



# Generating Impact

Turning frontier AI capabilities into  
frontline productivity and growth in Ireland

# Contents

|          |                                               |            |
|----------|-----------------------------------------------|------------|
| <b>1</b> | <b>Foreword</b>                               | Page 3-4   |
| <b>2</b> | <b>Authors</b>                                | Page 5     |
| <b>3</b> | <b>Introduction</b>                           | Page 6     |
| <b>4</b> | <b>The state of AI</b>                        | Page 7-19  |
| <b>5</b> | <b>The challenges still to overcome</b>       | Page 20-22 |
| <b>6</b> | <b>Art of reinvention</b>                     | Page 23-39 |
| <b>7</b> | <b>Taking action: From adoption to impact</b> | Page 40-45 |
| <b>8</b> | <b>About the research</b>                     | Page 46-47 |



# Foreword



**Hilary O'Meara**

Country Managing  
Director, Accenture  
in Ireland



Ireland has earned its place at the top table of global technology—built on visionary State policy, world-class investment and the talent to make it work. But the AI era demands something more from all of us. The question that occupies me is whether Irish business is ready to own that responsibility.

Over the past two years, AI adoption has accelerated rapidly. AI tools are now part of daily work for many people, confidence is rising and the technology itself has advanced at extraordinary speed. Ireland has responded with ambition—appointing a dedicated Minister for AI, investing in skills and research, establishing the infrastructure for responsible AI governance, and launching a National Digital and AI Strategy, sending a clear signal that AI will be central to our future competitiveness.

But intent is not impact. What will determine Ireland's future success is not solely adoption but the pace and scale at which we apply AI to change how work gets done, how organisations perform and how value is created across the economy. The decisions made in the next two years will shape Ireland's economic trajectory for the decade that follows. Opportunities of this scale do not wait.

AI is no longer just a tool that makes individuals more productive. The rise of agentic AI—systems that can coordinate, execute and optimise work end-to-end—represents a fundamental change in how work itself gets done. This is not a technology upgrade; it is a fundamental redesign of work. It challenges every senior leader to reimagine not just their technology stack but their entire operating model. The organisations that will thrive are those that embed AI into how they work, not just the tools they use.

Across Ireland, we see a widening gap between individual adoption and organisational transformation. Many people are using AI to work faster or better. But organisations have yet to put the corresponding effort into redesigning processes, rethinking roles or rewiring decision-making to capture the value. The economic impact remains fragmented, visible in pockets but not yet felt at scale.

The gap between those moving now and those still deliberating is already widening. Those who move first will set the standard. Everyone else will spend years catching up. Closing that gap is not easy. It cuts across teams, budgets and ways of working that have been in place for years. Delivering value requires senior leaders to engage directly, make clear choices on where they will invest, align AI strategy to business strategy and personally drive the redesign of work, workforce and technology systems.

And the stakes go beyond business performance. AI will change roles, skills and career paths across every sector. Leaders must invest in their people as much as they invest in technology, building the confidence and capability that turns AI from a tool into a way of working. If they do not, adoption will stall. And we risk leaving people behind in a transition they had no hand in shaping. That is not the Ireland we want to build.

Ireland has all the ingredients to lead in the age of AI: a skilled workforce, a public and private sector that is proven to deliver, deep connections with the global technology industry and a genuine national ambition. Now the question is whether Irish business will play its part.

That is what this research is about. It shows where Ireland stands today, where the barriers are, and what separates organisations already using AI to drive productivity and growth from those still waiting for the right moment.

That moment is now. The question this research puts to me and every leader is a direct one: will we act boldly enough, and quickly enough, to make the moment count?

# Authors

This report was a collaborative effort between our team based in Ireland, supported by our research team and assisted with AI:



**Denis Hannigan**

AI & Data Lead



**Liam Connolly**

Managing Director,  
AI & Data



**Audrey O'Mahony**

Managing Director,  
Talent & Organisation



**Austin Boyle**

Managing Director,  
Technology



**Donal Óg McCarthy**

Managing Director,  
Cybersecurity



**Mike Moore**

Research Director,  
UK and Ireland



**Babak Moussavi**

Research Senior Principal,  
UK and Ireland



# Introduction

AI presents a step-change opportunity for the Irish economy, for organisations and for individuals. But capturing it requires more than deploying the technology. Building on our 2024 report *Generating Growth*,<sup>1</sup> we repeated our survey of executives and employees in Ireland. 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 Ireland 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 Ireland 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.<sup>2</sup>

This expands the economic scope of AI significantly. The share of working hours in Ireland that AI could enable has risen from 42% to 82% since we last reported. We already see evidence that AI is being used for tasks that account for 20% of working hours in Ireland (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."




