Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

FOUNDATION

October 13, 2026 05:51 AM GMT

Semiconductors | North America

Highlighting 7 Interesting Semiconductor Privates

The growth in AI has reshaped the semiconductor landscape, creating new markets and opportunities where private companies are increasingly relevant. This note highlights 7 companies innovating around themes like: AI factory efficiency, disaggregated inference, physical AI, and Security/sovereign AI.

PrivateCompany

Private Companies Highlighted: Axiado, Ayar Labs, d-Matrix, Eliyan, SambaNova, Sima.ai, and Tenstorrent. Ayar Labs, SiMa.ai, and d-Matrix will be attending the Morgan Stanley Spark Private company conference in Dana Point later this week.

We are seeing the rate of change in semiconductor technology accelerate rapidly due to the proliferation of AI. After a long period of the semiconductor TAM being dominated by consumer devices like PC's and smartphones the emergence of the datacenter and more recently AI end markets has opened up increasing opportunities for new and smaller companies to quickly gain global scale and relevance. We think incumbents are well positioned to address changing needs in the market and the shift from training to inference to agenetic workloads - that said as the size of the market grows there's room for more winners.

Below we highlight 4 themes we are watching in the private company landscape, and 7 companies with various exposures to those drivers.

Theme #1: Improving AI factory efficiency: Getting the most tokens, and therefore the most revenue, out of a limited power budget is a growing focus for neoclouds and hyperscalers. **Axiado** is gaining traction with its Trusted Control/Compute Unit (TCU), a single chip that handles server management and hardware security and uses AI to tune power and cooling, so more of each rack's power goes to compute. **Ayar Labs** is pioneering scale-up CPO solutions that cut power consumption and increase bandwidth for AI factories. **Eliyan** is addressing the chiplet interconnect and copper bottlenecks: its die-to-die IP enables high-bandwidth, low-power chiplet links on standard packaging, while NuLight optical interconnects and dual-port custom HBM aim to extend bandwidth between and into XPUs. In XPUs themselves, upstarts **d-Matrix, SambaNova** and **Tenstorrent** are pushing the frontier in terms of performance per watt through architectural innovations around memory and scale-up/scale-out topologies.

Theme #2: Disaggregated and fast Inference: An emerging trend which involves breaking up certain portions of the AI inference workload and running them on different processors depending on that workloads requirements. Disaggregation is a preferred way for the AI ecosystem to work around the increasing costs of memory, especially HBM. The most common form of disaggregation is to run the pre-fill (compute bound) portion of the workload on an Nvidia GPU, and the de-code

MORGAN STANLEY & CO. LLC

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

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

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

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

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

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

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.

(memory bound) on an SRAM based accelerator like Cerebras or Groq. We have seen the TAM for disaggregation has accelerated sharply around the end of 2025 with the rise of agenetic coding and the largest players in the semis space are participating: NVIDIA acquired Groq around this time, AWS subsequently announced a disaggregated architecture pairing Trainiumalphawave for prefill with Cerebras for decode over EFA, and Cerebras has announced a similar approach with AMD.

In this note we outline the approaches to this market from: (1) **d-Matrix**, which recently joined Nvidia's NVLink Fusion ecosystem (link) and is targeting speculative decode as its entry point to disaggregation. In this model the GPU runs the large model while d-Matrix accelerates the smaller models that predict outputs in advance. **d-Matrix** next-generation product named Raptor stacks DRAM directly on top of compute logic, which the company says delivers SRAM-class bandwidth at roughly one-tenth of HBM's energy per bit with higher yields (2) **SambaNova**, whose RDUs (Reconfigurable Dataflow Units) and three-tier memory structure (on-chip SRAM, HBM and directly attached DDR) keep very large models closer to the chip and let customers switch between them quickly. The architecture is being used for on-prem fast inference in air-cooled racks, and in disaggregated inference environments alongside GPUs. (3) **Tenstorrent**, which is less focused on disaggregation, though it sees opportunities in that market, and is instead focused on building the best general-purpose system for AI performance. Its Tensix core pairs a matrix engine, a vector engine and local SRAM with small RISC-V cores that handle data movement, and hundreds of these tiles connect over an on-chip mesh network.

Theme #3: Embodied AI: Embodied AI is thought to be the next frontier in AI, spanning robotics, drones and smart infrastructure. While it is a much smaller market than the data center today, the TAM for manufacturing, autonomous systems and other applications that can be augmented by physical AI technologies creates a market opportunity of trillions of dollars over the longer term, in our view. **SiMa.ai** designs low-power inference silicon for physical AI. Its Modalix MLSoC combines a machine learning accelerator, a computer vision processor and Arm cores on a single chip, and runs physical AI models at under 10W. Those SoCs are enabled by its Palette Neat software, **SiMa.ai** is initially focused on robotics, drones and automotive. **Tenstorrent**'s RISC-V IP also addresses this market, as the company targets edge and automotive applications alongside the data center. It licenses its Ascalon CPU cores and Tensix AI IP so that customers can build custom silicon for robotics and autonomous systems.

Theme #4: Security, and sovereign AI: With heightened geopolitical tension around AI technologies and their growing importance in military technologies, we are seeing increased focus on two related areas: (1) security, including cybersecurity and IP protection for LLM builders, and (2) sovereign AI, where governments around the world want access to and control over their own AI infrastructure rather than relying on clouds in other regions. **Axiado**'s TCU (Trusted Control/Compute Unit) combines the baseboard management controller (BMC), hardware root of trust and TPM (Trusted Platform Module) functions with an on-chip AI engine in a single SoC. The product is vendor neutral, which makes it a possible building block for secure infrastructure across XPU suppliers. **Tenstorrent** is supplying RISC-V IP in a more

open model than vendors like Arm, allowing greater customization and control over the silicon with fewer licensing restrictions, giving customers more independence in designing their own chips.

Exhibit 1: We highlight 7 private companies targeting the AI TAM with a variety of hardware/software

Company	Market Opportunity / What They're Building	Notable Potential Competitors
Axiado	Trusted Control Unit (TCU) AI-driven hardware platform for servers & data centers — combines BMC, root-of-trust and TPM functions on one SoC.	ASPEED, Nuvoton, Lattice, Marvell, Microchip
Ayar Labs	Optical I/O chiplets (TeraPHY) + SuperNova light source Co-packaged optics replacing copper links to scale AI compute bandwidth and reduce power.	Broadcom, Marvell, Nvidia, Lumentum, Coherent, Celestial AI (Marvell), Lightmatter (private)
d-Matrix	Inference accelerators Digital in-memory compute purpose-built for low-latency, high-throughput generative AI inference in the data center.	Groq (Nvidia), Cerebras, SambaNova (private), Etched (private), AMD
Eliyan	Interconnect IP that moves more data into, and between, AI chips IP portfolio spans across optical, electrical, and memory to enable more efficent AI infrastructure	Synopsys, Cadence, Alphawave Semi (Qualcomm), Broadcom, Marvell, ams OSRAM
SambaNova	Inference accelerators + SambaNova Cloud Full-stack AI platform focused on fast, efficient inference for large models.	Nvidia, AMD, Groq, Cerebras, Etched (private), d-Matrix (private)
Sima.ai	Modalix MLSoC + Palette software Low-power platform for physical AI in robotics, automotive, drones and smart vision.	Nvidia, Qualcomm, Hailo (private), Ambarella
Tenstorrent	RISC-V CPUs + AI accelerators Open-source software stack, direct chip sales, and CPU/AI IP licensed to partners.	Nvidia, AMD, SiFive (private), Arm, Andes

Source: Company Data, Morgan Stanley

For Further Reading:

Semiconductors: How can the AI ecosystem work around memory bottlenecks? (5 Oct 2026)

Cybersecurity: Agents on the Loose: Cyber Faces an Identity Crisis (20 Aug 2026)

US Public Policy & Global Thematics: AI Sovereignty: Policy, Market, and Economic Implications (19 Aug 2026)

Technology: Bringing More Light to Scale-Up Networks: An Updated Scale-Up Primer (13 Jul 2026)

Cerebras Systems: Unique capability in fast inference drives large economic opportunity; Initiate at OW (8 Jun 2026)

Humanoids: A $5 Trillion Global Market (29 Apr 2025)

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



Axiado - AI Chips for Hardware Security

Who is Axiado: Founded in 2017, Axiado designs and engineers AI-driven hardware security processors, secure management controllers, and platform-security silicon that protect servers, 5G, and data center infrastructure from vulnerabilities embedded in existing processors. The core of the product portfolio is the AX3080 TCU (Trusted Control/Compute Unit), which anchors platform security directly in hardware to secure the control and management of heterogeneous infrastructure.

Company Mission: Axiado's mission is to secure digital infrastructure by embedding a new breed of AI-driven hardware based platform security in servers, 5G, network and power management controller infrastructure. Axiado believes data security should be the default solution, not an option.

Market opportunity

Axiado is focused on AI rack and platform security management. Cybersceurity is an increasing priority across enterprises globally, and with new threats from AI Axiado expects enterprise focus on this area to continue to increase. Axiado is building a new AI first security platform for the datacenter and see's an opportunity to replace fragmented TPM, BMC, LOM, firmware / software, and FPGA-based RoT approaches to a more purpose built and seamless product. The central idea is that rack and pod level infrastructure redefines what a management controller must see, secure, and coordinate - creating a shift from board / tray management toward centralized rack and pod management. It's also worth noting that adding accelerated compute capabilities can improve the overall performance of the system, with thermal management a compelling use case (see here).

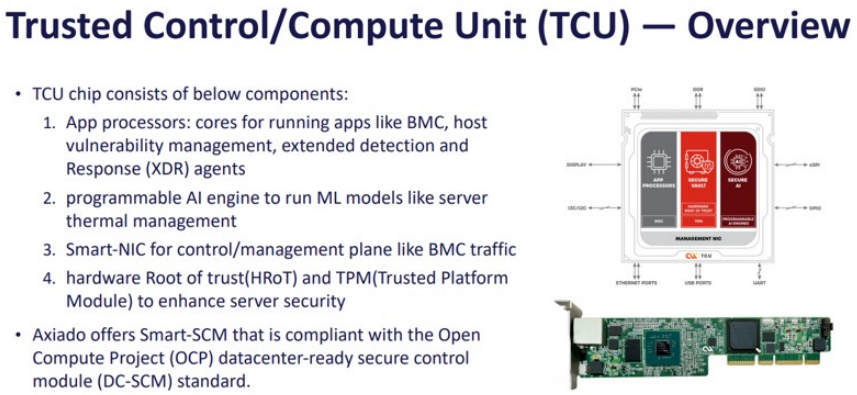
Product overview: Axiado's TCU combines multiple board/rack management solutions into one product. with a number of form factors outlined below. The company is vendor neutral, and sees an opportunity across XPUs as well as CPU servers.

Exhibit 2: Axiado product overview



Source: Company Data, Morgan Stanley Research

Exhibit 3: Axiado TCU overview



Source: Company Data [https://axiado.com/wp-content/uploads/2024/04/OCP24R-Increasing-Energy-Efficiency-using-DRL-BO.pdf], Morgan Stanley Research

Competitive differentiation

Axiado's differentiation is combining hardware, software, AI, and security in a vendor-neutral TCU platform. The company says it can attach across rack architectures and provide both root-of-trust and rack-level management.

Management has argued that rack vendors lack security expertise and that hyperscalers currently pay multiple vendors / people to handle functions Axiado can automate. The company also believes programmability is inappropriate for security, an advantage against FPGA-based solutions.

