Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

FOUNDATION

October 13, 2026 04:00 AM GMT

Global Thematics & Technology | North America

Spark Preview: What Will Privates Tell Us About the AI Cycle?

Private tech is a leading indicator of where capital, talent & tech are moving. At our Spark Private Company Conference (Oct. 13-15), we expect to speak to 100+ private companies across inference, vertical AI, security, semis/hardware, embodied AI, and more. We outline themes we're most focused on.

PrivateCompany

Key Takeaways

- Our Spark Private Company Conference will be held at Dana Point, CA, on October 13-15, bringing together founders and CFOs across private companies.
- We believe privates are a key signal of where the market is headed and can provide key insights on debates like power, compute efficiency, and agentic AI.
- Areas we're focused on: 1) rethinking the AI factory, 2) shift to agents, 3) inference-aaS, 4) security, 5) embodied AI, 6) defense, 7) healthcare, 8) fintech.
- We outline the key themes and companies we're watching in each of these buckets, as well as Morgan Stanley Research reports to get up to speed on these topics.

Tech Diffusion

A Morgan Stanley Research Key Theme of 2026

In a world of rapid innovation, the privates ecosystem has never been more important for public market investors. Some of the most consequential shifts in technology are showing up in startups first, acting as a leading indicator of where capital, talent, and technology is moving. That signal matters more than ever, as we navigate key debates like access to power, compute efficiency, AI shifting from assistant to agent, the role of cyber in an AI world, the rising value of inference, and implications for verticals like defense, logistics, and robotics. Each of these questions has direct implications for where value accrues across the stack, which categories of incumbent players face pressure, and how we can unlock enough durable ROIC from AI to justify the trillions of dollars spent on capex.

We look to unpack the most important debates in tech at our Spark Private Company Conference (Oct. 13-15). This conference brings together founders / CFOs from over 100 leading private companies, alongside investors and Morgan Stanley Research and Banking. In this note, we preview 1) the key themes we're watching, 2)

MORGAN STANLEY & CO. LLC

Stephen C Byrd
Equity Analyst
Stephen.Byrd@morganstanley.com

Sabrina E Hao
Equity Analyst
Sabrina.Hao@morganstanley.com

Adam Jonas, CFA
Equity Analyst
Adam.Jonas@morganstanley.com +1 212 761-1726

Brian Nowak, CFA
Equity Analyst
Brian.Nowak@morganstanley.com

Adam Wood
Equity Analyst
Adam.Wood@morganstanley.com +1 212 761-3656

Meta A Marshall
Equity Analyst
Meta.Marshall@morganstanley.com

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

Sanjit K Singh
Equity Analyst
Sanjit.Singh@morganstanley.com

Elizabeth Porter, CFA
Equity Analyst
Elizabeth.E.Porter@morganstanley.com +1 212 761-3632

William Tackett, CFA
Research Associate
William.Tackett@morganstanley.com

Jonathan Eisenson
Research Associate
Jonathan.Eisenson@morganstanley.com +1 212 761-2808

Gregory Gao
Research Associate
Greg.Gao@morganstanley.com

Nikhil Javeri
Research Associate
Nikhil.Javeri@morganstanley.com +1 212 761-3742

Lucas Cerisola
Research Associate
Lucas.Cerisola@morganstanley.com

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

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

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

companies doing firesides that we're excited to hear from, and 3) Morgan Stanley Research to get you up to speed on these topics. For a list of attendees, see Getting Up to Speed: Overview of Our Private Company Attendees.

Exhibit 1: Expected Attendees at Morgan Stanley's Spark Privates Conference



Source: Company Websites, Morgan Stanley Research

Exhibit 2: Eight themes we're focused on for this year's Spark Private Company Conference.



Source: Morgan Stanley Research

Getting Up to Speed: Overview of Our Private Company Attendees Getting Up to Speed: Overview of Our Private Company Attendees

1. The Rise of the AI Factory. As we wrote here, we are seeing new business models emerge around how to rethink the data center and coordinate increasingly complex infrastructure — power, compute, networking, cooling, and data — as part of a broader AI factory. We see 4 categories of companies in this theme: ***1) Power*** *(or accelerating access to it),* ***2) Compute*** *(increasing useful compute delivered per unit of power, given key bottlenecks),* ***3) Integration of Power + Compute*** *(redesigning power and compute as a single integrated system),* ***and 4) Reimagining the Data***

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

Center. Venture funding for this universe of startups has accelerated sharply in 2026, reaching $33.5 billion by early October, up 2x from 2025. Power and Compute — or the scarcest inputs in the ecosystem (see our work on the power shortfall here and here) — have accounted for more than 80% of 2026 capital. However, we see capital broadening to Integration of Power + Compute and Re-imagining the Data Center.

Who are we watching at Spark? We are interested in how **d-Matrix** sees the market for disaggregated inference solutions and the advantages of their Digital In-Memory Compute (DIMC) architecture vs. more common SRAM and HBM approaches. Within Networking, our team has written extensively about the value of optical in an AI world and the key debates around NPO/CPO adoption, the transition beyond pluggable optics, and the emergence of open networking architectures. We're interested in how **Lightmatter** is bringing optical connectivity closer to compute while reducing fiber requirements through bidirectional optics, **Upscale AI** is building open networking solutions that connect different accelerators across scale-up and scale-out networks, and **Ayar Labs** is integrating optical I/O directly into accelerator packages to overcome electrical connectivity limitations.

How to get up to speed on these themes...see relevant team research below:

US Thematics: Venture Vision: Data Center to AI Factory

Global Thematics & US Semiconductors: What Does the Power Shortfall Mean for Semis? (5 Oct 2026)

Powering AI: Racking Up the Power Shortfall (21 Sep 2026)

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

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

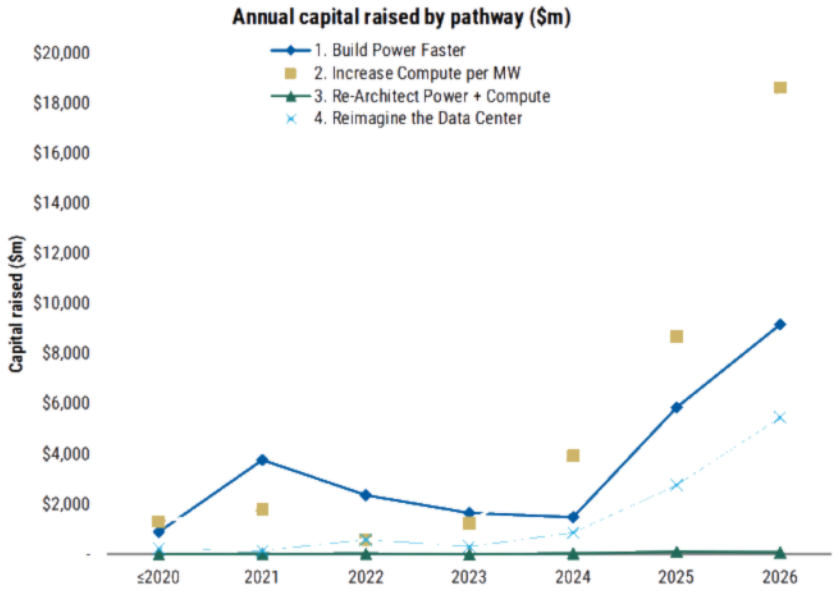
Technology: Bringing More Light to Scale-Up Networks: An Updated Scale-Up Primer (13 Jul 2026)

Video | Spring Training: Telecom & Networking Equipment and Cybersecurity (26 May 2026)

Global Technology: AI Transceivers: Growth Dominates Disruption (26 Feb 2026)

Telecom & Networking Equipment: Into the Spotlight: Optical Market Opportunities (23 Feb 2026)

Exhibit 3: Most of the capital raised to date has been focused on compute and power, but we see an increasing opportunity within rearchitecting power/compute and reimagining the data center.



Source: PitchBook, Morgan Stanley Research. Completed financings.

2. AI's Move from Assistant to Agent. Consumer and enterprise AI agents have come into focus after the launch of Muse, Google's Gemini AI agent, and others. On the enterprise side, our focus is on how AI-native software vendors are monetizing agents that execute work on a customer's behalf, a step beyond the assistant products that defined the first wave of enterprise AI. That step moves software's addressable market toward labor budgets, which are larger than IT budgets. We expect to learn more about three themes that will determine how much of that opportunity accrues to software vendors. The first is **value capture,** meaning how vendors frame the share of displaced labor cost they can charge for, and whether that share settles closer to 5% or 50% once customers keep their portion as margin. The second is **pricing model evolution,** as vendors move off per-seat toward usage- and outcome-based models. We will be listening for how outcomes are defined and verified in contracts, and what that means for revenue visibility. The third is **unit economics**, where gross margins depend on how efficiently vendors route each task to the lowest-cost capable model and shift workloads onto smaller or proprietary models. We will be watching this across a number of presentations at the conference (**Cognition, Replit, Legora, Mercor, Genspark, Elise AI, Cresta, Pallet**).

