

June 8, 2026 04:19 AM GMT

Cerebras Systems | North America

# Unique capability in fast inference drives large economic opportunity; Initiate at OW

Starting late last year we have seen increased demand for low latency inference solutions, a category that Cerebras has uniquely addressed with its wafer scale technology for several years. Contracts point to \$6 bn in 2028 revenues with potential for upside.

## Key Takeaways

- Low latency inference demand started to surge in late 2025 and remains a key focus that we think can be more than 10% of the inference market
- Cerebras has a long history in this space with wafer scale technology
- Contracted agreement that gets them to 1% AI processor market share - \$6 bn - by 2028
- Risks include competition - including NVIDIA's Groq products - and execution around building out cloud
- But with \$7 trillion in market cap for the two leading AI processor companies, there is significant value opportunity; Initiating OW with \$250 PT.

**We initiate coverage of Cerebras with an Overweight rating and \$250 price target.** We view Cerebras as one of the most differentiated AI infrastructure companies, built around the industry's only commercially deployed wafer-scale processor. As AI workloads become increasingly reasoning-intensive, demand for fast, low-latency inference is growing rapidly. Supported by a large contracted backlog and a 750MW committed capacity agreement, we believe Cerebras is well positioned to capture this opportunity. This is a unique chance to invest in an AI processor company with a first-mover advantage against NVIDIA, and offers substantial upside as the category evolves.

**Low latency inference is suddenly an important growth category that Cerebras is uniquely poised to address.** Increased usage of AI for high value added tasks has pressured response times. Fast tokens are more expensive than regular tokens, at least for now, but we could see the low latency category account for 10% or more of inference hardware sales over the next few years.

**Cerebras technology is uniquely designed to address this need.** Wafer scale technology - literally building a chip across the entire surface of the wafer, with connections crossing the scribe lines - allows for very high speeds. We think that competing with best in class GPU technology can be very difficult, but this is something different - a unique approach to idiosyncratic portions of the workload

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
- Nicole Kozhukhov**  
Research Associate  
Nicole.Kozhukhov@morganstanley.com +1 212 761-1636
- Mason Wayne**  
Research Associate  
Mason.Wayne@morganstanley.com +1 212 761-6012
- Shane Brett**  
Equity Analyst  
Shane.Brett@morganstanley.com +1 212 761-1022



**Cerebras Systems (CBRS.O, CBRS US)**

Semiconductors | United States of America

<b>Stock Rating</b>	<b>Overweight</b>
<b>Industry View</b>	<b>Attractive</b>
<b>Price target</b>	<b>\$250.00</b>
Shr price, close (Jun 5, 2026)	\$201.01
Mkt cap, curr (mm)	\$50,720
52-Week Range	\$386.34-196.73

Fiscal Year Ending	12/25	12/26e	12/27e	12/28e
EPS (\$)***	ND	(1.22)	0.88	5.44
Prior EPS (\$)***	-	-	-	-
P/E	-	NM	227.7	36.9
EPS (\$)§	-	-	-	-
Div yld (%)	-	-	-	-

Unless otherwise noted, all metrics are based on Morgan Stanley ModelWare framework

\*\* = Based on consensus methodology

§ = Consensus data is provided by Refinitiv Estimates

e = Morgan Stanley Research estimates

**QUARTERLY EPS (\$)**

Quarter	2025	2026e Prior	2026e Current	2027e Prior	2027e Current
Q1	-	-	(0.21)a	-	(0.21)
Q2	-	-	(0.40)	-	0.00
Q3	-	-	(0.49)	-	0.37
Q4	ND	-	(0.53)	-	0.65

e = Morgan Stanley Research estimates, a = Actual Company reported data

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.**

that create very high value.

**Contracted revenue opportunities get them to \$6 bn revenue - approximately 1% market share of overall processors - by CY28, and we are optimistic that there can be upside.** The two leading processing companies in the AI space, NVIDIA and Broadcom, have over \$7 trillion in market value between them. Even with lower margin potential of a niche solution, we think that every 1% share can be worth \$50-70 bn, and that's the base case here. Sales from Open AI - which should make up ~80% of revenues in the next few years - are contracted over the next 3 years, and talking to customers there is clearly demand for more hardware if they can get ahead of supply constraints.

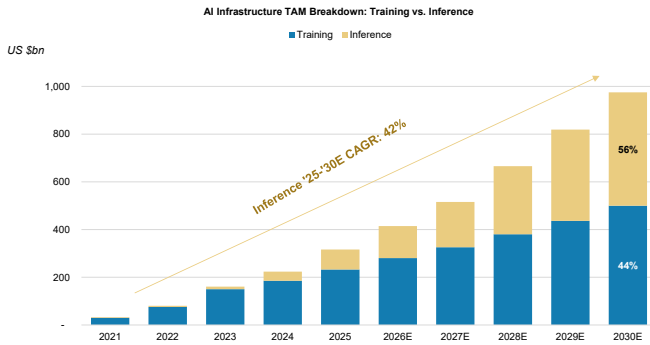
**We also expect competition in the category to intensify over time.** In particular, NVIDIA's acquisition of the assets of Groq - an AI processing company specializing in a low latency approach - is worth watching. While this deal does validate the category for Cerebras, and NVIDIA does not have wafer scale capability, they can bring the world class networking capabilities of its Mellanox business, as well as integration with its dominant GPU franchise, to bring a competitive offering. The size and profit of this opportunity will draw competition, but Cerebras first mover advantage and customer partnership is a good starting place.

We derive our \$250 price target using a 12x multiple on our 2028 revenue estimate adjusted for warrant contra revenue. The 12x sales multiple is in-line with SMID cap AI names with near-100% AI exposure that are outgrowing the overall market. While Cerebras operates a cloud-delivered infrastructure platform, we believe it is best compared to AI semiconductor peers given its differentiated hardware IP and economic structure. We believe Cerebras deserves a premium to the broader AI semiconductor peer group and a multiple in-line with high-growth SMID-cap AI companies, given its superior growth profile and exposure to one of the fastest-growing segments of AI infrastructure.

**Risks: With revenue under take or pay contracts, the potential challenges of the revenue ramp could be around execution.** We assume that most of the revenue will be Cerebras provided cloud services - though the contract could shift to become driven by hardware sales, it is more conservative, and probable, that it will remain cloud provision. This requires the companies to build out large scale cloud services which could cause some fits and starts along the way.

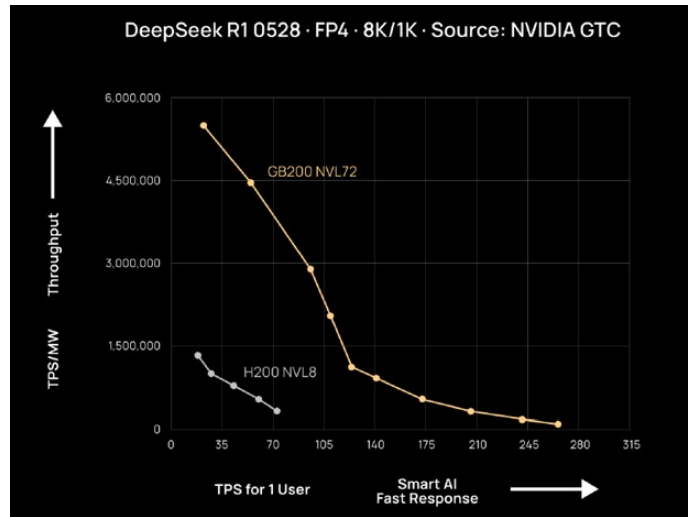
# Story in 6 Charts

**Exhibit 1:** We expect inference spend to outgrow training



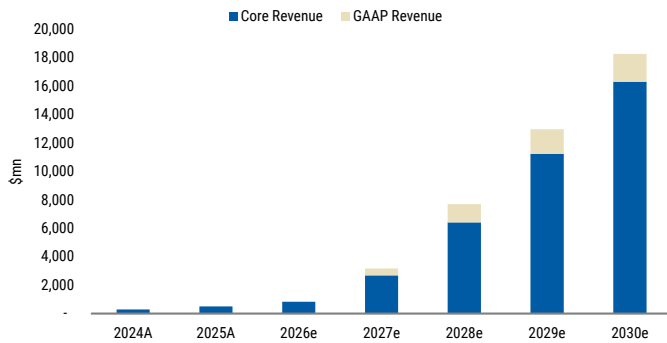
Source: BBG, Morgan Stanley Research

**Exhibit 2:** CBRS offers a differentiated accelerator optimized for inference



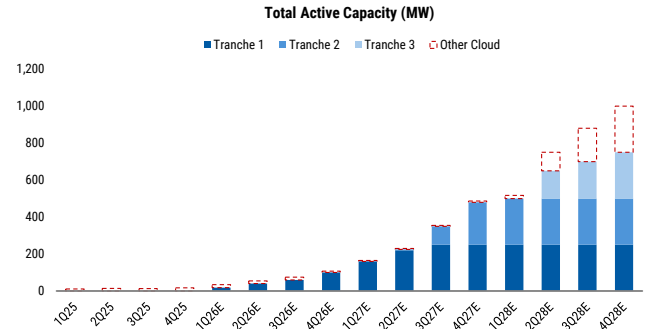
Source: Morgan Stanley Research

**Exhibit 3:** Multi-year agreement gives visibility to exponential growth for CBRS



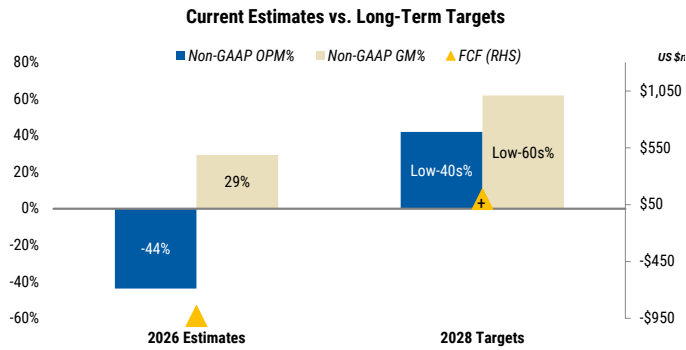
Source: Company Data, Morgan Stanley Research

