

April 19, 2026 08:01 PM GMT

Technology

Rise of the AI Agent – Global Implications

The agentic AI shift isn't only about more raw performance; it's about more coordination, with the CPU increasingly determining how well AI systems run. The next phase of AI infrastructure is smarter allocation across the whole stack, where CPU, foundry, memory, substrates and equipment become the bottleneck.

Key Takeaways

- Agentic AI widens the trade beyond GPUs, with CPUs becoming the control plane for multi-step workflows and system orchestration.
- We estimate \$32.5-60bn incremental CPU TAM by 2030 within a total server CPU TAM of \$100bn+, driven by agentic workloads.
- Memory becomes a key monetization leg, with agentic AI driving 26-77% of incremental DRAM demand by 2030.
- Supply-constrained enablers (foundry, ABF, BMC, interconnect) should capture outsized economics as system complexity rises.
- The beneficiaries are full-stack and global, spanning CPU, memory, substrates, infrastructure ICs and equipment.

Agentic AI marks a structural shift from compute to orchestration: While GPU demand remains strong, each model call now requires more coordination, memory and system-level compute, widening the AI spend pool beyond accelerators to CPUs and the broader infrastructure stack. As AI transitions from generation to autonomous action, the computing bottleneck is shifting towards CPU and memory, driving a step-change in general-purpose compute intensity.

Quantifying the incremental CPU TAM: We introduce a new framework to size the opportunity across compute and memory. CPU-side orchestration can account for 50-90% of total workload latency, materially increasing CPU intensity at the system level. We estimate \$32.5-60bn of incremental CPU TAM by 2030, within a total server CPU TAM of \$100bn+, and 15-45EB of incremental DRAM demand.

Architecture shift drives system-level content growth: Agentic workloads require CPU-centric or hybrid architectures to manage multi-step reasoning, tool execution and memory orchestration. The CPU-to-GPU ratio rises at the cluster level, and memory evolves from a passive storage layer into an active system component, supporting persistent context and continuous learning. As a result, content increases structurally across CPUs, DRAM and the broader infrastructure stack, including foundry, ABF substrates and interconnect layers.

Exposure to the theme: The opportunity extends beyond individual chips to the full AI system. The beneficiaries of this shift are global and full-stack – see [Exhibit 1](#).

MORGAN STANLEY & CO. INTERNATIONAL PLC+

Shawn Kim
Equity Analyst
Shawn.Kim@morganstanley.com +44 20 7677-1018

Lee Simpson
Equity Analyst
Lee.Simpson@morganstanley.com +44 20 7425-3378

MORGAN STANLEY & CO. LLC

Joseph Moore
Equity Analyst
Joseph.Moore@morganstanley.com +1 212 761-7516

MORGAN STANLEY TAIWAN LIMITED+

Charlie Chan
Equity Analyst
Charlie.Chan@morganstanley.com +886 2 2730-1725

MORGAN STANLEY & CO. LLC

Erik W Woodring
Equity Analyst
Erik.Woodring@morganstanley.com +1 212 296-8083

WHAT'S CHANGED		
SK hynix (000660.KS) Price Target	From W1,300,000	To W1,700,000
Samsung Electronics (005930.KS) Price Target	From W251,000	To W362,000
Samsung Electronics (005935.KS) Price Target	From W213,350	To W289,600
Samsung Electro-Mechanics (009150.KS) Price Target	From W500,000	To W710,000

Exhibit 1 : Beneficiaries of agentic AI take-up

Global Exposure Across the Stack ? Names by exposure					
CPU • NVDA • Intel • AMD • Arm	DRAM • Samsung • Hynix • Micron	NAND • Kioxia • SanDisk	HDD • Seagate • WDC • TDK	FOUNDRY • TSMC	IC-design • GUC • Egis
PCB/Substrate/CCL & Materials • SEMCO • Unimicron • NYPCB • Ibsiden • Nittobo • MEC			BMC, CPU & Memory interface • Aspeed • Renesas • Montage • WPG • AP Memory		
MLCC & CPU socket • Murata • TDK • Yageo • FIT Hon Teng • Lotes	ODM • Wiyynn • Hon Hai	SPE • ASML • ASMi • AMAT • Besi • KLAC • Tokyo Electron • Ulvac • Wonik			

Source: Morgan Stanley Research

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

+ = Analysts employed by non-U.S. affiliates are not registered with FINRA, may not be associated persons of the member and may not be subject to FINRA restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

Table of Contents

Table of Contents	
Our Thesis in 6 Charts	5
The Shift to Agentic AI in Numbers	6
Executive Summary	7
Changing Tides – From Generative AI to Agentic AI	8
Global Perspective	16
Growing the CPU Opportunity	23
Agentic AI Memory Opportunity	29
ABF Substrate – Still Early Innings	33
More CPU Sockets Will Be Needed	38
Passive Components TAM Will Also Grow	39
Greater China Semiconductor: Foundry, Design Service and IC Design – Seizing the CPU (Agentic AI) Opportunity	40
Earnings Revisions and Price Target Changes	50
Appendix – CPU vs. GPU, Inference vs. Agent	62

Contributors



MORGAN STANLEY & CO. INTERNATIONAL PLO+

Shawn Kim
Equity Analyst
+44 20 7677-1018
Shawn.Kim@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Charlie Chan
Equity Analyst
+886 2 2730-1725
Charlie.Chan@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Howard Kao
Equity Analyst
+886 2 2730-2989
Howard.Kao@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLO+

Nigel van Putten
Equity Analyst
+44 20 7425-2803
Nigel.Putten@morganstanley.com



MORGAN STANLEY ASIA LIMITED+

Duan Liu
Equity Analyst
+852 2239-7357
Duan.Liu@morganstanley.com



MORGAN STANLEY MUFG SECURITIES CO., LTD.+

Kazuo Yoshikawa, CFA
Equity Analyst
+81 3 6836-8408
Kazuo.Yoshikawa@morganstanleymufg.com



MORGAN STANLEY MUFG SECURITIES CO., LTD.+

Hidetaka Suzuki
Research Associate
+81 3 6836-8402
Hidetaka.Suzuki@morganstanleymufg.com



MORGAN STANLEY & CO. LLO

Nicole Kozhukhov
Research Associate
+1 212 761-1636
Nicole.Kozhukhov@morganstanley.com



MORGAN STANLEY ASIA LIMITED+

Daisy Dai, CFA
Equity Analyst
+852 2848-7310
Daisy.Dai@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Irene Yen
Research Associate
+886 2 2730-2869
Irene.Yen@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLO+

Lee Simpson
Equity Analyst
+44 20 7425-3378
Lee.Simpson@morganstanley.com



MORGAN STANLEY & CO. LLO

Erik W Woodring
Equity Analyst
+1 212 296-8083
Erik.Woodring@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Sharon Shih
Equity Analyst
+886 2 2730-2865
Sharon.Shih@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLO+

Cindy Huang
Equity Analyst
+44 20 7425-2915
Cindy.Huang@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLC, SEOUL BRANCH+

Ryan Kim
Equity Analyst
+82 2 399-4939
Ryan.G.Kim@morganstanley.com



MORGAN STANLEY MUFG SECURITIES CO., LTD.+

Yoshihito Hasegawa
Equity Analyst
+81 3 6836-8910
Yoshihito.Hasegawa@morganstanleymufg.com



MORGAN STANLEY & CO. LLO

Shane Brett
Equity Analyst
+1 212 761-1022
Shane.Brett@morganstanley.com



MORGAN STANLEY & CO. LLO

Ella Tulchinsky
Research Associate
+1 212 761-2222
Ella.Tulchinsky@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Tiffany Yeh
Equity Analyst
+886 2 7712-3032
Tiffany.Yeh@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Samantha Chen
Research Associate
+886 2 2730-2876
Samantha.Chen@morganstanley.com



MORGAN STANLEY & CO. LLO

Joseph Moore
Equity Analyst
+1 212 761-7516
Joseph.Moore@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Daniel Yen, CFA
Equity Analyst
+886 2 2730-2863
Daniel.Yen@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Derrick Yang
Equity Analyst
+886 2 2730-2862
Derrick.Yang@morganstanley.com



MORGAN STANLEY & CO. INTERNATIONAL PLO+

Amelia M Scicluna
Research Associate
+44 20 7425-6694
Amelia.Scicluna@morganstanley.com



MORGAN STANLEY MUFG SECURITIES CO., LTD.+

Shoji Sato
Equity Analyst
+81 3 6836-8404
Shoji.Sato@morganstanleymufg.com



MORGAN STANLEY MUFG SECURITIES CO., LTD.+

Suzune Tamura, CFA
Equity Analyst
+81 3 6836-8891
Suzune.Tamura@morganstanleymufg.com



MORGAN STANLEY & CO. LLO

Mason Wayne
Research Associate
+1 212 761-6012
Mason.Wayne@morganstanley.com



MORGAN STANLEY & CO. LLO

Dylan Liu
Research Associate
+1 212 761-4519
Dylan.Liu@morganstanley.com



MORGAN STANLEY TAIWAN LIMITED+

Lucas Wang
Research Associate
+886 2 2730-2875
Lucas.Wang@morganstanley.com

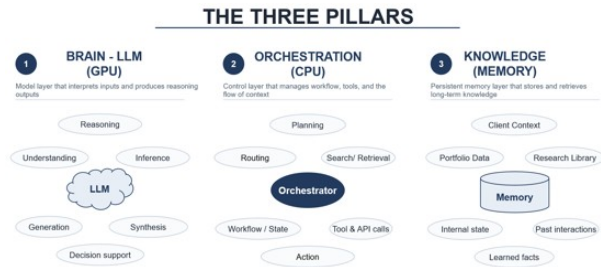


MORGAN STANLEY TAIWAN LIMITED+

Vivi Huang
Research Associate
+886 2 2730-2860
Vivi.Huang@morganstanley.com

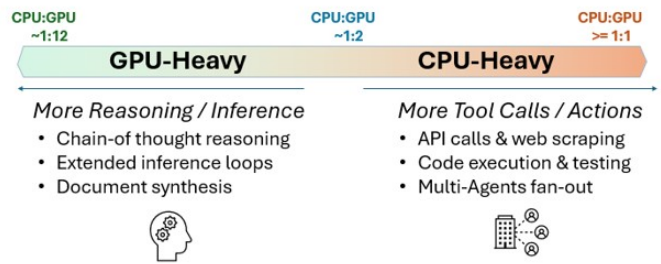
Our Thesis in 6 Charts

Exhibit 2: AI transitions from 'generation' to 'autonomous' action in agentic AI




Source: Morgan Stanley Research

Exhibit 3: Cluster-level CPU:GPU intensity rises as AI moves from reasoning to action




Source: Morgan Stanley Research

Exhibit 4: Agentic AI could add \$32.5-60bn CPU opportunity...



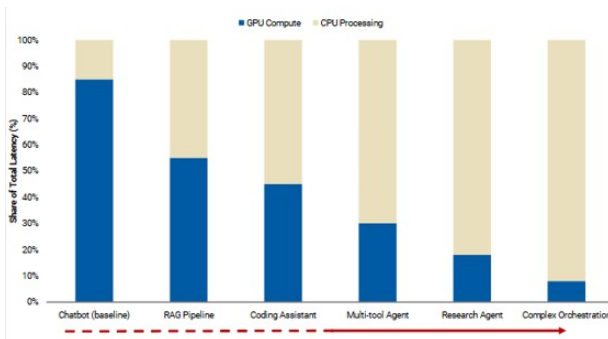
Source: Morgan Stanley Research estimates

Exhibit 5: ...and 15-45EB of DRAM demand by 2030




Source: Morgan Stanley Research estimates

Exhibit 6: Agentic AI shifts the latency bottleneck from GPUs to CPUs as workflows become more action-heavy




Source: Morgan Stanley Research, Georgia Tech and Intel paper. Note: Per-workload splits are our own directional estimates, not measured benchmarks.

Exhibit 7: We estimate agentic AI could drive 5-10% incremental ABF substrate demand upside by 2030 from the enlarged CPU TAM alone




Source: Prismark, Morgan Stanley Research estimates (e)

The Shift to Agentic AI in Numbers



Morgan Stanley's framework points to **\$32.5–60bn** incremental CPU TAM by **2030** within a **\$82.5–110bn** total server CPU TAM;



and **agentic orchestration** could drive **15–45EB** of incremental DRAM demand by **2030**, equivalent to roughly **26–77%** of **2027** DRAM supply.



CPU-side processing

can account for roughly **50–90%** of end-to-end latency in agentic workloads;



Cluster-level CPU:GPU

intensity rises from **1:12** to **> 1:1** of end-to-end latency in agentic workloads;



AI-driven ABF

substrate growth accelerates to a **17.9%** CAGR through **2030**



Source: Morgan Stanley Research estimates

Executive Summary



Tech Diffusion

A Morgan Stanley Research
Key Theme of 2026

How agentic AI reshapes the AI value chain

AI is entering a new phase. The first wave of generative AI was dominated by powerful chips running single tasks, such as answering queries or generating text. The next phase – agentic AI – goes further. These systems can plan, reason, remember past actions, and interact with tools to complete multi-step tasks. As a result, the economic value in AI is shifting away from any single chip and toward the overall system that supports these agents. Crucially, agentic AI does not reduce demand for GPUs, but it adds more work around them. Planning, coordinating tasks, managing memory and calling external tools all rely heavily on CPUs, memory, networking and storage.

CPU-side processing can account for 50-90% of total workload time in agentic systems. As these systems scale, the balance of infrastructure changes, with cluster-level CPU-to-GPU intensity rising. This has material implications for industry demand. We estimate agentic AI could create a \$32.5-60bn incremental total addressable market (TAM) for orchestration CPUs by 2030, within a total data center CPU market of \$82.5-110bn. Persistent memory requirements are equally important: our analysis suggests agentic workloads could drive 15-45 exabytes (EB) of additional DRAM demand by 2030, equivalent to 26-77% of 2027 annual DRAM supply.

The agentic AI shift broadens well beyond headline AI chips and we lay out below which stocks offer greatest exposure. Beneficiaries include CPU vendors, memory suppliers, storage companies, and advanced packaging and substrate providers, alongside foundries, equipment makers and server manufacturers. In short, agentic AI widens the AI investment landscape, shifting focus from owning the best accelerator to enabling the full system that makes intelligent agents work.

Global Exposure Across the Stack ? Names by exposure					
CPU • NVDA • Intel • AMD • Arm	DRAM • Samsung • Hynix • Micron	NAND • Kioxia • SanDisk	HDD • Seagate • WDC • TDK	FOUNDRY • TSMC	IC-design • GUC • Egis
PCB/Substrate/CCL & Materials • SEMCO • Unimicron • NYPCB • Ibiden • Nittobo • MEC			BMC, CPU & Memory interface • Aspeed • Renesas • Montage • WPG • AP Memory		
MLCC & CPU socket • Murata • TDK • Yageo • FIT Hon Teng • Lotes		ODM • Wiyynn • Hon Hai	SPE • ASML • ASMi • AMAT • Besi • KLAC • Tokyo Electron • Ulvac • Wonik		

Source: Morgan Stanley Research

Changing Tides – From Generative AI to Agentic AI

“The CPU is no longer simply supporting the model; it’s driving it,” NVIDIA CEO Jensen Huang during the GTC 2026 keynote, March 16, 2026.

We are moving from AI that answers questions to AI that takes action. AI is no longer just about making models bigger; it is now about building AI systems that operate continuously, making them smarter, more independent, and capable of operating at scale. Today’s AI computing is limited less by raw compute (GPU), and more by memory bandwidth, data movement, interconnect latency, and system coordination. Data centers today are no longer one single chip (GPU) but fully integrated stacks (system architecture), where each layer is designed to remove bottlenecks between compute, memory, and data flow.

The AI race is evolving from static models to agentic systems – autonomous, goal-driven entities that plan, reason, and act – where the concept of orchestration (CPU) and memory becomes central to intelligence itself. It is increasingly shifting towards control of the full stack of the AI factory – compute, networking, inference, orchestration and software, and no longer just raw compute power (GPU). LLMs (Large Language Models) excel in language generation, but they are passive tools requiring human prompting that lack persistent awareness of prior actions, user history, or environmental interpretation. Agentic AI changes this paradigm, enabling autonomous orchestration, multi-agent collaboration, and contextual recall. It introduces memory systems that persist across time, contexts, and agents, allowing continuity, learning, and adaptation.

Building AI agents is incrementally more CPU and memory heavy. While the AI hardware remains GPU-first dominant, co-existence with CPU becomes critical in an agentic AI world centered on reasoning, decision-making, and autonomy. Agentic AI represents the next frontier in enterprise automation, a shift from processing static data at scale to executing dynamic, multi-step tasks. These increasing workloads have become bottlenecks to today’s CPU architectures, driving a rethink for enterprise workloads and cloud architectures in balancing GPU investments (training/inference), with CPUs specifically engineered for reasoning tasks.

Understanding the agent AI stack

Modern AI breakthroughs are not just about algorithms but the incredible compute infrastructure built to train these models. Understanding this stack is critical for awareness of how large-scale AI actually works. Agentic AI consists of three distinct stacks that are interdependent layers: 1) the brain (LLM), 2) the system (orchestration) and 3) the knowledge (memory).

- **The Brain (LLM) is the GPU** that handles language understanding, generation, and complex inference. It processes information and action. LLMs use the memory layer to retrieve relevant context, and write new observations and conclusions back to the memory layer for long-term storage.
- **The System (orchestration) is the CPU** that acts as the control center managing the agent’s workflow, defines the sequence of actions (e.g., Plan -> Search -> Retrieve -> Reason -> Act), and connects the LLM to external tools and APIs. It

manages the flow of memory before feeding it to the LLM or CPU.

- **The Knowledge base is the memory**, an external brain of the agent, responsible for storing, organizing, prioritizing, and retrieving all long-term knowledge, including user profiles, past interactions, learned facts, and internal state. It is independent of the LLM and system.

The path to more capable, autonomous agents requires a fundamental shift in how we architect the AI stack, where a dedicated memory layer becomes the foundation that the next generation of agents will build on.

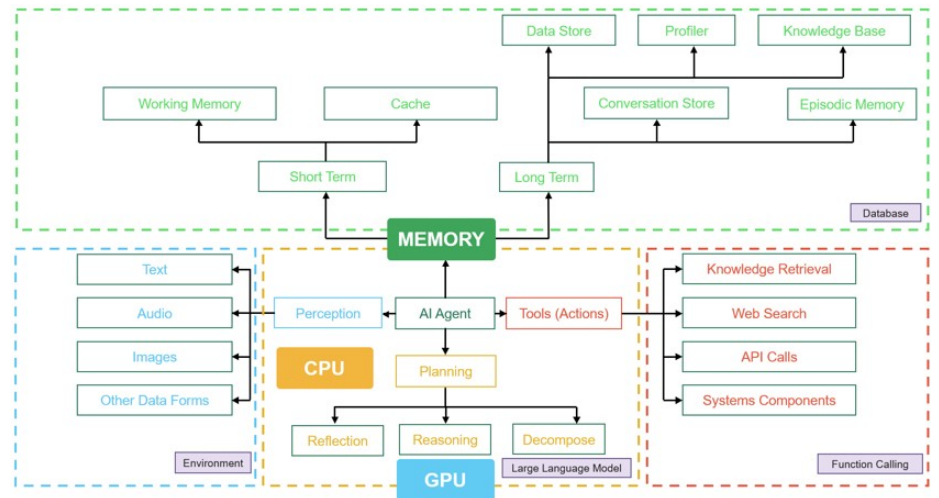
Exhibit 8: Agentic AI pillars



Source: Morgan Stanley Research

Orchestration is the new bottleneck

The entire data center is now treated as a single compute system. The CPU bottlenecks account for most (up to 90.6%) of execution time in agentic AI workloads, and establishes the necessity for holistic CPU-GPU co-optimization approaches. This is pure inference vs agent orchestration, where the CPUs manage multi-step tasks while GPUs focus on compute-intensive inference. GPUs handle token generation, CPUs are increasingly vital for tool invocation (API calls), data cleaning, and managing long context windows. As AI moves into production – especially with agentic systems – the CPU is becoming the control plane and no longer acting as a supporting component for the GPU. AI becomes a coordinated, real-time engine, not just a training problem. The CPU is coordinating thousands of concurrent processes, managing real-time inference and tool use and keeping massive GPU clusters fully utilized.

Exhibit 9: AI agents – CPU orchestration, memory and GPU execution

Source: Mongo DB

Future agentic AI roadmap

The birth of the cognitive anticipating AI agent. As agents become more autonomous, their memory must be elevated from a mere feature to a dedicated, foundational layer, separate from both the Large Language Model (LLM) and the Orchestration Framework. For memory, the current three-layer architecture for agentic AI – memorize, retrieve, evolve – adds an additional fourth 'Intention' layer that allows passive recall of memory to implement proactive AI, which predicts answers based on accumulated memories, proactively makes suggestions and prepares answers before any query.

The future of AI computing belongs to systems where CPUs and GPUs work as one brain. NVIDIA's GTC 2026 updated roadmap is a tightly choreographed system spanning GPUs, CPUs, LPUs, networking, storage, optics, and rack-scale design. Vera Rubin launch is 2H26, Rubin Ultra and the Kyber NVL144 platform from 2H27, and the Feynman generation with the new Rosa CPU and LP40 LPU in 2028. Future platforms are heavily reliant on co-design across compute, memory, networking, security, and interconnects. This roadmap is not just about future AI chips, but also about how the physical AI factory itself will need to evolve.


Optimization and efficiency. There is an increased focus on the capabilities of CPU hardware (like Arm) to handle AI workloads at the edge. Using CPUs for pre-processing keeps GPUs fully utilized on compute tasks, reducing idle time. It is now a balancing act for enterprise customers using GPU investments for training and inference against new CPU requirements for agentic workloads. The total cost of ownership calculations just got a lot more complex.

Competitive dynamics

The ground work is being laid for the next iteration of AI. NVIDIA is preparing and positioning itself exactly for that shift, with each piece of hardware and design having its strengths and weaknesses, depending on what it is exactly being used for, and with none being a one-size-fits-all. NVIDIA unveiled CPU processors specialized for agentic AI at the GTC conference, representing an addition to its pure GPU dominance. Both NVIDIA and AMD have recently expressed unprecedented demand for CPUs as AI workloads shift toward reasoning-intensive agentic applications.

Changing CPU-to-GPU ratio. Earlier Hopper AI server builds were often described in terms of one CPU servicing around 12 GPUs. However, more recent estimates for NVIDIA's Rubin place that closer to one CPU for every 2 GPUs, and some projections for Rubin Ultra go further toward 2 CPU for every 1 GPUs. Even the move from 1:12 to 1:8 represents a major increase in processor demand when applied across hyperscale deployments. That kind of volume increase can tighten the market quickly, especially when all major CPU suppliers are competing for the same high-value data center opportunities.

Exhibit 10: Cluster-level CPU:GPU intensity rises as AI moves from reasoning to action



Source: Morgan Stanley Research

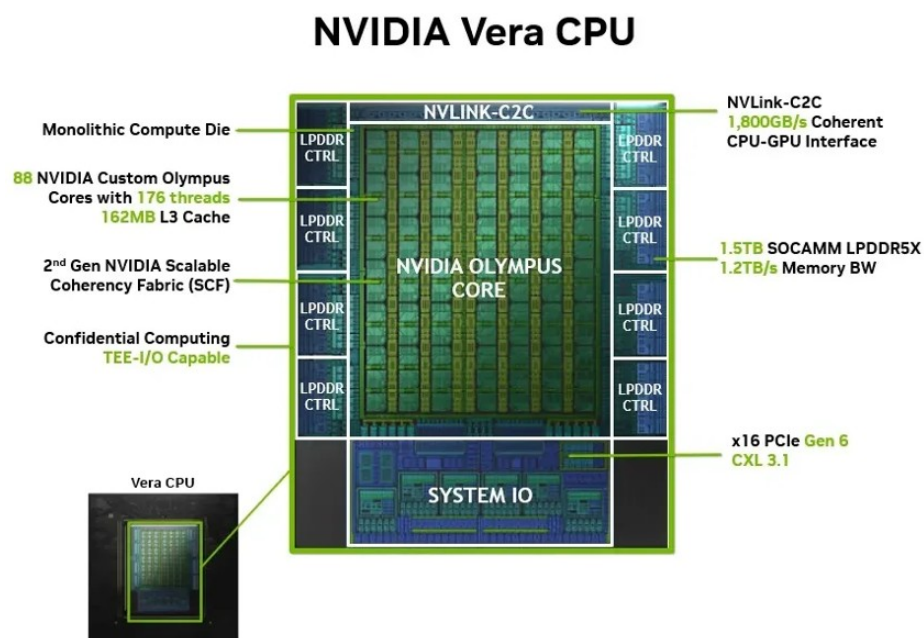
NVIDIA's and Arm's move into specialized CPU offerings brings them into more direct competition in the processor market with established players like **Intel** and **AMD**. Designing competitive CPU architecture demands different expertise, manufacturing processes, and go-to-market strategies than GPUs. The architectural requirements for agentic AI CPUs are different from traditional server CPUs that were not designed with these workloads in mind. These systems need processors that can handle long context windows efficiently, maintain state across extended reasoning chains, and switch rapidly between different types of computational tasks.

