

Private Equity International

July/August 2026 • privateequityinternational.com

EXTRA
This year's
Global Investor
150 ranking



**Private equity reckons with
software volatility**

Ignore

Accept




How AI is upending PE's software sweethearts



Cover story

*This year's 'SaaSpocalypse' has exposed the vulnerability and deepening uncertainty of outcomes in a much-beloved sector. Investors must work quickly to identify managers capable of thriving in the age of AI, writes **Alex Lynn***



Never send a human to do a machine's job." Those were the words of Agent Smith, an artificial intelligence system policing threats to the established order in 1999's *The Matrix*.

The line, immortalised by actor Hugo Weaving, was prescient. Less than three decades later, AI models – some of which are aptly known as 'agents' – threaten to subsume or even replace entire industries, with software among those most at risk.

The stakes are high: a February survey from placement firm Rede Partners found that 77 percent of LPs allocate between 10 percent and 30-plus percent of their private equity exposure to software. Software-as-a-service has historically been particularly sought after due to its predictable, reoccurring and high-margin revenues.

"For years, SaaS was the most investable asset class," Alex Baker, US TMT deals leader at PwC, tells *Private Equity International*. "You'd invest in a SaaS business because it was a SaaS business – it was a financial innovation as much as a product innovation. And so, you could throw a growth multiple on a business in almost any subsector and ride that wave."

This dynamic is now under threat. In early 2026, public markets experienced a major rout in software stocks prompted by mounting fears that AI could undermine business models. Forbes reported in February that the event – now widely referred to as the 'SaaS apocalypse' – wiped some \$1 trillion from software market caps in the first five weeks of the year alone.

Then, in early June, software stocks sharply rebounded from March and April lows, albeit with a bifurcation between strong performers and those who remain in the red. These

unpredictable highs and lows reflect the uncertainty of outcomes in a sector with the potential to be disproportionately impacted – both positively and negatively – by AI.

Such volatility also spells potential trouble for private equity-backed software assets, which are valued in part relative to their public market comparables, and places an onus on their owners to ensure these businesses end up on the right side of technological history.

"It's certainly something that we're concerned about; it's an issue for our LPs, and one they're seeking to understand better," says Garvan McCarthy, global alternatives chief investment officer at LP consultancy Mercer. "They're looking through existing portfolios and identifying meaningful exposure to this sector."

How to identify the managers that can navigate – or even benefit from – this pivotal moment has quickly become the multibillion-dollar question.

From SaaS to finish

To understand AI's potential impact on software portfolios, one must first understand precisely how and why the technology is changing the way these underlying companies operate.

SaaS businesses have traditionally operated on a per-seat basis. A super-market group might, for example, use software to track stock levels, manage reordering and co-ordinate store deliveries. Store managers and warehouse staff would all need to access the software, potentially requiring a subscription for each user.

Private equity firms have historically enjoyed this model because it creates sticky recurring revenues – sticky because subscriptions are paid annually, and because many of these products are 'mission critical'. A supermarket chain would struggle to function without the ability to efficiently manage its stock, for instance.

Coding this software can be complex, providing something of a competitive moat, and asset light, requiring minimal capital expenditure.

AI threatens to upend the status quo. “Fundamentally, the disruption that AI is causing to the software space is it’s forcing software players to rethink three categories: how their product is created, how their competition is positioned, and how users are actually engaging with their product,” says PwC’s Baker.

For starters, AI has made software much easier and quicker to develop. The supermarket chain could, in theory, use AI to write its own stock tracking software, thus eliminating the need for a third-party subscription.

Second, ‘frontier models’ – advanced, general-purpose models such as ChatGPT, Gemini and Claude – have many different applications in a broad swathe of industries, thus competing with niche incumbent software developers.

Lastly, AI is altering the value and the use of software. In some cases, AI becomes an agent, using software automatically to complete processes without human input. This might, for example, include data management of customer support chatbots.

Therein lies the rub: how does a software business charge users per seat if a single AI model is doing all the work?

“This is probably the biggest disruption, both downside and upside,” says Baker. “What that means is every software platform has to rethink how it’s built and how it’s priced. And that’s the existential threat that does exist for some portion of the market.”

