Goldman Sachs Research

Equity Research

21 August 2026 | 1:24AM PDT

AMERICAS TECHNOLOGY

Takeaways From Our 3rd Annual Silicon Valley AI Field Trip

George K. Tong, CFA
+1(415)249-7421 | george.tong@gs.com
Goldman Sachs & Co. LLC

Gabriela Borges, CFA
+1(212)902-7839 | gabriela.borges@gs.com
Goldman Sachs & Co. LLC

Matthew Martino
+1(212)902-0695 | matt.martino@gs.com
Goldman Sachs & Co. LLC

Sami Nasir, CFA
+1(415)834-7967 | sami.nasir@gs.com
Goldman Sachs & Co. LLC

Alex Lakritz
+1(415)249-7072 | alex.lakritz@gs.com
Goldman Sachs & Co. LLC

Max Gamperl
+1(415)249-7311 | max.gamperl@gs.com
Goldman Sachs & Co. LLC

Selina Zhang
+1(212)357-9979 | selina.zhang@gs.com
Goldman Sachs & Co. LLC

We hosted our 3rd annual Silicon Valley AI Field Trip on 8/18-8/19, featuring AI companies, VCs and researchers from Stanford and UC Berkeley/UCSF. Model capabilities are continuing to improve, agents are progressing from assistance to workflow execution and monetization is expanding beyond seats toward consumption, transactions and outcomes. The strongest businesses combine proprietary data, domain expertise, customer context, workflow ownership, distribution, verification and permission to act. Agentic adoption should advance fastest where outputs are verifiable and responsibility for errors is clear. Open models should gain token volume across routine tasks, and frontier models should retain value where reliability justifies higher costs, supporting continued token and compute growth. AI should also expand cybersecurity demand, and physical AI will require distinct architectures and commercialization strategies. For Business & Information Services, AI should increase the value of differentiated information assets and help platforms capture more customer spending through embedded workflows. For Software, select vendors are well positioned to help customers navigate an evolving intelligence curve that segments workflows into frontier vs. non-frontier tokens and enriches models with business-specific context. Our preferred AI-leveraged ideas are IRM, MCO, MSCI, SPGI and TRI in Business & Information Services and MSFT, MDB and RBRK in Software. The trip also reinforced our constructive long-term views on CRWD, NET, PANW and SNOW.

Trip Agenda and Background

We hosted our 3rd annual Silicon Valley AI Field Trip on Tues, 8/18, and Wed, 8/19, focused on major themes, developments and debates in AI. Day 1 stops on the field trip included private company presentations and product demos with the Co-Founder & CEO of Daloopa, VP of Software Engineering from Divergent, Founder & CEO of Clio and Co-Founder & CEO of Corgi, followed by venture capital visits to Lightspeed Venture Partners, Kleiner Perkins and Menlo Ventures, as well as presentations and a group dinner with researchers from Stanford University's Institute for Human-Centered AI (HAI) and Department of Aeronautics & Astronautics. Day 2 stops on the field trip included company presentations and product demos with the CFO of Harvey AI, CEO of Moody's Analytics at Moody's

(MCO), Founder & CEO of Vercel, CFO of ClickHouse, Co-Founder of Simile, a venture capital visit to Index Ventures, and a presentation by a Professor of Computational Precision Health, Statistics & Computer Science at UC Berkeley.

2026 Goldman Sachs Silicon Valley AI Field Trip

Time	Participants	Type	Speakers	Title
Day 1 – Tuesday, August 18, 2026				
8:30 AM – 9:15 AM	Daloopa	Private	Thomas Li	Co-Founder & CEO
9:15 AM – 10:00 AM	Divergent	Private	Sam Miller	VP of Software Engineering
10:00 AM – 10:45 AM	Clio	Private	Jack Newton	Founder & CEO
			John Foreman	Chief Product Officer
10:45 AM – 11:30 AM	Corgi	Private	Nico Laqua	Co-Founder & CEO
12:30 PM – 1:30 PM	Lightspeed Venture Partners	Venture Capital	Guru Chahal	Partner
			Sebastian Duesterhoeft	Partner
1:45 PM – 2:45 PM	Kleiner Perkins	Venture Capital	Aditya Naganath	Partner
3:00 PM – 4:00 PM	Menlo Ventures	Venture Capital	Rama Sekhar	Partner
			Matt Kraning	Partner
4:15 PM – 5:45 PM	Stanford University	University	Kiana Jafari	Postdoctoral Scholar in Aeronautics & Astronautics
			Duncan Eddy	Postdoctoral Scholar in Aeronautics & Astronautics
6:00 PM – 8:30 PM	Stanford University	Group Dinner	Kiana Jafari	Postdoctoral Scholar in Aeronautics & Astronautics
			Duncan Eddy	Postdoctoral Scholar in Aeronautics & Astronautics
Day 2 – Wednesday, August 19, 2026				
8:30 AM – 9:15 AM	Harvey	Private	Alan Ghelberg	CFO
			John LaBarre	General Counsel
			Madeleine Hyde	Legal Engineer
9:15 AM –10:00 AM	Moody's (MCO)	Public	Christina Kosmowski	CEO of Moody's Analytics
			Cihan Biyikoglu	Head of Product & Engineering at RMS
			Shivani Kak	Head of Investor Relations
10:00 AM – 10:45 AM	Vercel	Private	Guillermo Rauch	Founder & CEO
			Marten Abrahamsen	CFO
10:45 AM – 11:30 AM	ClickHouse	Private	Jimmy Sexton	CFO
11:45 AM – 12:45 PM	Simile	Private	Michael Bernstein	Co-Founder
			Madison Fox	Financial Services Vertical Lead
			Danny Wymer	Enterprise GTM Lead
1:15 PM – 2:15 PM	Index Ventures	Venture Capital	Nina Achadjian	Partner
			Nina Gerson	Partner
3:15 PM – 4:15 PM	UC Berkeley and UCSF	University	Adam Yala	Asst Prof of Computational Precision Health, Stats & CS

Source: Goldman Sachs Global Investment Research

The agenda included:

- **Company presentations and product demos** with venture-backed startups to discuss their AI strategies and products. Participants included Thomas Li, Co-Founder & CEO of Daloopa; Sam Miller, VP of Software Engineering at Divergent; Jack Newton, Founder and CEO of Clio; and Nico Laqua, Co-Founder & CEO of Corgi.
- **Venture capital visit to Lightspeed Venture Partners** to discuss the outlook for frontier intelligence, accelerating model capabilities, the division of workloads between open and closed models, enterprise AI monetization and the broadening of value creation across the AI stack, featuring Partners Guru Chahal and Sebastian Duesterhoeft.
- **Venture capital visit to Kleiner Perkins** to discuss the expansion of agentic AI

across knowledge-intensive industries, improving application-layer economics, the sources of durable AI application moats, disruption risks for legacy software and emerging opportunities across inference infrastructure, featuring Partner Aditya Naganath.

- **Venture capital visit to Menlo Ventures** to discuss the role of verifiable outputs and proprietary data in enterprise AI adoption, the evolution of software from systems of record to systems of action, cybersecurity opportunities and risks and the outlook for compute demand, featuring Partners Rama Sekhar and Matt Kraning.
- **Presentations by researchers at Stanford University** on world models for physical AI, the use of uncertainty to improve AI safety and the accountability, verification and trust infrastructure required to deploy enterprise agents, led by Kiana Jafari and Duncan Eddy, Postdoctoral Scholars in the Department of Aeronautics & Astronautics.
- **Company presentations and product demos** with venture-backed startups to discuss their AI strategies and products. Participants included Alan Ghelberg, CFO of Harvey; Guillermo Rauch, Founder & CEO of Vercel; Jimmy Sexton, CFO of ClickHouse; and Michael Bernstein, Co-Founder of Simile.
- **Company presentation and product demo with Moody's** focused on agentic AI products, connected intelligence across risk domains and the use of proprietary data and domain expertise to automate enterprise workflows, featuring Christina Kosmowski, CEO of MA, Cihan Biyikoglu, Head of Product & Engineering at RMS, and Shivani Kak, Head of Investor Relations.
- **Venture capital visit to Index Ventures** to discuss AI's impact on software and company creation, the characteristics of durable AI businesses, enterprise adoption and ROI, the outlook for incumbents and AI-native vertical applications and evolving infrastructure and application-layer economics, featuring Partners Nina Achadijian and Nina Gerson.
- **Presentation by a professor at UC Berkeley and UCSF** on the use of AI to predict patient outcomes, personalize screening and treatment decisions and accelerate clinical validation, including the shift from narrow imaging tools to broad medical foundation models, led by Adam Yala, Assistant Professor of Computational Precision Health, Statistics & Computer Science at UC Berkeley and UCSF.

Implications for Business & Information Services and Software

Business & Information Services Sector Implications

Proprietary information assets become more valuable when combined with domain expertise and workflow context. AI increases the utility of differentiated data by making it easier to retrieve, connect and apply inside customer decisions. That said, durable differentiation depends on more than owning a large dataset. Leading platforms combine authoritative content with validated models, specialized taxonomies, customer context and experts who understand how the information should be used. This combination can generate more accurate and verifiable outputs in financial analysis, legal services, insurance, risk management and healthcare. Business & Information Services companies with unique, continuously refreshed data and deep industry expertise should therefore strengthen their competitive positions as foundation-model capabilities become more accessible.

AI expands the sector from information delivery toward workflow execution. Business & Information Services companies have historically monetized data, analytics and professional expertise through subscriptions, reports and labor-intensive services. Agents can now structure unorganized information, perform analysis, monitor changing conditions and initiate downstream actions within a common workflow. Daloopa can update financial models and complete repeatable research tasks, Clio and Harvey can execute legal workflows and Moody's can convert insurance submissions into decision-ready data within seconds. These capabilities allow providers to address a larger share of customer spending instead of remaining one input into a broader process. The strongest platforms could evolve from systems of information into operating systems for knowledge-intensive industries.

