US Semiconductors

Profiling Private Semiconductor Companies - 5th Edition

Equities

Americas

Semiconductors

Timothy Arcuri
Analyst
timothy.arcuri@ubs.com
+1-415-352 5676

Natalia Winkler, CFA
Analyst
natalia.winkler@ubs.com
+1-415-352 4626

Seth Cho
Associate Analyst
seth.cho@ubs.com

Gianmarco Vella
Associate Analyst
gianmarco.vella@ubs.com
+1-415-352 4555

Aaryan Wadhwa
Associate Analyst
aaryan.wadhwa@ubs.com
+1-212-821 6481

Stacy Che
Associate Analyst
stacy.che@ubs.com
+1-212-713 1857

Semiconductor innovation focused on hardware constraints on computation, we profile 105 companies innovating to overcome these challenges.

We expand our private semiconductor company compendium to include 105 companies, adding 79 new profiles across AI inference, networking, edge computing, and advanced manufacturing versus prior editions. Many of these companies are attempting to address the hardware constraints revealed by AI computation (memory wall, interconnects, etc.) and may offer compelling advantages in terms of both performance and cost. Recent private-company activity highlights areas of the AI ecosystem where innovation remains particularly intense. See below for our detailed takes by high-growth segments:

- **Inference:** emerging accelerator architectures optimized for lower cost, reduced power consumption, and edge deployment broaden the opportunity set beyond a GPU-centric paradigm. Examples include Groq's low-latency LPU architecture, alongside companies such as Etched, SambaNova, and EnCharge AI that are developing purpose-built inference accelerators.
- **AI networking:** larger clusters and accelerating data movement requirements should support demand for switching, optical connectivity, and networking infrastructure, benefiting companies such as AMD, AVGO, MRVL, NVDA, and SMTC.
- **Physical AI:** increasing adoption of AI-enabled smartphones, PCs, robotics, and edge devices has the potential to drive higher semiconductor content per device, supporting QCOM as well as RF vendors like SWKS.
- **Manufacturing:** the migration toward chiplet architectures and more advanced packaging techniques continues to increase the strategic importance of packaging, testing, and metrology technologies.

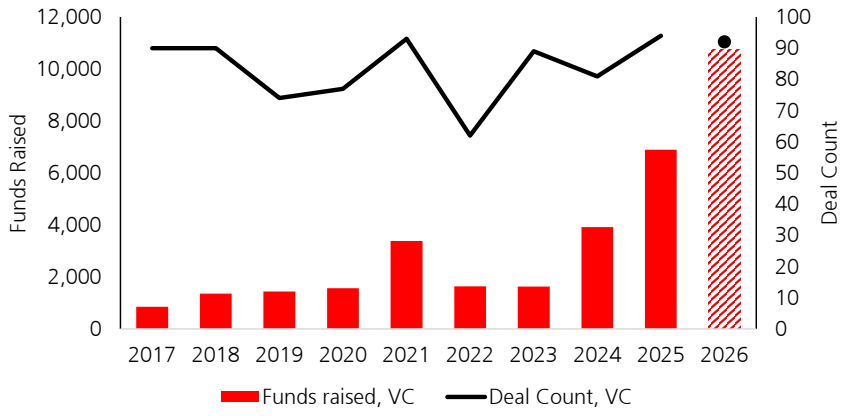
For prior compendiums, please see the following here: 1st edition, 2nd edition, 3rd edition, 4th edition

AI infrastructure remains in the early innings as venture investment becomes increasingly concentrated

Semiconductor VC investments have evolved from the broad-based funding surge of 2021 to a more concentrated, scale-up driven market today. While deal activity moderated following the 2022-23 correction, capital invested remains structurally above pre-AI-cycle levels, with funding increasingly concentrated in a smaller number of later-stage companies. Against that backdrop, 2026 has been characterized by a series of megarounds, including Cerebras ($1.1bn), Groq ($1.0bn), SambaNova ($1.0bn), Etched ($700m), Ayar Labs ($500m), MatX ($500m), SiFive ($400m), TensorWave ($350m), and xLight ($350m). In our view, the size and frequency of these financings underscore continued investor conviction in a multi-year AI infrastructure buildout and reflect a shift from early-stage experimentation toward commercialization, deployment, and scaled manufacturing. Further detail in back pages.

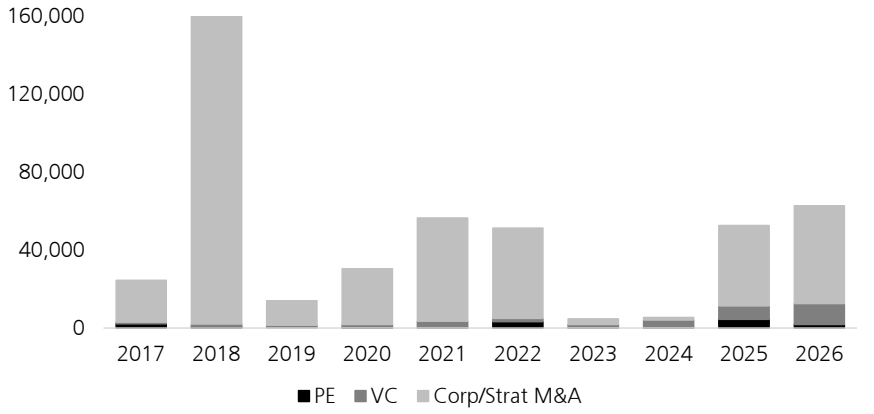
This report has been prepared by UBS Securities LLC. **ANALYST CERTIFICATION AND REQUIRED DISCLOSURES, INCLUDING INFORMATION ON THE QUANTITATIVE RESEARCH REVIEW PUBLISHED BY UBS, BEGIN ON PAGE 226.**

Figure 1: Overall Venture Capital Deal Flow



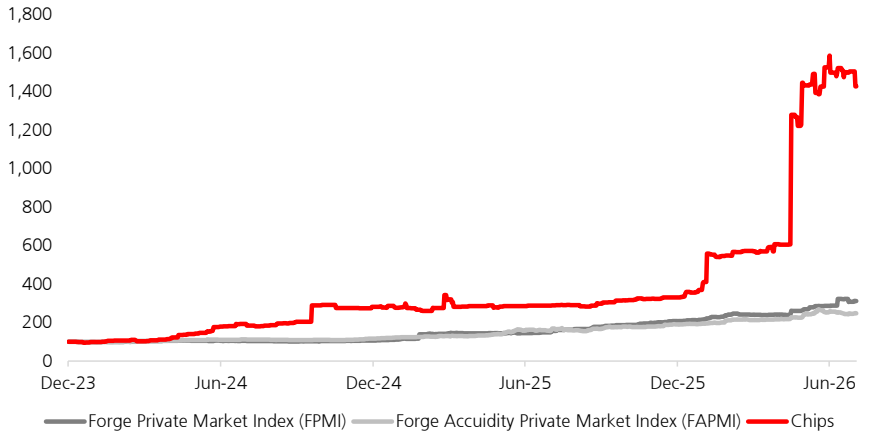
Source: Pitchbook, UBS

Figure 2: Investment Over Time



Source: Pitchbook, UBS

Figure 3: Secondary pricing "Chips" basket performance vs. broader secondary indices



Source: Forge

Figure 4: Public Semiconductor Peers

Companies	Public Peers
Achronix Data Acceleration	AMD, INTC, LSCC, CBRS
AheadComputing	Andes Technology, CEVA
Akash Systems	COHU
Ambature	AMSC, OXIG, BRKR
Ambra	NICE, GOOGL, AMZN, MSFT, CRM, FIVN
Analogix	2379-TW, HIMX, MegaChips, SYNA, SIMO, 6723-JP
AONDevices	AMBQ, KN, QCOM, SLAB
Atomic Machines	DM, PRLB
AvicenaTech	AVGO, COHR, INTC, LITE, MRVL
Axbio	ILMN, TMO, Roche Holding AG, BIO, QDEL
Axiado	SNPS, ARM, RMBS, IFX, NXPI
Ayar Labs	AVGO, COHR, INTC, LITE, MRVL
Baya Systems	CEVA
Bright Machines	JBL, FLEX, CLS, 2317 TW, ABB, FANUC, ROK, TER
Celera Semiconductor	CDNS, SNPS
ChipAgents	CDNS, SIE.DE, SNPS
CNEX Labs	MRVL, AVGO, MU, STX, WDC
Cognichip	CDNS, SNPS, ARM
Cornelis Networks	ANET, AVGO, CRDO, MRVL, NVDA
Crossbar	INTC, MU, WDC, Samsung, AMBQ
Density AI	AMD, CBRS, INTC
Dexterity	ABB, FANUC, ROC, TER, SYM, ZBRA
d-Matrix	AMD, CBRS, INTC, MU, NVDA
Driven Quantum Technologies	QBTS, IONQ, RGTI, INFQ, PSQL
EdgeQ	INTC, MRVL, NXPI, QCOM
Efficient Computer	ARM, NVDA, QCOM, INTC, ADI, CBRS
Eliyan	CDNS, RMBS, SNPS
Emerald AI	Schneider Electric, Siemens, ETN, GEV, ENEL, HON
EnCharge AI	AMBQ, MCHP, QCOM, TXN
Entanglement Consulting	IBM, QBTS, IONQ, RGTI, MSFT, GOOGL
Eridu	NVDA, AVGO, CSCO, ANET
Etched	AMD, INTC, NVDA
Expedera	CBRS, INDI
FemtoMetrix	KLAC, ONTO, NVMI, BRKR, 6146 T
Ferric	MPWR, TXN, ADI, Renesas, VICR
Figure AI	KSCP
Forge Nano	VECO, AMAT, LRCX, KLAC, ASM International
Fractile	CBRS, NVDA, AMD
Frore Systems	VRT
Furiosa	NVDA, INTC, CBRS
Gyrfalcon Technology	CBRS
Halo Industries	IPGP, LASR
HyperLight	COHR, LITE, AAOI, CSCO
Ideal Semiconductor	IFNNY, TXN, ON, STM, NXPI, AMAT
Kneron	AMD, QCOM
Lightmatter	AVGO, COHR, INTC, MRVL, NVDA
Lumotive	COHR, LITE, MVIS, OUST
Majestic	CBRS, NVDA
MangoBoost	NVDA, AMD, CBRS
MatterStack	COHR, LITE, MVIS, OUST
MatX	AMD, CBRS, INTC, NVDA
MemryX	AMD, INTC, NVDA, QCOM
Menlo Micro	ADI, IFX, QRVO, SWKS
Metalenz	WOLF, STM
Modal	AMZN, GOOGL, MSFT, ORCL
Mojo Vision	VUZI
Movandi	ADI, MCHP, QCOM, SWKS
Multibeam	AMAT, ASML, LRCX, KLAC, VECO
MultiLucent	GOOGL, TEM, ISMF, BFLY
Mythic	AMBQ, MCHP, QCOM, TXN
Neurophos	1879.HK
Nexthop AI	NVDA, CBRS
nEye.ai	CIEN, LITE
Normal Computing	CBRS
Omnitron Sensors	ADI, STM, TXN, IFX, NXPI
OpenLight	AVGO, COHR, LITE, MRVL
Path Robotics	LECO, ITW, ABB, 6954 T, 6506 T
Prahelika	AMD, INTC, MU, QCOM
Point2 Technology	AVGO, MRVL, CSCO, INTC, NVDA
Positron	AMD, INTC, NVDA
Power Lattice	SLAB, ADI, TXN, IFX
PSIQuantum	IBM, QBTS, RGTI
QC Ware	SDGR
Quadric	CBRS
QuEra Computing	INFQ, IONQ, RGTI, QBTS, IBM
Rain AI	CBRS
Reed Semiconductor	ADI, IFX, ON, STM, TXN, MPWR, Renesas
Recursive Superintelligence	GOOGL
Ricursive Intelligence	CDNS, SNPS
Relativity Networks	MSFT
Retym	CIEN, COHR, LITE
Sagence AI	NVDA, AMD, INTC
Salience Labs	AVGO, MRVL, CSCO, COHR, LITE
SambaNova Systems	AMD, CBRS, INTC, NVDA, MU
SiFive	ARM, CDNS, SNPS
SiFotonics	AVGO, COHR, CSCO, LITE, MRVL
SiLC Technologies	AEVA, INVZ, OUST
SiMa	AMBA, AMBQ, AMD, INTC, NVDA
Snowcap	QBTS, RGTI, IONQ, INFQ
Source Foundry	TSM, GFS, INTC, Samsung, 0981 HK
Substrate	GFS, INTC, TSEM, MU
SunRise Memory	Weebit Nano
Tachyum	AMD, CBRS, INTC, NVDA
Tacta Systems	ABB, 6954 T, 6506 T, TER, ROK
Tela Innovations	SNPS, CDNS, ARM
Tensordyne	CBRS
Tensorwave	NVDA, CBRS, CRWV
Tenstorrent	NVDA, CBRS
TeRAM	MU, WDC, STM
Unconventional AI	CBRS
Velaura	SU, CSCO, IBM, AVGO, NVDA
Volantis (Semiconductors)	ASML, COHR
xlight	ASML, VECO, NINOY
Xscape Photonics	AVGO, COHR, LITE, MRVL, CSCO
Zero ASIC	SNPS, CDNS, MRVL, AMD

Source: UBS

TABLE OF CONTENT

Inference

AheadComputing, Density AI, d-Matrix, Efficient Computer, EnCharge AI, Etched, Fractile, Furiosa, MatX, MultiLucent, Mythic, Neurophos, Normal Computing, Positron, Rain AI, Sagence AI, SambaNova Systems, SiFive, Snowcap, Tachyum, Tensordyne, Tenstorrent, Unconventional AI, Velaura

AI Networking

AvicenaTech, Ayar Labs, Baya Systems, Cornelis Networks, Eliyan Corporation, HyperLight, Lightmatter, MangoBoost, Menlo Microsystems, Nexthop AI, nEye.ai, OpenLight, Point2 Technology, Relativity Networks, Retym, Salience Labs, SiFotonics, Xscape Photonics

Edge AI

Achronix Data Acceleration, AONDevices, Axiado, EdgeQ, Expedera, Gyrfalcon Technology, Kneron, MemryX, Quadric, SiMa Technologies

Physical AI/Robotics

Atomic Machines, Bright Machines, Dexterity, Figure AI, Lumotive, Metalenz, Mojo Vision, Omnitron Sensors, Path Robotics, Tacta Systems

Semiconductor Manufacturing

ChipAgents, Cognichip, Crossbar, FemtoMetrix, Forge Nano, Halo Industries, Ideal Semiconductor, MatterStack, Multibeam, Prahelika, Recursive Superintelligence, Ricursive Intelligence, Source Foundry, Substrate, SunRise Memory, Tela Innovations, TeRAM, Volantis (Semiconductors), xlight, Zero ASIC

Quantum Computing

Driven Quantum Technologies, Entanglement Consulting, PSIQuantum, QC Ware, QuEra Computing

AI DC Power Management

Akash Systems, Ambature, Emerald AI, Ferric, Power Lattice, Reed Semiconductor

Outside Taxonomy

Ambra, Analogix, Axbio, Celera Semiconductor, CNEX Labs, Eridu, Frore Systems, Majestic, Modal, Movandi, SiLC Technologies, Tensorwave

Figure 5: Inference Companies & Public Peers

Companies	Public Peers
AheadComputing	Andes Technology, CEVA
Density AI	AMD, CBRS, INTC
d-Matrix	AMD, CBRS, INTC, MU, NVDA
Efficient Computer	ARM, NVDA, QCOM, INTC, ADI, CBRS
EnCharge AI	AMBQ, MCHP, QCOM, TXN
Etched	AMD, INTC, NVDA
Fractile	CBRS, NVDA, AMD
Furiosa	NVDA, INTC, CBRS
MatX	AMD, CBRS, INTC, NVDA
MultiLucent	GOOGL, TEM, ISMF, BFLY
Mythic	AMBQ, MCHP, QCOM, TXN
Neurophos	1879.HK
Normal Computing	CBRS
Positron	AMD, INTC, NVDA
Rain AI	CBRS
Sagence AI	NVDA, AMD, INTC
SambaNova Systems	AMD, CBRS, INTC, NVDA, MU
SiFive	ARM, CDNS, SNPS
Snowcap	QBTS, RGTI, IONQ, INFQ
Tachyum	AMD, CBRS, INTC, NVDA
Tensordyne	CBRS
Tenstorrent	NVDA, CBRS
Unconventional AI	CBRS
Velaura	SU, CSCO, IBM, AVGO, NVDA

Source: UBS

Inference: AheadComputing

COMPANY NAME	AheadComputing
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / RISC-V CPU Core IP
SIZE OF TOTAL ADDRESSABLE MARKET	AheadComputing develops high-performance RISC-V processor-core intellectual property for next-generation embedded systems. The company's technology is intended to improve performance, power efficiency, scalability, and time-to-market across a range of embedded applications.The company remains in the product-development stage.
GEOGRAPHIC PRESENCE	Beaverton, Oregon
NOTABLE COMPETITORS	**LISTED:** Andes Technology, CEVA **PRIVATE:** SiFive, Codasip
KEY SUPPLIERS / CUSTOMERS	AheadComputing does not disclose suppliers/customers
COMPETITIVE POSITIONING	AheadComputing is developing licensable RISC-V CPU core intellectual property for embedded and general-purpose computing applications. Key elements of its positioning include: • Performance-oriented RISC-V architecture: The company is focused on high-performance CPU cores rather than presenting itself solely as a low-power embedded RISC-V provider. • Efficiency and scalability: AheadComputing's architecture is intended to help customers improve system performance and reduce power consumption across a range of applications. • Time-to-market benefits: The company plans to provide processor-core IP that customers can incorporate into next-generation devices, potentially reducing internal development requirements and accelerating product development. • Microarchitecture development: Proceeds from the seed financing were designated partly for advancing the company's microarchitecture and demonstrating CPU-core performance.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://aheadcomputing.com
MANAGEMENT TEAM / FOUNDER(S)	**Deborah Marr, Ph.D. – Co-Founder, CEO, President & Board Member:** Deborah Marr leads AheadComputing and has served on the board since January 2024. Former Intel Fellow and Chief Architect of Intel's Advanced Architecture Development Group (AADG) with 33 years of experience at Intel. **Srikanth Srinivasan, Ph.D. – Co-Founder & Architect:** Former Intel Principal Engineer, CPU Architect, and microarchitecture inventor with over 25 years of industry. **Mark Dechene – Co-Founder & Microprocessor Architect:** Former Intel Principal Engineer and CPU Architect who led Memory Execution Architecture and core pathfinding.
SIZE OF COMPANY	~118 employees on PitchBook (accessed 6/17/2026)
GROWTH TRAJECTORY	AheadComputing remains in the product-development stage and is using its financing to expand its team, advance microarchitecture development, and demonstrate CPU-core performance. The company had 40 employees at the time of its January 21, 2026 seed financing and 118 employees as of June 17, 2026, representing an increase of 78 employees between the two reported observations. The company raised a combined $100.92 million through its seed and Series A financings. Its June 2026 Series A followed the January 2026 seed financing, while the post-money valuation increased from $238.09 million in the seed round to $378.00 million in the Series A.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	AheadComputing is venture capital-backed and completed a $47.83 million Series A financing on June 16, 2026. The Series A investors were undisclosed, and the financing brought total capital raised to $100.92 million. The company previously completed a $53.09 million seed financing on January 21, 2026. The round was led by Toyota Ventures, Eclipse Capital, and Cambium Capital Management. Other participating investors included Fundomo, Trousdale Ventures, Stata Capital Partners, Mesh Ventures, Corner Capital Management, Big Brain Holdings, Jim Keller, Maverick Capital, EPIQ Capital Group, and Conconi Growth Partners.
LATEST ROUND VALUATION	AheadComputing's June 16, 2026 Series A was completed at a $330.17 million pre-money valuation and a $378.00 million post-money valuation per Pitchbook.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For private, pre-revenue or product-development-stage semiconductor IP companies, the most directly observable valuation approach in the supplied source is the post-money valuation from the latest preferred-equity financing. When using a Growth Valuation Model, including drivers such as public and private market returns, average private-company growth, comparable performance, private and public comparable-company growth, company-specific signals, employee growth, previous valuation growth, company age, time since last financing, and baseline growth.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Density AI

COMPANY NAME	Density AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	DensityAI is developing AI accelerators designed specifically for frontier-scale large language model inference. The company's architecture is built around a novel memory-centric design that seeks to address memory bandwidth and utilization constraints that increasingly limit inference performance for large models. By integrating memory directly within the compute stack, DensityAI aims to combine the bandwidth advantages associated with near-memory compute architectures with the capacity benefits of high-bandwidth memory -based systems. The company is developing a full-stack platform spanning silicon, advanced packaging, systems software, compiler technologies, and inference-serving infrastructure. DensityAI operates within the broader AI accelerator and AI infrastructure market, benefiting from increasing demand for AI compute across cloud, enterprise, automotive, robotics, and industrial applications as TAM is expected to exceed $600B by 2033. The company's focus is specifically on inference workloads, where memory capacity, bandwidth, latency, and energy efficiency have become key determinants of system performance.
GEOGRAPHIC PRESENCE	Mountain View, California
NOTABLE COMPETITORS	**LISTED:** AMD, CBRS, INTC, NVDA **PRIVATE:** Groq, SambaNova, Tenstorrent, Etched, d-Matrix, Lightmatter
KEY SUPPLIERS / CUSTOMERS	While the company has not publicly disclosed customer names, DensityAI is targeting organizations deploying large-scale AI inference workloads, particularly within automotive, robotics, industrial AI, cloud infrastructure, and frontier-model deployments. The company's architecture is designed for customers seeking to improve the performance, utilization, and energy efficiency of LLM inference workloads at scale
COMPETITIVE POSITIONING	DensityAI's platform is centered on a memory-centric architecture designed specifically for frontier-scale LLM inference. Key differentiators include: • Memory-centric architecture: DensityAI integrates memory directly into the compute stack, seeking to combine the high bandwidth characteristics of SRAM-based architectures with the capacity advantages of HBM-based systems. Management believes this approach can reduce data movement and alleviate memory bottlenecks that constrain large-model inference performance. • Full-stack optimization: The company is co-designing silicon, packaging, systems software, compiler technologies, kernels, and serving infrastructure as a unified platform. This approach is intended to optimize performance across the entire inference stack rather than solely at the accelerator level. • Inference-focused design: Unlike many accelerator architectures originally developed for AI training workloads, DensityAI's platform is designed specifically for large-scale inference. Management believes this specialization can improve utilization, latency, throughput, and energy efficiency for production LLM deployments.
YEAR FOUNDED / ESTABLISHED	2025
WEBSITE	https://www.densityai.com
MANAGEMENT TEAM / FOUNDER(S)	**Ganesh Venkataramanan – Co-Founder & CEO:** Prior to founding DensityAI, Mr. Venkataramanan led Tesla's Dojo supercomputer program and served as Senior Director of Autopilot Hardware. Before Tesla, he held senior architecture and engineering roles at AMD, including Corporate Fellow and Chief Architect. He brings extensive experience designing large-scale AI compute infrastructure and custom silicon platforms. **Bill Chang – Co-Founder:** Mr. Chang previously worked within Tesla's AI and Dojo infrastructure organization and is part of the founding team developing DensityAI's next-generation inference platform. **Ben Floering – Co-Founder:** Mr. Floering previously held engineering roles within Tesla's Dojo development team and is helping lead the development of DensityAI's architecture and systems platform.
SIZE OF COMPANY	62 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	DensityAI is developing a full-stack AI inference platform spanning custom silicon, advanced packaging, systems software, compiler technologies, and rack-scale infrastructure. The company's architecture combines the near-memory compute topology of SRAM-based designs with the density advantages of HBM-based systems, with the goal of improving hardware utilization for frontier-scale LLM inference workloads. Management has stated that the architecture is designed to achieve greater than 90% compute utilization, compared with roughly 7% utilization for conventional GPU architectures on certain large-model inference workloads. DensityAI is actively expanding its engineering organization across silicon design, compiler development, advanced packaging, hardware systems, manufacturing, and infrastructure engineering.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Density AI has raised a seed round and early stage VC round for an undisclosed amount in September 2024 and 2025. Notable Investors: Dolby Family Ventures, Firestreak Ventures, NextGen Venture Partners, PEAK6 Strategic Capital, South Park Commons, imec.expand, Ignite Innovation
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: d-Matrix

COMPANY NAME	d-Matrix
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	d-Matrix develops hardware and software platforms designed for generative AI and large language model (LLM) inference. The company's flagship product, Corsair, is an AI inference platform built specifically for data center deployments. d-Matrix's architecture is based on its proprietary Digital In-Memory Compute (DIMC) technology, which integrates compute and memory resources to address memory bandwidth limitations that increasingly constrain AI inference performance. The company believes this approach can reduce data movement, lower latency, improve power efficiency, and reduce the total cost of deploying large-scale AI inference workloads. The global AI computing market is expected to continue growing as data center IT capex is expected to reach $3T by 2030.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** AMD, CBRS, INTC, MU, NVDA **PRIVATE:** Groq, SambaNova, Etched, Tenstorrent
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry) Customers are not publicly disclosed, but have commitments from high-profile hyperscalers, neoclouds, and frontier AI labs.
COMPETITIVE POSITIONING	d-Matrix's platform is centered on the Corsair architecture, which is designed specifically for generative AI inference. Key differentiators include: • Digital In-Memory Compute (DIMC) architecture: d-Matrix integrates compute directly with memory resources to reduce data movement and address the memory-bandwidth constraints that increasingly limit AI inference performance. Management believes this architecture can improve throughput, lower latency, and enhance energy efficiency relative to conventional accelerator designs. • High-bandwidth memory subsystem: The Corsair platform combines up to 4GB of integrated Performance Memory delivering 300 TB/s of bandwidth with up to 512GB of Capacity Memory in a dual-card configuration. The architecture is designed to support both interactive inference workloads and large-scale frontier models. • Inference-focused design: Unlike many accelerators that were initially developed for AI training workloads, Corsair was designed specifically for generative AI inference. The platform is optimized for token generation, latency-sensitive workloads, and large language model deployments. • Full-stack platform approach: d-Matrix is developing a vertically integrated platform that includes Corsair accelerators, JetStream networking products, and the Aviator software stack. Management's strategy focuses on optimizing the entire inference system, from silicon and packaging through software and rack-scale deployment.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	https://www.d-matrix.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Sid Sheth - Co-founder and CEO:** Prior to d-Matrix, Mr. Sheth held engineering leadership roles at NetLogic Microsystems, Intel, Inphi Corporation and was a co-founder of Aeluros. He holds an MS in Electrical Engineering from Purdue University. **Sudeep Bhoja - Co-founder and CTO:** Prior to d-Matrix, Mr. Bhoja was CTO of the Networking Business Unit at Inphi and held engineering leadership roles at Broadcom, Big Bear Networks, Lucent, and Texas Instruments. He holds an MS in Electrical Engineering from Purdue University.
SIZE OF COMPANY	324 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	d-Matrix's first commercial product platform, Corsair, entered customer sampling prior to commercial production and is being deployed for large-scale AI inference applications. The platform was publicly unveiled on November 19, 2024, and is built around the company's proprietary Digital In-Memory Compute (DIMC) architecture. Corsair is implemented as a PCIe accelerator card featuring four chiplets per package fabricated on TSMC's 6nm process technology and is designed to support both scale-up and scale-out deployments for large language model inference workloads. More recently, d-Matrix announced a strategic collaboration with NVIDIA that will integrate its next-generation inference accelerators into NVIDIA's NVLink Fusion and MGX rack-scale ecosystem, providing a pathway to deploy d-Matrix inference hardware alongside NVIDIA AI infrastructure at hyperscale. Management previously guided to $70-75 million of revenue and break-even profitability during 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	d-Matrix has raised a total of $446.3MM over 5 funding rounds. In its latest rounds of funding, the company received $275MM in November of 2025 for Series C and $10MM in March of 2026 as a later stage funding. Notable investors: Bullhound Capital, Triatomic Capital, Temasek, Qatar Investment Authority (QIA), EDBI, M12 (Microsoft's Venture Fund), Industry Ventures, Nautilus Venture Partners, Mirae Asset, Playground Global, Lam Capital, Palo Alto Networks, Samsung Ventures, SK hynix, Marvell, Entrada Ventures, TSVC, Amino Capital, and Foothill Ventures.
LATEST ROUND VALUATION	$2B post Series C
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Efficient Computer

COMPANY NAME	Efficient Computer
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Energy-Efficient Edge Processors
SIZE OF TOTAL ADDRESSABLE MARKET	Efficient Computer manufactures energy-efficient, general-purpose processors for edge computing and embedded systems. Its platform combines a spatial dataflow architecture with a proprietary compiler to reduce instruction-centric bottlenecks and support AI inference and autonomous workloads on battery-powered devices. The company was spun out of Carnegie Mellon University and is generating revenue.
GEOGRAPHIC PRESENCE	Pittsburgh, Pennsylvania
NOTABLE COMPETITORS	**LISTED:** ARM, NVDA, QCOM, INTC, ADI, CBRS **PRIVATE:** MatX, d-Matrix
KEY SUPPLIERS / CUSTOMERS	BrightAI is identified as a deployment partner for scalable intelligent-edge applications. Suppliers and named customers are N/A.
COMPETITIVE POSITIONING	Key differentiators include: • Spatial dataflow architecture: Designed to eliminate instruction-centric processing bottlenecks. • Energy efficiency: Targets complex AI inference and autonomous pipelines on battery-powered devices with minimal energy consumption. • Hardware-software integration: Combines general-purpose processors with a proprietary compiler. • Edge-computing focus: The Electron E1 targets industrial, defense, consumer-electronics, and emerging AI-driven applications.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://www.efficient.computer/
MANAGEMENT TEAM / FOUNDER(S)	**Brandon Lucia, Ph.D. – Co-Founder, Chief Executive Officer & Board Member:** Full Professor at Carnegie Mellon University (CMU) and Ph.D. from the University of Washington. **Graham Gobieski, Ph.D. – Co-Founder & Chief Technology Officer:** Ph.D.from Carnegie Mellon University (CMU). **Nathan Beckmann, Ph.D. – Co-Founder, Chief Architect & Board Member:** Associate/ Assistant Professor in Computer Science at Carnegie Mellon University and Ph.D. from MIT
SIZE OF COMPANY	48 employees as of February 18, 2026.
GROWTH TRAJECTORY	Efficient Computer is generating revenue and has launched the Electron E1 Evaluation Kit and Efficient Labs. Its August 2026 financing is intended to accelerate the product roadmap, expand engineering and developer headcount, and bring Electron E1 into additional AI-driven applications. Total capital raised increased to $169.29 million.
ANY MAJOR M&A / CORPORATE ACTIONS	Spun out of Carnegie Mellon University in 2022.
INVESTORS / INVESTMENT STAGE	The company raised $93.07 million of early-stage venture funding led by Triatomic Capital on August 24, 2026. Participants included Toyota Ventures, RTX Ventures, Overmatch Ventures, Eclipse Capital, In-Q-Tel, Overlap Holdings, Mana Ventures, Beyond Earth Ventures, What If Ventures, Union Square Ventures, and TO VC. The company has raised $169.29 million in total.
LATEST ROUND VALUATION	The August 2026 round valuation was undisclosed. The latest observed valuation was $225 million post-money in the December 2025 Series A.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.

SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: EnCharge AI

COMPANY NAME	EnCharge AI, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	EnCharge AI develops analog in-memory computing accelerators designed for AI inference across edge and client computing platforms. The company's architecture performs matrix multiplication operations directly within memory using charge-domain computing, reducing data movement and improving energy efficiency relative to conventional digital processor architectures. EnCharge AI is targeting AI deployments in laptops, workstations, edge servers, and other power-constrained environments where performance-per-watt, latency, and cost are critical. AI inference market TAM is expected to exceed $600B by 2033.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** AMBQ, MCHP, QCOM, TXN **PRIVATE:** Hailo, d-Matrix, Mythic, Groq
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry) Customers are not publicly disclosed, but management has referenced engagements with OEM and ODMs.
COMPETITIVE POSITIONING	EnCharge AI's core differentiator is its proprietary analog in-memory computing architecture, which integrates storage and computation within the memory subsystem. Key attributes include: • Analog in-memory computing architecture: The company's charge-domain computing technology performs matrix multiplication operations directly within memory, reducing the movement of data between memory and compute resources that constrains conventional digital architectures. • High energy efficiency: Management states that its architecture can deliver up to 20× higher efficiency (TOPS/W) than conventional AI accelerators by minimizing data movement and leveraging analog computation. • Edge inference optimization: Unlike cloud-oriented accelerators designed primarily for data center deployments, EnCharge AI's products target inference workloads in laptops, workstations, edge servers, and embedded systems operating under stringent power and thermal constraints. • Scalable hardware and software platform: The company's architecture incorporates custom silicon, software tools, compilers, and runtime infrastructure designed to support deployment across multiple form factors ranging from M.2 modules to PCIe accelerator cards.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://www.enchargeai.com/
MANAGEMENT TEAM / FOUNDER(S)	**Naveen Verma - Co-founder and CEO:** Prior to EnCharge AI, Dr. Verma has been a Professor of Electrical & Computer Engineering at Princeton since 2009, where six years of research on charge-domain computing underpinning EnCharge AI's core technology. He holds an MS and PhD in Electrical Engineering from MIT and a BASc from the University of British Columbia. **Kailash Gopalakrishnan - Co-founder and CTO:** Prior to EnCharge AI, Mr. Gopalkrishnan was an IBM Fellow who led worldwide AI-hardware R&D, compiler, and low-precision arithmetic programs. He earned a PhD from Stanford University. **Echere Iroaga - COO:** Prior to EnCharge AI, Mr. Iroaga was the VP and GM of MACOM's connectivity business unit, founder of Renovus (acquired by Ikanos Communications), and earlier held engineering roles at QCOM, TER, and TXN. He holds a PhD in Electrical Engineering.
SIZE OF COMPANY	~112 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	On May 29, 2025, EnCharge AI announced the EN100, the company's first commercial AI accelerator product built on its analog in-memory-computing architecture. The EN100 is available in both M.2 and PCIe form factors and delivers 200+ TOPS of AI compute within an 8.25W power envelope for client and edge computing platforms. The launch marked EnCharge AI's transition from research and development to commercial product deployment.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	EnCharge AI raised $144MM over 3 funding rounds. In its latest Series B round of funding, the company received $100MM in November of 2025. Notable investors: Tiger Global, Maverick Silicon, Capital TEN, SIP Global Partners, Zero Infinity Partners, CTBC VC, Vanderbilt University, Morgan Creek Digital, RTX Ventures, Anzu Partners, Scout Ventures, AlleyCorp, ACVC, S5V, VentureTech Alliance, Samsung Ventures, HH-CTBC, Constellation Technology Ventures, and In-Q-Tel (IQT).
LATEST ROUND VALUATION	$430MM post Series B
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and edge processor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and edge processor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Etched

COMPANY NAME	Etched
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	Etched develops AI inference hardware designed exclusively for Transformer-based neural networks. The company's flagship product, Sohu, is a custom application-specific integrated circuit (ASIC) that hardwires Transformer operations directly into silicon rather than executing them on general-purpose GPU architectures. By eliminating architectural flexibility and optimizing specifically for Transformer inference, Etched aims to deliver significantly higher throughput, lower latency, and improved economics relative to conventional GPU-based deployments. The global AI computing market is expected to continue growing as data center IT capex is expected to reach $3T by 2030.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:**AMD, INTC, NVDA **PRIVATE:** Groq, d-Matrix, Tenstorrent, SambaNova
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry). One of Etched's first customers was Jane Street who received and deployed the initial hardware rack for operational workloads in July 2026.
COMPETITIVE POSITIONING	Etched's core differentiator is its transformer-specific ASIC architecture. Key attributes include: • Transformer-only architecture: Sohu is purpose-built for Transformer inference, with Transformer attention mechanisms implemented directly in fixed-function silicon rather than executed on programmable compute resources. Management believes this specialization enables significantly higher efficiency and throughput than general-purpose AI accelerators. • Large memory subsystem: Each Sohu processor integrates 144GB of HBM3E memory to support large-model inference workloads. • Rack-scale system design: Etched designed the full inference platform, including chips, boards, cooling, networking, and rack infrastructure, enabling optimization across the entire hardware stack rather than at the silicon level alone. • Open software ecosystem: The company provides a software stack designed to support frontier AI models and large-scale inference deployments. Current systems have been demonstrated running models including Llama, Qwen, DeepSeek, and Mamba.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://www.etched.com
MANAGEMENT TEAM / FOUNDER(S)	**Gavin Uberti - Co-founder and CEO:** Prior to Etched, Uberti held various software engineering positions while studying at Harvard. At Harvard, Uberti was getting his master's degree concurrently with his bachelor's degree before dropping out to start Etched. **Robert Wachen - Co-founder and COO:** Prior to Etched, Wachen founded various companies, including a startup incubator, edtech startup, cake non-profit, and influencer marketing agency. While at Harvard University, he designed the Decision Science degree. **Chris Zhu - Co-founder:** Prior to Etched, Zhu held various software engineering internships and was a teaching fellow at Harvard.
SIZE OF COMPANY	~534 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Management previously indicated that the company's June 2024 Series A financing would support development and tape-out of the Sohu processor at TSMC. On June 30, 2026, Etched emerged from stealth and announced: First-pass (A0) silicon success on TSMC N4P process technology, 1 working Sohu inference processor and rack-scale inference platform, and more than $1 billion in signed customer contracts. Etched reported that Sohu systems are running models including DeepSeek, Qwen, Mamba, and Llama and that rack-scale inference systems are being deployed for customer evaluation including at Jane Street.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Etched has raised $1.7B in funding since inception. In August 2026, Etched announced a $700MM financing led by Jane Street at a reported $21 billion valuation. Notable investors: Sequoia Capital, Andreessen Horowitz, Stripes, Jane Street, SK hynix, VentureTech Alliance, Positive Sum, Ribbit Capital, Hudson River Trading, Jump Trading, Two Sigma, Radical Ventures, Primary Venture Partners, Peter Thiel, Andrej Karpathy, Geoffrey Hinton, Fei-Fei Li, Arthur Mensch, Scott Wu, and Stanley Druckenmiller.
LATEST ROUND VALUATION	Etched's most recent financing round in August 2026 reportedly valued the company at $21 billion, up from a reported $10.3 billion valuation in July 2026 and $5 billion in late 2025.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and ASIC/FPGA companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and ASIC/FPGA companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Fractile

COMPANY NAME	Fractile
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Inference Accelerators
SIZE OF TOTAL ADDRESSABLE MARKET	Fractile develops AI inference processors and systems designed to accelerate frontier-model deployment by addressing memory-bandwidth constraints that increasingly limit large language model (LLM) inference. Founded in 2022 as a spinout from the University of Oxford, the company is focused exclusively on AI inference rather than training. Fractile's architecture is based on a proprietary memory-compute fusion approach that integrates memory and computation more closely than conventional accelerator designs. Management believes this architecture can reduce data movement, improve throughput, lower latency, and significantly reduce the cost of serving large-scale AI workloads. According to the company, inference is becoming the primary bottleneck in AI deployment as models generate increasingly long outputs and require substantially more compute at runtime.
GEOGRAPHIC PRESENCE	London, Bristol, and Newbury, United Kingdom.
NOTABLE COMPETITORS	**LISTED:** CBRS, NVDA, AMD **PRIVATE:** SambaNova Systems, Groq, Etched, Salience Labs, MatX, Graphcore
KEY SUPPLIERS / CUSTOMERS	A September 2026 news item references a $250 million Anthropic deal. Suppliers are N/A.
COMPETITIVE POSITIONING	Fractile's platform is built around a proprietary memory-centric architecture designed specifically for frontier AI inference. Key differentiators include: • Memory-compute fusion architecture: Fractile performs matrix multiplication operations directly within memory resources, reducing the need to move model parameters between separate memory and compute units. The company believes this approach addresses the memory-bandwidth bottleneck that increasingly constrains inference performance and economics on conventional accelerator architectures. • Inference-only optimization: Unlike many accelerators originally developed for AI training workloads, Fractile's architecture is designed specifically for inference. Management argues that inference workloads require a fundamentally different hardware design philosophy centered on latency, throughput, and serving efficiency rather than training performance. • Long-context and reasoning workload focus: Fractile is targeting emerging workloads such as reasoning models, agentic AI systems, software engineering agents, scientific discovery, and other sequential-compute applications that may require generating millions of tokens per task. The company believes these workloads will place increasing importance on inference speed and cost. • Full-stack systems approach: Fractile is developing both silicon and systems intended for data center deployment rather than a standalone chip. The company describes its strategy as spanning AI systems research, chip architecture, and process-level innovation to optimize inference performance at scale.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://www.fractile.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Walter Goodwin, Ph.D. – Co-Founder & Chief Executive Officer:** Dr. Goodwin founded Fractile in 2022 while completing doctoral research at the University of Oxford's Robotics Institute. His academic background includes robotics, machine learning, AI systems, and semiconductor architecture. Dr. Goodwin's central thesis is that inference, rather than training, will become the primary infrastructure bottleneck limiting future AI deployment. **Matthew Arnold – Chief Financial Officer:** Prior to joining Fractile, Mr. Arnold was CFO and COO of Mind Foundry, an Oxford-based AI software company. Earlier in his career, he spent more than a decade at Goldman Sachs, ultimately serving as a Managing Director in Principal Strategic Investments and Fixed Income Analytics. He was also a founding investor and Principal at Oxford Science Enterprises, where he worked with numerous Oxford deep-tech spinouts including PQShield, Diffblue, BibliU, and OxfordVR. Mr. Arnold holds a Master of Engineering in Information Engineering from the University of Cambridge.
SIZE OF COMPANY	89 employees as of May 13, 2026

GROWTH TRAJECTORY	In February 2026, Fractile announced plans to invest £100 million in its UK operations over three years, including a new hardware engineering facility in Bristol dedicated to AI-system assembly and software testing. Reuters reported that the expansion is intended to support commercialization of Fractile's next-generation AI inference systems and scale the company's engineering organization. In May 2026, Fractile announced a Series B financing to accelerate development of its first commercial inference platform. The company stated that proceeds would be used to bring its first chips and systems to production and expand engineering operations across the UK, US, and Taiwan. Fractile's first commercial silicon is expected to enter production in 2027, with the company's architecture specifically designed to address inference latency and throughput limitations associated with large-scale frontier AI deployments. Media reports have indicated customer discussions with frontier AI developers, including Anthropic, although no commercial deployment has yet been publicly announced by Fractile.
ANY MAJOR M&A / CORPORATE ACTIONS	Spun out of the University of Oxford in 2022.
INVESTORS / INVESTMENT STAGE	Fractile completed a $303.72 million Series B on May 5, 2026, led by Founders Fund, Accel, and Factorial Funds. Other investors included Oxford Science Enterprises, Kindred Capital, NATO Innovation Fund, 8VC, Conviction Investment Partners, Gigascale Capital, O1A, Buckley Ventures, 01 Advisors, and Felicis. Total capital raised was $321.20 million following the completed Series B.
LATEST ROUND VALUATION	The completed May 2026 Series B valued Fractile at $1.05 billion pre-money and $1.35 billion post-money.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and ASIC/FPGA companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and ASIC/FPGA companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Furiosa

COMPANY NAME	FuriosaAI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Data Center AI Inference Accelerators
SIZE OF TOTAL ADDRESSABLE MARKET	Furiosa manufactures AI accelerators for data-center inference. Its neural processing units and software stack use a tensor-contraction architecture to execute large-language-model, generative-AI, and computer-vision workloads with high throughput and low power consumption.
GEOGRAPHIC PRESENCE	Seoul, Korea
NOTABLE COMPETITORS	**LISTED:** NVDA, INTC, CBRS **PRIVATE:** SambaNova Systems, Groq, Rebellions
KEY SUPPLIERS / CUSTOMERS	Broadcom is identified as a partner for a next-generation inference platform. Helikai and ByteBridge are also identified as commercial or technology partners. Named suppliers and customers are N/A.
COMPETITIVE POSITIONING	FuriosaAI's core differentiator is its inference-focused Tensor Contraction Processor (TCP) architecture, designed to improve compute utilization and energy efficiency across modern AI workloads. Key attributes include: • Inference-focused NPU architecture: Purpose-built for data-center AI inference. • Tensor-contraction architecture: Supports large-language-model and computer-vision workloads. • Full-stack offering: Combines neural processing units with a specialized software stack. • Power efficiency: Targets high-throughput inference with lower power consumption. • Next-generation roadmap: Series D proceeds are intended to support chips using HBM4.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://www.furiosa.ai/
MANAGEMENT TEAM / FOUNDER(S)	**June Paik – Co-Founder & Chief Executive Officer:** Prior to co-founding FuriosaAI in 2017, June Paik held engineering roles at Samsung Electronics and AMD, where he worked on storage systems and high-performance computing software. He founded FuriosaAI with the goal of developing purpose-built AI inference processors that can deliver competitive performance with lower power consumption than conventional GPU-based solutions. Paik holds a Master's degree in Electrical Engineering from the Georgia Institute of Technology. **Hanjoon Kim – Co-Founder & Chief Technology Officer:** Prior to co-founding FuriosaAI, Hanjoon Kim was an engineer at Samsung Electronics. At FuriosaAI, he leads AI accelerator architecture development and the design of the company's neural processing unit (NPU) platforms. His research interests include neural network accelerator architectures and interconnection networks. Kim received a Ph.D. in Computer Science from KAIST (Korea Advanced Institute of Science and Technology). **Jeehoon Kang – Co-Founder & Chief Research Officer:** Jeehoon Kang is a co-founder of FuriosaAI and currently serves as Chief Research Officer, leading the company's compiler, software architecture, and research efforts. Prior to assuming this role, he was a faculty member at KAIST School of Computing, where his research focused on parallel systems, compilers, programming languages, concurrency, and verification. He earned both his B.S. and Ph.D. in Computer Science from Seoul National University and has received distinguished paper awards at leading computer science conferences including PLDI and POPL.
SIZE OF COMPANY	213 employees as of September 3, 2026.
GROWTH TRAJECTORY	Furiosa is generating revenue. There are reports of revenue of $4.0 million for 2025 and $68.4 million for 2026. The company raised an undisclosed Series D in September 2026 to develop HBM4-based next-generation AI chips and expand South Korea's domestic AI semiconductor ecosystem. Total capital raised was $260.70 million.
ANY MAJOR M&A / CORPORATE ACTIONS	Furiosa was in discussions to be acquired by Meta Platforms, with the proposed transaction covering 100% of the company. The transaction was subsequently cancelled in February 2025 per Bloomberg.

INVESTORS / INVESTMENT STAGE	Furiosa completed an undisclosed Series D on September 3, 2026, with Premier Partners, AP Investment Finance, The National Growth Fund, and other undisclosed investors. The prior $124.17 million Series C included Invites Ventures, Intellectual Discovery, Xperix, Korea Investment Partners, Korea Development Bank, IMM Investment, DSC Investment, and other investors. Total capital raised was $260.70 million.
LATEST ROUND VALUATION	The September 2026 Series D valuation was undisclosed. The latest disclosed valuation was an estimated $606.25 million pre-money and $730.43 million post-money for the July 2025 Series C per Pitchbook.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and ASIC/FPGA companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and ASIC/FPGA companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: MatX

COMPANY NAME	MatX
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	MatX is a semiconductor startup developing purpose-built AI accelerators designed for large-scale training and inference of foundation models and generative AI workloads. Founded in 2023 by former Google TPU engineers Mike Gunter and Reiner Pope, the company is building a vertically integrated AI computing platform centered on its proprietary MatX One architecture. Unlike general-purpose GPUs, MatX's hardware and software stack is designed specifically for transformer-based AI workloads, with an emphasis on performance-per-dollar, system utilization, and large-scale cluster efficiency. MatX participates in the broader AI infrastructure market, targeting the growing demand for compute associated with large language models (LLMs), multimodal AI systems, and agentic AI applications with TAM expected to exceed $600B by 2033.
GEOGRAPHIC PRESENCE	Mountain View, California
NOTABLE COMPETITORS	**LISTED:** AMD, CBRS, INTC, NVDA **PRIVATE:** d-Matrix, Etched, Tenstorrent, Groq, SambaNova, Fractile, Positron, Encharge AI, Mythic, Tachyum, Cornami Relative to many competitors, MatX's differentiation centers on building an end-to-end platform optimized specifically for transformer workloads rather than adapting legacy compute architectures to AI applications.
KEY SUPPLIERS / CUSTOMERS	While MatX has not publicly disclosed a broad customer base, the company is targeting hyperscale cloud providers, frontier model developers, AI-native startups, sovereign AI infrastructure projects, enterprise AI deployments, and data center operators
COMPETITIVE POSITIONING	MatX's platform is centered on the MatX One architecture, which combines custom silicon, networking, systems software, and cluster-level optimization. Key differentiators include: • Transformer-native architecture: Unlike GPUs, which were originally designed for graphics workloads, MatX's architecture was developed specifically for transformer models and generative AI workloads. The design is optimized for matrix operations, attention mechanisms, and token-generation workflows that dominate modern AI training and inference. • Memory-optimized design: MatX One combines high-bandwidth memory (HBM) with substantial on-chip SRAM capacity to minimize data movement and improve hardware utilization. The architecture is designed to keep model weights, activations, and intermediate computations closer to compute resources, reducing external memory bottlenecks. • Full-stack optimization: The company is developing hardware, networking, and software as an integrated platform. Management believes improving cluster utilization and reducing inefficiencies is as important as increasing raw compute performance. • Focus on cost efficiency: MatX aims to deliver superior performance-per-dollar relative to existing AI infrastructure solutions by optimizing the entire system architecture for transformer workloads.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://matx.com
MANAGEMENT TEAM / FOUNDER(S)	**Mike Gunter – Co-founder & CEO:** Prior to founding MatX, Mr. Gunter spent more than a decade at Google, where he served as the head of TPU hardware engineering and played a significant role in the development of multiple generations of Google's Tensor Processing Units. He also previously served in engineering leadership roles at AMD. **Reiner Pope – Co-founder & CTO:** Mr. Pope previously worked at Google, where he was a founding member of the TPU architecture team and contributed to the development of Google's AI accelerator roadmap. Prior to Google, he worked on advanced processor architectures and machine learning infrastructure technologies.
SIZE OF COMPANY	~134 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	MatX announced a $500 million Series B financing, bringing total disclosed funding to $620 million. Management stated that the proceeds would support broader customer deployment of the MatX One platform, expansion of manufacturing capacity, and scaling of commercial operations. Following the financing, MatX announced plans to accelerate commercialization of the MatX One platform and expand engagement with hyperscale cloud providers, frontier model developers, enterprise AI customers, sovereign AI infrastructure projects, and data center operators. The company's growth strategy is centered on delivering a vertically optimized hardware and software platform capable of improving utilization and lowering the total cost of ownership for large-scale AI training and inference deployments.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	MatX has raised $620 million to date across multiple financing rounds. The company's latest funding round was a $500 million Series B completed in February 2026. Notable Investors: Spark Capital, Jane Street, NFDG, Andrej Karpathy, Marvell Technology. Harpoon VC
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: MultiLucent

COMPANY NAME	MultiLucent
SECTOR / SUB-SECTOR	TMT / Inference / Governed AI Systems
SIZE OF TOTAL ADDRESSABLE MARKET	MultiLucent develops governed AI systems for consequential applications in education and healthcare. The company's products and research focus on controlled intelligent systems that support governance considerations across multiple industry verticals. Its products are in beta testing.
GEOGRAPHIC PRESENCE	Mesa, Arizona
NOTABLE COMPETITORS	**LISTED:** Alphabet, TEM, ISLMF, BFLY **PRIVATE:** Numerion Labs, Lambda, Kneron, segMedix
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	MultiLucent's core differentiator is its governed AI architecture, designed to preserve human oversight and control in consequential applications. Key attributes include: • Governed AI: MultiLucent develops controlled AI systems that preserve context, organize evidence, expose uncertainty, and maintain explicit human authority over consequential actions. • Vertical orientation: The company targets education and healthcare, with its initial commercial application focused on AI-assisted workflows in higher education. • Broad technology scope: MultiLucent spans AI software and applied research, with the company participating across a broad set of technology categories including workflow and database software, healthcare systems and diagnostics, semiconductors, and educational software. • Early-stage product deployment: MultiLucent's flagship LEMz TA platform is currently in working beta, operating in controlled Canvas course environments with professor-controlled safeguards for consequential actions.
YEAR FOUNDED / ESTABLISHED	2020
WEBSITE	https://multilucent.com/
MANAGEMENT TEAM / FOUNDER(S)	**David Wallace, Ph.D. – Founder, Chief Executive Officer & Chief Scientist:** Prior to founding MultiLucent, David Wallace held technology development and innovation roles at Intel Corporation, where he developed and patented technology and later served as an Entrepreneur in Residence. He also founded an Intel spinout company that was subsequently reacquired by Intel. In addition to his corporate experience, Dr. Wallace has served as an instructor at Arizona State University and has worked in innovation, product development, human-systems engineering, and applied AI. At MultiLucent, he leads the company's scientific direction, product architecture, intellectual property strategy, and responsible AI initiatives. He holds a Ph.D. in Human Systems Engineering from Arizona State University.
SIZE OF COMPANY	~15 employees on PitchBook (accessed 9/16/2026)
GROWTH TRAJECTORY	MultiLucent is in beta testing and joined Microsoft for Startups on January 9, 2026, without exchanging equity or receiving funding. The source reports that the company is seeking capital to support commercialization, infrastructure, intellectual property development, compliance, and team growth.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	MultiLucent previously raised $2.20 million from Skysong Innovations and joined Microsoft for Startups as a non-equity accelerator participant.
LATEST ROUND VALUATION	The latest completed financing valued MultiLucent at $41.80 million pre-money and $44.00 million post-money. The source records a prospective 2027 Series A at a $3.80 billion pre-money and $4.20 billion post-money valuation, but this transaction is marked upcoming rather than completed.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	Latest completed preferred-equity financing valuation; prospective financing valuation; P/S or EV/S once measurable revenue is available.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Mythic

COMPANY NAME	Mythic, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	Mythic develops analog AI processors for inference acceleration and edge AI deployment. The company's product portfolio includes the M1076 Analog Matrix Processor, analog processing units (APUs), software development kits (SDKs), and the recently introduced Starlight sensing platform. Mythic targets a broad range of end markets, including defense, aerospace, robotics, automotive, industrial automation, machine vision, drones, smart cities, and AI infrastructure. Mythic's technology is built around analog compute, where neural network weights are stored directly within memory and computation is performed with minimal digital circuitry. Management believes this approach can address the power, latency, and cost limitations of conventional AI processors, particularly in edge AI deployments where power consumption and thermal constraints are critical. The AI inference market TAM is expected to exceed $600B by 2033.
GEOGRAPHIC PRESENCE	Austin, United States
NOTABLE COMPETITORS	**LISTED:** AMBQ, MCHP, QCOM, TXN **PRIVATE:** Hailo, d-Matrix, Groq
KEY SUPPLIERS / CUSTOMERS	Manufacturer: UMC Japan/USJC (foundry), Codasip Customers: Lockheed Martin, Honda, US DoD, Defense and Aerospace contractors, Automotive OEMs
COMPETITIVE POSITIONING	Mythic's platform is centered on its Analog Compute Engine (ACE), which combines analog computing, compute-in-memory (CIM), and dataflow architectures. Key differentiators include: • Analog compute architecture: Mythic stores neural network weights directly within analog memory and performs computation inside the memory array, reducing data movement and improving energy efficiency relative to conventional digital architectures. • Compute-in-memory (CIM) design: The company's Analog Compute Engine integrates memory and computation within the same architecture, helping alleviate memory bandwidth bottlenecks associated with AI inference workloads. • Energy-efficient inference: Management states that its latest analog processor architecture can deliver up to 100× greater energy efficiency than traditional GPU-based platforms for certain AI inference workloads, while reducing latency, power consumption, cooling requirements, and system cost. • Edge-to-system scalability: The architecture is designed to scale from sub-watt edge devices to larger AI inference systems, enabling deployment across cameras, drones, robotics systems, industrial equipment, and defense platforms.
YEAR FOUNDED / ESTABLISHED	2012
WEBSITE	http://www.mythic.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Taner Ozcelik - CEO:** Prior to Mythic, Mr. Ozcelik held many roles as a semiconductor executive, including EVP and GM at Luminar Technologies, SVP and GM of Intelligent Sensing Group at ON, and VP and GM of Automotive and Embedded Business at NVDA. **Dave Fick – Co-founder and CTO:** Mr. Fick recently transitioned back to CTO after being CEO for ~1.5 years. Prior to being CEO, he served as CTO for about ~10.5 years. Mr. Fick received his Ph.D. in Computer Science and Engineering from the University of Michigan. **Laura Fick - Chief Analog Compute Architect, Founding Engineer:** Prior to Mythic, Ms. Fick developed Mythic's Analog Compute technology as part of her PhD thesis at the University of Michigan. Since then, Ms. Fick has led the efforts on the development of new algorithms, circuits and silicon testing methodologies for analog compute.
SIZE OF COMPANY	~97 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	In December 2022, Mythic restructured operations amid a challenging funding environment. In March 2023, the company secured $13 million in financing to restart operations and continue commercialization of its analog compute technology. In December 2025, Mythic announced an oversubscribed $125 million funding round led by DCVC to accelerate development of its Analog Processing Unit (APU) platform across both edge and data center AI applications. The company also introduced its Starlight sensing platform, expanded into automotive, robotics, and defense markets, and broadened its roadmap to support AI inference infrastructure and 1T + parameter large language model (LLM) deployments.
ANY MAJOR M&A / CORPORATE ACTIONS	In May 2026, Mythic acquired Videantis, a digital processor IP company, to accelerate development of a hybrid AI compute platform spanning edge devices, autonomous systems, and data center AI deployments.
INVESTORS / INVESTMENT STAGE	Mythic has raised $300MM over multiple funding rounds. In its latest venture financing, the company raised $125MM in December 2025 in an oversubscribed round led by DCVC. Notable investors: DCVC, NEA, Atreides Management, Future Ventures, SBVA (SoftBank Ventures Asia), S3 Ventures, Linse Capital, One Madison Group, Catapult, UDC, Honda Motor, Lockheed Martin, and eCAPITAL.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and edge processor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and edge processor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Neurophos

COMPANY NAME	Neurophos
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / Optical AI Compute
SIZE OF TOTAL ADDRESSABLE MARKET	Neurophos manufactures optical processors designed to avoid resistive losses associated with electrical circuits. Its metasurface-based optical compute-in-memory system operates with sparse kernels in absolute and Fourier space, enabling larger kernel sizes without corresponding increases in power consumption or latency.
GEOGRAPHIC PRESENCE	Austin, Texas
NOTABLE COMPETITORS	**LISTED:** Lightelligence **PRIVATE:** Lightmatter, Optalysys, Mythic, Ayar Labs, Etched, Quadric
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	Neurophos's core differentiator is its metasurface-based optical computing architecture, which uses photonic processing to perform AI compute with substantially higher compute density and energy efficiency than conventional electronic architectures. Key attributes include: • Metasurface-based optical processing: Uses optical metamaterials to perform compute functions. • Optical compute-in-memory: Designed to reduce electrical data movement and associated resistive losses. • Sparse-kernel support: Operates in absolute and Fourier space. • Power and latency scaling: Intended to increase kernel size without increasing power consumption or latency. • Patent portfolio: Includes optical vector-matrix multiplication, optical convolution lenses, optical metamaterials, and tunable metamaterial elements.
YEAR FOUNDED / ESTABLISHED	2020
WEBSITE	https://neurophos.com
MANAGEMENT TEAM / FOUNDER(S)	**Patrick Bowen, Ph.D. – Co-Founder & Chief Executive Officer:** Prior to co-founding Neurophos in 2020, Patrick Bowen conducted advanced research in metamaterials and photonics at Duke University under Professor David Smith. Following his doctoral work, he co-founded Metacept, a metamaterials commercialization and engineering firm focused on defense, sensing, and communications applications. At Neurophos, Dr. Bowen leads the company's effort to commercialize photonic AI accelerators that leverage metamaterial-based optical processing for high-performance AI inference. He holds a Ph.D. in Electrical Engineering from Duke University and an M.S. in Micro-Nano Systems from ETH Zurich.
SIZE OF COMPANY	~63 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Neurophos is generating revenue and entered a commercialization and team-expansion phase following its Series A. The company raised $109.17 million in November 2025 and subsequently opened a Silicon Valley office. Total capital raised was $119.17 million.
ANY MAJOR M&A / CORPORATE ACTIONS	Neurophos acquired SiliconBee on August 11, 2025. SiliconBee is identified as a San Jose-based general-purpose semiconductor startup founded in 2024 and is listed as a Neurophos subsidiary.
INVESTORS / INVESTMENT STAGE	Neurophos completed a $109.17 million Series A on November 25, 2025, led by Gates Frontier. Participants included M12, Microsoft, Aramco Ventures, Bosch Ventures, Morgan Creek Capital Management, Carbon Direct Capital, Wonderstone Ventures, Silicon Catalyst Ventures, Beyond Earth Ventures, Vannevar Ventures, MetaVC Partners, Mana Ventures, Space Capital, Gaingels, Geometry, DNX Ventures, Alumni Ventures, and others. Total capital raised was $119.17 million.
LATEST ROUND VALUATION	The November 2025 Series A valued Neurophos at $200.83 million pre-money and $310.00 million post-money.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.

SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Normal Computing

COMPANY NAME	Normal Computing
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI-Native Semiconductor Design and Thermodynamic Computing
SIZE OF TOTAL ADDRESSABLE MARKET	Normal Computing develops a semiconductor and computational-engineering platform for dependable AI in high-stakes applications. Its offering combines a unified electronic-design-automation platform, probabilistic machine learning grounded in thermodynamic physics, and thermodynamic computing hardware.
GEOGRAPHIC PRESENCE	New York, New York
NOTABLE COMPETITORS	**LISTED:** CBRS **PRIVATE:** Mythic, SambaNova Systems, Groq, Extropic, Cognichip, Multiscale Technologies.
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	Normal Computing's core differentiator is its physics-based approach to AI computing, combining AI-native semiconductor design with thermodynamic computing architectures to improve the efficiency of AI inference and probabilistic workloads. Key attributes include: • AI-native semiconductor design: Normal Computing provides an AI-native EDA platform that combines simulation and synthesis engines with AI agents to accelerate semiconductor design, verification, and optimization. • Thermodynamic computing: The company develops physics-based computing architectures that use stochastic dynamics as a computational primitive, including thermodynamic processors designed for AI inference, linear algebra, sampling, and probabilistic workloads. • Dependable AI: Normal Computing targets advanced AI workloads including inference and reasoning, with its hardware roadmap prioritizing applications such as long-context agentic reasoning. • Energy efficiency: The architecture is designed to increase AI inference throughput per dollar and per watt by exploiting physical dynamics rather than relying exclusively on conventional deterministic computation. • Integrated platform: Normal Computing combines semiconductor engineering software with probabilistic computing research and purpose-built AI hardware, spanning chip design through deployment.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://normalcomputing.com/

MANAGEMENT TEAM / FOUNDER(S)	**Faris Sbahi – Co-Founder & Chief Executive Officer:** Prior to co-founding Normal Computing in 2022, Faris Sbahi held research and engineering roles at X (Alphabet's Moonshot Factory), Google, and Google Brain, where he worked on machine learning systems, Bayesian inference, decision-making under uncertainty, and large-scale production AI infrastructure. At Normal Computing, he leads the company's effort to develop thermodynamic computing hardware and probabilistic AI systems designed to improve the efficiency, reliability, and scalability of next-generation AI workloads. He studied at Duke University. **Matthias Tan – Co-Founder & Chief Operating Officer:** Prior to co-founding Normal Computing, Matthias Tan held product, partnership, venture capital, and technology leadership roles across the startup and enterprise technology ecosystem. Public profiles indicate experience with Alphabet's partnership initiatives, Pony.ai, Boston Consulting Group, and Intel Corporation. He also served as a Venture Partner at NextGen Venture Partners and a mentor at Alchemist Accelerator. Dr. Tan earned a Ph.D. from Purdue University. At Normal Computing, he oversees operations, strategic partnerships, and commercialization activities. **Antonio Martinez Jr. – Co-Founder:** Prior to co-founding Normal Computing, Antonio Martinez Jr. worked in quantum computing, machine learning, and probabilistic systems research, including roles at Google Quantum and X. His academic work focused on quantum information systems, quantum computing, and machine learning applications, and he has contributed to publications spanning quantum machine learning and thermodynamic computing. At Normal Computing, he has served as the company's chief architect and has played a central role in developing the firm's probabilistic computing architectures and next-generation thermodynamic hardware platforms.
SIZE OF COMPANY	~105 employees on LinkedIn (accessed 9/18/2026)
GROWTH TRAJECTORY	Normal Computing is generating revenue. Its March 2026 Series B proceeds are intended to advance AI-native semiconductor design and physics-based computing architectures, with a focus on improving intelligence per dollar and per watt. The company's post-money valuation increased from $28.52 million at seed, to $66.95 million at Series A, and to $195.00 million at Series B.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Normal Computing raised $50 million in a Series B led by Samsung Catalyst Fund on March 4, 2026. Other investors included North 50 Venture Capital, View Different, Micron Ventures, Galvanize, Brevan Howard Macro Venture Fund, ArcTern Ventures, Drive Capital, Celesta Capital, First Spark Ventures, Galvanize Climate Solutions, Hyperion Capital, and Eric Schmidt. Total capital raised was $75.47 million.
LATEST ROUND VALUATION	The March 2026 Series B valued Normal Computing at $145 million pre-money and $195 million post-money.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Positron

COMPANY NAME	Positron
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Inference
SIZE OF TOTAL ADDRESSABLE MARKET	Positron AI develops inference-optimized AI hardware, accelerator systems, and software designed specifically for deploying and serving large language models (LLMs), agentic AI systems, and generative AI applications. Unlike general-purpose GPUs, Positron's architecture focuses exclusively on AI inference, utilizing a memory-centric design that prioritizes memory bandwidth utilization, energy efficiency, and performance-per-dollar. The company believes that memory capacity and bandwidth, rather than raw compute, have become the primary bottlenecks for next-generation AI inference workloads. The company currently ships its Atlas inference platform and is developing its next-generation Titan platform powered by proprietary Asimov silicon. The AI inference market TAM is expected to exceed $600B by 2033.
GEOGRAPHIC PRESENCE	Reno, Nevada
NOTABLE COMPETITORS	**LISTED:** AMD, INTC, NVDA **PRIVATE:** Groq, d-Matrix, Etched, Tenstorrent Relative to many competitors, Positron differentiates itself through a memory-first architecture focused on maximizing bandwidth utilization and memory capacity for long-context, agentic, multimodal, and multi-trillion-parameter inference workloads.
KEY SUPPLIERS / CUSTOMERS	Manufacturer: U.S.-based design, fabrication, and assembly partners Customers: Cloudflare, Parasail, SnapServe, Jump Trading, AI cloud providers, enterprise AI deployments
COMPETITIVE POSITIONING	Positron's platform is centered on a memory-first architecture designed specifically for AI inference. Key differentiators include: • Inference-only architecture: Unlike GPUs designed for both training and inference workloads, Positron's technology is optimized exclusively for transformer inference, token generation, and serving large-scale AI models. • Memory-centric design: Management believes memory bandwidth and capacity are the primary bottlenecks limiting future AI inference performance. Atlas reportedly achieves approximately 93% of theoretical HBM bandwidth utilization, substantially higher than typical GPU utilization rates on memory-bound workloads. • Atlas production platform: Positron states that Atlas delivers approximately 3.5x better performance-per-dollar and up to 66% lower power consumption than NVIDIA's H100 on selected inference workloads. Atlas supports Hugging Face models, OpenAI-compatible APIs, and deployments in standard air-cooled data centers. • Next-generation Asimov silicon: Positron's custom Asimov architecture is designed around high-capacity LPDDR5X memory rather than conventional HBM. Management has stated that each Asimov device will support more than 2,304 GB of memory, compared with approximately 384 GB for competing next-generation GPU systems. • Titan platform: The upcoming Titan system will incorporate four Asimov accelerators and provide 8+ TB of memory per system, support models up to 16 trillion parameters, and enable context windows exceeding 10 million tokens.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://positron.ai

MANAGEMENT TEAM / FOUNDER(S)	**Mitesh Agrawal – CEO:** Prior to joining Positron in February 2025, Mr. Agrawal served as Chief Operating Officer at Lambda, where he helped scale the company from approximately $500,000 to nearly $500 million in annualized revenue run rate and supported funding rounds totaling more than $1 billion. He joined Positron to lead commercialization, scaling, and go-to-market execution as the company expanded deployment of its Atlas inference platform and developed its next-generation Asimov silicon roadmap. **Thomas Sohmers – Co-founder and CTO:** Mr. Sohmers co-founded Positron in 2023 and serves as Chief Technology Officer. A Thiel Fellow and serial semiconductor entrepreneur, he previously founded REX Computing and Terapute Technologies and later served as Director of Technology Strategy and Head of Distributed Systems at Groq. At Positron, he leads development of the company's inference-first hardware architecture, custom silicon roadmap, and large-scale AI infrastructure strategy. **Edward Kmett – Co-founder and Chief Scientist:** Mr. Kmett co-founded Positron and serves as Chief Scientist. A mathematician, programming language researcher, and software architect, he leads the company's research initiatives, software architecture, developer tooling, and compiler strategy. Prior to Positron, Kmett worked alongside Sohmers at Groq, helping shape software systems for large-scale AI infrastructure.
SIZE OF COMPANY	~90 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	In July 2025, Positron announced an oversubscribed $51.6 million Series A financing to support continued Atlas deployments and development of its next-generation hardware platform. In February 2026, Positron announced an oversubscribed $230 million Series B financing at a valuation exceeding $1 billion to accelerate development of its custom Asimov silicon and expand commercial deployment of Atlas systems. Management stated that the financing would support the company's roadmap from shipping Atlas today to tape-out of Asimov in late 2026 and production in 2027. The company also disclosed plans for its next-generation Titan platform, powered by four Asimov accelerators and designed to support up to 16 trillion parameters, 10 million+ token context windows, and more than 8 TB of memory per system, targeting frontier AI models and memory-intensive inference workloads.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Positron has raised $313MM in total funding. In its latest Series B round, the company raised $234MM in May 2026. Notable Investors: Valor Equity Partners, Atreides Management, DFJ Growth, Flume Ventures, Resilience Reserve, 1517 Fund, Unless, ARENA Private Wealth, Jump Trading, Qatar Investment Authority (QIA), Arm Holdings, and Helena.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Rain AI

COMPANY NAME	Rain AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / Compute-in-Memory AI Hardware
SIZE OF TOTAL ADDRESSABLE MARKET	Rain AI develops hardware intended to improve the energy efficiency of AI computing. The company combines compute-in-memory architecture, hardware-algorithm co-design, and machine-learning optimization tools for AI training and inference.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** CBRS **PRIVATE:** SambaNova Systems, Mythic, Groq, d-Matrix
KEY SUPPLIERS / CUSTOMERS	Andes Technology is identified as Rain AI's RISC-V partner. Named suppliers and customers are N/A.
COMPETITIVE POSITIONING	Rain AI's core differentiator is its Digital In-Memory Computing (D-IMC) architecture, which combines compute and memory with hardware-algorithm co-design to improve the efficiency of AI training and inference. Key attributes include: • Compute-in-memory architecture: Rain AI's proprietary D-IMC cores perform computation closer to where data is stored, reducing the data movement and storage inefficiencies associated with conventional computing architectures. • Hardware-algorithm co-design: The company co-designs its D-IMC hardware with numerical formats, quantization techniques, and training algorithms to optimize AI performance across the stack. • Machine-learning optimization: Rain AI develops proprietary quantization and fine-tuning techniques designed to improve computational efficiency while maintaining model accuracy. • Training and inference support: Unlike many compute-in-memory architectures focused primarily on inference, Rain AI's D-IMC cores are designed to support both AI training and inference workloads. • Patent portfolio: The company's intellectual-property portfolio includes memristive arrays, efficient matrix multiplication, AI accelerator architectures, and flexible compute-engine microarchitectures.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://rain.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Jack Kendall – Co-Founder, Chief Executive Officer & Board Member:** Prior to co-founding Rain AI (formerly Rain Neuromorphics) in 2017, Jack Kendall conducted interdisciplinary research spanning physics, neuroscience, materials science, and computing at the University of Florida. He is an inventor whose early work focused on neuromorphic computing architectures and nanowire-based neural networks designed to emulate biological brain function. At Rain AI, Mr. Kendall has led the company's evolution from early neuromorphic hardware concepts toward advanced AI accelerators designed to deliver significantly improved energy efficiency for AI training and inference workloads. He holds a bachelor's degree in Chemical Engineering and Physics from the University of Florida. **Juan Nino, Ph.D. – Co-Founder, Chief Scientific Advisor & Board Member:** Prior to co-founding Rain AI, Juan Claudio Nino served as a Professor of Materials Science and Engineering at the University of Florida, where his research focused on advanced functional materials, neuromorphic hardware, brain-inspired computing, semiconductors, and artificial intelligence applications. Dr. Nino's research has spanned electroceramic materials, energy systems, electronics, and AI-driven materials characterization. At Rain AI, he helped develop the company's foundational neuromorphic computing technology and continues to serve as Chief Scientific Advisor. He holds a Ph.D. in Materials Engineering from Penn State University and a B.S. in Mechanical Engineering from Universidad de Los Andes.
SIZE OF COMPANY	~45 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Rain AI is generating revenue and has raised $67.04 million. The company completed a $25 million Series A in May 2024 and subsequently added senior architecture talent.

ANY MAJOR M&A / CORPORATE ACTIONS	A proposed $150 million Series B was canceled. The company also completed private secondary transactions in which several investors sold stakes to undisclosed buyers, including a June 2023 transaction involving Aramco Ventures per Pitchbook.
INVESTORS / INVESTMENT STAGE	Rain AI's latest completed priced round was a $25 million Series A led by Prosperity7 Ventures and Grep VC in May 2024. Other participants included Explorer Equity Group, Gaingels, Metaplanet Holdings, Vannevar Ventures, Buckley Ventures, Pioneer Fund, Deepwater Asset Management, and individual investors. Total capital raised was $67.04 million.
LATEST ROUND VALUATION	$101 million pre-money and $126 million post-money in May 2024.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Sagence AI

COMPANY NAME	Sagence AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference
SIZE OF TOTAL ADDRESSABLE MARKET	Sagence AI develops analog in-memory computing hardware and software designed to accelerate AI inference workloads at significantly lower power and cost than traditional GPU-based architectures. The company's technology targets one of the largest opportunities in semiconductor infrastructure: enabling economically viable deployment of AI across datacenters, enterprise infrastructure, edge devices, and embedded systems. Management positions the company around the growing challenge of AI power consumption and inference economics. As generative AI adoption accelerates, inference workloads are expected to become an increasingly large portion of total AI compute demand, creating opportunity for specialized architectures that can improve performance-per-watt and reduce total cost of ownership relative to conventional GPU deployments.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** NVDA, AMD, INTC **PRIVATE:** d-Matrix, Encharge AI, Mythic, SambaNova, Positron Sagence competes broadly across the AI inference acceleration market, particularly against companies developing custom architectures designed to improve AI performance and efficiency relative to GPUs. The company's closest architectural peers are analog and in-memory computing startups such as EnCharge AI and Mythic.
KEY SUPPLIERS / CUSTOMERS	Sagence has not publicly disclosed manufacturing partners or commercial customers. Management has indicated that its solutions target a broad range of AI inference use cases including content generation, personalized recommendation systems, real-time anomaly detection, video analytics, enterprise AI systems, datacenter infrastructure, and edge AI deployments.
COMPETITIVE POSITIONING	Sagence AI's core differentiator is its analog in-memory computing architecture, which seeks to address the power and efficiency limitations of conventional digital processors. • Analog in-memory computing architecture: Unlike GPUs and most AI accelerators that rely on digital computation, Sagence performs computations directly within memory arrays using analog techniques. Management believes this significantly reduces energy consumption associated with moving data between memory and compute resources. • Elimination of the memory bottleneck: Traditional AI accelerators frequently move data between processors and memory, creating latency and power inefficiencies. Sagence's in-memory architecture colocates storage and computation, reducing data movement and improving efficiency. • Inference-focused optimization: The company is specifically focused on inference workloads rather than AI model training. Management believes inference will become the largest long-term compute market as AI applications scale across enterprises and consumer devices. • Power-efficient deployment model: Sagence states that its architecture is designed to enable datacenter-grade AI workloads at substantially lower power consumption than conventional accelerator platforms, potentially improving AI deployment economics. • Scalable hardware and software platform: In addition to semiconductor products, the company develops software tools and programming frameworks designed to deploy AI models across Sagence hardware platforms. • Complementary rather than replacement architecture: Management has stated that analog accelerators are intended to complement digital processors and GPUs in many deployments, accelerating specific AI workloads where efficiency is paramount.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://www.sagence-ai.com

MANAGEMENT TEAM / FOUNDER(S)	**Vishal Sarin - Founder and CEO:** Vishal Sarin founded the company in 2018 after a long career in semiconductor design, focusing on analog circuits, memory technologies, and AI hardware architectures. According to public interviews, Sarin identified growing power and economics challenges associated with GPU-centric AI computing and launched the company to develop analog in-memory alternatives. The company was originally established as Analog Inference before becoming Sagence AI.
SIZE OF COMPANY	45-75
GROWTH TRAJECTORY	Sagence AI was founded in 2018 to commercialize analog in-memory computing technology for AI inference applications. Its original focus was developing semiconductor intellectual property capable of overcoming the performance and economics limitations associated with conventional digital AI architectures. The company operated in relative stealth for several years while developing its technology platform. In 2021, the company raised funding from investors including TDK Ventures and other venture firms to support product development and commercialization efforts. In May 2023, Sagence announced a roughly $22 million financing round with participation from multiple strategic and venture investors. The funding was used to continue development of its analog in-memory inference platform. The company emerged from stealth publicly in late 2024 and disclosed details regarding its analog AI inference architecture, software stack, and product roadmap. Management stated that commercial products would target both datacenter and edge AI deployments.
ANY MAJOR M&A / CORPORATE ACTIONS	The company was originally established as Analog Inference before rebranding as Sagence AI.
INVESTORS / INVESTMENT STAGE	Sagence has raised approximately $50MM+ of disclosed funding across multiple venture rounds. Notable investors include TDK Ventures, Khosla Ventures, Aramco Ventures, Prosperity7 Ventures, Alumni Ventures, MFV Partners, Cambium Capital, Green Sands Equity.
LATEST ROUND VALUATION	Sagence completed its latest publicly disclosed financing round in 2023, raising approximately $22 million.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For AI accelerator, analog computing, and inference semiconductor companies that remain in development or early commercialization stages, EV/Sales and P/Sales methodologies are generally the most appropriate valuation frameworks given limited profitability and significant R&D investment requirements. As companies achieve scale and production deployment, investor focus typically transitions toward EV/EBITDA, P/E, and P/FCF metrics, similar to mature semiconductor and infrastructure providers.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: SambaNova Systems

COMPANY NAME	SambaNova Systems, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Fast Inference
SIZE OF TOTAL ADDRESSABLE MARKET	SambaNova Systems develops integrated AI computing platforms that combine custom AI accelerators, systems software, foundation models, and cloud infrastructure for enterprise, government, and sovereign AI deployments. The company's proprietary Reconfigurable Dataflow Architecture (RDA) and Reconfigurable Dataflow Units (RDUs) are designed to accelerate AI inference workloads more efficiently than traditional CPU- and GPU-based systems. SambaNova focuses on enterprise AI, sovereign AI, hyperscale inference, and agentic AI deployments, positioning itself as an alternative to GPU-centric infrastructure. The AI inference market TAM is expected to exceed $600B by 2033.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** AMD, CBRS, INTC, NVDA, MU **PRIVATE:**Tenstorrent, Groq, d-Matrix, Hailo Relative to many competitors, SambaNova differentiates itself through its dataflow architecture, memory-centric design, air-cooled rack-scale deployments, and integrated cloud and software platform.
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry). Customers: JPMorgan Chase, Argonne National Laboratory, Oak Ridge National Laboratory, Texas Advanced Computing Center, OTP Bank, Saudi Aramco, SoftBank
COMPETITIVE POSITIONING	SambaNova's platform is centered on its Reconfigurable Dataflow Architecture (RDA), which combines custom silicon, systems software, cloud services, and rack-scale infrastructure. Key differentiators include: • SN50 next-generation accelerator: In February 2026, SambaNova introduced the SN50 RDU, designed specifically for agentic AI and large-scale inference workloads. The company states that SN50 delivers up to 5x higher performance than competing solutions and 3x lower total cost of ownership than GPU-based inference platforms. • Memory-centric architecture: The SN50 combines 432MB of on-chip SRAM, 64GB of HBM, and a CXL-attached DDR memory tier, creating a three-tier memory hierarchy designed to improve memory-bandwidth utilization for increasingly memory-bound inference workloads. • Massive scalability: SambaNova states that SN50 can scale across 256 accelerators within a scale-up domain using a multi-terabyte-per-second Ethernet interconnect while supporting models up to 10 trillion parameters. • Air-cooled rack-scale systems: Unlike many competing AI infrastructure platforms that require liquid cooling, SambaNova's SN50 rack-scale systems are designed to operate using conventional air-cooled data center infrastructure. Management believes this reduces deployment complexity and retrofit costs for enterprise customers. • Full-stack AI platform: In addition to custom accelerators, SambaNova offers DataScale systems, SambaCloud, and SambaStack software, enabling customers to deploy AI workloads on-premise, in private clouds, or through managed cloud services.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	http://www.sambanova.ai

MANAGEMENT TEAM / FOUNDER(S)	**Rodrigo Liang – Co-founder and CEO:** Prior to co-founding SambaNova Systems in 2017, Mr. Liang spent more than two decades developing high-performance processors at Sun Microsystems and Oracle Corporation, where he served as Senior Vice President of SPARC Processor and ASIC Development. During his career, he led development of multiple generations of enterprise processors and ASICs. Mr. Liang holds both a B.S. and M.S. in Electrical Engineering from Stanford University. **Kunle Olukotun – Co-founder and Chief Technologist:** Mr. Olukotun is the Cadence Design Systems Professor of Electrical Engineering and Computer Science at Stanford University and is widely recognized as a pioneer of modern multi-core processor architectures. Prior to SambaNova, he founded Afara Websystems, which was acquired by Sun Microsystems in 2002, and has led numerous research initiatives in parallel computing and artificial intelligence. He continues to direct Stanford's Pervasive Parallelism Lab while serving as SambaNova's Chief Technologist.
SIZE OF COMPANY	~400 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	SambaNova launched SambaNova Cloud in 2024, providing developers and enterprises with access to large foundation models powered by its SN40L architecture. In February 2026, the company introduced its next-generation SN50 accelerator, announced a multi-year strategic collaboration with Intel, and raised more than $350 million in Series E financing to expand manufacturing capacity and scale its AI cloud infrastructure. The company also disclosed that SoftBank Corp. would become the first customer to deploy SN50 systems within next-generation AI data centers in Japan. More recently, JPMorgan Chase selected SN40L and SN50 systems for secure on-premises AI inference deployments. In July 2026, SambaNova completed the first close of a $1 billion Series F financing at an $11 billion valuation while continuing to expand deployments across enterprise, sovereign AI, neo-cloud, and cloud service provider customers worldwide.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	SambaNova has raised $2.5B over multiple funding rounds. In its latest Series F financing, the company raised $1B in July 2026, following a $350MM Series E round completed earlier in March 2026. Notable Investors: General Atlantic, Seligman Ventures, T. Rowe Price, SoftBank Vision Fund, Intel Capital, BlackRock, Temasek Holdings, GIC, Walden International, GV (Google Ventures), Qatar Investment Authority (QIA), Vista Equity Partners, Battery Ventures, Celesta Capital, Atlantic Bridge, and Cambium Capital.
LATEST ROUND VALUATION	Latest reported valuation: $11B following the July 2026 Series F financing.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: SiFive

COMPANY NAME	SiFive
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / RISC-V CPU and Semiconductor IP
SIZE OF TOTAL ADDRESSABLE MARKET	SiFive is a leading provider of RISC-V CPU IP – it licenses customizable processor core designs based on the open-source RISC-V instruction set. Its business rides the trend of companies adopting RISC-V to build their own chips without proprietary ISAs.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** ARM, CDNS, SNPS **PRIVATE:** Imagination Technologies SiFive competes in the processor IP market, where Arm remains the dominant provider of CPU architectures across mobile, embedded, automotive, and increasingly data-center applications. SiFive and other RISC-V vendors are seeking to capitalize on growing adoption of open-standard processor architectures that provide customers with greater customization flexibility and reduced licensing costs relative to proprietary instruction set architectures (ISAs). The company also competes indirectly against x86 providers such as AMD and Intel in data-center and high-performance compute markets, where customers are increasingly evaluating RISC-V-based alternatives for AI, cloud, and edge-computing workloads.
KEY SUPPLIERS / CUSTOMERS	Suppliers include INTC (foundry). Customers include NASA, MCHP, and Arkmicro.
COMPETITIVE POSITIONING	SiFive utilizes the RISC-V architecture to deliver customizable, high-performance processor IP across embedded, edge, mobile, automotive, and data-center applications. Key differentiators include: • SiFive Essential: A family of foundational RISC-V processor cores designed for embedded and microcontroller applications, giving customers flexibility to customize processor configurations. • SiFive Intelligence: AI-focused processor IP designed to accelerate machine-learning workloads through scalar, vector, and matrix computing capabilities. In September 2025, SiFive introduced its 2nd Generation Intelligence family, including the X160 Gen 2, X180 Gen 2, X280 Gen 2, X390 Gen 2, and XM Gen 2 products targeting AI workloads spanning edge devices through data-center infrastructure. • SiFive Performance: High-performance RISC-V processor IP designed for cloud, enterprise, and data-center applications, providing energy-efficient compute density and scalability. • SiFive Automotive: A family of automotive-grade processors supporting ADAS, powertrain, zonal architecture, and central-compute applications. • Broad processor portfolio: SiFive's portfolio spans low-power embedded MCUs, vector processors, accelerator control units (ACUs), AI accelerators, and high-performance application processors targeting workloads from IoT devices to large-scale AI inference and cloud infrastructure.
YEAR FOUNDED / ESTABLISHED	2015
WEBSITE	http://www.sifive.com

MANAGEMENT TEAM / FOUNDER(S)	**Patrick Little – Chairman, President & CEO:** Prior to joining SiFive, Mr. Little served as Senior Vice President and General Manager at Xilinx before becoming Senior Vice President and General Manager at Qualcomm. He holds an MBA from MIT. **Krste Asanović – Co-Founder & Chief Architect:** Mr. Asanović is Professor Emeritus and Professor of the Graduate School in the EECS Department at UC Berkeley. He co-created the RISC-V instruction set architecture and has been one of the leading academic figures behind open-standard processor development. He holds a Ph.D. from UC Berkeley and a bachelor's degree from the University of Cambridge. **Yunsup Lee – Co-Founder & CTO:** Mr. Lee received his Ph.D. in Computer Science from UC Berkeley, where he co-designed the RISC-V ISA and helped develop the first RISC-V microprocessors. **Andrew Waterman – Co-Founder & Chief Engineer:** Mr. Waterman co-designed the RISC-V ISA and the first RISC-V processors alongside Mr. Lee at UC Berkeley. He received his Ph.D. in Computer Science from UC Berkeley.
SIZE OF COMPANY	~517 employees on LinkedIn and in press reports (accessed 8/19/2026)
GROWTH TRAJECTORY	In May 2025, SiFive announced a collaboration with Red Hat to support Red Hat Enterprise Linux 10 on RISC-V platforms, further expanding the software ecosystem surrounding enterprise and cloud deployments. In September 2025, the company launched its 2nd Generation Intelligence family, introducing five new RISC-V AI processor IP products spanning edge, industrial, automotive, and data-center workloads. Management also disclosed adoption of the new X100 series by multiple Tier 1 semiconductor companies and previously highlighted expectations for a hyperscaler deployment of a SiFive-based SoC. The company continues to expand its RISC-V processor portfolio as adoption of open-standard architectures increases across AI, cloud, automotive, and enterprise computing markets.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	SiFive raised $765MM over 11 funding rounds. In its latest Series G round of funding, the company received $400MM in April of 2026. Notable Investors: Qualcomm Ventures, Intel Capital, Sutter Hill Ventures, Spark Capital, Osage University Partners, Huami Corporation, Prosperity7 Ventures (Aramco Ventures), Coatue Management, Moore Strategic Ventures, Temasek Holdings, SK hynix.
LATEST ROUND VALUATION	$3.25 billion pre-money and $3.65 billion post-money in the April 2026 Series G.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Snowcap

COMPANY NAME	Snowcap
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / Superconducting Computing
SIZE OF TOTAL ADDRESSABLE MARKET	Snowcap develops a superconducting computing platform intended to deliver high processing performance per watt. The platform uses superconducting technology, standard fabrication facilities, and proven cryogenic systems to support sustainable data-center scaling.
GEOGRAPHIC PRESENCE	Palo Alto, California.
NOTABLE COMPETITORS	**LISTED:** QBTS, RGTI, IONQ, INFQ **PRIVATE:** Oxford Quantum Circuits
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	Snowcap's core differentiator is its superconducting computing architecture, which replaces conventional semiconductor transistors with superconducting Josephson junctions to deliver substantially higher performance and energy efficiency. Key attributes include: • Superconducting architecture: Snowcap uses superconducting Josephson junctions in place of conventional semiconductor transistors, targeting significantly higher performance per watt for compute-intensive workloads. • Performance and energy efficiency: The architecture is designed to deliver orders-of-magnitude improvements in processing speed and efficiency relative to conventional CMOS-based computing. • Scalable manufacturing and cooling: The platform is designed to use existing semiconductor fabrication processes and high-volume fabs, while leveraging helium-based cryogenic infrastructure already used in quantum computing. • AI, HPC, and quantum applications: Snowcap targets AI, high-performance computing, and quantum-classical workloads, including next-generation quantum and low-temperature computing systems. • Speed and low-latency computing: The platform is engineered around high computational speed, energy efficiency, and low latency, positioning superconducting logic as an alternative pathway beyond conventional CMOS scaling.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://snowcapcompute.com
MANAGEMENT TEAM / FOUNDER(S)	**Michael Lafferty – Chief Executive Officer:** Prior to joining Snowcap Compute as Chief Executive Officer, Michael Lafferty spent more than 25 years in semiconductor physical design and advanced technology development. Most recently, he served as Director of Cadence Design Systems' "More than Moore" engineering group, where his work focused on pioneering superconducting and quantum computing technologies. At Snowcap, he is leading the commercialization of superconducting computing systems designed to deliver substantial improvements in performance and energy efficiency for artificial intelligence, quantum computing, and high-performance computing workloads. He holds a degree from Lehigh University. **Rick Hyman – Co-Founder & Board Member:** Prior to co-founding Snowcap Compute, Rick Hyman held executive leadership positions across the semiconductor and EDA industries. His career includes roles at NVIDIA, where he served as Vice President of Business Development; Synopsys; LSI Logic; National Semiconductor; and Silicon Spice, which was later acquired by Broadcom. He also served as CEO of Teraoptic Networks. At Snowcap, Mr. Hyman serves on the Board of Directors and contributes to the company's commercialization and industry partnership strategy. **Cameron Rogers – Co-Founder:** Prior to co-founding Snowcap Compute, Cameron Rogers spent more than 17 years at Apple in a variety of hardware and software product leadership roles. His work included product management responsibilities for Mac technologies, Apple Silicon initiatives, AR/VR technologies, artificial intelligence projects, and graphics software. He later served as Director of Advanced EV Development Software at Ford Motor Company. At Snowcap, he is helping develop next-generation superconducting computing platforms intended to serve future AI and data-center workloads.

SIZE OF COMPANY	~24 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Snowcap was founded in 2024 and is in product development. Its seed funding is intended to support next-generation quantum and low-temperature compute systems focused on speed, energy efficiency, and low latency.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Snowcap raised $23 million of seed funding led by Playground Global in June 2025. Cambium Capital Management, VSquared Ventures, Page One Ventures, and Osage University Partners also participated.
LATEST ROUND VALUATION	$40 million pre-money and $63 million post-money in June 2025.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	Latest preferred-equity financing; comparable private-company financing valuations.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Tachyum

COMPANY NAME	Tachyum
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Inference
SIZE OF TOTAL ADDRESSABLE MARKET	Tachyum develops processors, AI accelerators, and system architectures for artificial intelligence, high-performance computing (HPC), cloud computing, and hyperscale data centers. The company's flagship product, the Prodigy Universal Processor, is designed to unify CPU, GPU, and AI accelerator functionality on a single architecture, eliminating the need for separate processor types across AI, HPC, and enterprise workloads. Management believes a unified processor architecture can reduce infrastructure complexity, improve utilization, lower power consumption, and reduce total cost of ownership for large-scale AI deployments. Tachyum targets AI infrastructure, hyperscale cloud providers, telecommunications networks, sovereign AI deployments, and supercomputing centers. The company believes demand for next-generation AI infrastructure is creating an opportunity for alternative architectures that can deliver higher performance-per-watt and lower deployment costs than traditional CPU- and GPU-based systems with a TAM expected to exceed $600B by 2033.
GEOGRAPHIC PRESENCE	Las Vegas, United States
NOTABLE COMPETITORS	**LISTED:** AMD, CBRS, INTC, NVDA **PRIVATE:** Groq, SambaNova, Tenstorrent Relative to competitors, Tachyum differentiates itself through its Universal Processor strategy, which seeks to replace heterogeneous CPU- and GPU-based infrastructure with a single architecture optimized for AI, HPC, and cloud workloads.
KEY SUPPLIERS / CUSTOMERS	Manufacturer: Planned TSMC production partners Customers: Undisclosed; targeting hyperscale AI data centers, cloud providers, and supercomputing centers.
COMPETITIVE POSITIONING	Tachyum's platform is centered on its Prodigy Universal Processor architecture. Key differentiators include: • Universal Processor architecture: Prodigy is designed to integrate CPU, GPU, and AI accelerator functionality into a single processor architecture, enabling a common platform for AI training, AI inference, HPC, and cloud computing workloads. • High-core-count design: Tachyum disclosed in 2025 that each Prodigy compute chiplet is designed to contain 256 custom 64-bit cores, increasing from prior product specifications as the company optimized the architecture for AI and HPC workloads. • Unified AI and HPC platform: Management believes Prodigy can reduce the need for heterogeneous CPU-GPU architectures, enabling customers to consolidate workloads onto a single platform while improving utilization and reducing operating costs. • AI, cloud, and HPC focus: The architecture is designed to support large-scale AI training, inference, scientific computing, cloud services, and enterprise workloads using a common software and hardware stack.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	http://www.tachyum.com/
MANAGEMENT TEAM / FOUNDER(S)	**Radoslav Danilak – Co-founder and CEO:** Prior to founding Tachyum, Mr. Danilak founded and served as CEO of Skyera, which was acquired by HGST/Western Digital. Earlier, he co-founded and served as CTO of SandForce, which was acquired by LSI Corporation in 2011. He also previously worked as a Senior Architect at NVIDIA. **Igor Shevlyakov – Co-founder and VP of Software Engineering:** Prior to Tachyum, Mr. Shevlyakov held engineering roles at HGST and Skyera and spent approximately 13 years as a Compiler Engineer at Microunity. He leads development of Tachyum's software stack, compilers, and ecosystem tools.
SIZE OF COMPANY	~38 Employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	Tachyum continued advancing development of its Prodigy Universal Processor throughout 2025 and 2026 while expanding commercialization efforts. In October 2025, the company announced a binding $220 million Series C financing and disclosed a $500 million purchase order for Prodigy processors from the same investor. Management stated that the financing would support completion of Prodigy's tape-out process and preparation for commercial deployment. In June 2026, Tachyum reaffirmed that Prodigy development remained active, with engineering teams continuing software ecosystem development, product validation, and commercialization activities ahead of expected tape-out milestones. More recently, in August 2026, the company announced that the Prodigy ISA and platform would be open sourced, enabling developers and partners to build, port, and optimize software for the architecture.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Tachyum has raised $245MM across multiple funding rounds. In its latest Series C financing, the company announced a $220MM investment in October 2025, representing its largest funding round to date. Notable Investors: IPM Group, Across Private Investments, Government of the Slovak Republic
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Tensordyne

COMPANY NAME	Tensordyne (formerly Recogni)
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / Application-Specific Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Tensordyne develops high-performance, energy-efficient AI inference infrastructure for hyperscalers, cloud providers, and enterprises. Its portfolio includes custom AI chips, inference systems, interconnect technologies, proprietary mathematical architectures, and software development tools.
GEOGRAPHIC PRESENCE	Sunnyvale, California
NOTABLE COMPETITORS	**LISTED:** CBRS **PRIVATE:** Graphcore, Mythic, Groq, SambaNova Systems, d-Matrix Relative to many competitors, Tensordyne differentiates itself through its logarithmic mathematics approach, tightly integrated memory architecture, air-cooled rack-scale deployment model, and focus exclusively on AI inference workloads.
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A. Customers: N/A.
COMPETITIVE POSITIONING	Tensordyne's core differentiator is its logarithmic-math-based AI inference architecture, which replaces conventional multiplication-heavy computation with addition-based processing and co-optimizes compute, memory, and interconnect for large-scale inference. Key attributes include: • Logarithmic AI computing: Tensordyne's proprietary approach uses logarithmic mathematics to replace large-scale multiplication operations with addition-based computation, targeting higher compute density and performance per watt for AI inference. • Purpose-built inference processor: The Napier processor is designed specifically for AI inference, with Tensordyne focusing its platform on large-scale data-center inference rather than general-purpose acceleration. • Integrated memory architecture: Each Napier processor combines fast SRAM with HBM, with the architecture designed to reduce idle compute cycles and balance compute capacity with model-memory requirements. • High-speed scale-up interconnect: Tensordyne's proprietary TDN Link connects processors through a scale-up fabric designed for sub-microsecond communication latency, addressing interconnect bottlenecks at system scale. • Rack-scale system co-design: Tensordyne develops Napier as a full inference system spanning math, processors, memory, networking, and software rather than optimizing the accelerator chip in isolation. • MoE and agentic AI optimization: The company has optimized its platform around Mixture-of-Experts and agentic workloads and positions Napier for both standard and disaggregated inference configurations.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://tensordyne.ai
MANAGEMENT TEAM / FOUNDER(S)	**Marc Bolitho – CEO:** Prior to joining Tensordyne in July 2022, Mr. Bolitho spent 28 years at TRW and ZF in engineering and business leadership roles, most recently serving as Senior Vice President and General Manager of ZF's Electronics and Advanced Driver Assistance Systems business. **R. K. Anand – Co-founder and Chief Product Officer:** Mr. Anand founded Tensordyne, formerly Recogni, in 2017 and initially served as CEO, leading the company through its seed and Series B financings and the architecture, tape-out, and bring-up of its first chip. Prior to Tensordyne, Mr. Anand served as President and CEO of OttoQ, CEO of Kumu Networks, and Executive Vice President and General Manager of Juniper Networks' Data Center Business Unit.
SIZE OF COMPANY	~114 on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	In 2024, Tensordyne strategically shifted its focus toward data-center-scale AI inference and established a technology partnership with Juniper Networks. During 2025, the company optimized its architecture for Mixture-of-Experts (MoE) and agentic AI workloads. In June 2026, Tensordyne announced successful tape-out of its 3nm Napier processor, developed in partnership with Broadcom and manufactured by TSMC. The company also disclosed a sales pipeline exceeding $200 million, more than a dozen Letters of Intent tied to system evaluations, and ongoing customer engagements spanning hyperscalers, neo-cloud providers, and sovereign AI infrastructure operators as it advances toward benchmarking, beta deployment, and broader commercial rollout of the Napier platform.
ANY MAJOR M&A / CORPORATE ACTIONS	Formerly known as Recogni and subsequently renamed Tensordyne.
INVESTORS / INVESTMENT STAGE	Tensordyne has raised $208.9MM across multiple funding rounds. In its latest Series C1 financing, the company raised $35.0MM in July 2025 through a combination of $30.4MM of equity and $4.6MM of convertible debt. Notable Investors: Xora, Xora Innovation, Karmel Capital, Maiora Ventures, Premji Invest, Juniper, Celesta Capital, DNS Capital, GreatPoint Ventures, Mayfield Fund, BMW i Ventures, SAIC Capital, HSBC Holdings, Bosch Ventures, Toyota Ventures, and Volvo Cars Technology Fund.
LATEST ROUND VALUATION	The valuation for the July 2025 Series C1 was not disclosed. The latest disclosed transaction valuation was $510.0MM post-money in the February 2024 Series C.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Tenstorrent

COMPANY NAME	Tenstorrent
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / Application-Specific Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Tenstorrent develops AI processor architectures and open-source software for deep-learning workloads. Its portfolio includes RISC-V CPU cores, AI chiplets, neural-network processors, compute systems, and open-source compiler software.
GEOGRAPHIC PRESENCE	Toronto, Ontario
NOTABLE COMPETITORS	**LISTED:** NVDA, CBRS **PRIVATE:** Graphcore, SambaNova Systems, Mythic, Groq
KEY SUPPLIERS / CUSTOMERS	Suppliers: Samsung foundries. Customers: LG.
COMPETITIVE POSITIONING	Tenstorrent's core differentiator is its programmable Tensix architecture, which combines RISC-V control processors, specialized AI compute engines, and an open software stack to support scalable AI training and inference. Key attributes include: • Tensix compute architecture: Tenstorrent's processors are built around programmable Tensix cores that combine multiple RISC-V processors with matrix and vector compute engines and local SRAM, enabling software-controlled orchestration of compute and data movement. • Distributed memory architecture: Tensix cores integrate local SRAM and communicate through an on-chip network, allowing computation and data movement to be distributed across a many-core architecture rather than relying solely on centralized compute resources. • Scale-out connectivity: Tenstorrent's AI accelerators are designed to scale from individual accelerator cards to multi-chip systems, with products including PCIe cards and enterprise-class Galaxy servers. • Open software ecosystem: The company provides a software stack spanning the TT-Forge compiler, TT-NN high-level neural-network library, TT-MLIR compiler infrastructure, and low-level TT-Metalium SDK, supporting frameworks and model formats including PyTorch, JAX, and ONNX. • Training and inference support: Tenstorrent's platform supports both AI inference and training, with software resources covering LLM serving, fine-tuning, multi-device training, and training from scratch.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://tenstorrent.com
MANAGEMENT TEAM / FOUNDER(S)	**Jim Keller – CEO:** Mr. Keller has served as CEO of Tenstorrent since January 2023, after joining the company as Chief Technology Officer in December 2020. Prior to Tenstorrent, he spent two years as Senior Vice President of Intel's Silicon Engineering Group. Over a career spanning four decades, Mr. Keller has held senior engineering and architecture leadership roles at Tesla (Vice President of Autopilot Hardware), AMD (Corporate Vice President and Chief Cores Architect, where he led development of the Zen and K12 architectures), Apple (leading development of the A4 and A5 processors following Apple's acquisition of P.A. Semi), P.A. Semi (Vice President of Engineering and Chief Architect), Broadcom/SiByte, and Digital Equipment Corporation. He is widely recognized for his contributions to the AMD K7/K8 architectures, x86-64, HyperTransport, Apple's A-series processors, Tesla's self-driving chip, and other landmark semiconductor designs.
SIZE OF COMPANY	~1,306 employees on LinkedIn (accessed 9/30/2026)

GROWTH TRAJECTORY	In 2021, Tenstorrent raised capital to continue developing AI-focused hardware for developers. In 2023, the company accelerated the development of its AI chipsets and machine-learning software roadmap while shifting its business model from direct product sales toward licensing. The company began securing significant licensing agreements, including with LG. In 2024, Tenstorrent expanded its open-source AI software stack, hired additional developers, enlarged its global development and design-center footprint, and invested in AI systems and cloud infrastructure. In July 2025, the company acquired Blue Cheetah Analog Design, adding general-purpose semiconductor capabilities. In June 2026, Tenstorrent launched TT-Ascalon S, reported new performance records, and expanded its presence in Japan.
ANY MAJOR M&A / CORPORATE ACTIONS	Tenstorrent acquired Blue Cheetah Analog Design in July 2025; the transaction value was not disclosed.
INVESTORS / INVESTMENT STAGE	Tenstorrent has raised $2.32BN across multiple funding rounds. In its latest later-stage VC financing, the company raised $1.29BN from undisclosed investors in September 2026, representing its largest funding round to date. Notable Investors: Samsung Catalyst Fund, Samsung Venture Investment, Samsung Securities, LG Electronics, Export Development Canada, Baillie Gifford, Fidelity Management & Research Company, Hyundai Motor Group, Bezos Expeditions, Ontario Pension Board, Eclipse Capital, Real Ventures, Optimist Ventures, Golden Falcon Capital, GiantLeap Capital, Plug and Play, and SM Ventures.
LATEST ROUND VALUATION	The valuation for the September 2026 financing was not disclosed. The latest disclosed transaction valuation was an estimated $2.60BN post-money in the December 2024 Series D1 per Pitchbook.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Unconventional AI

COMPANY NAME	Unconventional AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Inference / Alternative AI Computing Architecture
SIZE OF TOTAL ADDRESSABLE MARKET	Unconventional AI develops a computing architecture intended to improve the efficiency and scalability of AI processing. The architecture is designed to bypass limitations associated with traditional von Neumann systems, reduce data bottlenecks, and improve energy efficiency.
GEOGRAPHIC PRESENCE	San Diego, California
NOTABLE COMPETITORS	**LISTED:** CBRS **PRIVATE:**SambaNova Systems, Groq, Ayar Labs, Etched, Mythic, Tenstorrent
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	Unconventional AI's core differentiator is its physics-based computing architecture, which uses physical dynamical systems rather than conventional digital computation to target substantially greater energy efficiency for AI workloads. Key attributes include: • Physics-based computing architecture: Unconventional AI is developing a new computing paradigm that harnesses physical dynamics as the computational substrate, seeking an alternative to conventional digital architectures for executing AI workloads. • Reduced data-movement bottlenecks: The company's research identifies memory and data movement as key constraints on AI energy efficiency and is exploring alternative computing and memory architectures designed to address these limitations. • Energy-efficient AI processing: Unconventional AI's architecture is designed around substantially reducing the energy required to run modern AI, with the company targeting approximately 1,000x lower energy consumption than today's machines. • Oscillator-based computation: The company's Un-0 model demonstrates its approach through a simulated system of coupled oscillators, using the dynamics of the physical system itself to perform computation for an image-generation workload. • Hardware-software co-design: Unconventional AI's research spans physical computing substrates, neural architectures, memory systems, system modeling, and an instruction-set architecture for dynamical-system hardware, reflecting a full-stack approach to developing unconventional AI compute.
YEAR FOUNDED / ESTABLISHED	2025
WEBSITE	https://unconv.ai/

MANAGEMENT TEAM / FOUNDER(S)	**Naveen Rao, Ph.D. – Co-Founder & Chief Executive Officer:** Prior to co-founding Unconventional AI in 2025, Naveen Rao served as Vice President of AI at Databricks following the company's acquisition of MosaicML, which he co-founded and led as CEO. Earlier, he co-founded Nervana Systems, an AI processor company acquired by Intel, where he subsequently served as Corporate Vice President and General Manager of the AI Products Group. Dr. Rao's career spans semiconductor architecture, machine learning systems, and computational neuroscience. He holds a Ph.D. in Neuroscience from Brown University and previously studied electrical engineering and computer architecture. At Unconventional AI, he is leading the development of alternative computing architectures designed to dramatically improve AI energy efficiency. **Sara Achour – Co-Founder & Research Fellow:** Prior to co-founding Unconventional AI, Sara Achour joined Stanford University as a jointly appointed Assistant Professor in Computer Science and Electrical Engineering. Her research focuses on programming languages, compilers, runtime systems, and software tools for emerging computing platforms, including analog, quantum, and neuromorphic systems. Dr. Achour earned a Ph.D. in Computer Science from the Massachusetts Institute of Technology in 2021. At Unconventional AI, she works on the software foundations required to program unconventional hardware architectures and emerging physical computing substrates. **MeeLan Lee – Co-Founder & Vice President of Engineering:** Prior to co-founding Unconventional AI, MeeLan Lee held senior chip design and engineering roles across the semiconductor industry, including positions at Google, Qualcomm, Atheros, and Intel. Her career includes leading development of Atheros' highest-volume Wi-Fi chipset, delivering Qualcomm's first flagship product capable of greater than 1 Gbps throughput, and helping deploy Google's Titan security chip into production. She later co-founded Chipletz, focused on advanced packaging technologies. Ms. Lee holds B.S. and M.Eng. degrees from the Massachusetts Institute of Technology and currently leads Unconventional AI's hardware research and engineering efforts.
SIZE OF COMPANY	~57 employees on LinkedIn (accessed 9/18/2026)
GROWTH TRAJECTORY	Unconventional AI was founded in 2025, is generating revenue, and raised $475 million in its first priced financing. The company joined Neo's non-equity accelerator program in October 2025 and introduced the oscillator-based Un-0 model series in June 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Unconventional AI raised $475 million of seed funding led by Andreessen Horowitz and Lightspeed Venture Partners in December 2025. Participants included Databricks, Wing Venture Capital, Eniac Ventures, Mesh Ventures, Decibel Partners, Greycroft, Lux Capital, Future Ventures, DCVC, Sequoia Capital, Felicis, Playground Global, and individual investors.
LATEST ROUND VALUATION	$4.03 billion pre-money and $4.50 billion post-money in December 2025.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	Latest preferred-equity financing; comparable private AI semiconductor financing valuations.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Inference: Velaura

COMPANY NAME	Velaura
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / /Inference / AI Compute and Web Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	Velaura, formerly known as Auradine, develops infrastructure spanning hardware, cloud software, blockchain, and AI. The company is expanding the development and commercialization of an AI compute portfolio alongside engineering and customer-facing operations.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** CSCO, IBM, AVGO, NVDA **PRIVATE:** SiMa.ai, Aurascape, Tenstorrent
KEY SUPPLIERS / CUSTOMERS	Velaura does not disclose suppliers, but states that it has multiple hyperscale engagements and is working with leading hyperscaler XPU partners across 3nm and 2nm process nodes.
COMPETITIVE POSITIONING	Velaura AI's core differentiator is its ultra-low-power silicon design and IP platform, which optimizes AI accelerator designs to reduce power consumption while maintaining functional equivalence with existing architectures. Key attributes include: • Broad infrastructure portfolio: Velaura AI operates across a broader infrastructure technology portfolio while expanding its focus on power-efficient AI compute. • Ultra-low-power silicon design: The company's Titan Core platform applies proprietary low-voltage libraries, circuit technology, and physical-design methodologies to optimize the power efficiency of AI accelerator silicon. • AI accelerator power efficiency: Velaura AI states that Titan Core can enable up to 2x lower overall chip power for AI accelerators, with its technology targeting the energy consumed by arithmetic operations including matrix multiplication. • Integration with existing designs: The platform starts from customers' existing RTL designs and applies Velaura AI's proprietary IP and physical-design methodology, enabling integration into existing SoC architectures and design flows while maintaining functional equivalence. • Patent portfolio: The company's intellectual-property portfolio covers areas including integrated-circuit power control, hardware acceleration, thermal management, security, and approximate computing.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://velaura.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Rajiv Khemani – Co-Founder, Chief Executive Officer, President & Board Member:** Prior to co-founding Velaura AI, Rajiv Khemani was the founder and Chief Executive Officer of Innovium, a cloud and data-center networking silicon company that was acquired by Marvell Technology. Earlier in his career, he held executive and engineering leadership roles at Cavium Networks, Intel, Atheros Communications, and Cisco Systems. Throughout his career, he has focused on high-performance networking, communications infrastructure, and semiconductor platforms. At Velaura, he leads the company's strategy to develop ultra-low-power AI compute infrastructure and silicon IP platforms. **Sanjay Gupta – Co-Founder, Chief Strategy Officer & President, Strategy and GTM:** Prior to co-founding Velaura AI, Sanjay Gupta held senior executive and go-to-market leadership positions across the semiconductor and communications industries. His career has spanned product strategy, corporate development, ecosystem partnerships, and commercialization of advanced infrastructure technologies. At Velaura, he oversees corporate strategy, market development, customer engagement, and go-to-market execution as the company expands commercialization of its AI compute portfolio. **YJ Kim – Co-Founder & President of Product:** Prior to co-founding Velaura AI, Kim held semiconductor product and engineering leadership roles focused on networking, communications, and infrastructure silicon. At Velaura, he is responsible for product strategy and development across the company's AI compute and silicon IP portfolio, including the recently introduced Titan Core platform designed to improve AI accelerator power efficiency.

