EXECUTIVE INSIGHTS

Revolution or Extinction? Rethinking SaaS in the Age of Agentic AI

Key takeaways

1. Agentic and generative AI are redefining software workflows, enabling autonomous, outcome-driven operations that reduce reliance on traditional user interfaces and manual inputs.
2. Traditional SaaS models face disruption on two fronts: agentic AI may bypass interfaces altogether, while generative AI empowers users to build customised software that mimics and potentially replaces SaaS platforms.
3. Three future scenarios are emerging. SaaS and AI may remain distinct, converge through integration, or see agentic AI eclipse SaaS entirely, with implications varying by industry and use case.
4. SaaS vendors must move faster — rethinking architecture and pricing for an AI-native world. For users with significant SaaS spend, agentic AI offers a chance to reinvent how work gets done.

Two of the most intriguing developments in the world of AI are the emergence of "agentic AI," systems capable of autonomously planning, reasoning and executing multi-step tasks, and "generative AI," capable of developing code and software. For SaaS companies, these trends present both a significant opportunity and a looming threat.

As AI shifts from being a supportive feature to becoming an autonomous actor in its own right, the traditional SaaS business model faces pressure to adapt or risks obsolescence.

In this *Executive Insights*, we examine how the rise of agentic and generative AI is reshaping software workflows and redefining what SaaS platforms must become to stay relevant.

Understanding the difference: Agentic workflows vs SaaS

SaaS platforms, by design, serve as systems of record. They're structured environments built to store, manage and present business-critical data through intuitive user interfaces. Examples include Salesforce, Zoom, Slack and Outlook. These platforms require users to input data manually and interact with software directly.

Agentic workflows, by contrast, introduce a more fluid and intelligent approach to interacting with systems and completing tasks. Rather than relying on human-driven inputs to operate static interfaces, agentic workflows enable AI agents to interpret high-level goals and autonomously determine the steps required to achieve them. These agents can access and orchestrate actions across tools, invoke APIs, query data sources and adapt their behaviour in real-time, reducing or even eliminating the need for direct user engagement.

The differences are foundational. SaaS platforms are inherently static and reactive, built around defined schemas and tightly bound logic. They're powerful, but inflexible. Agentic workflows are dynamic and proactive, capable of adapting to changing conditions and contextual signals. Where SaaS requires human action at every step, agentic workflows shift the burden of execution from user to machine.

While it may seem natural to combine the structured, data-rich environments of SaaS with the intelligent automation of agentic workflows, the two originate from fundamentally different design philosophies. SaaS platforms are built around structured interfaces, manual control and the integrity of stored data. Agentic workflows focus on interpreting and acting on data flexibly and autonomously. Integrating them into SaaS demands rethinking how data flows, decisions are made and tasks are carried out, shifting from static systems to more adaptive, intelligent operations.

The dual threat to traditional SaaS

Adding to this complexity are two emerging types of AI-driven disruption. First, agentic workflows have the potential to replace the user interface (UI) and user experience (UX) layers of SaaS entirely. In this scenario, instead of interacting with dashboards and forms, users simply prompt an AI agent that retrieves data from a consolidated database and executes actions autonomously. Here, SaaS persists only as a backend service, while the workflow is owned and driven by agents.

Second, even if users retain a preference for the SaaS interface model, generative AI is now enabling them to build those very interfaces themselves. People and companies can increasingly use AI to create bespoke SaaS-like applications on the fly. These products look and feel like conventional SaaS platforms, but they're customised, created without coding

expertise, and crucially, they bypass the need for traditional SaaS providers altogether. This disruption challenges the entire SaaS economic model, as software is generated rather than purchased.

These disruptions are already unfolding. OpenAI's Operator, for example, is a compelling example of agentic workflows in action. It allows users to plan and book a jazz concert with a single prompt, handling search, scheduling and payment autonomously.

On the generative side, tools like Replit empower users to build functional software, such as a customised chess game for a child, without needing to write complex code. It's not hard to imagine this extending to full-fledged CRMs or ERPs in the near future.

What does the future hold? Exploring three scenarios

Whilst agentic AI presents a clear threat to SaaS, it also offers significant potential to enhance and augment existing platforms. Given the contrasting foundations of SaaS and agentic AI, there are three potential scenarios for how their relationship might evolve:

- They may continue to operate distinctly
- They may converge, with agentic workflows embedded into, or layered on top of, SaaS platforms
- Agentic AI may eventually eclipse SaaS altogether

Below, we explore how each scenario might play out (see Figure 1).

