



The *adoption* decade

Closing the execution divide
and making AI work for Britain

OLIVER WYMAN
A MARSH BUSINESS

| **C | B | I** |

Co-chair foreword

The United Kingdom starts from a position of real strength. It has world-class universities, one of the deepest pools of AI research and talent anywhere, and an enviable technology sector. That strength is not in doubt, and it is a genuine reason for confidence.

Every few generations, a technology changes not only what an economy produces, but how it competes. Steam power, electrification, and the internet were each such a moment. Artificial intelligence is creating the next. History shows that invention alone is never enough. The largest gains come when a technology spreads through firms, workflows, infrastructure, and skills, reaching businesses of every size, in every sector. Adoption on that scale also opens a route to reindustrialise Britain, combining AI with advanced manufacturing to strengthen domestic supply chains, raise productivity, and create high-quality jobs in the places that need them most. Both prizes are well within reach. Reaching them is a matter of execution.

The question this report sets out to answer is a practical one: how does Britain turn its advantage in AI into prosperity that is felt in every place and by every person?

Today, an execution divide is opening between firms that are redesigning workflows, establishing governance, and measuring returns, and those still experimenting and piloting. The barriers are practical, including an unclear business case, fragmented guidance, shortages of skills and confidence, uncertain regulatory routes, and concerns about data, suppliers, and service continuity. They are also solvable.

Throughout this work, we heard from businesses already turning ambition into results: a medical technology company whose AI now acts as a GPS for surgeons, guiding them through minimally invasive repairs; a construction group teaching AI to identify the early warning signs of a safety incident before anyone is hurt; and a telecoms company that has taken a thousand of its people through AI apprenticeships in six months and opened its levy funding so smaller firms can do the same. Even in fire and rescue services, frontline employees far from any traditional technology function are retraining and building AI tools for their own operations. These examples show both the breadth of AI's potential and how its benefits are already beginning to be realised in Britain today.

That is why this report calls for AI adoption to become a national productivity mission, judged not by the number of pilots launched but by the difference it makes to people's lives. Wider adoption can help firms reduce the cost of doing business and, through stronger productivity, investment, and wage growth, contribute to easing pressures on the cost of living. The report highlights the economic value at stake, takes stock of where businesses and government stand today, and sets out how the UK can lead in the years ahead. This is the work of the taskforce that brought together British business leaders, government representatives, and AI experts from civil society to agree a set of practical, achievable steps that build on the government's ambition. At their heart is a simple proposition: both the supply and demand sides of the market need to work, and to work together, so that firms of every size have the infrastructure, skills, clear rules, and confidence to put AI to work at scale.

Timing matters. The government has made growth and rising living standards central to its mission, at a moment when the country is looking for evidence that its economy can move forward and that the gains can be felt everywhere. That evidence will not come from strategy alone. Nothing is as ready to deliver it as the widespread, confident, and responsible adoption of AI, showing up in everyday life: in better public services, in secure and well-paid work, and in opportunities for young people. The country has both the capability and the opportunity; what matters now is turning that shared ambition into delivery. That means joined-up action, with government, business, working people, and civil society moving together and with purpose. It will require a clear, time-limited mechanism to prioritise action, aligning existing initiatives, assigning accountable owners, and tracking delivery; immediate no-regret action by businesses to move from pilots to scale; and decisive government action to clear the practical barriers holding adoption back. This report sets out recommendations for how each of these priorities can be delivered.

We are grateful to the members of the taskforce for their time, expertise, and candour, and to the many businesses whose experience shaped this work. The prize is AI that works for people: an economy that is more productive, more resilient, and more widely prosperous, and that everyone can participate in, benefit from, and trust. We look forward to helping build it.



Lisa Quest

CEO, Marsh UK

Head of UK and Ireland, Oliver Wyman



Rain Newton-Smith

CEO of Confederation of British Industry (CBI)

Acknowledgements

This report has been shaped by the expertise and experience of a wide range of contributors to whom we are grateful for generously giving their time to participate in discussions and share their perspectives throughout the development of this report.

In particular, we would like to thank the members of the Accelerating AI Adoption Taskforce, whose leadership and expertise have been instrumental in shaping the report's findings and recommendations. Beyond individuals and organisations named below, perspectives were also received from representatives of the legal services and technology sectors.

Accelerating AI Adoption Taskforce members

Asha Kumar	Zurich Insurance Group
Ron Chambers and Hannah McIlroy	Kier Group
Tom Meakin and Daniel Wilson	BT
Leigh Chantal Pharand	AstraZeneca
Deepak Nath	Smith+Nephew
Adolfo Hernandez	Capita Group
Marcus Hunter	Evri
Chantelle de Villiers and Manav Thiara	Sage
Clare Schramm Fergus	Lloyds Banking Group
Matt Prebble	Accenture
Niels van den Brink	Asda
Sharon Hague	Pearson
Frank Millar and Arun Harish	CPI

We also acknowledge the valuable contributions of members of the CBI's Tech and Innovation Committee, AI and Data Protection working group, our research panel of AI experts from civil society, and government representatives engaged throughout the process. Their perspectives and input have helped ensure the report reflects a broad range of views from across business, civil society, and government.

Note: Quotes from taskforce members refer to representatives of Taskforce member organisations engaged in this work, and are not necessarily limited to the named individuals above.

Contents

Executive summary	6
-------------------	---

Recommendations	11
-----------------	----

The economic imperative	13
-------------------------	----

Enterprise demand and AI infrastructure	24
---	----

Governance and trusted adoption	48
---------------------------------	----

Skills and workforce readiness	69
--------------------------------	----

From recommendations to delivery	90
----------------------------------	----

Executive summary

The UK's greatest AI opportunity is to become the world's leading human-centric AI-enabled economy: adopting AI fastest, shaping its use responsibly, and building durable, resilient strength in the parts of the AI stack that matter most for competitiveness and service continuity.

01 | THE PRIZE

A more prosperous, resilient, and inclusive UK

AI adoption should ultimately be judged by the value it creates for people, businesses, and communities across the UK. With its services-led economy and strong digital base, the country is well placed to benefit from wider AI adoption. The report identifies four outcomes to guide the UK's approach, capturing both the benefits it should pursue and defining what successful AI adoption should deliver.

Higher living standards through broad-based productivity gains

Widespread AI adoption can raise productivity, improve services, and strengthen business performance across sectors, regions, and firms of all sizes, so that the gains are shared in every part of the UK.

Greater resilience through scaled adoption and a competitive UK AI stack

Wider adoption, supported by reliable infrastructure, data, compute, skills, and implementation capability, can strengthen business resilience while enabling the UK to capture more value from the parts of the AI stack where it has genuine strengths.

Trusted and responsible human-centric adoption that gives businesses, workers, customers, and the public confidence to engage

Clear governance, proportionate regulation, and credible assurance can help manage risk, reduce uncertainty, and support AI adoption at scale.

Broad participation in an AI-enabled economy

Stronger AI literacy, practical skills, and transition support can help people benefit from and adapt to technological change, as employees, as consumers, and as users of public services that are faster, more accessible, and more responsive.

02 | THE ADOPTION LANDSCAPE

Execution is becoming the differentiator

The evidence presents a mixed picture with strong momentum and clear areas of UK advantage, alongside uneven progress and persistent constraints.

Adoption is broad but selectively deep

Usage is widespread and often experimental, clustered in employee productivity, customer service, operational efficiency, and analytics, with a smaller group beginning to embed AI into core processes and move towards scaled or agentic use.

A clear return-on-investment gap is opening between leaders and laggards within industries

Firms securing the strongest returns tend to combine strategic leadership, organisational readiness, and disciplined implementation. The AI divide is no longer a technology divide — it is an execution divide.

Most AI value still accrues outside the UK, but the UK can capture more

Revenues tied to frontier models and hyperscale infrastructure largely accrue elsewhere, yet much of the value businesses seek is created closer to deployment, in workflow integration, applied models, trusted data, secure hosting, assurance, cyber resilience, and specialist implementation. These layers align with UK strengths in forward deployment in regulated sectors, research, services, assurance, cybersecurity, and finance, despite challenges around power, planning and physical infrastructure, alongside uneven real-world 5G coverage and capacity across locations.

Businesses approach sovereignty pragmatically, with resilience the priority

Most do not seek a fully domestic stack for every workload, favouring hybrid models that combine global infrastructure with domestic storage and tighter controls over data and access where sensitivity, security, or resilience require them. The priority is targeted domestic capability that reduces critical dependencies on a small number of global suppliers, rather than blanket localisation that raises costs, fragments services, and constrains adoption.

03 | FROM ADOPTION TO VALUE

What success requires

Whether the UK captures the opportunity will depend on whether firms can move from experimentation to scaled, trusted adoption, supported by effective implementation, practical guidance, navigable regulation, credible assurance, and an AI-ready workforce.

Value increasingly depends on implementation, not access to technology alone

AI delivers the greatest benefits when treated as business transformation rather than a technology roll-out. Successful adopters align AI with strategic priorities, redesign workflows, establish appropriate governance, design in security and resilience as deployments scale, invest in data and technology foundations, build the skills and capability their workforce needs to use AI well, and measure whether implementation is delivering value.

Many firms would value more practical, trusted guidance, with smaller businesses facing the greatest constraints

Some struggle to identify valuable use cases or find credible support relevant to their sector and circumstances within a fragmented landscape. Larger firms can absorb the costs of experimentation more readily, while trusted intermediaries such as universities, education and training providers, trade and professional bodies, catapults, banks, insurers, and government-supported programmes can translate national ambition into practical regional and sector-specific support.

Among more advanced adopters, governance is becoming an enabler of scale.

Boards set ambition and risk appetite, business units own the use case, controls and compliance follow the level of risk, and humans remain accountable. Responsibility is clearest when it follows control, with developers generally best placed to manage model design, safety, documentation, and upstream assurance, and deployers managing the data, integration, oversight, and operating environment they control.

Navigating regulators is becoming a barrier to scale

The challenge is not the absence of another approval mechanism, but the lack of a clear, informed point of entry into the existing system. A “no wrong door” model, supported by

a practical regulatory map, could help firms identify the relevant requirements and regulators without creating a new approval authority.

Shared norms can reduce duplicated governance, while assurance is an area of UK strength

Firms are often addressing similar governance challenges in isolation. Shared principles for responsible AI adoption, potentially supported by a voluntary Responsible AI Adoption Code and public pledge, could provide firms with a common reference point and build confidence among customers, employees, and supply chains. The UK's capabilities in standards, audit, conformity assessment, risk management, and professional services provide a strong basis for a globally recognised AI assurance market.

Skills are becoming a source of competitive advantage, while workforce confidence remains a limiting factor

Needs span baseline AI literacy, practical usage, judgement and oversight, leadership and transformation, and specialist capability. Confidence is stronger where employees understand how tools will be introduced and how work will change. The largest gaps sit beyond the reach of individual employers, particularly for people entering work, moving between roles or sectors, and later-career workers.

04 | FROM AMBITION TO DELIVERY

Priorities for action

Taken together, these findings point to one conclusion — the UK's challenge is no longer forming a strategy, but delivery at scale. Many of the building blocks are already in place, but they need to be better aligned and translated into coordinated action. The redistribution of AI, digital, and innovation responsibilities across the Department for Business, Innovation, Science and Trade (BIST), the Cabinet Office, and the Department for Digital, Culture, Media and Sport (DCMS) creates an opportunity to embed delivery more closely across business growth, public-service transformation, and digital policy, but also makes clear coordination and accountability essential.

The task now is to translate this opportunity for greater alignment into a prioritised programme of action, with clear ownership, measurable outcomes, and rapid learning.

Three immediate actions follow:

- **Establish a time-limited Business AI Adoption Delivery Group, a government and industry partnership feeding into the Prime Minister's AI Taskforce.** With the Taskforce leading on public sector adoption, this group should give equivalent focus to the private sector. Convened by the CBI with other partners and sponsored by the Minister for AI, prioritising the report's recommendations, aligning existing initiatives, and overseeing delivery through a shared roadmap and scorecard.
- **Businesses should pursue immediate no-regret actions.** Firms should assess their approach against emerging best practice, prioritise a small number of the highest-value use cases, including those that drive growth, and move them from pilot to scale. They should embed governance and resilience into deployments, build role-specific skills through implementation, and measure adoption through business outcomes rather than the number of tools deployed. They should also share practical experience through trusted forums where this does not compromise commercial advantage.
- **Government should clear the path for adoption.** Immediate priorities include using public sector adoption and procurement in high-volume citizen-facing services to stimulate demand, piloting the "no wrong door" model, mobilising trusted regional and sectoral diffusion channels, tackling data, mobile connectivity, infrastructure, and energy bottlenecks, and establishing a consistent national framework for measuring the depth, value, and distribution of AI adoption.

The CBI will play an active role in delivery. It will provide the business voice within the emerging delivery arrangements, convene firms and delivery partners, support practical knowledge sharing across sectors and regions, and continue gathering and supplying evidence on adoption to help government and industry refine priorities as implementation progresses. It will also work with an independent and credible convener to bring together business, workers, civil society, and government to explore shared principles for responsible AI adoption.

Recommendations

TARGET OUTCOME 1

Achieve productivity benefits through AI adoption to raise living standards across the country

Foundational recommendation

Make AI adoption a national productivity mission, focused on diffusion across sectors and firm sizes

TARGET OUTCOME 2

Scale AI adoption and build a competitive UK AI stack to underpin resilience

For business

Build the organisational foundations for scaled AI adoption

Build capability by engaging trusted local intermediaries and shared learning

For government

Strengthen the research and development ecosystem, including public compute, AI-ready data, models, chips, memory, and energy efficiency

Accelerate investment and regional deployment of AI stack infrastructure, including AI Growth Zones, power, planning, connectivity, cyber resilience, financing, and insurance

Accelerate the government's own adoption, and use procurement, AI Champions, and adoption support to diffuse AI nationwide, with priority given to high-volume public services where AI can demonstrably improve accessibility and citizen experience

TARGET OUTCOME 3

Enable trusted, responsible, and human-centric AI adoption to manage AI-related risks and give end users the confidence to engage

For business	Implement robust governance and responsible adoption practices
---------------------	--

For the CBI	Convene stakeholders to develop shared principles and explore a voluntary Responsible AI Adoption Code and public pledge
--------------------	--

For government	Evolve the regulatory model through assurance, clearer routes through the regulatory system, the AI Growth Lab, and growth-enabling regulatory simplification for firms
-----------------------	---

TARGET OUTCOME 4

Build an AI-ready workforce and society to ensure broad participation in an AI-enabled economy as consumers, employees, and users of public services

For government	Systematically infuse AI knowledge into education, employment support, entry-level jobs, workforce transition, and later-life retraining
-----------------------	--

For business	Sustain investment in AI skills at every level, from workforce literacy to role-specific and specialist capability, and leadership at board and executive levels
---------------------	--

01

The economic
imperative

Introduction

The current commentary and analysis on AI's impact on society and the economy is deeply polarised. Productivity gains, infinite abundance, new economic models on the plus side; joblessness, civilisational erasure, and systemic risks from cyber attacks and fraud for the naysayers.