Microsoft: Battle of the AI Stack: Built for a Multi-Model World (7 Oct 2026)

Meta Platforms Inc: 23 Days of Musing...Top Investor Questions and Our Answers (30 Sep 2026)

Internet: The Morgan Stanley Agentic Playbook (18 Sep 2026)

Software: The Moat & The Journey – Assessing Software Durability + Growth in the Age of AI (21 Jul 2026)

3. Model Customization and the Value of Inference-aaS Platforms. As enterprise AI moves from experimentation into production and the world becomes multi-model, we expect a larger share of value in the model stack to accrue to the layers that serve and customize models. **Fireworks AI** and **Baseten,** both inference platforms at our conference, should give us a read on key themes. The first is the **rising importance of the orchestration layer,** where platforms decide which model handles each request and how that workload scales across hardware. The value of this only rises as companies look to use both open and closed-weight models in order to rationalize token spend. We want to understand whether this layer can sustain pricing as model providers and hyperscalers offer competing serving options. The second is **efficiency in deployment,** as customers look to get more output from less compute through smaller distilled models and higher GPU utilization. The open question is how much of the resulting cost decline platforms keep as margin and how much they pass through to customers. The third is the **growing value of post-training relative to pre-training,** as customers use fine-tuning and reinforcement learning on proprietary data to adapt open models to specific tasks. We will be listening for whether post-training creates durable differentiation and switching costs. These themes also bear on the agent vendors above, since inference is a primary component of their cost of revenue and efficiency gains at this layer flow directly into their gross margins. The fourth is **token cost monitoring and budgeting.** Companies are increasingly looking for offerings to monitor their token spend and look more in a more granular fashion into ways to rationalize spend, and we're interested in how startups are addressing this.

Microsoft: Battle of the AI Stack: Built for a Multi-Model World (7 Oct 2026)

Internet and Software: Why Amazon Bedrock and Microsoft Foundry Should Become Household Names (7 Oct 2026)

Global Thematics: Weighing In: Open-Weights Models & 3 States of the World (3 Aug 2026)

Software: The Moat & The Journey – Assessing Software Durability + Growth in the Age of AI (21 Jul 2026)

4. Security Gets Rebuilt for the Agent Era. Security in the age of AI has become an increasingly important topic, given the litany of agent breaches seen over the past few months. We believe increasing autonomous agent attacks, agents taking on new behaviors, and an increasing speed/frequency of breaches require new technologies to identify and protect against these threats. This is against an increasingly attractive spending backdrop — according to IDC, the ~$300bn cybersecurity market (software and services) is set to grow at a 15% CAGR over the next couple of years, a number we think could accelerate to 23% in our base case as magnified threats heighten the need for next-generation security.

Who are we watching at Spark? Security companies in attendance range from pure software names to hardware names. An assortment of cyber/security companies are to be in attendance at the conference. Particularly interesting, we think, are the names on our panel, **Island** and **Verkada,** all at the front lines of cyber / security defense, with platforms that are quickly evolving. For Island, focus traditionally on an enterprise browser has expanded into a more robust platform as a widening set

of interactions with AI takes place. For Verkada, AI is moving the company towards becoming the OS for the physical world, with a growing set of control points to keep customers secure and connected.

Cybersecurity & Global Thematics: Attack of the Agents: Cybersecurity Scenarios In 3 States of the World (20 Sep 2026)
Cybersecurity: Agents on the Loose: Cyber Faces an Identity Crisis (20 Aug 2026)

Software: Security in Runtime: SIEM Market Evolution (1 Jun 2026)

5. Physical AI finally becoming mainstream. The implications are vast. Over the next 12-18 months, we expect a number of large new entrants into the physical AI debate, including large private labs as well as companies in the Mag 7, which will both increase competition among the startups, but also bring substantial credibility into the theme. To be more specific, we think the recent shift among the labs from competing purely on who is at the edge of the frontier, to instead focusing on productization will eventually lead into the physical world in search of large TAM expansion opportunities (LLMs > coding agents > agentic platforms > consumer devices > devices that move, i.e., robotics). We believe the ripple effects are highly underappreciated. The large gap in excitement for physical AI between our early stage VC contacts and public equity clients, to us, is an indication of what's to come.

Who are we watching at Spark? 1) Distributed compute: We believe the need for low-latency and data sovereignty for robotics could lead to a significantly greater % of inference compute deployed on the edge, both on-premise and on-device (key companies: **SiMa.ai, Liquid.AI**); 2) First-principles redesigned manufacturing platforms: AI is opening the door to entirely new ways of production as we face manufacturing shortfalls in the US for robotics and defense (key company: **Hadrian**); 3) Vision-capture: we believe a physical AI world will be adorned with a range of sensors to help AI understand the world around them (key companies: **Flock Safety, Verkada**); 4) Physical AI foundational models: The challenge with physical AI is the lack of internet-scale training data to train advanced general purpose models, many companies are taking unique approaches to circumvent this gap (companies to watch: **Genesis AI, Nuro**). *Separately, we are holding our Bay Area "RoboFest" bus tour October 12-13 visiting a number of disruptive physical AI startups including: Applied Intuition, Bright Machines, Genesis AI, Matter, Nuro, Physical Intelligence, and Waabi.*

A number of companies are also helping to automate the logistics behind physical work, such as **Flock Freight** that makes shared truckload software that pools compatible shipments onto one truck so shippers pay only for the space they use, while **Stord** pairs outsourced fulfillment with order and warehouse management software across ~100 sites and 1,000+ brands.

See relevant team research below:

Robotics: RoboFest California Preview: Bots by the Bay (16 Sep 2026)

Robotics: Actuators: Robot Muscle Needs American Hustle (13 Sep 2026)

Physical AI: The Morgan Stanley Robot Almanac (7 Dec 2025)

6. Autonomy and Software Are Redefining Defense. The line between commercial

technology and national security is blurring, as geopolitical tension, drone warfare, and the speed of modern threats push governments toward software-defined, AI-enabled systems and away from traditional platforms. We believe the shift toward autonomy, real-time intelligence, and lower-cost, scalable systems is creating a new class of defense and public safety technology companies.

Who are we watching at Spark? Companies in attendance span the stack from autonomous systems to intelligence to public safety. We're interested in how **Anduril** is building autonomous defense systems on a common software layer, and in how **Shield AI** is applying AI pilot technology to enable autonomous flight, including in GPS- and communications-denied environments. **Vantor** provides satellite imagery and geospatial intelligence, and we're curious how demand is evolving as customers look for faster, higher-resolution views of what is happening on the ground. **Dataminr** provides real-time event and risk detection, and we're interested in its ability to surface early signals from public data and turn them into actionable alerts for governments and corporate security teams. **Flock Safety** supplies public-safety cameras, sensors, and cloud software to police departments, businesses, schools, and neighborhoods, using AI to help identify vehicles and link evidence across incidents to speed up investigations. **Varda** makes recoverable capsules that manufacture pharmaceuticals in orbit and provide hypersonic reentry testing for US defense customers.

Defense: Dollars, Disruptors, and De(Consolidation); Identifying Investment Opportunities (6 Mar 2026)

Defense Tech: Anduril: a NeoPrime(r) (12 Jan 2026)

Robotics: The Robot Almanac (Vol. 5): Space & Defense (18 Dec 2025)

7. Healthcare Tech: AI Targets the Administrative Burden. Healthcare is a regulated, labor-intensive industry where a clear AI use case is administrative work, such as documentation, scheduling, intake, insurance verification, and revenue cycle. We see this as a close parallel to the shift from assistant to agent in software, with labor budgets rather than IT budgets as the addressable pool. Our team has written about the very large opportunity in healthcare revenue cycle management to tackle very significant inefficiencies facing the system today, with automation potential is highest in prior authorization, coding, and denials. Ambient scribes are another early AI use case, and we want to understand how vendors think about pricing, differentiation, and the risk of EHR vendors bundling competing products. We're watching if proprietary data becomes the durable moat as AI lowers the barrier to building workflow tools and how to think about competition between incumbents and new entrants.

Transforming Healthcare: AI Is Altering The Competitive Landscape In Revenue Cycle Management (5 Oct 2026)

Transforming Healthcare: Private Company Research Series: Latest Checks On Ambient AI Scribe and LLM Use In Health Systems (20 Jul 2026)

Transforming Healthcare: Private Company Research Series: Ambient AI Scribes Take Healthcare By Storm (24 Nov 2025)

8. Fintech Is Being Rebuilt Around Agents and New Rails. We see three themes in fintech. First, agentic finance: As AI agents increasingly access data, manage money, and transact on behalf of consumers and businesses, we believe the infrastructure for payments, identity, and settlement has to adapt. Second, digital assets and stablecoins are moving from the periphery toward mainstream payments and capital markets infrastructure. Third, vertical fintech companies are embedding payments, credit, and insurance directly into industry workflows, often serving customers that incumbents have underserved. Across all three, we want to understand how fraud and compliance costs evolve as AI raises the stakes on identity, and whether AI-driven automation can meaningfully lower the cost to serve. **Moonpay** (crypto-payment infrastructure) and **Mercury** (banking, payments, and financial software) are two companies we are watching within this theme, as well as **AtoB** (vertical fintech for trucking workflows).

Digital Assets: Digital Assets: From Disruption to Convergence (8 Sep 2026)

Global Banks & Exchanges: Digital Rails, Real Economics: Digital Assets & the Future of Wholesale Banking (2 Jun 2026)



Getting Up to Speed: Overview of Our Private Company Attendees

Exhibit 4: Expected Attendees at Morgan Stanley's Spark Privates Conference.



