October 5, 2026 06:27 AM GMT

Semiconductors | North America

How can the AI ecosystem work around memory bottlenecks?

We expect the memory shortage to be a constraint on AI builds through the duration of this AI cycle – but AI won't wait for DRAM fabs to be built. We look at ways to work around shortages.

Key Takeaways

- Intensity of memory shortages may fluctuate, but AI will basically use all supply for the foreseeable future
- CEO Jensen Huang has said that the industry needs to think differently to work around these shortages. We contemplate what he might be considering
- Among the potential beneficiaries would be MRVL and ALAB for CXL & larger scale-up, CBRS for disaggregation
- Remain OW MU and SNDK as we don't see the shortage ending

What was the catalyst for this note? In recent conversations with compute companies we have heard a great deal about the need to work around memory bottlenecks - NVIDIA CEO Jensen Huang among them. This was not a negative memory conversation - it's very clear that the industry anticipates multiple years of shortages.

But if memory supply cannot keep up with token growth, the industry will need to become more memory efficient. NVIDIA has talked about using its control over compute, networking, and storage to mitigate those bottlenecks.

In this note, we look at some of the ways that the industry might try to work around memory shortages, notably:

De-speccing, reducing the rack scale central memory (LPDDR5), and selectively reducing high bandwidth memory. This is in many ways the least favorable, but sometimes needed, approach. Our view is that there will continue to be products that maximize memory content, but that there will be selective de-speccing.

Potential beneficiaries: None directly, but allows compute companies to make their numbers - and shifts some of the bottleneck from DRAM to NAND. At the margin, de-speccing pushes for a larger scale up domain, which would help scale up plays such as ALAB and MRVL.

Disaggregation of AI workloads. This involves breaking up AI tasks into memory intensive and less memory intensive portions, where areas such as prefill which do not require as much memory.

We also note that as DRAM prices rise, some of the lower latency solutions from

MORGAN STANLEY & CO. LLC

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

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

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

Cate Folan
Research Associate
Cate.M.Folan@morganstanley.com +1 212 296-3520

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

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

SEMICONDUCTORS
North America
Industry View **Attractive**

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.

companies such as Groq (now within NVIDIA) or Cerebras or others use on chip SRAM, which has historically been much more expensive than DRAM, but currently is not.

Potential beneficiaries: compute solutions that can be a factor in disaggregated workloads, such as Cerebras, and NVIDIA through Groq

CXL memory centralization - CXL is a switched access to DRAM which allows for substantial efficiency advantages to centralized memory "pooling", allowing several processors to access the same memory. This has generally been considered a general purpose compute technology, but recent comments from chip suppliers are pointing to large opportunities within AI as well.

Potential beneficiaries: ALAB, MRVL.

Isn't this all bad for DRAM? In terms of near term earnings maximization, yes, a little. An environment in which the AI ecosystem comes to a halt and stops building because of the DRAM shortages probably leads to the best near term pricing outcomes.

But that was never realistic, and we have consistently argued for duration, not amplitude, as the primary driver of the cycle from here. And we would make the argument that reducing memory specs because of supply constraints - importantly, not because of pricing - leads to a situation where more supply will be very easily absorbed as we respec to higher levels.

While we don't take this for granted - we do extensive monthly checks into memory conditions at this point - our working hypothesis is that because of this pent-up demand, the risk to the memory cycle is the same as the risk to the compute space, that is a deceleration in AI itself. (We should introduce one risk factor here - that if AI slows, not because of demand, but because of land/power/shell bottlenecks, such a pause could cause memory disruption that we would not see with the compute names.

For that reason, we continue to be positive on MU and SNDK in our coverage.

The Brute Force reduction in memory specifications

Summary: Memory is a key aspect to compute performance, and generally there is reluctance to lower specifications. But the world is not going to wait for DRAM fabs to get built, so there are optimizations around lower memory levels. We would note that so far every reduction in specifications causes memory bottlenecks to surface in other parts of the compute spectrum.

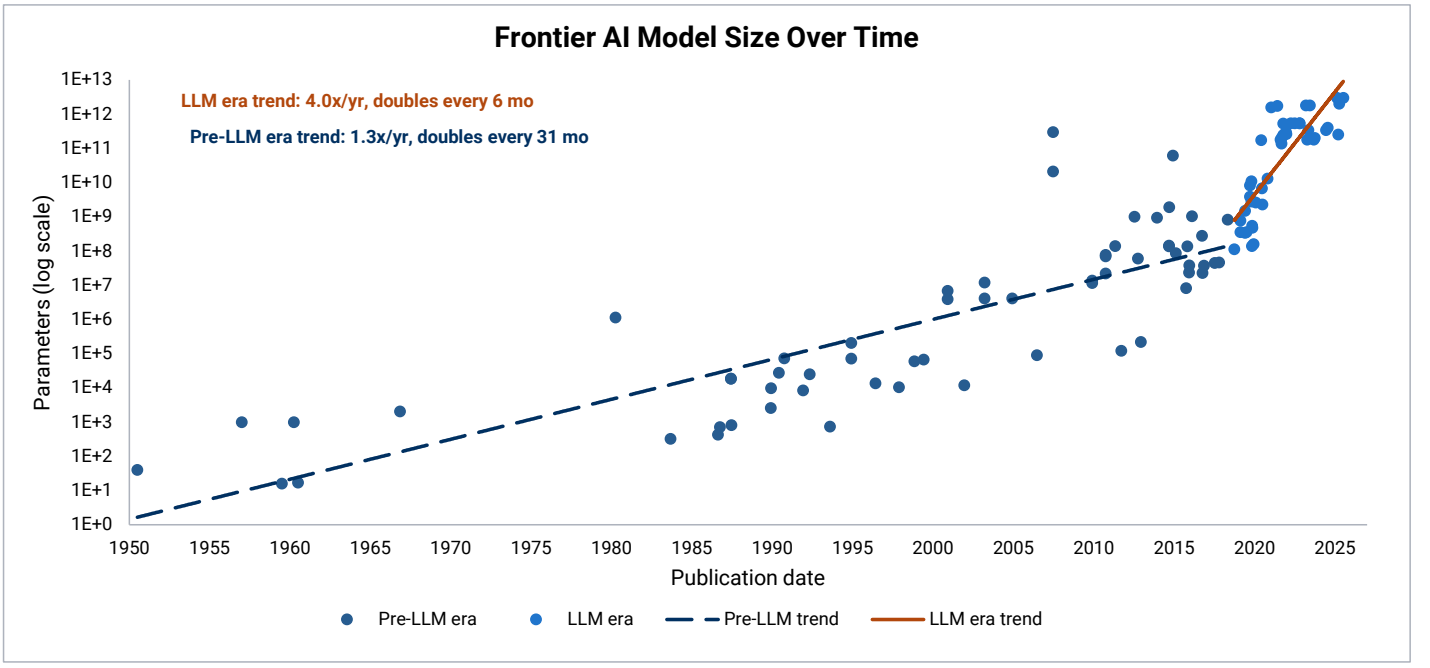
We are seeing Nvidia offer reduced memory content skews for rubin, including both HBM and main memory, in an attempt to offset rising prices and enable continued unit growth but that's less of a negative for memory companies than it appears. Increasing memory demand for AI a secular trend, and reducing memory content at one level of the memory hierarchy puts incremental pressure on the remaining tiers, shifting the memory bottleneck more than eliminating it.

