

Technology, Media & Telecommunications Practice

The AI-centric imperative: Navigating the next software frontier

With AI-native upstarts redefining speed and scale, software incumbents hoping to remain competitive must fully embrace the new technology to reinvent their value proposition and internal operations.

This article is a collaborative effort by Jeremy Schneider, Joshan Abraham, Matt Linderman, and Naveen Sastry, with Mónica Perestrelo and Shayan Salam, representing views from McKinsey's Technology, Media & Telecommunications Practice.



The software industry is entering a new era—and it may yet prove even more disruptive than the software-as-a-service (SaaS) revolution that preceded it. The emergence of gen AI and, more recently, agentic AI is not just another technology wave; it is a foundational shift redefining what software is, who builds it, who uses it, and how companies are organized and operate.

Gen AI alone is projected to unlock \$4.4 trillion or more in annual value across the global economy, with software companies poised to capture 10 to 15 percent of that total—and agentic AI may well accelerate the speed at which this value is realized. But capturing it is far from guaranteed, and incumbent companies will face heightened competitive intensity and complex new challenges. Some may not survive. Recent moves by AI players such as OpenAI underscore this urgency. By embedding their own AI-powered sales, support, and contract tools directly into workplace processes, these companies could end up competing with the very SaaS players they have been enabling—a shift that may further upend industry dynamics and intensify the pressure on those incumbents.

Gen AI and agentic AI may start by creating new ways for users to engage with software, but the technology should ultimately have a much greater impact. That may lie in enabling customers to tailor software that can act autonomously, make decisions, and interact with different users, software, and systems across workflows traditionally unsupported by legacy tools. These advances could usher in an era marked by equally large shifts in the competitive landscape, including an acceleration of [vendor switching and customer churn](#), a realignment of user segments and value pools, and an increase in corporate in-house or “citizen” development of software. New risks will need to be addressed, from algorithmic bias and data quality to explainability, IP infringement, and newfangled security threats.

As software companies navigate this AI-centric transformation, they face two fundamental paradigms: rethinking their value proposition holistically, from products to business models to go-to-market approaches; and reimagining operations end to end, spanning sales, customer success, support functions, and the infrastructure stack.

AI-native disruptors architected from day one around the technology are already making their presence felt. Start-ups such as Anysphere (developer of Cursor), Gamma, and Lovable are redefining speed and scale, achieving product–market fit in record time and scaling to hundreds of millions in annual recurring revenue (ARR) with teams of fewer than 100. These companies aren’t layering AI into legacy workflows; they’re rebuilding the entire software organization around AI-centric principles.

Meanwhile, incumbents are racing to catch up with large investments and company-wide imperatives. Salesforce, for instance, has launched Agentforce, declaring its ambition to become AI-centric and embedding AI into its development and operational backbone. Atlassian has launched a suite of intelligent agents while committing to organization-wide AI integration.

For incumbent software companies, the imperative is clear: becoming AI-centric is no longer optional—it is essential to remain competitive. Those that can successfully adapt and thrive will help define the next era of software.

This article explores why this transformation is so critical, what it means to become truly AI-centric, and the essential steps for software companies to achieve that goal. Among the sources informing this article is a recent survey of top executives from across the software landscape—including large enterprises, mid-size firms, and start-ups from around the world—offering a broad perspective of how industry leaders are navigating this shift.

What it means, and takes, to be an ‘AI-centric’ software company

Becoming AI-centric doesn't happen overnight or with incremental improvements. It is a foundational transformation across seven key shifts that span both sides of the dual paradigm, comprehensively reshaping how a software company operates and the organization required to drive it.

Software executives increasingly recognize the urgency and scale of this change. In our most recent survey, more than 60 percent of top industry executives identified these seven shifts as a top strategic priority. And the potential upside is compelling: 40 percent of software leaders expect AI to unlock more than 20 percent revenue growth beyond their current trajectory, with 11 percent anticipating gains of over 50 percent. At the same time, AI-centric organizations are reporting 20 to 40 percent reductions in operating costs and 12-to-14-percentage-point increases in EBITDA margins, driven by automation, faster cycle times, and more efficient allocation of talent and infrastructure (Exhibit 1).

Only by focusing their efforts on the seven areas can software incumbents hope to capture similar gains and stay competitive as the industry undergoes dramatic upheaval (Exhibit 2).

We believe software companies should begin investing across all these seven shifts today, even though the time to realize impact will be different for each. The operational shifts could drive a step change in velocity and productivity within a year, while the value proposition shifts will likely be felt in one to two years when the AI disruption to SaaS business becomes more pronounced. Lastly, the fundamental shifts to workforce and workplace will likely need two to three years to fully manifest.