(Relative) UK economic growth has been uninspiring for many years. Between 2008 and 2025, the UK has averaged economic growth of ~1.2% per year and (labour) productivity growth of ~0.4%. This compares with ~2% and ~1.3% respectively for the US, and ~1.1% and ~0.7% respectively for the EU 27.^{1,2} Because of compounding effects over decades, a significant gap has emerged between EU and UK wealth and income purchasing power and that of the US.

We examine whether AI could materially alter the UK's long-run economic trajectory through:

- 1. Business AI adoption and market demand:** Economic growth through productivity improvements and product, process, and service innovation as firms and consumers adopt AI tools.
- 2. Market supply:** Economic growth attributable to the AI "stack" across applications including forward deployment, models, compute (on premises and in the cloud), hyperscaler platforms, data centres, physical infrastructure, and energy. In other words, the market supply of AI-related products and services contributes to economic growth through capital expenditure, employment, and AI-related products and services.

This chapter is a simple synthesis of expert opinion on the expected economic impact of these distinct contributors to GDP growth. The synthesis reflects the expected positive impact and *assumes no negative shocks*. That is not because we are blind to the downsides, but because it is our belief that risks can be mitigated by appropriate industry, policy, and regulatory action.

It is clearly too early to observe GDP benefits from AI adoption thus far. Market demand for AI-related services is in its infancy but is clearly materialising at pace. This is observable in the valuation and revenues of the leading players in the AI stack. Current company guidance and analyst forecasts imply more than \$700 billion for 2026 and more than \$850 billion in 2027³ clearly signals confidence in future demand. **Industry growth at this pace is truly exceptional relative to technological discontinuities of the past.**

1 The World Bank: GDP Growth (annual %): Country official statistics, National Statistical Organizations and/or Central Banks; National Accounts data files, Organisation for Economic Co-operation and Development (OECD); Staff estimates, World Bank (WB).

2 Refers to SDG indicator 8.2.1: Annual growth rate of output per worker, GDP constant 2021 international dollars at PPP (%). Source: International Labour Organization. (2026). ILO modelled estimates database, ILOSTAT [ILOEST]. Available from <https://ilostat.ilo.org/data/>.

3 Oliver Wyman analysis of public equity analyst reports on Alphabet, Amazon, Meta, Microsoft, and Oracle.

Market demand for AI could lift UK productivity, but estimates vary

We examine conclusions reached by academics, think tanks, and official bodies. We do not include bodies that may have a natural bias, as these tend to be more positive (or negative) depending on their business model. We have mainly examined research for the UK, EU 27, and US given the depth of research in these economic systems. We also examined research on China, but this tends to be less extensive compared to the other economic blocks.

Analysis and synthesis of authoritative UK⁴ sources conclude:

Directional consensus on the basic qualitative conclusion

Greater AI adoption will raise productivity and in plausible scenarios raise GDP over the medium to long run. The dominant economic effect is positive (higher output per worker, new products and services, efficiency gains).

Sectoral concentration

Consensus identifies information-intensive sectors as primary beneficiaries (finance, professional and business services, government services, parts of healthcare, advertising and marketing, research and development, and information and communications technology). Manufacturing gains are more mixed and depend on automation and integration of AI with advanced robotics that underpin production processes across design, procurement, assembly, and delivery.

Complementarities are important

Consensus indicates that realising gains depends on complementary investments in skills/reskilling, digital infrastructure, data governance and access, organisational change, and capital (compute and software) spending.

⁴ Oliver Wyman analysis of public reports on AI and productivity; Office for Budget Responsibility (2024-2025); Institute for Fiscal Studies (2025); The Alan Turing Institute (2024).

Predominance of augmentation in the near term

Consensus expects the earliest benefits to come mainly from augmenting workers in existing roles (increasing speed, accuracy, and quality) rather than from job cuts. Over a longer horizon, a significant share of growth is expected to come from new products, services, and markets rather than from doing existing work more efficiently.

Timing

Consensus expects benefits to materialise over a multi-year horizon (5–10 years), similar to timelines experienced in major technological advances of recent history (such as the internet and smartphones) that took time to diffuse through the full economic system.

Magnitude

Not surprisingly, there is less consensus on this dimension. Conclusions present scenarios that are modest to material, with common central scenarios pointing to single-digit percentage increases in GDP over a decade (for example, ~2–5% level uplift over 5–10 years). In the research, GDP uplift is driven primarily by annual (labour) productivity growth in the order of 0.1–0.5% per year in plausible scenarios. Applying a simple short-term heuristic that translates (labour) productivity uplift into GDP growth, there is a one-to-one correspondence between productivity lift and GDP growth. **What is clear, however, is that given the UK's recent economic record, even the more pessimistic scenarios will have a measurable impact on economic growth.**

Two growth channels

The estimates above focus mainly on efficiency gains within existing firms and tasks. A second channel is growth from new products, services, business models, and markets that AI enables. This channel is harder to quantify and depends less on frontier-model access than on wide diffusion of AI into everyday business activity.

It is important to note that all these sources emphasise that precise economic predictions are impossible at this stage given major uncertainties in the analysis. Key uncertainties are a) assumptions on how saved labour time converts into higher measured output (more production, higher quality products, lower prices, or leisure) and therefore to GDP uplift; b) the unclear pace of adoption especially among small and medium-sized enterprises; and c) general equilibrium effects are not taken into consideration (wage responses, capital deepening, firm entry and exit) and these may materially affect outcomes.

Comparison to conclusions reached for China, the EU 27, and the US by similar authoritative sources⁵ yields some interesting insights:

There is broad consensus on the expectation that AI adoption will raise productivity and GDP potential. The biggest near-term winners are information-intensive sectors, and therefore services-heavy economies such as the UK should benefit more than others, all else being equal.

The UK has slightly higher potential relative to the EU 27, comparable to digitally advanced economies such as those in the Benelux region, France, Germany, and the Nordic countries, reflecting the UK's comparative strengths. A strong services sector, fintech and professional services, high digital adoption in certain industries, and a comparatively smaller SME ecosystem support significant potential benefits. However, the UK is facing some challenges that may limit immediate growth prospects, for example, regulatory divergence with the EU would increase barriers to services trade and stifle growth.

US analyses indicate roughly double the economic upside and faster diffusion. This reflects larger domestic technology champions, higher research and development intensity, a larger venture capital ecosystem, faster diffusion among large firms, and greater data availability in some sectors. The US also has better bottom-up evidence given its relatively advanced market. Projections are often supported by empirical analysis, randomised controlled trials, firm pilots, and measurable productivity experiments conducted within US firms or by US-based AI providers. This produces relatively robust estimates for specific tasks and firm contexts, including measured productivity gains of 10%–40% for certain cognitive tasks.

China has a relatively larger share of manual and physical tasks and manufacturing. Therefore, some AI applications, especially current generative models focused on cognitive tasks, may initially yield smaller measured labour productivity gains than in services-heavy economies where cognitive augmentation yields immediate productivity gains. AI-driven automation, process optimisation, and digital adoption can scale, but the pace of diffusion could be a challenge given the enormous diversity of ecosystems, regions, and firms. China's greatest potential gains may require more capital-intensive automation and take longer to appear in headline productivity statistics unless capital investment and robotics deployment accelerate. However, China's centralised industrial policy and large domestic AI champions can accelerate deployment and diffusion relative to many European countries. That supports faster scaling in targeted sectors, especially where government procurement can have a transformative effect on demand for AI-related services. Moreover, in many sectors China benefits from abundant large-scale datasets including digital payments, mobile platforms, and Internet of Things-enabled smart factories that support model training and product and service delivery at scale, a comparative advantage over many EU countries.

⁵ Oliver Wyman analysis of public reports on AI and productivity: European Commission (2025-2026); OECD (2025); European Investment Bank (2025-2026); Bruegel (2024); Congressional Budget Office (2024-2026); Federal Reserve Board/Federal Reserve Bank (2025-2026); National Bureau of Economic Research (2025-2026).

Five proposed building blocks for measuring AI adoption

The government has established the AI and the Future of Work Unit with a remit to consider the labour market implications of AI adoption. Its work programme has since been incorporated into the AI Economics Institute (AIEI), which retains a dedicated Future of Work workstream.

The AI Economics Institute (AIEI), established jointly by HM Treasury (HMT) and the former Department for Science, Innovation and Technology (DSIT), reflects the importance of bringing together technological and economic analysis. AIEI's work spans three areas: 1) building and analysing the evidence base (a core function); 2) developing models and scenarios (a core function); and 3) informing policy and engagement.

Within objective 1, AIEI will be working,⁶ *inter alia*, on:

- **Measurement and statistics:** Develop and promote improved measurement frameworks for AI's economic effects, while supporting HMT and other government departments to become informed customers and articulate the user needs those statistics should serve.
- **Empirical research:** Establish and maintain an evidence-based picture of AI's economic effects while deepening understanding of the mechanisms driving them. Priority questions include: How is AI being adopted across firms, sectors, and regions in the UK? What determines the pace, pattern, and depth of adoption? What are the effects on productivity, and through what channels do they operate?

Internationally, the measurement question has been considered by the Organisation for Economic Co-operation and Development (OECD), Eurostat, and the Brookings Institution, while in the UK, the former DSIT, the Office for National Statistics (ONS), and Cambridge Bennett have considered the issue.⁷ Independent analysis of these sources indicates that:

AI adoption is a multi-dimensional phenomenon

AI adoption should be measured as a multi-dimensional phenomenon, not a single binary indicator. It should go beyond “use” and “non-use” towards combinations of incidence, intensity, use cases, enablers, and outcomes, with firm size and sector built into the framework.

Survey data should be complemented by objective data sources

Survey data should be complemented by objective data such as provider telemetry, expenditure, and administrative and financial accounts. An additional overlay, developed by the ONS and used by Cambridge Bennett, is a management practice assessment scored on continuous improvement, key performance indicators, targets, and employment practices.⁸ This reflects a hypothesis that better-run firms are more likely to test and scale AI more quickly and efficiently.

⁶ For further detail refer to: HMT, DSIT (2026), [AI Economics Institute \(AIEI\) — Prospectus](#).

⁷ Oliver Wyman analysis of public reports from OECD, Eurostat, Brookings, DSIT, ONS, and Cambridge Bennett.

⁸ Refers to ONS' Management Practice Score. For further detail, refer to: ONS (2025) Office for National Statistics (ONS), released 24 March 2025, ONS website, article, [Management Practices and the Adoption of Technology and Artificial Intelligence in UK Firms: 2023](#) and Coyle, D., Nguyen, N., Lourenze Poquiz, J., and Riley, R. (2026). More than just plug and play: Early evidence on organisational capital and AI adoption. Bennett School of Public Policy, University of Cambridge. DOI: <https://doi.org/10.17863/CAM.12974>.

Job family and functional area should underpin assessment of labour market consequences

Job family and functional area should be the primary segmentation basis for determining the labour market consequences of AI. This is consistent with the former DSIT's AI adoption research,⁹ which proposed a survey-based methodology to measure current AI use, planned adoption, the AI technologies used, business areas' use of AI, piloting or live use, the degree of human oversight, readiness to scale or implement, budgeting, barriers, trust, and impacts on productivity, processes, and revenue. Another DSIT review¹⁰ argues that barriers differ by sector, firm size, and technology type. It identifies technology maturity and safety, alongside financial constraints, process-change costs, and access to talent, as key concerns.

Based on these considerations, a proposed data and methodology hypothesis for consideration by AIEI is set out below, built around five key building blocks:

Layer 01	Incidence A simple binary indicator of whether AI is used at all, by job family and functional area and, where relevant, by technology type. Key sources could include surveys, provider telemetry, and expenditure.
Layer 02	Maturity and intensity Pilot and live use, frequency, number of tools, budget, and the importance of AI to core processes. Key sources could include surveys, provider telemetry, expenditure, and administrative and financial accounts.
Layer 03	Organisational enablers Data infrastructure, skills (measured across both AI literacy and role-specific and specialist capability), management practices, finance, and external support. Key sources could include surveys and administrative and financial accounts.
Layer 04	Barriers and trust Cost, use-case identification, regulation, safety, ethics, and oversight. Key sources could include surveys and provider telemetry.
Layer 05	Outcomes Productivity, process change, revenue, innovation, and workforce effects. Key sources could include administrative and financial accounts.

⁹ DSIT (2026), [AI Adoption Research](#).

¹⁰ DSIT (2025), [Technology Adoption Review 2025](#).

Market supply will contribute to economic growth

In addition to the benefits accruing to users of AI products and services, economic systems will benefit from growth inherent in investment and activity in the AI stack as defined above.

Currently, the bulk of these benefits is accruing to US companies (and Chinese companies for their home market). A European or UK user of AI-related services generates revenues that are accruing to US-based companies such as hyperscalers and AI model specialists. Revenue growth that we are observing from US-based players is evidently generated globally (except China) over and above services to US clients. Therefore, the bulk of GDP growth generated by the AI stack is currently being captured in the US by firms that serve global demand. This concentration of the AI stack in a small number of US and Chinese suppliers is not only a question of where economic value is captured; it is also a resilience and supplier-dependency risk for the firms and economies that rely on them.

Many market observers are increasingly concerned that energy infrastructure is becoming a key constraint on this build-out, particularly the ability of grid, transmission, and distribution networks to connect new demand at pace. This is a different challenge from overall energy generation capacity, especially in the US, where the issue is less the availability of energy in aggregate and more the infrastructure, permitting, and connections needed to deliver it to high-demand sites. Authoritative US sources conclude that:

The AI stack is already a measurable component of GDP through investment (capital expenditure), services output (data processing and cloud services), and energy consumption. Its contribution is observable in national accounts, specifically through measured investment flows into data centres and servers, metered electricity consumption of data centres and particular large training runs, and local and regional economic impact studies.¹¹

A dominant share of the \$700 billion to \$800 billion capital expenditure announced by hyperscalers for 2026 is targeted at build-out in the US (to serve global needs).

Triangulating capital expenditure, energy consumption, and IT investment in national accounts, with considerable error margins, suggests that 2%–5% of US GDP may be directly attributable to the AI stack, or somewhere in the region of 0.2%–0.5% annual GDP growth in the last few years.¹² Recall that US GDP growth in 2025 was ~2.2%,¹³ compared to 1.3%–1.5% in Europe,¹⁴ indicating that 30%–40% of the “extra” US growth could be attributable to the AI stack.

11 Oliver Wyman analysis of public reports; Brookings (2023-2026); US Department of Energy (2026); International Energy Agency (2025); US Bureau of Economic Analysis (2026).

12 Oliver Wyman analysis.

13 US Bureau of Economic Analysis (2026), [GDP 4th Quarter and Year 2025](#).

14 Eurostat (2026), [GDP database](#).

However, global firms are recognising that a global service model is perhaps not suited to a world in which geopolitical tensions, system-level cyber threats, and regulatory divergence create material risks to the cross-border transfer of data, models, and services.