SIZE OF COMPANY	~67 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Velaura is generating revenue and reported $150 million of revenue for 2025. Total capital raised increased to $452.72 million following its August 2026 financing. The latest proceeds are intended to expand development and commercialization of the AI compute portfolio and grow engineering and customer-facing operations.
ANY MAJOR M&A / CORPORATE ACTIONS	Auradine was renamed Velaura.
INVESTORS / INVESTMENT STAGE	Velaura raised $110 million in an August 2026 financing led by Seligman Ventures. Samsung Catalyst Fund, MARA Holdings, Prosperity7 Ventures, Premji Invest, Maverick Silicon, Mayfield Fund, StepStone Group, and Capricorn Investment Group also participated. Total capital raised was $452.72 million.
LATEST ROUND VALUATION	$890 million pre-money and $1.00 billion post-money in August 2026.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Figure 6: AI Networking Companies & Public Peers

Companies	Public Peers
AvicenaTech	AVGO, COHR, INTC, LITE, MRVL
Ayar Labs	AVGO, COHR, INTC, LITE, MRVL
Baya Systems	CEVA
Cornelis Networks	ANET, AVGO, CRDO, MRVL, NVDA
Eliyan	CDNS, RMBS, SNPS
HyperLight	COHR, LITE, AAOI, CSCO
Lightmatter	AVGO, COHR, INTC, MRVL, NVDA
MangoBoost	NVDA, AMD, CBRS
Menlo Micro	ADI, IFX, QRVO, SWKS
Nexthop AI	NVDA, CBRS
nEye.ai	CIEN, LITE
OpenLight	AVGO, COHR, LITE, MRVL
Point2 Technology	AVGO, MRVL, CSCO, INTC, NVDA
Relativity Networks	MSFT
Retym	CIEN, COHR, LITE
Salience Labs	AVGO, MRVL, CSCO, COHR, LITE
SiFotonics	AVGO, COHR, CSCO, LITE, MRVL
Xscape Photonics	AVGO, COHR, LITE, MRVL, CSCO

Source: UBS

AI Networking: AvicenaTech

COMPANY NAME	AvicenaTech, Corp.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Networking
SIZE OF TOTAL ADDRESSABLE MARKET	The company operates in the rapidly expanding silicon photonics market. The global AI computing market is expected to reach over $380B per year by 2029.
GEOGRAPHIC PRESENCE	Sunnyvale, California
NOTABLE COMPETITORS	**LISTED:** AVGO, COHR, INTC, LITE, MRVL **PRIVATE:** Atomica, Ayar, Eliyan, Lightmatter Silicon photonics is very early in its lifecycle, with key SiPho companies primarily including publicly traded networking leaders AVGO, MRVL, CSCO, and COHR. Optical interconnects are primarily being developed by private peers Atomica, Ayar, and Lightmatter.
KEY SUPPLIERS / CUSTOMERS	Suppliers include OSRAM, TSMC, Nanosys, Lumileds. Customers not publicly disclosed.
COMPETITIVE POSITIONING	Avicena manufactures ultra-dense, ultra-low power petabit optical chip interconnects. Its optical interconnect inter-chip communications and vastly exceed the performance of legacy interconnects. Key differentiators include: • LightBundle technology: Avicena's LightBundle technology is based on microLEDs based on gallium nitride (GaN) chemistry, connected via multi-core fiber bundles to CMOS photo detector ICs. • On-Board Optics (OBO) solution: The company's OBO solution consists of an array of 1.6Tbps OBO transceivers placed near compute resources, reducing interconnect power consumption, boosting bandwidth, and facilitating scale-up.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	https://avicena.tech/
MANAGEMENT TEAM / FOUNDER(S)	**Marco Chisari – CEO:** Prior to Avicena, Mr. Chisari served as Executive VP at Samsung running the U.S. foundry business. Prior to this, he held various roles at Nokia, STMicroelectronics, GlobalFoundries, JPMorgan, and Bank of America, both in the semiconductor industry and in the financial and investment banking sectors. **Bardia Pezeshki – Founder**: Prior to Avicena, Mr. Pezeshki founded Santur and held roles at Kaiam, IBM's T. J. Watson Research Center and at SDL (now Lumentum). Mr. Pezeshki has a Ph.D. in electrical engineering from Stanford University. **Rob Kalman – Co-founder and Chief Scientist:** Prior to Avicena, Mr. Kalman founded Optivision, ONI Systems, and Luminous Networks and held engineering roles at Kaiam, where he led development of 100Gbps and 400Gbps optical transceivers for hyperscale data centers. Mr. Kalman has a BS in Electrical Engineering from UC Berkeley and a Ph.D. in Electrical Engineering from Stanford University. **Alex Tselikov – Co-founder and Principal Engineer:** Prior to Avicena, Mr. Tselikov developed optical technologies at Iolon, KLA Tencor, Opnext, Kaiam and Polar. Mr. Tselikov received his Ph.D. in Electrical Engineering from Texas A&M University in 1998.
SIZE OF COMPANY	~98 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	The company's most recent capital raise is intended "to ramp first product into production".
ANY MAJOR M&A / CORPORATE ACTIONS	In 2025, Avicena announced collaboration with TSMC to overcome copper connection distance limits and reduce energy use in AI GPU clusters by optimizing photodetector (PD) arrays for Avicena's LightBundle microLED-based interconnects using TSMC's CIS Sensor technology. In 2023, the company announced a partnership with ams-OSRAM to develop high-volume manufacturing of GaN microLED arrays for its industry-leading LightBundle architecture. In 2022, Avicena acquired a GaN microLED fabrication facility and its engineering team from Nanosys for direct control over manufacturing lines.

INVESTORS / INVESTMENT STAGE	Avicena has raised a total of $111.6MM. In its latest Series B round of funding, it raised $65MM in May 2025. Investors include Tiger Global, SK hynix, Hitachi Ventures, Lam Research, Clear Ventures, Micron Ventures, Cerberus Capital Management, Samsung Catalyst Fund, and the National Science Foundation.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Ayar Labs

COMPANY NAME	Ayar Labs, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / / AI Networking / Silicon Photonics
SIZE OF TOTAL ADDRESSABLE MARKET	Ayar Labs competes in the silicon photonics market, in support of AI computing devices. The global data center accelerator market is expected to reach over $380B per year by 2029.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** AVGO, COHR, INTC, LITE, MRVL **PRIVATE:** Atomica, Ayar, Eliyan, Lightmatter Silicon photonics is very early in its lifecycle, with key SiPho companies primarily including publicly traded networking leaders AVGO, MRVL, CSCO, and COHR. Optical interconnects are primarily being developed by private peers Atomica, Ayar, and Lightmatter.
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (Foundry), MACOM, Lumentum, Corning Customers include Lockheed Martin, Intel, HPE.
COMPETITIVE POSITIONING	Ayar Labs develops optical I/O networking solutions that replace traditional copper interconnects. The company's technology is designed to reduce latency, power consumption, and cost while improving compute efficiency. Key differentiators include: • TeraPHY optical I/O chiplet: Ayar Labs' TeraPHY optical I/O chiplet replaces conventional copper interconnects with optical connectivity for chip-to-chip communication. • SuperNova remote light source: The company's SuperNova solution serves as a multi-wavelength remote light source supporting its optical I/O architecture. • Co-packaged optics partnership with Alchip: In September 2025, Ayar Labs partnered with Alchip to co-package optical engines with AI ASICs using TSMC's COUPE and SoIC technologies, providing a pathway toward hyperscaler-class, multi-chip fabrics.
YEAR FOUNDED / ESTABLISHED	2015
WEBSITE	https://ayarlabs.com/
MANAGEMENT TEAM / FOUNDER(S)	**Mark Wade – Co-founder and CEO:** Prior to Ayar Labs, Mr. Wade held roles as CTO, SVP of Engineering and led the team that designed the optics in the world's first processor to communicate using light. Mr. Wade holds a Ph.D. from the University of Colorado. **Lisa Cummins Dulchinos – CFO:** Prior to Ayar Labs, Ms. Dulchinos was CFO of Penguin Computing, Adept Technology and has over 25 years of experience as a growth oriented financial executive in global high-tech organizations. Ms. Dulchinos holds a CPA (inactive status), an MBA from St. Mary's College, and a BA from UC Santa Barbara. **Vladimir Stojanovic – Co-founder and CTO:** Prior to Ayar Labs, Mr. Stojanovic was a professor of EECS at UC Berkeley, and led the team that designed the world's first processor to communicate using light. He was also one of the key developers of the Rambus high-speed link technology. Mr. Stojanovic holds a Ph.D. from Stanford University. **Chen Sun – Co-founder and Chief Scientist:** Prior to Ayar Labs, Mr. Sun led the group that designed the electronics and optical device drivers in the world's first processor to communicate using light. He holds a Ph.D. from MIT.
SIZE OF COMPANY	258 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Ayar's commercialization in 2026 centers on scaling volume production of its TeraPHY optical engine and SuperNova light emitters to replace copper interconnects and eliminate bottlenecks in rack-scale AI infrastructure.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Ayar Labs has raised a total of $872MM over 12 funding rounds. In its latest Series E round of funding, the company received $500MM in March of 2026. Investors include Intel Capital, Alchip Technologies, ARK Invest, Neuberger Berman, MediaTek, Light Street Capital, Playground Global, Alumni Ventures, IAG Capital Partners, NVIDIA, Boardman Bay Capital Management, 3M New Ventures, Advent International, Venture Tech Alliance, Global Foundries, Lockheed Martin Ventures, Autopilot Ventures, AXIAL Partners, AMD Ventures, Applied Ventures, Capital TEN, Hewlett Packard Pathfinder, Tyche Partners, Berkeley Frontier Fund, Infinitum Partners, Nautilus Venture Partners, Agave Partners Capital, Page One Ventures, Atreides Capital, Founders Fund, Alumni Ventures, BlueSky Capital, SGInnovate, Asuun, Downing Ventures, Castor Ventures, Parkway Venture Capital, DARPA, Silicon Valley Bank, TechU Ventures, The House Fund, Silicon Catalyst, National Science Foundation, CITRIS Foundry.
LATEST ROUND VALUATION	Press reports indicate that the Mar 2026 round was raised $500M at $3.75B valuation.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Baya Systems

COMPANY NAME	Baya Systems
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking / Chiplet Interconnect & Network-on-Chip IP
SIZE OF TOTAL ADDRESSABLE MARKET	Operates in the chiplet interconnect, network-on-chip, semiconductor system IP, heterogeneous compute, AI accelerator, automotive semiconductor, data-center, and HPC markets. Baya Systems develops software-defined data-movement and system-interconnect IP for modular SoCs and multi-chiplet systems.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** CEVA **PRIVATE:** Achronix Data Acceleration, SiFive, Eliyan
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: Openchip
COMPETITIVE POSITIONING	Baya Systems develops software-driven interconnect and data-movement IP designed to connect CPUs, GPUs, accelerators, memory, and I/O components across complex SoCs and chiplet-based systems. Its platform targets the bandwidth, latency, integration, and scalability requirements of AI infrastructure, automotive, and HPC designs. Key differentiators include: • WeaveIP software-defined fabric: Baya's platform provides a configurable fabric for moving data among heterogeneous compute, memory, accelerator, and I/O components within advanced semiconductor systems. • Chiplet-ready system architecture: The company extends interconnect functionality across modular chiplet designs, allowing multiple dies to operate as an integrated system. • Network-on-chip automation: Baya's intellectual property covers hierarchical NoC mapping, multi-instancing, packet movement across clusters, automated control-register generation, and shared-buffer configuration. • Cache-coherent data movement: The platform includes cache-coherency management for multi-level systems, supporting data consistency across increasingly complex processors and chiplet configurations. • Software-driven design exploration: Baya combines configurable IP with software-based architectural exploration to reduce integration complexity, shorten design cycles, and manage implementation risk.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://bayasystems.com/
MANAGEMENT TEAM / FOUNDER(S)	**Sailesh Kumar, Ph.D. - Founder & CEO:** Prior to founding Baya Systems, Dr. Kumar co-founded NetSpeed Systems and served as CTO until its acquisition by Intel, where Dr. Kumar subsequently became an Intel Fellow and worked on Xeon architecture and data-center compute protocols. Dr. Kumar holds a B.Tech. from IIT Kanpur and a Ph.D. from Washington University in St. Louis and has authored more than two dozen papers and over 150 patents. **Nandan Nayampally - Chief Commercial Officer:** Nayampally leads Baya's commercial activities and semiconductor ecosystem engagement, including its positioning across chiplet integration, AI scalability, and heterogeneous computing. **Joji Philip - Co-Founder & SVP of IP Engineering:** Philip co-founded Baya and leads engineering development of the company's hardware and semiconductor-IP portfolio.
SIZE OF COMPANY	~128 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In July 2025, Baya received Frost & Sullivan's Global Semiconductor IP Interconnect Solutions Technology Innovation Leadership recognition. In October 2025, Baya was named Andes Technology's Partner of the Year. In January 2026, the company expanded its Bengaluru engineering hub to support AI, automotive, and HPC development. In June 2026, Openchip selected Baya's data-movement platform for next-generation RISC-V intelligent-compute systems.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Baya Systems has raised approximately $47MM over two funding rounds. Before its latest $36MM Series B round in January 2025, the company raised an estimated $11MM in its Series A round in April 2023. Baya Systems investors include Maverick Silicon, Intel Capital, Synopsys, and Matrix. The company's current investment stage is Series B.
LATEST ROUND VALUATION	Baya Systems completed a $36MM Series B financing in January 2025 at a $133MM pre-money valuation and $169MM post-money valuation. Series B preferred shares were issued at $2.64 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For chiplet-interconnect, network-on-chip, and semiconductor system-IP companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary during customer qualification and design-win conversion. More mature semiconductor-IP companies are generally valued on P/E or P/FCF given their recurring licensing and royalty revenue, high incremental margins, and established profitability.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Cornelis Networks

COMPANY NAME	Cornelis Networks, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Networking
SIZE OF TOTAL ADDRESSABLE MARKET	Cornelis Networks specializes in high-performance interconnects for AI and HPC clusters, continuing Intel's Omni-Path technology. Its networking hardware links large numbers of AI chips with minimal latency.
GEOGRAPHIC PRESENCE	Wayne, Pennsylvania
NOTABLE COMPETITORS	**LISTED:** ANET, AVGO, CRDO, MRVL, NVDA **PRIVATE:** Duality Tech, Optalysys, Fabric Cryptography
KEY SUPPLIERS / CUSTOMERS	Suppliers: Not disclosed Customers: Lawrence Livermore National Lab, NASA, various research universities
COMPETITIVE POSITIONING	Cornelis Networks provides end-to-end purpose-built high-performance fabrics for commercial, scientific, academic, and government organizations. The company's offerings are designed to address key AI infrastructure requirements, including scalable architectures, high-bandwidth connectivity, and compatibility with accelerators and GPUs. Key differentiators include: • Cornelis Omni-Path Express fabric: Cornelis' Omni-Path Express fabric includes adapters, switches, gateways, cables, and software optimized to deliver high-performance networking for HPC applications such as computational fluid dynamics, computational chemistry, molecular dynamics, genomics, financial risk modeling, computer-aided engineering, weather prediction, and seismic imaging. • CN6000 Ethernet SuperNIC: Unveiled in November 2025, the CN6000 is an 800 Gbps Ethernet SuperNIC that extends the Omni-Path architecture to Ethernet RoCEv2 while maintaining full compliance with the Ultra Ethernet standard. • Ultra Ethernet roadmap: Cornelis plans to transition to Ultra Ethernet in 2027 with a 1.6 Tbps NIC and a high-density 72-port switch, expanding into intra-rack scale-up networking.
YEAR FOUNDED / ESTABLISHED	2020 (INTC spinoff)
WEBSITE	https://www.cornelis.com/
MANAGEMENT TEAM / FOUNDER(S)	**Lisa Spelman – CEO:** Prior to Cornelis Networks, Ms. Spelman served as Corporate Vice President in the Data Center & AI Group at Intel, where she was also GM of Xeon Products and Solutions. Ms. Spelman holds a BA in Business Administration from the University of Washington. **Philip Murphy Jr. – Co-founder and COO:** Prior to Cornelis Networks, Mr. Murphy Jr. was a director at Intel, where he managed fabric platform planning, product positioning, and business development. He also held engineering leadership roles at QLogic, SilverStorm, and Unisys Corporation. Mr. Murphy Jr. holds a BS in mathematics from St. Joseph's University and an MS in computer and information science from the University of Pennsylvania. **Vladimir Tamarkin – Co-founder and VP of Platform Engineering:** Prior to Cornelis Networks, Mr. Tamarkin was a senior principal engineer and director of fabric platform engineering at Intel Corporation. He also held leadership roles at QLogic, SilverStorm, and Unisys Corporation. Mr. Tamarkin holds a BS in mechanical engineering from Drexel University and an EMTM from the University of Pennsylvania. **Gunnar Gunnarsson – Co-founder and VP of Solutions Delivery/Support:** Prior to Cornelis Networks, Mr. Gunnarsson was director of OEM program management at Intel. He also held leadership roles at QLogic, SilverStorm, and Unisys Corporation. Mr. Gunnarsson holds a BS in electrical engineering from the University of Delaware, an MS in engineering from Penn State University, and an MBA from the Wharton School of the University of Pennsylvania.
SIZE OF COMPANY	247 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	In October 2024, the company announced a partnership with Lenovo through which its CN5000 Omni-Path high-performance networking products will be paired with Lenovo's ThinkSystem servers and storage product offerings. In November 2025, the company introduced the CN6000 800Gbps Ethernet SuperNIC, marking its expansion beyond traditional Omni-Path deployments into Ultra Ethernet-compliant AI networking infrastructure and broadening its addressable market across hyperscale AI environments.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Cornelis Networks has raised a total of ~$139MM to date. Investors: Foresight Group, SQN Venture Partners, Adit Ventures, KittyHawk, Intel Capital, Ridgeline Capital Partners, IAG Capital Partners, Chestnut Street Ventures, Global Brain Corporation, Downing Ventures, Alumni Ventures.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Eliyan Corporation

COMPANY NAME	Eliyan Corporation
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Networking
SIZE OF TOTAL ADDRESSABLE MARKET	Eliyan provides chiplet interconnect IP enabling high-speed communication between chiplets in advanced semiconductors. Its technology (e.g. NuLink) helps stitch together multiple silicon dies in a package.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** CDNS, RMBS, SNPS **PRIVATE:** Tenstorrent (Blue Cheetah Analog Design), Alphawave Semiconductor
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry). Customers are not publicly disclosed.
COMPETITIVE POSITIONING	Eliyan develops die-to-die interconnect architectures designed to provide high bandwidth and low latency with high area and power efficiency. The company's technology is focused on enabling efficient inter-die communication for advanced computing applications, including AI training and inference workloads. Key differentiators include: • NuLink die-to-die PHY: Eliyan's NuLink is a die-to-die interconnect PHY that delivers performance, power, and area efficiency across both silicon and organic substrates. • Simultaneous bidirectional signaling (SBD): NuLink's SBD capability enables data transmission in both directions simultaneously, effectively doubling bandwidth per interface and maximizing performance within a minimal interconnect area
YEAR FOUNDED / ESTABLISHED	2021
WEBSITE	https://eliyan.com/
MANAGEMENT TEAM / FOUNDER(S)	**Ramin Farjadrad – Co-founder and CEO:** Prior to Eliyan, Mr. Farjadrad co-founded Velio Communications (acquired by Rambus/LSI Logic) and Aquantia (acquired by Marvell). Mr. Farjadrad holds a Ph.D. from Stanford University. **Patrick Soheili – Chief Strategy and Business Officer:** Prior to Eliyan, Mr. Soheili was CEO of three early-stage technology companies and held leadership positions at eSilicon, Altera, AMD, and TRW. Mr. Soheili holds a BS in Math from UCSB and the University of Bristol, and an MBA from Purdue University. **Syrus Ziai – Chief Operating Officer:** Prior to Eliyan, Mr. Ziai held engineering leadership roles at Nuvia, Ikanos, Qualcomm, and PsiQuantum and has managed 100+ engineering teams across the semiconductor industry. Mr. Ziai holds an MS in electrical engineering from Stanford University and has developed 11 patents.
SIZE OF COMPANY	~101 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Chief Strategy and Business Officer Patrick Soheili indicated Eliyan expects revenue to grow from low millions of dollars in 2025 to targeting initial chiplet shipments and forecasting hundreds of millions of dollars in sales by the end of 2027. In November 2024, the company announced the successful tape out of its NuLink PHY in a x64 UCIe Advanced Package Module on Samsung Foundry's SF4X 4nm advanced manufacturing process.The interconnect solution utilizes advanced packaging micro bumps at 45-micron pitch and offers the industry's lowest power density. The company will use the technology for building its custom High Bandwidth Memory (HBM) base dies and offers it as IP to enable ASIC disaggregation approaches that leverage chiplets.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Eliyan raised a total of ~$252MM to date. In its latest Series C round of funding, the company received $145MM in July of 2026. Investors: AMD, AMD Ventures, Alphabet, Arm, Coherent, Meta, Seligman Ventures, Cisco Investment, Lumentum, VentureTech Alliance, MESH, Cleveland Avenue, Samsung Catalyst Fund, Tiger Global Management, Intel Capital, SK Hynix, Tracker Capital Management, Celesta Capital, Micron Ventures.
LATEST ROUND VALUATION	$252MM valuation as of Eliyan's July 2026 Series C funding round.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: HyperLight

COMPANY NAME	HyperLight
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Networking
SIZE OF TOTAL ADDRESSABLE MARKET	HyperLight is positioned at the intersection of AI networking and optical interconnects. The company's TFLN-based modulators, PICs, and optical chiplets target a critical enabling layer for 800G, 1.6T, and future co-packaged-optics deployments.
GEOGRAPHIC PRESENCE	Cambridge, Massachusetts
NOTABLE COMPETITORS	**LISTED:**COHR, LITE, AAOI, CSCO **PRIVATE:**LIGENTEC, PhotonIC, Ayar Labs, Celestial AI, OpenLight
KEY SUPPLIERS / CUSTOMERS	Key manufacturing partners include UMC and Wavetek, which provide high-volume 6-inch and 8-inch foundry manufacturing for HyperLight's TFLN Chiplet platform, and Jabil, which contributes electronics manufacturing, assembly, and systems-integration capabilities for AI optical infrastructure. HyperLight works with customers across the semiconductor, networking, telecom, and hyperscale data-center markets, though major end customers remain largely undisclosed.
COMPETITIVE POSITIONING	HyperLight develops thin-film lithium niobate (TFLN) photonics technology to address the bandwidth and power constraints of AI-scale optical networks. Its platform is designed to enable more power-efficient optical links as networks transition to higher-speed optical architectures. Key differentiators include: • TFLN Chiplet platform: HyperLight's TFLN Chiplet platform combines high electro-optic bandwidth, CMOS-level drive voltages, and low optical loss to enable more power-efficient optical links for 200G- and 400G-per-lane architectures. • Common photonics platform: The company's technology is designed as a common platform that supports IMDD pluggables, coherent datacom and telecom modules, and co-packaged optics. • High-volume foundry production: HyperLight has established high-volume 6-inch and 8-inch foundry production with UMC and Wavetek, supported by manufacturing ecosystem partnerships including Jabil.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	http://www.hyperlightcorp.com
MANAGEMENT TEAM / FOUNDER(S)	**Dr. Mian Zhang - Co-Founder and CEO:**Dr. Zhang earned his PhD in Applied/Engineering Physics from Cornell University before joining Harvard as a postdoctoral fellow from 2015–18, where he worked in Marko Lončar's group on the TFLN technology that became HyperLight's foundation. He has led HyperLight from Harvard spinout through development of its commercial foundry platform and current AI-infrastructure expansion. **Dr. Christian Reimer - Co-Founder and Chief Architect:**Dr. Reimer has played a central role in translating the technology from research into commercial products. He now leads architectural development of the company's integrated-photonics platform. **Dr. Kevin Luke - Co-Founder and Head of Technology Development:** Dr. Luke leads technology development and process innovation, helping translate TFLN device breakthroughs into scalable, manufacturable photonics products.
SIZE OF COMPANY	HyperLight has between 50-100 employees.
GROWTH TRAJECTORY	HyperLight has evolved from a 2018 Harvard spinout focused on TFLN research into a commercial photonics platform targeting AI and data-center infrastructure. Following its Series B funding in 2024, HyperLight achieved ISO 9001 certification and advanced customer qualification of its manufacturing platform. In 2026, HyperLight reached a key commercialization milestone by partnering with UMC and Wavetek for high-volume 6-inch and 8-inch foundry production of its TFLN Chiplet platform and deepening manufacturing relationships with Jabil, before raising an $80M Series C led by MediaTek. The company is now focused on scaling manufacturing and customer adoption as its product roadmap advances from 200G-per-lane solutions toward 400G-per-lane optical interconnects spanning pluggable, coherent, and co-packaged optics architectures.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	HyperLight has raised ~$135M since inception, including a $37M Series B in 2024 and an $80M Series C led by MediaTek in 2026. Investors: MediaTek, Summit Partners, UMC Capital, Jabil, Foxconn, EDBI, CDIB-TEN Capital, Qatar Investment Authority, Engine Ventures, Foothill Ventures and Xora Innovation.
LATEST ROUND VALUATION	Post Series C valuation is $867MM.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and edge processor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and edge processor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Lightmatter

COMPANY NAME	Lightmatter, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Silicon Photonics
SIZE OF TOTAL ADDRESSABLE MARKET	The company operates in the rapidly expanding silicon photonics market. The global AI computing market is expected to reach over $200B per year by 2028, and the advanced packaging market is ~$30B according to UBS estimates.
GEOGRAPHIC PRESENCE	Mountain View, California
NOTABLE COMPETITORS	**LISTED:** AVGO, COHR, INTC, MRVL, NVDA **PRIVATE:** Ayar Labs, Celestial AI, Lightelligence, Luminous Computing, Avicena Tech, Ranovus, OpenLight
KEY SUPPLIERS / CUSTOMERS	Key manufacturing and technology partners include GlobalFoundries, TSMC, Tower Semiconductor, Amkor, ASE, GUC, Cadence, Synopsys and Qualcomm. Lightmatter joined the NVIDIA NVLink Fusion ecosystem in 2026, allowing Passage CPO/NPO products to connect semi-custom AI accelerators with NVIDIA switching infrastructure. The company is working with hyperscalers and custom-silicon developers on AI interconnect deployments, although most major end customers remain undisclosed.
COMPETITIVE POSITIONING	Lightmatter develops a 3D silicon-photonics architecture designed to address the bandwidth, reach, and power constraints of electrical I/O in large-scale AI clusters. The company's technology integrates optical connectivity directly within semiconductor packaging to enable more efficient processor-to-processor communication. Key differentiators include: • Passage platform: Lightmatter's Passage platform integrates optical connectivity directly within advanced semiconductor packaging, enabling higher bandwidth density and more efficient processor-to-processor communication than traditional edge-I/O architectures. • 1.6 Tbps-per-fiber optical links: The company has demonstrated optical links capable of delivering 1.6 Tbps per fiber. • Photonic connectivity portfolio: Lightmatter is building a portfolio spanning photonic interposers, optical engines, and co-packaged optical connectivity. • Manufacturing partnerships: The company has established manufacturing partnerships with GlobalFoundries, TSMC, Amkor, and ASE to support high-volume deployment. • NVIDIA NVLink Fusion ecosystem: Lightmatter's participation in NVIDIA's NVLink Fusion ecosystem strengthens its position in next-generation AI networking.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://lightmatter.co/
MANAGEMENT TEAM / FOUNDER(S)	**Nicholas Harris, PhD - Co-founder and CEO:**Prior to founding Lightmatter, Mr. Harris was an Intelligence Community Postdoctoral Fellow and National Science Foundation Fellow at MIT where he received his PhD in Electrical Engineering and Computer Science. His research in integrated photonics and programmable optical systems helped establish the field of programmable photonics. He has authored 80+ publications, and holds 80+ patents. **Darius Bunandar - Co-founder and Chief Scientist:** Leads Lightmatter's long-term photonics and systems-architecture initiatives. Prior to founding Lightmatter, Mr. Bunandar received his PhD in physics at MIT, where he focused on quantum computation and communication using nanophotonic circuits. **Thomas Graham - Co-founder and Head of Machine Learning:** Leads machine-learning efforts and has played a key role in translating the company's photonic technologies into AI hardware and software architectures. Lightmatter currently lists him as Head of Machine Learning.
SIZE OF COMPANY	~367 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	Lightmatter has evolved from a 2017 MIT spinout focused on photonic AI computing into a leading optical-interconnect platform for hyperscale AI infrastructure. Manufacturing partnerships with GlobalFoundries, Amkor, and ASE have established a path to volume production, while recent product launches and demonstrations, including the Passage platform, M1000 photonic interposer, Passage L200, and 1.6 Tbps-per-fiber optical connectivity, have expanded the company's role in next-generation AI networking. Lightmatter's inclusion in NVIDIA's NVLink Fusion ecosystem in 2026 marked a significant commercial milestone, positioning the company to benefit as co-packaged and advanced optical interconnect technologies move toward deployment in large-scale AI clusters.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Lightmatter has raised $821M since inception, latest Series D amount in December 2025 was not disclosed. Investors: Bossa Invest, T. Rowe Price, Google Ventures, Fidelity, Eastward Capital Partners, Viking Global Investors, SIP Global Investors, AXIAL Partners, Trajectory Ventures, Hewlett Packard Enterprise, Matrix, Spark Capital, Fast - by GETTYLAB, Lockheed Martin Ventures.
LATEST ROUND VALUATION	$4.5B post-money valuation as of the last Oct 2024 Series D funding round.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and edge processor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and edge processor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: MangoBoost

COMPANY NAME	MangoBoost
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking / AI Infrastructure & Data Processing Units
SIZE OF TOTAL ADDRESSABLE MARKET	MangoBoost operates in the data-processing-unit, AI infrastructure, AI-factory networking, distributed-compute orchestration, intelligent tiered memory, storage acceleration, Ethernet connectivity, and hyperscale data-center markets. MangoBoost combines DPU hardware, networking, memory, storage, and orchestration software to manage AI workloads across distributed infrastructure.
GEOGRAPHIC PRESENCE	Bellevue, Washington
NOTABLE COMPETITORS	**LISTED:** NVDA, AMD, CBRS **PRIVATE:** Groq, Enfabrica, Axelera AI
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	MangoBoost develops full-stack AI infrastructure designed to orchestrate agent-based computing workloads across distributed environments. Its platform combines DPU-based acceleration, heterogeneous compute management, intelligent memory and storage, Ethernet networking, and workload-orchestration software. Key differentiators include: • Proprietary DPU architecture: MangoBoost develops data-processing-unit technology that offloads networking, storage, virtualization, and infrastructure-management tasks from host CPUs. • Multi-agent workload orchestration: The platform coordinates agent-based AI workloads across distributed environments through a unified infrastructure-management layer. • Heterogeneous compute management: MangoBoost manages resources across different processors and system architectures, targeting improved utilization in large-scale AI infrastructure. • Intelligent tiered memory and storage: The company's technology includes disaggregated storage acceleration, data reduction, storage virtualization, and tiered-memory management. • Ethernet-based AI networking: MangoBoost develops networking infrastructure for all-to-all communication, multi-node training, network-interface acceleration, and large GPU clusters. • Full-stack optimization model: The company combines hardware, systems software, and orchestration rather than addressing networking, storage, or compute bottlenecks as isolated functions.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://mangoboost.io
MANAGEMENT TEAM / FOUNDER(S)	**Jangwoo Kim, Ph.D. - Co-Founder & CEO:** Dr. Kim is a professor of Electrical and Computer Engineering at Seoul National University and previously taught at POSTECH. Before entering academia, Dr. Kim worked as a CPU and systems architect at Sun Microsystems and Oracle on SPARC T4 processors and scalable servers; Dr. Kim holds a Ph.D. from Carnegie Mellon University and undergraduate and master's degrees from Cornell University. **Junki Park, Ph.D. - Co-Founder, CFO & COO:** Dr. Park co-founded MangoBoost and oversees the company's financial and operating functions as it expands its DPU and AI-infrastructure platform across North America and Asia. **Dongup Kwon, Ph.D. - Co-Founder & CTO:** Dr. Kwon co-founded MangoBoost and leads development of the company's DPU architecture, distributed-compute infrastructure, and systems technology.
SIZE OF COMPANY	~139 employees on LinkedIn (accessed 9/16/2026)

GROWTH TRAJECTORY	In October 2024, MangoBoost and WorldTech IT joined the Linux Foundation's Open Programmable Infrastructure project. In April 2025, the company reported MLPerf Inference v5.0 results for Llama 2 70B running on AMD Instinct MI300X GPUs. In May 2025, MangoBoost and AIC signed an MoU to explore DPU-based server and storage solutions. In June 2025, MangoBoost reported a new multi-node LLM-training benchmark on AMD Instinct GPUs. The company continued presenting its DPU and data-center acceleration solutions during 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	In December 2025, Helios Ventures acquired a stake from an undisclosed seller for an estimated KRW 11.0B, equivalent to $7.5MM, through a private secondary transaction per Pitchbook.
INVESTORS / INVESTMENT STAGE	MangoBoost has raised $66MM across its completed Seed and Series A rounds. The company has also secured $34.7MM toward an announced Series B round that remained in progress as of September 2, 2026 and was expected to close by December 31, 2026. MangoBoost investors include Samsung Venture Investment, KT Investment, AJU IB Investment, Kiwoom Investment, BonAngels Venture Partners, Aion Asset Management, IMM Investment, Stonebridge Ventures, Premier Partners, DSC Investment, Shinhan Venture Investment, KB Investment, Korea Development Bank, Carnegie Mellon University, and Helios Ventures. The company's pending investment stage is Series B.
LATEST ROUND VALUATION	The announced $34.7MM Series B round did not disclose a valuation. The latest completed financing with a disclosed valuation was the September 2023 Series A round, which valued MangoBoost at $250MM pre-money and $305MM post-money. Series A preferred shares were issued at $14.74 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For DPU, AI-infrastructure, and distributed-compute companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may fluctuate with product development, hyperscaler qualification, and infrastructure deployment cycles. More mature semiconductor and networking-infrastructure companies are generally valued on P/E or P/FCF given established shipment volumes, recurring software and support revenue, and more predictable profitability.
SECONDARY PRICING DATA	Helios Ventures acquired an equity stake for an estimated KRW 11.0B, or approximately $7.5MM, in December 2025. The number and class of shares acquired and the transaction's implied company valuation were not disclosed.

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Menlo Microsystems

COMPANY NAME	Menlo Microsystems, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Networking / MEMS & Switching Technologies
SIZE OF TOTAL ADDRESSABLE MARKET	Menlo Micro focuses on MEMS switches, with a potential to capture a $100B+ market in RF communications, power switching, and protection devices.
GEOGRAPHIC PRESENCE	Irvine, California
NOTABLE COMPETITORS	**LISTED:** ADI, IFX, QRVO, RNECY, SWKS **PRIVATE:** AirMems, Pickering Electronics In the RF MEMS switch market, key players include ADI and QRVO, competing with Menlo Micro and start-ups like AirMems. Competition also arises from SWKS, IFX, and RNECY, which provide a range of RF components and power switching solutions.
KEY SUPPLIERS / CUSTOMERS	Menlo Micro's technology originated at GE Research, while Corning has been an early investor and strategic partner supporting its glass-based MEMS platform. The company is expanding manufacturing through its New York-based Ideal Fab. While most customers remain undisclosed, Menlo reports supporting 14 of the top 20 semiconductor manufacturers in test and measurement and has disclosed aerospace and defense programs spanning guided missiles, drone systems, and the U.S. Navy's advanced circuit-breaker initiative.
COMPETITIVE POSITIONING	Menlo Micro develops switching technology based on a proprietary MEMS architecture designed to combine the advantages of electromechanical relays with semiconductor-style manufacturing. The company's technology targets applications across RF, high-speed digital, power, test & measurement, and cryogenic markets. Key differentiators include: • Ideal Switch: Menlo Micro's proprietary Ideal Switch combines a metal-to-metal MEMS contact architecture with semiconductor-style wafer manufacturing. • Performance characteristics: The technology is designed to deliver the low loss, high isolation, and power handling of electromechanical relays while providing faster switching, longer operating life, and a smaller footprint. • Switching speed, lifetime, and frequency range: Menlo reports switching speeds below 10 µs, lifetimes exceeding 3 billion cycles, and operation from DC to at least 70 GHz. • Wafer-scale manufacturing: Menlo's wafer-scale manufacturing approach provides a potential cost, scalability, and consistency advantage relative to conventional relay technologies.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://menlomicro.com/
MANAGEMENT TEAM / FOUNDER(S)	**Richard "Rich" Simoncic — CEO**: Appointed CEO in August 2026, Simoncic brings more than three decades of semiconductor industry experience from MCHP, where he most recently served as COO. During his tenure, he helped build MCHP's analog and mixed-signal businesses into major growth drivers. His background in analog, mixed-signal, power, and manufacturing operations is highly relevant as Menlo scales commercialization of its Ideal Switch platform. **Chris Keimel — Co-Founder and CTO**: As CTO, he leads the company's technology roadmap, product architecture, and long-term innovation strategy. He previously spent more than a decade at GE Research, where he led development of the Metal MEMS technologies that became the foundation of Menlo's Ideal Switch. **Mark Czepiel - CFO:** Prior to Menlo Micro, Mr. Czepiel served as CEO of WiSpry, a RF MEMS start-up. He earned his MBA in finance from USC's Marshall School of Business and holds a bachelor's degree in economics from Loyola University of Chicago.
SIZE OF COMPANY	~70-80 employees according to various sources.

GROWTH TRAJECTORY	Menlo Micro has evolved from a GE Research-developed MEMS technology platform into a commercial supplier of switching solutions for RF, power, semiconductor test, aerospace & defense, and emerging quantum applications. Following a Series B financing in 2020 and a $150M Series C in 2022, the company expanded its product portfolio and scaled manufacturing through its Ideal Fab facility in New York. Menlo surpassed one million cumulative Ideal Switch shipments in 2025, providing an important commercialization milestone. In 2026, the company broadened its product footprint across high-speed AI test and validation, millimeter-wave communications, cryogenic quantum systems, and advanced power-switching applications, including the U.S. Navy's advanced circuit-breaker program.
ANY MAJOR M&A / CORPORATE ACTIONS	Menlo was spun out of GE Global Research in December 2016 to commercialize MEMS switching technology developed over more than a decade of GE R&D. In July 2023, Menlo completed acquiring a facility in Tompkins County, NY previously operated by Kionix (a ROHM Semiconductor subsidiary) — establishing its first domestic production facility, the Ideal Fab, with an initial three-year investment commitment of $50M and a planned expansion of up to $150M contingent on growth.
INVESTORS / INVESTMENT STAGE	Menlo Micro raised ~$235MM to date. In its latest Series C round of funding, the company received $150MM in March of 2022. Investors: GE Ventures, Corning, Microsemi, Paladin Capital Group, Vertical Venture Partners, Future Shape, Piva Capital, Fidelity Management & Research, DBL Partners, Adage Capital, Standard Investments, and PeopleFund.
LATEST ROUND VALUATION	Valued at $482MM post Series C in March 2022.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For RF and switching companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer RF and switching companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn



AI Networking: NextHop AI

COMPANY NAME	NextHop AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking
SIZE OF TOTAL ADDRESSABLE MARKET	NextHop AI operates in the AI data-center networking, Ethernet switching, optical connectivity, network operating systems, hyperscale cloud infrastructure, AI-cluster deployment, and custom networking-hardware markets. Lightspeed estimated a $75B total Ethernet-switching market by 2029, including a $35B cloud segment.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** NVDA, CBRS **PRIVATE:** Groq, Tenstorrent, Eridu
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: Undisclosed hyperscale cloud operators
COMPETITIVE POSITIONING	Nexthop AI develops customized AI-networking infrastructure for hyperscale cloud operators. Its offering combines networking hardware, software, optics, and operating-system capabilities intended to improve energy efficiency, reduce operating costs, increase reliability, and accelerate deployment of large AI clusters. Key differentiators include: • Joint-development manufacturing model: Nexthop works directly with hyperscale customers to co-design networking systems rather than offering only a standardized merchant-switch portfolio. • Integrated hardware, software, and optics stack: The company combines customized switch hardware, network software, optical components, and physical-system design within one infrastructure offering. • AI-cluster power efficiency: Nexthop targets lower energy use and operating cost per AI workload by optimizing network architecture, signal integrity, thermals, and power consumption per port. • Hyperscaler-specific customization: Products are designed around individual customers' GPU topology, networking requirements, deployment model, and supply-chain constraints. • Compressed development cycles: The company's joint-development approach is intended to shorten the time required for cloud operators to evaluate, design, and deploy specialized AI-networking systems. • Scale-out and scale-across networking: Nexthop's platform targets both connectivity within large AI clusters and networking across increasingly distributed compute environments.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://nexthop.ai
MANAGEMENT TEAM / FOUNDER(S)	**Anshul Sadana - Founder, CEO & Chairman:** Prior to founding Nexthop AI, Sadana spent 17 years at Arista Networks, including as Chief Customer Officer and COO, and helped scale the company across cloud and enterprise networking while overseeing hardware design, manufacturing, supply chain, sales, and support. Before Arista, Sadana spent eight years at Cisco and holds an M.S. in Computer Science from the University of Illinois and an MBA from the Wharton School.
SIZE OF COMPANY	~299 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Nexthop AI emerged publicly in March 2025 with customized networking infrastructure for large cloud operators. The company's Series A proceeds were designated to accelerate networking solutions for cloud and hyperscale AI clusters. By March 2026, Nexthop had introduced scale-out and scale-across switches engineered for hyperscalers and neo-cloud operators. The company's broader platform includes networking devices, software, operating-system technology, optics, and photonics tools.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Nexthop AI has raised approximately $610MM over three funding rounds. Before its latest $500MM Series B round in March 2026, the company raised $85MM through debt and equity in its Series A round in March 2025 and $25MM in Seed funding. Nexthop AI investors include Lightspeed Venture Partners, Andreessen Horowitz, Battery Ventures, WestBridge Capital, Altimeter Capital Management, Craft Ventures, Emergent Ventures, and Kleiner Perkins. The company's current investment stage is Series B.
LATEST ROUND VALUATION	Nexthop AI completed a $500MM Series B financing in March 2026 at a $3.7B pre-money valuation and $4.2B post-money valuation. Series B preferred shares were priced at an estimated $30.23 per share per Pitchbook.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For AI-networking and hyperscale infrastructure companies that are growing rapidly but remain sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, with valuation influenced by hyperscaler customer adoption, product-deployment visibility, and revenue concentration. More mature networking and AI-infrastructure companies are generally valued on P/E or P/FCF given established shipment scale, recurring software and support revenue, and more predictable margins.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: nEye.ai

COMPANY NAME	nEye.ai
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking / Optical Circuit Switching & Silicon Photonics
SIZE OF TOTAL ADDRESSABLE MARKET	Operates in the optical circuit switching, silicon photonics, AI data-center networking, composable infrastructure, optical interconnect, and fiber-optic equipment markets. nEye.ai develops chip-scale optical switches intended to connect and dynamically allocate compute, memory, and storage resources in AI data centers.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** CIEN, LITE **PRIVATE:** Salience Labs, iPronics, Xscape Photonics
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	nEye.ai develops chip-scale optical circuit switches designed to improve bandwidth density and energy efficiency in AI data-center networks. Its technology integrates silicon photonics, MEMS, and CMOS using a foundry-compatible architecture intended to replace larger mechanical optical-switch assemblies. Key differentiators include: • OCS-on-a-chip platform: nEye.ai integrates the core functionality of an optical circuit switch into a chip-scale platform, reducing the size and component complexity associated with traditional optical-switch systems. • Silicon photonics, MEMS, and CMOS integration: The company combines silicon-photonic waveguides, MEMS switching structures, and CMOS control technology within a unified architecture. • Composable AI infrastructure: The platform enables dynamic pooling and reconfiguration of CPUs, GPUs, memory, and storage resources across AI data-center environments. • Scale-up, scale-out, and scale-across connectivity: nEye.ai is positioning its optical-switching platform across multiple layers of AI fabrics, from local accelerator connectivity to larger distributed clusters. • Foundry-compatible manufacturing: The company uses a wafer-scale manufacturing approach intended to improve scalability and cost efficiency relative to complex mechanical optical-switch assemblies. • Pre-volume commercialization: nEye.ai is moving from technology development into high-volume manufacturing, with its 2026 expansion centered on meeting hyperscale performance and production requirements.
YEAR FOUNDED / ESTABLISHED	2020
WEBSITE	https://www.neye.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Ashish Vengsarkar, Ph.D. - CEO:** Prior to nEye.ai, Dr. Vengsarkar worked at Google on optical networks and spent approximately 20 years founding and building optical-networking companies Photuris and Nistica. Dr. Vengsarkar began his career at Bell Labs, holds 35 patents, and earned a B.Tech. from IIT Bombay, a Ph.D. in Electrical Engineering from Virginia Tech, and an MBA from the Wharton School. **Ming C. Wu, Ph.D. - Co-Founder & Chief Scientist:** Dr. Wu is the Nortel Distinguished Professor of Electrical Engineering and Computer Sciences at UC Berkeley and previously co-founded Berkeley Lights and Optical Micro-Machines. Dr. Wu began his career at AT&T Bell Laboratories and earned a B.S. from National Taiwan University and a Ph.D. from UC Berkeley. **Tae Joon Seok, Ph.D. - Co-Founder & CTO:** Prior to co-founding nEye.ai, Dr. Seok served as a professor at the Gwangju Institute of Science and Technology and as a postdoctoral scholar at UC Berkeley. Dr. Seok specializes in integrated photonics, silicon photonics, and optoelectronics and has authored more than 40 technical publications.
SIZE OF COMPANY	~48 employees on LinkedIn (accessed 9/16/2026)

GROWTH TRAJECTORY	In July 2025, nEye.ai joined Google, NVIDIA, and Microsoft in the Open OCS project. In April 2026, the company announced plans to accelerate development and high-volume manufacturing of its proprietary optical circuit switches. The company's commercialization focus has shifted from laboratory development toward foundry-based manufacturing and qualification for AI-infrastructure applications.
ANY MAJOR M&A / CORPORATE ACTIONS	nEye.ai was spun out of the University of California, Berkeley in September 2020.
INVESTORS / INVESTMENT STAGE	nEye.ai has raised approximately $153MM over three funding rounds. Before its latest $80MM Series C round in April 2026, the company raised $58MM in its Series B round in April 2025 and $14.5MM in its Series A round in September 2023. nEye.ai investors include Sutter Hill Ventures, CapitalG, M12, Micron Technology, NVIDIA, Socratic Partners, Raptor Group, UC Berkeley Chancellor's Fund, InnoLight, and TEDA Investment Holding. The company's current investment stage is Series C.
LATEST ROUND VALUATION	nEye.ai completed an $80MM Series C financing in April 2026 at a $340MM pre-money valuation and $420MM post-money valuation. Series C preferred shares were issued at $6.23 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For optical circuit switching and silicon-photonics companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary during customer qualification and manufacturing scale-up. Peer optical networking and connectivity companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: OpenLight

COMPANY NAME	OpenLight
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking / Integrated Silicon Photonics & Photonic ASICs
SIZE OF TOTAL ADDRESSABLE MARKET	OpenLight operates in the silicon photonics, photonic integrated circuit, optical interconnect, AI data-center connectivity, telecommunications, high-performance computing, sensing, and photonic foundry markets. OpenLight provides design IP, process-design kits, photonic components, and manufacturing support for custom photonic integrated circuits.
GEOGRAPHIC PRESENCE	Goleta, California
NOTABLE COMPETITORS	**LISTED:** AVGO, COHR, LITE, MRVL **PRIVATE:** HyperLight, Ayar Labs, Lightmatter, Luminous Computing, Rockley Photonics
KEY SUPPLIERS / CUSTOMERS	Suppliers: Tower Semiconductor. Customers: Advantest
COMPETITIVE POSITIONING	OpenLight develops an open silicon-photonics platform designed to integrate active and passive optical components on a common chip. The company combines indium-phosphide lasers and amplifiers with silicon-photonic modulators, detectors, and waveguides, enabling customers to design custom photonic ASICs. Key differentiators include: • OpenLight Photonics PDK: The company provides a process-design kit containing validated active and passive photonic components that customers can use to develop custom photonic integrated circuits. • Heterogeneously integrated on-chip lasers: OpenLight integrates indium-phosphide laser technology directly with silicon photonics, reducing reliance on separately attached external laser components. • 400G modulator technology: The company is expanding its PDK around a 400G modulator platform for high-bandwidth AI, cloud, and networking applications. • Active and passive component library: OpenLight combines lasers, amplifiers, modulators, detectors, waveguides, polarization components, and related photonic structures in a common platform. • Open-foundry business model: Customers can license OpenLight's PDK and intellectual property, use the company's design services, and access a production path for custom photonic ASICs. • Manufacturing and test ecosystem: The company is working with Tower Semiconductor on the PH18DA photonics ecosystem and Cadence-compatible design tools and with Advantest on silicon-photonics test solutions for high-volume manufacturing.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://openlightphotonics.com
MANAGEMENT TEAM / FOUNDER(S)	**Adam Carter, Ph.D. - CEO:** Prior to OpenLight, Dr. Carter served as Chief Commercial Officer of Foxconn Interconnect and Oclaro and was part of Oclaro's senior executive team before its $1.85B acquisition by Lumentum. Earlier, Dr. Carter led Cisco's Transceiver Module Group and held senior roles at Avago, Agilent, and Hewlett-Packard; Dr. Carter holds a B.Sc. in Applied Physics from the University of Portsmouth and a Ph.D. from the University of Wales, Cardiff. **James Dixon - CFO:** Dixon serves as Chief Financial Officer and oversees OpenLight's financial organization as the company expands its PDK, design-services, and photonic-manufacturing platform.
SIZE OF COMPANY	~51 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In March 2026, OpenLight showcased its III-V heterogeneous integrated silicon-photonics technology and production capabilities for AI, cloud, and high-speed networking. In June 2026, the company partnered with Advantest to develop silicon-photonics test solutions for high-volume manufacturing. In August 2026, OpenLight and Tower Semiconductor expanded the PH18DA photonics ecosystem and launched its PDK on Cadence design tools. The company is also expanding its active and passive component library, including its 400G modulator and heterogeneously integrated on-chip laser technology.