**Exhibit 4:** But will require CBRS to bring on significant cloud capacity



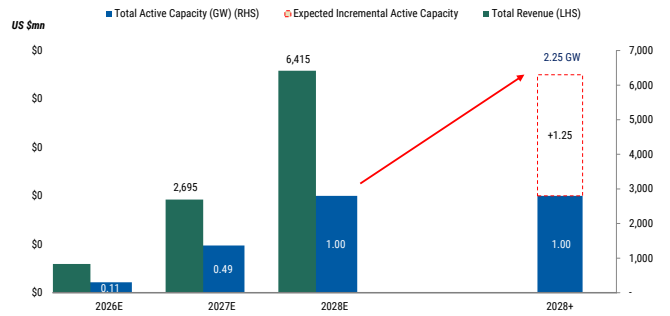
Source: Morgan Stanley Research

**Exhibit 5:** Execution matters and CBRS has ambitious long term financial targets



Source: Company Data, Morgan Stanley Research

**Exhibit 6:** Opportunity for Long-Term Upside - with options for an additional 1.25GW



Source: Morgan Stanley Research

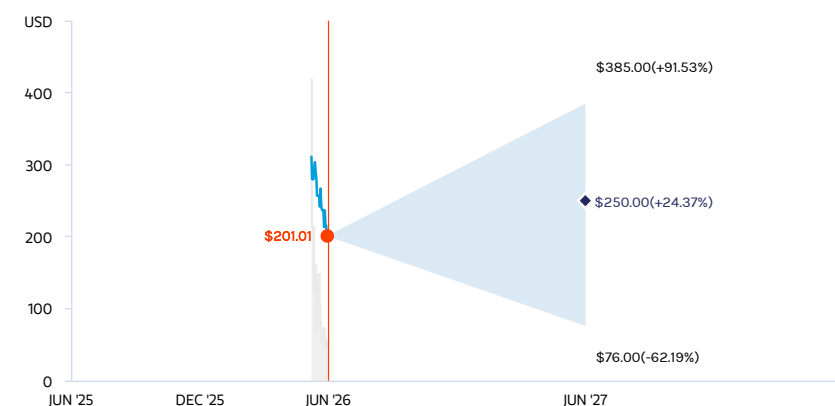
## Risk Reward – Cerebras Systems (CBRS.O)

Unique capability in fast inference drives large economic opportunity

### PRICE TARGET \$250.00

12x multiple on our 2028 adjusted revenue of \$6bn. We comp Cerebras to AI semis cohort and 12x is in-line with with smid cap AI names with near 100% AI exposure that are outgrowing the overall market.

### RISK REWARD CHART



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

Source: Refinitiv, Morgan Stanley Research

BULL CASE	\$385.00	BASE CASE	\$250.00	BEAR CASE	\$76.00
<b>12x CY28 revenue of \$10bn</b>		<b>12x 2028 adjusted revenue of \$6bn</b>		<b>7x CY28 adj revenue of \$2.7bn</b>	
- 26-28 revenue CAGR of 198%		- We comp Cerebras to AI semis cohort. 12x is in-line with with smid cap AI names with near 100% AI exposure that are outgrowing the overall market.		- 26-28 revenue CAGR of 95%	
- Hardware sales increase 21% y/y in CY27 and 203% y/y in CY28		- 26-28 revenue CAGR of 178%		- Hardware sales decline 9% y/y in CY27 and increase 37% y/y in CY28	
- Committed contract capacity comes online faster than expected, with 1GW+ deployed by 2028		- Ramps contract on time (750MW by 2028) using Cerebras Cloud		- Contract capacity deployments are delayed, resulting in slower revenue ramp and lower utilization.	
- Hardware demand also exceeds our base case and we get visibility to the additional 1.5GW being used		- Hardware sales increase 8% y/y in CY27 and 136% y/y in CY28 - Cerebras is able to ramp to 1GW of cloud capacity by 2028		- Hardware demand remains weak and they only ramp to 700MW.	

### OVERWEIGHT THESIS

Cerebras owns one of the most differentiated architectures in AI, built around the industry's only commercially deployed wafer-scale processor. We believe fast inference is one of the fastest-growing segments of AI infrastructure spending. Supported by a 750MW committed capacity agreement, Cerebras is well positioned to capitalize on this opportunity.

### Risk Reward Themes

Technology Diffusion: *Positive*

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

## Risk Reward – Cerebras Systems (CBRS.O)

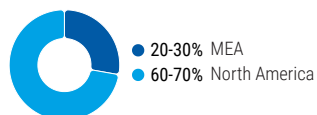
### KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
GAAP Revenue (\$, mm)	510	792	3,174	7,706
Non-GAAP Revenue (\$, mm)	510	831	2,695	6,415
Non-GAAP Gross Margin (%)	39.0	29.5	51.1	57.9
Non-GAAP EBITDA (\$, mm)	(62)	(158)	1,074	3,601

### INVESTMENT DRIVERS

- AI revenue growth
- Wafer Scale Engine Adoption
- AI Cloud Capacity Build

### GLOBAL REVENUE EXPOSURE



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

### RISKS TO PT/RATING

#### RISKS TO UPSIDE

- Contract Extension
- Additional customer
- Capacity deployment

#### RISKS TO DOWNSIDE

- Capacity deployment/infrastructure risk
- Customer concentration risk
- Competitive risk

### MS ESTIMATES VS. CONSENSUS

FY 2027e

**Sales / Revenue (\$, mm)** ◆ 2,695  
 Note: There are not sufficient brokers supplying consensus data for this metric

**EBITDA (\$, mm)** ◆ 1,074  
 Note: There are not sufficient brokers supplying consensus data for this metric

**EPS (\$)** ◆ 0.88  
 Note: There are not sufficient brokers supplying consensus data for this metric

◆ Mean ◆ Morgan Stanley Estimates

Source: Refinitiv, Morgan Stanley Research



# Inference is the Next AI Compute Battleground

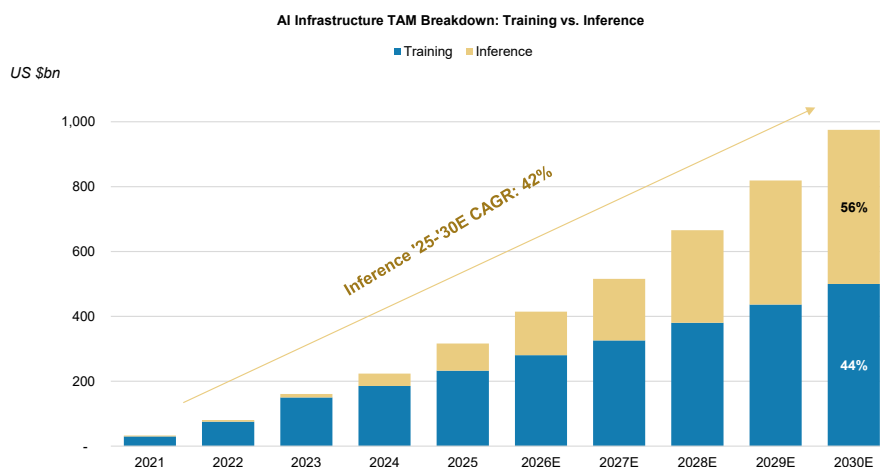
**Main takeaway:** We believe AI infrastructure spending is increasingly shifting from training toward inference. As reasoning models generate more tokens and perform more complex workflows, performance is increasingly constrained by memory bandwidth and communication rather than raw compute. Cerebras' Wafer-Scale Engine was designed specifically to address these bottlenecks, positioning the company as one of a small number of architectures optimized for the emerging market for ultra-fast inference. While we expect low latency inference to be a subsegment of inference - at least for now offering expensive but very fast tokens - it is a focus category for key customers.

## The Fast Inference Opportunity

**The next compute wave will be driven by inference.** While training remains critical to model development, recent advances in reasoning and agentic AI have shifted the industry's focus toward inference-time compute. Rather than generating responses in a single pass, modern frontier models increasingly perform multiple reasoning steps, verification loops, and tool calls before producing an answer. As a result, compute demand is no longer determined solely by model training requirements, but increasingly by the amount of compute consumed during deployment.

This shift has important implications for AI infrastructure providers. Inference demand scales not only with the number of users, but also with usage intensity and the amount of compute required per interaction. As reasoning models become more capable, compute consumption per query is increasing significantly, creating a rapidly expanding market for inference infrastructure. **We expect AI inference infrastructure spend to grow more than twice as fast as AI training infrastructure through 2029.**

**Exhibit 7:** Total AI infrastructure spend forecast



Source: BBG, Morgan Stanley Research

As inference becomes a larger portion of overall AI compute demand, speed is becoming

increasingly important. Coding assistants, research agents, conversational AI applications, and future autonomous systems all require near real-time responses to maintain user engagement. While users may tolerate longer runtimes during model training, they are significantly less tolerant of delays during inference. **As a result, latency is emerging as a key competitive differentiator across AI infrastructure.**

This shift also changes how AI infrastructure is evaluated. For model providers, the ability to generate tokens efficiently directly impacts revenue generation, making tokens-per-watt and inference efficiency increasingly important metrics. We believe this dynamic benefits architectures optimized for inference workloads, particularly those designed to reduce communication overhead and maximize memory bandwidth.

