accenture

Generating impact

Harnessing frontier AI capabilities to unlock frontline productivity and growth in the UK

Contents

Foreword

Matt Prebble

CEO of Accenture in the UK and Ireland

Every day, I see what happens when organisations in the UK commit seriously to reinvention: not more pilots, not more experimentation, but fundamental change in how they operate. The results can be remarkable. And they point to a bigger truth about the UK economy. Our problem has never been a lack of capability. It has been a failure to convert capability into impact at scale.

The UK has extraordinary strengths: world-class universities, leading businesses, deep pools of talent and a long tradition of innovation. Yet for too long, we have paired those strengths with weak productivity, fragmented investment and a hesitancy about growth. We do not need to invent a new foundation. We need to execute far more boldly on the one we already have.

AI gives us a chance to change that trajectory.

As we showed in our 2024 report, the UK stands to gain more from AI than any other G7 economy. And recent advances, particularly in agentic AI, have moved the technology much closer to practical deployment at scale. The real opportunity now is not simply to adopt new tools, but to put AI to work in ways that raise productivity on the frontline, reshape how organisations operate and create new sources of value.

That shift is already beginning. Individuals are adopting AI quickly and weaving it into their daily routines. Organisations are investing, building capability and learning by doing. But the full value is still not being captured, because too often AI is being applied at the margins: improving isolated tasks rather than transforming workflows, decision-making and performance end to end.

That is why the barriers to progress are no longer primarily technological. They are organisational and human. They lie in how work is structured, how decisions are made, how people are equipped, and how trust is built. This is, above all, a leadership moment. Technology creates the possibility. Leadership turns possibility into impact.

The organisations moving furthest and fastest understand this. Their leaders are engaged directly. They are not treating AI as a side project or delegating understanding downwards. They are using it themselves, asking harder questions of their businesses, and driving the redesign of processes, roles and ways of working. That is how value is captured: not through isolated experimentation, but through deliberate reinvention.

And none of this reduces the importance of people. In fact, it does the opposite. AI will increase the premium on judgement, adaptability, learning and trust. The UK now faces a once-in-a-generation reskilling challenge, as millions of people move into roles and industries reshaped by AI, automation and digital technology. If we lead this transition well, we have the chance to improve not only productivity, but opportunity: creating growth that is broader-based, more inclusive and more sustainable. If we get it wrong, we risk leaving too many people behind.

So this is the challenge in front of us. We must move beyond pilots and prove that technology-led reinvention can deliver growth at scale. We must connect innovation to investment, investment to execution, and execution to outcomes that people can see and feel. Above all, we must act with greater confidence and greater urgency.

The opportunity is real. The technology is available. The talent is here. What matters now is whether we have the ambition to lead.

Authors

This report was a collaborative effort between our AI & Data team based in the UK, supported by our Research team and assisted with AI:

Chris Lane

AI & Data Lead—
UK and Ireland

Kayur Rughani

Managing Director,
AI & Data

Bella Thornely

Managing Director,
AI & Data

Suhail Kapoor

Managing Director,
AI & Data

Nick Tate

Managing Director –
Talent & Workforce Lead

Mark Farbrace

Managing Director –
Gen AI Lead

Thomas Niven

Managing Director,
AI & Data

Ali Shah

Managing Director –
Responsible AI Lead

Fernando Lucini

Managing Director,
AI & Data

Mike Moore

Research Director—
UK and Ireland

Introduction

AI presents a step-change opportunity for the UK economy, for organisations and for individuals. But capturing it requires more than deploying the technology. Building on our 2024 report, Generating Growth,[1] we surveyed more than 450 UK executives and 1,800 employees, updated our economic modelling and analysed around 30 million job postings. The findings reveal a system under tension: rapid individual adoption, uneven organisational progress and an economic prize that remains largely uncaptured.

This report examines where the UK stands, what is holding organisations back and what it takes to bridge the gap between AI's potential and its performance. It sets out a framework for reinvention across five dimensions, with practical recommendations for turning frontier AI capability into frontline productivity and growth. Organisations that treat AI primarily as an efficiency tool risk misreading the moment: the deeper question is how AI changes who wins in your industry, not just how cheaply you can operate.

1. The state of AI

AI capability: Advancing faster than its economic impact

In our 2024 report, Generating Growth, we examined the economic potential of generative AI for the UK and argued that realising it would depend on organisations redesigning workflows, building skills and strengthening trust. Since then, adoption has accelerated and the technological frontier has continued to advance. As generative AI becomes embedded in daily use, attention has shifted to agentic systems.

The most important technological development since our last report is the emergence of agentic AI: systems that do not simply respond to prompts but can interpret goals, plan multi-step actions, interact with enterprise systems and execute autonomously within defined boundaries. This shifts AI from a tool that assists individuals to one that can execute and coordinate work. The viral rise of tools such as OpenClaw (an open-source AI agent that connects to a user's messaging apps, files and services to execute tasks autonomously around the clock) offers a glimpse of where this is heading: AI that acts as a persistent, proactive presence rather than a reactive assistant. The length and complexity of tasks these agents can perform autonomously has approximately doubled every seven months.[2]

This expands the economic scope of AI significantly. The share of UK working hours that AI could enable has risen from 47% to 82% since our 2024 report. We already see evidence that employees are using AI for tasks that account for 21% of their working hours in the UK (Figure 1).

Whether this translates into realised productivity gains depends on what organisations do next. But in capability terms, the boundary of what AI can do has moved decisively forward.

As Marc Warner, CEO of Faculty and CTO of Accenture, noted, "Even if we froze the technology where it is today, it would still have a profound impact".

Figure 1: More than four in five UK working hours could be reinvented with AI

Share of work time in scope of reinvention with AI
UK labour force. Accenture industries based on ISIC sectors

