
AI's Impact on Software as a Service

March 2026

Charles Phillips

Co-Founder and Managing Partner

Recognize

Contents

Summary	3
Introduction	3
The Pyramid of SaaS Applications	3
Evolution of M-3 Gardens	7
The Value of Agents for SaaS.....	8
The New Roles for Software Developers	9
CEO Priorities	10
CIO Priorities	11
Impacts on IT Services.....	11
Author's Note.....	12

Summary

Investors have real concerns about AI's impact on the future of Software as a Service (SaaS), but the risk is highly asymmetrical across companies.

At the bottom of the pyramid, simple applications serving small businesses are vulnerable to AI displacement, a view corroborated by our recent survey of IT professionals. At the top sit large, mission-critical applications that must be accurate, must scale, and must not fail — the “M-3” platforms. These applications are more insulated from AI by their size, complexity, critical data, and high switching costs.

M-3s face a mixed impact. AI threatens some functions, but creates opportunities to enhance and extend complex processes, adding value without ceding ground.

Meanwhile, senior executives, weighing limited time and resources and more strategic priorities, are unlikely to rush to replace routine SaaS vendors with in-house solutions just because it appears doable. The build vs. buy economic analysis still applies.

AI will lower the barrier to entry for some applications and could create more choice and competition, changing the build analysis. This could lower SaaS values in some, but not all, segments.

Introduction

Investors re-rated the SaaS sector down by \$300 billion in recent months. Although the sector saw frothy multiples for years, the advent of AI raises legitimate questions about the addressable market going forward.

Our recent survey of CTOs and mid-level IT practitioners captures the current mood. Of 265 respondents, more than 80% said they are likely to replace some commercial software applications with AI-generated alternatives. And 63% of respondents expect their IT organizations to be larger in two years with more budget. Of course, these respondents are career technologists who like to build more.

Meanwhile, our conversations with senior executives, CIOs and CEOs, suggested a more measured strategy for prioritizing for impact.

The Pyramid of SaaS Applications

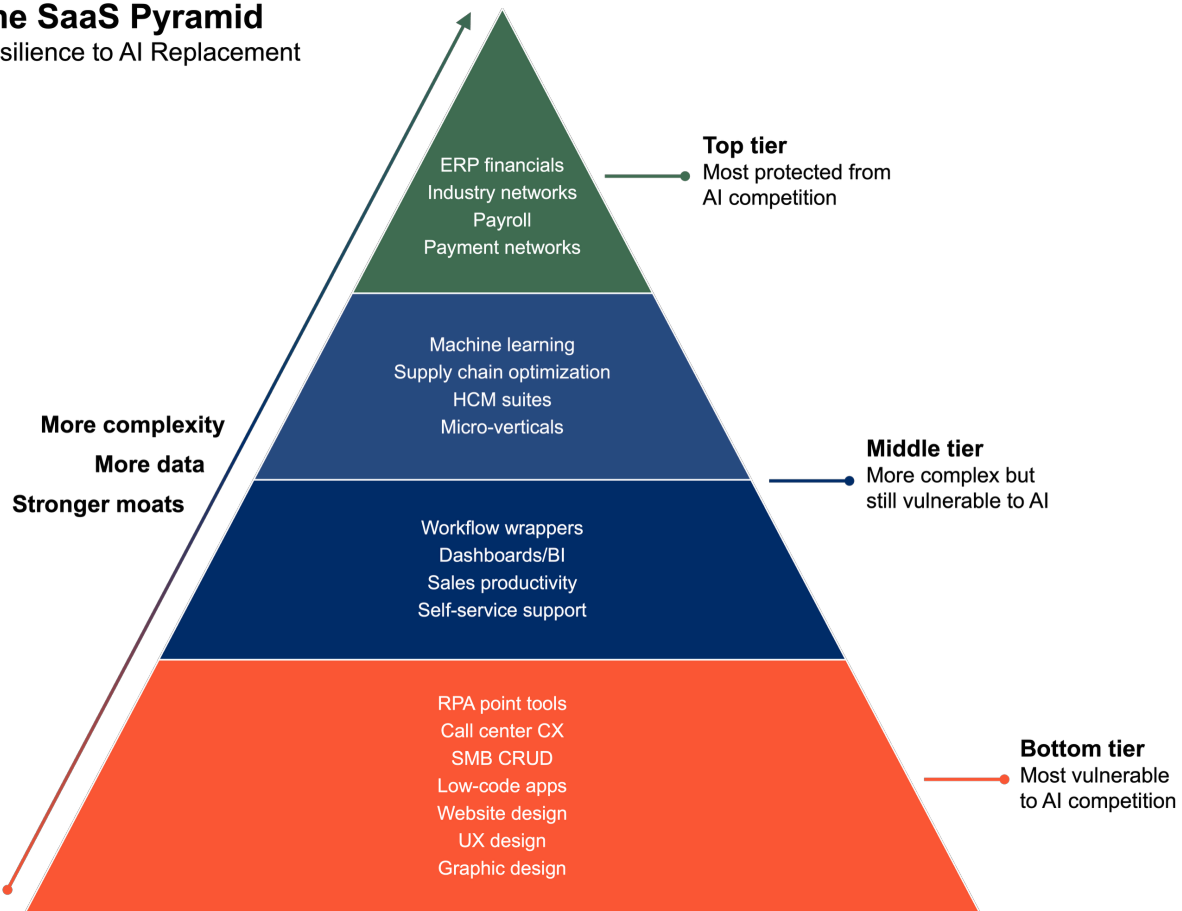
Not all SaaS platforms were created equally. The continuum is as wide as the gap between high school basketball and the NBA Finals, and investors will have to delineate which companies are playing small ball and which make the dream team.

