PitchBook
a Morningstar company

Institutional Research Group

Rudy Torrijos
Director, Industry Research
rudy.torrijos@pitchbook.com

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

pbinstitutionalresearch@pitchbook.com

Published on February 9, 2026

Contents

EMERGING TECH RESEARCH

SaaS Is Dead, Long Live SaS

The enterprise AI super-cycle begins now

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 "SaaS-pocalypse" is a false prophecy:** Incumbents are not standing still, they are becoming the largest AI companies and greatest consumers of tokens in the world. The market is pricing SaaS for obsolescence just as those vendors are transforming into "service-as-software" companies that stand to consume the largest total addressable market in history: the labor market.
- **The software TAM will merge with the labor TAM in the next 20-year enterprise AI super-cycle:** Outcome-based pricing shifts vendor revenue capture from an annual charge of around $1,200 per seat to an annual charge of about $10,000 per automated workflow, an 8x revenue uplift per unit. Software budgets will merge with payroll budgets as SaaS replaces the original value-creation workflow.
- **ARC-AGI-2 scores now surpass 60% at $1 to $10 per task:** 60% is widely regarded as the labor-replacement milestone. Costs per task are declining incrementally toward the sub-$1 threshold.[1]
- **Hyperscaler AI capital expenditure is approaching $1 trillion annually:** Amazon, Google, Meta, and Microsoft are building expert intelligence inference environments for SaaS incumbents to deploy agentic AI. Infrastructure investment precedes application-layer deployment, a pattern consistent with every prior technology super-cycle.
- **Fear is the greatest accelerator:** Market turmoil is forcing the SaaS-to-SaS transition roughly 18 months ahead of when agentic AI infrastructure will be ready for full-scale deployment. Fear of business-model collapse is driving restructuring faster than the allure of new capabilities ever could.
- **Incumbents retain 10 structural moats:** distribution, data context, switching costs, workflow integration, security, compliance & governance, AI as a revenue layer, outcome pricing, vendor consolidation, investment scale, and deep business logic. AI is the ultimate margin-accretive lever for those who own the customer and the data.
- **Investors must reassess every software position:** PE and credit investors should audit portfolios for workflow-wrapper risk and restructure covenants with gross margin triggers below 60%. VCs should avoid funding UI wrappers and pursue vertical agents with proprietary data moats. Public equity should overweight cybersecurity, payments, and critical infrastructure while avoiding professional services and seat-based value traps.

1: "ARC-AGI-2 Leaderboard," ARC Prize, n.d., accessed February 9, 2026.

Executive summary

The market is suffering from a fundamental misunderstanding of the AI technology investment cycle. A wise man once said, “the market is always right.” The same is true today and the disruption represents a decisive opportunity for investors and vendors to assess scalable technologies as enterprise AI begins a 20-year super-cycle in 2026.

SaS is the realization that the end service is the value-added deliverable and that software is the vehicle by which it is implemented

Software as a service (SaaS) is dead. Long live service as software (SaS).

SaS is the realization that the end service is the value-added deliverable and that software is the vehicle by which it is implemented. The paradigm shifts from enterprises renting productive tools on a subscription basis to value-based solutions centered on labor replacement. This is the promise of SaS. Software budgets will merge with payroll to enable automation of complete employee workflows. This increases operating margins for the enterprise and increases software TAM for vendors.

We are witnessing a seismic transformation in advanced computing and enterprise software, culminating in the simultaneous valuation implosion of public software company market caps and the explosion of capital expenditure on AI infrastructure from hyperscale computing providers (hyperscalers). The prevailing narrative is that large language models (LLMs) will commoditize application software, erode those margins, and end the recurring-revenue model that has underpinned the past two decades of technology investing. We declare that this is a false narrative.

Concurrently, Amazon, Google, Meta, and Microsoft reported dramatically increased capital investments for fiscal year 2026. Investors have become worried that hyperscaler capital expenditures on datacenters are growing too fast and will quickly eat through available operating cash flow and into capital balances. We believe the ongoing step-function growth in capital expenditures, after dramatic year-over-year increases since 2022, is solely driven by the enterprise. In specific domains (healthcare, life sciences, legal, finance, software coding), frontier models have achieved expert-level intelligence much faster than we expected just two months ago when we published the Q3 2025 Launch Report: Advanced Computing.

Just as DeepSeek shook AI expectations with the introduction of large-scale mixture-of-expert LLM architectures in early 2025, Anthropic similarly triggered investors to sell software positions first, possibly to ask questions later. Anthropic’s announcement last week of an AI lawyer-assistant-for-hire plugin seemingly convinced investors of the imminent death of enterprise software. Market sentiment turned sharply critical on software as a service, forcing PE, VC, and private credit investors as well as technology vendors, IT managers, and the public markets to reevaluate everything they thought they knew about AI to develop what we call SaS. Ultimately the “SaaS-pocalypse” was the fear-based catalyst we did not know we needed.

