



Beyond the code: How AI is redefining the software industry

Insights from 400 European software CTOs
2026



The better the question.
The better the answer.
The better the world works.



EY Parthenon

Shape the future with confidence

 **POTLOC**



Contents

Foreword: four structural shifts redefining the software industry	4
1. Software moat has moved up the stack	6
2. SDLC has been re-architected around AI	14
3. Engineering talent model has been refined	22
4. Software economics have been rewritten	28
Final words	36
Survey methodology and demographics	40
About EY-Parthenon Software Strategy Group	44

Foreword: four structural shifts redefining the software industry

Over the past year, software leaders, boards and investors have repeatedly asked us similar questions:

- Is software becoming easier to replicate?
- Will AI-native challengers disrupt incumbents?
- Do engineering teams need to become smaller?
- How much premium should we pay for AI maturity?

Strong opinions on these topics are easy to find. Data is harder.

This gap between opinion and evidence motivated us to conduct structured research on how AI is reshaping the software industry as a whole. Our findings confirm certain widely held assumptions, challenge others and, most importantly, reveal a more nuanced reality than either optimists or pessimists suggest.



Sander van Weele

Partner



For decades, software companies competed through the breadth and depth of product functionality, engineering excellence and the ability to deliver innovation at speed. Recent years have challenged many of the industry's long-held assumptions. Our 2025 research highlighted the changing Chief Technology Officer (CTO) mandate and the emergence of artificial intelligence (AI) as a strategic priority across the software sector.

This year, our findings point to something further: AI is no longer just changing how software is developed; it is changing what makes software companies valuable. It is reshaping the foundations on which they compete.

This is a broader structural shift in the sector, rather than a temporary technology cycle

Drawing on insights from 400 European software CTOs and our experience with technology leads and investors, this report examines four structural shifts that are redefining the software industry as a whole:

- Software moat has moved up the stack
- Software development lifecycle (SDLC) has been re-architected around AI
- Engineering talent model has been refined
- Software economics have been rewritten

1

Software moat has moved up the stack

For a generation, software was defended by how hard it was to build. But that defence is now disappearing. This is the first of the structural shifts examined in this report.

Code, architecture, years of effort once kept competitors out. That is no longer the case


A decade ago, a software leader would say that the scale of construction made their product hard to copy: millions of lines of code, an architecture refined over successive releases and an engineering organization that understood it all. That accumulated effort deterred entrants and locked in customers, because rebuilding something equivalent was simply too costly to attempt.

This form of defensibility is now actively eroding. As AI reduces the cost of building software, a large codebase that was once a fortress becomes a liability to maintain. What AI erodes is not the architecture itself, but the advantage of the patterns that have become common across the industry. The advantage that came from having built something large and complex fades, because the cost of building something comparable is falling for everyone – including the challenger who has not built it yet.

The survey bears this out at the product level. When leaders identify parts of their product most exposed to replication, they point to the standardized, visible layers, the interface, reporting and analytics, standard integrations and generic workflow logic. These are precisely the areas that are now the easiest to compress. What sits at the safe end is telling: proprietary data, compliance and regulatory logic, and customer-specific implementation. The construction of the software code is becoming replicable. The context around it is not.

Product, workflow and implementation areas most vulnerable to AI-driven replication over the next three years

% of respondents



Source: CTO survey by EY-Parthenon and Potloc

// We have seen examples of vibe-coded CRM applications reaching 80% feature parity compared to the target product in just a couple of weeks. Yet, without the underlying data, regulatory logic, customizations and a maintainable architecture to embed it all in, those examples are little more than empty shells.



Erik Oltmans
Partner




The real moat was never the product – it was the context around it

When leaders rank what defends their core product today, no single factor dominates. Model quality leads only narrowly, sitting alongside workflow ownership, proprietary data, integration depth and trusted brand. We read this clustering as an important signal: even now, at least four of the top five sources of advantage have nothing to do with how well the software is engineered – and everything to do with the position it holds in the customer's environment.

Model quality is the one product attribute that still ranks highly and, in our view, is the most fragile of the group. It is exactly the capability that AI makes easier to reach over time, which means today's model advantage is tomorrow's baseline. The durable advantages are the ones a competitor cannot download or reproduce: the dataset built from years of customer activity, the workflow a customer runs their business on, the integration that would be painful to unpick and the trust that took years to earn. These do not get cheaper as AI improves. If anything, they get scarcer.