AI can expand addressable markets by converting professional work into repeatable software. Financial analysis, legal research, insurance underwriting and market research contain expensive workflows that historically required substantial professional labor. AI can encode portions of that expertise into reusable agents, making existing work faster and previously uneconomic tasks commercially viable. Lower delivery costs can increase service accessibility, usage frequency and the number of questions customers can afford to answer. This creates an opportunity to capture professional-services spending instead of only competing for existing software budgets. Business models should benefit most where providers can automate structured portions of a workflow but preserve expert judgment for forecasts, consequential decisions and exception handling.

Monetization should broaden beyond seats toward usage, transactions and outcomes. Seat-based pricing becomes less aligned with customer value when agents complete more work independently and each professional can supervise greater activity. Business & Information Services platforms can monetize incremental data consumption, completed analyses, processed documents, transactions and agent usage. Daloopa cited roughly 100x growth in data consumption across certain integrated agent workflows, and Clio combines subscriptions with payments and other transaction-based revenue. Moody's is also beginning to charge incremental subscription fees for MCP access and separate fees for agent usage. These models can allow revenue to grow independently of customer headcount and give providers a direct way to participate in the productivity

gains their products create.

Incumbents retain meaningful advantages, but AI-native entrants can win through deeper workflow control. Established platforms benefit from trusted brands, embedded distribution, accumulated customer context and access to the applications where decisions are executed. These advantages give incumbents permission to create records, post transactions, process payments and trigger actions without requiring customers to transfer information between systems. AI-native entrants can still gain share through materially better products, specialized workflows or difficult-to-replicate infrastructure. Corgi combines AI with insurance licenses, regulatory capital and direct control of underwriting and claims, and Harvey combines legal expertise with workflow-specific agents and evaluations. The most exposed providers are those with commoditized data, shallow workflow integration or business models tied primarily to labor and user headcount.

Verification and accountability will determine how quickly AI reaches production. Business & Information Services workflows frequently involve financial, legal, regulatory or operational consequences that make unverified outputs difficult to deploy. Adoption should progress fastest where decisions are bounded, results can be checked, incorrect actions can be reversed and responsibility for errors is clear. Source-linked data, audit trails, confidence thresholds, independent evaluations and targeted human review can reduce the oversight burden without eliminating accountability. Providers with trusted content, validated domain models and established regulatory relationships should have an advantage over general-purpose applications in these markets. New opportunities may also emerge in certification, monitoring, benchmarking, incident reporting and insurance as enterprises require infrastructure that allows agents to act safely.

Software Implications

Open source vs. frontier models: the outcome is likely more nuanced than either camp suggests on the path to more granular workflow segmentation.

- The frontier camp argues that model progress is significantly underestimated because enterprise benchmarks are poor proxies for real capability gains. This makes step-function improvements in intelligence difficult to observe. If intelligence is the largest addressable market in history, there may be no shortage of problems for increasingly capable frontier models to solve. Interestingly, several early-to-market AI-native application companies will talk about model diversification, but still remaining largely dependent on frontier providers in production environments, because of total cost of ownership considerations (any trade-off in accuracy is not worth the model savings) and product priority considerations (focus on platform functions now, and cost later).
- At the same time, industry experts suggested that most enterprise workflows do not require frontier-level intelligence, and customers are becoming more willing to trade marginal performance for materially lower costs. One VC suggested 90% of tokens will go to open source in ~12-18 months. Early signs of pricing pressure, rising usage limits, and examples of customers reducing frontier model consumption suggest model providers are already competing on economics rather than capability alone. The more nuanced view is that frontier models retain the highest value workloads,

while open source captures the majority of volume.

As the Pareto-optimal curve of intelligence ebbs and flows, we are increasingly confident that independent software vendors like Microsoft, Cloudflare, Databricks and Vercel will be able to capture value by giving enterprises tools to orchestrate models while keeping their own content proprietary.

Rise of world models suggests that the next AI platform shift may look bigger than LLMs. A growing number of researchers are shifting focus beyond LLMs over the last 18 months, arguing that next-token prediction alone is unlikely to capture how the real world behaves. Instead, the next generation of AI will need to model environments, causality, physics, and real-world interactions. The competitive moat could also look different from the LLM era. Rather than training on largely shared internet-scale datasets, many world model applications depend on highly specialized data tied to physical systems, industries, and operating environments. That creates a fundamental debate around whether value accrues to broad foundational models or vertically specialized systems built on proprietary datasets. The economic opportunity could ultimately exceed that of language models. Physical, industrial, scientific, and robotic systems represent a much larger addressable problem set than text generation alone, with many workloads requiring substantially greater compute intensity. **Token demand from world models creates an S curve that will layer on top of our enterprise forecast for compute to support 24x more tokens in the next 5 years. This likely drives tighter supply/demand dynamics for longer, benefiting Microsoft, Oracle and CoreWeave.**

Durable differentiation will come from ownership of workflow context, not just data. The vertical vs. horizontal debate increasingly misses the real source of differentiation. Simple AI wrappers appear most vulnerable as model capabilities improve, while horizontal platforms benefit from scale but remain exposed to feature competition. The strongest positions may belong to workflow-native software companies that combine domain expertise, embedded processes, and proprietary context. In an AI world, owning the workflow may matter more than owning the application. The same principle applies to data. Access to customer data is often mistaken for a moat, but many forms of enterprise data access are effectively commoditized. Any vendor integrated into the same systems can often see the same underlying information. The more durable advantage comes from being embedded in the workflow that continuously creates, structures, and enriches that data over time. As software creation becomes cheaper and faster, unique workflow context may become one of the few remaining forms of defensible differentiation. **Shopify in our view is a good example of leveraging incumbency into a bleeding-edge product roadmap in a vertical-specific customer base. Samsara (IOT) illustrates the same dynamic, with its connected asset footprint continuously generating the operational context required to deliver differentiated AI products.**

Agents are becoming a new top of the funnel. Historically, developers and end users chose tools directly. In an agentic world, software increasingly needs to be discoverable, callable, and usable by agents acting on behalf of users. Vercel's commentary suggests that the fundamental building blocks of the internet are shifting from users and applications toward agents and tokens, creating a world where infrastructure is no longer optimized for human interaction alone. If agents become a meaningful source of traffic creation and workflow orchestration, the biggest beneficiaries could be the

infrastructure vendors that sit closest to where those agents operate.

The infrastructure stack is being optimized for agent-native workloads. Vercel and ClickHouse both noted that AI workloads place fundamentally different demands on infrastructure. Agents generate more requests, execute more actions, make more decisions, and interact with more systems than traditional software. That shifts the bottleneck away from application development and toward the infrastructure required to support autonomous execution. Real-time databases, low-latency compute, scalable data pipelines, and globally distributed systems become increasingly important because infrastructure is now serving machines operating at machine speed rather than humans operating at human speed.

Our industry conversations echoed the consensus view that Security should benefit from AI, and incumbents could benefit disproportionately, with a focus on SOC. AI is likely to expand the attack surface, similar to prior platform shifts like internet and cloud (see our lag between AI and AI security spend note). It should also make attacks more frequent and more sophisticated, especially as open-source models improve and put stronger capabilities in the hands of attackers. This will likely drive a near-term push toward people and process (e.g. for patching and hygiene) and product experimentation. There is a scenario where AI in the hands of the defenders shifts the odds of success permanently in favor of enterprises, as they may gain perfect visibility into their environments such that attackers only need to make one anomalous move to be caught. At the same time, AI also creates an opportunity to help enterprises with two core cyber problems: talent shortages and too many poorly integrated vendors. This is particularly true in SecOps, where there is an opportunity to rebuild with better automation, continuous remediation, and systems that are configured correctly from the start rather than stitched together. One VC noted that most organizations don't even have the basics of data aggregation and organization, such that they could introduce more sophisticated reasoning on top of data stores. This likely creates a push toward modernization built on top of a SIEM data store, while a new layer of agent-native tools emerges to reason over that data, enforce controls, and take action in formats built for agents rather than human dashboards.

Takeaways from Private Companies

Daloopa

Company Introduction

Daloopa is an AI-powered financial data and research platform that automates the extraction, standardization and validation of financial statement line items, KPIs and company guidance from regulatory filings, earnings materials and other corporate disclosures. The platform is designed for equity analysts, investors and financial institutions, providing source-linked and audit-ready financial data that can be integrated directly into research workflows and valuation models. It supports model maintenance, earnings analysis, financial forecasting and investment research workflows. The platform is used by hedge funds, asset managers and investment banks seeking to improve productivity, data quality and workflow efficiency.

Product Demo

Daloopa's product demo showcased its AI-powered financial data platform, which collects every historical financial statement line item, KPI, company guidance metric and non-GAAP adjustment from primary source documents and links each data point to the underlying disclosure for verification. The company demonstrated how its real-time data infrastructure can automatically map newly reported financial information into customer-specific Excel models regardless of their structure, reducing the manual work required to update models during earnings. Daloopa also showed a library of AI skills that uses its structured data to perform workflows such as earnings reviews, research note generation, guidance analysis, supply chain mapping and industry analysis. The demo illustrated a knowledge graph that organizes companies by sector and connects each company with previously completed analyses, creating a continuously updated research memory that can support additional queries without regenerating the underlying work. Daloopa also presented an automatically generated company model that incorporated historical financials, consensus estimates, segment forecasts and supporting analyses, with outputs linked to their original sources.

Meeting Takeaways

Financial data infrastructure is the core differentiator. Daloopa extracts, structures and sources disclosed financial data, KPIs and company-specific adjustments while preserving the terminology used by each company. Each data point links to its original source, allowing customers and AI systems to verify the information behind an output. Proprietary data pipelines, structured operating processes and a roughly 400-person validation team help address AI's inherent variability while meeting stringent accuracy and latency requirements. This combination of financial expertise, technology and human validation is more difficult to replicate than the LLM or agent layer.