What we are seeing is both healthy and necessary to propel the entire technology value chain forward. For technology companies, this moment represents a massive opportunity to focus on products that expand their total addressable markets (TAMs).

The software TAM is lining up to merge with the much more expansive TAM for human labor, and the SaaS-pocalypse market turmoil was the starting gun to align software companies on this trajectory. For investors, AI-era winners will be those companies that accelerate and focus their investments in AI-native architectures to leverage distribution, trust, and proprietary data to achieve outcomes no pre-AI software tool ever could achieve alone.

Enterprises wake up

Enterprises are now fully engaged. For specific business processes, AI has achieved expert, Ph.D.-level capabilities. Current AI intelligence benchmarks like ARC-AGI-2 show scores are now surpassing 60%, with visibility to 70% at costs per task between $1 and $10.[2] Below $1 per task, these scores have approached 30%. 60% is widely regarded as the labor-replacement milestone, especially as costs per task are reduced incrementally.

We are witnessing the precise moment when enterprise AI transitions from AI pilot-program curiosity to a competitive workflow imperative. The simultaneous explosion in hyperscaler capital expenditures and the rapid ascent of these intelligence benchmarks are directly correlated and portend the beginning of a grand enterprise AI super-cycle. Proof of this is evident in recent Q4 earnings reports.

- Advanced Micro Devices reported that the number of large businesses deploying its high-performance server chips on-premises more than doubled in 2025, and the company exited the year with record server sell-through.[3]

- SAP reported that more than two-thirds of its Q4 cloud order entry included its Business AI platform, with that figure rising to 90% for the company's 50 largest deals.[4]

- ServiceNow's centralized governance platform exceeded targets by more than 4x, and the company noted that customers are expanding entitlements by 13x upon renewal because the business value is now proven.[5]

- Amazon, Google, and Microsoft reported that non-AI workloads or traditional enterprise demand is increasing as enterprise customers accelerate their move to the cloud.

We expect this trend will accelerate in 2026 as enterprise software companies highlight increasing demand from customers investing in new AI features and capabilities for existing seats.

2: "ARC-AGI-2 Leaderboard," ARC Prize, n.d., accessed February 9, 2026.
3: "AMD Reports Fourth Quarter and Full Year 2025 Financial Results," Advanced Micro Devices, February 3, 2026.
4: "SAP Announces Q4 and FY 2025 Results," SAP, January 29, 2026.
5: "ServiceNow Reports Fourth Quarter and Full-Year 2025 Financial Results; Board of Directors Authorizes Additional $5B for Share Repurchase Program," ServiceNow, January 28, 2026.

The market's reaction to Anthropic's legal services plugin was based on the belief that it is the catalyst for corporations' acceptance of the virtual employee. It signals that specific AI applications are now believed to be robust enough to handle the system-of-record (SoR) tasks that enterprises require and that continued advancements are likely inevitable. The extreme fear that SaaS companies are being replaced by LLMs is the mandate for SaaS companies to transform into LLM token consumers. Going forward, software solutions without high-fidelity expert intelligence built in will be commoditized with no defensible moats.

Infrastructure precedes application

Hyperscalers are building expert intelligence inference environments for SaaS incumbents to deploy agentic AI. History has shown that infrastructure investment precedes application-layer deployment. The acceleration in AI capital expenditures lays the necessary physical foundation for the next generation of software development. The gap between infrastructure spending and application value realization is a deployment lag that has historically allowed end customers to test, pilot, adopt, and prioritize. The current market turmoil has decisively prioritized creation and implementation of AI-native architectures across the entire value chain, lest current software vendors risk future economic loss.

As such, we anticipate that demand for AI-enabled software will outstrip supply for the foreseeable future. The introduction of expert-level AI coding assistants accelerates software development, but we estimate that the demand for new digital business-process solutions will remain unsatisfied for the next 20 years. This is the enterprise AI super-cycle, and it is beginning in 2026. For the past 20 years, software has been created to help employees do a specific workflow function within strict boundaries and with limited flexibility. Over the next 20 years, software will be created to replace employee workflows entirely. This is the promise of AI, even as it dramatically increases employee capacities and capabilities across all sectors.

As datacenters are built out, costs per token decline, and as intelligence scales up, enterprise developers and engineers will enable high-fidelity intelligent agents to manage every IT process, from the mundane to the complex. The IT department will be free to focus on business process improvements and transition from a cost center to a revenue enabler. One IT manager will manage a team of AI agents to keep the enterprise working. Other IT managers will be assigned teams of AI agents to assist in marketing, sales, customer success, finance, strategy, manufacturing, and supply chain management to optimize business processes, minimize costs, and optimize go-to-market strategies. Business models will be optimized at any given level of head count.