Figure 1

Three potential futures for SaaS and agentic AI

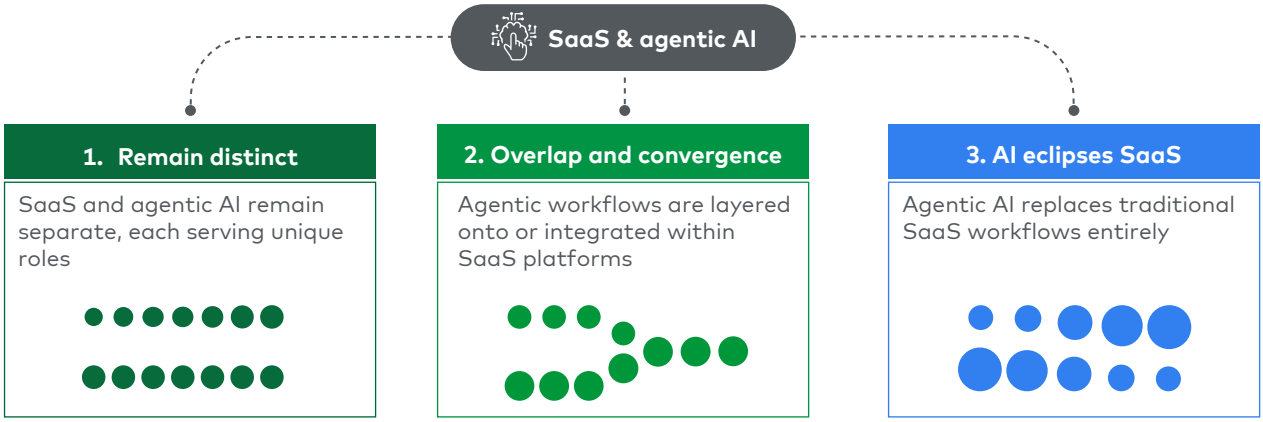

Source: L.E.K. research and analysis

1. Remain distinct

There are strong reasons to believe that SaaS and agentic workflows will remain separate in many contexts. For one, issues of user trust and control remain paramount. Many users, particularly in regulated industries, prefer manual oversight, especially where compliance and precision are critical. Allowing AI agents to make autonomous decisions introduces risks that many organisations cannot accept.

Secondly, many verticalised industries, such as healthcare or manufacturing, rely on highly specialised SaaS tools that are deeply embedded into existing workflows. These tools offer niche functionality that is difficult for general-purpose AI to replicate.

Moreover, SaaS platforms benefit from mature integration ecosystems. Structured API-based connections, such as those linking CRMs to ERPs, are robust, standardised and hard to replicate in dynamic, agent-led architectures.

Finally, privacy considerations are a major concern. SaaS companies often design their products in compliance with strict local policies. Agentic AI systems, being inherently decentralised and adaptive, face challenges in meeting these requirements without significant additional oversight.

2. Overlap and convergence

Despite these distinctions, we're already witnessing increasing overlap between SaaS and agentic workflows. In this scenario, the two paradigms complement one another, creating hybrid systems that enhance productivity and UX. For example, embedded AI agents are already transforming SaaS products:

- Salesforce's Einstein AI acts as a copilot in CRM systems, automating data entry and suggesting next steps.
- Adobe's marketing agents perform A/B testing and generate content, while still allowing human designers to oversee branding decisions.

These are both examples of the first type of disruption, where agentic AI replaces traditional UI/UX interactions.

Agentic systems are also being used to orchestrate tasks across multiple SaaS tools. AI agents can now schedule meetings in Google Calendar, update tasks in Asana and send messages through Slack, all without human input.

Additionally, SaaS interfaces themselves are evolving. Traditional dashboards and point-and-click interactions are giving way to conversational interfaces, with tools like Microsoft 365 Copilot enabling users to generate complex documents or analytics using simple natural language commands.

This convergence is already taking hold. According to the 2024 SaaS Benchmarks Report,[1] 56% of companies reported that they've launched or tested AI features in their products within the past year. Of this group, 41% are monetising AI features, a number that's up 9% from 2023. These overlaps suggest that SaaS and agentic AI are converging, with AI increasingly being layered on top of structured SaaS foundations to deliver more autonomous, user-friendly workflows.

However, the convergence between AI and SaaS may also prove disruptive to existing SaaS platforms. The rise of generative AI and low-code/no-code development means that companies can now build highly tailored solutions without relying on third-party SaaS vendors. In effect, companies and individuals are bypassing SaaS providers while still creating software that behaves like SaaS, just customised, more flexible and potentially cheaper.

This trend is already visible in the increasing number of custom CRM builds and individual no-code projects that mimic traditional applications. Companies will end up with a product which looks and feels like the SaaS platforms they're already used to using and trust but without the expensive price tag associated with traditional SaaS companies.

3. AI eclipses SaaS

As adoption accelerates, a third scenario emerges: the potential for agentic AI to eclipse SaaS in certain domains. This may occur in response to growing consumer demands for faster, simpler and more customised solutions, something agentic AI is well positioned to deliver.