One way to frame the risk is with a pyramid of SaaS applications. At the bottom of the pyramid, software is more exposed to the risk of competition and displacement from AI. At the top, software is more insulated from these risks.

But at each tier, applications share attributes that are useful in discerning their exposure to AI risk.

The SaaS Pyramid

Resilience to AI Replacement



The Bottom of the Pyramid

At the bottom of the pyramid — and exposed to more risk — are the many disparate applications used by small businesses or departments. If an AI model can generate 85% of the code with a prompt, that app is on shaky ground.

Attributes to look for:

- **Small code base** — One relevant (but not always dispositive) metric is the size of the code base, and apps with small code base may face more risk. The average web site has 500 to 1,000 lines of code; SAP, by contrast, has 400 million.
- **Simplicity** — Basic CRUD (create, read, update, delete) apps — for example, a project management app — are easier for AI to replicate. And if the bulk of the app is what the user

sees on the screen, with little algorithmic complexity operating behind the scenes, it may not be complicated to reproduce with AI.

- **Portal apps** — Apps that serve as front-ends for external services (e.g., payment rails or inventory/pricing databases) face more risk. Increasingly, AI agents can call those external services directly, reducing the need for intermediary interfaces.
- **Design tools** — AI excels at image generation, which will make design faster and more accessible. But it may also reduce the need for complex design tools; for example, website design will trend toward self-service.
- **Micro-vertical focus** — Legacy software for niche small and mid-size business sectors — like fitness and leisure, club management, moving and storage, dealerships, and justice systems — is in dire need of modernization, but large companies and top engineers don't typically work on apps in these markets. Now a new wave of startups can build much better software, and more domain experts can become software developers.
- **SMB-dependent customer base** — Apps serving smaller customers are more vulnerable, because these customers are more cost-sensitive (e.g., owners may personally fund IT) and churn easily. These customers can also tolerate more risk, and AI is lowering their switching costs by cleaning and formatting data and documenting processes.
- **Dashboard functionality** — Apps that create KPIs, charts, and tables from other apps can be easily duplicated by AI agents.

