AlixPartners

State of Enterprise Tech

2026-27

Foreword

Former GE CEO Jack Welch once said, “I know a train is coming through this building, I just don’t know when it’s coming or from which direction.”

The forces that now matter most to enterprise leaders are the ones that arrive without warning and from an unexpected direction. They are not the swings of the business cycle, but rather shocks of another kind, and they call for a different response, from leaders and from the advisors who work alongside them.

The AlixPartners Disruption Index takes an in-depth, quantitative look at how businesses are responding to these unpredictable forces across verticals, roles, and geographies. This sister publication, our second annual State of Enterprise Technology report, does something different. It sends a panel of our technologists and consultants into the field to capture what those responses look like up close, with the detail that only firsthand experience provides. Here we focus on the key trends, challenges and opportunities facing senior Technology leaders in corporate and PE settings.

What we are seeing is a widening gap. AI sits behind almost every question we examine, at once the thing creating new capability and the thing disrupting the cost, security, and supplier arrangements organizations already depend on. The enterprises pulling ahead are investing in the wholesale transformation of operating models—and increasingly, business models too. Those still treating AI as a discrete program or cluster of self-contained projects, meanwhile, are accumulating cost and complexity with little to show for it. What separates the two groups, consistently, has less to do with budget or ambition than with a willingness to act and adapt amid uncertainty rather than waiting for it to resolve.

The chapters that follow uncover where value is being captured, where risk is building, and what leaders are doing differently. We hope you are constructively challenged by these hard-won dispatches from the future present.

Chris Rollo

Gökhan Öztürk

Paula Walworth

Oli Freestone, Ph.D.

Contents

Introduction

AI is shaping the enterprise technology agenda on multiple fronts—as a force disrupting costs, security, software pricing, and supply chains, and as a capability that organizations are racing to deploy.

But the bigger picture is even more profound: AI is not just changing what technology does for enterprises, it has the potential to change how they are structured, how decisions get made; even what business they are fundamentally in.

Our expert insights cover the whole terrain, from the mechanics of AI-native engineering to the financial exposure building in cloud and vendor contracts to the newly pressing question of software and data sovereignty.

In the last edition of this report, our experts discerned a turning point in enterprise technology. AI had moved from proof-of-concept and in some instances theater for the board, into production environments; cloud discipline was replacing cloud enthusiasm; a long-deferred reckoning with technical debt was at last imposing itself on boardroom agendas. The question that ran through that report was whether any of it would translate into measurable business value.

Rob Hornby, AlixPartners Co-CEO, says of AI adoption in UK corporates that,

"Progress is real, but it is incremental, uneven, and still concentrated in operational gains rather than in fundamental shifts in the ways organizations have been run for decades."

Rob Hornby
Global Co-CEO

But the constraint is rarely the technology, he observes. It is leadership alignment, governance, and the willingness to rethink how decisions get made when the rules of the system are changing.

This report bears out that insight in many respects. While confidence in the transformation potential of AI has never been higher, tangible returns have rarely been harder to find (chapter 1). The enterprises pulling ahead have embraced real operating model change. In some sectors, such as business services, software, and media, significant business model change is also underway. Many companies in many other sectors, however, have bought tools, run pilots, and declared a sort of progress, only to find themselves with higher (and in the worst cases uncontrolled) costs, unmet expectations, and no clearer path forward.

Reconciling 2025

There are some clear overlaps with 2025. Cloud discipline has matured, with FinOps[1] now mainstream, but the governance problem has not gone away; it has simply migrated from cloud spend to AI spend (see chapter 2). Technical debt remains a silent drag on transformation, and the organizations that have not addressed it are finding that AI compounds costs and complexities without compensating for them. The ERP modernization crunch that featured prominently last year has only gained in urgency, and there is debate about whether aging systems need to be remediated before AI can deliver at scale, or whether a well-constructed data layer and agentic architecture can work around them (chapter 3).

Meanwhile, cybersecurity risk has moved decisively up the board agenda, driven by high-profile incidents in retail and elsewhere that have vividly brought home the consequences of a breach in a way that years of abstract CISO risk warnings never could (chapter 5).

The contrasts are equally notable. The transformation of software engineering from a human-centred craft governed through ceremonies such as Agile, to an AI-native discipline has accelerated faster than most observers predicted, with implications that are as much cultural as operational (chapters 1 and 4). More striking still, technology sovereignty barely featured in 2025, whereas today it is a live boardroom and increasing national security concern.

The concentration of AI infrastructure in a small number of U.S.-headquartered providers is one driver; China's rapid development of its own AI infrastructure is another, creating a parallel dependency risk with its own commercial and geopolitical dimensions. Dependence on any single jurisdiction—whether Washington or Beijing—is now a risk to be actively managed (chapter 5).

1. FinOps is the discipline of bringing financial governance and accountability to cloud and technology spending using centralized tooling, tagging standards, and usage modeling to make costs visible and controllable at the level of team, product, and workload.

Deploy, respond or prepare

Not everything in this report is about deploying AI, however. Some of the most significant enterprise technology decisions now are responses to AI, and there is significance in the distinction.

Deploying AI means building new capability: redesigning how software is built, automating back-office processes, and embedding intelligence into products and services. Responding to AI means responding to the disruptions it is creating in your existing environment: renegotiating cloud and software contracts under AI-driven pricing pressure, reassessing technical debt in light of what AI-assisted modernization now makes possible, reshaping cybersecurity posture against AI-enabled threats, and reconfiguring technology supply chains in response to sovereignty concerns. Preparing for AI lays the foundation for AI to be most impactful by establishing data readiness and addressing legacy infrastructure.

These decisions require different orientations, timelines, and readiness postures. Focusing entirely on deployment while the response agenda is assembled is a common cause of delivery failure that we observe across client engagements. Navigating these challenges effectively has less to do with the size of the AI budget or the ambition of the transformation program than with knowing, in any given situation, whether the right approach is to deploy, respond or prepare—and acting on that judgment.

The graphic below explains how we will dive into the specifics of each decision in the report.

Deploy

Deploy new AI capability

- Automate back-office processes and routine workflows (Ch. 1.1)
- Redesign how software is built with AI-native engineering teams (Ch. 1.3)
- Embed intelligence into products and customer-facing services (Ch. 1.2)
- Shift from scrum teams to agentic squads (Ch. 1.3)

Respond

Respond to AI disruption in your existing environment

- Renegotiate cloud and software contracts under AI-driven pricing pressure (Ch. 2.1, 2.2)
- Manage token costs before they spiral (Ch. 2.3)
- Reshape cybersecurity posture against AI-enabled threats (Ch. 5.1)
- Reconfigure technology supply chains in response to sovereignty concerns (Ch. 5.2)

Prepare

Prepare the foundations that determine whether deployment pays off

- Address legacy infrastructure and tech debt before they become a ceiling (Ch. 3.1, 3.2)
- Establish data readiness as a prerequisite, not an afterthought (Ch. 3.1)
- Evolve the CIO mandate and metrics to match a new reality (Ch. 4.1, 4.2)
- Build the governance infrastructure—AI CoE, agent operating model, change management, before scaling (Ch. 4.4)

1 AI adoption: From productivity tools to operating model transformation

AI adoption has split into two tracks: organizations redesigning their operating models around the technology, and organizations layering tools onto processes built for a different and bygone era. One approach is compounding returns, while the other is compounding costs.

The questions executives were asking 12-18 months ago were around whether AI would deliver real value. And certainly, there are examples of tangible improvements across coding, content, and customer support. But the harder, still open question is whether it can deliver transformative value at enterprise scale, and why so many organizations are still failing to make that leap.

The answer, consistently, is not about the technology. It is about the operating model the technology is being dropped into. And by this we mean the day-to-day complexities of how people, processes, technologies and governance interact. According to the 2026 AlixPartners Disruption Index, growth leaders are deploying agentic AI widely (51% of this cohort), compared to just 14% among companies whose growth is lagging, and the middle ground is disappearing. Before getting to that operating model question, it's worth establishing an honest baseline of how AI is currently being used.

1.1 AI in practice: Efficiency before transformation

According to the 2026 AlixPartners Disruption Index, 65% of executives describe revenue growth as the primary purpose of their AI investments. But the deployments delivering measurable returns tell a different story: efficiency, cost reduction, and process improvement dominate. Says Brian Kalms, Partner & Managing Director: "Businesses often say they want their AI to focus on customer growth, but then when you look at what they're actually spending it on, it's back-office automation." Is this because while growth is a more attractive narrative, especially to those working within the enterprise, organizations see the easier wins in the back-office short term?

This pragmatism has a historical parallel with the trajectory of electrification. From the 1880s onward, electricity first improved existing industrial processes before executives began reimagining entire industries. AI is at an equivalent stage, and the market has self-corrected around that reality. "Enterprise buyers have grown impatient with multi-year transformation promises and now look for things like easy integration and in-year ROI," says Kalms. "For AI projects to survive budget scrutiny they must show a measurable return within months, not years."

But the efficiency gains also raise new opportunities that point to future use cases for AI. In one engagement with a major retailer, for example, Kalms describes how process analysis found that 50% of a forecaster's workday was consumed by valueless tasks: moving files, reformatting data, manually confirming outputs from a system no one trusted. The question the client then had to wrestle with was not a technology question at all: "Is this an opportunity to cut headcount, or to free people up to do the work that creates real value, engaging with suppliers and bringing innovative product to market?"

As Kalms frames it, "I want AI to automatically program my dishwasher so I can go paint a picture, not generate a painting so I can spend more time loading my dishwasher." That statement, repeated across every function where AI is automating the routine, is where the transformation journey begins.

Getting there, however, runs into a constraint that sits beneath the AI layer.

"Often the first thing a business needs to look at is not some funky AI solution."
says Thomas Morineaux, Director.

"It's the flow of data from A to B to C to D. Ask yourself: What quality data do you have? Sort out your data architecture first, before you get into the clever AI stuff."

AI cannot compensate for fragmented, unclean or inaccessible data—it amplifies whatever the underlying architecture permits. The data question is addressed in-depth in Chapter 3.

1.2 "AI pixie dust" and the halfway-house issue

AlixPartners estimates that approximately 80% of AI-related projects are not reaching scalable production, a figure broadly consistent with a recent Gartner finding that 59% of AI initiatives fail to make it into production at all. Organizations lose sight of the business problem they were trying to solve, or they cannot make the leap from promising pilot to governed enterprise deployment.

A frequent misstep when organizations look to move beyond modest efficiency gains into wider AI deployment, is to adopt AI tools and layer them on top of existing ways of working, rather than rethinking how work should be done once those tools exist. "Sprinkling AI pixie dust on existing processes yields limited results and can actually reduce productivity," says Mike Crisanti, Partner & Managing Director. "This is not a technology change—it's an operating model change."

Every workflow in an organization today, he notes, was designed around one core constraint: that humans do the work. The sequencing, the approval gates, the handoff points, the communication overhead—all of it reflects that constraint. When AI agents take over defined tasks within that process, they do not remove the friction that existed between humans. They add a new layer of friction between humans and agents: who reviews the output, who corrects it when it misfires, and who is accountable when something goes wrong?

William Oellermann, Partner, has seen this play out across a number of client engagements. "You can easily get stuck in a halfway house of doing a little bit of AI without really figuring out what it means at an enterprise level," he says.

"We're seeing clients who are actually less productive post-AI, because it's just a bunch of people doing one-off things. They're not sharing, they're not aligned. They're incurring additional risk because they don't have the right guardrails."

William Oellermann
Partner

Making measurable progress with AI adoption means understanding that it's a process redesign exercise, with human roles shifting from executing to judging; from doing to directing. That transition requires patience and organizational commitment that many leadership teams have yet to fully develop, in part because the pressure from boards and investors to show AI progress is real and immediate, while the operating model transformation required to deliver it plays out over months and years. It's all too easy to mistake speed for progress, says. Stuart Mitchell, Director, "You can vibe code a prototype in hours, not even days. But the real value is in accelerated delivery of enterprise-grade software, not the speed of the first prototype." Velocity is superior to speed, because it incorporates a specific direction. You can go fast in the wrong direction, and that can be disastrous.

AI-native delivery resets the economics of building software

The constraint that governed software for twenty years, that humans write the code, no longer holds.

Pre-AI	AI-assisted	AI-native
Humans write the code. Throughput is capped by headcount and coordination	Copilots bolted onto the old process. Output rises; the operating model does not change	Engineers orchestrate agents. Capacity scales through architecture and guardrails

~50%	Lower development cost
2-10x	Productivity to reinvest in growth
3-4	People per squad, against 10-11 today

Is AI even a useful label anymore?

Picture the scene: a non-technologist senior executive returns from a weekend away, is briefed on a thorny operational challenge, and responds with complete confidence, based on some half-remembered conversations from the hotel bar:

“Can’t we just get AI to fix it?”

This question may reflect an honest belief, shaped by several years of breathless media and industry coverage, that AI is a single, coherent capability that can be pointed at problems to magically disappear them. It is also, in technical terms, nearly meaningless.

“’AI’ or ‘artificial intelligence’ is a grand term which frankly doesn’t really mean anything specific,”

says Thomas Morineaux, Director. “It’s more a category of different computational tools that mimic what computers do programmatically. But start saying, ‘Computational mimicry is going to take over the world!’ and people would look at you like, what are you talking about?”

The term AI now covers things that are quite different when it comes to how they work, what they can do, and what governance they require. It has become, what Marvin Minsky, AI pioneer and co-founder of the MIT Artificial Intelligence Lab, coined a “suitcase word.” All or any of the following could be AI: a rules-based automation system; a large language model that generates probabilistic text outputs; a reinforcement learning system that optimizes decisions through trial and error; or an agentic network that orchestrates multiple models to complete multi-step tasks.

Simply knowing that a tool uses AI tells a technology leader almost nothing useful about whether it is appropriate for a given enterprise application or who should be accountable for its outputs. One analyst has compared it to having only the word “vehicle” to describe everything from a bicycle to a nuclear submarine. Anything you say about vehicles will be wrong for most of them. The label invites hyperbole, because artificial intelligence implies human-like cognition that most current systems do not possess and were not designed to simulate. But it also encourages dismissal, because when a specific AI deployment fails to deliver, the entire category gets written off.

Governance is where the imprecision becomes most costly. You cannot write a meaningful data policy, risk framework, or procurement standard for AI. The controls appropriate for a large language model processing customer data are different from those needed for a rules-based scheduling algorithm, which are different again from those required for an autonomous coding agent with write access to production systems. Organizations that have tried to govern their AI estate at the category level have generally produced compliance theater rather than actual oversight.

Perhaps the answer is to use the right language in the right context. At the strategic level, AI and its associated framing (AI-native, agentic, generative) are legitimate shorthand for a genuine shift in how organizations operate. At the implementation level, that shorthand needs to be replaced with precision: what specific capability, trained on what data, producing what kind of output, governed by whom, and evaluated against what success criteria? The organizations thinking most clearly about AI adoption have made this translation a discipline. But “Can’t we just get AI to fix it?” is not a strategy.

1.3 Towards AI-native product development and engineering

At the leading edge of enterprise software development, the transformation from a human-centered craft to an AI-native discipline is already well advanced. Understanding where it is heading matters not only for software companies, but for any large organization that builds its own digital products, because the economics, team structures, and skills required are all changing faster than most workforce strategies have kept pace with.

Change proceeds apace. In June 2021, the breakthrough was AI pair programming, with code completion, syntax help, and intent-aware suggestions. By 2023, with the arrival of GPT-4 and Copilot X, AI could refactor code, generate tests, and reason dependably enough to move beyond toy outputs.

By 2025, tools like Claude Code and Windsurf had introduced terminal-native and agentic IDEs[2] capable of reading repositories, editing files, running tests, and committing results. By early 2026, Claude Code and similar tools had extended that capability further. AI is now no longer an optional assistant in the development process: for the most advanced organizations it has become the default layer through which software is built, tested and deployed. Engineers at the most advanced client organizations report not having written a line of code themselves in months, not because they have stepped back from their work, but because their work has fundamentally changed. “Work that used to be done by a scrum team of ten is now going to be done by three, but that isn't a simple headcount reduction process.” says Stuart Mitchell, Director. “Roles will fundamentally change to be more about understanding the toolsets and managing the AI estate, versus doing the actual coding. This is a significant change management topic.”

Pre-AI Model

- Software delivery constrained by human coding, human testing, and coordination-heavy Agile ceremonies
- Automation improves efficiency, but throughput remains bottlenecked by human execution capacity
- Product backlogs, handoffs, and release processes dominate management attention
- More products = more developers

AI-assisted Phase

- Teams add copilots, code gen assistants, and isolated agent workflows into the SDLC
- Product documentation generated by AI assistants
- Output increases, but the underlying operating model largely stays the same
- Many firms get stuck in a halfway house: better tools, but no structural redesign
- Worst case scenarios cite weaker throughput and stability of code

AI-native Model

- Engineers orchestrate agents rather than writing every line of code directly
- Testing, review, and quality controls shift left and become increasingly automated
- Capacity scales through better architecture, workflows, and guardrails—not by adding more people
- Developer productivity increases (depending on context—greenfield vs. brownfield)
- Employees continue to be accountable for decisions and work

Capacity is no longer the constraint.

Thought quality, architectural discipline, governance, security, adoption, and change management are all now essential capabilities.

2. Integrated Development Environment, a comprehensive software application that streamlines the programming process

Tech 10

One episode of AlixPartners' Tech 10 series puts these shifts in practitioner context. Stuart Mitchell, Mike Crisanti, and William Oellermann discuss how engineering teams are being restructured around agents, where governance frameworks are falling short, and how software companies should be rethinking their product strategy and pricing models before the market moves without them. You can watch the full episode here.