Key sources of product defensibility: today vs. in three years
% of respondents



Source: CTO survey by EY-Parthenon and Potloc

And the ground keeps moving toward context

This is not a static picture. Looking three years ahead, leaders expect the balance to tilt further toward context, not away from it. Model performance, evaluation data and AI quality are the clearest risers, followed closely by proprietary data. The advantages that depend on access, position and trust are expected to matter more than raw model quality.


The strategic implication is uncomfortable for product-led organizations – If your differentiation lies in features, you are defending the part of the stack that is depreciating fastest. If it lies in the data you own and the workflows you run, you are defending the part that is appreciating. Software companies that recognize this early will stop investing in being harder to build and start investing in being harder to leave.

This is not a forecast – it is already happening

It would be easy to dismiss all of this as a story leaders tell themselves about the future. The survey tells us otherwise. When asked what an external investor or board would recognize as evidence-backed proof of AI advantage, leaders point to moats reinforced by AI across workflow, integration, compliance, trust and data.

The pressure is just as visible on the demand side. More than four in ten technology leaders say that, as part of renewal conversations, customers routinely probe whether part of their product could be rebuilt or replaced using AI.

Areas where investors and boards see the strongest evidence of AI advantage
% of respondents



Source: CTO survey by EY-Parthenon and Potloc


The AI-native threat is loud, but it is not yet lethal

The prevailing narrative says AI-native challengers, unburdened by legacy architecture and technical debt, will sweep incumbents aside. However, the survey tells a more nuanced story. AI-native competitors are now visible in almost every segment – but just over half of CTOs say these entrants have not cost them a single deal or lost them usage. Only a minority of 14% report losing meaningful share.

Perhaps the market is confusing noise with impact. AI natives are excellent at generating attention – demos, funding rounds and buyer curiosity – but that attention has not yet converted into share at the scale the narrative implies. This is the single most important nuance in the AI-competition debate, and it is where incumbents still have room to act.

Competitive impact of AI-native entrants and drivers of their success

% of respondents



Source: CTO survey by EY-Parthenon and Potloc

When challengers do win, it is on speed – not on AI


When incumbents do lose, the reasons why reveal what the challenger's advantage is built on: faster time-to-value, shorter implementations and cleaner user experience. In other words, challengers win on how quickly and easily they deliver value, not on how sophisticated their technology is. That is an important distinction because speed and experience are gaps that an incumbent can close. A structural technology disadvantage would be far harder to answer.

Challengers take commodity edges. Incumbents hold defensible core

In cases where AI natives advance, they do so in the commoditized edges of the market: in data extraction and document automation, in self-serve and small-business segments, and in low-regulation workflows. They make the least progress in regulated and core enterprise systems, where domain expertise, trust, compliance and deep integration still decide the winner. The incumbent's strongest edge, CTOs claim, is domain expertise, followed closely by trust and regulatory approval. These are not features; they are context – the same context identified earlier as the durable moat.

Market segments and use cases where AI-native entrants gain share fastest


% of respondents



Source: CTO survey by EY-Parthenon and Potloc

Incumbents' strongest competitive advantage compared to AI-native entrants

% of respondents, ranked top 3




Source: CTO survey by EY-Parthenon and Potloc

Follow the money, and the strategy is clear

The clearest evidence of where the moat has moved is where leaders are spending. Yes, the single most-funded action is shipping AI features faster – the pressure to be seen keeping pace is real, even if much of it answers the hype rather than building durable advantage. But taken together, the money flowing into deepening workflows and integrations, building proprietary data assets, and strengthening governance, trust, and compliance, outweighs the spend on features alone. Whatever leaders say about the AI feature race, their budgets are quietly reinforcing the context moat.

Most funded initiatives in response to risk of AI replication and AI-native competition

% of respondents, ranked top 3



Source: CTO survey by EY-Parthenon and Potloc

The takeaway: stop defending the product, start owning the context

AI is making software code cheaper to write, by stripping the protection out of the codebase and architecture that once kept competitors out. At the same time, it is raising the value of everything the product sits on: proprietary data, embedded workflows, customer integrations, trust and compliance. AI-native challengers are winning at the commoditized edges on speed and experience, not seizing the defensible core on technology. The moat has moved up the stack, from what the software does to the context in which it operates.

That context is produced in how software is designed, built, tested and shipped. The moat may have moved up the stack, but it is still manufactured on the factory floor. This is why software development lifecycle (SDLC) becomes the central battleground. The next chapter turns to that factory floor: how the SDLC is being re-architected around AI and why that re-architecture decides whether the moats described here can be built and defended at all.