The Middle of the Pyramid

Many applications in the middle of the pyramid do a bit more than CRUD, but perhaps not enough to avoid new competition. Much depends on whether an AI native is aware of and interested in a particular micro-vertical sector and humble enough to learn the requirements. The TAM could shrink in this sector. Regardless, these middle-tier applications shouldn't be confused with the big players at the top of the pyramid — the more important tier on which to focus.

The Top of the Pyramid

At the top of the pyramid — and more buttressed from AI risk — are large, critical, highly integrated applications. Their complexity and reach mean they must be accurate, must scale, and must not fail — the “M-3” applications. M-3s derive about 75% of their revenue from large enterprise customers that have global distribution and premium service.

Attributes to look for:

- **Large scale** — Supporting an application for 50 users is completely different than an application serving millions of concurrent users. Delivering performant, secure software with low latency, data consistency across petabytes of data, and “five nines” (99.999% uptime) requires sustained engineering innovation across layers in the technology stack.

- **Critical data** — Companies are hesitant to change or replace software that generates and maintains important content, such as clinical trials, SEC financial reporting, general ledger transactions, inventory, shipments, medical records, or payments. Disturbing apps in these categories carries considerable risk with low returns. Software that directly makes or tracks money are in better shape than supporting-role software, like CRM or marketing. And migrating data is messy, since no two apps classify, store, and describe data the same way.
- **Critical processes** — Applications that keep the trains running at large companies — financial accounting, order to cash, payroll, demand deposits, and trading desks — may be too important to take unnecessary risks by moving to an unknown platform.
- **Many integrations** — M-3 applications have hundreds of thousands of third-party applications connected to them to extract data or extend processes, typically via APIs. Thousands of interconnected applications can resemble the maze of plumbing under the streets of Manhattan. Changing the central hub can therefore cause a cascade of other breakages.
- **Large code base** — SAP's nearly 400 million lines of code has been carefully enhanced with industry and regulatory requirements over 50 years.
- **Credibility** — Enterprise customers with worldwide operations need vendors who can support them globally. M-3s offer the global distribution and premium service that large enterprises require — the opposite of an app store.
- **Engineering complexity** — Some M-3s build software for complex embedded systems, like those in trucks, airplanes, drones, and mining or farm equipment. These customers are not going to jump to replace a vital, proven application with new AI.
- **Accountability** — Someone must be accountable when things go off the rails. M-3s accept the risks of, say, failed payrolls, corrupted data, and erroneous financials, and they engineer the necessary processes for quality control, rollback, and recovery.
- **Supervision** — Some processes are too critical to not have humans checking agents and agents checking humans. An AI agent could monitor your personal income sources and deductions, calculate your taxes, and automatically extract funds from your bank account to pay the IRS, but most people still want a human tax professional in the loop. Companies are the same; trust is currency.
- **Aggregators** — Applications that own and aggregate unique, industry-wide content are difficult to replicate (e.g., insurers sharing claims data to improve underwriting). Other aggregators tediously collect, interpret, and automate legislative mandates, such as tax rates in thousands of jurisdictions or tariffs, customs, and duty rates in hundreds of countries.
- **Networks** — Commerce networks allow buyers and suppliers to exchange orders and other data. Networks produce a special inertia, since no one wants to be on a new network unless everyone else is on it.
- **Breadth and frequency** — M-3s are in a stronger position if most employees rely on the application often. Broad usage increases the cost of change management.