- **Vera is an AI CPU – built for agentic and reinforcement learning workloads.** It is the control and orchestration plane for GPU clusters and, as agentic AI scales, that role becomes extremely complex and requires significant use. Within the Vera Rubin NVL72 platform, Vera CPUs are paired with GPUs via NVLink-C2C and tightly integrated into NVIDIA's Rubin GPUs, which enable faster data movement between compute layers – which is critical for workloads that dynamically shift between CPU orchestration and GPU execution. A single rack with 256 Vera CPUs can now run 22,500 AI agents at once. NVIDIA is entering the standalone CPU market for the first time this year and has announced Meta and CoreWeave as customers. Over time, the company expects standalone CPUs to grow to a multi-billion dollar business.
- **AMD, widening the gap vs the competition with Venice.** AMD is the market-share leader in data center CPUs, due in part to the ecosystem advantages of x86 and the per-core performance advantage versus Intel. That is something we expect to

extend further with 2nm Venice, with that product adding new AI data type and AI pipeline support and up to 256 cores (vs 192 with Turin). As it will take until 2028-2029 for Coral Rapids from Intel, and Arm remains early on its journey, we see AMD in an increasingly strong position to capitalize on the higher growth vector for the market.

- **Arm has just launched its first production data center CPU.** A major strategic shift for Arm, moving from enabler to competitor, is a bold step. The Meta partnership clearly de-risks the launch and signals strong hyperscaler demand for more vertically integrated, AI-optimized compute platforms. Their own chip, their own server rack, and claiming TWICE the performance per rack versus x86 platforms from Intel and AMD.
- **Intel's roadmap remains focused on incremental improvements to x86, but with increasing emphasis on AI-relevant features.** Granite Rapids and Diamond Rapids aim to enhance memory bandwidth through more channels and MRDIMM support, while integrating AI acceleration features such as AMX. However, Intel's approach remains largely evolutionary, centered on improving the existing roadmap for general-purpose CPUs rather than redesigning for agentic workloads. We would expect to hear more about how Intel is closing the gap with AMD on performance with Coral Rapids as SMT reenters the roadmap. This may limit its competitiveness in certain highly specialized AI systems until then, although its ecosystem and installed base remain significant advantages. Collaboration with NVIDIA on NVLink compatible CPUs will be an important driver of Intel's opportunity in head nodes.
- **Qualcomm is an emerging participant, leveraging its custom Arm-based Oryon cores to target power-efficient AI compute.** Its roadmap builds on mobile and edge AI experience, with strengths in performance-per-watt and custom CPU design. However, Qualcomm remains early in the data center market, with limited ecosystem and deployment scale compared to incumbents. Its role in agentic AI infrastructure is likely to develop over a longer time horizon as Arm-based adoption expands.

Exhibit 11: Vera CPU chip architecture features 88 custom, Arm-based “Olympus” cores, with each core able to run two tasks using NVIDIA Spatial Multithreading, LPDDR5X memory delivering up to 1.2 TB/s of bandwidth, a second-generation Scalable Coherency Fabric for multi-tenant performance



Source: <https://www.servethehome.com/nvidias-vera-cpu-in-detail-high-perf-chip-takes-aim-at-broader-ai-server-market/nvidia-vera-cpu/>

Stock implications


Within the CPU supply chain, we are more positive on memory, foundry, substrates, CPU & memory interface and MLCC/CPU sockets, which all benefit from content growth, pricing power and face significant capacity constraints going into 2H26 and 2027, per our channel checks. Below, we highlight our covered names globally with exposure to these areas, and discuss these in greater detail in the [Global Perspective](#) section.

- **US Semiconductors** . Agentic AI-driven CPU demand structurally favors **AMD** in cloud share gains, but we prefer exposure via AI enablers (**NVIDIA, Micron**) where token growth and capex translate more directly into earnings upside.
- **US Semiconductor Equipment** . Rising compute and CPU TAM drive incremental WFE demand, with DRAM and leading-edge logic (<5nm/2nm) capacity expansion supporting upside for equipment names such as **AMAT** and **KLA**.
- **US Hardware** . AI agent proliferation is a structural tailwind for HDD demand, as persistent storage is required to capture growing volumes of context, history, and system-level data, with ~80% of cloud data still residing on disks. We expect HDD demand to remain strong (~30% EB CAGR), with supply tightness supporting a “stronger for longer” pricing backdrop, benefiting **Seagate** and **Western Digital**.
- **European Semiconductors** . **Arm** is well positioned to capture incremental CPU TAM from agentic AI, leveraging its power-efficient ISA and new AGI CPU design, while broader European semi-cap (**ASML**), EDA (**Synopsys, Cadence**) and analog (**STMicro**) benefit from rising compute intensity and design complexity.
- **Korea Technology** . While we continue to see **Samsung** and **SK hynix** as key

beneficiaries of memory content growth from AI CPU penetration, **Samsung Electro-Mechanics** is another way to gain CPU exposure via its increasing pricing power in ABF substrates.

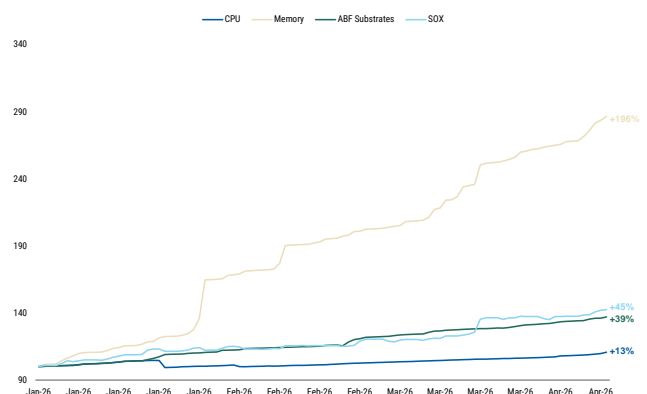
- **Greater China Semiconductors** . Rising CPU TAM is increasingly captured by **TSMC** (~70-75% share), supported by leading-edge node migration (2nm/3nm) across AMD, NVIDIA, and potentially Intel. Strong demand for Arm-based CPUs should drive incremental upside for **GUC**, with design wins across Google and Microsoft platforms. **Egis** is positioning into AI/HPC design services via Arm-based CPU and chiplet development, while **Aspeed** stands to benefit from higher CPU server deployments through increased BMC attach rates and ASP uplift. **Montage** is leveraged to the structural increase in memory content per server, with higher CPU and DRAM intensity driving demand for memory interconnect solutions.
- **Greater China Hardware** . Higher CPU demand benefits the server supply chain, including ABF substrates (**Unimicron**), PCBs (**GCE**), MLCCs (**Yageo**), and CPU sockets/connectors (**FIT**, **Lotes**). On the system side, **Wiwynn** and **Lenovo** gain from stronger general server demand driven by hyperscalers. Besides, agentic AI drives more power and advanced semi cap equipment demand, which helps on the power interconnect and subsystem assembly business for **Bizlink** and the factory automation component, e.g., linear guides, ball screws, pneumatic components, from both **Hiwin** and **AirTAC**.
- **Japan Semiconductors** . Increasing CPU complexity and memory scaling drive demand for advanced semi-cap (**Tokyo Electron**, **Ulvac**) and memory interface solutions, positioning **Renesas** as a beneficiary of higher DRAM bandwidth/capacity requirements. Agentic AI also supports NAND demand (**Kioxia**) and broader data center exposure across Japan semi names. Rising AI-driven packaging complexity is creating incremental upside for **Nittobo** and **MEC**, as higher ABF substrate demand, larger package sizes, and increased layer counts drive structurally tighter supply for low-CTE materials and adhesion solutions.
- **Japan Hardware** . Agentic AI drives demand for reliable power delivery and high-performance components, benefiting MLCCs, inductors, BBUs, and ABF substrates (**Murata**, **TDK**, **Ibiden**). Increased CPU/GPU intensity also supports storage (NL HDD) and next-gen packaging with embedded components.

Exhibit 12: Memory and ABF are leading the share price performance YTD



Source: Factset, Morgan Stanley Research

Exhibit 13: Memory stocks are leading the EPS earnings upward revision YTD



Source: Factset, Morgan Stanley Research

Exhibit 14:

AI agent beneficiaries

Ticker	Rating	Last Close (LC)	Company	Reason
CPU				
INTC.O	EW	68.50	Intel	Incremental CPU demand
NVDA.O	OW	198.35	Nvidia	Well-positioned AI compute solution provider + Vera CPUs
AMD.O	EW	278.26	AMD	Incremental CPU demand
ARM.O	EW	162.33	ARM	Incremental CPU demand
HDD				
WDC.O	OW	361.69	WDC	A rising AI tides lifts all boats for both HDD and eSSD
STX.O	OW	531.81	Seagate	A rising AI tides lifts all boats for both HDD and eSSD
NAND				
SNDK.O	OW	919.47	SanDisk	NAND Supercycle supported by agentic AI + inference demand
285A.T	OW	30,530.00	Kioxia	NAND Supercycle supported by agentic AI + inference demand
DRAM				
005930.KS	OW	216,000.00	Samsung Electronics	Better memory cycle driven by AI + HBM market share gain optionality
000660.KS	OW	1,128,000.00	SK hynix	Better memory cycle driven by AI
MU.O	OW	457.23	Micron	Better memory cycle driven by AI
ABF Substrates& Materials				
009150.KS	OW	679,000.00	SEMCO	Rising demand for high value-added ABF package substrates
3037.TW	OW	643.00	Unimicron	Better server CPU demand to support higher volumes and more favorable pricing of ABF substrates
8046.TW	OW	724.00	Nan Ya PCB	Better networking ABF substrate demand driven by higher general server volumes
4062.T	UW	9,235.00	Ibiden	Rising demand for high value-added ABF package substrates
3110.T	OW	27,510.00	Nittobo	Nittobo supplies low-CTE glass cloth used in semiconductor package substrates.
4971.T	OW	8,220.00	MEC	Larger ABF substrate sizes and a higher number of wiring layers bode well for MEC's CZ series adhesion promoters
Foundry				
2330.TW	OW	2,030.00	TSMC	Rising CPU TAM is increasingly captured by TSMC, supported by leading-edge node migration (2nm/3nm)
BMC/Memory Interface				
5274.TWO	OW	13,805.00	Aspeed	~70% of market shares in the CPU server BMC market
6809.HK	OW	209.40	Montage	Higher CPU and DRAM intensity driving demand for memory interconnect solutions
6723.T	OW	2,795.00	Renesas	A beneficiary of the agentic AI shift through its exposure to CPU- and DRAM facing memory interface
3702.TW	OW	98.80	WPG Holdings	Outgrowth from non-GPU components to drive distributors' ongoing outperformance
6531.TW	OW	644.00	AP Memory	WoW packaging technology could be an effective solution to improve memory bandwidth and power consumption
IP/Design Services				
3443.TW	OW	3,255.00	GUC	Strong demand for Arm-based CPUs drives incremental upside
6462.TWO	EW	122.50	Egis	Joining ARM's total design alliance with the aim of winning long-term ASIC business opportunities
Components/ODM				
6981.T	OW	4,600.00	Murata	Rising demand for highly durable and reliable electronic components/MLCC
6762.T	OW	2,616.00	TDK	Rising demand for highly durable and reliable electronic components/MLCC
6088.HK	OW	7.93	FIT Hon Teng	Clear beneficiaries of more CPU socket demand
2327.TW	OW	316.50	Yageo	Benefitting from increasing penetration of its products for both AI and general servers
3533.TW	EW	2,490.00	Lotes	CPU socket, PCIe/DRAM connector beneficiary
2368.TW	OW	1,175.00	Gold Circuit Electronics	Beneficiary of higher general server volumes and CPU MB PCB content growth
002916.SZ	EW	283.06	Shennan Circuits	Benefit from increasing server, networking and optical module PCB demand
600183.SS	EW	69.25	Shengyi	Benefit from increasing server, networking and optical module PCB/CCL demand
2317.TW	OW	206.00	Hon Hai Precision	Key ODM partner for brand/enterprise CPU and general servers
6669.TW	OW	3,805.00	Wiwynn	Key general server ODM
0992.HK	EW	11.30	Lenovo	General server supplier for select CSPs (ie. MSFT, ORCL)
000977.SZ	EW	68.23	Inspur	Clear ODM beneficiary of more AI and general server demand
SPE/Automation Equipment				
ASML.AS	OW	1,222.60	ASML	ASML benefits from increased EUV layer count
ASMI.AS	OW	768.00	ASMi	Beneficiary of increased advanced foundry node capacity.
BESI.AS	OW	220.50	Besi	Increasing accelerator demand drives greater adoption of advanced packaging technologies
VACN.S	EW	558.00	VAT Group	Positioned to benefit from strong WFE growth
240810.KQ	OW	121,200.00	Wonik IPS	Strong rebound on the back of Samsung 1c DRAM capex ramp up
AMAT.O	OW	389.90	Applied Materials	The most leverage to greenfield DRAM SPE beneficiary under US SPE coverage
KLAC.O	OW	1,734.85	KLA Corp	Rising compute and CPU TAM drive incremental WFE demand
8035.T	OW	44,010.00	Tokyo Electron	A potential incremental share gainer on Intel's capex recovery
6925.T	OW	3,153.00	Ulvac	MHM tool has seen solid adoption at Intel
3665.TW	OW	2,365.00	Bizlink	Increasing semi cap equipment demand positive to its power interconnect/subsystem business
2049.TW	OW	291.00	Hiwin	Increasing semi cap equipment demand positive to its industrial automation component business
1590.TW	OW	1,255.00	AirTAC	Increasing semi cap equipment demand positive to its industrial automation component business
Thermal/Power solution				
3017.TW	OW	2,400.00	AVC	Rising CPU/general server demand drives greater demand for advanced cooling solutions
3324.TWO	EW	1,035.00	Auras	Rising CPU/general server demand drives greater demand for advanced cooling solutions
2308.TW	OW	1,840.00	Delta	Elevated power consumption at data center level supports stronger growth for power supply players
2301.TW	EW	162.00	Lite-On	Elevated power consumption at data center level supports stronger growth for power supply players
EDA				
SNPS.O	EW	441.15	Synopsys	Benefiting from rising compute intensity and design complexity
CDNS.O	OW	306.96	Cadence	Benefiting from rising compute intensity and design complexity
Analog				
IFXGn.DE	OW	46.01	Infineon	Potential upside to outer year data centre estimates from the development of agentic systems with multi-step workflows
STMPA.PA	OW	34.95	STMicroelectronics	Potential upside to outer year data centre estimates from the development of agentic systems with multi-step workflows

Source: Morgan Stanley Research. Note: Priced at close on 16 April 2026.

Global Perspective

The rise of agentic AI is a long-duration infrastructure theme – and a global one. While deployment is led by US hyperscalers, value creation is globally distributed across CPUs, memory, substrates, foundries and equipment, with Asia critical on the supply side and Europe well positioned in industrial and infrastructure enablers. In this section, we highlight stock exposures to the broad agentic AI theme across regions and sectors.

US Semiconductors

Joseph Moore, Mason Wayne, Ella Tulchinsky, Nicole Kozhukhov


AMD in pole position to benefit from agent-driven CPU workload growth in the cloud.

As AI resets the growth in the server market higher, that should lead to outsized benefits for AMD. Having surpassed Intel in x86 cloud share, the incremental business is mostly AMD's to lose, we believe, as Intel struggles to deliver product against persistent supply constraints and a weak product line-up. Intel should see positive tailwinds as long as the market remains undersupplied, but, when it comes to capturing secular durable tailwinds from increasing demands on the CPU in the cloud, we think AMD is much better placed.

However, our preference for exposure to the growth of agentic is for AI leaders priced for a slowdown in cloud spending – NVIDIA, Broadcom, Micron, SanDisk.


While CPU is the cleanest story for exposure to new agentic AI workloads, we argue both Intel (foundry) and AMD (GPU) have more important elements of stock performance tied to other areas, and we are Equal-weight on both. We prefer exposure to agentic AI through the key enablers of the AI workloads, as more agents still means more tokens and more tokens means more processor demand. As excitement emerges around agentic AI workloads, hyperscaler AI capex is likely to continue to move higher, and we prefer exposure to this theme through NVIDIA and memory plays at, respectively, 18x FY27e P/E and 5-9x P/E, rather than through CPU with Intel and AMD at 23-64x P/E.

Exhibit 15: AMD has a strong lead in cloud CPU share



Source: IDC, Mercury Research, Morgan Stanley Research

Exhibit 16: However, Intel remains the standard in other markets



Source: IDC, Mercury Research, Morgan Stanley Research

US Semiconductor Equipment

Shane Brett

More compute equals more WFE. As we reach a point where the binding constraint to intelligence scaling is semiconductor production capacity, WFE should continue to see material upward revisions across DRAM/NAND and Logic. In the near term, DRAM spending is inflecting the most, driving our preference for **AMAT**.

We expect benefits of agentic AI to be widespread, but a larger CPU TAM in particular may drive upside to current expectations for 2027 leading logic. Our current 2026 forecast calls for 23% growth in WFE spending to \$14.3bn, with growth led by DRAM and foundry logic, and for 27% growth next year, with foundry logic WFE accelerating from 15% growth in 2026 to 28%. Growth in the CPU TAM likely translates to more aggressive capacity expansion at <5nm nodes, and at Intel and TSMC 2nm in particular. We have highlighted both of these areas as representing unique tailwinds for **KLA**.

US Hardware

Erik Woodring, Dylan Liu, Maya C Neuman, Rauf Ural

AI agent proliferation is a structural tailwind for HDD demand. Knowledge – one of the three core pillars of agentic AI – relies on persistent data storage to capture prior actions, user history, and environmental context. While this dynamic benefits memory (flash), it also supports HDD demand, given ~80% of cloud data continues to reside on disks. Although it remains uncertain how materially AI agents will change the long-term HDD exabyte (EB) growth trajectory, emerging applications, including AI agents, autonomous driving, and production optimization, are increasingly incremental drivers of HDD demand.

We reiterate our “stronger for longer” HDD thesis, with both STX and WDC positioned to benefit. Our recent checks suggest the HDD industry will likely deliver low- to mid-single-digit higher unit output than we expected entering the year; however, this upside has been more than absorbed by accelerating hyperscaler demand. As a result, we still expect industry shortages of ~200EB in CY26 (~10% of the market), widening to ~250EB in CY27. Greater automation via AI agents increases compute intensity, expands data generation, and ultimately drives incremental storage demand. We therefore expect HDD EB CAGR to track at least the long-term data generation CAGR of ~30% over the next several years, with the supply-demand imbalance potentially persisting longer than anticipated if emerging applications prove incremental. In our view, this backdrop remains supportive of HDD pricing under a rational oligopoly industry backdrop.

European Semiconductors

Lee Simpson, Nigel van Putten, Amelia Scicluna

Arm's AGI CPU in prime position for this opportunity. We see Arm as a primary beneficiary of a large new CPU market, as exemplified by its recent new chip design launch Arm AGI CPU (see our notes [here](#) and [here](#)). This 3nm, TSMC fabbed chip focuses on power efficiency, a single-threaded compute, as well as a performance design to handle most agentic tasks and scalability. Arm's RISC based instruction set (ISA), and its years of power-efficient CPU core design, make it well positioned to directly benefit from a large

new market. We think Arm's shift to a fabless chip designer is well timed and we see the strategic rationale as solid.

European semi-cap is positioned to benefit from a larger CPU TAM, as increased compute intensity drives higher layer counts and rising wafer fab equipment demand at advanced nodes. Within this, ASML remains the key bottleneck, as incremental CPU-driven demand translates into lithography intensity. We highlight potential for this to drive marginal upside to our outer year tool estimates. ASMI is also well positioned, given its exposure to advanced deposition steps that scale with increasing device complexity and are critical at leading-edge nodes. In addition, we see VAT Group as a key beneficiary of higher process intensity and tool complexity, with vacuum valves acting as an enabling component across deposition and etch steps; however valuation keeps us Equal-weight. Beyond the front-end, Besi is a direct beneficiary of a growing CPU TAM, as increasing accelerator demand drives greater adoption of advanced packaging technologies. The shift towards chiplet-based architectures and heterogeneous integration is increasing the need for hybrid bonding solutions, supporting higher packaging intensity per chip.

Analog is an indirect beneficiary of increased compute demands. We see potential upside to our outer-year data center estimates from the development of agentic systems with multi-step workflows. A material rise in CPUs will, in turn, require a new power architecture. While we think it is too early to quantify this opportunity, this would suggest upside to our outer year data center revenue estimates for Infineon and STMicro.

We also see **potential incremental upside to EDA**, where increased CPU and accelerator complexity drives higher design intensity, particularly at advanced nodes and in chiplet-based architectures. We think this could suggest additional growth for Cadence and Synopsys through greater tool adoption in areas such as verification and emulation; though we are Equal-weight on Synopsys pending visibility on profitability from the integration of Ansys.

Korea Technology

Shawn Kim, Ryan Kim, Duan Liu, Cindy Huang

Doing AI things requires more DRAM vs. thinking AI things (HBM). The memory cycle is now increasingly defined by the rapid shift to agentic reasoning, with OpenClaw accelerating that process. In agentic AI workloads, the bottleneck is increasingly shifting to the CPU and memory hierarchy as agentic AI takes what the GPU generates to turn it into actions. Markets tend to think linearly but the intelligence layer (agents, reasoning) is growing exponentially, where memory becomes the most difficult AI chip bottleneck.

Compute requirements multiply and memory becomes an architectural priority for AI agent infrastructure. Memory is what gives us context as AI agents need persistent memory across conversations, heterogeneous compute for orchestration and inference, and low-latency networking for inter-agent communication. They rely on a tiered architecture consisting of: 1) transitory cache for short-term working memory; 2) hot storage for active episodes that benefit DRAM; and 3) cold storage for archives favoring HDD and NAND. The storage demand significantly exceeds that of generative AI.

We think 2027 will be the year of the proof case for sustainable memory returns.

What matters for stocks from here, in our view, is when LTA agreements start to be priced

in, durability of the cycle upturn and sustainability of FCF returns, HBM share and where sustainable margins land vs. the near-term commodity memory pricing (earnings) outlook in the next 6 months – both likely higher, per our checks. For ABF substrates, we see price hikes throughout the year and think it is no longer about competition, but about how much companies can make and sell. We favor higher pricing power in DDR5 DRAM and NAND flash (**Samsung, SK hynix**), ABF substrates (**Samsung Electro-Mechanics**) and capex benefits via SPE (**Wonik IPS**) vs. downstream hardware-facing margin pressure.

Greater China Semiconductors

Charlie Chan, Daniel Yen, Daisy Dai, Lucas Wang, Tiffany Yeh

TSMC is the major CPU foundry vendor. We forecast overall CPU TAM for foundry to reach nearly \$33bn in 2026 and then \$37bn by 2028. Within the TAM, we estimate TSMC, serving as the major foundry, could reach around 70% market share in 2026 and keep growing to 75% in 2028, with the nodes for those CPU also continuing to migrate, like GPU and ASIC. For example, the upcoming server CPU from AMD (Venice) and NVIDIA (Vera) are adopting TSMC leading-edge 2nm and 3nm process, both a one-node migration from prior generation. We also expect Intel to outsource its server CPU production to TSMC by 2H27, as it is facing fierce competition from AMD in the server/data center CPU space, whereas TSMC is more available to provide time-to-market of its server CPU production with higher quality.

Strong Arm-based CPU demand could benefit GUC from 2026 and beyond. Within our Greater China Semiconductor coverage, GUC has the highest revenue mix from the CPU demand. Our recent industry checks suggest that inference and general server demand for Arm-based CPUs is even stronger, including Google CPU. GUC also helps with the Cobalt 200 CPU for Microsoft, an Arm-based design tailored for Microsoft Cloud. GUC management indicated that it has secured next-gen Arm-based CPU design wins from its major customers, including both Google CPU and Microsoft Cobalt CPU.