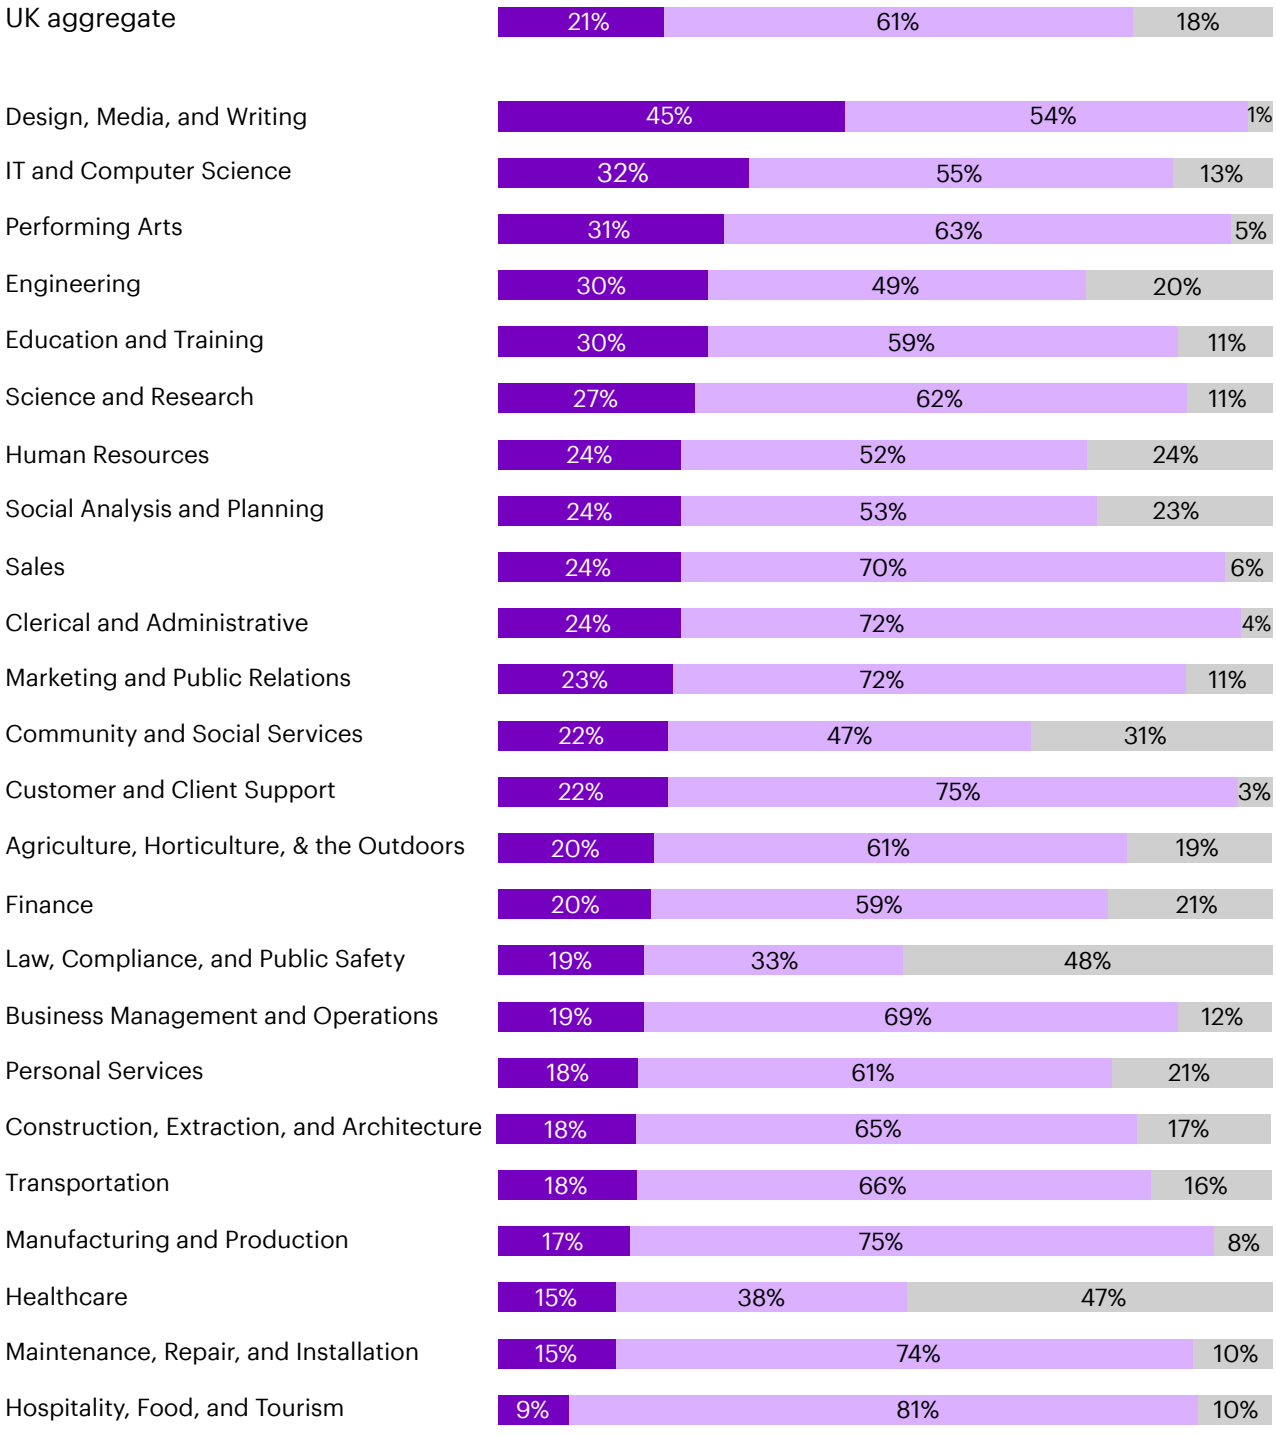

Notes: Accenture Research based on ONET task data and ONS employment data. ISCO-1 occupations ranked by share of work time with strong AI evidence (highest to lowest)

Employees: Both keen and cautious

Across the UK workforce, AI adoption has accelerated rapidly. Almost one in five (18%) employees now use generative AI tools in their work on a daily basis—triple the rate in 2024.

The perspectives of the workers we surveyed have matured, too: The sci-fi-inspired fears of robot revolutions in 2024 have largely given way to more practical, evidence-based concerns about quality, reliability and usefulness, alongside a growing awareness of environmental impact. Positive sentiment has shifted in the same direction, from hype-driven expectations to more grounded views of how AI can improve work (Figure 2).

But much of this adoption is happening ahead of, or outside, organisational decisions.

Only 39% of employees say they use AI tools provided by their employer; a further 24% have sourced tools independently, signalling an ongoing "shadow AI" phenomenon.

Figure 2: Attitudes to AI are shifting from hype and fear to practical concerns

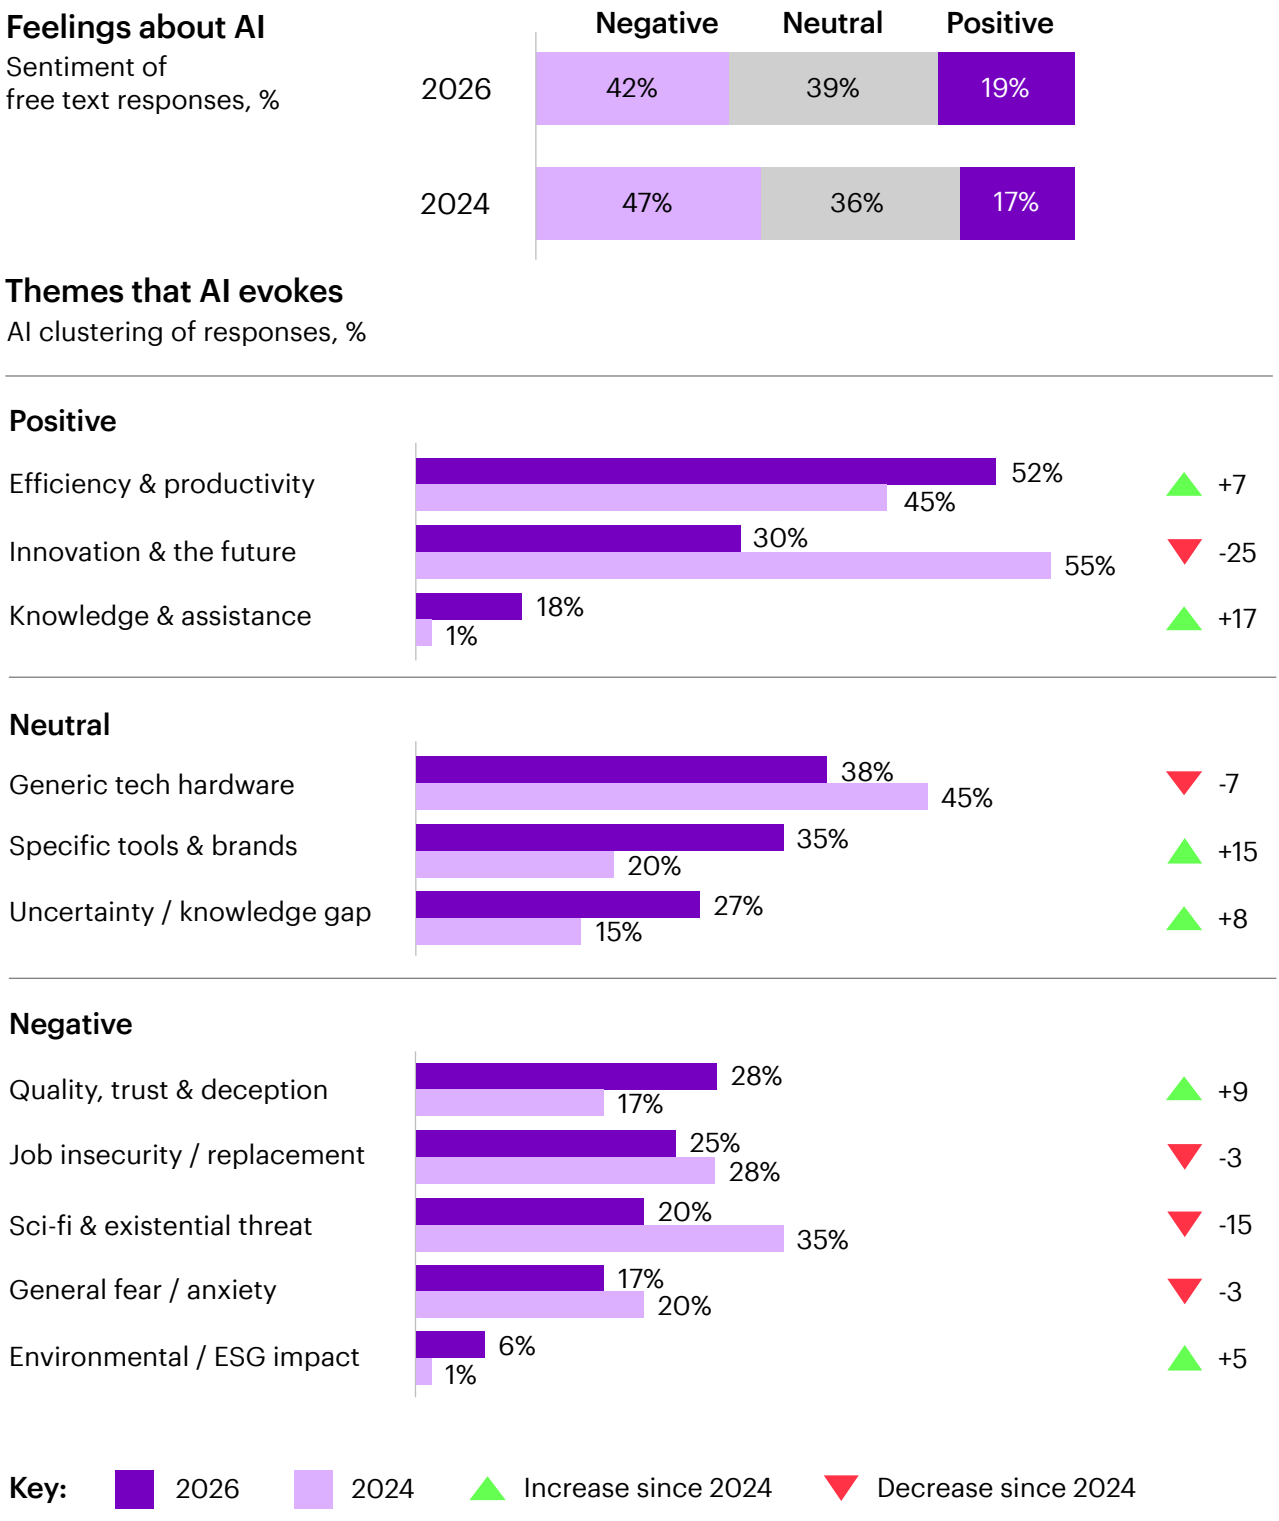

Note: We asked 2085 individuals this open-text question: "When you hear the word "Artificial Intelligence (AI)", what is the first thing (e.g., word, phrase, emotion) that comes to mind?" Analysis of free text responses was conducted using an LLM, to evaluate sentiment and generate thematic clusters. We asked this question in our previous study, Generating Growth, allowing for analysis of how sentiment and themes have shifted over time.

AI is already reshaping demand for skills. We analysed more than around 30 million UK job postings listed since the launch of ChatGPT. Demand for routine, structured cognitive skills is declining. AI-related skills are the second-fastest-growing category, but demand has also increased for skills rooted in people management, judgement, compliance, quality assurance, control and domain-specific process execution (Figure 3). These are often categorised as soft skills, but that framing underestimates what they represent. They reflect a shift in how work is defined: away from optimising tasks within today's workflows, and towards redesigning and governing how those workflows operate. As AI takes on execution, the human role expands towards judgement, oversight and system-level thinking across an evolving frontier of work.