Therefore, most firms are investing in some onshore capacity outside the US (for example, the Gulf, the UK, and the EU) to ensure that they can be proximate to their clients, assuage concerns over data movement across borders, mitigate regulatory risks, increase the resilience of services, and take advantage of government policy that favours the emergence of a domestic/onshore industry.

Authoritative European sources¹⁵ conclude that some degree of industry build-out in Europe would have an appreciable impact on economic growth:

Low estimate: Under this conservative scenario, most hardware is imported and margins are repatriated. European-based capital expenditure amounts to \$100 billion to \$300 billion, contributing around 0.1%–0.5% of GDP over five to 10 years, or approximately 0%–0.1% of annual GDP growth through AI stack investment and operating activity.

Mid estimate: A larger share of value is captured domestically, including construction, hosting, cloud services, and local managed services. European-based capital expenditure amounts to \$300 billion to \$1 trillion, contributing around 0.5%–2% of GDP over five to 10 years, or approximately 0.1%–0.5% of annual growth through data centre construction, hyperscaler expansion, and associated services.

High estimate: Under this more aggressive scenario, local chip fabrication plants, large hyperscaler operations, and associated services and supply chains are established domestically. European-based capital expenditure exceeds \$1 trillion, contributing around 2%–5% of GDP over five to 10 years, or more than 0.5% of annual GDP growth during periods of heavy capital expenditure, including the construction surge and commissioning of large campuses. This scenario is corroborated by a review of European Commission regional impact analyses suggesting that local GDP contributions can exceed 5% cumulatively over the investment cycle.

Given this expected economic impact, coupled with concerns over AI resilience and strategic independence, the EU has been active in attempting to build a home-grown AI stack, often involving US firms establishing onshore operations in Europe.

¹⁵ Oliver Wyman analysis of public reports: European Commission (2025-2026); OECD (2025); European Investment Bank (2025-2026); Bruegel (2024).

Over and above regulatory, certification, and standards initiatives that have been criticised by parts of industry (for example, the EU AI Act), government policy has provided six broad areas of support to the build-out of sovereign and domestic AI stacks with considerable commitment of taxpayer funds:

1. **Compute and hardware:** European High Performance Computing initiatives (EuroHPC) provide continent-scale supercomputing and exascale resources for AI research and training. The European Chips Act and proposed Chips Act 2.0 aim to strengthen EU capacity in advanced semiconductors through manufacturing, design, and supply resilience.
2. **Data and governance:** The EU Data Strategy, Data Governance Act, and Data Act, create rules and mechanisms for access, sharing, reuse, and governance of public and private data across the EU. With the independent Gaia-X, there is support for enabling a federated cloud infrastructure, and technical and operational rules to enable sovereign, interoperable European cloud services.
3. **Skills, research and development, and innovation:** Horizon Europe and Digital Europe provide funding for large-scale digital transformation, high-performance computing, edge computing, AI research, and pilot deployments. European Digital Innovation Hubs give small and medium-sized enterprises and public administrations access to tools, testbeds, and expertise.
4. **Power:** The EU Green Deal, REPowerEU, and decarbonisation targets support renewable generation like offshore wind and solar power, grid upgrades, and energy efficiency to meet energy demand, including growing data centre demand. Trans-European Networks for Energy (TEN-E) supports cross-border transmission, interconnectors, and priority energy infrastructure.
5. **Streamlining planning:** EU, national, and local governments are working to streamline consenting, planning, and grid connection for data centres.
6. **Government procurement:** Public procurement pilots and industrial initiatives collectively support AI adoption across many sectors including healthcare, transport, and defence.

It is clearly too early to assess the outcome of these policy initiatives. It is, however, clear that geopolitical factors will increase business focus on resilience, which is very likely to tilt investment into major economic blocs to minimise dependency on US- and China-based service provision.

For the UK, the implication is not that it should replicate the EU's approach wholesale or seek to build a fully sovereign AI stack in every layer. The opportunity is more targeted: to become a leading AI-enabled economy while capturing more value from the parts of the AI stack where the UK has credible strengths, including research, deployment services, assurance, trusted data, and secure hosting. This points to the need for a more selective national approach, focused on adoption at scale and on the layers where the UK is genuinely competitive. Realising this also means ensuring the benefits are felt across the UK, with AI investment, infrastructure, and adoption spread across the nations and regions rather than concentrating them in a few locations, so that the economic gains support growth and productivity nationwide.

Recommendation

Realising the productivity prize from AI will depend on diffusion nationwide. The challenge, and opportunity, for the UK is therefore to ensure that AI is adopted across sectors, regions, and firms in ways that raise economic output, improve services, strengthen business resilience, and contribute to raising living standards. This will not happen without concerted effort across government, industry, and civil society.

RECOMMENDATION

Make AI adoption a national productivity mission, focused on diffusion across sectors and firm sizes

BIST, the Cabinet Office, HMT, UK industry, and civil society should work jointly to confirm strategy, aspiration, and implementation governance and create a shared framework and timeline for action. A shared national framework would provide a common ambition, clear priorities, and coordinated implementation, aligning public and private efforts around the common objective of accelerating AI diffusion.

A joint measurement framework developed through the AIEI would help focus support where productivity gains are greatest and provide a consistent mechanism for tracking progress and maintaining momentum over time.

02

Enterprise
demand and AI
infrastructure

AI adoption in businesses is both broad and selectively deep, with many firms experimenting

Although public disclosures on AI investment remain limited and the definition of investment is highly variable, across sectors, companies of all sizes are making significant investments in AI. Comprehensive UK-specific spending data is unavailable. However, organisations consulted for this report consistently confirmed that a sustained increase in AI spending is expected. This finding is further supported by research covering a subset of organisations, which notes that AI-related investment in the UK is expected to grow.¹⁶

Embedding AI as a means of doing business is front and centre for many in industry as they deliver on their broader strategic objectives. This pattern is reflected in the experience of many taskforce members, who highlighted significant AI usage within their businesses resulting from encouraging experimental AI use in everyday tasks. Several taskforce members reported investing in priority use cases that support delivery of broader strategic goals. One survey noted that 61% of CEOs confirmed that they are preparing to move to scaled implementation of AI agents.¹⁷ However, there is also variation in who is adopting: large and mid-sized businesses are more likely than micro firms to be adopters, and adoption rates are highest within information and communication, finance and real estate, and business services and administration sectors.¹⁸

“We want to think of ourselves as truly an AI-first company in everything that we do”

— Taskforce member

“The greatest opportunity in terms of AI is where you can use it to augment work, create additional capacity, or enable people to work in a much more productive way and at a higher level”

— Taskforce member

Adoption is selectively deep, with use case application areas clustering in individual employee productivity tasks, customer sales and services, operational efficiency and back-office functions, and analytics and insights. Many firms are still in an exploration or piloting phase in major areas of AI application, but a sizeable number are now shifting to scaled or agentic AI deployment, indicating deeper penetration across specific business processes.¹⁹

¹⁶ Barclays (2025), [Business Prosperity research 2025](#).

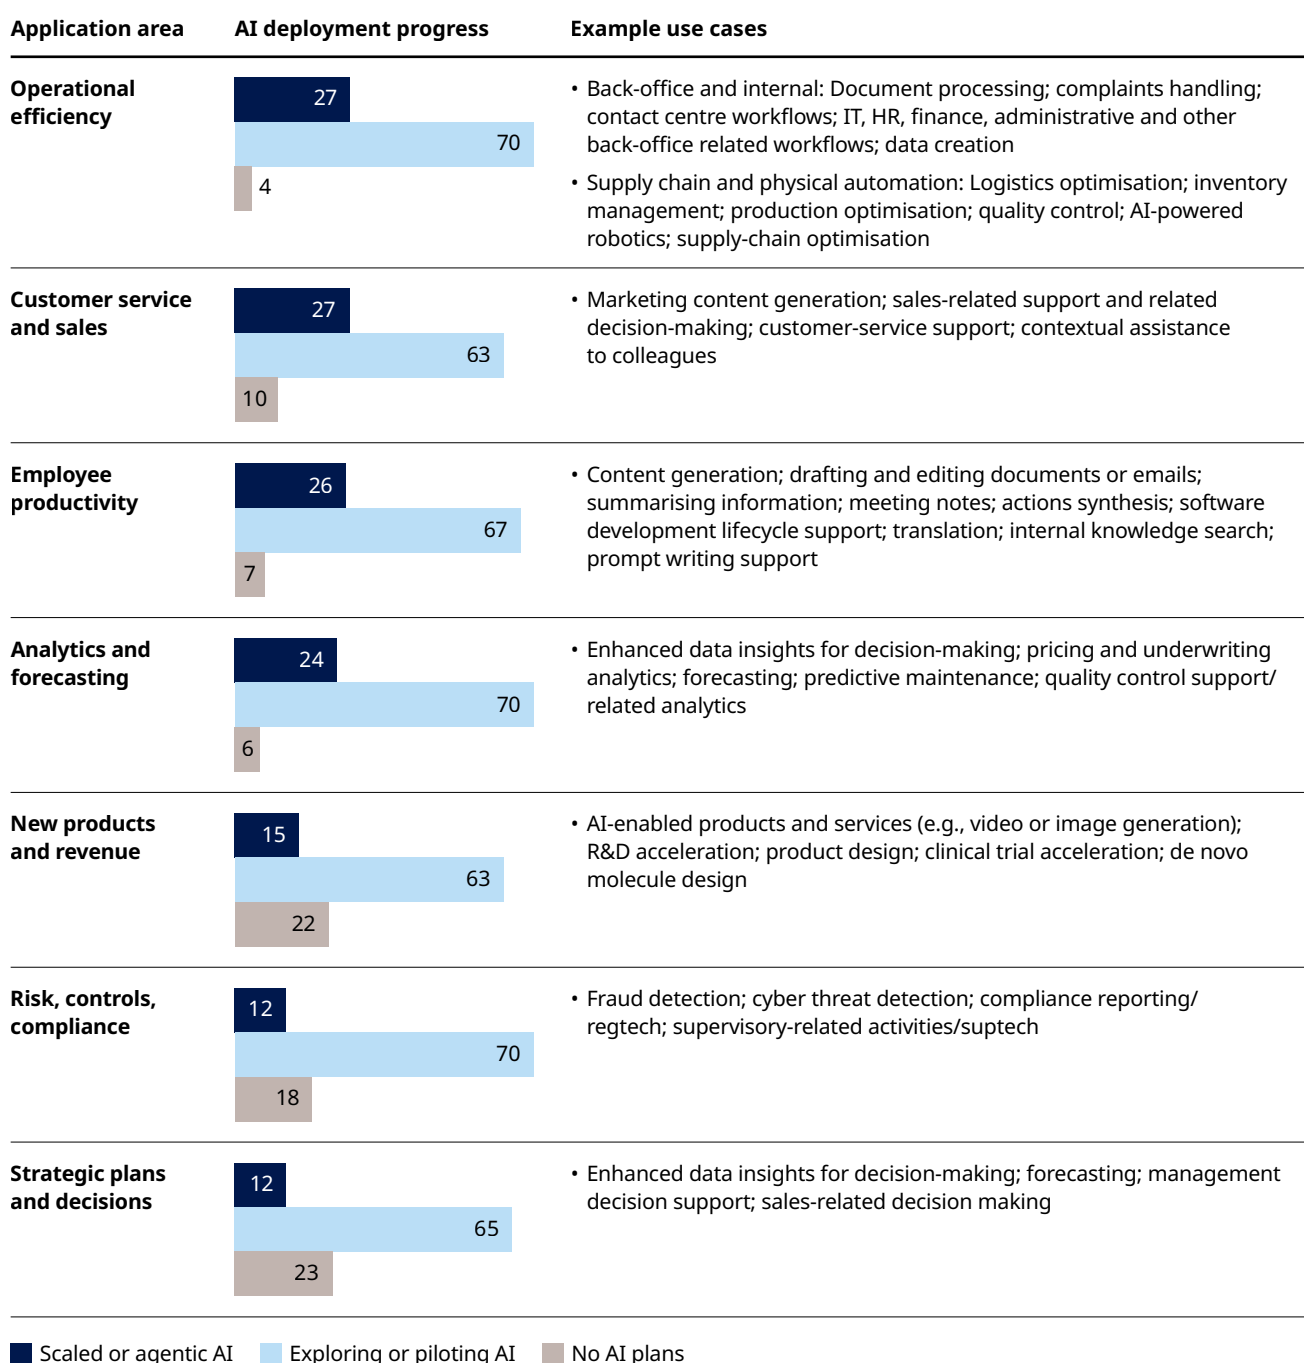
¹⁷ IBM and IBM Institute for Business Value (2025), [CEOs Double Down on AI While Navigating Enterprise Hurdles](#).

¹⁸ DSIT (2026), [AI Adoption Research](#).

¹⁹ Oliver Wyman Forum x NYSE Survey 2026. Refer to [The CEO Agenda 2026](#).

Across these applications, we note a range of use cases. See exhibit 1 below.

Exhibit 1: AI deployment progress by application area,²⁰ and example use cases²¹



²⁰ Oliver Wyman Forum x NYSE Survey 2026. Question: "What stage has your company reached in deploying AI?" (N=415). Figures may not sum to 100 due to rounding. Refer to [The CEO Agenda 2026](#).

²¹ Use cases reflect findings from independent analysis, interviews and synthesis of authoritative official sources: DSIT (2026), [AI Adoption Research](#); DSIT (2026), [AI Adoption Plan: Advanced Manufacturing](#); (Eurostat (2026), [The use of artificial intelligence technologies in the European Union — Key results — 2026 edition](#)).

There is a substantial ROI gap between adoption leaders and laggards within industries

This return on investment gap suggests that the AI divide is no longer a technology divide — it is an execution divide. While many leaders see AI as a major opportunity, only a relatively small group of firms have translated that objective into enterprise-wide transformation and measurable business and profitability outcomes. For example, the former DSIT found that while 35% of AI adopters reported delivering new or improved products or services, only 12% had realised measurable increases in revenue.²² The clear implication is that AI benefits are not being distributed evenly. Many firms' AI adoption remains characterised by pilots, isolated use cases, and bottom-up experimentation, with limited evidence of enterprise-wide deployment and structural process redesign. In contrast, profitability benefits are accruing disproportionately to adoption leaders that have successfully embedded AI into their operating model and core end-to-end business processes.²³

Forty-nine percent of AI deployment leaders are meeting or exceeding return on investment expectations, compared with just 15% for deployment laggards.²⁴

Put differently, rapid experimentation with frontier models and relatively unconstrained investment do not enable many organisations to deliver targeted benefits. There is value in linking use cases to clear business outcomes, supported by evidence of tangible benefits.


Leading firms enjoying returns from AI investment have some common characteristics.

Firms considered “value generators” or “transformers” in the adoption phase described in the schematic below are enjoying outcomes that result from the implementation of best practice.

²² DSIT (2026), [AI Adoption Research](#).

²³ Oliver Wyman Forum x NYSE Survey 2026. Refer to [The CEO Agenda 2026](#).

²⁴ Ibid.

