PitchBook
a Morningstar company

Institutional Research Group

Derek Hernandez
Senior Research Analyst,
Enterprise SaaS and
Infrastructure SaaS
derek.hernandez@pitchbook.com

pbinstitutionalresearch@pitchbook.com

Published on January 14, 2026

Contents

EMERGING TECH RESEARCH

Clash of the Titans

Incumbents versus challengers in the age of agentic AI

PitchBook is a Morningstar company providing the most comprehensive, most accurate, and hard-to-find data for professionals doing business in the private markets.

Key takeaways

- The enterprise SaaS sector is undergoing its most significant technological shift in a generation, driven by the maturation of artificial intelligence. This report provides a framework for private market investors to navigate this transformation, which is defined by a central dichotomy: AI-embedded incumbents versus AI-native challengers.

- While the enterprise adoption of AI is high (78% of organizations), meaningful business outcomes are not (95% of pilots are failing to accelerate revenue).[1] This execution gap creates a massive opportunity for a new class of vendors.

- This market is bifurcating. Incumbents are embedding AI copilots into legacy suites, leveraging their vast distribution and customer trust. In contrast, AI-native challengers are building intelligent systems from the ground up, creating new, highly defensible moats.

- For investors, we argue that the most durable value will not come from the AI models themselves, which are becoming commoditized. Instead, the new, defensible moats are being built on proprietary data pipelines that create a virtuous feedback loop, agentic orchestration that automates entire and complex workflows, domain-specific tuning that provides verifiable accuracy, and auditable control planes that satisfy enterprise governance and compliance needs.

- This report analyzes the market dynamics, investment landscape (including a $65 billion TAM set to grow to $190 billion by 2030), and competitive strategies of both incumbents and challengers. It concludes with an actionable playbook and diligence checklist for LPs, GPs, and founders to identify the true, defensible category leaders in the new AI-native world.

1: "State of AI in Business 2025," MIT NANDA, Aditya Challapally, et al., July 2025.

The radical transformation in SaaS

The enterprise software-as-a-service (SaaS) sector is undergoing its most significant technological shift in a generation, even larger than the transformation from perpetually licensed products to SaaS itself. This shift is driven by the ongoing maturation of AI, especially the advancements by major large language models (LLMs) and their rush to deployment across nearly every solution within enterprise SaaS. Excluding massive rounds by LLMs, we see this trend driving steady and sustained growth across enterprise SaaS investments since early 2024. All datasets herein exclude these recent and discrete mammoth AI funding rounds.

Enterprise SaaS VC deal activity by quarter

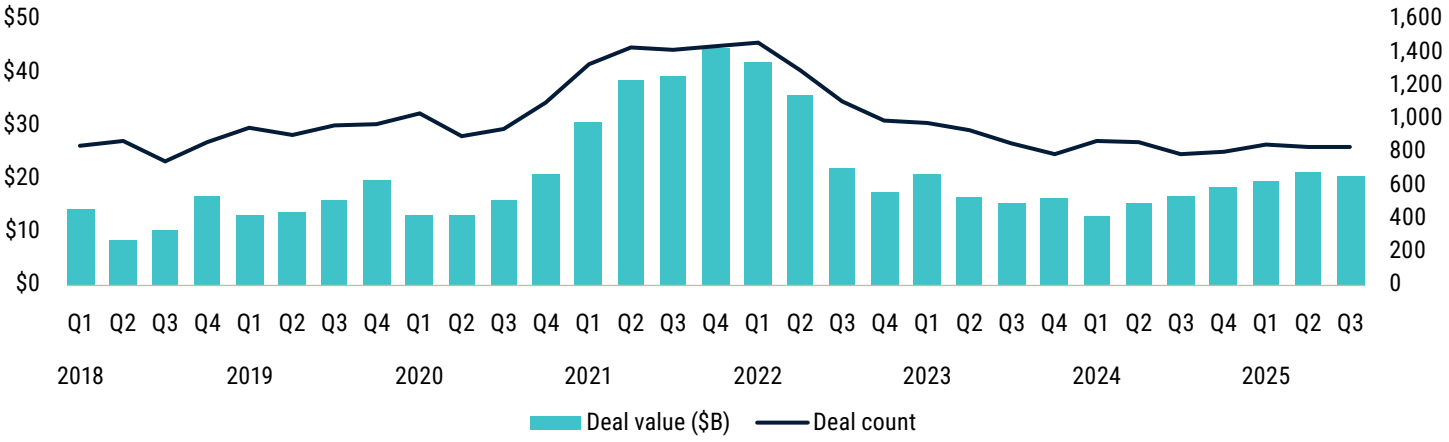

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Thus, this is the first in a series of reports on AI within enterprise SaaS, beginning with a high-level overview of the state of play today. In our future pieces, we will dive into AI in specific sectors and subsectors, including AI in customer relationship management (CRM), AI in HR tech, and others. We invite PitchBook clients to reach out with specific sectors of the enterprise SaaS landscape they would like us to cover in this series as well. Please direct suggestions to derek.hernandez@pitchbook.com.

The AI transformation of enterprise SaaS has sorted the segment into a critical dichotomy: the distinction between retrofitting existing systems with AI features (AI-embedded platforms) and designing new operations from the ground up around intelligent, autonomous processes (AI-native platforms). This choice is a fundamental strategic decision that will balance established distribution, defensibility, and compliance with competitive durability, operational efficiency, and evolving revenue streams for years to come.

The proliferation of generative AI has forced every enterprise software company to formulate an AI strategy. However, not all AI is created equal. The architectural approach an organization takes—either building new systems with AI at their core or layering AI onto existing platforms—is the single most important determinant of long-term success and value creation. This distinction creates a clear framework for market analysis and investment due diligence.

AI-embedded incumbents

AI-embedded solutions, thus far the domain of legacy SaaS incumbents, offer additional productivity gains within established solutions and platforms. These incumbents within enterprise SaaS are embedding copilots and task agents into existing suites, spanning CRM, enterprise resource planning (ERP), supply chain management (SCM), HR, collaboration, IT service management (ITSM), and cybersecurity. They include Salesforce's programmatic push to Agentforce, Microsoft's various Copilot integrations across Microsoft 365 and Dynamics 365, ServiceNow's Now Assist, and SAP's Joule. Across these cases, the copilot or agent is natively embedded, grounded in the customer's enterprise data, and increasingly able to take actions inside the system of record. This provides immediate benefits to established platforms with global distribution. According to Gartner, 80% of software vendors will have embedded AI into their applications in this manner by 2026, making it the largest and fastest-growing segment of AI capabilities.[2]

The AI-embedded approach keeps the user inside an incumbent system (ERP, CRM, ITSM, help desk, integrated development environments, and finance stacks) and inserts model-driven decision-making at a specific step. The difference is less about interface aesthetics and more about control over execution and ownership of a measurable outcome. In a copilot-only model, the user remains the primary orchestrator; in an agentic model, the system executes under guardrails and produces artifacts and logs suitable for audit.

However, the impact of this approach is inherently constrained by the underlying architecture of the legacy systems it augments. These systems are often burdened by technical debt, fragmented data silos, and rigid workflows, which work against the kind of transformative change clients are demanding today. Generally, AI-embedded software follows the traditional approach of building upon established systems that users are familiar with and enterprises are often already heavily invested in. The immediate advantages of leveraging legacy systems have driven incumbents to adopt this approach across a range of enterprise SaaS solutions. In nearly all cases, the underlying workflow of manual data entry and pipeline management remains unchanged but is augmented by AI-powered assistance in the form of an embedded LLM solution.

AI-native challengers

In sharp contrast, AI-native platforms are built on a different foundation. An AI-native workflow is a business process conceived and engineered around an agent or model from inception. The agent plans steps, selects tools, reads from and writes to systems of record, and either executes or proposes actions subject to policy. Human supervisors handle exceptions, provide feedback, and approve or revise actions at configured thresholds.

2: "Gartner CFO & Finance Executive Conference: Day 1 Highlights," Gartner, September 11, 2024.

In our view, and in the general view of the market so far, it is the AI-native vendors that are creating a new class of agile, hyperscalable, and possibly peak-efficiency businesses. In an AI-native workflow, a model or agent proposes or executes steps, while human users supervise exceptions and provide feedback, which is incorporated as further training. The consequence of this is that the decision logic, orchestration, and learning loop move into the model layer, and unit economics are measured in resolved tickets, processed invoices, reconciled transactions, or other discrete outcomes rather than the traditional pricing of seats deployed. This reframing guides how today's vendors build, price, and sell their product. It also changes how enterprises evaluate return on investment (ROI) and risk as well as how investors assess future defensibility and category leadership.