Figure 3: Skills that complement AI are in growing demand

Change in the number of 1,821 UK occupations with job postings mentioning each skill subcategory
2023-2025[1]

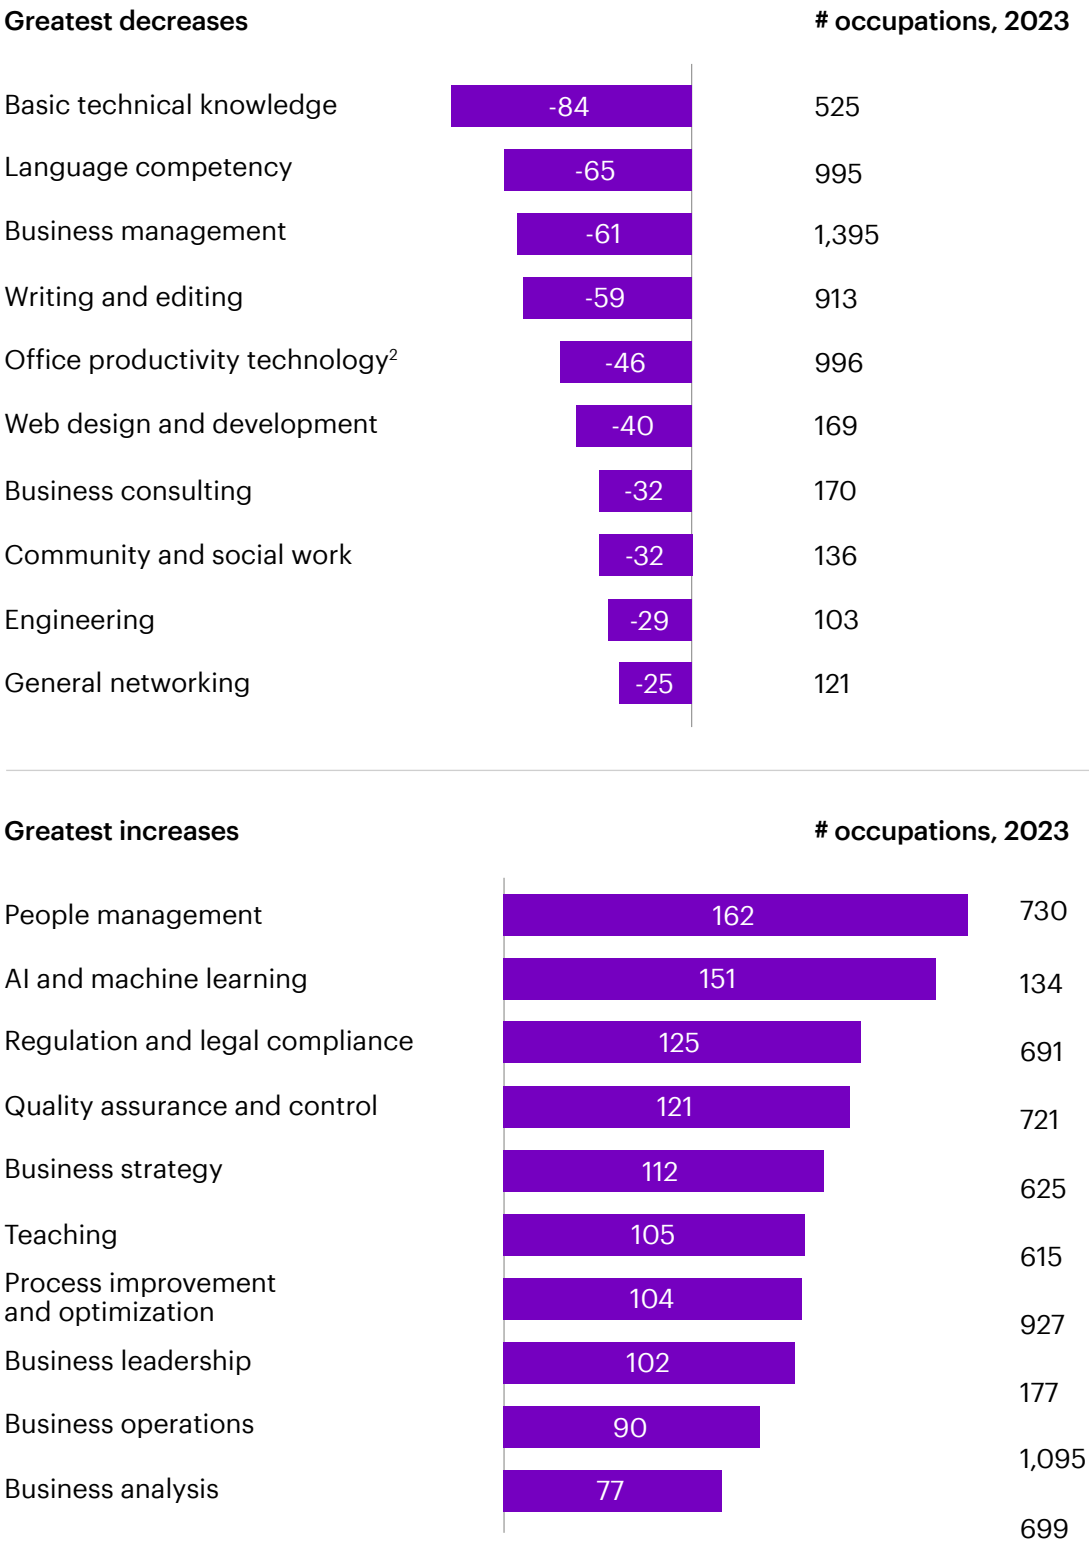

Source: Accenture Research analysis of Lightcast data. 1) At least one skill associated with the subcategory is listed in >5% of job postings for a given occupation. 2) Decline in this subcategory is associated with office software such as word processing tools and spreadsheets.

People recognise this shift, and its implications for their jobs, but they are prepared to respond and engage positively when supported. In our employee survey, the capabilities expected to rise in demand as AI spreads are communication and collaboration, critical thinking and leadership and change management. Almost one in three people (31%) expect their job either to be unrecognisable or to disappear completely by the end of the decade, twice the share who held that view just 18 months ago. Of those, 79% say they are likely to reskill and 55% say they are likely to change occupation.

The risk is that organisations are not yet providing enough clarity or support to facilitate this transition at scale. Almost one-third (32%) cite change management and workforce transition as one of their most significant AI-related skills gaps. Just 26% have conducted a skills audit to assess AI's impact on roles. Meanwhile, 27% do not provide AI-related training at scale, and only 30% are investing in reskilling and redeployment pathways for at-risk roles. This gap may be further widened by shifting leadership assumptions. In 2024, only one-third of executives believed AI would shrink the national workforce. Today, nearly half (49%) expect it to reduce national employment over the next decade, even as 70% expect it to drive higher productivity. If leaders assume displacement is inevitable, the incentive to invest in workforce transition weakens. Translating AI-driven efficiency into inclusive growth through reskilling, job creation and new forms of work remains a key challenge.

The implications are sharpest at the bottom of the career ladder. In 2024, 40% of executives expected AI to increase demand for entry-level roles. That has collapsed to 15%. The share expecting reduced demand has risen from 22% to 37%. If this expectation translates into hiring behaviour, AI risks hollowing out the traditional entry-level pipeline, with significant consequences for social mobility, equity and early-career skills development.

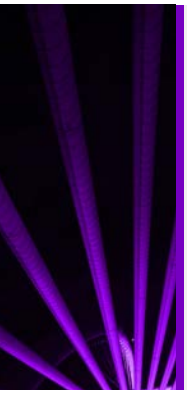

Almost one in three people (31%) expect their job either to be unrecognisable or to disappear completely by the end of the decade, twice the share who held that view just 18 months ago.

Organisations: Changing more slowly than their people

Organisational change is lagging behind. Employees are adopting AI at the level of individual tasks, but organisations have not yet redesigned the systems and workflows around them. Only around a quarter of employees say a major process in their team has been redesigned around AI in the past year. Without that redesign, productivity gains remain localised and fail to translate into enterprise performance.

25%

of executives say that if their organisation stopped using AI, some critical processes would slow down significantly.

This is not to say that organisations are standing still. A quarter (25%) of executives say that if their organisation stopped using AI, some critical processes would slow down significantly; a further 3% say some core operations would fail. Some banks are beginning to reclassify AI as "core infrastructure."[3] In functions such as transaction processing, fraud detection and supply chain management, AI is now embedded in how the business runs, not just in how individuals work within it.

But embedding AI in pockets is not the same as a systemic transformation. Siloed implementation creates a two-speed organisation, where some functions operate with AI-enabled productivity, while others work as they always have. Without enterprise-wide redesign, gains in one part of the business rarely compound into performance at the level of the whole.

The same pressure is building from the other direction. More than a quarter of UK consumers have already used AI tools to search, compare or buy on their behalf, and a further 29% say they would delegate the purchasing process to an agent if the tools were reliable. Yet only 21% of executives expect shopping agents to have significant influence on consumer decisions over the next three to five years. Consumers are starting to move first, while many organisations are not yet operationally ready for what that shift will mean.

The economic returns reflect this unevenness. Almost half of executives (46%) say AI has had little to no positive impact on profit and loss. Nearly a third (31%) say stopping AI tomorrow would have no material effect on their operations at all. Just under half (44%) believe at least some of their AI budgets are wasted, and nearly one in four (23%) say they cannot estimate whether their AI spending is effective at all.

Part of the problem is where organisations are looking for value. Much of AI deployment so far has focused on the denominator of productivity: improving efficiency. And twice as many executives see cost reduction as the greatest future benefit of AI rather than revenue growth. This stands in contrast to workers, who most commonly report using AI to enhance the quality of their output.

When everyone can automate routine processes, differentiation shifts elsewhere. The more enduring economic impact may therefore come from the numerator: growth.

Our analysis of AI's potential economic impact across 17 industries found that the opportunity for revenue uplift is more than twice the combined potential for labour cost avoidance and labour cost savings.