ANY MAJOR M&A / CORPORATE ACTIONS	OpenLight was formed as a joint venture between Juniper Networks and Synopsys in April 2022.
INVESTORS / INVESTMENT STAGE	OpenLight has raised $84MM over two funding rounds. Before its latest $50MM Series A1 and Series A2 financing in April 2026, the company raised $34MM in its initial Series A round in August 2025. OpenLight investors include Matter Venture Partners, Xora Innovation, Capricorn Investment Group, Mayfield Fund, Lam Capital, New Legacy Ventures, Catapult Ventures, Acclimate Ventures, IceNine Capital, K2 Access, Empire Angel Collective, Juniper Networks, and Synopsys. The company's current investment stage is Series A.
LATEST ROUND VALUATION	OpenLight completed a $50MM Series A1 and Series A2 financing in April 2026 at a $135MM pre-money valuation and $185MM post-money valuation. Series A1 and Series A2 preferred shares were issued at $0.40 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For integrated silicon-photonics and photonic-ASIC platform companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary during PDK adoption, customer qualification, and production ramp-up. Peer optical semiconductor and connectivity companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Point2 Technology

COMPANY NAME	Point2 Technology
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking / AI Interconnect & RF Waveguide Connectivity
SIZE OF TOTAL ADDRESSABLE MARKET	Point2 Technology operates in the AI data-center interconnect, high-speed SerDes, PAM4 connectivity, RF waveguide, hyperscale computing, 5G infrastructure, and board-to-board and chip-to-chip connectivity markets. Point2 Technology develops mixed-signal SoCs and RF-based interconnect technology for high-bandwidth data transmission.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** AVGO, MRVL, CSCO, INTC, NVDA **PRIVATE:** nEye.ai, XCENA, Oxide Computer
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Point2 Technology develops ultra-low-power, low-latency interconnect technology designed to meet the terabit-bandwidth requirements of AI data centers and 5G infrastructure. The company combines mixed-signal semiconductor design with RF transmission over dielectric waveguides as an alternative to conventional copper and optical links. Key differentiators include: • e-Tube RF-waveguide platform: Point2 uses RF signaling transmitted through dielectric waveguides to provide high-speed connectivity for AI data-center, automotive, and consumer-electronics applications. • UltraWire smart-retimer portfolio: The company's mixed-signal SoC portfolio includes retimer technology for extending transmission range and maintaining signal integrity across high-speed PAM4 links. • Terabit-scale connectivity: Point2 is developing 800 Gbps and 1.6 Tbps interconnect products intended to address the growing bandwidth requirements of accelerator clusters. • Low-power and low-latency architecture: The company positions its RF-waveguide approach around lower power consumption and latency than conventional interconnect architectures. • Chip-to-chip and board-to-board applications: Point2's intellectual property covers multi-channel chip interfaces, broadband RF communication, phase synchronization, and adjustable-bandwidth board-level waveguide interconnects. • Semiconductor ecosystem commercialization: The company is expanding commercialization with the support of strategic semiconductor and connectivity investors, including NVIDIA, Arm, UMC Capital, Bosch Ventures, and Molex.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://www.point2tech.com/
MANAGEMENT TEAM / FOUNDER(S)	**Sean Park, Ph.D. - Co-Founder, President & CEO:** Prior to founding Point2 Technology, Dr. Park founded and led 100G Ethernet transceiver company TeraSquare as CEO and CTO and held RF and analog semiconductor roles at Marvell, including Director. Dr. Park holds a B.S. in Electrical Engineering from Seoul National University and a Ph.D. in Electrical Engineering from the University of Washington. **HM Bae, Ph.D. - Co-Founder & Chairman:** Dr. Bae co-founded Point2 Technology in 2016 and serves as Chairman, supporting the company's technical and strategic development across semiconductor and high-speed networking applications. **Thomas Cho, Ph.D. - CTO:** Dr. Cho leads Point2's development of CMOS RF, analog, and mixed-signal integrated circuits for high-speed data-communications applications.
SIZE OF COMPANY	~54 employees on LinkedIn (accessed 9/16/2026)

GROWTH TRAJECTORY	In April 2026, Point2 announced a Series B extension led by NVentures to accelerate deployment of its AI-interconnect technology and was named a 2026 BloombergNEF Pioneer. In July and August 2026, the company showcased its e-Tube platform at OCP APAC and OCP Korea Tech Day. The company plans to use its Series B proceeds to accelerate development and commercialization, expand engineering operations, and strengthen semiconductor-ecosystem partnerships.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Point2 Technology has raised $146MM over two funding rounds. Before its latest $136MM Series B round in August 2026, the company raised $9.8MM in its Series A round in October 2021. Point2 Technology investors include LB Investment, Maverick Silicon, NVentures, NVIDIA, Arm, UMC Capital, Bosch Ventures, Molex, Samsung Securities, Time Folio Investment, Walden International, Nautilus Venture Partners, Quantum Ventures Korea, Shinhan Capital, Tigris Investment, and D3 Jubilee Partners. The company's current investment stage is Series B.
LATEST ROUND VALUATION	Point2 Technology completed a $136MM Series B financing in August 2026 at a $424MM pre-money valuation and $560MM post-money valuation. The financing included multiple Series B tranches with disclosed preferred-share prices ranging from $4.08 to $10.00 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For AI-interconnect, RF-waveguide, and high-speed connectivity companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary during product qualification and infrastructure deployment. Peer semiconductor connectivity and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Relativity Networks

COMPANY NAME	Relativity Networks
SECTOR / SUB-SECTOR	TMT / Communications Equipment / AI Networking / Fiber-Optic Networking Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	Relativity Networks develops hollow-core fiber networking infrastructure for distributed AI and inter-data-center connectivity. Its ChronoCore platform includes anti-resonant hollow-core fiber cables, proprietary couplers, connectors, and deployment support for integration with existing single-mode fiber networks.
GEOGRAPHIC PRESENCE	Winter Park, Florida
NOTABLE COMPETITORS	**LISTED:** MSFT **PRIVATE:** N/A
KEY SUPPLIERS / CUSTOMERS	Supplier: Prysmian. Customer: Undisclosed hyperscaler.
COMPETITIVE POSITIONING	Competitive Positioning: Relativity Networks' core differentiator is its hollow-core fiber networking architecture, which guides light primarily through air rather than solid glass to reduce propagation latency and increase the viable distance between distributed AI data centers. Key attributes include: • Air-guided optical transmission: ChronoCore uses an anti-resonant hollow-core structure that guides light through an air-filled core, reducing the propagation delay associated with conventional solid-glass fiber. • Lower-latency connectivity: ChronoCore operates at approximately 3.4 microseconds per kilometer, representing approximately 47% faster propagation and 32% lower latency than conventional fiber over the same distance. • Extended data-center reach: The lower propagation delay enables geographically separated AI data centers to operate within the same synchronization budget, expanding the area in which operators can locate infrastructure near available power. • Carrier-grade networking platform: Relativity Networks combines hollow-core fiber and cable manufactured with Prysmian with proprietary couplers and connectors designed for integration with existing single-mode networks. • Scale Across architecture: Rather than addressing optical connectivity solely within racks or individual data-center buildings, Relativity Networks targets connections between geographically distributed AI campuses that must operate as a synchronized system.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://relativitynetworks.ai
MANAGEMENT TEAM / FOUNDER(S)	**Jason Eichenholz – Co-founder and CEO:** Prior to co-founding Relativity Networks, Jason Eichenholz was a longtime photonics entrepreneur, inventor, and technology executive. He previously co-founded Luminar Technologies, where he helped commercialize lidar technologies for autonomous vehicles, and spent much of his career advancing optical, photonics, and laser-based technologies. Dr. Eichenholz holds a Ph.D. in Optical Sciences and has been involved in the commercialization of multiple deep-technology ventures. At Relativity Networks, he leads development and commercialization of the company's hollow-core fiber networking platform, which is designed to reduce latency between geographically distributed AI data centers. **Rodrigo Amezcua Correa – Co-founder and CTO:** Dr. Prior to co-founding Relativity Networks, Rodrigo Amezcua Correa established himself as a leading researcher in optical communications and hollow-core fiber technology. His work focused on optical-fiber physics, ultra-low-latency communications, and next-generation networking infrastructure. Dr. Amezcua Correa is widely recognized for his research contributions to hollow-core optical fibers, which form the foundation of Relativity Networks' ChronoCore platform. At Relativity Networks, he leads technology strategy, optical-fiber architecture, and development of the company's latency-optimized networking infrastructure.
SIZE OF COMPANY	~20 employees on LinkedIn (accessed 9/16/2026)

GROWTH TRAJECTORY	In February 2025, Relativity Networks emerged from stealth and introduced its hollow-core fiber technology for lower-latency data-center connectivity. In March 2025, the company entered a long-term manufacturing partnership with Prysmian to bring hollow-core fiber cables to market at scale, followed by initial deliveries to hyperscalers. In 2026, Relativity Networks and Prysmian produced a 24-fiber ChronoCore cable within a 10-millimeter diameter and completed installation testing in standard microducts with Dura-Line. In August 2026, the company disclosed a $40MM follow-on order from an undisclosed hyperscaler following successful testing of ChronoCore between two data centers. ChronoCore fiber is produced at Prysmian's facility in Eindhoven, with cabling completed in Claremont, North Carolina, and coupler assembly, fiber characterization, and installation training performed in Orlando.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Relativity Networks has raised $32.7MM across its disclosed pre-seed, seed, and SAFE financings. In its latest financing, the company announced a $22.0MM SAFE investment in August 2026, more than double its original target. Notable Investors: Rhapsody Venture Partners, Bell Ventures, Faster Than Glass, Prysmian, GOVO Venture Partners, and Fabric VC.
LATEST ROUND VALUATION	N/A. The August 2026 financing was structured as a SAFE investment, and no priced valuation, valuation cap, conversion discount, or ownership terms were publicly disclosed.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For communications equipment and optical networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer communications equipment and optical networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Retym

COMPANY NAME	Retym
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / AI Networking / Coherent DSP & Optical Networking Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Retym operates in the coherent optical DSP, data-center interconnect, optical networking semiconductor, cloud connectivity, telecommunications, high-speed serializer, DAC and ADC, and AI-infrastructure markets. Retym develops coherent DSP technology intended for shorter-reach and data-center optical links.
GEOGRAPHIC PRESENCE	Cupertino, California
NOTABLE COMPETITORS	**LISTED:** CIEN, COHR, LITE **PRIVATE:** nEye.ai, Xsight Labs, Scintil Photonics
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Retym develops coherent digital signal processors designed to support high-bandwidth optical connectivity in AI data centers, cloud networks, and telecommunications infrastructure. The company combines mixed-signal analog design, high-speed DAC and ADC technology, and DSP algorithms to extend coherent optics into shorter-reach applications. Key differentiators include: • Coherent DSP platform: Retym develops digital signal processors for coherent optical links, targeting data-center interconnect and shorter-distance applications that increasingly require coherent transmission. • Integrated DAC and ADC design: The company combines high-performance digital-to-analog and analog-to-digital converters with DSP processing to optimize system-level power, cost, and performance. • Short-reach coherent connectivity: Retym is extending coherent technology beyond traditional long-haul networks toward sub-100-kilometer and data-center interconnect applications. • High-speed transmitter architecture: The company's intellectual property includes common-mode- and supply-controlled serialization stages embedded within high-speed transmitter output circuits. • Clock recovery and frame alignment: Retym has developed mixed analog-digital technology for jointly managing clock recovery and frame alignment in high-speed optical communication systems. • Open ecosystem positioning: The company positions itself as an independent coherent-DSP supplier intended to avoid the channel conflicts that can arise when component suppliers also compete with their customers.
YEAR FOUNDED / ESTABLISHED	2021
WEBSITE	https://www.retym.com
MANAGEMENT TEAM / FOUNDER(S)	**Sachin Gandhi - Co-Founder & CEO:** Prior to founding Retym, Gandhi co-founded and served as COO of dryv.io, which was acquired by Synaptics, and Xpliant, which was acquired by Cavium. Gandhi subsequently served as Director of Marketing at Cavium and previously held engineering and marketing roles at Marvell and Cisco; Gandhi holds an engineering degree from the University of Mumbai and a master's degree from the University of South Florida. **Aharon El-Bahar, Ph.D. - Co-Founder & CTO:** Prior to founding Retym, Dr. El-Bahar served as CTO of Huawei's Israel Research Center and managed engineering organizations at Huawei, Samsung, and Intel. Dr. El-Bahar specializes in optics, electro-optics, cameras, chip design, and analog design and holds bachelor's, master's, and doctoral degrees in Electrical Engineering from the Technion.
SIZE OF COMPANY	~80 employees on LinkedIn (accessed 9/16/2026)

GROWTH TRAJECTORY	In March 2025, Retym emerged publicly and announced that it was finalizing production of its coherent DSPs, with plans to bring the products to market before the end of 2025. The company's product development centers on integrating high-performance DACs and ADCs with DSP algorithms to improve the power and cost profile of coherent links used in data-center interconnect and shorter-reach optical networks.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Retym has raised $180MM over four funding rounds. Before its latest $75MM Series D round in March 2025, the company raised an estimated $60MM in its Series C round in May 2023, an estimated $35MM in its Series B round, and an estimated $10MM in its Series A round per Pitchbook. Retym investors include Spark Capital, Mayfield Fund, Kleiner Perkins, Fidelity, and Arista Networks. The company's current investment stage is Series D.
LATEST ROUND VALUATION	Retym completed a $75MM Series D financing in March 2025 at a $397.5MM pre-money valuation and $472.5MM post-money valuation. Series D preferred shares were issued at $10.59 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For coherent-DSP and optical-networking semiconductor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary during product qualification, customer adoption, and initial production. Peer optical semiconductor and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Salience Labs

COMPANY NAME	Salience Labs
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Silicon Photonics / AI Networking Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	Salience Labs develops silicon-photonic optical switching solutions designed to address AI datacenter networking bottlenecks. The company targets the rapidly growing AI infrastructure market, specifically the networking and interconnect layer that connects large-scale GPU and AI accelerator clusters. As AI training and inference deployments continue to scale, networking bandwidth, latency, power consumption, and cluster utilization are becoming increasingly important constraints. Salience Labs believes all-optical switching architectures can significantly improve AI cluster performance by reducing latency, increasing bandwidth scalability, and lowering power consumption relative to conventional electrical switching architectures. The company is targeting next-generation AI datacenter deployments where optical networking is expected to play a larger role in cluster scaling.
GEOGRAPHIC PRESENCE	Oxford, England, United Kingdom
NOTABLE COMPETITORS	**LISTED:** AVGO, MRVL, CSCO, COHR, LITE **PRIVATE:** Ayar Labs, Lightmatter, Xscape Photonics The company's competitive landscape spans silicon photonics, optical interconnects, co-packaged optics, optical switching, and AI network infrastructure providers.
KEY SUPPLIERS / CUSTOMERS	The company leverages conventional semiconductor manufacturing processes to produce its silicon-photonic devices and has established relationships within the semiconductor ecosystem through investors including Applied Ventures (Applied Materials), Silicon Catalyst, and Oxford Sciences Enterprises. Salience has not publicly disclosed customer names. Management indicates that its products are being developed for hyperscalers, AI infrastructure providers, AI datacenter operators, and large GPU cluster deployments seeking lower-latency and lower-power networking fabrics.
COMPETITIVE POSITIONING	Salience Labs' primary differentiator is its silicon-photonic all-optical switching architecture designed specifically for AI datacenter connectivity. • All-optical switching architecture: Salience has developed optical circuit switch (OCS) technology that enables networking connections to remain in the optical domain, eliminating repeated optical-to-electrical-to-optical (OEO) conversions common in traditional networking infrastructure. Management believes this significantly improves bandwidth efficiency while reducing latency and power consumption. • Elimination of optical transceivers: The company's architecture removes the need for expensive and power-intensive optical transceivers typically required in electrical switching environments, potentially lowering both capex and operating costs. • AI-focused networking optimization: Salience's products are specifically designed to support large-scale AI training and inference clusters where networking increasingly represents a significant system bottleneck. Management believes lower-latency networking directly improves AI workload performance and token generation efficiency. • Silicon photonics platform: The company's switching technology is based on silicon photonics and leverages established semiconductor manufacturing processes, which management believes should support scalability, reliability, and high-volume manufacturing. • Integrated hardware and software stack: Beyond photonic integrated circuits, Salience provides control electronics, firmware, monitoring systems, and software interfaces designed to enable deployment within datacenter environments. • Commercial product introduction: In March 2026, Salience launched the OC-32M, which it describes as the industry's highest-performing 32-port all-optical silicon photonic switch. By September 2026, the company was demonstrating a 64-port optical circuit switch platform at ECOC 2026.
YEAR FOUNDED / ESTABLISHED	2021

WEBSITE	https://www.saliencelabs.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Vaysh Kewada - Founder and CEO:** Kewada serves as CEO and led the commercialization of the company's silicon-photonics technology. Prior to founding Salience Labs, he worked with Oxford Sciences Enterprises during the incubation and formation of the business. **Johannes Feldmann - Founder and CTO:** Feldmann leads technology development and is one of the principal inventors behind the company's photonic switching platform. Prior to founding Salience Labs, he conducted research in silicon photonics and optical computing technologies.
SIZE OF COMPANY	50-60
GROWTH TRAJECTORY	Salience Labs was formed based on more than a decade of research conducted at the University of Oxford and the University of Münster focused on silicon photonics and optical switching technologies. The company emerged with an $11.5MM seed round in 2022 co-led by Oxford Science Enterprises and Cambridge Innovation Capital. The financing supported early product development and commercialization efforts. In February 2025, Salience announced a $30MM Series A financing to accelerate development and commercialization of its optical switching platform for large-scale AI connectivity. Alongside the financing, the company expanded its executive team, including appointing Bonnie Tomei as CFO and adding Dr. William Jeffrey to the board. Product commercialization accelerated in 2026. In March 2026, the company launched its OC-32M all-optical switch platform and announced a collaboration with Keysight to develop an optical circuit switch testing environment. In September 2026, Salience demonstrated a 64-port optical circuit switch platform at ECOC 2026, illustrating continued scaling of its networking architecture.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Salience Labs has raised $42MM in disclosed funding across Seed and Series A financings. Notable investors include: Applied Ventures (Applied Materials), ICM HPQC Fund, Oxford Science Enterprises, Cambridge Innovation Capital, Braavos, Strategic Investment Fund, Silicon Catalyst, Jalal Bagherli, Amadeus Capital Partners, Atlantic Bridge, Passion Capital
LATEST ROUND VALUATION	Salience Labs announced a $30MM Series A financing in February 2025.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For silicon photonics, optical networking, AI interconnect, and optical-switching companies that are pre-scale and investing heavily in product development, we believe EV/Sales or P/Sales methodologies are generally the most relevant valuation frameworks, as profitability often remains depressed during commercialization and customer qualification phases. More mature networking, optical component, and semiconductor infrastructure companies are typically valued using P/E, EV/EBITDA, or P/FCF methodologies given established revenue bases and profitability profiles.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: SiFotonics Technologies

COMPANY NAME	SiFotonics Technologies Co., Ltd.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Capital Equipment / Silicon Photonics & Optical Interconnects &
SIZE OF TOTAL ADDRESSABLE MARKET	SiFotonics targets optical interconnects for AI data centers and telecom networks. LightCounting estimates that Ethernet optical transceivers and co-packaged optics (CPO) for AI clusters will grow from approximately $16.5B in 2025 to $26B in 2026, reflecting roughly 60% growth, while noting that supply-chain constraints remain the primary near-term bottleneck.SiFotonics addresses a portion of this opportunity through its Ge-on-Si photodetectors, silicon-photonic transmit and receive PICs, optical engines, and CPO solutions. The company also participates in adjacent markets including coherent optics, PON access networks, and optical sensing.
GEOGRAPHIC PRESENCE	Woburn, Massachusetts
NOTABLE COMPETITORS	**LISTED:**AVGO, COHR, CSCO, LITE, MRVL **PRIVATE:**Ayar Labs, Lightmatter, RANOVUS, HyperLight, Xscape Photonics, SCINTIL Photonics and Quintessent
KEY SUPPLIERS / CUSTOMERS	SiFotonics combines proprietary Ge-on-Si epitaxy, in-house testing and packaging, and Tier-1 CMOS foundry partnerships through an IDM-lite model. Backed by strategic investor UMC Capital, the company reports volume deployment across 5G xHaul, data-center, coherent-optics, and PON applications, and has collaborated with Anritsu on PCIe-over-optical connectivity.
COMPETITIVE POSITIONING	SiFotonics develops silicon photonics solutions based on its Ge-on-Si platform, supported by proprietary epitaxy and long-standing production experience. Its portfolio spans photodetectors, coherent PICs, optical engines, and AI-interconnect solutions. Key differentiators include: • Ge-on-Si platform: SiFotonics' Ge-on-Si platform is built on proprietary epitaxy and long-standing production experience. • Product portfolio: The company's portfolio spans Ge-on-Si photodetectors, coherent PICs, optical engines, and AI-interconnect solutions, including 800G and emerging 1.6T platforms. • Volume deployment: SiFotonics reports years of volume deployment and cumulative shipments exceeding 30 million chips and PICs. Key risks of the company: faces increasing competition from merchant silicon-photonics ecosystems and potential geopolitical and export-control exposure associated with its China-based manufacturing and R&D footprint.
YEAR FOUNDED / ESTABLISHED	2006
WEBSITE	https://www.sifotonics.com/en/
MANAGEMENT TEAM / FOUNDER(S)	**Dong Pan - Founder and CEO:** Pan founded SiFotonics and has led the company from its early Ge-on-Si technology development through high-volume commercialization. His technical background is in silicon photonics/optoelectronics, and he has overseen development of the company's proprietary Ge-on-Si device and integration platform.
SIZE OF COMPANY	~36 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Founded in 2007, SiFotonics has grown from an MIT spinout into a global silicon-photonics supplier with operations across the U.S., Taiwan, Hong Kong, and China. The company shipped more than 5M Ge/Si optoelectronic devices in 2019, when its Nanjing manufacturing facility entered qualified volume production, including coherent receiver chips for 100G, 200G, and 400G applications. By March 2023, cumulative shipments had exceeded 30M silicon-photonics chips and PICs as the company expanded into 800G optical interconnect solutions. The company delivered ~70M silicon photonics chips as of September 2024. Its roadmap now extends to 1.6T, co-packaged optics, and optical I/O for AI-scale systems.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	SiFotonics raised $42.7MM to date. Investors: UMC Capital, Jade Stone Venture Capital
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and edge processor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and edge processor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI Networking: Xscape Photonics

COMPANY NAME	Xscape Photonics
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Capital Equipment / AI Networking / Silicon Photonics & Optical Interconnects
SIZE OF TOTAL ADDRESSABLE MARKET	Xscape addresses the AI optical-interconnect market through external lasers, silicon-photonic engines, and optical-fabric connectivity. LightCounting estimates the AI-cluster optics market at ~$26B in 2026, up from $16.5B in 2025, while a narrower forecast for AI transceivers and LPO/CPO exceeds $10B in 2026 and approaches $20B by 2030. Growth is being driven by 1.6T/3.2T optical links and the continued migration of optics closer to compute and networking silicon.
GEOGRAPHIC PRESENCE	Fort Lee, New Jersey
NOTABLE COMPETITORS	**LISTED:**AVGO, MRVL, COHR, LITE, CSCO **PRIVATE:**Ayar Labs, Lightmatter, RANOVUS, Celestial AI, HyperLight, SCINTIL Photonics and Quintessent
KEY SUPPLIERS / CUSTOMERS	Tower Semiconductor is Xscape's manufacturing partner, with CombX built on Tower's production-qualified PH18 silicon-photonics platform. Designed for interoperability with existing silicon-photonics components across CWDM and DWDM architectures, the platform leverages established photonics ecosystems. Strategic investors include NVIDIA and Cisco Investments. While customers remain undisclosed, Xscape is working with hyperscale operators through evaluation and qualification programs as it commercializes its AI-fabric photonics platform.
COMPETITIVE POSITIONING	Xscape develops programmable multi-wavelength silicon photonics designed to increase bandwidth density while reducing power consumption and component count. Its technology is built around monolithic on-chip laser integration and targets the bandwidth constraints of AI fabrics. Key differentiators include: • ChromX platform, CombX laser engine, and FalconX eight-wavelength ELSFP module: Xscape's platforms are built around monolithic on-chip laser integration, eliminating the hybrid laser attachment used in most incumbent architectures. • Multi-wavelength scaling: The company is scaling its technology from 8 to 16, 32, and ultimately 128+ wavelengths to address AI-fabric bandwidth constraints. • Pre-volume commercialization: Xscape remains at the pre-volume stage, with commercialization centered on evaluation and qualification programs. Key risks include competition from larger optical-interconnect vendors pursuing alternative CPO, external-laser, and optical-I/O approaches.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://www.xscapephotonics.com
MANAGEMENT TEAM / FOUNDER(S)	**Dr. Vivek Raghunathan — Co-Founder and CEO:**Dr. Raghunathan has 18+ years of silicon-photonics experience spanning Broadcom, Rockley Photonics, and Intel. At Broadcom, he led integrated silicon-photonics programs for co-packaged optics and contributed to the industry's first 25.6T Ethernet switch co-packaged with silicon-photonics chiplets. Previously, he helped commercialize Intel's first silicon-photonics transceivers. **Professor Alexander Gaeta — Co-Founder and President**:Gaeta is a Columbia University professor and internationally recognized expert in nonlinear and quantum photonics. His pioneering work in silicon nonlinear optics and optical frequency-comb generation provides a key scientific foundation for Xscape's multi-wavelength photonics platform. He is a Townes Medal recipient and Fellow of Optica, APS, and IEEE.
SIZE OF COMPANY	~59 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Founded in 2022 as a Columbia University spinout, Xscape Photonics is developing multi-wavelength silicon-photonics platforms for AI data-center fabrics. Xscape raised a $44M Series A in 2024, bringing total funding at the time to $57M, and expanded the round by an additional $37M in March 2026, increasing total Series A financing to $81M. Alongside the financing, the company launched FalconX, an eight-wavelength external laser module, while advancing its CombX and ChromX multi-wavelength photonics roadmap. Current priorities include hyperscaler qualification, manufacturing scale-up, and deployment of multi-wavelength optical fabrics for AI infrastructure.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Xscape has raised nearly $91MM to date in disclosed funding, including an $81M Series A completed through a $44M financing in 2024. Investors: Addition, IAG Capital Partners, NVIDIA, Cisco Investments, Altair, Fathom Fund, Kyra Ventures, LifeX Ventures, and Oxford University-related investors.
LATEST ROUND VALUATION	Valued at $100MM post Series A in Nov. 2024.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and edge processor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and edge processor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Figure 7: Edge AI Companies & Public Peers

Companies	Public Peers
Achronix Data Acceleration	AMD, INTC, LSCC, CBRS
AONDevices	AMBQ, KN, QCOM, SLAB
Axiado	SNPS, ARM, RMBS, IFX, NXPI
EdgeQ	INTC, MRVL, NXPI, QCOM
Expedera	CBRS, INDI
Gyrfalcon Technology	CBRS
Kneron	AMD, QCOM
MemryX	AMD, INTC, NVDA, QCOM
Quadric	CBRS
SiMa	AMBA, AMBQ, AMD, INTC, NVDA

Source: UBS

Edge AI: Achronix Data Acceleration

COMPANY NAME	Achronix Data Acceleration
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Edge AI / FPGA & Data Center Acceleration
SIZE OF TOTAL ADDRESSABLE MARKET	Achronix Data Acceleration operates in the FPGA, data center acceleration, AI inference, HPC, networking infrastructure, and embedded programmable logic markets. The company develops high-performance FPGA products, accelerator cards, and embedded FPGA solutions for AI, networking, and compute infrastructure.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** LSCC, INTC, AMD, CBRS **PRIVATE:** Groq, SiFive
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: Amazon, Xilinx, Intel
COMPETITIVE POSITIONING	Achronix develops FPGA and embedded FPGA solutions designed for AI inferencing, networking infrastructure, data-center acceleration, and high-performance computing applications. Key differentiators include: • Speedster FPGA platform: Achronix offers standalone Speedster FPGA devices optimized for high-bandwidth networking, AI inference, and acceleration workloads. • Embedded FPGA IP portfolio: The company licenses embedded FPGA technology that allows customers to integrate programmable logic directly into custom ASICs and SoCs. • Data-center acceleration architecture: Achronix develops accelerator cards and hardware platforms designed for HPC, AI, and cloud infrastructure deployments. • High-bandwidth on-chip networking: The company's FPGA architecture is designed around high-throughput interconnects, distributed memory resources, and low-latency processing. • AI inferencing focus: Achronix is targeting AI inference applications where FPGAs can deliver superior efficiency relative to GPUs in edge and specialized deployments.
YEAR FOUNDED / ESTABLISHED	2004
WEBSITE	https://achronix.com/
MANAGEMENT TEAM / FOUNDER(S)	**Robert Blake - President & CEO:** Prior to Achronix, Blake was CEO of Octasic Semiconductor and held sales, marketing, and general-management positions at Altera, including Vice President of Product Planning. Earlier, Blake worked on ASIC technology at LSI Logic and Fairchild and earned degrees in applied physics, electronics, business, and microelectronics from Durham University. **John Holt - Co-Founder & Chairman:** Holt co-founded Achronix in 2004 and has served as Chairman since the company's formation. Achronix was spun out of Cornell University. **Kamal Choudhary, Ph.D. - SVP of Software Engineering:** Prior to Achronix, Dr. Choudhary was a Distinguished Engineer at Xilinx, where his roles included Virtex FPGA Architect. Dr. Choudhary holds degrees from IIT Bombay and a Ph.D. in Computer Science from the University of Cambridge.
SIZE OF COMPANY	~119 employees on LinkedIn (accessed 9/18/2026)
GROWTH TRAJECTORY	In May 2025, Achronix introduced its VectorPath 815 FPGA accelerator card. In November 2025, the company partnered with Mobiveil to offer high-speed controller IP and FPGA engineering services. In August 2026, Achronix announced availability of a high-performance JESD204C IP core for its Speedster7t FPGAs. The company has also collaborated with BigCat Wireless on FPGA acceleration for 5G and 6G applications.
ANY MAJOR M&A / CORPORATE ACTIONS	In September 2022, Achronix acquired FPGA IP assets from Accolade Technology. In January 2021, Achronix agreed to merge with ACE Convergence Acquisition at an indicated $2.1B valuation, but the reverse merger and related $150MM PIPE were subsequently cancelled. Private secondary transactions were completed in 2019 and June 2026.

INVESTORS / INVESTMENT STAGE	Achronix has raised $251MM across seven completed equity funding rounds, excluding the cancelled $150MM PIPE. The company's most recent completed primary round was a $51MM Series D2 financing in December 2016. Achronix investors and shareholders include New Science Ventures, Hercules Capital, Cornell University, Battery Ventures, Desjardins Capital, Argonaut Private Equity, Easton Capital Investment Group, Entrepia Ventures, George Kaiser Family Foundation, Mortimer Sackler, and Stata Capital Partners. Several legacy investors exited through secondary transactions.
LATEST ROUND VALUATION	The latest completed primary financing with a disclosed valuation was the December 2016 Series D2 round, which valued Achronix at $139MM pre-money and $190MM post-money. The cancelled 2021 reverse merger indicated a $2.1B valuation but did not close.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For FPGA and programmable-logic companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as product cycles, customer concentration, and R&D investment can create variability in profitability. More mature FPGA and semiconductor infrastructure companies are generally valued on P/E or P/FCF given established profitability, recurring IP revenue, and broader customer diversification.
SECONDARY PRICING DATA	A private secondary transaction was completed in June 2026, when Argonaut Private Equity, Entrepia Ventures, Easton Capital Investment Group, Mortimer Sackler, George Kaiser Family Foundation, Desjardins Capital, and Stata Capital Partners sold their stakes to an undisclosed buyer. Transaction pricing and implied valuation were not disclosed.

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: AONDevices

COMPANY NAME	AONDevices, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Edge AI
SIZE OF TOTAL ADDRESSABLE MARKET	AONDevices' competes in the microcontroller (MCU) portion of the analog market. According to UBS estimates and SIA, the MCU market is valued at ~$25B/year as of 2026.
GEOGRAPHIC PRESENCE	Irvine, California
NOTABLE COMPETITORS	**LISTED:** AMBQ (ultra-low-power always-on MCU), KN (Smart-mic DSP for wake word), QCOM (always-on audio/voice DSP), SLAB (MCU ML keyword spotting) **PRIVATE:**GreenWaves (GAP9 hearables/edge-AI processor)
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry) and ARM (architecture licenses). Customers include ORCL, MSFT, GOOG, Equinix Metal, Tencent Cloud, HPE.
COMPETITIVE POSITIONING	AONDevices develops ultra-low-power processors for battery-powered consumer IoT, automotive, and industrial devices. Its processors integrate with various sensors and support low-power sensor fusion. Key differentiators include: • Ultra-low-power neural processor: The AON1100/1120 family includes a proprietary ultra-low-power neural processor designed for battery-powered devices. • Pretrained models: The AON1100/1120 processors include pretrained models for sound classification, signal denoising, and voice sensing. • Integrated acoustic sensing: OEMs can use the processors to add acoustic sensing algorithms without requiring a separate companion MCU. • Always-listening mode: The processor supports an always-listening mode that draws 20 microamperes. • Low power consumption: The 20-microampere always-listening mode undercuts competing low-power neural processors with meaningful power draw.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	www.aondevices.com
MANAGEMENT TEAM / FOUNDER(S)	**Mouna Elkhatib – CEO/CTO and Co-founder:** Mouna spent 20 years at Qualcomm/ Conexant working on audio and digital signal processing products for mobile devices. She holds 13 US patents related to audio and artificial intelligence. Mouna serves on various industry groups including the Semiconductor Alliance Women's Leadership Council. **Dr. Adil Benyassine – VP of Technology:** Dr. Benyassine spent 25 years working at Intel and Conexant as the leader in speech coding and machine learning denoising algorithms. **Daniel Schoch – VP of Engineering:** Prior to AONDevices, Daniel worked at Broadcom/ Mindspeed as an ASIC architect. He spent >20 years working on mixed-signal SoC design.
SIZE OF COMPANY	~24 employees on LinkedIn (accessed 9/2/2026)
GROWTH TRAJECTORY	AONDevices received seed funding of $2.6MM in 2019, and started licensing IP to tier-1 chip vendors plus a mobile OEM before raising its Series A. The product roadmap is focused on "full-sensor fusion" allowing voice and inertial sensor gesture controls for earbuds, smart remotes, and automotive interfaces.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	AONDevices raised $5.3MM over 2 funding rounds. In July 2025, H&SHighTech invested an undisclosed amount in AONDevices, and in June 2023 AONDevices raised $5.3MM in a seed round from undisclosed investors. Investors: H&SHighTech.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For analog/microprocessor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer analog/microprocessor companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: Axiado

COMPANY NAME	Axiado
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Edge AI / AI Security Processors & Infrastructure Security
SIZE OF TOTAL ADDRESSABLE MARKET	Axiado operates in the hardware security processor, AI infrastructure security, data-center security, zero-trust infrastructure, network management, IoT security, and edge security markets. Axiado embeds AI-driven security and system-management capabilities into processors used across heterogeneous infrastructure.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** SNPS, ARM, RMBS, IFX, NXPI **PRIVATE:** Tenstorrent, Axelera AI, d-Matrix
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Axiado develops AI-driven hardware security processors designed to secure data centers, networking equipment, telecom infrastructure, and edge systems at the silicon layer. Its architecture combines platform security, system management, and AI-driven threat protection. Key differentiators include: • Hardware-rooted security processor: Axiado embeds security functions directly into infrastructure processors, addressing vulnerabilities at the hardware and firmware layers rather than relying solely on host software. • Secure-boot and memory protection: The company's intellectual property includes secure processor boot, securely updated executable code, processor-to-memory authentication, and protected memory transactions. • Control-plane threat protection: Axiado develops network-access proxy and real-time AI technologies intended to protect control and management planes from zero-day attacks. • AI-driven infrastructure management: The platform combines security with autonomous management capabilities for heterogeneous AI, networking, telecom, and edge infrastructure. • Inside-the-rack deployment model: Commercialization is centered on integrating cybersecurity into server and rack-level infrastructure, positioning security as part of the underlying compute platform.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://axiado.com
MANAGEMENT TEAM / FOUNDER(S)	**Gopi Sirineni - Founder, President & CEO:** Sirineni has more than 25 years of experience across the semiconductor, software, and systems industries. Earlier in his career, Sirineni worked at Qualcomm on wireless technologies associated with Bluetooth devices and secure Wi-Fi mesh networking. **Raghu Kondapalli - CTO:** Kondapalli has more than 25 years of systems, chip, and software-architecture experience across Intel, LSI, Marvell, and Nokia. Kondapalli has contributed to SoCs, AI/ML platforms, networking switches, and computing systems and holds more than 40 patents at various stages. **John Donovan - Strategic Advisor:** Donovan is a former CEO of AT&T Communications and advises Axiado on telecom, networking, and digital-infrastructure strategy.
SIZE OF COMPANY	~156 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Axiado raised capital in 2024 and 2025 to support continued development of hardware-anchored platform security, system management, and energy-efficiency solutions for AI data centers, networking, telecom, and edge infrastructure. During 2026, the company demonstrated autonomous AI infrastructure management and received industry recognition for zero-trust, enterprise security infrastructure, and AI-infrastructure cybersecurity.

ANY MAJOR M&A / CORPORATE ACTIONS	Axiado acquired APO'G, an IT consulting and outsourcing company, in July 2025. In a separate secondary transaction, 3Spoke Capital acquired stakes from Rocket Strategies, IIIT-H Foundation, and Richard Kramlich.
INVESTORS / INVESTMENT STAGE	Axiado has raised $236MM across five primary equity rounds. Before its latest estimated $40.5MM Series C4 round in April 2026, the company raised $100MM in a Series C3 round in December 2025 and $60MM in a Series C round in August 2024 per Pitchbook. Axiado investors include AMD Ventures, Maverick Silicon, Prosperity7 Ventures, Crosslink Capital, Orbit Venture Partners, Maverick Capital, Samsung Catalyst Fund, Atreides Management, NosTerra Venture Capital, SanTuri Ventures, David Welch, Dennis Uy, Qudit, and Western Technology Investment. The company's current investment stage is Series C.
LATEST ROUND VALUATION	Axiado completed an estimated $40.5MM Series C4 financing in April 2026 at a $684.5MM pre-money valuation and an estimated $725MM post-money valuation. Series C4 preferred shares were priced at $8.64 per share.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For hardware-security processor and AI-infrastructure security companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary while customers qualify and deploy new infrastructure platforms. More mature semiconductor-security peers are generally valued on P/E or P/FCF, while recurring-software cybersecurity peers may also be assessed on EV/ARR or EV/revenue.
SECONDARY PRICING DATA	3Spoke Capital acquired stakes from Rocket Strategies, IIIT-H Foundation, and Richard Kramlich in a private secondary transaction. Pricing and implied valuation were not disclosed.

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: EdgeQ

COMPANY NAME	EdgeQ Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Edge AI
SIZE OF TOTAL ADDRESSABLE MARKET	EdgeQ produces a software-defined base station-on-a-chip that converges 5G communications and edge AI processing on one platform. It targets telecom operators and enterprises deploying private 5G networks with built-in AI.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** INTC, MRVL, NXPI, QCOM **PRIVATE:** NXP (small cell RAN SoCs), Picocom (small cell SoC)
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry), ARM (Neoverse CPU cluster), Andes Technology (RISC-V IP), Arteris (NoC IP) Customers include Actiontec Electronics (small cell based on EdgeQ SoC), Airspan (small cell portfolio for MNOs), and other undisclosed customers.
COMPETITIVE POSITIONING	EdgeQ designs programmable RISC-V system-on-chips for base stations, integrating multiple communications and compute functions into a single platform. Key differentiators include: • Programmable RISC-V system-on-chips: EdgeQ's programmable RISC-V SoCs consolidate 4G/5G baseband processing, AI inference, and networking functions for base stations.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://www.edgeq.io/
MANAGEMENT TEAM / FOUNDER(S)	**Vinay Ravuri - Founder and CEO:** Prior to EdgeQ, Mr. Ravuri spent fifteen years at Qualcomm, rising to Vice President and launching its datacenter server and machine learning businesses. Mr. Ravuri holds an MS in Electrical & Computer Engineering from Georgia Tech and is inventor on twelve US patents. **Adil Kidwai - Chief Product Officer:** Prior to EdgeQ, Mr. Kidwai was lead modem architect and later Director of AI / ML Engineering at Intel, overseeing 4G / 5G IP and datacenter inference platforms. Mr. Kidwai received an MS in Electrical Engineering from UCLA, and an MBA from UC Berkeley. **Sriram Rajagopal - Head of Technology:** Prior to EdgeQ, Mr. Rajagopal was a Distinguished Engineer at Altiostar Networks and later Head of Physical Layer Systems at Beceem Communications. Mr. Rajagopal has 20 issued patents and 7 Conference Publications, holding a Masters Degree in Electrical Engineering from Stanford University.
SIZE OF COMPANY	~125 employees on LinkedIn (accessed 9/2/2026)
GROWTH TRAJECTORY	In February 2026, EdgeQ announced its baseband processor is used in AirUnity Small Cells, a small cell portfolio designed for MNOs and enterprises in 4G/5G coverage. In February 2025, EdgeQ partnered with BATS wireless on the ECHO 401 and ECHO 405 TAA-Compliant, Category A & B radios designed for indoor and outdoor CBRS deployments. In March 2024, EdgeQ's 4G / 5G base-station-on-a-chip entered revenue shipments and headlined MWC Barcelona with split-6 neutral-host demos for DenseAir / Radisys and a Wi-Fi 7-enabled small-cell platform from Actiontec, underscoring readiness for commercial private-network rollouts. In April 2023, the company strengthened its balance sheet with a $75 million Series B round to fund volume ramps and next-gen SoCs.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	EdgeQ raised $126MM over 2 funding rounds. In its latest Series B round of funding, the company received $75MM in April of 2023. Investors: Yaletown Partners, Strategic Development Fund, ST Engineering, Iron Grey, EDBI, ClearSky, 5G Ventures, Threshold, Fusion Ventures, AME Cloud Ventures
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: Expedera

COMPANY NAME	Expedera
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Edge AI IP & Neural Processing Units
SIZE OF TOTAL ADDRESSABLE MARKET	Expedera operates in the neural processing unit IP, edge AI inference, generative AI at the edge, automotive AI, robotics, computer vision, and semiconductor-IP licensing markets. Expedera supplies customizable neural-engine IP and deployment software for integration into customer SoCs.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** CBRS, INDI **PRIVATE:** Graphcore, SambaNova Systems, Mythic, Tenstorrent, Hailo, EnCharge AI
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Expedera develops customizable neural-engine semiconductor IP designed to improve inference performance, latency, power consumption, and implementation complexity across edge AI applications. The company combines licensable NPU architectures with a supporting software toolchain. Key differentiators include: • Origin neural-engine IP: Expedera licenses neural-processing architectures that semiconductor companies can integrate into custom SoCs rather than requiring customers to adopt a standalone accelerator chip. • Origin Evolution generative-AI platform: The company introduced Origin Evolution NPU IP in 2025 to extend its architecture into generative-AI workloads running on edge devices. • Integrated model-deployment toolchain: Expedera provides model import, quantization, automated completion, compilation, estimation, and profiling capabilities to simplify implementation. • Pipelined native execution: The architecture runs neural-network models natively through a pipelined design intended to reduce software complexity, latency, and data-movement overhead. • Automotive-qualified IP: Expedera achieved ISO 26262 ASIL-B automotive safety certification, supporting integration into safety-sensitive automotive semiconductor applications.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://www.expedera.com/
MANAGEMENT TEAM / FOUNDER(S)	**Siyad Ma, Ph.D. - Co-Founder & CEO:** Prior to becoming CEO, Dr. Ma served as Expedera's Vice President of Engineering and led development of its packet-based NPU architecture and software stack. Dr. Ma previously led ASIC and IP teams at Cisco and held senior positions at Spanslogic, Zettacom, Chameleon, and AMD; Dr. Ma holds a Ph.D. in Electrical Engineering from Stanford University. **Da Chuang, Ph.D. - Co-Founder & Executive Chairman:** Prior to Expedera, Dr. Chuang co-founded and served as COO of Memoir Systems, which was acquired by Cisco, and subsequently led data-center switch ASIC development at Cisco. Dr. Chuang also held roles at Nvidia and Abrizio and holds a B.S. from UC Berkeley and an M.S. and Ph.D. from Stanford University. **Sharad Chole - Co-Founder & Chief Scientist:** Prior to Expedera, Chole was an architect at Cisco, Memoir Systems, and Microsoft. Chole specializes in AI frameworks, power-aware neural-network optimization, and programmable dataflow architectures and holds a B.Tech. from IIT Kanpur.
SIZE OF COMPANY	~45 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In June 2024, Expedera opened an R&D office in India. In July 2024, the company achieved ISO 26262 ASIL-B automotive safety certification. In April 2025, Expedera achieved ISO 9001:2015 quality-management certification, and in May 2025 it introduced Origin Evolution NPU IP to support generative AI on edge devices.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Expedera has raised $48MM across three funding rounds. Before its latest estimated $21.4MM Series B round in March 2024, the company raised $18MM in its Series A round in December 2021 and $9MM in Seed funding per Pitchbook. Expedera investors include indie Semiconductor, Alibaba Entrepreneurs Fund, WI Harper Group, Ignite Innovation, Sehat Sutardja, and Weili Dai. The company's current investment stage is Series B.
LATEST ROUND VALUATION	Expedera completed an estimated $21.4MM Series B financing in March 2024 at an estimated $194MM pre-money valuation and $215.4MM post-money valuation.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For neural-processing IP and edge-AI semiconductor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, reflecting limited near-term profitability and the importance of licensing adoption. More mature semiconductor-IP companies are generally valued on P/E or P/FCF given their higher-margin royalty models and more predictable profitability.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: Gyrfalcon Technology

COMPANY NAME	Gyrfalcon Technology
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Edge AI Accelerators
SIZE OF TOTAL ADDRESSABLE MARKET	Gyrfalcon Technology operates in the edge AI inference, AI accelerator, embedded AI processor, IoT semiconductor, mobile AI, biometric processing, and edge-server markets. Gyrfalcon Technology develops processors intended to move AI inference from cloud infrastructure to local devices.
GEOGRAPHIC PRESENCE	Milpitas, California
NOTABLE COMPETITORS	**LISTED:** CBRS **PRIVATE:** Graphcore, Mythic, SambaNova Systems, Hailo, Kneron
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: LG Electronics
COMPETITIVE POSITIONING	Gyrfalcon Technology develops AI accelerator processors designed to extend cloud-based AI capabilities to local devices through power-efficient edge inference. Its architecture focuses on reducing data movement and memory bottlenecks in neural-network processing. Key differentiators include: • APLAI accelerator architecture: Gyrfalcon develops dedicated AI processors for executing neural-network inference across mobile, IoT, embedded, and edge-server environments. • Memory-centric processing: The company's technology is designed to reduce power-consuming data movement and address memory bottlenecks that constrain conventional AI processors. • High TOPS-per-watt performance: Gyrfalcon has emphasized inference efficiency, including a chip announced at 12.6 TOPS/W and the 2803 Plai development platform at 24 TOPS/W. • Chip-to-server product range: The company has expanded beyond processors into development hardware and edge-inference servers, including the Lacelli platform and a server developed with SolidRun. • Mobile and biometric commercialization: Gyrfalcon's AI chip was used in an LG smartphone, and the company partnered with Sensory to provide biometric AI functionality on edge devices.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://gyrfalcontech.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Kimble Dong - Co-Founder & Co-CEO:** Dong co-founded Gyrfalcon and serves as Co-Chief Executive Officer, leading the company's edge-AI processor commercialization activities. **Yasuo Nishiguchi - Co-CEO & Chairman:** Nishiguchi serves as Co-Chief Executive Officer and Chairman, with responsibility for the company's leadership and board oversight. **Lin Yang, Ph.D. - Co-Founder & Chief Scientist:** Yang co-founded Gyrfalcon and leads the company's AI-processing research and semiconductor architecture development.
SIZE OF COMPANY	~16 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In 2019, Gyrfalcon introduced AI accelerator hardware emphasizing 12.6 TOPS/W and 24 TOPS/W performance, launched its Lacelli edge-inference server, and supplied an AI chip used in an LG smartphone. The company also partnered with Sensory on edge biometric access. In February 2020, Gyrfalcon and SolidRun announced an edge-optimized AI inference server designed to provide lower power consumption than GPU-based alternatives.
ANY MAJOR M&A / CORPORATE ACTIONS	In January 2026, Alumni Ventures sold its stake in Gyrfalcon to an undisclosed buyer through a private secondary transaction.

INVESTORS / INVESTMENT STAGE	Gyrfalcon has raised $68MM across three primary equity rounds. Before its latest primary round, the company raised an estimated $8MM in Series B financing and an estimated $10MM in Series A financing. Gyrfalcon investors include LG Electronics, Sailing Capital, V-Capital, Plug and Play, and former investor Alumni Ventures. The company's latest disclosed primary investment stage is Series C.
LATEST ROUND VALUATION	Gyrfalcon completed an estimated $50MM Series C financing in August 2018 at a $300MM pre-money valuation and $350MM post-money valuation. Series C preferred shares were priced at $3.08 per share per Pitchbook.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For edge-AI accelerator and inference semiconductor companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary with customer adoption and device-design cycles. More mature AI semiconductor companies are generally valued on P/E or P/FCF once shipment volumes, gross margins, and customer demand become more predictable.
SECONDARY PRICING DATA	A private secondary transaction occurred in approximately January 2026 when Alumni Ventures sold its stake to an undisclosed buyer. Pricing and implied valuation were not disclosed.

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: Kneron

COMPANY NAME	Kneron
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Edge AI
SIZE OF TOTAL ADDRESSABLE MARKET	Kneron has developed a proprietary edge AI computing architecture centered on its Neural Processing Units (NPUs), enabling high-performance AI inference directly on devices with minimal power consumption. The company's integrated hardware and software platform supports computer vision, facial recognition, automotive, security, and AIoT workloads without requiring constant cloud connectivity. Kneron operates in the rapidly expanding edge AI semiconductor market, where demand is being driven by increased adoption of intelligent cameras, autonomous systems, industrial automation, smart devices, and sovereign AI applications.
GEOGRAPHIC PRESENCE	San Diego, California
NOTABLE COMPETITORS	**LISTED:** AMD, QCOM **PRIVATE:** Hailo, Axelera AI, Kinara, Mythic Kneron competes in the edge AI accelerator and neural processing unit (NPU) market, where power efficiency, cost, and on-device inferencing are critical requirements. Unlike many AI hardware companies targeting data centers, Kneron focuses primarily on computer vision, facial recognition, AIoT, security, automotive, and edge server applications. According to management, Kneron's processors are designed to deliver substantially lower power consumption, smaller form factors, and lower costs than competing edge AI solutions. The company competes against both traditional embedded semiconductor suppliers and newer AI accelerator startups focused on on-device inference.
KEY SUPPLIERS / CUSTOMERS	Customers: Chamberlain Group, NHN, Toyota, JVC Kenwood, Hanwha Vision, Vivotek, Chungwa Telecom, Comcast, Quanta
COMPETITIVE POSITIONING	Kneron's platform is centered on its proprietary Neural Processing Unit (NPU) architecture. Key differentiators include: • Edge AI focus: Kneron develops AI accelerators optimized for on-device inferencing, reducing latency, bandwidth requirements, and cloud dependency. • Power-efficient NPU architecture: Management states that Kneron's solutions can provide lower power consumption, reduced silicon area, and lower system costs compared with competing AI processors targeting similar workloads. • Computer vision specialization: The company focuses primarily on visual intelligence applications including facial recognition, security analytics, automotive driver monitoring, and machine vision. • Broad product portfolio: Current products include the KL520, KL530, KL720, KL730, and newer edge AI platforms, spanning from low-power embedded AI applications to edge servers capable of running larger language models. • Edge server strategy: Management disclosed that higher-end Kneron solutions can support models with up to approximately 120 billion parameters through multi-chip scaling, expanding the company's addressable market beyond traditional embedded AI use cases.
YEAR FOUNDED / ESTABLISHED	2015
WEBSITE	https://www.kneron.com

MANAGEMENT TEAM / FOUNDER(S)	**Albert Liu – Founder & CEO:** Prior to founding Kneron in 2015, Dr. Liu served as an Engineering Director at Qualcomm, where he led development of mobile AI, computer vision, and multimedia technologies. Earlier in his career, he held engineering and research positions at Samsung, Steering Solutions, and National Taiwan University. Dr. Liu received a Ph.D. in Electrical Engineering from the University of California, Los Angeles (UCLA) and has extensive experience in AI processor architecture, computer vision, and embedded systems. **Frank Chang – Co-Founder & Chairman:** Dr. Chang is a semiconductor industry veteran and former Chairman and CEO of TSMC. He co-founded Kneron to support development of next-generation edge AI processors and serves as Chairman of the company. Dr. Chang received the 2025 John Fritz Medal, one of the highest honors in engineering, recognizing his contributions to the global semiconductor industry. **Roger H. Chen – Co-Founder & President:** Prior to Kneron, Mr. Chen held senior leadership positions across the semiconductor and technology industry, including roles focused on business development, product strategy, and operations. At Kneron, he oversees global business operations, strategic partnerships, and commercialization of the company's edge AI platform.
SIZE OF COMPANY	~96 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Founded in 2015, Kneron has evolved from an edge AI startup focused on computer vision into a provider of full-stack AI hardware and software platforms for automotive, security, AIoT, and edge server deployments. In 2021, the company raised a $48.9 million Series B financing led by WRVI Capital to accelerate commercialization of its AI inference technology. In February 2024, Kneron completed a $102 million Series C financing co-led by Celesta Capital and GreatPoint Ventures to support development of next-generation AI inference systems and expand deployments across edge AI and enterprise applications. More recently, management disclosed that Kneron has shipped millions of AI processors and continues to expand adoption across automotive, security, AIoT, and edge infrastructure customers. The company has also broadened its product portfolio from low-power AI accelerators to higher-performance edge AI platforms and edge servers, while developing next-generation processors for larger AI workloads. Management expects continued growth driven by increasing demand for on-device AI inference, visual recognition, and sovereign AI deployments.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Kneron has raised a total of ~$521MM over 10 funding rounds. In its latest round of funding, the company received $300MM in Oct of 2024. Notable Investors: Horizons Ventures, Foxconn, Lite-On Technology, Adata, Qualcomm, Alibaba Entrepreneurs Fund, Sequoia Capital, Himax Technologies, Delta Electronics, DNS Capital, Mayfield Fund, Foundation Capital, MTS.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: MemryX

COMPANY NAME	MemryX, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Edge AI
SIZE OF TOTAL ADDRESSABLE MARKET	MemryX produces low-power AI inference chips designed for edge devices, emphasizing a memory-centric compute architecture. These chips aim to efficiently run vision and sensor neural networks in automobiles, cameras, and IoT devices. Overall, the edge AI hardware market is expected to grow to $59 billion by 2030.
GEOGRAPHIC PRESENCE	Ann Arbor, Michigan
NOTABLE COMPETITORS	**LISTED:** AMD, INTC, NVDA, QCOM **PRIVATE:** Hailo, SiMa.ai, Axelara AI, EdgeCortix, Kneron, Blaize
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry). Key customers include ASUS (AIoT), DYNICS, and Lenovo
COMPETITIVE POSITIONING	MemryX develops a streaming, near-memory NPU designed to deliver low-latency AI inference at very low power in compact form factors such as M.2. Key differentiators include: • Streaming, near-memory NPU: MemryX's NPU combines a streaming architecture with on-chip parameter storage to perform low-latency inference at very low power. • Compiler: The company's compiler automatically converts standard models into its dataflow programs, allowing developers to use ONNX models with minimal integration. • Scalability: MemryX's architecture can scale by adding and connecting chips. • Third-party verified performance: The platform enables developers to easily obtain third-party verified performance.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	www.memryx.com
MANAGEMENT TEAM / FOUNDER(S)	**Ross Jatou - CEO:** Prior to MemryX, Mr. Jatou was President of Alat Technologies. Before Alat Technologies, Mr. Jatou held various roles at Onsemi for 8 years and NVDA for 15 years. **Dr. Wei Lu - Co-founder and CTO:** Dr. Lu is a distinguished IEEE Fellow and has been a Professor in the EECS department at University of Michigan since 2005. He co-founded MemryX in 2019 and previously co-founded Crossbar in 2010, a resistive RRAM technology company.
SIZE OF COMPANY	~47 employees on LinkedIn (accessed 9/2/2026)
GROWTH TRAJECTORY	In June 2026, MemryX announced the expansion of its Cascade deployment platform family, with three new products built on the MX3 accelerator architecture. These products include the Cascade 100P PCIe Accelerator, 100U USB Accelerator, and 100R Raspberry Pi HAT+. The company also entered a strategic partnership with National Semiconductor Hub of Saudi Arabia to accelerate development of AI in May 2025. The company said more than a dozen companies have been selected to enter production with its MX3 platform as of the recent Series B funding round in March 2025.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	MemryX raised ~$44MM over 4 funding rounds. Before its latest venture round of undisclosed size, the company received $44MM in its Series B round of funding in March 2025. MemryX investors include NEOM Investment Fund, M Ventures, eLab Ventures, ARM, Motus Ventures, HarbourVest Partners, and the University of Michigan MINTS fund.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: Quadric

COMPANY NAME	Quadric
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Edge AI Processor IP
SIZE OF TOTAL ADDRESSABLE MARKET	Quadric operates in the edge AI processor-IP, NPU, automotive AI, autonomous-driving compute, robotics, industrial automation, AI PC, IoT, and embedded-AI semiconductor markets. Quadric licenses a unified processor architecture for integration into customer SoCs.
GEOGRAPHIC PRESENCE	Burlingame, California
NOTABLE COMPETITORS	**LISTED:** CBRS **PRIVATE:** Graphcore, SambaNova Systems, Mythic, Hailo, Tenstorrent, MIPS
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: TIER IV, NSITEXE.
COMPETITIVE POSITIONING	Quadric develops a General-Purpose Neural Processing Unit architecture designed to combine AI inference and general-purpose computing within one licensable processor platform. The technology targets automotive, robotics, industrial, AI PC, and embedded-edge applications. Key differentiators include: • GPNPU unified architecture: Quadric combines neural-network acceleration and general-purpose computing in a common processor architecture, reducing the need for separate CPU, DSP, and NPU domains. • Unified software model: The platform supports standard programming languages and AI graph formats, allowing AI and conventional workloads to be developed within a common environment. • Licensable processor IP: Quadric supplies semiconductor IP for integration into customer chips rather than depending exclusively on sales of standalone accelerator devices. • Configurable multicore execution: The company's patents cover configurable multicore arrays, directional operand broadcasting, multiply-accumulate execution, fetchless computation, and computation-aware memory transfers. • Automotive and edge design-win strategy: Commercialization is centered on automotive processors, Autoware deployment, robotics, enterprise edge AI, and local execution of large language and vision-language-action models.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://quadric.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Veerbhan Kheterpal, Ph.D. - Co-Founder & CEO:** Prior to Quadric, Dr. Kheterpal co-founded bitcoin-computing company 21, Inc. and semiconductor-software company Fabbrix, which was acquired by PDF Solutions. Dr. Kheterpal has experience spanning ASIC design, data-center systems, blockchain infrastructure, and consumer applications and holds a Ph.D. from Carnegie Mellon University. **Nigel Drego, Ph.D. - Co-Founder & CTO:** Prior to Quadric, Dr. Drego was Co-Founder and Chief Architect of 21, Inc., where Dr. Drego developed ASICs, HPC systems, and supporting software. Earlier, Dr. Drego worked at PDF Solutions and Intel and now leads Quadric's technology stack from software through silicon. **Daniel Firu - Co-Founder & CPO:** Firu co-founded Quadric and is responsible for product strategy and development of the company's processor-IP offering.
SIZE OF COMPANY	~74 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In January 2026, Quadric announced that its SDK had been selected by TIER IV for AI-processing evaluation and optimization supporting Autoware deployment. The same month, Quadric began joint development of high-performance automotive processor IP with NSITEXE and reported accelerating design wins across edge LLM, automotive, and enterprise applications. In May 2026, the company stated that it ran the Pi-0.5 vision-language-action model 6–10x more efficiently than NVIDIA Jetson Thor.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Quadric has raised $84MM across four funding rounds. Before its latest $46MM Series C financing, the company raised $21MM in Series B funding in December 2022, $15MM in Series A funding in May 2019, and $2.25MM in Seed funding in March 2017. Quadric investors include International Finance Corporation, BEENEXT Capital Management, Wanxiang America, Orchid Asia Group Management, Volta Energy Technologies, Pear, Uncork Capital, Silicon Catalyst Ventures, Gentree, Offline Ventures, Pivotal Ventures, Denso, NSITEXE, MegaChips America, Emerson Electric, Cota Capital, SV Angel, and Trucks Venture Capital. The company's current investment stage is Series C.
LATEST ROUND VALUATION	Quadric completed a $46MM Series C financing in January 2026 at a $254MM pre-money valuation and $300MM post-money valuation. The round included an initial $30MM close in January and a $16MM second close in July 2026.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For edge-AI processor-IP and semiconductor-licensing companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may remain limited during the design-win and customer-integration stages. More mature semiconductor-IP businesses are generally valued on P/E or P/FCF given recurring licensing and royalty revenue, high incremental margins, and more established profitability.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Edge AI: SiMa Technologies

COMPANY NAME	SiMa Technologies, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Edge AI
SIZE OF TOTAL ADDRESSABLE MARKET	SiMa.ai develops full-stack Physical AI solutions built around its proprietary Machine Learning System-on-Chip (MLSoC) architecture, enabling real-time AI inference for robotics, industrial automation, autonomous systems, automotive, aerospace, healthcare, and intelligent edge infrastructure. Rather than focusing solely on neural network acceleration, the company's platform is designed to accelerate the entire compute pipeline while maximizing performance-per-watt and simplifying deployment through an integrated hardware and software stack. SiMa.ai operates in the rapidly growing edge AI and Physical AI markets, where customers increasingly require low-latency, power-efficient AI processing outside traditional cloud data centers.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** AMBA, AMBQ, AMD, INTC, NVDA **PRIVATE:** Hailo, d-Matrix, Mythic, Groq, Encharge AI SiMa.ai competes in the embedded and edge AI semiconductor market, where power efficiency, ease of deployment, and real-time inferencing are critical requirements. Unlike many AI hardware vendors focused primarily on data-center acceleration, SiMa.ai targets Physical AI workloads that require devices to perceive, reason, and act in real time across robotics, industrial automation, autonomous systems, and intelligent infrastructure deployments.
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry), ARM (IP), SNPS, and Arteris. Customers include TRUMPF (laser systems company).
COMPETITIVE POSITIONING	SiMa.ai's platform is centered on its Modalix™ MLSoC™ architecture. Key differentiators include: • Machine Learning System-on-Chip (MLSoC™): SiMa.ai accelerates the entire compute pipeline rather than only the machine-learning workload, enabling developers to deploy AI applications through a unified hardware and software architecture. • Performance-per-watt optimization: The company focuses on maximizing AI performance while maintaining low power consumption, targeting robotics, industrial automation, autonomous vehicles, and other edge deployments where efficiency is critical. • Physical AI focus: SiMa.ai positions its platform for physical AI applications in which real-world systems must perceive, reason, and act in real time. The architecture supports increasingly complex workloads, including Large Language Models (LLMs) and Vision Language Models (VLMs) at the edge. • Integrated hardware and software stack: The company's Palette™ software platform enables customers to deploy AI applications across existing systems while simplifying the transition from prototype to production environments. • Production-ready modules: SiMa.ai offers System-on-Modules (SoMs) designed for rapid integration into robotics, industrial automation, and autonomous systems, reducing deployment complexity and time to market.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://sima.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Krishna Rangasayee – Founder & CEO:** Prior to founding SiMa.ai, Mr. Rangasayee held leadership positions at Xilinx, including Executive Vice President of Global Sales, before joining Groq as Chief Operating Officer. He earned a master's degree from Mississippi State University and an MBA from Stanford University.
SIZE OF COMPANY	~245 employees on LinkedIn (accessed 8/19/2026)

GROWTH TRAJECTORY	SiMa.ai has expanded from an edge AI semiconductor startup into a provider of Physical AI infrastructure spanning robotics, industrial automation, autonomous systems, and intelligent edge deployments. The company established partnerships with global technology providers including Arm, Synopsys, Cisco, Wind River, and TRUMPF while expanding adoption of its MLSoC platform across commercial markets. In April 2026, SiMa.ai announced a strategic investment from Micron Technology and expanded its collaboration around tightly integrated compute and memory architectures designed for Physical AI applications. The company also introduced production-ready System-on-Modules featuring Micron LPDDR5X memory and continues to position its Modalix™ platform as a scalable solution for robotics, autonomous systems, industrial automation, aerospace, healthcare, and next-generation edge AI deployments..
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Sima.ai raised ~$505MM over 5 funding rounds. In its latest Series C round of funding, the company received $150MM in September of 2026. Notable Investors: Fidelity Management & Research Company, Maverick Capital, MSD Partners, VentureTech Alliance, Point72, Jericho Capital, Dell Technologies Capital, Lip-Bu Tan, Storm Ventures, Micron Technology
LATEST ROUND VALUATION	$1.45B as of Series C
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and analog companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and analog companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Figure 8: Physical AI Companies & Public Peers

Companies	Public Peers
Atomic Machines	DM, PRLB
Bright Machines	JBL, FLEX, CLS, 2317 TW, ABB, FANUC, ROK, TER
Dexterity	ABB, FANUC, ROC, TER, SYM, ZBRA
Figure AI	KSCP
Lumotive	COHR, LITE, MVIS, OUST
Metalenz	WOLF, STM
Mojo Vision	VUZI
Omnitron Sensors	ADI, STM, TXN, IFX, NXPI
Path Robotics	LECO, ITW, ABB, 6954 T, 6506 T
Tacta Systems	ABB, 6954 T, 6506 T, TER, ROK

Source: UBS

Physical AI/Robotics: Atomic Machines

COMPANY NAME	Atomic Machines
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Physical AI / Advanced Manufacturing & MEMS Production Platforms
SIZE OF TOTAL ADDRESSABLE MARKET	Atomic Machines operates in the advanced manufacturing, MEMS fabrication, industrial automation, nanotechnology, semiconductor manufacturing equipment, and robotics-enabled production markets. The company develops robotic manufacturing platforms designed for micro-electromechanical systems (MEMS) and next-generation device production.
GEOGRAPHIC PRESENCE	Emeryville, California
NOTABLE COMPETITORS	**LISTED:**proto **PRIVATE:** Carbon, Formlabs, Figure, Physical Intelligence, Chef Robotics
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Atomic Machines develops robotic manufacturing platforms designed to create a new class of MEMS and micro-scale devices. Key differentiators include: • Robotic manufacturing architecture enabling rapid prototyping and scaled production. • MEMS-focused production platform designed to overcome conventional manufacturing constraints. • Nanotechnology and micro-device fabrication capabilities. • Flexible manufacturing approach intended to improve economics across prototyping and volume production. • Growing intellectual property portfolio covering MEMS actuators, electrical contacts, relays, and micro-manufacturing technologies.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	https://atomicmachines.com/
MANAGEMENT TEAM / FOUNDER(S)	**Jeff Holden - Co-Founder & CEO:** Prior to Atomic Machines, Holden was Chief Product Officer at Uber and previously led product development initiatives at Amazon. He founded Atomic Machines to build a new manufacturing paradigm for next-generation hardware production. **Prashant Patil, Ph.D. - Co-Founder & Advisor:** Dr. Patil is a serial entrepreneur and deep-technology executive with extensive experience in MEMS and advanced semiconductor manufacturing. **Klaus Zietlow - CTO:** Zietlow leads development of Atomic Machines' manufacturing platform and related hardware technologies.
SIZE OF COMPANY	~229 employees on LinkedIn (accessed 9/18/2026)
GROWTH TRAJECTORY	Atomic Machines continues developing robotic manufacturing technologies focused on MEMS and advanced hardware production. Recent patent activity includes MEMS actuators with magnetic latching, digitally controlled circuit breakers utilizing MEMS relays, and advanced electrical contact technologies. The company also appeared in PitchBook's Robotics & Physical AI, DefenseTech, and Advanced Computing market maps
ANY MAJOR M&A / CORPORATE ACTIONS	Countdown Capital exited its position through the company's February 2025 financing.
INVESTORS / INVESTMENT STAGE	Atomic Machines has raised $144MM across two disclosed venture rounds. Prior to its latest $75MM Series B financing completed in February 2025, the company raised approximately $69MM in Series A financing in November 2022. Atomic Machines investors include XTX Ventures, Sozo Ventures, Valor Equity Partners, Deep Ventures, E14 Fund, GigaFund, Jana Messerschmidt, Adrian Aoun, Erik Peterson, and other strategic investors. Current investment stage is Series B.