AI will not eliminate enterprise software, but it will fundamentally redefine where value is created. Software vendors that leverage proprietary data, support mission-critical or vertical-specific workflows, orchestrate ecosystems or provide regulatory-compliant systems of records will be best positioned to capture value in the near term.



David Canaple
Partner



2 Software development lifecycle has been re-architected around AI

The re-architecting of the software development lifecycle (SDLC) around AI is the second of four major structural shifts explored in this report.


Historically, software delivery relied on a series of handoffs between business stakeholders, product managers, architects and engineering teams. Business strategy was translated into product strategy, decomposed into specifications, refined into requirements and, ultimately, converted into software.

While the lifecycle steps remain largely unchanged, software organizations now report continuous use of AI across development activities, including code generation, coding assistance, debugging, testing and requirements creation. In some cases, AI agents sit among the stakeholders in the handoff chain.

AI has compressed software creation

This use, however, is no longer indiscriminate. Our findings point to software organizations moving from an early experimentation phase of applying AI everywhere toward a more structured, deliberate application across the SDLC - at points where it adds most value.

Most popular use cases for developer AI tools in software development process
% of respondents




Source: CTO survey by EY-Parthenon and Potloc

At the same time, one in three CTOs surveyed believe that, in three to five years, over 50% of new or materially modified production code will be AI-authored or agent-generated.

This is no longer about faster coding; it is a rebuilt production line. While the first wave of AI adoption focused on code generation, AI participation – now visible across the wider SDLC – is reshaping product management and engineering processes, reducing friction between ideas and implementation, and compressing the path from strategy to implementation.


Figure 1: AI-assisted software development lifecycle



Coding is no longer the bottleneck


Widespread adoption has settled the productivity debate. Organizations now report measurable efficiency increase in R&D processes – improved delivery time and release cadence, shortened approval times and fewer security reworks. Almost three in four respondents also report improved cost-to-ship.

Measured efficiency increase or decrease from AI coding assistants or coding agents in R&D process over the last 12 months
% of respondents



Source: CTO survey by EY-Parthenon and Potloc

R&D outcomes have changed since adopting AI coding tools or agents
% of respondents, excluding where outcome is not measured



Source: CTO survey by EY-Parthenon and Potloc

A year or two ago, using AI to code faster was itself a competitive advantage – early adopters moved faster than everyone else. Now, AI-assisted coding is an industry-wide expectation and rapid software creation is no longer a differentiator. The advantage today lies in what organizations do with the capacity AI frees up. Technology leaders recognize this and are redeploying existing capacity to execute more ambitious roadmaps with improved quality.

56% of CTOs report redeploying time saved by AI coding tools to quality assurance and code reviews, 55% to new features and roadmap items, 45% to modernization or refactoring


Notably, this capacity is not primarily flowing into cost reduction. It is being redirected toward quality assurance, product innovation and modernization efforts that were previously difficult to fund. As AI-assisted coding becomes a baseline capability across the industry, differentiation is shifting from how fast organizations can generate software to where they choose to deploy the capacity that speed creates.

AI didn't remove the bottleneck – it pushed it downstream

How organizations redeploy released capacity reveals new downstream constraints in the SDLC: in validation, governance, security and architecture. Organizations that modernize these downstream capabilities will capture disproportionate value from AI, while those focused solely on developer productivity risk creating faster pathways to existing operational bottlenecks.

Most common SDLC bottlenecks

% of respondents, ranked top 3




Source: CTO survey by EY-Parthenon and Potloc

As AI-generated code volumes increase, organizations are introducing additional controls to manage reliability, maintainability and risk. Mandatory human review has emerged as a common practice, alongside established governance mechanisms around approved tools, testing requirements and security controls.

Practices in place to ensure reliability, maintainability, and security of AI-generated or AI-assisted code

% of respondents




Source: CTO survey by EY-Parthenon and Potloc



The need for these practices is reinforced by the data: roughly one-third of respondents report having experienced at least one AI- or agent-related production incident in the past year. In many respects, software organizations are moving from a production problem to a trust problem.

AI-related incidents breakdown by type

% of respondents, 100% = 131



Source: CTO survey by EY-Parthenon and Potloc

Winning companies redesign the SDLC

These incidents point to new categories of operational risk extending beyond traditional software defects. Organizations increasingly need to manage issues related to model behaviour, security, compliance, explainability, auditability and autonomous decision-making. As a result, governance is moving from a peripheral activity into a core component of software delivery.

The SDLC no longer ends at deployment; it extends into monitoring, governance and continuous operational assurance. This is reshaping how technology leaders think about the engineering workforce required to capitalise on AI opportunities.



