

The Software Panic Is Two Years Late

What two independent datasets say about software, disruption, and where the risk actually sits

AUTHOR: ROSALIND BAZANY, PARTNER

Asset classes categorised as high-risk are increasingly being reconsidered in the context of risks created within supposedly stable allocations themselves. A large-cap public software allocation carries more directional risk than most portfolio frameworks acknowledge: a concentrated bet on specific architectural assumptions that are now being challenged at speed and at scale. Early-stage innovation exposure starts to look different in this context. Less a return-enhancement play, to my mind, and more a structural hedge against the kind of disruption that conventional allocations are now directly absorbing.

The question worth asking is not whether software is dying but which layer of the software stack is being replaced, and whether portfolio construction more broadly reflects a genuine view on that question.

Below is an attempt to answer it with data.

The fear, named precisely

The SaaS sell-off of early 2026 reflects a coherent set of concerns. AI agents are executing work that software seats used to be paid for, the interface layer of traditional software is being bypassed by tools that require no UI at all, and companies whose competitive moats were built on switching-cost friction rather than genuine product depth are being appropriately repriced.

The contagion has extended beyond public equity markets. Private credit lenders are applying the same logic to software companies that may have nothing structurally in common with the incumbents under pressure, compressing ARR multiples and tightening refinancing terms across a wider category than the underlying dynamics actually warrant. LP portfolios with software exposure in both equity and credit are absorbing pressure from two directions simultaneously, regardless of whether the specific companies they hold have anything to do with the problem.

None of this is an inaccurate description of what is happening to a specific type of software company. The problem lies in the extrapolation: that what is true of seat-based, interface-heavy enterprise software with legacy architecture tells you something meaningful about early-stage venture exposure to software broadly. The data suggests otherwise.

Two datasets, one direction

In 2023, for the first time, the majority of founders Antler backed in a given year were building AI companies. Not because we changed strategy or published a new thesis. It showed up in the data. 54% of that year's classified investments fell into companies where AI was a core or significant part of the product, up from 41% the year before. By 2025, that broad figure had reached 83%. The stricter number is more telling: 57% of 2025 investments were companies where AI is the core product itself, up from 9% in 2022.

This happened company by company, irreversibly, two years before the market began debating whether software was dying.

The software panic arrived two years after the transformation it is panicking about. In our portfolio, the disruption had already happened.

That shift is documented across two independent sources that use different methodologies and ask different questions. The first is Antler's internal portfolio classification: almost 1,800 companies assessed by Antler analysts across our full history from 2019 through 2025, with every company classified by AI type, business model, and outcome status. The second is our 2025 portfolio health check, in which 529 portfolio companies from pre-seed through Series B+ described their AI use and product direction in their own words. One is an

analyst-applied external view; the other is founders describing their own reality. They converge on the same picture.

Antler uses four AI categories, applied consistently since 2019. ‘Infrastructure and Developer Tools AI’ and ‘Application Layer AI’ are companies where AI is the core product. ‘AI Enabled’ covers companies where AI is a meaningful component but not the defining proposition. ‘Non AI’ is everything else.

Year	n	Infra & Dev Tools AI	Application Layer AI	AI Enabled	Non-AI
2019	86	1%	14%	15%	70%
2021	186	2%	11%	32%	56%
2022	290	1%	8%	32%	59%
2023	328	1%	21%	32%	46%
2024	423	13%	30%	22%	35%
2025	341	16%	41%	26%	17%

Source: Antler internal portfolio classification, 1,792 companies across all status categories (active, written off, and exited) as of end 2025. Figures reflect investment decisions by year, not current active portfolio composition. Percentages exclude companies with no recorded investment date or incomplete classification (n=0 to 8 per year). Totals may not sum to 100% due to rounding.

The headline shift has come from Application Layer AI (8% in 2022, 41% in 2025) and Infrastructure and Developer Tools AI (near-zero through 2022, 16% by 2025), while non-AI collapsed from 70% of the portfolio in 2019 to 17% last year. The portfolio is not simply doing more AI; it is doing progressively deeper AI.

Outcomes track the same gradient. Write-off rates run consistently lower as AI integration deepens, with the pattern holding across every cohort filter we apply.

