and seeks to do business with companies and/or trades on securities covered in this research reports. As a result, investors should be aware that the Firm may have a conflict of interest that could affect the objectivity of this report, however investors should also note that the Firm has in place organisational and administrative arrangements to manage potential conflicts of interest of this nature.

Securities recommended, offered, or sold by the Firm: (i) are not insured by the Federal Deposit Insurance Corporation; (ii) are not deposits or other obligations of any insured depository institution (including Citibank); and (iii) are subject to investment risks, including the possible loss of the principal amount invested. The Product is for informational purposes only and is not intended as an offer or solicitation for the purchase or sale of a security. Any decision to purchase securities mentioned in the Product must take into account existing public information on such security or any registered prospectus. Although information has been obtained from and is based upon sources that the Firm believes to be reliable, we do not guarantee its accuracy and it may be incomplete and condensed. Note, however, that the Firm has taken all reasonable steps to determine the accuracy and completeness of the disclosures made in the Important Disclosures section of the Product. The Firm's research department has received assistance from the subject company(ies) referred to in this Product including, but not limited to, discussions with management of the subject company(ies). Firm policy prohibits research analysts from sending draft research to subject companies. However, it should be presumed that the author of the Product has had discussions with the subject company to ensure factual accuracy prior to publication. Statements and views concerning ESG (environmental, social, governance) factors are typically based upon public statements made by the affected company or other public news, which the author may not have independently verified. ESG factors are one consideration that investors may choose to examine when making investment decisions. All opinions, projections and estimates constitute the judgment of the author as of the date of the Product and these, plus any other information contained in the Product, are subject to change without notice. Prices and availability of financial instruments also are subject to change without notice. Notwithstanding other departments within the Firm advising the companies discussed in this Product, information obtained in such role is not used in the preparation of the Product. Although Citi Research does not set a predetermined frequency for publication, if the Product is a fundamental equity or credit research report, it is the intention of Citi Research to provide research coverage of the covered issuers, including in response to news affecting the issuer. For non-fundamental research reports, Citi Research may not provide regular updates to the views, recommendations and facts included in the reports. Notwithstanding that Citi Research maintains coverage on, makes recommendations concerning or discusses issuers, Citi Research may be periodically restricted from referencing certain issuers due to legal or policy reasons. Where a component of a published trade idea is subject to a restriction, the trade idea will be removed from any list of open trade ideas included in the Product. Upon the lifting of the restriction, the trade idea will either be re-instated in the open trade ideas list if the analyst continues to support it or it will be officially closed. Citi Research may provide different research products and services to different classes of customers (for example, based upon long-term or short-term investment horizons) that may lead to differing conclusions or recommendations that could impact the price of a security contrary to the recommendations in the alternative research product, provided that each is consistent with the rating system for each respective product.

Investing in non-U.S. securities, including ADRs, may entail certain risks. The securities of non-U.S. issuers may not be registered with, nor be subject to the reporting requirements of the U.S. Securities and Exchange Commission. There may be limited information available on foreign securities. Foreign companies are generally not subject to uniform audit and reporting standards, practices and requirements comparable to those in the U.S. Securities of some foreign companies may be less liquid and their prices more volatile than securities of comparable U.S. companies. In addition, exchange rate movements may have an adverse effect on the value of an investment in a foreign stock and its corresponding dividend payment for U.S. investors. Net dividends to ADR investors are estimated, using withholding tax rates conventions, deemed accurate, but investors are urged to consult their tax advisor for exact dividend computations. Investors who have received the Product from the Firm may be prohibited in certain states or other jurisdictions from purchasing securities mentioned in the Product from the Firm. Please ask your Financial Consultant for additional details. Citigroup Global Markets Inc. takes responsibility for the Product in the United States. Any orders by US investors resulting from the information contained in the Product may be placed only through Citigroup Global Markets Inc.

The Citigroup legal entity that takes responsibility for the production of the Product is the legal entity which the first named author is employed by.

The Product is made available in **Australia** through Citigroup Global Markets Australia Pty Limited. (ABN 64 003 114 832 and AFSL No. 240992), participant of the ASX Group and regulated by the Australian Securities & Investments Commission. Citigroup Centre, 2 Park Street, Sydney, NSW 2000. Citigroup Global Markets Australia Pty Limited is not an Authorised Deposit-Taking Institution under the Banking Act 1959, nor is it regulated by the Australian Prudential Regulation Authority.