Notable Potential Competitors: Lattice Semiconductor/AMI, Aspeed, Nuvoton, Marvell, Microchip

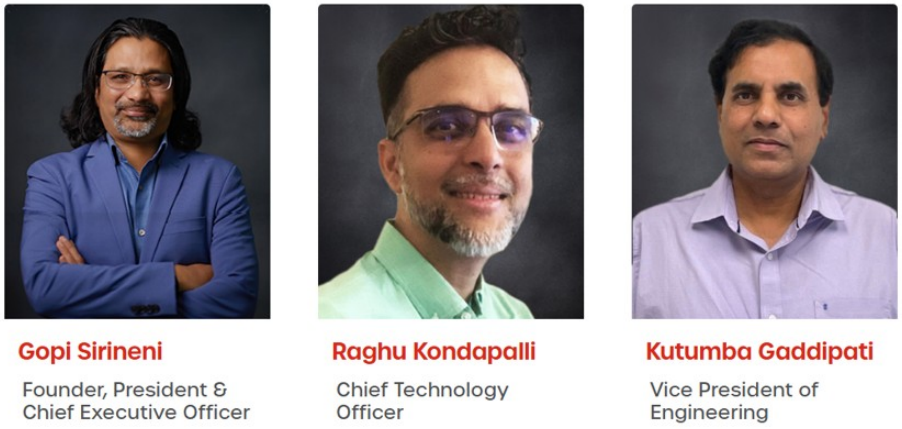
Recent developments

- **In November 2025 Bill Brennan, President and CEO of Credo Technology Group, joined Axiado's Board of Directors,** bringing over 25 years of leadership experience in the semiconductor and connectivity industries, including having guided Credo from a startup to a publicly traded company. His appointment reflects Axiado's push to build out its leadership bench with Tier-1 qualification and hyperscaler GTM experience.
- **In October 2025 at the OCP Global Summit, Axiado announced a collaboration with Inventec to integrate its TCU into Inventec's NVIDIA Blackwell GB200-based servers and Intel-based MGX platforms,** embedding hardware-anchored security and autonomous AI intelligence directly at the silicon level. The partnership pairs Axiado's SCM3002 card (built on the AX3000 TCU) for Intel/AMD platforms with the SCM3080-MT card (built on the AX3080 TCU) for NVIDIA platforms, unifying root-of-trust, BMC, TPM, secure networking, and AI inference into a single SoC. Embedded AI agents proactively detect anomalies, block attacks, and optimize system performance in real time. Capabilities the companies say software-only solutions cannot match, allowing threats to be stopped before the OS even boots.

Team

- **Gopi Sirineni , Founder, President & CEO,** founded Axiado in 2017 after serving as Vice President of Wi-Fi SON at Qualcomm (2012–2020) and as an Executive Board Advisor at Cognitive Systems Corp.
- **Raghu Kondapalli, Chief Technology Officer,** has served as Axiado's CTO since 2022, prior to which he was CTO of Intel's Data Center Group (2012–2020).
- **Kutumba Gaddipati,Vice President of Engineering,** joined Axiado in 2025, bringing prior experience as Director of SoC Engineering at SiMa.ai (2022–2025) and Director of SoC Architecture and Design at Akeana (2021–2022).

Exhibit 4: Axiado: Key Leadership



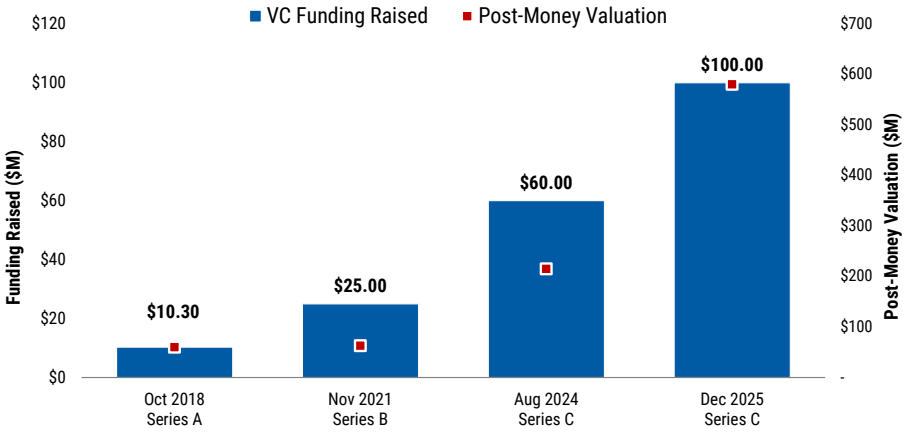
Source: Axiado

Investors and Funding

Axiado has raised a cumulative $195.3mn since 2018, the most recent round in December

2025 raised $100.0mn at a post money valuation of $0.58bn. Investors in the most recent round include Maverick Silicon (lead), Prosperity7 Ventures, Orbit Venture Partners, Crosslink Capital, and Nosterra Ventures.

Exhibit 5: Axiado Historical Funding Rounds (Sourced from PitchBook)



Source: Pitchbook, Morgan Stanley

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

Ayar Labs - Optical Interconnects for AI Data Centers

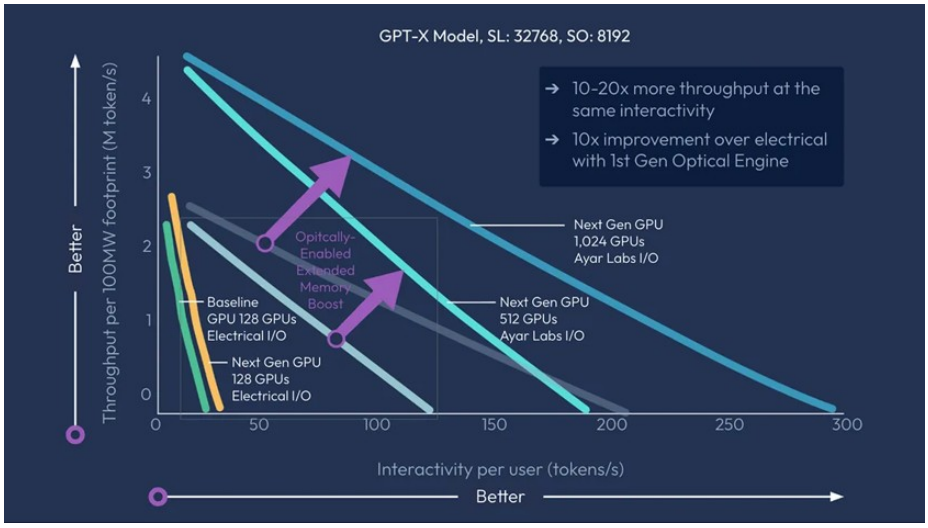
Who is Ayar Labs: Founded in 2015, designs and engineers optical interconnect chiplets and lasers that move data between chips using light instead of electrical signals, accelerating data movement for next-generation AI infrastructure. Ayar Labs Photonic Fabric delivers very high bandwidth at nanosecond latencies to break through the electrical I/O bottlenecks limiting AI data center compute scaling.

Company Mission: Ayar Labs' mission is to solve the foundational connectivity bottleneck in computing systems using a new generation of electronic/photonic integration and architectures.

Market opportunity

Ayar Labs is addressing the bandwidth and distance bottleneck in AI infrastructure through optical I/O and co-packaged optics. Management sees electrical I/O becoming the bottleneck, forcing more compute into racks and driving up power per rack. Removing the bandwidth-distance limitation could allow larger scale-up domains without the same constraints. The company sees optical engines, optical I/O, and CPO as increasingly central to AI infrastructure, with potential deployment in XPU-to-XPU, rack-to-rack, and future direct-attach optical systems. Compared to standard copper and pluggable optics Ayar sees 10x better bandwidth, 10x lower latency, and 3-5x the power efficiency. Ayar's solution is designed to be protocol agnostic, meaning the opportunity is applicable beyond any particular XPU or interface.

Exhibit 6: CPO value proposition for AI factories



Source: Ayar Labs

Product Overview: The company provides PIC + EIC solutions, external light sources, and optical I/O chiplets. The core product is the optical engine (TeraPHY), which integrates a

photonic integrated circuit (PIC) with an electronic integrated circuit (EIC) to convert electrical signals into light and back again directly at the chip package, replacing traditional copper I/O with in-package optical connectivity. This optical engine is powered by an external, remote light source (SuperNova), which keeps the laser off the hot compute die entirely — improving reliability and simplifying thermal management. Together, these components form Ayar Labs' full-stack optical interconnect platform that lets GPUs, CPUs, and switches communicate at high bandwidth and low latency across a chip, a rack, or multiple racks. Ayar's solution can deliver roughly 10–20x better performance-per-watt than copper interconnects while removing the reach limitations that constrain today's largest AI clusters.

Exhibit 7: Ayar Labs, Wiwynn Collaboration at OFC

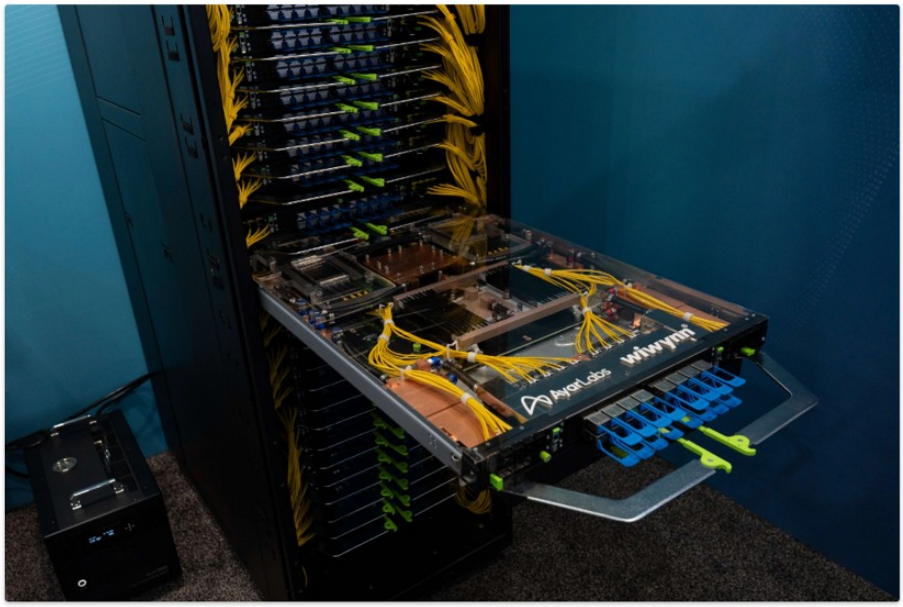


Figure 2. A L10 chassis designed and manufactured by Wiwynn that features Ayar Labs co-packaged optics and ELSFP light sources, as well as AI ASICs, advanced fiber management, and cold-plate liquid cooling.

Source: Ayar Labs [https://ayarlabs.com/blog/building-a-production-ready-optically-connected-rack-for-ai-scale-up/]

Competitive differentiation

Ayar labs brings optical engines directly onto advanced packages, integrating photonics with the XPU itself. This tight physical and integration creates a set of core competencies spanning photonics design, advanced packaging, and high-volume manufacturing that is difficult for a component-level optics supplier to replicate without also owning the packaging and system-level integration work. Ayar believes that the first generation of scale-up CPO will be dominated by TSMC's COUPE, and that mixing and matching EICs and PICs across different foundries or vendors will not be straightforward in the near term. Ayar Labs in one of just a few companies that have been public in building on TSMC's COUPE platform for several years.

On the technology roadmap itself, the company is skeptical that electro-absorption modulator (EAM)-based approaches can be productized in high volume within five years, and has instead bet on micro-ring-based optical systems paired with remote laser

architectures. Because the laser sits away from the compute die, this design sidesteps thermal constraints that would otherwise limit how memory and compute can be positioned relative to the optical engine, giving the company more architectural flexibility in how it lays out future systems, and reducing the risk that laser reliability issues affect the compute silicon.