AI expands Daloopa from data collection into financial analysis. Daloopa can automate model updates across large company universes, reducing manual data entry and allowing analysts to maintain broader coverage. Because the company captures a comprehensive financial history, individual customer models can be mapped to subsets of Daloopa's underlying data despite differences in model structure. Its workflows extend from earnings analysis and research generation to industry, supply-chain and

company-specific financial analysis. Daloopa automates repeatable and data-intensive work while leaving forecasts and investment conclusions with the user.

Open-source skills broaden adoption while proprietary data remains the moat. Daloopa open-sources standardized agents and analytical skills while keeping its data extraction, normalization and validation infrastructure proprietary. The company views common Wall Street analyses as a broad but finite canon that can be standardized and made available across models. Daloopa's data and skills can be integrated with AI labs and third-party platforms, embedding its capabilities in the interfaces where customers already work. This model expands distribution without requiring Daloopa to own the end-user interface.

A knowledge graph creates persistent and testable research memory. Daloopa runs standardized analyses across companies, sectors and regions, then connects those outputs within a continuously updated knowledge graph. The structure allows AI systems to retrieve completed research instead of recreating the same work every time a user asks a question. Research conclusions can be tested systematically against source-linked financial data and connected across related companies and industries. Over time, this knowledge graph could become an institutional research memory that improves consistency and reduces hallucinations as additional data and analyses are added.

Agentic workflows could increase financial data consumption by orders of magnitude. Traditional users access only a small portion of Daloopa's available data, while integrated agents can retrieve information continuously and apply it across large company universes. Daloopa cited roughly 100x growth in data consumption in certain integrated workflows, illustrating how agentic research can generate substantially greater usage than traditional analyst interfaces. The company is also seeing engagement shift from direct product use toward invisible consumption through AI labs and third-party platforms. This positions Daloopa as financial data infrastructure and could materially increase demand for accurate, source-linked proprietary data.

Divergent

Company Introduction

Divergent is an AI-enabled manufacturing technology company that has developed the Divergent Adaptive Production System (DAPS), an end-to-end digital manufacturing platform that combines AI-driven engineering software, additive manufacturing and automated assembly. The platform uses AI to translate product requirements into optimized designs that account for both engineering objectives and manufacturing constraints before producing components through industrial 3D printing and robotic assembly. Divergent's technology is used across aerospace, defense, automotive and industrial markets where performance, production flexibility and supply chain resilience are critical. By integrating design, production and assembly into a single software-defined manufacturing platform, the company can reduce development timelines and accelerate the commercialization of complex hardware products.

Product Demo

Divergent's product demo showcased complex, load-bearing structures produced through DAPS, its integrated platform for designing and manufacturing aerospace,

defense and automotive hardware. The demo walked through DAPS' four core stages of design, additive manufacturing, assembly and quality assurance, all connected through a common software backbone. Divergent's AI-enabled design process evaluates hundreds of structural alternatives before engineers select and refine a final configuration based on performance, manufacturability and cost requirements. The company also demonstrated how its Model One industrial metal 3D printers, automated assembly cells and closed-loop quality systems work together to manufacture a wide range of structures without requiring part-specific tooling or dedicated production lines. The demo highlighted how integrating AI, software, manufacturing equipment and factory operations can shorten product development timelines, reduce tooling costs and component mass and support repeatable production across multiple programs.

Meeting Takeaways

AI connects the full design-to-production workflow. Divergent embeds AI and software throughout DAPS, connecting design, additive manufacturing, assembly and quality through a common software backbone built with more than $1bn of investment. Its software evaluates hundreds of structural alternatives before engineers select and refine a final design. Machine learning predicts manufacturing distortion, while software adjusts production instructions and supports factory scheduling, robotics and retrieval of internal operating knowledge. These capabilities connect decisions that traditionally sit across separate engineering teams, suppliers and production systems. Divergent illustrates how vertical AI can create value by integrating an industry's complete design-to-production workflow.

Proprietary production data improves design and manufacturing over time. Divergent controls its software, Model One printers, manufacturing processes and factory operations, providing access to process data that competitors without a similarly integrated platform may struggle to reproduce. Historical production data helps predict how parts will deform during printing, while computational methods compensate for that distortion through changes to the manufacturing process. The resulting manufacturing insights also feed back into design, allowing engineers to evaluate whether a part can be produced reliably before it reaches the factory. Continued production creates additional information that can improve future design and manufacturing decisions. Divergent demonstrates how proprietary domain data can strengthen vertically integrated AI applications.

A modular architecture balances innovation with reliability. Divergent builds its AI and software from modular components instead of relying on a single black-box model. Individual optimization algorithms, simulation tools and AI models can be inspected or replaced as technology improves without rebuilding the entire platform. The company uses data-driven simulation for rapid early-stage assessments, physics-based simulation as designs mature and industry-standard finite-element analysis for final verification and validation. Engineers and customers retain control over critical design, certification and production decisions. This architecture allows Divergent to adopt improving AI capabilities while preserving the transparency, repeatability and accountability required in safety-critical markets.

Software-defined factories increase flexibility and reduce program concentration. DAPS can produce aerospace, defense and automotive structures using the same factories, printers and software without requiring dedicated tooling and production

lines for every program. Model One was designed for repeatable production of thousands of parts, extending additive manufacturing beyond prototypes and one-off components. Divergent's second Long Beach factory spans more than 400,000 square feet and is intended to house at least 64 Model One printers over the following 24 months. Once both factories are fully operational, Divergent indicated that annual capacity could support tens of thousands of missile airframes or hundreds of thousands of smaller parts depending on product mix. This flexibility can improve asset utilization and reduce dependence on any single customer, program or end market.

AI is producing measurable physical-world outcomes. A recent Divergent program moved from an initial whiteboard concept to first flight in 71 days, compared with conventional aerospace development cycles that can take years. Across its programs and applications, Divergent cited tooling reductions of up to 95% and mass reductions of approximately 34%. Certain munitions produced through DAPS could cost roughly $200,000-500,000 per unit compared with approximately $2-6mn for legacy systems, although economics vary by application. Divergent also has more than 20 active aerospace and defense programs and produces commercial automotive components used in vehicles today. These results demonstrate how integrated AI and manufacturing can improve development speed, tooling intensity, mass, cost and physical system performance.

Clio

Company Introduction

Clio is a legal tech platform that provides software for law firms to manage casework, client relationships and business operations. The platform incorporates AI capabilities by leveraging a firm's internal matter data, client records, communications and legal documents to automate legal workflows and improve attorney productivity. By combining LLMs with legal-specific context, Clio enables legal professionals to apply AI directly within legal workflows while leveraging firm-specific data and institutional knowledge. The company is expanding AI across practice management, client intake, document generation and legal research workflows, helping law firms improve efficiency and deliver services more effectively.

Product Demo

Clio's product demo showcased an end-to-end divorce matter workflow spanning client intake, case management, legal research, document drafting, court filing, docket monitoring, billing and payments within a unified platform. The workflow began with Clio's AI answering a prospective client's phone call, collecting case details, conducting a conflict check, scoring the lead, scheduling a consultation and automatically creating a legal matter populated with information from the call transcript. Clio then demonstrated its Flow capability, which used natural-language instructions to build an 18-step workflow combining deterministic automation, probabilistic AI tasks and human review, including sending questionnaires, adding relevant parties and communicating with the client. Within Clio Work, the platform recommended relevant legal research based on the matter's facts, generated responses grounded in Clio Library's primary legal content with links to cited authorities, and drafted required court forms through structured AI playbooks. The platform subsequently filed documents with the court, captured related expenses and billable time, monitored docket activity, analyzed opposing filings and

maintained an audit trail of actions completed by both attorneys and AI before generating and sending the client's bill.

Meeting Takeaways

Legal AI can expand the market for legal services. Legal is a strong AI use case because its inputs and outputs are predominantly text based, legal work is expensive and the category contains extensive public and private data. AI can reduce the time and cost required to research, draft and administer legal matters, making some previously uneconomic work viable. This creates an opportunity to expand legal-services consumption instead of only automating work already performed by lawyers. Clio is extending from practice-management software into research, drafting and other substantive legal workflows through Clio Work. The opportunity therefore includes both greater software penetration within law firms and broader access to legal services.

Context and authoritative legal content provide the competitive advantage. AI performance increasingly depends on the proprietary data, permissions and workflow context surrounding foundation models. Clio combines firm communications, documents, matter records and financial information with primary legal content from vLex across roughly 110 countries and more than 1bn court documents available through Clio Docket. Its AI can use firm-specific facts when executing work while grounding legal research in cited cases, statutes and regulations that lawyers can verify. This context remains available even as the underlying model layer evolves. Proprietary legal content, customer context and native workflow access could therefore remain durable differentiators as foundation-model capabilities become more widely available.

AI agents are moving from conversation into end-to-end legal workflows. Clio's product demonstration showed AI progressing from client intake and conflict checks through matter creation, research, drafting, filing, billing and payments. Clio Flows combines deterministic software actions, probabilistic AI tasks and human judgment within a single configurable process. Incoming calls, communications and documents can trigger actions automatically without requiring users to recreate context or move information between applications. Lawyers can inspect when AI enters a workflow and review consequential outputs, providing the governance required for legal work. Enterprise AI adoption is shifting from standalone chat interfaces toward event-triggered execution inside production workflows.

Systems of record can become systems of action. Clio serves more than 400,000 legal professionals across over 50,000 law firms, giving its AI access to the data, permissions and applications required to act for users. AI point solutions may require lawyers to upload documents, recreate matter context and transfer outputs into downstream systems. Clio can natively answer calls, create matters, conduct research, draft documents, file with courts, track time, issue bills and process payments. More than 300 integration partners extend these capabilities across broader law-firm technology environments. Existing workflow platforms can use installed-base distribution, native actions and accumulated customer context to compete with AI-native entrants.

AI expands monetization beyond user seats. Clio combines subscription software with AI products and transaction-based revenue streams that can grow independently of legal headcount. Clio Work became the fastest-selling new product in the company's history and generated more than 50% of overall ARR bookings in the latest quarter,

indicating that AI demand can be incremental to core practice-management spending. Clio customers bill nearly $50bn annually through the platform, while the company currently monetizes almost $9bn through Clio Payments. As AI executes more work inside legal workflows, monetization can increasingly reflect completed tasks and transaction volume rather than only licensed users. This gives integrated AI platforms multiple ways to capture value from productivity gains and higher legal-services consumption.