IT operations VC deal activity by quarter

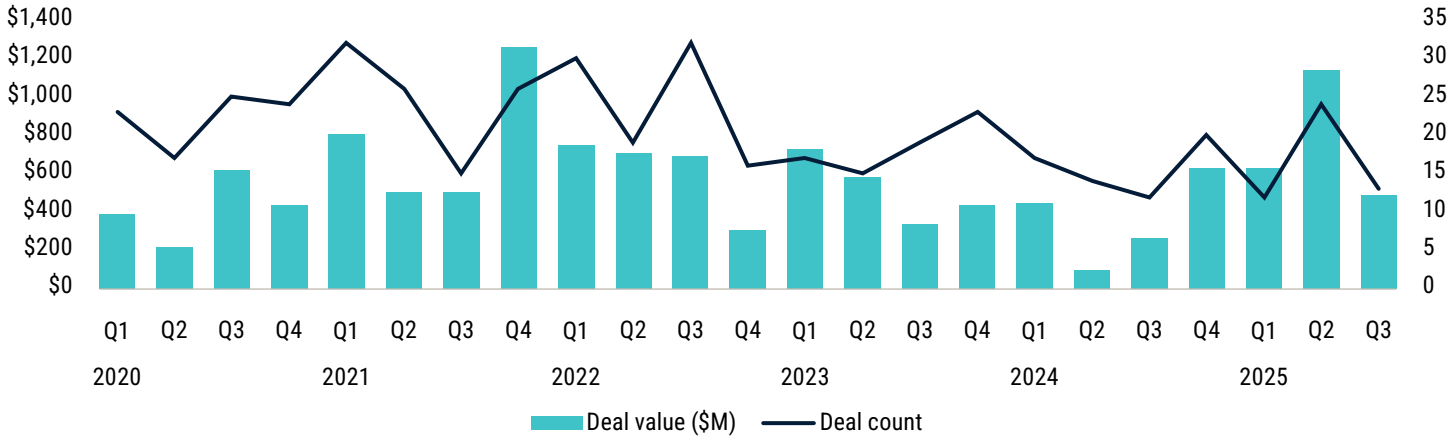

Source: PitchBook • Geography: Global • As of September 30, 2025

Productivity to replacement

AI reduces software production costs. For software companies, AI erodes defensive moats if the cost of coding is the only gating factor to an otherwise generic enterprise solution. With AI, coding is free and business process automation becomes the new critical moat. We call this "service as software," where the end service is the value-added deliverable and software is the vehicle by which it is implemented.

Investors must reassess their holdings in software companies, with a critical eye on those that do not or cannot develop contextual solutions for business. In traditional SaaS, financial models assume revenue built on employee productivity improvement: $100 per month per seat to save three hours a week of manual process, for example. The new AI paradigm is labor replacement: $10,000 per year to save 40 hours a week of routine workflows.

Currently, companies pay for an employee to use a tool, and the software vendor captures a small fraction of the cost saved by the company from using that tool. This is a paradigm born from the days of factories, machines, and widgets. With native AI architecture, we enter a new age of software deployed as agentic teams, automating both mundane and complex workflows. In this world, the software vendor charges an outcome-based fee for the work completed, cementing the return on investment (ROI) as part of its value-creation story. The employee would then be tasked with developing revenue-optimizing workflows, creating greater value for the enterprise, and becoming more valuable themselves.

This is highly deflationary for the company, generating the original value created by the SaaS workflow with a $10,000 SaS spend rather than the original $75,000 cost-center salary plus the $100 monthly SaaS cost. Equally, this is massively inflationary for the software vendor, increasing revenue capture from $1,200 ($100 per month) to $10,000 per year for workflow replacements. This is the promise of SaS. Software budgets will merge with the payroll budget as SaS replaces the original value-creation workflow.

This increases operating margins for the enterprise and increases the software TAM for vendors.

The transition to SaS is expected to be most lethal to industries based on billable hours. The labor arbitrage model—hiring armies of workers in lower-cost regions to do knowledge work—will collapse. If one professional with AI can do the work of 10 (or more) service providers, there is no need to employ human capital in this manner. As personal computers and word processing applications enabled law firms and other professionals to dispense with armies of typists in the late 20th century, so too can AI enable those former service providers to pursue more value-creating opportunities. Legal, consulting, and financial services will necessarily face a disruption of their entry-level workforce, as we are already observing from various anecdotes across these sectors. Work that previously cost $400 per hour is now replicable through a chat interface for $20 a month.