Pricing in the AI era will likely follow one of two models. The outcome model could, for example, involve charging based on how much the software improves a client’s customer churn and retention stats, a percentage of the

money it helps to save, or how much time it takes to resolve a specific problem. The action model would involve charging the agent for each time it completes an action using the software: instead of paying \$20 a month, the clients might pay 5 cents per action.

“The last [option] is hybrid,” says Baker, “which to be honest is probably the most common in the near term. The challenge here is outcome-based pricing is hard – you need to have a lot of confidence in the outcomes that you’re going to drive.”

What’s more, upending highly successful and longstanding business models is easier said than done.

“If you go back through history, that was the biggest obstacle for a lot of the leading companies, because they were afraid or moved too slow to disrupt themselves, and so their competition did it instead,” says Brian Lam, a Hong Kong-based portfolio manager in JPMorgan Asset Management’s private equity group. According to PEI data, JPMAM has backed software funds from the likes of London’s Hg,

August Equity and Axiom Equity, and Washington, DC’s Dcode Capital.


“While disrupting a business with a strong client base and sticky revenues is challenging in the short term, the key question is: will that business still be growing five or 10 years from now, or are they sitting on a declining pool of business that can’t really move anywhere else?” Lam adds.

Credit crunch

AI’s potential threat to software portfolios could be exacerbated by higher-for-longer interest rates – particularly for assets that have been held for extended periods.

In February, LP advisory Cambridge Associates noted that software assets acquired before 2022 rate hikes are most at risk as they were “deployed at entry values, growth rates and leverage assumptions reflecting a bygone era”. The adviser added: “It is not surprising that returns have come down and distributions have slowed for these vintages overall.”

The median holding period for the average PE-backed company rose to





four years in 2025, up from three years in 2022, according to data from special sits firm TREO Asset Management. The percentage of companies held for more than seven years has also risen from a decade low of 14 percent to nearly 20 percent.

According to Houlihan Lokey's Private Credit DataBank, a quarter of currently held software borrowers have a loan-to-value ratio above 50 percent, leaving them in a poor position to absorb valuation shocks.

The implications can be severe. Medallia, a struggling platform owned by Thoma Bravo, is a case in point. The business was valued at \$6.4 billion when it was acquired in October 2021; in April 2026, Reuters reported that Thoma Bravo was preparing to hand over the firm to lenders in a deal that would wipe out \$5.1 billion in equity. In June, a consortium led by Blackstone, Apollo and KKR agreed to take control of Medallia and inject \$150 million of new capital into the business, according to a statement.

SaaS hysteria

In spite of this disruption, software revenues at large have so far proven comparatively resilient. According to Houlihan Lokey, 69 percent of

software companies achieved positive year-on-year revenue growth in 2025, while almost as many (67 percent) increased their EBITDA margins.

Investment bank Lincoln International's Q1 Private Market Index, which tracks changes in the enterprise value of US privately held companies, found that 74.9 percent of software companies grew revenue in the first quarter, and 67 percent grew EBITDA. Year-on-year growth was 6.6 percent for both revenue and EBITDA.

"The performance in Q1 is the strongest we've had," says Jean-Baptiste Brian, co-chief executive of European software giant Hg. "Q2 is looking like it's very strong as well, and that is across top line and bottom line. AI [is] contributing to both top line, with additional bookings, and bottom line, with actual productivity gains that we can choose... to reinvest in product."

And yet, privately held US business values declined by 2.2 percent in the first quarter, Lincoln found, largely driven by a 7.8 percent decline for tech companies and an 8.8 percent fall amid software assets.

Data from PEI Group's SIPA – an index, benchmark and rating provider for private infrastructure and private equity markets – shows that an index of

subscription-based global information and communication businesses delivered a negative 2.44 percent annualised return on a USD basis over a three-year period and a 3.5 percent return over five years, compared with 12.93 percent over 10 years. This reflects a decline in valuations linked to higher rates and falling valuations, as well as a lukewarm exit environment, according to SIPA analysis.

Hg's marks were down around 5 percent from end-December to end-March, Brian notes, despite the portfolio being up 5 percent in past-12-months trading. This, he says, is partly a function of the firm's valuation methodology comprising a mix of precedent transactions and listed comps.