Source: Company Websites, Morgan Stanley Research

Company Profiles — By Fireside

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

A Conversation with Babak Siavoshy, CFO of Anduril

Anduril

Company Snapshot

Anduril	
Description	Autonomous defense systems and weapons
Founded	2017
Last raise / valuation	~$5b Series H (May 2026); ~$61b valuation
Key investors	Thrive Capital, Andreessen Horowitz (leads)
Operating scale	~$2.2b revenue (2025)

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

Source: Company Announcements, Anduril raises $5B, doubles valuation to $61B', TechCrunch (May 13, 2026)

- **Company description:** Anduril develops defense software and hardware (incl. autonomous vehicles, weapons, and sensors) primarily for US and allied governments. Its AI-enabled Lattice software integrates data from sensors, drones, and robotics to enable customers to detect, track and intercept unmanned aircraft systems (UAS) and autonomous drone systems. Beyond surveillance and counter-drone systems, Anduril also produces combat aircraft, missiles, solid rocket motors, underwater vehicles, and defense-ready compute, and communications infrastructure.

- **Product/ service verticals: Lattice**: Horizontal software operating platform connecting Anduril's hardware. It ingests sensor data from Anduril hardware, creates a common operating picture, enables command and control, and tasks autonomous systems. Its software development kit (SDK) also allows third-party hardware and applications to integrate with the platform. **Sensing, surveillance & counter-UAS:** Persistent UAS sensing and base and border protection through Sentry towers, Wisp infrared sensing, Spyglass/ Spark radar and Pulsar electromagnetic warfare. Anduril also produces UAS interceptors such as Anvil and Roadrunner. **Autonomous systems & precision strike:** Air strike through Ghost, Bolt, and Altius; autonomous strike via Fury; low-cost long-range strike via Barracuda family. Maritime underwater vehicles for surveillance, seabed monitoring, payload delivery and strike via Dive-LD, Dive XL, Copperhead and Seabed Sentry sensors. **Tactical infrastructure:** Deployable compute, communications, networking and soldier-level interfaces via products such as Menace, Voyager, EagleEye and Lattice Mesh. **Solid rocket motors (SRMs):** Produces SRMs, a critical missile component, for its own weapons and as a merchant supplier to third parties.
- **Listed company relationships:** Competes with **AeroVironment** (NASDAQ: AVAV) in portable strike drones and **Kratos** (NASDAQ: KTOS) in autonomous combat aircraft; its solid-rocket-motor business challenges **Northrop Grumman** (NYSE: NOC) and **L3Harris** (NYSE: LHX). **RTX/Raytheon** (NYSE: RTX) competes with Roadrunner-M in counter-drone interception but also partners with Anduril on advanced missile propulsion. **Boeing** (NYSE: BA) uses Anduril motors in its developmental Army interceptor, while **Saab** (NASDAQ Stockholm: SAAB B) selected Anduril to develop and produce motors for the Boeing–Saab Ground-Launched Small Diameter Bomb. **Palantir** (NASDAQ: PLTR) partners on military-data and command-and-control integration and leads the Army battlefield-intelligence program for which Anduril supplies hardware. L3Harris integrates night-vision equipment into Anduril's soldier systems. **Rheinmetall** (XETRA: RHM) partners on European versions of Fury aircraft and Barracuda missiles; **HD Hyundai Heavy Industries** (KRX: 329180) partners on autonomous surface vessels; and **Archer Aviation** (NYSE: ACHR) co-develops the Thunder hybrid-electric aircraft, which remains under development. **Rocket Lab** (NASDAQ: RKLB) is contracted to provide three hypersonic test launches under a ~$30m agreement funded by Anduril.
- **Key near-term milestones: 1)** Initial surface-launched Barracuda missile deliveries for the US Army targeting 1H27; **2)** Delivery of 600+ Bolt-M targeting April 2027; **3)** Long Beach campus (~1.18m sqft) opening targeting 2027; **4)** Air-launched Barracuda-500 deliveries targeting 2027.

The New Security Perimeter: Protecting the Digital and Physical Enterprise

Verkada

Company Snapshot

Verkada	
Description	Cloud-based physical security platform
Founded	2016
Last raise / valuation	CapitalG-led investment (December 2025) ~$5.8b
Key investors	CapitalG (Alphabet; NASDAQ: GOOGL) (lead), NVIDIA (NASDAQ: NVDA) (July 2026)
Operating scale	~$1b annualized bookings (December 2025), ~31k customers, ~2.4m devices

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

Source: Company Announcements, Verkada Reaches $5.8B Valuation with Investment from CapitalG, TechCrunch (Dec 3, 2025)

- **Company description:** Verkada is a physical-security company. Its cloud-based management platform, Command, consolidates disconnected security systems, centralizing monitoring and incident investigations across buildings.
- **Product/service verticals: Command** connects cameras with artificial intelligence (AI) footage search, door controllers and intercoms for entry, alarms and environmental sensors for alerts, and Workplace for visitor management. On-camera storage, combined with cloud management, removes the need for separate video recorders.
- **Listed company relationships: NVIDIA** (NASDAQ: NVDA) Partner on AI video-search development using Cosmos models, announced July 2026; **CDW** (NASDAQ: CDW) Reseller of Verkada licenses; **Motorola Solutions** (NYSE: MSI) Competitor through Avigilon Alta.
- **Key near-term milestones: 1)** Sound Systems speakers and audio players targeting Oct 2026; **2)** Transportation fleet camera availability in the US and Canada targeting November 30, 2026.

Island

Company Snapshot

Island	
Description	Enterprise browser and workplace security
Founded	2020
Last raise / valuation	~$400m Series F (September 2026); ~$6.4b
Key investors	Evolution Equity Partners (lead); Sequoia Capital, Coatue, Insight Partners
Operating scale	Millions of users; ~1,000 employees

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

Source: Company Announcements, Island Announces $400 Million Series F, Bringing Valuation to $6.4 Billion, Island (Sep 24, 2026)

- **Company description:** Island sells enterprise-browser and workplace-security software that controls application access and data sharing.
- **Product/service verticals: Enterprise Browser** uses Chromium's open-source foundation to enforce access, copy/paste and download controls locally. **Enterprise Platform** expansion extended shared policies to consumer browsers. **Enterprise Network** combines on-device checks with selective cloud inspection, reducing reliance on routing traffic through centralized security systems. Enterprise AI governs prompts and agent permissions, automates workflows and enables deployment of internal AI applications.
- **Listed company relationships: Amazon** (NASDAQ: AMZN) Customer; **CrowdStrike** (NASDAQ: CRWD) Integration partner, with Island feeding browser logs into Falcon security analytics; **Palo Alto Networks** (NASDAQ: PANW) Competitor via Prisma Browser.
- **Key near-term milestones:** N/A

What It Takes to Make AI Work in the Real World

Baseten

Company Snapshot

Baseten	
Description	AI model inference platform
Founded	2019
Last raise / valuation	~$1.5b Series F (June 2026); ~$13b
Key investors	Altimeter, Conviction, Spark Capital (leads), Sands Capital, Wellington (co-leads)
Operating scale	20x revenue and 40x inference-volume growth yoy (March 2026)

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

Source: Company Announcements, Announcing our Series F, Baseten raised a $1.5B Series F and achieved a $13B valuation, Baseten (Jun 22, 2026)

- **Company description:** Baseten is a model-serving platform for AI inference, enabling developers and enterprises to deploy open-source, custom and fine-tuned AI models with optimized runtimes, high availability and developer-oriented workflows, either on Baseten-managed infrastructure or on the customer's cloud. It focuses on the infrastructure required to take trained AI models from development into production at scale.
- **Product/ service verticals: Dedicated Inference:** Runs customers' proprietary, fine-tuned or open-source models. Sitting between models and underlying GPU/cloud infrastructure, Baseten optimizes execution and manages deployment, scaling, traffic routing, monitoring and reliability. Deployable on Baseten Cloud, customer-cloud or hybrid configurations. **Model APIs:** Provides pre-hosted, optimized models through APIs, without customers managing deployment. **Training/post-training:** Training Jobs runs customers' existing training scripts on managed GPUs.

Loops, its reinforcement learning training toolkit, remains in early access and supports moving trained models directly into production inference. **Baseten for Model Labs:** Enables closed-weight model developers to distribute and monetize models through branded APIs and Baseten's model library, with access controls, billing and supporting infrastructure.

- **Listed company relationships: NVIDIA** (Nasdaq: NVDA) – Supplier of GPUs and Dynamo inference software with chips accessed through cloud providers; **Alphabet** (Nasdaq: GOOGL) – Supplier / Competitor on cloud compute and fine-tuned model hosting respectively; **Amazon** (Nasdaq: AMZN) – Competes on managed inference via SageMaker AI; **HubSpot** (NYSE: HUBS) – Customer of inference platform

- **Key near-term milestones: 1)** 3x headcount this year; **2)** Next Base Labs research fellowship cohort starting Jan 2027.