Structural moats and the enemies at the gates

Despite the brave new world we envision, extreme valuation declines across the software industry ignore many structural moats that still safeguard enterprise software business models. AI is the ultimate margin-accretive lever for those who own the customer and the data. We believe the recent software "tech wreck" has been indiscriminate. Discernment is the opportunity, especially as real vulnerabilities exist to be exploited by AI-native challengers across every sector today. There are 10 critical reasons why software will not disappear, and why many incumbents remain well positioned to capture value, compete, thrive, and ultimately accelerate revenues in the age of enterprise AI:

1. Distribution

In a world where coding costs trend toward zero, distribution becomes the scarcest asset. Incumbents have sales channels and multi-year contracts based on established product suites. Large enterprise customers prefer to buy solutions over building their own. They cannot risk failure in a noncore competency and often require primary and backup vendors. A startup may have a better algorithm (for now), but they cannot replicate decades of vendor approval processes and business context workflows overnight.

AI-native vendors are redefining distribution by shifting the focus from contract ownership to measurable outcomes. Rather than waiting for multi-year procurement cycles, these challengers demonstrate ROI within weeks and enable team-by-team adoption without requiring full platform replacement. Companies like Sierra have acquired hundreds of enterprise customers in around 18 months, while Decagon has added over 100 corporate clients in rapid succession.

This speed-to-value approach allows challengers to bypass traditional gatekeepers, though incumbents' deeper customer relationships and compliance credentials continue to provide substantial defensive advantages for core systems of record.

2. Data context

LLMs are commodities, and context is the differentiator. Incumbents own the system of record that houses the data required to empower AI agents. Customer relationship management (CRM) and enterprise resource planning (ERP) data warehouses will become more valuable. Without access to the secure customer, inventory, pricing, or sales data that is housed within existing software platforms, AI agents are largely neutered. Incumbents must protect and leverage this advantage at every opportunity.

Challengers are pursuing a federated context model that positions data access as a layer above individual systems rather than a moat owned by any single platform. Companies like Glean explicitly position themselves as enterprise knowledge platforms that search and act across disparate applications including Salesforce, Microsoft Teams, and Zendesk, while Hebbia specializes in reasoning over large, unstructured internal document sets.

However, incumbents retain significant advantages through their ownership of structured transactional data, established security protocols, and the natural reluctance of enterprises to expose critical business data to third-party aggregation layers.

3. Switching cost

Customers will continue to pay premium prices for high quality, high availability, and especially peace of mind

Replacing a core SaaS platform is extraordinarily difficult and expensive. It is far safer for an enterprise to pay for an AI copilot add-on from a trusted vendor than to migrate to a startup. This is the toehold into the AI-era, which existing incumbents will use to drive customers into their future AI-native product solutions. Customers will continue to pay premium prices for high quality, high availability, and especially peace of mind.

Rather than attempting direct platform replacement, challengers are employing a “sidecar pattern” that seeks to neutralize switching costs by wrapping and augmenting existing systems. Customer support represents the clearest validation of this strategy as companies like Decagon and Sierra automate support flows through existing tooling, making the AI layer the primary interface while legacy platforms continue to house data.

While this approach provides market entry, incumbents can counter by accelerating their own AI feature development and leveraging their existing customer relationships to maintain strategic control of enterprise workflows.

4. Workflow integration

Real value lies in AI embedded within the workflow, not a chatbot. Incumbents begin with the advantage of controlling the user interface (UI) where work happens. Incumbents can deploy agents as friction-reducing features rather than destination apps, automating requests inside the CRM, ERP, knowledge-management-system or business-intelligence software.

While incumbents control specific interfaces, challengers increasingly orchestrate the actual work by coordinating actions across multiple systems including ticketing, CRM, billing, and communications platforms. Glean exemplifies this approach with its cross-application knowledge and action capabilities that extend beyond single-app copilots to function as a broader control plane.

However, this strategy faces significant headwinds: Enterprises remain cautious about granting broad system access to external parties, integration complexity increases exponentially with system count, and incumbents can potentially replicate successful cross-app patterns by leveraging their existing customer relationships and security credentials.

5. Security, compliance & governance

Enterprises are extremely risk-averse to data leaks. Incumbents offer the necessary SOC 2 compliance and role-based access-control permissions. When agents are deployed at scale, mistakes will happen at scale. Incumbents will offer security & compliance features to prevent agents from misusing data housed in secure systems.

AI-native companies assert that security & compliance have evolved from multi-year differentiators into table stakes that can be achieved relatively quickly. The rapid enterprise adoption in regulated verticals provides supporting evidence, as Harvey counts 50 of The American Lawyer 100 law firms as customers,[6] while Abridge achieved a $5.3 billion valuation as hospitals adopt its ambient clinical documentation tools.