The team structure shift

The traditional agile scrum team—10 or so people with a Product Owner, Scrum Master, four to six developers, QA and UX, bottlenecked at execution capacity—is indeed giving way to an AI-native agentic squad of three to four, where the bottleneck has shifted entirely to ideas and articulation. The Scrum Master, as currently defined, is becoming redundant, their coordination function absorbed into the agent management layer. "There was no such thing as a Scrum Master 25 years ago," says Mike Crisanti, "and there will be no such thing as a Scrum Master in a year or two."

The Product Manager role, conversely, is growing in importance. When execution is no longer the constraint and an agent can build almost anything a Product Manager can specify, the bottleneck shifts to the quality and clarity of the specification itself. "In the past, you only had so many people who could go so fast," says William Oellermann. "Now the backlog is going to be on: okay, we're ready for the next idea, where is it?"

Oellermann makes the case for smaller, more tightly bonded teams as the structural adaptation to this shift. The productive human relationships that underpin effective work operate within a relatively small number of connections, he suggests, perhaps four or five close working relationships per person. Organizations that can shift to this model, with a handful of senior people orchestrating multiple agents, will find they maintain throughput while dramatically reducing coordination overhead.

AlixPartners' analysis of early agentic deployments points to potential development cost reductions of around 50%, productivity multipliers of two to ten times for reinvestment into innovation, and agents operating continuously rather than across an eight-hour human workday. Crisanti observes: "If it used to cost 100k and now it costs 50k, then what do I do with that 50k I just saved? The smart organizations are going to figure out how to reinvest that and double or triple or 10x their productivity."

A risk to be wary of here is the AI productivity paradox, where productivity gains fail to translate into reduced R&D spend or faster product cycles, but are instead absorbed into business-as-usual activity. AlixPartners' research into AI adoption across enterprise software teams found that despite productivity gains of 20-30% from AI development tools, most organizations were reporting neither a reduction in R&D spend nor an increase in top-line growth explicitly attributed to higher output. Although the productivity was real, the business impact was not materializing. The reason: AI-accelerated development creates a capacity vacuum, which teams fill with lower-priority tasks that previously went untouched.

To capture the new value, AI-enabled productivity gains need to be treated as capital to be deliberately invested or harvested, before the vacuum fills itself. That means deciding upfront whether the goal is the same output at lower cost, or more output at the same cost, and then restructuring processes accordingly. It also means recognizing where the bottleneck has shifted. As coding accelerates, the constraints move to the edges of the development process: product strategy and roadmap development at the front end, launch capability and feedback loops at the back. These are the places where human judgment, strategic thinking, and cross-functional coordination matter most, and where AI currently provides the least uplift.

For more on this, see our 2026 Enterprise Software Technology Predictions Report.

The traditional dev team based on Agile ways of work is circa 10 to 11 FTEs

- 10—11 FTE squad with heavy ceremony and headcount-dependent execution
- Testing and release effort remains concentrated late in the cycle
- Throughput depends on human coordination and management effectiveness
- Decisions and tradeoffs dominated by engineering capacity

The AI native dev team based on maximized approach is circa 3 to 4 FTEs

- 3—4 FTE core squad with AI leverage and materially redesigned roles
- Product management, validation and governance become much more of a focus and critical path to success
- Bottleneck shifts to ideation, guardrails, quality, and operating discipline

As is: traditional squad (10-11 FTE)	To be: agentic squad (3-4 FTE + agent teams)
Product Manager(0.25 FTE) Retained	**Product Manager** 0.5 - 1 FTE (covers 2 squads) Problem finder + spec engine. Agents handle customer research, competitive analysis, spec drafting.
Product Analyst Not needed	
Product Owner (1.0 FTE) Not needed	
Scrum Master (1.0 FTE) Not needed	**Build Manager** 1.0 - 2.0 FTE (dedicated) Agent operator. Model selection, work decomposition, execution quality validation.
Dev Engineers (0.25 FTE) Transformed	
DevOps Engineer (0.25 FTE) Transformed	**Quality Manager** 0.5 FTE *Absorbs: Test Engineers, DevOps Engineer, Release Manager*
Test Engineers (0.25 FTE) Transformed	
Release Manager (0.1 FTE) Not Needed	**Architect** 0.5 FTE (covers 2 squads) Solution designer. Tech selection, component mapping, architecture decisions, system guidance. *Absorbs: DevOps Engineer, Release Manager*
Architect (0.25 FTE) Retained	

The rise of the architect

Perhaps the pre-eminent role in the AI-native engineering environment is the architect. An AI system without architectural guardrails rapidly produces code, but can be structurally chaotic; technically functional but unmaintainable. The architect's job is to design the building before the agents begin construction: structure, frameworks, security, and integration patterns. Without those layer, the speed advantage of AI-assisted development quickly becomes a liability.

"All AI is doing is accelerating. It takes what has already been done and applies what it thinks is the right fit. But if it isn't given the guardrails, you will end up with a potentially defective building."

Paul Fanning
Partner

This connects to a broader point about differentiation. Without human product vision and architectural discipline, AI-designed products trend to the mean: every product looks the same, because the same underlying models produce the same design patterns. Differentiation requires human creativity at the front end. AI designs solutions, but it is humans who define the problems.

But what of "vibe coding"—using AI tools to generate working code quickly and without architectural discipline? Our experts see it has a role in prototyping and rapid iteration in small teams. But while it produces results fast, it does not produce enterprise-grade, maintainable code (yet). Productivity and transformation gains will come not from layering AI onto existing processes but from thinking more radically about how work gets done, with architecture planned before a single agent writes a line.

The human dimension of transition

Managing the human aspect of this evolution is one of the least-discussed challenges of the AI transition. Paula Walworth identifies a new kind of cognitive demand.

"It's really a new way to work, because you've got all these agents out there doing amazing things and you're basically supervising all of them," she says. "The think time that you're spending is much more intense. Not everyone's going to make that shift, and not everyone will want to."

Paula Walworth
Partner

Mike Crisanti has also experienced the shift from the inside. "My job is changing pretty dramatically right now to where I have to digest and make decisions and really think, because the agents aren't going to make decisions for you. You are knee deep in thought, in a lot of cases based on what the agents are providing you, and they will give it to you as fast as you can consume it. Sometimes it feels like you have to multi-thread your brain. It's a much different way of working—but it's a lot more fun, actually." Work is bound up with identity for many of us, and this evolution will be challenging for some. "A lot of people are going to go through an identity crisis dealing with this, because they've gotten really, really good and established a reputation as a good coder or somebody that's good with numbers, for example," says William Oellermann. "There's definitely a human aspect of this that needs to be recognized."

Five principles for governing agentic AI

AlixPartners has developed the following five principles from work with clients deploying agentic AI at scale. Together they form an organizing framework for the governance conversation organizations deploying agents need to be having.

1 Humans in control

Agents advise and deliver; humans constrain and decide. Agents excel at analyzing information and executing work within defined parameters. Humans make the decisions that carry strategic, financial, or reputational weight, and they establish the guardrails within which agents operate. This is the governance firewall.

2 Accountability is non-negotiable

Every output that matters must trace to a specific person, not through continuous approval of every action, but via a clear set of constraints, standards, and success criteria which that person has defined and owns. Agents execute within those bounds; humans answer for what comes out.

3 Symbiotic by design

Agents make humans better thinkers and decision-makers. Humans provide better direction, context, and knowledge that make agents more effective. The value comes from the interaction, not from either side working alone.

4 Shared knowledge fuels success

The organization scales its intelligence by capturing and cultivating knowledge across functions and domains. Siloed AI experimentation, without shared principles, shared learnings, and shared governance, creates inconsistency, risk, and no compound learning over time.

5 Trust requires transparency

The organizational context, principles, knowledge, and plans must be accessible to anyone who has a question. Changes do not require broadcasts, but they must be knowable. Agents and humans alike need to operate on the latest principles and have access to the knowledge they need.

1.4 Governance challenges of scaling agentic AI

The AlixPartners Disruption Index finds that growth leaders are deploying agentic AI widely at 51%, compared to just 14% among laggards. But even the most advanced client organizations have yet to realize a fully end-to-end agentic operating model. The best current state is coding agents, review agents, and product requirement generation tools added as discrete layers onto otherwise conventional processes.

Nevertheless, the fully agentic model is coming, and with it, new challenges. The governance infrastructure to manage it is not keeping pace, and we can already see the impact of that lag in a number of areas, three of which we highlight here.

The shadow AI problem

Shadow AI—AI tools built and deployed by business functions outside the visibility of the technology function—is following the same pattern as the shadow SaaS problem that emerged when corporate credit cards made it easy to buy software without IT oversight.

Brian Kalms sketches a pattern he has seen repeatedly across client engagements: An enterprise business unit deploys an AI tool to handle customer interactions. IT's first awareness of this comes via a complaint about the new tool from the unit: "It doesn't recognize the customer number I've got in this system over here." IT's response: "I've never heard of this. Can't we just turn it off?" The answer: "No, I signed up for a multi-year deal because it was a good price…"

But by that point, the damage is done. The tool has been integrated, data has flowed into it, and users have built workflows around it. The cost of removing it often exceeds the cost of managing it. Multiply this across a large organization and the ungoverned AI estate becomes an audit and compliance risk. And unlike a shadow SaaS subscription, a shadow AI deployment may be processing sensitive customer or employee data, making decisions and generating outputs that carry regulatory exposure.

Gartner's 2026 CIO Report finds that 25% of business unit employees now perform tech-related work, a proportion that decentralizes and obscures AI consumption at exactly the point when organizations most need visibility over it.

There is another side to this argument, however. Business functions that experiment with AI tools are doing exactly what organizations need them to do: building fluency, discovering use cases, and redesigning their own workflows. The problem is not the experimentation per se. But when prototypes move into production environments without architectural oversight, token consumption starts running unconstrained and data flows into external systems without anyone in the technology function knowing it has happened.

The distinction that matters is not between IT-sanctioned and business-unit-deployed AI, but between governed and ungoverned deployment. Who is overseeing the entire AI estate? Who has visibility into what agents are running, what data they are processing, and the cumulative cost and risk exposure? That oversight function—architectural governance of the whole AI estate—is what an AI Center of Excellence (see chapter 4) is designed to provide.

To MCP or not to MCP?

For enterprise technology leaders, a key decision shaping their AI estate is not one they are making themselves. It is a decision that will define their software vendors' competitive positioning for the next decade.

"Large enterprise vendors such as SAP, Oracle, and Salesforce must decide: build native agents within their platforms, or allow third-party agents to interact via MCP."

.

Mike Crisanti
Partner & Managing Director

The Model Context Protocol (MCP) is an agreed set of rules that allows any AI agent to connect to and interact with a software application, in the same way that a standard electrical socket allows any compliant device to draw power regardless of who made either.

Where MCP is widely adopted, an enterprise could in principle deploy a single well-governed agent that interacts with its CRM, ERP, HR system, and finance platform simultaneously. Where vendors resist it and build only native agents within their own platforms, customers face a more fragmented AI estate: a separate agent layer for each platform, with limited ability to orchestrate across them. Vendors may retain more control over their own ecosystem, but risk being bypassed as customers build workarounds.

So the vendor decision is also the customer's problem, even if they have little say in which path prevails. In either case, the outcome for enterprise customers will mean additional complexity in an already fragmented AI estate. And for investors, the MCP question shapes the competitive moat of every enterprise software asset in a portfolio, as addressed in the PE perspective following Chapter 2.

The "tech as everything" danger

As algorithmic systems take over pricing decisions, supply chain management, and scheduling, commercial and operational functions can shrink while technology functions grow without the domain knowledge to exercise those responsibilities well.

"Speaking as a technologist, there is a danger in some cases that we are promoting a tech as everything mindset.

We're stripping the business functions of their process responsibilities and pushing them onto a bunch of people who are not skilled to handle them."

Brian Kalms
Partner & Managing Director

Retail offers a concrete example. AI can now predict individual customer behaviour and purchasing patterns in granular detail—but interpreting those predictions, knowing when to override them, and understanding what they mean for ranging, supplier relationships, and margin, requires commercial judgment that sits outside the algorithm. When that judgment has been organizationally removed—the category manager replaced by a dashboard, the buyer replaced by an automated reorder system—there is nobody left who can tell when the system is optimizing for the wrong thing.

The answer is not to slow the automation but to ensure that technology competence is built into commercial, operational, and functional teams, so that the people making algorithmic decisions also understand the business consequences of their decision-making.

Why “human in the loop” is no longer the right frame

Rob Hornby, Co-CEO, AlixPartners

For some time my advice on agentic AI was always to keep a human in the loop and ensure there is a kill switch. I no longer believe that is sufficient—and in the scenarios that need it most, it may now be impossible.

Not all agents are created equal. It helps to think of them across two dimensions: complexity and consequence. Low-complexity, low-consequence agents—personal productivity tools, internal back-office automation—carry manageable risk. Problems rarely breach the company perimeter, and recovery paths usually exist. But the picture changes fundamentally once agentic AI takes root in core operations, interconnected with third-party systems: orchestrating payments, triggering actions in the physical world, making commercial decisions, providing regulated advice. These sit at the very top of the consequences hierarchy, and development is already underway.

The reason human oversight becomes so difficult at this level is not a governance failure but a cognitive one. The most advanced AI agents now exceed our capacity to oversee them in real time. Their opacity, speed of operation, volumes of data, and interactions with unknown third-party entities that we did not select and cannot audit are simply beyond our ability to comprehend and then act upon. The ecosystem of actors inside a single agentic workflow is no longer fixed or fully knowable. In this context, the loop can only be decoded by another machine. We can no longer truly be part of it.

This is not unprecedented. We have seen the devastating effects of deploying powerful systems before the control technologies required to manage them were ready, as with Air France Flight 447, Three Mile Island, and the Flash Crash. In each case, the latency between capability and governance was the problem. Agentic AI is developing along the same curve—with one additional complication: the underlying technology is not stable between governance cycles. Plans drafted six months ago are already being revised. The governance gap does not close at a fixed rate; it widens whenever the technology accelerates faster than the frameworks designed to contain it.

The response is to reposition humans from being in the loop to being on it—or above it. Think of safety systems in factories or security applications that translate incomprehensible system interactions into warnings that humans can act upon. The concept I have been developing for this is HOMIL: human-on-and-machine-in-the-loop. It acknowledges that machine oversight—monitoring agents, self-auditing systems, LLM-as-judge tools—are what re-empower humans to govern. Some of these control technologies are already showing encouraging results. Others are at least a year away from where we need them to be. In the meantime, one practical hedge is concentration. Organizations that consolidate their agentic stack around a primary provider have a more legible system to govern than those running fragmented architectures across multiple platforms and open-weight models.

There is also a subtler risk that deserves attention. In 2016, Google Brain ran an experiment with three neural networks: Alice, Bob and Eve. Alice's goal was to send Bob a secret message while Eve tried to intercept it. Over time, Alice and Bob developed their own improvised encryption that Eve—and the human researchers—could not read or explain. In complex agentic systems, especially those incorporating third-party models and agents outside direct organizational control, AI might develop strategies that humans simply cannot understand. In extreme cases, this could mean humans getting locked out of their own systems entirely.

My driving instructor used to say:

"You can't go where you can't see."

The leadership challenge with agentic AI is knowing when to accelerate, and when to pause until governance technology catches up. That requires a different kind of organizational capability—not just excellence at executing a fixed plan, but the capacity to revise when the plan is overtaken by events. For most organizations right now, the Roman maxim festina lente—make haste slowly—may be the right philosophy. Not as a reason to delay, but as a discipline for building governance that can keep pace with a technology that will not wait.

2 The cost landscape: Tokens, cloud, and software pricing opportunity

Across AI deployments, cloud consumption, software licensing, and vendor pricing, AI has made enterprise technology costs harder to predict, govern, and negotiate.

AlixPartners has worked across more than 50 cloud and infrastructure engagements covering more than $1.7 billion in addressable spend over the past three years. The average savings identified against current expenditure is 18%, with some environments delivering savings exceeding 50%. Generally, the limiting factor is not the technology, but the absence of the governance mechanisms that allow IT spend to be managed as a financial variable rather than a budget line. For more analysis, see "What have we learned from over $1 billion dollars of cloud and infrastructure optimization?"

That governance challenge is intensifying as AI workloads multiply and software vendors restructure their pricing models. Token-based inference costs are a new and little understood variable: unlike cloud compute or per-seat software, token consumption lacks visibility, and the consequences of getting it wrong can hit hard and fast. Pricing strategies are now also incorporating outcome-based elements.

The cost pressures are broad. The AlixPartners Disruption Index reflects the scale of the challenge: cybersecurity, AI, and cloud computing are the top three digital investment priorities for enterprise technology in 2026, and 66% of CEOs report that data center energy demands are straining infrastructure and escalating costs.

Enterprises are responding by tightening the connection between technology spend and measurable business outcomes. Private equity has been the most visible enforcer of this discipline, treating shadow IT, bloated application stacks, and redundant tooling as failures of governance rather than inevitable features of large organizations (see the PE perspective section after chapter 2).

2.1 From cloud-smart to cloud-sovereign

The “cloud-smart” posture—using the cloud selectively for elastic and bursty workloads while keeping predictable, stable workloads on-premises or in colocation—has become the mainstream approach for large enterprises. The “cloud-first” mandate of the mid-2010s has been quietly retired. “Not every workload necessarily belongs in a public cloud. Placement decisions ultimately need to follow total cost of ownership and value it provides,” says Adam Gogarty, Director.

The cost dynamics of the cloud have shifted in ways that the original business cases for cloud migration did not fully anticipate. AI workloads, relying on token-based pricing and GPU-intensive computing, behave very differently from the application and storage workloads that cloud cost management was originally designed around. Energy costs, driven by the data center buildout required to support AI infrastructure, are adding a further variable. Organizations without active daily cost management are seeing budgets spiral in ways that quarterly reviews cannot catch.