"We've taken a big hit on the portfolio as of March... purely from the de-rating of the software comps during that period. We've communicated that to LPs in a very proactive way, saying there was no point in us changing anything in how we do normally do things, and as a result, they should expect valuations to come down despite the portfolio trading up strongly."

A contradiction between revenue growth and valuation decline could suggest that AI's threat is – for now, at least – more about perception than reality.

For the time being, software assets are trading at lower valuations. SIPA data shows that an index of 7,326 private equity-backed global computer programming, consultancy and related subsector businesses were transacting at an average 16x EV/EBITDA multiple as of March 2026, down markedly from 21.9x three years prior.

"Multiples paid in recent years were very high, often exceeding 20x beyond a certain size. Valuations were ultimately driven by relatively simple and consistent KPIs from one transaction to another," says Alexandre Armbruster, head of private equity and

infrastructure funds at France's Caisse des Dépôts et Consignations, which has more than €507 billion in assets under management.

"The sharp fluctuations observed between January and June... show that these valuations are complex to assess, with significant differences from one asset to another, and that the market has an uncertain view of the future and terminal value."

Though CDC does not disclose its private equity allocation, software


accounts for around 8-10 percent of its exposure to the asset class – most of which is in buyouts, Armbruster notes. According to *PEI* data, its recent commitments include Hg Genesis 11 and Hg Mercury 5, as well as FPE Fund IV from UK software specialist FPE Capital.

"It's also key to keep in mind for valuations that you need to invest more than used to be the case, especially for these super asset-light companies," notes Armbruster. "So, everything

equal, just because of the capex, valuations are not the same."

Uncertainty around valuations could hypothetically throw yet another spanner in the works for private equity's long-awaited exit recovery, with GPs unwilling to compromise on returns.

Hg, which has \$110 billion of assets under management and around 60 portfolio companies, is taking a pragmatic approach to realisations in this downmarket. The firm completed



seven liquidity events in Q1 2026 alone and has generated around \$24 billion of proceeds for funds and co-invest vehicles since the beginning of 2022.

“We do not time [the] market at Hg – never have, never will – so we buy and sell in good and bad markets,” Brian explains. “We have realisation budgets every given year, and we stick to those. Sometimes you quote-unquote ‘get lucky’ and... exit in a supportive exit environment where ratings are strong and [the] buyer’s appetite is strong, and frankly, sometimes you don’t... And we’re definitely [in] one of those markets right now.”

Factory settings

AI doesn’t have to spell disaster for software assets. Vista Equity Partners – one of the world’s largest private equity software specialists, with around \$110 billion of AUM – believes there is significant upside to be captured for companies with the ability to rapidly incorporate AI.

In its 2026 outlook, Vista cited Bain & Co data suggesting there could be \$5 trillion to \$7 trillion in cumulative value flowing to the software and application layer by 2030 – a more than 4x expansion from current levels. Agentic

solutions alone represent an estimated \$3 trillion of incremental value, linked to software’s expanding role in executing work that was previously done by humans.

“An agent has the ability for us to effectively create usage of our software to automate full workflows,” Monti Saroya, co-head of the Vista flagship fund, tells *PEI*. “We can go and attack actual physical labour and some of the low-level things that humans are doing from a workflow standpoint, like inputting and outputting, that should be done by agents, freeing the human to go do higher-value work.”

To that end, Vista in 2025 launched an ‘agentic factory’ designed to help scale agentic AI across its enterprise software portfolio. That was followed by a partnership in April 2026 with Google Cloud, which provides Vista portfolio companies with access to Google’s Gemini models, AI Hypercomputer and Gemini Enterprise – a platform for building and deploying AI agents.

“We basically put our portfolio companies through our own set of best practices with subject matter experts that are helping them deploy the latest and greatest generative AI,” Saroya explains. “We have 54 companies with

agent solutions [in] the marketplace, and there will be no software company left behind because all of them will go through this agentic factory.”


Hg launched its own AI incubator, Hg Catalyst, last year. The initiative embeds members of Hg’s more than 100-strong in-house AI team into its portfolio companies for a set period – typically three months – with a view to creating and producing agentic software products.

Hg Catalyst has around 1,600 AI projects running in its portfolio, representing about \$260 million of budgeted EBITDA impact across these businesses. Some 20 percent of its portfolio companies are generating over 10 percent of new bookings from AI. “That team is in very, very high demand,” says Brian.