LATEST ROUND VALUATION	Atomic Machines completed a Series B financing in February 2025 at an estimated $300MM pre-money valuation and estimated $375MM post-money valuation per Pitchbook.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For advanced manufacturing, robotics-enabled production platforms, and semiconductor manufacturing technology companies that are still scaling commercialization, we think it would be appropriate to value these companies on P/S or EV/S. More mature industrial automation and semiconductor equipment companies are generally valued on P/E or P/FCF given established profitability and shipment visibility.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Bright Machines

COMPANY NAME	Bright Machines, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Physical AI (SW only)
SIZE OF TOTAL ADDRESSABLE MARKET	Bright Machines provides a full-stack, software-defined automation solution for electronics manufacturing, delivered as modular 'microfactories' built from Bright Robotic Cells (BRCs) and orchestrated by its Brightware operating system. The company now positions itself as a next-generation manufacturer bringing AI and data center infrastructure production to the edge, rather than purely as an automation vendor. The addressable opportunity is the final assembly, inspection and test stage of electronics manufacturing - historically manual, labor-intensive and concentrated in low-cost geographies. Current demand is concentrated in AI infrastructure: GPU module installation, complex server assembly, fastening, labeling and end-of-line inspection. Demand drivers cited by the company are the reshoring movement, skilled labor shortages, supply chain disruption and geopolitical pressure to localize production, layered on top of the exponential build rate required for AI data center infrastructure.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** JBL (also an investor), FLEX, CLS, Foxconn/Hon Hai (2317.TW), ABB, FANUC (6954.T), ROK, TER **PRIVATE:** eBots, Tacta Systems, Path Robotics, Trener Robotics, Machina Labs, and internal automation programs at large EMS/ODM providers
KEY SUPPLIERS / CUSTOMERS	Strategic technology partners include NVIDIA (Bright Designer is built on NVIDIA Omniverse libraries and APIs) and Microsoft (integration and go-to-market partnership with Microsoft Azure delivering a full-stack environment across the manufacturing lifecycle, from product design through assembly to disassembly). Jabil is both an investor and an EMS partner. Restronics acts as a manufacturers' representative in the US Northeast and Southeast. Individual customers are not publicly named. The company reports serving more than 60 lifetime customers and strategic partners through both direct OEM relationships and EMS/ODM partners, and has cited hyperscale data center server output exceeding 300,000 units.

COMPETITIVE POSITIONING	Bright Machines' differentiator is a full-stack combination of modular robotic hardware, a manufacturing operating system and a unified production data thread - sold into the AI infrastructure build-out. Key attributes include: • Microfactory architecture: Modular Bright Robotic Cells combined with computer vision and automation software create programmable, reconfigurable assembly lines deployable faster than traditional fixed automation, without the long lead times of conventional line engineering. • Brightware operating system and Smart Skills: Brightware orchestrates robots, conveyors and inspection systems and gives operators the ability to run, monitor and control assembly lines. Smart Skills — described by management as the company's 'brain' — provide adaptable 3D assembly navigation and ML-based inspection. • Design-through-disassembly coverage: Bright Designer, launched March 2025 on NVIDIA Omniverse, extends the stack upstream into product design, providing Design for Automated Assembly (DFAA) recommendations and simulating automated assembly of CPU- and GPU-based server designs across CAD and PLM platforms. Flexible disassembly capability supports component harvesting and recycling at end of life. • Hybrid BRC (July 2026): Allows human operators to perform prescribed assembly steps inside a sensor-monitored robotic cell when manual intervention is required, preserving the production record end to end — addressing the industry's prior choice between stopping the line or moving units to a disconnected manual station. • Quantified customer outcomes: The company has cited up to 75% more units per hour and 65% fewer defects per production run, alongside potential carbon intensity reduction of 51% per unit. • Strategic investor alignment: NVIDIA, Microsoft, Jabil and BlackRock on the register provide technology, channel and capital alignment directly to the AI infrastructure supply chain.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://www.brightmachines.com
MANAGEMENT TEAM / FOUNDER(S)	**Chris Stori – Chief Executive Officer (appointed September 2024):** Prior to Bright Machines, Mr. Stori spent over 11 years at Meraki and Cisco, where he led strategy and development for enterprise networking. Earlier at McKinsey & Company he advised US manufacturing firms on entry into emerging markets and on renewable energy initiatives. **Lior Susan – Co-founder and Chairman:** Founding and Managing Partner at Eclipse Ventures, the company's earliest institutional backer. Served as interim CEO from December 2021 through the appointment of Mr. Stori. Eclipse also backed Cerebras Systems. **Amar Hanspal – Co-founder and former CEO:** Joined the founding team in May 2018 and shaped the company's growth strategy and product portfolio, taking it past $30MM in revenue in its first two years of commercial operation. Stepped down in December 2021, remaining briefly in an advisory capacity. **Michael Keogh – Chief Financial Officer:** 20-year technology and manufacturing industry veteran; joined in 2021 during preparation for the de-SPAC transaction.
SIZE OF COMPANY	~200–261 employees depending on source (PitchBook: 200 as of 2026; Tracxn: 261 as of April 2026). Headcount peaked at approximately 500 in 2021 before restructuring, including layoffs in 2022 and the closure of the Israel R&D center in 2024 (accessed 9/2/2026).

GROWTH TRAJECTORY	Founded in May 2018, Bright Machines closed a Series A in 2019 led by Eclipse Ventures and deployed its first Microfactory with modular Bright Robotic Cells that year, followed by the Brightware operating system in 2021. The company reached more than $30MM in revenue in its first two years of commercial operation. In May 2021 it agreed to go public via merger with SCVX Corp. (NYSE: SCVX) at a $1.6bn enterprise valuation, expecting up to $435MM in gross proceeds. The transaction was mutually terminated in December 2021 citing market conditions, on the same day co-founder Amar Hanspal stepped down as CEO. Headcount subsequently fell from approximately 500 and the company restructured through 2022 and 2024. A $132MM Series B (October 2022) led by Eclipse funded the Brightware roadmap and the Guadalajara Integration Hub. The company then repositioned around AI infrastructure: partnerships with NVIDIA and Microsoft in 2024, a GPU integration Microfactory product line, a $126MM Series C in June 2024 led by BlackRock, Chris Stori's appointment as CEO in September 2024, Bright Designer in March 2025, and the Hybrid BRC in July 2026. Current disclosed scale: more than 130 microfactories deployed across 10+ countries and 60+ customers
ANY MAJOR M&A / CORPORATE ACTIONS	No completed M&A. In May 2021 Bright Machines entered a business combination agreement with SCVX Corp. (NYSE: SCVX) to go public under the ticker 'BRTM' at a $1.6bn valuation, with up to $435MM in gross proceeds comprising $230MM in SPAC trust cash and a $205MM PIPE. PIPE investors included XN, Hudson Bay Master Fund, SB Management (SoftBank), Fidelity Management & Research and Alyeska Investment Group. The agreement was mutually terminated on December 13, 2021, citing the low likelihood of closing before the January 15, 2022 outside date and prevailing market conditions.
INVESTORS / INVESTMENT STAGE	Series C stage. Cumulative funding is reported inconsistently: the company's own June 2024 release states more than $400MM raised since 2018; PitchBook reports $502MM Disclosed rounds: $179MM Series A (2019, led by Eclipse Ventures); $132MM Series B (October 2022); $126MM Series C (June 2024) led by funds and accounts managed by BlackRock with participation from NVIDIA, Microsoft, Eclipse, Jabil and Shinhan Securities, plus $20MM venture debt from J.P. Morgan. Other investors include BMW i Ventures, Lux Capital, Hercules Capital and Reference Capital.
LATEST ROUND VALUATION	$1.10B post April 2024 Series C round.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale automation companies with a mixed hardware and software revenue model, we think it would be appropriate to value these companies on EV/S, with the multiple heavily dependent on the software attach rate — recurring Brightware and Smart Skills revenue should command a materially higher multiple than the underlying robotic cell hardware. Peer listed EMS and automation companies (JBL, FLEX, CLS, ABB, ROK, TER) are generally valued on P/E or EV/EBITDA given mature profitability, though EMS peers trade at structurally lower multiples reflecting thin assembly margins — a relevant anchor given Bright Machines' repositioning as a next-generation manufacturer rather than a pure tool vendor.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Dexterity

COMPANY NAME	Dexterity Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Physical AI (HW + SW)
SIZE OF TOTAL ADDRESSABLE MARKET	Dexterity develops physical AI software and integrated robotic systems that allow machines to perceive, reason and execute complex manipulation tasks in unstructured logistics and warehouse environments. Examples of this would be truck and container loading/unloading, parcel singulation, palletizing, depalletizing and fulfilment. The addressable opportunity is warehouse and parcel automation, driven by persistent labor shortage and rising parcel volume and variety. The company positions its systems as deployable without customers modifying existing warehouse infrastructure.
GEOGRAPHIC PRESENCE	Redwood City, California
NOTABLE COMPETITORS	**LISTED:** ABB, FANUC (6954.T), ROK, TER, SYM, ZBRA **PRIVATE:** Ambi Robotics, Pickle Robot, Anyware Robotics, GreyOrange, Addverb, Geek+, Righthand Robotics
KEY SUPPLIERS / CUSTOMERS	Dexterity's original model equipped commodity robot arms made by third-party OEMs with its own vision, touch and contextual intelligence software; the company has since moved toward integrated two-armed systems (DexR / 'Mech'). Disclosed customers and partners include FedEx (co-development partner on DexR trailer loading), GXO Logistics, UPS, Sagawa Express, Sumitomo Corporation (customer, investor and JV partner) and Dematic (deployment partner since 2022). The Sumitomo partnership targeted deployment of 1,500 warehouse robots across Japan by 2026.
COMPETITIVE POSITIONING	Dexterity's differentiator is an 'AI of AIs' architecture that orchestrates hundreds of physical AI models rather than relying on a single monolithic policy. Key attributes include: • Vision plus force sensing: Systems combine machine vision with force sensing and contextual intelligence to handle unstructured piles of goods with human-like dexterity and speed. • Hard-item generality: The company cites handling of 50,000+ SKUs spanning deformable polybags, delicate baked goods, poorly sealed cardboard, and crates and trays of consumer food - categories that defeat conventional suction-based picking. • Infrastructure-light deployment: Solutions are designed to operate without expensive auxiliary infrastructure or warehouse modification, lowering the customer's total deployment cost. • Production scale: The company reports more than 100MM autonomous actions completed in live production environments, and its DexR truck-loading system carries lift capacity of 60kg. • Strategic distribution: The Sumitomo relationship provides both capital and a Japanese go-to-market channel; the Dematic partnership provides systems-integrator reach.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://www.dexterity.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Samir Menon – Founder and CEO:** Founded Dexterity in 2017 to apply AI to physical manipulation tasks in logistics. Has led the company through Series A, B and subsequent rounds to unicorn status.
SIZE OF COMPANY	~185 employees per third-party trackers (as of June 2026)

GROWTH TRAJECTORY	Dexterity emerged from stealth in July 2020 with a $56.2MM Series A, followed by a $140MM Series B in October 2021 that took the company to unicorn status at a $1.4bn valuation while it deployed its first 1,000 robots into production. In March 2025 it closed a $95MM round at a $1.65bn valuation, and established a joint venture with Sumitomo to accelerate robot adoption in Japan. Positioning has shifted from robotic software-as-a-service on commodity arms toward integrated 'superhumanoid' Mech systems serving Fortune 50 and Fortune 500 parcel carriers, 3PLs and manufacturers.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Total funds of $139MM raised, latest round was Series C1, raised $104.9MM. Notable investors include: Lightspeed Venture Partners, Kleiner Perkins, Sumitomo Corporation, Obvious Ventures, B37 Ventures, Presidio Ventures, General Catalyst, Matter Venture Partners, Millennia Capital, GBM Ventures, Keybridge Venture Partners, Plug and Play Tech Center, and angels including Fidji Simo.
LATEST ROUND VALUATION	$1.65 billion as of the March 2025 round (up from $1.4 billion at the October 2021 Series B).
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale robotics and physical AI companies, we think it would be appropriate to value these companies on EV/S or on a per-unit-deployed basis, as gross margin profiles vary widely between hardware-led and software-led business models. Peer listed automation companies (ABB, FANUC, Rockwell, Teradyne) are generally valued on P/E or EV/EBITDA given more mature profitability profiles.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Figure AI

COMPANY NAME	Figure AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Humanoid Robotics & Physical AI
SIZE OF TOTAL ADDRESSABLE MARKET	Figure AI operates in the humanoid robotics, physical AI, warehouse automation, logistics automation, manufacturing automation, and intelligent robotics markets. The company develops autonomous humanoid robots for industrial and commercial deployment.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** KSCP **PRIVATE:** Boston Dynamics, Agility Robotics, Apptronik, 1X, Physical Intelligence
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Figure develops general-purpose humanoid robots designed to automate labor-intensive tasks across logistics, manufacturing, warehousing, and retail environments. Key differentiators include: • General-purpose humanoid platform: Figure develops humanoid robots designed to perform dynamic physical tasks that can be difficult to address with traditional fixed automation, combining manipulation and locomotion within a human-like form factor. • Proprietary robotics AI stack: Figure's robots are powered by its proprietary Helix vision-language-action models, with Helix 02 coordinating perception and whole-body control across the robot's hands, arms, torso, and feet. • End-to-end autonomous operation: The platform is designed to perceive changing physical environments and continuously adapt its movements rather than relying exclusively on fixed, hard-coded motions, enabling complex manipulation and logistics workflows. • Commercial and industrial applications: Figure targets real-world industrial workflows, with its robots deployed in automotive manufacturing and logistics applications at BMW Group Plant Spartanburg. • Strategic ecosystem: Figure's commercialization strategy includes partnerships with major industrial customers such as BMW, supporting the development and deployment of humanoid robots in production environments.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://www.figure.ai
MANAGEMENT TEAM / FOUNDER(S)	**Brett Adcock - Co-Founder & CEO:** Prior to Figure, Adcock founded Archer Aviation and Vettery. He is a repeat founder focused on transportation, automation, and AI-enabled platforms. **Gregg Hill - Co-Founder & Board Member:** Hill is also Co-Founder and General Partner of Parkway Venture Capital and has been involved with Figure since inception. **Jesse Coors-Blankenship - Co-Founder & Board Member:** Coors-Blankenship is Co-Founder and General Partner of Parkway Venture Capital and has supported Figure since its formation.
SIZE OF COMPANY	~909 employees on LinkedIn (accessed 9/18/2026)
GROWTH TRAJECTORY	Figure continues scaling deployment of humanoid robots across manufacturing, warehousing, logistics, and retail use cases. The company remains one of the largest funded humanoid robotics companies globally and has continued to develop robotics foundation models, humanoid hardware systems, robot calibration technologies, end effectors, and autonomous robot control systems. Recent patent activity accelerated throughout 2025 and 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Figure has raised $2.35B across multiple funding rounds. Prior to its latest financing activity, the company raised $1.5B in Series C funding in May 2025 and $675MM in Series B funding in February 2024. Earlier funding included a $70MM Series A round in 2023. Figure investors include NVDA, MSFT, OpenAI, AMZN, Bezos Expeditions, Intel Capital, Parkway Venture Capital, Qualcomm, LG Technology Ventures, Salesforce, Coatue, Brookfield Asset Management, Madrona, Prime Movers Lab, Align Ventures, and numerous strategic and financial investors. Current investment stage is Later Stage VC / Series C.
LATEST ROUND VALUATION	Figure completed its Series C financing in May 2025 at a $37.5B pre-money valuation and a $39.0B post-money valuation.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For humanoid robotics and physical AI companies that remain in commercialization mode, we think it would be appropriate to value these companies on P/S or EV/S given limited profitability and long deployment cycles. More mature automation and industrial technology companies are generally valued on P/E or P/FCF once utilization, software revenue, and margins become predictable.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Lumotive

COMPANY NAME	Lumotive, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Physical AI & Robotics
SIZE OF TOTAL ADDRESSABLE MARKET	Lumotive's TAM is the ~$1B automotive LiDar market.
GEOGRAPHIC PRESENCE	Redmond, Washington
NOTABLE COMPETITORS	**LISTED:** COHR, LITE, MVIS, OUST **PRIVATE:** Voyant Photonics, SiPh lidar-on-chip
KEY SUPPLIERS / CUSTOMERS	Lumotive does not disclose suppliers, but states LCMs are built on conventional silicon foundry processes. Key customers include Hokuyo (co-developed ships with 3D lidar with Lumotive LCM).
COMPETITIVE POSITIONING	Lumotive develops programmable Light Control Metasurfaces (LCM) for solid-state 3D sensing, electronically steering laser beams to enable compact LiDAR modules for automotive, robotics, and industrial applications. Key differentiators include: • Light Control Metasurfaces (LCM): Lumotive's programmable LCM technology electronically steers laser beams for solid-state 3D sensing, replacing mechanical or MEMS scanners and enabling compact LiDAR modules.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://lumotive.com/
MANAGEMENT TEAM / FOUNDER(S)	**Dr. Sam Heidari - CEO:** Prior to Lumotive, Heidari was previously CEO and Chairman of Quantenna Communications, CEO at Doradus Technologies, and CTO of Ikanos. He sits on the Global Semiconductor Alliance board and USC advisory boards, and holds a PhD and MS in Electrical Engineering from USC. **Dr. Gleb Akselrod - Co-founder and CTO:** Prior to Lumotive, Akselrod directed Optical Technologies at Intellectual Ventures and was a post-doctoral research fellow at Duke's Center for Metamaterials and Integrated Plasmonics. He earned a PhD in Physics at MIT and a BS in Engineering Physics from UIUC.
SIZE OF COMPANY	~83 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	In August 2023, Lumotive launched the LM10, its first production Light Control Metasurface chip for digital beam steering in commercial LiDAR modules. At SPIE Photonics West, Lumotive jointly announced the M30 Reference Design with Gpixel and planned for Meta-LiDAR integration.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Lumotive raised ~107.8MM over 2 main funding rounds. In its latest Series B round of funding, the company received $59MM in May 2025. Lumotive investors include Amazon, ITHCA Group, Eastbridge Group, EDOM Technology, Gates Frontier Fund, Grazia Equity, Himax Technologies, Hokuyo Capital, MetaVC Partners, Quan Funds, Swisscom Ventures, TSVC, USAA.
LATEST ROUND VALUATION	$199MM as of February 2025 Series B.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For analog companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer analog companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Metalenz

COMPANY NAME	Metalenz
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Metasurface Optics & Photonic Sensing
SIZE OF TOTAL ADDRESSABLE MARKET	Metalenz operates in the metasurface optics, 3D sensing, computer vision, photonics, mobile imaging, automotive sensing, industrial robotics, and optical semiconductor markets.
GEOGRAPHIC PRESENCE	Boston, Massachusetts
NOTABLE COMPETITORS	**LISTED:** WOLF, STM **PRIVATE:** Tunoptix, 2Pi Optics, Lightelligence, Lumus, Magic Leap
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: STMicroelectronics, UMC
COMPETITIVE POSITIONING	Metalenz develops metasurface optics designed to replace traditional multi-element lens stacks with semiconductor-manufacturable flat optical components. Key differentiators include: • Semiconductor-based metasurface optics: Metalenz uses planar metasurfaces composed of sub-wavelength nanostructures to manipulate light, with its optics designed for manufacturing using existing semiconductor foundry processes. • Flat-optics architecture: Metalenz's meta-optics can replace conventional stacks of refractive and diffractive lenses with a single flat optical element, simplifying optical system design and integration. • Multifunctional optical integration: A single Metalenz meta-optic can combine multiple optical functions, with the company stating that its technology can integrate the functionality of up to five optical elements into one surface. • Compact and scalable manufacturing: The technology reduces optical-module complexity and footprint while enabling mass production using semiconductor fabrication equipment and processes. • Broad sensing applications: Metalenz targets consumer, automotive, and industrial sensing applications, including smartphone biometrics, 3D sensing, LiDAR, robotics, gesture recognition, and object detection.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://metalenz.com/
MANAGEMENT TEAM / FOUNDER(S)	**Robert Devlin, Ph.D. - Co-Founder & CEO:** Devlin co-founded Metalenz after conducting metasurface optics research at Harvard University and has led commercialization of the company's flat-optics platform. **Federico Capasso, Ph.D. - Co-Founder:** Capasso is a renowned Harvard physicist widely recognized as a pioneer in metasurfaces and nanophotonics research. **Pawel Latawiec - CTO:** Latawiec leads development of Metalenz's optical platform and sensing technologies.
SIZE OF COMPANY	~37 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In 2025, Metalenz signed a licensing agreement with STMicroelectronics to accelerate metasurface optics adoption. In November 2025, the company announced that Polar ID entered mass production with UMC. In 2026, Metalenz launched Polar ID, an under-display face authentication solution, and introduced Polar 3D, a platform for relightable imaging and avatar generation from single-image capture.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Metalenz has raised $47MM across multiple funding rounds. Prior to its latest $30MM Series B financing completed in October 2022, the company raised $10MM in Series A financing in 2021 and $7MM in earlier financing in 2020.Metalenz investors include Neotribe Ventures, Intel Capital, TDK Ventures, M Ventures, Applied Ventures, 3M Ventures, Baidu Ventures, SGInnovate, Foothill Ventures, Osage University Partners, Global Brain, and Braemar Energy Ventures. Current investment stage is Series B.

LATEST ROUND VALUATION	Metalenz completed a $30MM Series B financing in October 2022 at a $100MM pre-money valuation and $130MM post-money valuation.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For photonics, sensing, optical semiconductor, and computer-vision hardware companies that are still scaling adoption, we think it would be appropriate to value these companies on P/S or EV/S. More mature optical component and semiconductor suppliers are generally valued on P/E or P/FCF given greater profitability and predictable end-market demand.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Mojo Vision

COMPANY NAME	Mojo Vision
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Micro-LED Displays & AR Optics
SIZE OF TOTAL ADDRESSABLE MARKET	Mojo Vision operates in the micro-LED display, augmented reality, virtual reality, AI infrastructure optics, advanced display technology, and optical semiconductor markets.
GEOGRAPHIC PRESENCE	Cupertino, California
NOTABLE COMPETITORS	**LISTED:** VUZI **PRIVATE:** Magic Leap, XReal, VueReal, Raxium, WaveOptics
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: Marvell, CY Vision
COMPETITIVE POSITIONING	Mojo Vision develops a micro-LED platform designed to enable next-generation display and optical technologies. Key differentiators include: • Silicon-based micro-LED architecture: Mojo micro-LED I The World's Smallest LED combines sub-micron GaN-on-silicon micro-LEDs with a 300mm semiconductor manufacturing approach designed for high-volume production. • Integrated optical platform: The architecture combines micro-LED emitters, high-performance quantum dots, optimized CMOS backplanes, wafer-to-wafer hybrid bonding, and custom micro-lens optics within an integrated technology platform. • Compact, high-performance displays: Mojo Vision's micro-LED technology is designed to provide high brightness and energy efficiency in compact form factors, including monolithic RGB displays for AI-enabled smart glasses. • AI infrastructure applications: Beyond displays, the company is applying its micro-LED platform to massively parallel optical interconnects for AI data centers, targeting higher bandwidth density and lower energy per bit. • Broad technology applicability: Mojo Vision positions its micro-LED platform across AI infrastructure and advanced display applications, leveraging a common semiconductor and optical technology base across different end markets.
YEAR FOUNDED / ESTABLISHED	2015
WEBSITE	https://www.mojo.vision/
MANAGEMENT TEAM / FOUNDER(S)	**Nikhil Balram, Ph.D. - CEO:** Balram leads Mojo Vision's transition toward micro-LED and AI infrastructure opportunities and previously held senior leadership roles within the company's core technology organization. **Michael Wiemer, Ph.D. - Co-Founder & CTO:** Wiemer co-founded Mojo Vision and leads development of its micro-LED, display, and optical technology stack. **Drew Perkins - Co-Founder & Chairman:** Perkins is a serial entrepreneur who previously founded Infinera and has extensive experience in optical networking and communications technologies.
SIZE OF COMPANY	~100 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In March 2026, Mojo Vision announced plans to accelerate commercialization of its micro-LED platform for AI infrastructure applications. The company also announced a collaboration with MRVL focused on next-generation high-density micro-LED connectivity solutions. Earlier, Mojo Vision expanded into automotive display applications through its work with CY Vision and continued development of micro-LED technologies for AR and advanced optical systems.
ANY MAJOR M&A / CORPORATE ACTIONS	Motorola Solutions Venture Capital exited its investment through a secondary transaction.

INVESTORS / INVESTMENT STAGE	Mojo Vision has raised $340MM across numerous financing rounds. Prior to its latest $17.5MM financing completed in March 2026, the company raised $75MM in Series B funding in July 2025 and $43.5MM in financing in October 2023. Earlier investors supported multiple rounds between 2017 and 2022. Mojo Vision investors include Future Ventures, Vanedge Capital, Khosla Ventures, NEA, Fusion Fund, Dolby Family Ventures, imec.xpand, TDK Ventures, Amazon Alexa Fund, Liberty Growth, Motorola Solutions Venture Capital, PTC, SVB Financial Group, Gradient Ventures, LG Electronics, and 8VC. Current investment stage is Later Stage VC.
LATEST ROUND VALUATION	Mojo Vision completed a financing round in March 2026 at a $340MM pre-money valuation and $357.5MM post-money valuation
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For display technology, micro-LED, AR optics, and optical semiconductor companies that remain in product-development or commercialization stages, we think it would be appropriate to value these companies on P/S or EV/S. More mature display, optical component, and semiconductor hardware companies are generally valued on P/E or P/FCF given established volume production and profitability.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Omnitron Sensors

COMPANY NAME	Omnitron Sensors
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / MEMS Sensors & Silicon Photonics
SIZE OF TOTAL ADDRESSABLE MARKET	Omnitron Sensors operates in the MEMS sensors, silicon photonics, LiDAR, autonomous vehicle sensing, augmented reality, and intelligent sensing semiconductor markets. The company develops semiconductor process technologies and MEMS platforms for next-generation sensing applications.
GEOGRAPHIC PRESENCE	Los Angeles, California
NOTABLE COMPETITORS	**LISTED:** ADI, STM, TXN, IFX, NXPI **PRIVATE:** Quintessent, Voyant, Scintil Photonics
KEY SUPPLIERS / CUSTOMERS	Suppliers: Silex Microsystems. Customers: N/A
COMPETITIVE POSITIONING	Omnitron Sensors develops a semiconductor process platform designed to improve performance, reliability, and cost efficiency for sensing applications. The company focuses on MEMS-based architectures for autonomous systems, intelligent devices, and photonic applications. Key differentiators include: • MEMS semiconductor platform: Proprietary MEMS technology designed to improve sensing performance while reducing manufacturing costs. • LiDAR and autonomous vehicle focus: The platform is designed to address reliability, size, and cost challenges associated with automotive LiDAR systems. • Silicon photonics integration: Omnitron combines MEMS fabrication expertise with silicon photonics capabilities. • Advanced MEMS manufacturing processes: The company's IP portfolio includes MEMS fabrication methods, scanning mirrors, optical arrays, and precision deployment mechanisms. • Scalable semiconductor production approach: The technology is intended for adoption across high-volume automotive, AI, and AR markets.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	https://omnitronsensors.com/
MANAGEMENT TEAM / FOUNDER(S)	**Eric Aguilar - Co-Founder & CEO:** Aguilar co-founded Omnitron Sensors in 2019 and has led the company's commercialization efforts in MEMS and autonomous sensing technologies. He was recognized as Latino Startup of the Year in 2021 following the company's early progress in automotive sensing and semiconductor innovation. **Trent Huang - Co-Founder & CTO:** Huang co-founded Omnitron Sensors and leads development of the company's MEMS architecture, semiconductor process platform, and silicon-photonics-enabled sensing technologies.
SIZE OF COMPANY	~20 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In 2023, Omnitron Sensors selected Silex Microsystems to support manufacturing of its MEMS scanner platform for LiDAR applications. The company has continued expanding its MEMS intellectual property portfolio, including scanning mirrors with dual-axis gimbal structures, optical arrays with self-aligned fibers and MEMS mirrors, and advanced MEMS fabrication technologies. The company's commercialization efforts remain focused on autonomous navigation, LiDAR, AR, and AI-driven sensing applications.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Omnitron Sensors has raised $25MM across multiple funding rounds. Prior to its latest $13.5MM Series A1 financing completed in January 2025, the company raised $6.7MM in Seed financing in March 2023, $2.4MM in Seed financing in October 2021, and $2.4MM in Seed financing in April 2020. Omnitron Sensors investors include Corriente Advisors, L'ATTITUDE Ventures, Gaingels, Makai VC, 10X Capital, Astor Perkins, Emerging Ventures, Interlock Capital, Plug and Play, Remapped Ventures, Stratminds VC, Lombardstreet Ventures, and 500 Global. The company's current investment stage is Series A1.
LATEST ROUND VALUATION	Omnitron Sensors completed a $13.5MM Series A1 financing in January 2025 at a $216.5MM pre-money valuation and a $230MM post-money valuation.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For MEMS, sensing, silicon photonics, and semiconductor platform companies that remain in product-development and commercialization stages, we think it would be appropriate to value these companies on P/S or EV/S, as profitability is typically not yet the primary value driver. More mature sensor and analog semiconductor companies are generally valued on P/E or P/FCF given established production scale, recurring customer demand, and stable profitability.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Path Robotics

COMPANY NAME	Path Robotics, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Physical AI & Robotics
SIZE OF TOTAL ADDRESSABLE MARKET	Path Robotics builds autonomous robotic welding systems powered by its Obsidian physical AI model. Systems use scanning, computer vision and AI to locate and weld seams autonomously, adapting to part-to-part variation rather than requiring parts to be perfectly positioned within a fixed cell. The demand driver is a structural shortage of skilled welders: the American Welding Society has cited a need for more than 320,000 new welders by 2029, and more than 50% of American goods require welded parts. End markets span general fabrication, heavy industry, construction and — increasingly — defense shipbuilding.
GEOGRAPHIC PRESENCE	Columbus, Ohio
NOTABLE COMPETITORS	**LISTED:** LECO, ITW, ABB, FANUC (6954.T), Yaskawa (6506.T) **PRIVATE:** Novarc Technologies, 1872, Anvil Robotics, and in-house automation programs at large fabricators
KEY SUPPLIERS / CUSTOMERS	Suppliers are not individually disclosed. Rove pairs the Obsidian model with a third-party quadruped robot platform. In February 2026 Path signed an initial memorandum of understanding with Huntington Ingalls Industries (HII), the largest US military shipbuilder, following an executive visit to Path's Columbus headquarters. In August 2026 this converted into a multi-year, performance-based automation agreement worth up to $900MM, of which Path's share is up to $600MM over seven years. Saronic plans to integrate Rove into its shipbuilding operations in Franklin, Louisiana.
COMPETITIVE POSITIONING	Path's differentiator is autonomous seam finding and adaptive welding — removing the programming and fixturing burden that limits conventional welding robots. Key attributes include: • Obsidian physical AI model: Enables robots to perceive the workpiece and adapt in real time rather than repeating a pre-programmed path, handling the imperfect sizing and highly reflective surfaces that make welding a hard robotics problem. • Rove mobile system: Launched April 16, 2026 at Sea-Air-Space, Rove pairs Obsidian with a quadruped platform to bring autonomous welding directly onto large, immovable assemblies — trading fixed-cell precision for mobile, adaptive reach in shipbuilding, heavy industry and construction. • Defense demand signal: The HII agreement provides a long-term, performance-based revenue commitment that funds sustained investment in facilities, workforce and the underlying architecture — an unusually strong contracted backlog for a company at this stage. • Vertical focus: Unlike general-purpose robotics platforms, Path has concentrated on one high-value, labor-constrained process, allowing deeper model specialization. • Concentration risk: The HII relationship represents a very large share of contracted future revenue. Performance-based structure means realization depends on delivery against Navy shipbuilding standards still under development.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://www.path-robotics.com/
MANAGEMENT TEAM / FOUNDER(S)	**Dr. Andy (Andrew) Lonsberry – Co-founder and CEO:** Co-founded Path with his brother while completing his PhD at Case Western Reserve University. Leads the company's physical AI strategy for manufacturing. **Dr. Alex Lonsberry – Co-founder and CTO:** Co-founder and Chief Technology Officer; has publicly framed the welder shortage as creating a vacuum that only intelligent robots can fill.
SIZE OF COMPANY	101–250 band, per Crunchbase

GROWTH TRAJECTORY	Commercial momentum accelerated sharply in 2026: an R&D agreement with HII in February, the launch of Rove in April at Sea-Air-Space, and in August a contract award of up to $600MM over seven years as Path's share of a $900MM HII automation agreement. The company has moved from stationary factory-cell deployments across US and Canadian fabricators into US Navy shipyard production work.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Path raised a total of $359.89M to date. The latest round Series D raised $107M in Oct. 2024, and added $68.7M later in Nov. 2025. Notable investors include BlackForest Ventures, Mana Ventures, RoboStrategy, and Asahi Kasei Corporate Venture Capital.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale robotics and physical AI companies, we think it would be appropriate to value these companies on EV/S or on a per-unit-deployed basis, as gross margin profiles vary widely between hardware-led and software-led business models. Given the contracted HII backlog, we would also consider a backlog-based or DCF approach for Path specifically, which is unusual for a company at this stage and distinguishes it from most private robotics comparables. Peer listed automation companies (ABB, FANUC, Rockwell, Teradyne) are generally valued on P/E or EV/EBITDA given more mature profitability profiles.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Physical AI / Robotics: Tacta Systems

COMPANY NAME	Tacta Systems, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Physical AI (HW + SW)
SIZE OF TOTAL ADDRESSABLE MARKET	Tacta develops 'Dextrous Intelligence' — a combined hardware, software and AI platform described by the company as a smart nervous system for robots, enabling them to sense, adapt and manipulate the physical world with human-like precision and speed. The company's thesis is that the hand, not the arm or the model, is the binding constraint in robotics, and that the hand problem cannot be solved without touch. Initial target markets are electronics, AI infrastructure and automotive manufacturing, with planned expansion into medical devices, aerospace and defense.
GEOGRAPHIC PRESENCE	Palo Alto, California
NOTABLE COMPETITORS	**LISTED:** ABB, FANUC (6954.T), Yaskawa (6506.T), TER, ROK **PRIVATE:** Shadow Robot, Sanctuary AI, Figure, Physical Intelligence, Skild AI, eBots, GelSight (tactile sensing).
KEY SUPPLIERS / CUSTOMERS	Suppliers are not publicly disclosed. Customers are not publicly disclosed. TactaBot is scheduled to ship in early 2027 for deployment with initial customers in the electronics, AI infrastructure and automotive manufacturing industries. Several strategic corporates are on the cap table — Sojitz Corporation, Yazaki Innovations and Woven Capital (Toyota) — which is a relevant signal on likely early industrial channels.
COMPETITIVE POSITIONING	Tacta's differentiator is proprietary tactile sensing hardware paired with a skill-capture data pipeline, rather than a vision-only manipulation policy. Key attributes include: • Sub-grain-of-sand tactile sensors: The company has developed tactile sensors smaller than a grain of sand that recognize a full dynamic range of human touch from 250 pascals to 700,000 pascals, sample at 400 times per second, and resolve temperature to within a tenth of a degree. • Skill Capture data pipeline: The Tacta Glove collects motion, force, video and temperature data from human workers performing manufacturing tasks; this data trains the Dexterous Intelligence models, improving adaptability and reducing training time for new tasks. • Vertically integrated stack: The company combines its own sensors, robotic hands (TactaBot) and AI models rather than layering software on third-party hardware. • Founder track record: Executive Chairman Andreas Bibl previously founded LuxVue Technology, acquired by Apple — a credential that carried weight with the investor base. • Timeline risk: TactaBot is not scheduled to ship until early 2027. The company remains pre-commercial, and its claims on sensor performance have not been independently validated in production.
YEAR FOUNDED / ESTABLISHED	Founded circa 2023–2024 (an investor described 18 months of progress as of mid-2025)
WEBSITE	https://www.tactasystems.com/
MANAGEMENT TEAM / FOUNDER(S)	**Andreas Bibl – Co-founder and Executive Chairman (previously CEO):** Prior to Tacta, Mr. Bibl founded and led LuxVue Technology, developer of low-power microLED displays for consumer electronics, which was acquired by Apple. **Vikram Pavate – Co-founder and CEO:** Serial entrepreneur; leads the company's commercialization of TactaBot for high-value manufacturing.
SIZE OF COMPANY	Headcount not publicly disclosed (accessed 9/2/2026)
GROWTH TRAJECTORY	Tacta emerged publicly on June 26, 2025, announcing $75MM in cumulative funding and unveiling its Dextrous Intelligence technology for the first time. In 2026 the company unveiled TactaBot, its robotic hand product for high-value manufacturing, incorporating its tactile sensors, Skill Capture system and Dexterous Intelligence models. TactaBot is scheduled to ship in early 2027. Progress to date has been rapid relative to hardware-robotics norms — an investor publicly described the company's 18-month progress as extraordinary — but the company has no disclosed revenue or production deployments.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Tacta has raised $86MM: a previously undisclosed $11MM seed round led by Matter Venture Partners, and a $64MM Series A co-led by America's Frontier Fund and SBVA. Notable investors: Matter Venture Partners, America's Frontier Fund, SBVA (SoftBank Ventures Asia), B Capital, EDBI, Sojitz Corporation, CDIB-TEN Capital, Yazaki Innovations, B5 Capital, Tyche Partners and Woven Capital (Toyota).
LATEST ROUND VALUATION	Valued at $220M post its Mar. 2025 deal, which raised $75M.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale robotics and physical AI companies, we think it would be appropriate to value these companies on EV/S or on a per-unit-deployed basis, as gross margin profiles vary widely between hardware-led and software-led business models. Peer listed automation companies (ABB, FANUC, Rockwell, Teradyne) are generally valued on P/E or EV/EBITDA given more mature profitability profiles.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Figure 9: Semiconductor Manufacturing Companies & Public Peers

Companies	Public Peers
ChipAgents	CDNS, SIE.DE, SNPS
Cognichip	CDNS, SNPS, ARM
Crossbar	INTC, MU, WDC, Samsung, AMBQ
FemtoMetrix	KLAC, ONTO, NVMI, BRKR, 6146 T
Forge Nano	VECO, AMAT, LRCX, KLAC, ASM International
Halo Industries	IPGP, LASR
Ideal Semiconductor	IFNNY, TXN, ON, STM, NXPI, AMAT
MatterStack	COHR, LITE, MVIS, OUST
Multibeam	AMAT, ASML, LRCX, KLAC, VECO
Prahelika	AMD, INTC, MU, QCOM
Recursive Superintelligence	GOOGL
Ricursive Intelligence	CDNS, SNPS
Source Foundry	TSM, GFS, INTC, Samsung, 0981 HK
Substrate	GFS, INTC, TSEM, MU
SunRise Memory	Weebit Nano
Tela Innovations	SNPS, CDNS, ARM
TeRAM	MU, WDC, STM
Volantis (Semiconductors)	ASML, COHR
xlight	ASML, VECO, NINOY
Zero ASIC	SNPS, CDNS, MRVL, AMD

Source: UBS

Semiconductor Manufacturing: ChipAgents

COMPANY NAME	ChipAgents
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / EDA & Design Automation
SIZE OF TOTAL ADDRESSABLE MARKET	ChipAgents builds an agentic AI platform that transforms specifications and code into production-ready RTL, verification assets and automated root-cause analysis, shortening design and verification cycles. The company distinguishes its approach from AI copilots: rather than assisting engineers with individual coding tasks, its specialized agents independently plan and execute multi-step design and verification workflows — compressing tasks that once took weeks or months into days or hours. The addressable market is design verification, which consumes a majority of engineering headcount and schedule on complex SoC programs, plus adjacent embedded systems workflows.
GEOGRAPHIC PRESENCE	Santa Clara, California (20,000 sq ft Silicon Valley headquarters opened 2026); originally Santa Barbara, California
NOTABLE COMPETITORS	**LISTED:** SNPS, CDNS, Siemens EDA (SIE.DE) — all now shipping their own agentic AI programs **PRIVATE:** Silimate, Diode Computers, Astrus, PrimisAI, and other agentic EDA entrants
KEY SUPPLIERS / CUSTOMERS	The company is an NVIDIA partner. Internal compute costs are described as substantial and were an explicit use of proceeds in the Series A. Deployed at more than 120 semiconductor companies as of mid-2026, including named customers MediaTek and Micron - both of which are also investors. The company describes this as one of the industry's largest commercial deployments of agentic AI for chip design and verification. Reported results include 15x faster specification reading and comprehension and a greater than 99% reduction in formal assertion generation time.
COMPETITIVE POSITIONING	ChipAgents' differentiator is that it shipped a production agentic product ahead of the incumbent EDA vendors and has converted that head start into broad deployment. Key attributes include: • Agents, not copilots: Specialized agents plan and execute complex, multi-step design and verification workflows autonomously rather than assisting with individual tasks. • Full-arc ambition: The roadmap targets agents handling the complete development arc - reading specifications, generating verification environments, running formal checks and debugging. • Verifying verification: The CEO has highlighted checking that human DV engineers' work is correct and exhaustive, particularly in defining corner cases, as a distinct and defensible use case. • First-mover in the category: The company had a product on the market for roughly a year before the large EDA vendors began publicizing their own agentic AI programs, and was the first agentic EDA startup to raise a Series A.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://chipagents.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Prof. William Wang – Founder and CEO:** Professor of machine intelligence at UC Santa Barbara. His AI work dates to 2011 and he began research on reasoning agents in 2017, which underpins the company's agentic architecture. **Advisors and board:** Advisors include EDA industry figures Wally Rhines, Raúl Camposano and Jack Harding. Sandeep Bharathi joined the Advisory Board in 2026.
SIZE OF COMPANY	Headcount not publicly disclosed

GROWTH TRAJECTORY	Founded in 2024, ChipAgents launched publicly in October 2024 alongside a $3.09 million pre-seed round led by ScOp Venture Capital. In October 2025, it closed a $21 million Series A with Bessemer Venture Partners and strategic investors, bringing cumulative funding to approximately $24 million. In February 2026, ChipAgents closed an oversubscribed $50 million Series A1 led by Matter Venture Partners, a TSMC-backed hardtech VC firm, bringing total capital raised to $74 million. In July 2026, during DAC in Long Beach, the company announced an additional $60 million Series A2 financing, with B Capital joining as a new investor, expanding its Series A financing to $134 million. The company reported 6x ARR growth in the first half of 2026 and deployments at more than 120 semiconductor companies, including MediaTek and Micron. These operating metrics are company-reported
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Cumulative funding of $137MM as of July 2026. Notable investors: Matter Venture Partners (Series A1 lead), B Capital (Series A2), Bessemer Venture Partners, Micron, MediaTek, Ericsson, ScOp Venture Capital (Ivan Bercovich), Amino Capital, and angels including Wally Rhines, Raúl Camposano and Jack Harding.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale enterprise software companies, we think it would be appropriate to value these companies on EV/ARR or EV/S given negligible near-term profitability, with the multiple driven by net revenue retention and growth durability. Peer listed EDA and design software companies (SNPS, CDNS) are generally valued on P/E or EV/FCF given mature profitability and high recurring revenue mix.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Cognichip

COMPANY NAME	Cognichip
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI-Based Semiconductor Design
SIZE OF TOTAL ADDRESSABLE MARKET	Cognichip operates in AI-enabled semiconductor design, electronic design automation and hardware-engineering software.
GEOGRAPHIC PRESENCE	Redwood City, California
NOTABLE COMPETITORS	**LISTED:** SNPS, CDNS, ARM **PRIVATE:** Graphcore, ChipAgents, ChipNexus, Silimate
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Cognichip develops foundational AI models designed specifically for semiconductor architecture and hardware engineering. Its physics-informed platform is intended to improve the speed, accuracy and efficiency of chip design. Key differentiators include: • Physics-informed foundation models: Cognichip's ACI platform is built around foundation models purpose-built for chip design and grounded in semiconductor physics, rather than relying on general-purpose language models for hardware-engineering tasks. • Full-stack AI-native design system: ACI Enterprise combines its foundation models with workflows, agent frameworks, project context and memory, compute and inference, and proprietary data to support semiconductor design activities across multiple stages of development. • AI-assisted hardware engineering: The platform integrates specifications, RTL, tests, constraints, verification, and power-performance-area optimization into AI-native workflows, enabling engineers to interact with and evolve designs across different design abstractions. • Engineering productivity: Cognichip positions ACI around accelerating semiconductor development while improving accuracy and convergence, with the goal of reducing manual handoffs and compressing design cycles. • Enterprise-grade security: Cognichip achieved SOC 2 Type II compliance in June 2026, with independently audited controls covering customer design data, intellectual property, and design workflows within its ACI infrastructure.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://cognichip.ai

MANAGEMENT TEAM / FOUNDER(S)	**Faraj Aalaei – Co-Founder, Chairman & CEO:** Prior to co-founding Cognichip, Faraj Aalaei built and led two semiconductor companies through public offerings as founder and CEO. He co-founded and led Centillium Communications, which developed broadband-access semiconductor solutions, and later founded Aquantia, a networking semiconductor company that was acquired by Marvell Technology. Following the acquisition, he led Marvell's Networking and Automotive businesses, a billion-dollar revenue segment. Earlier in his career, he held engineering and technology leadership roles at AT&T Bell Laboratories, Fujitsu Network Communications, and MA/COM. Mr. Aalaei holds a B.S. in Electrical Engineering from Wentworth Institute of Technology, an M.S.E.E. from the University of Massachusetts, an MBA from the University of New Hampshire, and an honorary Doctor of Engineering degree from Wentworth. **Ehsan Kamalinejad, Ph.D. – Co-Founder & CTO:** Prior to co-founding Cognichip, Ehsan Kamalinejad served as a machine-learning technology lead at Amazon Web Services, where he helped pioneer reinforcement learning from human feedback (RLHF) initiatives. Previously, he was a Senior Machine Learning Scientist at Apple, contributing to the Apple Photos "Memories" feature used by millions of customers. Earlier in his career, he held academic positions including Associate Professor at California State University, East Bay and conducted research in machine learning at UCLA and UC Riverside. At Cognichip, he leads development of the company's physics-informed foundation models for semiconductor design. Dr. Kamalinejad earned a Ph.D. in Applied Mathematics from the University of Toronto. **Simon Sabato – Co-Founder & Chief Architect:** Prior to co-founding Cognichip, Simon Sabato accumulated more than two decades of experience in chip architecture, networking silicon, and hardware-system design. His career includes architecture and engineering leadership roles at Google, Cisco Systems, Cadence Design Systems, Intel, F5, and other networking and semiconductor companies. His technical background spans ASIC design, FPGA accelerators, networking infrastructure, design automation, and error-correction technologies. Mr. Sabato holds a degree in Electrical Engineering and Computer Science from the University of California, Berkeley and is named as an inventor on 12 granted patents. At Cognichip, he leads the architecture of the company's Artificial Chip Intelligence (ACI®) platform and semiconductor-design infrastructure.
SIZE OF COMPANY	~91 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Cognichip has progressed from its 2024 formation into a 63-employee, revenue-generating developer of physics-informed AI for semiconductor design. Following its 2025 seed round, the company expanded its platform and established a Canadian subsidiary in Toronto. In 2026, Cognichip completed a global AI hardware hackathon program involving more than 220 students and 750,000 lines of RTL and achieved SOC 2 Type II compliance, supporting broader adoption by semiconductor-engineering teams.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Cognichip has raised a total of $93MM. In its latest Series A round of funding, it raised $60MM in April 2026. Investors include Seligman Ventures, Candou Ventures, VSC Ventures, Lux Capital, FPV Ventures, Mayfield Fund, Walden Catalyst, SBI Investment, Global Catalyst Partners Japan, and NewBuild Venture Capital.
LATEST ROUND VALUATION	N/A
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale AI software and semiconductor-design businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed EDA and semiconductor-IP providers such as SNPS, CDNS and ARM, together with private AI-native design companies such as ChipAgents, ChipNexus and Silimate.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Crossbar

COMPANY NAME	Crossbar
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / ReRAM Memory
SIZE OF TOTAL ADDRESSABLE MARKET	Crossbar operates in non-volatile memory, resistive random-access memory, embedded memory and integrated compute-and-storage systems.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** INTC, MU, WDC, Samsung Electronics, AMBQ **PRIVATE:** Nantero, InnoStar Semiconductor
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Crossbar develops ReRAM technology that can be embedded inside a system-on-chip or manufactured as a standalone memory device, enabling integrated compute and storage subsystems with non-volatile, reliable and energy-efficient memory. Key differentiators include: • ReRAM technology: Crossbar's resistive random-access memory provides non-volatile storage and can be implemented as either embedded memory or a standalone chip. • System-on-chip integration: The technology can be integrated directly inside an SoC, supporting closer coupling of compute and storage functions. • NVMe integration: Crossbar partnered with Mobiveil to apply NVMe SSD intellectual property to its ReRAM IP blocks. • Secure-computing applications: The company has expanded its technology into enhanced multiparty computation and open-source hardware for blockchain security. • Extensive IP portfolio: 180 patent families and 444 patent applications and grants covering two-terminal memory, selector devices, layered device structures and programming-current control. • Flexible deployment: The ability to deploy ReRAM as embedded IP or standalone memory differentiates Crossbar from conventional memory products limited to discrete storage implementations.
YEAR FOUNDED / ESTABLISHED	2008
WEBSITE	https://www.crossbar-inc.com/
MANAGEMENT TEAM / FOUNDER(S)	**Ker Zhang – CEO & Chairman:** Prior to becoming Chief Executive Officer and Chairman of Crossbar, Ker Zhang held executive leadership positions across the semiconductor and embedded-systems industries, including serving as Chairman and CEO of Greenliant Systems, a provider of NAND flash storage and embedded memory solutions. Earlier in his career, he held engineering, product, and business leadership roles at Cavium Networks, RMI Corporation, MIPS Technologies, and Silicon Graphics. At Crossbar, he leads commercialization of the company's resistive RAM (ReRAM), secure computing, and embedded-memory technologies. **Mark Davis, Ph.D. – President:** Prior to serving as President of Crossbar, Mark Davis spent decades in semiconductor memory research, development, and commercialization. He is widely recognized for his work in emerging non-volatile memory technologies and has played a key role in advancing the company's resistive memory platform from research into commercial deployment. At Crossbar, he oversees corporate operations, strategy execution, and commercialization initiatives surrounding the company's ReRAM technology portfolio. Crossbar's foundational memory technology originated from research conducted under his technical leadership. **Sung Jo, Ph.D. – CTO & Senior Vice President:** Prior to becoming CTO of Crossbar, Sung Jo built extensive expertise in memory-device physics, non-volatile memory architectures, and semiconductor process technology. At Crossbar, he leads technology development efforts related to ReRAM device architecture, embedded-memory integration, process scaling, and intellectual-property generation. His responsibilities include continued advancement of the company's non-volatile memory platform across automotive, industrial, AI, security, and embedded-computing applications.