**The share of working hours in Ireland that AI could enable has risen from 42% to 82% since we last reported.**

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

**Share of work time in scope of reinvention with AI**

Ireland labour force. Accenture industries based on ISIC sectors



**Key:**  Existing evidence of AI use (digital and physical)  Theoretical potential of AI (digital and physical)  Not in scope

**Notes:** Accenture Research based on ONET task data and Ireland Central Statistics Office employment data. ISCO-1 occupations ranked by share of work time with strong AI evidence (highest to lowest). Bars may appear slightly over or under 100% due to rounding.

Introduction

The state of AI

The challenges still to overcome

Art of reinvention

Taking action: From adoption to impact

About the research



## Employees: Both keen and cautious

Across the Irish workforce, AI adoption has accelerated rapidly. More than one in five (22%) employees now use generative AI tools in their work daily—slightly higher than in the UK (18%) and almost triple the rate in 2024 (8%).

Based on analysis of open-text responses from more than 2,000 individuals in both Ireland and the UK, we find that the perspectives of the workers we surveyed have matured, too. 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, as people have increased their familiarity with the tools (Figure 2).

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




**Only 34% of Irish employees say they use AI tools provided by their employer; a further 30% have sourced tools independently, signalling an ongoing “shadow AI” phenomenon.**

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

**Feelings about AI**


Sentiment of free text responses, %




**Themes that AI evokes**

AI clustering of responses, %


**Positive**



**Neutral**



**Negative**



**Key:** 2026 2024 Increase since 2024 Decrease since 2024

**Note:** We asked 2,085 individuals in the UK and Ireland 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.

Introduction

The state of AI

The challenges still to overcome

Art of reinvention

Taking action: From adoption to impact

About the research






AI is already reshaping demand for skills. Demand for routine, structured cognitive skills is declining. AI-related skills are the 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 toward redesigning and governing how those workflows operate. As AI takes on execution, the human role expands—toward 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 Irish occupations with job postings mentioning each skill subcategory 2023-2025<sup>1</sup>



**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 adoption increases are communication and collaboration, critical thinking, and leadership and change management. Almost two in five people (39%) expect their job either to be unrecognisable or to disappear completely by the end of the decade, more than in the UK (31%) and a rise from 31% who held that view just 18 months ago. Of those, 87% say they are likely to reskill and almost two thirds (63%) say they are likely to change occupation.

System-level initiatives are emerging in response. In Ireland, the government's National Digital and AI Strategy, published in February 2026, commits to a new national AI Skilling Platform for employers and individuals and a nationwide digital and AI skilling campaign, alongside an ambition for 100% of key public services to be "digitalised" by 2030.<sup>3</sup> But these do not substitute for organisational action. And the risk is that organisations are not yet providing enough clarity or support to facilitate the transition at scale. Two fifths (40%) of Irish executives indicate that change management and workforce transition is one of their most significant AI-related skills gaps. That said, there are encouraging signs of action: more than a third (35%) have conducted a skills audit to assess AI's impact on roles, 44% are investing in reskilling and redeployment pathways for at-risk roles, and only 11% say they do not provide AI-related training at scale.



**Almost two in five people (39%) expect their job either to be unrecognisable or to disappear completely by the end of the decade, more than in the UK (31%) and a rise from 31% who held that view just 18 months ago.**

Still more needs to be done, as many Irish executives continue to assume that AI's impact on jobs will not be wholly benign. In 2024, 39% of Irish executives believed AI would shrink the national workforce; in 2026, that share remains stubbornly unchanged, although it is lower than the 49% of UK executives that hold this view. Seven in 10, meanwhile, expect AI to drive higher productivity. If leaders assume displacement is inevitable, the incentive to invest in 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 sharper at the bottom of the career ladder. In 2024, almost half (49%) of executives expected AI to increase demand for entry-level roles. That has dropped to a third. By contrast, the share expecting reduced demand for such roles has risen from 21% to 37%. These trends echo movements in the UK, too. 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.

Where the disruption will be greatest remains uncertain. Recent analysis of US data by Anthropic finds that AI exposure is currently highest in more educated, higher-paid and white-collar roles, rather than being concentrated in entry-level positions.<sup>4</sup> Yet they also find that the earliest labour market effects appear to be showing up in hiring patterns rather than unemployment; there are signs of a slowdown in hiring for younger workers entering these more exposed occupations. This points to a potential divergence between where AI is most capable today and where its impact is first being felt: not through immediate displacement, but through changes to entry routes, job transitions and career progression. Even without widespread job loss, this dynamic could reshape the pipeline into high-skill work.