In an environment where the first type of disruption, agentic AI interfaces, is fully established, businesses may begin rationalising their back-end systems into unified data repositories. AI agents, being agnostic to the application layer, can access and act on this data directly, removing the need for distinct SaaS applications entirely.

Alongside this, the second disruptive force continues to grow. As AI becomes capable of generating functional applications on the fly, businesses may stop subscribing to SaaS altogether. Instead, they'll configure or request software tailored to their exact needs, delivered instantly by generative tools. Rather than buying licenses, they'll pay for outcomes.

This dual pressure on the interface and infrastructure risks relegating SaaS to a "middle layer," increasingly bypassed in favour of more agile, outcome-focused workflows. In marketing, for instance, agentic AI is already automating campaign management, content generation and customer insights, diminishing the role of traditional SaaS platforms.

Regulatory considerations

Regardless of which scenario ultimately prevails, agentic AI systems, whether standalone or embedded, will need to comply with local regulations, data protection laws and governance frameworks. This is especially critical as AI agents take on more decision-making authority.

Companies must ensure that agentic systems remain auditable, secure and aligned with ethical standards.

Building trust in these systems will be essential. Organisations should prioritise transparency, user consent and robust data controls as they adopt AI-driven workflows.

What should companies do?

For SaaS businesses, the imperative is clear: adapt or risk becoming obsolete. Companies should begin by embedding AI capabilities into their platforms as core components that support autonomous workflows. They must also consider moving towards agent-native architectures. This means rethinking how data flows through their systems, how users interact with platforms and how platforms themselves can take action on behalf of users.

Furthermore, pricing and product models may need to evolve. As customers shift from paying for access (e.g. per-seat licensing) to paying for outcomes, SaaS providers must adjust accordingly.

Strategic responses will vary depending on which disruption a company believes to be more imminent to their SaaS platform. If the future lies in agentic interfaces replacing SaaS UIs, then companies must invest in building agentic workflows today, and becoming the orchestrators of these AI-powered systems.

If the greater threat is generative SaaS built on the fly, then the goal should be to provide customers with flexible, build-your-own SaaS frameworks, effectively selling the tool that builds the tool. For example, rather than selling a static CRM platform, a company might offer an AI-powered builder that creates CRMs tailored to a client's needs.

Investors, too, should evaluate whether the businesses they support have a clear AI roadmap and a defensible value proposition in an AI-first market. Customers will increasingly prioritise tools that deliver real-world results, not just dashboards and features.

Conclusion

The future of SaaS in the age of agentic AI isn't black and white. In reality, all three scenarios — separation, convergence and eclipsing — will play out in parallel across different industries and company sizes.

In regulated or highly specialised sectors, traditional SaaS will likely remain dominant. But in others, we're already seeing the rise of "service-as-software" models, where outcomes matter more than interfaces and where AI agents orchestrate work with minimal human involvement.

Agentic AI doesn't necessarily spell the end of SaaS, but it certainly marks the end of SaaS as we know it. The companies that embrace this transformation and reimagine their role in an AI-driven world will shape the next generation of enterprise technology. Those who fail to evolve risk being outpaced by an entirely new class of competitors.

Contact us to see how we can help.

Endnote

[1]High Alpha, "2024 SaaS Benchmarks Report." https://www.highalpha.com/2024-saas-benchmarks-report

About the Authors

Rob Wild | Partner | R.Wild@lek.com

Rob Wild is a Partner in L.E.K. Consulting's London office and a member of the firm's Digital practice. With nearly two decades of experience in digital strategy and transformation, Rob specialises in advising businesses on strategic digital and AI initiatives to drive growth, enhance operations, and build competitive advantage.

Harsha Madannavar | Partner | H.Madannavar@lek.com

Harsha Madannavar is a Partner in L.E.K. Consulting's Technology Infrastructure practice globally (software, data, services and silicon systems), advising US and global corporates, private equity and hedge funds on a range of shareholder value issues including growth strategy, business model transformation, technology disruptions, product development, pricing and M&A. He is a member of the firm's global Board of Directors and also serves on the Board's Audit and Risk committees.

Romain Maitret | Partner | R.Maitret@lek.com

Romain Maitret is a Partner in L.E.K. Consulting's London office. With more than 10 years of consulting experience, he specialises in the technology sector, providing strategic advice to clients in fields including B2B software, infrastructure technology and IT services.

About L.E.K. Consulting

We're L.E.K. Consulting, a global strategy consultancy working with business leaders to seize competitive advantage and amplify growth. Our insights are catalysts that reshape the trajectory of our clients' businesses, uncovering opportunities and empowering them to master their moments of truth. Since 1983, our worldwide practice — spanning the Americas, Asia-Pacific and Europe — has guided leaders across all industries, from global corporations to emerging entrepreneurial businesses and private equity investors. Looking for more? Visit **lek.com**.