Egis is positioning itself as an emerging AI/HPC design services vendor, with increasing exposure to CPU and ASIC opportunities over the medium term. We believe its strategic pivot away from its legacy fingerprint business is beginning to gain traction, supported by investments in InPsytech and a ~20% stake in Alcor Micro. These moves expand its capabilities into CPU platform development and heterogeneous computing. In particular, participation in Arm's Total Design Alliance and ongoing hiring suggest a clear ambition to capture long-term ASIC design opportunities tied to AI and HPC workloads. That said, we remain Equal-weight given limited near-term visibility on large-scale project wins, with commercialization timelines for new CPU platforms extending into 2027.

Aspeed remains a key beneficiary of rising CPU intensity in agentic AI through its dominant BMC franchise. As CPU demand increases with more orchestration-heavy workloads, the need for server management and control functionality should expand proportionally, directly benefiting Aspeed's BMC portfolio. With ~70% market share in CPU server BMCs and a strong customer base, including major hyperscalers such as Meta, AWS, Alibaba and Microsoft, Aspeed is well positioned to capture this growth. We expect further upside from its next-generation AST2700 platform, which not only supports share gains at additional customers such as Dell and Google, but also drives 40-50% ASP uplift through specification upgrades.

Montage Technology is leveraged to the CPU and memory content upcycle through its leadership in memory interconnect chips. As agentic AI drives higher DRAM content per server and accelerates technology migration, demand for memory interface solutions should increase meaningfully. Montage, with a ~36.8% global market share, remains a key beneficiary of this trend, supported by its technology leadership and strong positioning with cloud customers. We expect continued outperformance driven by both cloud capex growth (c.30% CAGR over 2025-27) and rising memory complexity, which structurally increases interconnect content per server.

Greater China Hardware

Sharon Shih, Howard Kao, Derrick Yang

We believe **Unimicron** is a key beneficiary of the increases in CPU demand from agentic AI, as it is a key supplier to Intel, AMD, NVIDIA for their data center CPU substrates. **NYP** should also benefit, as more servers will require more networking demand as well, and **NYP** is a key ABF substrate supplier for Broadcom. **GCE** is another clear beneficiary of higher CPU server demand, as it produces server mainboard PCBs for most major ODMs, with ~70% of its revenues coming from servers (including ASIC servers).

Within our downstream component coverage, **FIT** and **Lotes** stand out to us as CPU beneficiaries as well, because both are major CPU socket suppliers to Intel and AMD with dominant supply shares, and will also have peripheral exposure from PCIe and DRAM connectors; however, Lotes' current valuation does not look sufficiently attractive to us, so we are Equal-weight.

We also believe **Yageo** would be a beneficiary of the increased passive components TAM, including MLCCs, resistors, inductors and tantalum capacitors. As a total solution supplier with all of these passive components, it is one of the key passive component beneficiaries in our Greater China coverage.

Within our ODM/OEM coverage, we believe **Wiwynn** benefits the most from rising CPU/general server demand, as it is a key general server ODM for Microsoft, Meta and Amazon. **Lenovo** should also benefit from higher general server demand, because it is a key supplier to Microsoft and Oracle; however, we are Equal-weight on Lenovo because of the cost headwinds from its PC business.

For hardware components, we think **Bizlink** should benefit from the agentic AI proliferation, as the increasing demand for more advanced semiconductor capital equipment will drive more demand for its power interconnects and subsystem assemblies, which it supplies to major customers like Lam Research, KLA, Applied Materials, ASML, etc. Moreover, better semiconductor capital equipment demand is also positive for the factory automation components suppliers like **Hiwin** and **AirTAC**, as their offerings, including linear guides, ball screws, robotic systems, pneumatic components, etc., are widely adopted in various kinds of equipment.

Thermal solution suppliers like **AVC** and **Auras** should benefit from greater demand for advanced cooling solutions (e.g., heat pipes, vapor chambers, and liquid cooling components) for rising CPU/general server demand. However, we are Equal-weight on Auras for unclear VR platform exposure at this point. At the same time, elevated power consumption at data center level supports stronger growth for power supply players such

as **Delta** and **Lite-On**, as they are also dominant suppliers for general servers to top hyperscalers. We are Equal-weight on Lite-on given headwinds to its legacy PC/IT business.

Japan Semiconductors

Kazuo Yoshikawa, Suzune Tamura, Yoshihito Hasegawa

Japan semi-cap stands to benefit further from rising CPU demand. Increasing device complexity should drive incremental demand for advanced process equipment. From a customer mix perspective, while Tokyo Electron is commonly viewed as a memory play, it has relatively high share in logic and within Intel (which used to be mid-teen customer for TEL 3 years ago). A potential recovery in Intel's capex, supported by stronger CPU demand, would position TEL as a potential incremental share gainer. In addition, Ulvac's MHM tool has seen solid adoption at Intel, which should provide an incremental positive for Ulvac's earnings outlook, although relative to peers we see lower medium-term growth prospects and stay Equal-weight.

Renesas would be positioned as a beneficiary of the agentic AI shift through its exposure to CPU- and DRAM-facing memory interface. Memory interface devices such as register-clock drivers (RCDs) and data buffers enable the entire memory subsystem to scale in speed and capacity, by allowing the clock, command, address, and data signals to be re-driven with much-improved signal integrity. As agentic AI requires higher DRAM content per server and technology migration of DRAM accelerates, demand for memory interface devices should increase meaningfully.

We estimate that data center sales accounted for approximately 15% of CY25 Q4 revenue, of which 50–60% was attributable to memory interface sales. We also estimate that the gross margin of the memory interface business is well above the overall gross margin of 61% for the Industrial, Infrastructure, and IoT segment. As a result, growth in memory interface sales should also contribute to an improvement in company-wide profitability.

We also expect KIOXIA to benefit from the agentic AI shift. The databases and documents referenced by agentic AI, such as those used in RAG (retrieval-augmented generation), require large density, high performance SSDs. The company has developed SSDs incorporating its 8th-generation BiCS FLASH 2Tb QLC chips and shipped the 122TB and 245TB models for customer qualification as planned at the end of 2025. It aims to begin volume shipments in 2026.

We believe Nittobo and MEC are well positioned to benefit from the broader adoption of agentic AI. Nittobo supplies low-CTE glass cloth used in semiconductor package substrates. The emergence of agentic AI as a new end-market should drive incremental demand for advanced semiconductor package substrates. While new entrants have recently increased capacity in the low-CTE glass cloth segment, we expect supply-demand conditions to remain tight, as incremental demand from AI-related applications adds a new layer of structural growth. At the same time, the trend toward larger ABF substrate sizes and a higher number of wiring layers should serve as a demand tailwind for MEC's CZ series adhesion promoters. The CZ series chemically modifies the surface of copper wiring layers to enhance adhesion with insulating resin layers. Given its critical role in ensuring the reliability of substrate materials that combine metals and resins with differing coefficients of thermal expansion, the CZ series has established a dominant market

position and remains an indispensable material in next-gen package substrates.

Japan Hardware

Shoji Sato

We expect the agentic shift to lift the earnings of major Japanese electronic component manufacturers, as it will drive increased demand for highly durable and reliable electronic components. It requires stable and reliable power delivery to CPUs and GPUs, which we believe will lead to a significant increase in demand for:

- MLCCs and power inductors capable of maintaining stable performance even under high-temperature conditions;
- Vertically powered supply systems; and
- LiB for battery backup units (BBUs).
- In addition, we expect continued growth in demand for NL HDDs, which provide a low-cost and stable data storage solution.

For ABF package substrates, we anticipate rising demand not only due to larger substrate sizes and increased layer counts, but also for package substrates that embed MLCCs and inductors within the substrate itself to ensure stable power delivery to CPUs and GPUs.

We expect **Murata Manufacturing** (6981.T) to further expand its already stable and high earnings base, driven by growth in sales of: (1) high-capacitance, highly durable MLCCs – where the company holds the world's leading market share – as well as highly reliable power inductors; (2) vertically integrated power supply systems, for which mass-production shipments are scheduled to begin in FY3/27; and (3) LiB for BBUs.

We also expect **TDK** to benefit in the medium to long term from increased sales of (1) power inductors, where the company holds the world's leading market share; (2) LiB for BBUs; and (3) HDD head and suspension business.

We expect **Ibiden** to continue expanding its earnings, benefiting from rising demand for high value-added ABF package substrates. It appears to be collaborating with Murata Manufacturing and other partners on the development and mass production of ABF package substrates that embed MLCCs and inductors. By integrating competitive technologies across ABF package substrates, MLCCs and inductors, we believe Ibiden will be able to stably supply highly competitive products. That said, we see a risk that near-term performance may fall short of overly high market expectations – due mainly to the constraint of material procurement such as T-Glass and CCL – and that market consensus will fall, so we are Underweight.

Growing the CPU Opportunity

Agentic AI adds multi-step orchestration – plan, retrieve, call tools, execute, iterate – lifting cluster CPU:GPU intensity and shifting incremental spend toward CPUs. Our analysis suggests \$32.5-60bn incremental CPU TAM by 2030 within \$100bn+ total server CPU TAM.

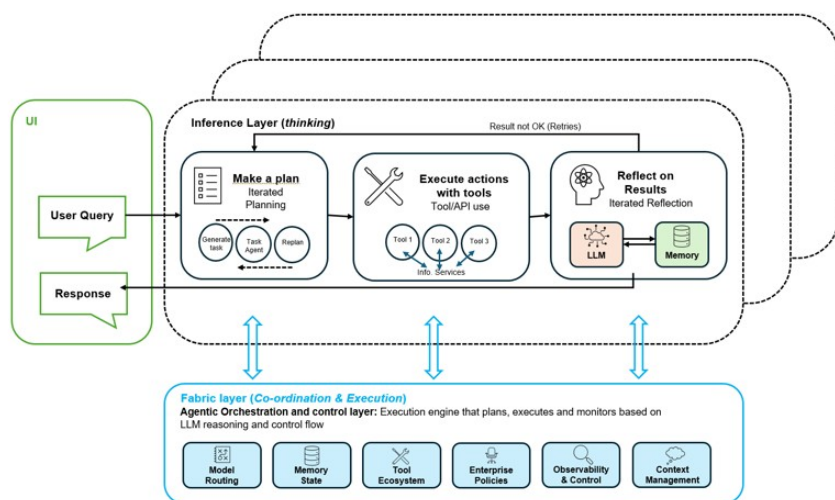
We believe agentic AI will increase the CPU-to-GPU mix in AI systems by adding more orchestration, memory, and tool-use work around each model call. This should not reduce GPU demand, but it does increase overall system complexity and shifts incremental infrastructure spend toward CPUs, networking, and memory. In this environment, the advantage is less about owning the accelerator and more about owning the system architecture.

What changes with agentic AI

The initial GenAI wave was dominated by GPU-centric model-serving, with relatively light control-plane overhead. Agentic inference introduces multi-step workflows (plan, retrieve, call tools, execute, iterate), greater reliance on persistent memory and external tools/APIs, and higher orchestration complexity (see our most recent Arm note [here](#)).

Net: more work per user request, and higher CPU, memory, and network intensity than in the first GenAI phase.

Exhibit 17: The inference fabric framework



Source: Morgan Stanley Research

Unpacking CPU/GPU ratio dynamics

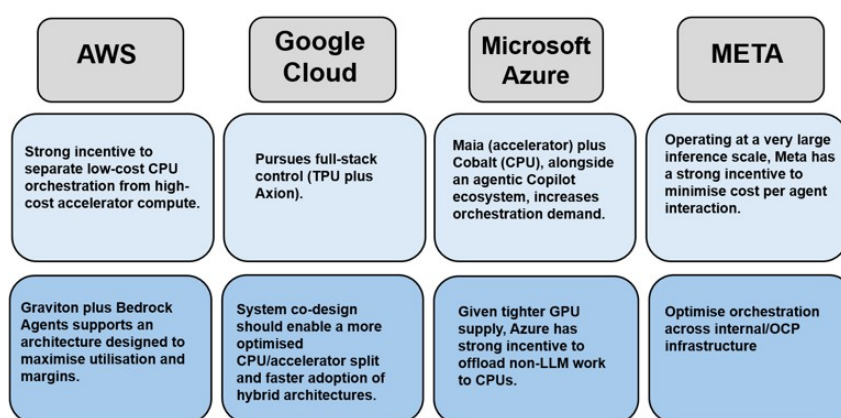
At the head-node level, we expect limited change. GPU servers should remain accelerator-dense, with companion CPUs (e.g., Grace, Vera) acting primarily as the host and orchestration layer.

At the cluster level, however, we expect additional CPU-heavy infrastructure. As agentic workloads scale, the effective CPU/GPU ratio should rise meaningfully.

Two drivers underpin this:

- Multi-step workflows. Agentic systems move beyond a single inference call. They add branching logic and latency-sensitive control flow (planning, routing, tool calls), which naturally runs on CPUs.
- Economics. As GPUs become more expensive, running orchestration on them is less efficient. The incentive is to keep GPUs focused on token generation and push non-LLM work to CPUs.

Exhibit 18: Strategic rationale for CPU deployments




Source: Morgan Stanley Research

CPU TAM: a \$100bn+ opportunity by 2030

We think agentic inference creates a new CPU TAM on top of traditional data center CPUs, driven by orchestration, tool use, and memory services. **We estimate this incremental TAM at c.\$32.5-60bn by 2030, within a total data center CPU market of \$100bn+ by 2030** (in line with recent Arm management remarks, [link](#)). In other words, the agentic-driven CPU opportunity could approach the size of the legacy CPU market.

We estimate the 2025 CPU server TAM as \$24bn using Mercury Research data. Mercury Research suggests a 2030 TAM of \$45.4bn; however, our top-down estimates of the market suggest that this does not consider agentic upside, as our estimate for the head node CPU TAM alone implies a \$50bn TAM.

Exhibit 19: The server CPU TAM has been about flat in both revenue and unit terms since 2018, but we see new agentic tailwinds causing durable growth from here



Source: Mercury Research, Morgan Stanley Research

As a top-down anchor, we use MS estimates for global AI data center infrastructure sales rising to ~\$1.2tn by 2030 (from ~\$242bn in 2025), spanning GPUs, ASICs and CPUs. The key shift is the CPU's role expanding from "host processor" to control plane, orchestration, and data-engine compute.

We break CPU TAM into three buckets: head-node CPUs, orchestration CPUs, and other supporting CPUs.

Head node CPUs

These CPUs are attached to GPU systems (rack control layer) and are typically tightly coupled to the accelerator platform. They host GPU clusters (e.g., NVL72 or TPU pods) and run scheduling, dispatch, and KV-cache coordination.

In deriving the head node CPU TAM, we think our estimate of \$1.2tn infrastructure sales by 2030 implies 5mn AI accelerators globally. Based on the assumption for two high-end CPUs per GPU, this suggests 10mn total CPUs. If these are next-gen CPUs (beyond Vera) and using 100 to 150 cores per CPU, we estimate an ASP of \$5k+/CPU. All in, this implies a c.\$50bn TAM by 2030.

Orchestration CPUs

This is the incremental TAM created by agentic workflows. It covers:

- Agentic runtimes (planning);
- Tools/API chaining (code execution, browser agents);
- Retrieval (and RAG pipelines);
- Memory services (KV-cache, vector DB); and
- Scheduling (policy, security, observability).

This category did not exist at meaningful scale before agentic AI.

We assume the number of CPUs to GPUs scales with agent complexity. We consider 2-3 additional CPU-heavy nodes per GPU. This implies 10-15 million CPUs of a higher core count than head node CPUs. The Arm AGI CPU is 136 cores, so we assume that, by 2030, core counts reach 200-300. Pricing is a mix of hyperscaler Arm CPUs (Graviton/Axion/Cobalt) and Merchant CPUs (AMD/Intel/Arm). This implies an ASP of \$3k by 2030. All in, applying the ASP to the total CPU units, this suggests an orchestration CPU TAM of \$30-45bn by 2030.

Other CPUs

This bucket includes CPUs used elsewhere in AI infrastructure, such as storage nodes (SSD-heavy clusters) and certain networking nodes.

Assuming 1-2 additional CPUs per GPU, this implies 5-10mn CPUs. As the ASP will be lower than the orchestration tier, we apply a range of \$500-1,500. This suggests a TAM of \$2.5-15bn.

Total CPU TAM

Our estimates for the three buckets – head node CPUs, orchestration CPUs and other CPUs – imply a CPU TAM \$82.5-110bn by 2030.

Exhibit 20: We estimate the CPU TAM by 2030 to reach \$82.5-110bn

	Lower Bound	Upper Bound
Head Node CPUs (million)		
AI Accelerators by 2030	5	5
CPUs per GPU board	2	2
Total CPUs	10	10
ASP (\$)	5,000	5,000
TAM (\$)	50,000	50,000
Orchestration CPUs (million)		
AI Accelerators by 2030	5	5
CPUs per GPU board	2	3
Total CPUs	10	15
ASP (\$)	3,000	3,000
TAM (\$)	30,000	45,000
Other CPUs (million)		
AI Accelerators by 2030	5	5
CPUs per GPU board	1	2
Total CPUs	5	10
ASP (\$)	500	1,500
TAM (\$)	2,500	15,000
Total CPU TAM by 2030	82,500	110,000

Source: Morgan Stanley Research estimates

We could be underestimating the size of this market. NVIDIA has spoken to AI data center infrastructure sales reaching \$3-5tn by 2030. If that is correct, this would imply significant upside to our CPU TAM estimates.

If we assume global AI data center infrastructure sales reach \$3tn by 2030, this would imply 12.5mn AI accelerators by 2030. In this instance, our lower-bound estimate for the CPU TAM is \$206.3bn and our upper-bound estimate \$275bn.

If we assume instead that global AI data center infrastructure sales reach \$5tn by 2030, this would imply 21mn AI accelerators by 2030. In this instance, our lower-bound estimate for the CPU TAM is \$344bn and our upper-bound estimate \$458bn.

Exhibit 21: Assuming global AI data center infrastructure sales reach \$3tn by 2030, we estimate the CPU TAM is \$206-275bn

	Lower Bound	Upper Bound
Head Node CPUs (million)		
AI Accelerators by 2030	12.5	12.5
CPUs per GPU board	2	2
Total CPUs	25	25
ASP (\$)	5,000	5,000
TAM (\$)	125,000	125,000
Orchestration CPUs (million)		
AI Accelerators by 2030	12.5	12.5
CPUs per GPU board	2	3
Total CPUs	25	37.5
ASP (\$)	3,000	3,000
TAM (\$)	75,000	112,500
Other CPUs (million)		
AI Accelerators by 2030	12.5	12.5
CPUs per GPU board	1	2
Total CPUs	12.5	25
ASP (\$)	500	1,500
TAM (\$)	6,250	37,500
Total CPU TAM by 2030	206,250	275,000

Source: Morgan Stanley Research estimates

Exhibit 22: Assuming global AI data center infrastructure sales reach \$5rn by 2030, we estimate the CPU TAM is \$344-458bn

	Lower Bound	Upper Bound
Head Node CPUs (million)		
AI Accelerators by 2030	21	21
CPUs per GPU board	2	2
Total CPUs	42	42
ASP (\$)	5,000	5,000
TAM (\$)	208,333	208,333
Orchestration CPUs (million)		
AI Accelerators by 2030	21	21
CPUs per GPU board	2	3
Total CPUs	42	63
ASP (\$)	3,000	3,000
TAM (\$)	125,000	187,500
Other CPUs (million)		
AI Accelerators by 2030	21	21
CPUs per GPU board	1	2
Total CPUs	21	42
ASP (\$)	500	1,500
TAM (\$)	10,417	62,500
Total CPU TAM by 2030	343,750	458,333

Source: Morgan Stanley Research estimates

Agentic AI Memory Opportunity

“The real breakthrough in AI isn’t better reasoning — it’s persistent memory,” Sam Altman, Co-Founder and CEO of OpenAI

Agentic AI shifts systems from reactive chatbots to proactive, multi-step agents, but durability hinges on memory, not models. Persistent memory enables continuity, reasoning and improvement across time. This drives a structural shift in infrastructure demand: we estimate 15-45EB of incremental DRAM demand by 2030, supported by 10-15m incremental orchestration CPUs with ~1.5-3TB DRAM per CPU, implying new demand vector equivalent to ~26-77% of the DRAM industry supply in 2027.

The future of agents is better reasoning with better memory. The agentic era is already here, but the memory era is just beginning. Turning AI from a reactive responder into a proactive agent cannot be done with clever prompts alone, and the key enabling technology is memory – and this time, memory is not an add-on but the layer that determines which agents actually last. Agentic AI transforms scattered information into connected intelligence enabling every interaction, or insight, to feed into a continuously learning ecosystem.

The competitive edge is shifting to context as foundational models begin to converge in performance. A chatbot is like a copilot for a single task, whereas an agent is more like an autopilot managing multi-step workflows with minimal oversight. To work effectively, these AI agents also rely on memory to carry context across sessions, authentication credentials to securely interact with external systems, and a persistent agent profile that allows them to remember users across time. AI agents rely on memory, preferences, documents, tone, and task history to perform effectively, and the more context an agent collects, the more personalized it becomes. Therefore, context, not model performance, is the true source of competitive edge, in our view.

Why does memory matter more in agentic AI?

Agentic AI doesn’t just do; it remembers, reasons, and improves. In agentic AI systems, memory is not about remembering a chat, but it is the mechanism that makes behavior consistent, enforceable, and economically scalable. An agent must actively manage its memory to ensure the most relevant facts are learned, which necessitates a system more nuanced than a simple history dump. Memory allows agents to behave with awareness, remembering facts, preferences, plans, past interactions, world state, and long-term learning outcomes. But just as humans selectively remember important details and let trivial ones fade, AI agents need clever strategies to remember what matters and forget what doesn’t. As such, memory infrastructure is a multi-layered retrieval and reasoning system rather than just storing/retrieving data.

The fundamental limitation. Language models have a fixed context window and once a conversation exceeds that window, older content falls off. However, this becomes a problem for a coding agent working on a large codebase over multiple sessions. A

production codebase may have thousands of files, complex inter-dependencies and decisions made weeks or months ago that still matter today. An agent that forgets all of that every time it starts a new session is nearly useless for serious work. There are other real-time concerns for agents:

- Remembering what it decided earlier in the conversation;
- Understanding the current state of files that have been modified; and
- Finding the relevant code without needing the user to point to every relevant file.

Memory allows the agent to access what has been happening over time, not just what is true about the world. Addressing all of these problems requires memory-tiered solutions to achieve a combination of persistent memory, short-term orientation and retrieval.

Agent multiplier. Agentic AI compounds knowledge over time, using knowledge from one agent and passing it on to others from general-purpose agents to specialized agents (knowledge no longer exists in silos). Multi-agentic systems use a series of agents, with a single coordinating agent, to work as a sort of AI team. The coordinating agent works to understand complex queries and delegate workflows to other agents, making multi-step, multi-system queries possible. In collaboration with people, who are essential for escalation, understanding significant ambiguity, and creative thinking, multi-agentic systems are becoming integral to digital-first workforces.

The fundamental solution – persistent memory. Memory cannot be an afterthought in agent systems, as the underlying issue isn't prompt quality (LLM) but the absence of a designed memory layer. In real-world applications, agents are expected to 1) run continuously, 2) serve multiple users, 3) coordinate across multiple agents and tasks, and 4) learn from past interactions without corrupting future behavior. In these environments, memory is no longer just context but becomes a long-term state, shared knowledge, and behavioral grounding. When a system can retain, understand and organize its past experiences, it gains continuity of thought – and continuity of thought gives rise to resolution. Without memory, there is no continuous thinking, no true forward-thinking. Giving AI a sense of self-determination means giving it a sense of time, situation awareness, better control of actions and their consequences.

Persistent memory improves output quality, which reduces redundancy as well as allows models to focus on its token budget, creating thoughtful production-ready UI instead of reconstructing context. Agents do not just consume memory but refine it. Memory shapes future behavior rather than simply persisting state. This is the transition from agents that execute tasks to agents that improve over time. In the future, the sophistication of memory systems will define the competitiveness of AI agents.

Agentic AI memory TAM

Agentic AI changes memory TAM in two ways: it raises the effective context per request and increases the amount of CPU-side memory provisioned per AI rack. As a result, the marginal memory demand shifts away from hot HBM alone and increasingly into host DRAM and rack SSD, where retained context, intermediate state and warm KV live closest to the CPU.

Using a top-down TAM model anchored in (i) incremental CPU shipments associated with

agentic AI, and (ii) higher DRAM per server for agentic roles, we estimate agentic AI could drive 15-45EB of incremental DRAM demand by 2030 (before considering adjacent upside from enterprise SSD, networking, edge device and CXL-based memory expansion).