As code generation becomes increasingly accessible, software development is no longer the primary bottleneck in translating ideas into tangible products. The time required to move from concept to a functional product demonstration has compressed significantly, enabling individuals and teams to test, iterate and communicate ideas with unprecedented speed.



Shaleen Shankar
Partner



3

Engineering talent model has been refined

The widespread adoption of AI is forcing software companies to rethink foundational assumptions about how engineering talent creates value. As code generation becomes increasingly automated, the activities that differentiate engineers are shifting toward architecture, judgment, validation and governance. Let us take a look at the third structural shift redefining the software industry.

Technical judgment is becoming the engineering differentiator

AI has expanded access to many activities that were previously gated by engineering capacity, from code generation and testing to documentation and debugging. Yet the most consequential decisions in software development remain fundamentally human.

Organizations still need people capable of making informed decisions about architectural trade-offs, evaluating technical risk, balancing short-term delivery against long-term maintainability and determining how technology decisions support broader business objectives. Available technology helps ideate and increasingly generates solutions, but it does not remove the need to choose between them – someone still needs to make the decision and stand by it.

As a result, the most sought-after individuals are not necessarily those who can produce the most code, but those who can provide direction, exercise judgment and guide increasingly AI-assisted delivery toward the right outcomes. This shift in what engineers are valued for has implications well beyond individual roles – it is reshaping hiring strategies, career paths, team structures and the mandate of technology leadership.

// What makes for a good software engineer is shifting. Factual technical knowledge and experience used to translate to faster and more effective production of code. But that is no longer sufficient. Deep insights in solution patterns, trade-offs, critical thinking and judgment are the traits that now make a difference, allowing teams to find the fastest path to robust implementation.



Prof. Dr. Ir. Joost Visser
External expert EY-Parthenon
Prof. Large Scale Software and Data Science at Leiden University

AI is reshaping the industry talent mix

Public discussion around AI has often focused on workforce reduction, with technology companies citing automation and productivity improvements as drivers of restructuring. Our survey points to a more nuanced reality within software R&D organizations: continued demand for technical talent, albeit with a changing profile, as the emphasis is moving from scaling the number of engineers to maximizing the impact of each one.

As software delivery becomes increasingly automated, organizations place a growing premium on expertise, context and decision-making capability. Over the past 12 months, respondents report the strongest growth in experienced engineering talent, AI and data-related roles, and platform, security, QA and compliance functions. Capacity allocated to junior engineering positions and outsourced delivery models has contracted materially over the same period.

Hiring plan for the next 12 months by seniority or role

	Reduce hiring % of respondents	Increase hiring % of respondents
Junior / graduate engineers	44%	5%
Experienced software engineers	3%	75%
AI / data, QA and security control roles	4%	73%
Contractors, near- and off-shore resources	51%	5%

Source: CTO survey by EY-Parthenon and Potloc

The pathway to engineering expertise is changing

As the composition of software engineering teams changes, so do the typical career pathways for industry talent. Technical expertise has long been considered a cumulative learning process. Senior engineers, architects and technology leaders rarely emerge from formal education alone: their expertise is built through years of hands-on delivery experience.

Many of the activities that formed the foundation of the engineering learning curve are now increasingly supported, if not replaced, by AI. Routine coding, documentation, testing and implementation of smaller features can often be completed faster and more efficiently through AI-assisted workflows. While this improves productivity, it creates an important tension: how do engineers develop the judgment required for senior roles when many of the traditional stepping stones to acquiring it are becoming less prevalent.

Impact on junior engineering roles in the past 12 months

% of respondents




Source: CTO survey by EY-Parthenon and Potloc

Many organizations are raising expectations for entry-level talent, from routine implementation to activities that complement AI-assisted development: testing, evaluation, specification, validation and effective use of AI tools. As a result, the challenge is no longer solely attracting junior talent but rethinking how they can acquire the experience required to become senior engineers, architects and technology leaders.

This is forcing organizations to reconsider the traditional progression from graduate engineer to senior technical leadership roles. In response, many CTOs report redesigning their apprenticeship model.

Change in junior-to-senior apprenticeship model

% of respondents



Source: CTO survey by EY-Parthenon and Potloc

At the same time, a notable minority report that today's trainees have fewer opportunities for hands-on learning along with increasing pressure on senior engineers responsible for coaching and review. This highlights one of the less discussed consequences of AI adoption: while technology accelerates software delivery, it can also disrupt the mechanisms through which future architects, engineering leaders and CTOs are developed.