Exhibit 2: Typical evolution of AI enablement in companies²⁵**AI wild west**

- High enthusiasm and energy
- First visible successes and quick wins
- Many ideas emerging from different teams

AI coordinated

- Clear priorities and focus on fewer, bigger topics
- Centralized governance and standards
- Ability to optimize funding and talent allocation

AI everywhere

- AI is embedded in everyday work; business-led, not project-led
- Decentralised innovation within clear guardrails
- Broad adoption across functions and levels

“So yes, they’ve got the experimentation, they proved out some ROI, but they can’t then scale that. And that’s just large transformation. Even without AI, [getting the foundations in place to enable them to scale] was a problem for many institutions”

— CBI AI and Data Protection working group member

“Just throwing out [specific AI tools] to everybody is a big no-no for me. Having a structured, designed, value-driven, metric-based deployment is how we’re going to do it”

— Taskforce member

“Let’s match up what the business priorities are, and let’s resource these things in a way that we go after it intentionally”


— Taskforce member

²⁵ Oliver Wyman | Quotient AI (2026).

The next phase of AI adoption will shift from discrete use cases, with benefits dependent on firms implementing best practice

The following components reflect a synthesis of emerging best practice as firms seek to scale responsible adoption and deliver AI-linked benefits. See exhibit 3.

Exhibit 3: Components of best practice



Align leadership and enterprise

Leadership alignment is a core enabler of scaled AI adoption, and AI is increasingly shaping leadership decisions and becoming central to business strategy.²⁶ Organisations that move beyond isolated pilots often have clear executive sponsorship, a company-wide vision for AI use (including workforce transformation), employee buy-in, and explicit priorities for investment, with business leaders actively shaping the approach rather than leaving AI deployment as purely the domain of IT.

²⁶ CBI (2026), [AI Skills Report: Building an AI Ready Workforce](#).

Emerging best practice is therefore to set the approach “from the top”, with leadership visibly sponsoring initiatives and embedding AI as a lever to better deliver its own strategic objectives. This also enables sequencing of AI-related initiatives and use cases.

“We have been going through this transformation for a couple of years now. We want to think of ourselves as truly an AI-first company”

— Taskforce member

“We have updated our 2030 ambition to include harnessing the power of AI”

— Taskforce member

Redesign workflows and operating models

Leading adopters view AI as a tool for changing how work is done, not simply a way to make existing tasks faster. Emerging best practice is therefore to treat AI adoption as a business transformation programme rather than a technology roll-out, starting from strategic business challenges — such as productivity, customer service, resilience, risk management, speed to market, or cost competitiveness — and identifying where AI can materially change how work is done.

“Artificial intelligence taking people’s jobs away? Actually, no — it’s about reimagining what you want your staff to be working on”

— Taskforce member

“Two years ago, AI was in a very different place. A real, live decision for us now is how we should be thinking about ERP system configuration and the scope of AI”

— Taskforce member

Establish governance

Governance is a critical foundation for moving from experimentation to scaled deployment. Leading adopters establish AI-related governance, spanning risk and compliance oversight, data governance, model governance, security and resilience controls, privacy protections, and guardrails for usage including responsible AI frameworks and human oversight for high-risk decisions. There is a range of resources available to support effective governance, further indicating its importance.²⁷ Leaders must be able to link AI use with business strategy and risk, and employees require the confidence to use tools safely and responsibly.²⁸

27 For example: [ISO/IEC \(2023\), ISO/IEC 42001: Artificial Intelligence Management Systems](#); NIST (2023), [Artificial Intelligence Risk Management Framework](#); Information Commissioner’s Office (2022), [AI and Data Protection Risk Toolkit](#).

28 CBI (2026), [AI Skills Report: Building an AI Ready Workforce](#).

Governance is a critical best practice for enabling the trusted adoption of AI, and is explored in greater detail in Chapter 3.

“We have to be 100% accurate all the time. We have guardrails so that ultimately humans always retain full authority to make sure it is accurate. That’s something we pay a lot of attention to”

— Taskforce member

“We are working very closely with our supply chain now on everything from cyber and AI to controls and policies”

— Taskforce member

Balance centralisation with core oversight

Leading adopters acknowledge the value of individuals experimenting with AI and identifying value in their workflows while also having an overall consistent strategic approach top-down across the organisation. Local implementation and oversight are therefore best for local innovation (for example, use case identification, process redesign, and adoption within business units). Core capabilities such as governance, security, enterprise architecture, approved platforms, data standards, and vendor management are typically best managed centrally to provide consistency, reduce duplication, and manage risk.

“This is as much a business priority as a technology priority. We have clear governance to enable us to prioritise AI use across the business with guardrails to measure and realise value”

— Taskforce member

“We have struck a balance between encouraging broad adoption and experimentation across the enterprise, letting a thousand wildflowers bloom and putting guardrails around it”

— Taskforce member

Invest in core enablers to allow AI scaling

Scaled adoption also depends on the foundations beneath the AI layer: high-quality and integrated data, secure and flexible technology architecture, reliable connectivity across sites, clear supplier management, common tooling, workforce skills (addressed in detail in Chapter 4), and workforce engagement and adoption metrics that focus on outcomes rather than usage alone. Leading adopters recognise that layering AI tools onto existing processes is unlikely to deliver sustained benefits, and investment in core enablers that allow AI to scale is key.

These investments are not optional extras; they determine whether AI systems can be deployed safely and reliably across the organisation. These investments are similar to past waves of technological innovation, where value was created when firms adapted how work was organised and invested in the complementary assets needed to scale: ICT infrastructure, workforce skills, management capability, organisational capital, and robust governance.²⁹

“The absolute fundamentals are to do with data. There is no point doing anything else”

— Taskforce member

Move away from ad hoc experimentation

A culture of experimentation is a common feature of leading adopters. Many of these firms are giving employees access to approved tools, creating forums to test and share ideas — such as hackathons, promptathons and competitions — and recognising teams that identify valuable use cases. Early experimentation is valuable because it builds confidence, exposes employees to new ways of working, and helps organisations understand which use cases are realistic. However, pilots can also become a substitute for transformation if they are not connected to business priorities, workflow redesign, data readiness, and value measurement. Leading firms therefore strike the balance between facilitating “bottom-up” experimentation while also channelling it towards a prioritised AI portfolio and credible AI roadmap, with some capabilities managed centrally.

“You need to create the culture for experimentation”

— Taskforce member

Measure value to refine implementation

Measuring the benefits of AI is important as organisations move from experimentation to scaled deployment and investment and seek to measure the value generated. A range of metric types are available, including those measuring active usage, outcomes, and impact. Early approaches often focused on activity-based metrics (for example, the number of employees using a tool), but leading adopters recognise the limitations of using these measures alone. Effective measurement also helps identify which use cases should be scaled, where additional process redesign may be required, and whether investments are generating the expected returns or productivity gains.

²⁹ OECD (2024), “Fostering an inclusive digital transformation as AI spreads among firms”, OECD Policy Briefs, No.8, OECD Publishing, Paris, <https://doi.org/10.1787/5876200c-en>.

Many organisations would value practical, trusted guidance on how to adopt AI

Most businesses want to adopt AI but lack practical, trusted guidance on how to do so.

Some struggle to identify where AI could add the most value to their operations. Others know what they want to achieve but cannot always find credible guidance on how to deliver it, amid a fragmented landscape of overlapping initiatives. Larger firms can absorb the cost of working this out through trial and error; smaller firms cannot. Especially in the case of smaller firms, which lack the teams or time to interpret complex guidance, trusted intermediaries can make AI adoption more accessible. Local universities, education and training providers, trade associations, catapults, professional bodies, banks, insurers, and government-supported programmes can help translate national AI strategy into sector-specific and regional support.³⁰

Firms can move faster by drawing on the support around them and learning from credible examples (in ways that do not require them to give up sources of competitive advantage).

The UK has a strong ecosystem of business organisations, universities, professional bodies, local institutions, and technology communities that can be used more deliberately as a diffusion engine that is particularly effective for current non-adopters. Knowledge sharing will be most valuable when it is linked to action, with use cases clearly demonstrating how value was delivered in a particular setting.

“There are no verified blueprints for us to rely on. Every business has some similarities in some way. It can’t be that hard to share that information”

— Taskforce member

“I would love to have guidance to say: these are the lessons learned, this is what good looks like, this is the right human oversight you need to have, this is where the human in the loop plays an important role, this is where accountability sits when an agentic AI makes a decision”

— Taskforce member

³⁰ CBI (2026), [AI Skills Report: Building an AI Ready Workforce](#).

The current debate on AI sovereignty does not reflect how many businesses approach the issue in practice

Across industry, there is a consistent view that sovereignty is not a binary choice between global ecosystems and fully domestic infrastructure, but a spectrum. Most enterprise workloads can sit on global cloud infrastructure, while a limited subset require tighter control over where data is held and who can access it. In many cases, the practical solution lies in a hybrid approach, combining global infrastructure with domestic storage and firm-level controls over sensitive elements such as data, encryption keys, and access.

“We’re not arguing for a world in which the UK can do everything on the AI stack, but I think there is a world where the partnerships set to deliver AI infrastructure could be widened, and there is a role for British companies to play in data centres”

— Taskforce member

The clearest exception to this view on sovereignty relates to resilience. From a critical national infrastructure perspective, there is a case for maintaining some domestic capability as a safeguard against the potential loss of access to foreign-controlled systems, particularly in light of recent export controls. At the same time, businesses emphasise the importance of continuing to use global solutions where they offer clear advantages. Data residency can also present a near-term constraint where infrastructure is not available within the UK or Europe, limiting adoption of some leading services.

“Some of this is basic supply chain resilience. Of course we can cooperate, and of course we can use solutions built elsewhere where it makes sense. But why would we not also build our own capabilities?”

— Taskforce member

For smaller firms, sovereignty is rarely a primary concern. Evidence consistently shows that barriers to AI adoption are far more immediate and practical, including skills gaps, cost constraints, and limited awareness. Many smaller businesses have little familiarity with the concept of AI sovereignty, and it does not shape their adoption decisions.

Overall, most businesses are not calling for a fully sovereign technology stack across every layer. However, many do see value in the UK building stronger domestic capability as a resilience safeguard. This appetite has grown as recent export controls have made the risk of losing access more tangible. What they favour is a pragmatic approach that strengthens the UK’s attractiveness as a destination for investment, including co-investment models, a competitive regulatory environment, and access to global-scale compute. The distinction that matters is between locating AI investment and infrastructure in the UK, which businesses welcome, and mandating that firms run their own workloads only on domestic infrastructure, which can constrain how they use AI and slow adoption.

Fragmentation is the underlying concern, since duplicated infrastructure, divergent rules, and country-specific service variants raise costs and narrow choice, particularly for smaller firms, while security, resilience, control, and interoperability do not require UK services to be separated from global ones. It is compulsory data-localisation requirements, not domestic investment, that businesses caution against.

This sits alongside the wider economic case for building more of the AI stack in the UK.

That case is about where investment, infrastructure, and value are located, whereas the caution above is about how firms run their own workloads. The two are complementary. A stronger domestic stack can improve resilience and capture more value, while businesses continue to use global-scale infrastructure for most of their needs.

Security and resilience are increasingly shaping how firms adopt AI. As AI moves into core processes, cyber and security considerations are becoming a more prominent factor in adoption decisions, influencing which systems firms deploy, how they architect them, and which vendors and infrastructure they are willing to rely on. AI raises the stakes on both sides: it expands the attack surface as more systems connect and more decisions are automated, while also becoming a tool for attackers and defenders alike. Some firms are also moving more inference to edge and on-premise infrastructure, which reduces dependency on a single external service, keeps sensitive data within their own environment, and supports operation if a supplier or connection is disrupted, though it places its own demands on device security.

For many businesses, confidence that AI can be adopted securely is now a precondition for scaling it.

“Cyber is moving just as quickly as data and AI in certain parts, and it only gets quicker because of AI and the third-party actors involved”

— Taskforce member

“We are learning the ways in which to secure an enterprise, making sure that the deployments we have meet all of the current security and other requirements. Having a centre of excellence for AI is helpful”

— Taskforce member

This pattern of adoption will increase the importance of the parts of the AI stack closest to practical deployment

As summarised above, AI adoption demand will shift from discrete AI use cases towards business process redesign, data integration, change management, risk, compliance, enterprise-wide implementation and a focus on cost, security, and resilience. These shifts will have significant implications for the AI services required. Most business adoption will not depend on frontier models alone; it will rely on smaller, task-specific, and diffusion-layer models that can be applied to firms' own data, processes, and workflows.

This shift also changes which parts of the stack matter most. Much of the value businesses are pursuing comes not from the most advanced models, but from applying capable, often smaller or task-specific models to their own data, processes, and workflows. For many use cases, a model tuned to a specific task and dataset is cheaper, easier to govern, and more than sufficient, and can increasingly run on local or edge infrastructure, which lowers the barrier to adoption for smaller firms. The parts of the stack that most shape adoption are therefore those closest to the point of use, deployment and integration, data quality and readiness, assurance, and the ability to run and adapt models cost-effectively. This plays to the UK's comparative strengths in services, standards, assurance, applied research, and sector-specific expertise, and to a particular strength in deploying AI safely in regulated environments such as financial services, insurance, health, justice, and defence. The priority is therefore to ensure that the wider AI stack supports this kind of adoption, while also strengthening the UK's long-term capability in public compute, AI-ready data, frontier research, sovereign AI capability, network connectivity, hardware innovation, and energy efficiency. Delivery of the government's ambition for standalone 5G across all populated areas by 2030,³¹ alongside Project Gigabit,³² is an adoption enabler as much as a connectivity objective. This would help ensure that the UK is not only an adopter of AI but also a competitive place to develop, test, and scale the technologies that underpin it.

This pattern also brings data provenance questions closer to the point of adoption.

As more firms train or fine-tune models on their own data, questions can arise about what that data contains and how it can be used, since proprietary datasets may include third-party licensed material, customer and supplier information, and content whose rights position is not always clear. Many businesses would benefit from clearer, practical information on data licensing, provenance, and rights to use, so they can adopt and adapt AI responsibly. Broader questions about copyright and AI are the subject of an ongoing government process, which concluded in March 2026 that more evidence is needed before major reform.

31 DSIT (2023), [UK Wireless Infrastructure Strategy](#). Ambition restated in HM Government (2026), [Digital and Technologies Sector Plan: Year One Update](#).

32 For further detail: House of Commons Library (2026), [Gigabit broadband in the UK: government targets, policy and funding](#).

Reflecting this, while the US and China have clear advantages across the full AI stack created by vast public and private spending that is out of reach for other countries, the UK does not need to compete across every layer. That said, maintaining research and sovereign capability where it matters for resilience and security remains a strategic priority.

As the industry matures and there is a greater focus on service resilience and trusted adoption, the UK has some important comparative strengths in several parts of the stack, concentrated in the same application and deployment layers where adoption demand is strongest.