The Product is made available in **Brazil** by Citigroup Global Markets Brasil - CCTVM SA, which is regulated by CVM - Comissão de Valores Mobiliários ("CVM"), BACEN - Brazilian Central Bank, APIMEC - Associação dos Analistas e Profissionais de Investimento do Mercado de Capitais and ANBIMA – Associação Brasileira das Entidades dos Mercados Financeiro e de Capitais. Av. Paulista, 1111 - 14º andar(parte) - CEP: 01311920 - São Paulo - SP.

This Product is available in **Chile** through Banchile Corredores de Bolsa S.A., an indirect subsidiary of Citigroup Inc., which is regulated by the Comisión Para El Mercado Financiero. Agustinas 975, piso 2, Santiago, Chile.

Disclosure for investors in the Republic of Colombia :This communication or message does not constitute a professional recommendation to make investment in the terms of article 2.40.1.1.2 of Decree 2555 de 2010 or the regulations that modify, substitute or complement it. Para la elaboración y distribución de informes de investigación y de comunicaciones generales de que trata este artículo no se requiere ser una entidad vigilada por la Superintendencia Financiera de Colombia.

The Product is made available in **Germany** by Citigroup Global Markets Europe AG ("CGME"), which is regulated by the European Central Bank and the German Federal Financial Supervisory Authority (Bundesanstalt fur Finanzdienstleistungsaufsicht BaFin). Börsenplatz 9, 60313 Frankfurt am Main, Germany.

Unless otherwise specified, if the analyst who prepared this report is based in Hong Kong and it relates to “securities” (as defined in the Securities and Futures Ordinance (Cap.571 of the Laws of Hong Kong)), the report is issued in Hong Kong by Citigroup Global Markets Asia Limited. Citigroup Global Markets Asia Limited is regulated by Hong Kong Securities and Futures Commission. If the report is prepared by a non-Hong Kong based analyst, please note that such analyst (and the legal entity that the analyst is employed by or accredited to) is not licensed/registered in Hong Kong and they do not hold themselves out as such. Please refer to the section “Research Analyst Affiliations / Non-US Research Analyst Disclosures” for the details of the employment entity of the analysts.

The Product is made available in **India** by Citigroup Global Markets India Private Limited (CGM), which is regulated by the Securities and Exchange Board of India (SEBI), as a Research Analyst (SEBI Registration No. INH000000438). CGM is also actively involved in the business of merchant banking (SEBI Registration No. INM000010718) and stock brokerage ((SEBI Registration No. INZ000263033) in India, and is registered with SEBI in this regard. Registration granted by SEBI and certification from National Institute of Securities Markets (NISM) in no way guarantee performance of the intermediary or provide any assurance of returns to investors. CGM’s registered office is at 1202, 12th Floor, First International Financial Centre (FIFC), G Block, Bandra Kurla Complex, Bandra East, Mumbai – 400098 & registered Tel: +91 22 61759999. Citi Research utilizes various AI tools and/or sourcesin providing research services. Citi maintains robust policies, procedures,controls, and training to ensure continued compliance with all applicable rulesand regulations. All recommendations contained herein are made by dulyqualified research analysts. CGM’s Corporate Identity Number is U99999MH2000PTC126657, and its Compliance Officer [Vishal Bohra] contact details are: Tel:+91-022-61759994, Fax:+91-022-61759851, Email: cgmcompliance@citi.com. The Investor Charter in respect of Research Analysts and Complaints information can be found at https://www.citivelocity.com/cvr/eppublic/citi_research_disclosures. The grievance officer [Nikita Jadhav] contact details are Tel: +91-022-42775089, Email: EMEA.CR.Complaints@citi.com. Investment in securities market are subject to market risks. Read all the related documents carefully before investing.

The Product is made available in **Indonesia** through PT Citigroup Sekuritas Indonesia. Citibank Tower 10/F, Pacific Century Place, SCBD lot 10, Jl. Jend Sudirman Kav 52-53, Jakarta 12190, Indonesia. Neither this Product nor any copy hereof may be distributed in Indonesia or to any Indonesian citizens wherever they are domiciled or to Indonesian residents except in compliance with applicable capital market laws and regulations. This Product is not an offer of securities in Indonesia. The securities referred to in this Product have not been registered with the Capital Market and Financial Services Authority (OJK) pursuant to relevant capital market laws and regulations, and may not be offered or sold within the territory of the Republic of Indonesia or to Indonesian citizens through a public offering or in circumstances which constitute an offer within the meaning of the Indonesian capital market laws and regulations.