Why agentic AI increases DRAM intensity

Agentic AI introduces more intermediate state and more “memory feet” around the GPU. Two mechanisms matter most:

1. **Agent concurrency and tool I/O drive larger CPU-side working sets** (more parallel agents, more cached results, more in-flight contexts). A recent Georgia Tech study highlights CPU-side orchestration/tool processing as a dominant contributor to latency and links throughput bottlenecks to CPU factors and GPU main memory capacity/bandwidth.
2. **Context windows and KV-cache management increasingly force the system to treat DRAM as a functional extension of HBM.** vLLM describes a KV offloading feature explicitly targeting offloading to CPU memory (DRAM) to improve inference throughput and manage memory pressure. NVIDIA’s Rubin/Vera platform materials go further, describing coherent CPU-GPU memory-enabling techniques such as KV-cache offload and multi-model execution efficiency – effectively turning LPDDR/DDR into an active tier for AI serving. The relevance rises as leading models push to ~1m token context windows, which mechanically expands cache/state requirements.

Modeling approach

We estimate incremental DRAM bit demand in 2030 driven by agentic AI orchestration workloads, focusing on CPU-attached memory. Our framework is intentionally simple and parameterized:

Agentic AI DRAM demand

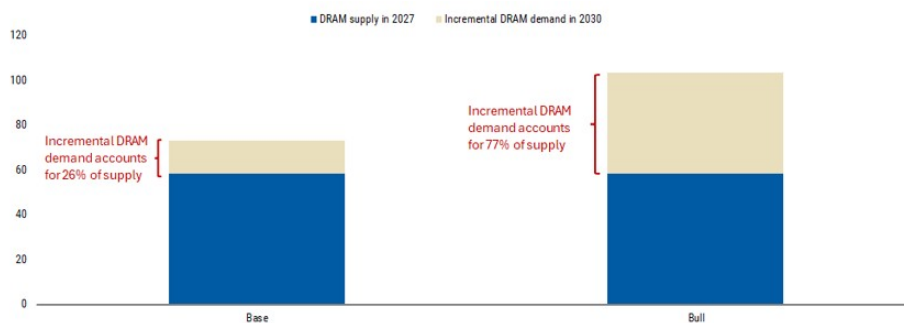
$$= (\text{Incremental orchestration CPU shipments}) \times (\text{average DRAM per CPU})$$

Key inputs

- **Incremental CPU shipment:** We derive incremental orchestration CPU demand from our CPU TAM analysis: 10mn units in base case and 15mn units in bull case. These represent incremental server CPU sockets deployed specifically for agentic AI orchestration workloads.
- **DRAM content per CPU (socket):** We anchor our assumptions on recent platform disclosures and architectural trends e.g., NVIDIA’s agentic-focused Vera CPU platforms scale to up to 1.5TB per CPU (and up to 400TB per rack). AMD’s EPYC 9005 architecture overview highlights that platform capacity can extend materially via CXL (example: 8TB max memory capacity conceptually framed as 6TB DDR + 2TB CXL). While x86 platforms can scale to multi-terabyte (4-8TB+) per socket, this reflects maximum capacity configurations. For agentic AI orchestration workloads, we expect a mix of: High-capacity x86 systems (memory-heavy coordination, databases) and more bandwidth-optimized architectures (e.g., NVIDIA

LPDDR-based systems). Hence, we assume 1.5TB per CPU content in a base case and 3TB per CPU in a bull case.

Exhibit 23: We expect incremental DRAM demand will take up 26-77% of supply by 2027



Source: Morgan Stanley Research estimates

Exhibit 24: DRAM specs for major AI CPUs

Company	CPU	Arch	Launch	Memory Type	Channels	Peak Bandwidth	Max Capacity (per socket/chip)
Intel	Xeon 6 (Granite Rapids)	x86	2024	DDR5 / MRDIMM	12	~600-840 GB/s	~4-6 TB
Intel	Xeon 7 (Diamond Rapids)	x86	2026 (est.)	DDR5 MRDIMM Gen2	16	~1.6 TB/s	~8-12 TB (est.)
AMD	EPYC 9004 (Genoa/Bergamo)	x86	2022-2023	DDR5	12	~460-500 GB/s	~6 TB
AMD	EPYC 9005 (Turin)	x86	2024	DDR5	12	~550-600+ GB/s	~6-8 TB
AMD	EPYC "Venice" (Zen 6)	x86	2H26 (est.)	DDR5	16 (rumored)	~1.2-1.6 TB/s	N/A
AWS	Graviton4	Arm	2024	DDR5	~8-12*	~300-500 GB/s*	~1-3 TB (est.)
Google	Axion	Arm	2025	DDR5	~8-12*	~300-500 GB/s*	~1-3 TB (est.)
Microsoft	Cobalt 100	Arm	2024	DDR5	~8-12*	~300-500 GB/s*	~1-3 TB (est.)
NVIDIA	Grace CPU	Arm	2023	LPDDR5X	Custom (wide bus)	~500 GB/s	~512 GB
NVIDIA	Vera CPU	Arm	2026 (est.)	LPDDR5X (SOCAMM)	Custom	~1.2 TB/s	~1-1.5 TB
ARM	AGI CPU	Arm	2026	DDR5	12	~816GB/s	~6TB


Source: Company data, Morgan Stanley Research; N/A = not applicable

ABF Substrate – Still Early Innings

ABF substrates are at the start of a powerful, AI-driven up-cycle set to last through the decade and supporting faster growth, pricing power and margin expansion. This cycle is structurally stronger and more sustainable than prior PC-driven cycles, with scope for sustained earnings growth in 2025-28. We see ABF demand shifting decisively toward AI, with GPUs, CPUs, ASICs and networking chips accounting for 75%+ of end demand by 2030. We project ABF value growth to accelerate to a 16.1% CAGR over 2025-30. Preferred exposures are Unimicron, Samsung Electro-Mechanics and Nanya PCB.

An AI-driven ABF substrate up-cycle is set to last to the end of the decade, driving growth to accelerate and margins to expand. Over the next few quarters, we expect ABF substrate suppliers to raise pricing to end customers to reflect higher costs, and, moving into 2027, we think there could be further price hikes, driven by a shortage of ABF substrate capacity. The stocks have started to move, but this cycle will be much stronger and more sustainable – fueling multiple expansion. This AI-driven cycle is structurally different and much stronger, driving significant earnings growth in 2025-28, we estimate. Our favored names for exposure are **Unimicron**, **Samsung Electro-Mechanics** and **Nan Ya PCB**, and least favored is **Ibiden**, although this down to valuation/overheated earnings expectations. See our recent ABF Substates note [here](#).

Exhibit 25: Entering the agentic AI era: ABF substrate market TAM to grow with AI




Source: Company data, Morgan Stanley Research estimates (e)

ABF substrate value growth is set to accelerate over the next 5 years amid a shift in demand to AI. Based on our bottom-up analysis, we project that ABF substrates will be in an up-cycle until the end of this decade, with industry undersupply starting from late 2026-27 onwards. What was once a PC-driven market – with PCs accounting for ~70% of it in 2015 – is undergoing a significant shift in demand, with AI-related applications (GPU, ASIC, server CPU, networking, etc.) capturing 75%+ of the end market by 2030, on our estimates. With incremental demand from agentic AI and the enlarged CPU TAM, we estimate this will drive ABF substrates' value to accelerate at a 17.9% CAGR over 2025-30 (vs 16.1% CAGR without the incremental CPU TAM), and vs. just 9.0% over 2020-25.

With agentic AI and the upside to CPU TAM of 10-15m units by 2030, this will result in the server CPU ABF substrate market to grow to ~\$4.7bn by CY30. In reality, the market upside should be even slightly higher, as more networking demand will be needed to connect all these CPUs together. But this incremental ~\$1.2bn of demand upside from server CPU ABF substrate alone would imply 5-10% upside to the overall ABF substrate TAM, all else equal.


For the ABF substrate market, without any new capacity additions, the enlarged CPU TAM would widen the supply-demand gap. In the scenario that we laid out back in February (see [here](#)), we estimate the under-supply gap will be ~7% by 2030, but with incremental 10-15m units of CPU TAM by 2030, this would cause the under-supply gap to widen to ~15%, which is a meaningful difference and would drive prices and margins up even more across the board for all ABF substrate suppliers.

Exhibit 26: Stronger ABF substrate market growth in the AI era



Source: Company data, Morgan Stanley Research estimates (e)

Exhibit 27: We estimate agentic AI could drive 5-10% incremental ABF substrate demand upside by 2030 from the enlarged CPU TAM alone




Source: Prismark, Morgan Stanley Research estimates (e)

Key agentic AI drivers. We believe the demand drivers will be mainly supported by CPUs, GPUs, ASICs, and networking chips used in data centers. Because of the increased demand from computing, we believe the demand for all of these semiconductors will grow and/or remain elevated in the coming years.


- For server CPUs, demand for general and storage computing highly correlates with overall AI demand. As such, we also forecast demand for server CPUs to rise in the coming years. In 2015, server CPUs only accounted for ~10% of total ABF substrate demand, but we estimate it has grown to ~37% in 2024 and will account for ~24% of the total ABF substrate market in 2030, still the largest application.
- For data center GPU units – primarily NVIDIA GPUs – we forecast AI GPU units to grow 48% year on year in FY26 (February 2026 to January 2027) and to grow an additional 2% in FY27. We do not have a forecast beyond that, but we model conservatively in the years following, and estimate the market demand to remain elevated, but flattish.
- For AI ASICs, demand for in-house designed chips appears strong in the near to medium term, so we project units to continue to grow throughout the rest of the decade – we also continue to see rising demand from Meta, Microsoft, Open AI, China, etc., where demand is small today but we bake in some growth over the next few years.
- Lastly, for networking chips, we believe this is highly correlated to computing demand. With GPU demand remaining elevated and growing ASIC volumes, we believe demand for networking chips will also continue to grow over the next several years.

Exhibit 28: ABF substrate demand by end applications, 2015-30e



Source: Prismark, Morgan Stanley Research estimates (e)

Exhibit 29: ABF demand led by server, AI GPU, AI ASIC, and networking-related ABF



Source: Prismark, Morgan Stanley Research estimates (e)

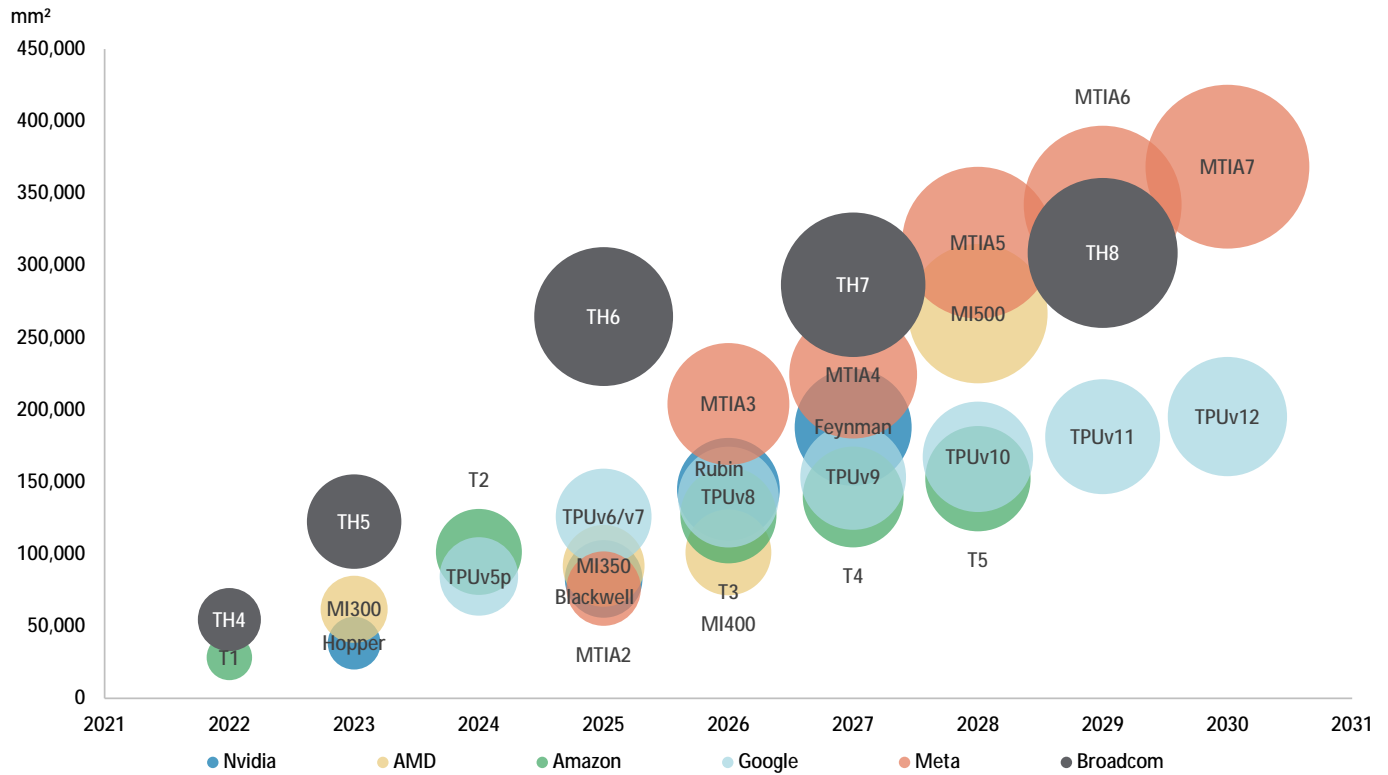
Larger package size and higher layer count are also a demand tailwind. Increasing demand for semiconductor circuit scale and functionality has been met with increases in the degree of integration of semiconductor dies and in package substrate mounting technology. With deterioration in cost performance owing to increases in the degree of integration of semiconductor dies and advancing process technology, moving to chiplets that contain multiple dies on one package substrate has become indispensable for raising total semiconductor performance.

Behind this is that once multiple functions are incorporated into a single semiconductor die, die size expands, and manufacturing complexity and cost rise. Closely watched methods for mounting multiple semiconductor elements on one package substrate and increasing total semiconductor performance are:

- 3D – where semiconductor dies are stacked vertically, and
- 2.5D – where semiconductor elements are connected via silicon or organic interposers on the package substrate.

Larger substrate size and higher layer count are particularly evident, in view of the increase in electrical connections to move data, control signals, and power between the AI accelerator dies and HBM memory stacks. This trend towards larger substrate size and higher layer count has been clear, and it does not appear to be reversing any time soon. So if we take into account the growth from substrate size and layer count, ABF substrate demand is set to grow even faster than unit growth.

Exhibit 30: Near-term ABF substrate trend points toward larger substrate sizes and higher layer count



Source: Company data, Morgan Stanley Research estimates. Note: Bubble size indicates ABF substrate demand, including area size and layer count.

More CPU Sockets Will Be Needed

This is another component that will benefit meaningfully from an enlarge CPU TAM, because every server CPU is accompanied by a CPU socket. Today, the server CPU market is mainly dominated by 3-4 suppliers, and we believe the enlarged server CPU TAM will mainly benefit FIT and Lotes under our coverage. We estimate FIT has around 40% global server CPU socket supply share, and Lotes ~35%.

On our estimates, every 1m CPU unit upside would drive 0.6% revenue upside to Lotes and 0.2% revenue upside to FIT, purely from increased sockets supplied; a 10-15m CPU unit upside would result in ~7% revenue upside to Lotes and ~2% revenue upside to FIT. However, we believe the revenue upside to these two companies should be even higher, because an enlarged CPU TAM means more general servers, DRAM connectors, more PCIe connectors, and other connectors and cables, all of which Lotes and FIT have exposure to. So we believe the overall revenue opportunity and upside to Lotes and FIT should be higher.

Exhibit 31: Every 1m CPU demand upside would imply 0.5% CPU socket revenue upside to Lotes

Every 1M Units of total CPU upside	
Server CPU Socket ASP (US\$)	21.9
Revenue upside to Lotes per 1M CPUs (US\$ M)	7.7
Revenue upside to Lotes per 1M CPUs (NT\$ M)	241
Lotes 2026 total revenue estimate (NT\$ M)	40,488
Lotes CPU Socket upside per 1M CPUs	0.6%

Source: Morgan Stanley Research estimates

Exhibit 32: ...and 10-15m CPU unit upside would drive ~7% CPU socket revenue upside to Lotes

If CPU TAM grows 10-15M units...	
Server CPU Socket ASP (US\$)	21.9
Revenue upside to Lotes if CPU TAM grows 10-15M units (US\$ M)	95.7
Revenue upside to Lotes if CPU TAM grows 10-15M units (NT\$ M)	3,015
Lotes 2026 total revenue estimate (NT\$ M)	41,065
Lotes CPU Socket upside if CPU TAM grows 10-15M units	7.3%

Source: Morgan Stanley Research estimates

Exhibit 33: Every 1m CPU demand upside would imply 0.2% CPU socket revenue upside to FIT

Every 1M Units of total CPU upside	
Server CPU Socket ASP (US\$)	21.9
Revenue upside to FIT per 1M CPUs (US\$ M)	8.8
FIT 2026 total revenue estimate (US\$ M)	5,483
FIT CPU Socket upside per 1M CPUs	0.2%

Source: Morgan Stanley Research estimates


Exhibit 34: 10-15m CPU unit upside would drive ~2% CPU socket revenue upside to FIT

If CPU TAM grows 10-15M units...	
Server CPU Socket ASP (US\$)	21.9
Revenue upside to FIT if CPU TAM grows 10-15M units (US\$ M)	109.4
FIT 2026 ABF revenue estimate (NT\$ M)	5,483
FIT CPU Socket upside if CPU TAM grows 10-15M units	2.0%

Source: Morgan Stanley Research estimates

Exhibit 35: We estimate FIT has ~40% global server CPU supplier share and Lotes ~35%

2025 Server CPU Socket Supply Share



Source: Morgan Stanley Research estimates

Passive Components TAM Will Also Grow

MLCC and passive components are also placed in general servers everywhere, so an enlarged server CPU TAM would drive upside to MLCC and other passive component demand. Different server specifications have different layouts of passive components and different volume and value content, so it is also extremely hard to quantify the content. But for simplicity's sake, we estimate ~\$30 MLCC content per general server, which houses on average 2x CPUs today. This was quantified in a report we published in 2024 ([Global Technology: MLCC – Quantifying the Cloud & Edge AI Opportunity](#)), where we quantified the MLCC market for different end segments.

In that report, we sized the general server MLCC TAM at \$400-500m. This year, we estimate the **total MLCC market globally is ~\$15-16bn** and, by 2030, we think it will reach ~\$20bn (assuming it remains at 5% CAGR from here on out). This would imply that a 10-15m CPU unit demand upside by 2030 would drive incremental ~\$500m MLCC value demand, or 2-3% upside to the overall MLCC market globally by 2030.

For Yageo specifically, we estimate that every ~1m CPU demand upside would result in 0.1% revenue upside, and a 10-15m CPU unit upside would result in a 1-2% revenue upside to Yageo, all else equal. However, Yageo also has exposure to inductors, resistors and tantalum capacitors, which are also all present in general servers, so the overall revenue upside should be higher.

Exhibit 36: Every 1m CPU demand upside would imply 0.1% MLCC revenue upside to Yageo

Every 1M Units of total CPU upside	
Server MLCC Content (US\$)	30.0
Revenue upside to Yageo per 1M CPUs (US\$ M)	4.5
Revenue upside to Yageo per 1M CPUs (NT\$ M)	142
Yageo 2026 total revenue estimate (NT\$ M)	155,443
Yageo MLCC revenue upside per 1M CPUs	0.1%

Source: Morgan Stanley Research estimates

Exhibit 37: 10-15m CPU unit upside would drive 1-2% MLCC revenue upside to Yageo

If CPU TAM grows 10-15M units...	
Server MLCC Content (US\$)	30.0
Revenue upside to Yageo if CPU TAM grows 10-15M units (US\$ M)	56.3
Revenue upside to Yageo if CPU TAM grows 10-15M units (NT\$ M)	1,772
Yageo 2026 total revenue estimate (NT\$ M)	158,893
Yageo MLCC revenue upside if CPU TAM grows 10-15M units	1.1%

Source: Morgan Stanley Research estimates

Greater China Semiconductor: Foundry, Design Service and IC Design – Seizing the CPU (Agentic AI) Opportunity


Strong server CPU demand from agentic AI and inferencing workload

The CPU plays a critical role in enabling agentic AI, serving as the central coordinator that executes instructions, manages system resources, and ensures reliable, real-time decision-making. While specialized hardware like GPUs and accelerators handle large-scale parallel computations, the CPU orchestrates the overall workflow – handling task scheduling, memory management, I/O operations, and the logic that underpins autonomous behavior. In agentic AI systems, which require continuous perception, reasoning and action, the CPU is essential for low-latency control loops, system stability, and integration across diverse components.

Therefore, we have been seeing strong server CPU demand since 2025, then with another jump in 2026. Besides server CPU demand, we have also seen the demand rebound for a few models of client CPU, such as Apple Mac Mini, to support the edge AI demand. As the result, we are forecasting overall CPU TAM for foundry to reach nearly \$33bn in 2026 and then to \$37bn by 2028. Within the TAM, we see TSMC, serving as the major foundry, as likely to reach around 70% of the market share in 2026 and continuing to grow to 75% in 2028 with the nodes for those CPU also continuing to migrate, like GPU and ASIC. For example, the forthcoming server CPU from AMD (Venice) and NVIDIA (Vera) are adopting TSMC's leading-edge 2nm and 3nm process, both one node migration from prior generation.

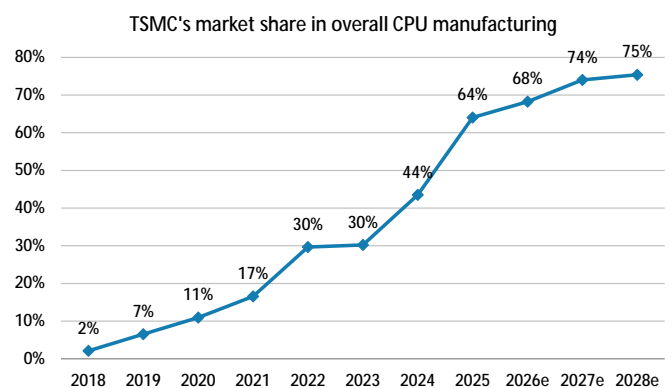
Lastly, we expect Intel will outsource its server CPU production to TSMC by 2H27, as it is facing fierce competition from AMD in the server/data center CPU space, whereas TSMC is more available to provide better time-to-market of its server CPU production with high-quality chips (suggesting better CPU performance).

Exhibit 38: Foundry addressable market size for CPUs




Source: Company data, Morgan Stanley Research estimates (e)

Exhibit 39: TSMC's market share in overall CPU manufacturing




Source: Company data, Morgan Stanley Research estimates (e)

Exhibit 40: TSMC CPU foundry revenue



Source: Company data, Morgan Stanley Research estimates

Exhibit 41: TSMC revenue contribution from Arm-based CPU



Source: Company data, Morgan Stanley Research estimates

Strong Arm-based CPU demand could benefit GUC from 2026 and beyond

Within our Greater China Semiconductor coverage, GUC has the highest revenue mix from the CPU demand. In our January notes [Benefiting from Google TPU/CPU demand upside](#) and [Stronger demand from automotive chips and Arm-based CPUs](#), we highlighted GUC's 2026 revenue upside from strong Google CPU demand. Our recent industry checks suggest that inference and general server demand for Arm-based CPUs is even stronger, including Google CPU. We now forecast Google CPU shipments to grow to 1.5mn units in 2026 (vs. 0.8mn units previously), which could contribute \$1.2bn of revenue to GUC, or >50% of its 2026 revenue. Our checks also suggest that gross margin could also be 1-2ppts better than our previous anticipation of 5% given the better order scale at TSMC.

In the meantime, we note that GUC also helps with the Cobalt 200 CPU for Microsoft, an Arm-based design tailored for Microsoft Cloud. Our checks also indicate strong demand for Cobalt at >1mn units in 2027 and onwards. We therefore forecast \$500-600mn revenue contribution from this in 2027-28.

Lastly, GUC management noted that it has secured next-gen Arm-based CPU design wins from its major customers, including both Google CPU and Microsoft Cobalt CPU. We therefore believe GUC can benefit from this agentic AI or inferencing demand at least until 2028.