These strengths can support the emergence of a national AI service or developer industry in areas of genuine comparative advantage. The table below sets out these relative strengths and weaknesses compared with other G20 countries excluding China and the US.³³

As described in Chapter 1, a significant portion of the capital expenditure associated with the development of the AI stack is expected to come in future years. Compared with a heavily US-centric investment pattern, future years are expected to see significantly more deployment in other economic blocs as concerns over cost, resilience, and cybersecurity increase, driven in part by unpredictable geopolitical developments, with the greatest impact on the choice of location for data centres (and the associated infrastructure). Therefore, most firms are investing in some onshore capacity outside the US (for example, the Gulf, the UK, and the EU) to ensure that they can be proximate to their clients, assuage concerns over data movement across borders, mitigate regulatory risks, increase resilience of services, and take advantage of government policy that favours the emergence of a domestic/onshore industry. The UK can therefore leverage its strengths in safe, large-scale forward deployment of AI in regulated and essential services (drawing on its strengths in financial services, insurance, health, defence, government, legal, and professional services), research, cybersecurity, compute capacity, assurance, and finance and insurance to capture a greater share of the capital expenditure expected in future years and develop a local AI service industry.

³³ Reflects qualitative indicative assessment from own analysis.

Exhibit 4: AI stack considerations

Element of the stack	Key future trends and market examples	Comparative UK position
Forward deployment/orchestration	Shift from experimental use cases to end-to-end process redesign and implementation, unlocking true productivity and efficiency. Competitive advantage to depend less on knowledge of AI models themselves and more on the ability to integrate AI into workflows, systems, and decision-making processes, including maximising the value and quality of data ecosystems and trusted data environments. Key organisations operating in this space will include model providers, hyperscalers, traditional business process outsourcers, specialist product providers, and advisory firms	Strong. Large digitally literate service industry anchored in regulated sector depth, managed services and operational transformation capabilities
AI models	While frontier model performance remains important, competition is increasingly shifting towards cost-per-token, efficiency, and the ability to operate within multi-model environments combining both open and proprietary systems with a view to deployment, integration, and domain-specific applications at a significantly lower cost-per-token	Strong. Strong AI research base and successful AI startup ecosystem, noting that clear leadership in scaled AI models is with the US (proprietary) and China (open)
Assurance/standards ecosystem	Growing importance as widespread AI adoption leads to adverse events: expected to emerge as key services including model validation/testing, monitoring, security, compliance, and independent assurance	Strong — early-mover advantage through the AI Security Institute in foundation-model safety and security, alongside the Centre for Data Ethics and Innovation's 2021 roadmap and subsequent application-assurance work, now embedded across GDS and other relevant teams
Data	Convergence of text and images, growth of synthetic data, increased focus on provenance and quality	Strong (latent advantage) — exceptional concentration of high-value public-sector (NHS), scientific (genomics, life sciences), and regulated-industry datasets (insurance, banking, legal services), with greater regulatory flexibility than the EU
Compute access and compute orchestration	Evolution towards a more distributed and heterogeneous industry structure: while hyperscale cloud providers are likely to remain dominant, businesses are increasingly exploring hybrid approaches that combine cloud, on-premise, and edge computing. Advances in specialised hardware, AI-enabled personal computing and, over the longer term, quantum-enabled workloads may further diversify compute technology.	Moderate — AIRR provides a credible sovereign-compute capability and research access, but UK remains dependent on international supply chains
Hosting infrastructure	Increasing emphasis on service continuity, resilience, and security. The emergence of national cloud offerings and greater scrutiny of operational resilience suggest that infrastructure decisions will increasingly be influenced by trust, control, and regulatory considerations alongside cost and performance	Mixed — strong cybersecurity capabilities and trusted institutions, but constrained by infrastructure economics and power availability
Physical infrastructure, connectivity and energy	Data-centre construction, power availability, grid connectivity, fixed and mobile network connectivity, cooling technologies, and energy economics are increasingly critical determinants of competitiveness as AI infrastructure scales. Standardisation, industrialisation, and shift to “behind the meter” integrated energy provision	Mixed — competitive network connectivity and cost, but uneven 5G coverage and capacity, slow permitting and energy cost disadvantage
Financing and risk mitigation	As AI infrastructure matures into a mainstream asset class, capital requirements are likely to diversify beyond specialist technology investors towards broader pools of institutional capital, accompanied by growing demand for risk mitigation and insurance solutions. This will lead to diversification from private markets into mainstream financing and risk mitigation techniques	Strong — world-class financing and insurance industries

Recommendations

For business

Firms should invest in the foundations needed to scale AI well and share practical lessons on what works through trusted knowledge-sharing forums.

RECOMMENDATION

Build the organisational foundations for scaled AI adoption

Align leadership and enterprise

AI-first organisations should review a credible AI roadmap regularly, asking whether the organisation is moving from experimentation to scaled adoption, whether investment is concentrated in the right processes, and whether the organisation is developing reusable capabilities rather than one-off solutions. Some use cases will create immediate productivity gains, while others will require upstream investment in data, workflow redesign, or operating model change before benefits can be captured.

A credible AI roadmap should therefore combine near-term deployments that build confidence with medium-term transformation bets that address the largest value pools.

Establish governance

Businesses should establish broader AI governance arrangements. Governance should define clear roles and responsibilities across business owners, technology teams, risk, compliance, legal, data, cybersecurity, and procurement, with proportionate oversight and escalation routes for higher-risk use cases. This should be embedded within business activity — such that colleagues across the organisation are held to the same standards — and linked to the proposed deployment playbook (see below). As shared norms develop, firms can align their governance with the voluntary Responsible AI Adoption Code and public pledge that the CBI is recommended to help convene (see Chapter 3), giving them a common reference point for responsible practice.

Balance centralisation with local oversight

Businesses should develop a practical deployment playbook that can be reused across functions and business units. The playbook should cover use-case selection, data readiness, procurement, security review, model evaluation, responsible AI assessment, employee training, change communications, benefits tracking, and post-deployment monitoring. A repeatable model reduces the cost and time required to move from pilot to production, while making it easier to apply consistent governance across different risk levels. Larger firms may establish a centre of excellence or federated AI transformation office; smaller firms may use a lighter playbook, external advisors, or trusted sector bodies. In all cases, implementation should be owned by the business function that will capture the value, not by technology teams alone.

Move away from ad hoc experimentation

Firms should move beyond disconnected use cases towards an implementation model that links AI investment to process ownership, operating model redesign, technology architecture, and change management. Such a portfolio should be linked to clear business owners, executive sponsors, data and technology requirements, risk ratings, resourcing, milestones, and benefit cases. High-potential pilots should move rapidly into proof-of-value stages with baseline metrics and explicit scale criteria. This enables low-value pilots to be stopped, consolidated, or reused as learning assets. This disciplined approach prevents initiative sprawl, focuses scarce specialist talent on priority use cases, and helps boards and investors distinguish genuine transformation from novelty-driven activity.

Redesign workflows and operating models

For priority processes, firms should map the end-to-end workflow, identify decision points, determine where AI can support analysis or action, and redesign roles around the combined strengths of people and systems. This includes defining human-in-the-loop or human-on-the-loop controls, updating standard operating procedures, changing escalation routes, redesigning quality assurance, and making sure employees understand how their responsibilities will change. Redesign works best when it starts from the needs of the people using the service and involves the employees who do the work, rather than being specified around the technology. Workflow redesign should also include customer, supplier, and regulator touchpoints, because the benefits of AI often depend on reducing hand-offs across organisational boundaries rather than automating a single internal step.

Invest in enablers

Businesses should invest in data quality, data lineage, metadata, integration layers, APIs, cloud and compute choices, security architecture, model operations, identity and access management, and robust vendor management. The core enabler of workforce skills is addressed in detail in Chapter 4. The appointment of data and process owners is also key, as is agreement on common tooling where possible and avoidance of a proliferation of unconnected AI solutions that create duplication, security risk, and operational fragility. A strong architecture preserves optionality across open and closed models and avoids unnecessary lock-in to individual suppliers.

Measure value to refine implementation

Firms should also measure adoption maturity, data and technology readiness, employee confidence, trust, process performance, customer outcomes, risk outcomes, and financial benefits. For major use cases, organisations should establish a baseline before deployment, define leading and lagging indicators, and review whether the process change is producing value. This may include productivity per employee, cycle times, service quality, error rates, customer satisfaction, revenue conversion, risk reduction, and employee experience. Measurement can then be used to refine workflows and investment priorities, and it should align with any government-led adoption measurement framework so that firms can benchmark progress against peers over time.

RECOMMENDATION

Strengthen business-to-business knowledge sharing

Build and engage with cross-sector learning networks nationwide

Businesses should engage with existing initiatives, local universities, education and training providers, professional bodies, sector councils, regional growth bodies, and CBI regional councils to cascade AI adoption insights into regions, cities, and AI Growth Zones. Trade associations can engage with TechUK's AI Adoption Coalition, for example. These networks should move beyond general awareness events and focus on practical implementation: common use cases, investment cases, operating model redesign, governance templates, supplier lessons, and benefits measurement. Cross-sector learning matters because some of the most valuable lessons will come from adapting methods used in one industry to another. For example, lessons from manufacturing process optimisation, financial-services assurance, or professional-services workflow redesign may be relevant far beyond those sectors.

Lean on trusted intermediaries to gain learnings

Businesses should treat these intermediaries as part of their adoption infrastructure and use them to reduce search costs and build confidence. Intermediaries should focus on practical questions: which use cases are credible, what data is required, what tools are available, how much implementation costs, what risks must be managed, and what skills are needed. Knowledge sharing will be most valuable when it is linked to action, with use cases clearly demonstrating how value was delivered in a particular setting.

Support the development of practical AI use-case libraries

Businesses should contribute anonymised examples of successful AI deployment to shared repositories of use cases and implementation guidance that are transferable to other organisations. Shared use-case resources should be developed to fill a demonstrable gap in existing vendor-, developer- or sector-led guidance. The default should be to signpost, federate, and improve access to high-quality existing material rather than recreate it.

A useful library should not simply list ideas; it should describe the business problem, process changes, data requirements, tools used, implementation costs, time to deploy, benefits realised, risks encountered, mitigations, and lessons learned. It should include examples from sectors beyond digital and technology, including manufacturing, retail, logistics, construction, agriculture, professional services, finance, health, and the creative industries. It should also include failed or paused pilots where appropriate, because these can help other firms avoid repeated mistakes. Ownership could sit with a trusted business body or a government-industry partnership, with sector curators maintaining quality.

Create peer-to-peer implementation sprints and mentoring cohorts

Business groups should create cohorts of firms working on similar use cases, such as customer-service automation, finance operations, predictive maintenance, professional-services knowledge management, or supply-chain forecasting. These cohorts could compare vendors, share procurement templates, test governance approaches, and benchmark benefits. Larger firms with mature AI programmes should be encouraged to mentor SMEs in their supply chains, where this does not compromise commercial confidentiality. This model can create diffusion effects through value chains and regions, while giving SMEs practical support that is more useful than generic guidance.

Share knowledge responsibly

Knowledge sharing should be structured to protect commercial confidentiality, personal data, intellectual property, and security. Firms should avoid sharing sensitive model details, proprietary datasets, or information that could create competition concerns. Instead, repositories and networks should focus on implementation patterns, governance methods, operating model lessons, and anonymised evidence. Clear rules will make firms more willing to participate. The governance foundations for these networks need to be worked through carefully at the outset, particularly where participants are competitors or operate under very different regulatory regimes, as the practical challenges of achieving meaningful exchange should not be underestimated.

A safe and trusted approach to sharing can help the UK build an evidence base for scaled adoption while reinforcing high standards of responsible AI use.

For government

The UK has established a strong policy foundation for AI through initiatives including the AI Opportunities Action Plan, UK Compute Roadmap, AI Growth Zones, National Data Library, AI Champions programme, and wider public-sector adoption initiatives. The next phase should focus on implementation: strengthening the UK's AI stack, accelerating deployment of secure, sustainable and financeable AI infrastructure, and using the government's own adoption, procurement, and convening power to diffuse AI adoption across sectors, firm sizes, and regions. Building on this existing momentum, we propose the following recommendations.

RECOMMENDATION

Strengthen the research and development ecosystem including public compute, AI-ready data, models, chips, memory, and energy efficiency

Expand and target the UK's public compute ecosystem, including the AI Research Resource (AIRR)

Government should align this recommendation explicitly with the UK Compute Roadmap, which sets out a modern public compute ecosystem, including expansion of the AIRR, mission-focused allocations, and a new national supercomputing service. BIST, UKRI, and HMT should ensure that AIRR access is targeted at areas of greatest strategic value for the UK, including AI for science, healthcare, climate, engineering biology, advanced manufacturing, defence, cybersecurity, financial services, and other high-productivity sectors. Access routes should remain open to researchers, universities, SMEs, and strategically important startups, with transparent allocation criteria, time-to-access targets, and clear support for users without large in-house compute teams. The AIRR should also remain flexible enough to support emerging quantum-AI workloads and should be coordinated with the National Quantum Computing Centre and wider quantum strategy.

Accelerate the National Data Library and trusted research environments

The National Data Library should move from discovery and kick-starter projects to a predictable pipeline of high-value, AI-ready datasets. Government should prioritise datasets where access can unlock economic or public-service value, while maintaining public trust and strong safeguards. This requires clear metadata, provenance, interoperability standards, secure access models, trusted research environments, synthetic data guidance, and routes for SMEs and researchers to access data without excessive friction. These environments should support the safe testing and deployment of AI in regulated settings such as health and financial services, not only research access. DCMS, working through the Government Digital Service (GDS), ONS, UK Research and Innovation (UKRI), and sector departments, should also provide greater clarity on data licensing, provenance, rights to use, and data-sharing governance, so that businesses can adopt and adapt AI on their own data with confidence. Wider questions about copyright and AI remain subject to the government's ongoing process, but the uncertainty is itself slowing adoption.

Government should urgently publish interim practical guidance for businesses on establishing the rights position of third-party material in their data holdings, and set a clear timetable for resolving the outstanding copyright and intellectual-property questions.

Build a stronger UK AI hardware and efficiency ecosystem

The AI stack strategy should more explicitly cover chips, memory, semiconductors, photonics, advanced packaging, cooling, edge inference, and energy-efficient AI hardware. This should form part of implementation of the Modern Industrial Strategy's Digital and Technologies Sector Plan, aligning innovation funding, procurement, skills, and investment support around the development and scaling of UK AI hardware firms. BIST, UKRI, UK Export Finance, and the British Business Bank should align the AI Hardware Plan, semiconductor policy, AI Growth Zones, and public procurement so that UK-developed technologies can access testbeds, finance, early customers, and export support. This should include support for energy-efficient inference, memory technologies, and system-level optimisation, because future competitiveness will depend not only on model quality but also on cost, resilience, and energy performance. Public compute and AI Growth Zones should be used as testbeds for UK hardware, software, and cooling innovations wherever this is commercially and operationally feasible.