Figure 4: AI's greatest impact may come from revenue growth, not cost reduction

Potential economic impact of AI by industry
Full-maturity (upper-bound) annual potential of digital and physical AI agents for the average UK G3K company, by industry (n=81)

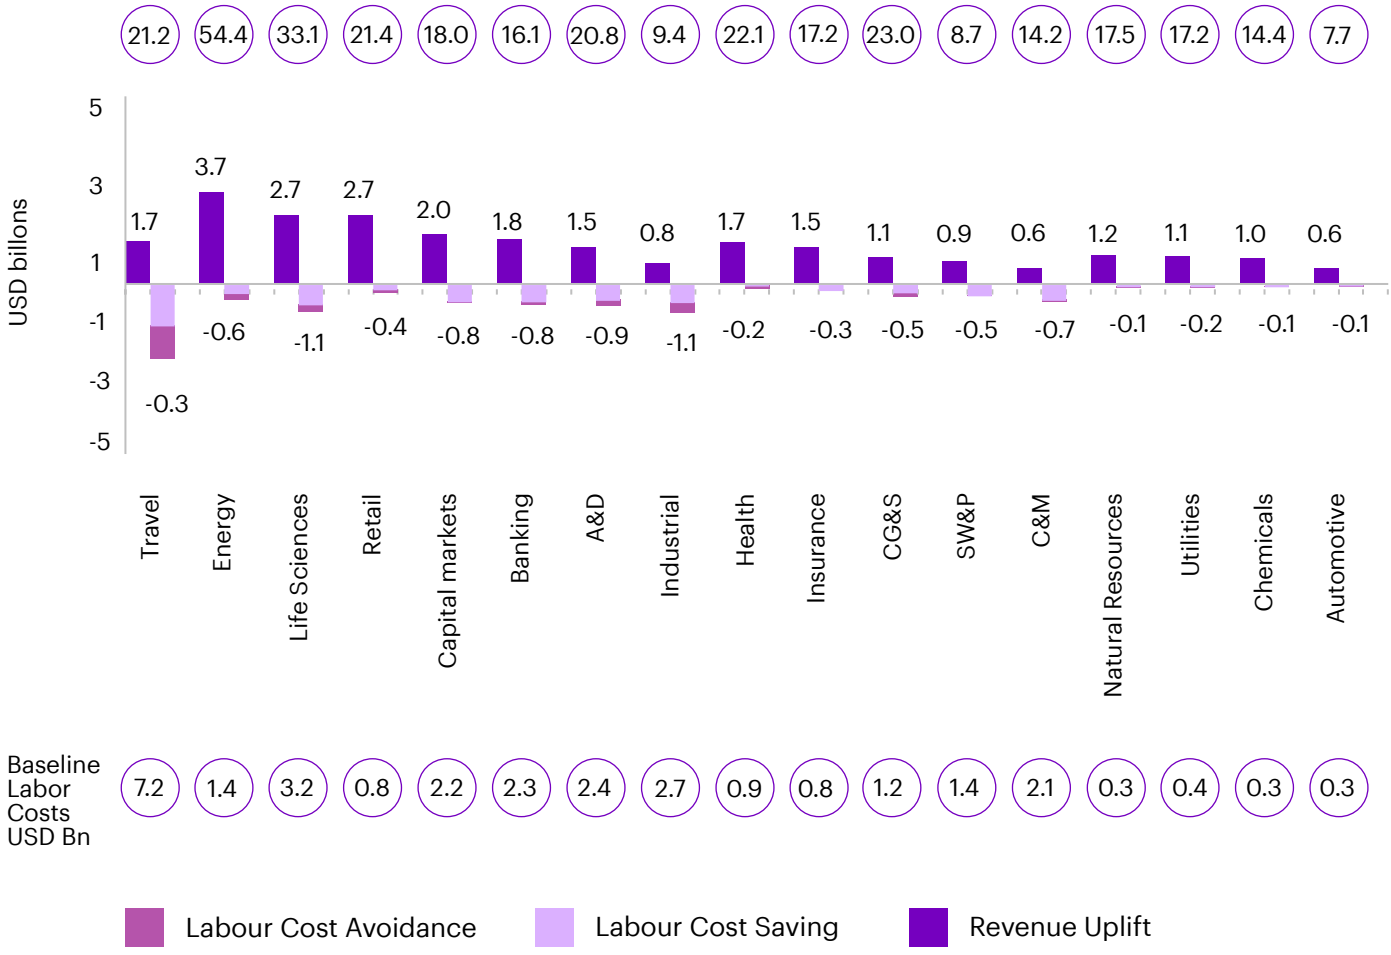

Source: Accenture Research analysis based on data from S&P Capital IQ, US ONET, US Bureau of Labor Statistics, International Labor Organization, and Lightcast. Baseline: 2024 revenues and labour costs for the average UK G3K firm in each industry, except for Automotive and Health for which we use a matched equivalent sample.

The economy: AI's promise has not yet appeared in productivity data

The stakes extend beyond individual firms. In her Mais Lecture in March 2026, UK Chancellor of the Exchequer, Rachel Reeves committed the UK to achieving the fastest AI adoption in the G7, backed by significant new investment. And the Office for Budget Responsibility now incorporates assumptions about AI-driven productivity growth into its medium-term fiscal forecasts.[4] If those gains materialise, they strengthen the UK's fiscal position. If they do not, weaker growth and lower revenues would tighten constraints on public spending. AI's productivity impact is no longer just a corporate performance question, it is a matter of national economic resilience.

Yet aggregate productivity data does not yet show a clear AI-driven acceleration. This is not unprecedented: General-purpose technologies from electricity to the internet historically required sustained complementary investment before delivering durable gains. But it underscores the urgency of moving from adoption to reinvention.

In her Mais Lecture in March 2026, UK Chancellor of the Exchequer, Rachel Reeves committed the UK to achieving the fastest AI adoption in the G7, backed by significant new investment.

2. The challenges still to overcome

In Generating Growth, we identified three structural tensions that explain why organisations struggle to turn AI capability into economic value: a delivery gap, a skills gap and a trust gap. Eighteen months on, none has been fully resolved, and advances in technology, particularly the emergence of agentic AI, have sharpened each of them.

The delivery gap persists. More than half (58%) of executives say their organisation is not ready to integrate AI agents with core enterprise systems, with legacy architectures continuing to block the move from pilot to production. Many companies have yet to scale even more established forms of AI: 32% are still piloting predictive AI and machine learning, and a further 26% are not using it at all.

Workforce readiness remains uneven. Only 7% of executives say their workforce is fully prepared for agentic AI. Yet more than half of employees (54%) say they expect to reskill in response to AI, signalling that many already see significant changes coming to their roles.

The trust gap endures. Compared with executives, almost twice as many employees expect a "significant decrease" in jobs as a result of AI. Data security (41%), reliability (32%) and trust and user acceptance (29%) remain the obstacles to AI adoption that executives cite most often (Figure 5). Over two-thirds (68%) of employees say they are uncomfortable with AI agents making higher-impact decisions.

Agentic AI also introduces new technical constraints that fall into three areas. Reliability: Longer, multi-step processes accumulate errors and models struggle to evaluate their own accuracy. Control: Specifying objectives precisely and ensuring they are interpreted correctly is harder than it sounds, and as autonomy increases, so does the cost of misalignment. Coordination at scale: Multi-agent systems add orchestration complexity, communication breakdowns and cascading errors, while production-grade deployment requires monitoring tools, fallback mechanisms, integration layers and audit trails that most organisations have not yet built.

Agentic AI introduces new technical constraints that fall into three areas: reliability, control and coordination at scale.

Figure 5: Risk, trust and quality concerns are now the primary barriers to scaling AI

Barriers to AI adoption
Proportion selecting this as a barrier, executive sample (2026 n=395, 2024 n=1085)

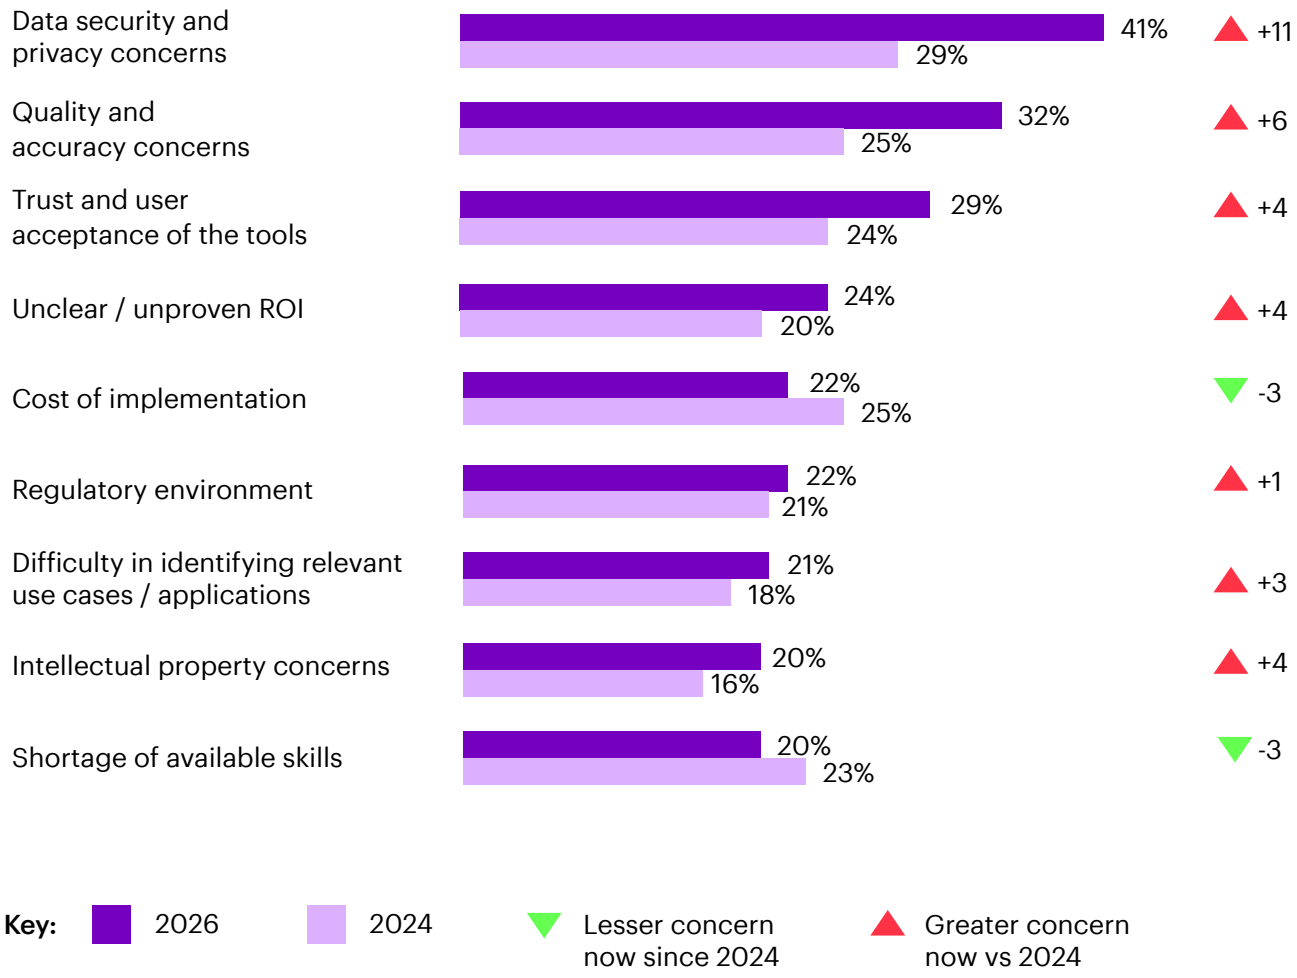

Question wording: What are the biggest barriers preventing your organisation from maximising value from AI and agentic AI? Note: Only options selected by 20% in 2026 or more are shown. The other categories were: Technology platform not ready to support scaling; Intellectual property concerns; Availability and/or interoperability of data; Lack of time to work on development; Lack of cross-organisation coordination of execution; Lack of executive alignment; None of these; Other.