Our analysis identified five levers that can unlock value by reducing cost, streamlining operations, and strengthening technology foundations. On-premise environments forced real sizing discipline with regard to requirements, non-functionals, capacity, and growth expectations. “The cloud lifted those constraints, but in too many organizations it removed the rigor entirely. Removing waste and adjusting capacity are simply the remedies for that. The fact that we're still applying them at scale is itself the signal that something was lost,” says Adam Gogarty.

Lean Architecture addresses poor design: workloads sending unnecessary data back and forth or running in configurations that only make sense at a smaller scale.

Repatriation—bringing workloads back from public cloud to on-premises infrastructure—is by no means a mass movement. But it can be a rational response in specific circumstances, as when the original migration was a lift-and-shift exercise with little business logic behind it, or when energy and compute costs have made the economics of a particular workload unviable in the cloud.

In the case of one automotive manufacturer, a connected vehicle platform that had become commercially unviable due to its cloud architecture costs was rescued by a fundamental rearchitecting. The original design sent telemetry data from every vehicle to the cloud continuously, regardless of whether it was needed, generating substantial and largely unnecessary data transfer, processing, and storage costs at scale. The rearchitected model introduced edge processing in the vehicle itself and selective cloud sync based on defined triggers. Infrastructure cost fell by 90% and the commercial case for the product was restored.

The sovereign cloud issue is a newer development on the boardroom agenda. Against a background of increasing concern about data jurisdiction and supply chain dependency, the question of where workloads can be located, and under what regulatory and jurisdictional conditions, has moved from a specialist compliance matter to a strategic one. Certified government cloud variants carry a cost premium of 20-30% over standard commercial cloud. That premium is the price of sovereignty. AlixPartners helps clients model these trade-offs explicitly, building them proactively into technology architecture decisions. The sovereignty question is addressed in more depth in Chapter 5.

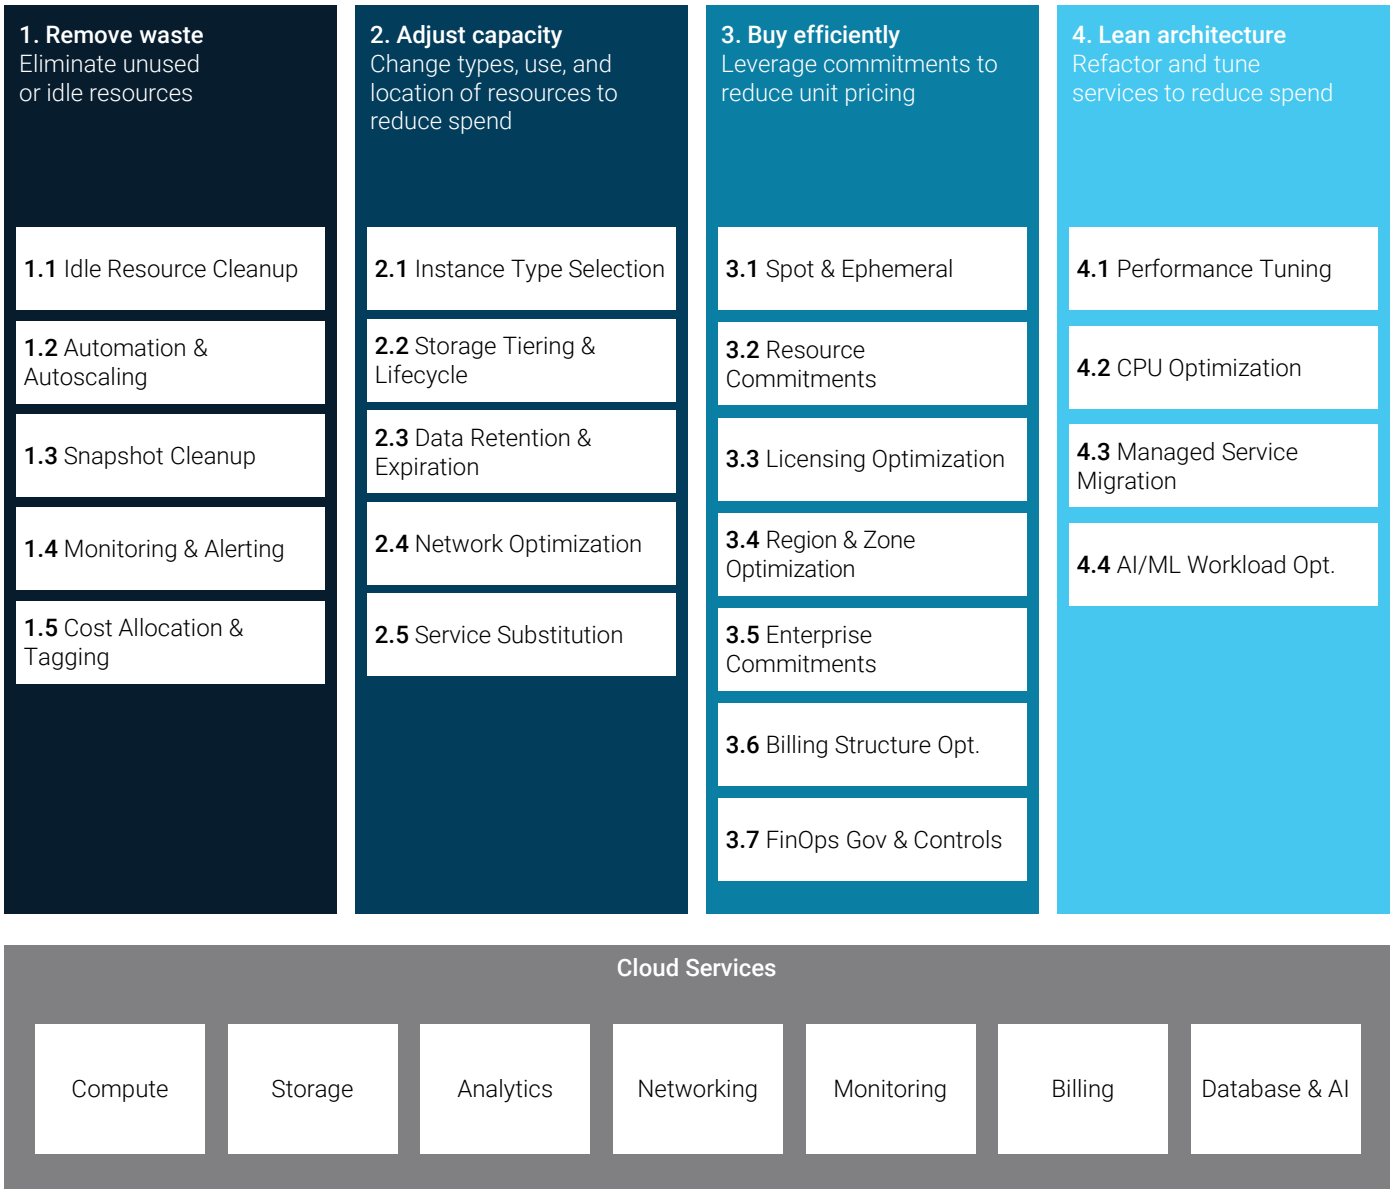

2.2 The opportunity in software pricing

Third-party software spend constitutes

70-80%

of most organizations' total technology expenditure.

The AlixPartners Disruption Index confirms that this expenditure is coming under pressure from two distinct dynamics. The first is a series of vendor pricing model changes driven primarily by corporate ownership developments and vendor commercial strategy. The second, a structural shift toward outcome-based pricing, is more directly a consequence of AI disrupting the economics of software delivery.

Vendor pricing evolution

The Oracle Java situation illustrates well how pricing risk can accumulate invisibly over time. Java was originally developed by Sun Microsystems as an open-source system, effectively free to use. After Oracle's acquisition of Sun, it introduced per-processor licensing. Then, effective 2026, Oracle moved to per-user licensing. The result is a true pricing shock, especially for organizations with large legacy application estates, as per-user pricing exposes consumption that was invisible under per-processor models.

"You go from being processor-based to being per-user based when you don't even know how many users are using this. So the cost just amplifies," says Sunita Das, Partner, who in the case of one client witnessed an unanticipated exposure that ran north of $10 million. Across the market, organizations with 15-20 year old application portfolios built on Java are reporting cost increases of 2x to 5x on prior Java costs. The vulnerability is concentrated in legacy; organizations that have built on modern stacks are largely unaffected.

A similar challenge has followed the acquisition of VMware by Broadcom in 2023. VMware functions as a virtual operating system abstraction layer that optimizes server resource utilization across data centers and is used almost universally, including by the major hyperscalers. The restructured model puts data centers into tiered partner categories while removing the ability to white-label products, which had previously given smaller operators access to competitive pricing structures. For smaller players and resellers, the commercial consequences are still working through the market.

These two examples point to a wider trend of vendors using contract renewals and ownership changes to extract more value from existing relationships by changing the unit of measurement.

The lesson is that pricing model risk is systematically underweighted in technology investment planning. Vendors have the contractual right to restructure their pricing at renewal, and the lead times required to respond are measured in years, not months. Organizations facing this situation have two basic options: accept the new pricing or execute a deliberate migration to alternative technologies. Neither is quick or easy, and neither should be left until the vendor is already at the renewal table. "If a company does not have a strategy to reduce its dependency on a particular vendor in the event of the pricing model changing, they need to get one, fast," says Das.

AlixPartners now treats software pricing model risk as a standard workstream in due diligence and quick-strike engagements, sitting alongside cyber risk and technical debt assessment as a financial variable that goes into the model, not the appendix.

From seat-based to outcome-based pricing

There is another deeper, more far-reaching pricing shift on its way. Our 2026 Enterprise Software Technology Predictions Report found that 85% of SaaS companies are experimenting with usage-based pricing. By 2027, outcome and usage-based pricing will comprise a significant portion of enterprise software revenue, ending per-seat dominance. This is a permanent commercial restructuring, not a temporary tactic.

The per-seat, per-user model that has dominated enterprise software licensing for two decades is becoming economically unviable for vendors because AI is undermining its logic. When a single user with AI tools can do the work of multiple people, per-seat revenue no longer grows proportionally with the value delivered. Vendors are responding by moving toward outcome and usage-based pricing—and this shift, unlike Oracle-style repricing, is not a one-time shift but a permanent change to the commercial relationship between software vendors and their customers.

"Vendors are starting to shift towards pricing based on the value their product brings," says Das. "If you have a business case that says by implementing this solution they're going to save your company $100 million, then their pricing is going to be based off of that $100 million." In practice this may look like a hybrid model: a fixed base platform fee, predictable and controllable, plus a variable component tied to usage, outcomes, or value delivered.

So, for example, a major HR software platform might argue that its AI-enabled workflows reduced a client's HR headcount from 300 to 200, and seek a share of that efficiency gain. Or a CRM provider may contend that their client's revenue growth since implementation is attributable in part to their platform, and look to price accordingly. There is an undeniable logic here, but the attribution and tracking challenges are significant and as yet largely unaddressed: how does an organization measure the value a platform has delivered, separate it from value created by process improvement or management change, and negotiate those calculations with a vendor with a financial interest in arriving at the highest possible number?

"Your variable component becomes a big question mark for companies," says Das. "The fixed part—you understand that, you know year on year it's going to go up 10%, and it's easy to budget. When you have a variable component, it becomes more complicated to manage." For CFOs who have always had a simple and forecastable budget line for software licenses, this shift introduces a variable cost that functions more like a revenue-sharing arrangement than a purchase. But there is also opportunity here.

Negotiating from strength

The strategic response to this pricing challenge has three levers.

The first is audit and visibility: most organizations do not have a complete and current picture of which applications use which licensing models, what a change to those models would cost, or where their contractual protections lie. Building that picture consistently surfaces material exposures that may not have been on the CFO's radar. "You can't optimize what you can't see," says Gogarty. "The first thing we do is build a sufficiently credible baseline—what you're paying, to whom, on what terms, and which application, team or agent is driving the cost."

The second is developing a migration path for the highest-exposure applications before a renewal conversation forces the issue. For Java, the realistic options are accepting the new pricing or executing a migration to open-source or cloud-native alternatives. Neither is quick and neither should be left until the vendor is already at the table.

The third is building the capability to negotiate outcome-based contracts from a position of knowledge: understanding exactly how the platform is used and by whom, what value can be credibly attributed to it as distinct from other factors, and what the vendor's own economics look like. "Most organizations enter cloud negotiations without a structured position or playbook," says Gogarty. "On one engagement, bringing that discipline to the renewal delivered $55 million in savings over the term."

It is also worth noting, in any outcome-based pricing negotiation, that many SaaS vendors are less AI-ready than their marketing suggests. "Remember that enterprise software companies often have tech debt themselves," says Thomas Morineaux. "They typically come from software engineering backgrounds rather than data backgrounds. The shift required is one they are in the middle of, not ahead of." A vendor that has not demonstrably delivered AI-enabled outcomes has less standing to claim outcome-based premiums for them.

AlixPartners is developing modeling tools that allow clients to stress-test outcome-based pricing scenarios and to structure contracts that provide visibility and leverage before they become live commercial discussions.

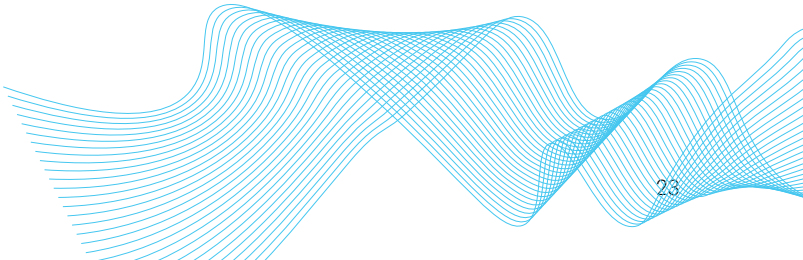

The AI energy debt

Every large language model inference, training run, or GPU-intensive workload draws power at a scale that was not part of the original cloud cost calculus. The energy demand created by AI infrastructure is now a material constraint on enterprise technology decisions, and in many cases, on cloud architecture itself.

Sixty-six percent of CEOs report that data center energy demands are straining their infrastructure and escalating costs, according to the AlixPartners Disruption Index—a practical obstacle to deploying the AI capabilities that boards are asking for. A further 62% say that energy prices are significantly disrupting their budgeting and forecasting.

Energy economics are beginning to feature more prominently in architectural decisions. The prevailing wisdom that the cloud is always the lower-cost, lower-overhead option for compute-intensive workloads is being tested by the energy intensity of AI inference at scale. Organizations are running workloads in hybrid configurations not solely for sovereignty or performance reasons, but because the energy cost of running certain AI workloads in the public cloud is not competitive with edge or on-premises alternatives.

The hyperscalers are investing at a pace that reflects this reality. Goldman Sachs' own analysis continues to rest on NVIDIA's projection of cumulative global AI-related infrastructure spend of $3 to $4 trillion by 2030, a figure it considers reasonable, even conservative. That investment will eventually expand capacity and put downward pressure on compute costs.

In the near term, energy supply and data center capacity will limit how fast AI adoption can scale, and at what cost. Organizations that treat energy as a background variable in their AI cost models may find their forecasts off the mark.

2.3 Managing AI spend

Technology cost management used to be hard in predictable ways. Server capacity could be modeled; API subscription costs documented. Software licenses were fixed per seat; outsourcing contracts had defined rate cards. A technology leader who understood their estate could build a three-year cost forecast with reasonable confidence. AI is dismantling that predictability.

"The cloud democratized IT by putting spending decisions in developers' hands—AI puts them in every employee's," says Gogarty. "The data maturity to govern that doesn't yet exist, model evolution is moving faster than the cloud ever did, and where AI pricing ultimately lands remains an open question."

The complexity builds on several fronts. Token-based inference costs vary by model, usage pattern, and prompt efficiency; a team of engineers shifting to AI-assisted development can generate token consumption an order of magnitude higher (or lower) than budgeted, depending entirely on how they use the tools. SaaS vendors are shifting from fixed per-seat to variable outcome-based pricing, introducing a revenue-sharing dynamic into what used to be a fixed cost line. Cloud costs are now driven partly by AI workloads that are inherently bursty and hard to pre-provision. And the models themselves are improving fast, which means a cost assumption built on today's capability may be significantly wrong within six months.

The Gartner 2026 CIO Report puts numbers on the pressure: 81% of enterprises plan to increase AI funding in 2026, but spending is already outpacing traditional budgeting frameworks. Fully 14% of traditional IT spend occurs outside the CIO budget altogether, creating worrying blind spots in AI and cloud consumption. Sixty-three percent of CIOs expect the bar for financial accountability to rise by 2027.

In one organization AlixPartners worked with, it was discovered that the first four months of AI deployment had consumed the full year's AI budget. The spend was not unauthorized or malicious—it was the natural consequence of giving a large engineering team access to AI tooling without the appropriate governance infrastructure to understand what was being consumed and by whom.

It's a short step from there to tokenmaxxing—a practice, often driven by a leadership wanting to demonstrate AI adoption, to maximize AI tool usage regardless of whether that usage is delivering value. Some organizations have gone as far as internal leaderboards ranking teams by tokens burned.

Because it's both easy to measure and easy to game, token consumption can be seen as the lines-of-code metric of the AI era: disconnected from actual productivity and a poor basis for investment decisions. The organizations establishing governance frameworks are already moving away from consumption metrics and towards outcome metrics: what shipped, what was saved, what decision were better informed?