## Organisations: Changing more slowly than their people

Organisational change is lagging behind. AI is being adopted at the level of individuals and tasks, but not at the level of systems and workflows. Only 29% 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.

# 26%

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. More than a quarter (26%) of executives say that if their organisation stopped using AI, some critical processes would slow down significantly; a further 4% say some core operations would fail. Some banks are beginning to reclassify AI as “core infrastructure.”<sup>5</sup> 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. Close to half (41%) of Irish consumers have already used AI tools to search, compare or buy on their behalf, and 37% say they would delegate the purchasing process to an AI agent if the tools were reliable. Yet half (49%) of executives don't 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 operationally ready for what that shift will mean.

The economic returns reflect this unevenness. Over one third of executives (35%) say AI has had little to no positive impact on profit and loss. While this is lower than the share of UK executives that believe this (46%), Irish executives are simultaneously more likely to believe that at least some of their AI budgets are wasted—with 72% saying this, compared with 44% in the UK.

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 many more executives see cost reduction as the greatest future benefit of AI rather than revenue growth (37% vs 26%). This stands in contrast to workers, more than half of whom 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.




**Half (49%) of executives don't expect shopping agents to have significant influence on consumer decisions over the next three to five years.**

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

The stakes extend beyond individual firms. Ireland's economic model is dependent on a combination of globally competitive multinationals, a large base of indigenous SMEs and the effectiveness of its public services. AI creates an opportunity to strengthen each of these pillars—and to position Ireland as a leading hub in the next wave of AI-driven growth, building on its success in attracting digital investment and its strong base of skills. But capturing that opportunity depends on whether productivity gains are realised across the economy as a whole.

Ireland is not starting from scratch. The National Digital and AI Strategy sets out a comprehensive programme across enterprise adoption, public services, infrastructure, skills, regulation and security. It is not a vision statement but an execution framework, with clear ownership and timelines. The foundations are in place: investment in research and compute infrastructure, new institutional capacity through the AI Office and regulatory sandbox, and a commitment to fully digital public services.



**AI creates an opportunity to position Ireland as a leading hub in the next wave of AI-driven growth, building on its success in attracting digital investment and its strong base of skills.**

But strategy does not guarantee outcomes. The challenge for the country is execution at scale. For multinationals, this means sustaining Ireland's role as a location for high-value, AI-enabled activity. For local firms and start-ups, it means translating capability into scalable businesses. For SMEs, it means closing a persistent adoption gap that risks becoming a competitiveness gap. And for the public sector, it means delivering measurable improvements in productivity and service quality through genuine redesign, not incremental digitisation. Investments across the economy—from skills availability to infrastructure and energy capacity—will shape how far and how fast AI-driven growth can materialise.

If productivity gains are realised at scale, they reinforce Ireland's competitive position and fiscal base. If they are delayed, fragmented or captured elsewhere, the consequences for a small, open economy are amplified. This matters not just because of growth, but because AI may not generate employment at the same scale as previous technology waves. Without broad-based productivity gains, the model becomes harder to sustain.

The need for wider investment and change is not unprecedented. General-purpose technologies, from electricity to the internet, have historically required sustained complementary investment before delivering durable gains. But it underscores how much needs to be done at the level of systems, not just firms. A national strategy creates the conditions but does not buy time.

AI's productivity impact is no longer just a question of corporate performance. It is an economic imperative for the country as a whole—one that will shape Ireland's growth, competitiveness and resilience over the next decade.

## 2. The challenges still to overcome

Three structural tensions persist



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. 54% 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: 37% are still piloting predictive AI and machine learning, and a further 18% are not using it at all.

Workforce readiness remains uneven. Only 5% of executives say their workforce is fully prepared for agentic AI. Yet more than half of employees (64%) say they are likely to reskill in response to AI, signalling that many already see significant changes coming to their roles.

The trust gap endures. More employees (24%) than executives (14%) expect a “significant decrease” in jobs as a result of AI. While explicit worries about trust and user acceptance are less of a concern now, data security (37%) is the obstacle to AI adoption that executives cite most often—and this has risen in salience since 2024 (Figure 4). Regulatory concerns have risen as well, with almost a third (32%) pointing to this as one of the biggest barriers. Meanwhile, more than half (52%) of employees say they are uncomfortable with AI agents making higher-impact decisions. This suggests that trust issues are limiting how far workers are willing to adopt and rely on AI.

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.