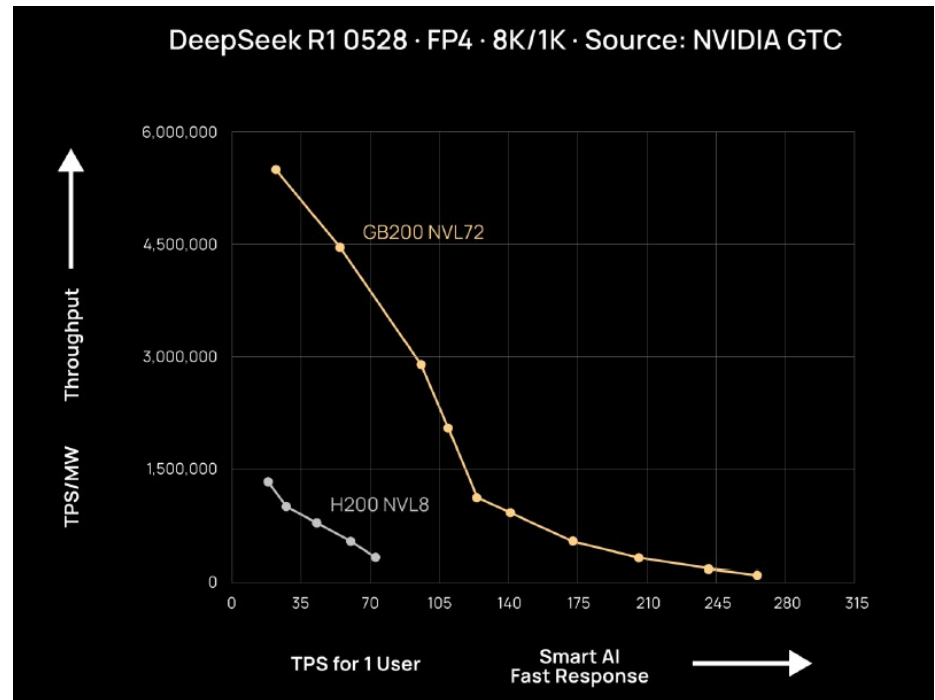
Importantly, not all stages of inference are equally constrained. While the prefill stage can generally be parallelized across large GPU clusters, the decode stage (token generation) is inherently sequential and increasingly limited by memory bandwidth and communication overhead. As reasoning models generate more tokens and perform longer chains of thought, decode latency becomes a larger determinant of overall application performance and user experience.

**Cerebras is targeting this emerging market for low-latency inference.** Applications such as coding agents, deep research tools, conversational AI, enterprise copilots, search, and future agentic systems place a premium on response speed, creating demand for architectures optimized around fast token generation. Similar to NVIDIA's Vera Rubin + Groq architecture, Cerebras is also positioning its wafer-scale systems alongside XPU infrastructure, with XPU's handling highly parallelizable workloads and Cerebras accelerating the latency-sensitive decode phase.

Looking at [Exhibit 2](#) [Exhibit 2](#), AI infrastructure increasingly faces a tradeoff between throughput and responsiveness. Historically, the industry optimized for throughput and cost efficiency. Going forward, we believe a growing subset of AI workloads will prioritize response speed, larger context windows, and more sophisticated reasoning capabilities.

**This premium speed segment of the inference market is where Cerebras is best positioned to compete.**

**Exhibit 8:** The AI inference tradeoff between throughput and responsiveness. Cerebras/Groq is positioned at the far right. As AI workloads become increasingly interactive and reasoning-intensive, a growing subset of inference demand values faster response times over maximizing total token output.

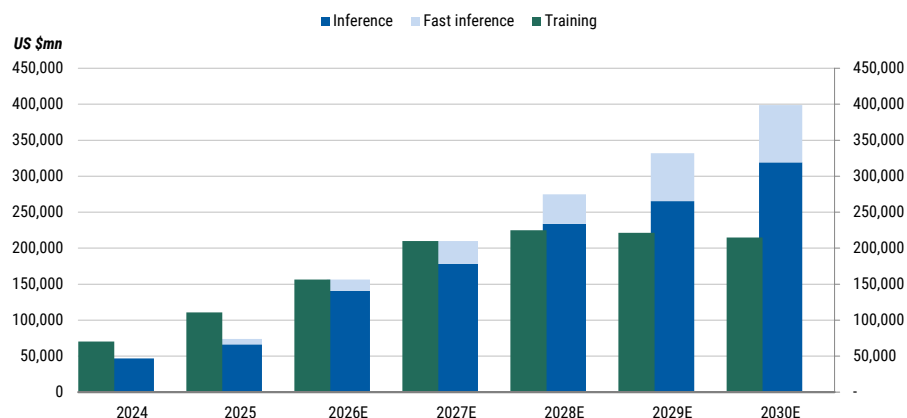


Source: NVIDIA 2026 GTC

### How Large is the Fast Inference Opportunity?

We believe Cerebras' addressable market is best defined as latency-sensitive inference workloads, where response speed directly impacts user experience and economic value. While the broader AI inference market continues to expand rapidly, not all inference workloads require ultra-fast response times. Batch processing, offline analytics, and many enterprise workloads can tolerate higher latency and are therefore less likely to prioritize specialized inference infrastructure. In contrast, applications such as coding agents, conversational AI, search, deep research, and future agentic systems place a premium on responsiveness, creating a distinct market for ultra-fast inference.

Sizing this opportunity remains challenging given the early stage of the market. However, at GTC 2026, NVIDIA CEO Jensen Huang estimated that ultra-fast inference applications represent approximately 25% of the processors in a data center that incorporates low latency inference solutions, and in other forums they have said it could be 10% of the inference market overall. Using this framework, we estimate the fast-inference segment is 10% of the inference market near-term, but could grow to 20% by 2030, representing an ~\$80bn opportunity and approximately 15% of the broader cloud semiconductor opportunity.

**Exhibit 9:** AI Cloud Semiconductor Forecast

Source: Morgan Stanley Research estimates

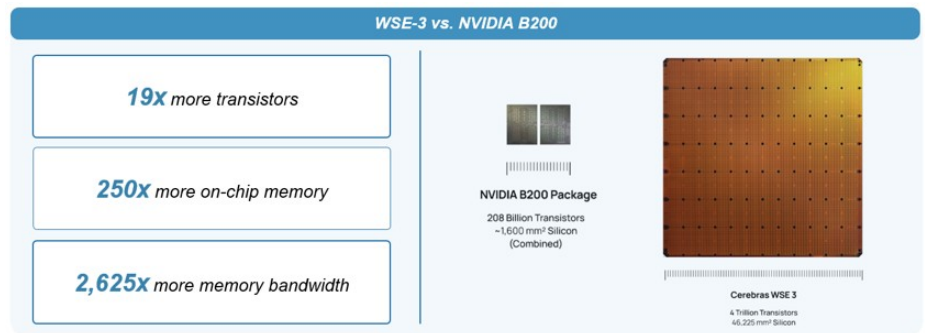
While NVIDIA's estimate should not be viewed as a direct measure of Cerebras' opportunity, it provides a useful framework for understanding the size of the premium tier of inference infrastructure where speed is a key differentiator. We believe Cerebras is one of a small number of architectures purpose-built for this segment, positioning the company to participate in one of the fastest-growing and highest-value areas of AI infrastructure spending as reasoning models, AI agents, and real-time AI applications continue to proliferate.

## The Wafer Scale Engine

Traditional AI systems are built by connecting hundreds or thousands of processors together. While advances in networking and software have enabled increasingly larger AI clusters, scaling performance requires moving growing amounts of data between processors, memory, and networking equipment. As models become larger and reasoning workloads generate more tokens, communication overhead increasingly becomes a limiting factor in system performance.

**We believe Cerebras' Wafer-Scale Engine (WSE) is one of the architectures best suited for ultra-fast inference, as it was designed to address the memory bandwidth and communication bottlenecks that increasingly constrain modern AI workloads.** Cerebras was founded around a different approach: rather than connecting many small chips together, the company built a single processor from an entire silicon wafer. The result is the Wafer-Scale Engine (WSE), the largest processor ever constructed. Cerebras' latest generation WSE-3 contains 4 trillion transistors, 900,000 AI-optimized compute cores, 44GB of on-chip SRAM memory, and 21 petabytes per second of memory bandwidth. According to the company, the WSE-3 is approximately 58x larger than NVIDIA's Blackwell GPU and provides more than 2,500x greater memory bandwidth.

**Exhibit 10:** The WSE is a completely new architecture vs. other GPU/XPUs

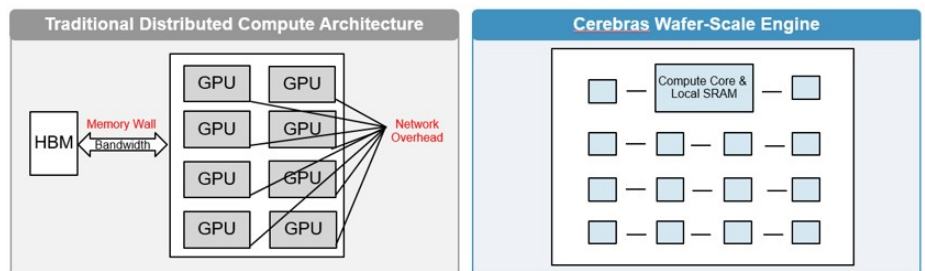


Source: Company Filings, Morgan Stanley Research

The primary advantage of wafer-scale computing is not simply size, but the ability to keep compute, memory, and communication on a single piece of silicon. Traditional GPU clusters require data to move across multiple processors and networking fabrics, introducing latency and consuming power. Cerebras instead places memory adjacent to compute across the wafer, reducing the need for chip-to-chip communication and allowing data to remain on-chip for a larger portion of the workload.

