



ViPARTNERS

Re-framing “SaaS is dead” in the Era of Intelligent Software

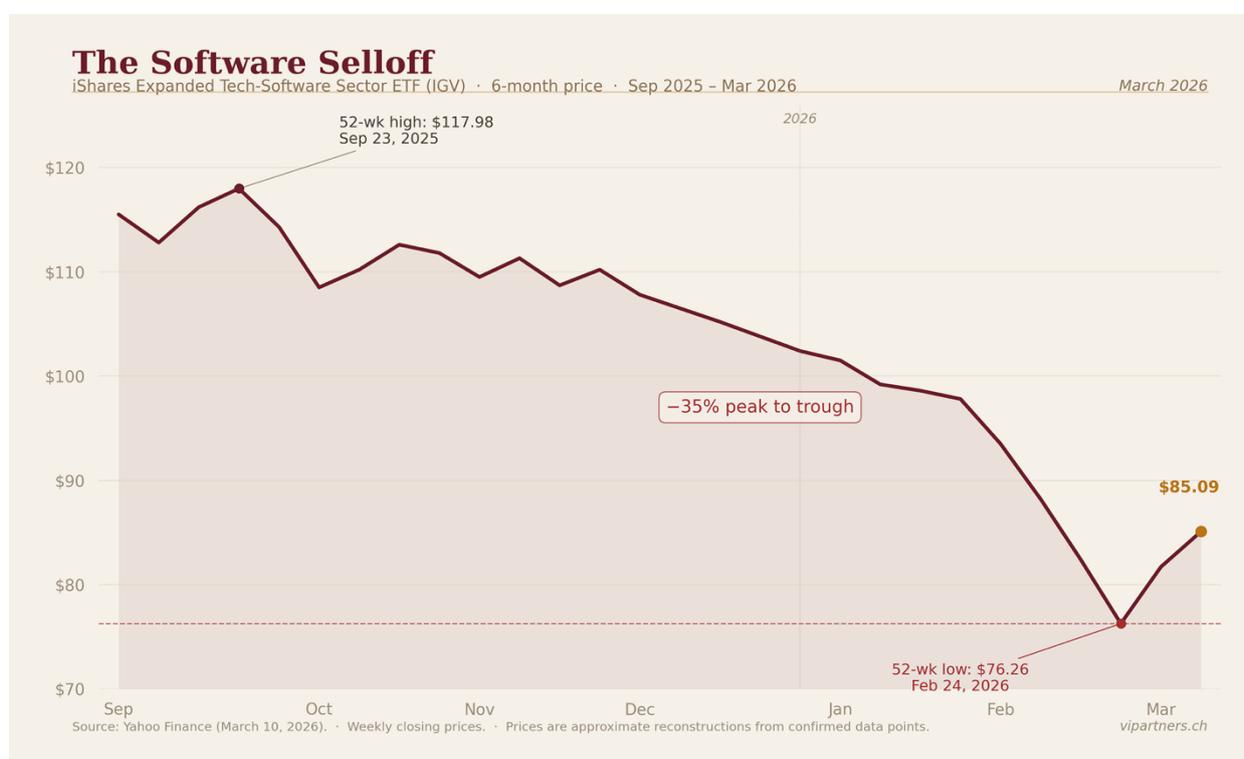
A Vi Partners Perspective on the 2026 SaaS Reckoning - navigating disruption, repricing, and opportunity in European software



From SaaS subscriptions to agentic workflows, the software pricing model of the last two decades is being repriced in real time.

The 2026 SaaS Reckoning

In [Unlocking Europe's AI Revolution](#), we outlined why we believe that AI is a massive opportunity and where it will impact industries the most. That perspective was forward-looking. This one is in reaction to what has been described a “SaaSapocalypse” unfolding in public markets in the recent months. The numbers are stark. By early February 2026, the [IGV ETF](#) (North American software) was down nearly 20% year-to-date and approximately 30% off its September 2025 peak. The sector’s [forward P/E collapsed from roughly 35x to around 20x](#), back to levels not seen since 2014. Public SaaS median EV/Revenue stands at 5.1x as of December 2025, down from a pandemic-era peak of 18-19x. Median revenue growth decelerated to 12.2% by Q4 2025, with further deceleration forecast through mid-2026.



The public market repricing has a direct knock-on effect for private software, compounded by structural problems specific to the PE-backed SaaS ecosystem. In fact, debt repricing is spreading across software portfolios as AI disruption reshapes credit markets, illustrated by the repricing of Thoma Bravo’s Verint deal and rising default expectations in private credit. At the same time, large PE firms are sitting on tens of billions of unexited, peak-valuation software investments from 2021, with IPO and M&A windows still largely closed and capital returns lagging. With public software multiples down sharply, exit opportunities remain compressed, leaving even strong companies struggling to offset valuation declines and pushing secondary pricing lower.