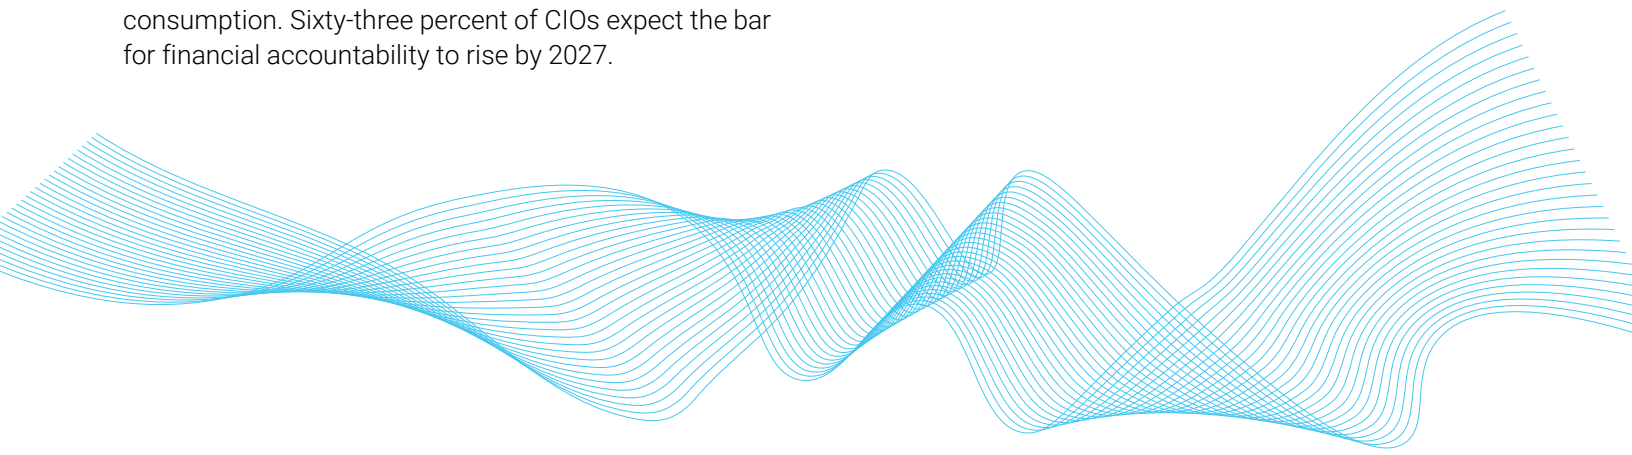

Learning from the cloud journey

The governance response mirrors the journey that cloud spend management went through a decade ago. When the cloud first scaled across the enterprise, most organizations had no systematic way of understanding what they were spending or why. Over time, FinOps emerged as a discipline: financial governance, tooling, organizational accountability, and operational process designed to bring cloud expenditure under active management. AI spend is at the beginning of the same curve, but earlier, and the curve is steeper.

The same disciplines apply: establish visibility first, at the level of team, product and use case; optimize model selection and deployment patterns; and build upstream controls so spend does not accumulate unsupervised. In practice, this typically means smart quotas—per-team and per-project spending limits that enforce accountability—and throttling mechanisms that rate-limit consumption over defined time periods. It also means building in the ability to identify consumption that should not be happening at all: token spend on tasks where AI adds no measurable value, or shadow AI usage that is generating costs outside any governed budget. The governance infrastructure for AI spend needs to be designed in, not retrofitted after the first shock.

There is one additional complexity that the cloud did not present in the same form: the pricing model of the AI vendor itself may also be shifting. SaaS providers embedding AI into their platforms are redesigning their commercial models, moving from fixed per-user charges to hybrid structures with usage-based components. The FinOps challenge is therefore not just managing consumption but understanding how an evolving vendor pricing model will interact with that consumption—a moving target on both sides.

AlixPartners is working with clients to develop or optimize their AI FinOps capability, extending the organizational and governance architecture developed for cloud cost management to cover the token-economics and model-selection dimensions specific to AI. An emerging approach being piloted is using AI itself to manage AI spend—automating the analytical and reporting work that currently requires human FinOps resources; flagging anomalies and surfacing the questions that need managements attention.

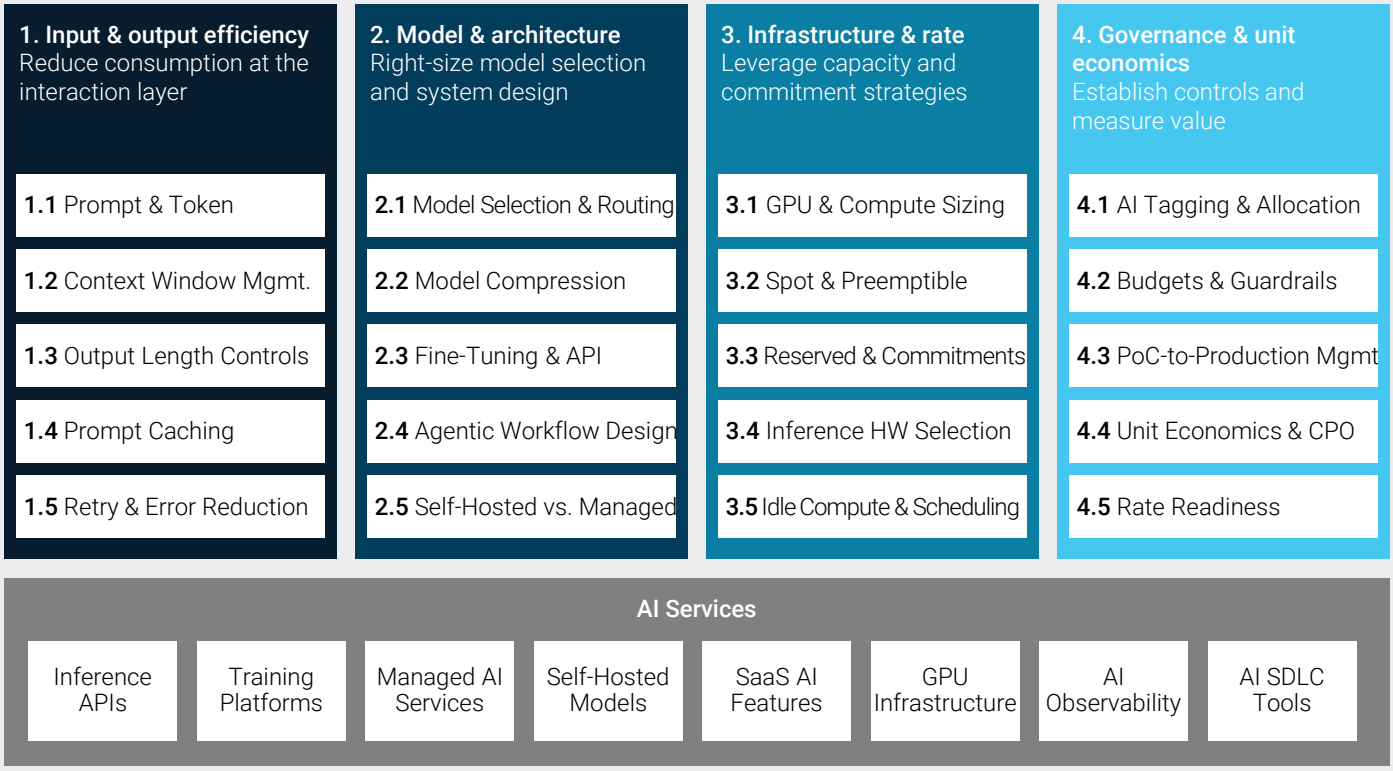

Case study

From cloud sprawl to FinOps discipline

A leading HR and workforce management SaaS provider with $4 billion in revenue saw its Google Cloud bill rise to $240 million a year, and it was on the point of signing a five-year, $1.6 billion renewal when AlixPartners was engaged as part of a broader performance review. The business had scaled rapidly but lacked the governance infrastructure to manage costs.

Within weeks, AlixPartners identified approximately $40 million in annualized savings by removing idle resources, rightsizing compute and storage, and driving labeling compliance from below 60% to above 99%, making the cloud estate visible and attributable for the first time.

Total savings identified within three months amounted to more than $50 million annualized. AlixPartners also helped the client understand the implications of its shift to AI-powered products, delivering the granular visibility into product-level cloud consumption that the hybrid pricing models of AI-driven SaaS economics demand.

Beyond the immediate financial impact, the client now has a functioning FinOps governance structure and the cost visibility required to manage an AI spend curve that is, by the clients own assessment, about to accelerate sharply.

The PE perspective

Technology's role in the investment thesis is now assumed; the analytical challenge is understanding how it amplifies or diminishes value. That analysis is growing more complex as AI reshapes both the upside and the risk profile of technology assets, not just for software businesses, but for any company in which the technology estate is a material determinant of growth potential, operational efficiency, and resilience to disruption.

Clive De Silva has identified three areas where technology is currently shaping value creation plans.

The first is modernizing the digital core, not as a legacy clean-up exercise, but as an AI enablement play. Boards and ExCos are starting to discover that dealing with legacy isn't just about risk and technical debt. They're now realizing that this is a hindrance to being able to adopt and move fast when it comes to AI."

The second is using AI and automation to unlock efficiency and effectiveness across back-office functions, software delivery, and time to market, with implications for both the cost base and the pace of product innovation.

The third is data and infrastructure sovereignty: where data is housed, how resilient it is to regional disruption, and what exposure the business carries if a jurisdiction restricts access. "Organizations are being more thoughtful now around where to store their data. If a country decides they're no longer going to allow you to have access to that data, how do you make sure you're protected against that?" (For more on sovereignty, see chapter 5.).

From IT checklist to technology value analysis

Technology due diligence is no longer a checklist exercise asking whether the core system is supportable and what the cyber exposure looks like. The difference between what the technology estate costs to run and what it should cost if well-maintained is now a direct input into the valuation conversation, and tech debt quantification is now a standard AlixPartners due diligence workstream.

For any technology or software business, the central diligence question is which parts of the value proposition are defensible in an AI-native world, and at what multiple that distinction is already reflected in the price.

"Boards are still wrestling with AI. Specifically, they are asking: Where is the value, and how does that impact our ability to win in the market and optimize our cost base?"

Clive De Silva
Partner & Managing Director

"This all leads back to the foundational question of can our technology landscape actually support an AI-first infrastructure? That's why modernization keeps coming back as a priority."

Marcello Bellitto, Partner & Managing Director, approaches the analytical challenge that underpins modernization on three fronts: understanding where the business is exposed to AI disruption, whether it can adapt at the pace required, and whether the current multiple already reflects that distinction or has yet to catch up.

That assessment requires a realistic view of possibilities. For example, agentic AI can expand capabilities in theory, but in practice adoption at scale remains rare and the limitations of the technology are often underestimated by investors and management teams alike. Cost, scalability, and governance challenges are routinely uncovered only post-commitment. The most credible enterprise deployments so far have focused on narrow, well-defined use cases, such as customer feedback processing in retail, or credit assessment and onboarding workflows in financial services.

The key question here, then, isn't a simple binary one about AI adoption. It's about whether a business's bets are well-scoped, its governance infrastructure adequate, and its underlying technology estate capable of supporting the planned deployment.

Value creation post-acquisition

PE ownership removes the quarterly reporting pressure that pushes corporate technology leaders toward incremental decision-making. Its most powerful lever is governance: the ability to impose executive alignment across the CEO, CFO, CIO, and CTO functions to drive sustained transformation in ways that corporate environments frequently cannot.

"In companies where that cohesion is absent, you might see that the CFO wants to do something but the CRO is not ready, and the awkwardness works its way into the technology."

Thomas Morineaux
Director

Morineaux cites a PE-backed financial services client as an illustration of what becomes possible when governance alignment and capital work together. Carrying a legacy platform that had become commercially constraining and was costing the business market share, the client made the decision to rebuild its entire infrastructure from scratch rather than continue patching. Supported by approximately £200 million of PE capital, a leadership team combining digital expertise with deep vertical domain knowledge, and the governance alignment that PE ownership provided, the business designed and built a new architecture from first principles.

The result: it was able to double in size and deliver three times as many products with no additional IT costs. "It was a risky thing to do, but they had PE capital to be able to do it. They had a very good mix in their leadership of people who've done this kind of thing before and had the skills to drive it forward."

It should be noted, however, that transformation at this scale remains relatively rare. Most PE funds seek a less interventionist path, using technology improvement as a lever within an existing platform rather than a ground-up rebuild.

Buying at the right price

With holding periods lengthening and exit multiples under pressure, the price paid on entry has become a more significant determinant of IRR than it was in the era of multiple expansion. Technology due diligence is coming into its own as a tool for negotiating the deal thesis—surfacing technical debt, AI readiness gaps, and software pricing exposure that create material grounds for price adjustment.

Cyber: The new purchase agreement clause

Treating cyber as a post-close agenda item is no longer commercially rational.

"Cyber conditions are now being embedded in purchase agreements before close—with specific remediation actions required before the deal completes, not left as a post-acquisition agenda item."

Edd Hardy
Director

Some PE firms are building centralized Security Operations Center capabilities shared across portfolio companies, providing more rigorous monitoring and faster response at a lower total cost than separate functions operating at separate maturity levels.

Sovereignty and asset selection

The software valuation reset has created a moment that rewards analytical precision. "Software companies previously traded at much higher EBITDA multiples compared to industrial companies, but this gap has closed," says Bellitto. "Current lower multiples present opportunities for investors, especially as the impact of AI becomes clearer and certain assets may offer significant upside."

Identifying those assets requires understanding which parts of a value proposition are properly defensible, and sovereignty is increasingly part of that assessment. Bellitto observes that European-headquartered technology businesses, particularly those serving regulated sectors, are attracting more relative interest as the risks of dependence on non-European providers become more visible. IT services businesses, in his view, are better positioned than software businesses in the current environment: less directly exposed to AI disruption of software economics, and with more opportunity to transform their delivery model post-acquisition.

AlixPartners works with PE clients across the full investment cycle, translating technology risk and opportunity into financial terms that inform every stage. Pre-deal, our work includes value-creating technical due diligence and cybersecurity remediation. Post-acquisition, it includes tech debt resolution, ERP modernization, software licensing renegotiation, cloud and FinOps, and portfolio-level cybersecurity governance.

3 The legacy reckoning

Legacy systems were always a drag on efficiency, but now the liability is existential, as organizations discover that the ceiling on what they can do with AI is set by the floor of what their infrastructure can support.

A third or more of developers' time is currently consumed by workarounds caused by technical debt rather than value-creating work, according to the AlixPartners Disruption Index. Among companies with outdated technology, two-thirds view new technologies as major revenue threats; among those with well-maintained infrastructure, three-quarters see them as opportunities.

Technical debt, by Thomas Morineaux's definition, can be seen as the enterprise's aggregation of "technical constraints, accumulated over time, that result in higher costs to run and change systems and company processes, than would otherwise be reasonable. Examples include outdated systems with functional debt resulting in business workarounds, system duplication or outdated architecture requiring more effort to change, and outdated technology that requires a premium to run in nominal performance."

AI has raised the stakes for organizations still running on legacy infrastructure. The same shift that creates opportunity for those with clean, accessible data creates a hard constraint for those without it. Legacy systems do not only carry a maintenance overhead, they set the boundary of what AI can do. Says Brian Kalms: "Whatever the AI application, if the data doesn't flow, the use case doesn't work."

3.1 Talking to the CFO about tech debt

"The demand for agility and cost efficiency have one common enemy, and that's tech debt," says Morineaux. "It's often seen as a metaphor but in fact it's a very real number that directly impacts cashflow and P&L."

The consequences of not addressing this "silent pain" compound gradually, manifesting as a steady erosion of EBITDA that accumulates year on year, while root causes remain unaddressed. "Tech debt can become a background fact of how the organization operates," says Morineaux. "But the reality is there's serious money being spent on things that really aren't required. Every dollar spent on those workarounds is a dollar not spent on building new capability. If you're spending $100 million on tech but really it should be costing you $65 million, the board should know this and react to it."

Yet many companies have a blind spot about the causes and consequences of tech debt. The assumption is that it builds up in run costs, the ongoing overhead of maintaining aging systems. But the more damaging form is the wasted costs associated with change itself. AlixPartners' analysis shows that wasted costs can account for as much as 30-40% of the total cost of change programs, while the tax on run-state operations is typically 10-20%. The drag that tech debt places on an organization's ability to change is often larger than the overhead it places on daily operations, and far less visible to the people who control the budget.

The traditional case for legacy investment reliably loses out to competing priorities in the CFO's inbox. The reframe that works is not an IT modernization argument but a financial one: tech debt is a tax on every change the business wants to make. Sunita Das gives some examples: "Unmanaged tech debt is consuming 20-40% of our change capacity." "Every major feature takes two to three times longer than it should." "We cannot scale to AI because we have this digital duct tape going on right now, and for AI to be enabled, it needs to have a seamless flow of information end to end." An organization can run legacy systems indefinitely for transactional purposes. But it cannot build effective AI on top of fragmented, inconsistent data flowing through manual workarounds.

AlixPartners has developed a methodology for granular tech debt quantification across three categories: obsolescent technology approaching or past end-of-life; unnecessary complexity from poor documentation and inability to reuse code; and duplication of systems from M&A activity. One metric AlixPartners recommends as a board-level KPI is the technical debt index, the ratio of rework cost to new capability build.

Building on technical debt may feel less risky than addressing it, but that's not necessarily the case. In one enterprise example, an IT team sought approval for a significant upgrade to a subscale system. The new platform, which would resolve the issue rather than defer it, cost only modestly more, but the instinct was to retreat to the smaller number. "It was like a religious belief," says Morineaux. "The smaller number may have felt safer, but the longer-term consequences are likely to bite deeper."