Today, AI-native workflows are transforming core enterprise functions, including Adept's agents that use software, Runway's multimodal creative automation, and Glean's enterprise knowledge assistant.

In addition to these horizontal approaches, there are strictly vertical solutions, particularly function-specific or embedded agents. These target narrow, high-leverage decision points inside incumbent systems, including ticket triage, invoice coding, spending anomaly detection, reconciliations, knowledge retrieval, and action drafting. These solutions follow a typical approach of integrating within established systems, owning a measurable step, showing a speedy payback, and expanding to adjacent steps. These include Ramp's Intelligence features, Forethought's support triage and autoresponse, Zendesk's agentic additions to its service stack, Intercom's Fin AI agent, and Vic.ai's autonomous accounts payable solution. We continue to see these investments focus on North America, particularly the US in 2025, where this region represented 68% of total investment. This is the highest percentage since 2018, even excluding raises by the foundational LLM companies.

Share of enterprise SaaS VC deal value by region

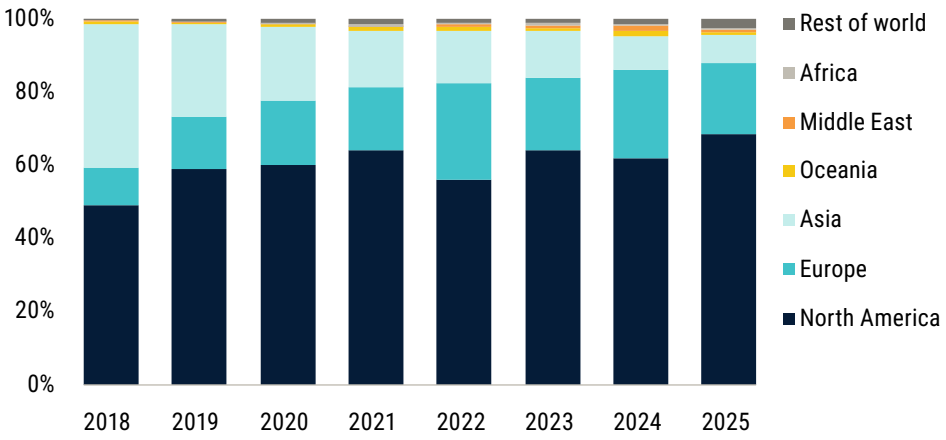

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

AI-embedded versus AI-native platforms overview

	AI-embedded platforms (incumbents)	AI-native platforms (challengers)
Core architecture	Legacy systems and platforms with AI solutions deeply integrated	Platforms where AI is the foundational, operational layer, built from the ground up
Data strategy	Data may be accessed across fragmented departmental silos within existing enterprise frameworks	Requires developing a unified "data fabric" as a strategic asset that continuously feeds learning models
Workflow impact	Automates discrete tasks within a human-led process, learning and adopting best practices to develop further autonomy	Ideally replicates entire outcomes with autonomous, adaptive workflows, first including human-in-the-loop reinforcement
Primary value proposition	AI solutions deployed within familiar frameworks to drive efficiency gains, cost reductions, and greater adoption	AI solutions with competitive differentiation, intended to produce new revenue streams and ultimately create market disruption
Performance and scalability	Potentially limited by legacy constraints, though deliver incremental ROI from day one and can offer speedier adoption	Designed for exponential scale and continuous learning with performance gains 2x to 5x greater for AI workloads
Key risks	Accumulating technical debt, falling behind the competitive curve, and being disrupted by nimbler challengers	High initial investment, implementation complexity, "black box" opacity, and a higher failure rate for projects

Source: PitchBook

Market development and investment

This new market is defined by a paradox: explosive growth in investment and adoption set against a backdrop of significant technical and practical hurdles. Enterprise adoption of AI has accelerated dramatically. According to Stanford's 2025 AI Index Report, 78% of organizations reported using AI in at least one business function in 2024, a significant increase from 55% the previous year.[3] A McKinsey survey from early 2025 confirms this momentum,[4] finding that over three-quarters of respondents' organizations are now using AI.

However, this widespread adoption has also exposed a significant execution gap. A sobering 2025 report from MIT reveals that a staggering 95% of business attempts to integrate generative AI are failing to achieve meaningful revenue acceleration.[5] This aligns with McKinsey's finding that, while 92% of companies plan to increase their AI investments, only 1% of leaders describe their organizations as "mature" in their AI deployment, where the technology is fully integrated into workflows and driving substantial business outcomes.[6]

This gap between high ambition and poor outcomes has not deterred investors; in fact, it has fueled a surge in funding for solutions that promise to deliver real value.

3: "The 2025 AI Index Report," Stanford HAI, Yolanda Gil and Raymond Perrault, November 2025.
4: "The State of AI in 2025: Agents, Innovation, and Transformation," McKinsey & Company, Alex Singla, et al., November 5, 2025.
5: "State of AI in Business 2025," MIT NANDA, Aditya Challapally, et al., July 2025.
6: "Superagency in the Workplace: Empowering People to Unlock AI's Full Potential," McKinsey & Company, Hannah Mayer, et al., January 28, 2025.

This is fueling major investment in these solutions, as AI-related VC deals accounted for 64.3% of all VC deal value in 2025 through Q3, up from just 12% in 2017. Private equity, while more cautious, is aggressively targeting the essential data infrastructure required to power this revolution.

AI SaaS deal activity has accelerated from 2023 through 2025, with significant later-stage financings signaling the emergence of potential category leaders. Adept announced a $350 million Series B in March 2023. Runway extended its Series C by $141 million in June 2023. Glean raised a $150 million Series F at a $7.2 billion valuation in June 2025. On the embedded-agent side, we would note Ramp's July 2025 round at $22.5 billion and Cresta's November 2024 $125 million Series D. These demonstrate that growth-equity and strategic interests are aligned, especially where outcomes are measurable. We see this driving investments across all major sectors within enterprise SaaS, as figures from 2025 through Q3 nearly matched or surpassed full-year 2024 numbers.

Analytic platforms VC deal activity

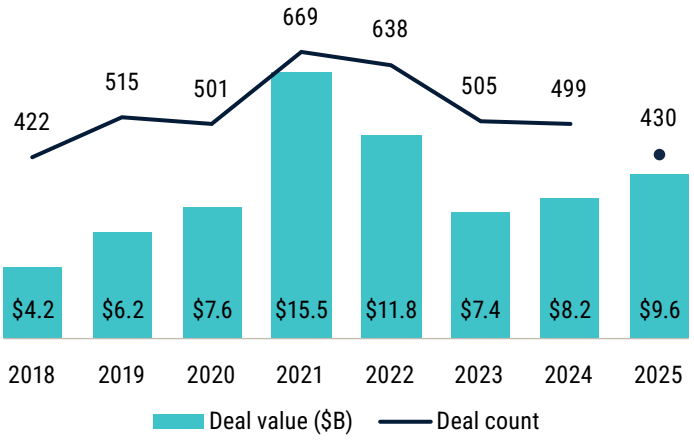

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Customer relationship management VC deal activity

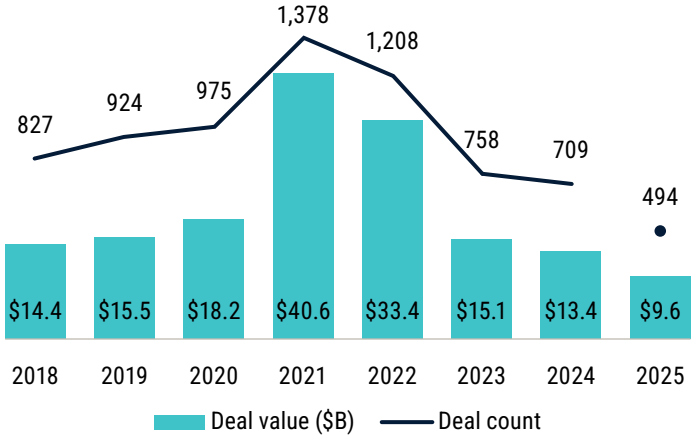

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Enterprise resource planning VC deal activity

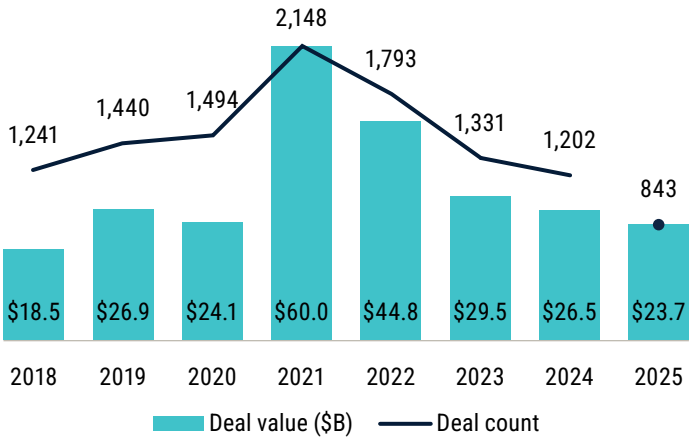

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Supply chain management VC deal activity