Ayar labs is positioned as an alternative to switch/networking vendors who bundle optics into a broader portfolio, being able to bundle is an advantage on paper, but management believes that customers building next-generation GPU/XPU fabrics will prefer to avoid being locked into any single vendor's ecosystem. Management does expect large GPU/AI-accelerator vendors to build some internal CPO capability over time given how strategically important optical I/O is becoming.

NPO approaches should prove complementary rather than directly competitive, and sees them as more likely to absorb overflow demand during periods of tight co-packaged-optics supply than to displace the tightly integrated approach in mainstream scale-up deployments.

Notable Potential Competitors: Broadcom, Marvell, Nvidia, Lumentum, Coherent, Celestial AI (Marvell), Lightmatter (private)

Team

- **Mark Wade, Co-Founder and Chief Executive Officer**, leads Ayar Labs' strategy and growth as the company commercializes optical I/O chiplets for AI infrastructure; he previously served as Ayar Labs' CTO and SVP of Engineering, and co-invented the company's core photonics technology at MIT and UC Berkeley between 2010 and 2015, holding a PhD in Electrical Engineering from the University of Colorado Boulder.
- **Lisa Cummins Dulchinos, Chief Financial Officer,** oversees financial strategy and has led multiple successful funding rounds for the company, building the financial infrastructure to support Ayar Labs' scale-up as a VC backed hardware company.
- **Vladimir Stojanovic, Co-Founder and Chief Technology Officer**, led the team that designed the world's first processor to communicate using light prior to founding Ayar Labs, and previously was a professor of EECS at UC Berkeley and a key developer of Rambus's high-speed link technology, holding a PhD from Stanford University.

Exhibit 8: Ayar Labs: Ley Leadership



Mark Wade

Co-Founder and
Chief Executive Officer



Lisa Cummins Dulchinos

Chief Financial Officer



Vladimir Stojanovic

Co-Founder and
Chief Technology Officer

Source: Ayar Labs

Recent developments

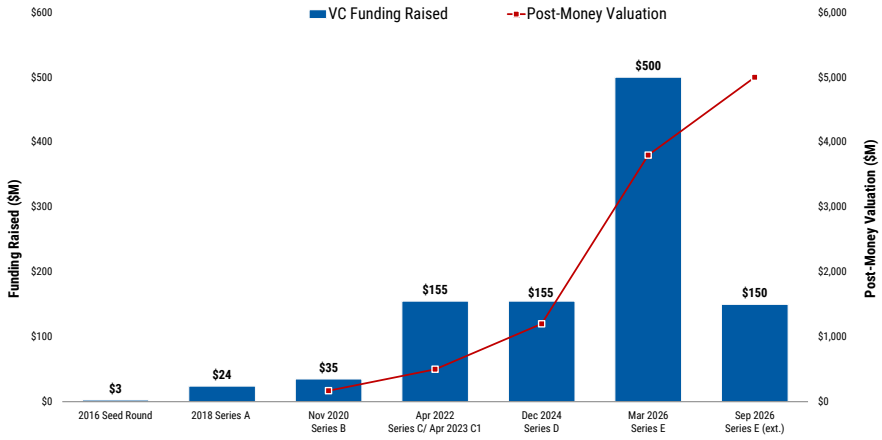
- **In June 2026 Ayar Labs joined the NVIDIA NVLink Fusion ecosystem**, making its co-packaged optics products optically and electrically compatible with NVIDIA's optical and SerDes technologies — enabling hyperscalers to deploy heterogeneous compute, including custom silicon, within NVIDIA rack-scale platforms while preserving NVLink-based architecture investments. The announcement built on Ayar Labs' close of a $500M Series E round (which included NVIDIA as a participant).
- **In March 2026 Ayar Labs and Wiwynn announced a strategic partnership to build optically connected, rack-scale AI systems, unveiling a joint reference design at OFC 2026** featuring a 100% liquid-cooled, HVDC-enabled architecture that scales to 1,024 accelerators per rack using Ayar's TeraPHY optical engines and SuperNova remote light sources — a design the companies plan to further showcase at Computex Taipei.
- **In September 2025 Ayar Labs and Alchip unveiled co-packaged optics for AI datacenters**. The companies showed working silicon, debuting a multi-chip package (MCP) built with Alchip that integrates eight optical engines, using TSMC's COUPE platform.
- **In October 2025 Ayar Labs formed a strategic partnership with GUC** to integrate co-packaged optics (CPO) into GUC's advanced ASIC design services, targeting hyperscaler AI infrastructure.
- I**n September 2025 Ayar Labs and Alchip unveiled co-packaged optics for AI datacenters**, enabling ultra-high bandwidth, low latency, and efficient multi-rack AI cluster scale-up.
- **In July 2025 Ayar Labs expanded its U.S. and Taiwan operations** and appointed Vivek Khanzodé as VP of Engineering to support high-volume CPO growth.
- **In March 2025 Ayar Labs unveiled the Universal Chiplet Interconnect Express (UCIe) optical interconnect chiplet,** aimed at maximizing AI infrastructure

performance and efficiency.

Fund-raising

Ayar Labs has raised a cumulative ~$1bn since 2018, the most recent rounds in 2026 raised $650.0mn at a post money valuation of between $4-5bn+ . Investors in the most recent round include Neuberger Berman (lead), ARK Invest, Artisan Partners, Greycroft, Insight Partners, Qatar Investment Authority, Sequoia Global Equities, 1789 Capital, AMD, Alchip Technologies, MediaTek, and NVIDIA.

Exhibit 9: Ayar Labs: Historical Equity Funding



Source: Pitchbook, Company Data, Morgan Stanley Research

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

d-Matrix - In-Memory Compute for Fast AI Inference

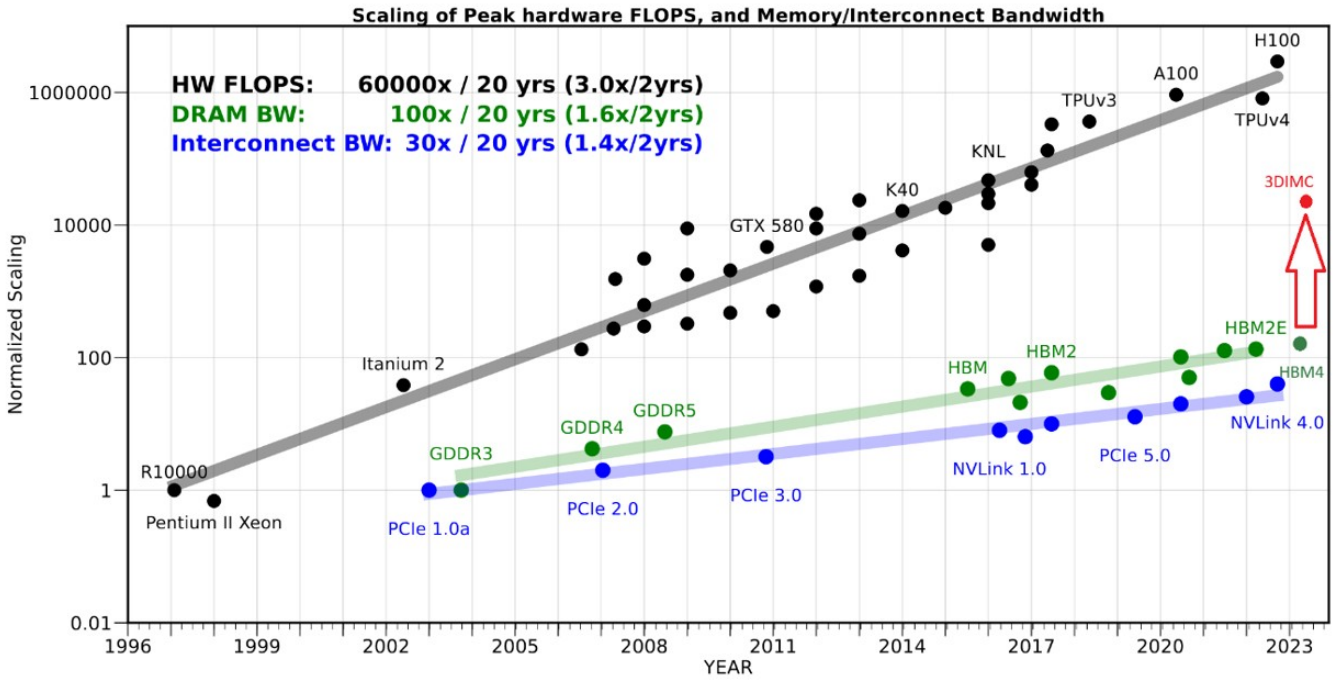
Who is d-Matrix: Founded in 2019, d-Matrix builds an XPU platform for generative AI inference, using digital in-memory computing (DIMC) techniques and a chiplet based architecture to reduce the memory compute bottleneck.

Company Mission: d-Matrix's mission is to dramatically reduce the cost, power consumption, and latency of deploying large language models and other transformer-based workloads at scale, making generative AI commercially viable for enterprises everywhere.

Market opportunity

d-Matrix is targeting low-latency inference and the emerging premium token market, particularly as workloads become more interactive, agenetic, and memory-bound. Management views disaggregated inference as the path forward and highlighted three forms: prefill/decode disaggregation, expert disaggregation, and speculative decoding. The company is targeting speculative decode to start .

Exhibit 10: The growing demand for inference should increase the value of solutions to the memory wall



Memory bandwidth improvement has lagged behind compute performance gains over the past 30 years. d-Matrix 3DIMC technology promises to help close the gap.

Source: d-Matrix [https://www.d-matrix.ai/scaling-ai-inference-with-3dimc/]

Product overview: d-Matrix builds XPUs around a memory-centric architecture designed

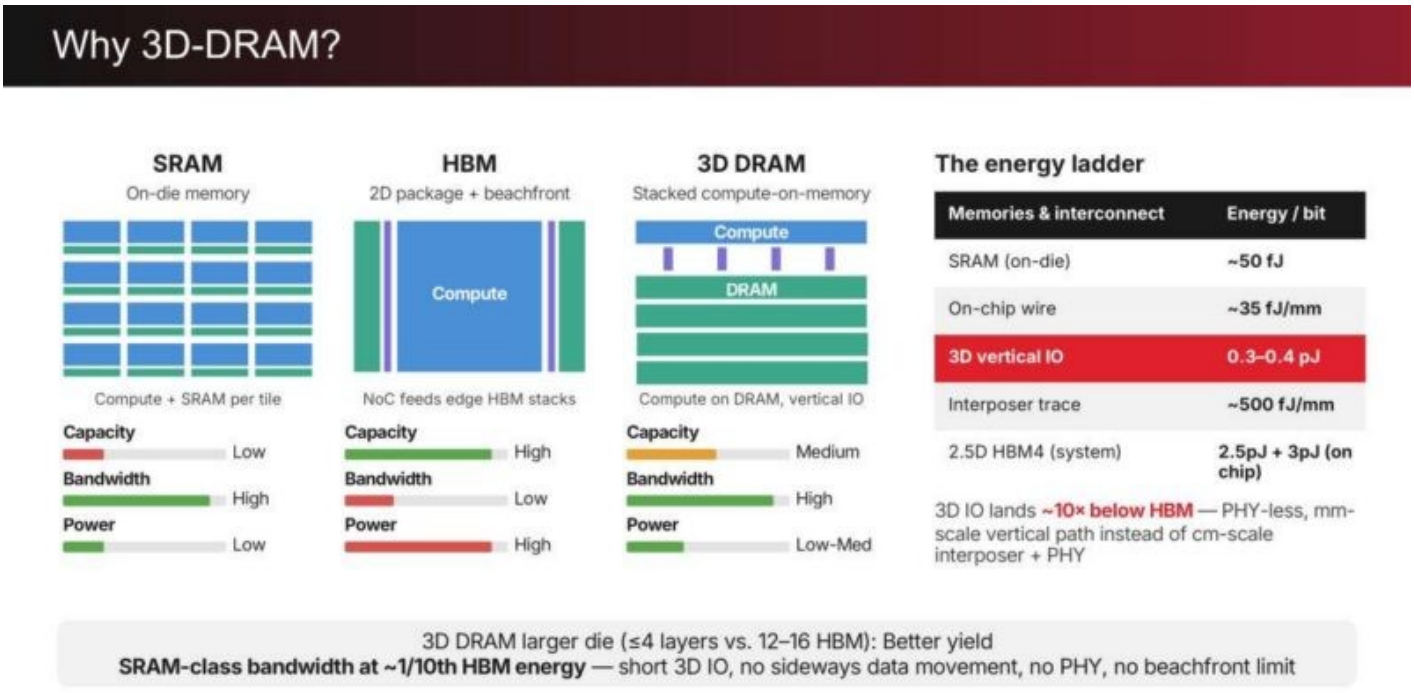
to improve token throughput and latency by reducing memory bottlenecks. This is enabled by Digital In-Memory Compute (DIMC), d-Matrix's proprietary technology that embeds multiply-accumulate logic directly into the SRAM bit cell itself, rather than moving weights back and forth between separate memory and compute units as in a conventional architecture. Because computation happens where the data already lives, DIMC delivers higher memory bandwidth than HBM, while retaining full digital precision (unlike earlier analog in-memory compute attempts).