Corgi

Company Introduction

Corgi is an AI-native insurance platform and licensed carrier that applies AI across underwriting, policy administration, claims processing and insurance operations. The company has built its technology platform around AI models that automate portions of the insurance workflow traditionally performed through manual processes. Its platform is designed to accelerate quoting, underwriting and policy issuance while improving operational efficiency. Corgi also underwrites, prices, issues and services its own policies through a digitally native insurance platform.

Meeting Takeaways

AI unlocks language-intensive industries. Insurance represents a compelling AI opportunity because policies, applications, claims and regulatory documents require extensive natural-language interpretation. Earlier fintech innovation focused on numerical databases and general ledgers, while foundation models can now process the unstructured language that defines insurance products and workflows. Corgi has automated most of the underwriting production process, although regulators may still require human review of a policy sample. In trucking insurance, this automation can reduce quote times from roughly two weeks at legacy carriers to approximately two minutes. Foundation models can therefore address language-based workflows that traditional financial software could not reliably automate.

AI-native platforms can compete through a better customer experience. Corgi uses AI primarily to complete insurance workflows faster and more effectively, with cost reduction serving as a secondary benefit. Automation across underwriting, regulatory documentation and claims can compress processes that traditionally take weeks into minutes or days. During catastrophic events, AI voice agents can scale first-notice-of-loss intake dynamically when claim volumes exceed the staffing capacity of legacy carriers. Faster claims resolution can improve customer satisfaction while reducing reserve uncertainty and regulatory pressure. Corgi illustrates how AI can support growth by improving speed, reliability and responsiveness instead of competing primarily on price.

Agentic AI is becoming viable in regulated production workflows. Recent model advances have increased the reliability of multistep automation and workflow hand-offs, enabling capabilities that Corgi indicated were not practical roughly six months earlier. Agents can interpret insurance documents, initiate downstream processes and coordinate tasks across underwriting, documentation and claims. Structured software and regulatory controls remain part of the workflow, while humans participate in physical inspections, selected judgments and required oversight. Critical actions can be

repeated or checked by additional models because errors involving payments, coverage or regulatory requirements carry significant consequences. Agent adoption in regulated industries will therefore combine greater autonomy with validation, auditability and targeted human review.

Regulatory infrastructure creates a more durable moat than model access. Corgi spent multiple years and tens of millions of dollars establishing the licenses, capital and operating infrastructure required to participate at the insurance-carrier and reinsurance layers. The company controls underwriting, pricing, reserving, regulatory documentation and claims while frequently relying on brokers, agencies and other partners for distribution. This structure allows Corgi to redesign and automate regulated workflows without depending on a legacy carrier's systems or priorities. Regulatory approvals, capital requirements and insurance operating expertise are substantially more difficult to reproduce than access to foundation models. AI-native companies can build durable positions by combining technology with control of regulated infrastructure.

Model commoditization shifts value toward applications and workflows. Open-source models increasingly provide sufficient performance for many application-layer tasks, while Corgi can select among providers based on performance, convenience and cost. Lower inference costs allow the company to use more model calls, repeat critical actions and validate outputs before completing regulated transactions. Proprietary frontier models can remain useful when their performance or ease of deployment justifies the higher expense. This supports a multi-model architecture rather than dependence on a single provider. As model capabilities become more broadly available, differentiation should migrate toward workflow design, customer experience, distribution and regulated operating infrastructure.

Harvey

Company Introduction

Harvey is an AI platform built specifically for legal workflows that combines foundation models with legal-specific data, documents and institutional knowledge to support drafting, summarization, diligence and legal research. Harvey enables lawyers to query large volumes of legal and business information using natural language and generate work product grounded in both internal firm knowledge and external legal sources. The platform is designed to streamline complex legal workflows, reduce time spent on manual document review and improve the speed and quality of legal analysis.

Product Demo

Harvey's product demo showcased its agent and Vaults capabilities across two investment-management legal workflows. The company first demonstrated how a lawyer can use natural-language instructions to create a no-code agent that monitors substantial shareholding and position-reporting requirements across jurisdictions, runs on a recurring schedule and generates an executive briefing, interactive dashboard and PowerPoint presentation. The agent produced pinpoint citations linked to the underlying regulatory sources, allowing lawyers to verify each conclusion while reducing a research process that could require more than 60 hours to minutes. Harvey then demonstrated Vaults, which can analyze up to 10,000 documents simultaneously, by extracting fund and investor rights from more than 1,000 side letters into a structured review table.

Users could query the table to identify existing investors eligible to elect into the economic terms of a proposed side letter, with each answer linked to supporting documents for review.

Meeting Takeaways

Vertical AI platforms can become industry operating systems. Harvey combines AI with the documents, institutional knowledge, permissions and tools required to complete legal work within a common platform. Users can research, draft, review documents and analyze large datasets without repeatedly transferring context between disconnected applications. Shared workspaces also support collaboration within organizations and between law firms and in-house legal teams while maintaining appropriate access controls. The platform brings together research, drafting, agents, Vaults and knowledge sources rather than treating AI as a standalone chatbot. Vertical AI leaders could become operating systems for knowledge-intensive industries by combining intelligence, context, collaboration and workflow execution.

AI agents turn professional expertise into repeatable software. Harvey allows legal professionals to create agents through natural-language instructions instead of code, converting recurring research, drafting and review processes into reusable workflows. One demonstrated agent monitored substantial-shareholding requirements across jurisdictions, cited regulatory developments and generated a summary, interactive dashboard and PowerPoint presentation. The agent can substantially reduce work that historically required dozens of hours while preserving links to underlying sources for professional verification. Harvey's Vaults product can similarly analyze up to 10,000 documents and extract structured answers across workflows such as side-letter review. Agentic applications can encode specialized professional methods in repeatable software while allowing experts to inspect and refine the results.

Open-source models improve application-layer economics. The narrowing performance gap between open-source and frontier models allows applications to route simpler tasks toward lower-cost models while reserving frontier intelligence for complex work. Harvey evaluates models across roughly 1,200 legal tasks and approximately 60 private evaluation sets to determine which models perform best for specific workflows. The company can also post-train smaller models using proprietary legal evaluations and task environments without absorbing the substantially higher cost of pretraining a foundation model. A multi-model architecture can optimize accuracy, latency and expense as model capabilities evolve. Application-layer companies can therefore benefit from model commoditization while differentiating through evaluations, workflow design and customer context.

Domain expertise is critical to enterprise AI adoption. Deploying a capable model does not ensure that professionals will redesign entrenched workflows or understand which outputs require additional review. Harvey employs more than 240 lawyers within a workforce of approximately 1,200 people, integrating legal expertise across product development, customer engagement and other functions. Adoption typically progresses from experimentation to inserting AI into selected workflow steps and ultimately redesigning broader processes around AI. Usage and time savings increase as professionals become more familiar with the system and identify additional tasks it can perform. Vertical AI companies may need substantial domain expertise and workflow-transformation capabilities to convert technical performance into sustained

production usage.

Data access can matter as much as data ownership. Harvey combines customer-controlled internal information with external legal and enterprise knowledge sources through a mix of native capabilities, integrations and partnerships. Customers retain ownership of their data and outputs, Harvey does not train shared models on customer information and customer-specific models remain exclusive to the relevant customer. Connections with enterprise systems and third-party data providers give AI the context required to perform legal work without requiring Harvey to own every underlying dataset. This approach can broaden coverage while preserving customer trust, security controls and flexibility across jurisdictions. AI application moats may therefore depend on permissioned access, workflow integration and trusted data orchestration as much as outright ownership of proprietary content.

Vercel

Company Introduction

Vercel is a cloud platform that provides developers with tools to build, deploy and scale web applications, with a focus on AI-native software development. The company offers infrastructure, developer frameworks and AI tooling that enable developers to build applications powered by LLMs without managing complex back-end infrastructure. Vercel helps bridge the gap between foundation models and end-user applications by providing the infrastructure needed to deploy, manage and scale AI-native software applications. Please see our detailed takeaways from Vercel's June customer event here.

Meeting Takeaways

Agentic infrastructure is emerging as a new category of the software stack and becoming the control point through which agents build, deploy, and operate software. AI has catalyzed the largest acceleration in Vercel's history, driven by a shift where the fundamental building blocks of the internet are no longer just users and applications, but agents and tokens. As AI agents generate, deploy, and interact with software autonomously, demand will grow for the underlying infrastructure that manages distribution, execution, and observability. Vercel already sits at the center of this trend, supporting 1.5tn monthly requests, 45tn monthly tokens, and 230mn deployments (a majority of which are initiated by agents). Ultimately, the proliferation of AI-generated software will increase demand for a new class of "agentic infrastructure," and Vercel is positioning itself as the backbone through which an increasing share of internet activity, both tokens and user experiences, ultimately flows.

Infrastructure is moving from human-first to agent-first, and "agentic infrastructure" could disrupt platforms purpose-built for the cloud-era. Much of today's infrastructure platforms, from hyperscalers to cloud era infrastructure platforms, were designed for dashboards and workflows consumed by humans. In contrast, Vercel believes future infrastructure will increasingly be consumed directly by agents. This vision extends beyond deployment into token routing, observability, security, cost controls, and agent orchestration. As such, Vercel is positioning itself to be the "agentic infrastructure," which it sees as analogous to what AWS became during the cloud transition. Said differently, the company believes that platforms like Snowflake, Datadog, and much of the current infrastructure stack were designed for humans

operating software, while Vercel believes the next generation will be designed for software operating software.