This architectural approach is particularly relevant for inference workloads. During the decode phase of inference, models generate tokens sequentially and repeatedly access model parameters stored in memory. As a result, inference performance increasingly depends on memory bandwidth and communication efficiency rather than raw compute throughput. Cerebras' architecture is designed specifically to address this bottleneck by maximizing on-chip memory bandwidth while minimizing data movement.

**Exhibit 11:** Relying on off-chip memory and distributed clusters, requiring constant data movement. Cerebras solves this with wafer-scale integration where compute and memory on a single chip eliminates data movement bottlenecks.



Source: Morgan Stanley Research

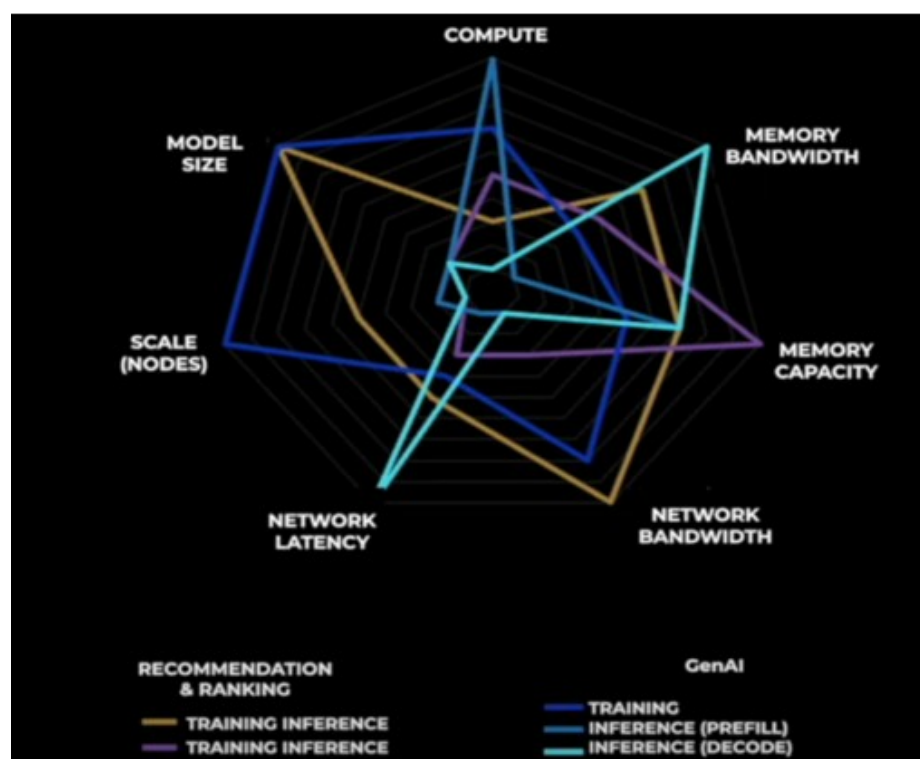
The company's CS-3 system packages the WSE-3 into a complete AI platform that can be deployed as a standalone inference engine or alongside GPU infrastructure. Rather than competing directly with GPUs across every workload, Cerebras is increasingly positioning the WSE as a specialized accelerator for latency-sensitive inference, where response speed and token generation rates are critical determinants of user experience.

We believe Cerebras' wafer-scale architecture represents one of the most differentiated approaches in AI infrastructure. While the company faces substantial competition from NVIDIA, AMD, and other AI accelerator vendors, its focus on memory bandwidth, communication efficiency, and low-latency inference provides a distinct value proposition as AI workloads become increasingly reasoning- and agent-driven.

## How does the WSE compare?

In our view, there is no universally optimal architecture for LLM inference. Performance depends heavily on workload characteristics such as batch size, sequence length, and model scale, which creates opportunities for specialized accelerators to outperform traditional GPUs in specific use cases. Against a backdrop of persistent AI infrastructure constraints and rapidly growing demand, we are seeing customers show increasing willingness to evaluate alternative architectures. As a result, the market is beginning to fragment across distinct performance tradeoffs, with emerging platforms sacrificing some degree of flexibility in exchange for superior latency, throughput, or efficiency in targeted workloads.

**Exhibit 12:** AI workloads optimize for different performance bottlenecks, no single architecture is optimal across all workloads.



Source: Meta, Morgan Stanley Research

**Cerebras is designed to excel in communication-bound inference workloads, where data movement rather than raw compute increasingly constrains performance.**

Cerebras' architecture is particularly well suited to low-batch inference workloads that fit within a single wafer-scale processor. By minimizing off-chip communication, the WSE reduces the latency and energy costs of data movement, a growing bottleneck in modern AI systems. The resulting performance advantages can be significant for communication-bound workloads, although outcomes remain highly dependent on model size and workload characteristics. In the chart shown in [Exhibit 6](#), Cerebras performs very well for inference decode.

**Groq is designed to excel in latency-sensitive inference workloads, where execution efficiency and deterministic performance are the primary constraints.** Groq takes a different approach to high-performance inference, optimizing for deterministic execution.

Its LPU architecture combines high-bandwidth SRAM with a compiler-driven execution model that precisely schedules operations ahead of time, reducing the need for dynamic control logic and congestion management. The result is a highly predictable inference engine capable of delivering very low latency for individual requests. However, the architecture's advantages appear most pronounced in smaller-scale, latency-sensitive workloads, where performance is determined by execution efficiency rather than the ability to scale across large pools of compute resources.

Groq is perhaps the most widely recognized alternative architecture to Cerebras. NVIDIA plans to incorporate Groq's technology into its roadmap, with solutions expected to begin shipping as part of the Vera Rubin platform in 2H26. Given the early stage of deployment, it is difficult to fully assess the competitive positioning of the offering before seeing it in production at scale. More broadly, the size of the inference opportunity suggests there is likely room for multiple architectures to coexist, each optimized for different workload requirements. At this stage, we believe both Cerebras and Groq have credible paths to compete in the market.

**GPUs are designed to excel across a broad range of AI workloads, where flexibility, scalability, and software ecosystem support are more important than optimizing for a single performance bottleneck.** Traditional GPUs remain the dominant platform for AI training and inference due to their flexibility, mature software ecosystem, and ability to scale across a broad range of workloads. Unlike purpose-built accelerators such as Cerebras and Groq, GPUs are designed as general-purpose parallel processors capable of supporting diverse model architectures and deployment environments. While this flexibility comes at the cost of greater scheduling overhead, more complex memory hierarchies, and higher communication requirements, GPUs remain the default choice for organizations seeking a versatile platform that can support both training and inference without workload-specific optimization. We don't see these specialized accelerators as threatening traditional GPUs leadership position.

# Multi-Year Contract Drives Growth and Scale

**Main takeaway:** We view OpenAI adoption as both a commercial validation of Cerebras' technology and the catalyst that transforms the company from a niche hardware vendor into a scaled AI infrastructure platform. While customer concentration creates inherent risk, the contract provides a credible path toward multi-billion-dollars of revenue and demonstrates growing demand for specialized inference infrastructure for Cerebras.

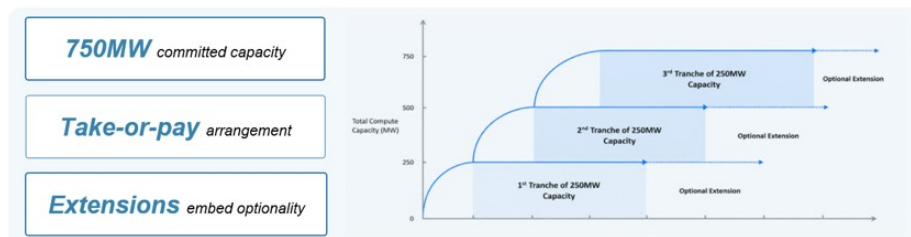
**Validation of the fast inference thesis.** We believe this contract represents the most important commercial validation of Cerebras' technology to date. The magnitude of the deal is huge, and the fact that OpenAI has chosen to deploy Cerebras' technology at this scale is a significant indicator of the importance of the market Cerebras operates in and the value of their WSE. In our view, the economics Cerebras was able to secure is also indicative of how compute constrained the entire ecosystem is.

## Terms of the Deal and Modeling Implications

Cerebras committed to a take-or-pay agreement with OpenAI for 750MW of inference capacity, deployed in three 250MW tranches scheduled over the next three years. The contract is structured around dedicated capacity rather than usage, with pricing based on installed MW over the contractual term.

Tranche 1 (250MW) is delivered through the Cerebras Cloud and carries a 3-year base term with two optional one-year extensions (maximum 5 years total). Tranche 2 (250MW) and Tranche 3 (250MW) may be delivered through the Cerebras Cloud or as on-prem hardware purchases. Tranche 3 has a 4-year committed term with a one-year extension option. All tranches can extend to a maximum of approximately 5 years.

**Exhibit 13:** Multi-year agreement with OpenAI provides near-term revenue visibility. Additional capacity options create embedded upside.



Source: Company Filings, Morgan Stanley Research

OpenAI also holds an option to purchase an additional 1.25GW of capacity by 2029, bringing total potential deployment to approximately 2GW. The option is linked to warrant incentives and is not included in the base contractual commitment.