Exhibit 42: GUC: Major customer breakdown

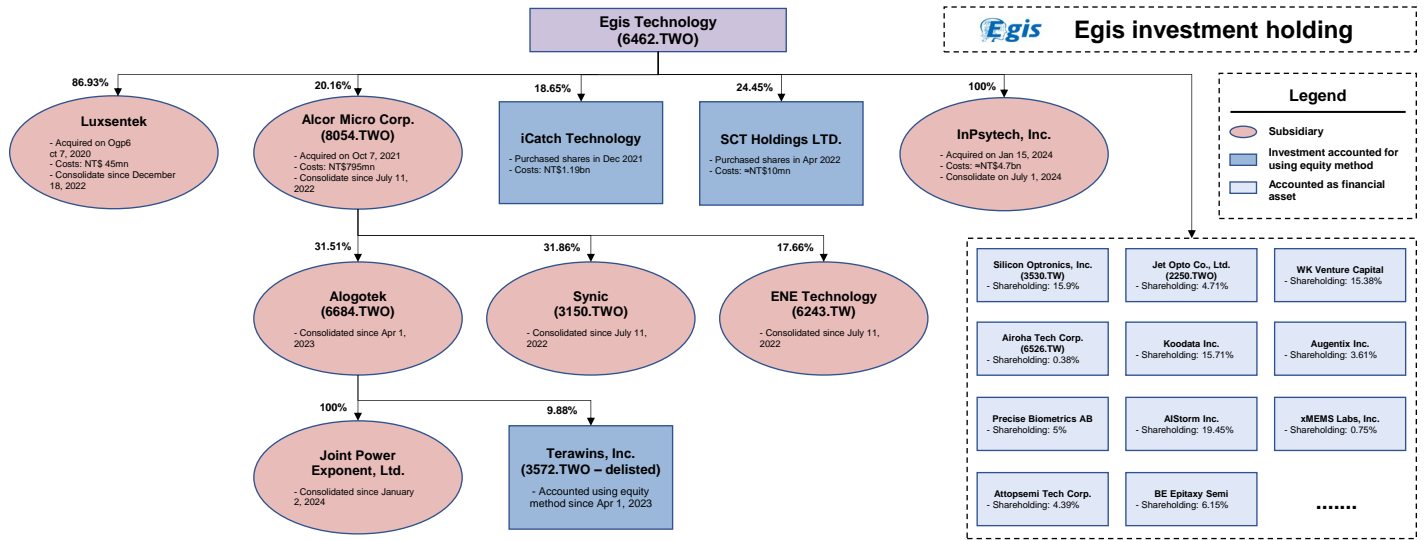
Revenue (US\$ mn)	2025e	2026e	2027e	2028e	2025 Y/Y	2026 Y/Y	2027 Y/Y	2027 Y/Y
NRE								
Normal NRE projects	76	170	200	220	-49%	124%	18%	10%
Ship-outs	185	108	110	120	68%	-42%	2%	9%
Total NRE	261	278	310	340	0%	7%	12%	10%
Turnkey								
Microsoft	60	80	150	400	100%	33%	88%	167%
Google	150	1,200	1,600	1,600	NA	700%	33%	0%
Korean AI Start-up	0	20	80	80	NA	NA	300%	0%
xAI/Tesla	0	125	800	1,200	NA	NA	540%	50%
Rivos	0	30	40	40	NA	NA	33%	0%
Amazon Lab 126	0	50	100	100	NA	NA	100%	0%
SK Hynix	50	10	10	10	-38%	-80%	0%	0%
Sandisk	0	20	250	250	NA	NA	1150%	0%
China ADAS	0	20	100	200	NA	NA	400%	100%
Suzhou Centec	54	27	14	7	-50%	-50%	-50%	-50%
Macronix	56	68	81	98	30%	20%	20%	20%
DJI	35	10	5	5	-31%	-71%	-50%	0%
Aspeed	105	121	139	150	50%	15%	15%	8%
Sony	29	26	23	21	-10%	-10%	-10%	-10%
Other turnkey	292	251	208	202	221%	-14%	-17%	-3%
Total turnkey	831	2,058	3,600	4,362	65%	148%	75%	21%
Ex-US CSP turnkey	621	653	1,050	1,162	31%	5%	61%	11%
IP								
Google (GLink)	2	5	5	5	NA	150%	0%	0%
Microsoft (HBM3E)	4	5	5	5	0%	25%	0%	0%
Others	2	7	7	7	-80%	250%	0%	0%
Total revenue	1,100	2,353	3,927	4,719	41.2%	113.9%	66.9%	20.2%

Source: Company data, Morgan Stanley Research estimates (e); Y/Y = year on year

Within our Greater China semiconductor coverage, among other design service vendors, Egis is working on diversifying into the AI and HPC supply chain. The company has made several investments into non-fingerprint business in the past few years, and these are now coming to fruition. We believe the two most important are its acquisition of InPsytech and a ~20% stake in Alcor Micro (not covered) (Exhibit 43). Egis has also expanded its staff and joined Arm's total design alliance with the aim of winning long-term ASIC business opportunities.

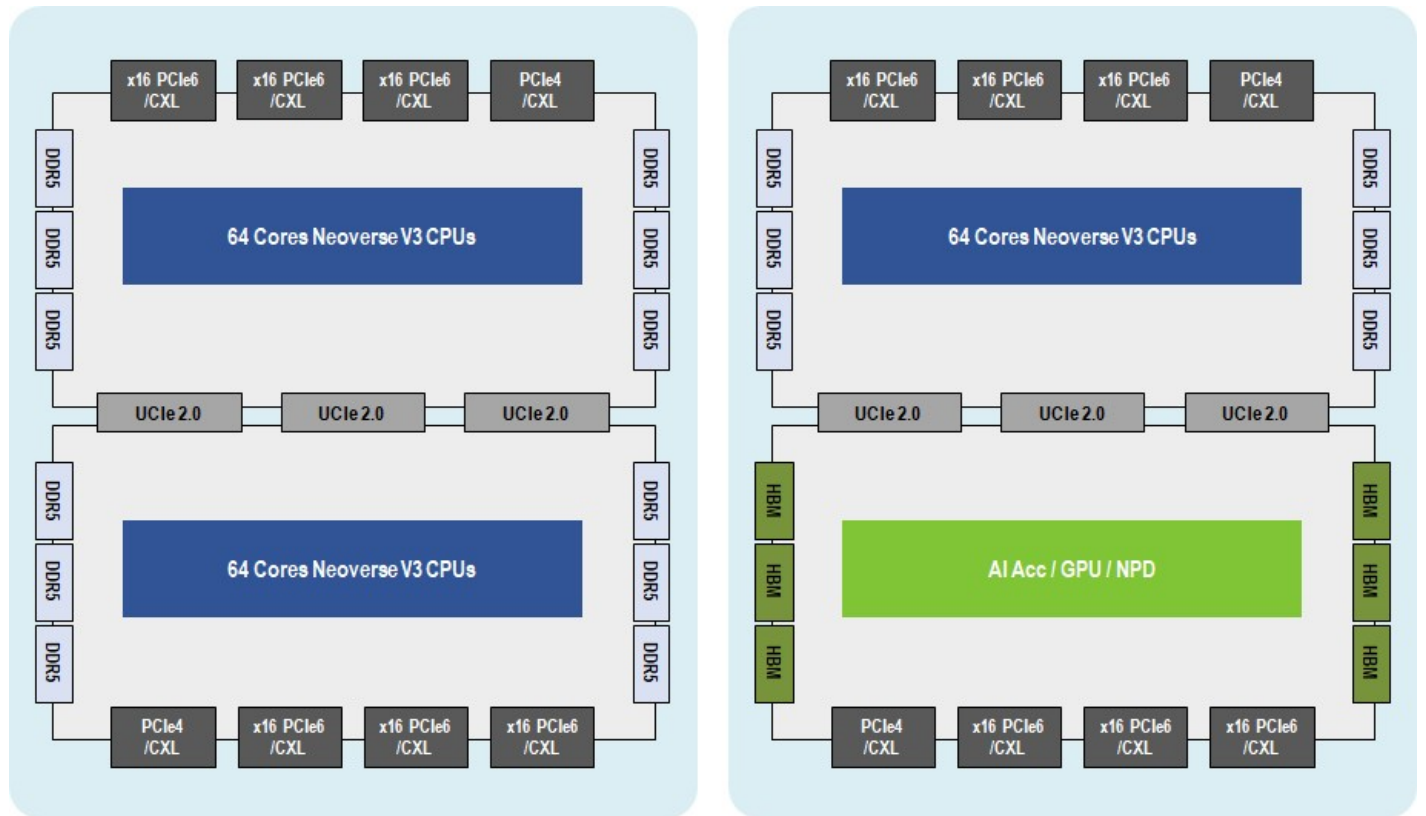
Alcor provides an Arm-based CPU platform, with two cores connected by UCI-E named Mobius 100 (Exhibit 44), which have been developed and built internally within the Egis alliance instead of licensing from a third party. According to Alcor, Mobius 100 could also provide heterogeneous chiplets (CPU with AI Acc/GPU/NPD). We think this could be suitable for neocloud or supercomputing clients outside of the US, given the ongoing data center buildout there (Exhibit 46), according to our supply chain checks. Alcor previously announced the timeline for its Arm-based CPU Mobius 100, with tape-out in 2Q26 and mass production in 4Q27. We are Equal-weight on Egis, however, as we are still waiting for future project wins.

Exhibit 43: Egis: Organizational structure



Source: Company data, Morgan Stanley Research

Exhibit 44: Alcor Arm-based CPU Mobius 100 (CCS V3) with CSA compliance



Source: Alcor

Exhibit 45: Spec summary of Mobius 100: single vs dual die configuration

List of Feature	70 Cores with Single CPU die	140 Cores Single package with Two CPU dies
Project / Product Name	Mobius100	
Product Application	HPC Server	
CSS V3 Version	CSS V3 V1 EAC2 (The latest Version, support CPU to CPU and CPU to AI Accelerator)	
Process Node	TSMC N3P	
Single Die Size	420 mm ²	
Processor Architecture Specifications	Supports RME (Realm Management Extension) and Arm Confidential Compute Architecture (CCA) Supports SVE2 (Scalable Vector Extension x 2) with 128-bit vector length	
Core Voltage	0.75v	
CPU Core Max Clock Frequency	Normal Mode: 3.0GHz @0.75v Turbo Mode: 3.6GHz @0.85v	
Power Consumption	-170W	-350W
Core Numbers	One Die is 70Cores, 6Cores for redundant	One package is 140Cores, 12Cores for redundant
DRAM Capacity	DDR5@12800, MRDIMM x6 DDR5 BW : 614.4 GB/s	DDR5@12800, MRDIMM x12 Capacity up to 1,536 GB
PCIe Slot Numbers	PCIeGen6 x 3 PCIeGen4 x 1	Double to single CPU die
D2D Feasibility	UCIe2.0-S x 3 6 x CCG by 64B CXS I/F 3 x CCG_CML_SMP, 3 x CCG_CHIC2C	
Multichip Expansion	1. Up to 4 CPU Chiplets 2. Accelerator or memory expansion	1. Through CXL/PCIe and CML_SMP UCIe 2. Through CXL type2/3 or CHIC2C UCIe

Source: Alcor

Exhibit 46: Cloud Service Providers continuing their data center construction plans in Europe

Company	Country	Location	Available	Upcoming	Sub Total
Google	Belgium	Farciennes St. Ghislain	1	1	21
	Denmark	Fredericia	1	1	
	Finland	Hamina	6	1	
	Germany	Hanau	1		
	Ireland	Dublin	2		
	Netherlands	Eemshaven	2		
		Groningen		1	
		Middenmeer	1		
		Winschoten		1	
		Norway	Skien		
	UK	Waltham Cross		1	
AWS	Sweden	Stockholm	3		27
	Germany	Frankfurt	3		
	Ireland	Ireland	3		
	UK	London	3		
	Italy	Milan	3		
	France	Paris	3		
	Spain	Spain	3		
	Switzerland	Zurich	3		
	EU	European Sovereign Cloud		3	
Meta	Denmark	Odense	1		3
	Ireland	Clonee	1		
	Sweden	Lulea	1		
Oracle	UK	London	2		17
		Newport	2		
	France	Paris	1		
		Marseille	1		
	Germany	Frankfurt	2		
	Italy	Milan	1		
		TBD		1	
	Netherlands	Amsterdam	1		
	Serbia	Jovanovac	1		
	Spain	Madrid	2		
		TBD		1	
	Sweden	Stockholm	1		
	Switzerland	Zurich	1		
MSFT	UK	London	3		49
		Cardiff	1		
	Switzerland	Zurich	3		
	Sweden	Gävle and Sandviken	3		
	Spain	Madrid	3		
	Poland	Warsaw	3		
	Norway	Oslo	3		
	Italy	Milan	3		
	Greece			3	
	Germany	Frankfurt	3		
	France	Paris	3		
	Finland	Helsinki		3	
	Ireland	Ireland	3		
	Netherlands	Netherlands	3		
	Denmark	Copenhagen		3	
Belgium	Brussels		3		
Austria	Vienna	3			
Coreweave	UK	London	1	1	7
		Crawley	1		
	Spain	Barcelona	2		
	Sweden	Stockholm	1		
	Norway	Vennesla		1	
Total			98	26	124

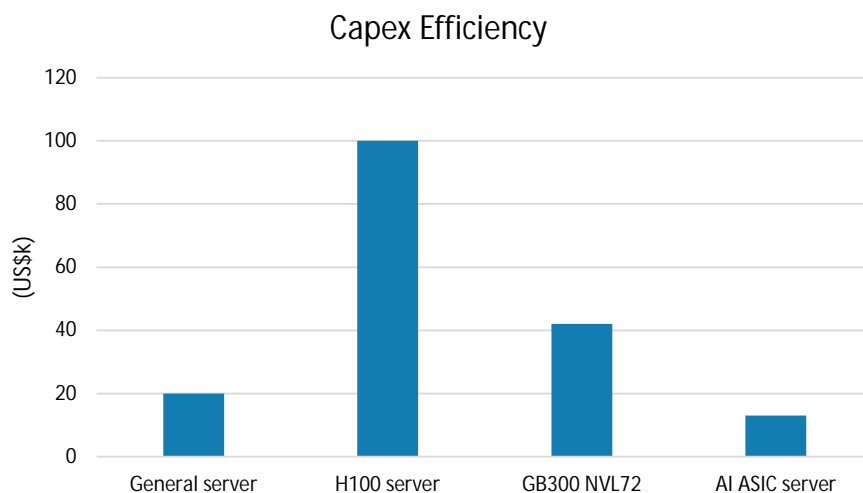
Source: Company data, Morgan Stanley Research

BMC to benefit from strong CPU demand

Baseboard management IC (BMC) is the key component on the server board, with a view to enabling the remote control for the server firmware system and maintaining the stability of the data center. Previously, Aspeed indicated that general server and AI ASIC server actually generate higher efficiency among the capex spending from hyperscalers ([Exhibit 47](#)). As a result, as the CPU plays a critical role in enabling agentic AI, serving as the central coordinator that executes instructions, manages system resources, and ensures reliable, real-time decision-making, we expect the growing demand for CPU/general server to further lift Aspeed's BMC revenue in the coming years.

Currently, Aspeed has ~70% market shares in the CPU server BMC market, with major US and Chinese hyperscalers being its end customers, including Meta, AWS, Alibaba and Microsoft. However, as it migrates into the new generation BMC product AST2700, we expect it to further expand its market shares in customers including Dell and Google, which would help expand its revenue as spec migration also brings 40-50% of ASP increase.

Exhibit 47: General server and AI ASIC server generate the highest capex efficiency for BMC



Source: Company data, Morgan Stanley Research. Note: Capex efficiency is defined as the amount of server spend that will contribute to one BMC demand. The lower capex efficiency, the better for Aspeed.


Aspeed recently announced a rise in the company's BMC TAM from 46.5mn to 65.77mn in 2030, mainly driven by AI-related general server demand, which we think agentic AI would be the key driving force behind. General server growth should stay stable at 6% from 2025-2030, while AI-related general servers should account for: (1) 25% of total 2026 general server demand, (2) 45% in 2027, and (3) 20% throughout over 2028-2030 ([Exhibit 48](#)).

On the order visibility front, customers began providing forecasts 7-8x its single-month revenue. Although Aspeed could not roll out any double booking, the current dynamic is for shipments way beyond customers' needs. According to the company, customers also expect double capacity year on year in 2027.

Exhibit 48: Aspeed expects its BMC TAM to reach 65.7mn units in 2030

BMC TAM Forecast Revised

- We see huge demand for **AI-related General Server** for agentic AI workload and simple inference tasks.
- BMC TAM is estimated to be **65.77 million** units in 2030.



Note : 1) BMC for General Servers will grow at 6% CAGR.
 2) BMC for AI-related General Servers is estimated to be 8.5 million in 2026, with 45% growth rate in 2027 and 20% thereafter.
 3) BMC for GPU servers and AI ASIC servers in 2026 is based on the estimate of GPU and AI ASIC shipment and the server architecture.
 4) Assuming BMC for GPU Servers grow at 45% in 2027 and 20% afterwards. We also include BMCs for ICMS here.
 5) Assuming BMC for AI ASIC Servers grow at 45% in 2027 and 20% afterwards.

Source: Aspeed

Exhibit 49: Aspeed's product roadmap

Delivering quantum-resilient **OCP Caliptra security** across the entire **ASPEED** product portfolio.



Source: Aspeed

Montage set to benefit from CPU demand through memory interface chips

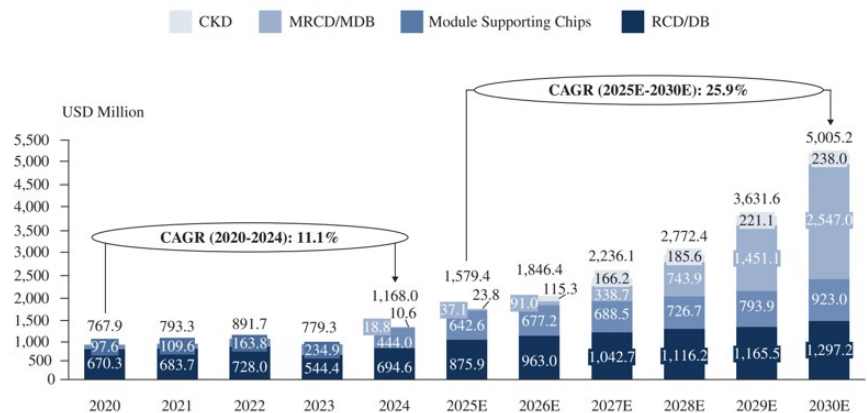
Core business DRAM interconnect business levered to cloud capex/memory supercycle: Montage is a cloud capex proxy through its main memory interconnect business. Globally, cloud capex industry growth is projected to remain high, at a 30% CAGR over 2025-2027, we estimate, with ongoing investment from the top CSPs. The memory supercycle from 2H25 also suggests higher memory interface content per general server. We believe the strong cloud/memory cycle will drive faster technology migration, which is positive for Montage, a technology leader. As of 2024, the company had a revenue market share of 36.8% in the global memory interconnect chip market, followed by Rambus (36.0%) and Renesas (20.5%). Montage has been an incumbent leader in the segment, and it continues to maintain technology advantages over its competitors. As of 2025, the memory interconnect revenue generated by Montage was ~1.7x that of the product revenue of Rambus.

Exhibit 50: Application and deployment of memory interconnect chips across different types of memory modules

Memory Module	RCD	DB	SPD	TS	PMIC	MRCD	MDB	CKD
Servers RDIMM (DDR4)	1	-	1	-	-	-	-	-
LRDIMM (DDR4)	1	9	1	-	-	-	-	-
RDIMM (DDR5)	1	-	1	2	1	-	-	-
LRDIMM (DDR5)	1	10	1	2	1	-	-	-
MRDIMM (DDR5)	-	-	1	2	1	1	10	-
PCs UDIMM/SODIMM (DDR4)	-	-	1	-	-	-	-	-
UDIMM/SODIMM (DDR5)	-	-	1	-	1	-	-	-
CUDIMM/CSODIMM/CAMM (DDR5)	-	-	1	-	1	-	-	1
LPCAMM (DDR5)	-	-	1	-	1	-	-	-

Source: Frost & Sullivan. Note: UDIMM: Unbuffered Dual In-line Memory Module; (2) CUDIMM: Clocked Unbuffered Dual In-line Memory Module; (3) SODIMM: Small Outline Dual In-line Memory Module; (4) CSODIMM: Clocked Small Outline Dual In-line Memory Module; (5) CAMM: Compression Attached Memory Module; (6) LPCAMM: Low Power Compression Attached Memory Module.

Exhibit 51: Market size of memory interconnect chips (by revenue), breakdown by product types, 2020-30E



Source: Frost & Sullivan (E = estimates)

Legacy memory players still need time to catch up

We have turned selective on legacy memory: The old memory supercycle began with DDR4. We double-upgraded Winbond to Overweight in March 2025 (link) and called out the DDR supercycle in June of that year (link). NOR and MLC NAND followed from 4Q25 (link). The old memory supercycle could run from 9 months to 1.5 years, slightly shorter than the mainstream cycle given the smaller TAM.

DDR4 – narrowing supply-demand gap: On the supply side, Winbond noted 100% year-on-year bit shipment growth in 2026 (link), while Samsung kept DDR4 shipments flattish quarter on quarter in 2Q (note). Also, we believe PSMC will license 1y or even 1z process on DDR4 production, likely ramping from 1H27. On the demand side, 3C product demand continues to shrink. We now expect undersupply to narrow to <20% in 2H, below our prior estimate. Into 2027, CXMT may add 10kwpm of DDR4 capacity (on top of its current 8kwpm) to meet strong China local demand, which could put pressure on DDR4 pricing. CXMT’s DDR5 (Gen4B) capacity expansion in Shanghai should also accelerate in 2027.

Legacy Flash – likely the next DDR4: We think MLC and legacy TLC NAND could face a deeper shortage into 2H, with **undersupply potentially reaching ~40%**. Global capacity for 2GB to 64GB NAND has virtually disappeared. Customer inventory generally covers only another 6-9 months. Macronix remains the only supplier able to fill the gap. We expect MLC and legacy TLC pricing to rise more than 200% from 1Q26 to 4Q26.

Earnings Revisions and Price Target Changes

SK hynix (Overweight)

What's changed? We update our 2026 ASP assumptions according to our latest channel check data. As most of the suppliers are in talks with key customers on 3-5 year long-term agreements (LTAs), which provide downside protection to the pricing trend, we now expect pricing to remain elevated throughout 2027, compared with our previous model which assumed prices would start to decline in 2H27. For 2028, we now model -19% and -23% for DRAM/NAND ASP year on year, reflecting our cautious view on when capacity catches up.

As a result, our 2026-28 EPS estimates for SK hynix rise by 24%, 37% and 78%, respectively, and we increase our residual income based price target by 31% to W1,700,000, implying 4x 2027 P/E, on our estimates, and c.50% potential upside from the current share price. We increase our bull case valuation by 61% to W2,580,000, reflecting stronger earnings momentum, while our bear case remains unchanged at W400,000 as we assume weaker macro conditions could slow data center growth.