The AI Replacement Matrix

Risk of replacement in 1-3 years, 4-6 years, and 7+ years

	1-3 years	4-6 years	7+ years
Low risk	Low-code development	HCM suites Call center CX	ERP financials Payroll Payment systems
Moderate risk	SMB CRUD apps Self-service support tools Website design	CRM suites Supply chain planning Budgeting and planning	Data lakehouse Industry networks Supply chain execution
High risk	Dashboards / BI Graphic design UX design	RPA point tools Workflow wrappers Vertical micro-SaaS	Machine learning

Evolution of M-3 Gardens

The attributes that make M-3 applications difficult to displace also create opportunity. M-3 customers are heavily invested in these platforms, with high exit costs. That investment creates a walled garden — an ecosystem customers want or need to stay in.

Strong incumbents sometimes evolve and adapt. When the internet emerged, retailers were pronounced irrelevant, but companies like Walmart and Costco added e-commerce services that leveraged their existing store footprints for fulfillment and returns. Switching costs in retail were low, because customers could easily shop elsewhere, and yet most stayed with their trusted brands as those brands evolved.

M-3 applications are positioned for a similar shift. Many are rapidly building AI agents within their walled gardens — agents with proprietary access to the data and logic that underlie these platforms. Enterprise customers will likely see advantages and lower risk by rolling with their incumbent M-3 partners on their AI journey, rather than starting over.

But M-3 vendors must also decide how to leverage this positioning.

Options that M-3s will experiment with:

- **Toll booth** — M-3s can control access to their garden. Some will charge for access; others may close the gate. Salesforce, for example, updated Slack's API terms to block third-party AI platforms — a decision that may change if customers demand more third-party agent access.
- **Consumption-based pricing** — Agents can use SaaS apps around the clock and tax SaaS infrastructure to a degree that would be impossible for a human user. Per-user or per-agent

pricing is no longer optimal, since fewer users on the system no longer means less work completed. More companies will experiment with consumption pricing (Snowflake, Twilio), outcomes (Palantir), and other workload-based pricing models.

- **Data sovereignty** — In an AI-first world, the value mix shifts modestly from the code toward the proprietary data trapped within the garden. External agents can't replicate data they can't reach.
- **Control tower** — Established protocols allow agents to collaborate: Agent2Agent (agent communication), Model Context Protocol (standards to access data and resources), and NANDA (global agent discovery protocol). But given the coming wave of billions of agents, enterprises still need a control tower to manage permissions, compliance, governance, policies, security, and monitoring. M-3s are attempting to position themselves as the control tower, but it's uncertain if customers will prefer a more open or neutral platform.
- **Inference platforms** — Inference platforms deploy and scale LLMs, adding value through cost optimization, lower latency, and reliable performance. These platforms sit between the prompt and the model. Bringing the model to the data instead of the data to the model is a current best practice for security, control, and performance. M-3s that embed inference directly inside the garden can potentially offer customers a frictionless, integrated experience.

The Value of Agents for SaaS

AI agents are likely to augment SaaS applications much faster than competitors can replace these applications. SaaS companies, then, can bolster themselves by layering on agentic capabilities.

Examples of where agents can add value:

- **Repetitive tasks** — Robotic Process Automation (RPA) tools record data entry screens for repetitive tasks and then automate that entry with brittle rules. More adaptable agents could replace RPA and automatically evolve as applications change.
- **Friction** — Agents can remove friction points that slow down applications. For example, AI could clean and format data for supply chain planning apps, resolving inconsistent units of measure — a major headache. Or AI could add real-time external data to supplement historical time-series data.
- **Gaps** — Agents can supplement SaaS apps by removing manual steps connecting long-running processes, like anomaly detection and transaction-matching for general ledger reconciliation.
- **Sensing and responding** — Agents can sense changes and then invoke process in enterprise applications. An agent could monitor inventory levels and demand (sales and operations planning) to kick off replenishment orders via applications already wired to suppliers.

- **Self-service** — It's no secret that enterprise applications sport painful user interfaces. Agents can provide a natural language interface for employees to select benefits, request time off, or ask for records.

The New Roles for Software Developers

Software development crossed the Rubicon last year with the release of Opus 4.5 and GPT 5.2. Developers now realize LLMs can generate 90% or more of the code for many applications. Guiding a project no longer means writing every, or even any, line of code.