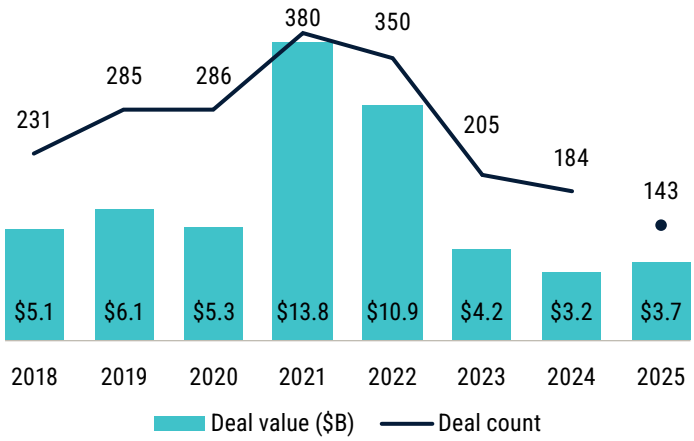

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Knowledge management systems VC deal activity

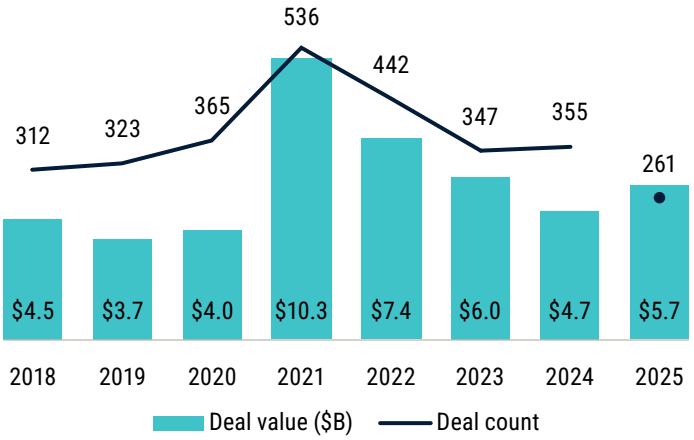

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Other application software VC deal activity

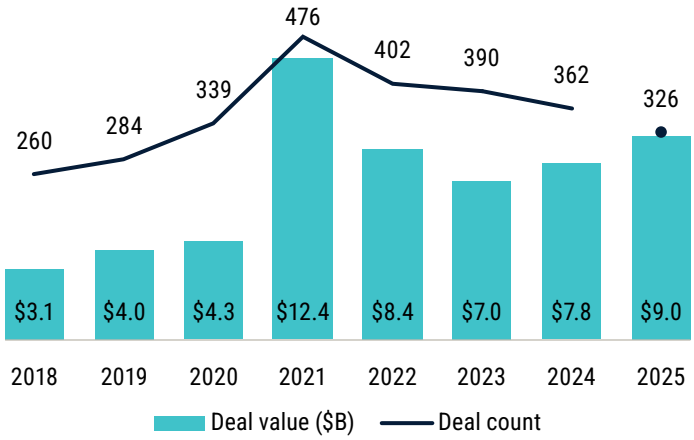

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Market sizing and growth

Market sizing for AI-native and AI-embedded workflows is tricky to pin down because these AI products are often replacing existing solutions within vendors. Additionally, pricing is itself shifting from seats to usage and outcomes, though not in every case. We size this market through estimates of existing software spend per function (across analytic platforms, CRM, ERP, KMS, SCM, and other application software), replacement or augmentation rates through 2030, and net-new spend created by agents that own actions (for example, per document, per resolution, or per transaction). We estimate a 2025 total addressable market (TAM) across six segments at $65 billion, growing to about $190 billion by 2030. Growth rates range from around 18% to 27% over the five years, with the fastest growth in vertical agentic AI and AI-native productivity, where vendors replace full workflows over aiding existing human-centered workflows.

Enterprise SaaS AI workflow TAM

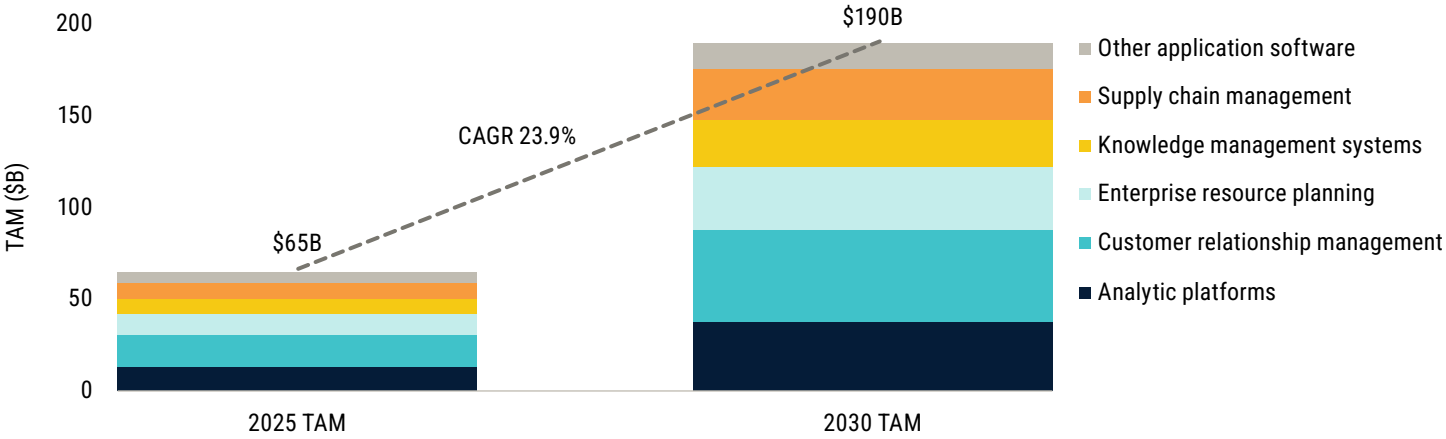

Source: PitchBook • Geography: Global • As of September 30, 2025

Today, we have already seen concrete announcements by major incumbents sketching out much of this growth path. SAP has publicly discussed hundreds of embedded AI use cases around Joule across S/4HANA and SuccessFactors,[7] suggesting a pipeline of agentic tasks being pushed into existing enterprise contracts. ServiceNow continues to treat the Generative AI Controller as a core layer, implying attach-rate revenue as models route through Now Platform flows. Salesforce is now on its third Agentforce iteration in 12 months (with its June 2025 update centered on observability and control), indicating the scale of Salesforce's investment in easing deployment friction. While not revenue disclosures, these examples show a sustained roadmap by incumbents that control large footprints across the enterprise landscape, which we would expect to drive AI attach rates over time even if near-term adoption is uneven and at times doubted, as in the reporting around the aforementioned MIT study of a 95% failure rate.[8]

Enterprise SaaS AI workflow TAM by sector

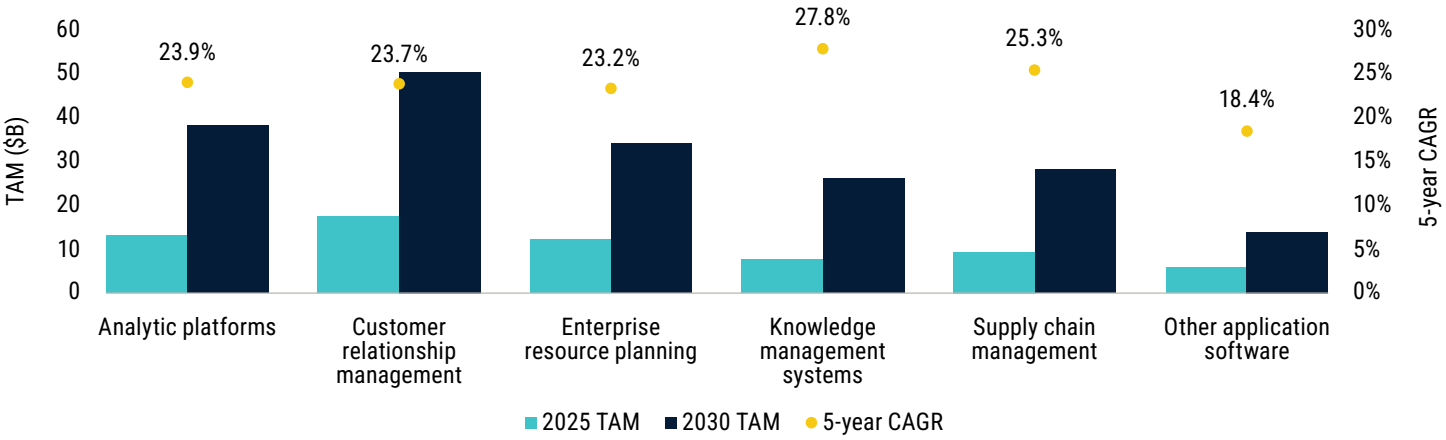

Source: PitchBook • Geography: Global • As of September 30, 2025

Currently, we identify two forces that are most likely to restrain this growth. First, pricing confusion across copilot products and agent skills can slow procurement. Salesforce's June 2025 list-price updates and analyst commentary about "decision fatigue" after Agentforce releases show that enterprises will pause purchases if value capture is unclear.[9] Second, evaluation and audit readiness remain uneven. Enterprises are becoming stricter about lineage, model traceability, and fallbacks. Vendors that package evaluation harnesses and admin policies reduce the time to production today and we believe will likely be overrepresented in 2026-2027 enterprise wins.