1. Reinvent core products and launch new AI offerings

Our research suggests three distinct models are emerging for how software will incorporate agents to reinvent its product offerings (Exhibit 3).

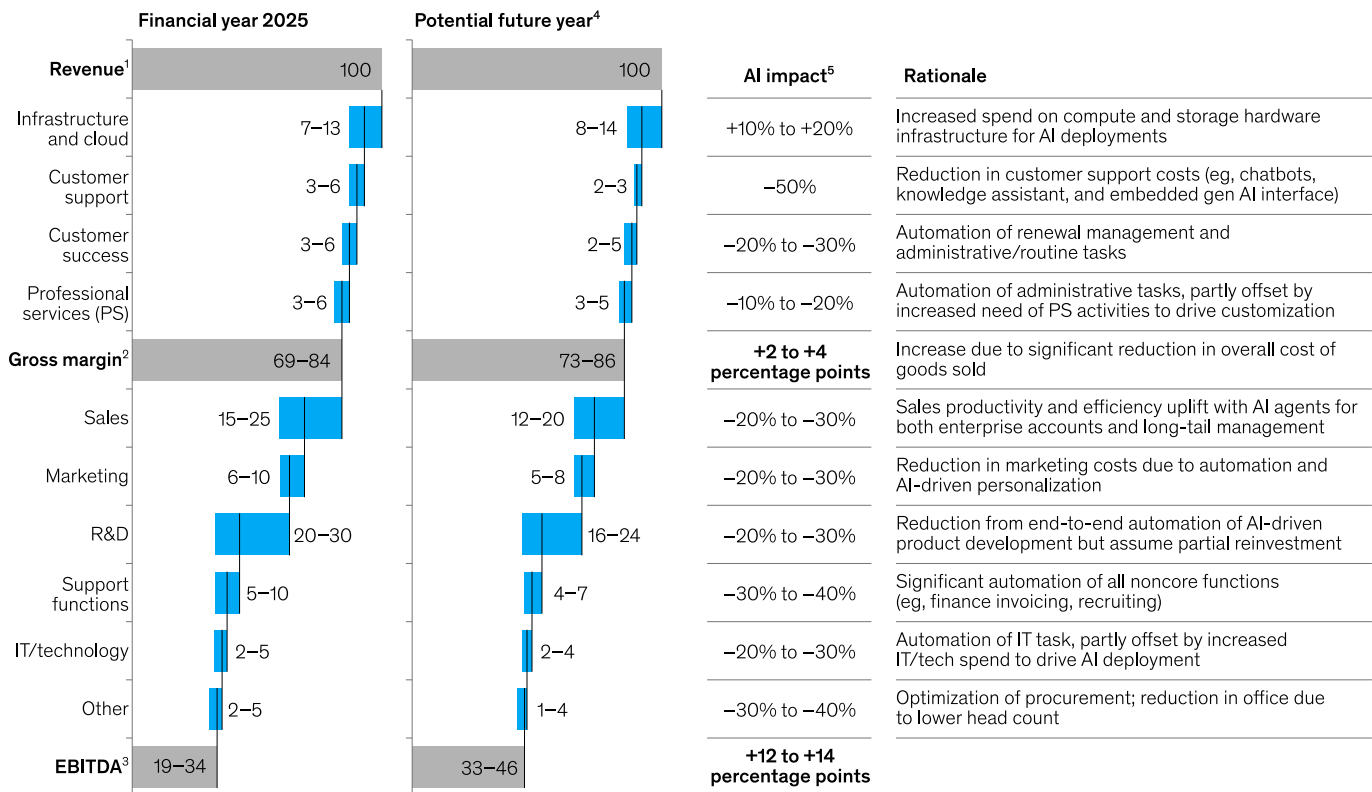
Archetype 1: Agents as users/augmentation

In this model, agents primarily automate the repetitive activities of human users, serving as users of existing SaaS software and following the business process that it enables. Software vendors continue to sell their core SaaS products along with agents, which convert enterprise

Exhibit 1

The AI-centric software company will likely have a markedly different cost structure than that of today's organizations, driving potentially higher margins and earnings.

Financials of software company in 2025, % (illustrative)



¹Revenue held constant for modeling purposes, but upside from new AI products/services could bring baseline up ~30–40%. ²Figures may not sum, due to rounding. ³EBITDA margins are computed based on a best-case scenario, including only the lower bound of operating cost ranges. ⁴Based on the lower bound of the AI impact range. ⁵Based on the highest potential cost reduction or increase over the next 5 years as cited by at least 10% of survey respondents.