The reckoning has not fully arrived in private markets. Early-stage valuations lag public repricing by 6 to 18 months, which means the markdowns ahead of us are larger than the ones behind us. That

creates both risk and opportunity: existing holdings face further compression, but entry points are improving for disciplined buyers willing to act before the market catches up. As VC investors, we need to ask what this means for every company we hold and every new commitment we make.

The trigger of the SaaS reckoning was not a macro shock. It was a succession of agentic AI product launches, including tools capable of orchestrating entire white-collar workflows, raising a structural question: which software subscriptions survive when AI agents run the workflows? None of this is without precedent. Every technology transition of the past two centuries has produced a version of this correction, and in every case, the sell-off was more indiscriminate than the underlying disruption warranted. Re-reading Alasdair Nairn's *Engines that moves markets*, a study of every major technology cycle from the railway mania to the dot-com bubble, made it impossible to look at the current AI disruption without recognising the pattern: transformational technology, capital flood, valuation disconnection, reckoning, then concentration of value in a small number of durable franchises. We are in the reckoning phase. What follows is a framework for navigating it.

Our central claim is that the market's reaction is too undifferentiated and that it is not asking the right questions to make sense of the situation. The consensus question - "is SaaS dead?" - assumes AI replaces software and that enterprises will stop paying for software on a per seat basis. We believe in something more nuanced. First, we think that software will be the primary mechanism through which AI diffuses into the enterprise and second that "intelligent software" will capture larger chunks of the Opex of every enterprise so that they will spend more on "software" in aggregate rather than less.

In our view, GenAI (LLMs) and software are complementary, not substitutes. Enterprises run on predictability. They have near-zero tolerance for error, and their software reflects this: deterministic, auditable, high-throughput by design. Generative AI is the opposite - architecturally probabilistic, variable in output, and opaque in its reasoning. It lacks the process logic needed to run enterprise operations at scale. Even though causal AI models and other model architectures might solve some of those limitations, we think that a large part of agentic capabilities can be built on top of existing software architectures.

Software vendors that don't embrace AI capabilities to rearchitect their product for the "agentic" era will certainly be disrupted. However, AI is not necessarily a deadly competitor to software. AI generates intelligence; software processes, validates, stores, and acts on it. They are complementary layers, not substitutes. The real question is not whether enterprise software survives AI - it will - but which software becomes more valuable as AI is embedded deeply within it, and which AI renders redundant.

This paper provides a framework for telling the two apart at early stage. We start by dissecting the "SaaS is dead" argument to understand what this slogan entails. We then argue that software is the vehicle through which AI diffuses into the economy, and that this reckoning is partly cyclical partly structural, provided incumbents can reinvent themselves as intelligent software companies rather than

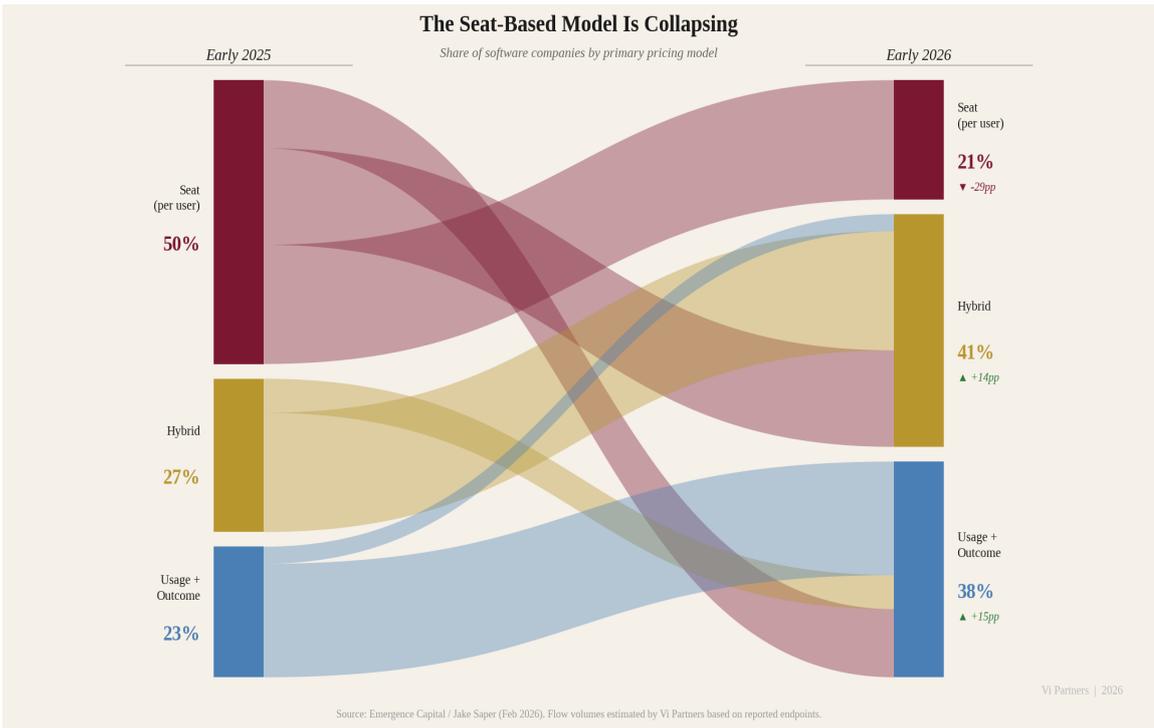
static repositories of workflows. Finally, we propose a triage framework to sort the winners from the disrupted, along with new metrics for measuring success in the intelligent software era.