SIZE OF COMPANY	~67 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Crossbar has evolved from a memory-device startup into a revenue-generating developer of embedded and standalone ReRAM technology. The company expanded its intellectual-property portfolio to 180 patent families and has pursued applications spanning IoT hardware, NVMe storage and secure computing. In 2023, Crossbar partnered with Mobiveil to combine NVMe SSD IP with ReRAM blocks, and in 2025 it worked with CertiK on enhanced multiparty-computation and open-source hardware for blockchain security.
ANY MAJOR M&A / CORPORATE ACTIONS	In December 2022, Orfin Ventures sold its entire stake in Crossbar to an undisclosed buyer.
INVESTORS / INVESTMENT STAGE	Crossbar has raised a total of $140.3MM. In its latest later-stage VC round of funding, it raised $5.2MM in January 2024. Investors include Innovation Partnerships, Lam Capital, Artiman Ventures, Kleiner Perkins, Northern Light Venture Capital, Friendly Hill Capital, TriplePoint Capital, VMS Group, Tyche Partners, Korea Investment Partners, Correlation Ventures, China Broadband Capital Partners, and the University of Michigan Endowment.
LATEST ROUND VALUATION	N/A
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale memory-semiconductor businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed memory and semiconductor providers such as INTC, MU, WDC, Samsung Electronics and Weebit Nano, together with private memory companies such as Nantero and InnoStar Semiconductor.
SECONDARY PRICING DATA	A private secondary transaction was completed in December 2022, when Orfin Ventures sold its stake to an undisclosed buyer. PitchBook does not disclose the consideration, number of shares or per-share price.

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: FemtoMetrix

COMPANY NAME	FemtoMetrix, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Process Control & Metrology
SIZE OF TOTAL ADDRESSABLE MARKET	Femtometrix develops optical metrology systems for the detection of Non-Visual Defects (NVDs) — surface, buried-interface and structural irregularities that conventional in-line inspection tools cannot resolve. The Harmonic series platform is built on Optical Second Harmonic Generation (SHG), a non-destructive, contactless characterization technique applied to surfaces, interfaces, thin films and bulk material properties. Applications span bulk silicon as well as SOI, GaN, SiC and other engineered substrates, across high-volume manufacturing, pilot production, development and research environments.
GEOGRAPHIC PRESENCE	Los Angeles, California
NOTABLE COMPETITORS	**LISTED:** KLAC, ONTO, NVMI, BRKR, 6146.T (Disco) - broader process control incumbents **PRIVATE:** Park Systems (KOSDAQ-listed in Korea), Exnodes, Nearfield Instruments
KEY SUPPLIERS / CUSTOMERS	Customers are not publicly disclosed. The company has stated tool placements globally and has publicly co-presented process control study results with Tokyo Electron (TEL). Strategic relationships include a patent and technology license agreement with The Boeing Company underpinning elements of the optical metrology platform, and strategic investor relationships with Samsung and SK Hynix.
COMPETITIVE POSITIONING	Femtometrix's core differentiator is its patented application of Optical Second Harmonic Generation to in-line semiconductor process control. Key attributes include: • Non-visual defect detection: Detects surface and buried metallic and organic contamination and other yield-limiting non-uniformities that are invisible to conventional optical and e-beam in-line tools. • Non-destructive and contactless: Measurement does not damage or consume the wafer, enabling in-line deployment rather than destructive offline analysis. • Buried-interface access: The Harmonic F1x system characterizes defects located at buried interfaces, a capability positioned as complementary to, rather than substitutive for, incumbent inspection platforms. • Advanced substrate coverage: Applicable to bulk silicon, SOI, GaN, SiC and other advanced engineered substrates, with tool configurations spanning HVM fabs, pilot lines and research labs. • Strategic investor alignment: Samsung Ventures and SK hynix are both on the register, providing early access to leading-edge memory and logic customers.
YEAR FOUNDED / ESTABLISHED	2011
WEBSITE	https://www.femtometrix.com/

MANAGEMENT TEAM / FOUNDER(S)	**Hovik Nazaryan – Founder:** Founded Femtometrix in 2011 around the commercialization of SHG-based optical metrology. **Alon Raphael – President and CEO:** Prior to Femtometrix, Mr. Raphael co-founded Elumium Network Technologies (internet infrastructure, operating under the AOL Optimized and NetZero Platinum brands) and Tamarisc Diagnostics, a nanobiotechnology-based medical diagnostics firm commercialized out of UC Santa Barbara. Earlier roles at DreamWorks Interactive and Load Media Networks. **Tyler Rubin – CFO:** Prior to Femtometrix, Mr. Rubin was an investor at an internal JPMorgan hedge fund covering technology and telecom equities and debt, and co-founded Tysam Networks, which merged with EQAL (acquired by Everyday Health, NYSE: EVDY, for $25MM in 2012). A.B. in Physics, Harvard. **Board and advisory bench:** Includes Brian Modoff (former EVP of Strategy, M&A and Venture Investments at Qualcomm; former Managing Director, Deutsche Bank equity research), Magnus Ryde (former President and CEO, TSMC North America; 18 years at KLA-Tencor), Itzik Goldberger (former VP Marketing / GM, Applied Materials and KLA-Tencor) and Richard Hong (former President, KLA-Tencor Korea; 2008 SEMI Bob Graham Award recipient). Samsung Ventures and SK hynix hold board observer seats.
SIZE OF COMPANY	~11–50 employees per LinkedIn
GROWTH TRAJECTORY	Femtometrix has followed a long, capital-light commercialization path relative to peers. The company entered a patent and technology license agreement with The Boeing Company that enabled it to incorporate enhanced non-destructive metrology capability into its product line. Product line has progressed from the Harmonic F1x inspection system to the F3x process control tool, with the company co-presenting results at FCMN 2022 with Tokyo Electron on the use of F3x for sub-nanometer critical dimension measurement in advanced devices. The company reports tool placements globally and maintains a South Korea office supporting the Samsung and SK hynix relationships. Growth has been incremental rather than step-function; the company remains sub-scale relative to listed process control peers.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Latest reported financing was $5M in October 2025. Notable investors include Samsung Venture Investment, SK Telecom Venture Capital, AVACO, and the Center for Technology Transfer & Commercialization. from Samsung Venture Investment.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For semiconductor capital equipment companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer semiconductor capital equipment companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Forge Nano

COMPANY NAME	Forge Nano
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Atomic Layer Deposition Equipment
SIZE OF TOTAL ADDRESSABLE MARKET	Forge Nano operates across atomic layer deposition equipment, semiconductor manufacturing, battery materials and nanoscale surface coatings.
GEOGRAPHIC PRESENCE	Thornton, Colorado
NOTABLE COMPETITORS	**LISTED:** VECO, AMAT, LRCX, KLAC, ASM International **PRIVATE:** ATLANT 3D, Coreshell, Elementum 3D
KEY SUPPLIERS / CUSTOMERS	Disclosed commercial counterparty: An unnamed photonics-technology company placed an order for Forge Nano semiconductor wafer-fabrication equipment in June 2026. Strategic partners: Samsung SDI and Voltaiq. Forge Nano formed a partnership with Samsung SDI for US production of advanced battery cells and expanded its battery-development and manufacturing partnership with Voltaiq. Suppliers: N/A
COMPETITIVE POSITIONING	Forge Nano develops semiconductor equipment and advanced-materials technologies based on atomic layer deposition. Its Atomic Armor platform applies nanoscale coatings to semiconductor structures, battery materials and other critical components to improve performance at the material surface. Key differentiators include: • Atomic Armor platform: Forge Nano's core platform applies precise atomic-scale surface coatings to semiconductor and battery materials. • Semiconductor ALD equipment: The company provides ALD equipment for commercial semiconductor manufacturing, including processes addressing high-aspect-ratio structures. • Battery-material coatings: Atomic Armor is used for nano-engineered coatings on anodes, cathodes and solid-state electrolytes. • Cross-market platform: The same core deposition technology addresses both AI-era semiconductor manufacturing and advanced battery applications. • Commercial photonics application: In June 2026, Forge Nano secured an equipment order from an unnamed photonics-technology company for commercial integrated-photonics manufacturing. • Battery-manufacturing expansion: Forge Nano partnered with Samsung SDI on US production of advanced battery cells and expanded its North Carolina lithium-ion battery facility in 2026. • IP portfolio: 30 patent families and 129 patent applications and grants spanning ALD, particle coatings, battery materials and additive manufacturing.
YEAR FOUNDED / ESTABLISHED	2011
WEBSITE	https://forgenano.com

MANAGEMENT TEAM / FOUNDER(S)	**Paul Lichty, Ph.D. – Co-Founder & CEO:** Prior to co-founding Forge Nano in 2013, Paul Lichty worked as a Research Engineer at SunDrop Fuels, where he helped develop renewable-energy and biofuels technologies. Earlier, he served as a Research Engineer at Copernican Energy and as an engineering consultant. While completing doctoral research at the University of Colorado Boulder, Dr. Lichty helped develop and patent a scalable atomic layer deposition (ALD) process that became the foundation of Forge Nano. Under his leadership, the company has expanded from a nanocoatings technology provider into commercial battery manufacturing, semiconductor-process equipment, and advanced materials markets. He holds a B.S. in Mechanical Engineering and a Ph.D. in Chemical Engineering from the University of Colorado Boulder. **James Trevey, Ph.D. – Chief Technology Officer:** Prior to becoming CTO of Forge Nano, James Trevey co-founded the company and helped develop its core Atomic Armor™ platform and ALD-based materials technologies. He oversees technology development, product scaling, commercial engagements, intellectual-property strategy, and engineering execution across Forge Nano's battery, semiconductor, and nanocoating initiatives. Dr. Trevey is an expert in electrochemistry and solid-state battery technologies and has authored numerous technical publications and patents. He holds degrees in Materials Science, Mechanical Engineering, and Engineering Physics from the University of Colorado Boulder.
SIZE OF COMPANY	~129 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Forge Nano has evolved from a Northwestern University spinout into a revenue-generating supplier of ALD equipment and nano-engineered materials for semiconductor and battery applications. In 2026, the company secured a commercial integrated-photonics equipment order, partnered with Samsung SDI on US battery-cell production and began expanding its North Carolina lithium-ion battery facility. Revenue declined from $8.3MM in 2024 to $7.4MM in 2025, while the company continued scaling semiconductor-equipment and battery-manufacturing capacity.
ANY MAJOR M&A / CORPORATE ACTIONS	Forge Nano reached a definitive agreement in April 2026 to complete a $1.2B reverse merger with Archimedes Tech SPAC Partners II. The proposed combined company is expected to trade under the ticker NANO, but PitchBook classifies the transaction as announced and in progress. A related $123MM PIPE was announced in April 2026. Forge Nano acquired ALD NanoSolutions in February 2020. Forge Nano formed Forge Battery as a subsidiary in 2023.
INVESTORS / INVESTMENT STAGE	Forge Nano has raised a total of $179.6MM. In its latest completed Series D round of funding, it raised $87.6MM through a combination of $82.5MM of equity and $5.1MM of debt in April 2026. Investors include Ascent Funds, Rock Creek Capital, Top Material, Breton Capital Ventures, Orion Infrastructure Capital, Hanwha Aerospace, GM Ventures, LG Technology Ventures, Mitsui Kinzoku-SBI Material Innovation Fund, SBI Investment, Sumitomo, Volkswagen Group, and the U.S. Department of Energy.
LATEST ROUND VALUATION	N/A
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor-equipment and advanced-materials businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed semiconductor-equipment providers such as VECO, AMAT, LRCX, KLAC and ASM International, together with private atomic-scale materials companies such as ATLANT 3D and Coreshell.
SECONDARY PRICING DATA	PitchBook reports a proposed PIPE price of $10.00 per common share. The PIPE and reverse merger remain announced and in progress, so this price is not presented as a completed secondary-market transaction.

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Halo Industries

COMPANY NAME	Halo Industries
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Silicon Carbide Wafering Equipment
SIZE OF TOTAL ADDRESSABLE MARKET	Halo Industries operates in silicon-carbide wafer manufacturing, laser-based semiconductor production equipment and advanced materials processing. Its end markets include AI data centers, power electronics, augmented reality, e-mobility, defense, industrial motors and electric-grid applications.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** IPGP, LASR **PRIVATE:** TRUMPF, Laserline, Fabmatics, Sphere Semi, HighLine Technology
KEY SUPPLIERS / CUSTOMERS	Equipment supplier: Revasum supplied Halo Industries with 7AF-HMG grinding equipment, including a second system announced in 2023. Manufacturing-software provider**:** Eyelit Technologies was selected in 2026 to provide a composable manufacturing-execution system for Halo's scalable silicon-carbide wafering production. Customers:N/A.
COMPETITIVE POSITIONING	Halo Industries develops a proprietary laser-based manufacturing platform for producing high-quality silicon-carbide wafers. Its in-house designed equipment is intended to replace traditional mechanical, chemical and thermal wafer-processing methods while reducing cost and supporting volume manufacturing. Key differentiators include: • Laser-based wafering platform: Halo uses laser energy to create controlled subsurface damage and enable precise material cleavage. • Controlled crack propagation: The company's patents cover incident-radiation-induced subsurface damage and capacitive clamping processes for controlled wafer separation. • Silicon-carbide specialization: Halo's platform is optimized for SiC wafers used in power electronics, AI data centers, electric vehicles, defense and grid applications. • In-house production equipment: Halo designs and builds its own laser-production systems rather than relying solely on third-party processing equipment. • Volume-manufacturing focus: The platform is designed to move beyond laboratory processing into scalable SiC wafer production. • Manufacturing infrastructure: Halo purchased wafer-grinding equipment from Revasum and selected Eyelit Technologies to provide its manufacturing-execution system. • Stanford-originated technology: Halo was spun out of Stanford University and has continued developing its laser-cleaving intellectual property since its formation.
YEAR FOUNDED / ESTABLISHED	2014
WEBSITE	https://halo-industries.com/

MANAGEMENT TEAM / FOUNDER(S)	**Andrei Iancu, Ph.D. – Co-Founder, CEO & Board Member:** Prior to co-founding Halo Industries in 2014, Andrei Iancu conducted multidisciplinary research at Stanford University spanning aeronautics, electrical engineering, and mechanical engineering. Halo emerged as a Stanford spinout commercializing laser-based manufacturing technologies developed during his doctoral research. Under his leadership, the company has developed laser-processing systems that improve silicon-carbide wafer yield and quality while reducing material waste for power-electronics applications. Dr. Iancu earned B.S. degrees in Electrical Engineering and Aerospace Engineering from Iowa State University and subsequently completed M.S. and Ph.D. degrees at Stanford University. He is also a Venture Partner at Prime Movers Lab. **Philip Van Stockum – Co-Founder & CTO:** Prior to co-founding Halo Industries, Philip Van Stockum conducted doctoral research at Stanford University and performed research at Los Alamos National Laboratory. He also served as a founding partner of Innova Science. As Halo's CTO, he leads development of the company's laser-production equipment, controlled wafer-cleaving technology, and manufacturing systems for semiconductor substrates. Dr. Van Stockum earned a bachelor's degree from Princeton University and a Ph.D. from Stanford University. **Charlie Rudy – Co-Founder & Chief Scientific Officer:** Prior to joining Halo Industries, Charlie Rudy worked as a Laser Scientist at IPG Photonics and previously conducted optics and photonics research at Stanford University. His scientific background centers on laser systems, fiber optics, nonlinear optics, and materials processing, reflected in a substantial body of academic publications and citations. At Halo, Dr. Rudy leads scientific development of the company's laser-induced subsurface-damage, crack-propagation, and wafer-cleaving technologies. He earned B.S., M.S., and Ph.D. degrees through studies at University of California San Diego and Stanford University.
SIZE OF COMPANY	~128 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Halo Industries has evolved from a Stanford University spinout into a 128-employee manufacturer of laser-based silicon-carbide wafering equipment. The company progressed from early laboratory development into volume-oriented production systems serving power electronics, AI data centers and e-mobility applications. Following its Series B, Halo expanded production capacity, purchased additional wafer-grinding equipment from Revasum and selected Eyelit Technologies in 2026 to support scalable manufacturing execution.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Halo Industries has raised a total of $122MM. In its latest Series B round of funding, it raised $80MM in May 2024. Investors include US Innovative Technology Fund, 8VC, Ilham Capital Ventures, Atento Capital, Fontinalis Partners, Prime Movers Lab, W Ventures, Stanford University, Harmonix Fund, California Energy Commission, and the U.S. Department of Energy.
LATEST ROUND VALUATION	N/A
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor-equipment and advanced-manufacturing businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed laser-equipment providers such as IPGP and LASR, together with private semiconductor-production companies such as TRUMPF, Fabmatics and HighLine Technology.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: iDEAL Semiconductor

COMPANY NAME	iDEAL Semiconductor
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Power Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Ideal Semiconductor operates in power semiconductors, silicon MOSFETs, electronic power conversion, electric vehicles and data-center power efficiency.
GEOGRAPHIC PRESENCE	Bethlehem, Pennsylvania
NOTABLE COMPETITORS	**LISTED:** Infineon Technologies AG, TXN, ON, STM, NXPI, AMAT **PRIVATE:** Metalenz, Retym
KEY SUPPLIERS / CUSTOMERS	Manufacturing partner: Polar Semiconductor manufactures Ideal Semiconductor's SuperQ power devices. Distribution partner: Mouser Electronics distributes SuperQ-based power devices globally. Technology partner: Richardson Electronics entered a technology-partner agreement covering SuperQ MOSFETs. Customers: N/A.
COMPETITIVE POSITIONING	Ideal Semiconductor develops energy-efficient silicon power devices for electronic power conversion. Its proprietary SuperQ platform is intended to improve power efficiency in electric vehicles, high-voltage batteries, battery-powered devices and data centers while retaining a silicon-based architecture. Key differentiators include: • SuperQ platform: Ideal Semiconductor's proprietary silicon power-device architecture is designed to deliver improved power-conversion efficiency. • Charge-balance technology: The company's patents cover simplified charge balancing, charge compensation and edge-termination structures in semiconductor devices. • Silicon-based approach: SuperQ seeks to improve power-device performance without requiring customers to transition entirely to silicon-carbide or gallium-nitride devices. • Automotive-qualified product: Ideal achieved automotive qualification for SuperQ and launched the iS20M028S1CQ MOSFET in October 2025. • High-voltage battery safety: In November 2025, the company introduced SuperQ MOSFET technology intended to improve protection in high-voltage battery systems. • Manufacturing partnership: Polar Semiconductor manufactures Ideal Semiconductor's SuperQ power devices. • Commercial distribution: Mouser Electronics distributes SuperQ-based products globally, while Richardson Electronics is a disclosed technology partner. • Data-center relevance: The technology is intended to improve compute density and reduce cooling requirements through more efficient power conversion.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	https://idealsemi.com/

MANAGEMENT TEAM / FOUNDER(S)	**Mark Granahan – Co-Founder, CEO, President & Board Member:** Prior to co-founding iDEAL Semiconductor in 2016, Mark Granahan spent more than three decades in semiconductor technology and business leadership. Earlier in his career, he held management and technology leadership positions at Texas Instruments and Lucent Microelectronics (Bell Labs). He later founded Ciclon Semiconductor, a developer of high-performance power semiconductors that was acquired by Texas Instruments in 2009. Following a career spanning engineering, product management, business development, acquisitions, and executive leadership, Mr. Granahan founded iDEAL Semiconductor to commercialize its SuperQ™ silicon power-device architecture. He holds a B.S. in Materials Science and Engineering from Pennsylvania State University and a graduate degree in Management Science from the University of Texas at Austin. **David Jauregui – Co-Founder & CTO:** Prior to co-founding iDEAL Semiconductor, David Jauregui accumulated more than 30 years of experience in power semiconductors, device architecture, analog ICs, and power-management technologies. His career included senior engineering, technical, and business leadership positions at Texas Instruments, Ciclon Semiconductor, and International Rectifier. Drawing on extensive experience with both silicon and wide-bandgap power devices, he co-developed iDEAL's SuperQ™ architecture as an alternative path to improving high-voltage power-semiconductor performance. At iDEAL, he leads development of the company's charge-balanced semiconductor architecture, edge-termination technologies, and MOSFET product portfolio. **Michael Burns – Co-Founder & Executive Chairman:** Prior to co-founding iDEAL Semiconductor, Michael Burns built a career spanning technology entrepreneurship, private equity, and industrial leadership. He has held senior executive positions at numerous technology and semiconductor companies and was a co-founder, investor, or executive at organizations including Agere Systems, Ciclon Semiconductor, and Traffic.com. Earlier in his career, he served as Chairman and CEO of Dana Corporation, President of General Motors Europe, Managing Director at Guggenheim Partners, and a director of United Parcel Service. He holds engineering and business degrees from the University of Pennsylvania, including studies at the Wharton School. At iDEAL Semiconductor, he provides strategic guidance and governance as Executive Chairman.
SIZE OF COMPANY	~58 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Ideal Semiconductor has progressed from power-device development into commercial production of its SuperQ silicon MOSFET platform. After launching SuperQ in 2023, the company signed Mouser as a global distributor, Richardson Electronics as a technology partner and Polar Semiconductor as its manufacturing partner in 2025. Ideal also achieved automotive qualification for SuperQ, launched the iS20M028S1CQ MOSFET and introduced devices targeting improved safety in high-voltage battery systems.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Ideal Semiconductor has raised a total of $76.6MM. In its latest Series C round of funding, it raised $40MM in May 2023. Investors include Applied Ventures and Ben Franklin Technology Partners of Northeastern Pennsylvania. PitchBook does not disclose the other Series C investors.
LATEST ROUND VALUATION	$240MM as of May 2023 Series C.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale power-semiconductor businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed power-semiconductor providers such as Infineon Technologies AG, TXN, ON, STM and NXPI, together with private semiconductor companies such as Retym.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: MatterStack

COMPANY NAME	MatterStack
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI-Driven Materials Discovery
SIZE OF TOTAL ADDRESSABLE MARKET	Matterstack operates across microelectronics materials research, computational materials science, materials simulation and AI-driven materials discovery.
GEOGRAPHIC PRESENCE	Austin, Texas
NOTABLE COMPETITORS	**LISTED:** COHR, LITE, MVIS, OUST **PRIVATE:** Voyant Photonics, SiPh lidar-on-chip
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Matterstack develops AI tools for discovering, analyzing and simulating materials used in microelectronics. The company is positioned at the intersection of materials science, semiconductor research and AI-based modeling. Key differentiators include: • AI models for materials discovery: Matterstack develops models intended to help researchers identify and study materials relevant to microelectronics. • Microelectronics materials analysis: The platform focuses specifically on analyzing materials used in semiconductor and related electronic applications. • Computational research system: Matterstack combines computational materials research with AI-driven analytical tools. • Materials data modeling: The technology includes systems for organizing and modeling materials data. • Simulation platform: Researchers and engineers can use the company's tools to simulate and evaluate material behavior. • Integrated materials intelligence: Matterstack brings discovery, analysis, modeling and simulation into a single AI-based materials-research offering.
YEAR FOUNDED / ESTABLISHED	2025
WEBSITE	https://matterstack.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Ellad Tadmor – Co-Founder:** Prior to co-founding Matterstack, Ellad Tadmor built a distinguished academic career in computational materials science and multiscale modeling. He serves as a professor at the University of Minnesota and is widely recognized for research spanning atomistic simulation, computational mechanics, and materials design. His work has focused on developing computational frameworks that connect materials behavior across length scales, from atomic interactions to macroscopic performance. At Matterstack, he is associated with development of the company's computational materials and AI-driven modeling platform. **Stefano Martiniani – Co-Founder:** Prior to co-founding Matterstack, Stefano Martiniani established a research career in statistical physics, machine learning, optimization, and computational modeling. He serves as an Assistant Professor at New York University, where his research focuses on complex systems, machine learning methods, scientific computing, and data-driven discovery. His work combines artificial intelligence with physics-based modeling to study high-dimensional systems and design challenges. At Matterstack, he is associated with development of AI-driven tools for materials research and discovery.
SIZE OF COMPANY	~9 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Matterstack was founded in 2025 to develop AI models and simulation tools for microelectronics materials research. The company has built a nine-person team and joined the NYU Entrepreneurial Institute through a non-equity accelerator program. Its platform spans materials discovery, analysis, data modeling and simulation, although PitchBook does not disclose product launches, named customers, revenue or external financing.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	Matterstack has raised no disclosed external funding. Its latest disclosed institutional program was participation in the NYU Entrepreneurial Institute on an undisclosed date, with no equity or funding exchanged. Institutional supporters include the NYU Entrepreneurial Institute.
LATEST ROUND VALUATION	N/A.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale AI software and materials-informatics businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference private AI and research-software providers such as Apprentice, Modal Labs, Human Archive, Centralize and JustAI.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing:Multibeam

COMPANY NAME	Multibeam
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Maskless E-Beam Lithography
SIZE OF TOTAL ADDRESSABLE MARKET	Multibeam operates in semiconductor lithography equipment, maskless direct-write patterning, electron-beam processing and production tools for wafers and advanced packaging panels.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** AMAT, ASML, LRCX, KLAC, VECO **PRIVATE:** xLight, Substrate
KEY SUPPLIERS / CUSTOMERS	Disclosed customer: An unnamed customer in Taiwan placed the first Taiwan order for Multibeam's newly launched MBX platform in July 2026. Commercial partner: Marketech partnered with Multibeam to accelerate sales of its multiple-column E-Beam Lithography systems in Taiwan. Suppliers: N/A.
COMPETITIVE POSITIONING	Multibeam develops maskless semiconductor lithography equipment based on multicolumn electron-beam technology. Its systems are designed to eliminate expensive photomasks while supporting faster and more flexible patterning for integrated circuits, wafers and advanced packaging. Key differentiators include: • MBX platform: Multibeam's newly launched MBX system applies multiple-column electron-beam lithography to maskless semiconductor production. • Multicolumn EBL architecture: The platform uses multiple electron-beam columns to increase processing throughput relative to a single-column system. • Maskless direct-write patterning: Designs can be written directly onto the production substrate without requiring a dedicated photomask. • Miniature e-beam column arrays: Multibeam uses arrays of compact charged-particle columns as a core element of its lithography architecture. • 300mm wafer and panel-level support: Series B proceeds were designated for a next-generation platform serving 300mm wafers and panel-level maskless lithography. • Precision positioning and electron optics: The company's intellectual property covers processing methods, charged-particle positioning, modular beam columns, precision stages and multibeam inspection. • Commercial validation in Taiwan: Multibeam secured the first Taiwan order for MBX and established a local sales partnership with Marketech.
YEAR FOUNDED / ESTABLISHED	2001
WEBSITE	https://www.multibeamcorp.com/
MANAGEMENT TEAM / FOUNDER(S)	**David Lam, Ph.D. – Founder, CEO & Chairman:** Prior to founding Multibeam Corporation, David Lam was one of the most prominent entrepreneurs in the semiconductor equipment industry. He founded Lam Research in 1980 and built it into one of the world's leading wafer-fabrication equipment suppliers. Earlier in his career, he worked at Texas Instruments and held engineering and management positions in semiconductor manufacturing technology. In addition to founding Lam Research, he later established and invested in multiple semiconductor and technology ventures. Dr. Lam holds a Ph.D. in Chemical Engineering from the Massachusetts Institute of Technology and has spent decades commercializing advanced semiconductor processing technologies. At Multibeam, he has overseen development and commercialization of the company's maskless electron-beam lithography platform.
SIZE OF COMPANY	70 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Multibeam has evolved from a developer of electron-beam technology into a revenue-generating supplier of maskless lithography systems. Following its 2025 Series B, the company accelerated development of a multicolumn EBL platform for 300mm wafers and panel-level applications. In 2025 and 2026, Multibeam expanded sales activities in Taiwan through Marketech, launched the MBX platform, secured its first Taiwan order and signed a $140MM letter of intent with the U.S. Department of Commerce.

ANY MAJOR M&A / CORPORATE ACTIONS	In April 2021, Booga Ventures, DynaFund Ventures and Global Strategy Investment Management sold their stakes to an undisclosed buyer.
INVESTORS / INVESTMENT STAGE	Multibeam has raised a total of $41.6MM. In its latest later-stage VC round of funding, it raised an undisclosed amount in August 2026. Investors include Super Capital Group, Onto Innovation, Lam Capital, UMC Capital, MediaTek, In-Q-Tel, DynaFund Ventures, Global Strategy Investment Management, and Booga Ventures.
LATEST ROUND VALUATION	N/A.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor-equipment businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed semiconductor-equipment providers such as AMAT, ASML, LRCX, KLAC and VECO, together with private manufacturing-technology companies such as xLight and Substrate.
SECONDARY PRICING DATA	A private secondary transaction was completed in April 2021, but PitchBook does not disclose the transaction value, number of shares or per-share price.

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Prahelika

COMPANY NAME	Prahelika
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Advanced Packaging and Heterogeneous Integration
SIZE OF TOTAL ADDRESSABLE MARKET	Prahelika is developing semiconductor integration technology for advanced electronic systems. Its technical focus includes heterogeneous integration of scaled gallium nitride dielets with glass, silicon CMOS, and single-crystal diamond interposers, alongside advanced packaging and RF, analog, and mixed-signal semiconductor technologies.
GEOGRAPHIC PRESENCE	Boston, Massachusetts
NOTABLE COMPETITORS	**LISTED:** AMD, INTC, MU, QCOM **PRIVATE:** Avicena, Ferric, Monolithic 3D
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A. Customers: N/A.
COMPETITIVE POSITIONING	Prahelika's core differentiator is its focus on heterogeneous semiconductor integration, combining compound-semiconductor dielets with silicon CMOS and advanced interposer materials to develop highly integrated electronic systems. Key attributes include: • Heterogeneous semiconductor integration: Prahelika is focused on integrating scaled gallium nitride dielets with silicon CMOS and other substrate and interposer technologies. • Multiple interposer materials: The underlying technical work includes integration with glass, silicon CMOS, and single-crystal diamond interposers. • Gallium nitride expertise: Prahelika's technical foundation includes gallium nitride processing, GaN-on-silicon and GaN-on-silicon-carbide technologies, and millimeter-wave monolithic microwave integrated-circuit design. • Advanced packaging capabilities: The company is positioned around semiconductor integration and advanced packaging rather than a single standalone semiconductor component. • Domestic integration layer: Prahelika describes its mission as building a domestic integration layer for advanced electronics.
YEAR FOUNDED / ESTABLISHED	2026
WEBSITE	https://prahelika.com/
MANAGEMENT TEAM / FOUNDER(S)	**Pradyot Yadav – Founder and CEO:** Pradyot Yadav is a Ph.D. candidate at the Massachusetts Institute of Technology's Microsystems Technology Laboratories and an NDSEG Fellow. Before founding Prahelika, Pradyot Yadav conducted work at Raytheon on 200mm W-G-band GaN-on-silicon and GaN-on-silicon-carbide MMICs, at IBM on cryogenic CMOS RF design, modeling, packaging, and AI chiplets, and at Georgia Tech on GaN MMIC designs for 5G applications using Qorvo research technology. Pradyot Yadav's research includes heterogeneous integration of GaN, silicon CMOS, glass, and single-crystal diamond interposers, advanced packaging, and RF, analog, and mixed-signal integrated-circuit design.
SIZE OF COMPANY	1
GROWTH TRAJECTORY	In August 2026, Pradyot Yadav founded Prahelika to commercialize semiconductor integration technology based on heterogeneous integration, gallium nitride devices, silicon CMOS, advanced interposers, and packaging. In September 2026, Prahelika, Inc. completed its Massachusetts registration as a Delaware corporation. As of September 2026, the company remained in its initial development phase and described itself publicly as "currently building." No product launch, commercial customer engagement, manufacturing milestone, or revenue disclosure was identified.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	N/A
LATEST ROUND VALUATION	N/A

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For semiconductor companies that are pre-revenue or at an early stage of technology development, we think it would be appropriate to value these companies using precedent venture financings, probability-adjusted milestone analysis, or comparable transaction values, as revenue-based valuation may not yet be meaningful. Peer semiconductor companies with established commercial revenue are generally valued on P/S or EV/S, while more mature and profitable semiconductor companies are generally valued on P/E or P/FCF.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Recursive Superintelligence

COMPANY NAME	Recursive Superintelligence
SECTOR / SUB-SECTOR	TMT / Software / Self-Improving Artificial Intelligence
SIZE OF TOTAL ADDRESSABLE MARKET	Recursive Superintelligence is developing AI systems capable of recursive self-improvement, where AI autonomously identifies weaknesses, develops improvements, and accelerates its own advancement. The company's long-term objective is to automate scientific discovery and AI research itself before extending into fields such as biology, chemistry, and physics. The opportunity spans frontier foundation models, AI research automation, scientific discovery, and broader artificial general intelligence infrastructure markets.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** GOOGL **PRIVATE:** OpenAI, Anthropic, Cohere, AI21, Ineffable Intelligence, Prometheus, Advanced Machine Intelligence
KEY SUPPLIERS / CUSTOMERS	Cloud-infrastructure partner: AWS. Recursive entered a $410MM multi-year collaboration with AWS in July 2026 to scale its self-improving AI platform. Customers: N/A.
COMPETITIVE POSITIONING	Recursive develops a self-improving AI platform designed to automate the generation of new knowledge rather than only retrieve or synthesize existing information. The platform uses open-ended evolutionary processes to maintain and expand an archive of diverse discoveries. Key differentiators include: • Open-ended evolutionary platform: Recursive applies continuously operating evolutionary processes rather than limiting development to a predetermined target. • Discovery archive: The platform maintains an expanding archive of diverse discoveries that can support subsequent reasoning and experimentation. • Darwinian evolution mechanisms: Recursive incorporates selection and variation mechanisms inspired by biological evolution. • Cultural evolution mechanisms: The system also uses processes inspired by the cumulative development and transfer of knowledge. • Iterative self-improvement: The platform is designed to improve its reasoning and pattern-recognition capabilities over successive iterations. • Automated knowledge discovery: Recursive targets scientific and intellectual discovery workflows rather than only general-purpose conversational use cases. • Large-scale compute collaboration: The company entered a $410MM multi-year AWS arrangement to scale the infrastructure supporting its self-improving AI systems.
YEAR FOUNDED / ESTABLISHED	2025
WEBSITE	https://www.recursive.com/
MANAGEMENT TEAM / FOUNDER(S)	**Richard Socher, Ph.D. – Co-Founder & CEO:** Dr. Socher co-founded Recursive in 2025 and serves as Chief Executive Officer, leading the company's efforts to develop recursive self-improving AI systems for scientific discovery and knowledge generation. Prior to Recursive, he co-founded and has served as CEO of You.com, an AI-powered search and AI infrastructure company. Earlier, Dr. Socher founded MetaMind, an artificial intelligence startup acquired by Salesforce in 2016, after which he served as Chief Scientist and Executive Vice President of Salesforce. Dr. Socher is widely recognized for his pioneering research in deep learning and natural language processing, including contributions to recursive neural networks and co-authoring the GloVe word-embedding framework. He earned a Ph.D. in Computer Science from Stanford University, where he later served as an adjunct professor.
SIZE OF COMPANY	~60 employees on LinkedIn (accessed 9/23/2026)

GROWTH TRAJECTORY	Recursive was founded in 2025 and has progressed into a revenue-generating, 25-person developer of self-improving AI systems. The company assembled a founding team focused on evolutionary AI, reasoning and pattern recognition and raised a $650MM Series A in May 2026. In July 2026, Recursive entered a $410MM multi-year collaboration with AWS to scale the compute infrastructure supporting its knowledge-discovery platform.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Recursive has raised a total of $655MM. In its latest Series A round of funding, it raised $650MM in May 2026. Investors include GV, Greycroft, AMD Ventures, Zoom Ventures, South Park Commons, NVIDIA, DCVC, Earlybird Venture Capital, Alumni Ventures, Headline, Wing Venture Capital, AIX Ventures, Invus Opportunities, Glasswing Ventures, Agi House, and The House Fund.
LATEST ROUND VALUATION	$4,500MM as of May 2026 Series A.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale AI software and research-platform businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference private AI model and research-platform providers such as OpenAI, Anthropic, Cohere, AI21, Google DeepMind and Ineffable Intelligence.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Ricursive Intelligence

COMPANY NAME	Ricursive Intelligence
SECTOR / SUB-SECTOR	TMT / Software / Self-Improving Artificial Intelligence
SIZE OF TOTAL ADDRESSABLE MARKET	Ricursive Intelligence develops AI systems designed to automate and accelerate semiconductor design. The company's long-term vision is to create a recursive feedback loop in which AI designs increasingly capable chips that, in turn, accelerate future AI development. Management believes semiconductor design is one of the primary bottlenecks constraining AI progress, creating a large opportunity at the intersection of EDA software, semiconductor IP, and AI infrastructure. The addressable market includes semiconductor design software, custom silicon development, and the broader AI infrastructure market.
GEOGRAPHIC PRESENCE	Palo Alto, California
NOTABLE COMPETITORS	**LISTED:** CDNS, SNPS **PRIVATE:** d-Matrix, Deep Silicon
KEY SUPPLIERS / CUSTOMERS	Suppliers and customers have not been publicly disclosed. Management has stated that the platform is intended for semiconductor companies broadly, including GPU vendors, AI accelerator companies, hyperscalers, and custom silicon developers. NVIDIA, AMD, and other chipmakers have been identified as potential customers rather than direct competitors.
COMPETITIVE POSITIONING	Ricursive's strategy is centered on applying frontier AI models to the semiconductor design process. • AlphaChip heritage: Founded by the creators of AlphaChip, a reinforcement learning-based chip floorplanning technology deployed across four generations of Google's TPU accelerators and adopted by external semiconductor companies. • AI-driven chip design automation: The company is developing systems intended to reduce chip development timelines from years to significantly shorter design cycles by automating complex engineering decisions. • Recursive improvement framework: Management's vision is to create a continuous feedback loop whereby AI accelerates semiconductor development, enabling more capable compute platforms that subsequently improve AI performance. • Full-stack semiconductor design approach: Ricursive is building technology that spans the semiconductor design workflow rather than focusing on a single design stage such as floorplanning.
YEAR FOUNDED / ESTABLISHED	2025
WEBSITE	https://ricursive.ai/
MANAGEMENT TEAM / FOUNDER(S)	**Anna Goldie, Ph.D. - Co-Founder and CEO**: Prior to founding Ricursive Intelligence, Dr. Goldie was a Senior Staff Research Scientist at Google DeepMind and an early employee at Anthropic, where she worked on Gemini and Claude. Earlier, she co-founded and led the ML for Systems team at Google Brain and co-led the AlphaChip project, which applied reinforcement learning to semiconductor floorplanning and was deployed across multiple generations of Google's TPU accelerators. Dr. Goldie holds a Ph.D. in Computer Science from Stanford University and bachelor's degrees in Computer Science and Linguistics, along with a master's degree in EECS, from MIT. **Azalia Mirhoseini, Ph.D. - Co-Founder and CTO:** Prior to Ricursive, Dr. Mirhoseini held senior AI research roles at Google DeepMind and Anthropic. She co-created AlphaChip alongside Dr. Goldie and has been a leading researcher in applying machine learning to computer systems and semiconductor design. At Ricursive, she leads technology development and platform architecture.
SIZE OF COMPANY	~33 employees on LinkedIn (accessed 9/23/2026)
GROWTH TRAJECTORY	Ricursive publicly launched in late 2025 and rapidly became one of the most highly funded AI infrastructure startups. The company raised a $35 million seed round led by Sequoia Capital before announcing a $300 million Series A financing led by Lightspeed Venture Partners at a $4 billion post-money valuation less than two months after launch. Management intends to use the capital to expand research, engineering, and compute infrastructure while developing a platform to automate semiconductor design and accelerate hardware development cycles.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Ricursive has raised approximately $335 million across seed and Series A financings. Notable investors include Lightspeed Venture Partners, Sequoia Capital, DST Global, NVentures, Felicis Ventures.
LATEST ROUND VALUATION	$4.0 billion post-money valuation per Pitchbook
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale AI software and research-platform businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference private AI model and research-platform providers such as OpenAI, Anthropic, Cohere, AI21, Google DeepMind and Ineffable Intelligence.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Source Foundry

COMPANY NAME	Source Foundry
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Semiconductor Manufacturing Technology
SIZE OF TOTAL ADDRESSABLE MARKET	Source Foundry operates in semiconductor production equipment, process automation and chip-manufacturing technology.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** TSM, GFS, INTC, Samsung Electronics, Semiconductor Manufacturing International **PRIVATE:** SirenOpt, Nearfield Instruments, Frore Systems
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Source Foundry is developing semiconductor manufacturing technology intended to increase the speed and scale of chip production. The company remains in stealth mode, limiting the availability of product-specific or technical information. Key differentiators include: • Faster semiconductor manufacturing: Source Foundry is developing manufacturing technology designed to increase the speed of semiconductor production. • Manufacturing scalability: The company's technology targets greater production scale, addressing manufacturing capacity as a key constraint in semiconductor supply. • Semiconductor production technology: Source Foundry operates at the semiconductor manufacturing and equipment layer, differentiating its focus from chip-design software and fabless semiconductor companies. • Process automation: The company's technology focus includes semiconductor process automation and process technology, supporting a more automated approach to chip manufacturing. • Manufacturing automation orientation: Source Foundry combines semiconductor-production technology with an emphasis on automation, positioning the company around improvements to physical chip-manufacturing workflows. • Stealth-stage technology development: Source Foundry remains in stealth mode, with detailed product specifications, manufacturing-process architecture, node capabilities, and customer qualification status not publicly disclosed.
YEAR FOUNDED / ESTABLISHED	2025
WEBSITE	https://sourcefoundry.com/
MANAGEMENT TEAM / FOUNDER(S)	**Abdulmalik Obaid – Co-Founder, CEO & CFO:** Prior to co-founding Source Foundry, Abdulmalik Obaid conducted research in materials science, semiconductor fabrication, and advanced manufacturing at Stanford University. He earned a Ph.D. from Stanford after previously studying physics, computer science, and mathematics at Wake Forest University. Publicly available profiles indicate that his doctoral research focused on microfabrication and advanced materials technologies, and he subsequently worked on wafer-related technology initiatives prior to founding Source Foundry. As CEO and CFO, he leads the company's efforts to develop next-generation semiconductor manufacturing technologies and advanced lithography systems intended to address bottlenecks in AI-chip production.
SIZE OF COMPANY	~50
GROWTH TRAJECTORY	Source Foundry was founded in 2025 to develop technology for faster and more scalable semiconductor manufacturing. The company remains in stealth mode but has raised $500MM across two financing rounds, including a $400MM Series A in August 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Source Foundry has raised a total of $500MM. In its latest Series A round of funding, it raised $400MM in August 2026. Investors include Situational Awareness and Sequoia Capital.
LATEST ROUND VALUATION	$5,000MM as of August 2026 Series A

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor-manufacturing and equipment businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed semiconductor manufacturers such as TSM, GFS and INTC, together with non-US-listed manufacturers such as Samsung Electronics and Semiconductor Manufacturing International and private equipment companies such as SirenOpt and Nearfield Instruments.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Substrate

COMPANY NAME	Substrate
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Next-Generation Semiconductor Fabrication
SIZE OF TOTAL ADDRESSABLE MARKET	Substrate operates in leading-edge semiconductor manufacturing, wafer production and cost-efficient fabrication services.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** GFS, INTC, MU, TSEM **PRIVATE:** SkyWater Technology, xLight, Unconventional AI, Positron AI
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Substrate is developing next-generation semiconductor manufacturing services intended to reduce the capital and operating costs of leading-edge chip fabrication. The company positions its proprietary technology as a way to extend Moore's Law while making advanced wafer production more economical. Key differentiators include: • Cost-efficient leading-edge fabrication: Substrate targets reductions in both semiconductor-fab capital expenditure and operating expense. • Next-generation manufacturing service: The company plans to provide fabrication capabilities rather than operate solely as a semiconductor-design business. • Moore's Law extension: Substrate's proprietary innovations are intended to support continued increases in semiconductor performance and density. • Lower-cost wafer production: The platform is designed to enable customers to manufacture more powerful chips at lower production cost. • US manufacturing focus: Substrate states that its objective is to restore US leadership in semiconductor fabrication. • Vertically focused operating model: The business combines proprietary process innovation with manufacturing services for advanced chips.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://substrate.com/
MANAGEMENT TEAM / FOUNDER(S)	**James Proud – Founder & CEO:** Founded Substrate and serves as Chief Executive Officer. James Proud is the company's primary contact and lead management for its October 2025 seed round. James Proud leads Substrate's development of cost-efficient, next-generation semiconductor-fabrication services.
SIZE OF COMPANY	~21 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Substrate was founded in 2022 and has progressed into a revenue-generating, 50-person semiconductor-manufacturing company. The company is developing proprietary fabrication technology intended to reduce the capital and operating costs of leading-edge fabs and wafer production. Following its $100MM seed round in October 2025, Substrate continued developing its manufacturing platform.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Substrate has raised a total of $100MM. In its latest seed round of funding, it raised $100MM in October 2025. Investors include Allen & Company of Florida, Valor Equity Partners, Founders Fund, General Catalyst, In-Q-Tel, Long Journey Ventures, Gallery Ventures, and MIT Investment Management Company.
LATEST ROUND VALUATION	$1,000MM as of October 2025 Seed Round.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor-manufacturing businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed semiconductor manufacturers such as GFS, INTC, MU and TSEM, together with private or acquired manufacturing-technology companies such as SkyWater Technology and xLight.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: SunRise Memory

COMPANY NAME	SunRise Memory
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Next-Generation Non-Volatile Memory
SIZE OF TOTAL ADDRESSABLE MARKET	SunRise Memory operates across semiconductor memory, enterprise storage, non-volatile memory and ferroelectric memory technologies.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** Weebit Nano **PRIVATE:** Ferroelectric Memory Company, Sofics
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	SunRise Memory develops next-generation non-volatile semiconductor-memory technologies. Its platform combines patented ferroelectric-memory architectures, silicon validation and integrated system engineering, targeting reliable, energy-efficient and scalable storage products. Key differentiators include: • 3D ferroelectric memory architecture: SunRise Memory Corp. is developing 3D ferroelectric RAM using low-voltage ferroelectric memory cells, targeting an alternative memory architecture for data-intensive AI and high-performance computing. • Three-dimensional memory scaling: The company's technology includes vertically scaled 3D memory structures and thin-film ferroelectric transistors, with patented architectures covering three-dimensional NOR memory arrays and vertical memory structures. • Silicon validation and device engineering: SunRise Memory combines patented ferroelectric-memory architectures with silicon-level validation as part of its development approach. • System-level memory architecture: The company's capabilities extend from semiconductor devices and processes through memory design and system architecture, including a design that manages hundreds of independent memory banks in parallel. • Manufacturing compatibility: The 3D FeRAM architecture is being designed for production using existing high-volume memory fabs with established 3D process capabilities, such as 3D NAND manufacturing infrastructure. • Extensive intellectual-property portfolio: The company has 79 patent families covering 457 patent documents, including 351 active patents, with prominent technologies spanning electronic memory devices and semiconductor-device architectures. • Industry patent citations: SunRise Memory Corp.'s patent portfolio has received forward citations from companies including Monolithic, Taiwan Semiconductor Manufacturing Company, Micron Technology, Samsung Electronics, and Macronix International.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://sunrise3d.com/
MANAGEMENT TEAM / FOUNDER(S)	**Eli Harari, Ph.D. – Chief Executive Officer & Chairman:** Leads SunRise Memory and has served as CEO and chairman since January 2016. **Richard Chernicoff – Chief Financial Officer:** Serves as the company's chief financial officer and is based in San Jose, California. **Stephen Luczo – Vice Chairman:** Serves as vice chairman and has held that position since June 2017.
SIZE OF COMPANY	~18 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	SunRise Memory was founded in 2016 and has raised $101.99 million. The company progressed through Series A, Series B and Series C financings, with its most recent disclosed funding round being an estimated $74.80 million Series C in October 2019. PitchBook identifies 80 employees as of June 2022 and records continuing patent activity through September 2026.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	SunRise Memory is venture-capital-backed and has raised approximately $101.99 million. Its disclosed institutional investors include Lam Research, Tokyo Electron and Seagate Technology. The latest disclosed financing was an estimated $74.80 million Series C completed on 7 October 2019 per Pitchbook.
LATEST ROUND VALUATION	The October 2019 Series C financing valued the company at approximately $900 million pre-money and $974.80 million post-money.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For semiconductor memory and semiconductor IP companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary as the company scales. Peer semiconductor companies are generally valued on P/E or P/FCF, as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Tela Innovations

COMPANY NAME	Tela Innovations
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Semiconductor Design IP & Standard Cell Libraries
SIZE OF TOTAL ADDRESSABLE MARKET	Operates in the semiconductor design IP, EDA enablement, and lithography optimization market, serving advanced-node semiconductor manufacturers and chip designers. The company's technologies are focused on standard cell libraries, physical design optimization, design-for-manufacturability, and process-node migration solutions.
GEOGRAPHIC PRESENCE	Los Gatos, California
NOTABLE COMPETITORS	**LISTED:** SNPS, CDNS, ARM **PRIVATE:** Silvaco, Empyrean Technology, SiFive, Pliops, Skorpios Technologies
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Tela Innovations develops lithography-optimized standard cell libraries designed to reduce die area and improve integrated circuit performance. The company focuses on helping semiconductor customers address restrictive design-rule challenges at advanced process nodes. Key differentiators include: • Lithography-optimized standard cell libraries: Tela develops proprietary standard-cell architectures intended to improve chip density and performance. • Patented layout methodology: The company's patented layout approaches help optimize integrated circuit physical implementation while complying with manufacturing constraints. • IP porting and customization services: Tela helps customers migrate and adapt existing semiconductor IP across technology nodes. • Design-for-manufacturability capabilities: The platform is designed to address restricted design-rule challenges associated with advanced semiconductor manufacturing. • Patent-backed technology portfolio: Tela's intellectual property covers mask-layout optimization, FinFET structures, interconnect architectures, and dynamic array designs.
YEAR FOUNDED / ESTABLISHED	2005
WEBSITE	https://www.tela-inc.com/
MANAGEMENT TEAM / FOUNDER(S)	**Scott Becker - Co-Founder, CEO & Board Member:** Mr. Becker co-founded Tela Innovations in 2005 and currently serves as Chief Executive Officer and board member of the company. **James Hogan - Co-Founder & Board Member:** Mr. Hogan co-founded Tela Innovations in 2005 and serves on the company's board of directors. **Dhrumil Gandhi - President & COO:** Mr. Gandhi serves as President and Chief Operating Officer of Tela Innovations.
SIZE OF COMPANY	~9 employees (accessed 9/16/2026)
GROWTH TRAJECTORY	Tela Innovations has continued developing its semiconductor design and lithography optimization technology portfolio. Recent patent activity includes dynamic array architectures, linear FinFET structures, and cross-coupled transistor layouts designed for restricted gate-level layout architectures.
ANY MAJOR M&A / CORPORATE ACTIONS	Tela Innovations acquired Blaze DFM in February 2009. In October 2021, Christopher Lucas's firm, Black Diamond Ventures, exited its investment through a private secondary transaction involving an undisclosed buyer.

INVESTORS / INVESTMENT STAGE	Tela Innovations has raised $81.7MM over five disclosed funding rounds. Before its latest financing round, the company raised $65MM in a Later Stage VC round in November 2015. Prior financings included a $4.75MM Series C round, a $5.54MM Series B1 round, a $3.86MM Series B round, and a $2.50MM Series A round. Tela Innovations investors include Taiwan Semiconductor Manufacturing Company (TSMC), Intel Capital, Qualcomm Ventures, Lightspeed Venture Partners, Opus Capital, Telos Venture Partners, Sand Hill Finance, KT Venture Group, AsiaTech Management, Cadence Design Systems, Teton Capital Partners, Western Technology Investment, and David Price. The company is venture-capital backed.
LATEST ROUND VALUATION	The latest disclosed primary financing round was a $65MM Later Stage VC financing completed in November 2015. The round valued Tela Innovations at $195MM pre-money and $260MM post-money.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For semiconductor design IP, EDA, and semiconductor infrastructure companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability may vary as the company scales. More mature semiconductor IP and EDA peers are generally valued on P/E or P/FCF given their recurring revenue profiles and established profitability.
SECONDARY PRICING DATA	The most recent disclosed transaction was a private secondary sale completed in October 2021, when Black Diamond Ventures sold its stake to an undisclosed buyer. Transaction value and implied valuation were not disclosed. The last disclosed valuation was $260MM post-money.