Agent-led distribution is emerging as a powerful new top-of-funnel acquisition channel. Historically, Vercel's growth was driven by developers adopting open-source frameworks such as Next.js and bringing the platform into their organizations. Management argued that AI is amplifying this motion by introducing agents as a new source of distribution and product adoption. Customers often begin with web development workloads, but increasingly expand into AI and agent-based applications, or vice-versa. This creates a self-reinforcing cycle in which open-source adoption attracts developers, developers attract agents, and agents generate more software and infrastructure consumption. This dynamic is helping drive broader enterprise standardization, with >4,000 $10k+ ARR customers, growing +100% y/y in June 2026, which boasts 150%+ NDR. At the same time, Vercel continues to gain share among the next generation of software companies, with penetration across Y Combinator cohorts increasing from 55% to 71% (Fall 2025 to Winter 2026), supporting the idea that Vercel is becoming the default application platform for the AI era, similar to AWS's role as the default infrastructure platform during the cloud transition.

ClickHouse

Company Introduction

ClickHouse is an open-source analytical database and real-time data platform that enables organizations to store, process and analyze large volumes of data. Originally developed for analytics workloads, the platform has expanded into machine learning and AI-related use cases that require fast access to large datasets. ClickHouse supports real-time analytics, data processing and retrieval across structured and unstructured information. The platform is used by organizations to support AI agents, retrieval systems, customer-facing applications and other latency-sensitive workloads.

Meeting Takeaways

Real-time analytics remains the core platform, while observability is becoming an increasingly important wedge into the enterprise. ClickHouse was built to power large-scale, low-latency analytics workloads, and management continues to view real-time analytics as the company's foundational use case. However, some of the strongest recent commercial momentum is coming from observability-related workloads, particularly log analytics and AI application telemetry, where customers are looking to reduce spend on platforms such as Splunk and Elastic. Management emphasized that winning observability budgets is often a more effective entry point than pursuing wholesale warehouse migrations, positioning observability as an increasingly important path to enterprise expansion.

AI is creating a new class of data-intensive workloads. Many of ClickHouse's fastest-growing customers sit at the center of the AI ecosystem, including companies involved in model development, AI observability, agent infrastructure, and developer tooling. Management noted that AI applications generate substantial volumes of telemetry, usage, and operational data that must be captured, analyzed, and queried at scale. As AI adoption expands, so does the need for the data infrastructure underlying those workloads. Rather than creating an entirely new category for ClickHouse, AI is

increasing demand for the capabilities the platform was originally built to provide.

Near-term opportunities may come from new and specialized workloads rather than platform consolidation. ClickHouse notes that enterprises are becoming increasingly comfortable running multiple specialized data platforms, each optimized for a different job. Rather than competing head-on for enterprise warehouse migrations, the company is focused on workloads such as observability, telemetry, and real-time analytics, where existing platforms were not built to operate. As a result, future market share gains may come less from platform replacement and more from becoming the system of choice for emerging operational and AI-native workloads.

Database engineering remains the foundation of differentiation, though ClickHouse is selectively expanding the platform around it. Management attributes the company's performance advantages to years of engineering work around compression, scalability, elasticity, and low-latency query performance. At the same time, the company is broadening its product surface area through initiatives such as ClickStack, enterprise observability functionality, and a recently launched managed Postgres offering. Notably, management highlighted Postgres as a particularly attractive opportunity given strong developer adoption and the rapid growth of platforms such as Supabase. While ClickHouse continues to evaluate opportunities higher in the stack, the company appeared focused on strengthening its position as the underlying data platform before aggressively expanding into more application-oriented categories.

Simile

Company Introduction

Simile is an AI simulation platform that uses AI agents and foundation models to predict how individuals and groups may respond to products, policies and business decisions. The company creates digital representations of real-world populations using behavioral data to simulate potential outcomes and interactions. Its technology is designed to help organizations evaluate strategies, test scenarios and generate insights before making decisions. Simile serves enterprises seeking alternatives to traditional surveys, focus groups and other market research methods.

Product Demo

Simile's product demo showcased its AI-driven human behavior simulation platform, which combines interviews, surveys and behavioral data to model how individuals and customer populations may respond to changes in products, pricing, messaging or policy. The company demonstrated how users can select or filter a population, enter an open-ended research question and simulate responses across thousands of digital representations grounded in interviews with real people. The demo featured an autonomous shopping use case that evaluated whether existing AI users would permit an agent to complete purchases, the maximum transaction value they would authorize and the product categories for which they would restrict agent activity. Simile converted the results into an executive-level report containing key findings, response distributions, visualizations and recommended actions, while allowing users to examine the underlying interview questions, participant quotations and modeled rationale supporting each conclusion. The platform also displayed confidence and groundedness scores intended to quantify prediction reliability and showed how customers can create proprietary

populations using their own data while keeping that information separate from the broader foundation model.

Meeting Takeaways

Simulation shifts AI from prediction to decision-making. AI is increasingly being used to answer counterfactual "what if" questions, helping organizations evaluate how changes in pricing, products, messaging or policy could alter outcomes before committing resources. Traditional forecasting models often rely on historical correlations and can become less reliable when the underlying environment changes. Simulation systems instead seek to model the causal drivers behind human behavior and estimate how people may respond under alternative conditions. Simile's platform supports scenarios at both the individual and population level. This extends AI from predicting likely outcomes toward evaluating how specific decisions could change them.

Proprietary behavioral data may become a durable moat. As foundation models become more widely available, differentiation is shifting toward specialized data that competitors cannot easily reproduce. Simile combines behavioral data with attitudinal information gathered through roughly 20,000-30,000 completed interviews and surveys each week across more than 40 global panel partners. The model learns jointly from what people say, why they make decisions and what they ultimately do. Simile can redirect collection toward domains where predicted accuracy is weaker while relying on accumulated learning in areas the model already understands. The resulting combination of proprietary data, recruitment infrastructure and behavioral modeling may prove more defensible than the underlying open-weight models.

Simulation can make market research more frequent and accessible. Traditional surveys, focus groups and consulting engagements can be slow, expensive and difficult to repeat, limiting how frequently organizations test new ideas. Simulation platforms allow users to evaluate pricing strategies, product concepts, customer segments and marketing initiatives through iterative counterfactual analysis. Users can filter populations, ask follow-up questions and generate updated reports without commissioning a new study for every scenario. Early customers are also using Simile to conduct research that previously would not have justified the time or budget required by traditional methods. Simulation could therefore expand research demand while materially reducing the incremental time and cost of testing additional scenarios.

Trust and uncertainty remain critical adoption barriers. Enterprise users need to understand not only what a simulation concludes but also how much confidence to place in the answer. Simile pairs outputs with predicted-accuracy estimates and uncertainty categories derived from continuing model evaluations. Users can also trace findings to underlying interviews, survey answers and observed behavior, providing evidence for why the model reached a conclusion. Simile validates outputs against private historical studies, post-training-cutoff events and preregistered predictions where possible. As AI enters higher-stakes decisions, transparent uncertainty, traceability and rigorous evaluation may become as important as raw model capability.

Human behavior may emerge as a new AI frontier. Much of the AI industry focuses on coding, mathematics and general reasoning, while modeling how people think, decide and respond represents a distinct problem requiring different data and objectives. Simile post-trains open-weight models on proprietary behavioral and attitudinal

information to develop a foundation model of human behavior. Applications span customer research, investing, healthcare, public policy and economic decision-making, where understanding reactions to potential actions may matter more than forecasting a single outcome. A shared foundation model can also support customer-specific populations and private data without exposing that information to other customers. If performance continues to improve, behavioral foundation models could become a new category of decision-support infrastructure across enterprises and governments.

Takeaways from Venture Capital Visits

Models and Infrastructure

Capability gains continue. AI capabilities are improving through better data curation, training techniques, model architectures, inference methods and infrastructure efficiency. AI is also helping researchers write code, optimize kernels and develop subsequent model generations, creating a recursive improvement cycle. These advances allow leading labs to generate substantially more intelligence from the same compute.

Frontier intelligence retains value. Complex and economically valuable workloads continue to favor frontier models given their higher accuracy, reliability and task-completion rates. Model economics should therefore be measured by the cost of successfully completing a task, including token consumption and success rates, not nominal token prices. The frontier market could support several large providers with differentiated positions, resembling hyperscale cloud computing.

Open models serve distinct workloads. Open models offer attractive economics for constrained, localized and cost-sensitive workloads, while frontier models retain an advantage where reliability justifies higher costs. Applications can route work across frontier, open and internally tuned models, although optimization generally becomes a priority only after reaching meaningful scale. Token share should not be confused with revenue or profit share.

Compute demand outpaces efficiency. Falling token prices are unlikely to eliminate compute constraints because usage and intelligence requirements can expand faster than efficiency improves. Premium models can coordinate fleets of cheaper models, routing each task based on required quality, latency and cost. Long-running research and autonomous workflows create additional sources of compute consumption.

Inference expands hardware TAM. AI infrastructure investment is broadening from training toward inference chips, edge computing and hardware-enabling software. Specialized architectures can compete on cost, throughput and power efficiency, although most alternatives have not reached meaningful production volumes. On-device models represent another underexplored opportunity by enabling low-latency applications without continuous cloud access.

Agents and Enterprise Adoption

Enterprise monetization is clearer. Enterprise AI offers a clearer monetization path than consumer AI because companies can measure returns through labor savings, revenue acceleration and higher productivity. AI spending can be compared directly with employee costs or the value of incremental output. Consumer AI can generate substantial usage, but willingness to pay remains lower outside answer engines, entertainment and personal assistants.

Agents expand application TAM. Agents can capture a larger share of customer spending by automating complete workflows instead of only assisting individual employees. Software development, legal and healthcare are attractive early markets because they employ highly paid professionals, face labor constraints and produce work suited to LLMs. Monetization is shifting from seat-based subscriptions toward

consumption, credits and outcome-based pricing.

Agents reshape software and labor. Enterprise software is moving from systems of record that store information toward systems of action that execute work. Agent architectures combine model routing, tools, protocols and trust controls, allowing employees to supervise substantially more activity than they could perform directly. Adoption will remain uneven because organizations differ in their ability to instruct, supervise and integrate agents.