**Only 5% of executives say their workforce is fully prepared for agentic AI.**

**Figure 4:** Data security, privacy and regulatory concerns are now the primary barriers to scaling AI

### Barriers to AI adoption

Proportion selecting this as a barrier, Ireland **executive sample** (2026 n=57, 2024 n=128)



**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.

Introduction

The state of AI

The challenges still to overcome

Art of reinvention

Taking action: From adoption to impact

About the research

# 3. The art of reinvention



Introduction

The state of AI

The challenges  
still to overcome

Art of reinvention

Taking action:  
From adoption  
to impact

About the  
research



# 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 enable growth.

04



## Digital Core

Connecting, securing and scaling the systems AI needs to 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 far deeper than most organisations have acknowledged. Two fifths 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.



## Strategy: Start with the outcome

### Where organisations are today.

AI strategies are often developed separately from business strategy, driven by technological momentum rather than clearly defined business outcomes, producing a proliferation of 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, Scalars 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 of 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 (e.g. cost or speed per decision or revenue impact) alongside a mechanism to track how capacity is redeployed. 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 and 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.



## 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 (52%) 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:**

### Redesign workflows from first principles.

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. In practice, organisations making the fastest progress tend to concentrate this capability in a dedicated function, whether a reinvention hub, centre of excellence or cross-functional design authority, that owns the discipline of redesigning work around AI rather than leaving it to individual teams.



## Simulate, monitor and continuously adapt.

Once a workflow has been zero-based (designed from first principles), 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.



## Workforce: Build human–agent teams

### Where organisations are today.

Workforce readiness is low. Only 5% 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 (those at the highest stage of maturity), 76% of executives say they are very confident in the quality of their AI training, compared with just 6% at Experimenter (lowest stage) 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 34% of employees in Ireland that possess the skill of creativity but do not use it at work—this was the highest rate of underutilisation of any skill we evaluated.

# 63%

of executives say employees are expected to be transparent about AI use, only 17% of employees say they always disclose it.

## Measure and manage for evolving human-agent performance.

Existing performance frameworks were built for a purely human workforce. Ireland appears to be making progress here: 56% of Irish executives say their organisation has updated performance measures for AI-enabled ways of working, compared with 43% in the UK. Meanwhile 68% say managers can assess performance fairly when outputs are partly AI-assisted—against just 40% in the UK. But employees are slightly more likely to say AI-assisted work is valued less rather than more (28% vs 25%), creating conditions for “quiet AI”: while 63% of executives say employees are expected to be transparent about AI use, only 17% say they always disclose it. Across the full sample, 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.






## Digital Core: Connect systems so AI can act

### Where organisations are today.

Across almost every aspect of the digital core we assessed, from cloud to data and security, executives tend to say they are not ready for agentic AI (Figure 5). Only 3% of executives in Ireland say all aspects of their technology and data stack are fully ready. The foundations require particular attention: fewer than half of Irish executives say their cloud infrastructure is ready to support scalable agentic workloads, and readiness is lower still on core system integration and multi-agent workflow orchestration, where around six in 10 organisations remain at partial readiness or below. Significant investment in the digital core will be needed if Irish organisations are to move from experimenting with AI to embedding it in operations.




**The digital core is the technology foundation that makes AI actionable. It brings together the data, systems, integrations, APIs, security controls and governance infrastructure that agents need to access information, execute tasks and coordinate across the organisation. Without it, AI operates in isolation; with it, AI can be embedded into the workflows and decisions that drive performance.**



Data remains a key 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. Rather than attempting to fix the entire data estate, organisations should prioritise the data domains linked to high-value use cases, avoiding “boiling the ocean.” 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. As AI systems combine data and instructions, data governance becomes not just a quality issue but a control mechanism for how systems behave.

**Figure 5:** The technology and data foundations for agentic AI are not yet in place across most organisations

**Digital core readiness for agentic AI**  
%, Ireland executive sample (n=57)



**Key:**  Net ready  Net not ready  Net positive  Net negative

**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".

Introduction

The state of AI

The challenges still to overcome

Art of reinvention

Taking action: From adoption to impact

About the research

**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. This also requires making relationships between data entities explicit, so agents can reason across information rather than simply retrieve it. 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.



## 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. As systems operate across changing contexts and evolve through interaction with data, people and other systems, their behaviour becomes less predictable and harder to reproduce, limiting the effectiveness of point-in-time governance. As agents act across multi-step workflows involving multiple systems, accountability becomes harder to trace, risking fragmentation across functions and weakening governance in practice.

# 40%

of executives report having the tools to monitor and audit agent behaviour at scale (versus 28% in the UK), more than half (56%) say they are not ready to supervise agents in real time.