The current gen product,"Corsair", is an SRAM-based product in production now, the product is a chiplet-based PCIe Gen5 accelerator card combining eight 6nm chiplets, 2GB of on-chip SRAM, and up to 256GB of off-chip LPDDR5X capacity memory, purpose-built for generative AI inference.

The next-gen product, Raptor, adds DRAM stacked directly on top of TSMC 4nm logic dies and is capable of running 2.8T-parameter Kimi 3 on a single rack. The compute is d-Matrix's proprietary Digital In-Memory Compute (DIMC); Raptor will use RISC-V IP from Andes for the control core, which also handles vector compute and activation-function offload. d-Matrix 3D-DRAM can achieve better yield than HBM, with SRAM class bandwidth at 1/10th the power consumption.

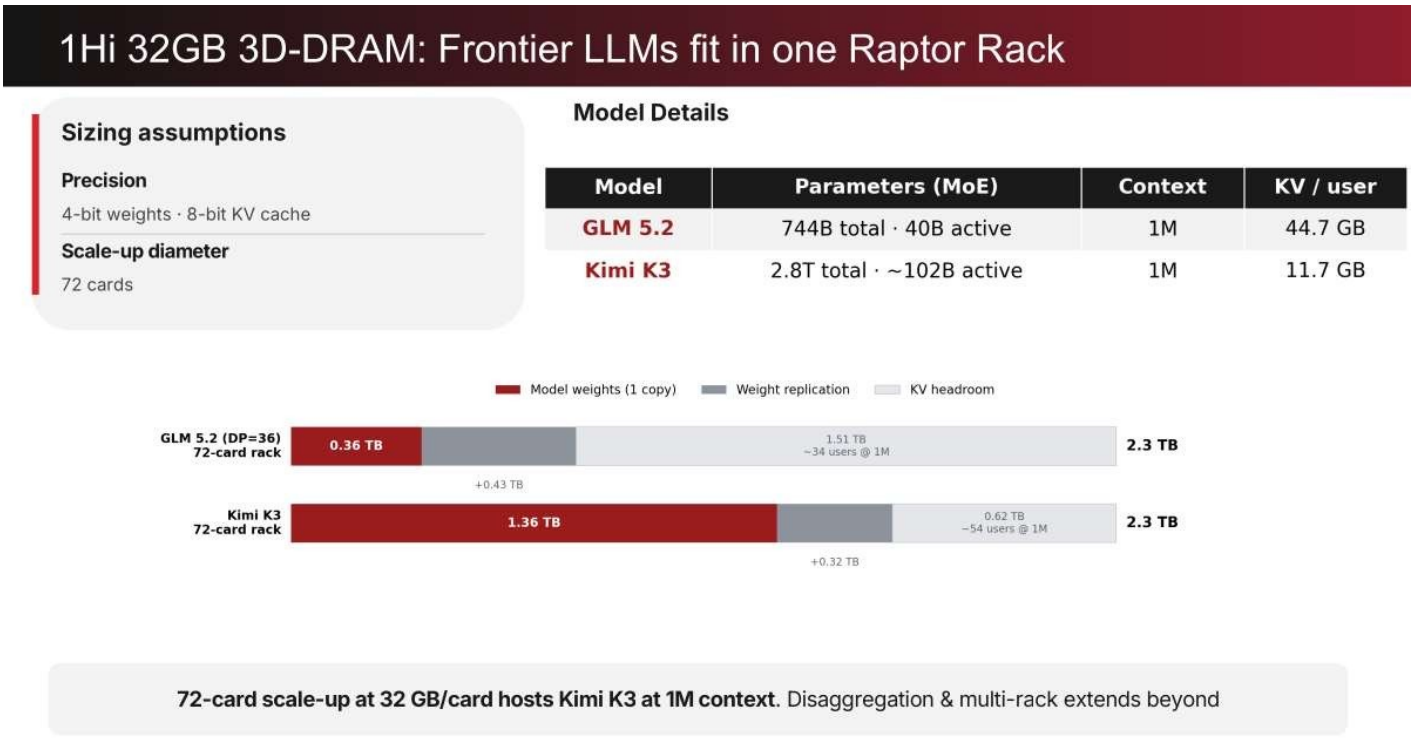
For scale-up, d-Matrix has developed its own die-to-die and card-to-card interconnects, DMX Link connects chiplets within a package at roughly 1 TB/s, while DMX Bridge cards link multiple Corsair cards together into a single all-to-all cluster (up to 16 chiplets). For scale-out d-Matrix invested JetStream, an ethernet-based NIC that connects multiple Corsair-equipped servers into a larger inference cluster. In early 2026, d-Matrix acquired the data center business of GigaIO, adding FabreX-based composable PCIe fabric technology to extend rack-scale scale-up further — building on the companies' prior joint "SuperNODE" reference system, which paired Corsair with GigaIO's fabric to demonstrate GPU-free inference at rack scale.

Exhibit 11: 3D DRAM can offer SRAM bandwidth at 1/10th HBM energy



Source:d-Matrix [https://www.servethehome.com/d-matrix-raptor-3d-dram-accelerator-for-generative-inference-at-hot-chips-2026/]

Exhibit 12: d-Matrix "raptor" can run frontier LLMs on a single rack



Source: d-Matrix [https://www.servethehome.com/d-matrix-raptor-3d-dram-accelerator-for-generative-inference-at-hot-chips-2026/]

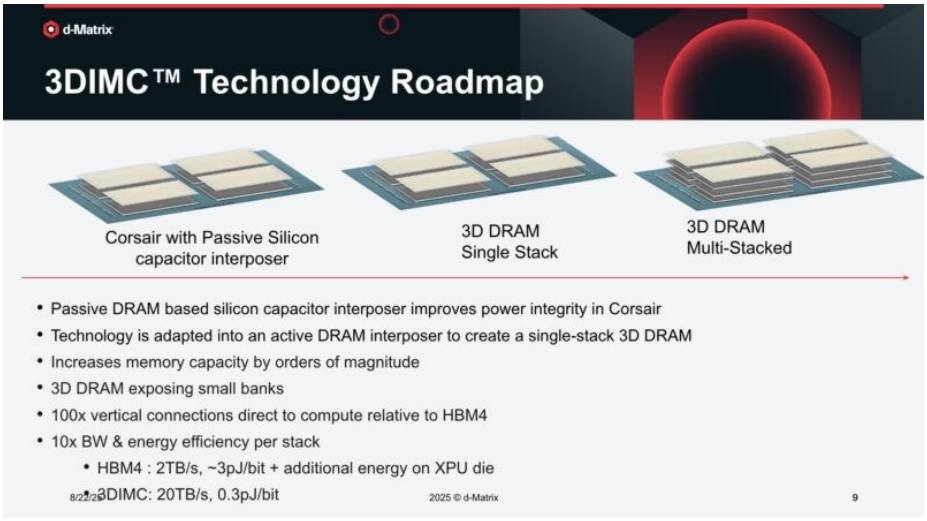
Competitive differentiation

d-Matrix's core differentiation is the memory architecture, with 3D DRAM an industry first technology per the company, positioning the company between HBM enabled GPUs and SRAM based approaches - offering customers higher speeds than HBM and greater capacity than SRAM. This allows d-Matrix to enable disaggregated inference with fewer processors, which becomes more important as memory demands from new models continue to scale. The company has also built its own ethernet-based scale-out interconnect as an alternative to the NvLink networking standard, aiming to give customers flexibility across vendors.

On the software side, the company is investing heavily in its own kernel and SDK strategy to shorten the time it takes to port new models onto its hardware. Kernels are written in C, a widely-understood language that lets frontier AI labs get models running on the hardware quickly, with AI coding tools now capable of accelerating that process further. Management believes this SDK-first approach differentiates the company from other specialized inference hardware players.

Notable Potential Competitors: Broadcom, Cerebras, SambaNova (private), Nvidia (groq), AMD, Etched (private)

Exhibit 13: d-Matrix 3DIMC technology roadmap



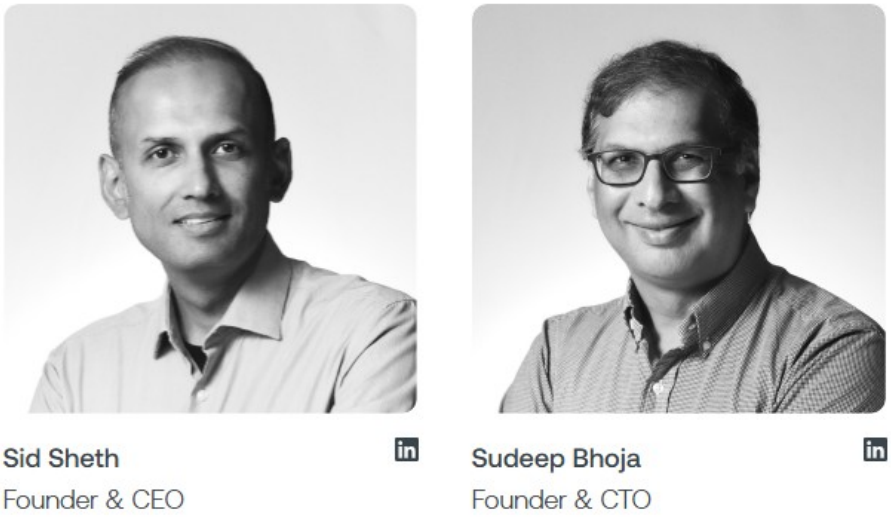
Source: d-Matrix, Morgan Stanley

Team

- **Sid Sheth, Co-Founder, President and Chief Executive Officer,** co-founded d-Matrix in 2019 to address the growing memory bottleneck in AI inference, betting early on the shift toward transformer-based generative AI workloads; he previously served as SVP/GM of the Networking Business Unit at Inphi Corp. (2011–2019) and as Director of Marketing for Network Connectivity at Broadcom (2007–2011).
- **Sudeep Bhoja, Co-Founder and Chief Technology Officer,** leads d-Matrix's chip architecture and technical strategy, having co-founded the company alongside Sheth after years of working together on networking silicon; he previously served as CTO of Networking Interconnect at Inphi Corp. (2012–2019) and as Technical

Director at Broadcom (2006–2012).