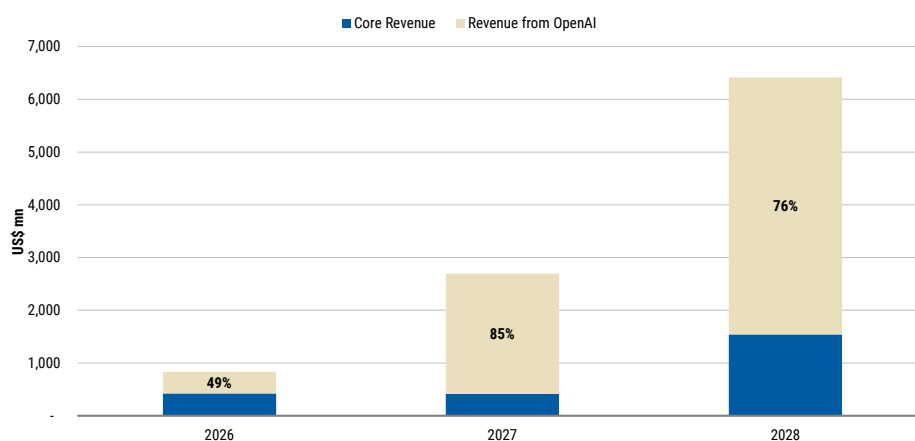
## Modeling Framework

**We model this agreement using a deployed-capacity framework based on the implied**

**lifetime contract value per GW.** We estimate the agreement is worth approximately ~\$30–40bn per GW over the life of the contract; using the low end of this range and a blended 4-year contract duration implies approximately ~\$7.5bn of annual revenue per GW, or roughly ~\$1.875mn of quarterly revenue per deployed MW. Revenue is therefore driven primarily by deployed capacity rather than contracted capacity, with quarterly revenue increasing as additional MW comes online. The disclosed \$21.6bn RPO includes the firm 750MW commitment and estimated pass-through revenue for the initial 250MW deployment, but excludes future pass-through associated with later tranches as well as the optional 1.25GW expansion. We estimate the total economic value of the total deal is closer to ~\$30bn in our base case scenario.

The agreement provides substantial embedded upside beyond the initial 750MW commitment. Each tranche includes extension options that can increase the effective contract duration from 3-4 years to as much as 5 years, extending revenue generation from deployed capacity beyond the current forecast horizon. In addition, OpenAI holds an option to purchase an incremental 1.25GW of capacity by 2029, increasing total potential deployment to approximately 2GW. We only model to the base case 750 MW, so the exercise of extension options could materially increase revenue to Cerebras beyond what we forecast.

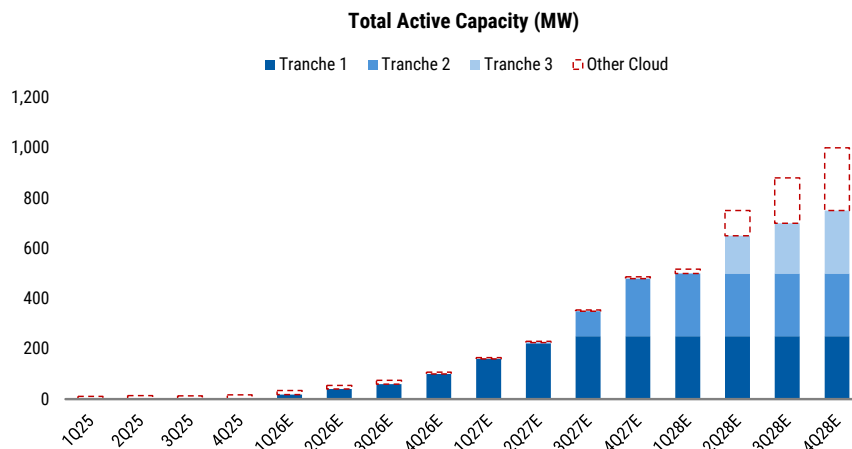
**Exhibit 14:** Revenue from the OpenAI contract at the peak will be 85% of revenue concentration in our base scenario



Source: Morgan Stanley Research

**Under this framework, deployment timing becomes the primary driver of quarterly revenue.** In our base case, we assume the first 250MW tranche becomes fully deployed by mid-2027, the second tranche comes online in early 2028, and the final tranche is deployed by the end of 2028. We assume deployment efficiency improves over time as Cerebras gains experience scaling its infrastructure footprint. As a result, any acceleration or delay in the deployment schedule would represent a meaningful source of upside or downside to our forecast.

**Exhibit 15:** MSe capacity schedule



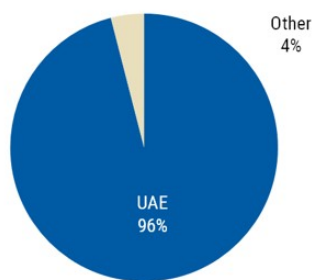
Source: Morgan Stanley Research

### Customer Concentration

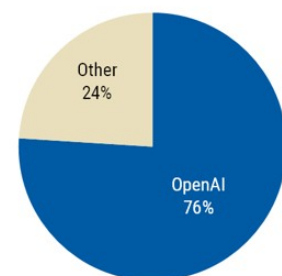
The agreement is expected to drive the majority of Cerebras' growth over the next several years and underpins a substantial portion of our revenue forecasts. We estimate revenue will increase from approximately \$0.8bn in 2026 to \$3.2bn in 2027 and \$6.4bn in 2028, with the agreement representing approximately 49%, 85%, and 76% of revenue, respectively. Related revenue is expected to increase by approximately \$4.5bn between 2026 and 2028, accounting for the majority of Cerebras' revenue growth during the period. While this concentration undoubtedly increases risk, we believe it should be viewed in the context of the company's stage of development, and see other fast-growing AI companies have similar amounts of customer exposure.

**Exhibit 16:** High dependence on small customers introduces volatility and risk

**Cerebras 2025 Customer Mix**



**Cerebras 2028E Customer Mix**



Source: Company data, Morgan Stanley Research

We do not assume any material contribution from new large customers during our forecast period, despite management indicating active discussions with prospective customers. In our view, available manufacturing capacity and cloud infrastructure will be

largely dedicated to fulfilling the contract to OpenAI through 2028, limiting Cerebras' ability to support additional large-scale deployments in the near term. As a result, we view 2026 as primarily supply constrained, with much of the capacity expansion through 2027 already allocated. Our forecasts therefore assume existing customers, including G42 and MBZUAI, continue purchasing systems and services at a steady pace, while excluding any meaningful revenue contribution from new customer wins over the next two years. **We begin incorporating revenue from new customers in 2028, when we estimate approximately 6% of revenue will be generated from customers not currently under contract**, implying potential upside should Cerebras secure additional capacity or accelerate customer acquisition ahead of our assumptions.

**Note on AWS deal:** Cerebras also has a multi-phase partnership with AWS. The companies have completed the proof-of-concept stage and are progressing toward initial deployments, but unlike the agreement with OpenAI, this is not a committed take-or-pay contract and remains significantly smaller in scale. As a result, we do not assume a meaningful revenue contribution from AWS over the next two years. When we begin modeling incremental customer additions in 2028, AWS represents the most likely source of upside given their partnership, although there is no guarantee.

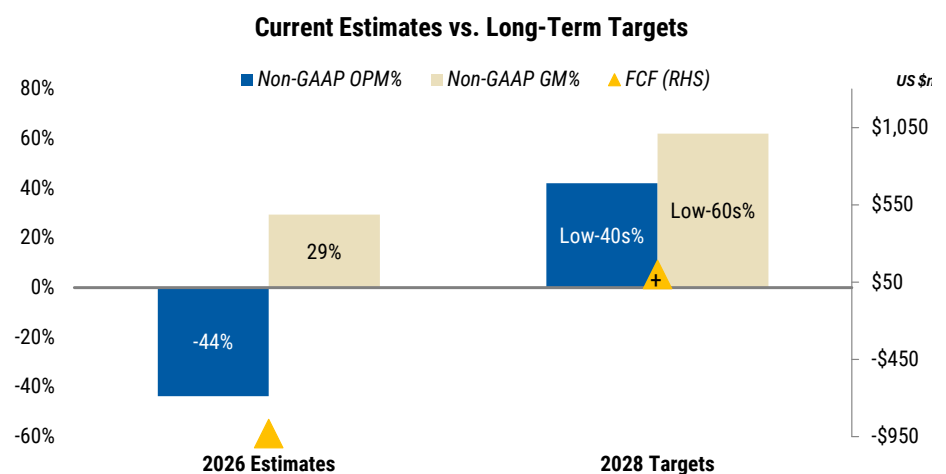
# Execution Matters: Capacity, Commercialization & Financial Complexity

**Main takeaway:** The demand side of the story appears largely de-risked; the challenge will be execution. Cerebras' ability to deploy capacity and monetize their technology will determine whether it can bridge the gap between current financial results and its long-term targets. Investors should also be mindful of the increasing divergence between core and GAAP results as cloud deployments scale and accounting adjustments become more significant.

Given that the OpenAI contract is take-or pay and driving all of the growth near-term, execution and commercialization is key. Factors that are important to understand beyond the technology are: the company's ability to secure and deploy cloud capacity, the mix between cloud and hardware revenue, and the differences between reported GAAP and Core (non-GAAP) results. In our view, how effectively Cerebras converts its technological advantage into sustainable revenue growth and profitability will be an important area to watch.

Near-term financials are quite far below targets. We expect Cerebras to approach their long-term target by the end of 2028, assuming the fulfillment of the base 750 deal with OpenAI and that they can secure nearly 1GW in capacity by 2028. Their ability to execute is something they will need to prove out over time.