Legora

Company Snapshot

Legora	
Description	AI software for legal teams
Founded	2023
Last raise / valuation	~$550m Series D (March 2026) ~$5 5b
Key investors	Accel (lead), Atlassian (NASDAQ: TEAM), NVIDIA (NASDAQ: NVDA) via NVentures (extension)
Operating scale	~$200m+ ARR, ~2 100 firms, ~130k monthly users (September 2026)

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

Source: Company Announcements, Legora raises $550 million Series D to fuel US growth", Legora (Mar 10, 2026)

- **Company description:** Legora provides AI software to law firms and in-house legal teams. It combines internal documents and legal sources to reduce manual research, review and drafting, while retaining lawyer oversight.
- **Product/service verticals: Agent** executes multi-step assignments; **Tabular Review** answers standardized questions across document sets in source-linked tables. **Legal Research** retrieves cited legal authorities; **Word Add-in** drafts documents and proposes tracked edits. Shared firm precedents, reusable instructions and document-system integrations connect these tasks into repeatable workflows.
- **Listed company relationships: Salesforce** (NYSE: CRM) Customer adopting Legora across its global legal organization, announced September 2026; **Alphabet** (NASDAQ: GOOGL) Supplies cloud/models and partners on a beta Gemini Enterprise for Legal connection announced August 2026; **Thomson Reuters** (NASDAQ: TRI) Competitor in research, review and drafting through CoCounsel Legal/Westlaw; **RELX** (NYSE: RELX) Competitor via LexisNexis' Lexis+ with Protége.
- **Key near-term milestones: 1)** Legal-research ontology targeting Q4 2026

Mercor

Company Snapshot

Mercor	
Description	Expert human data for AI
Founded	2023
Last raise / valuation	~$350m Series C (October 2025); ~$10b; in talks at ~$20b (July 2026)
Key investors	Felicis (lead); Benchmark, General Catalyst, Robinhood Ventures (NASDAQ: HOOD)
Operating scale	~$2b ARR (June 2026)

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

Source: Company Announcements, Unlocking Human Potential in the AI Economy, Mercor (Oct 27, 2025)

- **Company description:** Mercor supplies human expertise and training data to artificial intelligence (AI) laboratories and enterprises. Specialists translate professional judgment into tasks and scoring rules, helping customers train and test models on complex workflows rather than generic questions.
- **Product/service verticals: Human Data** recruits and manages expert teams, monetized through hourly billing and a percentage-based fee tied to expert pay. Its "Open Box" approach provides direct expert access and lets customers retain data on their own platforms. **AI Productivity Index (APEX)** benchmarks measure professional-task performance. **Mercor Enterprise** uses the same expert network and evaluation infrastructure to deploy custom agents. **Deeptune** adds sandboxed, simulated reinforcement learning, training agents in dynamic environments with tools and interfaces.
- **Listed company relationships: NVIDIA** (NASDAQ: NVDA) – Reported customer; **Innodata** (NASDAQ: INOD) – Competitor
- **Key near-term milestones:** N/A

Lunch with Scott Wu, Co-Founder & CEO of Cognition

Cognition

Company Snapshot

Cognition	
Description	Autonomous AI software-engineering agents
Founded	2024
Last raise / valuation	~$2b Series E (September 2026); ~$48b
Key investors	Andreessen Horowitz, Accel (leads); Founders Fund, General Catalyst
Operating scale	~$1b ARR (September 2026)

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

Source: Company Announcements, Don't call me Devin: Announcing our Series E, Cognition (Aug 9, 2026)

- **Company description:** Cognition builds autonomous AI software engineers for enterprise engineering teams.
- **Product/service verticals: Devin Cloud** is its flagship product that can autonomously plan, write, test and ship production code inside the customer's existing codebase. Related Devin interfaces include **Desktop** for IDE interfacing, **CLI** for command-line interfacing, and **Review** for code-review. **Fusion** is an execution architecture inside Devin.
- **Listed company relationships: Amazon** (NASDAQ: AMZN) – Strategic distribution partner; **Cognizant** (NASDAQ: CTSH) – Partner and Devin customer; **Microsoft** (NASDAQ: MSFT) - Competitor via GitHub Copilot's cloud agent
- **Key near-term milestones:** N/A

Building More Efficient AI Infrastructure From the Data Center to the Edge

Liquid AI

Company Snapshot

Liquid AI	
Description	Compact on-device AI models
Founded	2023
Last raise / valuation	~$250m Series A (December 2024), ~$2b
Key investors	AMD (NASDAQ: AMD) (lead)
Operating scale	~48.7m model downloads

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

Source: Company Announcements, 'Liquid AI just raised $250M to develop a more efficient type of AI model' TechCrunch (Dec. 13, 2024)

- **Company description:** Founded as an MIT spinout, Liquid AI develops compact AI foundation models for automakers, device manufacturers and enterprises. Its models run locally on devices, addressing cloud latency, privacy and power constraints.
- **Product/service verticals: Liquid Foundation Models (LFMs)**, currently LFM2.5, process text, images and speech on-device; the **Liquid Edge AI Platform (LEAP)** customizes and embeds them in applications. **Liquid Context** supplies permissioned user memory to Liquid Agent, a task-executing assistant offered for manufacturer evaluation.
- **Listed company relationships: Shopify** (NASDAQ: SHOP) – Customer under a multi-year, multimillion-dollar license, with search in production by November 2025; **Mercedes-Benz Group** (XETRA: MBG) – Partner for multi-year embedded-speech for selected North American vehicles; **Qualcomm** (NASDAQ: QCOM) – Partner optimizing Context and Agent for Snapdragon Hexagon; **AMD** (NASDAQ: AMD) – Series A lead investor and partner optimizing LFMs for Ryzen AI processors.
- **Key near-term milestones: 1)** Mercedes-Benz in-vehicle deployment targeting 2H26; **2)** LFM integration into MacPaw's Eney assistant for macOS targeting later in 2026.

SiMa.ai

Company Snapshot

SiMa.ai	
Description	Edge AI chips and software
Founded	2018
Last raise / valuation	~$150m Series C (September 2026), ~$1.45b
Key investors	Fidelity, Amplify Partners (co-leads); Dell Technologies (NYSE: DELL) via Dell Technologies Capital
Operating scale	~150 customers (September 2026)

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

Source: Company Announcements, SiMa.ai Reaches $1.45B Valuation with $500 Million in Total Funding to Scale Physical AI in Humanoids, Automotive and Drones' SiMa.ai (Sep 28, 2026)

- **Company description:** SiMa.ai supplies chips and software modules to robotics, drone, automotive and industrial-equipment manufacturers. It runs AI models locally, reducing cloud dependence within device power and latency constraints.
- **Product/service verticals: Modalix machine-learning system-on-chip (MLSoC)** runs vision and language workloads; **Palette** prepares and deploys models; open-source **Palette Neat** adds natural-language application development. Third-generation accelerator IP can be integrated into high-volume customers' chips.

- **Listed company relationships: Synopsys** (NASDAQ: SNPS) – Partner on an automotive chip-design blueprint; **Micron Technology** (NASDAQ: MU) – Supplies low-power memory for Modalix modules; **TSMC** (NYSE: TSM) – Manufactures Modalix on its 6-nanometer process; **NVIDIA** (NASDAQ: NVDA) – Competes through Jetson Orin NX and Nano in embedded AI; SiMa modules are compatible with existing carrier boards.
- **Key near-term milestones: 1)** Third-generation hardware targeting 1H28 focusing on machine learning IP, chiplets, and system-on-chips (SoCs).

d-Matrix

Company Snapshot

D-Matrix	
Description	In-memory AI inference chips
Founded	2019
Last raise / valuation	~$275m Series C (November 2025). ~$2b
Key investors	BullhoundCapital. Triatomic Capital. Temasek (co-leads). Microsoft (NASDAQ: MSFT) via M12
Operating scale	~250+ employees (November 2025)

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

Source: Company Announcements, d-Matrix Raises $275 Million to Power the Age of AI Inference", D-Matrix (Nov. 12, 2025)

- **Company description**: d-Matrix supplies AI inference hardware to cloud providers and enterprises. Its chips compute within memory, designed to reduce data movement, response latency and power consumption when running trained models.
- **Product/service verticals**: **Corsair** is its flagship AI inference accelerator platform, using memory-centric compute to run the inference stage of pre-trained AI models at low latency. **JetStream** networking connects accelerators across servers; **Aviator** software compiles and executes models; **SquadRack** integrates all three into rack systems. The acquisition of **Wallaroo.ai** in August 2026 adds a deployment orchestration software.
- **Listed company relationships**: **NVIDIA** (NASDAQ: NVDA) – Partnered to integrate next-generation, stacked-memory Raptor accelerators into NVLink Fusion/MGX racks, targeting 4Q27 availability; **Astera Labs** (NASDAQ: ALAB) – Developing custom connectivity for D-Matrix-integrated NVIDIA racks; **TSMC** (NYSE: TSM) – Manufactures Corsair on its N6 process; **Alchip** (TWSE: 3661) – Corsair design/ manufacturing partner and Raptor co-developer.
- **Key near-term milestones: 1)** Raptor tape-out–final design handoff for fabrication targeted by year-end 2026, ahead of planned NVIDIA rack integration.

A Conversation with Amjad Masad, Co-Founder & CEO of Replit

Replit

Company Snapshot

Replit	
Description	AI app-building platform
Founded	2016
Last raise / valuation	~$400m Series D (March 2026). ~$9b
Key investors	Georgian (lead). Andreessen Horowitz. Coatue
Operating scale	~60m users (September 2026). ~$150m ARR (September 2025)

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

Source: Company Announcements, "The Future is Actually Very Human", Replit (Mar. 11, 2026)

- **Company description:** Replit provides AI software-creation tools for entrepreneurs and enterprise teams. Its platform combines coding and deployment, reducing manual setup for prototypes, internal tools and customer-facing applications.
- **Product/service verticals: Agent 4** writes, tests and deploys applications; **Replit Design** generates interactives mockups and live apps. **Database** is a fully-managed SQL database that can be added to Replit App to provide persistent data storage.
- **Listed company relationships: Alphabet** (NASDAQ: GOOGL) – Gemini-model supplier; **Microsoft** (NASDAQ: MSFT) – Partner via integrating with Azure and Rayfin and deploying Replit apps into Fabric; **Visa** (NYSE: V) – Customer; **Wix** (NASDAQ: WIX) – Competitor via Base44.
- **Key near-term milestones: 1)** New London office headcount could reach 100 by early 2027 (vs. ~20 in September 2026).