AI Classification	Write-off Signal	Caveat on interpretation
Infrastructure & Dev Tools AI	Lowest observed	Predominantly 2024–2025 vintage; least time to reach write-off outcome
Application Layer AI	Lower	Growing fastest as share of new investment; cohort skews recent
AI Enabled	Moderate	AI as component, not core product; older average vintage
Non AI	Highest observed	Largest cohort; includes pre-AI era companies from 2019–2022

Source: Antler internal portfolio classification, 1,792 companies across all status categories (active, written off, and exited) as of end 2025. Write-off signal reflects relative ordering only, not absolute rates. Classification applied by Antler analysts using a consistent taxonomy since 2019. Infrastructure and Dev Tools AI companies are disproportionately 2024–2025 cohort investments and have had less time to write off, which flatters their headline rate.

The gradient is a correlation, not a causal finding, and vintage effects are likely to explain part of it. Non-AI companies are disproportionately concentrated from 2019 to 2022 vintages and have had four to seven years to reach a write-off outcome; infrastructure AI companies are mostly 2024 to 2025 investments with one to two years at most. We present it as a signal worth monitoring rather than a conclusion.

The second dataset arrives from a different angle. Our 2025 health check asked 529 portfolio companies to describe their AI use directly. Across the full sample, 98% describe themselves as using AI actively or moving toward it. That figure is not surprising, but it is ours: first-party data across 529 companies, rather than a survey of intent or a market estimate. The more revealing cut is how AI integration changes across funding stages:

Stage	Building Own AI	3rd Party Integration	Exploring	No AI
Pre-Seed (n=446)	37%	47%	14%	2%
Seed (n=71)	48%	37%	14%	1%
Series A (n=13)	8%	54%	31%	8%
Series B+ (n=3)	67%	33%	—	—

Source: Antler Portfolio Sustainability Health Check 2025, pre-seed through Series B+. Self-reported by founders; responses reflect a point-in-time snapshot. Distributed to all active portfolio companies meeting eligibility criteria at time of completion; 529 responses received, representing a 49% response rate across the eligible global portfolio.

Nearly half of seed-stage companies are building proprietary AI models rather than integrating a third-party API, and the Series B+ cohort skews further still toward builders. When read alongside the write-off gradient from the internal dataset, the two sources are telling the same story from opposite directions: depth of AI integration correlates with durability, and the founders building at the deepest layer are the ones pressing into that bet hardest as they scale.

Two independent datasets using different methodologies, different samples, and different questions arrive at the same composition shift. That convergence is the story here.

What is at risk

None of this means early-stage software is insulated from the pressures driving the sell-off. The companies most directly in the line of disruption are not simply those without AI integration but those whose revenue model places a human being at the point of consumption, because the seat-based assumption is what the current shift is dissolving, replacing per-user pricing with token consumption, outcomes, and machine-driven workflows that neither require nor recognise the old unit of sale. Some of it is directly exposed for exactly the same reasons as the incumbents being repriced: point solutions without data moats, single-feature tools that an agent can replicate, workflow products without meaningful integration depth carry the same risks regardless of stage.

The non-AI share of the portfolio is not automatically a liability. Some categories demand determinism over probabilistic output, particularly in regulated industries and domains where auditability matters more than speed. The 17% of 2025 investments classified as non-AI includes companies making a deliberate design choice, not ones that missed the memo.

Within early-stage venture, diversification still matters, and the form it takes is more textured than simply spreading across asset classes. Antler's portfolio spans 27 markets, which in practice means exposure to companies operating inside materially different regulatory environments: the EU AI Act shapes product decisions for European founders in ways that US-based operators do not face, and APAC markets are seeing different incumbent software stacks being disrupted on different timelines. When geopolitical volatility spikes, as it has repeatedly across 2025 and into 2026 with export restrictions and shifting tariff regimes, the effect on an early-revenue company in Oslo is structurally different from its effect on a US-listed software ETF.

The defensive instinct running through markets right now, rotating toward capital-intensive sectors on the basis that physical assets are harder to displace than digital ones, is a coherent response to a specific pressure. But several of the sectors drawing capital on that basis are among the most dependent on the digital infrastructure supposedly being rotated away from. The competitive differentiation in advanced manufacturing, grid operation, and industrial logistics is increasingly sitting in the software and AI layer built on top of the physical asset, not in the asset itself. Owning what disruption cannot easily reach is not the same as seeing where it is going next, and for a long-duration portfolio, that distinction deserves at least as much attention as the near-term rotation trade.