**Exhibit 17:** 2026 MSe Financials vs. Long-Term Targets



Source: Company data, Morgan Stanley Research

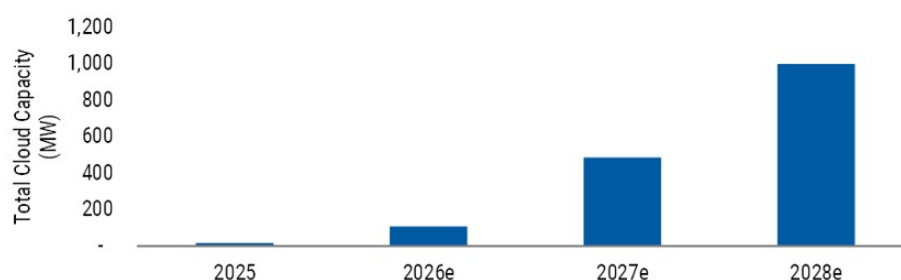
## Capacity Deployment Is the Critical Execution Variable

A key factor to execution relies on Cerebras securing enough capacity, and do so in an efficient manner to meet the terms and milestones of the OpenAI deal. Under our

forecasts, total deployed cloud capacity increases to nearly 1GW by 2028, including the 750MW committed to OpenAI as well as capacity supporting additional cloud customers. Achieving this scale will require Cerebras to secure leases, which are competitive, and make the necessary investments in power and other infrastructure costs. Future revenue growth is directly tied to the pace at which they can bring incremental capacity online.

This creates both opportunity and risk. Successful deployment supports substantial revenue growth, margin expansion, and increased utilization of the installed base. Conversely, delays in capacity deployment could impact revenue recognition and create pressure on customer commitments. We therefore view capacity deployment an important metric to track. **We believe that Cerebras has so far secured roughly ~240MW capacity** through their announced Bell Canada deal, the recent DGXX deal, and others, though it will take time for that to come online.

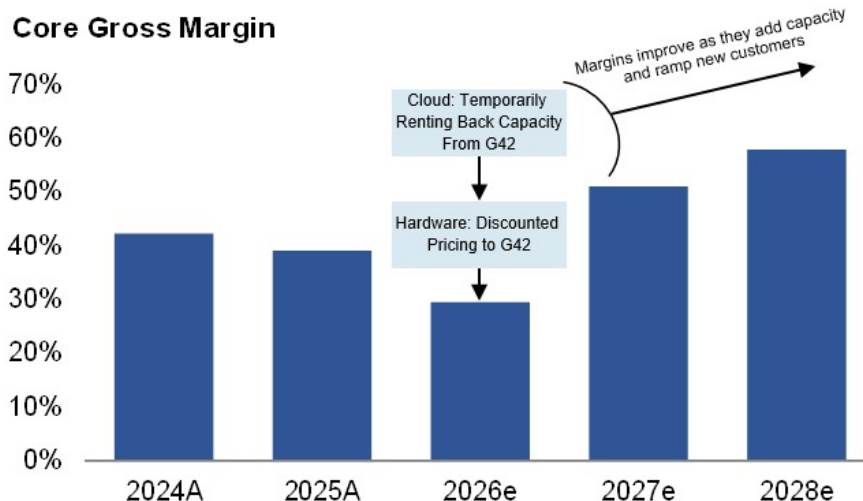
**Exhibit 18:** MSe Active Cloud Capacity Ramp



Source: Morgan Stanley Research

**We are already seeing the financial impact of the challenges associated with securing sufficient cloud capacity to support current contracts.** To meet their initial ramp, Cerebras is renting back capacity from G42, its largest historical customer, rather than serving all workloads through its own infrastructure. As a result, Cerebras incurs a temporary drag on cloud gross margins. We estimate total GAAP gross margins of 15% and cloud GAAP gross margins of 13% for 2026 (and total core gross margin of 29% and core cloud gross margin of 26%), well below management's long-term target of 60%+. We believe that reliance on G42 should gradually fall off by 2H 2027 as Cerebras brings additional owned capacity online. Accordingly, we assume cloud margins improve meaningfully thereafter, although the pace of expansion will ultimately depend on execution and the timing of capacity deployments.

Exhibit 19: MSe Gross Margin Forecast

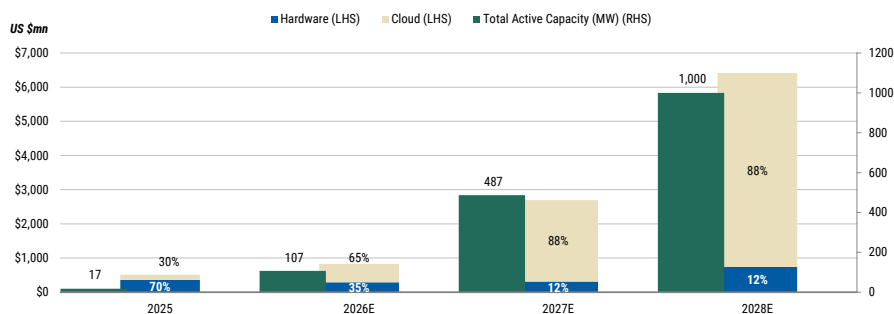


Source: Company Data, Morgan Stanley Research

### Cloud vs. Hardware Business Model

**It is still to be determined what the long-term final business model will be.** We model increasing contribution from cloud revenue over time, as we assume that the entirety of the contract with OpenAI is delivered through the cloud in our base case, and that cloud revenue becomes the majority of revenue by 2028. Given the size of the contract with OpenAI relative to the rest of the business, this assumption has a significant impact on the overall revenue mix and revenue recognition. Increased adoption of on-premise systems could drive a higher hardware mix than we currently assume, accelerating revenue recognition and improving near-term cash generation. As a result, we view the hardware business as a source of upside optionality.

Exhibit 20: Cloud revenue scales materially in our base case assumption



Source: Company Data, Morgan Stanley Research

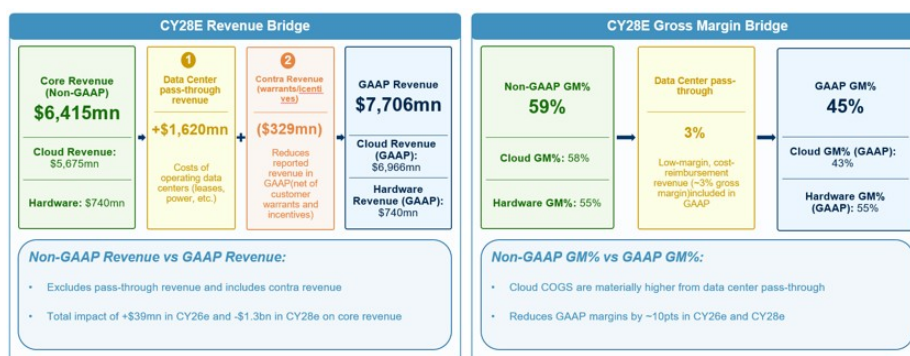
**The biggest difference would be on the balance sheet and cash flow timing as long-term margin economics of cloud vs. hardware sales are not assumed to be materially different.** The assumed mix between cloud and hardware deployments is the most important driver of cash flow and capital intensity. A greater cloud mix increases capex requirements and data center pass-through revenue, while a greater hardware mix pulls revenue forward and reduces capital requirements. The cloud-versus-hardware decision also materially impacts the bridge between core and GAAP financials due to the inclusion

of pass-through revenue and related costs (discussed further below). We expect both hardware and cloud margins to be near ~60% when the business is fully ramped.

## Understanding the Financials: GAAP vs Non-GAAP

Cerebras' financials require careful interpretation, as Cerebras uses non-GAAP (core) metrics that incorporate complex adjustments, including (1) contra-revenue associated with customer warrants and (2) pass-through data center revenue/costs related to reimbursements from OpenAI for operating the cloud. Over time, the data center pass through becomes the larger effect as the cloud business scales materially.

**Exhibit 21:** GAAP to Core Adjustments in CY28



Source: Company Data, Morgan Stanley Research

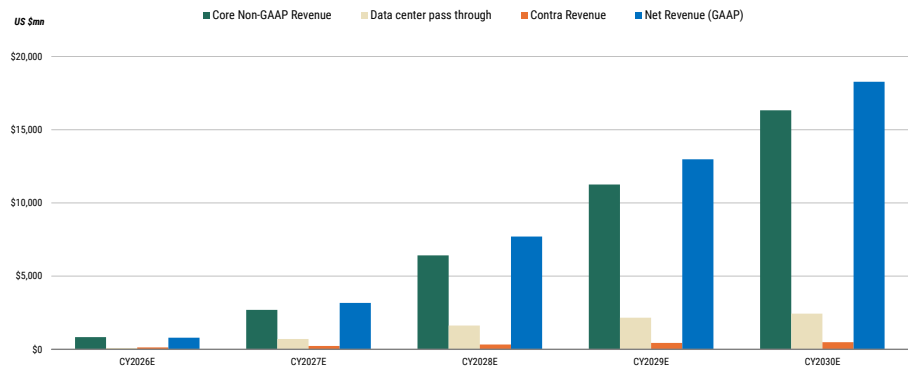
**Warrants:** Under GAAP, customer warrants are recorded as contra revenue and reduce reported revenue as the associated purchases are achieved, effectively functioning as a customer discount with no impact on cash flow. While we expect warrant-related contra revenue to decline to a LSD% of revenue over time assuming no new warrant issuances, we view it as a real economic cost. Consistent with peers that report revenue net of contra revenue, we include the impact of warrants in our valuation to maintain a conservative and comparable framework (see Valuation section for more detail). We would also note that if OpenAI opts into additional extensions, there would be additional dilution. Our model assumes the base 750MW.

**Data center pass through:** The agreement with OpenAI includes a "pass-through" mechanism for data center costs, where OpenAI will reimburse Cerebras for the costs to build and operate their cloud at a low markup. Pass-through revenue consists of two components: (1) monthly recurring costs, including rent, power, and operating expenses that scale with deployed capacity and (2) non-recurring engineering (NRE) costs associated with data center fit-out, including power, cooling, racks, and supporting infrastructure. While NRE is funded upfront by Cerebras and recorded as capital investment, the associated reimbursement from OpenAI is recognized ratably over the contract term. **As a result, pass-through revenue can be volatile and back-end loaded, but in general we estimate it to be ~30% of cloud revenue.** Management has decided to exclude pass-through revenue and associated costs when reporting core financial metrics. **As the cloud business scales, this will drive a significant difference between GAAP vs non-GAAP revenue, causing lower non-GAAP revenue but at higher profitability.**

Under our forecasts, the cumulative impact of these adjustments increases from approximately \$36 million in CY26 to roughly \$1.5 billion in revenue by CY28. This also

impacts the gross margin line. We estimate CY28 non-GAAP gross margins of approximately 55% versus GAAP gross margins of approximately 45%. We understand the rationale for focusing on core metrics, but believe it is important for investors to understand the differences between core and GAAP results as these adjustments create a meaningful divergence in reported financial performance.