Opportunities and constraints

This market growth is supported by key technological enablers but is also held in check by significant, practical constraints.

7: "SAP Business AI: Release Highlights Q1 2025," SAP, Philipp Herzig, April 7, 2025.
8: "State of AI in Business 2025," MIT NANDA, Aditya Challapally, et al., July 2025.
9: "Salesforce AI Agent Adoption Hits a Hurdle. Customers Have 'Decision Fatigue,' Analyst Says," MSN, Tae Kim, August 18, 2025.

On the opportunity side, several factors are accelerating adoption. Model access and cost curves have moved to where multimodel routing and task-specific fine-tuning are viable within enterprise budgets, allowing users to route by task rather than choosing a single model. Data access patterns have also matured, with retrieval-augmented generation, lineage tracking, and permission-granting moving from demos to production. This is exemplified by vendors like ServiceNow, which now offer a "Generative AI Controller" layer. Finally, interfaces have shifted to intent capture, making the "ask the system to do the work" concept familiar and auditable.

Adoption is also supported by measurable time-to-value metrics in specific functions. In developer tooling, GitHub has published enterprise-focused studies with Accenture indicating strong adoption and at least directional productivity and satisfaction gains.[10] Meanwhile, Google and other external surveys during 2025 found 90% of engineering teams report incorporating AI coding tools and GitHub Copilot is the modal choice.[11]

Conversely, we identify several meaningful constraints that recur in due diligence. A primary hurdle is governance and lineage, as enterprises need to know which data was used, which model, and how the system behaves under policy guidance. Reliability and regression are also major concerns; without evaluation harnesses, model changes can silently degrade outcomes. Change management is another barrier, as agentic actions require redesigning roles and incentives. Pricing clarity also remains a point of friction, with "copilot sprawl" creating confusion. Vendors that price per action, like Intercom's explicit "per resolution" fee, are having fewer procurement debates because the payback math is straightforward. Lastly, vertical constraints in areas like law and healthcare mean that generalized assistants rarely suffice, requiring domain-specific tuning and compliance.

Defensibility and the tech stack

To win in this new market, both incumbents and challengers must build on a new, AI-first foundation. Understanding this new architecture is key to identifying durable value.

A useful way to interpret vendor roadmaps is to map the four layers that now define an AI-enabled enterprise application. At the base is the data plane—connectors, permissions, governance, lineage and retrieval. This layer supplies the grounding to keep outputs anchored in enterprise reality. Above it is the model/agent layer, which plans, chooses tools, evaluates, and executes under policy. The third layer is the workflow layer—composable steps, observability, rollback, service-level agreements, and escalation interfaces. At the top is the experience layer, which captures intent (such as chat, forms, and API functions) and presents results and artifacts with the right context. In a bolt-on world, the AI sat near the top. In an AI-native world, the model/agent and workflow layers are the core of the product, and the control plane (encompassing identity, policy, lineage, and evaluation) becomes a moat.

10: "Quantifying GitHub Copilot's Impact in the Enterprise With Accenture," GitHub, Ya Gao, May 13, 2024.
11: "Google Says 90% of Tech Workers Are Now Using AI at Work," CNN, Lisa Eadicicco, September 23, 2025.

Salesforce's public materials framed the shift from Einstein Copilot to Agentforce as moving from answers to actions—retrieving data, reasoning, planning, then executing tasks—directly in Salesforce. ServiceNow's Now Assist highlights domain-specific modeling and ties directly into ServiceNow's workflow engine, while its Generative AI Controller exposes model routing and governance as first-class features. SAP's Joule describes a system of collaborative AI agents operating across supply chain, procurement, and finance, "grounded in trusted SAP knowledge,"[12] which aligns with the described data-to-agent-to-workflow stack. Microsoft's Copilot in Dynamics 365 and Microsoft 365, backed by updated documentation for specific apps (finance and operations), further normalize agentic actions within the record systems where audits matter most.

From an investment perspective, incumbents' advantages are distribution, single sign-on, and permissions integration, and credibility with information security and audit committees. Their constraints are legacy data models and user-interface paradigms that can slow the transition from assistive to agentic usage. The near-term revenue expression is add-on copilot products and uplift on existing modules. The medium-term strategic question is whether these vendors can demonstrate outcome-based pricing and real task ownership at scale—for example, measurable reductions in time-to-resolve metrics, close cycles, or exception rates—without overwhelming enterprises with overlapping products.

The new moats

In our view, both AI-native and AI-embedded products suffer much of the same risks around moats and defensibility. It is simply not possible to rely on access to a popular model to defend hard-won ground. Nor may a category leader build a simpler wrapper that proves easily replicable today. Instead, the most defensible are building control-plane moats with four major anchors emerging.

The first and arguably most important anchor today is data capture and rights, where systems capture high-fidelity task data, including inputs, intermediate states, decisions, outcomes, and secure use rights to improve models. This creates a proprietary improvement loop that others cannot easily replicate and has the potential to create a virtuous cycle of gaining customers and creating a better product through greater use, therefore gaining more customers. Specifically, the AI (native or embedded) application is used within a specific workflow, which generates proprietary data that is then used to provide context to the AI model, improving the quality and relevance of its output. This improved output makes the application more valuable and stickier, which under the right circumstances encourages more usage by the employee base. Importantly, a new entrant, even with a similar or more advanced technological solution, would not have access to the accumulated historical data, and an incumbent in an adjacent market would struggle to capture the same workflow-specific context.

The second anchor is agentic orchestration, including planning, tool-use, safety policies, and rollback under audit, supported by a test harness that simulates workflows against regressions, particularly when models, prompts, or tools change. The third is domain

12: "Joule," SAP, n.d., accessed January 6, 2026.

tuning, meaning the retrieval schemata, prompt libraries, and selective fine-tuning actions or adapters that encode domain semantics and produce measurable accuracy deltas and lower human-in-the-loop load. Examples of specific domain schematics include legal citation chains or healthcare risk-adjustment models. The fourth and final anchor is, as always, observability, including traces, cost and latency telemetry, and policy compliance views that are now part of the product instead of just in the development toolchain.

Incumbents and challengers meet

These new moats are the weapons in a strategic battle between legacy incumbents and AI-native challengers. Both are leveraging the new stack but in fundamentally different ways.

For enterprises, there are several critical needs for these solutions to succeed and gain traction. Data access and lineage must be explicit, retrieval and fine-tuning require observability, agent behavior demands evaluation suites, and all actions need to be fully auditable by users.

Incumbents are uniquely positioned to meet this need, turning governance into a core product feature. Recent releases highlight this observability pivot. Salesforce's Agentforce announcement emphasized a Command Center for "visibility and control."[13] ServiceNow's Generative AI Controller documents retention and integration patterns. SAP's Joule framing stresses "grounding in business data."[14] These materials describe the minimum product surface enterprises are asking for before granting an agent access to write a system of record. Salesforce, ServiceNow, and SAP have all publicly framed their strategies as moving from simple answers to governed, auditable actions within their core platforms. This has included massive acquisitions by many incumbents. These have kicked off a resurgence in exit value reported in Q3 2025.