What's happening with de-speccing today? The current shortages and rising prices of both DRAM and NAND has put significant pressure on memory buyers, and in response companies throughout the ecosystem are looking for ways to help ease those constraints. The most straight forward response, especially for DRAM, is to reduce the amount of memory per device or per server. We're seeing this happen out of necessity in most cases, where reducing the amount of DRAM is a measure taken to enable higher unit shipments. We believe Nvidia is offering additional rubin skews to customers, with LPDDR5 per rack content of 28TB vs the original 54TBs by asking DRAM vendors to supply 96GB SOCAMM2 modules vs 192GBs. For HBM Rubin was originally expected to be 288GB of HBM per GPU (8 stacks of 12hi HBM4), but we now expect skews with 192GB (8 stacks of 8hi HBM4) to be offered. For Rubin ultra content was originally debuted at 1TB HBM4e per 4 die accelerator, the move to keep rubin ultra at a 2 die design makes 512GB the relevant comparison vs the originally announced specs. Versus that 512 we could see Nvidia stay at HBM4 rather than move to 4e, as well as offer 8hi vs 12hi alternatives lowering overall capacity to the 192GB-384GB range.

What drives the need for more memory capacity in AI workloads? Before understanding what's lost when memory capacity is reduced, we think its important to look at why there has been a persistent push to add more memory over time. We highlight three factors below:

Model size: Over decades of machine learning history there has been a consistent desire for larger and larger models, as research has consistently shown the larger the model the better its performance. Model weights are almost always stored in HBM, increasing the need for larger HBM capacity over time.

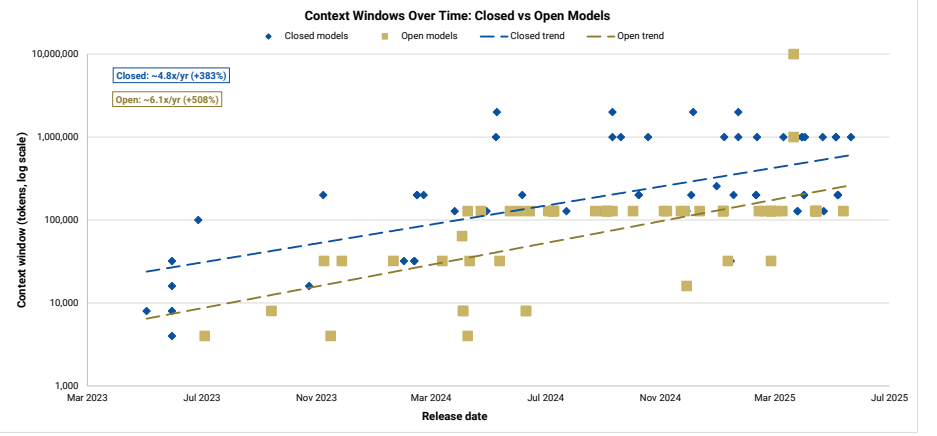
Exhibit 1: Frontier AI model sizes are doubling every 6 months



Source: EPOC.ai, Morgan Stanley Research

Context length: LLMs can only take into account a certain number of tokens at one time, including the prompt, uploaded documents, conversation history, and the model output. That limit is called the context length or context window, with most leading models today at about 1M tokens. As AI use cases increasingly revolve around long horizon agentic tasks the context window is a critical bottleneck for the amount of work an agent can accomplish.