3. The art of reinvention

Crossing the productivity gap

Closing the gap between adoption and impact requires reinvention across five dimensions in sync:

01 

Strategy

Using AI to advance your business strategy, while asking whether AI is changing what that strategy should be.

02 

Work

Redesigning workflows so AI can execute and coordinate processes end-to-end.

03 

Workforce

Building confidence and competence in AI, whilst building the right operating model to drive growth.

04 

Digital Core:

Connecting data and systems so AI can act.

05 

Safety and Security

Establishing governance and guardrails so AI systems can be trusted and scaled with confidence.

Improving performance in just one of these dimensions increases the likelihood of scaling AI by 2.2x. Acting across all five increases it more than fourfold.

To understand where organisations stand, we classified them into four stages of AI maturity:

Experimenters
(27%) use AI primarily for pilots and isolated tasks

Adopters
(16%) deploy AI tools in multiple teams but struggle to move from prototypes into production.

Integrators
(48%) have AI in production and are beginning to embed it in workflows, but are not yet delivering material enterprise-level value.

Scalers
(10%) have AI embedded in core operations with measurable improvements in financial performance and operational efficiency

The analysis highlights where organisations get stuck. Early-stage organisations are most often held back by workforce readiness. As deployment broadens, the bottleneck moves to the digital core. Among Scalers, the limiting factor is increasingly strategy and how organisations prioritise and coordinate a scaled portfolio. The elements most strongly associated with reaching Scaler status are work and workforce: The greatest opportunity lies in how organisations redesign work and prepare people to operate within AI-enabled systems.

The scale of change required runs deeper than most organisations have acknowledged. Nearly one-third of executives believe minimal or limited process change will be sufficient to benefit from agentic AI. But adding AI to existing processes produces incremental gains at best. Like the digital wave before it, AI demands simultaneous change across strategy, processes, people, data and governance. Organisations that treated digital transformation as a technology purchase rather than an organisational one consistently underdelivered. Those who approach AI the same way will reach the same conclusion.

01

Strategy: Start with the outcome

Where organisations are today

Organisations often separate AI strategy from business strategy, letting technology momentum drive decisions and spawning pilots that fail to shift the metrics that matter most. Investment decisions are shaped by what can be justified quickly, particularly cost reduction, biasing organisations toward incremental efficiency over more transformative opportunities. Often, business units push to adopt AI before value is defined, while technology teams are asked to justify value retrospectively. In contrast, Scalers are more than five times as likely to have embedded processes for identifying where AI can deliver the greatest impact. Without this discipline, AI remains a collection of initiatives rather than a driver of strategic change.

The risk of getting this wrong extends beyond wasted investment. AI is beginning to redraw the competitive boundaries of many industries, and organisations that focus primarily on efficiency while others use AI to redefine their markets may find the gap difficult to close. Moving too slowly carries its own cost, one that compounds over time and becomes harder to recover from than the cost of a poorly executed initiative.

To realise value, organisations must move from AI-led experimentation to strategy-led deployment:

Think big. Build in steps.

Define a bold strategic destination tied to core business outcomes, then sequence initiatives that build toward it. Every AI investment should be linked to a top-level KPI. The nature of the problem should determine whether agentic AI is the right approach: agentic systems are best suited to objectives requiring multi-step reasoning, high volumes of structurally repetitive but contextually variable tasks and coordination across multiple systems or teams. Customer service illustrates the progression: from generative AI tools that help agents retrieve knowledge and draft responses, to agents that classify issues and resolve low-complexity cases autonomously, to full process redesign where AI detects issues proactively and communicates directly with customers while humans move into oversight and exception management.

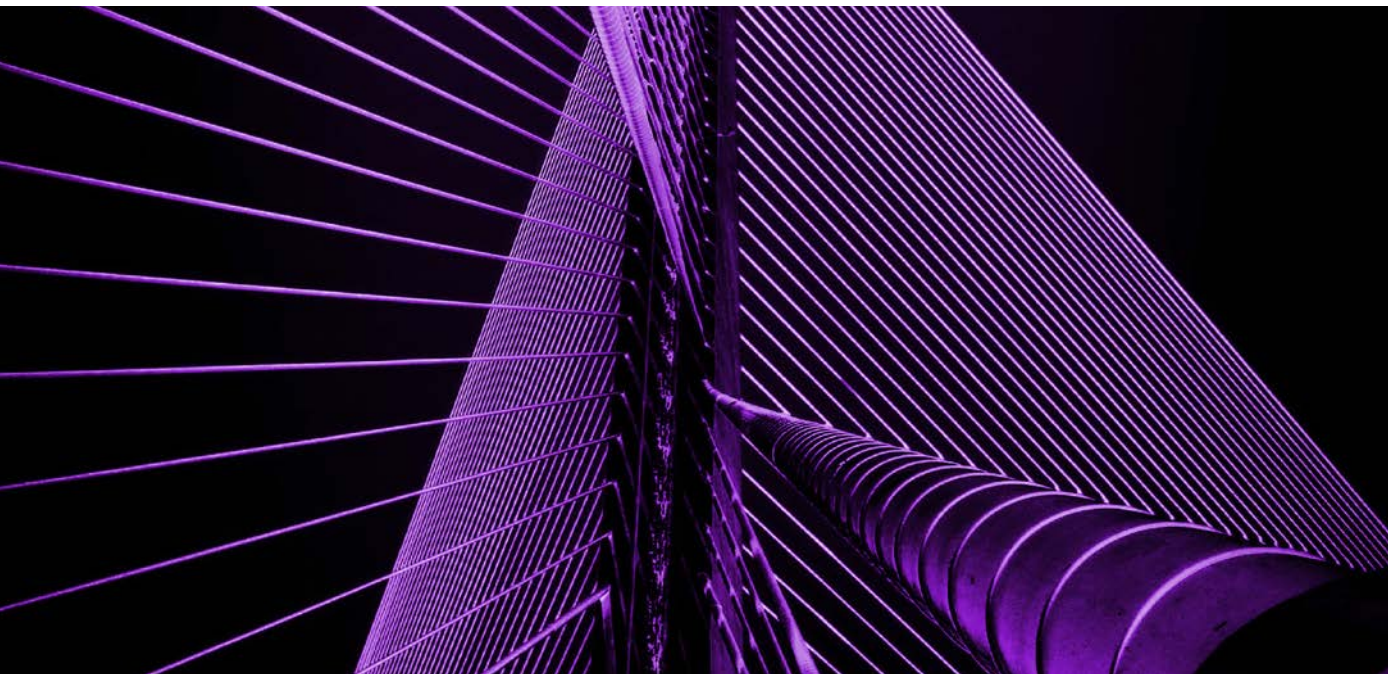

Update your investment strategy, not just your AI strategy.

Most finance teams still assess AI through static, upfront business cases focused on cost savings and labour reduction. That lens will never justify the scale of transformation AI makes possible, and it struggles to capture the real economics of agentic systems, where value is realised over time through improved decisions, faster execution and new sources growth. Instead, organisations need to adopt hypothesis-led investment models, with funding staged and tied to defined outcomes. At present, a significant share of AI-driven productivity gains appears as freed capacity rather than direct cost reduction. Unless leaders deliberately channel that capacity into growth—through faster product iteration, new offerings or smarter capital allocation—productivity gains remain trapped as efficiency improvements.

The unit of measurement itself needs to change. Existing frameworks measure what individuals do: task completion, headcount, activity volumes. When humans and agents share a workflow, the relevant metrics become end-to-end (like cost or speed per decision or revenue impact) alongside mechanisms to track redeployed capacity. In general, this all calls for closer integration between finance, technology and the business. Finance can no longer be a gatekeeper of spend, but a partner in scaling value, providing the discipline and flexibility needed to turn AI from isolated initiatives into sustained performance.

Build the AI-native version of your business.

A new generation of AI-native companies are reaching $100 million in recurring revenue within 12 months with teams of 20 people. For established organisations, the question is not simply how to keep pace, but what it would mean to build a version of their business unconstrained by existing architecture, processes and organisational habits. The most practical way to answer that question is to try: take a business unit or service line and build it from first principles, as if starting with a blank page. That experiment will reveal not only what an AI-native model looks like, but what it would genuinely take to transform the existing organisation: the gaps in data, the constraints in governance, the depth of process redesign required. Existing customer relationships, proprietary data, regulatory licences and brand trust remain real advantages, but only if they are actively carried into the new model rather than left behind in the existing one.

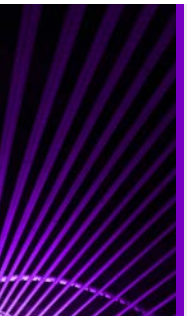

Understanding what the greenfield version requires is the clearest guide to what the brownfield transformation actually demands.