The potential upside can be huge. Vista, for example, has reportedly agreed to sell healthcare administration business Alegeus Technologies to UnitedHealthcare in a deal worth around \$3 billion, per Bloomberg. The sale is expected to generate a 4.6x gross multiple of invested capital, according to a letter to Vista investors seen by *PEI*.

“We had this phenomenal SaaS business in its own right, organically growing with really strong retention,

TMT funds’ proportion of global private equity fundraising



Source: Private Equity International

and then we started using these agents to make ourselves more profitable on everything from code development to how we support our customers,” says Saroya, declining to comment on investment performance.

“We started winning more customers because we were using it in sales. If you can start to agentify your own company and create these agents internally, the large players in our space are looking at it as an acquisition opportunity and being able to compound on that.”

Inference cost


Nevertheless, cost remains a key challenge for many GPs hoping to incorporate AI into their portfolios. AI model usage is measured in ‘tokens’, representing the amount of data units consumed to solve a problem – also known as inference. ‘Inference cost’ refers to the amount a company has to pay each time an AI model generates a response.

There is plenty of noise around the direction of travel for inference costs. Advisory firm Gartner reckons inference costs on large language models will cost AI providers at least 90 percent less in 2030 than it did last year, with improvements driven by semiconductor and infrastructure efficiency improvements, model design innovations and higher chip utilisation.

Vista’s Saroya believes the picture is more complicated in the near term. “In 2026, the cost of inference is going up, not down. The issue is that usage of inference is being done by AI agents, rather than humans, so the maths around how many tokens they use is a little bit fuzzier. The cost of inference is incredibly important because that impacts our gross margins.”

In May, for example, Axios reported that one AI consultant had an unnamed client accidentally spend \$500 million in a single month after failing to limit employee usage of Claude.


Valuations for computer programming, consultancy and related businesses (EV/EBIDTA, x)



Data based on an index of global PE-backed companies that operate in the global computer programming, consultancy and related sub-sectors

Source: SIPA, PEI Group

Annualised returns among privately owned ICT businesses (%)



Data based on an index of global PE-backed companies that operate in the information and communication class. Returns calculated on a USD-basis

Source: SIPA, PEI Group

Vista believes it has found a solution: the firm teamed up with Cambium Capital in February to co-lead a reported \$350 million funding round for AI chip business SambaNova. The company is said to produce inference chips that are faster than those produced by its competitors and that run at a lower cost. “It’s allowing us to reduce the cost of inference for our end customers and our companies,” says Saroya.

Hg, for its part, sees the price of each token that a model processes falling – albeit partly offset by newer

models becoming more capable and consuming more tokens. It considers inference costs to be something for Catalyst to help businesses manage, rather than a barrier to adoption.

“Start with the best available model to prove the concept and get a product into customers’ hands, then optimise the cost base once you know what ‘good’ looks like,” says Brian. “That means measuring everything from day one, so you can see exactly which parts of a workflow drive cost for a given level of accuracy, and then engineering those costs down through

better architecture and smarter model choices over time.”

Winners and losers

Some GPs may get left behind in this technological space race.

“Disruption creates winners and losers,” says Mercer’s McCarthy. “[This] may play out quite quickly, and more quickly than maybe other changes in the past. We may not – as investors, and as perhaps GPs on the other side – have the luxury of this being a slow trend that will evolve and that people can pivot away from over time.”

This pace of change could create a headache for LPs pondering which GPs to back.

That uncertainty is being reflected in due diligence processes. According to Rede’s survey, LPs are increasingly focused on understanding the composition of software GPs’ portfolios by product category, end-market, business model and AI vulnerability, as well as the leverage levels supporting those assets.

“Whenever I engage with GPs in our portfolio, I ask them, what are you doing right now with AI? Give me specific examples on where you have either generated more revenue or you have saved cost. How did you do that? How do you talk to your portfolio management teams about this, and how do you roll that out across your portfolio companies?” says JPMAM’s Lam.

“In every single conversation I have, the effective use cases are distinct and [improve] over time. It keeps me up at night, both from an opportunity and threat perspective, because if you’re slow to this, you risk getting leap-frogged by the competition.”