Exhibit 52: SK hynix: Earnings revisions

(W bn)	FY26E			FY27E			FY28E		
	Previous	Revised	Change	Previous	Revised	Change	Previous	Revised	Change
Sales	281,377	335,034	19%	378,771	487,189	29%	329,365	505,430	53%
DRAM (& HBM)	227,536	249,923	10%	311,339	370,194	19%	272,861	387,771	42%
NAND	52,528	83,797	60%	66,119	115,682	75%	55,190	116,346	111%
Operating Profit	222,498	276,791	24%	298,316	407,567	37%	226,146	403,298	78%
DRAM (& HBM)	191,816	214,203	12%	260,803	319,658	23%	208,254	323,164	55%
NAND	30,682	62,588	104%	37,513	87,909	134%	17,892	80,134	348%
OP Margin	79%	83%	4pp	79%	84%	5pp	69%	80%	11pp
DRAM (& HBM)	84%	86%	1pp	84%	86%	3pp	76%	83%	7pp
HBM	72%	72%	0pp	69%	69%	0pp	64%	64%	0pp
NAND	58%	75%	16pp	57%	76%	19pp	32%	69%	36pp
Net Profit	172,174	214,166	24%	232,037	316,738	37%	174,699	310,625	78%
EPS for consensus	241,579	300,498	24%	325,574	444,418	37%	245,121	435,841	78%

Source: Morgan Stanley Research estimates

Exhibit 53: SK hynix: Key assumptions

	FY26E			FY27E			FY28E		
	Previous	Revised	Change	Previous	Revised	Change	Previous	Revised	Change
DRAM Bit (1Gb Eq, mn)	114,607	114,607	0.0%	143,134	143,134	0.0%	175,210	175,210	0.0%
DRAM Bit Growth	20.9%	20.9%	0.0pp	24.9%	24.9%	0.0pp	22.4%	22.4%	0.0pp
DRAM ASP (1Gb eq; US\$)	1.6	1.8	12%	1.51	1.89	25%	1.05	1.6	49%
DRAM OPM	84.3%	85.7%	1.4pp	83.8%	86.3%	2.6pp	76.3%	83.3%	7.0pp
HBM Bit (1Gb Eq, mn)	13,657	13,657	0%	18,144	18,144	0%	20,769	20,769	0%
HBM ASP (1Gb eq; US\$)	2.1	2.08	0%	2.14	2.14	0%	2.04	2.04	0%
HBM OPM	72.0%	72.0%	0.0pp	68.7%	68.7%	0.0pp	64.1%	64.1%	0.0pp
Conv.DRAM ASP	1.5	1.7	14%	1.4	1.9	30%	0.9	1.5	62%
Growth YoY	191.1%	232.6%	22%	(5.1%)	8.2%	nm	(34.6%)	(18.5%)	nm
NAND Bit (1GB, mn)	213,755	213,755	0.0%	257,515	257,515	0.0%	312,764	312,764	0.0%
NAND Bit Growth	14.9%	14.9%	0.0pp	20.5%	20.5%	0.0pp	21.5%	21.5%	0.0pp
NAND ASP (1GB eq; US\$)	0.2	0.3	70%	0.18	0.3	84%	0.12	0.3	111%
NAND OPM	58.4%	74.7%	16.3pp	56.7%	76.0%	19.3pp	32.4%	68.9%	36.5pp
USD/KRW rate	1,450	1,450	0.0%	1,341	1,341	0.0%	1,312	1,312	0.0%
Number of Shares (mn)	712,702	712,702	0.0%	712,702	712,702	0.0%	712,702	712,702	0.0%

Source: Morgan Stanley Research estimates; nm = not meaningful

Exhibit 54: SK hynix: Residual income model

RIM													
Fiscal period	FY24A	FY25A	FY26E	FY27E	FY28E	FY29E	FY30E	FY31E	FY32E	FY33E	FY34E	FY35E	Terminal value
Forecast year	0	1	2	3	4	5	6	7	8	9	10	11	
Total Shareholder Equity	73,916	113,916	325,468	640,264	953,684	1,175,303	1,358,816	1,564,869	1,801,530	2,067,566	2,357,277	2,673,851	2,754,066
Core Net Profit	19,797	42,948	213,690	315,865	314,578	222,778	184,670	207,328	237,935	267,310	290,985	317,848	327,384
Residual income	32,148	188,426	260,335	222,926	100,361	38,959	39,216	44,367	44,837	36,557	28,558	29,415	
ROE	45.73%	97.27%	65.41%	39.47%	20.93%	14.57%	14.18%	14.14%	13.82%	13.15%	12.64%	12.06%	
Discount period		0.00	0.50	1.50	2.50	3.50	4.50	5.50	6.50	7.50	8.50	9.50	
Discount factor		1.00	0.95	0.85	0.76	0.68	0.61	0.55	0.49	0.44	0.40	0.36	0.36
PV of Equity Capital		32,147.6	178,444.3	221,116.0	169,814.1	68,565.2	23,870.8	21,550.1	21,866.1	19,818.5	14,492.1	10,153.6	10,458.2
Beginning Equity Capital			325,468										
PV of Equity Capital			760,149										
PV of Continuing Value			126,729										
Total Equity Value			1,212,346										
FD shares outstanding ('000)			728,002										
Per share value			1,700,000										

Source: Company data, Morgan Stanley Research estimates (E)

Samsung Electronics (Overweight)

What's changed? We update our 2026 ASP assumptions according to our latest channel check data. As most of the suppliers are in talks with key customers on 3-5 year LTAs, which provide downside protection to the pricing trend, we now expect pricing to remain elevated throughout 2027, compared with our previous model which assumed prices would start to decline in 2H27. For 2028, we now model -14% for DRAM/NAND ASP year on year, reflecting our cautious view on when capacity catches up.

As a result, our 2026-28 EPS estimates for Samsung rise by 45%, 74% and 121%, respectively, and we increase our residual income based target price by 44% to W362,000, implying 5x 2027 P/E, on our estimates, and 68% potential upside from the current share price. We increase our bull case valuation by 34% to W430,000, reflecting stronger earnings momentum, and our bear case PT by 18% to W120,000 as we assume weaker macro conditions could slow data center growth.

For Samsung Electronics preferred shares, we now apply a 20% discount (based on the past 2-year average) to the common shares vs. a 15% discount previously (based on the past 10-year average). The reason for the change is that most of the new ETF/basket products on AI and memory segments only include the common share, which leads to larger discrepancy structurally. Our updated price target is now W289,600, and our bull and bear valuations are W344,000 and W96,000, respectively.

Exhibit 55: Samsung Electronics: Earnings revisions

(W bn)	FY26E			FY27E			FY28E		
	Previous	Revised	Change	Previous	Revised	Change	Previous	Revised	Change
Sales	583,348	735,259	26%	719,814	1,010,929	40%	698,805	1,074,793	54%
Semiconductor (DS)	386,042	514,881	33%	485,922	749,058	54%	427,946	774,028	81%
Memory	350,390	479,122	37%	435,727	698,862	60%	375,001	721,084	92%
Display	38,926	39,036	0%	40,402	40,402	0%	44,528	41,510	-7%
Mobile/Network	125,837	148,800	18%	153,359	181,338	18%	180,429	213,352	18%
Harman	19,822	19,822	0%	22,310	22,310	0%	22,979	22,979	0%
Operating Profit	299,008	428,085	43%	367,446	630,601	72%	271,322	609,886	125%
Semiconductor (DS)	265,550	397,738	50%	327,480	592,017	81%	245,041	585,307	139%
Memory	267,164	399,361	49%	327,495	592,032	81%	243,587	583,853	140%
Display	5,569	1,589	-71%	5,523	2,020	-63%	6,424	2,075	-68%
Mobile/Network	3,969	4,837	22%	11,686	13,807	18%	14,672	17,319	18%
Harman	1,982	1,982	0%	2,231	2,231	0%	2,298	2,298	0%
OP Margin	51.3%	58.2%	7.0ppt	51.0%	62.4%	11.3ppt	38.8%	56.7%	17.9ppt
Semiconductor (DS)	68.8%	77.2%	8.5ppt	67.4%	79.0%	11.6ppt	57.3%	75.6%	18.4ppt
Memory	76.2%	83.4%	7.1ppt	75.2%	84.7%	9.6ppt	65.0%	81.0%	16.0ppt
Display	14.3%	4.1%	-10.2ppt	13.7%	5.0%	-8.7ppt	14.4%	5.0%	-9.4ppt
Mobile/Network	3.2%	3.3%	0.1ppt	7.6%	7.6%	0.0ppt	8.1%	8.1%	0.0ppt
Harman	10.0%	10.0%	0.0ppt	10.0%	10.0%	0.0ppt	10.0%	10.0%	0.0ppt
Net Profit	242,832	353,057	45%	275,994	480,008	74%	217,225	480,929	121%
EPS for consensus	36,027.7	52,371.8	45%	40,967.0	71,227.7	74%	32,091.3	71,049.1	121%

Source: Morgan Stanley Research estimates

Exhibit 56: Samsung Electronics: Key assumptions

	FY26E			FY27E			FY28E		
	Previous	Revised	Change	Previous	Revised	Change	Previous	Revised	Change
Mobile shipment (mn)	309	353	14.2%	347	397	14.5%	357	409	14.5%
Mobile blended ASP (US\$)	258	269	3.9%	305	317	3.9%	350	364	3.9%
OPM	3.2%	3.3%	0.0ppt	7.6%	7.6%	0.0ppt	8.1%	8.1%	0.0ppt
DRAM OP	209,299	259,020	23.8%	257,887	382,923	48.5%	184,396	389,254	111.1%
DRAM OPM	83.3%	85.9%	3.2%	82.1%	86.7%	5.6%	73.4%	83.8%	14.1%
DRAM Bit (1Gb Eq, mn)	137,787	143,117	3.9%	164,635	176,205	7.0%	189,036	216,297	14.4%
DRAM Bit Growth	21.7%	26.4%	4.7%	19.5%	23.1%	3.6%	14.8%	22.8%	7.9%
DRAM ASP (1Gb eq: US\$)	1.26	1.45	15.5%	1.42	1.87	31.4%	1.01	1.64	61.7%
DRAM ASP Growth	223.4%	273.4%	50.1%	13.1%	28.7%	15.6%	(28.9%)	(12.4%)	16.4%
HBM OP	26,092	26,351	1.0%	39,375	42,039	6.8%	35,535	42,493	19.6%
HBM OPM	68.0%	68.6%	1.0%	68.3%	68.9%	0.9%	64.2%	64.9%	1.0%
HBM Bit (1Gb Eq, mn)	13,075.2	13,055	(0.16%)	20,121.7	21,340	6.05%	20,366.8	24,108	18.37%
HBM Bit Growth	149.66%	149.26%	(0.39%)	53.89%	63.47%	9.57%	1.22%	12.97%	11.76%
HBM ASP (1Gb eq: US\$)	2.02	2.03	0.16%	2.14	2.13	(0.22%)	2.07	2.07	(0.01%)
HBM ASP Growth YoY	20.07%	20.26%	0.19%	5.52%	5.13%	(0.39%)	(3.10%)	(2.90%)	0.20%
NAND OP	57,866	140,341	142.5%	69,608	209,109	200.4%	59,191	194,599	228.8%
NAND OPM	58.4%	79.0%	35.2%	57.3%	81.3%	42.0%	47.8%	75.8%	58.8%
NAND Bit (1GB, mn)	375,583	375,583	0.0%	450,858	460,809	2.2%	517,079	548,252	6.0%
NAND Bit Growth	17.6%	17.6%	0.0%	20.0%	22.7%	2.6%	14.7%	19.0%	4.3%
NAND ASP (1GB eq: US\$)	0.18	0.33	79.4%	0.20	0.42	106.9%	0.18	0.36	95.2%
NAND ASP Growth	110.3%	277.3%	166.9%	10.6%	27.6%	17.0%	(9.2%)	(14.3%)	(5.1%)
USD/KRW rate	1,450	1,450	0.0%	1,341	1,341	0.0%	1,312	1,312	0.0%
Number of Shares (mn)	6,607	6,607	0.0%	6,607	6,607	0.0%	6,607	6,607	0.0%

Source: Morgan Stanley Research estimates

Exhibit 57: Samsung Electronics: Residual income model

RIM													
Fiscal period	FY25A	FY26E	FY27E	FY28E	FY29E	FY30E	FY31E	FY32E	FY33E	FY34E	FY35E	FY36E	Terminal
Forecast year	0	0	1	2	3	4	5	6	7	8	9	10	value
Total Shareholder Equity	429,422	754,785	1,202,481	1,652,237	2,009,719	2,303,187	2,557,400	2,811,212	3,127,541	3,514,111	3,992,943	4,575,127	4,112,731
Residual Income	284,965	367,465	316,783	316,783	178,093	76,650	5,903	(23,710)	6,025	35,849	78,350	120,694	69,992
Core Net Profit	45,225	353,057	480,008	480,929	388,656	324,642	285,387	284,985	347,503	417,744	510,005	613,358	525,306
ROE		59.63%	49.05%	33.69%	21.23%	15.05%	11.74%	10.62%	11.70%	12.58%	13.59%	14.32%	13.2%
Discount period		0.00	1.00	2.00	3.00	4.00	5.00	6.00	7.00	8.00	9.00	10.00	
Discount factor		1.00	0.90	0.80	0.72	0.65	0.58	0.52	0.47	0.42	0.38	0.34	0.38
PV of Equity Capital		284,965.2	329,565.4	254,807.3	128,476.4	49,592.1	3,425.2	(12,339.0)	2,812.0	15,006.6	29,414.7	40,638.3	26,276.8
Beginning Equity Capital		1,202,481											
PV of Equity Capital		867,676											
PV of Continuing Value		318,413											
Total Equity Value		2,388,570											
FD shares outstanding ('000)		6,606,854											
Per share value		362,000											

Source: Company data, Morgan Stanley Research estimates (E)

Samsung Electro-Mechanics (Overweight)

What's changed? We raise our SEMCO MLCC margin assumption from mid-teens to mid-high teens over 2026-28 on the back of stronger AI MLCC mix as well as IT MLCC recovering from its previous trough. We increase ABF revenue growth to ~50% level in 2026-27 due to stronger ASP amid industry-wide tightness, while in 2028 we model ABF revenue growth of 57% as new capacity ramps up. We lift our 2026-28 margin assumptions for ABF from mid-20% to low 30% due to better pricing. As a result, our 2026-28 EPS forecasts increase by 0%, 19% and 31% respectively.

As MLCC and ABF exposure to AI is increasing, the market tends to re-rate the stocks amid the structural tailwind. Hence, the previous cycle peak valuation reference is no longer suitable for SEMCO, we have change our methodology from an average of residual income and P/B approaches to just RI, reflecting the intrinsic value of the company.

We lift our residual income based price target by 42% to W710,000, implying 4x 2027 P/B, on our estimates, and c.5% potential upside. The change reflects our earnings upgrades and higher P/B multiple, given SEMCO's stronger growth prospects. While we increase our bull case value by 38% to W730,000 on stronger earnings momentum, our bear case only moves up by 9% to W200,000, as we assume the MLCC localization trend intensifies, with more competition from Chinese suppliers resulting in increased pricing pressure.

Exhibit 58: Samsung Electro-Mechanics: Earnings revisions

(W bn)	FY26E			FY27E			FY28E		
	Previous	Revised	Change	Previous	Revised	Change	Previous	Revised	Change
Sales	13,432	13,915	4%	16,649	17,894	7%	19,486	22,257	14%
Optics	4,036	4,036	0%	4,113	4,113	0%	4,039	4,039	0%
Substrate	2,933	2,951	1%	3,777	4,016	6%	4,785	5,765	20%
Components (MLCC)	6,463	6,928	7%	8,759	9,765	11%	10,662	12,453	17%
Operating Profit	1,753	1,802	3%	2,504	3,003	20%	3,111	4,110	32%
Optics	188	188	0%	192	192	0%	202	219	8%
Substrate	462	492	7%	605	924	53%	756	1,488	97%
Components (MLCC)	1,102	1,121	2%	1,707	1,886	11%	2,154	2,402	12%
OP Margin	13.0%	12.9%	-0.1pp	15.0%	16.8%	1.7pp	16.0%	18.5%	2.5pp
Optics	4.7%	4.7%	0.0pp	4.7%	4.7%	0.0pp	5.0%	5.4%	0.4pp
Substrate	15.8%	16.7%	0.9pp	16.0%	23.0%	7.0pp	15.8%	25.8%	10.0pp
Components (MLCC)	17.1%	16.2%	-0.9pp	19.5%	19.3%	-0.2pp	20.2%	19.3%	-0.9pp
Net Profit	1,482	1,484	0%	2,038	2,420	19%	2,477	3,244	31%
EPS for consensus	18,733	18,759	0%	25,766	30,596	19%	30,572	40,033	31%

Source: Morgan Stanley Research estimates

Exhibit 59: Samsung Electro-Mechanics: Residual income model

RIM	FY25A	FY26E	FY27E	FY28E	FY29E	FY30E	FY31E	FY32E	FY33E	FY34E	FY35E	Terminal value
Forecast year	0	1	2	3	4	5	6	7	8	9	10	
Total Shareholder Equity	9,471	10,778	13,020	16,086	20,262	25,261	31,162	37,754	45,605	53,900	63,414	66,585
Residual Income		472	1,230	1,788	2,543	2,907	3,264	3,331	3,868	3,504	3,834	3,475
Core Net Profit		1,484	2,420	3,244	4,360	5,183	6,086	6,776	8,035	8,479	9,699	10,184
ROE		14.66%	20.34%	22.29%	23.99%	22.77%	21.57%	19.67%	19.28%	17.04%	16.54%	15.7%
Discount period	0.00	1.00	2.00	3.00	4.00	5.00	6.00	7.00	8.00	9.00	10.00	
Discount factor	1.00	0.91	0.83	0.75	0.68	0.62	0.56	0.51	0.47	0.42	0.39	0.39
PV of Equity Capital	0.0	428.8	1,016.9	1,343.6	1,736.9	1,805.1	1,842.7	1,709.1	1,804.2	1,486.0	1,478.1	1,339.6
Beginning Equity Capital	9,016											
PV of Equity Capital	15,991											
PV of Continuing Value	28,132											
Total Equity Value	53,139											
FD shares outstanding ('000)	75,547											
Per share value	710,000											


Source: Company data, Morgan Stanley Research estimates (E)

Risk Reward – SK hynix (000660.KS)

Robust commodity supercycle supported by AI in 2026

PRICE TARGET W1,700,000

Our base case price target is derived from our residual income valuation model. We assume a cost of equity of 11.5% (5% risk-free rate, 6.5% risk premium, 1.0 beta) and a terminal growth rate of 3%.




Source: Refinitiv, Morgan Stanley Research

OVERWEIGHT THESIS

- The commodity cycle is likely to stay robust in 2026 and into 2027, we believe, thanks to AI inference demand growth.
- We do not think this is the time to rotate out of outperformers such as SK hynix.
- Our price target implies c.4x 2027e P/E, in line with the stock's commodity cycle peak.


RISK REWARD CHART



Key: — Historical Stock Performance ● Current Stock Price ◆ Price Target

Source: Refinitiv, Morgan Stanley Research

Consensus Rating Distribution



● MS Rating

Source: Refinitiv, Morgan Stanley Research

Risk Reward Themes

Pricing Power: *Positive*
Secular Growth: *Positive*
Self-help: *Positive*

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

W2,580,000

Memory TAM grows; AI HBM leadership is extended

We see a stronger-than-expected commodity cycle for DRAM and NAND, supported by AI inference demand moving into 2027. Our bull scenario valuation implies 14x P/E on HBM contribution and 5x P/E on commodity memory based on 2027e EPS.

BASE CASE

W1,700,000

Commodity up-cycle

We see a stronger-than-expected commodity cycle for DRAM and NAND, supported by AI inference demand in 2026 and into 2027. HBM competition risks remain but are being increasingly recognized by the market. Our price target implies 3.8x 2027e P/E, underpinned by Hynix's leadership in HBM and a sharp rebound in commodity memory pricing.

BEAR CASE

W400,000

Weaker macro climate, HBM competition intensifies

We assume a period of AI computing digestion in 2026, leading to more aggressive margin erosion for the HBM segment owing to competitors' aggressive capacity expansion and pricing competition, as well as faster DRAM pricing drops with muted demand recovery. In our long-term assumptions for the bear case, we assume margin erosion for the HBM segment amid rising competition. Our bear case valuation implies 1.3x 2026e P/B.

Risk Reward – SK hynix (000660.KS)


KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Revenue Growth (%)	125.9	650.8	1,386.9	663.6
OPM (%)	48.6	82.6	83.7	79.8
CAPEX (%)	(2,840,890.0)	(6,030,605.9)	(7,307,828.6)	(7,581,454.7)

INVESTMENT DRIVERS


- Supply and demand outlook for DRAM and NAND
- Global demand for data center products (server DRAM, enterprise SSD)

GLOBAL REVENUE EXPOSURE



Source: Morgan Stanley Research Estimate
View explanation of regional hierarchies [here](#)

MS ALPHA MODELS



Source: Refinitiv, FactSet, Morgan Stanley Research; 1 is the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING


RISKS TO UPSIDE

- ASP growth re-acceleration
- Preemptive production cuts or disruption
- Sharp demand improvements
- Significant increase in capital returns

RISKS TO DOWNSIDE

- End demand turning weaker than expected
- Rising competition for DDR5, causing overspending on the supply side.
- Inventories remaining elevated at cloud and Chinese smartphone customers.





OWNERSHIP POSITIONING



Source: Refinitiv, Morgan Stanley Research

MS ESTIMATES VS. CONSENSUS

FY Dec 2026e



Source: Refinitiv, Morgan Stanley Research

Risk Reward – Samsung Electronics (005930.KS) Top Pick


Improvement in HBM + commodity up-cycle = stock re-rating

PRICE TARGET **W362,000**


Our base case price target is derived from our residual income valuation model. At our price target, the 2027e P/B multiple would be c.2x, in line with the stock's commodity cycle peak of 2.0x. We assume an 11.5% cost of equity, based on a beta of 1.0, and our terminal growth rate assumption is 3%.

Consensus Price Target Distribution

Source: Refinitiv, Morgan Stanley Research



RISK REWARD CHART




Key: — Historical Stock Performance ● Current Stock Price ◆ Price Target

Source: Refinitiv, Morgan Stanley Research

OVERWEIGHT THESIS

- We expect pricing to move well beyond historical mean reversion from here, on a price per gigabit basis, which should create further incremental tailwinds to forecasts for 2026-27. The commodity cycle is better than feared thanks to AI inference demand.
- We expect HBM improvement and potential market share gains from 2026.
- Rising prices and falling inventories, along with supply cuts and improving demand, are positives for margin recovery in the memory business.
- The implied 2027e P/B at our price target is 2x, aligning with the stock's commodity cycle peak of 2.0x.

Consensus Rating Distribution



MS Rating

Source: Refinitiv, Morgan Stanley Research

Risk Reward Themes

Secular Growth: *Positive*
Self-help: *Positive*

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

W430,000

HPC drives sustained memory pricing

High-performance computing drives sustained memory pricing strength and capital returns exceed expectations: We assume continued robust demand growth for memory and OLED with stabilization in the macro environment and China's failure to successfully enter the memory semiconductor industry. Our bull case valuation is based 14x HBM, 5x conventional memory and 10x P/E for the rest of the group profitability.

BASE CASE

W362,000

Commodity cycle recovery

We expect favorable commodity pricing trends to bode well for Samsung's share price with improving HBM market share and new product execution. Our base case reflects our relatively conservative view on HBM progress this year, but sequential improvement in 2026 and beyond support Samsung's earnings growth and stock re-rating amid low expectations. 2027e implied P/B at c.2x.

BEAR CASE

W120,000

Slower memory consumption, competition from China

Prolonged consumption slowdown and Chinese competition: Weaker macro conditions stall near-term global consumption of IT products. Chinese competition disrupts OLED and memory profitability. We also assume demand for server content slumps amid a slowdown in data center growth. DRAM and NAND ASP price hikes are smaller than expected in 2026. Our bear case valuation is at 1.1x 2026e P/B, closer to historical trough levels.

Risk Reward – Samsung Electronics (005930.KS)


KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Revenue Growth (%)	10.9	120.4	37.5	6.3
OPM (%)	13.1	58.2	62.4	56.7
CAPEX (%)	(4,730,271.0)	(14,641,899.4)	(13,561,744.5)	(13,229,277.6)

INVESTMENT DRIVERS

- Supply-demand outlook for memory and OLED
- Smartphone market share and margins
- Capital returns of 50% FCF

GLOBAL REVENUE EXPOSURE



Source: Morgan Stanley Research Estimate
View explanation of regional hierarchies [here](#)

MS ALPHA MODELS

1/5
MOST
3 Month
Horizon

Source: Refinitiv, FactSet, Morgan Stanley Research; 1 is the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- New technological developments, especially in memory and foldable displays
- Longer-lasting memory up-cycle from AI and hyperscale data center growth

RISKS TO DOWNSIDE

- Product and memory cycle, including Apple and new Chinese smartphone competition
- Earnings growth concentration in semiconductors





OWNERSHIP POSITIONING

Inst. Owners, % Active 78.3%

Source: Refinitiv, Morgan Stanley Research

MS ESTIMATES VS. CONSENSUS

FY Dec 2026e



◆ Mean ◆ Morgan Stanley Estimates
Source: Refinitiv, Morgan Stanley Research

Risk Reward – Samsung Electronics (005935.KS)

Overweight with discount to common shares to narrow

PRICE TARGET W289,600


Our price target for the preferred shares is based on our PT for the common shares and our assumption that the preferred shares' trade at a 20% discount to the common shares, based on the past 2-year average. We assume an 11.5% cost of equity, based on a beta of 1.0, and our terminal growth rate assumption is 3%.

Consensus Price Target Distribution

Source: Refinitiv, Morgan Stanley Research



RISK REWARD CHART




Key: — Historical Stock Performance ● Current Stock Price ◆ Price Target

Source: Refinitiv, Morgan Stanley Research

OVERWEIGHT THESIS

- We expect pricing to move well beyond historical mean reversion from here, on a price per gigabit basis, which should create further incremental tailwinds to forecasts for 2026-27. The commodity cycle is better than feared thanks to AI inference demand.
- We expect HBM improvement and potential market share gains from 2026.
- Rising prices and falling inventories, along with supply cuts and improving demand, are positives for margin recovery in the memory business.
- Common shares' implied 2027e P/E is 5x. We assume that the preferred shares' discount to the common shares is 20% (the past 2-year average).