Verification controls adoption. AI adoption is advancing fastest where outputs can be objectively validated, with coding remaining the clearest breakout category. Progress is slower in mission-critical workflows where errors carry material consequences and human review becomes the primary bottleneck. Better verification and lower oversight costs are critical to expanding AI into higher-stakes applications.

Early revenue can mislead. Rapid expansion from zero to $50mn of revenue may reflect enterprise experimentation budgets instead of durable product-market fit. Customer diligence should establish a clear ROI, recurring budget source and reason to renew after the initial trial period. Retention and expansion will separate products delivering measurable value from those benefiting primarily from AI testing.

Application Moats and Software Disruption

Usage, workflows and data define moats. Durable AI applications combine habitual usage, strong distribution, superior product experiences and ownership of critical workflows. Proprietary data matters when it is unique, continuously refreshed and unavailable through equivalent customer relationships. Thin software layers without deep engagement or compounding feedback loops face increasing replication risk from foundation model labs and new competitors.

Software outcomes will diverge. Software durability depends on workflow depth, customer ownership, switching costs and exposure to foundation-model substitution. Vertical software leaders and core systems of record have more time to adapt given their distribution, embedded data and operational importance. Products with weak retention or functionality that directly overlaps with foundation models face greater disruption risk.

Physical AI, Defense and Cybersecurity

Physical AI needs a distinct playbook. Robotics lacks internet-scale training data and introduces hardware, deployment and real-world data-collection challenges that differ from software AI. Attractive companies may control the model, hardware and deployment for a specific task, creating proprietary operating data through each installation. Additional opportunities exist in AI-enabled hardware design, industrial control software and system integration.

Defense becomes a major AI market. Geopolitical pressure and autonomous systems are expanding opportunities across drones, counter-drone technologies, space infrastructure and defense software. AI improves the capabilities and economics of lower-cost systems, while supply-chain localization increases demand for new domestic providers. The strongest opportunities combine differentiated technology with credible manufacturing and deployment capabilities.

AI expands the cybersecurity opportunity. Companies need new security products to protect AI models and applications from prompt injection, model theft and adversarial attacks. AI agents can also strengthen existing defenses by detecting vulnerabilities, remediating systems and automating labor-intensive security workflows. The most durable opportunities may sit in operational remediation, where agents must act across complex production environments.

Top VC predictions for AI in 2027

- **Another capability breakthrough emerges.** Model progress could produce another market-surprising capability over the next 12-18 months as advances continue beneath the surface.
- **Open models gain inference volume.** Open models could process a substantial share of inference tokens as applications adopt model routing, although frontier models may retain a disproportionate share of economic value.
- **Agents move beyond prompts and coding.** Agents could expand into more enterprise workflows, operate continuously and proactively present completed work. This transition could drive another step-up in token consumption.
- **AI accelerates scientific discovery.** AI could generate meaningful advances across life sciences, mathematics and physics. An AI-discovered drug could provide an important proof point for faster and more efficient research.
- **Robotics enters real-world deployment.** Initial robotics deployments could become more visible in daily life, although adoption will remain gradual and operationally uneven.
- **AI infrastructure gets redesigned.** Data centers, compute architectures and infrastructure software could evolve as inference workloads diversify and model requirements change.
- **A major AI-enabled cyberattack emerges.** Open-weight models could contribute to a major security incident over the next one to two years, accelerating enterprise investment in continuously operating defensive agents.
- **Legacy SaaS disruption accelerates.** More public software companies may need to pursue fundamental AI transformations as incremental product additions prove insufficient.
- **AI consolidation continues.** Strategic acquisitions could increase as large technology companies compete for models, talent, applications and distribution.

Takeaways from Stanford AI Lab and Group Dinner

Biography

Kiana Jafari is a Postdoctoral Researcher at Stanford University working in the Stanford Intelligent Systems Laboratory (SISL). She also serves as Executive Director of the Stanford Center for AI Safety, lead Human Factors Research and Lab Operations at the Stanford Flight Simulator Laboratory and Assistant to the Editors-in-Chief at JAIR. Prior to joining Stanford, she completed her Ph.D. in Systems and Information Engineering at the University of Virginia.

Duncan Eddy is a Research Fellow in the Department of Aeronautics and Astronautics at Stanford University, Executive Director of the Stanford Center for AI Safety and a postdoctoral researcher with appointments in Mineral-X and the Stanford Intelligent Systems Laboratory (SISL). Prior to Stanford, he started and led the Spacecraft Operations Group at Capella Space and the Constellation Operations and Space Safety Groups at Project Kuiper. Most recently, he was a Principal Applied Scientist at Amazon Web Services working on software services for edge computing applications. He received a Ph.D. in Aerospace Engineering from Stanford University.

Research

Kiana Jafari's research focuses on Human-Agent Teaming and Human-Centered AI, with an emphasis on designing intelligent systems that enhance collaboration between humans and machines. Her work explores how AI systems can support human decision-making while preserving autonomy, transparency and alignment with human values. Her research examines how intelligent systems can be developed to complement human strengths across diverse domains and operational environments.

Duncan Eddy's research focuses on decision-making in safety-critical, climate and space systems where operational decisions must be made quickly and correctly in complex environments. His work spans space operations, autonomous systems and AI safety, with an emphasis on ensuring decision-making processes remain understandable to human stakeholders. His research examines how advanced decision-making systems can improve the reliability, safety and performance of complex systems operating under uncertainty.

Meeting Takeaways

World models offer a new path for physical AI. World models learn how a system evolves by observing states, actions and subsequent outcomes, allowing AI to answer, "If I do this, what happens next?" They derive system behavior algorithmically from data instead of relying entirely on handwritten physical equations. The architecture compresses observations into an internal representation, predicts how that representation changes after an action and can reconstruct the resulting future state. This learned system dynamics may be reused across multiple goals more effectively than task-specific approaches that require new training data for each objective. World models could expand AI beyond language into robotics, autonomous vehicles, manufacturing, industrial inspection and space exploration.

World models could ease physical AI's data bottleneck. Physical systems generate

large volumes of sensor data while still encountering only a small fraction of the environments, objects and tasks they may face. World models address this challenge by compressing observations into the features most relevant to predicting future outcomes. In one example, six hours of unstructured drone flight data produced a model that predicted movement through a room while preserving walls, windows and doors and excluding less relevant detail. This compression can reduce data and compute requirements, although discarding too much information can lower fidelity while retaining too much can impair generalization. Effective deployment will therefore require matching model capacity and representation detail to the specific task.

World models provide a native signal for uncertainty and potential failure. Probabilistic world models generate a distribution of possible outcomes, allowing the system to detect when predictions become uncertain or move beyond conditions represented in training data. A robotic system can use this signal to pause, escalate or request additional information before completing an incorrect action. In one example, uncertainty increased immediately before a robot failed to connect a cable, creating an opportunity for human intervention. The signal does not create a complete safety system by itself because thresholds, fallback processes and escalation authority must still be defined. World models could nevertheless combine autonomous decision-making with continuous monitoring in mission-critical physical systems.

Physical AI will likely remain application specific in the near term. Broad world models spanning all physical environments remain difficult to build because each application requires observations, actions and outcomes that represent its unique dynamics. Narrow models can already learn useful behaviors from modest datasets, including several hours of drone flight or recordings of industrial workflows. Near-term commercial opportunities are likely to emerge in contained environments where tasks are repetitive, labor is scarce and human safety risks are high, including mines, refineries, factories and warehouses. Companies with domain expertise and access to relevant operating environments should be well positioned, although limited data requirements may also enable specialized entrants. Transfer across unrelated applications and the development of broadly reusable base world models remain open research questions.

Enterprise agent adoption is constrained by accountability, not capability. The gap between organizations piloting AI agents and deploying them in production remains wide because enterprises have not determined who bears responsibility when autonomous systems fail. Most deployments remain in copilot or supervised modes where humans execute decisions or approve consequential actions despite agents being technically capable of greater autonomy. Delegated and fully autonomous workflows concentrate legal, operational and financial risk with the enterprise. General counsel, risk functions, insurers and auditors may therefore influence adoption as much as the teams selecting the underlying model. Enterprise deployment should accelerate as agent reliability becomes independently measurable and liability becomes assignable, transferable and insurable.

Bounded, checkable and reversible workflows will reach production first. The strongest initial agentic use cases involve constrained decisions, observable outcomes, reversible actions and clear thresholds for escalation. Invoice processing provides one example, with AI extracting structured fields, applying automated checks, routing low-confidence cases for review and posting transactions into a reversible ERP workflow. Documentation, customer service and selected administrative tasks may offer

similar starting points where errors can be identified and corrected without disproportionate consequences. Workflows involving irreversible financial, clinical or safety outcomes require more rigorous evaluation and control systems. Near-term value creation should concentrate in narrow processes where risk can be measured and accountability can be assigned explicitly.

A trust infrastructure layer could unlock enterprise AI adoption. Enterprise agents need certification, independent evaluation, monitoring, incident reporting and insurance frameworks comparable to those used in aviation and other safety-critical industries. Human review alone may be insufficient because attention can deteriorate when systems are consistently correct and reviewers may defer excessively to AI recommendations. A credible framework must define which decisions agents can make, when humans intervene and which party accepts liability for each failure. This creates opportunities in pre-deployment testing, agent-specific insurance, reliability benchmarking, incident reporting and standardized certification. Durable value may accrue not only to models and applications but also to companies that measure reliability, enforce controls and assume or transfer risk.

Takeaways from UC Berkeley and UCSF AI Lab

Biography

Adam Yala is an Assistant Professor of Computational Precision Health, Statistics and Electrical Engineering & Computer Science at UC Berkeley and UCSF. His breast cancer prediction tool, Mirai, has been deployed across 66 hospitals in 30 countries and now supports prospective clinical trials. His work has been featured in major outlets. He earned his BS, MEng, and PhD in Computer Science from MIT, where he was affiliated with the Jameel Clinic and Computer Science & Artificial Intelligence Laboratory (CSAIL).