Realisations – limited though they may be in private equity’s present exit drought – will provide some clues as to which GPs have executed effectively on the AI opportunity.


“LPs should look for evidence in the realised data... to show either revenue protection in the sense that it’s not being disrupted, or indeed, revenue growth if AI has been used to enhance businesses,” says McCarthy. “[That] will tell a lot about who the winners and losers may be over the next number of quarters and years ahead, because I think the margin expansion, the gross margins, will bifurcate pretty quickly from here.”

Sector specialists, inevitably, seem likeliest to benefit from this dynamic. “What we’re finding is maybe some of the more generalist managers are likely

to be the ones that may be left behind by the specialists in this sort of environment,” McCarthy notes.

But how does one identify a true specialist? Mercer, for its part, suggests looking at a GP’s operational capacity, network, proprietary access to dealflow and evidence of AI implementation playbooks.

“That’s the space where you could actually get some confidence in the GP’s understanding of what they do, the effect of change and why they’re going to be able to manage change in a value creation sense – not just in a general narrative around either efficiency



acquisition or sort of margin expansion,” adds McCarthy.

Software by AI other name

SaaS M&A reached new heights last year, with 2,698 completed deals, according to Software Equity Group. Private equity and venture capital buyers participated in nearly 58 percent of all transactions. Notably, however, this activity was driven predominantly by add-ons, with new platform investments falling to their lowest level – 7.7 percent – since 2021.

What’s more, fundraising for tech, media and telecommunications-focused private equity funds has declined in recent years. Between 2023-25, TMT-focused funds raised on average \$112.3 billion each year, representing 15.4 percent of global private equity fundraising per annum, according to *PEI* data. By comparison, these strategies raised on average \$179.4 billion per annum in the preceding three years, equating to a yearly average of 24.3 percent.

Nevertheless, the potential upside from AI integration could prompt the most bullish to lean further into software.

“You need to be risk-on, because the future of the world lies in tech

spending and AI; the worst reaction is to be totally risk-off,” says CDC’s Armbruster. “If you want to back GPs dealing with [these kinds] of opportunities over the coming years, you need to do it now.”

CDC’s risk appetite has limits. “It’s moving so fast and the impacts are so big that it’s wise to be risk-off in the very short term on potential co-investments and potential CVs, because here you are timing the market much more,” Armbruster concedes. “We just passed on a software co-investment alongside a great GP because it’s hard for us to assess the potential disruption of the business model, and it’s super hard to have a clear view on valuations.”

Though some specialist funds will inevitably thrive in these conditions, there is no guarantee that a software focus will directly translate into fundraising success.

“There will be a considerable flow to software, but it may not be software-only focused funds,” notes McCarthy. “I think it’s going to be part of a digital AI-focused strategy where software and scale and the cost-effective elements of delivery that come with that become an AI-enabled theme as much as a software-only theme.”

As such, LPs may prioritise sector experts that possess the ability to look beyond software into the tech sphere more broadly.

“You cannot deny that software specialists are also tech specialists, and [are] some of the most insightful, well-placed and equipped investors to benefit from AI opportunities going forward,” says Armbruster. “You still need to be close to the specialists, but make sure they are flexible enough and opportunistic enough not to deal only with their super niche software strategy that they had until now. They need to adapt.”

This transition is already changing the way some GPs position themselves. As recently as 29 January, Hg’s website homepage retained the firm’s longstanding pledge to “build enduring software and services leaders”, according to Web Archive. By 9 February, this had been broadened to “enduring technology and services leaders”. Then, by 15 March, it had evolved into something new altogether: “Leading private equity in AI transformation.”

The transition begs an important question: does Hg now consider itself a software firm or an AI firm?

“Both,” hedges Brian. “If you are – like we are – a software and tech-enabled services investor, you have to be AI-first.

“This is how we’ve adapted. We’ve run a big project over actually nine months... to adapt our investment strategy to the new world. In everything we do, it is the lens through which we look.”

Weaving’s Agent Smith could replicate himself by overriding the consciousness of any human inhabiting the synthetic Matrix universe. Just as these individuals all became indistinguishable in his trademark black suit and tie, so too may software and other tech strategies begin to converge under one dominant and inescapable theme. ■