Source: Bank of America Institute; IDC; Mizuho Securities; McKinsey Software CxO AI Impact Survey, Q2 2025, n = 81

McKinsey & Company

spend from labor to software. The value in this model is in the access to and management of underlying data from the SaaS software and the corresponding business process that it enables.

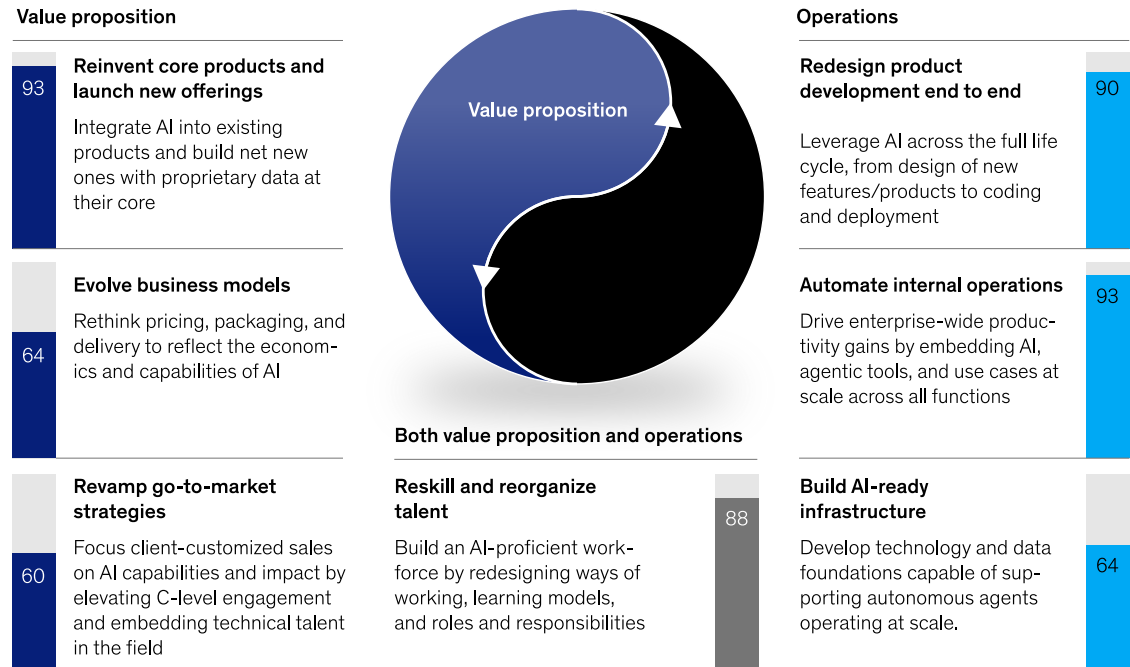
Archetype 2: Agent-centric architecture

In this post-SaaS model, a human employee primarily interacts with one agent interface while accomplishing their core workflows. That agent, in turn, interacts with several back-end agents and uses back-end APIs to act directly on data repositories. This should accelerate the commoditization of these elements, with more of the value moving to the agent layer. The value in this model is in the user experience (UX) and the ability of the agent to use tools that a human would rely on to accomplish its daily workflows.

Exhibit 2

Software companies need to make several strategic shifts in both their value proposition and operations to become truly AI-centric.

Software companies that have invested in AI-centric strategic shifts (past 1–2 years), % of respondents