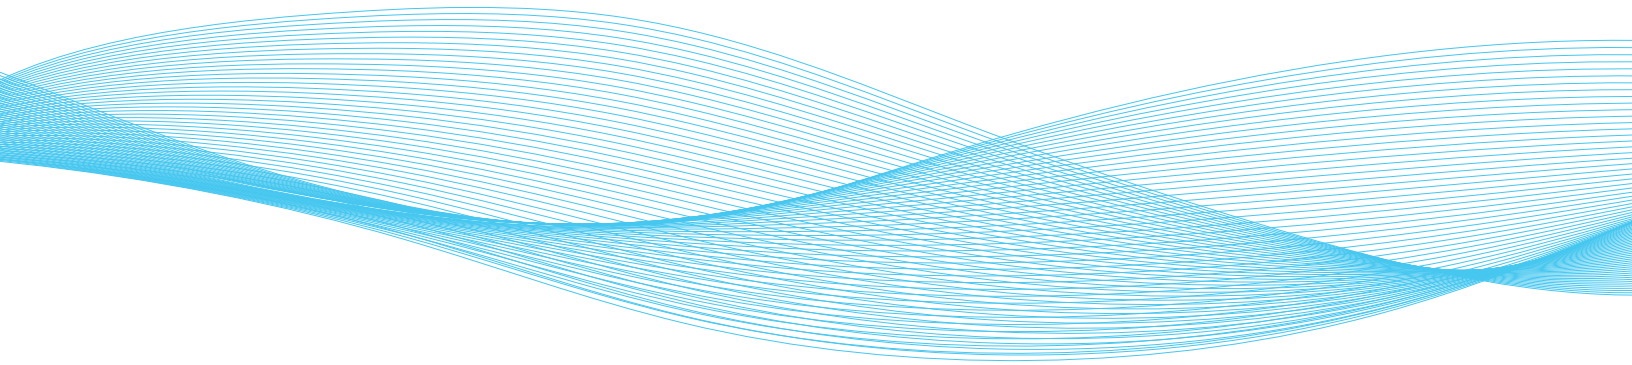

Should you even bother addressing legacy?

A debate has emerged within technology leadership circles about what to do with legacy systems in the AI era. For investors, the question has a sharper edge. As Paul Kelly, Partner & Managing Director, frames it: "The question is less 'fix the tech or not?' and more 'which tech path maximises risk-adjusted returns over the hold period?'"

The arguments fall into three broad camps:

Camp 1: Ignore it and live with it

The reasoning: We have always managed. Legacy systems run the core business reliably. The cost and disruption of change exceeds the cost of continuity, and as long as AI tools can be layered on top via a data extraction layer, the underlying estate can be left alone.

This is a common position, but it underestimates how much tech debt constrains AI adoption in practice. AI cannot compensate for fragmented, inconsistent data. The extraction layer itself becomes a load-bearing structure, and the underlying system that was supposed to be retired is still there ten years later, with a middleware layer *and* an agent layer on top.

Camp 2: Address it and modernize

The reasoning: the cost of inaction compounds. Every year the estate is not modernized, the tax on change grows, competitors pull ahead, and the eventual modernization becomes more expensive.

Kelly characterises this as "tech modernization with embedded AI": Replacing or reshaping key platforms, fixing data quality at the source, and embedding AI into the operational core. It requires more capital and time, but "you pay to transform once, reduce structural risk, and create a compounding asset where each new AI-enabled workflow strengthens the equity story and supports higher exit multiples."

Modernization must still be scoped analytically rather than ideologically: The key is identifying which parts of the estate need replacing before the constraint becomes critical. Whole-scale ERP transformation is not always the answer. And AI is now changing what the answer looks like.

One approach gaining traction is using AI tools to build a working replica (or twin) of the legacy system in parallel—testing and validating the new version against live conditions while the existing system continues to run, then switching over once the replacement is production-ready. It sidesteps the high-stakes cutover risk that has historically made large-scale migration so costly and, as AI coding tools extend their reach into existing codebases, is becoming viable at a cost and speed that was not realistic even 18 months ago.

Camp 3: Wrap it and shrink it

The reasoning: Reframe legacy systems as systems of record. Wrap them in a data extraction layer, expose the data to AI agents, and build new capability on top without touching the underlying estate.

This "data layer on old tech" play captures both its appeal and its risk: It can deliver faster uplift in EBITDA and cash with lower initial spend but, says Kelly, "it defers underlying technical debt and can create an additional layer of complexity that future buyers will price in unless there is a credible plan for the core."

The argument has force for smaller, less complex systems and as a short-term tactical response. But it tends to underestimate both the level of data quality within the legacy system and the permanence of the wrapper. A system that cannot flow clean, structured, consistent data into an extraction layer cannot be effectively wrapped. And the history of enterprise technology is full of integration layers that started as temporary bridges and became permanent load-bearing structures, adding cost and complexity without resolving the original problem.

Wrapping is a legitimate tactic, not a strategy. "The firms that will win are the ones that explicitly choose between 'modernize with embedded AI' and 'data layer on old tech' deal by deal, tie that choice to the value creation and exit thesis, and then execute with discipline." Fixing legacy tech, he adds, "is no longer a credible reason to delay performance improvements": The tools now exist to do both at once.

3.2 ERP modernization in an AI world

The tech debt argument is most consequential in ERP, where the systems are largest, the customization deepest, and the cost of inaction highest. The 2027 SAP ECC end-of-life deadline has put thousands of companies in the middle of S/4HANA migrations, and many are discovering that what looked like a technology project is in fact a business transformation that happens to involve technology. The resources, governance, and leadership attention required are completely different from those of a typical IT implementation. “The crux of the things that give people problems around ERPs is basically data,” says Morineaux. “How it flows across different systems, all the way from sales through to financing, accounting, and payments. That is a question that to this day seldom gets asked.” The hard part is not usually the associated technology, however, but getting the business to agree on how data should be structured, who owns it, and how it should move between functions. That conversation tends to surface every long-standing disagreement about process, reporting, and accountability in the organization.

Why architectures go wrong

Our experts identify three common factors that recur in failing ERP programs. The first is at the integration level. How systems will integrate at a data model level is rarely examined upfront. The CFO wants a finance system, the CRO wants a CRM, but the question of how a customer record in the commercial system becomes an account in the finance system without manual intervention gets treated as a technical detail rather than a design principle. Plugin culture makes this worse. “Salesforce has the biggest plugin library in the world,” Morineaux observes. “People want to use plugins and love the idea that things just plug together so you don’t actually have to think about the integration. That’s part of the issue.”

The second factor is customization complexity. Years of custom code accumulated to accommodate the preferences of individual functions rather than follow overarching architecture principles, particularly in sales and CRM, have led to programs built on unstable foundations. When migration comes, the organizational rationale for that code may have been lost, and it must now be unpicked, rebuilt, or abandoned.

The third factor is leadership misalignment. ERP transformation requires the CEO, CFO, CIO, CTO, and CRO to be aligned on a single plan and committed to it for the long run. “Getting that cohesion in the leadership is very important, and they need to be in it for the long run,” says Morineaux. “Otherwise, a lack of leadership alignment will be reflected in the tech.”

Assessing whether to rebuild or patch should be based on the total cost of each option over a realistic horizon, rather than on which number looks smaller in the first year. The PE perspective following Chapter 2 includes a case study showing what becomes possible when the governance alignment and capital exist to make the bolder choice—in this case, rebuilding a FS platform from first principles and ending up with infrastructure that can support twice the business at no additional IT cost.

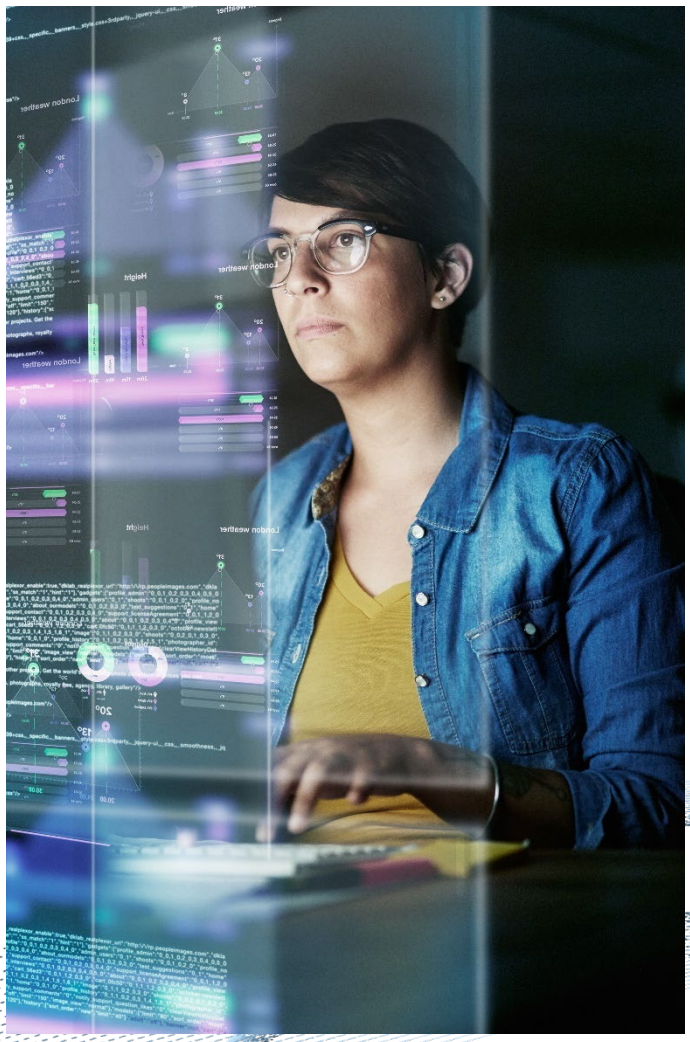

Case study

ERP transformation in six months

A major European industrial goods company, facing the SAP ECC end-of-life deadline and operating on a system that was constraining operational efficiency and the ability to scale, engaged AlixPartners to lead an S/4HANA transformation program.

The program was structured around four dimensions of readiness in parallel rather than in sequence:

Business readiness

The target operating model and core processes were defined and agreed to before any system configuration began.

Data and system readiness

The data migration approach was validated upfront, not deferred, with data quality issues surfaced and resolved before they became program risks.

Program governance

The business case was owned and sponsored by the business, not IT, with financial accountability sitting with the CFO rather than the CIO.

Organizational readiness

Change management and capability building ran from day one, not as an afterthought in the final quarter.

The program was completed in six months. The contribution to EBITDA over the following three years was 34 million euros, with 18 million euros in cost savings over the same period.

The parallel rebuild option

The rate of change in AI-led modernization capability surprises even seasoned observers. “Twelve months ago, the option to rebuild was considered too expensive and/or high risk; there are many examples of failed modernization programs.” says Stuart Mitchell.

“However, now, even with older code bases it is becoming increasingly viable for AI to understand that code base, and reverse engineer specifications; and then rebuild using pure Agentic SDLC at a fraction of the cost.”

Stuart Mitchell
Director

In other words, rather than trying to remediate the tech debt/constraints and having to make compromises due to time and budgets, the opportunity to completely rebuild is becoming increasingly viable.

The parallel rebuild approach—constructing a new version of a system alongside the old one, validating it, then migrating across—is now a realistic option, at a cost and speed that was previously prohibitive. For PE-backed portfolio companies carrying significant technical debt, this represents a material change in available options and should prompt a reconsideration of the rebuild-or-patch question across portfolios.

That capability changes not just what modernization costs, but what the resulting architecture can look like. And with it comes a new set of governance questions.

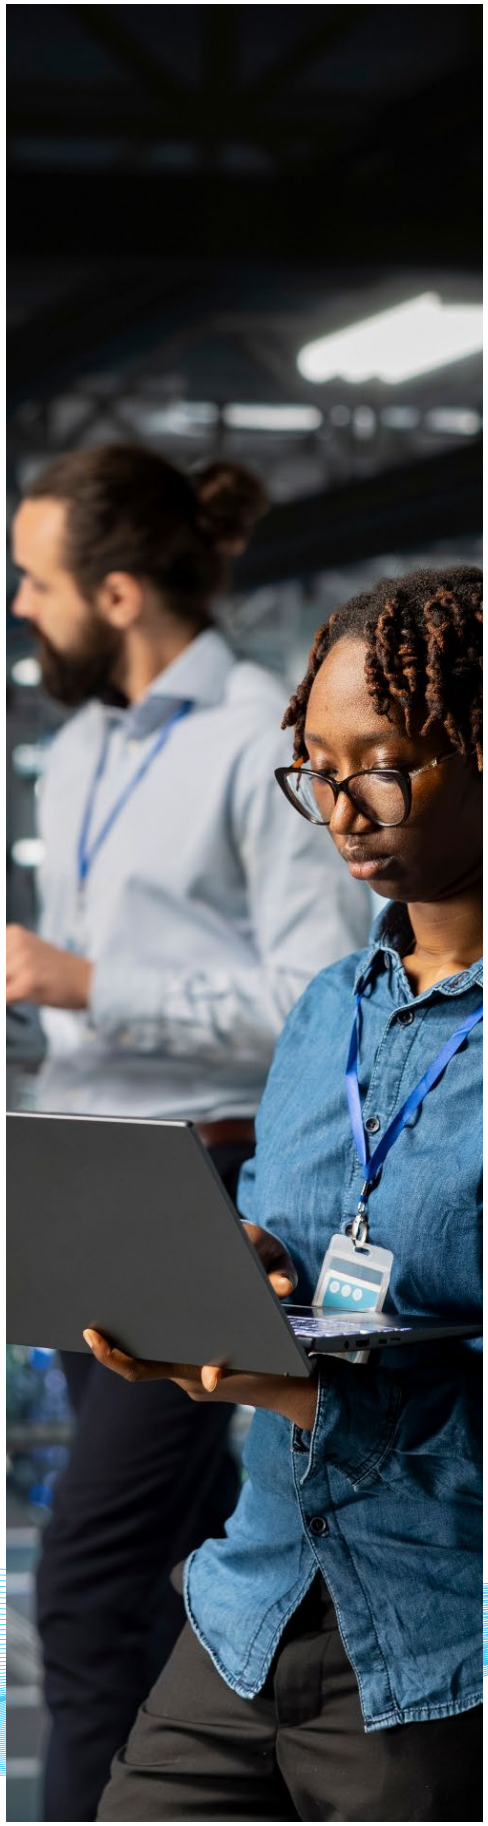

3.3 The next-generation stack

The IT architecture that most large enterprises are running today did not arrive by design. It accumulated. Thin, basic ERP systems were progressively customized over decades, layer upon layer, until organizations ended up with heavy, complex, bespoke backbones encoding accumulated business logic that today may no longer be fully understood but that nobody wants to disturb. The question worth asking now is what the architecture of the next decade will look like, and whether it will be more or less legible than the one it replaces.

The direction of travel is toward what Marcello Bellitto describes as a "balkanization" of the IT stack. At the base, he envisages a thin, standardized ERP layer with minimal customization, where the company adapts to what the software delivers rather than demanding that the software adapt to the company. On top of that base, there would be a thick, proliferating layer of AI agents, built and maintained by multiple parties, system integrators, internal teams, and individual business functions, using a range of tools and to varying standards.

The potential gains look powerful. An architecture that can be changed at speed, without the cost and risk of modifying bespoke code that is poorly documented and tightly coupled, offers a significant competitive advantage. "If the company adapts to what the software delivers rather than the other way around, you get less customization and a more standard approach," says Bellitto.

The governance challenges are significant, however. "The stack of the future will not be cheaper. It will simply be more agile; faster to be changed. But the governance could be a mess." Unlike legacy ERP customizations, which at least sat within a managed application portfolio with known owners and documented dependencies, the agent layer will be diffuse, multi-authored, and difficult to audit. Who is responsible for an agent that has been running for two years and is producing outputs no one is monitoring? What happens when two agents from different business functions interact in ways their respective owners did not anticipate?

The shift to probabilistic systems compounds this. "Until now, the approach to developing IT stacks was very deterministic," says Bellitto. "I have a system, I have inputs, and thanks to the system's processes I have a given output. I know the input, I know the mechanics, I know the output. Today is a completely different world because the system is a black box. And the outcome is probabilistic."

Application maintenance in a probabilistic world means something different from fixing bugs too; it means ensuring that models deliver consistent, non-hallucinating outputs over time, that outputs do not drift as the model is updated, and that the business consequences of probabilistic error are within acceptable bounds.