02

Work: Redesign workflows end-to-end

Where organisations are today

More than four in 10 working hours are still tied up in coordination and transaction costs. When AI is introduced into these fragmented processes, it tends to optimise individual steps rather than resolve underlying bottlenecks. Indeed, half (53%) of employees say AI helps them do the same work faster rather than changing how work is done. In contrast, organisations reporting significant or transformational process redesign are five times more likely to report substantial P&L uplift than those making incremental changes.

To realise value, organisations must move from task automation to workflow reinvention:

Zero-base your workflows.

Before introducing agents, organisations should resist the temptation to map existing workflows and automate them as they stand. The more valuable discipline is to ask what the workflow would look like if designed from first principles with AI in mind, starting from the outcome it needs to deliver, not from how work currently flows.If the process is fragmented or full of unmanaged exceptions, agentic AI will amplify the disorder rather than remove it. Standardisation and consolidation must come before agent deployment. Process mining and knowledge-mapping tools, increasingly powered by AI, can reconstruct how work actually flows at a speed and level of accuracy that was not previously possible. The goal is a workflow designed for what agents can do, not a digitised version of what humans currently do. An important discipline is treating agent deployment more like onboarding a new employee rather than deploying software: agents should be given clear “job descriptions” that define their scope, decision boundaries and escalation points.

Simulate, monitor and continuously adapt.

Once a workflow has been zero-based, the next step is to test it before deploying it in production. Digital twins of priority workflows enable leaders to simulate redesigns in low-risk environments before introducing change into live systems, testing alternative sequencing, decision rights and coordination models with the assumption that AI handles portions of execution. One global manufacturer built a digital twin of its production plants, collapsing what was a four-week physical validation process into three days without disrupting live production.

Unlike process redesign of the past, which produced a fixed blueprint and moved on, workflow design in an AI-enabled organisation is never finished. Agent capabilities improve, new integrations become possible and the boundary between humans and machines shifts continuously. Organisations that build reusable patterns for common workflows (standard ways of structuring decisions, handoffs and governance) are better placed to absorb these changes at pace, scaling what works rather than redesigning from scratch each time something shifts.

Let agents coordinate the workflow, not just complete the tasks.

At higher maturity, agentic systems manage coordination layers across workflows in real time. Rather than a single all-purpose agent, organisations deploy networks of task-specific agents coordinating through shared protocols such as Model Context Protocol (MCP) and Agent-to-Agent (A2A). When agent networks connect across organisational silos, linking demand signals to supply chain to finance to fulfilment, the result is a fundamentally different operating system in which decisions that previously required human coordination across functions can be sensed, evaluated and executed in real time.

03

Workforce: Build human–agent teams

Where organisations are today

Workforce readiness is low. Only 7% of executives say their workforce is fully prepared for agentic AI. Trust is the most immediate constraint: People trust AI only after building familiarity, and familiarity only comes from use. Yet many organisations deploy AI without transparently explaining what it does, how accountability works and how it affects individuals. Many leaders lack direct competence and confidence themselves. Executives at Scaler organisations spend more than 20 times longer each week using AI than those at organisations still in pilot stages. Among Scalers, 76% of executives say they are very confident in the quality of their AI training, compared with just 6% at Experimenter organisations.

Of all five dimensions, workforce readiness is the most common constraint for organisations at the earliest stage of maturity, and it remains the dimension most directly within leadership's control. Technology can be procured and processes can be redesigned, but the confidence, capability and culture of the people inside the organisation cannot be bought or deployed overnight.

To realise value, organisations must move from human-in-the-loop rhetoric to humans leading AI-enabled work:

Design the human–agent relationship. Not just the hand-off.

Organisations need to define how humans and AI work together inside workflows, establishing autonomy tiers as a starting point: which tasks agents handle independently, which require human approval and which remain human-led. These distinctions must be visible to everyone involved. As agents take on more coordination, the basic unit of delivery increasingly becomes a small, outcome-focused team of people supervising a larger network of specialised agents. A group of two to five people can already oversee dozens of agents running an end-to-end process. Decision rights also need to flatten. Approval chains designed for human processing speeds become bottlenecks when agents execute in real time.

Build the skills to create, supervise and innovate, not just use.

As agents take on execution, the nature of human work shifts from doing the work to specifying it well, delegating it effectively and judging the result. Subject matter expertise is the decisive advantage. Experts give better instructions, evaluate output faster and spot errors sooner. New roles are emerging to support this shift: AI supervisors who monitor agent performance, algorithm auditors who verify behaviour aligns with policy and workflow orchestrators who design and refine human-AI interactions. As routine coordination is automated, people are freed to focus on designing better processes, identifying new opportunities and experimenting with new ways of creating value. This is good news for the 29% of UK employees that possess the skill of creativity but do not use it at work—the highest rate of underutilisation of any skill we evaluated.

Measure and manage for evolving human-agent performance.

Existing performance frameworks were built for a purely human workforce. Fewer than half of executives say their organisation has updated performance measures for AI-enabled ways of working (43%), and only 40% say managers can assess performance fairly when outputs are partly AI-assisted. Employees are roughly twice as likely to say AI-assisted work is valued less rather than more, creating conditions for "quiet AI": While 64% of executives say employees are expected to be transparent about AI use, only 14% say they always disclose it. Employees who say they would need significant training are around 2.4 times more likely to hide their use of AI. Performance management in an AI-enabled world means measuring outcomes and cycle time rather than effort and inputs. It should create the conditions for and expectations of growth, while ensuring equitable access to tools and setting clear norms for disclosure and accountability.

64%

of executives say employees are expected to be transparent about AI use, only 14% say they always disclose it.

04

Digital Core: Connect systems so AI can act

Where organisations are today

For every aspect of the digital core we assessed, from cloud to data and security, most executives say they are not ready for agentic AI (Figure 6). Only 3% say all aspects of their technology and data stack are fully ready. Data remains a foundational constraint. Many organisations treat cloud migration as the end goal rather than a step toward a more complete data platform, leaving data products poorly owned, inconsistently defined and not easily consumable by humans or agents. Beyond data, the challenge is integration. Agents can only act when connected to enterprise systems through application programming interfaces (APIs) and secure interfaces; without clean integration architecture, autonomy remains theoretical.

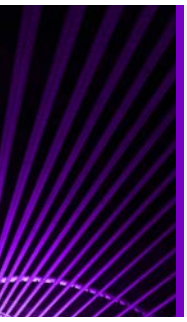

The digital core is a technology capability that brings together key components – like cloud, data, AI and security – to drive reinvention.

Figure 6. The technology and data foundations for agentic AI are not yet in place across most organisations