Enterprise SaaS VC exit activity by quarter

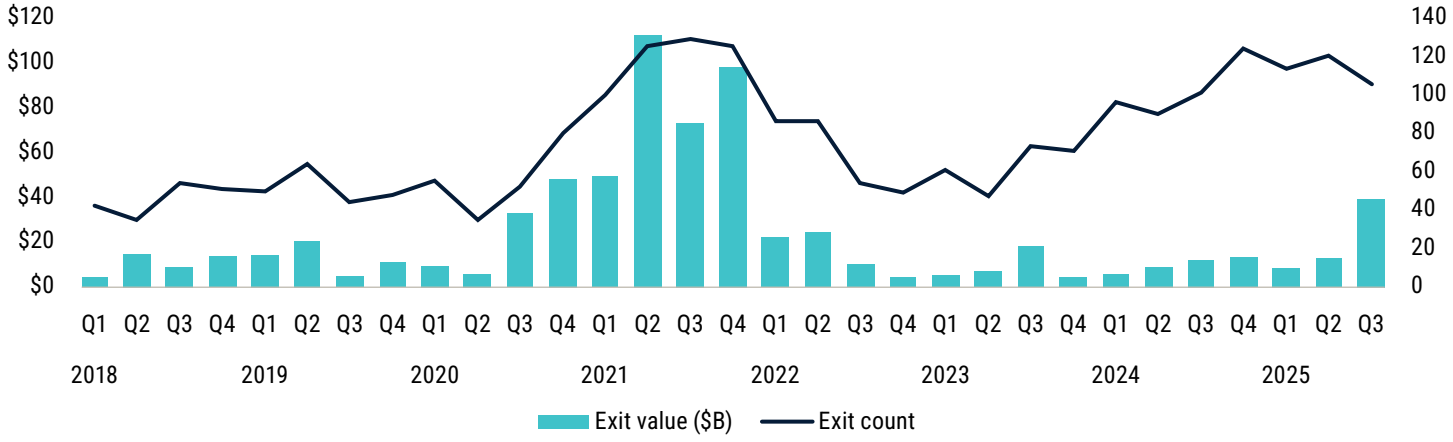

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

13: "Salesforce Launches Agentforce 3 to Solve the Biggest Blockers to Scaling AI Agents: Visibility and Control," Salesforce, June 23, 2025.
14: "Build, Deploy, and Extend AI Agents With Joule Studio," SAP, Bharat Sandhu, May 20, 2025.

AI-native challengers, by contrast, invert the stack. They start with the model/agent layer as the product's core and build only the user experience required to capture intent and supervise exceptions, relying on connectors and APIs for data ingress/egress. Adept's positioning illustrates a design choice to operate across applications, not only within one. Runway's multimodal toolchain treats content generation and editing as agent-coordinated tasks. Glean now positions itself as a work assistant that operates over enterprise knowledge graphs. These approaches trade immediate distribution for speed and clarity of architecture; how quickly they scale into risk-sensitive enterprises depends on how fast they can expose control-plane features.

Competitive interactions are intensifying at the edges. Incumbents have extended native agents into spaces once occupied only by startups (for example, Zendesk's autonomous agents and Microsoft's Copilot in CRM/ERP). Startups respond by deepening scope (owning more steps), tightening governance (shipping evaluation and policy features), and moving to outcome pricing where the payback math is transparent. The result is fewer "copilot curiosities" and more agent deployments tied to specific key performance indicators. Where the wedge is narrow (such as a single step in autonomous platforms), startups must either move horizontally across adjacent steps or vertically into orchestration and analytics or risk being boxed in by the platform vendor's roadmap.

Enterprise SaaS VC exit count by type

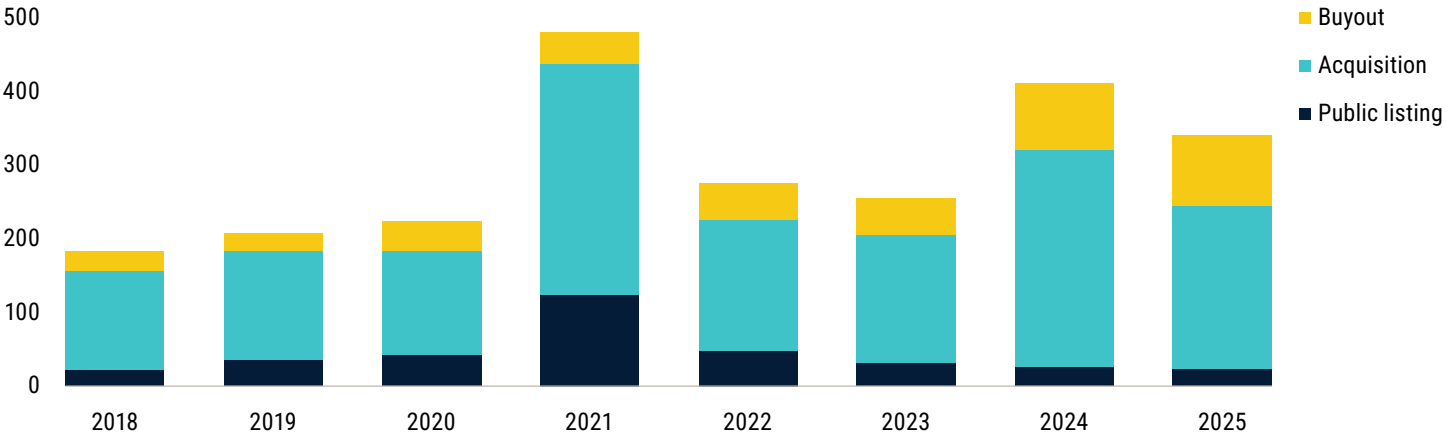

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

The investor perspective

For private markets professionals, this complex and fast-moving landscape requires a new playbook for sourcing, diligence, and value creation.

For investors in the new AI world today, the mandate must be to identify companies that either replace a workflow end-to-end with an AI-native backbone and credible governance or, as a second option, insert an AI-embedded agent into a large incumbent platform and price to measurable outcomes. In the first example, recent financings at Adept and Glean show that growth-stage capital is concentrating on agentic action across apps and work assistants over enterprise knowledge. In the

second group, momentum around Ramp in finance operations, Zendesk and Intercom in support, and Cresta in contact centers shows that embedded agents are gaining market share. Ultimately, these dual approaches both require the same basic building blocks for success, step ownership, control-plane maturity, and distribution leverage.

As always, due diligence must be designed to fully challenge the product. Evaluation suites must be both vendor-supplied evaluation suites and your own. Production telemetry anonymized and aggregated over three months, or more, is an important evaluation, including intervention rates by intent or document class, latency distributions, and cost per action. As mentioned previously, data rights must be verified. Admin consoles for identity and policy mapping should be inspected, and if the product writes to a system of record, a demo of rollback and audit is required. It is useful to see the incumbents' approach as a benchmark. Salesforce's Agentforce Command Center, ServiceNow's model governance, and SAP's Joule positioning all signal that larger enterprises expect visibility and control. We believe that startups without those surfaces are at a market disadvantage.

A pragmatic defensibility checklist for investor diligence across today's AI landscape now includes:

1. Task ownership, specifically the percentage of steps executed autonomously, with human-override thresholds.
2. The human intervention rate and its trend over the last several model releases.
3. Evaluation coverage, meaning the pass/fail outcomes tied to real workloads, not just synthetic prompts.
4. Lineage and audit documentation, including which data, model, retrieval chunk, and policy were utilized.
5. Policy and identity mapping across tool permissions and environment scoping.
6. Distribution hooks, including technology and service integration partner channels. For vertical-SaaS-style AI vendors, this would also include regulatory systems, such as federal privacy laws for healthcare or privilege chains and enforceable citations for legal applications.

Enterprise SaaS median VC pre-money valuation ($M)

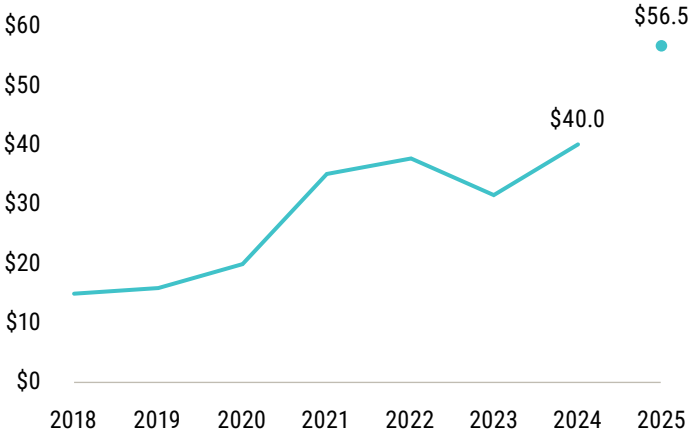

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

Enterprise SaaS median VC deal value ($M)

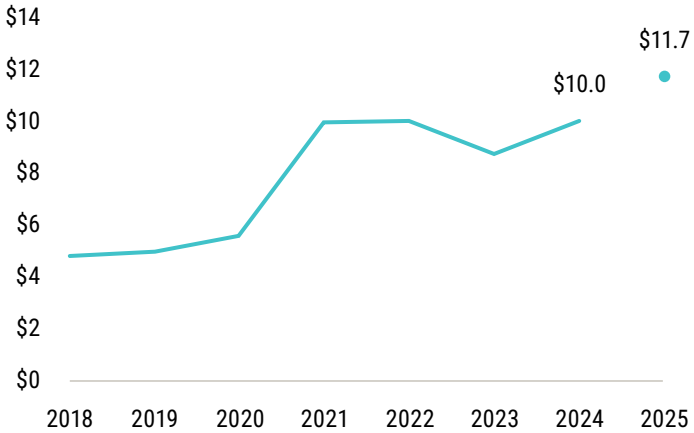

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Data excludes deals for foundational LLM companies (OpenAI, Anthropic, and others).