Mentorship matters more than ever

Technology leaders need to focus on mentoring and career development models, so that the next generation of engineers can still acquire the right technical judgment.

When junior engineers no longer engage in large volumes of implementation work, judgment must be taught more explicitly and intentionally: through structured code review, deliberate exposure to architectural decisions, and direct coaching from senior engineers on the trade-offs AI cannot make on its own. Organizations that treat mentoring as a casualty of efficiency gains risk hollowing out their own future leadership pipeline.


The CTO mandate continues to expand

Workforce shifts outlined in this chapter extend well beyond hiring decisions. R&D leadership needs to address a new set of organizational challenges that did not exist a few years ago.

Leaders must determine which activities should be automated and which should remain human-led, establish governance around AI usage, define new productivity expectations, redesign career progression frameworks and preserve the mechanisms through which future engineers acquire technical judgment – all while ensuring software quality is maintained despite rapidly changing development practices.

The result is a growing leadership burden. AI reduces the effort required to build software, but it increases the complexity of managing the people behind it and adds new AI-related accountabilities under the CTO scope. Over 60% of CTOs report moved or increased accountability for AI product roadmaps, and almost half cite AI governance and AI vendor selection.

AI-related accountabilities moved to or expanded under the CTO mandate % of respondents



Source: CTO survey by EY-Parthenon and Potloc

This reinforces a trend identified in our 2025 research: the remit of the CTO continues to expand. Last year's survey highlighted the growing involvement of technology leaders in business strategy and broader organizational decision-making. The 2026 findings show that mandate evolving further – CTOs are becoming the operational owners of AI transformation.

CTOs already hold primary decision rights for internal R&D coding agents across the vast majority of organizations. Now, they also play a leading role in internal AI productivity initiatives, AI governance, model controls, security and AI value measurement. Rather than sitting at the periphery of AI adoption, technology leaders increasingly find themselves at the center of the decisions that determine how AI is deployed, governed, measured and ultimately translated into business value.

Primary decision rights owner by domains

Area	Role / function (% of respondents)
Customer-facing AI features	CPO / product (83%)
Internal R&D coding agents	CTO / engineering (69%)
Internal enterprise productivity agents	CTO / engineering (53%)
AI governance and security	CISO / security (34%)

Source: CTO survey by EY-Parthenon and Potloc

All this is beginning to alter the economics of software creation itself. Understanding whether AI acts primarily as a cost reduction lever, a growth multiplier or an entirely new economic model is becoming one of the most important strategic questions facing software companies today.

4

Software economics have been rewritten

AI is separating value creation from value capture. Software companies are embedding AI deeper into their products than ever before, driving revenue. But the ability to price it, charge for it and protect margin on it has not kept pace with the ability to build it. This gap is the defining economic challenge of the current phase of AI adoption and the final structural shift investigated in this report.

AI is already contributing meaningfully to software revenue


More than one-third of respondents report that over 25% of software or data-enabled revenue is linked to AI capabilities, AI-generated outputs, autonomous workflows or AI agents. AI is already influencing purchasing decisions, retention, adoption and expansion.

AI is no longer treated as a standalone capability. It is increasingly embedded directly within customer workflows as a recommendation, search, summarization, decision-support or automation layer, which can make it difficult to distinguish where software ends and AI begins.

So how do software companies capture value when AI becomes part of every product?

Role of AI in the main customer-facing product today

% of respondents



Source: CTO survey by EY-Parthenon and Potloc

The industry has not yet figured out how to charge for AI

Our data shows that companies have already started experimenting with a range of monetization models - from volume- and outcome-based pricing to token-based models. However, the majority of respondents (~57%) still package AI functionality within existing product subscriptions, rather than charging for it separately.

This approach may be practical in the short term, but it risks creating a gap between the value AI delivers and the revenue companies can capture, given that different AI capabilities lend themselves to different economics. AI copilots align with seat-based models. AI agents that execute workflows support transaction- or task-based pricing. Infrastructure-oriented offerings such as Model Context Protocol (MCP) servers, AI platforms and AI runtimes are often better suited to usage, API, token or compute-based models. In some cases, AI decision engines support outcome-based pricing if the value they generate can be measured directly.

//
Right now, AI is a cost line before it is a margin line.

Most organizations are carrying real spend on infrastructure, talent and model consumption, while the efficiency gains are still largely prospective. And, in the meantime, customers already expect AI to be part of what they are paying for today. That gap between committed cost and realized value is the central economic challenge of the next 18 months, and pricing is the only lever that closes it.