1. Dissecting the “SaaS Is Dead” Argument

1.1 What the Public Market Is Pricing In

The sell-off is not random. It reflects five distinct, reinforcing fears and theories included in the slogan “SaaS is dead”:

- **The seat-count crisis.** AI boosts productivity per human, which means fewer humans needed, which means fewer seats licensed. [Atlassian](#) is the poster child: its stock hit a 52-week low on February 23, 2026, losing more than a third of its value over the course of February alone. If AI allows two people to do the work of twenty, per-seat revenue contracts mechanically, regardless of product quality. Any SaaS company whose revenue scales with headcount is exposed. Moreover, if we move from seat-based from outcome-based pricing, there is a large amount of uncertainty around the future contract value of SaaS companies.



- **AI labs attacking the application layer directly.** [Anthropic’s Claude Cowork](#), [OpenAI’s “Frontier Alliances”](#) with McKinsey, BCG, and Accenture, and Palantir CEO Alex Karp’s February 2026 statement that AI has made many SaaS companies irrelevant all point in the same direction. The fear: foundation model providers bypass existing software vendors

entirely and become the new enterprise application layer, relegating incumbents to legacy infrastructure.

- **Vibe coding and DIY in-house replacement.** Enterprises are exploring building internal tools with AI coding assistants rather than buying packaged SaaS. Reports have circulated of Fortune 50 companies planning to cut Salesforce and ServiceNow spend by 60%, replacing licensed software with raw API credits from model providers. The actual feasibility of this at enterprise scale is questionable (shipping a v1 is perhaps 2% of the work), but the market does not wait for nuance.
- **Budget reallocation from software to AI infrastructure.** Hyperscalers are spending over \$600bn on infrastructure in 2026, with roughly 75% directed at AI. CIOs are consolidating vendors: the average number of SaaS applications per organisation dropped from 112 to 106 by mid-2025, with 82% of organisations actively reducing their vendor count. Every dollar redirected to AI infrastructure, AI tooling, or AI headcount is a dollar not going to another SaaS seat or module. As Jason Lemkin at SaaStr framed it: AI is not eating the product, it is eating the budget.
- **Structural re-rating of terminal value.** The market is simultaneously raising the discount rate (higher perceived structural risk) and haircutting future cash flow expectations. Pre-AI, quality SaaS names were priced for 15-20% forward revenue growth; implied growth for many is now mid-single-digits. Hedge funds have generated approximately \$24bn in profits from short positions in software stocks in 2026 alone. The newspaper-to-internet analogy (95% value destruction between 2002 and 2009) is circulating in sell-side research as a worst-case framing. Public SaaS growth rates have declined every single quarter since the 2021 peak; the AI narrative gave the market permission to finally re-rate what the numbers had been signalling for three years.

As we can see, “SaaS is dead” encompasses many different theories about how SaaS is supposed to be declining. In our view, each of these arguments has merit on its own, but they should be applied separately and concretely when assessing individual software companies and the industries they serve. In the next section, we aim to examine in greater detail how to distinguish the survivors from the rest and what it takes to thrive in the era of intelligent software.

2. Who Survives the Reckoning and evolves as “ intelligent software”?

The market debate has crystallised around two questions. The first is whether the correction is cyclical or structural. The second is which categories of software survive.

2.1 Cyclical Panic or Structural Regime Change?

The bull case for incumbent SaaS rests on precedent market position and distribution, valuation, and timing. If full AI agent replacement of SaaS workflows is a post-2028 story and that today’s enterprise AI is copilot-style augmentation, not wholesale substitution, it leaves time for the incumbents to evolve to agentic “intelligent” software. Leading SaaS companies still carry gross retention rates above 90%. Valuations are at multi-year lows: Adobe trades at roughly 12x forward earnings versus a five-year average of 30x; ServiceNow at 28x versus a historical average of 67x. SaaS panics in 2016 and 2022 recovered within months. If this is cyclical and that those companies find a way to position themselves in the new paradigm of agentic workflows, these are generational entry points.