We would highlight the Harvey-LexisNexis alliance as an example of content and workflow defensibility and Suki's integration with MEDITECH Expanse APIs as an example of system-of-record depth. Beyond specific regulated industries, we also note the rise of highly specialized AI solutions targeting specific, high-value workflows within larger markets. These include Prediko for e-commerce demand planning and Seemplicity for cybersecurity remediation, which are prime examples of companies building deep moats by solving a narrow but critical problem with a novel AI workflow.

Top VC investors in enterprise SaaS companies since 2023

Investor	Deal count	Pre-seed/seed	Early-stage VC	Late-stage VC	Venture growth	Investor type
Alumni Ventures	185	62	51	58	14	VC
Gaingels	173	36	42	79	16	VC
Sequoia Capital	159	59	47	38	15	VC
Andreessen Horowitz	140	29	60	39	12	VC
General Catalyst	135	36	49	32	18	VC
FJ Labs	120	27	42	41	10	VC
Insight Partners	106	8	27	51	20	VC
Accel	101	20	39	30	12	VC
Lightspeed Venture Partners	101	19	45	26	11	VC
Bessemer Venture Partners	93	14	32	37	10	VC
Mana Ventures	79	19	20	31	9	VC
Index Ventures	69	16	33	16	4	VC
Khosla Ventures	68	13	28	20	7	VC
Calm Ventures	68	15	18	31	4	VC
Soma Capital	66	35	25	5	1	VC
Salesforce Ventures	61	9	20	25	7	Corporate VC
BoxGroup	61	23	28	10	0	VC
Menlo Ventures	60	13	23	21	3	VC
GV	60	10	20	22	8	Corporate VC
Flybridge	58	20	20	13	5	VC
Bain Capital Ventures	58	10	30	16	2	VC
SV Angel	53	18	23	9	3	VC
8VC	52	17	17	14	4	VC
Seedcamp	51	24	15	10	2	VC
Liquid 2 Ventures	51	19	15	16	1	VC

Source: PitchBook • Geography: Global • As of September 30, 2025

Top VC investors in enterprise SaaS companies since 2023 (cont.)

Investor	Deal count	Pre-seed/seed	Early-stage VC	Late-stage VC	Venture growth	Investor type
New Enterprise Associates	50	15	13	16	6	VC
Endeavor Catalyst	48	0	10	32	6	VC
First Round Capital	47	18	17	10	2	VC
Pioneer Fund	45	27	14	4	0	VC
Redpoint Ventures	44	3	22	14	5	VC
Craft Ventures	43	9	21	8	5	VC
Peak XV Partners	43	10	18	13	2	VC
Citi Ventures	43	3	15	18	7	Corporate VC
Spark Capital	41	4	12	18	7	VC
Kleiner Perkins	41	6	9	15	11	VC
Samsung NEXT Ventures	41	9	11	20	1	Corporate VC
Tiger Global Management	41	3	7	20	11	VC
Speedinvest	40	18	12	10	0	VC
QED Investors	39	8	18	13	0	VC
Abstract Ventures	39	13	14	12	0	VC
Lux Capital	39	8	12	11	8	VC
Antler	39	10	17	11	1	VC
Headline	39	7	13	19	0	VC
Felicis	39	8	21	8	2	VC
K5 Global	39	16	11	9	3	VC
Bossa Invest	38	2	6	19	11	VC
Founders Fund	37	4	8	13	12	VC
Recall Capital	36	13	15	8	0	VC
F-Prime Capital	36	4	13	14	5	VC

Source: PitchBook • Geography: Global • As of September 30, 2025

This new reality places a distinct mandate on every participant in the private markets ecosystem. For LPs, the imperative is to seek out GPs who possess the technical acumen to perform due diligence these new architectural paradigms. For GPs, the challenge is to build teams that can look beyond traditional financial metrics and identify the durable technical moats of the AI era. For founders, the mandate is clear: build natively.

Top PE investors in enterprise SaaS companies since 2023

Investor	Deal count	Primary investor type
Hg	32	PE/buyout
Thoma Bravo	28	PE/buyout
TA Associates Management	27	PE/buyout
Blackstone	24	PE/buyout
Vista Equity Partners	20	PE/buyout
Aquiline	17	PE/buyout
PSG	15	Growth/expansion
Francisco Partners	15	PE/buyout
Warburg Pincus	15	PE/buyout
Kohlberg Kravis Roberts	15	PE/buyout
TPG	14	PE/buyout
Silver Lake	12	PE/buyout
Accel-KKR	11	PE/buyout
JMI Equity	11	PE/buyout
The Carlyle Group	11	PE/buyout
Advent International	11	PE/buyout
Mainsail Partners	10	PE/buyout
Clearlake Capital Group	10	PE/buyout
General Atlantic	10	Growth/expansion
Genstar Capital	10	PE/buyout
EQT	10	PE/buyout
Sixth Street Partners	9	Growth/expansion
BGF	8	Growth/expansion
Bridgepoint Group	8	PE/buyout
New Mountain Capital	8	PE/buyout
Battery Ventures	8	Growth/expansion

Source: PitchBook • Geography: Global • As of September 30, 2025

Top PE investors in enterprise SaaS companies since 2023

Investor	Deal count	Primary investor type
Nordic Capital	8	PE/buyout
FTV Capital	8	Growth/expansion
Hellman & Friedman	7	PE/buyout
STG Partners	7	PE/buyout
Lead Edge Capital	7	Growth/expansion
Summit Partners	7	PE/buyout
Bain Capital	7	PE/buyout
The Riverside Company	7	PE/buyout

Source: PitchBook • Geography: Global • As of September 30, 2025

Top strategic acquirers in enterprise SaaS companies since 2023

Investor	Deal count	Investor type
Descartes Systems Group	9	Corporation
Fullsteam Operations	8	PE-backed company
Salesforce	8	Corporation
Shift4 Payments	6	Corporation
Atlassian	5	Corporation
Ai Software	5	VC-backed company
Visma	5	PE-backed company
International Business Machines	5	Corporation
Check Point Software Technologies	5	Corporation
Sage Group	5	Corporation
Bending Spoons	4	PE-backed company
HubSpot	4	Corporation
Snowflake	4	Corporation
Workday	4	Corporation
Motorola Solutions	4	Corporation
SPS Commerce	4	Corporation
Health Catalyst	4	Corporation
Valsoft Corporation	4	PE-backed company
Corpay	4	Corporation
nCino	4	Corporation

Source: PitchBook • Geography: Global • As of September 30, 2025

Top strategic acquirers in enterprise SaaS companies since 2023

Investor	Deal count	Investor type
Cox Enterprises	3	Corporation
SoundHound AI	3	Corporation
Datasite	3	PE-backed company
Software Circle	3	Corporation
Nearmap Australia	3	PE-backed company
VitalHub	3	Corporation
Cvent	3	PE-backed company
Blue Yonder	3	Corporation
Link Mobility Group Holding	3	Corporation
ServiceNow	3	Corporation
OpenAI	3	VC-backed company
PDI Technologies	3	PE-backed company
Thomson Reuters	3	Corporation
Fortinet	3	Corporation
Nayax	3	Corporation
Dropbox	3	Corporation

Source: PitchBook • Geography: Global • As of September 30, 2025

Conclusion and outlook

We expect that in the coming quarters and years, category leaders will demonstrate a shift from suggestion to accountable action and they will be priced accordingly. The market has already begun this shift as Intercom and Zendesk are putting "per resolution" at the center of their customer experience monetization, Ramp emphasizes agentic fraud detection and review in finance operations, and Salesforce is branding a Command Center to give executives visibility and control over agent behavior. These are all steps toward an operating model in which both AI-native and AI-embedded solutions are evaluated through observable output and cost.

Control planes themselves have the potential to contribute to moat building. Identity-aware tool permission-granting, policy stores, lineage, and observability dashboards all require significant investment and tuning against real customer workloads. This creates switching friction. These companies turn control into products, and both incumbents and startups alike have the potential to win deals even if their base model choice is undifferentiated.