Exhibit 14: d-Matrix: Key leadership



Source: d-Matrix, Morgan Stanley

Recent developments

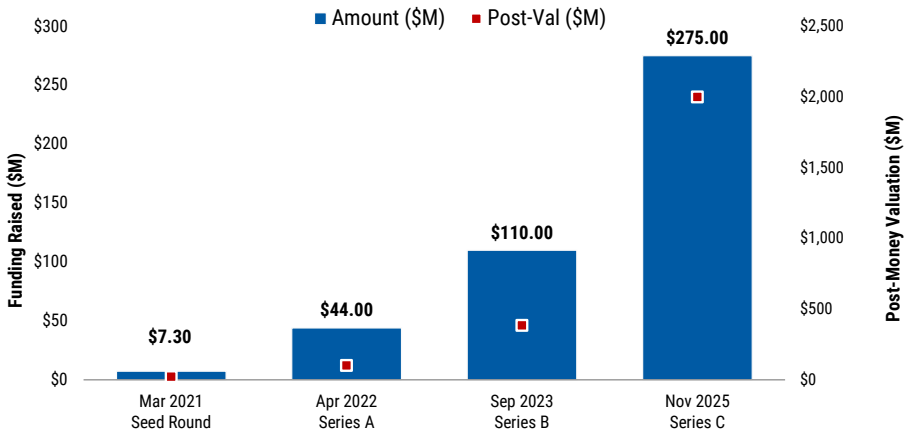
- **In September 2026, d-Matrix announced a multi-year collaboration with NVIDIA to connect its next-generation Raptor XPUs to NVIDIA's AI infrastructure through NVLink Fusion**, building on NVIDIA's MGX rack architecture alongside Vera CPUs, BlueField-4 DPUs, ConnectX-9 SuperNICs and Spectrum-X networking. The racks are designed to run next to NVIDIA GPU systems such as Vera Rubin NVL72 for disaggregated inference, and d-Matrix is working with Astera Labs on connectivity. Integrated MGX systems are expected initially in the fourth quarter of 2027.
- **In August 2026 d-Matrix acquired Wallaroo.ai**, adding AI inference deployment and orchestration software to its platform
- **In August 2026 d-Matrix presented its Raptor 3D-DRAM accelerator** at Hot Chips 2026, showcasing a stacked memory architecture delivering roughly 100 TB/s of bandwidth at a fraction of HBM4's power consumption.
- **In July 2026 d-Matrix and Parasail announced a deployment combining the Corsair platform with NVIDIA Hopper and Blackwell GPU architectures** within Parasail's cloud data center.
- **In April 2026 d-Matrix acquired the data center business of GigaIO,** including the SuperNODE platform and FabreX PCIe-based memory fabric, extending its platform beyond silicon into full system-level rack-scale infrastructure and adding a new engineering hub in Carlsbad, California.
- **In March 2026 d-Matrix and Alchip announced a collaboration** on what the companies describe as the world's first 3D DRAM solution for AI inference, aimed at delivering up to 10x faster inference than HBM4-based solutions via the upcoming Raptor accelerator.
- **In November 2025 d-Matrix partnered with Andes Technology** to integrate Andes' RISC-V CPU IP into Raptor, d-Matrix's next-generation accelerator.

- **In October 2025 d-Matrix announced SquadRack**, a rack-scale reference architecture developed with Broadcom and Arista, offering 3x cost-performance and 10x token speed for AI inference at the Open Compute Summit.

Fund-raising

d-Matrix has raised a cumulative $446.3mn since 2021, the most recent round in November 2025 raised $275.0mn at a post money valuation of $2.00bn. Investors in the most recent round include Bullhound Capital, Triatomic Capital, and Temasek (co-leads), with Qatar Investment Authority, EDBI, M12, Nautilus Venture Partners, Industry Ventures, and Mirae Asset.

Exhibit 15: d-Matrix Historical Funding Rounds



Source: Pitchbook, Morgan Stanley

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

Eliyan - Breaking down the memory wall with chiplet interconnects

Who is Eliyan? Founded in 2021, Eliyan designs and engineers interconnect technology that moves more data into, and between, AI chips. The portfolio spans three areas: NuLink electrical die-to-die (D2D), chip-to-chip (C2C) and chip-to-module (C2M) PHY IP; NuLight 2D µVCSEL optical interconnects; and a dual-port custom HBM (cHBM) architecture. Simultaneous bi-directional (SBD) signaling raises bandwidth per electrical connection, NuLight targets the "Copper Wall" in AI scale-up fabrics, and dual-port cHBM increases bandwidth available to memory and I/O.

Company Mission: Eliyan's mission is to addresses the data movement limits that constrain AI systems.

Market opportunity

Eliyan is targeting the memory and I/O bottlenecks in AI infrastructure, including die-to-die, chip-to-chip, rack-to-rack, and HBM connectivity. The company believes electrical links are reliable and low-cost, but that as AI scale-up fabrics grow, copper faces simultaneous limits in bandwidth density, power and reach, which it refers to as the Copper Wall. It sees a large opportunity in reducing hop count, lowering AI pod power and improving bandwidth density, with opportunities across ASIC IP, chiplets, memory expansion, scale-up EIC chipsets and optical engines.

Product overview:

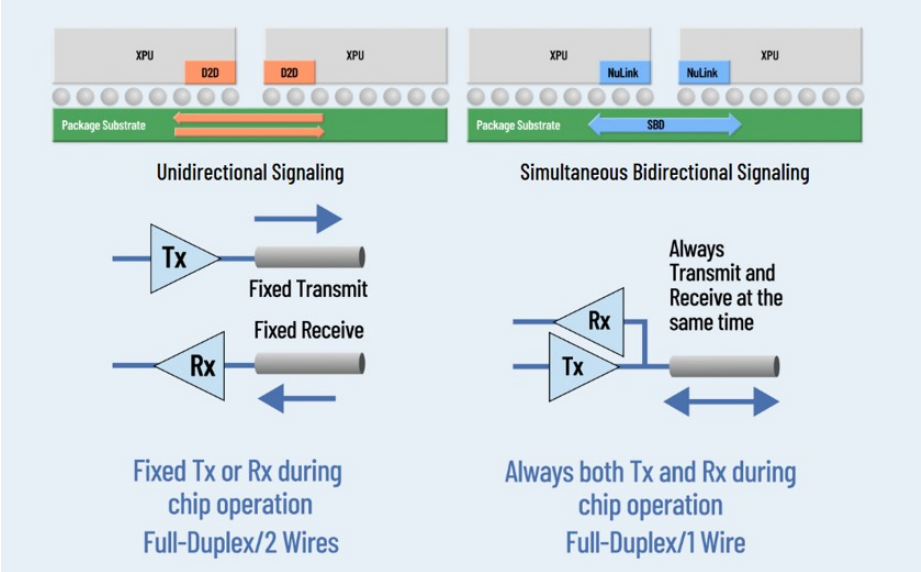
NuLight (optical interconnect): NuLight uses two-dimensional arrays of micro vertical-cavity surface-emitting lasers (2D µVCSELs) to move data optically. Eliyan states that NuLight delivers 8x the bandwidth density, 5x the power efficiency (one-fifth the energy per bit) and 50x the reach of copper links, which would allow more bandwidth between accelerators over larger physical systems without a commensurate increase in power. NuLight's slow-n-wide approach sends data over four 56Gbps NRZ optical lanes instead of one 224Gbps PAM4 lane.

Dual-port cHBM: Eliyan's dual-port custom HBM architecture allows memory and I/O chiplets to use all four edges of the XPU, so memory and scale-up I/O can share the package perimeter rather than competing for scarce die-edge bandwidth.

NuLink (electrical interconnect): Eliyan designs and licenses chiplet interconnect PHY IP that lets multi-die AI chips achieve advanced packaging performance without requiring an expensive silicon interposer. The NuLink family covers D2D (and, via NuLink-X, C2C and C2M) links built around patented SBD signaling, which transmits and receives over the same connection at the same time, which can double the aggregate bandwidth per connection. Because NuLink achieves this on standard organic substrates rather than

costly, size-limited silicon interposers, it extends the reach between chiplets by up to 10x, allowing more HBM or compute chiplets on a single package. NuLink also underpins Eliyan's Universal Memory Interface (UMI), which lets a single ASIC design connect interchangeably to DDR, GDDR or HBM memory chiplets, while NuGear chiplets adapt standard HBM PHY interfaces to NuLink. The technology is protocol-agnostic, supporting both UCIe and BoW. The latest product, NuLink-XD, is a 3nm 224G PAM4 SerDes optimized for chip-to-chip and chip-to-module signaling in markets including NPO and CPO.

Exhibit 16: Eliyan Bidirectional Signaling



Source: Eilyan, Morgan Stanley

Competitive differentiation

In optical Eliyan's NuLight's 2D μVCSEL arrays use less constrained GaAs supply (compared to InP lasers), enable high bandwidth density, with low optical-engine power, and high reliability. In comparison with other optical approaches NuLight offers higher bandwidth density, lower optical-engine power and better reliability than linear-drive optics (LDO) and DWDM silicon photonics (OCI), with the reliability advantage stemming from its patented optical-array repair technology.

Eliyan's differentiates in electrical interconnects by delivering silicon-interposer-class bandwidth on standard, low-cost substrates, removing the need for more expensive advanced packaging process. This is enabled by a patented simultaneous bidirectional signaling architecture that effectively doubles bandwidth per connection at lower power than conventional PHY approaches. A technical approach that took years of dedicated architectural work to develop and is protected by a large patent portfolio.

Because the company's interconnect works on widely available substrates rather than a narrow set of advanced packaging suppliers, it also offers customers a more diversified, lower-cost supply chain, and can improve manufacturing yield and time-to-market versus interposer-based approaches.

Notable Potential Competitors: Synopsys, Cadence, Alphawave Semi (Qualcomm),

Broadcom, Marvell, ams OSRAM

Team

- **Ramin Farjadrad, Co-Founder and Chief Executive Officer**, founded Eliyan in 2021 after recognizing that multi-die packages needed a fundamentally different PHY architecture; he previously served as CTO and VP of Networking/Automotive at Marvell (2019–2021) and as Co-Founder and CTO of Aquantia Corp. (2005–2019).
- **Patrick Soheili, Co-Founder and Chief Strategy and Business Officer**, leads Eliyan's go-to-market and partnership strategy; he has served as Managing Director at Ridgeview Venture since 2008 while also VP of Product Management and Corporate Development at eSilicon (2008–2020).
- **Syrus Ziai, Co-Founder and Chief Operating Officer**, oversees Eliyan's operations and execution; he previously served as VP of Engineering at Nuvia (2019–2021) and SVP of Engineering at PsiQuantum (2016–2019).

Exhibit 17: Eilyan: key leadership



Source: Eilyan

Recent developments

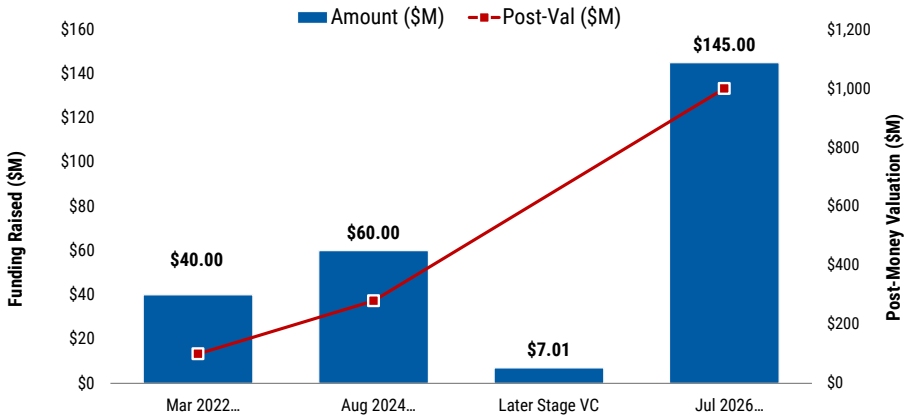
- **In July 2026, Eliyan introduced NuLink-XD, a power-reach configurable 224G PAM4 SerDes** built on TSMC's 3nm process, purpose-built for chip-to-chip and chip-to-module connections feeding into next-generation optical interconnects (including near-package optics and co-packaged optics) — delivering up to 40% lower power than prior designs to support AI scale-up fabrics spanning multiple racks.
- **In November 2025, Eliyan joined the Arm Total Design ecosystem**, supporting the Foundation Chiplet System Architecture (FCSA) that Arm contributed to the Open Compute Project, while also leading a memory-interconnect contribution extending OCP's BoW 2.0 chiplet interconnect specification to support higher-bandwidth memory applications.