**Exhibit 22:** Scaling of GAAP vs non-GAAP revenue



Source: Morgan Stanley Research

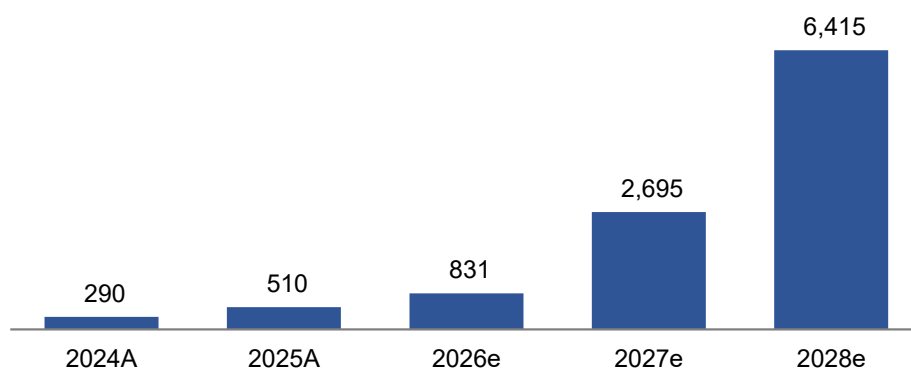
# Financial Outlook

## Revenue: CY26E-CY28E CAGR of 178%

We forecast revenue to inflect sharply over the next several years as Cerebras deploys the committed 750MW contract and scales its cloud infrastructure. Revenue growth is primarily driven by capacity deployment rather than customer additions, with their signed contract representing the majority of revenue through our forecast period. Our base case assumes the full 750MW commitment is delivered through Cerebras Cloud, resulting in cloud revenue becoming the dominant contributor by 2028. We do not assume meaningful revenue from new customers and assume steady purchasing activity from existing UAE customers.

**Exhibit 23:** CBRS Core Revenue Forecast

### Core Revenue (\$M)



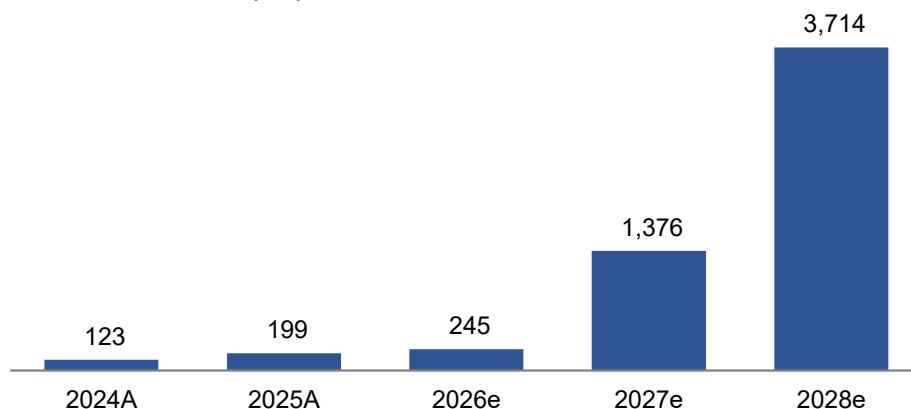
Source: Company Data, MSE

## Gross Margins: Temporary Pressure Before Expansion

We forecast gross margins to remain below long-term targets through 2026 as Cerebras relies on rented capacity from G42 to support the initial contract ramp. This creates a temporary drag on cloud profitability as the company incurs both the underlying infrastructure cost and the margin charged by G42. As owned cloud capacity comes online and reliance on rent-backs declines through 2027, we expect margins to improve meaningfully. Over the longer term, we believe margins will grow to their long term target of ~60%.

**Exhibit 24:** CBRS Core Gross Profit Forecast

**Core Gross Profit (\$M)**

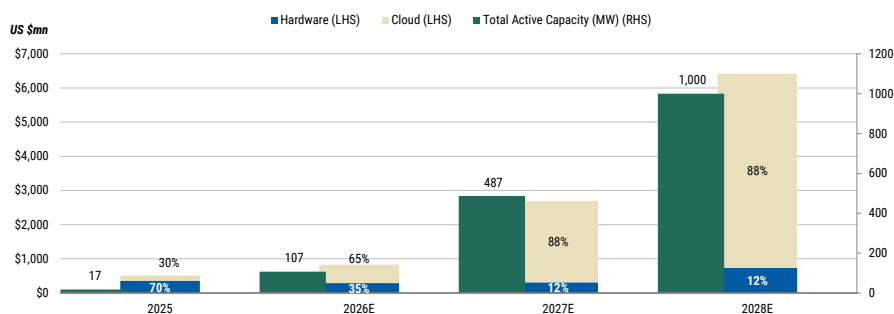


Source: Company Data, Morgan Stanley Research

**Cloud Deployment Is the Key Execution Variable**

The pace of capacity deployment remains the most important driver of financial performance. Under our forecasts, total deployed cloud capacity approaches ~1GW by 2028, including the 750MW already committed and additional capacity supporting other cloud customers. Future revenue growth is directly linked to Cerebras' ability to secure data center capacity, power, and infrastructure required to support this expansion. Any acceleration or delay in deployments would have a meaningful impact on revenue recognition, margins, and cash flow.

**Exhibit 25:** Cloud revenue mix expected to increase significantly

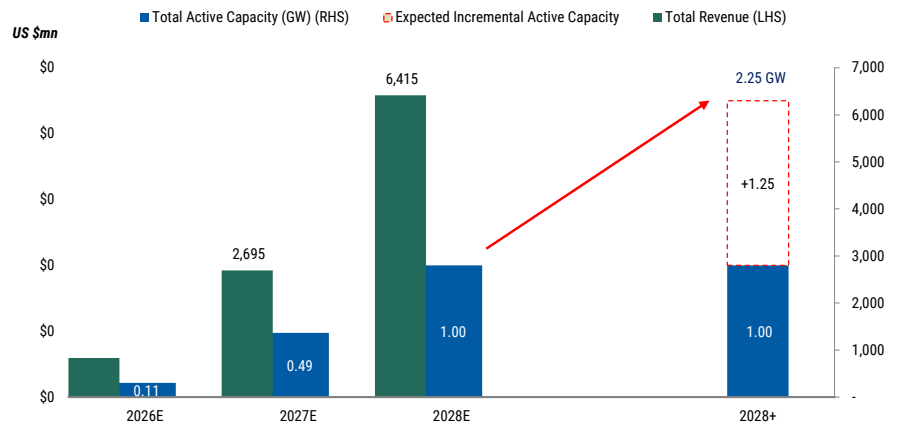


Source: Company Data, Morgan Stanley Research

**Long-Term Upside Remains Significant**

Our forecasts assume only the base 750MW commitment from OpenAI and does not include the exercise of extension options or the incremental 1.25GW capacity option available to OpenAI. Successful deployment of the initial commitment, combined with contract extensions or incremental capacity purchases, could materially increase revenue visibility and long-term contract value beyond our current forecast horizon.

**Exhibit 26:** Options for additional capacity imply a significant amount of revenue upside for CBRS



Source: Morgan Stanley Research

# Valuation Methodology

**Our Price Target is \$250.** We derive our \$250 price target using a 12x multiple on our 2028 revenue estimate adjusted for warrant contra revenue. While Cerebras operates a cloud-delivered infrastructure platform, we believe the company is best compared to AI semiconductor peers given its ownership of differentiated hardware IP and its economics, which remain more closely aligned with semiconductor companies than traditional cloud providers. The AI semiconductor peer group currently trades at approximately 9x 2028 EV/Sales at the median, with small and mid-cap AI companies trading above that level.

**We believe Cerebras deserves a premium to the broader AI semiconductor peer group and a multiple in-line with high-growth SMID-cap AI companies.** The company is expected to deliver substantially faster revenue growth than most AI infrastructure peers, supported by a contract which provides an unusual degree of revenue visibility through committed capacity purchases. In addition, our estimates do not fully reflect the potential upside from extensions, additional customer wins, or incremental capacity deployments. Also, we would note that all of their revenue is from AI compute. As a result, we believe a 12x 2028 EV/Sales multiple is appropriate and consistent with other high-growth AI peers.

Applying a 12x multiple to our 2028 revenue estimate of approximately \$6.0 billion yields our \$250 price target. The \$6bn 2028 revenue we use for valuation is lower than non-GAAP core revenue of \$6.4bn since it incorporates approximately \$329 million of contra-revenue associated warrants. We feel this is the most conservative revenue number and is comparable with semi peers with warrants.

- **\$385 Bull Case:** In our bull case, committed contract capacity comes online faster than expected, with the full 750MW deployed by 2028. Hardware demand also exceeds our base case, driving 21% growth in 2027 and 203% growth in 2028. Cerebras successfully expands cloud capacity to 1.5GW by 2028 and secures additional large customers. Together, these factors drive higher revenue growth, better utilization, and greater operating leverage than in our base case.
- **\$76 Bear Case:** In our bear case, contract capacity deployments are delayed, resulting in slower revenue ramp and lower utilization. Hardware demand remains weak, declining 9% in 2027 and increasing only 37% in 2028. We also assume cloud capacity reaches just 700MW by 2028, limiting the company's ability to support additional customers. Together, these factors result in lower revenue growth, reduced operating leverage, and weaker profitability than in our base case.