Consensus Rating Distribution



Source: Refinitiv, Morgan Stanley Research

Risk Reward Themes

Secular Growth: *Positive*
Self-help: *Positive*

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

W344,000

20% discount to bull case value for common shares

High-performance computing drives sustained memory pricing strength and capital returns exceed expectations: We assume continued robust demand growth for memory and OLED with stabilization in the macro environment and China's failure to successfully enter the memory semiconductor industry.

BASE CASE

W289,600

20% discount to base case value for common shares

We expect favorable commodity pricing trends to bode well for Samsung's share price with improving HBM market share and new product execution. Our base case reflects our optimistic view on commodity upcycle and stock re-rating with LTA to potentially smooth the cyclical of memory sector.

BEAR CASE

W96,000

20% discount to bear case value for common shares

Prolonged consumption slowdown and Chinese competition: Weaker macro conditions stall near-term global consumption of IT products. Chinese competition disrupts OLED and memory profitability. We also assume demand for server content slumps amid a slowdown in data center growth. DRAM and NAND ASP price hikes are smaller than expected in 2026/2027.

Risk Reward – Samsung Electronics (005935.KS)


KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Revenue Growth (%)	10.9	120.4	37.5	6.3
OPM (%)	13.1	58.2	62.4	56.7
CAPEX (%)	(4,730,271.0)	(14,641,899.4)	(13,561,744.5)	(13,229,277.6)

INVESTMENT DRIVERS

- Supply-demand outlook for memory and OLED
- Smartphone market share and margins
- Capital returns of 50% FCF

GLOBAL REVENUE EXPOSURE



Source: Morgan Stanley Research Estimate
View explanation of regional hierarchies [here](#)

MS ALPHA MODELS

1/5
MOST

3 Month
Horizon

Source: Refinitiv, FactSet, Morgan Stanley Research; 1 is the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- New technology developments, especially in memory and foldable displays
- Longer-lasting memory up-cycle from AI and hyperscale data center growth
- Further narrowing of discount to common shares

RISKS TO DOWNSIDE

- Product and memory cycle, including Apple and new Chinese smartphone competition
- Earnings growth concentration in semiconductors
- Deepening of discount to common shares


OWNERSHIP POSITIONING

Inst. Owners, % Active 76.8%

Source: Refinitiv, Morgan Stanley Research

MS ESTIMATES VS. CONSENSUS

FY Dec 2026e



Source: Refinitiv, Morgan Stanley Research

Risk Reward – Samsung Electro-Mechanics (009150.KS)

Dual AI opportunity – ABF and MLCC

PRICE TARGET W710,000


Our base case price target is derived from our residual income valuation model. We apply a 10.5% cost of equity (beta 1.0, equity risk premium of 6.5% and risk-free rate of 4%) and a 5% terminal growth rate. Our base year is 2026 and our forecast period is 2026-35.

Consensus Price Target Distribution

Source: Refinitiv, Morgan Stanley Research



RISK REWARD CHART




Key: — Historical Stock Performance ● Current Stock Price ◆ Price Target

Source: Refinitiv, Morgan Stanley Research

OVERWEIGHT THESIS

- We see long-term improvement in SEMCO's operating profile – multi-layer core ABF capacity sold out from ASIC customers and accelerating AI penetration in MLCC, driving substantial content from rising complexity.
- With the AI computing opportunity and the prospect of a cyclical recovery from 2H26, earnings expectations appear to be lagging significantly based on the current outlook.
- We think the share price will look to discount accelerating EPS estimate revisions.
- NTM P/B is below its mid-cycle average, indicating continued valuation upside on top of earnings.
- Our price target implies 2027e P/B of 4.0x.

Consensus Rating Distribution



Source: Refinitiv, Morgan Stanley Research

Risk Reward Themes

Secular Growth: *Positive*

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

W730,000

4.1x 2027e P/B

ABF opportunities from major US CSPs drive meaningful content growth, and we estimate future ASIC content and volume growth will continue to outpace. Customer diversification into non-mobile companies also lowers earnings volatility and improves profitability. Our bull case implies a 2027e P/B at historical peak of 4.1x. We believe this premium is supported by what we view as an unprecedented supercycle in the underlying industry dynamics.

BASE CASE

W710,000

4.0x 2027e P/B

We see a number of catalysts ahead for SEMCO, including a large ABF opportunity in ASIC chip wins and FCF inflection. Although the stock has moved higher, our long-term conviction in custom silicon opportunities via ABF substrates has grown substantially.

BEAR CASE

W200,000

1.1x 2027e P/B

MLCC localization trend intensifies with more competition from Chinese suppliers resulting in increased pricing pressure. Substrate business remains soft amid weak macro climate and tepid end demand. China's smartphone demand slows, coupled with growing geopolitical tension/macro-economic uncertainty driven by the US imposition of tariffs. We reflect slower content growth for MLCC and substrates.

Risk Reward – Samsung Electro-Mechanics (009150.KS)


KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Revenue Growth (%)	9.9	23.0	28.6	24.4
OPM (%)	8.1	12.9	16.8	18.5
CAPEX (%)	8.7	10.7	9.7	10.4

INVESTMENT DRIVERS


- AI servers, high-speed networking, Auto 2.0, and smartphone content growth
- Customer diversification away from Samsung Electronics, targeting major CSPs
- MLCC pricing/margin trend

GLOBAL REVENUE EXPOSURE



Source: Morgan Stanley Research Estimate
View explanation of regional hierarchies [here](#)

MS ALPHA MODELS



Source: Refinitiv, FactSet, Morgan Stanley Research; 1 is the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- Price hikes driven by MLCC shortage
- Better-than-expected smartphone demand
- Higher MLCC margin from better yield and product mix
- Strong tailwind from China's consumption-friendly policies

RISKS TO DOWNSIDE

- Significant pullback in Samsung's mobile flagship product cycle and SEMCO's lower market share
- Execution risk in penetrating Chinese smartphone customers
- Consumer demand headwinds in 2026



OWNERSHIP POSITIONING



Source: Refinitiv, Morgan Stanley Research

MS ESTIMATES VS. CONSENSUS

FY Dec 2026e



◆ Mean ◆ Morgan Stanley Estimates
Source: Refinitiv, Morgan Stanley Research

Appendix – CPU vs. GPU, Inference vs. Agent

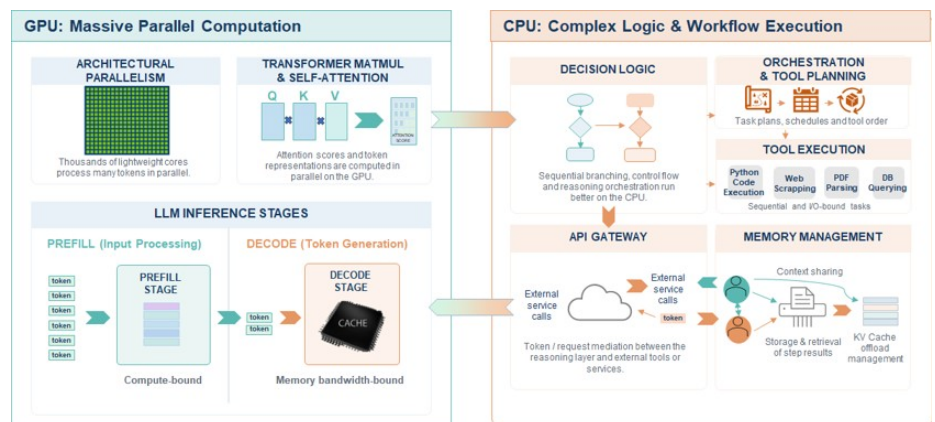
CPU vs. GPU in a nutshell

CPUs handle sequential reasoning and multi-step processes with branching logic, while GPUs execute parallel processing to analyze millions of data, images or pictures simultaneously. Agentic AI systems need CPUs that efficiently manage long context windows, maintain state across reasoning chains, and rapidly switch between computational tasks. Agentic AI systems are autonomous and require immense CPU power, not just GPU muscle.

GPU workloads are thousands of cores doing matrix math but agentic AI is coordination-heavy – agents talking to agents, tools calling tools, decisions made in sequence – which is CPU-heavy. The introduction of multiple data center CPU offering signals a shift in AI data center design as orchestration, inference, and real-time execution move to the center of next-generation workloads. The AI hardware landscape is entering a new phase of complexity as the industry learns that no single processor architecture can handle every workload efficiently for agentic AI systems.

GPU does the math while CPUs manage the rapidly expanding and more intricate agentic control layer. Where traditional AI models excel at pattern recognition and generation – tasks that leverage the parallel processing power of GPUs – agentic AI systems require the sequential reasoning and branching logic that CPUs handle more efficiently. In simple terms, this is the difference between analyzing millions of data simultaneously and methodically working through a multi-step business process with conditional branches and decision trees. GPUs are often excessive for simple inference, whereas CPUs efficiently handle tasks requiring lower latency and less parallel computation, reducing costs. Autonomous agents require more multi-step workflows, which are heavily dependent on CPU processing for tasks such as searching and validating data. But GPUs aren't going anywhere anytime soon – for general-purpose AIs, they are most efficient for many applications.

Exhibit 60: The LLM agent workflow: GPU “thinking” vs. CPU “doing”

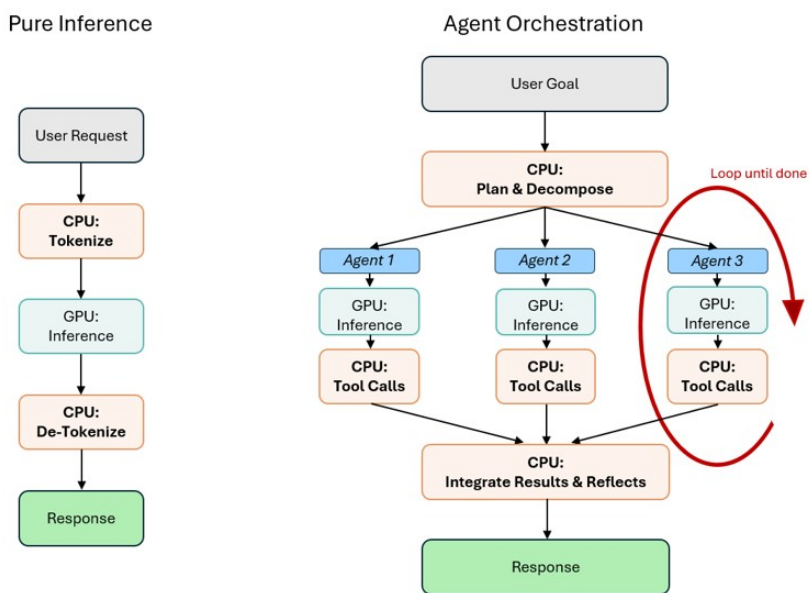


Source: Morgan Stanley Research

Pure inference vs. agent orchestration

In pure inference, the CPU's role is limited to tokenizing the request, passing it to the GPU, and de-tokenizing the result. The GPU does the heavy lifting. In agent orchestration, the CPU becomes the command layer: it plans, decomposes goals into sub-tasks, fans out to parallel sub-agents, manages tool calls and API requests, monitors token streams, integrates results, and runs reflection loops. CPU performance ultimately dictates how much of the GPU is utilized. That is why agentic AI should create structural demand not only for GPUs, but also for higher-performance server CPUs and the broader memory and I/O stack that keeps those agents running at machine speed.

Exhibit 61: Pure inference vs. agent orchestration workflow



Source: Morgan Stanley Research

Morgan Stanley is acting as exclusive financial advisor to AP Grange Holdings, LLC ("AP Grange") in relation to a definitive agreement with Intel Corporation ("Intel"), for Intel to repurchase the 49% equity interest in the joint venture related to Intel's Fab 34 in Ireland not held by Intel, as announced on April 1, 2026. The transaction is subject to customary closing conditions. AP Grange has agreed to pay fees to Morgan Stanley for its financial advisory services. Please refer to the notes at the end of this report.

Risk Reward Reference links

1. View explanation of Options Probabilities methodology - [Options_Probabilities_Exhibit_Link.pdf](#)
2. View descriptions of Risk Rewards Themes - [RR_Themes_Exhibit_Link.pdf](#)
3. View explanation of regional hierarchies - [GEG_Exhibit_Link.pdf](#)
4. View explanation of Theme/Exposure methodology - [ESG_Sustainable_Solutions_External_Link.pdf](#)
5. View explanation of HERS methodology - [ESG_HERS_External_Link.pdf](#)

Disclosure Section

Share prices for other stocks mentioned, as of 16 April 2026: Broadcom Inc \$398.47, Winbond Electronics Corp TWD 87.80.

The information and opinions in Morgan Stanley Research were prepared or are disseminated by Morgan Stanley Asia Limited (which accepts the responsibility for its contents) and/or Morgan Stanley Asia (Singapore) Pte. (Registration number 199206298Z) and/or Morgan Stanley Asia (Singapore) Securities Pte Ltd (Registration number 200008434H), regulated by the Monetary Authority of Singapore (which accepts legal responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research), and/or Morgan Stanley Taiwan Limited and/or Morgan Stanley & Co International plc, Seoul Branch, and/or Morgan Stanley Australia Limited (A.B.N. 67 003 734 576, holder of Australian financial services license No. 233742, which accepts responsibility for its contents), and/or Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813, which accepts responsibility for its contents), and/or Morgan Stanley India Company Private Limited having Corporate Identification No (CIN) U22990MH1998PTC115305, regulated by the Securities and Exchange Board of India ("SEBI") and holder of licenses as a Research Analyst (SEBI Registration No. INH000001105); Stock Broker (SEBI Stock Broker Registration No. INZ000244438), Merchant Banker (SEBI Registration No. INM000011203), and depository participant with National Securities Depository Limited (SEBI Registration No. IN-DP-NSDL-567-2021) having registered office at Altimus, Level 39 & 40, Pandurang Budhkar Marg, Worli, Mumbai 400018, India; Telephone no. +91-22-61181000; Compliance Officer Details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: tejarshi.hardas@morganstanley.com; Grievance officer details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: msic-compliance@morganstanley.com which accepts the responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research, and their affiliates (collectively, "Morgan Stanley"). Morgan Stanley India Company Private Limited (MSICPL) may use AI tools in providing research services. All recommendations contained herein are made by the duly qualified research analysts.

For important disclosures, stock price charts and equity rating histories regarding companies that are the subject of this report, please see the Morgan Stanley Research Disclosure Website at www.morganstanley.com/eqr/disclosures/webapp/generalresearch, or contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY, 10036 USA.

For valuation methodology and risks associated with any recommendation, rating or price target referenced in this research report, please contact the Client Support Team as follows: US/Canada +1 800 303-2495; Hong Kong +852 2848-5999; Latin America +1 718 754-5444 (U.S.); London +44 (0)20-7425-8169; Singapore +65 6834-6860; Sydney +61 (0)2-9770-1505; Tokyo +81 (0)3-6836-9000. Alternatively you may contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY 10036 USA.

Analyst Certification

The following analysts hereby certify that their views about the companies and their securities discussed in this report are accurately expressed and that they have not received and will not receive direct or indirect compensation in exchange for expressing specific recommendations or views in this report: Shane Brett; Charlie Chan; Daisy Dai, CFA; Yoshihito Hasegawa; Cindy Huang; Howard Kao; Ryan Kim; Shawn Kim; Duan Liu; Joseph Moore; Nigel van Putten; Shoji Sato; Sharon Shih; Lee Simpson; Suzune Tamura, CFA; Erik W Woodring; Derrick Yang; Tiffany Yeh; Daniel Yen, CFA; Kazuo Yoshikawa, CFA.

Global Research Conflict Management Policy

Morgan Stanley Research has been published in accordance with our conflict management policy, which is available at www.morganstanley.com/institutional/research/conflictolicies. A Portuguese version of the policy can be found at www.morganstanley.com.br

Important Regulatory Disclosures on Subject Companies

The analyst or strategist (or a household member) identified below owns the following securities (or related derivatives): Cindy Huang - ASML Holding NV (common or preferred stock); Erik W Woodring - NVIDIA Corp. (common or preferred stock).

As of March 31, 2026, Morgan Stanley beneficially owned 1% or more of a class of common equity securities of the following companies covered in Morgan Stanley Research: Advanced Micro Devices, AirTAC International, AP Memory Technology Corp, Applied Materials Inc., Arm Holdings plc, Asia Vital Components Co. Ltd., ASML Holding NV, Auras Technology Co Ltd, Bizlink, Broadcom Inc., Cadence Design Systems Inc, Delta Electronics Inc., Global Unichip Corp, Gold Circuit Electronics Ltd., Hiwin Technologies Corp., Ibiden, Infineon Technologies AG, Intel Corporation, KIOXIA Holdings, KLA Corp, Lotes Co. Ltd., MEC, Micron Technology Inc., Montage Technology Co Ltd, Nan Ya PCB, Nitto Boseki, NVIDIA Corp., Samsung Electro-Mechanics, SanDisk Corporation., Seagate Technology, SK hynix, STMicroelectronics NV, Synopsys Inc., TDK, Tokyo Electron, TSMC, Ulvac, Unimicron, Western Digital, Winbond Electronics Corp, Wiywynn Corp, Wonik IPS Co Ltd, WPG Holdings.

Within the last 12 months, Morgan Stanley managed or co-managed a public offering (or 144A offering) of securities of Broadcom Inc., Intel Corporation, KIOXIA Holdings, Micron Technology Inc., Montage Technology Co Ltd, Samsung Electronics, Seagate Technology, SK hynix, Wiywynn Corp.

Within the last 12 months, Morgan Stanley has received compensation for investment banking services from Advanced Micro Devices, ASM International NV, BE Semiconductor Industries NV, Broadcom Inc., Intel Corporation, KIOXIA Holdings, Lenovo, Micron Technology Inc., Montage Technology Co Ltd, Seagate Technology, SK hynix, STMicroelectronics NV, Tokyo Electron, Wiywynn Corp, Yageo Corp..

In the next 3 months, Morgan Stanley expects to receive or intends to seek compensation for investment banking services from Advanced Micro Devices, AirTAC International, AP Memory Technology Corp, Applied Materials Inc., Arm Holdings plc, Asia Vital Components Co. Ltd., ASM International NV, ASML Holding NV, BE Semiconductor Industries NV, Bizlink, Broadcom Inc., Cadence Design Systems Inc, Delta Electronics Inc., FIT Hon Teng Ltd, Gold Circuit Electronics Ltd., Hon Hai Precision, Ibiden, Infineon Technologies AG, Intel Corporation, KIOXIA Holdings, KLA Corp, Lenovo, Lite-On Technology, MEC, Micron Technology Inc., Montage Technology Co Ltd, Murata Manufacturing, Nitto Boseki, NVIDIA Corp., Renesas Electronics, Samsung Electronics, SanDisk Corporation., Seagate Technology, SK hynix, STMicroelectronics NV, Synopsys Inc., TDK, Tokyo Electron, TSMC, Ulvac, Unimicron, VAT Group AG, Western Digital, Winbond Electronics Corp, Wiywynn Corp, WPG Holdings, Yageo Corp.

Within the last 12 months, Morgan Stanley has received compensation for products and services other than investment banking services from Advanced Micro Devices, Applied Materials Inc., ASM International NV, BE Semiconductor Industries NV, Broadcom Inc., Cadence Design Systems Inc, Hon Hai Precision, Infineon Technologies AG, Intel Corporation, Lenovo, Micron Technology Inc., NVIDIA Corp., Seagate Technology, Synopsys Inc., Tokyo Electron, TSMC, Winbond Electronics Corp, Yageo Corp..

Within the last 12 months, Morgan Stanley has provided or is providing investment banking services to, or has an investment banking client relationship with, the following company: Advanced Micro Devices, AirTAC International, AP Memory Technology Corp, Applied Materials Inc., Arm Holdings plc, Asia Vital Components Co. Ltd., ASM International NV, ASML Holding NV, BE Semiconductor Industries NV, Bizlink, Broadcom Inc., Cadence Design Systems Inc, Delta Electronics Inc., FIT Hon Teng Ltd, Gold Circuit Electronics Ltd., Hon Hai Precision, Ibiden, Infineon Technologies AG, Intel Corporation, KIOXIA Holdings, KLA Corp, Lenovo, Lite-On Technology, Micron Technology Inc., Montage Technology Co Ltd, Murata Manufacturing, NVIDIA Corp., Renesas Electronics, Samsung Electronics, SanDisk Corporation., Seagate Technology, SK hynix, STMicroelectronics NV, Synopsys Inc., TDK, Tokyo Electron, TSMC, Ulvac, Unimicron, VAT Group AG, Western Digital, Winbond Electronics Corp, Wiywynn Corp, WPG Holdings, Yageo Corp..

Within the last 12 months, Morgan Stanley has either provided or is providing non-investment banking, securities-related services to and/or in the past has entered into an agreement to provide services or has a client relationship with the following company: Advanced Micro Devices, Applied Materials Inc., ASM International NV, BE Semiconductor Industries NV, Broadcom Inc., Cadence Design Systems Inc, Hon Hai Precision, Infineon Technologies AG, Intel Corporation, Lenovo, Micron Technology Inc., Montage Technology Co Ltd, NVIDIA Corp., Renesas Electronics, Seagate Technology, Synopsys Inc., Tokyo Electron, TSMC, Winbond Electronics Corp, Yageo Corp..

The equity research analysts or strategists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality of research, investor client feedback, stock picking, competitive factors, firm revenues and overall investment banking revenues. Equity Research analysts' or strategists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

Morgan Stanley and its affiliates do business that relates to companies/instruments covered in Morgan Stanley Research, including market making, providing liquidity, fund management, commercial banking, extension of credit, investment services and investment banking. Morgan Stanley sells to and buys from customers the securities/instruments of companies covered in Morgan Stanley Research on a principal basis. Morgan Stanley may have a position in the debt of the Company or instruments discussed in this report. Morgan Stanley trades or may trade as principal in the debt securities (or in related derivatives) that are the subject of the debt research report.

Certain disclosures listed above are also for compliance with applicable regulations in non-US jurisdictions.

STOCK RATINGS

Morgan Stanley uses a relative rating system using terms such as Overweight, Equal-weight, Not-Rated or Underweight (see definitions below). Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold and sell. Investors should carefully read the definitions of all ratings used in Morgan Stanley Research. In addition, since Morgan Stanley Research contains more complete information concerning the analyst's views, investors should carefully read Morgan Stanley Research, in its entirety, and not infer the contents from the rating alone. In any case, ratings (or research) should not be used or relied upon as investment advice. An investor's decision to buy or sell a stock should depend on individual circumstances (such as the investor's existing holdings) and other considerations.

Global Stock Ratings Distribution

(as of March 31, 2026)

The Stock Ratings described below apply to Morgan Stanley's Fundamental Equity Research and do not apply to Debt Research produced by the Firm.

For disclosure purposes only (in accordance with FINRA requirements), we include the category headings of Buy, Hold, and Sell alongside our ratings of Overweight, Equal-weight, Not-Rated and Underweight. Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold, and sell but represent recommended relative weightings (see definitions below). To satisfy regulatory requirements, we correspond Overweight, our most positive stock rating, with a buy recommendation; we correspond Equal-weight and Not-Rated to hold and Underweight to sell recommendations, respectively.

Stock Rating Category	Coverage Universe		Investment Banking Clients (IBC)			Other Material Investment Services Clients (MISC)	
	Count	% of Total	Count	% of Total IBC	% of Rating Category	Count	% of Total Other MISC
Overweight/Buy	1534	42%	461	50%	30%	698	43%
Equal-weight/Hold	1573	43%	372	40%	24%	716	44%
Not-Rated/Hold	4	0%	1	0%	25%	1	0%
Underweight/Sell	568	15%	89	10%	16%	209	13%
Total	3,679		923			1624	

Data include common stock and ADRs currently assigned ratings. Investment Banking Clients are companies from whom Morgan Stanley received investment banking compensation in the last 12 months. Due to rounding off of decimals, the percentages provided in the "% of total" column may not add up to exactly 100 percent.

Analyst Stock Ratings

Overweight (O). The stock's total return is expected to exceed the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Equal-weight (E). The stock's total return is expected to be in line with the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Not-Rated (NR). Currently the analyst does not have adequate conviction about the stock's total return relative to the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Underweight (U). The stock's total return is expected to be below the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Unless otherwise specified, the time frame for price targets included in Morgan Stanley Research is 12 to 18 months.

Analyst Industry Views

Attractive (A): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be attractive vs. the relevant broad market benchmark, as indicated below.