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: TeRAM

COMPANY NAME	TeRAM
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Advanced Memory & AI Infrastructure Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	TeRAM operates in the advanced memory, AI infrastructure, machine learning acceleration, high-performance computing (HPC), and datacenter semiconductor markets. The company is developing customized 3D SRAM memory technologies designed to accelerate AI and computational workloads.
GEOGRAPHIC PRESENCE	Los Altos, California
NOTABLE COMPETITORS	**LISTED:** MU, WDC, STM **PRIVATE:** NEO Semiconductor, EnCharge AI, Untether AI, Rain AI, Groq
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: Initial datacenter customer (undisclosed)
COMPETITIVE POSITIONING	TeRAM develops advanced memory technologies designed to accelerate machine learning and high-performance computing applications. The company is focused on enabling faster computational performance and improved energy efficiency through next-generation memory architectures. Key differentiators include: • 3D memory architecture: TeRAM integrates vertically layered thyristor cells into three-dimensional memory array clusters. • AI and HPC optimization: The technology is designed specifically for machine learning and high-performance computing workloads. • Memory-centric acceleration: The platform seeks to improve compute efficiency through innovations at the memory layer rather than relying solely on processor improvements. • Customized 3D SRAM development: The company is developing customized memory chips tailored for datacenter deployments. • Energy-efficient computing: TeRAM's architecture is designed to improve power efficiency for intensive data-processing applications.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://teram.com/
MANAGEMENT TEAM / FOUNDER(S)	**Charlie Cheng - Co-Founder & CEO:** Mr. Cheng co-founded TeRAM and currently serves as Chief Executive Officer, leading the company's commercialization and business strategy. **Harry Luan Ph.D. - Co-Founder & CTO:** Dr. Luan co-founded TeRAM and serves as Chief Technology Officer, leading development of the company's advanced memory technologies and architecture.
SIZE OF COMPANY	~8 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	In September 2026, TeRAM announced the completion of its first institutional financing round and indicated that proceeds will be used to validate its technology and scale its engineering team. The company also disclosed plans to build customized 3D SRAM memory chips for its initial datacenter customer. TeRAM's commercial focus is centered on machine learning, high-performance computing, and AI infrastructure applications.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	TeRAM has raised $37MM over one funding round. The company's latest and only disclosed financing was a $37MM Seed round completed in September 2026. TeRAM investors include Hyperion, SemiAnalysis, B Capital Group, Primary Venture Partners, U.S. Venture Partners, LDV Partners, Lightscape Partners, Alumni Ventures, Bradley Tusk, Thomas Dohmke, John Donovan, Frankwell Lin, and James Proud. The company's current investment stage is Seed.
LATEST ROUND VALUATION	N/A

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For advanced memory, AI infrastructure, and next-generation compute semiconductor companies that are pre-scale or pre-volume production, we think it would be appropriate to value these companies on P/S or EV/S, as profitability is typically not yet meaningful and commercialization timelines remain uncertain. More mature semiconductor and memory companies are generally valued on P/E or P/FCF given established manufacturing scale, revenue visibility, and profitability.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Volantis Semiconductor

COMPANY NAME	Volantis Semiconductor
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Silicon Photonics & Optical Interconnects
SIZE OF TOTAL ADDRESSABLE MARKET	Volantis Semiconductor operates in the silicon photonics, optical interconnect, AI infrastructure, datacenter networking, and high-performance computing semiconductor markets. The company develops photonic semiconductor hardware and optical chip-to-chip communication technologies for AI and HPC systems.
GEOGRAPHIC PRESENCE	Portland, Oregon
NOTABLE COMPETITORS	**LISTED:** ASML, COHR **PRIVATE:** Ayar Labs, Lightmatter, Rain AI, Vaire, OptiCore
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Volantis develops photonic semiconductor technology for optical chip-to-chip communication in AI and high-performance computing systems. Key differentiators include: • Optical interconnect architecture: The company develops photonic hardware designed to replace traditional electrical interconnects. • AI infrastructure focus: The platform is targeted at AI and HPC workloads where bandwidth and power efficiency are increasingly critical. • High-speed data transfer: Volantis seeks to improve communication between chips through optical connectivity. • Energy-efficient computing: The company's technology is designed to reduce power consumption while increasing data-transfer performance. • Photonics intellectual property: The company has developed a growing patent portfolio covering optical communication, wafer-scale integration, packaging, and power-delivery technologies.
YEAR FOUNDED / ESTABLISHED	2022
WEBSITE	https://volantissemi.ai
MANAGEMENT TEAM / FOUNDER(S)	**Tapabrata Ghosh - Founder & CEO:** Mr. Ghosh founded Volantis and leads the company's development of photonic interconnect technologies for AI and computing infrastructure. He has been the primary executive leading the company since inception.
SIZE OF COMPANY	~25 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In June 2025, Volantis unveiled its photonic compute platform targeting AI infrastructure applications. The company continues to expand its optical interconnect and photonic computing roadmap and has recently developed intellectual property around wafer-scale integration, modular power delivery, laser-assisted bonding, and advanced packaging technologies. Recent patent activity accelerated throughout 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	A private secondary transaction occurred in which IceNine Capital sold its stake to an undisclosed buyer.
INVESTORS / INVESTMENT STAGE	Volantis has raised $100MM across multiple funding rounds. Prior to its latest $80MM Series A / Series A1 financing completed in July 2026, the company raised approximately $11.3MM in 2024 and $9MM in 2025. The company also received pre-seed funding in 2022. Volantis investors include YO Ventures, Quain Investments, Techni Ventures, Entropia Capital, Mana Industries, VXI Capital Management, Unpopular Ventures, IceNine Capital, Vibe Cap, Sam Altman, Alex Wang, Trevor Blackwell, Jun Hwang, Corey Nobile, and other strategic investors. Current investment stage is Series A.
LATEST ROUND VALUATION	Volantis completed an $80MM Series A financing in July 2026 at a $281MM pre-money valuation and a $361MM post-money valuation.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For silicon photonics, optical interconnect, and advanced AI infrastructure semiconductor companies that remain sub-scale, we think it would be appropriate to value these companies on P/S or EV/S, as profitability is typically secondary to technology adoption and hyperscaler deployment opportunities. More mature networking, optical, and semiconductor infrastructure companies are generally valued on P/E or P/FCF once revenue visibility and profitability are established.
SECONDARY PRICING DATA	A private secondary transaction was completed involving IceNine Capital and an undisclosed buyer. Transaction value and implied valuation were not disclosed.

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: xLight

COMPANY NAME	xLight
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Semiconductor Lithography & Light Source Systems
SIZE OF TOTAL ADDRESSABLE MARKET	xLight operates in the semiconductor lithography equipment, advanced manufacturing, semiconductor process technology, EUV light source, and wafer-patterning markets. The company develops advanced light-source systems used in semiconductor manufacturing.
GEOGRAPHIC PRESENCE	Palo Alto, California
NOTABLE COMPETITORS	**LISTED:** ASML, VECO, NINOY **PRIVATE:** Cymer, GIGAPHOTON, Ayar Labs, Snowcap, Substrate
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	xLight develops advanced light-source systems designed for semiconductor manufacturing. The technology is focused on improving lithography performance while reducing energy and water consumption. Key differentiators include: • Programmable light characteristics: The company develops configurable light-source technologies optimized for semiconductor manufacturing. • Lower operating costs: xLight's systems are designed to reduce power and water usage by up to eight times per wafer. • Yield and productivity improvements: The technology aims to improve lithography patterning precision and fab productivity. • Advanced manufacturing focus: The platform is designed for next-generation semiconductor production processes. • Proprietary light-source technology: The company has developed intellectual property covering laser systems, particle accelerators, undulator systems, plasma technologies, and X-ray-based light-source architectures.
YEAR FOUNDED / ESTABLISHED	2021
WEBSITE	https://xlight.com/
MANAGEMENT TEAM / FOUNDER(S)	**Nicholas Kelez - CEO & CTO:** Mr. Kelez serves as both Chief Executive Officer and Chief Technology Officer of xLight and leads development of the company's advanced semiconductor light-source platform. **Patrick Gelsinger - Executive Chairman:** Mr. Gelsinger previously served as CEO of Intel and VMware and is one of the most prominent executives in the semiconductor industry. He has served as Executive Chairman of xLight since 2022. **Kevin Heidrich - Chief Commercial Officer:** Mr. Heidrich leads commercialization and go-to-market activities for xLight's semiconductor manufacturing platforms.
SIZE OF COMPANY	~59 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	In December 2025, xLight signed a $150MM letter of intent with the U.S. Department of Commerce supporting development of next-generation lithography light-source technology. In June 2026, the Department of Commerce finalized CHIPS Act incentives supporting xLight's platform. The company also announced advances in EUV laser technology and continues developing programmable semiconductor light-source systems focused on reducing power and water consumption in semiconductor fabrication.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	xLight has raised $202MM across multiple funding rounds. Prior to its latest announced $350MM Later Stage VC financing in June 2026, the company raised approximately $191MM in a Series B financing completed in December 2025 and approximately $11MM in a Series A financing completed in December 2022. xLight investors include Bain Capital, Boardman Bay Capital Management, IAG Capital Partners, Marvel Capital, Open Field Capital, Kape Capital, Playground Global, Morpheus Ventures, Lam Capital, Roadster Capital, Schox Venture Capital, U.S. Department of Commerce, Alchemist Accelerator, and Stephen Cole. Current investment stage is Later Stage VC.
LATEST ROUND VALUATION	The latest disclosed completed valuation occurred during the December 2025 Series B financing, which valued xLight at approximately $439MM pre-money and $630MM post-money. The announced June 2026 financing did not disclose an updated valuation.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For semiconductor equipment, lithography, and advanced manufacturing technology companies that remain in commercialization mode, we think it would be appropriate to value these companies on P/S or EV/S, as profitability is often impacted by product-development cycles and customer qualification timelines. More mature semiconductor equipment companies are generally valued on P/E or P/FCF given stronger profitability, recurring service revenue, and established shipment volumes.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Semiconductor Manufacturing: Zero ASIC

COMPANY NAME	Zero ASIC
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semiconductor Equipment / Manufacturing
SIZE OF TOTAL ADDRESSABLE MARKET	Zero ASIC is developing a composable chiplet platform and automated silicon design infrastructure intended to dramatically reduce the cost, complexity, and development time associated with custom semiconductor design. The company is targeting the growing market opportunity created by custom silicon, chiplet-based architectures, AI accelerators, edge computing devices, defense applications, and domain-specific processors. Management believes traditional ASIC development remains prohibitively expensive for most organizations, often requiring years of engineering effort and development budgets exceeding tens of millions of dollars. The company's mission is to "democratize silicon" by enabling semiconductor developers to configure and assemble custom systems from standardized chiplets in a manner analogous to assembling software or catalog components. As the semiconductor industry transitions toward chiplet-based architectures, advanced packaging, heterogeneous integration, and custom AI silicon, Zero ASIC is positioning itself as a platform provider that reduces barriers to custom hardware development.
GEOGRAPHIC PRESENCE	Cambridge, Massachusetts
NOTABLE COMPETITORS	**LISTED:** SNPS, CDNS, MRVL, AMD **PRIVATE:** Efabless, ChipFlow, Baya Systems, Arteris, Tenstorrent Zero ASIC competes across multiple categories, including electronic design automation (EDA), chiplet ecosystems, embedded FPGA IP, design automation, and custom silicon infrastructure.
KEY SUPPLIERS / CUSTOMERS	The company has not publicly disclosed commercial customers. Zero ASIC's platform is designed for semiconductor companies, defense contractors, silicon developers, hardware startups, embedded system providers, and organizations seeking to develop custom ASICs without maintaining large in-house design teams. The company leverages industry-standard semiconductor manufacturing ecosystems and supports integration with multiple foundry technologies and design environments.
COMPETITIVE POSITIONING	Zero ASIC's primary differentiator is its attempt to combine standardized chiplets, automation software, and open hardware standards into a unified silicon development platform. • Composable chiplet architecture: The company has developed a platform based on standardized chiplets that can be assembled into customized silicon systems. Management describes this approach as enabling large numbers of potential system configurations from reusable building blocks. • Design automation focus: Zero ASIC develops software platforms that automate multiple stages of semiconductor design, packaging, validation, and implementation. Management believes automation can significantly compress custom chip development timelines. • Active silicon interposer technology: The company's eFabric architecture separates networking and compute resources through an active 3D silicon substrate designed to improve chiplet composability and communication efficiency. Management reports communication efficiency below 0.1 pJ/bit. • Open-source ecosystem strategy: The company has released several open-source tools including SiliconCompiler, Wildebeest FPGA synthesis software, static timing analysis tools, and other automation technologies intended to reduce dependence on proprietary design environments. • Embedded FPGA product portfolio: In 2025 and 2026, Zero ASIC introduced its Platypus family of embedded FPGA products, which management describes as the industry's first open-standard eFPGA platform designed to avoid vendor lock-in and FPGA obsolescence. • AI-driven chip design tools: In 2026 the company introduced Tardigrade, an ASIC synthesis engine designed specifically for increasingly automated and AI-assisted silicon development workflows.

YEAR FOUNDED / ESTABLISHED	2008
WEBSITE	https://www.zeroasic.com/
MANAGEMENT TEAM / FOUNDER(S)	**Andreas Olofsson - Founder and CEO:** Andreas Olofsson founded the company in 2008 as Adapteva. Prior to founding the company, Olofsson worked as a processor architect and led development efforts on Analog Devices' TigerSHARC DSP product line. He established Adapteva with the goal of improving computing energy efficiency through highly parallel processor architectures.
SIZE OF COMPANY	10-15
GROWTH TRAJECTORY	The company was founded in 2008 as Adapteva, initially focused on developing highly efficient many-core processors based on its proprietary Epiphany architecture. During this period, the company released multiple generations of processors and gained recognition for energy-efficient parallel computing. Adapteva achieved several technical milestones, including development of 16-core, 64-core, and 1024-core processors, shipment of products to thousands of customers, participation in DARPA-related programs, and completion of a successful Kickstarter campaign for its Parallella development platform. In 2020, the company rebranded as Zero ASIC and shifted its strategic focus toward chiplets, advanced packaging, semiconductor automation, and custom silicon development platforms. Management believed chiplets represented a more scalable path toward democratizing semiconductor design. In October 2023, Zero ASIC emerged from stealth and announced its automated chiplet design and emulation platform alongside its eFabric active interposer technology. Management outlined a roadmap targeting substantial reductions in ASIC development cost and complexity. Product momentum accelerated during 2025 and 2026 with the launch of the Platypus open-standard eFPGA platform, Wildebeest FPGA synthesis software, PackageCompiler automated packaging software, and Tardigrade ASIC synthesis technology.
ANY MAJOR M&A / CORPORATE ACTIONS	Rebranding from Adapteva to Zero ASIC
INVESTORS / INVESTMENT STAGE	The company raised early venture financing following its formation, including $1.5MM in Series A funding during its Adapteva period. Notable investors include Lockheed Martin Ventures and Viola Ventures.
LATEST ROUND VALUATION	N/A
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For semiconductor infrastructure, chiplet platform, EDA, and FPGA IP companies with limited commercial revenue, EV/Sales and P/Sales are generally the most relevant valuation frameworks given significant ongoing R&D investment. More mature semiconductor IP, design software, and infrastructure providers are typically valued using EV/EBITDA, P/E, and free cash flow methodologies once recurring licensing and royalty streams become established.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Figure 10: Quantum Computing Companies & Public Peers

Companies	Public Peers
Driven Quantum Technologies	QBTS, IONQ, RGTI, INFQ, PSQL
Entanglement Consulting	IBM, QBTS, IONQ, RGTI, MSFT, GOOGL
PSIQuantum	IBM, QBTS, RGTI
QC Ware	SDGR
QuEra Computing	INFQ, IONQ, RGTI, QBTS, IBM

Source: UBS

Quantum Computing: Driven Quantum Technologies

COMPANY NAME	Driven Quantum Technologies
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Quantum Computing
SIZE OF TOTAL ADDRESSABLE MARKET	Driven Quantum Technologies operates across quantum computing, quantum communications, neuromorphic hardware and dual-use quantum systems for defense, intelligence, medical, assistive and law-enforcement applications.
GEOGRAPHIC PRESENCE	Chevy Chase, Maryland
NOTABLE COMPETITORS	**LISTED:** QBTS, IONQ, RGTI, INFQ, PSQL **PRIVATE:** Delft Circuits, eleQtron, Biobab AiBIO
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Driven Quantum Technologies develops periodically driven and nonequilibrium quantum technologies for specialized operational and analytical applications. Key differentiators include: • Floquet-engineered quantum systems: The company applies periodic driving techniques to engineer and manipulate quantum-system behavior, differentiating its approach from quantum architectures centered primarily on static system properties. • Nonequilibrium quantum control: Its technology focuses on controlling periodically driven quantum systems, using dynamic interactions as a mechanism for creating and managing desired quantum behavior. • Dual-use technology platform: The architecture is positioned for both government and civilian applications, providing a common quantum-technology foundation across specialized end markets. • Defense and intelligence applications: Military and intelligence use cases represent an important part of the company's application focus, reflecting its orientation toward mission-specific quantum technologies. • Broader specialized applications: The platform also has exposure to medical, assistive, and law-enforcement applications, extending its potential use cases beyond defense. • Cross-technology architecture: The company's technology positioning spans quantum communications as well as adjacent areas including neuromorphic hardware and human-computer interaction.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	https://drivenquantumtechnologies.com/
MANAGEMENT TEAM / FOUNDER(S)	**Jerome Strombosky – Founder, CEO & Board Member:** Founded Driven Quantum Technologies and has served as a board member since January 2019. Jerome Strombosky is the company's primary contact and lead executive for its completed angel financing and proposed Series A. Jerome Strombosky leads the development and commercialization of the company's Floquet-engineered and nonequilibrium quantum technologies.
SIZE OF COMPANY	~5 employees (accessed 9/16/2026)
GROWTH TRAJECTORY	Driven Quantum Technologies has progressed from an angel-backed startup into a revenue-generating developer of Floquet-engineered quantum systems. The company is developing periodically driven quantum-control methods and dual-use platforms for defense, intelligence, medical, assistive and law-enforcement applications.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Driven Quantum Technologies has raised an undisclosed amount of funding. In its completed angel round of funding, it raised an undisclosed amount in April 2019. Investors: N/A.
LATEST ROUND VALUATION	N/A

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale quantum-computing and quantum-technology businesses, P/S or EV/Sales may be appropriate once revenue is generated. Comparable-company analysis may reference listed quantum-computing providers such as QBTS, IONQ, RGTI, INFQ and PSQL, together with private quantum-hardware companies such as Delft Circuits and eleQtron.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company reports and website

Quantum Computing: Entanglement Consulting

COMPANY NAME	Entanglement Consulting
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Quantum Computing
SIZE OF TOTAL ADDRESSABLE MARKET	McKinsey estimates the overall quantum-computing technology market could reach $4.4B by 2028 vs $1B+ of industry revenue in 2025. Entanglement's near-term addressable opportunity spans substantially larger AI, cybersecurity, optimization, defense and HPC markets.
GEOGRAPHIC PRESENCE	Miami, Florida
NOTABLE COMPETITORS	**LISTED:**IBM, QBTS, IONQ, RGTI, MSFT and GOOG **PRIVATE:**Quantinuum, 1QBit, Multiverse Computing, QC Ware, Classiq, Q-CTRL and SandboxAQ
KEY SUPPLIERS / CUSTOMERS	Key partners include Groq, Oracle, and Cassava Technologies. Oracle deployed seQure's Ground-Truth platform on Oracle Cloud Infrastructure in 2025, while Cassava signed a strategic MoU with Entanglement in November 2025 to expand AI, cybersecurity, and quantum-inspired solutions across Africa and the Middle East. The U.S. Army is Entanglement's most prominent disclosed government customer, having validated elements of the company's hybrid quantum/ classical cybersecurity technology. Entanglement is deploying its Quantum Logic platform across defense, logistics, healthcare, telecommunications, and other sectors through affiliated operating companies, although most end customers remain undisclosed.
COMPETITIVE POSITIONING	Entanglement operates as a platform company that commercializes its Quantum Logic Core (QLC) through specialized businesses rather than as a standalone software vendor. Its technology combines quantum-inspired methods, advanced mathematics, and AI while running on conventional computing infrastructure. Key differentiators include: • Quantum Logic Core (QLC): Entanglement's QLC combines quantum-inspired methods, advanced mathematics, and AI while running on conventional computing infrastructure, enabling near-term deployment without requiring fault-tolerant quantum hardware. • Specialized businesses: The company commercializes QLC through specialized businesses including seQure for cybersecurity, Maritech Intelligence Systems for maritime applications, and Hyper Intelligence for AI infrastructure. • U.S. Army work: Entanglement reported materially faster anomaly-detection performance in its work with the U.S. Army using its technology stack deployed on Groq hardware. • Quantum and quantum-inspired optimization IP: The company has secured IP covering quantum and quantum-inspired optimization applications in logistics, routing, scheduling, and supply-chain management.
YEAR FOUNDED / ESTABLISHED	2017
WEBSITE	www.entanglement.com
MANAGEMENT TEAM / FOUNDER(S)	**Jason Turner — Founder, Chairman and CEO:** Jason has more than two decades of experience spanning advanced computing, digital media, entrepreneurship and investing. He is also CEO of private investment company Sensate Holdings. Turner was a founding participant in the U.S. QED-C and Quantum Industry Coalition and has advised the White House on workforce-development policy. **Dr. Fred Glover — Chief Scientific Officer:** Glover is one of the foundational researchers in modern combinatorial optimization and the creator of Tabu Search, as well as major contributions to scatter search and metaheuristics. He received the John von Neumann Theory Prize and was elected to the U.S. National Academy of Engineering. His optimization research forms an important intellectual foundation for QLC. **Christopher Verlinden — Chief Technology Officer:** Verlinden leads the company's technology strategy and product development. In addition to his role at Entanglement, he has been involved in ocean and maritime technology initiatives, including Applied Ocean Sciences, supporting the company's expansion into maritime intelligence and autonomous systems applications.
SIZE OF COMPANY	~45 employees according to various sources

GROWTH TRAJECTORY	Entanglement has evolved from an advanced-computing and quantum-inspired R&D company into a commercialization platform built around its QLC. Early work with U.S. government organizations provided external validation of the technology, which the company has since expanded through partnerships with Groq, Oracle, and Cassava Technologies. Entanglement has commercialized QLC through specialized operating businesses, including seQure (cybersecurity), Hyper Intelligence (AI infrastructure), and Maritech Intelligence Systems, which was launched following the acquisition of Applied Ocean Sciences in 2026. A key milestone came in May 2025, when seQure's Ground-Truth cybersecurity platform became available on Oracle Cloud Infrastructure. Entanglement states that QLC is deployed across applications spanning cybersecurity, defense, logistics, maritime, and healthcare, although most end-customer relationships remain undisclosed.
ANY MAJOR M&A / CORPORATE ACTIONS	Entanglement expanded its maritime capabilities through the acquisition of Applied Ocean Sciences (AOS) in 2026, adding ocean-acoustics, oceanographic modeling, and large-scale sensing expertise. The company subsequently launched Maritech Intelligence Systems, combining AOS's maritime-domain capabilities with Entanglement's Quantum Logic Core (QLC) and AI platform. More broadly, Entanglement has adopted a portfolio-company model, commercializing QLC through specialized operating businesses spanning cybersecurity (seQure), maritime intelligence (Maritech), and AI infrastructure (Hyper Intelligence), with additional initiatives in healthcare, biosurveillance, and other advanced-computing applications.
INVESTORS / INVESTMENT STAGE	Funding amounts, valuation, and institutional investors have not been publicly disclosed.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company reports and website

Quantum Computing: PSIQuantum

COMPANY NAME	PSIQuantum
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Quantum Computing
SIZE OF TOTAL ADDRESSABLE MARKET	PSIQuantum is developing a fault-tolerant quantum computer using photonic qubits (single photons transmitted through silicon photonic chips). They aim for a million-qubit quantum computer capable of solving chemistry and optimization problems beyond classical supercomputers. The company operates in the ~$1B quantum computing market.
GEOGRAPHIC PRESENCE	Palo Alto, California
NOTABLE COMPETITORS	**LISTED:** IBM, QBTS, RGTI **PRIVATE:** Quantinuum, Pasqal, QuEra, Atom Computing
KEY SUPPLIERS / CUSTOMERS	Suppliers include GFS (foundry).
COMPETITIVE POSITIONING	PSIQuantum is building a utility-scale quantum computer based on photonic qubits, targeting fault-tolerant quantum computing while leveraging existing semiconductor manufacturing, packaging, and cryogenic cooling technologies. Key differentiators include: • Photonic qubits: PSIQuantum utilizes photonic qubits that do not require exotic materials or manufacturing processes, instead using repurposed integrated photonics technology. • Fault-tolerant architecture: The company aims to build useful, fault-tolerant quantum computers capable of running algorithms with an advanced architecture, using entangled states to ensure error-correcting codes. • Conventional packaging and cooling: The majority of the company's packaging utilizes conventional cooling technologies, while PSIQuantum is testing cryogenic cabinets that offer better cooling power than dilution refrigerators.
YEAR FOUNDED / ESTABLISHED	2015
WEBSITE	https://www.psiquantum.com/
MANAGEMENT TEAM / FOUNDER(S)	**Jeremy O'Brien - Co-founder and CEO:** Prior to founding the company, O'Brien was a Professor of Physics and Electrical Engineering at Stanford and Bristol Universities as well as Director of the Center for Quantum Photonics. **Mark Thompson - Co-founder and CTO:** Prior to founding the company, Thompson worked at Corning Cable, Bookham Technology, and Toshiba alongside being part of the faculty at Cambridge University and Bristol University. **Pete Shadbolt - Co-founder and CSO:** Prior to founding the company, Shadbolt obtained his PhD in experimental photonic quantum computing at the University of Bristol in 2014 and postdoc at Imperial College while researching photonic quantum computing theory. **Terry Rudolph - Co-founder and Chief Architect:** Rudolph is the inventor for the quantum computing approach at PsiQuantum. Prior to founding the company, he held positions at the University of Toronto and was the Professor of Quantum Physics at Imperial College.
SIZE OF COMPANY	~660 employees on LinkedIn (as of 8/24/2026)
GROWTH TRAJECTORY	The company is building the world's first utility-scale fault-tolerant quantum computer in Brisbane, Australia set to be operational by late 2027. Furthermore, the company recently announced a utility-scale, fault-tolerant quantum computer in Illinois Quantum and Microelectronics Park (IQMP).
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	PSIQuantum raised $1.7B in private funding since inception, including a $1B Series E in 2025 at a $7B valuation. Including the AUD 940M (~$620M) of Australian equity, grants and loans committed to its Brisbane project, total financing secured is roughly $2.3B. Investors: Queensland Government, Australian Government, M12 - Microsoft's Venture Fund, Blackrock, Blackbird Ventures, Baillie Gifford, Temasek Holdings, SevenTrainVentures, Scale-Up, TEL Venture Capital, Stanton Green, New Ground Ventures, MFV Partners, WorldQuant Ventures LLC, Redpoint, Open Field Capital, A&E Investments, The House Fund, Founders Fund, Atomico, Roadster Capital, Playground Global, C4 Ventures, SciFi VC.
LATEST ROUND VALUATION	$7B as of the Sept 2025 Series E funding round.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For quantum computing companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer quantum computing companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company reports and website

Quantum Computing: QC Ware

COMPANY NAME	QC Ware
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Quantum Computing
SIZE OF TOTAL ADDRESSABLE MARKET	In a Dec 2025 interview, CEO Matt Johnson called out a ~$2B current TAM for QC Ware's Promethium SaaS offering. As Promethium expands across drug discovery, materials science and broader molecular simulation workflows, management believes the TAM could expand into the tens of billions of dollars.
GEOGRAPHIC PRESENCE	Palo Alto, California
NOTABLE COMPETITORS	**LISTED:** SDGR **PRIVATE:**SandboxAQ, QunaSys, QSimulate, HQS Quantum Simulations
KEY SUPPLIERS / CUSTOMERS	Key technology partners include NVIDIA and AWS, with Promethium leveraging NVIDIA GPU acceleration and AWS cloud infrastructure. Disclosed enterprise and government customers/ partners include JPMorgan Chase, Goldman Sachs, BMW, Airbus, Covestro, Roche, Itau Unibanco, AISIN, POSCO Holdings, Boehringer Ingelheim, TotalEnergies, Equinor and the US Air Force Research Labs.
COMPETITIVE POSITIONING	QC Ware has shifted its focus from broad quantum-algorithm development to computational chemistry, using GPU acceleration to deliver near-term value while preserving upside from future quantum hardware. Key differentiators include: • Promethium platform: QC Ware's Promethium platform accelerates high-accuracy electronic-structure calculations for pharmaceutical, materials, and chemical R&D, providing a nearer-term software monetization path than approaches dependent on fault-tolerant quantum computing. • GPU-native algorithms and quantum-inspired methods: The company differentiates from traditional computational-chemistry vendors through GPU-native algorithms and quantum-inspired methods. • Today's accelerated-computing infrastructure: QC Ware differentiates from quantum-software peers through a product that runs on today's accelerated-computing infrastructure.
YEAR FOUNDED / ESTABLISHED	2014
WEBSITE	https://qcware.com
MANAGEMENT TEAM / FOUNDER(S)	**Matt Johnson – Co-founder and CEO:**Leads the business and serves on the governing board of the Quantum Economic Development Consortium (QED-C). Previously served as a US Air Force captain and a Managing Director at Apollo. **KJ Sham – Co-founder and COO:**Oversees day-to-day operations and QC Ware's three principal business units: Promethium, Quantum Services and Q2B. KJ has previously founded and exited multiple medical devices and high-tech consumer product businesses. **John Hynes – President and CFO:** John oversees the finance, fundraising, treasury and corporate strategy functions at QC Ware.
SIZE OF COMPANY	QC Ware has slightly under 50 employees (as of 8/24/2026)
GROWTH TRAJECTORY	QC Ware has evolved from a quantum algorithms and services business into a product-led software company centered on Promethium, its GPU-accelerated computational chemistry platform launched in 2023. Promethium allows the company to monetize its quantum expertise on today's NVIDIA GPU infrastructure while preserving a pathway to hybrid and quantum-accelerated workflows as hardware matures. The platform has expanded to 25+ capabilities across electronic structure, molecular interactions, reaction modeling, molecular properties, and drug discovery, targeting pharmaceutical and materials-science customers. QC Ware has also built relationships with enterprises including JPMorgan Chase, Goldman Sachs, BMW, Airbus, Covestro, and Roche, while deepening its alignment with NVIDIA's accelerated and quantum computing ecosystem.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	QC Ware has raised ~$32M across a $25M Series A equity round led by Citi and Goldman Sachs in 2018, and a $6.5M Series B equity round led by Koch Disruptive Technologies and Covestro in 2021. Investors: Citi, Goldman Sachs, Koch Disruptive Technologies, Covestro, Airbus Ventures, D.E. Shaw, Samsung Ventures, Pegasus Tech Ventures, and Fenox Venture Capital.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For quantum computing companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer quantum computing companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company reports and website

Quantum Computing: QuEra Computing

COMPANY NAME	QuEra Computing
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Quantum Computing
SIZE OF TOTAL ADDRESSABLE MARKET	McKinsey estimates the overall quantum-computing technology market could reach $4.4B by 2028 vs $1B+ of industry revenue in 2025. QuEra's directly addressable market is a subset of this opportunity centers on neutral-atom quantum computing and, increasingly, fault-tolerant quantum system.
GEOGRAPHIC PRESENCE	Boston, Massachusetts
NOTABLE COMPETITORS	**LISTED:**INFQ, IONQ, RGTI, QBTS, IBM **PRIVATE:**Pasqal, Atom Computing, PsiQuantum, planqc
KEY SUPPLIERS / CUSTOMERS	QuEra's largest disclosed customer is AIST, which awarded a ¥6.5B (~$41M) contract for a neutral-atom quantum computer integrated with the NVIDIA-powered ABCI-Q supercomputer. The company has also been selected to provide a quantum testbed for the UK's national quantum program. Aquila has been available on Amazon Braket since 2022. Key partners include AWS for cloud distribution and future access to QuEra's Libra fault-tolerant system, and NVIDIA for CUDA-Q-enabled hybrid quantum-classical computing.
COMPETITIVE POSITIONING	QuEra develops neutral-atom quantum computing systems designed around large, reconfigurable qubit arrays, effectively all-to-all connectivity, and a scalable path toward fault-tolerant quantum computing. The company increasingly focuses on quantum error correction and logical-qubit performance rather than raw physical-qubit counts. Key differentiators include: • Neutral-atom architecture: QuEra's architecture combines large, reconfigurable qubit arrays with effectively all-to-all connectivity and room-temperature operation with relatively low infrastructure requirements. • Quantum error correction: The company's architecture is designed to support advanced error-correction schemes that could reduce physical-qubit overhead, with increasing emphasis on logical-qubit performance. • Aquila: QuEra's Aquila system has been available through Amazon Braket since 2022 as part of the company's longstanding AWS partnership. • Libra: QuEra is targeting its 256-logical-qubit Libra fault-tolerant system for 2028.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://quera.com
MANAGEMENT TEAM / FOUNDER(S)	**Andy Ory – CEO:**Veteran technology entrepreneur and former CEO/co-founder of Acme Packet, which was acquired by Oracle. He leads QuEra's commercialization and scaling efforts. **Vladan Vuletić – Co-founder and CTO:**MIT physics professor and a pioneer in neutral-atom quantum computing. He is one of the principal architects of QuEra's technology platform and fault-tolerance roadmap. **Dr. Fred Furtek – Co-founder and CSO:** Harvard physics professor and a leading researcher in quantum information science whose work underpins QuEra's neutral-atom approach.
SIZE OF COMPANY	QuEra has ~200 employees
GROWTH TRAJECTORY	QuEra has rapidly evolved from a research-driven neutral-atom startup into a commercial quantum-hardware company, advancing from its 256-qubit Aquila system on Amazon Braket in 2022 to the 260-qubit Gemini platform and its first on-premises deployments, including Japan's AIST installation alongside the NVIDIA-powered ABCI-Q supercomputer. Backed by $230M+ raised in 2025, the company is expanding manufacturing capacity and supply-chain capabilities ahead of fault-tolerant systems. Its roadmap targets Libra in 2028 with 256+ logical qubits and 10,000+ physical qubits, followed by a 1,000+ logical-qubit, gigaquop-class successor. An expanded AWS partnership will bring Libra to Amazon Braket, while collaborations with NVIDIA and other ecosystem partners support the commercialization of fault-tolerant quantum computing.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A

INVESTORS / INVESTMENT STAGE	QuEra has raised $247M in disclosed equity funding, including $17M in 2021, an undisclosed round in 2023, and >$230M in convertible financing in February 2025. Subsequently, the company increased funding through an undisclosed strategic investment from NVIDIA's NVentures in September 2025. Investors: Google Quantum AI, SoftBank Vision Fund 2, NVentures, Valor Equity Partners, QVT Family Office, Safar Partners, Rakuten, and Day One Ventures.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For quantum computing companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer quantum computing companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company reports and website

Figure 11: AI DC Power Management Companies & Public Peers

Companies	Public Peers
Akash Systems	COHU
Ambature	AMSC, OXIG, BRKR
Emerald AI	Schneider Eletric, Siemens, ETN, GEV, ENEL, HON
Ferric	MPWR, TXN, ADI, Renesas, VICR
Power Lattice	SLAB, ADI, TXN, IFX
Reed Semiconductor	ADI, IFX, ON, STM, TXN, MPWR, Renesas

Source: UBS

AI DC Power Management: Akash Systems

COMPANY NAME	Akash Systems
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Diamond-Based Thermal Management
SIZE OF TOTAL ADDRESSABLE MARKET	Akash Systems operates across AI data-center thermal management, semiconductor cooling, high-power RF devices and satellite communications.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** COHU **PRIVATE:** Diamond Foundry, Infleqtion, TeraXion, Castelion, Accelsius, Cambium
KEY SUPPLIERS / CUSTOMERS	Akash Systems does not disclose suppliers. Disclosed customer: NxtGen AI. Akash Systems delivered diamond-cooled NVIDIA GPU servers to NxtGen AI in February 2026.
COMPETITIVE POSITIONING	Akash Systems develops thermal-management hardware that combines synthetic-diamond substrates with semiconductor materials to improve heat dissipation and system efficiency. Key differentiators include: • Diamond-based cooling: Synthetic-diamond materials are used to dissipate heat from high-power semiconductor devices. • GaN-on-diamond technology: The platform integrates gallium nitride semiconductor layers with diamond substrates. • AI data-center applications: Akash applies its thermal technology to high-performance GPU servers. • Aerospace and satellite exposure: The technology also supports RF and millimeter-wave devices used in communications systems. • Reduced thermal constraints: Improved heat dissipation is intended to reduce system size and power consumption while supporting higher performance.
YEAR FOUNDED / ESTABLISHED	2015
WEBSITE	https://akashsystems.com/
MANAGEMENT TEAM / FOUNDER(S)	**Felix Ejeckam, Ph.D. – Co-Founder, Co-CEO & Board Member:** Co-founded Akash Systems and has served as a board member since December 2015. Felix Ejeckam has led the development and commercialization of Akash's GaN-on-diamond semiconductor and thermal-management technology. **Soham Khanna – Co-CEO:** Serves as Co-Chief Executive Officer of Akash Systems. The role supports the company's commercialization of diamond-cooled semiconductor and AI data-center infrastructure products. **Daniel Francis, Ph.D. – Chief Technology Officer:** Serves as Chief Technology Officer. The role is directly relevant to Akash's development of GaN-on-diamond materials, high-power RF devices and diamond-cooled GPU-server technology.
SIZE OF COMPANY	~177 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Akash Systems has evolved from a developer of GaN-on-diamond semiconductor materials into a supplier of thermal-management hardware for aerospace, satellite communications and AI infrastructure. The company received an $18.2MM U.S. Department of Commerce grant in 2024 and subsequently expanded its diamond-cooling technology into GPU-server applications. In February 2026, Akash delivered diamond-cooled NVIDIA GPU servers to NxtGen AI, providing a disclosed commercial deployment of its technology in AI infrastructure.
ANY MAJOR M&A / CORPORATE ACTIONS	In May 2026, Social Capital, DCVC, Ruvento Ventures, Backstage Capital, BoomStartup and Noveus Capital sold their stakes in Akash Systems to an undisclosed buyer. The consideration and per-share price were not disclosed. In September 2019, Polymath Capital Partners, Sego.VC and two individual shareholders sold their stakes to Sax Capital for an undisclosed amount.

INVESTORS / INVESTMENT STAGE	Akash Systems has raised a total of $83.6MM. In its latest Series B round of funding, it raised approximately $66MM in July 2025. Investors include ACME Capital, Gaingels, Embedded Ventures, Realm Capital Ventures, Cade Ventures, Banyan Ventures, Broadshade Investments, Kearny Jackson, Solaris Ventures, Hypersonic Ventures, 2468 Ventures, TLI Bedrock, Invariantes Fund, Founders Fund, Khosla Ventures, and the U.S. Department of Commerce.
LATEST ROUND VALUATION	$537.5MM as of July 2025 Series B.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor and thermal-management businesses, P/S or EV/Sales may be appropriate once revenue is generated. Comparable-company analysis may reference COHU, together with private advanced-materials and data-center cooling companies such as Diamond Foundry, Accelsius and Cambium.
SECONDARY PRICING DATA	Private secondary transactions were completed in September 2019 and May 2026.

Source: Pitchbook and Company's reports and website, LinkedIn

AI DC Power Management: Ambature

COMPANY NAME	Ambature
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Superconducting Materials & Devices
SIZE OF TOTAL ADDRESSABLE MARKET	Ambature operates across high-temperature superconducting materials, quantum devices, energy transmission, magnetic sensing and quantum-computing infrastructure.
GEOGRAPHIC PRESENCE	Scottsdale, Arizona
NOTABLE COMPETITORS	**LISTED:** AMSC, OXIG, BRKR **PRIVATE:** Luvata, Bluefors, Absolics, PsiQuantum, Ayar Labs.
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Ambature develops high-temperature superconducting materials and quantum devices for energy-efficient computing and power applications. Key differentiators include: • A-axis superconducting systems: Ambature develops superconducting structures intended to improve device performance. • High-temperature superconducting materials: The technology is designed to support energy-efficient electronic and power applications. • Quantum-device applications: Ambature develops Josephson junction technologies relevant to quantum-computing systems. • Power and sensing exposure: Potential applications include energy transmission and magnetic sensing. • Low-resistance materials: The company's intellectual property covers devices formed from extremely low-resistance materials.
YEAR FOUNDED / ESTABLISHED	2007
WEBSITE	https://ambature.com/
MANAGEMENT TEAM / FOUNDER(S)	**Ron Kelly – CEO, President & Board Member:** Serves as Chief Executive Officer, President and a board member. Ron Kelly leads the commercialization of Ambature's superconducting materials and quantum-device technologies. **Mitch Robson, Ph.D. – Chief Technology Officer:** Serves as Chief Technology Officer. Mitch received a PhD in Engineering Physics from McMaster University, where his research focused on semiconducting nanowires. **Peter Kastelic – Chief Financial Officer:** Serves as Chief Financial Officer, with prior experience in both private and public corporations. Peter has facilitated revenue growth from $30MM to $400MM over four years and the exit from two companies for a value of $350MM.
SIZE OF COMPANY	~39 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Ambature has evolved into a revenue-generating developer of high-temperature superconducting materials and quantum devices for energy transmission, sensing and scalable quantum infrastructure. The company has expanded its intellectual-property portfolio to 17 patent families and continued advancing its a-axis Josephson-junction and extremely low-resistance material technologies through 2026.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Ambature has raised a total of $44.3MM. In its latest disclosed Series B round of funding, it raised approximately $41MM in October 2015. Investors include Zoic Capital and MFP Ventures.
LATEST ROUND VALUATION	$211MM as of October 2015 Series B.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale superconducting-materials and quantum-device businesses, P/S or EV/Sales may be appropriate once revenue is generated. Comparable-company analysis may reference listed superconducting and scientific-equipment providers such as AMSC, OXIG and BRKR, together with private quantum and advanced-materials companies such as PsiQuantum, Ayar Labs and Bluefors.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI DC Power Management: Emerald AI

COMPANY NAME	Emerald AI
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	Emerald AI targets the emerging market for power-flexible AI data centers, operating at the intersection of AI infrastructure, energy management, and grid-services software. The company estimates that nearly 50 GW of new U.S. AI data-center capacity is expected to seek grid interconnection over the next three years, while industry forecasts suggest only about half may connect on schedule due to power constraints. Emerald's thesis is that flexible, grid-responsive AI workloads can help alleviate this bottleneck: the company estimates that power-flexible AI infrastructure could unlock up to 100 GW of additional capacity on the existing U.S. power system, or ~20% of annual U.S. electricity consumption. This creates a substantial opportunity for software that enables AI data centers to connect faster, optimize energy consumption, and participate as grid resources.
GEOGRAPHIC PRESENCE	Washington, DC
NOTABLE COMPETITORS	**LISTED:**Schneider Electric, Siemens, ETN, GEV, Enel and HON **PRIVATE:**Voltus, CPower, GridBeyond, Lancium, Phaidra
KEY SUPPLIERS / CUSTOMERS	NVIDIA is Emerald AI's most important tech partner and an investor through NVentures. Emerald's Conductor platform integrates with NVIDIA's DSX Flex architecture to enable AI factories to respond dynamically to grid conditions while maintaining AI workload performance. Major ecosystem partners include Digital Realty, EPRI, PJM Interconnection, Portland General Electric, National Grid, and Silicon Valley Power. Emerald has also announced collaborations with leading power-sector companies, including AES, Constellation, Invenergy, NextEra Energy, Nscale Energy & Power, and Vistra, to advance power-flexible AI infrastructure.
COMPETITIVE POSITIONING	Emerald AI develops software-defined solutions to address AI data centers' power constraints by transforming data centers from inflexible electricity consumers into dynamically controllable grid assets. Key differentiators include: • Conductor platform: Emerald AI's Conductor platform integrates directly with AI compute infrastructure and grid systems to orchestrate temporal workload shifting, geographic workload migration, and onsite energy resources while maintaining customer-defined AI performance requirements. • GPU-intensive AI infrastructure: Unlike traditional demand-response platforms primarily designed around industrial or commercial loads, Emerald AI is purpose-built for GPU-intensive AI infrastructure and can make workload-level decisions based on latency, model performance, and compute requirements. • NVIDIA DSX Flex integration: Integration with NVIDIA DSX Flex - combined with demonstrations across multiple grids and commercial data centers - provides early technical validation.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://www.emeraldai.co
MANAGEMENT TEAM / FOUNDER(S)	**Dr. Varun Sivaram - Founder and CEO:** Prior to founding Emerald AI, he served as Chief Strategy & Innovation Officer at Ørsted, and previously as CTO of ReNew Power. He later served as Managing Director for Clean Energy and Innovation at the U.S. State Department, where he helped launch the First Movers Coalition. Sivaram holds a Ph.D. in condensed-matter physics from the University of Oxford, earned undergraduate degrees from Stanford University, and was both a Rhodes Scholar and Truman Scholar. **Prof. Ayse Coskun - Chief Scientist:**Professor of Electrical & Computer Engineering at Boston University, Director of BU's Center for Information and Systems Engineering and a leading researcher in grid-aware/flexible high-performance computing. Her research into dynamically controlling data-center power demand forms an important technical foundation for Emerald. Previously worked at Sun Microsystems and holds a PhD from UC San Diego.
SIZE OF COMPANY	~55 employees according to LinkedIn (accessed 10/1/26)
GROWTH TRAJECTORY	The company has built a broad ecosystem spanning NVIDIA, Digital Realty, National Grid, PJM, and leading power producers including AES, Constellation, Invenergy, NextEra Energy, Nscale Energy & Power, and Vistra. Emerald is now moving from pilot projects to commercial-scale implementations, including the 96 MW Aurora AI Factory in Virginia.

ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Emerald AI has raised $67.5M in publicly announced funding, including a $24.5M Seed round, an $18M Seed extension, and a $25M Strategic Expansion Round announced in March 2026. Investors: Radical Ventures, NVentures, Lowercarbon Capital, Energy Impact Partners, Salesforce Ventures, National Grid Partners, GE Vernova, Siemens, In-Q-Tel, CRV, Amplo and Neotribe, alongside angels Jeff Dean, Fei-Fei Li, John Doerr, John Kerry and Malcolm Turnbull.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI DC Power Management: Ferric

COMPANY NAME	Ferric, Inc.
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Power Delivery
SIZE OF TOTAL ADDRESSABLE MARKET	Ferric designs, manufactures and sells chip-scale DC-DC power converters, marketed as Integrated Voltage Regulators (IVRs), for high-performance digital processors. The company's IVRs integrate power transistors, thin-film magnetic core inductors, capacitors, feedback control, and telemetry into a single device that can be placed adjacent to — or inside the package of — the processor it powers. The addressable opportunity is driven by the shift to vertical power delivery for AI accelerators: modern processors can demand in excess of 5kW per chip, more than ten times CPU and GPU requirements of a few years ago, forcing a rearchitecture of the power delivery network (PDN).
GEOGRAPHIC PRESENCE	New York, NY
NOTABLE COMPETITORS	**LISTED:** MPWR, TXN, (IFX, ADI (via Empower Semi), Renesas, VICR **PRIVATE:** Enphase-adjacent IVR entrants, and internal power teams at hyperscalers and processor vendors
KEY SUPPLIERS / CUSTOMERS	Suppliers include TSMC (foundry). Ferric's thin-film magnetic process is developed by Ferric and licensed to TSMC, which has announced availability of fully integrated voltage regulator functions on a CMOS chip using that process. Inductors in the flagship Fe1766 are fabricated using CMOS BEOL layers on TSMC silicon. In June 2025 Ferric and Marvell (MRVL) announced a collaboration under Marvell's package-integrated voltage regulator (PIVR) initiative to provide pre-validated IVR power solutions to Marvell's custom XPU customers. The company has publicly referenced deployment in Apple products since 2021. Non-dilutive support has come from DARPA SBIR/STTR, NSF SBIR Phase II and the DOE HPC4Mfg program.
COMPETITIVE POSITIONING	Ferric's core differentiator is its patented integration of thin-film magnetic power inductors with CMOS, which eliminates the bulky external passives that constrain conventional voltage regulators. Key attributes include: • On-chip integrated magnetics: Thin-film magnetic core inductors are integrated on-chip or on-interposer using CMOS back-end-of-line layers, allowing power FETs and control circuitry to sit directly beneath the inductor stack and producing solutions as thin as 0.2mm. • Power density leadership: The flagship Fe1766 delivers 160A from 35.5mm² of silicon (8 × 4.4mm) at greater than 4.5 A/mm² current density — described by the company as 3x power per area and more than 20x per volume versus competitive solutions, and more than 10x smaller than conventional approaches. • Regulation bandwidth: A 16-phase interleaved buck converter architecture delivering greater than 10MHz control bandwidth, over 10x faster than conventional multiphase regulation, at 89% efficiency for 1.8V to 0.75V conversion at full thermal design current. • Vertical power delivery and scaling: Supports a scalable vertical PDN architecture capable of kilowatt-class delivery, scaling beyond 10kW with 64 devices, and positioned as a 'power chiplet' for co-packaging with high-performance silicon. • Foundry-embedded moat: The TSMC process licensing relationship embeds Ferric's magnetics in a production foundry flow, a structural barrier for competitors attempting to replicate integrated magnetics at scale.
YEAR FOUNDED / ESTABLISHED	2011
WEBSITE	https://www.ferric.com/
MANAGEMENT TEAM / FOUNDER(S)	**Dr. Noah Sturcken – Co-founder, President and CEO:** Prior to Ferric, Dr. Sturcken worked at AMD's R&D lab, where he developed integrated voltage regulator technology. He holds more than 40 issued patents and 15 publications on IVRs and leads Ferric's business development, marketing and new technology development. Ph.D., Columbia University. **Prof. Kenneth Shepard – Co-founder:** Professor at Columbia University; co-founded Ferric around the commercialization of integrated thin-film magnetics research.

SIZE OF COMPANY	~50 employees on LinkedIn (accessed 10/1/26)
GROWTH TRAJECTORY	Ferric has followed a long research-to-production arc, funded early by non-dilutive US government awards (NSF SBIR Phase II, DARPA SBIR/STTR, DOE HPC4Mfg) before transitioning to venture capital. The company's TSMC partnership brought integrated magnetics into production foundry manufacturing. Product progression moved from the Fe1736 package voltage regulator family, exhibited at ISSCC 2025, to the flagship Fe1766 IVR launched August 25, 2025. Commercial validation accelerated through 2025: the June 2025 Marvell PIVR collaboration positioned Ferric within a major custom-silicon vendor's power ecosystem, and the company has cited deployment in Apple products since 2021. Ferric is positioning IVRs as power chiplets co-packaged with AI accelerators as rack power densities move toward multi-kilowatt processors.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Ferric raised $32.1MM in a Series B closing reported in a December 2024 SEC filing; the total offering was for $50.9MM with eight investors participating in that close. Cumulative reported equity funding is approximately $41MM across nine rounds, inclusive of early grant funding beginning May 2012. Investor disclosure is limited; the company has not publicly named its institutional syndicate. Non-dilutive support has come from the NSF (SBIR Phase II), DARPA and the US Department of Energy.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale analog and power semiconductor companies, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales and as design wins convert to production volume. Peer listed analog and power management companies (MPWR, TXN, ADI, IFX, VICR) are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI DC Power Management: Power Lattice

COMPANY NAME	Power Lattice
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Power Delivery Chiplets
SIZE OF TOTAL ADDRESSABLE MARKET	PowerLattice operates in AI data-center power delivery, semiconductor voltage regulation, high-performance computing and processor-package power management.
GEOGRAPHIC PRESENCE	Camas, Washington
NOTABLE COMPETITORS	**LISTED:** SLAB, ADI, TXN, IFX **PRIVATE:** Empower Semiconductor, Ferric, Endura Technologies
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	PowerLattice develops power-delivery chiplets for AI, data-center and high-performance computing processors. Key differentiators include: • In-package voltage regulation: Power regulation is positioned inside the processor package rather than on the printed circuit board. • Proximity to the compute die: The voltage regulator is placed a few hundred micrometers from the compute die instead of several inches away. • Reduced distribution losses: Shorter delivery distances are designed to reduce electrical power losses. • Higher power density: The architecture is intended to support processors where power density and cooling constrain performance. • Chiplet architecture: The technology can be integrated into advanced processor packages used by AI-chip developers. • Location-mapped power management: PowerLattice's patents cover power-management techniques based on the position and configuration of power-delivery chiplets.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://powerlatticeinc.com/
MANAGEMENT TEAM / FOUNDER(S)	**Peng Zou – Co-Founder, CEO, President & Board Member:** Co-founded PowerLattice and has served as a board member since January 2023. Prior to founding PowerLattice, he held technical leadership roles at QCOM/NUVIA, Huawei, and Intel. **Gang Ren – Co-Founder, Head of Engineering & Board Member:** Co-founded PowerLattice and serves as Head of Engineering and a board member, with the board role recorded from May 2025. Prior to founding PowerLattice, he held various leadership roles at QCOM/NUVIA, Maxim (ADI), and Huawei. **Sujith Dermal – Co-Founder & Board Member:** Co-founded PowerLattice and has served as a board member since January 2023. Prior to founding PowerLattice, he held various technical roles at QCOM/Nuvia, Dialog, and Freescale.
SIZE OF COMPANY	~20 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	PowerLattice has progressed from its 2023 formation into a revenue-generating developer of in-package power-delivery chiplets for AI and high-performance computing systems. The company developed an architecture that places voltage regulation within the processor package and secured active U.S. patents covering location-mapped chiplet power management in 2025 and 2026. Its Series A funding is being used to address power-efficiency constraints that limit AI processor performance and cooling.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	PowerLattice has raised a total of $30.8MM. In its latest Series A round of funding, it raised approximately $30.8MM in November 2025. Investors include Celesta Capital, Playground Global, and G Vision Capital.

LATEST ROUND VALUATION	$94MM as of November 2025 Series A.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale power-semiconductor and semiconductor-chiplet businesses, P/S or EV/Sales may be appropriate once revenue is generated. Comparable-company analysis may reference listed power-semiconductor providers such as SLAB, ADI, TXN and Infineon Technologies AG, together with private power-management companies such as Empower Semiconductor, Ferric and Endura Technologies.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

AI DC Power Management: Reed Semiconductor

COMPANY NAME	Reed Semiconductor
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Analog and Power Management
SIZE OF TOTAL ADDRESSABLE MARKET	Reed Semiconductor competes in the power analog market. Power analog represents roughly one-third of the approximately $100B analog semiconductor market, implying a market size of about $30B, based on UBS market segmentation and industry estimates.
GEOGRAPHIC PRESENCE	Warwick, Rhode Island
NOTABLE COMPETITORS	**LISTED:**ADI, IFX, ON, STM, TXN, MPWR, Renesas **PRIVATE:**uPI Semiconductor, pSemi, Nengda
KEY SUPPLIERS / CUSTOMERS	As a fabless semiconductor company, Reed relies on TSMC as its publicly disclosed foundry partner, with distribution through WT Microelectronics and DigiKey. Its power-management portfolio targets AI/data-center, server, networking, communications, and automotive-HPC applications, supporting NVIDIA, AMD, Intel, and ARM processor architectures. Public court filings identify HPE Aruba (through contract manufacturer Accton) as a downstream OEM customer. Other hyperscaler and AI-accelerator customers have not been publicly disclosed.
COMPETITIVE POSITIONING	Reed Semiconductor designs and manufactures power management ICs and power modules for data center, automotive, communications, and industrial markets. Its portfolio spans multiphase controllers, smart power stages, power modules, IBC, eFuses, Power Mux, and POL DC-DC controllers. Key differentiators include: • Power management ICs and power modules: Reed Semiconductor specializes in multiphase controllers, smart power stages, power modules, IBC, eFuses, Power Mux, and POL DC-DC controllers. • Power delivery applications: The company's portfolio supports applications including 48V/12V Server Power Delivery, Network Switch Power Delivery, 12V PCIe Card power solutions, and Accelerator Card power solutions.
YEAR FOUNDED / ESTABLISHED	2019
WEBSITE	https://www.reedsemi.com/
MANAGEMENT TEAM / FOUNDER(S)	**Wenkai Wu - President and CEO:** Founder and CEO of Reed Semiconductor since 2019. Prior to Reed, he held leadership roles at Texas Instruments and International Rectifier, focusing on power management solutions for servers, data centers, communications, and high-performance computing.
SIZE OF COMPANY	~65 employees on LinkedIn (accessed 8/19/2026)
GROWTH TRAJECTORY	Reed has expanded rapidly since its 2019 founding, evolving from a supplier of multiphase controllers and smart power stages into a broader power-delivery vendor targeting AI, data-center, and HPC applications. Its portfolio now spans eFuses, IBCs, power modules, and POL converters, enabling increasingly comprehensive 48V-to-load power architectures for next-generation AI servers. Product introductions accelerated in 2025-26 with high-current eFuses, 48V IBCs, and vertical power-delivery modules, reflecting rising accelerator power densities and growing power-semiconductor content per system. Reed's June 2026 $100M funding round is expected to support product development and market expansion.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Reed has raised over $140M in total funding since inception, culminating in an upsized $100M financing in June 2026 backed by undisclosed global semiconductor companies. Earlier rounds have not been publicly confirmed by the company. Investors: UMC Capital
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For analog companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer analog companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.

SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Figure 12: Outside Taxonomy Companies & Public Peers

Companies	Public Peers
Ambra	NICE, GOOGL, AMZN, MSFT, CRM, FIVN
Analogix	2379-TW, HIMX, MegaChips, SYNA, SIMO, 6723-JP
Axbio	ILMN, TMO, Roche Holding AG, BIO, QDEL
Celera Semiconductor	CDNS, SNPS
CNEX Labs	MRVL, AVGO, MU, STX, WDC
Eridu	NVDA, AVGO, CSCO, ANET
Frore Systems	VRT
Majestic	CBRS, NVDA
Modal	AMZN, GOOGL, MSFT, ORCL
Movandi	ADI, MCHP, QCOM, SWKS
SiLC Technologies	AEVA, INVZ, OUST
Tensorwave	NVDA, CBRS, CRWV

Source: UBS

Outside taxonomy: Ambra

COMPANY NAME	Ambra
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Compute Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	Ambra Compute develops edge AI infrastructure for deploying modular GPU compute capacity at renewable energy sites. Its addressable markets include AI compute infrastructure, edge computing, modular data centers and AI neoclouds.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** NICE, GOOG, AMZN, MSFT, CRM, FIVN **PRIVATE:**Earth Wind and Power, Cyanis AI, Enerhash, 639 Solar, Built Robotics.
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	Ambra Compute develops modular GPU compute nodes deployed at renewable energy sites, allowing customers to run AI inference workloads using curtailed or otherwise underutilized renewable electricity. Key differentiators include: • Energy-linked compute infrastructure: Compute nodes are deployed directly at renewable energy sites rather than exclusively in centralized data centers. • Use of curtailed electricity: The company seeks to convert unused renewable electricity into productive AI compute capacity. • Modular deployment model: Modular GPU nodes may be deployed across grid-edge locations based on available energy resources. • Inference-focused architecture: The infrastructure is intended specifically to support AI inference workloads. • Sustainability positioning: The business model combines AI infrastructure with renewable energy utilization.
YEAR FOUNDED / ESTABLISHED	2026
WEBSITE	https://www.ambracompute.com
MANAGEMENT TEAM / FOUNDER(S)	**Bianca G. – Co-Founder & CEO:** Prior to co-founding Ambra Compute in 2026, Bianca G. worked as an Associate at Citadel. Earlier in her career, she held roles spanning quantitative research, renewable-energy investing, and finance, including positions at Intersect Power, Bank of America Merrill Lynch, Goldman Sachs, and Fondo Italiano d'Investimento. Her background combines quantitative analysis, energy markets, and infrastructure investment. She holds a Master's degree in Operations Research from Columbia University, studied Economics at Imperial College London, and earned a degree in Finance from Bocconi University. At Ambra Compute, she leads the deployment of modular AI-compute infrastructure powered by renewable-energy resources and located at grid-edge energy sites.
SIZE OF COMPANY	2 employees
GROWTH TRAJECTORY	Ambra Compute was founded in 2026 to deploy modular GPU compute nodes at renewable-energy sites, using curtailed electricity to support AI inference workloads. The company is classified as generating revenue and completed its first venture financing in September 2026. With two employees, Ambra remains at an early commercialization stage.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Ambra Compute raised an undisclosed amount of early-stage venture funding in September 2026. Notable investor: Transpose Platform Management. Investment stage: Early-stage VC.
LATEST ROUND VALUATION	Undisclosed valuation as of latest round of funding.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For compute and networking companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer compute and networking companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.

SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: Analogix

COMPANY NAME	Analogix
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Display & Connectivity Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Analogix manufactures application-specific semiconductors for digital multimedia applications. Its addressable markets include display-interface semiconductors, timing controllers, signal retimers, DisplayPort connectivity, VR/AR display chips and related mixed-signal products.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** 2379-TW, HIMX, MegaChips Corporation, SYNA, SIMO, 6723-JP **PRIVATE:** Imagination Technologies, InnoPhase.
KEY SUPPLIERS / CUSTOMERS	Product collaborations with LG and MediaTek, but does not explicitly identify either company as a customer or supplier.
COMPETITIVE POSITIONING	Analogix manufactures display and connectivity semiconductors used in smartphones, notebooks, VR head-mounted displays, televisions and high-end graphics cards. Its portfolio includes VR/AR display chips, timing controllers, retimer chips and DisplayPort-related products. Key differentiators include: • Display and connectivity specialization: Analogix focuses on display interfaces, signal conditioning, timing control and multimedia connectivity. • Broad end-market exposure: Its products address smartphones, notebooks, televisions, VR devices and graphics systems. • Established intellectual property: 44 patent families and 105 patent applications and grants. • International operating footprint: The company has offices across the United States, mainland China, South Korea and Taiwan. • Profitable operating profile: Analogix is classified as profitable in the latest profile. • Product collaborations: References DisplayPort and related product collaborations with LG and MediaTek
YEAR FOUNDED / ESTABLISHED	2002
WEBSITE	https://www.analogix.com

MANAGEMENT TEAM / FOUNDER(S)	**Patrick Li – CEO / Lead Management:** Prior to becoming Chief Executive Officer of Analogix Semiconductor, Patrick Li served as President of Greater China, having joined the company in 2019 to drive expansion across China and international markets. Before joining Analogix, he served as China Chief Executive of Ingram Micro and previously served as CEO of Orange Business Services, where he led the company's transition from a network-service provider to an ICT infrastructure provider. Earlier in his career, he spent more than 14 years in senior leadership positions at Intel and Seagate Technology. Mr. Li brings more than 30 years of experience in technology, communications, storage, networking, automotive, and enterprise markets. He holds a B.S. in Electrical Engineering from Shanghai Jiao Tong University. At Analogix, he oversees the company's strategy and commercialization efforts, including its proposed STAR Market listing and continued expansion in China. **Ning Zhu – Chief Technology Officer:** Prior to becoming CTO of Analogix Semiconductor, Ning Zhu served as Vice President of Engineering and has been with the company since its founding in 2002. Throughout his tenure, he has led development of several of Analogix's key technologies, including high-speed SerDes, HDMI, and DisplayPort interface products. Earlier in his career, he held engineering and technical leadership positions at Mindspeed Technologies, Conexant Systems, and Synopsys. Dr. Zhu is widely recognized as an expert in high-speed digital interfaces and holds multiple patents in data communications. He earned B.S., M.S., and Ph.D. degrees in Microelectronics from Tsinghua University. At Analogix, he leads development of display controllers, signal-conditioning devices, retimers, and digital multimedia semiconductor products. **Kewei Yang – Co-Founder**: Prior to co-founding Analogix Semiconductor in 2002, Kewei Yang served as Vice President of Engineering at Mindspeed Technologies, where he led development of high-speed transceivers and broadband switching devices. Earlier, he served as Chief Scientist at HotRail and subsequently Conexant following Conexant's acquisition of HotRail. He also held engineering roles at Rendition and Hewlett-Packard's Computer Technology Laboratory. Dr. Yang earned a B.S. in Electrical Engineering from Tsinghua University and M.S. and Ph.D. degrees in Electrical Engineering from Johns Hopkins University. He founded Analogix to commercialize high-speed mixed-signal semiconductor technologies and was instrumental in establishing the company's leadership in DisplayPort, HDMI, timing-controller, and connectivity solutions.
SIZE OF COMPANY	~169 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Analogix has developed a portfolio spanning VR/AR display chips, timing controllers, retimers and DisplayPort connectivity for smartphones, notebooks, televisions, graphics cards and head-mounted displays. Revenue increased from $130.1MM in 2021 to $132.9MM in 2022, while net income rose from $12.4MM to $16.8MM. The company continued developing signal-repeater and display-control technologies through 2026, although its proposed STAR Market IPO was withdrawn in 2024 and Analogix was subsequently acquired on undisclosed terms.
ANY MAJOR M&A / CORPORATE ACTIONS	Analogix was acquired in a transaction covering 100% of the company. The acquisition date, buyer and consideration were not disclosed. Analogix filed for an IPO on the Shanghai Stock Exchange STAR Market. The application was withdrawn on 9 August 2024. In March 2022, UNIC Capital and Beijing Shanhai Kunlun sold stakes in Analogix to Zhongyuan Hechuang Investment, Guangdong Aoao Semiconductor Investment and Orient Huijia. A $500 million buyout was completed on 30 March 2017. A proposed $547 million merger or acquisition announced in March 2019 was subsequently cancelled.
INVESTORS / INVESTMENT STAGE	Analogix has raised a total of $235.7MM. In its latest disclosed PE Growth / Expansion round of funding, it raised $235.7MM in January 2022. Investors include Shenzhen Capital Group, China Merchants Fund, GF Xinde Investment Management, Tsinghua Holdings Capital, China Universal Asset Management, TCL, Hengxu Capital, H-young Investment, Lanhor Capital, and Zhongyuan Hechuang Investment.
LATEST ROUND VALUATION	$1,020MM as of January 2022 PE Growth / Expansion.

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For analog companies that are sub-scale, we think it would be appropriate to value these companies on P/S or EV/S as profitability may vary as the company scales. Peer analog companies are generally valued on P/E or P/FCF as they have a more mature profitability profile.
SECONDARY PRICING DATA	Analogix completed private secondary transactions in January 2021, January 2022 and March 2022. January 2021 secondary transaction: estimated transaction amount of $30.61 million, per Pitchbook. January 2022 secondary transaction: estimated transaction amount of $99.70 million. March 2022 secondary transaction: transaction consideration and per-share price were not disclosed. The disclosed transaction amounts do not establish a current secondary share price.

Source: Pitchbook and Company's reports and website, LinkedIn

Outside Taxonomy: Axbio

COMPANY NAME	Axbio
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Bio-Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Axbio develops semiconductor nanotechnology biochips, detection reagents and a microfluidic Bio-CMOS platform for molecular diagnostics and high-throughput sequencing. Its addressable markets include molecular in-vitro diagnostics, genomic sequencing, early cancer detection, rapid pathogen detection and diagnostic instruments.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:** ILMN, TMO, Roche Holding AG, BIO, QDEL **PRIVATE:** Freenome, Delfi Diagnostics, SniperDX.
KEY SUPPLIERS / CUSTOMERS	Tower Semiconductor announced medical-chip development with Axbio in 2019.
COMPETITIVE POSITIONING	Axbio develops semiconductor nanotechnology biochips and a microfluidic Bio-CMOS platform capable of supporting high-throughput sequencing and molecular diagnostics. Key differentiators include: • Semiconductor-enabled diagnostics: The company combines semiconductor technology with molecular diagnostic applications. • Integrated Bio-CMOS platform: The platform brings together microfluidics, sequencing and molecular detection. • Broad instrument range: The technology can support low-cost handheld devices as well as automated instruments for hospitals and independent laboratories. • Multiple target applications: Intended use cases include early cancer screening, pathogen detection and drug development. • Biochip and reagent integration: Axbio develops both detection reagents and semiconductor-based biochips. • US-China footprint: The company maintains a presence in both Santa Clara and Shenzhen.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://www.axbio.com/
MANAGEMENT TEAM / FOUNDER(S)	**Hui Tian, Ph.D. – Founder, Chairman & CEO:** Prior to founding Axbio, publicly available sources indicate that Hui Tian developed expertise in semiconductor-enabled biotechnology, genomics, and molecular-diagnostics technologies. Since founding the company, he has overseen development of Axbio's semiconductor biochip platforms, DNA sequencing technologies, and molecular-diagnostics systems, positioning the company at the intersection of semiconductor engineering and life sciences. Based in Shenzhen, Dr. Tian serves as Chairman and Chief Executive Officer and remains the company's principal executive and primary public representative.
SIZE OF COMPANY	~46 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Axbio has progressed from product development into a revenue-generating provider of semiconductor biochips, detection reagents and a microfluidic Bio-CMOS platform for sequencing and molecular diagnostics. Following its Series B, the company worked to optimize its sequencing technology, construct sequencer production capacity and accelerate commercialization. Revenue increased from $0.5MM in 2024 to $1.2MM in 2025, although the company remained loss-making and its proposed Hong Kong IPO was postponed.
ANY MAJOR M&A / CORPORATE ACTIONS	Axbio filed to list on the Hong Kong Stock Exchange on 30 April 2025 under an undisclosed ticker. The offering was subsequently postponed. The company raised $70 million in Series B financing in June 2022 to optimize its sequencing technology, construct sequencer production capacity and accelerate commercialization.

INVESTORS / INVESTMENT STAGE	Axbio has raised a total of $154.6MM. In its latest Series B round of funding, it raised $70MM in June 2022. Investors include Wuxi AstraZeneca CICC Venture Capital Partnership, Yunfeng Capital, CBC Group, CMG SDIC Fund, I-Bridge Capital, 5Y Capital, Kaitai Capital, MSA Capital, Suzhou Oriza Bioventures, Green Pine Capital Partners, and ZhenFund
LATEST ROUND VALUATION	$347.5MM as of June 2022 Series B
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	Early-commercial diagnostics and sequencing companies may be valued using recent financing rounds, comparable-company analysis and probability-adjusted DCF. Potential comparables include sequencing and diagnostic companies such as Illumina, Freenome, Delfi Diagnostics and SniperDX. Valuation adjustments would be required for regulatory progress, product maturity, commercial scale, installed base and research-and-development expenditure.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: Celera Semiconductor

COMPANY NAME	Celera Semiconductor
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Semiconductor Design Automation
SIZE OF TOTAL ADDRESSABLE MARKET	Celera Semiconductor develops an AI-based platform for automating the design of custom analog and mixed-signal integrated circuits. Its addressable markets include electronic design automation, analog and mixed-signal IC design, custom silicon and semiconductor IP development.
GEOGRAPHIC PRESENCE	Santa Clara, California
NOTABLE COMPETITORS	**LISTED:**CDNS, SNPS **PRIVATE:**SiFive, Imagination Technologies, Empower Semiconductor, Ceremorphic, Axiado
KEY SUPPLIERS / CUSTOMERS	N/A
COMPETITIVE POSITIONING	Celera develops an AI-based software platform designed to automate the development of custom analog and mixed-signal integrated circuits, enabling semiconductor customers to design chips more quickly and at lower cost. Key differentiators include: • AI-based circuit development: The platform automates portions of analog and mixed-signal IC design. • Nestos digital twin library: Celera's patented Nestos platform uses a digital twin-based analog design library to support rapid, high-precision IC development. • Analog and mixed-signal specialization: Celera is more narrowly focused on analog design than broad electronic design automation providers. • Custom-silicon focus: The platform is intended to lower cost and development-time barriers associated with custom chips. • Patent portfolio: Four patent families and 36 patent applications and grants. • Demonstrated design output: Celera announced that it had sampled an analog IC designed entirely by software.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://celerasemi.com
MANAGEMENT TEAM / FOUNDER(S)	**Patrick Brockett – Co-Founder & CEO:** Co-founded Celera Semiconductor and serves as Chief Executive Officer. Patrick Brockett was identified as lead management for the company's June 2025 Series A, August 2026 Series B and May 2026 acquisition of SiliconGate. Patrick Brockett has led the company during the commercialization of its AI-based analog design platform and its expansion through acquisition. **Helen Song – CTO & Vice President of Engineering:** Serves as Chief Technology Officer and Vice President of Engineering. The combined technology and engineering role is directly relevant to Celera's development of AI-based automation tools for custom analog and mixed-signal integrated circuits. **Jim LoCascio – Co-Founder & Architect:** Co-founded Celera Semiconductor and serves as Architect. The architect role is directly connected to the company's semiconductor-design platform and automated circuit-generation technology.
SIZE OF COMPANY	~20 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Celera Semiconductor has developed an AI-based platform for automating custom analog and mixed-signal IC design. Following its 2025 Series A, the company began bringing to market analog ICs developed using Nestos, its patented digital twin-based analog design library. In 2026, Celera acquired SiliconGate to expand its technology footprint and raised a Series B to accelerate its design-automation roadmap.
ANY MAJOR M&A / CORPORATE ACTIONS	Celera acquired SiliconGate on 29 May 2026. SiliconGate is based in Porto, Portugal and was founded in 2008. The transaction value was not disclosed.

INVESTORS / INVESTMENT STAGE	Celera Semiconductor has raised a total of $64MM. In its latest Series B round of funding, it raised $30MM in August 2026. Investors include Maverick Silicon, Maverick Capital, Augment Ventures, and Valley Capital Partners.
LATEST ROUND VALUATION	$125.5MM as of August 2026 Series B.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale semiconductor design-automation companies, relevant approaches may include latest-round valuation, EV/Revenue and comparable-company analysis. Potential comparables include Cadence Design Systems, Synopsys and selected custom-silicon companies. Material adjustments would be required for Celera's smaller scale, private status, earlier commercial stage and specialized analog focus.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside Taxonomy: CNEX Labs

COMPANY NAME	CNEX Labs
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Storage Controllers
SIZE OF TOTAL ADDRESSABLE MARKET	CNEX Labs develops semiconductor storage controllers for cloud and enterprise data centers. Its addressable markets include data-center storage controllers, solid-state storage, software-defined storage, storage fabrics and cloud data infrastructure.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** MRVL, AVGO, MU, STX, WDC **PRIVATE:** Alif Semiconductor, SiFive, Pliops.
KEY SUPPLIERS / CUSTOMERS	Industry participants that have invested in CNEX include Dell Technologies Capital, MaxLinear, AMD Ventures, Micron Technology, Cisco Investments, Samsung Venture Investment, Seagate Venture Fund and UMC Capital.
COMPETITIVE POSITIONING	CNEX develops semiconductor storage controllers for cloud and enterprise data centers. Its products are intended to improve the performance and scalability of big-data computing, analytics and software-defined storage. Key differentiators include: • Data-center storage focus: CNEX develops semiconductor and software solutions for data-center storage, with its controller architecture designed around high-performance NVMe-based flash storage. • Low and predictable latency: CNEX's SSD controllers are designed to deliver high performance with low and predictable latency, including hardware acceleration of key flash-translation functionality. • Software-defined storage architecture: The controllers support Open-Channel and LightNVM operating modes that move greater control over data placement and I/O scheduling to the host, enabling host-based FTLs and workload-specific storage optimization. • Storage-fabric scalability: CNEX integrates native Ethernet I/O and NVMoE capabilities into its controller architecture to support low-latency scaling across Ethernet-based storage fabrics. • Storage-controller specialization: The company focuses on SSD controller ASICs and supporting software, positioning its technology as an enabling layer for storage infrastructure rather than as a complete storage-system offering. • Strategic semiconductor backing: CNEX has been funded through venture capital and strategic investments from large companies in the storage and networking industries.
YEAR FOUNDED / ESTABLISHED	2013
WEBSITE	https://www.cnexlabs.com/
MANAGEMENT TEAM / FOUNDER(S)	**Alan Armstrong, Ph.D. – Co-Founder, CEO & Board Member:** Co-founded CNEX Labs and has served as a board member since January 2013. Alan Armstrong is the company's Chief Executive Officer and primary contact and has led CNEX through its development from a seed-stage business to a Series G venture-backed semiconductor company. The role is directly relevant to CNEX's development and commercialization of storage-controller technology for cloud and enterprise data centers. **Joseph DeFranco – Co-Founder & Vice President of Sales:** Co-founded CNEX Labs and serves as Vice President of Sales. The position is directly related to the commercial development of CNEX's semiconductor storage-controller products.
SIZE OF COMPANY	~50 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	CNEX Labs has evolved from a startup into a revenue-generating supplier of semiconductor storage controllers for cloud and enterprise data centers. Its products support low and predictable latency, host-based flash translation layers and Ethernet-connected storage fabrics for big-data computing and analytics.
ANY MAJOR M&A / CORPORATE ACTIONS	In April 2022, Mellanox Capital and Xilinx Ventures sold their stakes in CNEX to an undisclosed buyer. The transaction consideration and per-share pricing were not disclosed.

INVESTORS / INVESTMENT STAGE	CNEX Labs has raised a total of $168.8MM. In its latest Series G round of funding, it raised approximately $32.7MM in May 2022. Investors include MaxLinear, Dell Technologies Capital, Gravity Fund, Claridge Venture Partners, AMD Ventures, Micron Technology, M12, Celesta Capital, Cisco Investments, Samsung Venture Investment, Seagate Venture Fund, Sierra Ventures, UMC Capital, and Walden International.
LATEST ROUND VALUATION	$132.7MM as of May 2022 Series G.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	Semiconductor storage-controller companies may be valued using EV/Revenue, EV/EBITDA or P/E, depending on scale and profitability. Potential comparables include Marvell, Broadcom, Micron, Seagate, Western Digital and Pliops. Adjustments would be required for CNEX's smaller scale, private ownership, product concentration and historical valuation trajectory.
SECONDARY PRICING DATA	A private secondary transaction was completed on 4 April 2022, when Mellanox Capital and Xilinx Ventures sold their stakes to an undisclosed buyer.

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: Eridu

COMPANY NAME	Eridu
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Networking Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Eridu develops networking hardware and silicon for scale-up and scale-out AI infrastructure. Its addressable market includes AI networking silicon, data-center switches, high-speed interconnects and the associated systems and software used to connect large AI compute clusters.
GEOGRAPHIC PRESENCE	Saratoga, California
NOTABLE COMPETITORS	**LISTED:**NVDA, AVGO, CSCO, ANET **PRIVATE:**Juniper Networks, Nexthop AI, Enfabrica, Etched
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Eridu develops networking hardware and silicon designed to remove infrastructure bottlenecks in large-scale AI workloads. The company uses a clean-sheet approach across silicon, systems and software and is developing network switches and interconnect products for both scale-up and scale-out architectures. Key differentiators include: • Clean-sheet architecture: Eridu is developing its silicon, systems and software together rather than optimizing only one element of the networking stack. • Scale-up and scale-out coverage: The company's products are intended to support connectivity within tightly integrated AI computing systems as well as networking across larger clusters. • Networking-radix optimization: Eridu's architecture is designed to improve network radix and efficiency as AI clusters increase in size. • AI-specific product design: The platform is designed specifically for AI infrastructure rather than traditional enterprise networking workloads. • Full-stack positioning: Eridu combines networking silicon with switches, interconnects and supporting software. • Patent development: Seven patent families and eight patent applications and grants, with disclosed technologies covering switches, transceiver error correction, packet-buffer management, three-dimensional processing architectures and semiconductor packaging.
YEAR FOUNDED / ESTABLISHED	2024
WEBSITE	https://eridu.ai/

MANAGEMENT TEAM / FOUNDER(S)	**Drew Perkins – Co-Founder, CEO & Board Member:** Prior to co-founding Eridu in 2024, Drew Perkins built a distinguished career as a networking entrepreneur and technology pioneer. He was a principal architect of the Point-to-Point Protocol (PPP), an important early Internet protocol, and later co-founded several networking companies including Lightera Networks, which was acquired by Ciena, Infinera, which became a public optical-networking company, and Gainspeed, which was acquired by Nokia. Most recently, he co-founded Mojo Vision. He holds a B.S. degree in Electrical Engineering, Computer Engineering, and Mathematics from Carnegie Mellon University. At Eridu, he leads development of AI networking infrastructure designed to address communication bottlenecks within large-scale AI clusters. **Omar Hassen – Co-Founder, Chief Product Officer & Board Member:** Prior to co-founding Eridu, Omar Hassen accumulated more than three decades of experience in networking silicon, communications semiconductors, and infrastructure products. His career has included executive leadership roles at Ventana Micro Systems, where he served as SVP of Business Development, AppliedMicro, where he was VP and General Manager of the Connectivity Business Unit, International Rectifier, Marvell Technology, and Broadcom. Earlier in his career, he co-founded Entridia, a semiconductor startup developing WAN network processors. He holds a bachelor's degree in Computer Engineering from California State University, Long Beach and an MBA from the University of Southern California. At Eridu, he leads product strategy and development spanning AI networking silicon, switches, interconnects, and software platforms. **Mike Capuano – Co-Founder, Chief Business Development Officer & Chief Marketing Officer:** Prior to co-founding Eridu, Mike Capuano built a career spanning more than 25 years in networking, cloud infrastructure, and communications technology. He held leadership roles at Motorola, Juniper Networks, Cisco Systems, Infinera, and Pluribus Networks. At Cisco, he led global marketing for a multi-billion-dollar routing, switching, and optical networking portfolio. He later served as Vice President of Global Marketing at Infinera and Chief Marketing Officer at Pluribus Networks before its acquisition by Arista Networks. He studied at Colgate University. At Eridu, he leads business development, partnerships, commercialization, and go-to-market strategy for the company's AI networking platform.
SIZE OF COMPANY	~119 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Eridu has progressed from stealth development into a 100-employee AI networking company developing an integrated platform across silicon, systems and software. Its products include network switches and interconnects for scale-up and scale-out AI architectures, supported by patents covering transceivers, packet-buffer management, 3D processing and advanced packaging. In 2026, the company expanded its commercial organization while continuing to address networking bottlenecks in large AI clusters.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Eridu raised $18.25MM in June 2024 at an $86MM post-money valuation, followed by $11.75MM in March 2025 at an estimated $180MM post-money valuation. In March 2026, the company raised $200MM through a combination of $193MM of Series A equity and $7MM of debt at a $442MM post-money valuation. The company's disclosed post-money valuation increased from $86MM in June 2024 to $180MM in March 2025 and $442MM in March 2026. PitchBook reports a 1.34x valuation step-up between the March 2025 seed round and the March 2026 Series A.
LATEST ROUND VALUATION	$442MM as of March 2026 Series A.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale companies without profitability, P/S or EV/Sales may be more relevant than earnings-based multiples once revenue becomes available. Comparable-company analysis may reference listed AI networking and semiconductor companies such as NVDA, AVGO, CSCO and ANET, together with private companies such as Enfabrica. Material adjustments would be required for Eridu's earlier commercial stage, smaller scale, product-development risk and lack of disclosed financial metrics.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: Frore Systems

COMPANY NAME	Frore Systems
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Thermal Management Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Frore Systems develops solid-state thermal-management chips for compact electronic devices, edge AI systems and dense AI servers. Its addressable markets include semiconductor thermal management, data-center cooling, edge-device cooling and cooling systems for consumer electronics.
GEOGRAPHIC PRESENCE	San Jose, California
NOTABLE COMPETITORS	**LISTED:** VRT **PRIVATE:** CoolIT Systems, Asetek, Thermal Grizzly Corporation, NanoLambda, Jetcool, ZutaCore, Alif Semiconductor
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Frore Systems develops thermal-management chips intended to remove heat from compact consumer-electronic and AI devices and maintain sustained processor performance. Key differentiators include: • Solid-state active cooling: Frore's cooling chips are designed to provide active thermal management without conventional rotary fans. • Thin and silent form factor: The technology is designed for compact devices where space and acoustic performance are important. • High back pressure: Frore's chips are designed to move air through restrictive device enclosures and dense systems. • Dust-resistant design: The solid-state architecture supports compact and dustproof device configurations. • Data-center and edge coverage: Frore targets both consumer and edge devices and denser AI-server environments. • Multiple cooling architectures: Frore's technology supports liquid-assisted and PAK-based systems for dense AI servers. • Thermal-throttling mitigation: The products are intended to remove heat more effectively so processors can maintain higher sustained performance. • Patent portfolio: Five patent families and 15 patent applications and grants, including patents covering MEMS-based active cooling systems, active heat sinks and engineered actuators.
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://froresystems.com/

MANAGEMENT TEAM / FOUNDER(S)	**Seshagiri Madhavapeddy, Ph.D. – Co-Founder & CEO:** Prior to co-founding Frore Systems in 2018, Seshu Madhavapeddy was a repeat semiconductor and communications entrepreneur with multiple successful exits. Earlier in his career, he co-founded Spatial Wireless, which was acquired by Alcatel-Lucent, and later led Sipera Systems, which was acquired by Avaya. He also held senior leadership positions at Qualcomm, Samsung Mobile, and Texas Instruments. At Qualcomm, he served as VP/GM of the company's in-display fingerprint sensor business and previously managed major IoT and mobile-computing product initiatives. Dr. Madhavapeddy holds a B.Tech. from the Indian Institute of Technology Kharagpur and a Ph.D. in Computer Science from the University of Texas at Dallas. At Frore, he has led commercialization of the AirJet® and LiquidJet® thermal-management platforms and guided the company into edge AI and data-center cooling markets. **Suryaprakash Ganti, Ph.D. – Co-Founder & CTO:** Prior to co-founding Frore Systems, Surya Ganti built a career focused on MEMS technologies, sensors, displays, manufacturing, and advanced materials. Before founding Frore, he served as a technology executive at Qualcomm, where he led development of the Mirasol MEMS display program and ultrasonic fingerprint-sensing technologies from proof-of-concept through mass production. Earlier, he spent approximately seven years as a Research Scientist at General Electric, where he worked across aircraft engines, power generation, healthcare, and consumer technologies. Dr. Ganti holds a Bachelor of Technology degree from the Indian Institute of Technology Madras and a Ph.D. in Mechanical Engineering from the Massachusetts Institute of Technology. He is the principal architect behind Frore's AirJet® solid-state cooling technology and leads development of the company's MEMS-based cooling chips, active heat sinks, and thermal-management platforms.
SIZE OF COMPANY	~96 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	Frore Systems has evolved into a revenue-generating supplier of solid-state cooling chips for consumer electronics, edge AI devices and dense AI servers. The company expanded its AirJet platform from compact-device cooling into liquid-assisted and PAK-based data-center applications. In 2026, Frore collaborated with Lenovo on an ultra-slim notebook concept powered by AirJet technology and began scaling its thermal platforms globally across edge and data-center markets.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Frore Systems has raised a total of $373.9MM. In its latest Series D round of funding, it raised $143MM in March 2026. Investors include MVP Ventures, Global Alternative Investment Management, Qualcomm Ventures, Top Tier Capital Partners, Bow River Capital, Clear Ventures, Alumni Ventures, Mayfield Fund, Addition, Fidelity Management & Research Company, Rigby Technology Investments, StepStone Group, Aramco Ventures, Prosperity7 Ventures, and Sustana Cooling Partners.
LATEST ROUND VALUATION	$1,643MM as of March 2026 Series D.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For a revenue-generating but sub-scale semiconductor business, EV/Revenue or P/S may be more applicable than P/E or P/FCF if profitability is not mature. Once sufficient financial information is available, valuation could also consider product shipments, design wins, content per device, gross margin and customer concentration. Comparable-company analysis may include thermal-management, data-center cooling and semiconductor component businesses. Relevant reference points include VRT and private or acquired cooling companies such as CoolIT Systems, Asetek, Jetcool and ZutaCore. Adjustments would be required for differences in product architecture, revenue scale, profitability and exposure to consumer, edge and data-center applications.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside Taxonomy: Majestic

COMPANY NAME	Majestic
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / Memory-Dense AI Systems
SIZE OF TOTAL ADDRESSABLE MARKET	Majestic Labs develops AI infrastructure designed to address memory-capacity and memory-bandwidth constraints in large-scale AI workloads. Its addressable markets include AI servers, accelerated computing, data-center infrastructure and memory-intensive AI systems.
GEOGRAPHIC PRESENCE	San Francisco, California
NOTABLE COMPETITORS	**LISTED:** CBRS, NVDA **PRIVATE:** Graphcore, Lambda, SambaNova Systems, Physical Intelligence, Positron AI, UnifabriX
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Majestic Labs develops an AI infrastructure platform designed to address memory limitations in large-scale AI workloads. The company is developing servers that combine advanced compute with substantially expanded memory capacity. Key differentiators include: • Memory-dense architecture: The platform is designed to increase memory capacity available to AI workloads. • Memory-bottleneck focus: Majestic targets the "memory wall" that can constrain the processing of increasingly large models and datasets. • Integrated server platform: The company is developing complete servers rather than offering only an individual semiconductor component. • Full software stack: Series A proceeds are intended to support development of the platform's complete software stack. • Large-scale AI focus: The systems are intended for enterprises and AI developers running increasingly demanding computational workloads. • Pilot-based commercialization: Majestic plans customer pilot deployments to validate performance and cost benefits before broader general availability. • Cross-functional technical development: The company is expanding its AI data science, software, silicon and systems teams.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://majestic-labs.ai/

MANAGEMENT TEAM / FOUNDER(S)	**Ofer Shacham – Co-Founder & CEO:** Prior to co-founding Majestic Labs in 2023, Ofer Shacham served as Vice President and Head of Silicon at Meta Reality Labs, where he led development of the custom silicon platform powering Meta's AR initiatives, including the Orion AR-glasses program. Before Meta, he led Silicon Design and Implementation for Consumer Hardware at Google, where he initiated and directed development of the Pixel Visual Core, Pixel Neural Core, portions of the Tensor platform, and multiple generations of Google's Titan security processors. Earlier, he founded Chip Genesis, a Stanford spinout that was acquired by Google, and worked in IBM's research organization in Israel. Dr. Shacham holds a B.Sc. in Electrical Engineering and Computer Science from Tel Aviv University and M.Sc., Ph.D., and Postdoctoral degrees from Stanford University. At Majestic Labs, he leads development of memory-dense AI infrastructure designed to address the AI "memory wall." **Shahriar Rabii, Ph.D. – Co-Founder & President:** Prior to co-founding Majestic Labs, Sha Rabii served as Vice President and Head of Silicon Engineering at Meta, where he established silicon-engineering and technology-engineering organizations focused on AI, AR/VR, and advanced computing systems. Before Meta, he founded Google's silicon engineering organization and technology engineering group, overseeing development of products including the Pixel Visual Core, Pixel Neural Core, Tensor-related programs, Titan secure elements, and large-scale video-transcoding infrastructure. Earlier in his career, he founded Arda Technologies, co-founded Aeluros, and was a member of the founding team at Atheros Communications. Dr. Rabii holds a Ph.D. in Electrical Engineering from Stanford University. At Majestic Labs, he helps lead development of the company's AI servers, silicon architectures, systems design, and software platform. **Masumi Reynders, J.D. – Co-Founder & COO:** Prior to co-founding Majestic Labs, Masumi Reynders spent more than two decades in leadership positions spanning operations, product management, business development, partnerships, and silicon strategy. She held senior leadership roles at Meta, including Senior Director for Advanced Technologies and Senior Director for Silicon, where she led product strategy, partnerships, go-to-market activities, and IP licensing efforts across custom silicon, optics, and sensing technologies. Before Meta, she served as Director of Product Management for Google's custom silicon organization, helping drive programs including Pixel Visual Core, Titan security products, machine-learning accelerators, and advanced hardware technologies. Earlier in her career, she held business-development and partnership leadership roles at Motorola Mobility. She earned her undergraduate degree from the University of Michigan and is identified publicly as holding legal training (J.D.). At Majestic Labs, she oversees operations, team scaling, partnerships, and market expansion.
SIZE OF COMPANY	~59 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Majestic Labs has progressed from an early-stage startup into a 40-employee developer of memory-dense AI servers. Following its 2024 seed round and 2025 Series A, the company expanded its AI data science, software, silicon and systems teams and advanced its platform toward general availability. Majestic is also developing a full software stack and preparing customer pilot deployments to validate platform performance and cost benefits.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Majestic Labs has raised a total of $100MM. In its latest Series A round of funding, it raised $90MM in September 2025. Investors include Bow Wave Capital Management, Grove Ventures, QPVentures, Aidenlair Capital, Lux Capital, Tal Ventures, Hetz Ventures, Upfront Ventures, SBI Investment, and Pipeline Capital Partners.
LATEST ROUND VALUATION	N/A

VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale compute and networking businesses, P/S or EV/Sales may be appropriate once revenue is available. Additional operating benchmarks could include installed server capacity, memory capacity per system, utilization, system performance, customer pilots, design wins and contracted backlog. Comparable-company analysis may reference AI compute and systems providers such as CBRS and NVDA, together with private infrastructure companies such as SambaNova Systems, Lambda, Positron AI and UnifabriX. Significant adjustments would be required for Majestic's limited operating history, early commercialization stage, smaller workforce and lack of disclosed revenue.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside Taxonomy: Modal

COMPANY NAME	Modal
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Compute Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	Modal Labs operates in serverless cloud computing, AI infrastructure and infrastructure SaaS.
GEOGRAPHIC PRESENCE	New York, New York
NOTABLE COMPETITORS	**LISTED:** N/A **PRIVATE:** RunPod, Spell, Determined AI, Together AI, Baseten, Replicate.
KEY SUPPLIERS / CUSTOMERS	Suppliers: AMZN, GOOGL, MSFT, ORCL Customers: N/A
COMPETITIVE POSITIONING	Modal Labs provides a serverless cloud platform for deploying and scaling compute-intensive machine-learning and data-processing workloads. Its platform supports code-defined container environments, automatic scaling from zero to hundreds of GPUs, real-time monitoring and job scheduling. Key differentiators include: • AI-native serverless infrastructure: Modal provides cloud infrastructure purpose-built for AI workloads, abstracting underlying server and infrastructure management so developers can focus on applications and models. • Automatic GPU scaling: The platform automatically provisions and scales GPU-backed containers with workload demand, including scaling resources back to zero when idle. • Broad AI workload support: Modal supports inference, model training and fine-tuning, and large-scale batch processing within a common infrastructure platform. • Code-defined environments: Developers can define application logic, container environments, and hardware requirements directly in code through the Modal SDK, reducing reliance on separate infrastructure configuration. • Integrated orchestration and observability: The platform combines workload scheduling with integrated logging and visibility across functions, containers, and production deployments.
YEAR FOUNDED / ESTABLISHED	2021
WEBSITE	https://modal.com/
MANAGEMENT TEAM / FOUNDER(S)	**Erik Bernhardsson – Co-Founder, CEO & Board Member:** Prior to co-founding Modal Labs in 2021, Erik Bernhardsson served as Chief Technology Officer of Better.com, where he led an engineering organization of approximately 300 employees. Before Better, he spent six years at Spotify, first leading analytics engineering and later building and managing the company's machine-learning organization. He is widely recognized for developing early versions of Spotify's music recommendation systems, including technologies that powered Related Artists, Radio, and Discover Weekly. He also created and open-sourced the widely used Luigi workflow engine and Annoy approximate-nearest-neighbor library. Mr. Bernhardsson holds an M.Sc. in Physics from the KTH Royal Institute of Technology in Stockholm and previously interned at Google. At Modal, he has overseen the company's evolution into a serverless AI infrastructure platform serving machine learning, inference, and cloud-computing workloads. **Akshat Bubna – Co-Founder & CTO:** Prior to co-founding Modal Labs, Akshat Bubna was a Staff Software Engineer at Scale AI, where he worked on core engineering systems, natural-language-processing initiatives, and quality infrastructure. Earlier, he studied Mathematics and Computer Science at the Massachusetts Institute of Technology. Mr. Bubna is also recognized for his accomplishments in competitive programming, having earned a Gold Medal at the 2014 International Olympiad in Informatics after winning Bronze in 2013. At Modal, he leads engineering and product development and has helped build the company's underlying platform, including its custom runtime, scheduler, storage systems, orchestration infrastructure, automatic GPU scaling capabilities, and code-defined cloud architecture.
SIZE OF COMPANY	~243 employees on LinkedIn (accessed 9/22/2026)

GROWTH TRAJECTORY	Modal Labs has evolved into a revenue-generating serverless AI infrastructure provider supporting large-scale inference, training and batch-processing workloads. The company expanded from 14 employees in 2023 to 170 employees in 2026 and added automatic GPU scaling, monitoring, scheduling and cloud-tool integrations to its platform. PitchBook reports that revenue increased from $50MM in 2025 to $300MM in 2026, while Modal also expanded through the acquisitions of Dottxt, Tidbyt, Twirl and Jamsocket.
ANY MAJOR M&A / CORPORATE ACTIONS	Acquired Jamsocket in July 2025. Acquired Twirl in May 2025. Acquired Tidbyt in November 2024. Invested in Dottxt through a $9MM seed round in August 2024.
INVESTORS / INVESTMENT STAGE	Modal Labs has raised a total of $465MM. In its latest Series C round of funding, it raised $355MM in May 2026. Investors include Redpoint Ventures, General Catalyst, Definition, Accel, Scale Angels Fund, Menlo Ventures, Bain Capital Ventures, Transpose Platform Management, Lux Capital, Amplify Partners, Positive Sum, Dimension, AGI House, and Anti Fund.
LATEST ROUND VALUATION	$4,650MM as of May 2026 Series C.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale AI infrastructure and infrastructure software businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference private AI infrastructure providers such as RunPod, Together AI, Baseten and Replicate.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: Movandi

COMPANY NAME	Movandi
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / RF & Millimeter-Wave Semiconductors
SIZE OF TOTAL ADDRESSABLE MARKET	Movandi operates in 5G, millimeter-wave communications, fixed wireless access, satellite communications and integrated RF and antenna systems.
GEOGRAPHIC PRESENCE	Irvine, California
NOTABLE COMPETITORS	**LISTED:** ADI, MCHP, QCOM, SWKS **PRIVATE:** Echodyne, Vistance Networks, Siklu, Pivotal Commware, Airspan Networks, EdgeQ, Tarana Wireless
KEY SUPPLIERS / CUSTOMERS	Suppliers: N/A Customers: N/A
COMPETITIVE POSITIONING	Movandi develops wireless-network technology, integrated RF systems and antennas for high-frequency and millimeter-wave communications. Key differentiators include: • Integrated RF platform: Movandi combines internally developed RF chipsets, integrated antenna modules, algorithms, software, and complete system designs, enabling optimization across the RF stack rather than at the component level alone. • Smart-repeater architecture: The company's BeamXR platform uses active smart repeaters to amplify and redirect mmWave signals, addressing propagation, blockage, and coverage limitations across indoor, outdoor, and mobile environments. • Beamforming and phased-array technology: Movandi integrates phased-array antennas with beamforming, beam steering, multi-beam operation, and software-defined beam networking to dynamically manage high-frequency RF coverage. • High-frequency coverage expansion: The platform is designed to extend the range and coverage of mmWave networks while overcoming physical barriers and line-of-sight limitations associated with high-frequency wireless signals. • Broad wireless applications: Movandi's RF systems support 5G mmWave, satellite communications, and fixed wireless access, spanning terrestrial and non-terrestrial network applications. • Proprietary technology portfolio: The company describes its platform as built around patented RF technology, custom mmWave chipsets, MIMO/phased-array antennas, smart-repeater modules, and software-defined beam networking.
YEAR FOUNDED / ESTABLISHED	2016
WEBSITE	https://movandi.com

MANAGEMENT TEAM / FOUNDER(S)	**Maryam Rofougaran – Co-Founder, Co-CEO & Board Member:** Prior to co-founding Movandi in 2016, Maryam Rofougaran co-founded Innovent Systems in 1998 alongside her brother, Reza Rofougaran. Innovent pioneered RF CMOS wireless technologies and was acquired by Broadcom, where she subsequently played a key role in building Broadcom's wireless business into a multi-billion-dollar operation. At Broadcom, she served as Senior Vice President of Wireless Radios and helped drive development of wireless technologies deployed in billions of devices worldwide. She holds B.S. and M.S. degrees in Electrical Engineering from UCLA. A prolific inventor, she has been credited with more than 300 patents and numerous technical publications. At Movandi, she leads commercialization of the company's RF chipsets, phased-array antennas, smart-repeaters, satellite-connectivity solutions, and fixed-wireless infrastructure platforms. **Reza Rofougaran – Co-Founder, CTO & Board Member:** Prior to co-founding Movandi, Reza Rofougaran was one of the principal architects behind the wireless technologies that enabled modern smartphones and RF connectivity. He co-founded Innovent Systems in 1998, which was acquired by Broadcom in 2000. At Broadcom, he served as Senior Vice President of Engineering and later became a Broadcom Fellow, leading development of Bluetooth, Wi-Fi, NFC, GPS, 60GHz wireless, cellular radios, and integrated RF CMOS architectures. He earned B.S., M.S., and Ph.D. degrees in Electrical Engineering from UCLA. Dr. Rofougaran is an IEEE Fellow, a member of the National Academy of Engineering, and one of the world's most prolific inventors, with more than 800 issued patents. At Movandi, he drives the company's technology roadmap spanning RF integrated circuits, phased-array antenna systems, active repeaters, beamforming platforms, and millimeter-wave communications infrastructure. **Scott Burri – Co-Founder, CFO & Chief Administrative Officer:** Prior to helping launch Movandi, Scott Burri founded Huntington Ventures and served as its Managing Partner, investing in and advising early-stage semiconductor, communications, software, and technology companies. Throughout his career, he has served as a director, investor, and advisor to numerous venture-backed technology businesses including Sequoia Communications, Sensys Instruments, Vector Software, and other communications and semiconductor companies. He holds an undergraduate degree from Oregon State University and an MBA from the University of Southern California. At Movandi, he oversees finance, administration, fundraising, investor relations, and corporate operations as the company expands across 5G, fixed wireless access, and satellite communications markets.
SIZE OF COMPANY	~49 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	Movandi has evolved into a revenue-generating supplier of integrated RF, antenna and smart-repeater technologies for 5G and high-frequency wireless networks. The company expanded its product scope from active routers and repeaters into fixed wireless access, satellite communications and wireless infrastructure for AI applications. In 2025 and 2026, Movandi announced collaboration with Kyocera, smart-repeater deployments, new frequency synthesizers and a strategic partnership with ACES to expand next-generation 5G and 6G infrastructure in Saudi Arabia and the Middle East.
ANY MAJOR M&A / CORPORATE ACTIONS	N/A
INVESTORS / INVESTMENT STAGE	Movandi has raised a total of $114.6MM. In its latest Series C1 round of funding, it raised $41.2MM in October 2025. Investors include ITHCA Group, Murata Manufacturing Company, VT Alliance, Sierra Ventures, Cota Capital, DNX Ventures, Celesta Capital, 5G Ventures, Crescent Capital, Alumni Ventures, KT Investment, UCLA Investment Company, and WNC.
LATEST ROUND VALUATION	$227.5MM as of October 2025 Series C1
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale wireless semiconductor and communications-equipment businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference private wireless-infrastructure providers such as Pivotal Commware, EdgeQ, Tarana Wireless and Echodyne.
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: SiLC Technologies

COMPANY NAME	SiLC Technologies
SECTOR / SUB-SECTOR	TMT / Semiconductors & Photonics / FMCW LiDAR & 4D Vision
SIZE OF TOTAL ADDRESSABLE MARKET	SiLC Technologies develops silicon-photonics-based 4D vision systems utilizing Frequency Modulated Continuous Wave (FMCW) LiDAR. The company's products are designed to provide range, velocity, motion, and polarization data for industrial automation, physical AI, robotics, mobility, security, and machine vision applications. Management believes high-performance machine perception is a foundational enabler of robotics and autonomy. The company participates in the broader LiDAR, machine vision, robotics perception, autonomous systems, and physical AI markets.
GEOGRAPHIC PRESENCE	Monrovia, California
NOTABLE COMPETITORS	**LISTED:** AEVA, INVZ, OUST **PRIVATE:** Xscape Photonics, Scintil Photonics, Ayar Labs
KEY SUPPLIERS / CUSTOMERS	Strategic investors and partners include: Honda, Sony, Epson, and ROHM Semiconductor. Publicly disclosed customers include:AutoX (robotaxi deployment)
COMPETITIVE POSITIONING	• Silicon photonics platform: Proprietary silicon-photonics architecture integrates photonic functions on-chip to improve scalability, cost, and power consumption. • FMCW LiDAR technology: Unlike conventional time-of-flight LiDAR systems, SiLC's coherent sensing platform simultaneously measures distance and instantaneous velocity. • 4D vision architecture: Eyeonic systems capture range, velocity, motion, and polarization information, enabling richer scene understanding than traditional sensing approaches. • Physical AI focus: Management positions the company as a key enabler of robotics and autonomous systems requiring human-like machine perception
YEAR FOUNDED / ESTABLISHED	2018
WEBSITE	https://silc.com/
MANAGEMENT TEAM / FOUNDER(S)	**Mehdi Asghari, Ph.D. – Founder & CEO:** Prior to founding SiLC Technologies in 2018, Dr. Asghari spent more than two decades developing silicon photonics technologies and founded multiple silicon-photonics companies. He is a recognized pioneer in optical communications and integrated photonics. At SiLC, he leads development of the company's FMCW LiDAR and 4D vision platform, which combines coherent sensing with silicon photonics to enable advanced machine perception for robotics, industrial automation, security, and mobility applications.
SIZE OF COMPANY	86 employees on LinkedIn (accessed 9/22/2026)
GROWTH TRAJECTORY	SiLC has commercialized its Eyeonic family of products and expanded into industrial automation, robotics, perimeter security, and mobility applications. Notable milestones include deployment with AutoX, strategic investment from Honda, and continued expansion of its Eyeonic product line. The company has also demonstrated sensing ranges exceeding one kilometer using its coherent FMCW platform.
ANY MAJOR M&A / CORPORATE ACTIONS	Strategic investment from Honda (2024) to support next-generation mobility sensing solutions.
INVESTORS / INVESTMENT STAGE	Notable investors include Dell Technologies Capital, Andreessen Horowitz, Intel Capital, Sony Innovation Fund, Honda
LATEST ROUND VALUATION	N/A
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	Photonics and sensing companies are typically valued using: • EV/Revenue • EV/Gross Profit • Strategic technology premiums • Comparable transactions in LiDAR, photonics, and robotics perception markets
SECONDARY PRICING DATA	N/A

Source: Pitchbook and Company's reports and website, LinkedIn

Outside taxonomy: TensorWave

COMPANY NAME	TensorWave
SECTOR / SUB-SECTOR	TMT / Semiconductors & Semi Cap Equipment / AI Compute Infrastructure
SIZE OF TOTAL ADDRESSABLE MARKET	TensorWave operates in AI cloud computing, GPU infrastructure, large-scale model training and real-time inference.
GEOGRAPHIC PRESENCE	Las Vegas, Nevada
NOTABLE COMPETITORS	**LISTED:** NVDA, CBRS, CRWV **PRIVATE:** Graphcore, SambaNova Systems, Lambda, Crusoe.
KEY SUPPLIERS / CUSTOMERS	Key supplier / technology provider: AMD. TensorWave's cloud infrastructure uses AMD accelerators, and the Series B proceeds are intended to deploy additional AMD chips across data centers. Disclosed customer: Luma AI.
COMPETITIVE POSITIONING	TensorWave provides bare-metal AI cloud infrastructure for large-scale training and real-time inference workloads. Key differentiators include: • AMD-based AI infrastructure: TensorWave builds its AI cloud around AMD Instinct accelerators, with infrastructure optimized specifically for memory-intensive AI training and inference workloads. • Bare-metal deployment: The platform provides dedicated bare-metal servers with direct hardware access and no virtualization overhead, giving customers end-to-end control and predictable performance. • High-capacity memory and networking: TensorWave leverages high-memory-capacity AMD Instinct accelerators alongside high-bandwidth interconnects to support large-model training and inference at scale. • Flexible AI workload scaling: The platform supports training, fine-tuning, and inference workloads, with deployment options spanning bare-metal infrastructure and managed Kubernetes and Slurm environments. • Open AI ecosystem: TensorWave supports widely used AI frameworks and libraries including PyTorch, JAX, TensorFlow, and vLLM, positioning its AMD-based infrastructure as an alternative to more vertically integrated GPU ecosystems. • Enterprise-grade deployment: The platform combines infrastructure monitoring with enterprise security capabilities, including SOC 2 Type II and ISO 27001 certification, to support production AI workloads.
YEAR FOUNDED / ESTABLISHED	2023
WEBSITE	https://www.tensorwave.com

MANAGEMENT TEAM / FOUNDER(S)	**Darrick Horton – Co-Founder, CEO & Chairman:** Prior to co-founding TensorWave in 2023, Darrick Horton was a technology entrepreneur focused on cloud infrastructure, semiconductors, and accelerated computing. Most notably, he co-founded and led VMAccel, an FPGA-focused cloud computing company, serving in leadership positions including CTO and CEO. Earlier in his career, he worked on Lockheed Martin Skunk Works projects involving compact nuclear fusion research and participated in NASA-funded plasma physics research. Public biographies also reference involvement in astrophysics-related research through the LIGO program. Mr. Horton studied mathematics, physics, and mechanical engineering at Andrews University. At TensorWave, he has led development of one of the largest AMD-based AI cloud platforms in North America, expanding the company into large-scale AI training infrastructure and hyperscale data-center deployments. **Piotr Tomasik – Co-Founder, President, COO & Board Member:** Prior to co-founding TensorWave, Piotr Tomasik spent more than two decades building software, SaaS, fintech, and AI-enabled technology businesses. He is best known as co-founder and former CTO of Influential, an AI-driven influencer-marketing platform that was later acquired by Publicis. He also founded or led several other technology companies, including Lets Rolo, On Guard Data, and ActiveSide, while holding senior technology leadership positions at CARD.com and Marker Trax. In addition, he serves as a General Partner at 1864 Fund and co-founded StartUp Vegas. Mr. Tomasik earned a B.S. in Computer Science with a Mathematics minor from the University of Nevada, Las Vegas. At TensorWave, he oversees operations, infrastructure deployment, power procurement, customer execution, and data-center scaling initiatives.
SIZE OF COMPANY	~171 employees on LinkedIn (accessed 9/16/2026)
GROWTH TRAJECTORY	TensorWave has progressed from a 2023 startup into a revenue-generating AMD-based AI cloud provider supporting large-scale model training and real-time inference. The company generated $100MM of revenue in 2025 and used its Series A to accelerate deployment of an AMD Instinct MI325X training cluster. In 2026, TensorWave expanded through AMD Helios rack-scale infrastructure, supported production inference for Luma AI and entered a $6.5B AI data-center lease with Fermi America to increase future capacity.
ANY MAJOR M&A / CORPORATE ACTIONS	FundNV sold its entire stake to an undisclosed buyer through a private secondary transaction. Pricing was not disclosed. TensorWave entered a reported $6.5B AI data-center lease with Fermi America in August 2026.
INVESTORS / INVESTMENT STAGE	TensorWave has raised a total of $499.9MM. In its latest Series B round of funding, it raised $350MM in May 2026. Investors include AMD Ventures, Magnetar Capital, Western Frontier, Fifth Growth Fund, Nexus Venture Partners, Maverick Silicon, Gimmel Ventures, Prosperity7 Ventures, Aramco Ventures, Atreides Management, TransLink Capital, Javelin Venture Partners, Cadenza Capital, SuRo Capital, and StartupNV.
LATEST ROUND VALUATION	$1,550MM as of May 2026 Series B.
VALUATION APPROACHES / METHODOLOGIES USED IN THE SECTOR	For sub-scale compute and AI cloud infrastructure businesses, P/S or EV/Sales may be appropriate once revenue is available. Comparable-company analysis may reference listed AI compute and cloud providers such as NVDA, CBRS and CRWV, together with private infrastructure companies such as SambaNova Systems, Lambda and Crusoe.
SECONDARY PRICING DATA	FundNV sold its stake to an undisclosed buyer through a private secondary transaction. PitchBook does not disclose the transaction date, consideration, number of shares or per-share price.

Source: Pitchbook and Company's reports and website, LinkedIn

Valuation Method and Risk Statement

We use various valuation techniques such as P/E, EV/FCF for valuing the companies in this report. Risk factors include but are not limited to macroeconomic factors such as a downturn in the economy, a disruption of international trade, technological disruption due to new inventions, or business model innovation whereby structural changes in the industry alter the future course of unit sales, ASPs, and revenues.