The Product is made available in **Japan** by Citigroup Global Markets Japan Inc. ("CGMJ"), which is regulated by Financial Services Agency, Securities and Exchange Surveillance Commission, Japan Securities Dealers Association, Tokyo Stock Exchange and Osaka Securities Exchange. Otemachi Park Building, 1-1-1 Otemachi, Chiyoda-ku, Tokyo 100-8132 Japan. In the event that an error is found in an CGMJ research report, a revised version will be posted on the Firm's Citi Velocity website. If you have questions regarding Citi Velocity, please call (81 3) 6270-3019 for help.

The product is made available in the **Kingdom of Saudi Arabia** in accordance with Saudi laws through Citigroup Saudi Arabia, which is regulated by the Capital Market Authority (CMA) under CMA license (17184-31). 2239 Al Urubah Rd – Al Olaya Dist. Unit No. 18, Riyadh 12214 – 9597, Kingdom Of Saudi Arabia.

The Product is made available in **Korea** by Citigroup Global Markets Korea Securities Ltd. (CGMK), which is regulated by the Financial Services Commission, the Financial Supervisory Service and the Korea Financial Investment Association (KOFIA). The address of CGMK is Citibank Center, 50 Saemunan-ro, Jongno-gu, Seoul 03184, Korea. KOFIA makes available registration information of research analysts on its website. Please visit the following website if you wish to find KOFIA registration information on research analysts of CGMK. http://dis.kofia.or.kr/websquare/index.jsp?w2xPath=/wq/fundMgr/DISFundMgrAnalystList.xml&divisionId=MDIS03002002000000&serviceId=SDIS03002002000. The Product is made available in Korea by Citibank Korea Inc., which is regulated by the Financial Services Commission and the Financial Supervisory Service. Address is Citibank Center, 50 Saemunan-ro, Jongno-gu, Seoul 03184, Korea. This research report is intended to be provided only to Professional Investors as defined in the Financial Investment Services and Capital Market Act and its Enforcement Decree in Korea.

The Product is made available in **Malaysia** by Citigroup Global Markets Malaysia Sdn Bhd (Registration No. 199801004692 (460819-D)) ("CGMM") to its clients and CGMM takes responsibility for its contents as regards CGMM's clients. CGMM is regulated by the Securities Commission Malaysia. Please contact CGMM at Level 43 Menara Citibank, 165 Jalan Ampang, 50450 Kuala Lumpur, Malaysia in respect of any matters arising from, or in connection with, the Product.

The Product is made available in **Mexico** by Citi México Casa de Bolsa, S.A. de C.V., Grupo Financiero Citi México which is a wholly owned subsidiary of Citigroup Inc. and is regulated by Comision Nacional Bancaria y de Valores. Prolongación Reforma 1196, 24 floor, Colonia Santa Fe, Alcaldía Cuajimalpa de Morelos, C.P. 05348, Ciudad de México.

The Product is made available in **Poland** by Biuro Maklerskie Banku Handlowego (DMBH), separate department of Bank Handlowy w Warszawie S.A. a subsidiary of Citigroup Inc., which is regulated by Komisja Nadzoru Finansowego. Biuro Maklerskie Banku Handlowego (DMBH), ul.Senatorska 16, 00-923 Warszawa.

The Product is made available in **Singapore** through Citigroup Global Markets Singapore Pte. Ltd. ("CGMSPL"), a capital markets services and Exempt Financial Advisor license holder, and regulated by Monetary Authority of Singapore. Please contact CGMSPL at 8 Marina View, 21st Floor Asia Square Tower 1, Singapore 018960, in respect of any matters arising from, or in connection with, the analysis of this document. This Product is intended for recipients who are accredited, expert and institutional investors as defined under the Securities and Futures Act 2001. For Citi Private Bank, the Product is made available in Singapore by Citi Private Bank through Citibank, N.A., Singapore Branch. Citibank N.A., Singapore Branch is a licensed bank in Singapore that is regulated by the Monetary Authority of Singapore. Please contact your Private Banker in Citibank N.A., Singapore Branch if you have any queries on or any matters arising from or in connection with this document. The Product is intended for recipients who are accredited, expert and institutional investors as defined under the Securities and Futures Act 2001. For Citibank Singapore Limited ("CSL"), the Product is distributed in Singapore by CSL to selected Citigold/Citigold Private Clients. CSL provides no independent research or analysis of the substance or in preparation of the Product. Please contact your Citigold//Citigold Private Client Relationship Manager in CSL if you have any queries on or any matters arising from or in connection with this document. The Product is intended for recipients who are accredited investors as defined under the Securities and Futures Act (Cap. 289).