How software development is changing:

- **Assembly** — Developers have long assembled pre-written code from Stack Overflow and GitHub and leveraged open-source libraries. Open-source companies such as Red Hat built large enterprises packaging, supporting, testing, and marketing software that they didn't build.
- **Automation** — Infrastructure-as-code, continuous integration, automated testing, low code, containers, and other innovations have progressively automated the stages of the development life cycle. Coding assistants take this journey to its logical next step.
- **Non-coding time** — Developers historically spent only 15–20% of their day writing code. The balance was spent on code reviews, debugging meetings, planning, testing, security assessments, and gathering requirements.
- **System engineering** — The software development role is evolving from coding to system engineering — verifying code, orchestrating components into a system, and then designing for performance, security, lower cost, quality, and testability. Engineers must also master prompts, agents, inference, context, permissions, and observability.
- **More code** — More code means more bugs, tech debt, performance issues, security breaches, and optimization of overly verbose code. Once code is in production, change requests surface from users, statutory requirements change, and cloud platforms evolve. Someone must manage this constant change.


What this means for SaaS overall:

- **Elephant in the room** — This increased coding capacity might derive a higher ROI by refactoring custom, expensive, and brittle legacy apps (e.g., 250 billion lines of COBOL, according to IBM) versus replicating working payroll. AI is getting better at analyzing old code and generating the underlying process maps.
- **Hiring continues** — Software engineering job openings in tech companies peaked in 2021 at 115,000, driven by pandemic hiring. Job posts then fell sharply and bottomed out at 40,000 in 2024. According to tech jobs site TrueUp, openings are now trending back up at 65,000, despite the explosion of coding assistants. More of those jobs mention AI skills and

data engineering. Engineering job openings at tech companies are up 72% from the low, which suggests we need more engineers as we generate more code.

Estimated Job Openings in Engineering

Early 2023 to Early 2026



Data Source: TrueUp

CEO Priorities

CEOs whose companies rely on SaaS have different priorities than IT executives. They are less likely to ask a question like, “Can we replace our payroll system?” and more likely to ask, “How can we make better decisions, lower costs, be more productive, or service customers better using AI?”

CEOs make tough decisions with material consequences and limited information. They are contending with the chaotic, unpredictable nature of business in the digital economy. The current term in vogue is VUCA — Volatility, Uncertainty, Complexity, and Ambiguity — which applies to companies doing business in multiple countries against a backdrop of changing tariffs, volatile component pricing, and faster competitive changes.

CEOs derive some context from systems of record, but that data often reflects historical performance. The more current information they need synthesized in real time is the unstructured data around the world — news, customer sentiment, social media, regulatory proposals, interest rates, investor feedback, the bond market, banker advice, emails, white papers, and other sources — as well as correlations across systems.

AI is tailor-made for the challenge of deriving insight from disparate sources like these. And in the face of finite resources, CEOs may be less interested in using AI to replace systems of records and more focused on critical competitive context in the global economy.

Specialization

Companies have long outsourced functions that are necessary but more efficiently delivered by outside specialists.

OpenAI, Google, and Anthropic each has an incentive to prove it could generate a human resources application, but all use Workday to automate HR. Even if they could replicate the application, the long tail of maintenance and support would be a distraction. Reverting to the age of vertical integration seems unlikely to happen quickly.

Any company could task their general counsel with litigation responsibility, but normally they rely on external counsel who litigate full-time. Companies could also run their own logistics, warehousing, and delivery operations, but they choose to hire third-party logistics vendors who can spread delivery costs across multiple customers.

Hospitals outsource medical coding and billing, banks outsource card issuing and mortgage processing, product companies outsource to contract manufacturers, and many companies outsource their call centers to contact centers.