Fund-raising

Eliyan has raised approximately $395 million to date and was most recently valued at $1 billion with the close of its $145 million Series C in July 2026 (led by Seligman Ventures). Strategic investors include AMD, Arm, Coherent, Meta, Cisco Investments, Lumentum,

Intel Capital, Samsung Catalyst Fund, Micron Ventures, and SK hynix.

Exhibit 18: Eliyan:



Source: Pitchbook, Company Data, Morgan Stanley

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

SambaNova - Memory Centric Computing for Fast Inference

Who is SambaNova? Founded in 2017, SambaNova designs and engineers XPUs, and XPU systems and software platform for enterprise AI, enabling organizations to develop and deploy AI models cloud and on-prem. The core of the product portfolio is the Reconfigurable Dataflow Unit (RDU), which uses a software-defined, dataflow-driven architecture to optimize performance and efficiency for both AI training and inference.

Company Mission: SambaNova's mission is to improve AI hardware by letting software dictate hardware configuration, delivering AI systems that make the most advanced capabilities accessible without requiring unlimited resources.

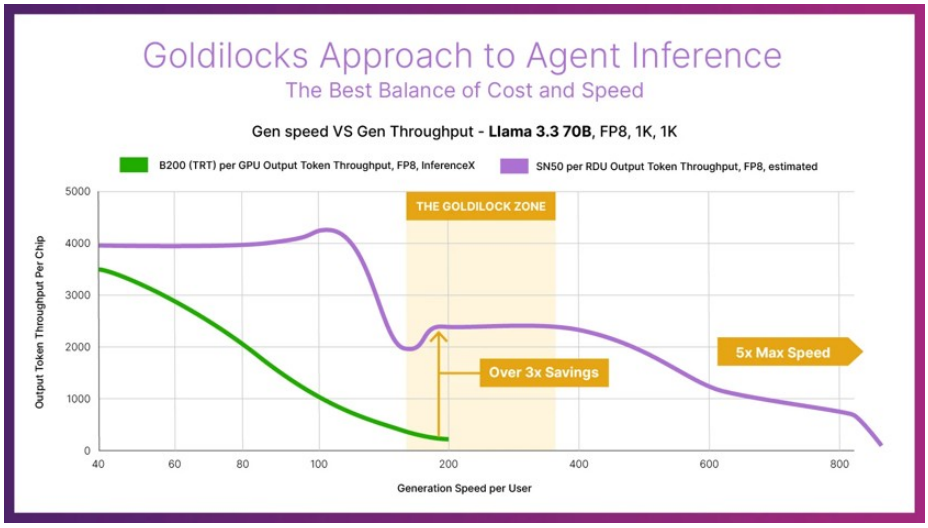
Market opportunity

SambaNova is focused on low-cost, low-latency inference for large models. . As models surpass 1T-parameters and premium token markets are emerging around speed, quality, and always-on agentic workloads. The company sees a need for new hardware solutions to best serve those workloads.

Product overview: SambaNova designs and engineers XPUs built around the Reconfigurable Dataflow Unit (RDU), a processor architecture that's different from the kernel-by-kernel execution model used by GPUs. Rather than moving data between memory and compute units for each operation, the RDU is configured to stream data through a pipeline of compute and memory connected by an on-chip network. Fusing hundreds of operations into a continuous dataflow graph rather than separate kernel calls.

The core of the product portfolio is the dataflow architecture combined with a three-tier memory system (on-chip SRAM, HBM, and directly-attached DDR DRAM), which lets the RDU store large AI models across tiers and switch between them quickly, an advantage in agentic workloads that call on multiple models across a task. The company's software stack (SambaFlow/SambaNova Suite) helps map standard PyTorch and TensorFlow models onto this hardware without requiring hand-written kernels. Commercially, this is packaged into SambaRack (a 16-RDU, air-cooled rack-scale system), SambaCloud (a managed inference platform), and SambaManaged (a turnkey inference deployment for third-party data centers).

Exhibit 19: SambaNova sees significant cost per token advantages vs GPU only solutions



Source: SambaNova, Morgan Stanley

Competitive differentiation

SambaNova's core differentiation is its memory architecture, its RDUs use SRAM, HBM, and DDR memory, reducing dependency on the CPU and PCIe bus. Management contrasted this with GPU architectures and CUDA's kernel-by-kernel temporal mapping, while SambaNova spatially maps workloads.

SambaNova's dataflow architecture fuses many operations into a single streaming pipeline and keeping intermediate results on-chip to minimize memory bottlenecks.

The company also points to power efficiency as a differentiator, positioning its systems as viable in standard air-cooled data centers rather than requiring the liquid cooled infrastructure.

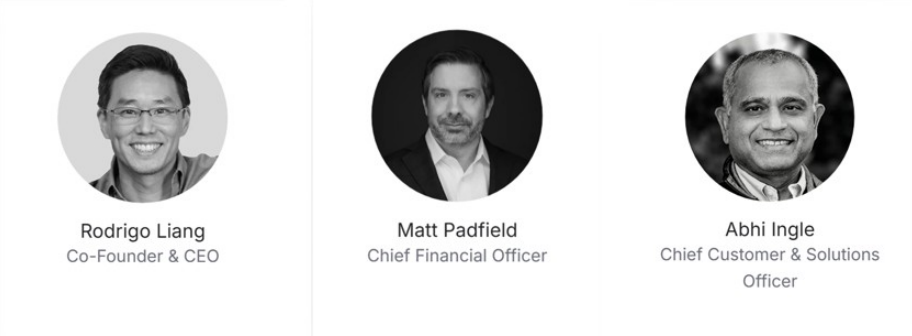
Notable Potential Competitors: NVIDIA, AMD, Nvidia (Groq) , Cerebras, Etched (private), d-matrix (private)

Team

- **Rodrigo Liang, Co-Founder and Chief Executive Officer**, co-founded SambaNova in 2017 after leading a nearly 1,000-person chip design team at Oracle, where he was responsible for SPARC processor and ASIC development; he began his career at Hewlett-Packard and holds BS and MS degrees in Electrical Engineering from Stanford University.
- **Abhi Ingle, Chief Product and Strategy Officer**, leads global product strategy, innovation, and corporate development at SambaNova; he previously built multiple billion-dollar businesses at AT&T and served as Chief Strategy and Corporate Development Officer at Qualtrics through its pre-IPO scale-up, and holds an MBA from Harvard Business School.
- **Matt Padfield, Chief Financial Officer,** joined SambaNova in June 2026 to build the company's financial infrastructure for its next phase of growth; he previously

served as Group Vice President of Corporate FP&A at Coherent and as VP of FP&A and Corporate Development at Synaptics, where he also served as acting CFO and led over $1 billion in M&A activity.

Exhibit 20: Sambanova; Key Leadership



Source: Sambanova

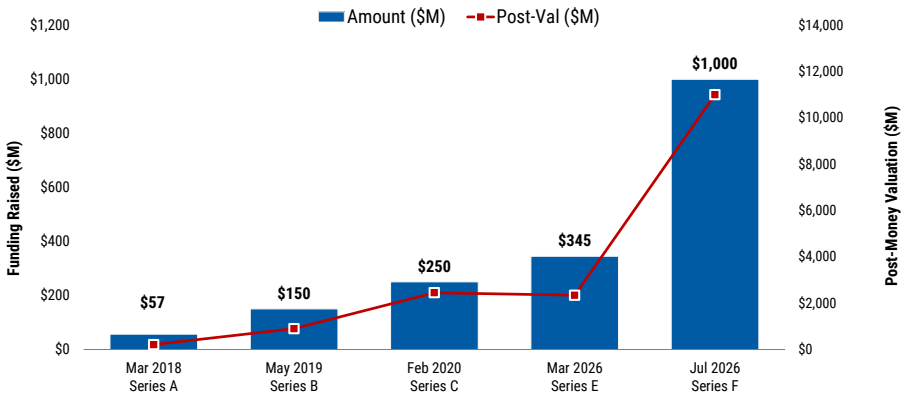
Recent developments

- **In June 2026, Vista Equity Partners and Cambium Capital launched Vector Core Compute (VC2),** the first commercially available enterprise inference cloud for disaggregated inference, backed by a $3.5 billion compute commitment to SambaNova; the platform combines Intel Xeon CPUs, SambaNova SN40 RDUs, and Nvidia Blackwell GPUs, with Together.ai as its first commercial customer and early access extended across Vista's 90+ portfolio companies.
- **In June 2026, at Computex, SambaNova, Intel, and Foxconn demonstrated the first live disaggregated inference system for AI agents** — pairing Nvidia GPUs for prefill with SambaNova's SN40 RDU for decode — running at Vector Core Compute's new data center with Together.ai as the first commercial customer, and announced plans for production rack-scale systems combining Intel Xeon processors with the upcoming SN50 RDU.
- **In April 2026, SambaNova and Intel outlined plans for an integrated hardware system** combining Intel Xeon 6 processors with SambaNova's SN50 RDU for large-scale agentic AI workloads, alongside a distribution agreement with TEPCO Systems to deploy AI infrastructure in Japan; Intel also moved to increase its ownership stake in SambaNova to approximately 9%.
- **In February 2026, SambaNova introduced its fifth-generation SN50 chip** — claimed to offer higher performance than Nvidia's B200-based systems at a lower total cost of ownership — alongside a $350 million Series E round and a multi-year strategic collaboration with Intel, following the collapse of earlier acquisition talks between the two companies.

Fund-raising

SambaNova has raised a cumulative $2.48bn since 2018, the most recent round in July 2026 raised $1,000mn at a post money valuation of $11bn. Strategic investors include Intel, Google Ventures, Vista Equity Partners, and SoftBank

Exhibit 21: SambaNova: Funding History



Source: Pitchbook, Morgan Stanley Research

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

SiMa.ai - Purpose-built for Physical AI

Who is SiMa.ai? Founded in 2018, SiMa.ai is an AI software company that builds it's own chips. The company designs and engineers Machine Learning System-on-Chip (MLSoC) SOCs and companion agenetic AI software that bring computer vision, transformer, and generative AI workloads to physical AI devices across industrial, automotive, robotics, and aerospace applications. The core of the product portfolio is the Modalix MLSoC paired with the Palette Software Suite, which together deliver a software-centric "ONE Platform for Physical AI" supporting any model, or modality at industry-leading performance per watt.

Company Mission: SiMa.ai's mission is to help customers deploy high-performance AI/ML physical AI applications with a high bar for security and privacy with partners, radically simplifying physical AI deployment through one unified platform.