Source: McKinsey Software CxO AI Impact Survey, Q2 2025, n = 81

McKinsey & Company

Archetype 3: Agents as experts

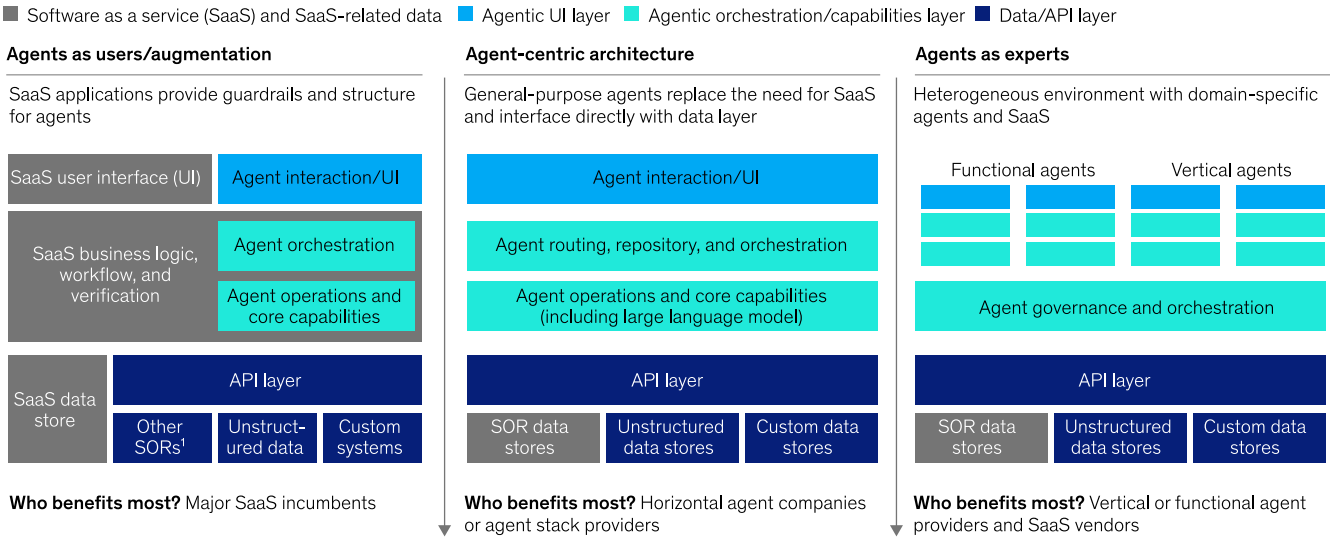
In this hybrid model, which combines elements of both SaaS and agent-centric architecture, the agents are distinguished by the functional or industry-specific knowledge the software vendor imparts to them (for instance, a legal agent that is trained by a team of lawyers). The value comes from the domain knowledge and proprietary data used to train the agent to perform in ways the other two archetypes can't.

These different models are already being used. Several workflows in critical business processes (such as sales or human capital) are being captured by archetype 1 players that have launched agents that primarily function as an add-on to the existing system of record or database

Exhibit 3

Software providers may have three distinct archetypes to choose from as they leverage agentic AI to reinvent their product offerings.

Archetypes by system architecture



¹Systems of record.

McKinsey & Company

software. Archetype 2 is seeing traction in individual productivity workflows, with horizontal agents capturing significant shares in knowledge work, data analysis, and content creation. Many of the most successful and well-funded app layer AI companies are in niche verticals (for example, legal tech, healthcare) where they position their domain-specific expertise and performance as primary differentiators.

Each of these models represents a viable route, albeit with distinct trade-offs. Augmentation agents can allow incumbents to make some rapid AI progress in an existing tool, but they are constrained by host APIs and vendor road maps. The more comprehensive and ambitious agent-centric architecture takes more time and effort to build and maintain, while expert agents involve some integration overhead.

But regardless of which archetype vendors use, AI is altering the source of value creation in software products. Competitive advantage in software is shifting from features to proprietary data access and control. Companies that own or integrate deeply with high-quality, domain-specific data can train more performant models, deliver more personalized outcomes, and ultimately justify premium pricing. Unlike features, which are easily replicated, data confers defensibility and monetization shifts from licensing functionality to charging for insight,

prediction, and automation, often through usage-based pricing or differentiated intelligence tiers.

2. Evolve business models

As AI-centric companies evolve their core products and launch new AI offerings, they will also need to reinvent their business models for the AI era across monetization, sources of differentiation, and delivery models. In fact, 63 percent of software leaders said they believe AI will fundamentally change their business model in the next three to five years.

AI is poised to do just that by shifting from per-seat (or user) pricing toward consumption pricing (whether usage/token-based, output-based, or outcome-based). This shift is critical for several reasons. As agentic AI takes over tasks traditionally performed by humans, the number of active users engaging with software decreases, rendering seat-based monetization on core SaaS software increasingly challenging.

What's more, with AI inference introducing new variable costs, particularly regarding compute and infrastructure, independent software vendors (ISVs) need to embrace a scalable pricing metric that protects margin as customer usage scales. In response, software companies are moving toward models that align revenue with the value delivered, charging based on outcomes achieved, actions taken, or compute resources consumed. Between 2015 and 2024, the number of consumption-based software companies more than doubled, and leaders such as Salesforce, Zendesk, Intercom, and LexisNexis are already monetizing their AI capabilities through these models, often unlocking significantly higher per-customer revenue than with traditional SKUs (Exhibit 4).