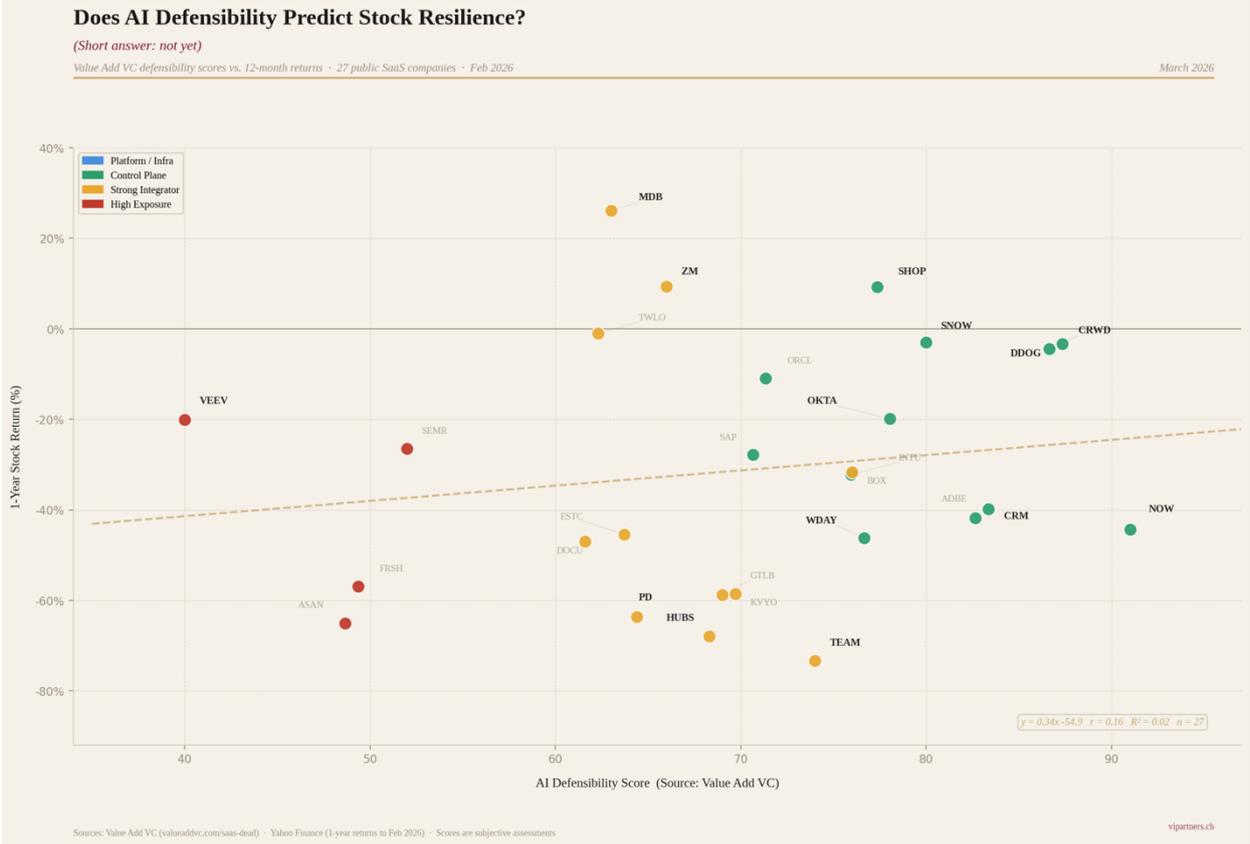
The bear case rests on mechanism, not analogy. AI agents do not need to replace enterprise software to impair it; they need only reduce the headcount that uses it. Public SaaS growth rates have declined every single quarter since the 2021 peak; the AI narrative probably gave the market permission to re-rate what the numbers had been signaling for three years. Software is the most net-sold subsector year-to-date, with hedge fund net exposure at record lows. An estimated [35% of point-product SaaS tools could be replaced by agents by 2030](#). The bear case is not that software disappears; it is that the seat-based pricing model is structurally broken and terminal values will never recover.

There is a paradox at the center of this debate: investors are simultaneously punishing hyperscaler stocks because large AI capex might not generate sufficient returns (“[the 600B question](#)”), and punishing software stocks because AI will destroy their business. Both cannot be true. If AI is powerful enough to disrupt established software categories, the infrastructure capex is early-stage investment that will compound over a decade and capture larger shares of the Opex of companies by replacing large labor pockets in the opex will justify it. If the capex is a write-off, the disruption threat to software recedes. The market has not yet resolved this contradiction, and that irresolution is itself evidence of indiscriminate pricing.

It is worth naming what the 2021 peak was. Technology bubbles do not emerge from technology alone; they require easy money, prior prosperity, and an uncritical press. 2020 to 2021 supplied all three. The correction is could be a normalization, not a fundamental verdict on software’s long-term value.