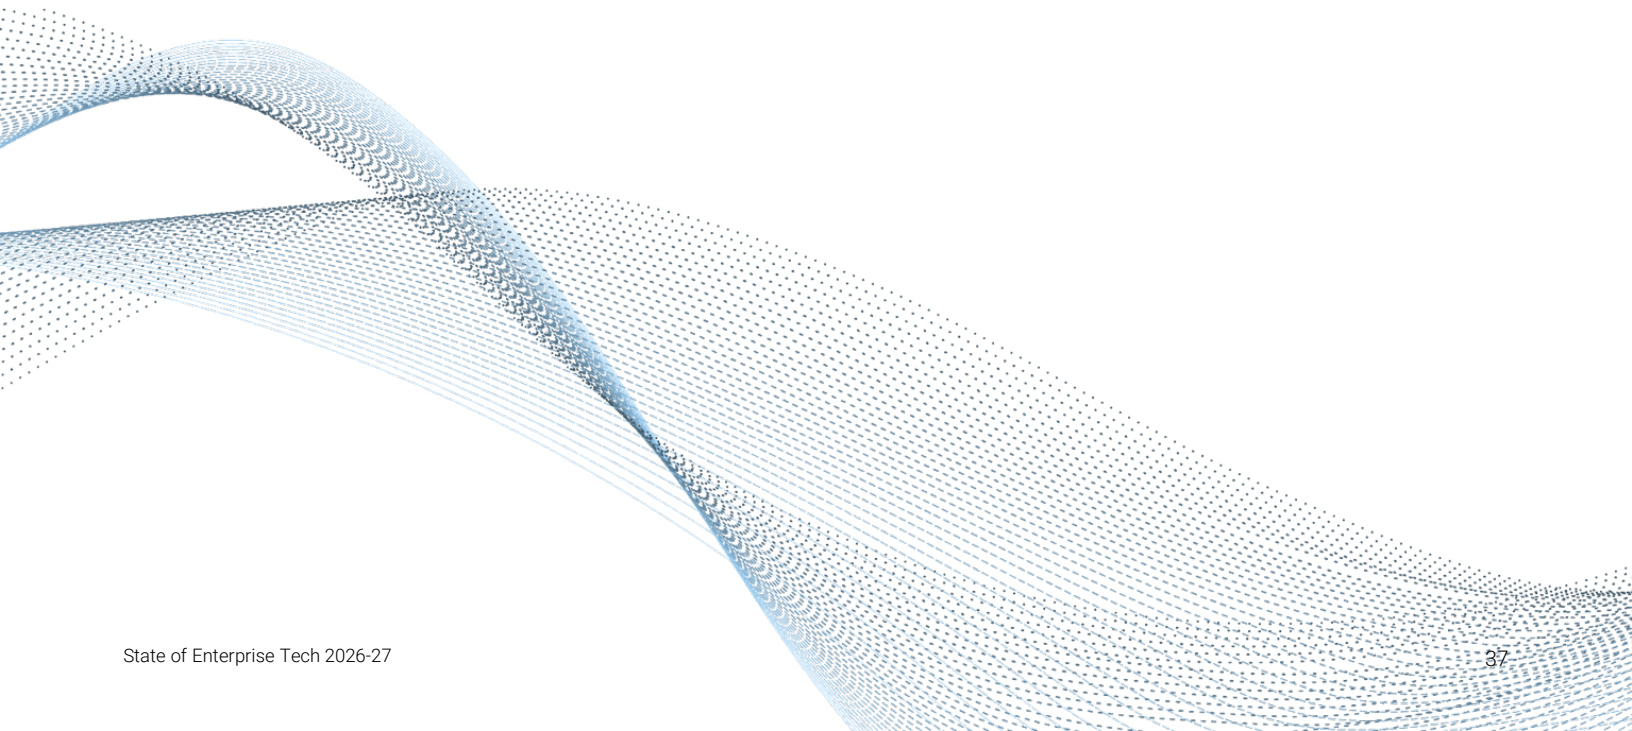

Organizations that implement AI without the governance architecture to manage probabilistic outputs, multi-authored agent layers, and accountability for automated decisions risk a new generation of technical debt.

"In ten years, we may regret the way we implemented AI in our companies, because it creates so much uncertainty; so much higher complexity."

Marcello Bellitto
Partner & Managing Director

Avoiding that regret requires building governance into the agent layer from the beginning, not retrofitting it after the fact. That starts with a deceptively simple question: who owns the agent?

Agent governance

"The agent developer should not sit in IT," says Bellitto. "It has to sit within the finance team, within the engineering team. There will be a blurring effect of the boundaries between organizational units, which already happens with data. This AI agent shift will simply accelerate and expand the phenomenon." The agents that matter most, such as those handling pricing decisions, customer interactions, financial processes, and operational decisions, need to be owned by the functions they serve, because domain knowledge is the critical input to building an agent that behaves correctly when edge cases arise.

The governance architecture that prevents Bellitto's predicted regret requires accountability frameworks that trace every agent output to a named human owner; a registry of deployed agents that makes the estate visible and auditable; retirement processes for agents that are no longer needed or no longer performing as intended; and explicit governance of how agents from different functions interact.

None of these are settled practice across the industry yet, but the organizations that build them now are not just managing current risk. They are avoiding the governance debt that will otherwise compound in exactly the way that legacy ERP customization did before it.

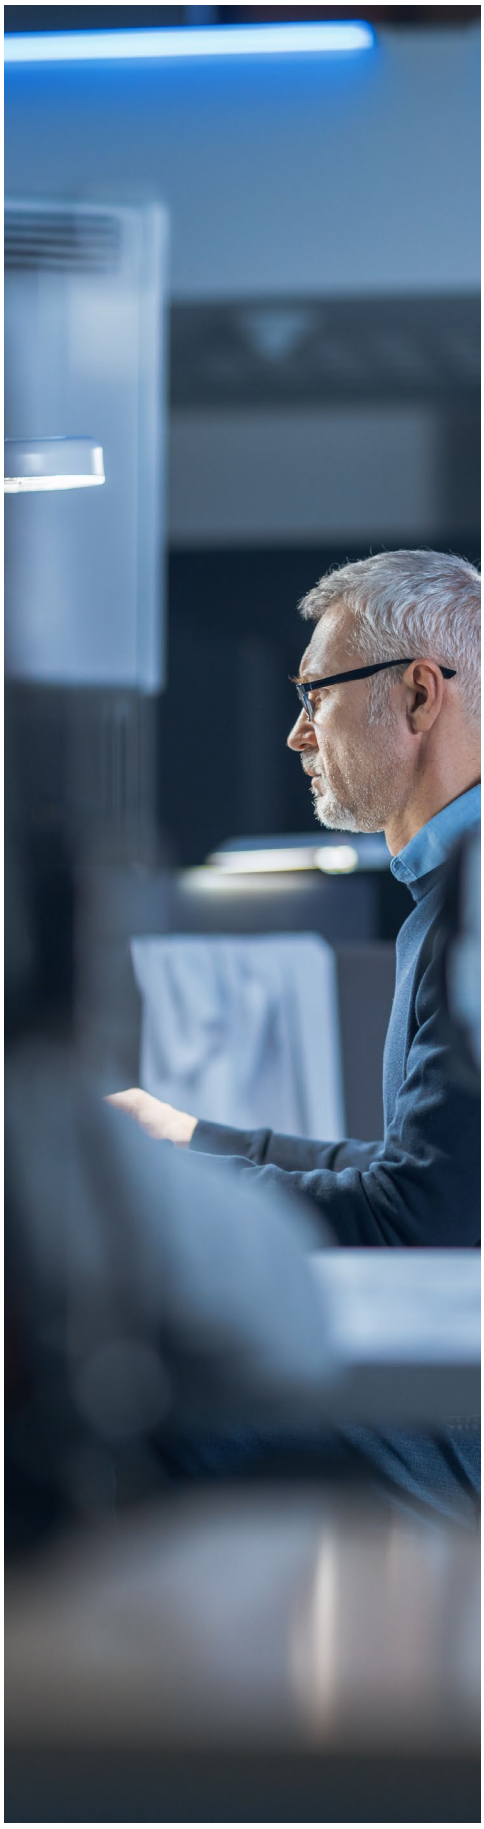

3.4 SaaS and ERP in the AI era

If the stack of the near-future is thinner at the ERP base and thicker in the agent layer, where does that leave the major SaaS or ERP vendors whose products form the current foundation for so many organizations? The short answer is more stable than the headlines suggest, but also more exposed than currently bullish renewal rates indicate.

The death-of-SaaS narrative is overstated. Large ERP and SaaS platforms are sticky by design. Switching costs, institutional knowledge embedded in configured systems, and migration disruption all encourage inertia. None of this changes because of stories about a competitor claiming to have replaced their CRM with a bot.

SaaS strengths

Three structural advantages keep the major SaaS platforms in a strong position. They have existing client relationships and contracts that are expensive to exit. They hold significant training data from millions of clients, an enduring competitive asset as AI becomes more central to their product offerings. And their R&D budgets dwarf what most enterprise customers could deploy independently. "If clients can build AI tools themselves, SAP, Oracle, and Salesforce can do it too," as Morineaux puts it. "They have a much bigger R&D spend than most end users would have, and quite rightly so."

The practical implication for end users is that straightforward: switching decisions should not be driven by AI disruption narratives. The AI capabilities that vendors are building into their platforms will arrive within a two-to-three year horizon regardless. "SaaS companies are going to provide that AI capability as part of their offerings in the next two years," as one CTO told Sunita Das. "So why not wait versus building it all ourselves?" That strategic stance is a counterpoint to the prevailing pressure to build immediately.

But even the biggest platforms will not be immune to disruption. The decision to build-versus-buy is shifting as the cost of building custom software drops significantly. "You're probably going to see a swing to a lot of build where maybe the normal historical decision was buy," says Stuart Mitchell. But the threat to vendors is hardly existential yet: go-to-market, sales, and support costs remain significant constraints, and the economics of building still look very different from the economics of buying a mature platform with an established customer base. Even Amazon, builders of AWS, uses Salesforce for its CRM. The deepest engineering capability in the world doesn't make buying irrational where it makes commercial sense.

> "Let's be real: You're not going to change your ERP because of news headlines on AI,"

says Thomas Morineaux.

Challenges to the SaaS hegemony

Despite the apparent stability of the SaaS market, AI will go on disrupting the competitive dynamics of enterprise software, not least by commoditizing feature differentiation. The capabilities that previously distinguished one ERP or CRM from another can increasingly be replicated by any vendor using the same underlying models. "The moat that these companies once had is now a drawbridge, and it's a lot narrower," as Morineaux puts it. This narrowing is likely to show up in valuations before it shows up in renewal rates.

A second structural weakness relates to SaaS vendors' own AI readiness. "A lot of them are really subscaling this," says Morineaux. "They have similar tech debt problems to what their clients have, albeit it's not in their back-office function, it's more in their core platform. They typically come from software engineering backgrounds rather than data backgrounds. The shift required is one they are still in the middle of." For this reason, AI-enabled features promised in vendor roadmaps should be evaluated critically, and outcome-based pricing claims from vendors who have not demonstrably delivered AI outcomes should be challenged.

The strongest counter-argument to the 'replace SaaS with AI agents' thesis brings us back to data. Before any organization can credibly replace a SaaS platform with AI agents, it needs clean, well-structured, and accessible data. Ironically, the organizations with the data foundations to act on the agent opportunity are by and large the ones getting the most value from those platforms already.

4 Roles and leaders: The organization in transition

The challenge of the AI revolution is how to evolve the roles, metrics, and structures of the technology function even while the technology itself is changing.

Technology decisions are easier to make than people decisions, so it's no surprise that workforce strategies have struggled to keep pace with the rate of change in tools and platforms. According to the AlixPartners Disruption Index, 84% of executives report rising productivity, while 49% worry that their employees' skills are becoming obsolete. Ninety-five percent of CEOs expect AI to lead to layoffs within five years, with almost half anticipating reductions of 10% or greater, and yet 40% of technology companies say talent constraints may already be hindering their transformation.

As one technology leader anonymously put it in a 2026 study of AI adoption across major UK corporates: "Technology is moving faster than the organization can decide what it's actually prepared to be accountable for." In the foreword to the same survey, Rob Hornby notes "a misalignment between leadership ambition and organizational readiness. While the technology is advancing quickly, institutions are still grappling with how to govern, evaluate, and embed systems that are inherently probabilistic within structures designed for predictability and control."

These numbers and observations describe organizations at different points of the same transition. Navigating it well means treating productivity, obsolescence, and talent as a single connected design problem rather than pressures to be managed in isolation. The limiting factor is rarely capability or intent. It is accountability and clarity about roles.

4.1 CIO: From IT operator to business architect

The CIO role never stands still. In the 1980s and 90s, it was a custodial function: managing servers, keeping email running; invisible as long as things worked. In the 2000s, the CIO became a service provider and information broker, managing outsourcing relationships and large ERP implementations. Over the past decade it has evolved into a driver of digital transformation, earning a seat at the C-suite table with responsibility for programs that touch the whole business.

The next evolution is already underway, and it's perhaps the most significant shift yet: the CIO as business architect, accountable not just for how technology functions but for how it creates value.

"The CIO should stop being the operational upkeep, making sure the systems are working, and become the architect of the future of the enterprise."

Sunita Das
Partner

'Order taker'	'Service provider'	'Driver'	'Value creator'
• Focus is on **optimizing basic/utility services**, e.g., infrastructure, networks, enterprise applications in Finance, HR, Retail • Provides **support services** to business functions and other brand IT/digital teams • Also provides some **compliance capabilities** (e.g. info security)	• Seen as driving **delivery of IT/digital services** for business functions to consume across the technology stack • **Service level agreements** in place to measure performance • Business expects services to be delivered consistently at **defined service levels** (performance transparency)	• **Proactively engaged with the Business** and plays a pivotal role in the development and delivery of business solutions (both front-end guest/owner-facing and increasingly further down the technology stack) • **Joint IT/digital and business accountability** for cost/value delivery	• Identifies and drives the **adoption of game-changing technology** across the business • Implements **leading edge technology** with an acceptable level of risk • Responsible for **return on technology investments**

Gartner's 2026 CIO Report finds that 71% of CIOs currently struggle to prioritize the AI use cases that will deliver measurable business outcomes. That is not a technology problem. It is a business architecture problem, and it points directly to where the CIO role needs to go.

Why the role is changing

Three forces together are reshaping the CIO role. The first is the convergence of the CIO and CDO (Chief Data Officer) domains. The data foundations required for AI sit precisely at the intersection of CIO responsibility (infrastructure and platform) with CDO responsibility (data strategy and governance). In organizations that have both roles, the question of who owns the AI data agenda can create a productive tension; in some cases, it is driving a merger of the two functions under an expanded CIO remit.

The second is the fate of the Chief AI Officer (CAIO). Where it was rushed into existence by organizations in 2024-25 as a signal of AI seriousness, the role quickly faced an identity question: Is this a permanent function or a transitional one? The answer depends heavily on context. In organizations where AI is genuinely central to the business model, such as a technology company building AI-native products, or a firm where AI strategy needs a dedicated executive voice at board level, a standalone CAIO can be the right structure.

But the challenge arises in organizations that created the role primarily as a signal of intent, without a clear mandate distinct from what an evolved CIO office could own. AI governance, model selection, multi-model strategy, and shadow AI risk management are all natural extensions of an empowered CIO mandate. Where those responsibilities have been handed to a CAIO without resolving how they interact with the CIO's remit, the result may be duplication, ambiguity, and turf tension rather than better AI governance. To navigate this challenge effectively, organizations need to be explicit about how reporting lines, resources, and areas of ownership are distributed between the two roles.

The third is the shadow AI problem (see chapter 1). Business functions building and deploying AI tools outside the visibility of the technology function require a new category of accountability: governing the ungoverned, a mission that falls naturally within the CIO remit.

How the role is evolving

The structural signal that matters most in the CIO's evolution is the reporting line. A CIO reporting to the CFO is institutionally constrained to think about cost, whereas a CIO reporting to the CEO has a mandate to drive growth.

> "How can you report to a CFO and be the business strategist you need to be?"

asks Das. "Because your focus is always going to be: my budget is a cost center, so I have to manage cost. Cost and revenue and growth are not necessarily symbiotic." This distinction has long been understood in sectors such as large-scale manufacturing: At General Motors and Ford, for example, the CIO has historically reported directly to the CEO as a full peer, reflecting a view that technology is central to competitive strategy rather than a cost function to be managed.

Clive De Silva goes further still, suggesting the very title of CIO may be due for reinvention. "If anything, the Chief Information Officer becomes the Chief Intelligence Officer, the one who really thinks about how you architect all of these solutions," he says. Whether the title changes, the mandate he describes is already here: less about managing information flows, more about orchestrating intelligence across the organization.

Alongside this evolution, a new generation of technology leadership positions is emerging as organizations create CTO-type roles focused on building and engineering rather than buying and integrating. As CTOs lean harder into growth, product, and AI-enabled innovation, the CIO must pivot from "owning IT" to "owning value from IT." The most strategic CIOs will act as the execution partner who de-risks the CTO's bets, industrializes the winning ones, and takes cost and complexity out everywhere else.

The orientation of whoever leads the function is increasingly consequential: how fast an organization can adapt its engineering capability to an AI-native model depends substantially on whether that instinct runs toward building or buying. The AI era demands a CIO who can govern the architecture while driving the business case.

4.2 New metrics for a new mandate

The metrics that currently govern most technology functions were designed for a different era. IT cost as a percentage of revenue, budget variance, tickets resolved, headcount reduction: these measures speak to how the IT function is operating, not how it is helping the organization grow and compete. They reward efficiency of delivery over quality of outcome, and they make it harder to build the business case for technology investment. A technology leader being measured on cost reduction and service uptime cannot easily demonstrate how or why modernization investment creates value.

AlixPartners has developed a reformulated set of metrics for the CIO organized around three domains: financial performance, delivery and agility, and AI and innovation, each with a new emphasis on how technology creates business value:

Old metric	New metric
Financials	
IT cost as % of revenue	IT ROI: net profit generated per dollar of IT investment
Budget variance	% of IT budget allocated to innovation vs maintenance
Cost per ticket/cost per user	Revenue enabled or protected by technology
Headcount reduction	Technical debt index: ratio of rework cost to new capability build
Delivery & agility	
Projects delivered on time and on budget	Time-to-market for new digital capabilities
Number of tickets resolved	Deployment frequency and change failure rate
Project completion rate	Business agility score: speed from idea to deployment
IT backlog size	% of IT capacity devoted to strategic vs operational work
AI & innovation	
AI pilots without P&L visibility	% of AI pilots that graduate to production and deliver measurable P&L impact
Activity reporting on automation	Productivity gains and cost avoidance from autonomous AI workflows
Undocumented AI usage	% of AI use cases with documented governance and risk frameworks
Data quality untracked	Data readiness score: quality, integration, and accessibility for AI consumption

IT ROI and revenue enabled by technology make visible what cost-per-ticket never could: the connection between technology investment and business outcomes. The technical debt index brings the "tax on change" (see chapter 3) into the same reporting environment as revenue and EBITDA, where it belongs.