From the Rack to the Real World

Lightmatter

Company Snapshot

Lightmatter	
Description	Photonic interconnects for AI chips
Founded	2017
Last raise / valuation	~$400m Series D (October 2024), ~$4.4b
Key investors	T. Rowe Price-advised funds (lead), Fidelity, GV, Lockheed Martin (NYSE: LMT), HPE (NYSE: HPE)
Operating scale	~310 employees (March 2026), evaluation platforms shipping

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

Source: Company Announcements, Lightmatter Raises $400M Series D, Quadruples Valuation to $4.4B as Photonics Leader for Next-Gen AI Data Centers, Businesswire (Oct 15, 2024)

- **Company description**: Lightmatter develops photonic (light-based) interconnects for AI chipmakers and hyperscale data centers. Optical links address the bandwidth, power and distance limits of electrical connections, allowing larger processor clusters to communicate efficiently.
- **Product/service verticals**: **Passage L200** stacks optical and electronic chips, spreading connections across chip surfaces rather than only edges; **Passage M1000** is a photonic interposer (light-routing base) beneath multiple chips. Forthcoming **Passage L20** places optics beside processors, with bidirectional links that halve fiber requirements; **Guide** modules supply the laser light.
- **Listed company relationships**: **NVIDIA** (NASDAQ: NVDA) – NVLink Fusion partner (June 2026) adapting Lightmatter optics to connect custom processors to NVIDIA switches; **GlobalFoundries** (NASDAQ: GFS) – Fabrication partner for M1000; **Marvell** (NASDAQ: MRVL) – Competitor in optical interconnects via Celestial AI.
- **Key near-term milestones:** 1) Guide DR liquid-cooled laser-module sampling targeting Q4 2026; 2) Passage L20 CPX evaluation-kit shipments targeting Q1 2027

Upscale AI

Company Snapshot

Upscale AI	
Description	AI data-center networking silicon
Founded	2025 (launched)
Last raise / valuation	~$190m Series A-1 (June 2026); ~$2b
Key investors	Premji Invest (lead), NVIDIA (NASDAQ: NVDA), Intel (NASDAQ: INTC), Qualcomm (NASDAQ: QCOM)
Operating scale	~200 employees (March 2026); pre-revenue

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

Source: Company Announcements, Upscale Adds $180 Million in Extension to Series A, Reaching Half-Billion Dollars in Total Funding, Upscale AI (Jun. 22, 2026)

- **Company description**: Upscale AI develops networking silicon, systems and software for AI data centers. Its open-standard fabric connects accelerators, memory and storage, targeting network bottlenecks in large, synchronized AI clusters.
- **Product/service verticals**: **SkyHammer** is its scale-up switch silicon, linking accelerators, memory and storage within a rack at low, deterministic latency. Scale-out Ethernet systems built on NVIDIA Spectrum-X switch silicon connect racks and clusters. Both run on an AI-optimized version of SONiC, the open-source network operating system.
- **Listed company relationships**: **NVIDIA** (NASDAQ: NVDA) – Investor (June 2026) and supplier of Spectrum-X silicon for scale-out systems; **Intel** (NASDAQ: INTC) and **Qualcomm** (NASDAQ: QCOM) – Series A investors; **Broadcom** (NASDAQ: AVGO), **Cisco** (NASDAQ: CSCO) and **Arista Networks** (NYSE: ANET) – Incumbent competitors in AI datacenter switching.
- **Key near-term milestones: 1)** Product launch event on October 8, 2026; 2) First product shipments targeting late 2026 to early 2027

A Conversation with Lin Qiao, Co-Founder & CEO of Fireworks AI

Fireworks AI

Company Snapshot

Fireworks AI	
Description	AI model inference and customization cloud
Founded	2022
Last raise / valuation	~$1.5b Series D (July 2026); ~$17.5b
Key investors	Atreides, Index Ventures, TCV (leads), NVIDIA (NASDAQ: NVDA), AMD (NASDAQ: AMD)
Operating scale	~$1b+ ARR; ~40t+ tokens daily (July 2026)

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

Source: Company Announcements, Fireworks Secures $1.5 Billion in Series D Funding, Fireworks AI (Jul. 15, 2026)

- **Company description:** Fireworks AI provides cloud infrastructure to customize and run AI models. Enterprises and developers use it for coding assistants and industry-specific applications without managing underlying compute.
- **Product/ service verticals:** Inference runs models on shared or dedicated GPU capacity, priced per token or GPU hour. Managed Training and the Training API customize models with customer data and feed directly into Inference. Nexus routes coding requests across open and proprietary models, with spending controls.
- **Listed company relationships: Microsoft** (NASDAQ: MSFT) – Distribution partner via Microsoft Foundry, billed through Azure; **NVIDIA** (NASDAQ: NVDA) – Investor

and GPU supplier via cloud providers; **AMD** (NASDAQ: AMD) – October 2025 multiyear agreement for MI325X/MI355X access through cloud providers; **Amazon** (NASDAQ: AMZN) – AWS cloud and co-sell partner, but competes via Bedrock custom-model inference.

- **Key near-term milestones: 1)** 3–4x compute footprint expansion targeting October 2026.

Technology in Public Safety, National Security, and Defense

Dataminr

Company Snapshot

Dataminr	
Description	Real-time AI event and risk intelligence
Founded	2009
Last raise / valuation	~$300m financing (September 2025), valuation undisclosed, last priced: ~$475m Series F (March 2021) at ~$4.1b
Key investors	Fortress, NightDragon (leads, 2025), Eldridge, Valor Equity Partners
Operating scale	~$200m ARR (April 2025), 18 national militaries

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

Source: Company Announcements, Dataminr Announces $100M Investment From Fortress to Accelerate Gen AI and Agentic AI Product Innovation, and to Expand its Reach to Enterprises & Governments Worldwide", Dataminr (Apr 24, 2025)

- **Company description:** Dataminr provides real-time event and risk intelligence to enterprises, governments and newsrooms. Its AI turns public text, imagery and sensor data into alerts, reducing manual monitoring of threats to people, assets and operations.
- **Product/service verticals:** First Alert Advanced and Dataminr Advanced for Corporate Security, launched August and September 2026 respectively, add corroborating sources, context and near-term forecasts for government and corporate security teams. Dataminr for Cyber Defense (built on acquired ThreatConnect) combines external threat intelligence with internal security data; Dataminr for News flags breaking stories.
- **Listed company relationships: Microsoft** (NASDAQ: MSFT) – Integration partner feeding Dataminr alerts into Sentinel security workflows; **Amazon** (NASDAQ: AMZN) – Cloud host and distribution channel via AWS Marketplace; **Mastercard** (NYSE: MA) – Competitor in cyber-threat intelligence via Recorded Future.
- **Key near-term milestones:** N/A

Flock Safety

Company Snapshot

Flock Safety	
Description	Public-safety cameras and software
Founded	2017
Last raise / valuation	~$275m raise (2025), ~$7.5b
Key investors	Round lead undisclosed; Andreessen Horowitz (led March 2025), Tiger Global
Operating scale	~$500m+ ARR (1H26); 12k+ customers (July 2026)

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

Source: Company Announcements, Accelerating Innovation: Flock Secures $275 Million to Advance Crime-Solving Technology" Flock Safety (Mar 13, 2025)

- **Company description:** Flock Safety supplies public-safety cameras, sensors and cloud software to police departments, businesses, schools and neighborhoods. Its platform links vehicle sightings and incident feeds to generate investigative leads for responders.

- **Product/service verticals: License plate recognition (LPR) cameras** capture plates and vehicle attributes; video cameras stream incidents; acoustic sensors locate gunfire; docked drones provide aerial views. **FlockOS** combines feeds on a command-center map; **FreeForm** searches footage using ordinary language. Its cross-agency network extends searches beyond a customer's cameras, subject to customer-controlled sharing permissions.
- **Listed company relationships: Motorola Solutions** (NYSE: MSI) – Partner integrating Avigilon video with Flock plate alerts; **Axon** (NASDAQ: AXON) – Key listed competitor via Outpost LPR cameras and Fusus real-time crime center software; **Amazon** (NASDAQ: AMZN) – Cloud supplier via AWS.
- **Key near-term milestones: 1)** Mandatory search-misuse monitoring (Audit Assistance) and case codes for law-enforcement searches targeting year-end 2026

Vantor

Company Snapshot

Vantor	
Description	Satellite imagery and geospatial intelligence
Founded	1992 (as DigitalGlobe)
Last raise / valuation	No standalone raise. Maxar take-private (May 2023) at ~$6.4b EV
Key investors	Advent International (owner), British Columbia Investment Management (minority)
Operating scale	3.5m+ sq km daily imaging capacity, 1.2m+ platform users

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

Source: Company Announcements, Maxar Technologies To Be Acquired by Advent International for $6.4 Billion", Businesswire (Dec 16, 2022)

- **Company description:** Formerly Maxar Intelligence, Vantor provides satellite imagery, 3D mapping and AI-enabled geospatial software to defense, intelligence and commercial customers.
- **Product/service verticals: WorldView** satellites collect imagery on demand; **Vivid** supplies maps and three-dimensional (3D) terrain. **Tensorglobe** fuses sensor feeds using AI and powers Sentry automated site monitoring; **Raptor** matches drone imagery to 3D terrain for GPS-denied navigation.
- **Listed company relationships: BAE Systems** (LSE: BA.) – Selected June 2026 to build next-generation Vantage satellites; **CACI** (NYSE: CACI) – Selected September 2026 to supply Pulse optical payloads; **Alphabet** (NASDAQ: GOOGL) – Partner integrating Google Earth AI into Tensorglobe, including for classified environments; **Planet Labs** (NYSE: PL) – Competitor in high-resolution satellite imaging and site monitoring through SkySat/Pelican.
- **Key near-term milestones: 1)** Initial Pulse 40cm-class imaging satellite deployments targeting 2027–2029; **2)** Vantage 20cm-class satellites entering service before 2030.