Exhibit 2: LLM context windows are going ~5x annually



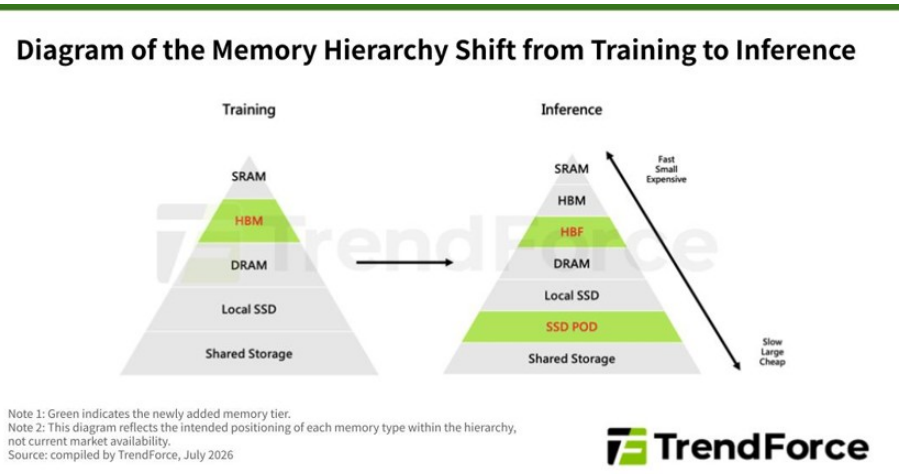
Source: EPOC.ai, Morgan Stanley Research

Inference concurrency: The number of requests an AI system serves at the same time. LLM inference has two steps, prefill and decode. Prefill is usually limited by compute, while decode is usually limited by memory. More simultaneous requests require more memory capacity, because each session stores its own context (the KV cache), and more memory bandwidth, because every decode step needs to process data from the KV cache.

We expect over time LLMs will become smarter (model size), need more in-context information (context length), and serve more users/agents over time (inference concurrency). That means de-speccing memory is likely to only prove temporary as each of those vectors continue to require greater amounts of memory across the hierarchy.

When capacity is reduced at one level of the memory hierarchy it puts pressure on the remaining tiers, de-speccing HBM/LPDDR is creating opportunities for networked DRAM (CXL) and NAND. Given the factors outlined above creating incremental demand for memory capacity for AI, when one tier is reduced (i.e HBM) the data needs to move elsewhere creating new pain points in memory and IO. We are seeing this in real time, as Nvidia moves to de-spec LPDDR5 racks we are seeing incremental demand for NAND as KV cache storage demands increase for that level of the memory tier. We are also seeing lower DRAM content putting increasing weight on networking, as bits that were stored in HBM need to be spread across more GPUs (putting pressure on scale-up bandwidth), or read from main memory (putting pressure on CPU IO).

Exhibit 3: De-speccing HBM/LPDDR should open up opportunities lower in the memory hierarchy



Source: TrendForce, Morgan Stanley

There will also be opportunities that emerge for new tiers in the memory hierarchy; we discuss CXL and SRAM based architectures below but going forward, optically connected memory is an interesting avenue for the industry as well as new technologies such as HBF.

Is cutting HBM capacity a "free lunch" for bandwidth bound workloads? If HBM content is being reduced by lowering HBM stack heights, like the rubin de-spec (moving from 12hi to 8hi), bandwidth isn't affected, as HBM bandwidth is function of the interface and pin speeds. That can work well for small context and smaller models, where there is lower capacity requirements but it creates risks. If trends outlined above (long context, agentic workloads) continue the number of workloads that fall into that category is likely to shrink, requiring more GPUs and/or increasing latency as more data if off-loaded to lower memory tiers. That said bandwidth is a key bottleneck for LLM decode, and lowering stack heights is a straight forward way to reduce costs. Going forward the key to watch if HBM vendors can charge a premium for higher bandwidth as HBM generations progress - that has been true historically. However, as custom base die's are introduced, the driver of HBM bandwidth (the DRAM PHY) can be designed by the compute vendor,

which may shift the value away from DRAM vendors.

The scale-up domain is also a factor to consider with Rubin Ultra and beyond. An effective scale-up mechanism effectively creates a rack that acts as a single GPU, by using a multiplexed switching mechanism to connect all to all within the rack. We don't know what Rubin Ultra looks like at this point - the company has said they have altered the Kyber form factor and will have a higher scale up domain - ie more GPUs per rack - and this would allow the total memory content at the rack level to stay constant with less memory per GPU.

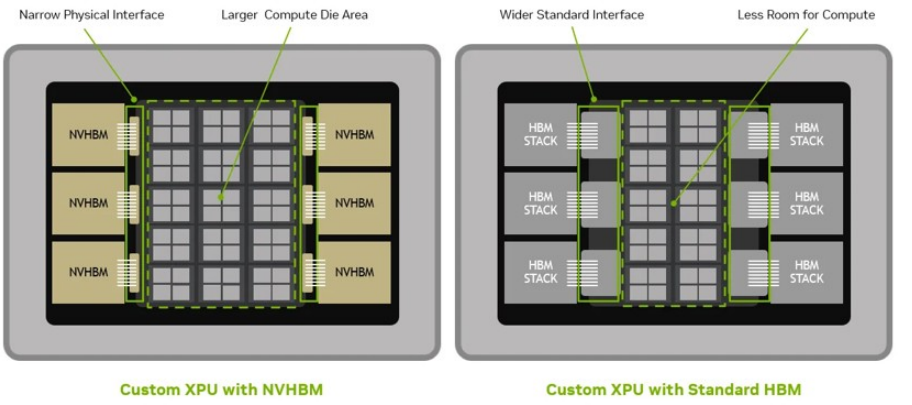
Exhibit 4: Nvidia custom HBM4e offers 30% more bandwidth vs JDEC standard HBM4e

Feature	NVHBM **benefit**
Bandwidth	Up to 30% more memory bandwidth compared with standard HBM4e
Area	More efficient interface connections allow up to 25% more compute die area for additional XPU capabilities
Power	Up to 15% lower HBM power usage compared with standard HBM4e adds up savings across thousands of XPUs

Table 1. NVHBM brings three main platform-level advantages to AI accelerator programs: higher memory bandwidth, more package and silicon area, and lower HBM power usage