Roughly half of executives cite loss of human oversight or excessive autonomy (47%) as the biggest risk of agentic AI, with a further 37% pointing to incorrect or harmful decisions. While 40% report having the tools to monitor and audit agent behaviour at scale (versus 28% in the UK), more than half (56%) 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 across the combined UK and Ireland samples have any of them fully operational. At the same time, a gap is emerging between governance intent and technical implementation, with controls often defined at a policy level but not consistently embedded into system architecture.

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.

**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—controls designed for human-in-the-loop systems are no longer sufficient. When executives select agentic AI vendors, security, privacy and compliance guarantees are the factors cited most often (40%).

## **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 and the ability to reconstruct any decision to audit standards and lifecycle. Governance must be maintained across the full operational life of a system rather than discontinued after deployment. At the core of this control plane is digital identity. Every agent must have a verifiable identity, with clearly defined permissions, accountability and auditability. Without strong identity, organisations cannot reliably attribute actions, enforce least-privilege access or meet regulatory requirements. With it, autonomy becomes governable: every decision, action and escalation can be traced, constrained and—when necessary—stopped. Identity is what turns agentic AI from an unmanaged risk into a trusted part of the enterprise.

# 4. Taking action: From adoption to impact



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 and 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?
- Can you clearly identify, constrain and audit every autonomous action your agents take in the same way we would a human employee?
- 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 Ireland. It is broadly balanced across genders and spans a wide age range, with around 30% under 35. We compared results to a survey of 3,752 the UK workers and 409 Irish 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 453 respondents from the UK, and 57 from Ireland, with a strong spread across the mid-market and large enterprise. We compared results to a survey of 1,085 UK executives and 128 Irish executives fielded in July–August 2024.

## 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 Scalars. 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 Ireland's reinvention", December 2024.
2. METR, "How does time horizon vary across domains?", July 2025.
3. Government of Ireland, "Digital Ireland – Connecting our People, Securing our Future", February 2026.
4. Anthropic, "Labor market impacts of AI: A new measure and early evidence", March 2026.
5. Banking Exchange, "JP Morgan Chase Reclassifies AI Spending as 'Core Infrastructure'", January 2026.

## 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](https://www.accenture.com)

## About Accenture Research

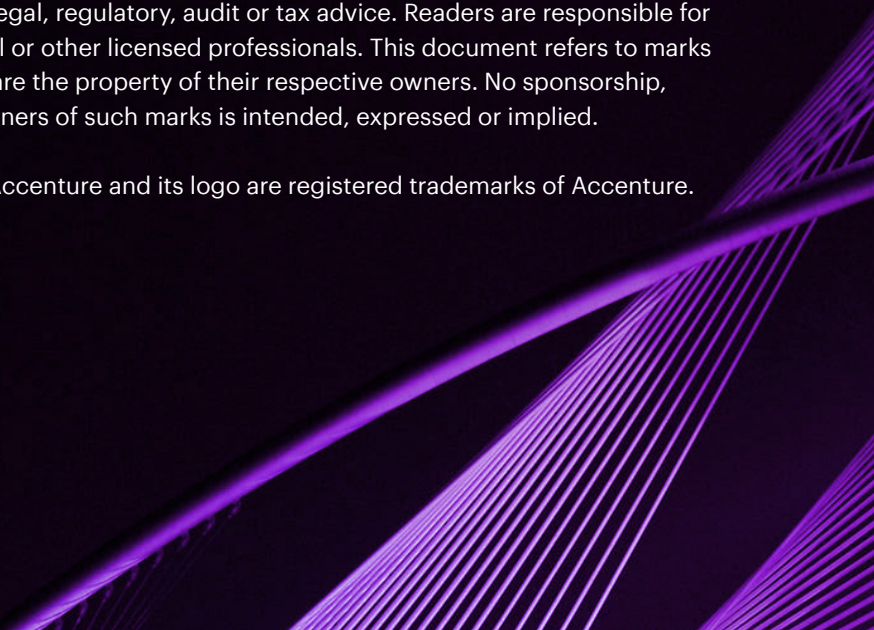
Accenture Research creates thought leadership about the most pressing business issues organizations face. Combining innovative research techniques, such as data-science-led analysis, with a deep understanding of industry and technology, our team of 300 researchers in 20 countries publish hundreds of 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](https://www.accenture.com/research)

---

Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date provided on the front page, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information. Accenture does not provide legal, regulatory, audit or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals. This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.

Copyright © 2026 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.

A decorative graphic in the bottom right corner consisting of numerous thin, parallel purple lines that curve and fan out, creating a sense of motion and depth.