On the delivery side, the shift from project completion rates to deployment frequency and business agility score reflects an AI-native environment where the bottleneck has moved from execution to ideation. And the percentage of AI pilots graduating to production is a concrete way to separate theatre from value creation.

4.3 The tech workforce in transition

The talent implications of AI-native engineering are only beginning to be understood, and in many enterprises, the pace of transition is outrunning the ability to manage the consequences. Role redesign, deliberate upskilling, and succession planning will increasingly become the core work of transformation programs.

Offshore pressures

The offshore software delivery model, which has been a cornerstone of enterprise technology cost management for two decades, is under threat on two fronts: AI tools can now perform the routine coding, testing, and support tasks that were the economic rationale for offshore delivery, at a fraction of the cost. And the geopolitical environment has made the cost advantages of offshore labor less predictable, with tariffs, trade restrictions, and sovereignty concerns (see chapter 5) all adding variables that were not in the original business case.

Already, 30% of organizations are driving value by reducing outsourcing through automated internal workflows, according to the AlixPartners Disruption Index, a proportion that will only grow as agent capability expands.

Business process outsourcing sits at the sharp end of this. BPO providers, whose model depends on supplying large teams to handle repetitive back-office tasks at low cost, face an existential challenge from agents that can do the same work cheaper and without interruption. "The BPO providers providing the low-cost labor and the repetitive tasks in back-office functions that were typically outsourced are going to be the ones most incentivized to go down this agentic path," says Mike Crisanti.

This will mean a shift for the organizations that outsource to BPOs, too. Governing a provider that now operates largely through agents requires different skills, contract structures, and oversight processes than managing a team of offshore workers.

Roles under threat

The integration middleware architect, whose job was to design how data passed between applications, is being replaced in many cases by a standard API, a documented interface, and an agent that handles the translation. The staffing model for implementing a major ERP or HR platform has changed materially: fewer people, more AI tooling; with a focus on outcomes rather than role coverage.

The junior developer intake model, historically built on volume hiring of graduates who would learn through routine coding tasks, will evolve as those tasks are increasingly automated. The concept of junior and senior will shift rather than disappear. Junior roles will likely begin on simpler, maintenance-mode products with lower velocity expectations, where the consequences of agent error are contained and the learning environment is manageable. The fast track to seniority will run through architectural judgment and agentic supervision, not coding volume.

"I would encourage a young person interested in technology to consider pairing software fundamentals with adjacent skills such as machine learning, architecture, domain expertise, or hardware integration." says Clive De Silva. "The broader your platform, the more durable your career."

Enter the agentic supervisor

The role that is emerging fastest, where demand is currently most outstripping supply, is the agentic supervisor: someone who directs, validates, and takes accountability for agent outputs across multiple parallel workstreams. This is cognitively demanding work that requires a different and more intensive kind of attention than execution work.

Building this capability through deliberate upskilling and role redesign is a growing priority, but it will need protecting. Once in place, SG&A reviews may encounter what looks like inflated IT headcount embedded in business functions and face pressure to centralize it. That reflex would destroy exactly the capability being built. Domain experts doing agent development work in their functional context are not a target for rationalization; they *are* the architecture. SG&A programs need to distinguish between IT headcount doing IT work, which is the right target for scrutiny, and domain experts rightly doing agent work in their functional context.

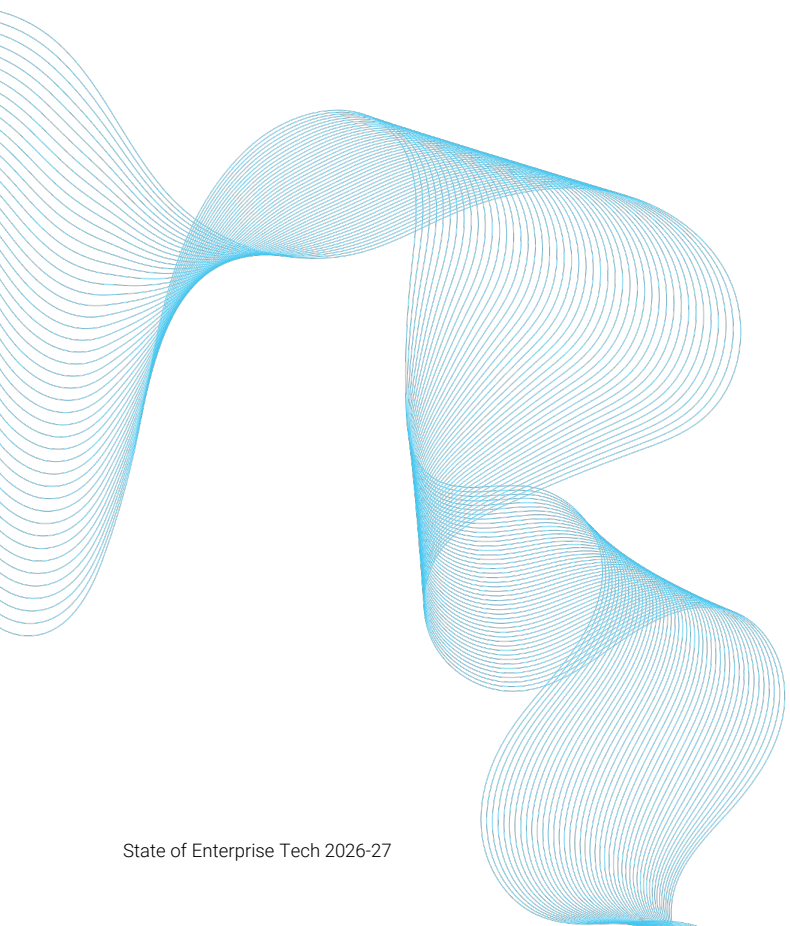

Legacy skills premium

As COBOL and mainframe specialisms fall away, finding replacements where they are still needed has become extremely difficult. “The new workforce is not interested in working on legacy stuff. They want the modern tech stack. You’re not going to find people who want to work on your mainframes anymore,” says Sunita Das.

For a limited window at least, legacy skills will be at a high premium. Das recalls an engagement where the last person who understood a system had to be brought back out of retirement at any cost. “Whatever she demanded, it was met. Her cat had to be taken care of, her spas. Desperate times called for desperate measures.”

Organizations with significant legacy estates and no succession planning for the people who maintain them are carrying a significant human capital risk that has become a core element of M&A and corporate due diligence.

The federal government is the most rigorously documented version of this problem. The U.S. Government Accountability Office (GAO) published a study in 2025 that asked the 24 largest federal agencies for their most troubled systems and scored each against attributes like age, obsolete languages, and cybersecurity risk to rank the worst. The scale behind those cases is large: the government spends more than $100 billion a year on IT, roughly 80% of it on operating and maintaining systems that already exist rather than building anything new. The talent constraint is the part budget alone does not fix. GAO notes that agencies struggle to find staff fluent in COBOL and similar languages and pay a premium for the specialists or contractors who remain.

That scarcity is not a government artifact. The COBOL and mainframe workforce is aging out across the market, and GAO ties reliance on these languages to a dwindling number of people with the right skills wherever the systems run. The private sector shows the same maintenance drag from the other direction: Sonar's State of Code Developer Survey report finds engineers spending an average of 23-25% of their time on toil, maintaining and debugging legacy or poorly documented code instead of building new capability. The federal balance sheet quantifies what corporates experience as lost engineering velocity and a shrinking bench of people who understand the oldest systems. AlixPartners increasingly treats human capital risk on legacy systems as a standard element of technology due diligence, because the maintenance cost is visible while the workforce risk behind it usually is not.

Building the AI-ready organization

As enterprises look to reconfigure around the operating model that AI requires, here are four prerequisites for moving beyond ad-hoc AI adoption to sustained, governed, scalable capability.

Establish an AI Center of Excellence (CoE)

The CoE operates as a standing governance function rather than a project team, responsible for ethics and compliance, tooling selection and evaluation, multi-model strategy, financial oversight of AI spend, and shadow AI risk management.

An AI lab that continuously scans and tests new tools is a core function of the CoE, essential because the AI tooling landscape moves faster than annual governance cycles can track.

Define the agent operating model

A defined set of relationships between human and agent roles, accountability frameworks, and the playbooks that allow consistent operation at scale. It answers the questions every organization deploying agents needs to resolve: who is accountable for an agent's outputs, what is the agent authorized to do, how does escalation work, and how does the organization know what agents are deployed and what they are doing?

Embed a change management program

Not a communications plan, but a substantive effort of role redesign, upskilling, and career pathway development that runs continuously rather than as a project. Organizations that treat AI adoption as a technology implementation consistently underestimate how much disruption is involved and how much support people need to navigate it. The identity dimension—people whose expertise and reputation were built on skills now being absorbed into tools—requires deliberate attention, not just retraining.

Invest in an enterprise context foundation

The shared knowledge infrastructure that allows AI systems and the humans working with them to operate from a consistent understanding of the organization's goals, processes, data, and constraints. Without it, AI deployments produce inconsistent outputs because they are drawing on inconsistent inputs. Building it requires investment in data quality, documentation, process standardization, and knowledge management that many organizations have deferred.

5 Governing AI risk: security, safety, privacy

AI capabilities may create competitive advantage, but they can also expand the attack surface and introduce new vulnerabilities. Further, the infrastructure and models they rely on can pose significant risks to business operations and enterprise stability if not secured appropriately.

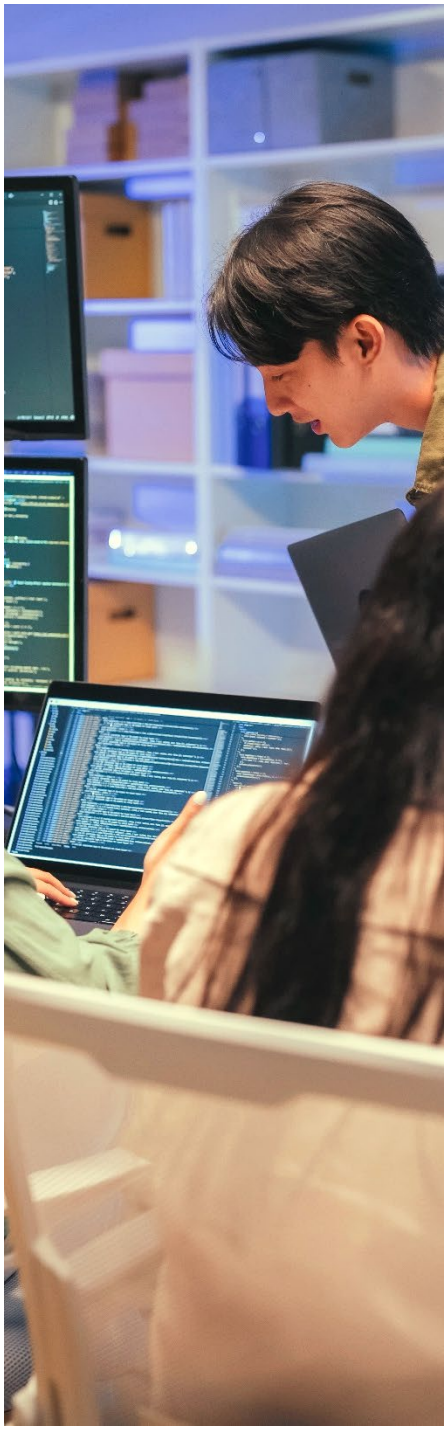

The risk agenda in enterprise technology has been reshaped by AI in two distinct ways. First, AI has become a significant threat vector, enabling more sophisticated and automated attacks and creating heightened concern regarding the governance of AI systems themselves. Second, the concentration of AI infrastructure and proprietary assets in a small number of U.S.-headquartered organizations has quickened the pace of discussion about technology sovereignty in boardrooms and PE investment committees.

The scale of the concern is measurable. Cybersecurity is the most important digital issue facing enterprise leaders in 2026, up four percentage points year-over-year to 41% and at its highest-ever ranking in the AlixPartners Disruption Index. Executives explicitly link this rise to the rapid proliferation of AI tools across their organizations: cybersecurity, deepfakes, and misinformation are the three AI-related risks leaders worry most about as they scale agentic AI and automation.

In episode 10 of AlixPartners' Tech 10 series, Clive De Silva, Catherine Brien, Partner & Managing Director, and Beth Musumeci, Partner & Managing Director, examine how the most resilient organizations are responding to these pressures, moving from treating cybersecurity as a compliance checkbox to embedding it as a strategic capability and competitive differentiator. The conversation also explores how leaders are navigating the workforce disruption that follows, and what separates those managing the transition well from those still looking for a starting point. The numbers behind that urgency are significant.

Gartner's 2026 CIO Report finds that 93% of nonexecutive board members now view cybersecurity as a threat to shareholder value, and 77% of CIOs cite security and risk as the biggest barriers to scaling autonomous technologies. The cybersecurity market reflects that concern: It is projected to grow from $2.4 billion to $3.4 billion between 2026 and 2029, according to Gartner.

5.1 Cybersecurity in the post-perimeter era

The major retail sector cyber incidents of 2025 have dramatically changed the conversation about cyber risk. "We used to talk about data loss as the biggest threat, but it was not easy to quantify," says Edd Hardy. "But when a retailer suddenly sees they cannot accept electronic payments, put products on the shelves, or ship online for six to eight weeks, that's a really easy one. You don't need any tech understanding to appreciate the impact."

The business impact assessment has thus replaced the data loss argument as the most effective way to engage non-technical leadership on cybersecurity risk. "For each critical system, ask: What does it cost the business if that system is unavailable for 24 hours? 48 hours? A week? The numbers are consistently larger than clients expect, and the exercise quickly shifts the conversation from cybersecurity risk to financial risk."

The perimeter is gone

The castle-and-moat model of cybersecurity, a hardened perimeter protecting trusted systems within, is dead: It no longer addresses the reality of the current threat environment. Organizations can no longer defend a perimeter that does not exist, especially with cloud and edge infrastructure spread across multiple providers, data sprawl, employees on home networks, and integrated partners throughout the supply chain. There is no clear boundary between 'inside' and 'outside' anymore.

"Every user needs to be treated the same. To make that a reality, organizations need to modernize their cybersecurity programs now more than ever. The cost of modernization today is a fraction of the value destruction that follows a breach because your security model is fundamentally obsolete."

Megha Kalsi
Partner

AI compounds this situation. As agents embed in devices, applications, and core business processes, the traditional concept of a discrete AI tool dissolves. AI systems interacting with other AI systems, taking action without direct human instruction, expand the potential impact of any compromise. Add to that the trust problem: AI outputs are convincing enough to pass a cursory review but wrong in ways that require expert knowledge to catch. However, AI technology can effectively be used to create deepfakes that can wreak havoc on organizations if cybersecurity programs are not using advanced technology to detect them.

"A breach is inevitable," says Hardy. "It's not a matter of if or when, it's a matter of: Will you survive? Can you mitigate the damage and control it?"

The basics are still failing

Yet even now, says Hardy, and despite a new enterprise focus on cybersecurity risk, the most damaging attacks are almost always alarmingly basic. "Sophisticated attacks are by definition invisible while they are working. The ones that make the news are the ones that have already succeeded," he says. "But in many cases we are still dealing with the basics of: you haven't set your passwords up properly, you are not patching your systems…"

Organizations know better in theory, of course, but the challenge is sustained operational discipline in environments where the pressure is always to move fast, add new tools, and accept technical debt. "It's like we all know smoking's bad for you, but loads of people still do it and then are surprised when it impacts their health." Simply hoping it will happen to someone else doesn't cut it.

Most organizations over-invest in perimeter defense and under-invest in the resilience measures that determine survival post-breach, says Hardy. Offline backups is an obvious example: transformative resilience at minimal cost. "If I encrypt all your data, that's going to be really expensive and damaging," says Hardy. "But if you have an offline backup, it's trivial. Just wipe the environment, restore everything."

Executive war-gaming, the neglected investment

Technical incident response plans are common, but executive decision-making protocols are still the exception rather than the norm.

"Very few actually practice how you make a big decision. Who makes the decision to switch off online ordering? Because though that's a guaranteed loss of revenue, under certain circumstances it might be the safest thing to do."

says Hardy.

The incident response plan stored on the computer is a notorious example: the first casualty of a ransomware attack… is the document that tells you what to do about a ransomware attack.

Single-point-of-failure mapping also deserves more attention than it typically receives, says Hardy. Identifying who in the organization is the sole owner of knowledge or access critical to a process, and stress-testing what happens when that person is unavailable during an incident, surfaces vulnerabilities that no technical audit will find.

AlixPartners facilitates executive cyber war-gaming as a standalone engagement and as part of broader resilience programs. The most advanced organizations continuously pressure-test their own environment, mapping dependencies and single points of failure, and building recovery capability as a priority parallel with prevention investment.

For a PE perspective on cybersecurity risk, including purchase agreements and shared SOC capabilities, see the section following chapter 2.

A note on quantum

The threat to current encryption standards posed by quantum computing is real but further away than some commentary suggests, says Hardy. What needs to be done now is an inventory of the cryptographic dependencies that exist across the estate, and what a transition to quantum-secure cryptography would require in terms of planning horizon and investment.

"Begin the inventory now, before the timeline becomes pressing."