Research

Adam Yala's research focuses on developing machine learning methods to personalize healthcare and bring them into clinical practice, with an emphasis on predicting cancer risk and designing individualized screening strategies. His work aims to deliver the right medical interventions, such as screenings or treatments, tailored to each patient's unique risks and preferences. His research centers on three major themes, including modeling comprehensive patient records using multi-modal data to better understand outcomes, optimizing clinical decisions through AI-driven predictors and translating these tools into real-world healthcare settings.

Meeting Takeaways

Medical guidelines could become continuously improving algorithms. Many clinical guidelines still rely on a few variables such as age and family history despite their limited ability to identify which patients will develop disease. AI can extract additional signals from mammograms, CT scans, MRIs and longitudinal records containing billions of tokens of patient context. Imaging models have already outperformed traditional risk calculators in forecasting breast and lung cancer. Better predictions could personalize screening, identify patients requiring additional testing and reduce late diagnoses and unnecessary procedures. Healthcare's standard of care could therefore improve at the pace of model development instead of depending on guidelines updated every several years.

Medical AI is shifting toward broad foundation models. Earlier imaging models addressed one modality, disease or task, while emerging models can span imaging formats, anatomical regions and time periods. Pillar-0, an open-source model trained on roughly 200,000 exams, evaluated approximately 350 conditions and performed about ten points better on average than five competing models. Broad pre-training also reduces the data needed to model pediatric diseases, brain cancers and other rare conditions. Pillar-1 extends this approach across mammography, CT and MRI, using tens of millions of exams to generate full radiology reports and disease-specific forecasts. The definition of a narrow model could consequently expand from a single disease detector to an entire medical specialty.

Data access is becoming a commercial advantage in healthcare AI. The primary barrier to developing leading medical models has often been assembling sufficiently large and usable clinical datasets rather than inventing a new model architecture. Academic researchers historically had clinical expertise and limited data but insufficient compute, while startups often had capital and compute without access to clinical

records. Specialized vendors can now work with health systems and make clinical data available to well-funded developers. Data pricing has also become more rational as providers recognize that inaccessible archives have limited standalone value and the marginal value of additional samples declines. Capital, data access and the ability to combine clinical expertise with compute should therefore become increasingly important competitive advantages.

Personalized care requires AI to optimize actions, not only predict outcomes. Identifying elevated cancer risk creates limited value unless the healthcare system can determine whether a patient should receive an MRI, return in six months or follow another pathway. This can be framed as a control problem where clinical context represents the state, the intervention represents the action and the reward balances outcomes against cost and procedural harm. Historical patient trajectories and clinical knowledge can support simulations of alternative screening or treatment schedules. The resulting policy can also incorporate individual preferences and different tolerances for testing. AI could therefore replace fixed population schedules with dynamic pathways that personalize both disease risk and intervention trade-offs.

Clinical evaluation must improve as quickly as medical models. Traditional clinical trials can take years while medical models may improve every few months, creating a mismatch between technical progress and evidence generation. Adaptive trials can reuse evidence when a new model agrees with a validated predecessor and focus incremental enrollment on cases where their recommendations differ. Greater agreement between versions can reduce trial size, cost and duration without eliminating prospective validation. Evaluations must also reflect full clinical workflows because isolated images and captions may not measure performance on complete imaging volumes and diagnostic reports. Medical AI adoption will depend as much on cumulative and clinically representative validation as on model capability.

Takeaways from Moody's

Company Introduction

Moody's (MCO) is a global provider of credit ratings, research and risk analysis. It operates through two primary segments: Moody's Investors Service, which issues credit ratings for debt instruments and entities, and Moody's Analytics, which offers data, software, and advisory services for financial risk and economic analysis. The company's products are used by financial institutions, corporations, governments, and regulators to assess creditworthiness and manage risk. Moody's supports decision-making in areas such as investment analysis, regulatory compliance, and economic forecasting. Its services are grounded in quantitative models, proprietary data, and expert research.

Product Demo

Moody's product demo walked through AI-powered insurance workflows built on its Intelligent Risk Platform, combining proprietary data, catastrophe models and domain expertise to improve underwriting and portfolio management decisions. The first workflow transformed unstructured insurance submissions, including PDFs, spreadsheets and policy documents, into standardized data for underwriting, reducing a process that can take hours or days manually down to seconds while maintaining source-level traceability and high accuracy. Beyond extracting information, the workflow interpreted complex policy language and automatically mapped risks, occupancy types and coverage details into standardized classifications required for downstream pricing and risk models. The second workflow showcased a natural-language analytics interface that enabled users to evaluate portfolio exposures, catastrophe losses and reserve requirements without requiring programming expertise. Using an Australian wildfire scenario, the platform analyzed loss exposures across multiple return-period scenarios and identified portfolio concentrations and potential reserve impacts in real time, illustrating how AI can help insurance professionals move from data gathering to decision-making.

Meeting Takeaways

Connected intelligence breaks down data silos. Risk decisions increasingly require connecting credit, KYC, supply-chain, climate, cyber and insurance data rather than analyzing each domain independently. Companies must understand second- and third-order dependencies across suppliers, infrastructure, counterparties and ownership relationships. Moody's combines proprietary data, entity linkages and specialized models around the decision a customer needs to make. A common data structure can also create consistent answers across insurers, brokers, reinsurers and other market participants. AI is shifting enterprise software from delivering isolated information toward connected and actionable intelligence.

Proprietary data and domain expertise remain the primary AI moat. Foundation models are increasingly accessible, shifting differentiation toward unique data, validated models, scientific expertise and trusted customer relationships. Decades of credit, claims, climate and entity information can be combined with specialized risk models to generate outputs that software-only entrants may struggle to replicate. Continuing industry activity and customer relationships help Moody's refresh its data and improve its models over time. Regulatory experience and thousands of domain experts add

credibility in markets where analytical errors carry financial consequences. AI appears to be increasing the value of proprietary data franchises rather than diminishing them.

AI is compressing complex risk workflows from days to seconds. Insurance submissions arrive through inconsistent spreadsheets, PDFs and free-form documents that must be converted into structured data before pricing and portfolio analysis can begin. Complicated submissions may require days of manual work, while Moody's demonstrated AI reducing the process to seconds. The system must interpret nonstandard policy terms, classify exposures and encode information into insurance-specific taxonomies rather than simply extract text. Errors can affect pricing, loss estimates and portfolio concentration, making domain expertise and validation essential. High-value enterprise AI adoption is increasingly centered on automating complex workflows rather than conversational assistance.

Vertical AI benefits from specialized context and validated domain models. Moody's insurance platform represents information through risk-native concepts such as policies, portfolios, contracts and exposures instead of generic data fields. Agents can combine this context with established catastrophe, pricing and simulation models to analyze portfolios and loss scenarios. Narrowing the problem to a defined risk domain may reduce hallucinations, resolve ambiguity and lower token consumption, although Moody's is still evaluating these benefits. Generative AI therefore complements rather than replaces the scientific models used in high-stakes risk decisions. Enterprise AI should increasingly reward deep contextual understanding and model integration over foundation-model breadth alone.

Enterprise analytics is shifting from point-in-time information toward continuous decision support. Customers increasingly want continuous monitoring, proactive recommendations and actions embedded inside existing workflows instead of static reports. AI agents can connect new information with relevant models and help users determine what decision should follow. This transition makes existing data and analytics assets more actionable without requiring customers to leave the applications where they already work. It also expands the role of analytics providers from delivering model outputs toward supporting recurring operational decisions. Systems of analysis could therefore evolve into systems of decision support and embedded action.

Focus Investment Ideas Leveraged to AI

Iron Mountain (IRM, Buy, $153 12M PT, $37bn market cap, covered by George Tong): IRM's growth portfolio spanning data centers, ALM and digital solutions increased more than 50% y/y in 2Q and now represents 35% of total revenue. Data center revenue increased 39% y/y, supported by 110MW of YTD leasing and 325MW of capacity scheduled to energize over the next 24 months. Signed leases support data center growth of 25%+ in 2026 and 20%+ in 2027, providing visibility as hyperscale AI investment expands. AI infrastructure investment also supports ALM, where revenue increased 88% y/y as hyperscalers decommission equipment, and digital solutions, which maintained 25%+ growth as customers digitize information and adopt IRM's AI-enabled DXP platform. IRM therefore provides exposure to multiple stages of the AI investment cycle, including compute infrastructure, equipment lifecycle management and enterprise information workflows. Durable records cash flows, a visible data center pipeline and accelerating growth businesses create an attractive combination of downside support and secular AI exposure.

Moody's (MCO, Buy, $593 12M PT, $87bn market cap, covered by George Tong): MIS revenue increased 25% y/y in 2Q as AI investment, infrastructure funding, private credit and refinancing supported issuance, with hyperscalers issuing more debt YTD than during the prior three years combined. In MA, ARR growth accelerated to 9% y/y from 8% in 1Q as customers embedded Moody's proprietary data, models and analytics into banking, insurance, KYC and risk workflows. Differentiated information should generate greater value when combined with domain expertise, validated models, customer context and direct workflow integration, all of which are core MCO strengths. MCO has established more than 100 MCP and Smart API connections and deployed AI-enabled compliance solutions that reduced false-positive alerts by approximately 50%. Incremental subscription fees for MCP access and separate charges for agent usage should allow MCO to participate directly in higher data consumption and workflow activity. MCO offers differentiated AI leverage through both increased financing activity in MIS and recurring, consumption-driven monetization in MA.

MSCI (MSCI, Buy, $671 12M PT, $41bn market cap, covered by George Tong): Recurring subscription revenue growth accelerated to 9.0% y/y in 2Q from 8.6% in 1Q, supported by Index, Private Assets and expanding demand across the broader trading ecosystem. Hedge funds generated 19% subscription run-rate growth and 75% recurring net new sales growth as trading-oriented customers increased adoption of custom indices, factor data and systematic investment tools. AI is accelerating MSCI's product-development cycle, with more than 80 launches over the past two quarters, approximately twice the number introduced during all of 2025. The contribution from new products to sales increased approximately 40% y/y in 1H 2026, and IndexAI Insights surpassed 1,000 clients following its February 2026 launch. AI can increase the value of proprietary information by improving discovery, accelerating product creation and embedding data more directly into customer decisions. Broader distribution and faster innovation provide a credible path for recurring subscription growth to accelerate modestly from the high-single-digits over the next 2-4 quarters.