RECOMMENDATION

Accelerate investment and regional deployment of AI stack infrastructure, including AI Growth Zones, power, planning, cyber resilience, financing, and insurance

Expand the AI Growth Zone model and integrate power, planning, and local delivery

BIST, the Department for Energy Security and Net Zero, and the Ministry of Housing, Communities and Local Government should build on the AI Growth Zones programme by offering a more coordinated route for AI infrastructure deployment. This should bring together the key enablers of AI infrastructure deployment, including planning support, grid connections, power availability, fixed and mobile connectivity, permitting, and wider local infrastructure requirements, while working with industry and infrastructure providers to address delivery challenges. The model should support sustainable energy solutions, storage, and grid innovations where they help unlock infrastructure delivery without displacing wider industrial demand or undermining whole-economy energy priorities. It should also link to the AI Energy Council's work on energy demand forecasting and to the Compute Roadmap's sustainability objectives. AI Growth Zones should not become isolated data centre locations; they should be regional platforms for compute, adoption, skills, supplier development, and public-service innovation.

Strengthen cybersecurity, resilience, and national security oversight

Building on the significant work already underway, the National Cyber Security Centre (NCSC) and Government Communications Headquarters, working with DCMS, BIST, the Cabinet Office, and relevant regulators, should consolidate and coordinate emerging cybersecurity and resilience initiatives for UK-based AI infrastructure and service delivery. The aim is a clear, coherent, and proportionate set of expectations that avoids duplication and gives businesses a single, navigable set of requirements. This should draw together the Cyber Security and Resilience Bill, which brings data centres into scope by classifying data infrastructure as a relevant sector and qualifying data centres as an essential service, and the AI Cyber Security Code of Practice implementation guide, which addresses risks such as data poisoning, model obfuscation, and prompt injection and is intended to support global standard-setting. Critically, AI infrastructure receiving public support should be required to demonstrate resilience by design, including incident reporting, continuity planning, supply-chain security, physical security, operational monitoring, and recovery arrangements.

Position the UK as the leading market for AI infrastructure financing, insurance, and risk-transfer solutions

Government, through BIST, HMT, and the relevant regulators, should invite relevant trade associations, the British Business Bank, UK Export Finance, the insurance market, and institutional investors to develop proposals that make the UK a global hub for financing and insuring AI infrastructure. This should include data centres, grid and power assets, cooling, security, hardware, software, sovereign AI companies, and adoption enablers. The work should align with the Modern Industrial Strategy, the Financial Services Growth and Competitiveness Strategy, and current reforms to risk-transfer markets. London and the wider UK financial ecosystem have an opportunity to lead in project finance, specialist insurance, risk analytics, captives, insurance-linked securities, and export finance for AI infrastructure in the UK and internationally.

Ensure place-based infrastructure creates wider economic benefits

AI infrastructure can deliver regional benefits beyond construction jobs. AI Growth Zones and related infrastructure projects should be linked to local skills provision, apprenticeships, university partnerships, SME supplier opportunities, heat-reuse schemes, community benefits, and adoption programmes for local firms and public services. Government should ask each zone to publish an adoption and local-benefits plan alongside its infrastructure plan. This will strengthen public consent, support regional growth, and ensure that infrastructure investment contributes to a sustainable AI ecosystem rather than only adding compute capacity.

RECOMMENDATION

Accelerate government's own adoption and use procurement, AI Champions, and adoption support to diffuse AI nationwide, with priority given to high-volume public services where AI can demonstrably improve accessibility and citizen experience

Accelerate government's own adoption through reusable components and clear standards

Government should use the AI Exemplars programme, the AI Playbook for the UK Government, and the roadmap for modern digital government to move from pilots to reusable AI capabilities. Departments and public bodies should adopt common standards, use open-source reusable components where appropriate, avoid duplication, and publish lessons from implementation. Government's own adoption should prioritise high-volume citizen-facing services where AI can improve outcomes, reduce backlogs, improve accessibility and the quality of citizen experience, and improve staff productivity, while maintaining transparency, human oversight, and public trust. Government should act as an anchor customer for responsible AI and use its experience to inform private-sector adoption guidance.

Adopt a layered national measurement framework

The AI Economics Institute should lead the development of a layered measurement framework for AI adoption, working with ONS, the CBI, and sector bodies. This should combine incidence, intensity, maturity, organisational enablers, barriers, trust, workforce impacts, and productivity outcomes. It should build on the evidence and methodology developed through the former DSIT's 2026 AI Adoption Research, alongside ONS instruments such as its Business Insights and Conditions Survey (BICS), while improving comparability across sectors, firm sizes, and regions. Over time, the framework should support public adoption benchmarks so that firms can understand how their maturity compares with peers and government can target support where it is most needed.

Use AI Champions' plans and new adoption support to drive sector-specific implementation

BIST and sector departments should incorporate delivery of AI Champions' adoption plans into reporting on delivery and implementation of the government's Industrial Strategy. This should include sector-specific use cases, worker engagement, skills, data, regulatory issues, procurement, and local diffusion. SME adoption should be a priority delivery channel rather than an indirect beneficiary of broader activity. Sector Champions should be asked to report on how their plans support SMEs and supply chains, and how they are aligning with regional institutions, universities, trade bodies, and local growth plans.

Use procurement as a lever for responsible AI adoption, especially for SMEs

Cabinet Office, working through the Government Commercial Agency, should update procurement practices so that AI can be used to make procurement more efficient and so that procurement encourages responsible AI practices among suppliers. Procurement should also seek proportionate evidence of user-centred design, workforce and frontline engagement, meaningful human oversight, and measurable improvements in service outcomes, alongside interoperability, assurance, explainability, data protection, and cybersecurity.

If developed and adopted, adherence to the shared principles of responsible adoption could in time be recognised as positive evidence, but requirements must remain proportionate and SME-friendly.

Government has previously used procurement to accelerate adoption of cloud, digital services, and cybersecurity standards across the public sector and its supply chains and should do the same for AI by scaling proven solutions, spreading best practice, supporting supplier diversity, and creating routes for innovative UK firms to win public contracts without excessive friction. Public procurement should reward responsible AI capability without creating a new compliance barrier.

Create local diffusion channels that link infrastructure, skills, and business support

AI adoption support should be delivered through trusted local channels, including combined authorities, universities, colleges, business representative organisations, local growth bodies, AI Growth Zones, and sector clusters, actively backing the strong convening role these local partners play, particularly for reaching SMEs. Local programmes should combine awareness, diagnostics, implementation support, skills, finance signposting, and peer learning. This is especially important for SMEs, which often need help specifying a use case, selecting suppliers, and changing processes rather than a generic explanation of AI. Local delivery should be linked to national standards and measurement so that support is consistent but responsive to regional strengths, while avoiding adding layers of complexity or duplication for firms operating in more than one area.

03

Governance and
trusted adoption

The significant economic potential inherent in AI adoption and trusted deployment across the economy will not be delivered by technology availability alone. It will depend on whether firms across every sector of the economy have the confidence, capability, and regulatory clarity to move from experimentation to scaled deployment. Good governance is one of the key best practices identified among adoption leaders and is therefore not a defensive add-on to the adoption agenda. It is, in fact, one of the conditions that turns business demand into investment, process redesign, and productivity gains. The strongest adopters are treating AI not as a standalone technology project, but as a business transformation agenda with governance designed to enable responsible use at scale.

Yet member discussions show that the immediate barriers to trusted adoption are often practical. These include data-readiness gaps, uncertainty over accountability, vendor dependency, service resilience, and limited organisational confidence.

Businesses broadly support the UK's sector-based approach to AI regulation. Member discussions showed backing for a model in which sector regulators interpret obligations in the context of their markets, and for the UK's position between the lighter-touch approach of the US and the more prescriptive, horizontal approach of the EU. That support rests on an important fact: the UK does not need to build a legal framework for AI from scratch. AI-enabled decisions are governed by existing law, including data protection, equality, consumer protection, product safety, employment, and health and safety law. The 2026 King's Speech reinforced this direction through the proposed Regulating for Growth Bill, including a stronger growth focus for regulators and cross-economy sandboxing powers to support controlled testing of innovative technologies.

The issue is not that the UK has a regulatory vacuum for AI; it is whether the system is clear, agile, and coordinated enough for businesses to navigate as AI use cases increasingly cut across regulatory boundaries.

"If you look at the UK, we have got quite a lot of rules and regulations which are actually capable of biting in this area"

— Taskforce member

"If you look at the two extremes — the US, very light on regulatory constraints, and the EU, much more horizontal and enforced — the fact that the UK strategy sits somewhere in the middle is absolutely the right one"

— Taskforce member

Trusted adoption is a business confidence issue

Adoption at scale depends on two mutually reinforcing forms of confidence: the confidence of businesses deciding whether to invest, and the confidence of citizens, customers, and employees who interact with AI-enabled products and services. Public confidence is currently fragile. ONS polling in 2025 found that only 41% of British adults believed AI would benefit them. The Ada Lovelace Institute and The Alan Turing Institute found that 72% of the public would be more comfortable with AI if there were laws and regulations around its use.³⁴ Confidence is strongest where AI supports rather than replaces human judgement in high-impact decisions. The implication is clear: well-designed governance is not a brake on adoption, but one of its enablers.

The same point comes through from business. Firms want regulation that is clear, proportionate, and predictable enough to support investment. Experimentation with AI is now widespread, but the harder work in the next phase of adoption is moving it into core business processes, measuring value, building workforce confidence, setting risk appetite, and deciding who is accountable when AI affects business outcomes.

“There is too much information right now. It is hard to distinguish what is real from what is not, and to keep up to date”

— Taskforce member

This is particularly important for smaller firms and mid-market businesses, which lack the teams or time to interpret complex regulatory guidance, and, as discussed in Chapter 2, have lower rates of AI adoption overall. Larger organisations in regulated sectors are increasingly building formal governance infrastructure: approval routes, model inventories, cross-functional committees, audit processes, and board reporting. Smaller firms are less likely to have the capacity to build this from scratch. For these firms, the governance challenge is not a lack of willingness to act responsibly, but a lack of accessible capacity, templates, and confidence to translate high-level principles into day-to-day business decisions. That makes practical support, including toolkits, worked examples, self-assessment, and shared frameworks important.

³⁴ Office for National Statistics (August 2025), [Public opinions and social trends, Great Britain](#); The Ada Lovelace Institute and The Alan Turing Institute (March 2025), [How Do People Feel About AI? Wave two of a nationally representative survey of UK attitudes to AI](#).

Managing risk is part of confident AI adoption

As with any new technology, AI introduces risks in how it is developed, deployed, and used. Many of these are familiar operational, consumer, and reputational risks, including misleading outputs, unfair or discriminatory decisions, misuse of data, cyber vulnerabilities, and loss of customer trust. However, they can be magnified by the speed, scale, and complexity of AI systems.

For most use cases, the priority is therefore not to create a separate risk framework for AI, but to apply existing protections and business controls effectively as AI becomes embedded in core processes.

The causes of harm vary, and so must the response. Some risks arise from deliberate misuse, where existing criminal, fraud, cyber, and data-protection rules already apply. Others arise despite good intentions, where governance has not kept pace with a fast-moving use case, or where a model's limitations in training data, testing, or transparency produce unreliable or biased outputs. It is this second group that proportionate governance and assurance are designed to prevent before harm occurs.

Businesses are alert to these risks, and in many cases the caution is commercial as much as ethical. Firms know that trust, once lost, is hard to rebuild, whether with customers, employees, regulators, or investors. Members described being deliberate about where and how AI is deployed so that it strengthens rather than undermines trust, including holding back from uses that, while technically possible, they judge inappropriate.

“You could have a chatbot on your app answering questions such as ‘Where’s my order?’, and nobody expects a human at the other end; that is fine. But if you have a video of a customer service agent with a heavy accent, and you use AI to reduce that accent, that starts to edge into really dangerous, ethical territory. It is a really fine, nuanced line”

— Taskforce member

That awareness is translating into concrete practice. Firms described establishing cross-functional AI governance committees, assessing and monitoring the risk of individual use cases, maintaining approved-supplier lists, vetting vendors for data handling and security, strengthening data governance, requiring training before tool access, and keeping humans accountable for decisions in higher-risk settings.

“People wanted a clear workplace policy as they weren’t sure what they were allowed to use AI for — a policy that says, ‘yes, you can use it for these tasks but not for these’”

— Taskforce member

Where governance is still catching up with adoption, the constraint is usually confidence and capability rather than any lack of willingness to act responsibly.

The practical priority is to help governance keep pace, through proportionate, business-led approaches that make responsibilities easier to understand, apply, and evidence. Adverse events underline the importance of proportionate governance and assurance as enablers of confident AI adoption.

Governance is becoming an enabler of scaled adoption

Taskforce interviews and roundtables point to a clear pattern. The most advanced adopters are building governance approaches that are risk-based, business-owned, and designed to enable adoption rather than block it.

They are not applying the same controls to every use case. Instead, they are calibrating governance to the inherent risks. Heavier controls are needed where AI affects customers, regulated outcomes, safety-critical processes, or sensitive data. Lighter controls can work for lower-risk productivity tools, provided they are backed by monitoring and clear rules of use.

Some observations around governance best practice:

Boards set the ambition and risk appetite

AI governance needs clear senior ownership. Boards do not need to approve every use case, but they should understand where AI is being used, what strategic outcomes it is intended to deliver, what risks and dependencies it creates, what the control and risk mitigation architecture is, and how the organisation will respond if systems fail. Good board oversight should cover AI risk appetite, senior accountability, high-impact use cases, escalation routes, workforce implications, service resilience, assurance arrangements, and how AI generally nests within the overall risk and compliance architecture. It should also test whether investment is linked to business transformation rather than simply technology deployment.

“AI has to be a board- and CEO-driven change initiative. Its strategic importance is bigger than the arrival of the internet in the late 1990s, because it fundamentally shapes how the business operates, its economics, culture, and operating model”

— Taskforce member

Business units own the use case

Once the ambition and risk appetite are set, ownership needs to sit close to the process being redesigned. The strongest approaches avoid leaving AI decisions solely to technical teams. Rather, business units own the use case, the intended outcome and day-to-day risk management — within the guardrails set by the second line — because they understand how the process works, where value will be created, and where operational risks are likely to arise.

“It is being centrally orchestrated through a cross-functional team with legal, compliance, and risk. But the idea is going out into the business to give people empowered decision-making ability, so they understand the risks and where they can lean in and take a bit more risk and where they can’t”

— Taskforce member

Controls and compliance follow the risk

Risks are inherent in any business process, and the key risk is that AI tools may magnify these risks, as well as create new ones. A “three lines of defence” approach is a useful organising principle that is applied to all key risks that are relevant to an organisation (for instance, cyber or operational). The first line is the business owner deploying the AI system and managing the use case day to day. The second line is the risk, legal, compliance, data protection, cyber, and technology functions that set policy, risk taxonomy, and mitigation strategies and generally challenge or validate decisions taken by the first line. The third line is independent assurance or internal audit, testing whether the framework is operating as intended and typically reporting to the board. For smaller firms, the same logic should be translated into a lighter-touch discipline. This could include clear ownership, challenge proportionate to risk, and periodic review of whether systems are performing safely and effectively.

“It is critical that we identify what the risk is, quantify the size of the risk, and identify what the mitigation strategies are if something was to go awry. And then the most important part is how do we monitor that risk?”