Nevertheless, incumbents retain meaningful advantages through their existing compliance track records, established relationships with enterprise security teams, and the deep trust that comes from years of reliable operation in mission-critical environments.

6. AI is a revenue layer

For incumbents, AI is a feature that will justify future price hikes. This will help subsidize current inference costs until greater cost efficiencies are developed and additional AI use cases and workflow automation products generate even greater value.

While incumbents want AI to function as a premium add-on, enterprise customers expect AI to reduce seat counts and services spending, creating direct revenue cannibalization. AI-native vendors face no such conflict, they can price aggressively to capture budget reallocation from incumbent solutions, using AI as a competitive wedge.

6: "Harvey's Momentum in Q3," Harvey AI, September 29, 2025.

However, this challenger advantage may prove temporary. Incumbents with diversified product portfolios and strong customer relationships can likely absorb some margin pressure while they transition pricing models, whereas AI-native startups face pressure to demonstrate sustainable unit economics to justify their venture valuations.

7. Outcome pricing

Software will be priced strictly by value creation over time, with labor costs reduced as AI solutions replace manual workflows

Incumbents have recognized the vulnerabilities of seat-based pricing and will monetize agents by shifting to outcome pricing. Software will be priced strictly by value creation over time, with labor costs reduced as AI solutions replace manual workflows. This shift allows incumbents to capture a portion of the labor-market revenue, significantly expanding their TAM.

Shifting from seat-based to outcome-based models compresses existing annual recurring revenue and complicates multi-year contract renewals, creating organizational resistance and financial reporting challenges. Companies like Decagon and Sierra can price support agents based on deflection or resolution metrics without legacy revenue conflicts, making this transition natural rather than disruptive.

Incumbents counter that their established customer relationships provide negotiating leverage for pricing transitions, their diversified product portfolios allow for experimental pricing in some areas while maintaining traditional models elsewhere, and their financial stability enables them to absorb transition friction that might sink venture-backed startups.

8. Vendor consolidation

Enterprises favor platforms over point solutions, which enables them to compete on pricing discounts. Point-solution startups are thus disadvantaged. Incumbents provide multiproduct solutions and will bundle AI capabilities to preserve market share.

AI agents operate indifferently to application boundaries and can seamlessly execute across best-of-breed tools, thus the economic value of suite bundling diminishes significantly. In this scenario, consolidation might occur around the agent orchestration layer rather than around traditional ERP or CRM suites.

Still, enterprises value consolidated vendor relationships for reasons beyond UI efficiency (including procurement simplicity, support escalation, and strategic alignment), and incumbents can build or acquire agent orchestration capabilities to defend their platform positions.

9. Investment scale

Incumbents have balance sheets to acquire winning startups and secure compute capacity. Incumbents can outspend and outwait venture-backed challengers, which are meanwhile burning runway at a higher pace than ever before.

Today, the private market is directing massive capital flows toward AI-native enterprise leaders at a scale that diminishes traditional fundraising advantages by incumbents. Today's outsized funding rounds provide challengers with sufficient resources to secure compute capacity, hire elite talent, and build distribution infrastructure.

Incumbents retain strategic advantages beyond raw capital. Their existing profitable operations generate ongoing cash flow for sustained investment, their established customer relationships reduce customer acquisition costs, and their track records enable more favorable M&A terms. The question is less about absolute financial resources and more about deployment efficiency and strategic focus.

10. Deep business logic

The structural moat of the incumbent is deep business logic. Enterprises run on rules, compliance, and complex, multistep workflows. Incumbents have accrued this knowledge over thousands of customer interactions. Enterprises use AI to accelerate and optimize workflows, but confine AI in the safety of deterministic business logic that incumbents have built into their product.

AI-native vendors argue that deep business logic is being rebuilt through superior feedback mechanisms rather than inherited from legacy systems. Harvey provides a compelling example: Deployment across elite law firms generates dense feedback loops from expert users that accelerate domain logic acquisition. Similarly, companies like Fundamental demonstrate that purpose-built AI models for enterprise data can potentially outperform legacy logic engines.

Meanwhile, incumbents can implement similar feedback systems while leveraging their existing logic as a starting point. In addition, regulated industries require extensive validation before replacing proven systems, and enterprises demonstrate a strong preference for battle-tested logic over theoretically superior learning systems when operational risk is high.

The future, today

We believe that the ultimate outcome will likely reflect market segmentation rather than winner-take-all dynamics. Incumbents appear best positioned to defend mission-critical systems of record where switching costs, compliance requirements, and operational risk favor established vendors. Challengers may successfully capture emerging AI-first workflows, customer support automation, and cross-application orchestration where their architectural advantages matter most. Enterprises will likely maintain hybrid environments, leveraging incumbent platforms for core operations while selectively adopting AI-native solutions for specific high-value use cases.