Our [analysis](#) based on the work of [Value-Add VC](#), confirms this empirically. We correlated AI defensibility scores against trailing stock returns for 27 public SaaS companies, measuring defensibility across five dimensions: proprietary IP, AI embedding depth, data moat strength,

distribution lock-in, and displacement risk. The headline result: a R-squared of 0.024. In other words, defensibility explains roughly 2% of one-year return variance. The market is not distinguishing between defensible and exposed companies in a systematic way.



The grade-bucket analysis tells a more actionable story than the raw regression. A-grade companies (scores 80-94, including CrowdStrike, Datadog, Snowflake) held up materially better, averaging -18.5% on a one-year basis with a median of just -4.4%. B-grades averaged -29.1%. F-grades averaged -47.4%. The dispersion narrows at the top: strong defensibility does not guarantee positive returns in a sector-wide selloff, but it provides a meaningful floor. This gap between fundamental defensibility and market pricing is precisely where patient investors should focus. It is also worth noting that in every technology transition, the losers have been easier to identify than the winners. The categories that are structurally impaired today, seat-based tools in agentic blast zones, are visible now. The winners will clarify only as the cycle matures. Our analysis tells us that that this might be cyclical for some of the stocks and for some, it might be structural (AI will render them fully irrelevant). Which camp a company fall into depends on his ability to transition to intelligent software in industries and workflow where enterprise customers need it.

2.2 Software as a Trojan horse for AI

The five fears outlined in Section 1.1. share a common assumption: that AI replaces software. The sell-off is real, but its bluntness creates mispricing. Companies with strong workflow lock-in, compliance requirements where responsibility is at stake and judgement is required, and deterministic execution needs (Veeva Systems in life sciences, Temenos in core banking, Procore in construction) have been sold off as aggressively as exposed SaaS feature-layers. The deterministic vs. probabilistic divide becomes critical to triage the casualties of AI: LLMs cannot substitute systems where 90% reliability equals 100% failure. Finance, healthcare, legal, and regulated manufacturing require deterministic, auditable execution. These categories require “judgement” (and identifiable accountability). In our view, these categories are mispriced by a market that has conflated all software into a single risk bucket without doing the triaging work required.

Foundation models trained on public data lack knowledge of optimized, large-scale private architectures developed over decades by incumbent vendors. Even with strong code, new entrants face daunting barriers: enterprise software requires years demonstrating 99.999% uptime, error-free operations across diverse IT environments, security, trusted brands, and enterprise sales forces. Switching costs remain prohibitively high given the risks of revenue disruption, productivity loss, and system failures during platform replacements.

What we are seeing instead is AI being “domesticated” within the application stack via domain-specific agents. We observe that many software companies have their existing software and data repositories as a trojan horse to diffuse intelligence to “solve” workflows. Consider [ServiceNow](#), which has embedded AI agents within its existing IT Service Management workflows rather than being replaced by them, or [Veeva’s Vault platform](#), where AI accelerates clinical trial documentation without displacing the regulatory-grade system of record. Same of Salesforce that is already generating half a billion alone [with its Agentforce](#). The franchise value of software companies is their domain expertise, which cannot be easily replicated. Across economic cycles and technology transitions (on premises to cloud, mainframe to PC), software companies with deep domain expertise embedded in customer operations have proven capable of sustained growth. The pattern holds further back still. GE did not win the electricity era on the strength of its generator; it won by controlling the installed base that could not switch. Microsoft did not win the PC era with a superior operating system; it won by owning the workflow layer enterprise customers had built upon. In every technology transition, it is what you do with the technology, not the technology itself, that determines where value accrues. 2026 marks the kick-off year for AI monetization within software, not a replacement of it. Obviously, it will require unwavering attention and courage from executive teams to execute the transition to “intelligent software,” sometimes sacrificing short-term profits for a longer-term prize to stay in the race. Some incumbents will be left behind given the speed and the depth of the evolution required.

The implication for investors is clear: the relevant distinction is not between “SaaS” and “AI” but between software that becomes more valuable as AI is embedded within it and software that AI renders redundant. Many incumbents already have their Trojan horse in place: deep workflow

integration, proprietary data, and entrenched customer relationships. The question is whether they can deploy AI within those footholds to move from recording transactions to setting them up and executing them end to end.

3. Measuring Success in the Age of Intelligent Software

3.1 Triaging criteria

If the structural thesis is even partially correct - terminal value is decreasing independent of the cycle - the question becomes which software sits on which side. Five criteria are worth monitoring closely:

- i. workflow depth (workflow-critical vs. workflow-adjacent)
- ii. business model (seat-based vs. usage/outcome)
- iii. deterministic requirements (regulated vs. general-purpose)
- iv. data lock-in (proprietary operational data vs. static repositories)
- v. budget line (labour budget vs. IT software budget).