Source: Nvidia, Morgan Stanley

Exhibit 5: Nvidia custom HBM reduces PHY and support area 67% vs JDEC standard HBM4e



Source: Nvidia, Morgan Stanley

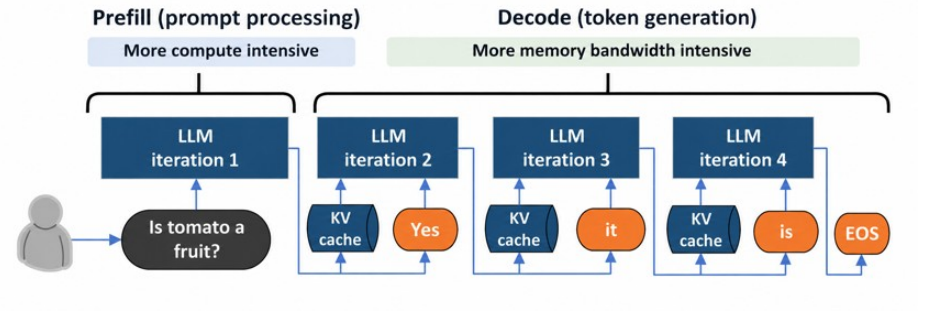
Not every step requires so much memory: Disaggregation

The opportunity in disaggregation comes from optimizing hardware around the different requirements of each stage of inference. Not all parts of inference place the same demands on compute and memory. Prefill processes the input prompt in parallel and is relatively compute intensive. Decode generates output one token at a time and repeatedly accesses model weights and the growing KV cache, making memory bandwidth and capacity more important. The same accelerator is therefore being asked to handle two workloads with very different resource requirements.

Disaggregation separates these stages so each can run on hardware better suited to the task. Compute dense systems can handle prefill, while architectures optimized for memory bandwidth and low latency can handle decode. This allows operators to scale the two resources independently and improve utilization, rather than provisioning the same mix of compute and memory for every stage of inference.

From a memory perspective, disaggregation allows the industry to use different types of memory where they are most valuable. Memory bandwidth intensive portions of inference can move to architectures built around high bandwidth SRAM or other approaches that reduce data movement, while HBM and larger capacity memory can remain focused on workloads that need their capacity and flexibility. Disaggregation does not eliminate memory demand, but heterogeneous architectures can improve how efficiently scarce and expensive memory is used as inference scales.

Exhibit 6: AI inference consists of two distinct stages: prefill and decode



Source: Adapted from Patel et al., "Splitwise: Efficient Generative LLM Inference Using Phase Splitting," ISCA 2024; Morgan Stanley Research.

Why now? The memory demands of inference are increasing as models process longer context and generate more output tokens. Reasoning models, coding agents and other agentic applications can require substantially more inference time than traditional question and answer workloads, increasing the importance of both memory capacity and bandwidth. At the same time, faster interconnects and improvements in inference software are making it more practical to move the KV cache between separate prefill and decode engines, helping address one of the main implementation challenges of disaggregation.

Interest in disaggregated inference accelerated sharply around the end of 2025. NVIDIA acquired Groq around this time, AWS subsequently announced a disaggregated architecture pairing Trainium for prefill with Cerebras for decode over EFA, and Cerebras has since announced a similar approach with AMD. Meanwhile, d Matrix has moved Corsair

into volume production. Taken together, these developments point to a broader move toward heterogeneous inference, where different processors are optimized for different parts of the workload.

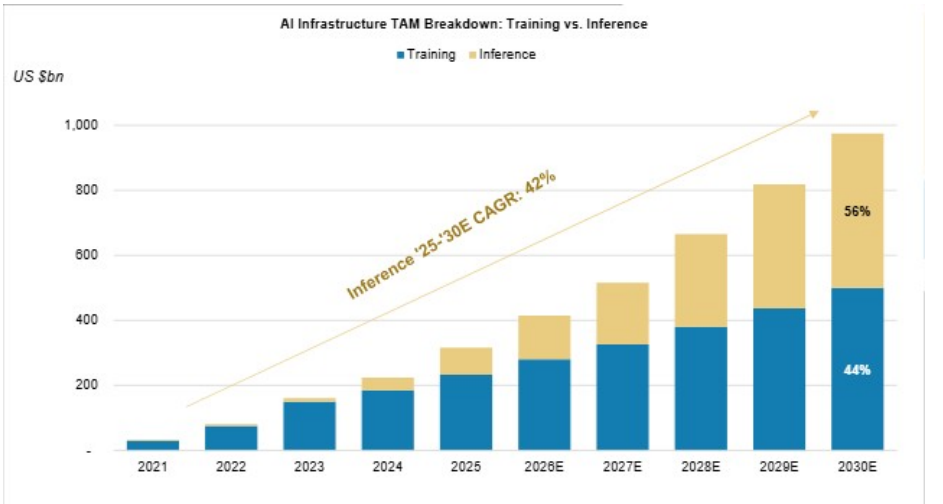
The emerging solutions take different approaches to the memory problem, ranging from replacing external memory with high bandwidth SRAM to combining SRAM, HBM and DDR within the same inference system:

- **CBRS:** Cerebras uses its wafer scale processor to combine a large amount of SRAM and compute on the same piece of silicon, reducing the need to move data between processors and external memory. The architecture provides very high memory bandwidth, which is particularly useful during decode when model weights must be accessed repeatedly for each generated token. Cerebras can run both prefill and decode, while its disaggregated architectures with AMD and AWS use GPUs or Trainium for the more compute intensive prefill stage and Cerebras for the memory bandwidth intensive decode stage.
- **NVIDIA Groq:** Groq's LPU uses high bandwidth on chip SRAM and a compiler scheduled architecture designed for low latency inference. NVIDIA is integrating Groq technology alongside its Vera Rubin platform, combining HBM based GPUs with SRAM based LPUs. Rubin GPUs handle prefill and the attention portion of decode, including workloads that require large KV caches, while Groq handles the feed forward and mixture of experts portions of decode. NVIDIA Dynamo coordinates the workload and transfers activations between the two systems.
- **D Matrix:** d Matrix uses a digital in memory compute architecture that places compute directly alongside high bandwidth on chip SRAM, allowing model weights to be processed where they are stored rather than repeatedly moved between separate memory and compute resources. This is designed to reduce the data movement bottleneck that becomes particularly important during decode, where model weights are accessed repeatedly for each generated token. d Matrix combines this high bandwidth local memory with larger capacity memory for data that does not need to remain on chip, and is positioning its accelerators alongside GPUs in heterogeneous systems where GPUs handle prefill and d Matrix handles decode.
- **SambaNova:** SambaNova uses a reconfigurable dataflow architecture that maps the model's computation directly onto the processor, allowing data to flow from one operation to the next without repeatedly returning to memory between each step. The architecture is paired with a three tier memory hierarchy spanning on chip SRAM, HBM and DDR, with the fastest memory used for the hottest local data and larger memory tiers providing capacity for model weights, KV cache and other less latency sensitive data. SambaNova is using this architecture in heterogeneous inference systems where GPUs handle prefill, its RDUs handle decode and CPUs manage orchestration and agentic tool execution.

We expect inference spending to grow faster than training as AI usage expands and each interaction consumes more compute. Disaggregated inference should be particularly relevant for workloads where prefill and decode place very different demands on the underlying hardware. Coding assistants and agentic applications are good examples: they can ingest large amounts of context, such as codebases, documents or prior interactions, and then generate substantial output across repeated model calls. This

creates a mix of compute intensive prefill and memory bandwidth intensive decode, making it more attractive to separate the two stages and scale each independently. Long context workloads can also benefit because a large prefill request can otherwise interfere with token generation for other users. As these workloads become a larger share of inference, disaggregation provides another way for the ecosystem to scale AI output without requiring every stage to be provisioned with the same mix of compute and high performance memory.

Exhibit 7: Total AI infrastructure spend forecast



Source: BBG, Morgan Stanley Research

Potential Stock Beneficiaries: CBRS, NVDA

We see Cerebras as one of the clearest pure play beneficiaries of greater adoption of disaggregated inference. The company is already working with both AMD and AWS on systems that pair third party processors for prefill with Cerebras for the memory bandwidth intensive decode stage. With AMD, Helios handles prompt processing and large context windows while the Cerebras Wafer Scale Engine handles decode, with the solution expected to enter production in 4Q26. AWS is taking a similar approach with Trainium for prefill and Cerebras for decode, with the combined solution expected to reach Amazon Bedrock in 1Q27. These partnerships give Cerebras a path to participate in heterogeneous systems even where another accelerator remains the primary compute engine

Disaggregation can also materially improve the economics of Cerebras Cloud. Cerebras and AMD have disclosed that their combined system can increase throughput by up to 5x while maintaining Cerebras inference speed, allowing each Cerebras system to produce substantially more tokens from the same underlying Cerebras capacity. Because Cerebras operates its own cloud infrastructure, lower cost per token can accrue directly to the company through higher gross margin, lower customer pricing, or some combination of the two. This makes disaggregation more than an incremental hardware opportunity for Cerebras, it can increase the amount of revenue generated from each deployed system while lowering the cost of producing each token.

Other potential beneficiaries include compute solutions that can participate in disaggregated inference workloads. NVIDIA is incorporating Groq technology into Vera

Rubin, pairing HBM based GPUs with SRAM based LPUs for different portions of inference. As disaggregated inference becomes more common, we see an opportunity for a broader set of processors to participate in AI infrastructure by specializing around the portions of the workload where their architecture is most efficient.



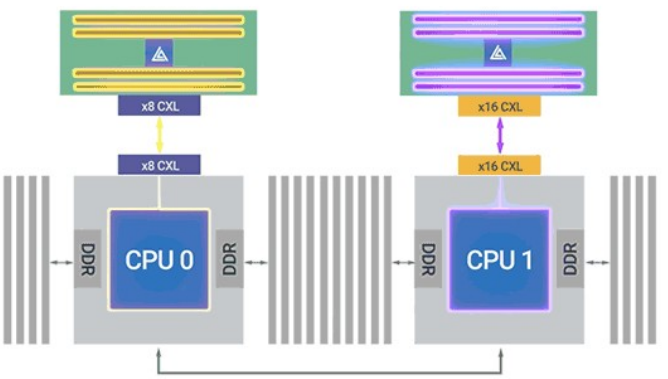
CXL: Using networking/interconnect technologies to optimize memory usage

The CXL Opportunity

CXL, or Compute Express Link, is a high speed interconnect designed to give processors more flexible access to memory. It uses the same underlying physical infrastructure as PCIe, but adds specialized memory and cache protocols that allow external memory to behave much more like memory directly attached to the processor. The key benefit is that memory no longer needs to be rigidly tied to a specific CPU or server, allowing data center operators to use available DRAM more efficiently and add capacity without necessarily adding more compute.