For incumbents, the imperative is clear: Accelerate AI capability development, demonstrate willingness to evolve pricing models, and leverage existing customer relationships before challengers establish themselves as the default AI layer. For challengers, success requires proving sustainable unit economics, navigating enterprise security requirements, and capturing sufficient market share before

incumbents close capability gaps. The AI transformation of enterprise software has created genuine competitive opportunity, but incumbent moats remain substantial—and reports of their demise are significantly premature.

Recommendations to investors and vendors

The transition to SaS demands distinct strategies for each class of investor and operator. Each requires specific navigation strategies depending on their position in the market. Here are the specific execution frameworks for five key stakeholders:

PE and private credit

For private equity and private credit investors, the risk is more real than for public equity investors given the inherent obstacle of illiquidity. Any investment funded on the basis of seat growth or retention is now an opportunity for reassessment and reinvestment. The transition to SaS challenges the predictability of the cash flows that underpin leverage models. Investors must move from passive monitoring to active restructuring using a four-phase approach:

1. **Audit portfolios:** Investors must categorize their assets. SoR companies holding proprietary data carry low risk with caveats. The action here is to restrict API access to scrapers and mandate that they build the governance infrastructure to protect data misuse and misappropriation by agents at all layers within and above the SoR. Conversely, startups developing workflow wrappers are high-risk investments. These companies facilitate processes without owning data and require immediate consolidation or aggressive repricing.

2. **Reprice offerings:** Management must lead a shift from seat-based to outcome-based pricing. Pricing should align with work completed rather than user access. Outcome- and usage-based pricing will increasingly dominate the market. Additionally, consumption models with usage meters must be implemented to ensure revenue scales with the required AI compute costs.

3. **Restructure covenants:** Credit agreements must be updated. Gross margin triggers are essential to monitor for margin compression due to AI inference costs. Investors should trigger reviews if gross margins drop below 60%. Furthermore, investors must monitor monthly logo churn in the small and medium-size business (SMB) segment. SMBs are the early adopters of good-enough AI replacements, making them a leading indicator of competitive disruption.

4. **Consolidate:** Sponsors should merge point solutions to create platforms with high switching costs. Horizontal software consolidators are statistically safer than vertical specialists in this environment. Merging companies pools proprietary data, creating a larger dataset for training defensive internal models and more aggressive future solutions.

Enterprise SaaS PE deal activity by type

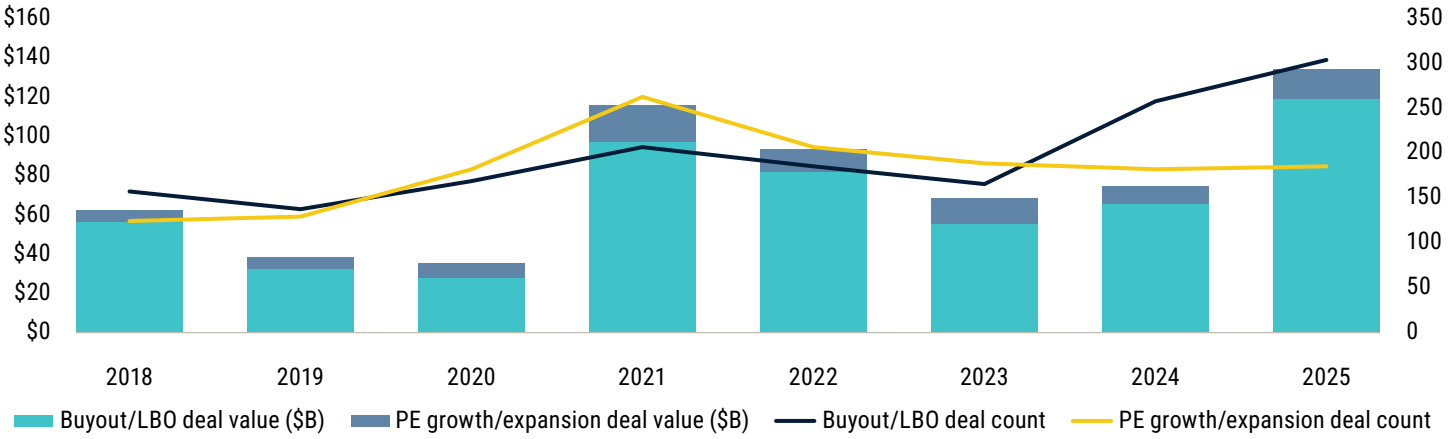

Source: PitchBook • Geography: Global • As of December 31, 2025

VC and corporate VC

The error of the 2020-2022 VC vintage was over-indexing on productivity tools. The error of the 2025 vintage will be over-indexing on wrappers.