In our view, vertical depth will be a greater priority over horizontal veneer. Legal agents that cannot guarantee citation chains at scale will not displace incumbent research workflows. In healthcare, ambient documentation will have to write

directly to electronic health records with appropriate audit functionality to maintain competitiveness. These integrations are vital for annual contract value and renewal likelihood. Distribution leverage will separate companies with similar technology. Cresta's arrangement with Accenture demonstrates why services partners matter in complex contact-center environments. Winners align with the go-to-market physics of their category rather than fighting them.

Finally, we predict that pragmatic model choices by vendors will surpass the singular model basis. Performance advantages will emerge from how well vendors select and evaluate models for each unique task rather than from the superiority of any single foundational LLM provider. Ultimately, practical benefits will accrue, including lower costs per action and fewer fallbacks. In our view, enterprises will reward reliability, auditability, and improving unit economics.

Appendix: Scoring key incumbents and challengers

Below, we provide a brief summary of 15 key incumbents and 30 challengers across the agentic AI space today. This is by no means an exhaustive list, as nearly every SaaS company and SaaS startup is deploying AI and AI agents as their frontier emerging technologies.

Our AI Score methodology represents a rigorous shift from evaluating AI based on conversational fluency to operational efficacy. Unlike traditional benchmarks that test static reasoning (such as Massive Multitask Language Understanding scores), our scoring assesses an agent's capacity to function as an autonomous, fiduciary operator within a complex enterprise environment. The core philosophy is that our AI Score is not defined by a single breakthrough model today, but by a system's ability to close the loop between perception (ingesting real-time data), reasoning (planning a course of action), and execution (using tools to change the state of the world), all while maintaining resilience when things go wrong. This approach explicitly penalizes fragile intelligence (systems that are smart but brittle) by demanding that an agent be as reliable and governable as the human employee it aims to augment.

We score these companies' offerings across 10 dimensions represented by three critical layers: Cognitive, Operational, and Governance. At the Cognitive layer, metrics like Goal Decomposition and Causal Reasoning measure the agent's ability to navigate ambiguity and solve novel problems rather than just matching patterns. The Operational layer (Tool Proficiency, Active Perception, and Adaptive Resilience) validates that the agent can manipulate software and recover from errors without crashing. Finally, the Governance layer (Bounded Autonomy and Verifiable Explainability) ensures the agent is safe to deploy, acting within strict guardrails to prevent runaway automation. This holistic view prevents a company from scoring highly just because they have a powerful LLM; they must also demonstrate the infrastructure (memory, orchestration, and safety) to turn that LLM into a functional employee.

The final score is calculated using the geometric mean, not the arithmetic average. This is a deliberate choice to penalize single points of failure. In an arithmetic average,

a company could score 100 in Reasoning and 0 in Security (Bounded Autonomy) and still get a passing grade (50). In this geometric framework, a zero or near-zero score in any critical dimension (like Contextual Continuity or Explainability) drastically pulls down the entire score. This reflects the reality of enterprise IT: An agent that is a brilliant coder but deletes the wrong database (low safety) or cannot remember the project specs (low memory) is functionally useless. Therefore, a high score (85 or more) indicates a balanced, mature system where no single dimension is a deal-breaker for real-world deployment today.

Key incumbents among AI-embedded platforms

Company	AI Score	Key strength (the "moat")	Agentic product set
Palantir Technologies	91.4	Reasoning (98): "Ontology" maps real-world physics/logic to prevent hallucinations.	AIP Agents: Built in AIP Agent Studio. Features "AIP Logic" (reasoning blocks) and "Automate" to run headless agents for logistics, defense, and supply chain.
Salesforce	89.5	Autonomy (92): "Atlas" engine creates multistep plans autonomously to resolve service cases.	Agentforce: Includes Service Agent (resolves tickets), SDR Agent (inbound/outbound sales), and Agentforce Studio for low-code custom agent building.
Snowflake	86.4	Active Perception (98): "Cortex" agents sit directly on top of the Data Cloud and have zero-latency access.	Snowflake Intelligence: A platform for building data agents that can autonomously query, analyze, and visualize data.
ServiceNow	85.6	Perception (98): Complete visibility into IT assets via CMDB, allowing agents to fix infrastructure.	Now Assist and AI agents: Specialized agents for ITSM (resolve incidents), CSM (order disputes), and HRSD. Built/managed via AI Agent Studio.
UiPath	84.8	Tool Dexterity (97): Computer vision allows agents to control legacy apps without APIs.	Agentic automation: Formerly Autopilot. Uses Clipboard AI and UI automation to let agents "see" screens and click buttons in legacy software (SAP, mainframes).
Microsoft	83.5	Integration (95): Agents have native read/write access to Graph data (Email, Teams, SharePoint).	Copilot agents: Built in Copilot Studio. Includes Autonomous Agents for Dynamics 365 (sales, service, and finance) that run background triggers.
Alphabet	82.9	Perception (96): Gemini 1.5's massive context window lets agents "read" entire codebases instantly.	Vertex AI Agents: The platform builder includes customer agents, employee agents, and creative agents. Deeply integrated with BigQuery and Gemini.
Workday	80.4	Governance (94): "Safe" financial/HR agents designed to prevent sensitive data leaks.	Illuminate: The platform powering Recruiting Agent (sourcing), Expenses Agent (auditing), and Succession Agent (planning).
Oracle	79.8	Integration (90): Deeply embedded supply chain agents that track physical goods via ERP.	Fusion AI Agents: Role-based agents for shift scheduling, opportunity qualification, and project management. Built via Oracle AI Agent Studio.
HubSpot	78.5	Goal Orientation (90): Specialized SMB agents focused on "prospecting" and "blogging."	Breeze Agents: Includes Prospecting Agent (researches leads), Social Media Agent (posts/analyzes), Content Agent (blogs), and Customer Agent.
SAP	77.2	Integration (85): Strong agents inside the "SAP walled garden" (finance/supply chain).	Joule: The primary copilot/agent. Joule Studio enables custom skills. It features specialized agents for consulting and coding (ABAP), and HR (SuccessFactors).

Source: PitchBook • Geography: Global • As of September 30, 2025

Key incumbents among AI-embedded platforms (cont.)

Company	AI Score	Key strength (the "moat")	Agentic product set
Amazon	77	Tool Dexterity (85): Bedrock Agents connect easily to AWS Lambda for backend execution.	Bedrock Agents: Infrastructure for building agents. Connects LLMs to AWS Lambda functions to execute API calls and retrieve data from knowledge bases.
IBM	76.8	Governance (90): Focused on "explainable AI" agents for regulated industries (banks/ government).	Watsonx Orchestrate: A library of "prebuilt skills" and agents for HR and sales. Focuses heavily on audit trails and Granite models for safety.
Atlassian	76.1	Memory (85): "Rovo" agents build a knowledge graph from Jira tickets and Confluence documents.	Rovo: Includes Rovo Agents that can clean up backlogs, draft release notes, or review code. Users can build custom "Teamwork" agents.
Adobe	73.5	Tool Dexterity (80): "Firefly" agents execute complex creative workflows (for example, bulk image edits).	Firefly Services: API-based agents for bulk image editing, generative fill automation, and custom model training for brand consistency.

Source: PitchBook • Geography: Global • As of September 30, 2025

Key challengers among AI-native startups

Company	VC ($M) raised to date	AI Score	Key strength (the "moat")	Agentic product set	IPO probability	M&A probability	No exit probability
Anthropic	$45,554.0	94.1	Computer Use (99): Modeled to "see" a screen and click buttons like a human would on various tasks.	Claude 4.5 Sonnet (computer use): An API capability allowing Claude to control a desktop (mouse/ keyboard) to use any software, not just APIs.	85%	10%	5%
Cognition	$738.0	92.8	Autonomy (98): "Devin" can plan, code, debug, and deploy software completely independently.	Devin: The "AI Software Engineer" takes a Jira ticket or GitHub issue, sets up the environment, codes, fixes the errors, and deploys the fix.	90%	6%	4%
OpenAI	$63,916.5	91.0	Reasoning (97): Current models use chain-of-thought reasoning to plan complex sequences before acting.	ChatGPT Agent: A browser-based agent that navigates websites to book travel, buy food, or research. Powered by OpenAI's current reasoning models.	97%	1%	2%
Ramp	$2,832.4	86.5	Bounded Autonomy (98): Financial agents operate with strict, audit-proof controls.	Ramp Intelligence: Procurement Agent negotiates SaaS renewals; Accounting Agent categorizes ledger entries; and Expense Agent audits receipts.	97%	1%	2%
Anysphere	$3,376.0	88.9	Tool Dexterity (95): "Cursor" agents edit code across multiple files simultaneously with deep context.	Cursor (Composer): An AI code editor where the "Composer" feature acts as an agent to write/ refactor code across the entire project directory at once.	95%	1%	4%

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Probability data is based on PitchBook VC Exit Predictor methodology.