Section 4 and 5 below operationalizes these into a detailed investment selection framework. Technical superiority does not appear in this list, deliberately. The winners of prior technology cycles, from automobiles to PCs, were rarely those with the best technology. They were those with the clearest market vision and the strongest embedded position from which to defend it.

The transition will be faster than the bulls suggest and less complete than the bears predict. It is structural for seat-based, feature-thin, horizontally exposed software (project management tools like Monday.com, basic CRM layers, standalone analytics dashboards) and cyclical for workflow-critical, compliance-heavy, data-rich systems of action (Veeva in pharma, NICE Systems in compliance, Temenos in banking). Regulated industries everywhere, whether financial services, healthcare, or critical infrastructure, require deterministic, auditable systems. The regulatory bar is rising globally, which reinforces the moats of companies built to those standards from day one.

The outliers in our defensibility analysis illustrate the distinction sharply. Atlassian (-73% one-year), HubSpot (-68%), and PagerDuty (-64%) all underperformed their peers by wide margins. The common thread: all three sit in the “agentic blast zone,” where the core value proposition (helping humans shuttle data, coordinate tasks, manage tickets) is precisely what AI agents are designed to replace. Conversely, MongoDB (+26%), Shopify (+9%), and Zoom (+9%) outperformed, each with differentiated positioning: database infrastructure with developer lock-in, a commerce platform with merchant ecosystem stickiness, and deep undervaluation reversion, respectively. The market will eventually sort these categories; the opportunity exists because it has not done so yet.

3.2 Software pricing and NRR as leading indicators

The structural shift in how software is priced is not a commercial detail; it is a leading indicator of whether a company understands its own value delivery mechanism and is evolving into an intelligent software company. [Seat-based pricing fell from 21% to 15%](#) of companies in just twelve months during 2025. Hybrid pricing (usage + seat) increased from 27% to 41%. Companies retaining seat-based pricing for AI products see 40% lower gross margins and 2.3x higher churn than peers on usage or outcome-based models. The logic is straightforward: per-seat pricing made sense when software augmented workers. It breaks down when agents replace the workflow entirely. Salesforce’s shift toward consumption credits for its Agentforce platform, and Palantir’s outcome-based contracting with government agencies, illustrate how even the largest incumbents are adapting.

Pricing Model	How AI Changes Things	Near-Term Setup
Per-seat	AI boosts productivity, fewer humans needed, seats shrink, revenue under direct pressure	Facing both multiple compression and earnings risk
Usage / consumption	More agents, more queries, API calls, logs, compute. Usage-based vendors collect a tax on AI prosperity	Natural AI beneficiaries; better insulated from disruption
Data / infra-linked	Not tied to headcount; AI amplifies data volumes and infrastructure needs	Fundamentals more stable; sentiment still weighs on multiples
Outcome / per-FTE	Directly aligned to value delivered; replaces labour budget, not software budget	Highest value-capture alignment; attribution complexity is the risk

[Traditional SaaS metrics are becoming unreliable](#) under these new models. ARR and NRR as currently defined assume revenue locks in at point of sale. Under usage-based pricing, revenue unfolds over time, shaped by how customers engage with the product. A company showing \$3M ARR under a usage model has fundamentally different revenue quality than the same number under an annual contract. New metrics are emerging (CARR, UARR, AI ARR, usage ramp rate) but standardisation is months to years away. Gross margins are compressing from 80-90% toward 50-60% as model infrastructure becomes a first-class cost line.

For early-stage investors, NRR trajectory is now the single most important leading indicator of business quality. Enterprise SaaS benchmarks provide context: [median NRR for companies with ACVs above \\$100K is 118%](#), with top-quartile performers exceeding 130%; [median GRR for the same segment is 90%+, with best-in-class above 95%](#). [Top-quartile B2B SaaS companies \(NRR above 113%\) sustain higher valuations](#) through both bull and bear markets. A company growing 200% with sub-100% NRR is a treadmill: growth is constantly refilling a leaking bucket. The critical question is whether retention is occurring at the workflow level (sticky, expanding) or at the feature level (churny, replaceable). This

distinction is observable even pre-\$1M ARR and becomes the most reliable early filter for separating durable businesses from AI wrappers before scale obscures the signal.

4. Portfolio Construction Implications

Our base case is selective disruption over a 5–10-year horizon. AI capability continues to advance extremely rapidly but unevenly: some white-collar tasks are automated within 2-3 years, while full workflow replacement in regulated, complex and “judgement” heavy domains takes significantly longer.