S&P Global (SPGI, Buy, $528 12M PT, $128bn market cap, covered by George Tong): Hyperscaler issuance reached approximately $169bn in 1H 2026, prompting SPGI to increase its full-year outlook to $250-300bn from an initial assumption of approximately

$200bn. AI adoption is also accelerating across SPGI's data businesses, with MCP-connected and LLM-ready API customers surpassing 500, up more than 70% q/q. API call volume exceeded 1.1bn, more than 5x q/q, demonstrating how agentic workflows can materially increase consumption of accurate, machine-readable information. Customers using AI solutions are generating ACV growth approximately 60% faster than the broader MI customer base and approximately three times faster in Energy, with approximately 15% of LLM-ready API customers representing new or returning clients. Revenue conversion continues to trail usage growth as customers evaluate returns and SPGI develops consumption-based pricing, incremental licensing and other monetization models. AI-related issuance, expanding customer acquisition and rapidly increasing data consumption create several paths for engagement to translate into organic growth and margin expansion.

Thomson Reuters (TRI, Buy, $120/C$166 12M PT, $46bn market cap, covered by George Tong): AI-enabled products represented 32% of ACV in 2Q 2026, up from 30% in 1Q and 22% in 2Q 2025, as Westlaw Advantage and CoCounsel drove upgrades across law firms of all sizes. Legal organic revenue growth accelerated to 10% y/y and remained 11% excluding government, providing evidence that customer AI spending is primarily incremental to core subscriptions. TRI's authoritative legal and tax content, professional expertise and integration into established workflows differentiate its applications from horizontal AI alternatives. These assets allow TRI to generate verifiable outputs, preserve customer context and embed AI into consequential research, drafting and compliance processes. Big 3 organic growth accelerated to 10% y/y from 9% in 1Q, prompting TRI to increase its full-year outlook to 9.5%-10% despite continued AI investment and near-term severance costs. TRI demonstrates that incumbent information platforms can defend their franchises and capture incremental spending when authoritative content is paired with deeply embedded, verifiable workflows.

Microsoft (MSFT, Buy, $640 12M PT, $3.6tn market cap, covered by Gabriela Borges): We see Microsoft as well positioned as the industry focus shifts from training and infrastructure to the early stages of 'how to make AI work in enterprises.' Microsoft delivers the AI solutions and integration that enterprises want and need, leveraging on its 50-year legacy of being the premier enterprise Tech company in the world. Microsoft can provide the layer that marries AI models with business-specific context, and dynamically route across AI models to drive cost in any given performance envelope. June quarter results provided a meaningful step in reversing the stock's multi-quarter period of underperformance, as Microsoft delivered tangible proof points across several key investor debates, including: Azure acceleration; improving AI unit economics; and increasing evidence of Copilot monetization. Look for improved monetization to drive EPS growth acceleration over the next three years, from 12% in FY27E to >20% in FY29E.

MongoDB (MDB, Buy, $453 12M PT, $34bn market cap, covered by Matt Martino): We view MongoDB as a well-positioned data platform with multiple paths to sustain Atlas growth above current market expectations over the next several years. The core business is benefiting from stronger enterprise execution, several quarters of record Atlas customer additions and improving developer activity, while an expanding cohort of AI-native customers and model labs provides an incremental source of high-growth consumption. MongoDB's document architecture allows developers to manage complex and rapidly changing data within a flexible operational database, while integrated

search, vector search and retrieval capabilities expand the platform's role across modern application development. This positioning should become increasingly valuable as AI moves into production and applications generate more queries, manage more context and write more information back into enterprise systems. We see MongoDB as a later-cycle beneficiary of the AI investment cycle, with AI-related consumption broadening from early adopters into the company's much larger enterprise installed base and supporting the durability of Atlas growth over time. Earnings preview here.

Rubrik (RBRK, Buy, $106 12M PT, $20bn market cap, covered by Matt Martino): We view Rubrik as a high-growth cyber resilience platform with a strong core recovery franchise and an expanding set of growth vectors across identity, data security and AI governance. Continued displacement of legacy backup products supports durable execution in the core, while increasing identity attach and early engagement around Agent Cloud broaden Rubrik's relevance across CIO and CISO budgets. AI strengthens this opportunity by increasing data creation, accelerating the proliferation of non-human identities and enabling attacks and operational errors to occur at machine speed. Rubrik's ability to understand relationships across enterprise data, identities and applications positions the platform to help customers identify risk, enforce controls and recover cleanly following an incident. Identity Resilience and Agent Cloud remain earlier in their development, but both extend naturally from Rubrik's architecture and established role in recovery. We believe this combination of resilient core execution, scaling adjacent products and increasing strategic relevance in an AI-driven threat environment provides several paths for Rubrik to sustain premium growth as the business scales. Earnings preview here.

Valuation & Risks

Iron Mountain (IRM, Buy, covered by George Tong): Our 12-month price target of $153 is based on 23.0x our NTM + 1YR AFFO/shr estimate of $3.22. Our target multiple is in-line with the median multiple of 23.3x for self-storage and industrial REIT peers reflecting attractive organic growth trends as the company's growth portfolio continues to expand, mitigated by IRM's transformation from a legacy records management company into a data and technology company that comes with execution risk and medium-term financial uncertainty. Risks to the downside for IRM shares include competing technologies in records management, acquisition integration risk, financial leverage and FX risk.

Moody's (MCO, Buy, covered by George Tong): Our 12-month price target of $593 is based on 29.0x our NTM + 1YR EPS estimate of $20.43. Our target NTM P/E multiple comes above our Info Services peer group median of 18.1x, reflecting Moody's wide competitive moat, leverage to strong secular tailwinds in ratings, strong positioning with AI and attractive growth outlook in the MA segment. Risks to the downside include macro uncertainty, a slowdown in M&A volumes, higher interest rates, buy-side and sell-side budget pressures, ineffective execution of an AI strategy, and FX risk.

MSCI (MSCI, Buy, covered by George Tong): Our 12-month price target of $671 is based on 27.0x our NTM + 1YR EPS estimate of $24.86. Our target NTM P/E multiple comes above the Info Services peer group median of 18.1x due to MSCI's stronger revenue growth and higher EBITDA margins. Risks to the downside include AUM-based fee declines in a contracting market, a rotation of fund flows back to the US, declines in buy-side budgets, reliance on third-party data, market share loss due to competition in Index and Analytics, and the potential for customers to in-source spend on index, analytics, ESG and real estate solutions.

S&P Global (SPGI, Buy, covered by George Tong): Our 12-month price target of $528 is based on 24.0x our NTM + 1YR EPS estimate of $21.99. Our target multiple is higher than the Info Services peer group median of 18.1x reflecting S&P's strong track record of execution, wide competitive moat due to its proprietary and essential data, and higher EBITDA margins relative to peers. Risks to the downside include a moderation in debt issuance volume growth, financial services budget and headcount pressures that could weigh on Market Intelligence growth, a market pull-back that could put downward pressure on Indices growth, and FX risk.

Thomson Reuters (TRI, Buy, covered by George Tong): Our 12-month price target of $120 is based on 22.0x our NTM+1YR EPS estimate of $5.45. Our TRI.TO 12-month price target of C$166 is based on our TRI price target and a USD/CAD exchange ratio of 1.3842. Our target multiple comes above the median Info Services NTM P/E multiple of 18.1x due to TRI's resilient business growth characteristics, wide competitive moat and strong track record of execution. Risks to the downside include intensifying competition from AI-native companies, macro headwinds that could pressure customer budgets, a slower-than-expected pace of customer adoption for the company's AI-enabled products, elevated investment spending that could cause an EBITDA margin shortfall, and FX risk.

Microsoft (MSFT, Buy, covered by Gabriela Borges): Our 12-month price target of

$640 is based on a 28x target P/E multiple of adjusted NTM net income. Risks include: execution in Copilot, longer ramp for internal silicon which can limit market share gains or gross margin expansion, greater investments in projects outside of expectations (e.g., non-Azure), and a more meaningful shift to custom software that could negatively impact its applications business.

MongoDB (MDB, Buy, covered by Matt Martino): Our 12-month price target of $453 is based on a two-pronged valuation framework with 50% weight to a target EV/Sales multiple and 50% weight to a target EV/FCF multiple. Our EV/FCF valuation is based on 52x our Q5-Q8 FCF estimate and EV/S based on 10.5x our Q5-Q8 Revenue estimate. Key risks include: increased competitive pressure from relational database incumbents, particularly from Hyperscalers like Oracle, Microsoft, AWS, and GCP; broader standardization around Postgres impacts MongoDB's win rates and expansion opportunities; slowing pace in cloud migration, multi-cloud adoption, digital transformation, and application modernization; AI cycle proves to be a less meaningful tailwind than initially contemplated; margin degradation due to required reinvestment or unanticipated headwinds from the EA business; adverse changes in the IT spending environment.

Rubrik (RBRK, Buy, covered by Matt Martino): Our 12-month price target of $106 is based on a two-pronged valuation framework with 50% weight to a target EV/S multiple and 50% weight to a target EV/FCF multiple. We apply a 48x EV/FCF multiple to our Q5-Q8 FCF estimate and 12x EV/S multiple to our Q5-Q8 Revenue estimate. Our EV/FCF and EV/S multiples are informed by a blend of high-growth platform and security peers. Key risks include: heightened competition from Cohesity, Commvault and Veeam, among others, limits Rubrik's ability to capture a sizable portion of the legacy Backup and Recovery market; industry consolidation, akin to Cohesity combining with Veritas' data protection business, leading to better-resourced peers with improved distribution footprints and more competitive product offerings, impacting Rubrik's win rates; elevated product development and sales and marketing expenditures hinder further operating leverage; slower-than-expected adoption of next-generation backup and recovery solutions limits Rubrik's addressable market opportunity.