Founder Breakpoints: Navigating the New Fundraising Playbook

Varda Space Industries

Company Snapshot

Varda Space Industries	
Description	In-orbit pharmaceutical manufacturing and reentry
Founded	2021
Last raise / valuation	~$251m Series D (September 2026); ~$1.6b
Key investors	Lux Capital, Natural Capital (leads), Founders Fund, Khosla Ventures, General Catalyst
Operating scale	Six missions flown, 28 launches booked through 2029

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

Source: Company Announcements, 'Varda Announces $250 Million Series D to Scale Space-Based Pharmaceutical Processing', PR Newswire (Sep 30, 2026)

- **Company description:** Varda builds recoverable spacecraft that manufacture pharmaceuticals in microgravity and return them to Earth. The same capsules provide hypersonic reentry testing for US defense customers.
- **Product/service verticals:** W-Series spacecraft crystallize drug compounds in orbit and return them in detachable capsules, without a crewed station. Pharmaceutical services span ground screening, orbital processing and analysis, with future drug royalties targeted. Government contracts use ~Mach 25 reentries to test sensors and heat shields; government work is ~70% of 2027 payload demand.
- **Listed company relationships: United Therapeutics** (NASDAQ: UTHR) – Multi-mission pulmonary-drug formulation partner, announced May 2026; **SpaceX** (NASDAQ: SPCX) – Launch provider (W-6 on Transporter-16, March 2026) and competitor via its Starfall reentry program; **Rocket Lab** (NASDAQ: RKLB) – Former spacecraft-bus supplier, since replaced by Varda's in-house design; **Redwire** (NYSE: RDW) – Competitor in orbital drug crystallization via SpaceMD.
- **Key near-term milestones: 1)** Nine cumulative commercial reentries targeting end-2026; **2)** First human clinical dosing of a microgravity-enabled formulation targeting ~2030.

Stord

Company Snapshot

Stord	
Description	Outsourced e-commerce fulfillment and software
Founded	2015
Last raise / valuation	~$250m Series F (May 2026); ~$3b
Key investors	Existing investors (no single lead), Kleiner Perkins, Founders Fund, Franklin Templeton (NYSE: BEN)
Operating scale	1,000+ brands; ~$15b+ annual GMV; ~100 fulfillment sites (May 2026)

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

Source: Company Announcements, 'Stord Raises $250M Series F at $3B to Advance the Physical Intelligence Layer for Commerce', Stord (May 26, 2026)

- **Company description:** Stord provides outsourced fulfillment and commerce software to consumer brands selling online and through retailers. Its shared order-and-inventory platform connects warehouse operations and sales channels, reducing coordination across disconnected logistics providers and systems.
- **Product/service verticals:** Fulfillment covers storage, picking, packing and shipping through third-party carriers. Its order management system (OMS) routes orders to fulfillment sites; warehouse management system (WMS) manages inventory and warehouse workflows; Consumer Experience software adds delivery estimates, tracking and returns. Software integrates with Stord's physical network

but can also run customers' own warehouses. StordAI Chat, launched in beta in March 2026, uses artificial intelligence (AI) to answer order and inventory questions. Revenue combines logistics fees and software subscriptions.

- **Listed company relationships: UPS** (NYSE: UPS) – Fulfillment partner since Stord acquired Ware2Go from UPS (May 2025); **Procter & Gamble** (NYSE: PG) – Fulfillment customer through Native; **Amazon** (NASDAQ: AMZN) – Competitor via Multi-Channel Fulfillment for off-Amazon orders.
- **Key near-term milestones:** N/A

Ayar Labs

Company Snapshot

Ayar Labs	
Description	Optical interconnect chiplets for AI
Founded	2015
Last raise / valuation	~$500m Series E (March 2026). ~$3.75b. ~$150m extension (September 2026) at >$5b (reported)
Key investors	Neuberger Berman (lead). NVIDIA (NASDAQ: NVDA). AMD (NASDAQ: AMD). MediaTek (TWSE: 2454). Alchip (TWSE: 3661). Wiwynn (TWSE: 6669). Intel (NASDAQ: INTC) via Intel Capital
Operating scale	~200 employees

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

Source: Company Announcements, Ayar Labs Closes $500M Series E, Accelerates Volume Production of Co-Packaged Optics' Ayar Labs (Mar. 3, 2026)

- **Company description:** Ayar Labs develops optical interconnects for AI chipmakers. Its silicon-photonics chiplets move data between processors using light rather than copper, addressing bandwidth and power bottlenecks in large-scale AI training and inference.
- **Product/ service verticals:** Ayar offers two complementary hardware products. **TeraPHY** is a silicon-photonics chiplet designed to bring optical connectivity directly into the processor package, converting electrical data into light signals for high-bandwidth communication across servers and racks. **SuperNova** is an external laser source that supplies light to TeraPHY.
- **Listed company relationships: NVIDIA** (NASDAQ: NVDA) – Strategic investor and partner integrating Ayar's co-packaged optics (CPO) into NVIDIA rack-scale architecture; **AMD** (NASDAQ: AMD), **MediaTek** (TWSE: 2454) and **Alchip** (TWSE: 3661) – Strategic investors; **Lumentum** (NASDAQ: LITE) – External laser supplier under a 2022 agreement; **Marvell** (NASDAQ: MRVL) – Competitor via Celestial AI.
- **Key near-term milestones: 1)** Optical component qualification for volume production targeting end-2027, ahead of customer product ramps in 2028–2029.

The Machines Reshaping the Physical World

Nuro

Company Snapshot

Nuro	
Description	Licensed autonomous-driving technology
Founded	2016
Last raise / valuation	~$203m Series E (August 2025); ~$6b
Key investors	Uber (NYSE: UBER); NVIDIA (NASDAQ: NVDA); Baillie Gifford; Fidelity; T. Rowe Price; Tiger Global
Operating scale	~100 test vehicles in Bay Area and Houston (August 2026)

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

Source: Company Announcements, Nuro Closes $203 Million Series E Financing to Advance Its AI-First Self-Driving Technology and Commercial Partnerships" Nuro (Aug 27, 2025)

- **Company description:** Nuro licenses autonomous-driving technology to automakers and mobility operators for robotaxis, delivery fleets and consumer vehicles. Having shifted to licensing in 2024, it lets partners deploy driving automation without building the full system in-house.
- **Product/service verticals:** Nuro Driver combines AI software, sensors and onboard compute for driverless (Level 4) operation within defined conditions. Driver Assist targets hands-off, eyes-on consumer driving; Nuro Toolkit adds simulation, validation and remote-fleet tools.
- **Listed company relationships: Uber** (NYSE: UBER) – Investor (~$500m committed including milestone payments) and exclusive ride-hailing and fleet-operations partner for the planned robotaxi service; **Lucid** (NASDAQ: LCID) – Integrates Nuro into Gravity and future Midsize vehicles under Uber's April 2026 purchase commitment for at least 35,000 vehicles; **NVIDIA** (NASDAQ: NVDA) – Investor and supplier of DRIVE AGX Thor onboard compute.
- **Key near-term milestones: 1)** Bay Area robotaxi launch with Uber targeting late 2026; **2)** Houston launch targeting mid-2027

Hadrian

Company Snapshot

Hadrian	
Description	Automated aerospace and defense manufacturing
Founded	2020
Last raise / valuation	~$1.37b Series D (August 2026); ~$7.87b
Key investors	WCM; Washington Harbour; Valor Equity Partners; 137 Ventures; Baillie Gifford (leads); RTX (NYSE: RTX) via RTX Ventures
Operating scale	~3m sq ft capacity across four facilities

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

Source: Company Announcements, Hadrian Raises $1.5/B Series D to Build Highly Automated Factories to Accelerate America's Industrial Renewal, PR Newswire (Aug 6, 2026)

- **Company description:** Hadrian is a software-enabled contract manufacturer and factory operator that produces precision metal components to specifications for aerospace and defense customers. Its proprietary Opus software coordinates automated machining, production scheduling and inspection. Customers purchase finished parts or engage Hadrian to build and operate dedicated production lines or entire factories, including within their own facilities.
- **Product/ service verticals: Precision Components**: Manufacture precision parts, including components produced using computer numerical control (CNC) machining, for aerospace and defense customers from a product's prototype to production stage. **Manufacturing-as-a-Service**: Provide dedicated manufacturing

capacity or inspection cells, deployed either in Hadrian facilities or customer sites, supporting customer-specific production requirements. **Factories-as-a-Service**: Designs, builds and operates dedicated production infrastructure for a customer or program, to service entire production cycle

- **Listed company relationships: RTX** (NYSE: RTX) purchases components through Collins Aerospace and Raytheon, including Javelin missile parts, and is a strategic investor through RTX Ventures. **Lockheed Martin** (NYSE: LMT) partners under a December 2025 MOU to embed machining and inspection cells inside LMT's Missiles and Fire Control facility in Grand Prairie, Texas.
- **Key near-term milestones: 1)** Progress on additional factories and munitions/ autonomous-system production lines targeting August 2027; **2)** Fourth Alabama Factory phase-one targeting full-rate production within 24 months of contract award.