Isabel de Dios
Partner



Customer-facing AI capabilities pricing

% of respondents, rounded



Source: CTO survey by EY-Parthenon and Potloc

Why value creation and margin are no longer the same thing

Even where companies successfully embed AI and drive adoption, that success no longer automatically converts into better margins. This is the part of the value-capture gap that is easiest to overlook.

Three forces are moving at once, not in sequence: i) inference and compute costs scale directly with usage, in a way traditional software costs never did; ii) pricing models remain largely unchanged, absorbing that cost growth rather than passing it through; and iii) customer expectations are settling around AI being included "for free" within existing subscriptions, making it harder to introduce separate charges later.


A company can grow AI adoption, grow revenue linked to AI, and still see their margins being compressed – because the cost of delivering AI value is not fixed the way the cost of delivering traditional software was. Closing the value-capture gap requires treating pricing, cost architecture and customer expectations as one problem, not three separate ones to be solved independently.

Market dynamics make AI monetization even harder

Designing a sustainable pricing model is only part of the problem. Even if companies identify an economically rational pricing model, will customers and competitors allow them to charge that way? The survey shows that market dynamics are an even greater constraint. Respondents most frequently cite customer expectations, competitor pricing behavior and willingness-to-pay constraints as the primary barriers to monetizing AI.

Biggest challenges in monetizing AI capabilities

% of respondents selecting as first or second biggest challenge




Source: CTO survey by EY-Parthenon and Potloc

It is perhaps not surprising, then, that many organizations are still waiting to see how the market develops before making major changes to their pricing models.

AI-driven pricing and packaging changes in the past 12 months

% of respondents




Source: CTO survey by EY-Parthenon and Potloc

This uncertainty is understandable. However, one challenge remains difficult to ignore: someone ultimately needs to pay for the costs associated with delivering AI. Today, many vendors are absorbing a share of those costs, with nearly 60% of respondents reporting passing through up to a quarter of AI feature or agent usage costs to customers.

Share of AI feature or agent usage cost passed through to customers today

% of respondents




Source: CTO survey by EY-Parthenon and Potloc

Monetization is an acknowledged problem, but it is not yet treated as the primary battle

Technology leaders already recognize this challenge: when asked where they are furthest behind peers in AI readiness, more respondents point to commercial model and monetization than to technical or organizational factors.

Areas where organizations are furthest behind peers in AI readiness

% of respondents ranking as the biggest gaps




Source: CTO survey by EY-Parthenon and Potloc

However, monetization does not feature at the top of the 2026 CTO agenda; in fact, pricing and packaging redesign sits at the end of their priority list for the year ahead. This contradiction reflects where the market is today. The industry recognizes monetization as an unresolved challenge, but chooses to focus on adoption and product development, waiting for clearer market norms for AI monetization to emerge.

Top priority topics for 2026

% of respondents selecting as one of their top five priorities



Source: CTO survey by EY-Parthenon and Potloc

The monetization question is more urgent for incumbents than they currently assume

AI-native challengers can design products, architecture and commercial models around AI from inception. Established vendors face a more complex balancing act: introducing AI capabilities while simultaneously protecting existing revenue streams, absorbing AI-related costs and modernizing products built in a pre-AI world. For incumbents, monetization is more than a pricing issue. Revenue-critical products still require meaningful re-architecture before they can fully support AI-native workflows, and companies that capture more AI-driven value are expected to have more capacity to reinvest in that transition.

56% of incumbents say half or more of their revenue-critical products would require significant re-architecture to support AI-native workflows

Winners do not underestimate the urgency of AI monetization

We believe the next phase of AI is moving beyond proving that AI works to proving that AI pays. Organizations that can convert AI capabilities into measurable revenue and margin improvement will be better positioned to finance the next phase of transformation.

The AI race will be won not by those who build AI features first, but by those who establish the economics to sustain them. The most successful companies will use AI to increase productivity, improve product quality and create commercial models that allow them to capture and reinvest the value they generate.



The ubiquitous seat-based (Per User Per Month) Software as a Service pricing model is fundamentally challenged by the addition of costly AI components to the product offering, with typically fewer human users, but far more token use and cost per user. This is prompting CFOs to send already-stretched product teams back to the pricing drawing board.

In practice, finding new pricing models that align with customers' changing value perception has proven difficult. After years of operating with an easily understood model, strengthening commercial pricing skill sets should be high on the agenda.



Per Flaata
Partner



Final words: one technology, four industry shifts

The four shifts explored in this report are not independent consequences of AI – rather, they are different perspectives on the same underlying transformation. Together, they describe a fundamental reset in how software companies create, defend and capture value.