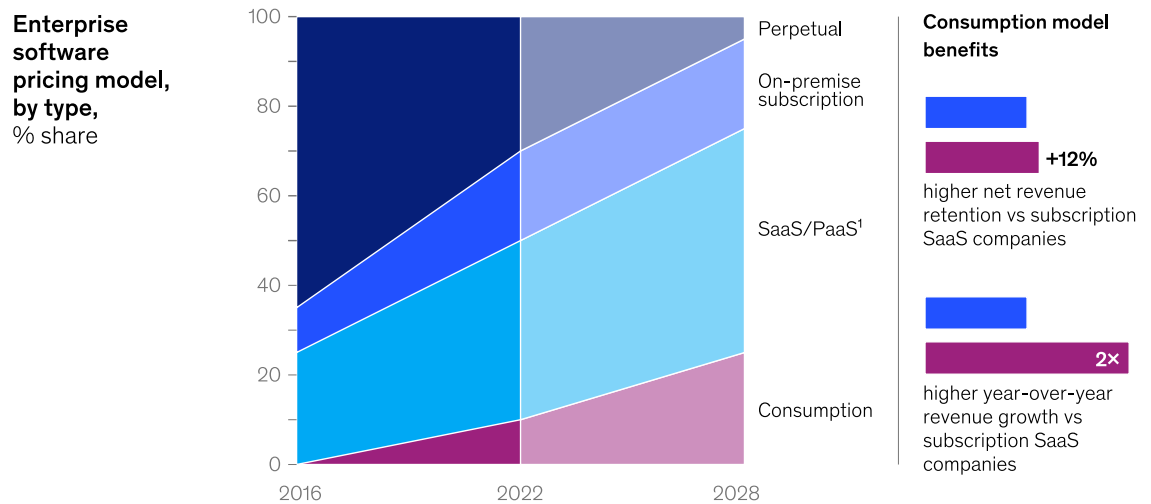
Delivery models are also evolving to support greater verticalization and customer specificity. For AI to deliver its full potential value, it needs access to relevant data, workflows, and context. As a result, one-size-fits-all applications are likely to be replaced with highly tailored applications, fine-tuned to specific industries, customer environments, and pain points. This is leading to the rise of service-as-software—a new delivery model in which software companies bundle platforms, AI agents, automation, and expert support into integrated, outcome-oriented solutions. Rather than simply providing tools, these offerings embed domain expertise directly into the product, solving end-to-end problems in vertical domains such as healthcare, retail, or financial services.

3. Revamp go-to-market strategies

Selling AI is not the same as selling traditional software. Its value is more complex, more contextual, and more variable. It requires new motions, new roles, and new partnerships. Reflecting this, nearly 70 percent of software executives rank go-to-market (GTM) transformation as a top investment priority for the next few years, making it one of the fastest-rising areas of strategic focus in the shift to AI.

Exhibit 4

As AI introduces new variable costs, more software vendors are turning to consumption-based pricing that enhances growth as customer usage scales.



¹Software as a service/platform as a service.
Source: Andreessen Horowitz; OpenView Venture Partners; Sequoia Capital; McKinsey SaaS Company Pricing Database

McKinsey & Company

Five shifts anchor this new commercial approach.

First, GTM strategies must now directly engage the C-suite. Since the emergence of gen AI, AI has been elevated from an IT decision to a board-level priority, a trend that has only been amplified by agentic AI. To position AI successfully, sales and marketing must speak to CEOs, CFOs, and chief technology officers (CTOs), framing the software not just as a tool, but as a partner in transformation.

Second, the sales motion must become more verticalized and personalized. As AI solutions become increasingly tailored to each customer's workflows, data, and use cases, generic sales pitches and broad messaging are no longer sufficient. Sales teams must adopt industry-specific language, showcase real outcomes, and align to customer success metrics that reflect AI's impact—not just usage.

Third, the nature of frontline roles is evolving rapidly. As AI products become more complex and customized, traditional sales engineers and customer success teams often lack the technical depth to help clients unlock real value. This is why leading firms are increasingly turning to a

new hybrid role known as forward-deployed engineers (FDEs), technical experts embedded directly on-site with clients to drive successful implementation, integration, and scaled adoption of AI. These teams bridge product and service, helping customers pilot gen AI, tailor agents to their workflows, and continuously optimize deployments. In consumption-led models—where usage drives revenue—this role becomes even more critical. FDEs not only accelerate time to value but also increase the likelihood of sustained adoption and deeper expansion within client organizations. As a result, forward engineering is emerging as a core part of modern GTM, especially for companies delivering AI-centric solutions that require high-touch, technically informed support.