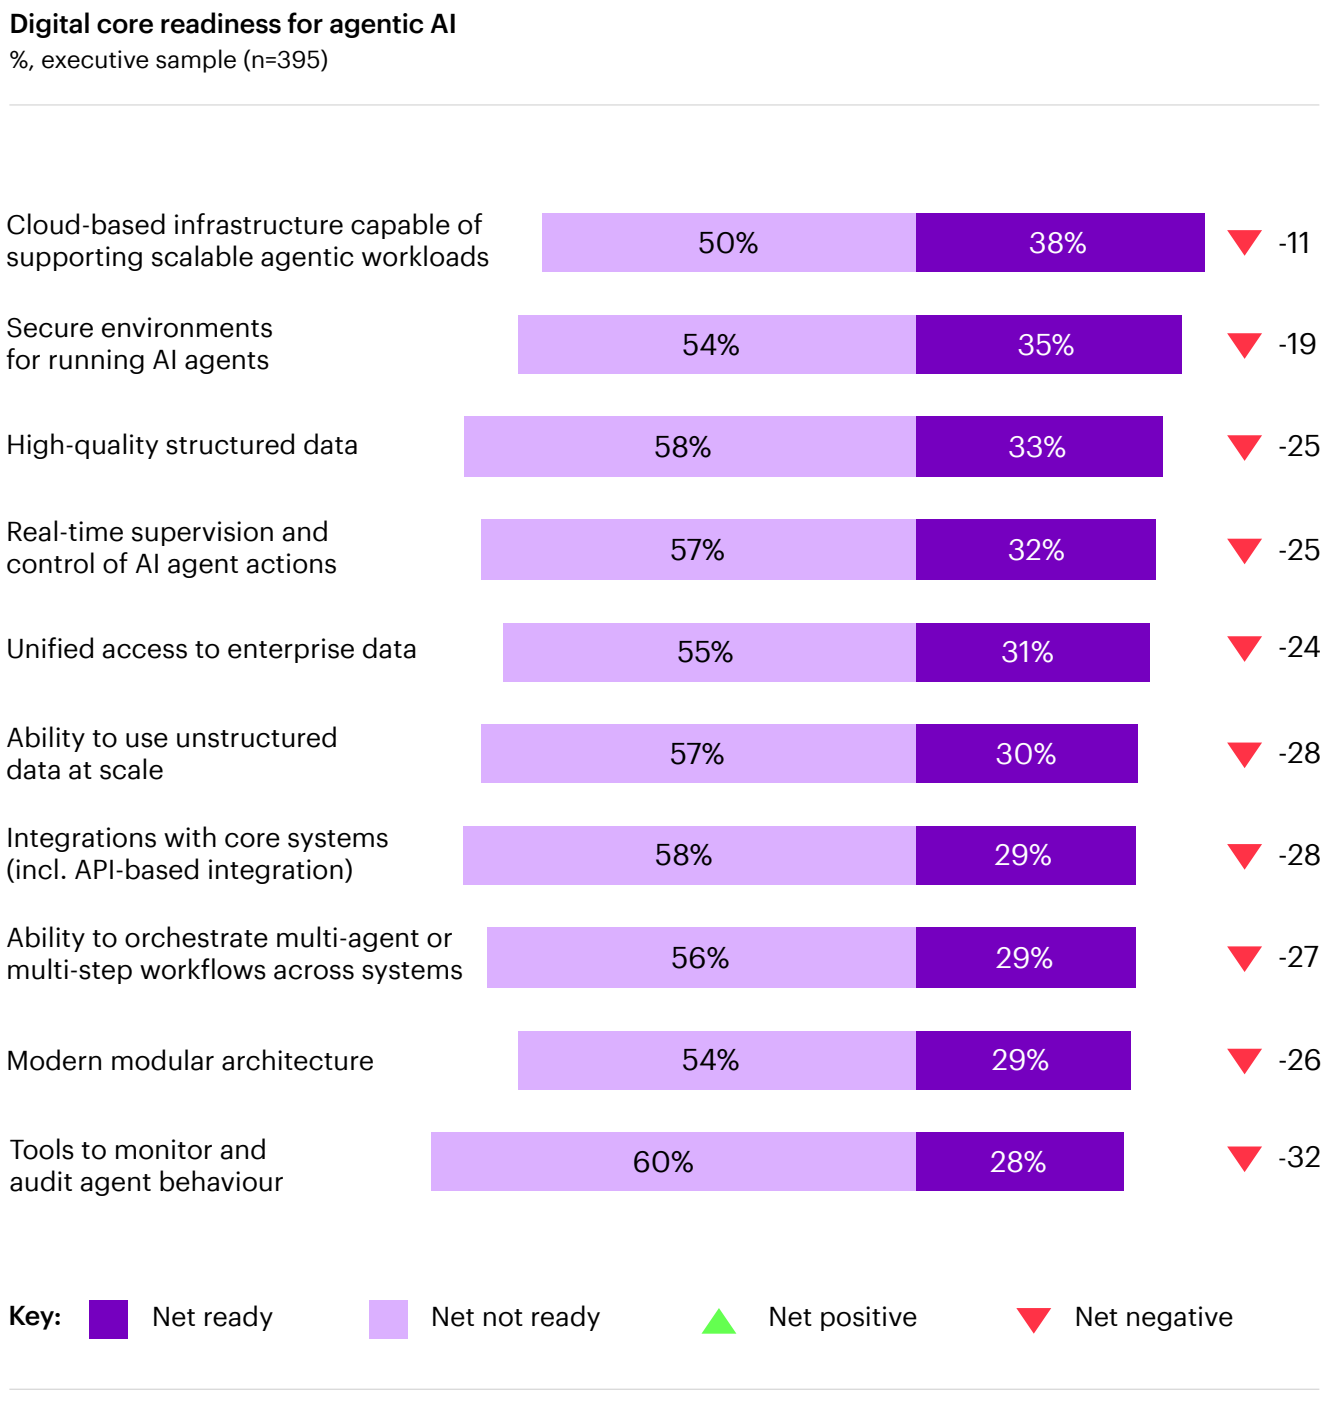

Question wording: How ready is your organisation across the following technology and data capabilities needed to support agentic AI? Answer options: "Not at all ready", "Early stage", "Partially ready", "Mostly ready", "Fully ready", "Don't know". "Net ready" is calculated as the aggregate of "Mostly ready" and "Fully ready." "Net not ready" is calculated as the aggregate of "Not at all ready", "Early stage", and "Partially ready." **Note:** Results shown exclude respondents saying "don't know." Numbers showing the net changes may not appear equal to the differences due to rounding.

To realise value, organisations must move from AI layered on systems to AI embedded within them:

Build the context layer agents actually need.

Agentic AI does not require a perfect enterprise data estate before it can deliver value. It does require trusted context for the workflow in question: the right process data, governed knowledge, clear business rules and enough history for the agent to act reliably. Organisations also need to capture the "indirect context" layer, the informal discussions and operational knowledge that informs human decisions but is currently lost or fragmented. The most effective approach is to curate high-quality data around the workflows that matter most and expand coverage progressively, building reusable knowledge and context layers that allow agents to work consistently across multiple processes.

Standardise the action layer.

Generative AI tools primarily read data: they summarise, draft and analyse. Agentic systems read and write. They retrieve information from one system, make decisions, update records in another, trigger actions in a third and maintain state across the sequence. Many organisations underestimate how much of the scaling challenge sits in this action layer: APIs, permissions, write-back rules, authentication and workflow orchestration. If each agent is wired into enterprise systems as a one-off, scaling becomes slow, expensive and difficult to govern. The digital core needs standard patterns for tool access, system integration and action execution

Operationalise oversight and adaptability.

Once agents begin operating inside live workflows, organisations need operational control: logging every significant agent decision, monitoring for drift from expected behaviour, tracing multi-agent interactions and establishing automated alerts when agents approach their autonomy boundaries. More advanced organisations are implementing graduated containment, monitoring systems that automatically constrain an agent's scope when they detect anomalous behaviour rather than waiting for human intervention. Model capabilities are improving on a time scale of months. The digital core needs to be modular enough to incorporate new capabilities, flexible enough to avoid lock-in and observable enough to absorb greater capability without losing control.

The concern that fast-moving capability cycles create unmanageable technical debt is real but overstated. Modular, API-first architecture limits lock-in, and AI tooling itself is increasingly being used to manage migration and integration work. The greater risk is the debt that accumulates from delay.

05

Safety & Security: Embed governance into the system

Where organisations are today

Most governance structures were built for a different phase of AI maturity, designed around the assumption that humans initiate decisions and systems execute predefined tasks with predictable outcomes. Agentic AI invalidates those assumptions. Roughly half of executives cite loss of human oversight or excessive autonomy (48%) and incorrect or harmful decisions (45%) as the biggest risks of agentic AI. But only 28% report having the tools to monitor and audit agent behaviour at scale, and 57% say they are not ready to supervise agents in real time. And across nine governance mechanisms we assessed, from autonomy limits to audit trails, fewer than one in four organisations on average have any of them fully operational. Organisations that are mostly or fully ready to monitor and audit agent behaviour are almost 10 times more likely to be Scalers than those that are not. This underlines that governance is one of the conditions that makes scaling possible in the first place.

28%

report having the tools to monitor and audit agent behaviour at scale, and 57% say they are not ready to supervise agents in real time.

To realise value, organisations must move from governance as policy to governance as an operational capability:

Decide where AI is allowed to act.

Establish clear autonomy thresholds, decision rights and escalation rules before deploying agents into production. Explainability must be defined operationally: what evidence must exist to explain how a decision was reached, who can access it and how it will be used in audit or regulatory review. Build governance into the workflow itself, with checkpoints, override rules and evidence capture defined before deployment.

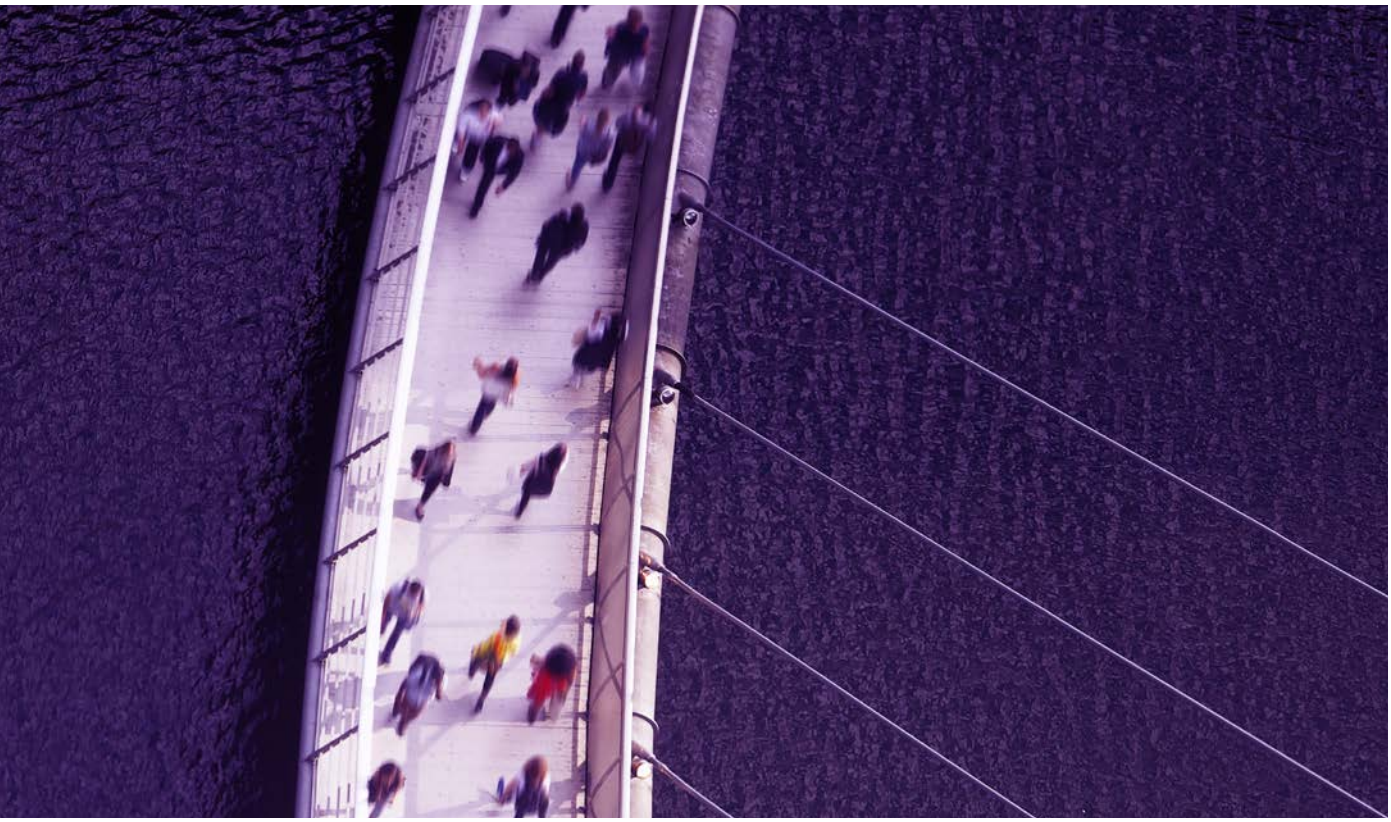

Be human-led: Empower both operators and overseers.

Employees inside redesigned workflows need clarity on when to intervene and how accountability works. Governance functions, including risk, compliance, audit and legal, need dedicated tools and capability to supervise agentic systems whose behaviour they may not yet know how to interrogate. Without focused investment in this group, human oversight remains a stated safeguard rather than an operational one.

Build guardrails into the system itself.