— Taskforce member

Resilience needs to be designed into critical AI deployments

As AI becomes embedded in business-critical processes, service continuity and resilience are moving from technical concerns to governance issues. Firms need to understand not only whether an AI system performs well in normal conditions, but how critical services would continue if access is restricted, a model is withdrawn, a supplier changes its terms, a cyber incident disrupts a dependency, or token and compute costs alter the economics of a use case or broader enterprise-wide usage.

In practice, this is reflected in a business continuity plan that is consistent with wider plans covering enterprise-wide service discontinuities (for example, internet outage).

“If something were to happen and country X decides they’re going to shut off the servers, medium and larger businesses do think about service continuity”

— Taskforce member

Humans stay accountable

Members consistently emphasised that human oversight cannot be reduced to a box-ticking “human in the loop” requirement. The real question is who has the authority, training, and confidence to challenge or override an AI output when it affects a customer, worker, regulated outcome, or critical operation. For higher-impact use cases, accountability stays with people who understand the operating context and know when to escalate, intervene, or stop a process.

“We have an absolute rule within the governance: it’s human at the helm, not human in the loop. These tools are very useful, but it is a human decision”

— Taskforce member

Standards can help translate these disciplines into repeatable responsible-adoption practice

Standards can help translate governance principles into repeatable responsible-adoption practice by giving firms a common language and structured approach. Standards can also be fully integrated into the three-lines-of-defence construct by setting parameters and guardrails for the way in which tools are used, and as a more intuitive way in which the business (first line of defence) actually employs the tool in service and production. ISO/IEC 42001, the international standard for AI management systems, is increasingly cited as a useful reference point, and the UK Accreditation Service (UKAS) granted the British Standards Institution (BSI) the first accreditation for certification to the standard in January 2026.³⁵ However, awareness and uptake remain uneven, particularly beyond larger regulated firms. Standards are most valuable in driving good practice and outcomes when they are used to embed monitoring, evaluation, and accountability into how AI systems operate in practice. They are less effective if treated as a tick-box certification exercise.

³⁵ For further detail, refer here: UKAS (2026), [UKAS grants first accreditation for artificial intelligence management systems](#).

Navigating regulators is becoming a barrier to scaled adoption

As AI moves from experimentation into core business processes, the regulatory challenge is becoming increasingly practical. Businesses need to understand how existing obligations apply to fast-evolving AI use cases, where to seek advice, and how questions spanning more than one regulatory remit will be addressed. Depending on the application, this can involve data protection, copyright, consumer protection, health and safety, employee rights, and sector-specific regulation. Data protection illustrates the point. For many firms, the UK General Data Protection Regulation, Data Protection Act 2018, and the Data (Use and Access) Act 2025 are already among the most immediate legal frameworks shaping AI deployment, particularly where AI systems use personal data, support customer-facing decisions, operate in the workplace, or interact with data held across multiple systems.

The challenge is not whether data protection applies, but how firms can apply requirements on lawful processing, fairness, transparency, impact assessment, and automated decision-making confidently to fast-evolving AI use cases.

The Information Commissioner's Office (ICO) offers substantial guidance for the deployments firms are running today, including on AI and data protection, its risk toolkit, and guidance on explaining decisions made with AI. For most firms, the immediate need is support to put this into practice consistently rather than new rules. In March 2026, the ICO also published updated draft guidance on automated decision-making and profiling, following the changes made by the 2025 Act in this area. Its forthcoming statutory code of practice on AI and automated decision-making, and its developing work on agentic AI, are expected to provide further clarity as systems become more autonomous.

AI use cases increasingly cut across sectors and business functions, making regulatory coordination more important. The Digital Regulation Cooperation Forum (DRCF) brings together the Competition and Markets Authority, ICO, Ofcom, and the Financial Conduct Authority to support cooperation on digital regulation. Its AI and Digital Hub pilot demonstrated the value of providing innovators with informal cross-regulatory advice where products or services engage more than one remit.³⁶ The DRCF's newer Thematic Innovation Hub builds on that model through focused engagement and support on priority topics, beginning with agentic AI, to identify practical challenges and areas where businesses require greater regulatory clarity.

However, this coordination is not always visible to businesses at the point of deployment. Firms operating across sectors can struggle to know which regulator to approach first, how different expectations interact, and whether they will receive a joined-up response. This matters particularly for platform and multi-sector businesses, where the relevant regulation for certain AI use cases can cut across the remits of multiple regulators simultaneously, from environmental impact to sector-specific safety rules. The problem therefore spans the whole regulatory landscape, not only the four digital regulators.

Evidence suggests that the principal gap is not the absence of another approval mechanism, but the lack of a clear and informed point of entry into the existing system.

AI is a general-purpose technology, and many use cases do not fit neatly within a single regulatory remit. Businesses therefore value an approach in which the first regulator they contact can provide meaningful initial orientation, even where responsibility is shared.

“There are structures that exist, like the Regulatory Innovation Office or DRCF for the main digital regulators, bringing them together so you get one answer and it’s the same answer”

— Taskforce member

This is often described as a “no wrong door” model of regulatory navigation. Under such a model, the first regulator contacted would be able to explain which aspects of a query fall within its remit, identify other relevant regulators, and indicate how existing rules and guidance are likely to apply. For established use cases, this initial response may be sufficient to give businesses greater confidence about how to proceed.

³⁶ For further detail, refer here: DRCF (2025), [AI and Digital Hub pilot report](#).

Novel or cross-sector applications may require a more coordinated process. In those cases, regulators may need to identify the relevant bodies and establish how a joined-up position will be reached. Even where a definitive regulatory view is still developing, greater clarity about where to begin, which organisations are involved, and how the question will be considered can reduce uncertainty for businesses. For particularly complex cases, a named case lead could provide continuity and coordinate engagement, without becoming an additional approval authority or assuming the statutory responsibilities of individual regulators.

Existing institutions provide a potential foundation for this model. The DRCF already performs a coordination role among the principal digital regulators. The Regulatory Innovation Office's emerging cross-regulator front-door function³⁷ may provide a route for applying similar principles across a wider range of sector regulators, although an economy-wide role in relation to AI would represent an extension of its current remit. In both cases, the function would be one of coordination and navigation rather than regulatory decision-making.

The Regulating for Growth Bill may also provide a mechanism for embedding more responsive engagement across the regulatory system. Its strengthened Growth Duty, together with accompanying ministerial strategic steers, could encourage regulators to provide substantive assistance when businesses first raise a question, without requiring the creation of a separate statutory duty or institution. In this context, supporting growth would not mean lowering regulatory standards. It would mean helping firms understand and meet those standards so that confidence, responsible adoption, and growth reinforce one another.

A longer-term opportunity lies in the use of AI to support regulatory navigation itself.

For example, a publicly owned AI tool could help draw together legislation, regulatory guidance, previous regulatory positions, and relevant case examples. This could assist businesses and regulators in identifying the regimes engaged by a particular use case and understanding how established expectations may apply. Regulatory judgement and accountability would remain with the relevant regulators, but such tools could strengthen public-sector capability and make engagement with the regulatory system more accessible.

³⁷ GOV.UK (2026), [Regulatory Innovation Office: Front Door pilot](#) published 16 January 2026, last updated 20 April 2026.

Responsibility is clearest when it follows control across the AI value chain

A recurring source of uncertainty is where responsibility sits between developers and deployers. The obligations appropriate to an upstream model or tool provider are different from those appropriate to a business deploying a third-party AI tool. Getting this wrong creates risks in both directions. If too much responsibility is pushed downstream, deployers can be overburdened with obligations they cannot meaningfully discharge, including technical risks that only developers can assess. If deployers are treated as passive users, real risks introduced by the deployment context, data, integration, fine-tuning, or downstream workflow can be missed.

Responsibility is clearest when it follows meaningful control. This logic is already familiar in UK data protection law, where controllers and processors have different obligations because they exercise different forms of control. **A similar distinction can help explain how responsibility may be allocated across the AI value chain.** Developers are generally best placed to assess and mitigate risks relating to model design, safety evaluation, documentation, and upstream assurance. Deployers are generally better placed to manage the use case, data inputs, integration, human oversight, monitoring, and the operating environment they control.

This matters for adoption because it lets assurance be reused across the value chain rather than repeated at each stage.

For that to work, a deployer needs to know what has already been tested upstream and what still needs to be assured in its own context, which depends on practical transparency, standardised model documentation, technical testing standards, and, over time, independent certification.

The Trusted Third-Party AI Assurance Roadmap and the AI Assurance Stakeholder Consortium³⁸ are building the market and the professional infrastructure for this, and the AI Security Institute evaluates the most capable models upstream. The piece still missing is making the results of that upstream testing legible to the deployers who are expected to rely on them. Upstream assurance can give deployers confidence in the model or tool itself, but it does not remove the need to assure how that tool is applied in a specific business context. Deployers remain responsible for assuring what they control: the use case, data inputs, integration, oversight, and live monitoring. Where a deployer fine-tunes or materially adapts a tool, the boundary between developer and deployer responsibility may become less clear. The UK Jurisdiction Taskforce's legal statement on the liability of AI harms under the private law of England and Wales is a welcome step forward in providing further clarity on legal treatment.³⁹

The challenge sharpens as AI becomes more agentic. Systems that act more autonomously, initiate actions across platforms, or change behaviour after deployment cannot be governed through one-off sign-off alone. This points towards continuous assurance, including ongoing monitoring, incident reporting, model evaluation, audit trails, and iterative updates that track how systems behave in real-world use. For businesses, the key question is not only who was responsible at deployment, but who remains responsible as the system operates, interacts with other tools, and creates new risks in practice.

Taken together, the evidence points towards a proportionate approach in which risks are addressed at the point in the value chain where they can most meaningfully be controlled. For most use cases, existing legal obligations remain the baseline.

The practical challenge is making those responsibilities easier to interpret, apply, and evidence as AI systems become more complex, embedded, and autonomous.

³⁸ In June 2026 the Government and BCS launched an AI Assurance Stakeholder Consortium to begin professionalising the market, with industry embedded directly through an advisory group.

³⁹ For further detail refer here: LawtechUK (2026), [Liability for AI Harms under the private law of England and Wales — LawtechUK](#).

Shared governance norms can reduce duplicated efforts across firms

Across the economy, firms of all sizes are solving the same governance challenges in isolation, and shared norms can change that. Large businesses are building their own internal frameworks, but many still want a clearer way to compare practice, share lessons, and avoid duplicating effort. Smaller and mid-market firms face the same questions with less capacity. Taskforce discussions highlighted the potential value of a credible, industry-led mechanism to help firms understand what responsible AI adoption looks like in practice, while remaining voluntary, proportionate, and focused on enabling adoption.

“What would be helpful is if that body helped businesses implement existing frameworks, because at the moment we are all kind of doing it ourselves and there is not a way to talk to each other and learn best practices”

— Taskforce member

A voluntary responsible adoption code and an accompanying public commitment, intended to provide firms with a shared reference point for responsible practice, helping them interpret and apply existing expectations more consistently without creating new legal obligations, represent one possible response. This responds to a practical gap in the current landscape. International principles and standards, such as the OECD AI Principles, the G7 Code of Conduct, and ISO/IEC 42001, already provide important foundations for responsible AI governance, but many firms, particularly smaller businesses, can find them difficult to translate into day-to-day adoption decisions. The opportunity is therefore not to create a parallel UK standard, but to make established expectations easier for businesses to understand and use.

A common set of principles can help reduce duplicated governance challenges across business, support greater consistency across firms, and give customers, supply chains, and other stakeholders greater confidence in how AI is being adopted, building trust, and accelerating adoption. Its value depends on remaining practical and proportionate, and on supporting rather than substituting for existing regulation, standards, and assurance mechanisms.

AI assurance is an integral element of AI adoption and can become a UK growth driver and export strength

AI assurance is an industry in its own right, and one the UK is well placed to lead. The country already has deep capability in the adjacent professions, including audit, standards, conformity assessment, professional services, finance, and risk management, and AI assurance draws on the same disciplines. Strong assurance will be critical to support trusted and sustained adoption of AI, giving businesses, regulators, and the public confidence that AI systems are reliable, safe, and fit for purpose. The UK has a credible basis to build a globally recognised assurance sector. It can span both process assurance, evidenced against management standards, and the harder performance assurance of how AI systems behave in context, rather than treating assurance purely as a domestic compliance cost.

The commercial opportunity is real and largely still to come. The UK AI assurance market was estimated at around £1 billion in gross value added in 2024, and the government's Trusted Third-Party AI Assurance Roadmap estimates it could reach £18.8 billion by 2035 if barriers to adoption are addressed.⁴⁰ There is fast-growing demand for assurance of AI applications in context, where deployers most need confidence. The Roadmap, the new AI Assurance Stakeholder Consortium, and the AI Security Institute's evaluation capability together provide the foundations to build on. Realising that opportunity depends on shaping the market around the firms that will actually rely on it, so that assurance becomes a competitive asset rather than an overhead.

International interoperability is the second prize, particularly for firms operating across the UK, EU, and US. For firms, AI adoption is increasingly shaped by global supply chains, data and security requirements, model access, service resilience, and the risk of regulatory divergence. Firms selling into the EU already have to meet the AI Act's requirements, which is a compliance reality for exporters rather than a case for replicating that model here. Standards, assurance, model evaluation, and incident reporting are practical routes to interoperability even where full regulatory alignment is difficult, and they are areas where UK leadership can shape emerging global norms and turn domestic assurance capability into an exportable strength.

As AI regulation is increasingly shaped by geopolitics, economic security, and strategic competition, the UK should pursue international collaboration wherever possible on these same foundations. The aim should be business-enabling global guardrails that support trust and reduce unnecessary duplication, while keeping the UK's domestic framework clear, dependable, and usable for firms.

⁴⁰ DSIT (2025), [Trusted third-party AI assurance roadmap](#); DSIT (2025), [Roadmap to Trusted Third-Party AI Assurance Written Ministerial Statement HCWS903](#).

Recommendations

For business

RECOMMENDATION

Implement robust governance and responsible adoption practices

Effective governance is a critical enabler of AI adoption and trusted use. As AI becomes embedded into core business processes, governance must evolve beyond general technology risk controls. Firms need practical mechanisms that allow teams to innovate while giving boards, employees, customers, and regulators confidence that AI systems are being used responsibly.

Governance should not be seen as a brake on adoption — done well, it accelerates adoption by making decision-making clearer, reducing uncertainty, and preventing avoidable failures.

The central recommendation is that businesses should evolve risk and compliance frameworks so that they fully reflect the inherent risks in AI-enabled service and production processes, and so that AI-related risks are connected to traditional risk categories such as operational resilience, conduct, safety, privacy, cyber, financial crime, competition, intellectual property, and reputational risk.

This requires board oversight, clear accountabilities, approved tools, data and privacy controls, risk-based use-case assessment, human oversight for high-impact decisions, ongoing monitoring of AI outputs and suppliers, and practical mechanisms such as AI policies, use-case registers, escalation routes, and role-specific training. Other stakeholders outside the organisation also have a role to play in encouraging firms to adopt responsibly, such as investors (for instance pension funds) or civil society actors. Governance should be proportionate: low-risk productivity tools should not face the same controls as safety-critical, customer-facing, or regulated decision systems. However, even low-risk tools need clear rules on data handling, tool approval, and acceptable use.