Key challengers among AI-native startups (cont.)

Company	VC ($M) raised to date	AI Score	Key strength (the "moat")	Agentic product set	IPO probability	M&A probability	No exit probability
Glean	$765.3	85.5	Memory (97): Enterprise search that indexes everything a company knows to power agents.	Work AI Agents: No-code agents grounded in company data. Examples: Sprint Planning Agent, Competitor Analysis Agent, HR Benefits Agent.	97%	1%	2%
Imbue	$245.6	84.2	Reasoning (95): "Agents that code agents"—focused on logic/math over creative writing.	CARA: An agent specializing in reasoning and coding tasks that require long chains of logic (for example, solving math proofs or complex refactors).	2%	85%	13%
Celonis	$1,367.5	84.0	Perception (96): "X-rays" business processes to find inefficiencies agents can fix.	Process Intelligence Graph: Identifies process bottlenecks (for example, late payments) and triggers Action Flows (agents) to resolve them automatically.	78%	20%	2%
Cresta	$282.1	83.1	Collaborative Orchestration (95): Excels at "human-in-the-loop" handoffs, where the agent coaches the human.	Cresta Opera: A real-time command center where Virtual Agents handle tier-1 calls and Agent Assist guides humans using behavioral cues.	86%	9%	5%
Scale AI	$15,902.9	82.6	Adaptability (90): Provides the "Test & Eval" infrastructure to ensure agents are not lying.	SGP (Scale GenAI Platform): A platform for building, testing, and evaluating enterprise agents. Focuses on RLHF (human feedback) to fix agent errors.	92%	6%	2%
Databricks	$24,744.0	82.4	Governance (90): "Mosaic AI" allows strict control over the data agents can access/retrieve.	Mosaic AI Agent Framework: A toolset to build "Compound AI Systems" (agents) using Unity Catalog for data governance and vector search.	96%	2%	2%
Harvey AI	$988.0	81.8	Goal Orientation (94): Legal agents trained on case law to act as associates, not just search.	Harvey: A dedicated legal agent that can draft filings, research case law, and review contracts. Trained specifically on legal datasets.	54%	21%	25%
Hebbia	$159.2	81.5	Reasoning (90): "Matrix" agents analyze millions of documents for M&A due diligence.	Matrix: An "AI analyst" interface. You give it a grid of questions and documents (such as 50 contracts), and it fills cells with cited answers.	25%	62%	13%
Sierra	$635.0	81.2	Adaptability (88): Customer service agents that prioritize brand voice and empathy.	Sierra Agents: Consumer-facing support agents. Includes an Agent SDK for defining strict policies (such as "never refund over $50 without check.")	77%	19%	4%

Source: PitchBook • Geography: Global • As of September 30, 2025

Note: Probability data is based on PitchBook VC Exit Predictor methodology.

Key challengers among AI-native startups (cont.)

Company	VC ($M) raised to date	AI Score	Key strength (the "moat")	Agentic product set	IPO probability	M&A probability	No exit probability
Adept	$413.9	80.5	Tool Proficiency (96): Adept's model pioneered "pixel agents" that navigate any graphical user interface.	Adept Experiments: Browser-based agents that can "use" tools like Salesforce, Redfin, or Excel by clicking and typing rather than using APIs.	N/A	N/A	N/A
Moveworks	$308.2	80.3	Goal Orientation (90): Autonomous IT/HR helpdesk agents. (Note: An acquisition by ServiceNow is pending)	Moveworks Copilot: An autonomous agent for employee support. Uses Reasoning Engine to resolve issues like "grant software access" or "fix WiFi."	N/A	N/A	N/A
Cohere	$1,640.0	80.1	Tool Dexterity (92): Models fine-tuned specifically for reliable API calling (RAG) in enterprise.	Command R+: An LLM built for agents. excels at "Tool Use" (selecting the right API) and "RAG" (citing sources reliably).	93%	5%	2%
Writer	$326.0	79.5	Memory (85): Full-stack enterprise platform that guarantees no training on customer data.	Palmyra: The model family powering Writer Apps. Agents can generate medical summaries, financial reports, or marketing copy with strict accuracy.	84%	14%	2%
Perplexity AI	$1,713.7	79.2	Perception (95): "Search Agents" that scour the live web to synthesize answers.	Pro Search: A research agent that breaks a question into multiple search queries, reads the results, and synthesizes a cited answer.	91%	3%	6%
Runway AI	$543.9	78.4	Active Perception (90): "General World Models" (Gen-3) that understand physics and lighting.	Act-Two: A generative character agent that maps human performance to animated characters, with complex camera moves and object consistency.	55%	43%	2%
Magic	$1,016.2	78.2	Memory (98): Aims for massive context windows to hold entire operating systems in RAM.	LTM-2-Mini: A model with a 100-million-token context window, designed to be an agent that "never forgets" any code or document it has seen.	53%	41%	6%
CrewAI	$24.5	77.9	Reasoning (85): Leading open-source framework for multiagent orchestration.	CrewAI Enterprise: A platform to design "crews" of agents (researchers, writers, and editors) that assign tasks to each other.	3%	80%	17%
Anduril Industries	$6,840.1	76.5	Autonomy (99): Autonomous defense systems/ drones; high agency, hardware constrained.	Lattice: The operating system for defense agents. Controls Ghost drones and Altius loitering munitions, enabling swarm autonomy without human pilots.	97%	1%	2%

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Probability data is based on PitchBook VC Exit Predictor methodology.

Key challengers among AI-native startups (cont.)

Company	VC ($M) raised to date	AI Score	Key strength (the "moat")	Agentic product set	IPO probability	M&A probability	No exit probability
Mistral AI	$2,713.1	75.8	Adaptability (80): Efficient open-weight models for on-premises/local agent deployment.	Le Chat/La Plateforme: Offers "agents" in their chat interface. Models like Mistral Large are optimized for reasoning and coding tasks.	51%	9%	40%
H Company	$220.1	75.5	Autonomy (85): Founded by DeepMind alumni; focused on "action models" over language models.	Action Models: Developing agents capable of multimodal action (controlling computers and robots) rather than just text generation.	N/A	N/A	N/A
Please.ai	N/A	75.0	Tool Dexterity (90): Browser agents that can book flights/tables by controlling Chrome.	Agent API: A browser automation agent. Developers use it to let their apps "control" websites (for example, "go to Amazon and buy X").	2%	72%	26%
Zapier	$1.4	74.5	Tool Dexterity (92): "Canvas" connects agents to over 6,000 apps, though logic is linear.	Zapier Central: An experimental workspace where you teach agents to trigger “Zaps” based on natural language logic.	13%	70%	17%
Lindy	$53.6	74.0	Goal Orientation (85): "Employee in a box" agents for specific roles (for example, medical scribe).	Lindy: Prepackaged "AI employees" (for example, medical scribes, executive assistants) that have their own emails and calendars.	1%	80%	19%
Figure AI	$2,345.0	73.0	Perception (90): Humanoid robotics agents that interact with the physical world.	Figure 03: Humanoid robots powered by OpenAI models. They can see, hear, and perform physical tasks (such as making coffee) autonomously.	67%	18%	15%
LangChain	$160.0	72.5	Reasoning (80): The primary infrastructure developers use to build reasoning loops.	LangGraph: A library for building stateful, multiagent applications. It defines the "loops" and "memory" for custom agents.	13%	85%	2%

Source: PitchBook • Geography: Global • As of September 30, 2025
Note: Probability data is based on PitchBook VC Exit Predictor methodology.

PitchBook provides actionable insights across the global capital markets.

Additional research:

Q3 2025 Enterprise SaaS VC Trends

Download the report here

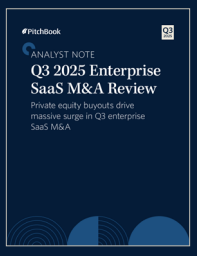

Q3 2025 Enterprise SaaS M&A Report

Download the report here

PitchBook Insights is an online compendium of in-depth data, news, analysis, and perspectives that shape the private capital markets.

PitchBook subscribers enjoy exclusive access to a comprehensive suite of private market insights, including proprietary research, news, data, tools, and more on the PitchBook platform.

Nizar Tarhuni
Executive Vice President of Research and Market Intelligence

Paul Condra
Senior Director, Global Head of Private Markets Research

James Ulan
Director, Industry & Technology Research

Report created by:

Derek Hernandez
Research Analyst

Oscar Allaway
Senior Data Analyst

Chloe Ladwig
Graphic Designer

Learn more about PitchBook's Institutional Research team.

Click here for PitchBook's report methodologies.