**Exhibit 27:** We compare Cerebras to other AI semis

Company	Ticker	Price	Market Cap		Sales CAGR	EV/Sales		EV/EBIT	
			(\$Bn)	EV(\$Bn)		26-28	2027	2028	2027
<b>Cerebras</b>	<b>CBRS</b>	<b>\$ 215.4</b>	<b>\$ 47</b>	<b>\$ 48</b>	<b>177.8%</b>	<b>2,695</b>	<b>6,415</b>	<b>263</b>	<b>2,046</b>
<b>AI Semis</b>									
NVIDIA	NVDA	\$ 219	\$ 5,292	\$ 5,266	31.9%	9.8x	8.1x	14.70x	12.43x
Broadcom	AVGO	\$ 419	\$ 1,983	\$ 2,088	39.4%	11.7x	9.3x	17.30x	13.77x
AMD	AMD	\$ 523	\$ 853	\$ 855	42.6%	11.2x	8.4x	34.68x	25.35x
Marvell	MRVL	\$ 316	\$ 277	\$ 284	40.8%	17.7x	12.8x	45.13x	31.76x
Astera	ALAB	\$ 358	\$ 61	\$ 64	34.3%	28.9x	22.9x	71.96x	54.96x
Credo	CRDO	\$ 218	\$ 40	\$ 40	43.1%	12.5x	9.4x	24.37x	18.37x
<b>Semis Mean</b>					<b>38.7%</b>	<b>15.3x</b>	<b>11.8x</b>	<b>34.69x</b>	<b>26.11x</b>
<b>Semis Median</b>					<b>40.1%</b>	<b>12.1x</b>	<b>9.3x</b>	<b>29.53x</b>	<b>21.86x</b>
<b>Cloud Providers</b>									
Coreweave	CRWV	\$ 108	\$ 48	\$ 90	77.5%	3.6x	2.2x	23.64x	12.05x
Oracle	ORCL	\$ 236	\$ 680	\$ 817	39.6%	7.3x	5.2x	19.13x	14.11x
Nebius	NBIS	\$ 260	\$ 57	\$ 80	143.4%	7.0x	3.9x	N/A	35.38x
<b>Cloud Mean</b>					<b>86.8%</b>	<b>6.0x</b>	<b>3.8x</b>	<b>21.4x</b>	<b>20.5x</b>
<b>Cloud Median</b>					<b>77.5%</b>	<b>7.0x</b>	<b>3.9x</b>	<b>21.4x</b>	<b>14.1x</b>
<b>Mean</b>					<b>38.7%</b>	<b>12.2x</b>	<b>9.1x</b>	<b>31.36x</b>	<b>24.24x</b>
<b>Median</b>					<b>40.1%</b>	<b>11.2x</b>	<b>8.4x</b>	<b>24.01x</b>	<b>18.37x</b>

Source: Factset, Morgan Stanley Research.  
Note: CBRS metrics are MSE. Data As of 6/5/25.

### Catalysts

- **Contract Extension:** Additional 1.25GW option is much larger than existing 750MW commitment. We would watch timing of tranche elections, extension decisions, and any discussion of capacity beyond the base commitment.
- **Second hyperscaler (or other large customer) beyond current pilot:** A meaningful hyperscaler deployment would validate the architecture and reduce customer concentration concerns.
- **Capacity deployment/margin inflection:** Getting capacity online faster than expected, bringing faster revenue recognition and higher gross margins from reduced reliance on current rent-backs from G42. This is the key operational catalyst.

### Risks

- **Capacity deployment/infrastructure risk:** One of the largest near-term risks is Cerebras' ability to secure, build, and deploy sufficient cloud capacity to meet the current contract ramps. Delays in manufacturing, power availability, data center buildout, or infrastructure deployment could push out revenue recognition, delay margin expansion, and impair the company's ability to meet contractual milestones.
- **Customer concentration risk:** The business remains highly concentrated over the next three years. While the take-or-pay structure provides visibility, any change in customer deployment plans, timing, or future capacity elections would have an outsized impact on Cerebras' financial performance.
- **Competitive risk:** While Cerebras appears to have differentiated technology, long-term success depends on converting that advantage into a broader customer base. If customers adopt competing solutions, revenue growth and utilization could fall short of expectations.





## 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

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., and/or Morgan Stanley Canada Limited. 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., Morgan Stanley Canada Limited 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](http://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: Joseph Moore.

### 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](http://www.morganstanley.com/institutional/research/conflictolicies). A Portuguese version of the policy can be found at [www.morganstanley.com.br](http://www.morganstanley.com.br)

### Important Regulatory Disclosures on Subject Companies

As of April 30, 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, Allegro Microsystems Inc, Ambarella Inc, Analog Devices Inc., Arm Holdings plc, Astera Labs Inc, Broadcom Inc., Cadence Design Systems Inc, Intel Corporation, IonQ Inc, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., 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 Analog Devices Inc., Broadcom Inc., **Cerebras Systems**, GlobalFoundries Inc, Intel Corporation, NXP Semiconductor NV, ON Semiconductor Corp., Semtech Corp..

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**, Intel Corporation, IonQ Inc, Micron Technology Inc., NXP Semiconductor NV, ON Semiconductor Corp., Semtech Corp., 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., 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 has received compensation for products and services other than investment banking services from Advanced Micro Devices, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Broadcom Inc., Cadence Design Systems Inc, GlobalFoundries Inc, Intel Corporation, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., Qualcomm Inc., Silicon Laboratories Inc., 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., 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 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, Ambarella Inc, Amkor Technology Inc, Analog Devices Inc., Broadcom Inc., Cadence Design Systems Inc, GlobalFoundries Inc, Intel Corporation, IonQ Inc, Marvell Technology Group Ltd, Microchip Technology Inc., Micron Technology Inc., NVIDIA Corp., NXP Semiconductor NV, ON Semiconductor Corp., Qualcomm Inc., Semtech Corp., Silicon Laboratories Inc., Synopsys Inc., Texas Instruments, Wolfspeed, INC.

Morgan Stanley & Co. LLC makes a market in the securities of Aeva Technologies Inc, Ambarella Inc, **Cerebras Systems**, Silicon Laboratories 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 May 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	1542	42%	465	51%	30%	707	43%
Equal-weight/Hold	1571	43%	369	40%	23%	723	44%
Not-Rated/Hold	3	0%	0	0%	0%	1	0%
Underweight/Sell	551	15%	86	9%	16%	201	12%
Total	3,667		920			1632	

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](http://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.. 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](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](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](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](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 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. 23374-2, 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](mailto:tejarshi.hardas@morganstanley.com); Grievance officer details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: [msic-compliance@morganstanley.com](mailto: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 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 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. 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.

## INDUSTRY COVERAGE: Semiconductors

COMPANY (TICKER)	RATING (AS OF)	PRICE* (06/05/2026)
<b>Joseph Moore</b>		
Advanced Micro Devices (AMD.O)	E (06/09/2024)	\$466.38
Aeva Technologies Inc (AEVA.O)	E (07/19/2021)	\$23.01
Allegro Microsystems Inc (ALGM.O)	O (02/13/2026)	\$46.39
Ambarella Inc (AMBA.O)	O (03/29/2016)	\$63.52
Amkor Technology Inc (AMKR.O)	E (11/08/2023)	\$64.95
Analog Devices Inc. (ADI.O)	O (11/16/2023)	\$401.39
Astera Labs Inc (ALAB.O)	O (05/11/2025)	\$317.06
Broadcom Inc. (AVGO.O)	O (06/09/2024)	\$385.73
Cerebras Systems (CBRS.O)	O (06/08/2026)	\$201.01
GlobalFoundries Inc (GFS.O)	E (10/28/2024)	\$75.53
Intel Corporation (INTC.O)	E (02/22/2023)	\$99.17
IonQ Inc (IONQ.N)	E (04/25/2023)	\$56.78
Marvell Technology Group Ltd (MRVL.O)	E (09/14/2015)	\$263.47
Microchip Technology Inc. (MCHP.O)	E (07/10/2024)	\$88.34
Micron Technology Inc. (MU.O)	O (10/06/2025)	\$864.01
Navitas Semiconductor Corp (NVT.S.O)	U (04/06/2025)	\$25.08
NVIDIA Corp. (NVDA.O)	O (03/16/2023)	\$205.10
NXP Semiconductor NV (NXPI.O)	O (02/11/2025)	\$295.96
ON Semiconductor Corp. (ON.O)	E (05/11/2025)	\$117.26
Qorvo Inc (QRVO.O)	E (10/28/2025)	\$98.28

Qualcomm Inc. (QCOM.O)	U (02/10/2026)	\$215.94
SanDisk Corporation. (SNDK.O)	O (03/03/2025)	\$1,559.32
Semtech Corp. (SMTC.O)	E (04/06/2025)	\$151.02
Silicon Laboratories Inc. (SLAB.O)	E (01/19/2021)	\$218.11
Skyworks Solutions Inc (SWKS.O)	E (11/28/2018)	\$73.57
Texas Instruments (TXN.O)	U (04/13/2020)	\$285.06
Wolfspeed, INC (WOLF.N)	NR (04/06/2025)	\$55.06
<b>Lee Simpson</b>		
Arm Holdings plc (ARM.O)	E (04/07/2026)	\$342.93
Cadence Design Systems Inc (CDNS.O)	O (02/14/2024)	\$376.19
Synopsys Inc. (SNPS.O)	E (02/27/2026)	\$464.85

Stock Ratings are subject to change. Please see latest research for each company.

\* Historical prices are not split adjusted.