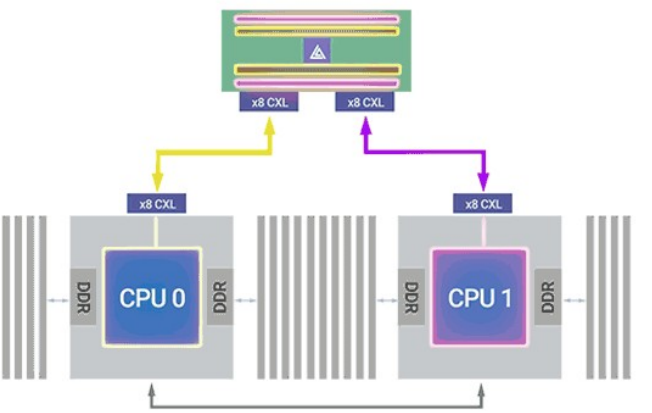
We see three main use cases: **Expansion** allows a server to add memory beyond the capacity supported by its local memory channels. **Sharing** allows multiple processors to access the same memory resources, reducing duplication and improving overall memory utilization. **Pooling** creates a common pool of memory that can be dynamically allocated across multiple processors or servers, reducing stranded capacity and overprovisioning.

Exhibit 8: Memory Expansion: Adds additional memory capacity to a single processor beyond its locally attached DRAM.



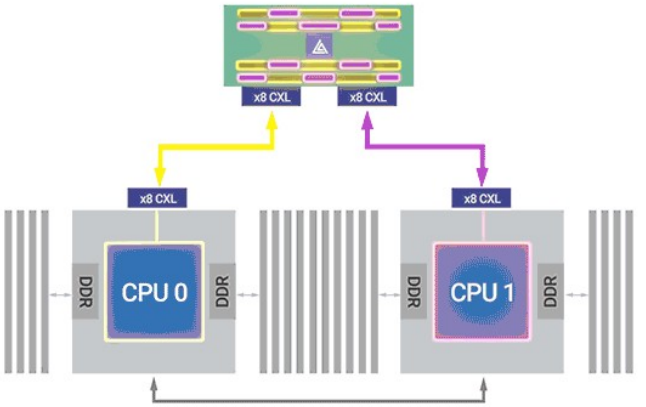
Source: Astera Labs

Exhibit 9: Memory Sharing: Allows multiple processors to access the same memory resource and data.



Source: Astera Labs

Exhibit 10: Memory Pooling: Creates a common memory pool that can be dynamically allocated across multiple processors.



Source: Astera Labs

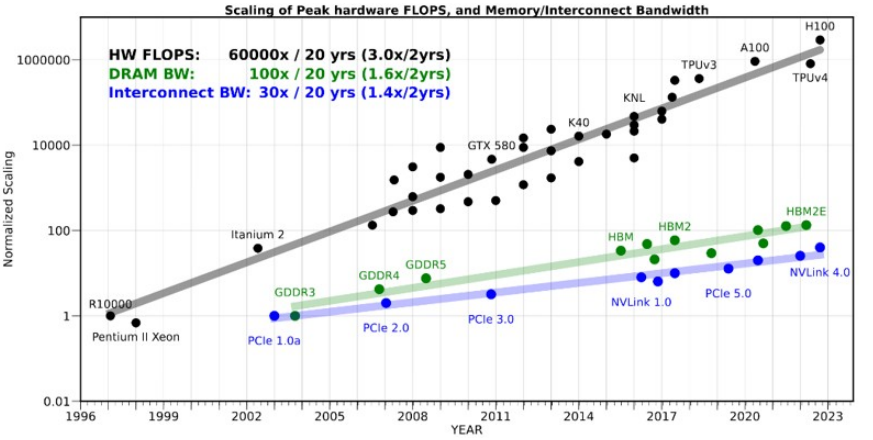
Why CXL for AI?

In the past, CXL has been used primarily for traditional CPU based compute, where workloads often benefit more from additional memory capacity than from optimizing for the absolute lowest latency and highest bandwidth. CXL introduces an additional hop between the processor and memory, making CXL attached DRAM slower than locally attached DRAM. **This has historically limited CXL's role in the most latency and bandwidth sensitive AI workloads, where GPUs rely on HBM for extremely high bandwidth**.

The broader backdrop is that compute performance has scaled much faster than memory and interconnect bandwidth, as shown in Exhibit 11. This widening gap makes memory placement and utilization increasingly important to overall system performance. It also increases the value of a tiered memory architecture, where HBM is reserved for the most bandwidth sensitive data while larger pools of DRAM can support workloads that are more capacity constrained.

More recently, CXL has attracted growing interest for AI inference as longer context windows, agentic workloads, and growing KV caches create much larger memory capacity requirements, while much of that data does not necessarily need to reside in HBM at all times. **CXL attached DRAM can therefore serve as a lower cost, higher capacity memory tier, with the most performance sensitive data remaining close to the accelerator and less sensitive data moved into external memory.**

Exhibit 11:



Source: "AI and Memory Wall," Amir Gholami, Zhewei Yao, Sehoon Kim, Coleman Hooper, Michael W. Mahoney, Kurt Keutzer.

CXL TAM broadens as memory constraints intensify

CXL has typically been sized as a CPU-attached memory expansion market. Astera, for example, has previously framed the CXL memory controller opportunity at over $4bn based on the volume of CXL-capable server CPUs, a market addressed by products such as Astera's Leo and Marvell's Structera X. While adoption has been slower than initially thought, we see conditions improving. Higher DRAM pricing and tighter supply are pushing hyperscalers to add capacity more efficiently and reuse memory they have already deployed. Astera has pointed to renewed customer interest on the back of these memory supply dynamics.

We now see AI expanding the opportunity to roughly $6bn by 2030, largely by creating new places to deploy CXL in AI systems. Within AI servers, host-attached CXL memory adds capacity for inference and agentic workloads. Beyond the host, suppliers are also building memory tiers that accelerators can reach directly over a fabric. Astera's Leo X targets this accelerator-side memory demand, and Astera has a custom KV-cache offload design win that ramps in 2027. As these tiers scale, sharing and pooling memory across larger groups of CPUs and accelerators also adds CXL switching content, a use case Marvell's Structera S targets at rack scale.