In-Line (I): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be in line with the relevant broad market benchmark, as indicated below.

Cautious (C): The analyst views the performance of his or her industry coverage universe over the next 12-18 months with caution vs. the relevant broad market benchmark, as indicated below.

Benchmarks for each region are as follows: North America - S&P 500; Latin America - relevant MSCI country index or MSCI Latin America Index; Europe - MSCI Europe; Japan - TOPIX; Asia - relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Stock Price, Price Target and Rating History (See Rating Definitions)

Samsung Electro-Mechanics (009150.KS) - As of 04/17/26 GMT in KRW
 Industry : S. Korea Technology



Stock Rating History: 4/1/21 : O/A; 5/12/21 : E/A; 7/19/21 : E/I; 8/12/21 : E/C; 6/1/22 : U/C; 10/4/22 : U/A; 11/9/22 : O/A; 7/15/24 : E/A; 7/21/24 : E/I; 9/15/24 : E/C; 12/18/24 : O/C; 4/7/25 : E/C; 6/13/25 : E/I; 7/31/25 : O/I; 9/21/25 : O/A


Price Target History: 3/1/21 : 250000; 5/12/21 : 180000; 10/19/21 : 170000; 10/27/21 : 160000; 1/24/22 : 180000; 3/18/22 : 170000; 6/1/22 : 140000; 11/9/22 : 180000; 10/11/23 : 170000; 10/27/23 : 160000; 7/15/24 : 150000; 12/18/24 : 170000; 4/7/25 : 133000; 7/31/25 : 198000; 9/10/25 : 220000; 10/10/25 : 236000; 11/5/25 : 300000; 12/1/25 : 309000; 2/23/26 : 430000; 4/2/26 : 500000

Source: Morgan Stanley Research Date Format : MM/DD/YY Price Target --- No Price Target Assigned (NA)
 Stock Price (Not Covered by Current Analyst) — Stock Price (Covered by Current Analyst) —
 Stock and Industry Ratings (abbreviations below) appear as ♦ Stock Rating/Industry View
 Stock Ratings: Overweight (O) Equal-weight (E) Underweight (U) Not-Rated (NR) No Rating Available (NA)
 Industry View: Attractive (A) In-line (I) Cautious (C) No Rating (NR)

Effective January 13, 2014, the stocks covered by Morgan Stanley Asia Pacific will be rated relative to the analyst's industry (or industry team's) coverage.

Effective January 13, 2014, the industry view benchmarks for Morgan Stanley Asia Pacific are as follows: relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Samsung Electronics (005930.KS) - As of 04/17/26 GMT in KRW
 Industry : S. Korea Technology



Stock Rating History: 4/1/21 : O/A; 7/19/21 : O/I; 8/12/21 : O/C; 10/4/22 : O/A; 7/21/24 : O/I; 9/15/24 : O/C; 6/13/25 : O/I; 9/21/25 : O/A


Price Target History: 2/25/21 : 115000; 5/18/21 : 93000; 6/8/21 : 96000; 8/12/21 : 89000; 9/15/21 : 95000; 12/3/21 : 97000; 3/18/22 : 95000; 4/28/22 : 85000; 6/10/22 : 80000; 7/5/22 : 75000; 7/22/22 : 70000; 9/17/22 : 68000; 3/21/23 : 70000; 5/30/23 : 90000; 7/7/23 : 95000; 3/22/24 : 97000; 4/16/24 : 101000; 6/6/24 : 105000; 9/15/24 : 76000; 12/18/24 : 65000; 3/19/25 : 70000; 8/1/25 : 86000; 9/21/25 : 97000; 10/6/25 : 111000; 10/23/25 : 120000; 10/30/25 : 139000; 11/5/25 : 144000; 1/6/26 : 170000; 1/30/26 : 210000; 2/24/26 : 248000; 3/18/26 : 251000

Source: Morgan Stanley Research Date Format : MM/DD/YY Price Target --- No Price Target Assigned (NA)
 Stock Price (Not Covered by Current Analyst) — Stock Price (Covered by Current Analyst) —
 Stock and Industry Ratings (abbreviations below) appear as ♦ Stock Rating/Industry View
 Stock Ratings: Overweight (O) Equal-weight (E) Underweight (U) Not-Rated (NR) No Rating Available (NA)
 Industry View: Attractive (A) In-line (I) Cautious (C) No Rating (NR)

Effective January 13, 2014, the stocks covered by Morgan Stanley Asia Pacific will be rated relative to the analyst's industry (or industry team's) coverage.

Effective January 13, 2014, the industry view benchmarks for Morgan Stanley Asia Pacific are as follows: relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Samsung Electronics (005935.KS) - As of 04/17/26 GMT in KRW
 Industry : S. Korea Technology



Stock Rating History: 4/1/21 : O/A; 7/19/21 : O/I; 8/12/21 : O/C; 10/4/22 : O/A; 7/21/24 : O/I; 9/15/24 : O/C; 6/13/25 : O/I; 9/21/25 : O/A


Price Target History: 1/12/21 : 90000; 5/18/21 : 80000; 6/8/21 : 84000; 8/12/21 : 77000; 9/15/21 : 82000; 12/3/21 : 83000; 3/18/22 : 82000; 4/28/22 : 73000; 6/10/22 : 69000; 7/5/22 : 65000; 7/22/22 : 60000; 9/17/22 : 58000; 3/21/23 : 60000; 5/30/23 : 77000; 7/7/23 : 82000; 1/20/25 : 52000; 3/19/25 : 60900; 8/1/25 : 74820; 9/21/25 : 84390; 10/8/25 : 96570; 10/23/25 : 102000; 10/30/25 : 120930; 11/5/25 : 125280; 1/8/26 : 147900; 1/30/26 : 182000; 2/24/26 : 210000; 3/18/26 : 213350

Source: Morgan Stanley Research Date Format : MM/DD/YY Price Target -- No Price Target Assigned (NA)
 Stock Price (Not Covered by Current Analyst) — Stock Price (Covered by Current Analyst) —
 Stock and Industry Ratings (abbreviations below) appear as ♦ Stock Rating/Industry View
 Stock Ratings: Overweight (O) Equal-weight (E) Underweight (U) Not-Rated (NR) No Rating Available (NA)
 Industry View: Attractive (A) In-line (I) Cautious (C) No Rating (NR)

Effective January 13, 2014, the stocks covered by Morgan Stanley Asia Pacific will be rated relative to the analyst's industry (or industry team's) coverage.

Effective January 13, 2014, the industry view benchmarks for Morgan Stanley Asia Pacific are as follows: relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

SK hynix (000660.KS) - As of 04/17/26 GMT in KRW
 Industry : S. Korea Technology



Stock Rating History: 4/1/21 : O/A; 7/19/21 : O/I; 8/12/21 : U/C; 12/3/21 : E/C; 2/11/22 : O/C; 7/22/22 : E/C; 10/4/22 : O/A; 7/21/24 : O/I; 9/15/24 : U/C; 3/19/25 : E/C; 6/13/25 : E/I; 9/21/25 : O/A


Price Target History: 2/25/21 : 174000; 5/18/21 : 146000; 6/8/21 : 156000; 8/12/21 : 80000; 9/15/21 : 86000; 12/3/21 : 110000; 12/23/21 : 125000; 1/24/22 : 130000; 1/28/22 : 136000; 2/11/22 : 155000; 3/18/22 : 150000; 4/27/22 : 130000; 6/10/22 : 120000; 7/5/22 : 110000; 7/22/22 : 105000; 10/4/22 : 130000; 12/7/22 : 120000; 3/21/23 : 110000; 5/30/23 : 140000; 7/7/23 : 170000; 9/21/23 : 210000; 3/22/24 : 230000; 4/16/24 : 270000; 6/6/24 : 300000; 7/25/24 : 260000; 9/15/24 : 120000; 10/24/24 : 130000; 12/18/24 : 135000; 1/20/25 : 150000; 3/19/25 : 230000; 4/7/25 : 210000; 6/13/25 : 250000; 7/24/25 : 260000; 9/21/25 : 410000; 10/8/25 : 480000; 10/23/25 : 570000; 10/29/25 : 630000; 11/5/25 : 730000; 1/2/26 : 840000; 1/30/26 : 1100000; 3/18/26 : 1300000

Source: Morgan Stanley Research Date Format : MM/DD/YY Price Target -- No Price Target Assigned (NA)
 Stock Price (Not Covered by Current Analyst) — Stock Price (Covered by Current Analyst) —
 Stock and Industry Ratings (abbreviations below) appear as ♦ Stock Rating/Industry View
 Stock Ratings: Overweight (O) Equal-weight (E) Underweight (U) Not-Rated (NR) No Rating Available (NA)
 Industry View: Attractive (A) In-line (I) Cautious (C) No Rating (NR)

Effective January 13, 2014, the stocks covered by Morgan Stanley Asia Pacific will be rated relative to the analyst's industry (or industry team's) coverage.

Effective January 13, 2014, the industry view benchmarks for Morgan Stanley Asia Pacific are as follows: relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Wonik IPS Co Ltd (240810.KQ) - As of 04/17/26 GMT in KRW
Industry : S. Korea Technology



Stock Rating History: 4/1/21 : O/A; 7/19/21 : O/I; 8/12/21 : O/C; 10/4/22 : O/A; 7/21/24 : O/I; 9/15/24 : O/C; 4/3/25 : E/C; 6/13/25 : E/I; 8/8/25 : O/I; 9/21/25 : O/A

Price Target History: 1/27/21 : 65000; 10/26/21 : 55000; 3/18/22 : 50000; 8/9/22 : 43000; 1/12/23 : 35000; 2/21/24 : 34000; 3/28/24 : 48000; 4/3/25 : 25000; 8/8/25 : 48000; 10/21/25 : 72000; 3/23/26 : 160000

Source: Morgan Stanley Research Date Format : MM/DD/YY Price Target --- No Price Target Assigned (NA)
Stock Price (Not Covered by Current Analyst) — Stock Price (Covered by Current Analyst) ■
Stock and Industry Ratings (abbreviations below) appear as ♦ Stock Rating/Industry View
Stock Ratings: Overweight (O) Equal-weight (E) Underweight (U) Not-Rated (NR) No Rating Available (NA)
Industry View: Attractive (A) In-line (I) Cautious (C) No Rating (NR)

Effective January 13, 2014, the stocks covered by Morgan Stanley Asia Pacific will be rated relative to the analyst's industry (or industry team's) coverage.

Effective January 13, 2014, the industry view benchmarks for Morgan Stanley Asia Pacific are as follows: relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Important Disclosures for Morgan Stanley Smith Barney LLC Customers

Important disclosures regarding any material conflict of interest that can reasonably be expected to have influenced Morgan Stanley Smith Barney LLC's choice of a third-party research provider or the subject company of a third-party research report, are available on the Morgan Stanley Wealth Management disclosure website at www.morganstanley.com/online/researchdisclosures. For Morgan Stanley specific disclosures, you may refer to <https://www.morganstanley.com/eqr/disclosures/webapp/generalresearch>.

Each Morgan Stanley research report is reviewed and approved on behalf of Morgan Stanley Smith Barney LLC. This review and approval is conducted by the same person who reviews the research report on behalf of Morgan Stanley. This could create a conflict of interest.

Other Important Disclosures

A member of Research who had or could have had access to the research prior to completion owns securities (or related derivatives) in the Micron Technology Inc., NVIDIA Corp., Western Digital. This person is not a research analyst or a member of research analyst's household.

Morgan Stanley Research policy is to update research reports as and when the Research Analyst and Research Management deem appropriate, based on developments with the issuer, the sector, or the market that may have a material impact on the research views or opinions stated therein. In addition, certain Research publications are intended to be updated on a regular periodic basis (weekly/monthly/quarterly/annual) and will ordinarily be updated with that frequency, unless the Research Analyst and Research Management determine that a different publication schedule is appropriate based on current conditions.

Morgan Stanley is not acting as a municipal advisor and the opinions or views contained herein are not intended to be, and do not constitute, advice within the meaning of Section 975 of the Dodd-Frank Wall Street Reform and Consumer Protection Act.

Morgan Stanley produces an equity research product called a "Tactical Idea." Views contained in a "Tactical Idea" on a particular stock may be contrary to the recommendations or views expressed in research on the same stock. This may be the result of differing time horizons, methodologies, market events, or other factors. For all research available on a particular stock, please contact your sales representative or go to Matrix at <http://www.morganstanley.com/matrix>.

Morgan Stanley Research is provided to our clients through our proprietary research portal on Matrix and also distributed electronically by Morgan Stanley to clients. Certain, but not all, Morgan Stanley Research products are also made available to clients through third-party vendors or redistributed to clients through alternate electronic means as a convenience. For access to all available Morgan Stanley Research, please contact your sales representative or go to Matrix at <http://www.morganstanley.com/matrix>.

Any access and/or use of Morgan Stanley Research is subject to Morgan Stanley's Terms of Use (<http://www.morganstanley.com/terms.html>). By accessing and/or using Morgan Stanley Research, you are indicating that you have read and agree to be bound by our Terms of Use (<http://www.morganstanley.com/terms.html>). In addition you consent to Morgan Stanley processing your personal data and using cookies in accordance with our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html), including for the purposes of setting your preferences and to collect readership data so that we can deliver better and more personalized service and products to you. To find out more information about how Morgan Stanley processes personal data, how we use cookies and how to reject cookies see our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html). Please use the provided link to review the Terms and Conditions and Most Important Terms and Conditions for Morgan Stanley India Company Private Limited (https://www.morganstanley.com/assets/pdfs/about-us-global-offices/india/Terms_and_conditions.pdf) and the following link to review the audit report (https://ny.matrix.ms.com/eqr/research/webapp/researchdocs/MSICPL_Morgan_Stanley_Research_Audit_Report.pdf).

If you do not agree to our Terms of Use and/or if you do not wish to provide your consent to Morgan Stanley processing your personal data or using cookies please do not access our research. Morgan Stanley Research does not provide individually tailored investment advice. Morgan Stanley Research has been prepared without regard to the circumstances and objectives of those who receive it. Morgan Stanley recommends that investors independently evaluate particular investments and strategies, and encourages investors to seek the advice of a financial adviser. The appropriateness of an investment or strategy will depend on an investor's circumstances and objectives. The securities, instruments, or strategies discussed in Morgan Stanley Research may not be suitable for all investors, and certain investors may not be eligible to purchase or participate in some or all of them. Morgan Stanley Research is not an offer to buy or sell or the solicitation of an offer to buy or sell any security/instrument or to participate in any particular trading strategy. The value of and income from your investments may vary because of changes in interest rates, foreign exchange rates, default rates, prepayment rates, securities/instruments prices, market indexes, operational or financial conditions of companies or other factors. There may be time limitations on the exercise of options or other rights in securities/instruments transactions. Past performance is not necessarily a guide to future performance. Estimates of future performance are based on assumptions that may not be realized. If provided, and unless otherwise stated, the closing price on the cover page is that of the primary exchange for the subject company's securities/instruments.

The fixed income research analysts, strategists or economists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality, accuracy and value of research, firm profitability or revenues (which include fixed income trading and capital markets profitability or revenues), client feedback and competitive factors. Fixed Income Research analysts', strategists' or economists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

The "Important Regulatory Disclosures on Subject Companies" section in Morgan Stanley Research lists all companies mentioned where Morgan Stanley owns 1% or more of a class of common equity securities of the companies. For all other companies mentioned in Morgan Stanley Research, Morgan Stanley may have an investment of less than 1% in securities/instruments or derivatives of securities/instruments of companies and may trade them in ways different from those discussed in Morgan Stanley Research. Employees of Morgan Stanley not involved in the preparation of Morgan Stanley Research may have investments in securities/instruments or derivatives of securities/instruments of companies mentioned and may trade them in ways different from those discussed in Morgan Stanley Research. Derivatives may be issued by Morgan Stanley or associated persons.

With the exception of information regarding Morgan Stanley, Morgan Stanley Research is based on public information. Morgan Stanley makes every effort to use reliable, comprehensive information, but we make no representation that it is accurate or complete. We have no obligation to tell you when opinions or information in Morgan Stanley Research change apart from when we intend to discontinue equity research coverage of a subject company. Facts and views presented in Morgan Stanley Research have not been reviewed by, and may not reflect information known to, professionals in other Morgan Stanley business areas, including investment banking personnel.

Morgan Stanley Research personnel may participate in company events such as site visits and are generally prohibited from accepting payment by the company of associated expenses unless pre-approved by authorized members of Research management.

Morgan Stanley may make investment decisions that are inconsistent with the recommendations or views in this report.

To our readers based in Taiwan or trading in Taiwan securities/instruments: Information on securities/instruments that trade in Taiwan is distributed by Morgan Stanley Taiwan Limited ("MSTL"). Such information is for your reference only. The reader should independently evaluate the investment risks and is solely responsible for their investment decisions. Morgan Stanley Research may not be distributed to the public media or quoted or used by the public media without the express written consent of Morgan Stanley. Any non-customer reader within the scope of Article 7-1 of the Taiwan Stock Exchange Recommendation Regulations accessing and/or receiving Morgan Stanley Research is not permitted to provide Morgan Stanley Research to any third party (including but not limited to related parties, affiliated companies and any other third parties) or engage in any activities regarding Morgan Stanley Research which may create or give the appearance of creating a conflict of interest. Information on securities/instruments that do not trade in Taiwan is for informational purposes only and is not to be construed as a recommendation or a solicitation to trade in such securities/instruments. MSTL may not execute transactions for clients in these securities/instruments.

Certain information in Morgan Stanley Research was sourced by employees of the Shanghai Representative Office of Morgan Stanley Asia Limited for the use of Morgan Stanley Asia Limited. Morgan Stanley is not incorporated under PRC law and the research in relation to this report is conducted outside the PRC. Morgan Stanley Research does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC. PRC investors shall have the relevant qualifications to invest in such securities and shall be responsible for obtaining all relevant approvals, licenses, verifications and/or registrations from the relevant governmental authorities themselves. Neither this report nor any part of it is intended as, or shall constitute, provision of any consultancy or advisory service of securities investment as defined under PRC law. Such information is provided for your reference only.

Morgan Stanley Research is disseminated in Brazil by Morgan Stanley C.T.V.M. S.A. located at Av. Brigadeiro Faria Lima, 3600, 6th floor, São Paulo - SP, Brazil; and is regulated by the Comissão de Valores Mobiliários; in Mexico by Morgan Stanley México, Casa de Bolsa, S.A. de C.V. which is regulated by Comisión Nacional Bancaria y de Valores. Paseo de los Tamarindos 90, Torre 1, Col. Bosques de las Lomas Floor 29, 05120 Mexico City; in Japan by Morgan Stanley MUFG Securities Co., Ltd. and, for Commodities related research reports only, Morgan Stanley Capital Group Japan Co., Ltd; in Hong Kong by Morgan Stanley Asia Limited (which accepts responsibility for its contents) and by Morgan Stanley Bank Asia Limited; in Singapore by Morgan Stanley Asia (Singapore) Pte. (Registration number 199206298Z) and/or Morgan Stanley Asia (Singapore) Securities Pte Ltd (Registration number 200008434H), regulated by the Monetary Authority of Singapore (which accepts legal responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research) and by Morgan Stanley Bank Asia Limited, Singapore Branch (Registration number T14FC0118); in Australia to "wholesale clients" within the meaning of the Australian Corporations Act by Morgan Stanley Australia Limited A.B.N. 67 003 734 576, holder of Australian financial services license No. 233742, which accepts responsibility for its contents; in Australia to "wholesale clients" and "retail clients" within the meaning of the Australian Corporations Act by Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813, which accepts responsibility for its contents; in Korea by Morgan Stanley & Co International plc, Seoul Branch; in India by Morgan Stanley India Company Private Limited having Corporate Identification No (CIN) U22990MH1998PTC115305, regulated by the Securities and Exchange Board of India ("SEBI") and holder of licenses as a Research Analyst (SEBI Registration No. INH00001105); Stock Broker (SEBI Stock Broker Registration No. INZ000244438), Merchant Banker (SEBI Registration No. INM000011203), and depository participant with National Securities Depository Limited (SEBI Registration No. IN-DP-NSDL-567-2021) having registered office at Altimus, Level 39 & 40, Pandurang Budhkar Marg, Worli, Mumbai 400018, India; Telephone no. +91-22-61181000; Compliance Officer Details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: tejarshi.hardas@morganstanley.com; Grievance officer details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: msic-compliance@morganstanley.com. Morgan Stanley India Company Private Limited (MSICPL) may use AI tools in providing research services. All recommendations contained herein are made by the duly qualified research analyst; in Canada by Morgan Stanley Canada Limited; in Germany and the European Economic Area where required by Morgan Stanley Europe S.E., authorised and regulated by Bundesanstalt fuer Finanzdienstleistungsaufsicht (BaFin) under the reference number 14-9169; in the US by Morgan Stanley & Co. LLC, which accepts responsibility for its contents. Morgan Stanley & Co. International plc, authorized by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, disseminates in the UK research that it has prepared, and research which has been prepared by any of its affiliates, only to persons who (i) are investment professionals falling within Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (as amended, the "Order"); (ii) are persons who are high net worth entities falling within Article 49(2)(a) to (d) of the Order; or (iii) are persons to whom an invitation or inducement to engage in investment activity (within the meaning of section 21 of the Financial Services and Markets Act 2000, as amended) may otherwise lawfully be communicated or caused to be communicated. RMB Morgan Stanley Proprietary Limited is a member of the JSE Limited and A2X (Pty) Ltd. RMB Morgan Stanley Proprietary Limited is a joint venture owned equally by Morgan Stanley International Holdings Inc. and RMB Investment Advisory

(Proprietary) Limited, which is wholly owned by FirstRand Limited. The information in Morgan Stanley Research is being disseminated by Morgan Stanley Saudi Arabia, regulated by the Capital Market Authority in the Kingdom of Saudi Arabia, and is directed at Sophisticated investors only.

Morgan Stanley Hong Kong Securities Limited is the liquidity provider/market maker for securities of ASM International NV, FIT Hon Teng Ltd, Hon Hai Precision, Lenovo listed on the Stock Exchange of Hong Kong Limited. An updated list can be found on HKEx website: <http://www.hkex.com.hk>.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (DIFC Branch), regulated by the Dubai Financial Services Authority (the DFSA) or by Morgan Stanley & Co. International plc (ADGM Branch), regulated by the Financial Services Regulatory Authority Abu Dhabi (the FSRA), and is directed at Professional Clients only, as defined by the DFSA or the FSRA, respectively. The financial products or financial services to which this research relates will only be made available to a customer who we are satisfied meets the regulatory criteria of a Professional Client. A distribution of the different MS Research ratings or recommendations, in percentage terms for Investments in each sector covered, is available upon request from your sales representative.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (QFC Branch), regulated by the Qatar Financial Centre Regulatory Authority (the QFCRA), and is directed at business customers and market counterparties only and is not intended for Retail Customers as defined by the QFCRA.

As required by the Capital Markets Board of Turkey, investment information, comments and recommendations stated here, are not within the scope of investment advisory activity. Investment advisory service is provided exclusively to persons based on their risk and income preferences by the authorized firms. Comments and recommendations stated here are general in nature. 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.

The following companies do business in countries which are generally subject to comprehensive sanctions programs administered or enforced by the U.S. Department of the Treasury's Office of Foreign Assets Control ("OFAC") and by other countries and multi-national bodies: Samsung Electronics.

The trademarks and service marks contained in Morgan Stanley Research are the property of their respective owners. Third-party data providers make no warranties or representations relating to the accuracy, completeness, or timeliness of the data they provide and shall not have liability for any damages relating to such data. The Global Industry Classification Standard (GICS) was developed by and is the exclusive property of MSCI and S&P.

Morgan Stanley Research, or any portion thereof may not be reprinted, sold or redistributed without the written consent of Morgan Stanley.

Indicators and trackers referenced in Morgan Stanley Research may not be used as, or treated as, a benchmark under Regulation EU 2016/1011, or any other similar framework.

The issuers and/or fixed income products recommended or discussed in certain fixed income research reports may not be continuously followed. Accordingly, investors should regard those fixed income research reports as providing stand-alone analysis and should not expect continuing analysis or additional reports relating to such issuers and/or individual fixed income products.

Morgan Stanley may hold, from time to time, material financial and commercial interests regarding the company subject to the Research report.

A family member of one of the authors of this research report serves as a representative of a corporate director of Visual Photonics Epitaxy Co Ltd (2455.TW). This author of the research report and their family members have no control or influence over voting decisions made by the company's board members.

Registration granted by SEBI and certification from the National Institute of Securities Markets (NISM) in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Investment in securities market are subject to market risks. Read all the related documents carefully before investing.