1. **Avoid startups focusing on user interfaces:** Do not continue to fund companies that primarily offer a UI layer over a foundation model. The agent layer above them absorbs their interface, while the systems of record below offer their own workflow. If a startup's core value is a dashboard, it is going to zero.

2. **Fund agents:** Seek companies building AI-native tools that execute workflows autonomously. The metric of success is not daily active users but work units completed and rapid logo adoption. Look for startups that price on outcomes with deep domain knowledge and accelerating uptake.

3. **Pursue governance:** There is massive opportunity in the mundane layer of governance. Security, monitoring, evaluation, and compliance tools that make AI agents safe to deploy at scale will mint the next generation of decacorns as enterprises scramble to control their agent fleets.

4. **Focus on verticals:** Generalist agents will be dominated by Google, Anthropic, and OpenAI. Vertical agents (for example, an agent for dental practice billing) will thrive. Look for startups leveraging unique, nonpublic datasets to fine-tune their agents, creating a data moat that generic models cannot breach.

Enterprise SaaS VC deal activity

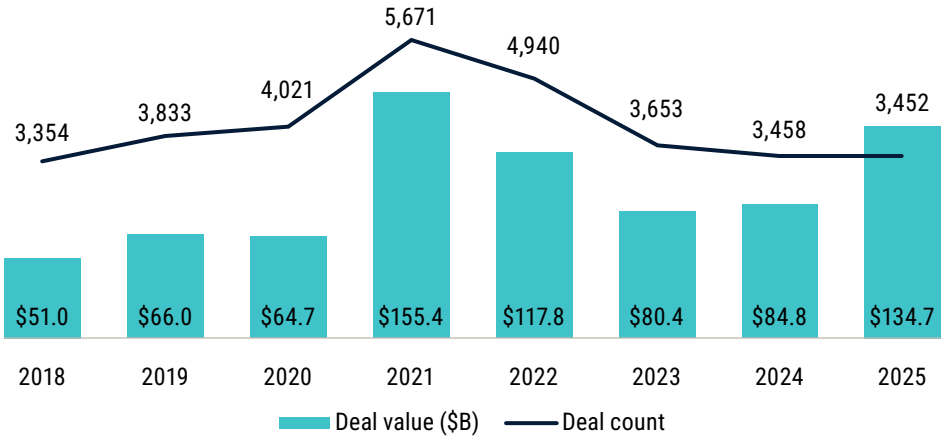

Source: PitchBook • Geography: Global • As of December 31, 2025

Public equity

Public markets have been dumping everything indiscriminately. This is where the opportunity lies.

1. **Identify compounders:** Buy incumbents who are successfully pivoting to outcome-based pricing. Look for the best of suite consolidators. Companies offering a broad, growing product portfolio will continue to outperform singular legacy solutions. These companies will be positioned to capture much of the service revenue currently flowing to consulting firms.

2. **Avoid value traps:** Sell companies rigidly stuck on seat-based pricing with low switching costs. If a company's primary growth lever is adding head count at their customer's organization, they are short the AI revolution.

3. **Overweight security:** Overweight sectors where failure is not an option—cybersecurity, payments, and critical infrastructure. Deterministic systems where precision is critical are more likely to be insulated. These sectors will command pricing power even in a deflationary AI environment.

4. **Avoid services:** Professional services firms (consultancies, business-process outsourcing firms) that rely on head count growth face existential headwinds. Look instead for platforms that automate their work.

SaaS founders

If you are building today, you cannot build using the 2025 playbook.

1. **Do not build a dashboard:** Do not show the user a dashboard; do the work for them. Your product should be an invisible worker, not a visible tool. Minimizing time-in-app is now a positive metric, not a negative one.

2. **Own the data layer:** If you do not have a system of record, build one. You need to own the proprietary data that your agents act upon. Agents come and go, but the data layer is forever.

3. **Price on outcomes:** Do not start with seat-based pricing. Start with value-based pricing (per report, per filing, per lead). It should be easier to sell initially as it requires less upfront commitment by the customer and aligns you with their long-term incentive to reduce head count. The downside is less guaranteed cash flow once the contract is signed, so produce or perish.

4. **Grow services:** This is counterintuitive for software vendors, but as the cost to build software drops, companies will attempt more complex customization. Building a high-touch service layer to help enterprises implement your agents allows you to capture the budget previously allocated to systems integrators.