Supplier commentary supports this direction. Marvell has framed CXL as a >$1bn revenue contributor around CY2028 within its XPU-attach business, a significant portion of our 2028 estimate. We therefore see the broader CXL silicon opportunity expanding from a > $4bn memory controller market toward a roughly $6bn opportunity for total CXL and related memory-attach semi content by 2030, as AI drives higher memory requirements, greater CXL content per system and increasing adoption of pooled memory architectures.

Exhibit 12:



Source: Morgan Stanley Research

Sizing this market with precision is difficult. Many of the AI use cases are still in early deployment, the boundaries between CPU-attached, accelerator-attached and pooled memory are still being defined, and hyperscalers are pursuing different architectures, including some built on non-CXL fabrics, larger HBM stacks or flash-based context tiers. Both the mix across use cases and the timing of adoption could look meaningfully different from our assumptions. But our central view is that AI has expanded the opportunity materially beyond the >$4bn controller market that has framed CXL to date, and we expect to refine it as deployments scale.

Potential Stock Beneficiaries: ALAB and MRVL

Astera Labs. We see Astera as one of the most direct public beneficiaries through its Leo memory controller family. Management expects both standard and custom Leo products to ramp in volume in 2027 across two US hyperscalers, spanning traditional compute and AI inference, including a custom KV cache offload design. We do not fully reflect a broad CXL adoption cycle in our Astera estimates today, making a stronger Leo ramp largely incremental to our model.

Marvell. Marvell has a broad CXL portfolio spanning memory expansion through Structera X and rack level pooling through Structera S. Management has previously framed CXL as a greater than $1bn revenue contributor around CY2028, while more recent commentary points to additional design wins and stronger demand as inference and memory constraints increase. Some of this opportunity is already contemplated in Marvell's longer term framework, but continued expansion in CXL could provide additional upside. We expect to learn more at Marvell's upcoming Investor Day, where we expect more detail on its memory expansion opportunity.

The content addressing private companies is being provided for informational purposes only and does not constitute a solicitation or imply future research coverage if the company goes public. Content is based on unaudited information. No investment recommendation is provided as there is limited public information available for private companies. Investors should conduct their own due diligence and be aware that additional or different

information may be provided by the company.

Disclosure Section

The information and opinions in Morgan Stanley Research were prepared by Morgan Stanley & Co. LLC, and/or Morgan Stanley C.T.V.M. S.A., and/or Morgan Stanley Mexico, Casa de Bolsa, S.A. de C.V.. As used in this disclosure section, "Morgan Stanley" includes Morgan Stanley & Co. LLC, Morgan Stanley C.T.V.M. S.A., Morgan Stanley Mexico, Casa de Bolsa, S.A. de C.V. and their affiliates as necessary.

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; Cate Folan; Nicole Kozhukhov; Joseph Moore; Ella Tulchinsky; Mason Wayne.

.

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/conflictpolicies. A Portuguese version of the policy can be found at www.morganstanley.com.br

Important Regulatory Disclosures on Subject Companies

As of August 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, Aeva Technologies Inc, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Arm Holdings plc, Astera Labs Inc, Broadcom Inc., Cadence Design Systems Inc, Cerebras Systems, Intel Corporation, IonQ Inc, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., PDF Solutions Inc, Qorvo Inc, Qualcomm Inc., SanDisk Corporation., Semtech Corp., Silicon Laboratories Inc., Skyworks Solutions Inc, Synopsys Inc., Texas Instruments, Wolfspeed, INC.

Within the last 12 months, Morgan Stanley managed or co-managed a public offering (or 144A offering) of securities of Advanced Micro Devices, Aeva Technologies Inc, Analog Devices Inc., Broadcom Inc., Cerebras Systems, GlobalFoundries Inc, Intel Corporation, NVIDIA Corp., ON Semiconductor Corp., PDF Solutions Inc, Quantinuum, Semtech Corp., SiTime Corporation.

Within the last 12 months, Morgan Stanley has received compensation for investment banking services from Advanced Micro Devices, Aeva Technologies Inc, Allegro Microsystems Inc, Amkor Technology Inc, Analog Devices Inc., Broadcom Inc., Cerebras Systems, GlobalFoundries Inc, Intel Corporation, IonQ Inc, Micron Technology Inc., Navitas Semiconductor Corp, NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., PDF Solutions Inc, Quantinuum, SanDisk Corporation., Semtech Corp., SiTime Corporation, Texas Instruments.

In the next 3 months, Morgan Stanley expects to receive or intends to seek compensation for investment banking services from Advanced Micro Devices, Aeva Technologies Inc, Allegro Microsystems Inc, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Arm Holdings plc, Astera Labs Inc, Broadcom Inc., Cadence Design Systems Inc, Cerebras Systems, GlobalFoundries Inc, Intel Corporation, IonQ Inc, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., Navitas Semiconductor Corp, NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., PDF Solutions Inc, Qualcomm Inc., Quantinuum, SanDisk Corporation., Semtech Corp., Silicon Laboratories Inc., SiTime Corporation, Skyworks Solutions Inc, Synopsys Inc., Texas Instruments, Wolfspeed, INC.

Within the last 12 months, Morgan Stanley has received compensation for products and services other than investment banking services from Advanced Micro Devices, Allegro Microsystems Inc, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Broadcom Inc., Cadence Design Systems Inc, Cerebras Systems, GlobalFoundries Inc, Intel Corporation, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., Qorvo Inc, Qualcomm Inc., SanDisk Corporation., Semtech Corp., Silicon Laboratories Inc., SiTime Corporation, Synopsys Inc., Texas Instruments.