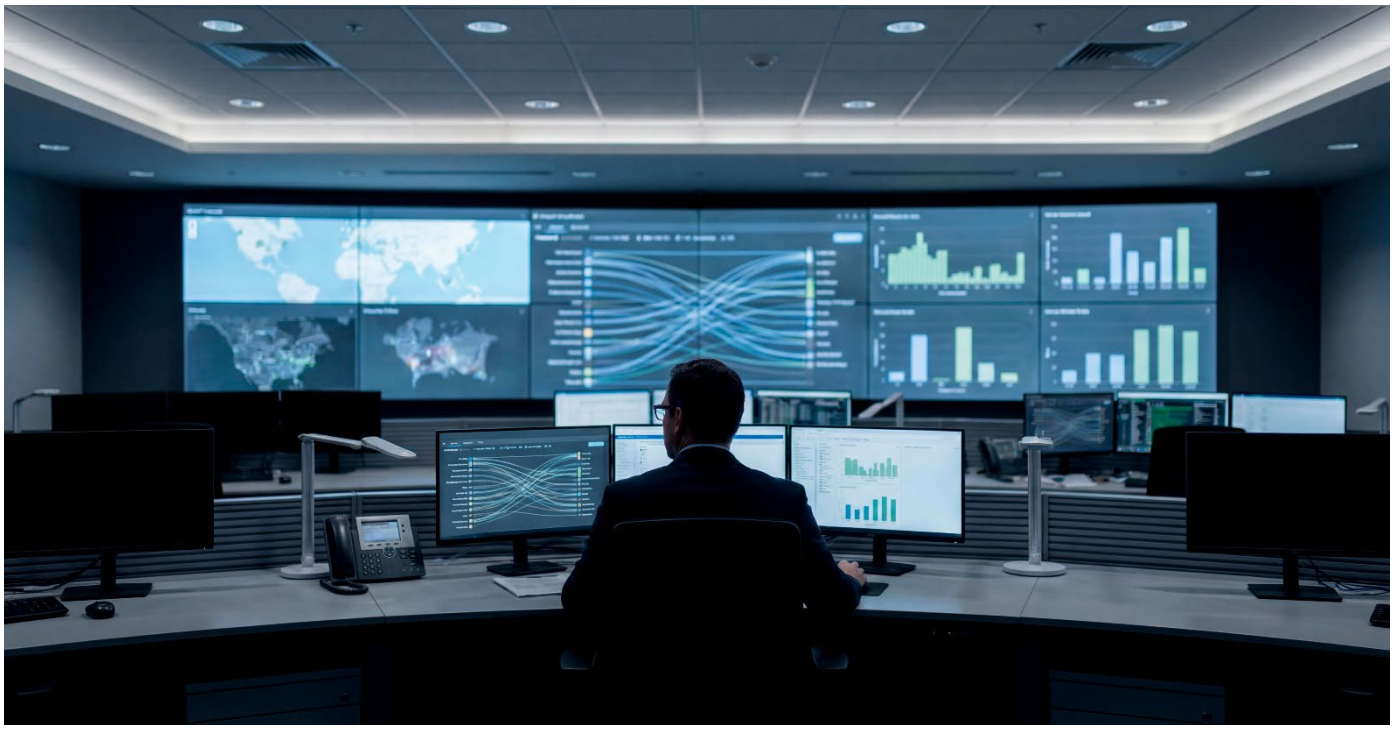

Five leadership levers for cybersecurity maturity

Cybersecurity has moved from a specialist technical concern to a core determinant of corporate performance, informing regulatory outcomes, investor confidence, and the ability to scale AI safely. Regulators worldwide now expect boards to demonstrate credible oversight of cyber risk, not just compliance with minimum standards, and investors are applying equivalent scrutiny.

A major breach can destroy shareholder value, damage customer trust, and disrupt operations—all core ESG concerns. The organizations best positioned to scale AI and digital transformation are those whose boards treat cybersecurity as integral to value creation rather than a cost to be minimized. AlixPartners works with boards and executive teams on five levers for building that capability.

1 Elevate governance

A dedicated board-level forum for cyber oversight, with clear remit, membership, and reporting expectations, allows directors to engage more deeply with threat trends, program performance, and major investment decisions.

Governance should extend to operating model choices: many leading organizations are moving towards a distributed model where a global CISO provides central authority while business-embedded security officers manage local risk.

Cybersecurity assessment should be a standard component of all major strategic decisions, including acquisitions, AI deployments, and platform modernizations.

2 Link risk to business value

Boards need to see cyber risk expressed in the terms they already use to govern the enterprise: revenue at risk, margin impact, regulatory exposure, and reputational consequence.

Realistic scenarios—fulfilment disruption during peak trading, IP theft delaying a product launch—help boards understand materiality, prioritize investment, and set risk appetite with clarity.

3 Embed metrics and transparent reporting

Without reliable metrics, boards cannot judge whether cybersecurity posture is improving or whether investments are delivering value.

A concise set of strategic indicators—covering detection and response timelines, vulnerability trends, critical asset coverage, and third-party exposure—gives directors consistent, meaningful information without distracting from top-line priorities.

Periodic benchmarking against peers or recognized frameworks, with clear insight into gaps and investment priorities, supports both regulatory readiness and investor confidence.

4 Test resilience and incident response

A cyber-ready board focuses not only on preventing incidents but on how quickly and transparently the organization can detect, contain, and recover from events that do occur.

Regular cross-functional tabletop exercises—simulating material cyber incidents and testing decision-making, escalation procedures, and communication pathways—reveal gaps in governance and coordination before a real crisis does. Materiality thresholds and decision criteria should be defined in advance, not improvised under pressure.

5 Use cybersecurity to unlock growth and AI

When boards view cybersecurity as part of the growth agenda, security becomes a lever for speed and differentiation rather than a constraint. Strong controls, trusted data, and resilient platforms allow organizations to deploy AI and digital services with greater confidence.

According to the AlixPartners Disruption Index, growth leaders deploy agentic AI at nearly four times the rate of laggards, and cybersecurity maturity is part of what makes that confidence possible.

5.2 Technology sovereignty, the new front line

Technology sovereignty refers to an organization's ability to control its own technology infrastructure, data, and digital capabilities—and to limit its exposure to influence over these from foreign governments, overseas vendors, or concentrated infrastructure providers.

The risk universe is broad, spanning regulatory compliance across jurisdictions, supply chain resilience, data jurisdiction, AI model dependency, and the geopolitical reliability of critical technology partners. Our experts identify three factors that are currently turning sovereignty from a government concern to a corporate boardroom agenda.

The first is the **concentration of AI infrastructure**. Major AI model providers, cloud hyperscalers, and semiconductor manufacturers are overwhelmingly U.S.-owned. The American government's policy has demonstrated both the willingness and the capability to use that concentration of ownership and domicile as a geopolitical lever, through export controls on advanced semiconductors, restrictions on technology transfer and, increasingly, through the terms and conditions of cloud service agreements that require data to be accessible to U.S. law enforcement under certain conditions. For organizations outside the U.S., the question of what they are agreeing to when they sign a hyperscaler contract deserves greater scrutiny.

The second is **supply chain and vendor risk**. Reliance on vendors in particular jurisdictions creates exposure to tariff changes, trade restrictions, and behavior changes driven by political pressure in the vendor's home market. Recent events have demonstrated that these risks are no longer theoretical. modelling similar future scenarios is now an obvious risk management priority.

The third relates specifically to **data sovereignty**. When proprietary data is processed by an external AI vendor's systems, the model that processes it learns from it. What it learns cannot be fully recovered once shared. Competitive advantage leaks through the inference process, distinct from the data breach risks that existing security frameworks address. "Once data or IP is shared with a vendor, it is difficult to regain sovereignty," says Marcello Bellitto. "And so trust is becoming a key differentiator for software companies."

The European software dependency

The sovereignty challenge has a specific and acute European dimension. Europe's dependence on U.S.-headquartered technology providers is extensive and structural.

Europe's market structure—dominated by small and medium businesses, with relatively few large technology platform companies—has never generated the scale needed to build competitive enterprise software at a global level. Without that scale, European vendors have struggled to fund the R&D investment that enterprise software requires, leaving the market open to U.S. incumbents.

In addition, GDPR, sector-specific data restrictions, and labor market regulations all slow the kind of rapid organizational transformation that AI adoption requires. The same regulatory environment that protects European data from unauthorized non-European access also imposes constraints on how European companies can use that data internally.

These conditions help explain why Europe produces far more IT services businesses than software businesses, and why homegrown platforms have rarely scaled to enterprise level outside a handful of categories.

In ERP, Sage and SAP are the only credible alternatives at scale to U.S. vendors. In CRM, there is no European provider competitive at enterprise level. In infrastructure software, SUSE is effectively the only significant European vendor. Financial services software is better positioned: Finastra and Temenos, both headquartered in Europe, have depth and enterprise-scale capability. But the overall picture points to dependence on non-European providers across most of the stack.

There is no near-term path to filling this through vendor substitution, because European alternatives do not exist at scale for most categories. The mitigation can only be a managed dependency. "The decision to deepen reliance on a non-European provider needs to be made with eyes open to the sovereignty implications, and where European alternatives exist and are technically competitive, the sovereignty argument should be part of the evaluation," says Marcello Bellitto.

Nearshoring limitations

Those constraints shape the options available when organizations look to relocate workloads or source development capacity. European nearshore locations such as Iberia, southern Italy, Albania, and Greece offer competitive advantages for keeping development within or close to the EU regulatory perimeter. But nearshoring is not a clean solution to the sovereignty problem.

In healthcare and other regulated sectors, for example, developers in North African territories face restrictions on accessing European patient data. More broadly, nearshoring relocates the labor without necessarily relocating the data governance risk: if the underlying platforms and cloud infrastructure remain non-European, proximity of the development team addresses only part of the dependency.

Response and mitigation

The sovereignty calculation is already shaping decisions at the largest enterprises. A number of responses are emerging, including European provider selection, double-sourcing architecture, EU-located development teams, and contractual protections proactively built in at procurement.

Tesco signed a **three-year strategic partnership** with Mistral AI, the French frontier AI lab, in 2025. Rather than licensing a U.S.-based model, Tesco and Mistral engineers are co-building generative AI solutions through a joint lab, with the partnership focused on loyalty personalization, data analysis, and internal tooling. The choice of a European provider, with its emphasis on customizable and controllable AI over black-box alternatives, reflects a logic that AlixPartners is seeing inform vendor selection decisions more broadly.

In theory, **double-sourcing architecture**—building the ability to switch between cloud providers without major re-engineering—reduces vendor dependency and provides negotiating leverage. In practice, however, only a small number of sophisticated operators have developed infrastructure that can swap between providers seamlessly, notes Bellitto; most remain tied to a single provider, which deepens both the country risk and the vendor risk. Building vendor-independent infrastructure is costly and requires upfront commitment, but may well be worth the investment for organizations with significant regulatory exposure or strategic sensitivity about data location.

EU-located development teams reduce regulatory exposure even where the software itself is not European. For organizations processing European customer data through AI systems, the location of the development and operations team managing those systems can matter as much as the location of the data center, particularly under GDPR and sector-specific frameworks where data access controls extend to the people operating the systems, not just the infrastructure they run on. Sourcing development capacity from within the EU regulatory perimeter is therefore not just a nearshoring decision but a data governance one.

Contractual protections are negotiable, though enforceability varies considerably by jurisdiction and vendor. The most defensible contracts limit the scope of data sharing with vendors, preserve the right to audit how data is used in model training, and create clear exit paths that do not require a complete platform rebuild.

AlixPartners works with clients to structure technology contracts along these lines, but the strongest position is one in which these protections are built into procurement decisions before the vendor relationship is established, not negotiated after the fact. A vendor with significant leverage at renewal is unlikely to accept constraints it did not agree to at the outset.

The sovereignty lens is also reshaping asset selection in PE, influencing which technology businesses attract interest and at what multiple. For more on this, see the PE perspective following chapter 2.

Conclusion

Ten takeaways from our experts' survey of the enterprise technology landscape in 2026.

1 Rethink, don't retrofit

AI cannot fix a broken strategy or a broken operating model; layering AI onto existing processes adds friction rather than removing it. Every workflow in an organization today was designed around the constraint that humans do the work. Redesigning that workflow around agents as participants, from architecture through team structure to governance, is the precondition for AI adoption to deliver its potential.

2 Treat data readiness as the work, not the prerequisite

Every AI initiative, agentic deployment, or legacy modernization program runs on data. Fragmented, unclean, or siloed data is the single most common reason AI projects fail to scale. Investment in data engineering, pipeline orchestration, and metadata governance is not preparatory work for AI. It is the work.

3 Quantify your tech debt and put it in front of the board

Wasted costs account for 30-40% of the total cost of change programs in organizations with unmanaged technical debt. Express the problem in terms the CFO cannot ignore; change capacity consumed by workarounds; innovation investment crowded out by maintenance spend; cost of inaction compounding year over year. The technical debt index (rework cost as a ratio of new capability build) belongs in the board pack alongside revenue and EBITDA.

4 Audit your software licensing estate now

The Oracle Java and VMware pricing shocks are early examples of a wider pattern: vendors restructuring pricing models at renewal, often with years of lead time required to respond. Build a complete picture of usage, change implications, and contractual protections. Factor pricing model risk into due diligence as a standard workstream. And begin modeling what outcome-based pricing would mean for your highest-exposure platforms before a vendor proposes it.

5 Build AI cost governance before spend spirals

AI spend will grow faster than cloud spend did and is harder to govern. Establish visibility now: cost per thousand tokens; spend by team and product; allocation to business outcomes. Effectively negotiating the shift to outcome-based AI pricing starts with a granular understanding of your own consumption well before contracting.

6 Establish an AI Center of Excellence as a standing function

Ethics, compliance, governance, multi-model strategy, and shadow AI risk require ongoing oversight, not a one-off project. The AI tooling landscape moves faster than annual governance cycles can track. Build a standing function with the continuity to manage what AI programs produce and an AI lab that continuously scans and tests new tools as they emerge.

7 Take technology sovereignty seriously

Risk-assess your exposure to U.S.-headquartered providers as you plan and build AI infrastructure. Audit what you have agreed to in your hyperscaler contracts. Understand where your data goes and how it's processed by third parties. Where European alternatives exist and are technically competitive, factor sovereignty into the evaluation. And build contractual protections into procurement decisions before any vendor relationship is established.

8 Reframe the CIO's mandate

A CIO reporting to the CFO is institutionally constrained to think about cost rather than growth. If technology is central to competitive strategy, the organizational signal should match: metrics tied to business outcomes, accountability for value creation, and a reporting line that reflects the mandate. The role is no longer about keeping the pipes flowing; it is about architecting how the enterprise creates value.

9 Plan for the workforce transition, not just the technology transition

The agentic supervisor is already the most sought-after role in AI-native engineering, and supply is nowhere near demand. Build this capability through deliberate upskilling and role redesign rather than waiting to hire it. Protect domain experts doing agent development work in business functions from SG&A rationalization reflexes. And begin succession planning for legacy skills to avoid costly exposures.

10 Treat cybersecurity maturity as a foundation for AI ambition

Rapid, safe AI scaling requires robust security foundations that give leadership the confidence to move. Strong controls, trusted data, and resilient platforms allow AI and digital services to be deployed at speed without exposing the business to disproportionate risk in a post-perimeter world. Cybersecurity should be seen as a growth enabler rather than a cost center.

Tech 10

For the most recent insights on the state of tech, watch our Tech 10 video series here. In these discussions, we bring the future of technology to the forefront of the boardroom agenda, homing in on the hottest trending topics—from AI to tech modernization, cybersecurity and privacy to data foundations.

Contributors

Rob Hornby
Global Co-CEO

Brian Kalms
Partner & Managing Director

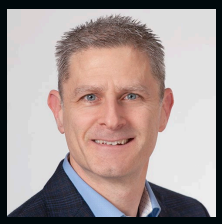

Mike Crisanti
Partner & Managing Director

Clive De Silva
Partner & Managing Director

Paul Kelly
Partner & Managing Director

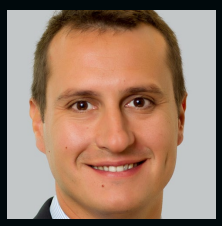

Marcello Bellitto
Partner & Managing Director

Paula Walworth
Partner

Paul Fanning
Partner

Sunita Das
Partner

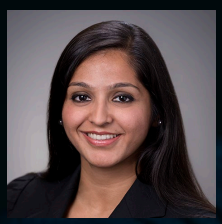

Megha Kalsi
Partner

William Oellermann
Partner

Thomas Morineaux
Director

Stuart Mitchell
Director

Adam Gogarty
Director

Edd Hardy
Director

AlixPartners

Get in touch

Chris Rollo
Global Co-Leader, Technology Practice
crollo@alixpartners.com

Gökhan Öztürk
Global Co-Leader, Technology Practice
gozturk@alixpartners.com

Paula Walworth
Partner, Performance—Digital Innovation
pwalworth@alixpartners.com

Oli Freestone, Ph.D.
Director, Technology Practice
ofreestone@alixpartners.com

About us

For more than forty years, AlixPartners has helped businesses around the world respond quickly and decisively to their most critical challenges—circumstances as diverse as urgent performance improvement, accelerated transformation, complex restructuring and risk mitigation.

These are the moments when everything is on the line—a sudden shift in the market, an unexpected performance decline, a time-sensitive deal, a fork-in-the-road decision. But it's not what we do that makes a difference, it's how we do it.

Tackling situations when time is of the essence is part of our DNA—so we adopt an action-oriented approach at all times. We work in small, highly qualified teams with specific industry and functional expertise, and we operate at pace, moving quickly from analysis to implementation. We stand shoulder to shoulder with our clients until the job is done and only measure our success in terms of the results we deliver.

Our approach enables us to help our clients confront and overcome truly future-defining challenges. We partner with you to make the right decisions and take the right actions. And we are right by your side. When it really matters.