Q4 2025 enterprise SaaS VC deal activity by segment

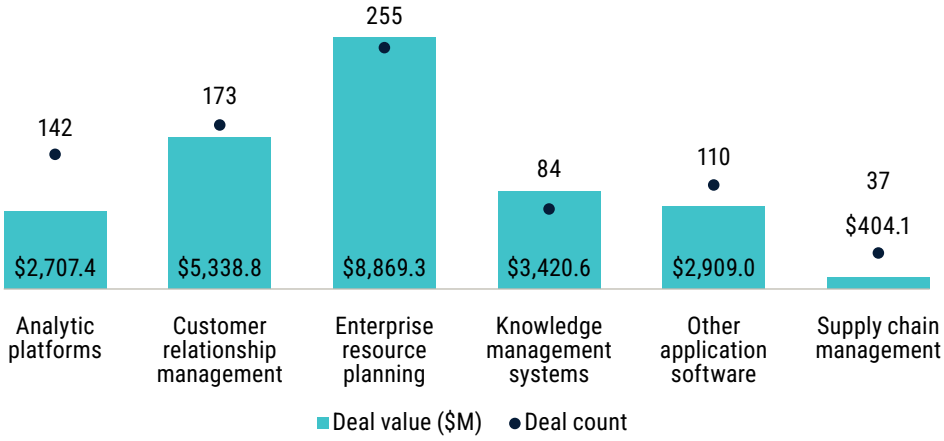

Source: PitchBook • Geography: Global • As of December 31, 2025

IT managers and chief information officers

The “shadow IT” aspect of the cloud era—wherein employees used unauthorized IT tools—is becoming “shadow AI” in the agent era. If you do not provide a sanctioned, secure agent fleet, your employees will paste proprietary intellectual property into a free, open-source model.

1. **Shift from buying tools to hiring agents:** Procurement processes must change. You are no longer evaluating software uptime and UI; you are evaluating work quality and error rates. Treating software vendors like staffing agencies will become the norm.

2. **Block the scrapers:** Aggressively audit your system-of-record exposure. Ensure that unauthorized AI agents cannot scrape your proprietary data via API to train their models.

3. **Demand walled gardens:** Do not allow open-ended LLM access to your core data. Demand private instances from your vendors. The risk of data leakage is the primary blocker to adoption. Prioritize vendors who offer enterprise-grade permission processes and isolation.

4. **Prepare for human-capital redeployment:** Your internal help desk and support functions will be the first to be automated. Plan for fewer Level 1 support staff, and more agent operations (AgentOps) engineers responsible for monitoring and fine-tuning the digital workforce.

Conclusion: Fear is the accelerator

Market turmoil is forcing the transition to a service-as-software economy. This is likely 18 months ahead of when generative AI and agentic AI tools and infrastructure will be ready for full-scale deployment. The SaaS-pocalypse is a false prophecy. The incumbents are not standing still, they are fast transitioning to become the largest AI companies and the greatest token consumers in the world.

Fear is a far more powerful force than the allure of success. The market is forcing a transition not through the beckoning call of new capabilities and efficiencies, but through the fear of imminent business-model collapse. This fear will drive the necessary restructuring of the software economy and by extension increase long-term corporate margins. We are entering an era where software not just supports work but actually performs it. This expands the revenue potential for incumbents who can successfully transition from selling tools to selling digital labor.

The market is currently pricing software assets for near-term obsolescence just as those same vendors are about to serve the largest TAM in history—the labor market. Investors should use this opportunity to reassess software investments. Look for the systems-of-record providers aggressively pivoting to systems of workflow automation with deep domain expertise embedded in their products.

The market is always right: Enterprises will stop buying SaaS. However, they will transition to employing digital labor to perform service as software. Many workflow process steps have already been replaced at today's level of LLM intelligence with simple AI chatbots. With AI capital expenditures approaching $1 trillion annually across the hyperscalers and their datacenter peers, we expect existing software incumbents; enterprise-focused LLM companies; and secure, highly governed agentic software technology startups to continue to show accelerating revenue growth. This will continue as next-generation compute platforms come online and the cost of expert intelligence declines radically through 2028.

The SaS era has begun. Enterprises and software companies that adapt business models and cost structures to align with AI labor-replacement processes will capture trillions of dollars of future enterprise value. The current SaaS-pocalypse marks the beginning of the enterprise AI infrastructure super-cycle, creating extraordinary opportunities for investors and technologies vendors.

PitchBook provides actionable insights across the global capital markets.

Additional research:

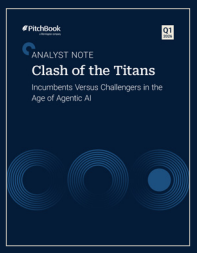

Q1 2026 Analyst Note: Clash of the Titans

Download the report here

Q3 2025 Launch Report: Advanced Computing

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:

Rudy Torrijos
Director, Industry Research

Derek Hernandez
Senior Research Analyst, Enterprise SaaS and Infrastructure SaaS

Josie Doan
Graphic Designer

Learn more about PitchBook's Institutional Research team.

Click here for PitchBook's report methodologies.