Citigroup Global Markets (Pty) Ltd. is incorporated in the **Republic of South Africa** (company registration number 2000/025866/07) and its registered office is at 145 West Street, Sandton, 2196, Saxonwold. Citigroup Global Markets (Pty) Ltd. is regulated by JSE Securities Exchange South Africa, South African Reserve Bank and the Financial Services Board. The investments and services contained herein are not available to private customers in South Africa.

The Product is made available in the **Republic of China (Taiwan)** through Citigroup Global Markets Taiwan Securities Company Ltd. ("CGMTS"), 14F, 15F and 16F, No. 1, Songzhi Road, Taipei 110, Taiwan, subject to the license scope and the applicable laws and regulations in the Republic of China (Taiwan). CGMTS is regulated by the Securities and Futures Bureau of the Financial Supervisory Commission of Taiwan, the Republic of China (Taiwan). No portion of the Product may be reproduced or quoted in the Republic of China (Taiwan) by the press or any third parties [without the written authorization of CGMTS]. Pursuant to the applicable laws and regulations in the Republic of China (Taiwan), the recipient of the Product shall not take advantage of such Product to involve in any matters in which the recipient may have conflicts of interest. If the Product covers securities which are not allowed to be offered or traded in the Republic of China (Taiwan), neither the Product nor any information contained in the Product shall be considered as advertising the securities or making recommendation of the securities in the Republic of China (Taiwan). The Product is for informational purposes only and is not intended as an offer or solicitation for the purchase or sale of a security or financial products. Any decision to purchase securities or financial products mentioned in the Product must take into account existing public information on such security or the financial products or any registered prospectus.

The Product is made available in **Thailand** through Citicorp Securities (Thailand) Ltd., which is regulated by the Securities and Exchange Commission of Thailand. 399 Interchange 21 Building, 18th Floor, Sukhumvit Road, Klongtoey Nua, Wattana ,Bangkok 10110, Thailand.

The Product is made available in **Turkey** through Citibank AS which is regulated by Capital Markets Board. Tekfen Tower, Eski Buyukdere Caddesi # 209 Kat 2B, 23294 Levent, Istanbul, Turkey.

In the **U.A.E**, these materials (the "Materials") are communicated byCitigroup Global Markets Limited, DIFC branch ("CGML"), an entityregistered in the Dubai International Financial Center ("DIFC") andlicensed and regulated by the Dubai Financial Services Authority ("DFSA",license #CL0221) to Professional Clients and Market Counterparties, as definedin DFSA regulations, only and should not be relied upon or distributed toRetail Clients. Financial products and/or services to which the Materialsrelate will only be made available to Professional Clients and MarketCounterparties. Citigroup Global Markets Limited DIFC Branch registered addressis Level 3, Gate District Building 02, Dubai International Financial Centre andcan be contacted on +971 4 509 97 90.

The Product is made available in **United Kingdom** by Citigroup Global Markets Limited, which is authorised by the Prudential Regulation Authority ("PRA") and regulated by the Financial Conduct Authority ("FCA") and the PRA. This material may relate to investments or services of a person outside of the UK or to other matters which are not authorised by the PRA nor regulated by the FCA and the PRA and further details as to where this may be the case are available upon request in respect of this material. Citigroup Centre, Canada Square, Canary Wharf, London, E14 5LB.

The Product is made available in **United States** and **Canada** by Citigroup Global Markets Inc., which is a member of FINRA and registered with the US Securities and Exchange Commission. 388 Greenwich Street, New York, NY 10013.

Unless specified to the contrary, within EU Member States, the Product is made available by Citigroup Global Markets Europe AG ("CGME"), which is regulated by the European Central Bank and the German Federal Financial Supervisory Authority (Bundesanstalt für Finanzdienstleistungsaufsicht-BaFin).