Genesis AI

Company Snapshot

Genesis AI	
Description	General-purpose robots and robot AI
Founded	2024
Last raise / valuation	~$105m seed (July 2025)
Key investors	Eclipse, Khosla Ventures (co-leads)
Operating scale	Dozens of robots built, ~60 employees (June 2026)

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

Source: Company Announcements, Genesis AI launches with $105M seed funding from Eclipse, Khosla to build AI models for robots', TechCrunch (Jul 1, 2025)

- **Company description:** Genesis AI develops general-purpose robots and robotics foundation models for manufacturing, logistics and laboratories. Its integrated hardware and learned control target variable, dexterous tasks that fixed automation cannot handle.
- **Product/service verticals: Eno,** unveiled June 2026, pairs a wheeled, height-adjustable body and human-like dexterous hands with GENE, its AI model for planning and controlling multistep tasks. Its data engine uses glove-recorded human actions and physics simulation to train and test models, aiming to reduce costly robot demonstrations.
- **Listed company relationships: LG CNS** (KRX: 064400) – Partner evaluating Eno for manufacturing and logistics at LG's US operations (June 2026); **NVIDIA** (NASDAQ: NVDA) – Competing robot-training stack via Isaac GR00T models and Isaac Sim; **Alphabet** (NASDAQ: GOOGL) – Competitor in robot foundation models via Google DeepMind's Gemini Robotics.
- **Key near-term milestones: 1)** Eno production start and initial customer deployments targeting end-2026.

A Conversation with Brandon Tseng, Co-Founder & President of Shield AI

Shield AI

Company Snapshot

Shield AI	
Description	AI pilots and autonomous aircraft
Founded	2015
Last raise / valuation	~$1.5b Series G (June 2026), ~$12.7b post-money
Key investors	Advent (lead), JPMorgan Chase (NYSE: JPM) (co-lead), Blackstone (NYSE: BX) funds (preferred equity)
Operating scale	Not disclosed

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

Source: Company Announcements, Shield AI to Acquire Software Simulation Company Aechelon and Raise $2B at $12.7B Valuation, Shield AI (Mar 26, 2026)

- **Company description:** Shield AI develops AI pilot software and unmanned aircraft systems (UAS) for US and allied militaries to operate aircraft with limited human intervention.
- **Product/ service verticals: Hivemind:** Flagship AI-powered pilot software that allows for centralized control of a wide range of unmanned systems to execute complex missions in disrupted and GPS/ comms-denied environments. **V-BAT:** Combat-proven unmanned aircraft system (UAS) designed for contested GPS/ comms-denied environments. It offers unassisted vertical take-off (VTOL) technology and swarming capabilities. **X-BAT:** Development-stage AI-piloted, vertical take-off and landing (VTOL) stealth fighter jet designed to operate without a traditional runway. It can carry air-to-air, air-to-surface weapons, and large strike weapons. **Vision Systems**: Shield AI's computer-vision hardware and software portfolio where its Visual Detection and Ranging (ViDAR) electro-optical and infrared (EO/IR) sensors and Tracker software automate detection and tracking for surveillance and counter-drone missions. **Aechelon:** Simulation software for pilot training and autonomous-system testing; acquired June 22, 2026.

- **Listed company relationships:** Competes with **AeroVironment** (NASDAQ: AVAV) and **Textron** (NYSE: TXT) in tactical unmanned aircraft/ISR, and with **Boeing** (NYSE: BA)/**Insitu**, **RTX** (NYSE: RTX)/**Collins**, **Northrop Grumman** (NYSE: NOC) and **Lockheed Martin** (NYSE: LMT) in mission-autonomy software. **L3Harris** (NYSE: LHX) licenses Shield's Tracker counter-drone software and partners on electronic-warfare autonomy; **RTX** (NYSE: RTX), **Airbus** (Euronext Paris: AIR) and **Northrop Grumman** (NYSE: NOC) integrate Hivemind into weapons, sensors or aircraft, although **RTX** (NYSE: RTX) and **Northrop Grumman** (NYSE: NOC) also field competing autonomy software. **Palantir** (NASDAQ: PLTR) partners on command-and-control integration and supplies manufacturing software, while **GE Aerospace** (NYSE: GE) supplies X-BAT's F110 propulsion system. Shield also works with **Mitsubishi Heavy Industries** (TSE: 7011) and **Korea Aerospace Industries** (KRX: 047810) to develop sovereign AI-pilot capabilities using Hivemind.
- **Key near-term milestones: 1)** X-BAT's first vertical take-off and landing (VTOL) flights targeted for 2026, production for 2029; **2)** Polish Navy V-BAT deliveries targeted for later in 2026; **3)** Cooperative multi-aircraft operations with Republic of Singapore Air Force using Hivemind targeted for 2027 following manned/ unmanned teaming in 2026.

The Financial Agent Economy: Building the Infrastructure for Agentic Finance & Commerce

MoonPay

Company Snapshot

MoonPay	
Description	Crypto payments and on/off-ramp infrastructure
Founded	2019
Last raise / valuation	~$555m Series A (November 2021). ~$3.4b; reported talks at ~$5b with ICE (December 2025)
Key investors	Coatue, Tiger Global, Intercontinental Exchange (NYSE: ICE) (reported talks); Galaxy (March 2025 debt)
Operating scale	~$108m revenue (2021); 160+ countries; ~440 employees (May 2026)

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

Source: Company Announcements, "MoonPay closes historic funding round", MoonPay (May 11, 2022); "NYSE Owner Is Said in Talks to Invest in Crypto Firm MoonPay", Bloomberg (Dec. 18, 2025).

- **Company description:** MoonPay supplies crypto-payment infrastructure to wallets, fintechs and merchants. It connects bank/card payments with blockchains, bundling identity checks, liquidity and settlement so partners need not build these capabilities separately.
- **Product/service verticals: Ramps** embeds cryptocurrency purchases and cash-outs within apps; **Trade** routes token exchanges across blockchains. **Commerce** provides merchant checkout with crypto or traditional-currency settlement. **Enterprise**, launched August 2026, converts bank deposits into stablecoins and supports cross-border payouts and branded stablecoin issuance. Shared infrastructure reduces vendor integration; availability remains jurisdiction-dependent. **Institutional** supplies wallet-security and trading infrastructure to financial firms.
- **Listed company relationships: PayPal** (NASDAQ: PYPL) – Payment partner enabling purchases and cash-outs in eligible markets; **Paysafe** (NYSE: PSFE) – Customer using Commerce and Virtual Accounts for US online-gaming payments, launched April 2026, with upgraded checkout live September 2026; **Coinbase** (NASDAQ: COIN) – Competitor serving wallet/app developers through Onramp and Offramp.
- **Key near-term milestones: 1)** Rollout of Ether.fi partnership in next few months with Ether.fi moving core payments infrastructure onto MoonPay's integrated stack

Mercury

Company Snapshot

Mercury	
Description	Banking and finance software for startups
Founded	2017
Last raise / valuation	~$200m Series D (May 2026). ~$5.2b
Key investors	TCV (lead), Sequoia Capital, Andreessen Horowitz, Coatue
Operating scale	300k+ customers (September 2026). ~$650m annualized revenue (September 2025)

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

Source: Company Announcements, "The 'future of banking' Today, we're announcing Mercury's $1.00M Series D at a $5.2B valuation, led by TCV", Mercury (May 20, 2026).

- **Company description:** Mercury provides banking, payments and financial software to startups and small businesses. It combines accounts, cards, bill pay and accounting in one platform, reducing reconciliation across separate systems.

- **Product/service verticals:** Free business accounts connect IO corporate cards, bill payments, invoicing and expense controls; Treasury invests surplus cash in short-term funds. Venture debt and working-capital loans finance startups and ecommerce businesses, respectively. Books, launched September 2026, categorizes and reconciles transactions; Command uses artificial intelligence (AI) for customer-approved financial tasks.
- **Listed company relationships: State Street** (NYSE: STT) – Partner supplying an exclusive government money-market fund share class; **Intuit** (NASDAQ: INTU) – Competitor via QuickBooks, also integrated with Mercury; **Capital One** (NYSE: COF) – Competitor in corporate cards and spend management through Brex.
- **Key near-term milestones: 1)** Embedded payroll via acquired Central targeting 2026; **2)** Treasury Ladders (individual US Treasury portfolios) launch targeting 2026.