Fourth, companies are rethinking their partner ecosystem to align with the distinct demands of the AI-centric era, where buying behavior, partner engagement, and value delivery diverge sharply from the traditional SaaS model. Today, CIOs, CTOs, and Chief Data and Analytics Officers lead nearly 60 percent of AI platform buying decisions, and a similar share of buyers rely on a hybrid approach with external partner support (such as system integrators [SIs] for AI advisory and implementation), giving partners significant influence over solution design and vendor selection. As a result of this closer, ongoing collaboration, companies are narrowing their partner mix to a smaller, more focused group of ISVs, SIs, and AI-native start-ups that not only sell solutions but co-develop and help implement them—often creating domain-specific agents. One leading enterprise software platform, for example, revamped its partner program to elevate strategic SIs with deep AI expertise, introduced industry- or function-specific sales playbooks across functions such as customer experience and tech risk, and launched an AI agent gallery to support scalable agent development. These shifts reflect a broader trend: Partner involvement now starts earlier in the sales cycle and extends post-sale to scale adoption and deliver sustained outcomes. Such a model increasingly depends on consumption-based incentives, deep technical training, and iterative, engineering-led deployment models that enable long-term value creation.

Fifth, agentic marketplaces will likely soon emerge, in which agents are discovered (potentially with agents), trialed, and consumed, all with minimal human input. These will initially be suited for specific tasks such as synthesis or data interpretation, fact gathering, or translation services. Over time, what is discovered and procured on marketplaces will increasingly be more specialized, “upstack” services.

4. Redesign product development end to end

One of AI’s primary applications has been software development. But while companies report 30 to 50 percent improvement in developer productivity in pilots, these gains rarely translate into equivalent impact on the top or bottom lines.

To unlock the full potential of AI-centric product development, companies must extend AI across the entire product development life cycle (PDLC) while also fundamentally redesigning the product development process itself and rethinking talent models.

Extending AI across [the PDLC](#) means embedding AI end to end—from discovery and viability, through build, test, release, monitor, and operate—to improve efficiency, quality, and experience. For example, during discovery, AI can automate requirement generation by helping to surface customer pain points, ideate potential solutions, and translate the best ideas into concrete requirements. In the build phase, agentic coding tools can generate detailed technical specs, write comprehensive tests, and iteratively generate code that both passes test cases and follows company-specific coding practices. During monitor and operate, AI-powered IT-ops assistants can triage incidents, perform root cause analysis, and provide predictive insights to prevent future issues. Together, these capabilities illustrate how deeply embedding AI across the life cycle can unlock full product development potential.

However, injecting these tools into the existing process will only go so far—a more fundamental redesign toward an AI-centric PDLC process is critical. This involves moving to a more systemized product development process, with consistent steps, artifacts, tools, and participants at each stage, and fewer stage gates and handoffs between them.

This new product development process and technology landscape will also require shifts in team structures, roles, and skills. Roles are being redefined to align with new process flows, and “one pizza” teams are emerging as role boundaries blur and skill requirements change, with specialized practitioners becoming agent managers.

This broader transformation has been applied successfully by a leading global IT services provider, which scaled gen AI solutions across more than 100 accounts and codified AI “recipes” across functions such as application development, testing, and SAP migration. Their programmatic transformation included rigorous pilot tracking, productivity-linked monetization, and deployment of gen AI agents—ultimately upskilling some 2,000 engineers on an AI-centric development process, team structure, and roles.

5. Automate internal operations

As software companies become AI-centric, they automate internal operations across the enterprise—not just in traditional back-office functions such as HR or finance but also in high-impact, front-facing areas including sales, customer support, marketing, and professional services. This broad-based transformation reflects a growing recognition that AI is a strategic lever for unlocking productivity at scale. In fact, alongside reimagining core offerings, automation of internal operations ranks as the top investment priority for software leaders today, with 93 percent citing it as a high-priority area.

Leading companies are already deploying gen AI and agentic AI to reduce costs, boost productivity, and free up employees to focus on higher-value work (Exhibit 5). In sales, for example, AI enables a full reimagination of commercial activity. Sellers are supported by virtual assistants and copilots that identify leads, prioritize outreach, draft responses, and provide real-time coaching. AI systems automate key steps in the sales cycle—from lead enrichment and

nurturing to smart pricing and proposal generation—helping teams reduce prep time by up to 80 percent and increase customer-facing time by 50 percent.

A recent case study with a leading HR software platform illustrates the power of this approach: Across seven AI/gen AI use cases, the company unlocked \$600 million to \$700 million in projected revenue uplift. This included deploying AI sales development representative (SDR) agents to autonomously manage long-tail accounts, resulting in a 25 percent productivity gain from freeing seller capacity.

Customer support is undergoing a similarly deep transformation. AI agents are being deployed to handle how-to requests, gather account information, and even resolve issues before the customer is aware of them. In one B2B software case, external-facing AI agents were trained to resolve simpler cases autonomously, while internal agents collaborated with human engineers on more complex issues. This hybrid model could lead to time savings of 50 to 60 percent for routine issues, 20 to 30 percent for account management, and 15 to 20 percent in troubleshooting and root cause analysis.