Governance cannot be retrofitted. Architectural guardrails, including policy gateways, traceability mechanisms and decision checkpoints, are the structural conditions under which autonomy remains operationally viable. Deploy agentic monitoring systems to govern agentic operations: purpose-built monitoring agents that observe production behaviour, detect deviations and initiate escalation workflows without waiting for manual review.

Treat agentic security as a governance discipline, not an IT perimeter problem.

An agentic system holds credentials, executes actions, calls external services and operates across organisational boundaries. Prompt injection can redirect system behaviour in ways that are difficult to detect. Credential misuse can grant access well beyond its intended scope. Security controls must be embedded at the architecture stage alongside governance controls. When executives select agentic AI vendors, security, privacy and compliance guarantees are the factors cited most often (42%).

Build the technical and operational harness that connects governance intent with agentic technical capabilities.

Organisations that scale agentic AI with confidence do so because they have established an integrated control plane that supervises AI behaviour across the enterprise: connecting monitoring, traceability, incident response and governance reporting into a continuous operational function. This requires a live inventory of all agentic systems with clear ownership and risk classification, continuous monitoring against safety indicators, defined incident management workflows, the ability to reconstruct any decision to audit standard and lifecycle governance maintained across the full operational life of a system rather than discontinued after deployment.

4. Next steps

Scaling AI requires crossing a series of capability thresholds. Moving from experimentation to adoption requires a clear strategy, workflows that have been redesigned from first principles and initial workforce readiness. Moving to integration requires embedded governance and processes rebuilt around what agents can do, and continuously adapted as those capabilities evolve. Reaching scale requires AI to become part of the operating model itself.

The question is no longer whether AI can deliver value. It is whether organisations can translate frontier capability into frontline productivity and growth. That depends on reinvention, not just adoption. And reinvention has to start somewhere.

The most common reason organisations stall is the assumption that everything must be in place before meaningful progress is possible. It does not. The organisations making the most progress pick a specific place to begin, build it properly, and use that experience to understand what genuine transformation of the wider organisation actually requires. For each of the five dimensions, the questions and immediate actions below are designed to help you do exactly that.

Strategy

Key questions

- Have you mapped how AI is likely to affect your organisation over the next two to three years, not just in efficiency terms, but in terms of how your industry's competitive contours are being redrawn?
- Are your AI initiatives aligned to your strategic vision, or have they accumulated as a series of independent pilots with no clear line of sight to your most important KPIs?
- Are you sequencing your AI investments to build capability toward a strategic destination, or approving them case by case as they arise?

Where to start

Pick one business problem that matters, something tied directly to a top-level KPI, and make that the anchor for your first AI investment. Everything else follows from that choice: which workflow to redesign, which people to equip, which data to prepare. The pace of change in AI capability is not a reason to wait; modular, well-governed foundations limit lock-in, and the cost of delay compounds faster than the cost of adaptation.

Work

Key questions

- Are your AI investments making existing processes faster, or are they redesigning how work gets done?
- Have you identified which workflows, if redesigned around agents, would move your most important performance metrics?
- Do you have a way to test workflow redesigns before committing them to production, and a plan to keep adapting them as agent capabilities evolve?

Where to start

Within the business unit or service line you have chosen, pick one workflow and create a dedicated space to redesign it from first principles with AI in mind. Before redesigning anything, establish a clear understanding of the inputs, outputs, constraints and risk management requirements. AI-powered process mining tools can accelerate this significantly. With that foundation in place, zero-base the process: design it as it should be, not as it currently is. Then simulate the redesign before committing it to production. The goal is not a perfect workflow but a working prototype that generates evidence and builds confidence. Getting the process design right is only half the challenge; bringing the people inside that workflow with you, and equipping them to operate in the new model, is equally critical to success.

Workforce

Key questions

- Do your people understand how their roles are changing, and do they feel part of the journey, with the support they need to develop the skills that will matter?
- Are your managers equipped to supervise AI-enabled work and assess performance fairly when outputs are partly AI-assisted?
- Are you creating the conditions for employees to use AI openly and productively, or are unclear expectations inadvertently encouraging them to hide it?

Where to start

As you redesign the chosen workflow, conduct a skills audit of the people inside it, understand what capabilities they will need to operate effectively in the new model, and close that gap before asking them to work in a new way. Set clear expectations about how AI should be used and disclosed. The experience those early adopters develop, what works, what requires human judgement and where the boundaries need to be drawn, becomes the most practical curriculum you have for bringing the rest of the organisation with you.

Digital Core

Key questions

- Do you have the data, integrations and access controls needed to support the specific workflow you want to redesign first?
- Have you sequenced your data and infrastructure investment around where AI will be deployed first, rather than attempting to modernise everything at once?
- Do you have a roadmap that builds your digital foundations in the sequence your organisation needs, rather than in the order that is easiest to approve?

Where to start

You do not need a fully modernised data estate to begin. Running alongside the workflow and workforce steps, identify the data, integrations and access controls needed to make your chosen workflow function reliably and get that slice right. Build modularly: what you put in place for the first workflow, the data pipelines, permissions and write-back controls, becomes the foundation for the next one. The goal is not a perfect platform for everything, but a working foundation for one thing that can be extended without starting from scratch.

Safety and Security

Key questions

- Have you defined what your AI agents are allowed to access, decide and do, and who is accountable when something goes wrong?
- Is your approach to risk calibrated to the actual risk profile of each deployment, or are uniform constraints slowing everything down equally?
- Is risk management embedded at every stage of your AI programme, proportionate to your organisation's risk appetite and inclusive of the risk of moving too slowly?

Where to start

Before deploying your first agent into production, define its autonomy limits, escalation rules and accountability structure explicitly. Assign a named owner for every autonomous workflow. Work with your risk, compliance and legal functions to agree proportionate guardrails for the first deployment, creating a reusable governance template rather than solving the problem from scratch each time. The goal is a governed space in which experimentation can happen safely, not a set of controls that prevent it from happening at all.

About the research

Employee survey

We partnered with YouGov to survey 2,085 employees across the UK and Ireland in February–March 2026. The sample includes 1,891 respondents from the UK and 194 from the Republic of Ireland. It is broadly balanced across genders and spans a wide age range, with about 30% under 35. We compared results to a survey of 3,752 UK workers fielded in July–August 2024.

Executive survey

We surveyed 510 business executives across the UK and Ireland in February–March 2026. The sample comprises 395 respondents from Great Britain, 58 from Northern Ireland and a further 57 from the Republic of Ireland, with a strong spread across the mid-market and large enterprise. We compared results to a survey of 1,085 UK executives fielded in July–August 2024. Unless otherwise stated, executive survey results shown above are reported for Great Britain (England, Scotland and Wales).

Executive survey analysis: Maturity segmentation, capability indexing and predictive modelling

We classified organisations into four AI maturity levels based on three signals: whether AI is embedded in operations, whether it is delivering substantial measurable value and whether any form of AI has been scaled across multiple business areas. Level 4 “Scalers” meet all three criteria. Levels 1–3 capture progressively earlier stages, from piloting through to AI in production but not yet at scale.

To assess organisational readiness, we created a five-dimension index covering Strategy, Work, Workforce, Digital Core and Safety and Security. Each dimension is built from survey questions converted to a consistent 0–100 scale.

We then built a predictive model to estimate how performance across the five dimensions relates to the likelihood of reaching Level 4, simulating the effect of improving one or all dimensions by one standard deviation. To identify where organisations get stuck, we ran a binding constraint analysis for each maturity level, identifying the dimension where organisations scored lowest relative to their performance across the others. This revealed that Workforce is the most common constraint for Experimenters, Digital Core for Adopters and Integrators and Strategy for Scalers. Finally, we used permutation importance to identify which dimensions most strongly differentiate Level 4 organisations from the rest.

Estimating impacts of AI on work hours and economic value by industry

We estimate the proportion of working time in scope for AI augmentation using a task-based framework that integrates occupational data, AI capability mapping and empirical usage evidence. Tasks are assessed against critical human inputs (judgement, unpredictability, accountability, implicit knowledge and risk) and mapped to relevant AI modalities. Task-level estimates are then aggregated to the occupational level using task frequency weights.

To estimate potential economic value by industry, we connect task-level AI performance to financial outcomes across two channels: time savings and quality improvements. These are translated into industry-level estimates using occupational composition, wage benchmarks and revenue data. We model adoption trajectories using diffusion-style curves calibrated to observable readiness signals, reflecting that value accrues over time rather than immediately. Results are presented across two lenses: AI-driven labour impact and AI-driven revenue uplift. All outputs are scenario-based and should be interpreted as potential value under stated assumptions; actual outcomes depend on execution, governance and sustained adoption.

Expert insight

Interviews with practitioners were used to ground the narrative in real organisational experience and informed the development of the report's recommendations.

References

1. Accenture, Generating growth: How generative AI can power the UK's reinvention, October 2024

2. METR, "How does time horizon vary across domains?," July 2025

3. Banking Exchange, "JP Morgan Chase Reclassifies AI Spending as 'Core Infrastructure' " January 2026

4. UK Office for Budget Responsibility, Briefing paper No.9: Forecasting productivity, 2025

About Accenture

Accenture is a leading solutions and global professional services company that helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed across the enterprise, bringing together the talent of our approximately 786,000 people, our proprietary assets and platforms and deep ecosystem relationships. Our strategy is to be the reinvention partner of choice for our clients and to be the most AI-enabled, client-focused, great place to work in the world. Through our Reinvention Services we bring together our capabilities across strategy, consulting, technology, operations, Song and Industry X with our deep industry expertise to create and deliver solutions and services for our clients. Our purpose is to deliver on the promise of technology and human ingenuity and we measure our success by the 360° value we create for all our stakeholders.

Visit us at accenture.com

About Accenture Research

Accenture Research creates thought leadership about the most pressing business issues organizations face. Combining innovative research techniques, such as data-science-led analysis, with a deep understanding of industry and technology, our team of 300 researchers in 20 countries publish hundreds of reports, articles and points of view every year. Our thought-provoking research developed with world leading organizations helps our clients embrace change, create value and deliver on the power of technology and human ingenuity.

Visit us at accenture.com/research