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, Aeva Technologies Inc, Allegro Microsystems Inc, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Arm Holdings plc, Astera Labs Inc, Broadcom Inc., Cadence Design Systems Inc, Cerebras Systems, GlobalFoundries Inc, Intel Corporation, IonQ Inc, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., Navitas Semiconductor Corp, NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., PDF Solutions Inc, Qualcomm Inc., Quantinuum, SanDisk Corporation., Semtech Corp., Silicon Laboratories Inc., SiTime Corporation, Skyworks Solutions Inc, Synopsys Inc., Texas Instruments, Wolfspeed, INC.

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, Allegro Microsystems Inc, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Broadcom Inc., Cadence Design Systems Inc, Cerebras Systems, GlobalFoundries Inc, Intel Corporation, IonQ Inc, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., Qorvo Inc, Qualcomm Inc., SanDisk Corporation., Semtech Corp., Silicon Laboratories Inc., SiTime Corporation, Synopsys Inc., Texas Instruments, Wolfspeed, INC.

Morgan Stanley & Co. LLC makes a market in the securities of Aeva Technologies Inc, Ambarella Inc.

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 September 30, 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.

	Coverage Universe		Investment Banking Clients (IBC)			Other Material Investment Services Clients (MISC)	
Stock Rating Category	Count	% of Total	Count	% of Total IBC	% of Rating Category	Count	% of Total Other MISC
Overweight/Buy	1599	43%	455	48%	28%	783	44%
Equal-weight/Hold	1563	42%	394	42%	25%	763	43%
Not-Rated/Hold	2	0%	1	0%	50%	1	0%
Underweight/Sell	570	15%	96	10%	17%	227	13%
Total	3,734		946			1774	

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.

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 Analog Devices Inc., Broadcom Inc., Microchip Technology Inc., NVIDIA Corp., Qualcomm Inc.. 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.

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 Comision 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 T14FC0118J); 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. 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. 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; 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 149169; 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. Morgan Stanley Research is disseminated in Canada by Morgan Stanley Canada Limited and Morgan Stanley Wealth Management Canada Inc., members of the Canadian Investment Regulatory Organization (CIRO). 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.

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 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.

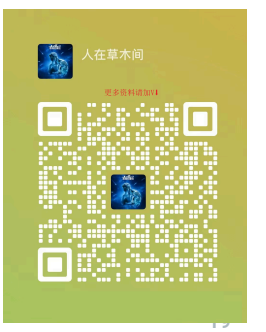
Morgan Stanley may hold, from time to time, material financial and commercial interests regarding the company subject to the Research report.

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.

References to any prediction market or event contract do not constitute an offer, solicitation, recommendation, or endorsement to participate in any such market or contract. Morgan Stanley Research may not be distributed or made available in any jurisdiction where such distribution would be unlawful or otherwise prohibited by applicable law.

INDUSTRY COVERAGE: Semiconductors

COMPANY (TICKER)	RATING (AS OF)	PRICE* (10/02/2026)
Joseph Moore		
Advanced Micro Devices (AMD.O)	E (06/09/2024)	$633.91
Aeva Technologies Inc (AEVA.O)	E (07/19/2021)	$14.76
Allegro Microsystems Inc (ALGM.O)	O (02/13/2026)	$39.53
Ambarella Inc (AMBA.O)	O (03/29/2016)	$68.65
Amkor Technology Inc (AMKR.O)	E (11/08/2023)	$56.02
Analog Devices Inc. (ADI.O)	O (11/16/2023)	
Astera Labs Inc (ALAB.O)	O (05/11/2025)	
Broadcom Inc. (AVGO.O)	O (06/09/2024)	
Cerebras Systems (CBRS.O)	O (06/08/2026)	
GlobalFoundries Inc (GFS.O)	E (10/28/2024)	
Intel Corporation (INTC.O)	E (02/22/2023)	
IonQ Inc (IONQ.N)	E (04/25/2023)	
Marvell Technology Group Ltd (MRVL.O)	E (09/14/2015)	



Microchip Technology Inc. (MCHP.O)	E (07/10/2024)	$81.32
Micron Technology Inc. (MU.O)	O (10/06/2025)	$1,074.89
Navitas Semiconductor Corp (NVTS.O)	U (04/06/2025)	$12.40
NVIDIA Corp. (NVDA.O)	O (03/16/2023)	$233.95
NXP Semiconductor NV (NXPI.O)	O (02/11/2025)	$243.66
ON Semiconductor Corp. (ON.O)	++	$84.89
Qorvo Inc (QRVO.O)	E (10/28/2025)	$114.17
Qualcomm Inc. (QCOM.O)	E (06/24/2026)	$184.87
Quantinuum (QNT.O)	E (06/29/2026)	$46.08
SanDisk Corporation. (SNDK.O)	O (03/03/2025)	$1,719.99
Semtech Corp. (SMTC.O)	E (04/06/2025)	$194.88
Silicon Laboratories Inc. (SLAB.O)	E (01/19/2021)	$222.76
SiTime Corporation (SITM.O)	O (09/16/2026)	$709.44
Skyworks Solutions Inc (SWKS.O)	E (11/28/2018)	$85.03
Texas Instruments (TXN.O)	U (04/13/2020)	$293.80
Wolfspeed, INC (WOLF.N)	NR (04/06/2025)	$35.28
Lee Simpson		
Arm Holdings plc (ARM.O)	E (04/07/2026)	$307.49
Cadence Design Systems Inc (CDNS.O)	O (02/14/2024)	$351.35
PDF Solutions Inc (PDFS.O)	O (09/08/2026)	$55.50
Synopsys Inc. (SNPS.O)	O (09/08/2026)	$489.90

Stock Ratings are subject to change. Please see latest research for each company.
* Historical prices are not split adjusted.

© 2026 Morgan Stanley