How AI Agents Are Expanding What's Possible Across the Enterprise

EliseAI

Company Snapshot

EliseAI	
Description	Vertical AI agents for housing and healthcare operations
Founded	2017
Last raise / valuation	~$350m (September 2026) ~$4b
Key investors	Andreessen Horowitz Bessemer Venture Partners. Ontario Teachers' Pension Plan (co-leads) Sapphire Ventures. Navitas Capital
Operating scale	$200mm+ ARR (2026) 1 in 6 apartments in the U.S. (2026). 30M unique renters reached since launch (Sep 2026)

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

Source: Company Announcements,'EliseAI Raises $350 Million at $4 Billion Valuation to Bring AI Deeper Into Housing and Healthcare Operations', GlobeNewswire (Sep 29, 2026), 'EliseAI hits $200m ARR After Five Consecutive Years of Growth Across Housing and Healthcare', EliseAI (Jun 10, 2026)

- **Company description:** EliseAI is a New York company founded in 2017 by CEO Minna Song and CTO Tony Stoyanov. Its vertical AI agents automate leasing, resident, maintenance and patient-access work for rental-housing operators and specialty physician groups through integrations with customers' workflow systems.
- **Product/service verticals: Housing** combines LeasingAI, ResidentAI, VoiceAI and EliseCRM, a customer relationship management (CRM) hub: it answers renters across phone, text, email and webchat, books tours, routes work orders, manages renewals and collections, and centralizes records. Apollo, launched in September 2026, is a prompt-based agent that executes permissioned tasks across the platform. **Healthcare** automates calls, scheduling, intake, referrals, insurance verification, chart preparation and follow-up.
- **Listed company relationships: Zillow Group** (NASDAQ: Z, ZG) – partner on AI Assist embedded in Zillow Rentals listings; **Vivmark Residential** (NYSE: VMRK) – successor to Equity Residential and AvalonBay, which EliseAI named as customers in August 2024; post-merger contract status is undisclosed; **AppFolio** (NASDAQ: APPF) – direct competitor in AI leasing, maintenance and resident operations through Realm-X; **Phreesia** (NYSE: PHR) – competitor in healthcare voice, scheduling, intake and payment workflows.
- **Key near-term milestones:** N/A

Genspark

Company Snapshot

Genspark	
Description	Agentic AI workspace for knowledge workers
Founded	2023
Last raise / valuation	~$100m Series B extension (June 2026), ~$2.6b post-money
Key investors	Emergence Capital (led March 2026 extension), Sozo Ventures, Mirae Asset, UpHonest Capital (extension), LG Technology Ventures, SBI, Pavilion Capital
Operating scale	~$250m ARR (June 2026)

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

Source: Company Announcements,

- **Company description:** Genspark, a wholly owned subsidiary of MainFunc, is an agentic AI workspace for knowledge workers and enterprises. Users state an objective; the service gathers context, routes work across models and tools, then returns editable deliverables or executes actions, reducing manual switching among research, office creation and business applications.
- **Product/service verticals: SecondBrain** retains work context; **Super Agent's** Mixture-of-Agents routes subtasks across models; **Office, Build and Content suites** create documents, spreadsheets, presentations, code, databases and media; **GenTeam** and **Claw** add human-agent collaboration and cloud/desktop execution. September 2026's Gen-1 Slides is Genspark's first model further trained for presentation work.
- **Listed company relationships: Microsoft** (NASDAQ: MSFT) – Agent 365 launch partner with Word/Excel/PowerPoint integration; **Amazon** (NASDAQ: AMZN) – preferred cloud and Bedrock supplier for agents; **Twilio** (NYSE: TWLO) – Programmable Voice supplier for Call for Me in 40+ countries; **Alphabet** (NASDAQ: GOOGL) – Genspark incorporates Gemini models and is listed on Google Cloud Marketplace; Gemini Enterprise is a potential workplace-agent substitute.
- **Key near-term milestones:** N/A

Cresta

Company Snapshot

Cresta	
Description	AI platform for contact centers
Founded	2017
Last raise / valuation	~$125m Series D (November 2024), ~$1.6b valuation in Series C 2022 raise
Key investors	WiL, QIA (leads), Accenture, EnvisionX, LG Technology Ventures, Qualcomm (NASDAQ: QCOM) via Qualcomm Ventures, Workday (NASDAQ: WDAY) via Workday Ventures; Andreessen Horowitz, Greylock, J.P. Morgan (NYSE: JPM), Sequoia Capital, Tiger Global
Operating scale	~$100m ARR (April 2026, Sacra estimate)

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

Source: Company Announcements, 'Cresta Closes $125M Series D to Accelerate Adoption of Human-Centric AI in the Contact Center', PR Newswire (Nov. 19, 2024), 'Cresta Closes $125M Series D to Accelerate Adoption of Human-Centric AI in the Contact Center', PR Newswire (Nov. 19, 2024), 'Cresta Raises Series C, Quadruples Valuation to $1.6B on its Mission to Revolutionize the Contact Center with Real-Time Intelligence', PR Newswire (Mar. 17, 2022).

- **Company description:** Founded at Stanford's AI Lab, Cresta sells AI software to large travel, financial-services, telecommunications, healthcare and retail contact centers. Its platform automates suitable voice and digital contacts, guides human agents on complex interactions and analyzes conversation data, consolidating automation, agent assistance and quality management on one system.
- **Product/service verticals: AI Agent** handles voice and digital conversations; **Agent Assist** supplies live guidance, source-backed knowledge and summaries;

Conversation Intelligence analyzes 100% of interactions for quality, coaching and automation. **Opera** lets customers configure workflows without code; **Conductor** helps developers build, test and optimize agents using customer-conversation data.

- **Listed company relationships: United Airlines Holdings** (NASDAQ: UAL) – customer using AI Agent, Agent Assist and Conversation Intelligence; **TELUS Corporation** (TSX: T; NYSE: TU) – its wholly owned TELUS Digital unit is a preferred implementation and managed-services partner; **Amazon.com** (NASDAQ: AMZN) – cloud infrastructure and marketplace provider for Cresta's SaaS offer; **NiCE** (NASDAQ: NICE) – direct competitor through CXone's AI agents, Copilot and interaction analytics.
- **Key near-term milestones:** N/A

The Freight and Logistics Infrastructure Behind Modern Commerce

Flock Freight

Company Snapshot

Flock Freight	
Description	Shared-truckload freight brokerage
Founded	2015
Last raise / valuation	~$60m Series E (May 2025); valuation undisclosed. >$1b at Series D (January 2022)
Key investors	O'Neil Strategic Capital (lead); SoftBank Group (TSE: 9984) via Vision Fund 2 (Series D lead)
Operating scale	Thousands of shippers and carrier partners (November 2025)

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

Source: Company Announcements, Flock Freight Secures $60 Million Series E Funding to Continue Scaling Shared Truckload", Flock Freight (Mar 14, 2025)

- **Company description:** Flock Freight is a US freight broker for retail, food-and-beverage and industrial shippers. Its patented shared truckload (STL) software pools compatible shipments onto one truck, so shippers pay only for space used and avoid terminal transfers.
- **Product/service verticals:** FlockDirect sells STL capacity at spot or contracted rates based on pallets or linear feet; conventional full-truckload and less-than-truckload (LTL) services share quoting, booking and tracking tools. STL AddOns fills carriers' remaining truck space with compatible loads. Transport is procured from third-party carriers.
- **Listed company relationships: Oracle** (NYSE: ORCL) – Integration partner for FlockDirect quotes within Oracle Transportation Management; **WiseTech Global** (ASX: WTC) – Integration partner via e2open; **Academy Sports + Outdoors** (NASDAQ: ASO) – STL customer; **C.H. Robinson** (NASDAQ: CHRW) – Competitor in US truckload and LTL brokerage.
- **Key near-term milestones:** N/A

Pallet

Company Snapshot

Pallet	
Description	AI agents for logistics operations
Founded	2022
Last raise / valuation	~$27m Series B (May 2025), valuation undisclosed
Key investors	General Catalyst (lead); Bain Capital Ventures, Activant, Bessemer
Operating scale	~2m annual agent runs at Lineage (single customer)

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

Source: Company Announcements, "We've Raised $27M to Automate the Backbone of Global Supply Chain", Pallet (May 27, 2025)

- **Company description:** Pallet provides AI agents for logistics providers, manufacturers, distributors and retailers. Its agents read emails and documents and execute tasks in existing operating systems, reducing manual coordination without requiring replacement of core software.
- **Product/service verticals: CoPallet** agents automate quoting, order entry, tracking and billing; Atlas combines customer-system data with agent activity, letting operators ask plain-English questions about routes, capacity and customer opportunities. **Forge** generates and tests agents; **Memory** supplies customer-specific operating rules.
- **Listed company relationships: Lineage** (NASDAQ: LINE) – Customer automating receiving, freight-forwarding and warehouse orders; **WiseTech Global** (ASX: WTC) – CargoWise integration used to create orders in Lineage's system; **Manhattan Associates** (NASDAQ: MANH) – Competitor via embedded Manhattan Active agents.
- **Key near-term milestones: 1)** 3m+ Lineage shipments processed in 2026, with expansion across US and international sites.

AtoB

Company Snapshot

AtoB	
Description	Fuel cards and trucking payments
Founded	2019
Last raise / valuation	~$130m Series C, equity and debt (September 2024), valuation undisclosed
Key investors	General Catalyst, Bloomberg Beta (leads); Mastercard (NYSE: MA)
Operating scale	35k+ fleets (undated)

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

Source: Company Announcements, AtoB Announces New Funding and Growth Milestones" Businesswire (Sep 19, 2024)

- **Company description:** AtoB provides fuel cards, payments and working-capital tools to trucking companies and fleets. Its cards and payment accounts help control fuel spending and fund immediate operating costs before freight invoices are paid.
- **Product/service verticals:** Flex fuel card provides weekly credit; Unlimited is prefunded. Both offer fuel discounts, transaction reporting and dashboard spending controls. Business Expense cards finance broader operating expenses in weekly installments; the digital wallet receives freight payments and supports round-the-clock driver payouts. Vehicle-location and fuel-level checks block suspicious purchases.
- **Listed company relationships: Mastercard** (NYSE: MA) – Investor and exclusive network for AtoB's commercial card program; **Uber Technologies** (NYSE: UBER) – Partner via Uber Freight's AtoB-powered carrier fuel card; **Murphy USA** (NYSE: MUSA) – Partner via the AtoB-powered Murphy USA Business Fuel Card (2026);

WEX (NYSE: WEX) – Competitor in fleet cards and SecureFuel fraud controls; **Corpay** (NYSE: CPAY) – Competitor in fuel and expense cards through Comdata and Corpay One.

- **Key near-term milestones:** N/A