This base case produces clear implications for what to avoid and what to favour:

AVOID	FAVOUR
Seat-based SaaS in displacement-heavy verticals where AI agents substitute the human the licence was sold to (e.g. Asana, Monday.com)	Vertical SaaS with proprietary data moats and regulatory lock-in (e.g. Veeva Systems, Tyler Technologies, Procore)
Long-tail horizontal SaaS without a clear answer to “why can’t the next frontier model do this?” (e.g. Zapier, standalone BI dashboards)	Compliance and regulatory workflow software, structurally AI-resistant (e.g. NICE Systems, Varonis, SailPoint)
Interchange-dependent fintech and consumer fintech with friction-based moats	HealthTech: demand driven by demographics, not employment or macro cycle (e.g. Veeva, Doximity, Tempus)
Expensive SaaS at pre-2026 ARR multiples without clear AI defensibility	AI infrastructure and tooling: core banking, payment rails, data platforms, observability (e.g. Temenos, Adyen, Datadog)
CRM/sales enablement priced per rep (sales teams are shrinking)	Outcome/usage-based pricing models insulated from headcount reduction

Several structural implications follow for fund deployment and exiting SaaS portfolios. First, maintain dry powder: the PE overhang and European VC fundraising constraints mean secondaries and restructurings will create entry opportunities in good assets that hit capital limits. Second, HealthTech serves as a portfolio anchor: ageing demographics and underfunded public health systems create demand independent of white-collar labour dynamics. Third, pre-2026 renewal rate assumptions are no longer valid for underwriting SaaS ARR; revenue quality under new pricing models requires fundamentally different diligence.

5. Where We See the Opportunities

The analysis in the preceding sections points to five investment themes where AI is additive rather than substitutive, where pricing models align with value delivery, and where moats deepen rather than erode as AI capabilities advance.

AI Infrastructure and Tooling

Inference, orchestration, observability, and data tooling represent the essential plumbing of the AI era. Gross margins are recovering toward traditional software levels as the layer matures. Europe, and Zurich specifically, is already producing global-scale companies in this category: vector databases, AI-native ETL, and evaluation frameworks emerging from the research ecosystem (e.g. [Kadoa](#), [Weaviate](#)). [Datadog's expansion into AI observability](#) and Weights & Biases' MLOps platform illustrate how tooling companies are capturing the growing complexity of AI operations. The primary moat here is proprietary architecture and switching cost built into data pipelines. The risk is commoditization pressure from hyperscalers; this layer is only defensible where tooling creates genuine lock-in through integration depth and data gravity, not through features alone.

Vertical AI in Regulated Industries

Vertical AI has the potential to eclipse even the most successful legacy vertical SaaS markets. [By some measures, vertical AI captured \\$3.5bn in 2025](#) alone, triple the previous year's investment. The analytical logic is clear: domain-specific reliability is a domain problem, not a model problem. Foundation models cannot encode how a specific team at a specific firm does their specific job. Target verticals consistent with our prior thesis include precision manufacturing, legal/professional services, and financial services. In these domains, the moat stack is strongest: deterministic execution requirements create a regulatory barrier, proprietary process data deepens over time, and the full operational context of how work gets done across systems is nearly impossible for a general-purpose tool to replicate. [EvenUp](#) in legal (AI-powered demand letters built on proprietary case outcome data), [Abridge](#) in clinical documentation and [Unique.ai](#) in financial services workflow automation exemplify the pattern: AI-native entrants that own the full workflow in trust-sensitive verticals.

Financial Infrastructure

FinTech sits at the intersection of our strongest conviction themes: regulatory complexity, deterministic execution requirements, and proprietary data lock-in. The opportunity falls into two distinct categories.

First, *compliance and regulatory infrastructure*. [RegTech investment reached \\$4.8bn in 2024](#), with venture funding increasing 340% over three years, driven by the fact that global financial services compliance costs reached \$206bn across major markets. AI is transforming AML transaction monitoring, sanctions screening, and regulatory reporting, but these are domains where explainability, auditability, and zero tolerance for false negatives are non-negotiable. Companies like

Unit21 (AI-powered fraud and AML operations) demonstrate how AI augments compliance workflows without replacing the deterministic controls regulators demand. The regulatory bar is rising: the EU AI Act classifies credit scoring and fraud detection as high-risk AI systems requiring explainability and bias mitigation; sponsor banks are demanding real-time AML monitoring from fintech partners before any deal proceeds. This tightening regulatory environment deepens the moat for compliance-native platforms.

Second, *financial infrastructure and payment rails*. Core banking, payment processing, and settlement infrastructure are structurally insulated from the seat-count crisis because their revenue scales with transaction volume, not headcount. AI amplifies throughput (more automated payments, more real-time fraud checks, more cross-border transactions) rather than shrinking it. Temenos in core banking and Adyen in payment infrastructure exemplify this: deeply embedded, high switching costs, and AI-additive by design.