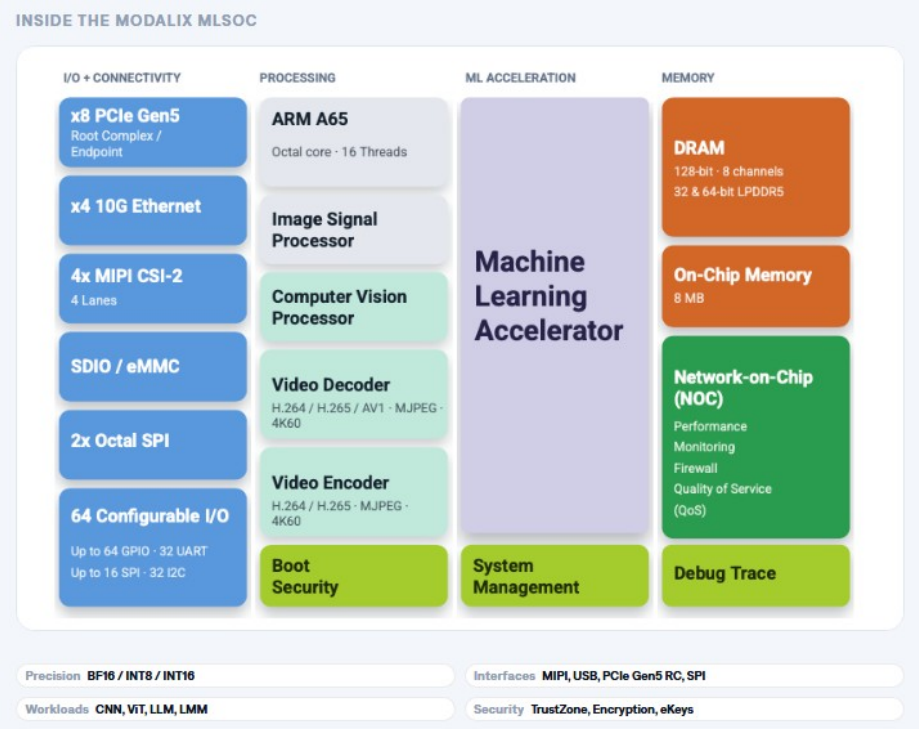
Market opportunity

SiMa.ai is targeting physical AI markets including drones, robotics, automotive, manufacturing, AI-enabled appliances, and healthcare. The company's view is that reasoning models will enable a transition from rules-based computer vision to LLM-based physical AI. Management has argued that the market needs AI-first silicon and software rather than legacy embedded architectures.

Automotive use cases include in-car AI assistants that avoid cloud dependency, run under 10W, support easy LLM upgrades, and deliver low latency. In drones, SiMa.ai positions against incumbents by offering a new class of intelligent drones with real-time Physical AI — integrating vision, sensor fusion, and multimodal AI into a unified architecture for full end-to-end drone applications.

Product overview: SiMa.ai designs and engineers SOCs called Machine Learning Systems-on-Chip (MLSoC) purpose-built for physical AI. SiMa.ai's MLSoC combines a dedicated accelerator, a computer vision processor, and an application compute unit (Arm cores) on a single chip. SiMa.ai expects to release its next generation MLSoC in the first half of 2028.

Exhibit 22: SiMa.ai MLSoC



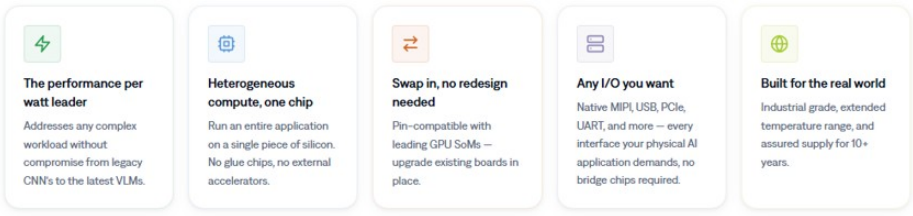
Source: SiMa.ai, Morgan Stanley Research

Competitive differentiation

SiMa.ai's key differentiation is low-power reasoning on-device for physical AI applications. SiMa.ai expects to be able to run a 1,000 dense TOPS engine (2x better than the competition) at 80W and has argued it can match Nvidia Orin SoM pricing while consuming meaningfully less power. The company built pin-to-pin replacement SoMs and believes Nvidia is moving toward higher-end Thor platforms while much of the market needs lower-power solutions to enable physical AI applications.

Software scalability is also a major differentiator versus legacy embedded vendors and FPGA models. SiMa.ai recently launched its agentic software called Palette Neat for physical AI that allows developers to describe sensors, cameras, radar, and systems in English, drastically reducing development time.

Exhibit 23: SiMa's key differentiation



Source: SiMa.ai, Morgan Stanley

Notable Potential Competitors: Nvidia, Qualcomm, Hailo (private), Ambarella

Team

- **Krishna Rangasayee, Founder and CEO**, founded SiMa.ai in 2018 after serving as COO of Groq and spending 18 years at Xilinx, where he rose to SVP and GM of the overall business and grew it to $2.5 billion in revenue at 70% gross margin; he holds an MSEE from Mississippi State University and completed the executive management program at Stanford's Graduate School of Business.
- **Gopal Hegde, Senior Vice President of Engineering and Operations**, joined SiMa.ai in 2020 to lead engineering and operations after more than 30 years in the industry, including as VP/GM of the Server Processor Business Unit at Marvell/ Cavium, COO at Calxeda, and senior engineering leadership roles at Cisco and Intel; he holds an MSECE from the University of Massachusetts Amherst.
- **Harry Kroeger, President of Automotive Business**, joined SiMa.ai in 2023 to lead its expansion into automotive after three decades in the industry, most recently as a member of the board of management at Bosch overseeing a $25 billion automotive electronics business, and 21 years at Mercedes-Benz; he has also served on the boards of Tesla and Rivian.

Exhibit 24: SiMa.ai: Key Leadership



Source: SiMa.ai, Morgan Stanley

Recent developments

- **In October 2026, SiMa.ai partnered with Wobot.ai to make Wobot's reasoning-based video analytics, available on the Modalix MLSoC as a PCIe card or system-on-module.** The solution processes video on-site at the edge rather than in the cloud, which SiMa.ai says helps address data-privacy requirements in regions such as Europe, Japan and Korea. Wobot.ai's software is deployed across QSR, retail, hospitality and manufacturing sectors.
- **In August 2026, SiMa.ai and ARK Electronics launched the ARK-SiMa.ai Modalix MLSoC Bundles.** The partnership creates production ready, drone-oriented hardware combining SiMa.ai's Modalix System-on-Module with ARK's made-in-the-USA flight controllers and sensors. The systems are designed to be form-factor compatible with the industry's leading GPU-based edge platform, giving drone OEMs a lower power, lower cost drop-in alternative.
- **In June 2026, SiMa.ai released Palette Neat, the first agentic development environment for Physical AI** applications as part of its Palette Software Suite. With Palette Neat, customers can build new or port existing applications seamlessly within days/hours instead of months.
- **In April 2026, SiMa.ai partnered with ARC Aerospace to power its Ruby FC and**

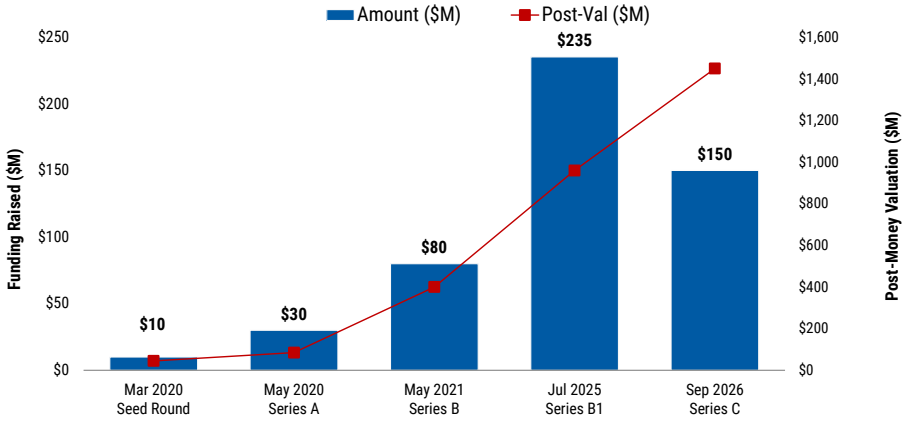
Ruby CEP flight computer series with the Modalix MLSoC family, supporting domestic U.S. production of next-generation flight computers for autonomous systems and counter-drone missiles, including navigation in GPS-denied environments and real-time target acquisition and tracking.

- **In March 2025, SiMa.ai accelerated the release of its second-generation MLSoC Modalix family**, built on TSMC's 6nm process, unifying computer vision, transformers, and generative AI into a single software-centric platform for the first time.

Fund-raising

SiMa.ai has raised about $500mn in total since its founding in 2018, most recently a $150mn Series C announced on September 28th 2026 at a $1.45bn valuation (up from $960mn after the prior round in mid-2025, per PitchBook). The round was co-led by Fidelity Management & Research Company and Amplify, with participation from Dell Technologies Capital, Maverick Capital, Point72, StepStone Group and others. Strategic investors include Micron Technology, which made an undisclosed investment in April 2026, and Dell Technologies Capital.

Exhibit 25: SiMa.ai: Equity funding history



Source: Pitchbook, Morgan Stanley

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

Tenstorrent - RISC-V Processors for Open AI Compute

Who is Tenstorrent? Founded in 2016, Tenstorrent designs and engineers AI processors, high-performance RISC-V CPUs, configurable chiplets, and open-source software for AI and HPC workloads, spanning everything from desktop development boxes to scalable data center compute. The core of the product portfolio is the Tensix core architecture, which combines an array math unit with RISC-V "baby" processors and a network-on-chip to deliver general-purpose, highly networked AI compute rather than narrowly specialized silicon.

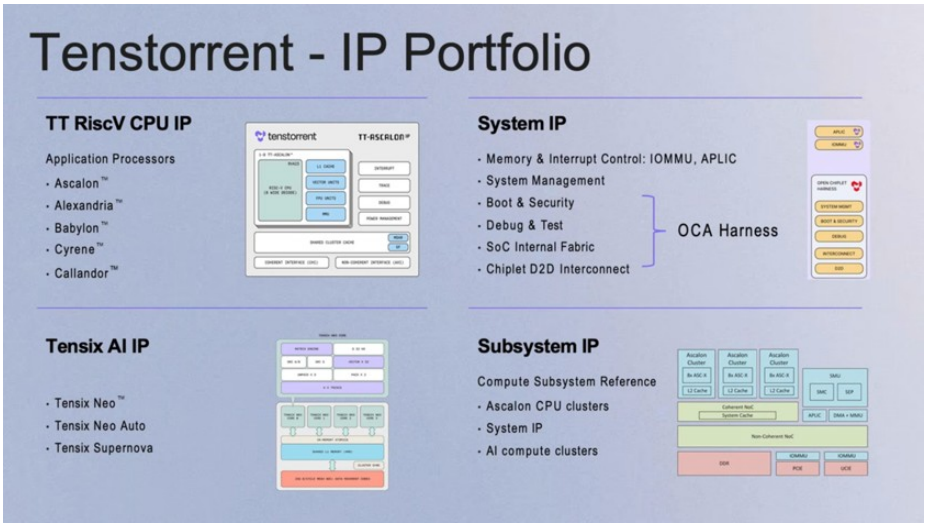
Company Mission: Tenstorrent's mission is general-purpose AI computing, the company positions itself as "anti-specialized," betting that as AI models evolve, hardware built for flexibility across chip types and workloads will outlast narrowly specialized accelerators.

Market opportunity

Tenstorrent is focused on building open AI compute, RISC-V CPU IP, chiplets, and AI systems for edge, automotive, and datacenter. The company sees a large opportunity in providing alternatives to proprietary AI infrastructure. Management has argued that customers increasingly want architectural flexibility and lower-cost licensing options.