6. Build AI-ready infrastructure

While most incumbent software companies are already “tech-forward,” very few are truly AI-ready. Building products in the gen AI era requires a fundamentally different infrastructure foundation—one that supports not just cloud-based delivery and analytics, but the high-performance, real-time needs of autonomous agents operating at scale. For many software leaders, this means developing a new muscle: investing in core CIO-level capabilities and materially evolving their platform architecture to support AI-centric workloads.

Supporting agentic AI at scale involves more than plugging in a large language model (LLM). It requires retooling foundational infrastructure across five major components: the data layer, governance, security, developer tool chains, and agent operations. Some elements of these components, such as the LLM layer itself within developer tool chains, may only require light upgrades. However, others represent major renovations or even net-new innovation.

For example, FinOps and context awareness need to be updated for AI. Companies will need significantly more transparency into token and cloud resource consumption—at the account, user, and agent level—to manage cost and provide billing clarity. Ensuring that large models have sufficient context awareness using proprietary data also demands new tools, governance, and compute capacity. Meanwhile, AI orchestration, agent evaluation, and agentic mesh—which didn’t exist as core capabilities under previous AI paradigms—will need to be built to support agent routing, performance tracking, and life cycle management.

The economics of AI infrastructure heighten the urgency of this shift. One of the most immediate and visible impacts of AI adoption is rising infrastructure costs, driven by the high compute intensity of inference workloads. Our survey found that 80 percent of software executives

Exhibit 5

Software leaders are deploying high-impact use cases across multiple functions as they work to become AI-centric organizations.

Software company gen AI/agent AI adoption, top use cases, and impact, by business function

Current AI adoption by function, %	Top use cases ¹	AI-driven cost reduction, ² %
<div> <div></div> <div>0100</div> </div> <div> R&D engineering <div>89</div> </div>	<ul style="list-style-type: none"> AI-driven code generation and review AI-driven prototyping 	~20–30%
<div> <div></div> <div>0100</div> </div> <div> Customer support <div>74</div> </div>	<ul style="list-style-type: none"> Conversational AI Virtual knowledge assistance and augmented contact center 	~50%
<div> <div></div> <div>0100</div> </div> <div> Product management <div>67</div> </div>	<ul style="list-style-type: none"> AI-powered requirements generation New feature suggestions based on usage insights 	~20–30%
<div> <div></div> <div>0100</div> </div> <div> IT <div>64</div> </div>	<ul style="list-style-type: none"> IT agent to address/solve help desk queries AI-powered cybersecurity threat detection 	~20–30%
<div> <div></div> <div>0100</div> </div> <div> Sales <div>58</div> </div>	<ul style="list-style-type: none"> Automated bespoke pre-meeting prep/account plan generation for sales reps AI-driven next-best opportunity/action 	~20–30%
<div> <div></div> <div>0100</div> </div> <div> Back office <div>54</div> </div>	<ul style="list-style-type: none"> HR: HR assistant for self-serve HR functions; Finance: revenue forecasting and scenario modeling; Legal: rapid, automated legal document review 	~30–40%
<div> <div></div> <div>0100</div> </div> <div> Marketing <div>53</div> </div>	<ul style="list-style-type: none"> AI-generated marketing collateral Hyperpersonalized nurture and outreach campaigns 	~20–30%
<div> <div></div> <div>0100</div> </div> <div> Customer success <div>51</div> </div>	<ul style="list-style-type: none"> AI-curated customer success insights Curated customer onboarding and adoption recommendation 	~20–30%
<div> <div></div> <div>0100</div> </div> <div> Professional services <div>26</div> </div>	<ul style="list-style-type: none"> Automated synthesis of customer interactions Knowledge agent to drive implementation process flow 	~10–20%
<div> <div></div> <div>0100</div> </div> <div> Procurement <div>20</div> </div>	<ul style="list-style-type: none"> Automated invoice matching Automated creation of core docs 	~30–40%

¹Based on survey respondents who are currently using AI for specific use case or planning to do so in the next 12 months.

²Based on the highest potential cost reduction over the next 5 years as cited by at least 10% of survey respondents.

Source: McKinsey Software CxO AI Impact Survey, Q2 2025, n = 81

McKinsey & Company

expect cloud and infrastructure costs to increase by at least 10 percent, with over a third anticipating increases above 20 percent. As these costs surge, companies will need to optimize across their stack not only by rearchitecting agent workloads for efficiency but also by rationalizing non-AI infrastructure and improving telemetry to better understand system-level performance and spend.