HealthTech

HealthTech deserves separate treatment because its demand drivers are structurally independent of the AI disruption cycle. Ageing demographics, underfunded public health systems, and chronic workforce shortages create demand regardless of what happens to white-collar employment or software budgets. Within HealthTech, AI is overwhelmingly additive: the ambient scribe market alone reached \$600M in 2025 (up 2.4x year-on-year), minting new unicorns like Abridge and Ambience alongside the market leader, Nuance's DAX Copilot. AI augments diagnostics, accelerates drug discovery, optimizes hospital operations, and automates clinical documentation. The moat hierarchy is particularly favorable: clinical data is among the most regulated and proprietary data categories in existence, brand trust is essential (patients and providers do not switch lightly), and network effects emerge naturally as more providers adopt a shared platform. The regulatory bar for new entrants is high.

Service-as-software firms

From the early days of the AI wave, we have believed in selling work, not software. As LLMs take over larger chunks of tasks, AI-enabled service companies that deliver better outcomes at lower cost are becoming venture-scale opportunities, given the immense TAM of global business services and AI's potential to expand margins. The natural entry points are close to traditional outsourcing (call centers, IT managed services, payroll, accounting), but the scope will expand rapidly as model capabilities improve and enterprises adapt. This category is also well suited to AI roll-ups: aggregating service firms and equipping them with AI infrastructure. In sectors where trust requirements are lower and price is the primary purchasing driver, very large "services-as-software" companies will emerge.

6. Filters for Early-Stage Investments

Synthesizing across the analysis in this paper and our own diligence experience, we organise our selection criteria into six filters. These are not independent; the strongest investments pass all six.

Filter	What to Look For	Red Flags
1. Workflow Depth	End-to-end ownership of a compliance-heavy, deterministic workflow where the product encodes domain-specific process knowledge. The company is the system of action, not a layer on top of one.	Feature layer over a general-purpose LLM. No switching cost. The product provides a UI for tasks AI can now perform directly.
2. Business Model	Outcome, usage, or data-linked pricing tied to measurable customer value. Credible gross margin path to 65%+. NRR trajectory exceeding growth rate.	Seat-based pricing with an AI wrapper. Permanent sub-60% gross margins accepted as structural. High growth masking sub-100% NRR (treadmill economics).
3. Deterministic Requirements	Operates in a domain where 90% reliability equals 100% failure: financial reporting, clinical decisions, regulatory compliance, safety-critical manufacturing. LLMs cannot bridge this gap today.	General-purpose productivity tool where probabilistic output is acceptable. No regulatory requirement for auditability or explainability.
4. Data Lock-In	Proprietary, operationally embedded data that deepens over time. The data loop creates compounding defensibility: more usage generates more proprietary data, which improves the product, which drives more usage.	Static data repositories. Data sourced entirely from public or licensable datasets. No proprietary training signal from customer operations.
5. Strategic Position	Selling into the labour/operations budget (CFO as buyer). Replacing headcount, not software. Domain-expert founders with credibility to sell into regulated verticals. Clear, specific answer to “why not GPT-5?”	Selling into the IT software budget (competitive, price-sensitive, being actively consolidated). No domain differentiation from general-purpose AI.

The Swiss ecosystem is a structural advantage here. [ETH Zurich and EPFL's 1,100+ spin-offs](#), the 700+ AI-focused PhD researchers in Zurich and Lausanne, and the country's [fourth-place global ranking for AI research intensity](#) create a pipeline of founders who combine deep technical capability with domain expertise.

Conclusion

The SaaS reckoning is real; however, the market's response has been too blunt. The opportunity for private investors is in the gap between indiscriminate public market panic and the structural durability of workflow-critical, compliance-heavy, deterministic software in regulated verticals. History suggests the window is finite: as the cycle matures and the survivors become identifiable, the entry premium closes.

This is precisely where our prior AI thesis pointed: full-workflow ownership, trust-sensitive sectors, and technical capability translating into domain-specific products. The core beliefs we articulated all hold. What has changed is the urgency. The market correction accelerates the timeline and sharpens the selection criteria. In every prior technology transition, the investors who acted on a framework built before the consensus formed captured the vintage. That moment is now.

Our central argument bears repeating: the market is asking the wrong question. Software is not the casualty of AI; it is the trojan horse through which AI diffuses in enterprises. AI models generate intelligence; software makes that intelligence deterministic, auditable, compliant, and operationally embedded. We are in an evolution toward intelligent software. The companies that pass our six filters (workflow depth, business model alignment, deterministic requirements, data lock-in, strategic position, and market timing) are not threatened by AI. They are the diffusion engine. As in other technological transition, it will require AAA teams who can adapt to the extremely rapid pace of change and who are bold enough to iterate rapidly around their product vision. Two centuries of technology transitions confirm the logic: the infrastructure layer never wins; the application layer that makes the infrastructure indispensable does. This will be also true in this new era of intelligent software.