The more honest structural concern is the pace of change relative to the visibility horizon. A company that looks defensible today may look exposed in eighteen months as AI capabilities extend into adjacent problem spaces. The disruption is moving faster than most portfolio review cycles, which is part of why proximity to the earliest cohort of founders provides a signal that later-stage investing cannot replicate by design.

What the data leaves open

The sell-off will find a floor, as these things do, but the sequence it has exposed will not reverse. The companies in this data were building the replacement architecture before the incumbents had named the problem, which means the disruption was not a response to a market signal but a precondition for one. The market did not lead; it caught up, two years late and at speed.

For allocators, the more uncomfortable implication sits not in what the data says about early-stage venture but in what it implies about the portfolios sitting adjacent to it. A large-cap software allocation was not, in my view, a neutral position in 2023. It was a concentrated bet on design assumptions that were already being dismantled, at pre-seed, in markets most institutional investors were not watching. The labels attached to those allocations, stable, defensive, core, described their historical behaviour rather than their forward exposure, and the distinction between those two things is what the current repricing is forcing into the open.

For allocators without early-stage exposure, the return case is its own conversation. What I find more interesting is the informational dimension. An institution with genuine proximity to company formation sees the next disruption cycle while it is still a founder's working assumption rather than a market consensus. That visibility does not guarantee outperformance at the fund level, and it would be intellectually dishonest to suggest it does. But it provides an earlier read on which architectural bets are compounding, which categories are being vacated, and where the pressure on existing allocations is likely to arrive before it shows up in public market data. In a period when the most consequential risks have been migrating from the periphery of institutional portfolios toward their core, I think that vantage point carries a value somewhat independent of the financial return it may or may not generate.

Composition has always mattered in venture. The current environment has made something legible that portfolio construction frameworks were not designed to surface: composition risk exists everywhere, and the tools most institutions use to assess software exposure were built for a technology cycle that has already closed. That is not necessarily a comfortable conclusion for anyone sitting across a portfolio review.

ABOUT ANTLER INTELLIGENCE

Antler is one of the most active early-stage venture firms globally, with US\$1.3B+ AUM, operations across 27+ cities, and a portfolio of 1,800+ companies spanning every major technology market. We invest from pre-seed and seed, where the majority of our portfolio sits, through to Series C via regional early-stage funds and global follow-on vehicles. That structure gives us something few investors have: a longitudinal, cross-geography dataset of founder behaviour, category formation, and technology adoption that begins before most institutional investors are in the room.

Antler Intelligence is our research and market insight function designed for an LP audience. The views expressed here are the author's own, informed by Antler's proprietary portfolio data. They represent an analytical perspective on market trends, not the firm's investment recommendations or the position of any Antler fund. We publish this analysis because the early stage vantage point is underrepresented in market discourse, and because our LPs and ecosystem partners benefit from understanding what we are seeing before the market names it.

DATA AND DISCLOSURE

This article may contain forward-looking statements based on current assumptions, which are subject to change. No obligation is undertaken to update this material. It is not intended for distribution in any jurisdiction where such distribution would be contrary to applicable law.

Internal portfolio data reflects Antler's Hub dataset as of end 2025 across 1,792 companies. AI classification is applied by Antler analysts using a consistent four-category taxonomy applied across all cohorts since 2019: Infrastructure/Dev Tools AI, Application Layer AI, AI Enabled, and Non AI. A small number of companies (zero prior to 2024; eight in 2025) are unclassified and excluded from all percentage calculations. Write-off rates are directional signals and should not be treated as precise absolute figures or as indicative of future results. Infrastructure and Dev Tools AI companies are predominantly 2024–2025 cohort investments with less time to write off. Health check data covers respondents to Antler's 2025 annual portfolio sustainability survey across pre-seed through Series B+; responses are self-reported and reflect a point-in-time snapshot. Portfolio composition figures reflect classified investments only, by recorded investment date.

For further discussion on portfolio composition or Antler's research, contact investor.relations@antler.co or visit our website [Antler.co](https://antler.co).