CIO Priorities

The Recognize CIO Council is comprised of Fortune 200 CIOs. CIOs we've spoken with don't seem eager to build and maintain custom code for standard or commodity applications. They spent years moving away from that. They would rather reserve their limited resources for applications that differentiate and accelerate the core mission. Generating millions of lines of code for mundane processes that will eventually mature into layers of tech debt may not be a high priority for overtaxed IT departments. They have a mandate to exploit AI, but they want to deliver transformative projects that weren't possible before AI.

Impacts on IT Services

If companies aren't building in-house, they're buying from SaaS vendors, as well from the digital services companies that implement, integrate, and maintain those platforms — the sector on which Recognize focuses.

This industry — comprised of systems integration, IT managed services, software development and maintenance contractors — is now valued at more than \$1 trillion, according to the International Data Corporation.

In the CTO/IT survey mentioned, 59% respondents anticipate needing more third-party contractors to help build and manage AI applications. The complexity and rate of change of AI appears to create more of a need for external specialists in the short term.

AI impacts the IT services companies differently from the SaaS vendors. Delivering quality, scalable, and performant code — the historical skill set for software companies — will get easier. Systems integrators (SIs) may be able to build more differentiated intellectual property to shorten time to value.

SIs may also be able to deliver projects more predictably if workflows can be quickly and accurately understood. Likewise, applications may become easier to integrate via agents communicating with agents.

The business model will need to adapt, however. SIs mostly charge by time and materials — billable headcount. They will have to learn how to charge for digital workers or switch to outcomes-based pricing, as Palantir has done. But tech services are a topic for the next white paper.

Author's Note

The firm I co-founded, Recognize, invests in digital and AI services. This paper captures ideas on how SaaS, in its many forms, may fare in the AI era. It reflects conversations with enterprise technology leaders, insights from our portfolio companies, and perspectives from Recognize's CIO Council, whose members kindly shared their views in a recent survey.

I've watched software evolve over the last few decades and made calculated bets from several vantage points — as an engineer, IT practitioner, Wall Street analyst, software industry CEO, and now as an investor. Per Yogi Berra, predictions are tough, especially about the future. These ideas will need updating as the landscape shifts, but this is what I'm seeing now.

— Charles Phillips, Co-Founder and Managing Partner, Recognize
charles@recognize.com

Disclosures

Charles Phillips is a Co-Founder and Managing Partner of Recognize Partners LP (“Recognize” or the “Firm”), a private equity firm investing exclusively in digital services companies. This opinion piece (the “Information”) reflects his own views as of the date hereof and does not necessarily reflect the views of any entity he represents; neither Mr. Phillips nor Recognize undertake to advise you of any future changes in the views expressed herein. The Information is being furnished to the recipient (“you”) solely for informational purposes. This Information is not, and may not be relied on, in any manner, as legal, tax, investment, accounting or other advice. This Information is not, and should not be construed as, an offer of investment advisory services by Recognize, nor as an offer to sell (or a solicitation of an offer to buy) an interest in any investment sponsored by or affiliated with Recognize. The Information is not an endorsement of any particular security or investment opportunity.

Certain information contained in this Information has been obtained from published and non-published sources prepared by other parties, which in certain cases has not been updated through the date hereof. While such information is believed to be reliable for the purposes of this Information, neither Mr. Phillips, nor Recognize, nor their affiliates assume any responsibility for the accuracy or completeness of such information, and such information has not been independently verified. Statements contained in this Information (including those relating to current and future market conditions and trends in respect thereof) that are not historical facts are based on current expectations, estimates, projections, opinions and/or beliefs. Such statements involve known and unknown risks, uncertainties and other factors, and undue reliance should not be placed thereon. In addition, no representation or warranty is made with respect to the reasonableness of any estimates, forecasts, illustrations, prospects or returns, which should be regarded as illustrative only, or that any profits will be realized. You should make your own investigations and evaluations of the contents of the Information and should note that such information may change materially.