The Product is not to be construed as providing investment services in any jurisdiction where the provision of such services would not be permitted.
Subject to the nature and contents of the Product, the investments described therein are subject to fluctuations in price and/or value and investors may get back less than originally invested. Certain high-volatility investments can be subject to sudden and large falls in value that could equal or exceed the amount invested.The yield and average life of CMOs (collateralized mortgage obligations) referenced in this Product will fluctuate depending on the actual rate at which mortgage holders prepay the mortgages underlying the CMO and changes in current interest rates. Any government agency backing of the CMO applies only to the face value of the CMO and not to any premium paid. Certain investments contained in the Product may have tax implications for private customers whereby levels and basis of taxation may be subject to change. If in doubt, investors should seek advice from a tax adviser. The Product does not purport to identify the nature of the specific market or other risks associated with a particular transaction. Advice in the Product is general and should not be construed as personal advice given it has been prepared without taking account of the objectives, financial situation or needs of any particular investor. Accordingly, investors should, before acting on the advice, consider the appropriateness of the advice, having regard to their objectives, financial situation and needs. Prior to acquiring any financial product, it is the client's responsibility to obtain the relevant offer document for the product and consider it before making a decision as to whether to purchase the product.
Citi Research product may source data from dataCentral. dataCentral is a Citi Research proprietary database, which includes the Firm's estimates, data from company reports and feeds from Thomson Reuters. The source for all referenced prices, unless otherwise stated, is DataCentral. Past performance is not a guarantee or reliable indicator of future results. Forecasts are not a guarantee or reliable indicator of future performance. The printed and printable version of the research report may not include all the information (e.g. certain financial summary information and comparable company data) that is linked to the online version available on the Firm's proprietary electronic distribution platforms.

Card Insights. Where this report references Card Insights data, Card Insights consists of selected data from a subset of Citi's proprietary credit card transactions. Such data has undergone rigorous security protocols to keep all customer information confidential and secure; the data is highly aggregated and anonymized so that all unique customer identifiable information is removed from the data prior to receipt by the report's author or distribution to external parties. This data should be considered in the context of other economic indicators and publicly available information. Further, the selected data represents only a subset of Citi's proprietary credit card transactions due to the selection methodology or other limitations and should not be considered as indicative or predictive of the past or future financial performance of Citi or its credit card business.

Where included in this report, MSCI sourced information is the exclusive property of Morgan Stanley Capital International Inc. (MSCI). Without prior written permission of MSCI, this information and any other MSCI intellectual property may not be reproduced, redisseminated or used to create any financial products, including any indices. This information is provided on an "as is" basis. The user assumes the entire risk of any use made of this information. MSCI, its affiliates and any third party involved in, or related to, computing or compiling the information hereby expressly disclaim all warranties of originality, accuracy, completeness, merchantability or fitness for a particular purpose with respect to any of this information. Without limiting any of the foregoing, in no event shall MSCI, any of its affiliates or any third party involved in, or related to, computing or compiling the information have any liability for any damages of any kind. MSCI, Morgan Stanley Capital International and the MSCI indexes are services marks of MSCI and its affiliates. Part of this product may contain Sustainalytics proprietary information that may not be reproduced, used, disseminated, modified nor published in any manner without the express written consent of Sustainalytics. Sustainalytics, its affiliates and any third party involved in, or related to, computing or compiling the information hereby expressly disclaim all warranties of originality, accuracy, completeness, merchantability or fitness for a particular purpose with respect to any of this information. Any information attributed to Sustainalytics is provided solely for informational purposes and on its own should not be considered an offer to buy or sell a security. Neither Sustainalytics nor all its third-party suppliers provide investment advice (as defined in the applicable jurisdiction) or any other form of (financial) advice. The information is provided "as is" and, therefore Sustainalytics assumes no responsibility for errors or omissions. Sustainalytics cannot be held liable for damage arising from the use of this product or information contained herein in any manner whatsoever. Where data is attributed to Morningstar that data is © 2025 Morningstar, Inc. All Rights Reserved. That information: (1) is proprietary to Morningstar and/or its content providers; (2) may not be copied or distributed; and (3) is not warranted to be accurate, complete or timely. Neither Morningstar nor its content providers are responsible for any damages or losses arising from any use of this information.

The Firm accepts no liability whatsoever for the actions of third parties. The Product may provide the addresses of, or contain hyperlinks to, websites. Except to the extent to which the Product refers to website material of the Firm, the Firm has not reviewed the linked site. Equally, except to the extent to which the Product refers to website material of the Firm, the Firm takes no responsibility for, and makes no representations or warranties whatsoever as to, the data and information contained therein. Such address or hyperlink (including addresses or hyperlinks to website material of the Firm) is provided solely for your convenience and information and the content of the linked site does not in any way form part of this document. Accessing such website or following such link through the Product or the website of the Firm shall be at your own risk and the Firm shall have no liability arising out of, or in connection with, any such referenced website.

ADDITIONAL INFORMATION IS AVAILABLE UPON REQUEST