We estimate that for many companies, 30 percent of today's shared platform capabilities will require meaningful change to support AI agents at scale. Companies that fail to make these investments risk poor performance, unchecked costs, limited agent impact, and growing technical debt.

7. Reskill and reorganize talent

AI is changing not only what work gets done, but who does it and how. To unlock its full potential, companies will need to cultivate AI fluency across their organizations by redesigning the “human + agent” workforce. This will include reinventing learning models by creating custom AI training pathways, such as evolving roles and skills, flattening structures, and reinforcing change through new support mechanisms.

At the center of this transformation is a fundamental shift in roles and capabilities (Exhibit 6). AI is already propelling a 20 to 30 percent net impact on workforce composition, with agents largely replacing some roles—renewal managers, support engineers, and SDRs—and freeing up those employees to be redeployed. Other roles, such as software engineers, account managers, and digital marketers, are being reshaped to integrate AI into their daily work, requiring new skill sets regarding agent interaction, prompt design, and oversight and intervention as warranted.

Entirely new roles, such as prompt engineers, agent coaches, and gen AI safety leads, are essential to operating and scaling agentic systems. However, with qualified talent scarce, companies often have to rely on internal upskilling to fill the positions. A leading B2B software company is now formally integrating agents into its organizational structure, with managers having both employees and agents reporting to them.


This transformation also requires rethinking team structures. Organizations are moving toward leaner models with fewer layers, redistributing work across flatter, cross-functional teams where humans and agents collaborate in real time. Junior and especially midlevel roles are shrinking as automation takes hold, while experienced talent is refocused on strategic, creative, and problem-solving tasks. Traditional apprenticeship and mentorship pathways must also be redesigned for the AI era, with structured, role-specific onboarding focused specifically on AI, faster cycles of ownership, and targeted upskilling to accelerate the development of future leaders.

To ensure sustained adoption, companies can embed change management by leveraging nudges, leadership modeling, and formal interventions to reinforce new behaviors. Cultivating an

Exhibit 6

AI is already meaningfully affecting the size and scope of the software workforce of the future.

AI-driven impact on software sector workforce makeup, by function, % share



McKinsey & Company

AI-fluent workforce means investing in technical training and a broad organizational rewiring—one that aligns roles, teams, and culture with the realities of the human + agent future.

In an industry where competitors and upstarts are rapidly embedding generative and agentic AI into both products and operations, piloting AI at the edges is no longer sufficient. To have a chance of thriving in this new technological landscape, software companies must take the steps to become truly AI-centric organizations. Such an overhaul revolves around two critical fronts: the company's value proposition and its internal operations. That means reinventing its products, business models, and GTM approach from the ground up, while also fundamentally reimagining how the company itself operates—which involves scaling AI across internal workflows and building the infrastructure and ways of working that unlock significant productivity gains.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



For software leaders, this shift begins with a clear AI-centered vision and enterprise-wide mandate anchored to ambitious, measurable targets. Deployment must be strategic, not scattershot. The most successful players will be those who focus AI where it creates true differentiation—whether by building proprietary, customizable capabilities, reshaping core workflows for scale, or leveraging existing solutions smartly instead of reinventing the wheel.

As part of this transformation, resources will need to be reallocated decisively, with capital and top talent shifted toward high-impact AI priorities, including the IT, data, and infrastructure backbone necessary to support AI-centric workloads. It will be equally critical to scale the workforce by hiring and upskilling for AI fluency while also redesigning teams, workflows, and culture to ensure long-term adoption and embedding sustained change management from the outset. Finally, a reinvention of this scale and scope requires accountability. Software leaders can establish centralized tracking of AI performance against pragmatic outcomes—spanning productivity, speed, and business value—to guide execution in real time and adjust rapidly as technology and use cases evolve.

Given AI's dizzying pace of innovation and evolution, the ability to adapt to rapidly shifting developments will be critical, if also challenging. Indeed, the current pace of disruption may seem daunting for software incumbents that have gotten used to the steady rhythms and riches of SaaS. But the biggest risk may be for those providers that don't at least try to make their AI-centric journey, and who soon find this new software era, and its massive value potential, rapidly passing them by.

[Jeremy Schneider](#) is a senior partner based in McKinsey's New York office, where [Joshua Abraham](#) is a partner, [Mónica Perestrelo](#) is a consultant, and [Shayan Salam](#) is an associate partner; [Matt Linderman](#) is a partner in the Connecticut office; and [Naveen Sastry](#) is a senior partner in the Bay Area office.

The authors wish to thank Andrea Suriano Simoncini, Angie Selzer-Wong, Lari Hämäläinen, Martin Harrysson, Roger Roberts, and Tejas Shah for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

Copyright © 2025 McKinsey & Company. All rights reserved.