Software's protection has moved from the codebase to the context around it – the data, workflows, integrations and trust that AI does not make easier to replicate, even as it makes code itself easier to write. That same force creates a new production challenge: as AI removes friction from coding, it exposes validation, governance and trust as the real constraints on delivering software at scale. Talent scarcity is shifting for the same reason – from coding capacity to judgment, since the organizations that win are the ones who can direct AI effectively, not simply deploy it.

Underlying these themes is an economic question that the industry has not yet answered: value creation is outpacing value capture, and the winners of the next phase will be those that can monetize it.

**The moat
has moved
from product
to context**

**AI did not eliminate
the software develop-
ment bottleneck, it
moved it downstream**

**The next talent
shortage will be
judgement, not
coding capacity**

**AI adoption is no
longer a challenge,
turning it into
economics is**

For software leaders: stop scaling output, start redesigning the system behind it

AI can no longer be viewed merely as a productivity tool; it is a catalyst for broader organizational change.

The strongest software companies will need to redesign multiple dimensions of their business simultaneously: what differentiates the product, how software is delivered, how engineering talent develops and how value is monetized. Optimizing a single function in isolation risks solving yesterday's problems while missing tomorrow's transformational value.

Today's CTOs must be willing to rethink the assumptions that have shaped software companies for decades, from how products create defensibility to how engineering organizations scale and how value is captured. Organizations that treat AI as a feature or productivity initiative will realize incremental benefits. Those that use it as an opportunity to redesign how they operate will shape the next generation of software businesses.

AI adoption may still dominate today's agenda, but software leaders should consider elevating monetization higher on their strategic priority list. Organizations that solve the economics challenge earlier will be better positioned to fund innovation, product transformation and competitive advantage.

For software investors: seek the moat AI cannot copy and the economics that can fund it

As software becomes easier to create, technology differentiation alone is a weaker predictor of long-term success. Increasing attention should be paid to the assets that remain difficult to replicate: proprietary datasets, workflow ownership, customer distribution, regulatory positioning, trust and ecosystem integration.

This does not make technology less important. Rather, technology acts as an amplifier of existing advantages. The most resilient businesses use AI to strengthen defensible positions they already have, and do not rely solely on AI capabilities for differentiation.

Equally important is the ability to sustain continued innovation in an environment where both AI capabilities and AI cost structures continue to evolve rapidly. Model performance is improving, pricing continues to shift and the economics of AI deployment are being rewritten in real time. Investors that recognize which companies can adapt as these underlying economics change, while still funding innovation, will capture the most value.

The most compelling equity stories in the software industry will combine durable competitive advantage and a resilient business model, with AI capabilities as a baseline that drives stronger moats, faster innovation cycles, and economic resilience.

For lenders, M&A advisors and capital markets participants: look beyond the AI hype

AI adoption is rapidly becoming a baseline expectation across the software industry. Its presence alone reveals little about future performance, credit quality or long-term enterprise value.

The more relevant questions concern the durability of the business model. How exposed is the company to replication? Can productivity gains be translated into sustainable operating leverage? Is management successfully converting AI investment into recurring revenue, margin improvement and stronger customer retention? Does the business possess the governance, architecture and operating discipline required to scale AI without introducing disproportionate operational risk?

Our findings challenge long-held assumptions about software economics. Productivity improvements do not automatically translate into lower R&D spending, nor do they guarantee stronger profitability. Many organizations are reinvesting released capacity into product development, modernization, governance, security and AI initiatives while simultaneously adapting their pricing and packaging models.

For lenders and advisors, diligence increasingly requires understanding how AI affects the quality of growth, resilience of cash flows, sustainability of margins and future funding requirements. The strongest software businesses will distinguish themselves through defensibility, monetization discipline, adaptability and financial resilience. Those characteristics are likely to remain valuable regardless of how fast the underlying technology evolves.

Looking to 2027: from AI-assisted software to programmable markets?

As agents become capable of executing increasingly complex workflows, software will evolve from assisting decisions to initiating them. Activities that still require human coordination today will increasingly be delegated to systems capable of evaluating options, interacting with other systems and completing transactions autonomously.

Commercial models will become progressively more granular and programmable, with task-based, outcome-based and consumption-based pricing existing alongside traditional subscriptions. Assets, permissions, incentives and ownership rights could become increasingly machine-readable, allowing software systems to negotiate, transact and exchange value directly with one another.

When commercial interaction is executed by software rather than coordinated through people, entirely new market structures become possible. Questions of interoperability, governance, trust, incentives and value distribution will become as important as the software itself.