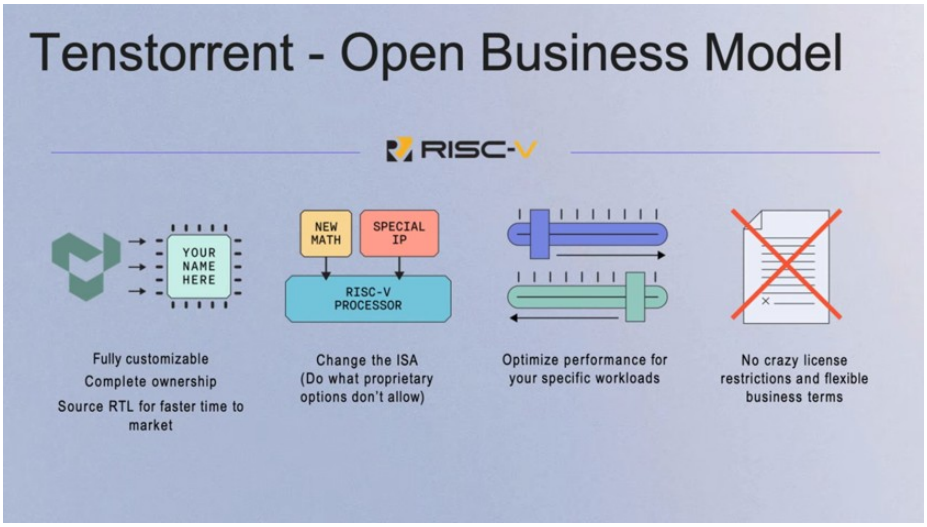
Product overview: Tenstorrent designs and engineers AI accelerators, high-performance RISC-V CPUs, and configurable chiplets. The core of the product portfolio is the Tensix core — a compute tile combining a matrix engine, a vector engine, local SRAM, and five small embedded RISC-V cores. Hundreds of these tiles are connected via an on-chip mesh network, letting any core access any other core's local memory directly rather than routing everything through a shared cache or DRAM — a design intended to reduce the memory-latency bottleneck. The current-generation Blackhole chip pairs 120 Tensix cores with 16 larger, Linux-capable RISC-V cores that manage the chip independently of a host CPU, and ships with an on-chip Ethernet mesh for scale-out rather than a proprietary interconnect. The accompanying software stack, TT-Metalium, is fully open-source under an MIT license, in contrast to Nvidia's proprietary CUDA. Beyond finished hardware, Tenstorrent also licenses its Tensix and Ascalon (RISC-V CPU) IP directly to customers building their own custom silicon, giving the company multiple ways to monetize the same underlying architecture across finished accelerators, chiplets, and licensable IP.

Exhibit 26: Tenstorrent IP Portfolio



Source: Tenstorrent's 10-year journey with RISC-V | embedded world 2026 [https://www.youtube.com/watch?v=A3Z3JIZcMvo]

Exhibit 27: Tenstorrent Open Business Model



Source: Tenstorrent's 10-year journey with RISC-V | embedded world 2026 [https://www.youtube.com/watch?v=A3Z3JIZcMvo]

Competitive differentiation

Tenstorrent's differentiation is its open platform: RISC-V CPU cores, AI compute, chiplets, servers, and open-source software. Management argued that it licenses the whole stack, enabling customers to build what they want. This is positioned against Arm's more restrictive architecture licensing model.

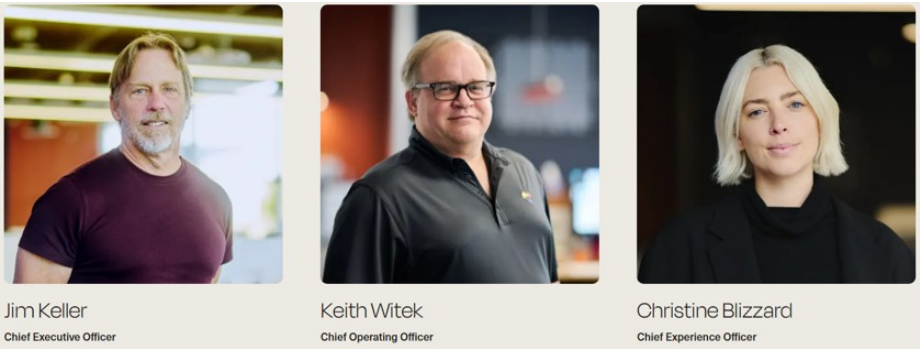
The company has emphasized that chiplets allow customers to move faster and that Tenstorrent can provide CPU, AI, and I/O building blocks. Larger incumbents such as AMD, Intel, and Xilinx have similar technology, but Tenstorrent is making it standard and broadly available.

Notable Potential Competitors: Nvidia, AMD, SiFive (private), Arm, Andes

Leadership Team

- **Jim Keller, Chief Executive Officer**, joined Tenstorrent as CTO and President in 2020 before becoming CEO in 2023; before Tenstorrent, he was Senior Vice President of Silicon Engineering at Intel (2018–2020) and previously led chip architecture at AMD (Zen), Apple (A4/A5), and Tesla's Autopilot hardware team, making him one of the most recognized chip architects in the industry.
- **Keith Witek, Chief Operating Officer**, joined Tenstorrent in January 2023 to scale global operations, finance, and IP licensing; he previously led Strategic Alliances for Google's consumer electronics group, served as SVP of Corporate Development and General Counsel at SiFive, held technology-enablement and legal roles at Tesla, and spent 13 years at AMD, culminating as Corporate VP of Strategy and Corporate Development.
- **Christine Blizzard, Chief Experience Officer**, leads Tenstorrent's brand, marketing, and communications; she previously founded y3k, a brand marketing studio, and worked as a Senior Creative Producer for Google's Events and Experiences Team.

Exhibit 28: Tenstorrent: Key Leadership



Source: Tenstorrent

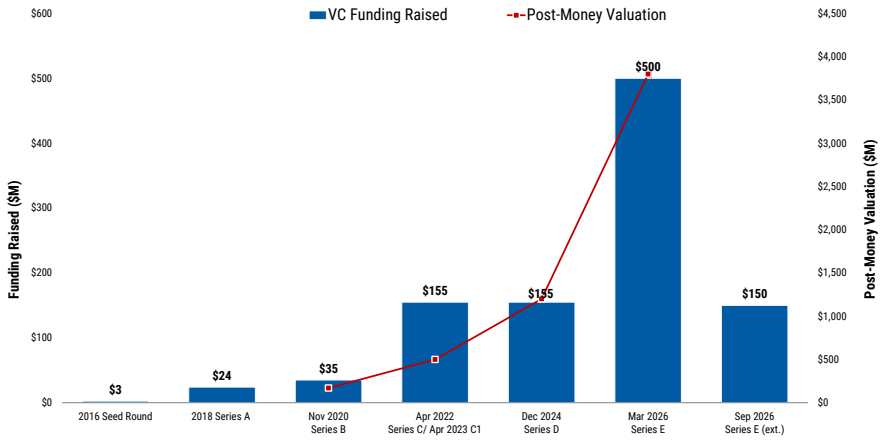
Recent company developments

- **In September 2026, Tenstorrent and ai& launched JapanFold**, a sovereign drug-discovery platform running on Tenstorrent Galaxy superclusters, giving Japanese researchers and pharmaceutical companies access to leading open-source structural biology models (including Boltz-2, OpenFold3, and RFdiffusion 3) with all computation kept entirely within Japan and near-linear performance scaling on Tenstorrent hardware.
- **In July 2026, Tenstorrent partnered with Stealthium** to deliver runtime observability for AI infrastructure running on its hardware.In July 2025, Tenstorrent acquired Blue Cheetah Analog Design, after previously licensing its die-to-die interconnect IP, to strengthen its in-house chiplet interconnect capabilities.
- **In June 2026, Tenstorrent launched TT-Ascalon S, a new RISC-V CPU core**, and expanded its presence across Japan.
- **In May 2026, Tenstorrent launched TT-QuietBox 2**, a liquid-cooled desktop AI workstation built on four Blackhole chips with a fully open-source software stack, able to run models up to 120 billion parameters locally starting at $9,999.
- I**n April 2026, Tenstorrent's Blackhole chip reached general availability**, built on the Tensix core architecture with on-chip Ethernet-based scale-out in place of a proprietary interconnect.

Fund-raising

Tenstorrent has raised approximately $2.4 billion in cumulative funding since 2016. Strategic and institutional investors across its financings include Samsung Securities, Samsung Catalyst Fund, Hyundai Motor Group, LG Technology Ventures, Bezos Expeditions, Fidelity, Baillie Gifford, and others.

Exhibit 29: Tenstorrent: Equity Funding history



Source: Pitchbook, Company Data, Morgan Stanley Research

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

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 September 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, 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, Qualcomm Inc., 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 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, D-MATRIX CORPORATION.

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., AXIADO CORPORATION, AYAR LABS, INC., D-MATRIX CORPORATION, ELIYAN CORPORATION, SAMBANOVA SYSTEMS, INC., TENSTORRENT 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., 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., AXIADO CORPORATION, AYAR LABS, INC., D-MATRIX CORPORATION, ELIYAN CORPORATION, SAMBANOVA SYSTEMS, INC., TENSTORRENT 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., 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* (10/09/2026)
Joseph Moore		
Advanced Micro Devices (AMD.O)	E (06/09/2024)	$608.10
Aeva Technologies Inc (AEVA.O)	E (07/19/2021)	$13.66
Allegro Microsystems Inc (ALGM.O)	O (02/13/2026)	$35.57
Ambarella Inc (AMBA.O)	O (03/29/2016)	$67.43
Amkor Technology Inc (AMKR.O)	E (11/08/2023)	$49.12
Analog Devices Inc. (ADI.O)	O (11/16/2023)	$406.19
Astera Labs Inc (ALAB.O)	O (05/11/2025)	$342.21
Broadcom Inc. (AVGO.O)	O (06/09/2024)	$361.54
Cerebras Systems (CBRS.O)	O (06/08/2026)	$164.17
GlobalFoundries Inc (GFS.O)	E (10/28/2024)	$48.24
Intel Corporation (INTC.O)	E (02/22/2023)	$104.70

IonQ Inc (IONQ.N)	E (04/25/2023)	$39.89
Marvell Technology Group Ltd (MRVL.O)	E (09/14/2015)	$275.28
Microchip Technology Inc. (MCHP.O)	E (07/10/2024)	$75.55
Micron Technology Inc. (MU.O)	O (10/06/2025)	$1,029.00
Navitas Semiconductor Corp (NVTS.O)	U (04/06/2025)	$10.78
NVIDIA Corp. (NVDA.O)	O (03/16/2023)	$229.28
NXP Semiconductor NV (NXPI.O)	O (02/11/2025)	$227.52
ON Semiconductor Corp. (ON.O)	++	$77.53
Qualcomm Inc. (QCOM.O)	E (06/24/2026)	$175.50
Quantinuum (QNT.O)	E (06/29/2026)	$41.30
SanDisk Corporation. (SNDK.O)	O (03/03/2025)	$1,581.82
Semtech Corp. (SMTC.O)	E (04/06/2025)	$187.30
Silicon Laboratories Inc. (SLAB.O)	E (01/19/2021)	$222.04
SiTime Corporation (SITM.O)	O (09/16/2026)	$612.37
Skyworks Solutions Inc (SWKS.O)	E (11/28/2018)	$76.35
Texas Instruments (TXN.O)	U (04/13/2020)	$283.74
Wolfspeed, INC (WOLF.N)	NR (04/06/2025)	$30.61
Lee Simpson		
Arm Holdings plc (ARM.O)	E (04/07/2026)	$266.28
Cadence Design Systems Inc (CDNS.O)	O (02/14/2024)	$352.84
PDF Solutions Inc (PDFS.O)	O (09/08/2026)	$54.40
Synopsys Inc. (SNPS.O)	O (09/08/2026)	$509.99

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

© 2026 Morgan Stanley