The next question for technology leaders and investors will therefore move beyond how much software AI can produce or how many tasks it can automate. It may instead become: what new businesses and market structures emerge when production, decision-making, and exchange become programmable?

We look forward to exploring this question in next year's research. Assuming, of course, that technology does not outpace expectations yet again.

Survey methodology and demographics

Potloc conducted a comprehensive survey with chief technology officers (CTOs) from June 16 to July 10, utilizing computer-assisted telephone interviewing (CATI) and computer-assisted web interviewing (CAWI). The study engaged 400 CTOs from software companies headquartered across 13 European countries. Quotas were applied to ensure representation by company size and annual turnover, providing balanced insights across different organizational profiles. The average completion time stood at 43 minutes, demonstrating Potloc's commitment to engaging hard-to-reach targets and fostering meaningful interactions.

About Potloc

Potloc is a survey platform used by leading consulting and private equity firms to collect primary insights, back up decisions and boost expertise. With its end-to-end, AI-powered tools and guidance from market research experts, Potloc streamlines the process of creating surveys, selecting samples, analyzing data and visualizing results.

Contact




Charles Gouiran
Account executive
charles.gouiran@potloc.com



Survey sample
(100% = 400)


Breakdown by headquarters location




For selected analyses, respondents were classified as AI-native (n = 84) or incumbent (n = 316) based on product characteristics, AI revenue exposure, autonomy and commercial model. A respondent was classified as AI-native if they met at least three of the following four criteria:

- AI is the core product, decision layer or primary workflow engine
- More than 25% of revenue is linked to AI capabilities
- Production AI exhibits meaningful autonomy
- Core pricing is based on usage, tokens, transactions, workflows or outcomes rather than traditional subscription


Breakdown by ownership type




Breakdown by revenue range and employee count



Breakdown by software vertical



Industries served by respondents



About EY-Parthenon Software Strategy Group

EY-Parthenon Software Strategy Group (SSG) is a specialized advisory practice that provides guidance to companies and investors on growth strategies, operational improvements, product and technology performance optimization, and M&A within the software industry. Software-driven and software-enabled companies and investors look to us for guidance on how to help drive value across commercial, product, technology and R&D strategy.

Making the right decisions about software strategy and investments has never been more important. Many companies, both technology-native and historically non-technology, are now part of the software economy, where software drives most - if not all - of a company's differentiation and value proposition.

SSG serves, among others:

- Investors, including private equity firms
- Publicly listed software companies and software-enabled services businesses
- Non-tech companies that are transforming into software companies

We help companies and investors solve their toughest challenges, including identifying investment opportunities and assessing risks across commercial, product, technology and R&D lenses through due diligence and value creation services. From developing growth strategies to helping to execute major technology changes and helping to improve R&D organizational effectiveness, SSG provides services that help companies to expand market share and enhance their core offering.

EY-Parthenon Software Strategy Group Contacts



Sander van Weele
Partner EY-Parthenon
Partner EY Adviseurs B.V.
sander.van.weele@parthenon.ey.com



Mark Reich
Partner EY-Parthenon
Partner EY Adviseurs B.V.
mark.reich@parthenon.ey.com



Erik Oltmans
Partner EY-Parthenon
Partner EY Adviseurs B.V.
erik.oltmans@parthenon.ey.com



Nick Lush
Partner EY-Parthenon
Partner Ernst & Young LLP
nick.lush@parthenon.ey.com



Per Flaata
Partner EY-Parthenon
Partner Ernst & Young AS
per.flata@parthenon.ey.com

Special thanks to Joanna Adamska, Sebastian Romero, Andreas Kroon, Irene Wang, Pei Kwok and Constantin Yogeshwar for their contributions to this publication.



EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multidisciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

About EY-Parthenon

Our unique combination of transformative strategy, transactions and corporate finance delivers real-world value – solutions that work in practice, not just on paper.

Benefiting from EY's full spectrum of services, we've reimagined strategic consulting to work in a world of increasing complexity. With deep functional and sector expertise, paired with innovative AI-powered technology and an investor mindset, we partner with CEOs, boards, private equity and governments every step of the way – enabling you to shape your future with confidence.

EY-Parthenon is a brand under which a number of EY member firms across the globe provide strategy consulting services. For more information, please visit www.ey.com/parthenon.

© 2026 EYGM Limited.
All Rights Reserved.

ED none
15501126

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, legal or other professional advice. Please refer to your advisors for specific advice.

ey.nl/parthenon