More broadly, the evolution of risk and compliance frameworks should include the following actions.

Establish clear leadership, ownership, and accountability for AI

Boards should set the organisation's ambition, risk appetite, and overall governance approach for AI. Executive committees should translate this into funding choices, policy, controls, and performance measures. Responsibility for individual AI use cases should sit with the business functions deploying them, because they own the process, the customer outcome, and the operational risk. Technology, data, legal, compliance, and risk functions should provide enabling expertise and independent challenge. Firms should define who is accountable for the model, the data, the process, supplier oversight, monitoring, and incident response. A clear accountability model reduces ambiguity and helps ensure that AI remains aligned to business outcomes rather than becoming a fragmented technology experiment.

Implement a proportionate, risk-based approach to governance and oversight

Businesses should classify AI use cases by risk, considering factors such as customer impact, safety, legal, or regulatory implications, sensitivity of data, degree of autonomy, explainability, third-party dependency, scale of deployment, and reversibility of harm. Low-risk tools should benefit from fast-track approval and clear usage rules; medium-risk tools should require structured assessment and monitoring; high-impact tools should require senior approval, documentation, testing, human oversight, and post-deployment review. This tiered model helps organisations avoid both extremes: excessive bureaucracy for simple productivity tools and insufficient scrutiny for AI systems that affect people, markets, or safety. The process should be agile enough to support innovation while robust enough to demonstrate responsible adoption.

Recognised standards such as ISO/IEC 42001 can serve as practical reference points for building repeatable governance, provided they are used to embed real accountability rather than as a tick-box certification exercise.

Embed resilience, assurance, and continuous monitoring alongside AI deployment

AI governance should focus on the full lifecycle of systems, not only one-off approval before deployment. Businesses should monitor performance, accuracy, drift, bias, security vulnerabilities, incidents, user feedback, and supplier changes. They should define thresholds for escalation, rollback and human intervention. Where AI is embedded in business-critical operations, firms should assess continuity risk, vendor concentration, data portability, model dependency, and cyber resilience. Assurance should combine internal controls, technical testing, audit, supplier assurance and, for high-impact systems, independent third-party review where appropriate. Continuous monitoring is particularly important for agentic systems and models that change in use, because risk can evolve after deployment.

Strengthen procurement and third-party AI oversight

Many businesses will adopt AI through third-party tools rather than building models internally. Procurement processes therefore need to ask new questions about model provenance, training data, data residency, security, intellectual property, audit rights, service continuity, subcontractors, model updates, interoperability, exit rights, and liability. Firms should require suppliers to provide appropriate documentation and assurance evidence, but should also recognise when deployment-specific testing is needed because the same tool may behave differently in different contexts. Contracting should cover how incidents will be reported, how model changes will be communicated and how customer or employee data will be protected. This is especially important for SMEs, which may rely heavily on suppliers and need simple templates for due diligence.

Embed a responsible AI ethos throughout organisational culture

Responsible adoption should extend beyond formal governance structures. Employees need practical guidance on appropriate AI use and clear routes to raise concerns. Managers should encourage teams to identify both opportunities and risks, and should make clear that human judgement can and should override AI-generated outputs where necessary. Where AI changes how work is done, firms should draw on employee insight and existing engagement channels to support effective implementation and identify practical risks early. Firms should also design around the people affected, considering the implications for customers, suppliers, and workers, including transparency, accessibility, fairness, and contestability. A responsible culture helps prevent shadow AI, over-reliance, and misuse. It also complements the proposed voluntary Responsible AI Adoption Code by giving employees and managers a common set of principles that can be translated into day-to-day decisions.

RECOMMENDATION FOR THE CBI

Convene business, workers, and civil society to explore a mechanism to articulate and scale trusted AI adoption

As AI adoption spreads, firms are asking a practical question: what does responsible adoption actually look like day to day? Larger organisations are building their own governance frameworks, but many firms, particularly SMEs, lack the capacity to design them from scratch, and the result is duplicated effort and uneven practice. A common set of trusted adoption principles would give firms a shared reference point, reduce duplication, and strengthen confidence among customers, employees, suppliers, and investors. It would also help the UK lead on trusted adoption, ensuring its pro-innovation regulatory model is matched by credible, industry-led practice rather than a patchwork of competing standards.

Explore a voluntary code and pledge for responsible AI adoption

We recommend exploring, with an independent and credible convener, the development of a voluntary Responsible AI Adoption Code and public pledge, bringing together business, workers, and civil society alongside government. A proven model already exists in cyber, where the Cyber Governance Code of Practice and its accompanying Cyber Resilience Pledge give firms clear, board-level actions to adopt and commit to publicly. A first phase should be deliberately light, covering shared principles, plain-language guidance, self-assessment, and signposting to existing standards such as ISO/IEC 42001, and focused on the issues firms most want help with: accountability, data governance, transparency, human oversight, and incident response. Greater Manchester's Good Employment Charter offers a further precedent, pairing an initial written commitment to improve with assessed membership for those meeting defined criteria, and is open to employers of any size or sector.

The aim is a common language for trusted adoption understood by boards, managers, employees, suppliers, and customers. Later phases could add sector toolkits, worked examples, and assurance templates. If the code gains traction, it could evolve or inform a more formal industry-led assurance mechanism, with recognition in procurement and by investors and insurers creating demand-side incentives for take-up.

For government

RECOMMENDATION

Evolve the regulatory model through assurance, regulatory navigation, the AI Growth Lab, and growth-enabling regulatory simplification for firms

The UK's principles-based, sector-led approach remains foundational, particularly when combined with the Regulatory Innovation Office, the Digital Regulation Cooperation Forum, and the proposed AI Growth Lab. The main risk for firms is not only over-regulation; it is uncertainty, fragmentation, and slow navigation across multiple regulators. The regulatory model should therefore preserve flexibility while giving businesses clearer routes to advice, testing, assurance, and scale. We envision six significant enhancements.

Develop the global standard on AI assurance in partnership with industry

BIST, working with the Cabinet Office, should build on the Trusted Third-Party AI Assurance Roadmap, the AI Assurance Innovation Fund, the National Physical Laboratory's Centre for AI Measurement, and the AI Security Institute's evaluation capability to create a globally recognised assurance ecosystem. Businesses draw an important distinction between process assurance, evidenced against management standards such as ISO/IEC 42001, and performance assurance of how a model behaves in context. The UK's opportunity lies in both, developed as complementary layers rather than a single standard, recognising that performance assurance is significantly harder and that assurance should remain voluntary and market-led. Industry should shape the approach, because buyers and deployers need assurance evidence that is usable in real procurement and operating contexts. Over time, assurance should become a UK export strength analogous to professional services, audit, standards, and accreditation.

Create a no-wrong-door model for AI regulatory navigation.

The Cabinet Office, working with BIST and DCMS, should establish a protocol under which businesses raising cross-cutting AI queries with any relevant regulator receive an informed first response and a clear route through the system. The DRCF should coordinate queries involving digital regulators, while the Regulatory Innovation Office (RIO) should extend coordination to wider sector regulators, building on the front door it has piloted. In complex or novel cases, a named case lead could coordinate input without becoming a new approval authority. This should be supported by sufficient AI expertise and coordination capacity within regulators and framed as a means of delivering their growth duties while maintaining safety, fairness, and trust.

Publish and maintain a practical regulatory navigation map

The DRCF, working with RIO and wider regulators, should publish and maintain a map showing which regulators are likely to be relevant to common AI use cases, including platform models, foundation-model-enabled services, agentic systems, and multi-sector deployments. The map should include anonymised case studies, links to sandboxes, typical evidence requirements, contact points, and expected timelines. It should be treated as a living product and updated as regulators publish new guidance or as AI systems create new cross-regulatory issues.

Build on the ICO's automated decision-making work and extend it to agentic AI

Following the changes made by the Data (Use and Access) Act 2025 to automated decision-making, the ICO published updated draft guidance on automated decision-making and profiling in March 2026. The ICO should finalise this guidance and its planned statutory code of practice on AI and automated decision-making and extend that clarity to more agentic systems that can act without a clear human decision point and adapt in use. This should include practical examples for common business deployments, including customer service agents, HR tools, finance operations, procurement agents, and systems that interact with external platforms, to reduce uncertainty and support responsible adoption.

Clarify responsibility across the AI value chain

Government, through BIST and the Cabinet Office, should work with the AI Assurance Stakeholder Consortium to establish a common approach to assurance that reflects where control sits between developers, providers, and deployers. This should include standardised model documentation, common testing standards, and a usable route for upstream model evaluation results, including from the AI Security Institute, to reach deployers. This would reduce duplication, make clear what has already been assured upstream, and help deployers focus on the use case, data, integration, oversight, and monitoring they control.

Pursue international interoperability through standards, assurance, and shared practice, while preserving UK flexibility

The UK should preserve its principles-based, sector-led model and pursue interoperability where it is practical and reduces cost for firms. This should include continued development of international standards such as ISO/IEC 42001, common approaches to model evaluation and incident reporting, and stronger sharing of good practice and lessons learned between governments and regulators, building on the work of the OECD, the G7 and the US National Institute of Standards and Technology. Interoperability at this level will reduce compliance costs for UK firms selling internationally without importing regulatory models designed for different systems.

04

Skills and
workforce readiness

The productivity and transformational benefits of AI are realised only when the workforce can effectively deploy, oversee, and reshape processes around the technology. Businesses consistently identify workforce capability, rather than the technology itself, as the factor that separates organisations generating value from AI from those whose pilots do not scale.

“Tech adoption, in itself, does not drive productivity. The skills element is vital”

— Taskforce member

Skills and capabilities are therefore central to converting AI adoption into economic value.

As AI moves from experimentation into everyday business processes, the skills challenge extends across the workforce, from leaders making strategic decisions to managers redesigning teams and processes, employees using AI tools safely and effectively, and workers whose roles may change as adoption scales. This chapter therefore treats skills not as a supporting theme, but as a determinant of whether the UK captures the economic opportunity of AI.

AI skills are becoming a source of competitive advantage for businesses

As AI adoption moves from experimentation into business process redesign, the skills challenge for firms is becoming more specific. Businesses are investing in the capabilities needed to apply AI safely and effectively to real workflows: redesigning processes, strengthening governance, improving data foundations, managing change, and equipping employees to use AI tools with judgement. The limiting factor is increasingly whether firms can embed AI into business processes, operating models, and decision-making, not simply whether employees understand the technology.

This points to four distinct but overlapping skill needs, drawing on published skills taxonomies⁴¹ and evidence from business engagement.

41 These frameworks include: Skills England (2026), [AI Foundation Skills for Work](#); The Alan Turing Institute/BridgeAI (2024), [AI Skills for Business Framework](#); UNESCO (2024), [AI Competency Framework for Students](#); SFIA, [AI Skills Framework \(v9\)](#) FutureDotNow (2026), [The Essential Digital Skills Framework](#); Pearson, [Mind the Learning Gap](#) (Pearson D.E.E.P Learning Framework).

Exhibit 5: Skills categorisation

Category 1	Baseline AI literacy and practical day-to-day usage Understanding the basics of AI and what it can do versus not; prompting and usage of tools; integrating AI into everyday tasks and knowing when to use it
Category 2	Oversight and judgement of AI tools Maintaining responsible usage of AI; critically assessing AI outputs; managing risks and governance associated with AI usage
Category 3	Leadership and strategic transformation Developing and implementing AI strategy; increasing AI adoption across the organisation; change management and re-engineering of business processes to embed AI into workflows; strategic workforce planning and retraining/reskilling
Category 4	Specialist AI capabilities Building AI engineering and systems design, deployment and monitoring; technical implementation; machine learning expertise

Baseline AI literacy, as well as oversight and judgement of AI tools, are foundational skill sets. These first two skills categories apply to almost every worker, with universal policy and safe-use training often required in organisations before employees are given access to AI tools. Beyond this baseline, businesses see the greatest competitive advantage in more targeted capability-building, focused on how AI changes specific roles, levels of responsibility, sectors, and business processes. Many businesses report investment in training and reskilling their existing workforce,⁴² but the strongest impact comes when that investment moves beyond generic courses and is embedded in the realities of how people do their jobs.

“We have different training based on whether you’re in an executive role, a manager role, or an individual contributor role, depending on how AI is changing your role and the skills required”

— Taskforce member

The pace at which AI tools and use cases are evolving is also changing how businesses deliver training. Formal courses remain important for baseline literacy, responsible use, and governance, but they are no longer sufficient on their own. Businesses are increasingly supplementing them with role-based learning, peer sharing, internal showcases, hackathons, promptathons, and hands-on use-case development, so that employees learn how to apply AI to the actual processes, data, and decisions they work with day to day.

“Learning needs to be embedded into the flow of work. The technology is changing so quickly that people cannot continuously come out of work to do courses”

— Taskforce member

⁴² CBI (2026), [AI Skills Report: Building an AI Ready Workforce](#).

Effective AI adoption depends on workforce confidence and practical support

There is broad agreement that skills are a limiting factor to scaled implementation, but the most important skills are not limited to prompt engineering or deep technical model expertise. As AI becomes embedded into ordinary business processes, organisations need a broader capability mix: practical AI literacy across the workforce, data and process engineering skills to redesign how work is done, transformation and delivery capability to take redesigned processes into live operation, specialist AI and assurance skills for higher-risk deployments, and leadership skills to sponsor change responsibly and effectively.

As AI systems take on more routine technical execution, complementary human skills are becoming more important.

Judgement, critical thinking, communication, and the ability to challenge AI outputs are increasingly the capabilities that determine whether AI is used safely and productively.⁴³ Businesses are still defining the AI-specific knowledge each role will require, but they are increasingly clear that these human capabilities will sit at the centre of effective adoption.

“It’s not about tech. The biggest part of the transformation is a human transformation. People bring value through the human aspect: the emotional intelligence, critical judgement, and critical thinking they bring”

— Taskforce member

Employee confidence to adopt AI increases when employees understand how tools will be introduced and how workflows are to be redesigned. Businesses report that individual openness and psychological safety are often stronger predictors of adoption than role or background.

⁴³ OECD (2026), “AI and skills: What we know so far”, OECD Publishing, Paris, <https://doi.org/10.1787/f843b352-en>.

Moreover, frontline employees with deep process knowledge are frequently best placed to identify useful AI applications, test them in practice, and share what works across the organisation. This is especially important where AI is being applied to operational processes, where workers closest to the task understand where data is unreliable, where processes break down, and where automation could remove friction without introducing new risks.

“There is appropriate consultation with the people who have the knowledge to decide how it’s deployed”

— Taskforce member

The need for practical, context-specific support is particularly acute for smaller businesses.

Smaller firms consistently cite skills, cost, and awareness as principal barriers to AI adoption, alongside constraints around finance, regulatory complexity, and knowing where to turn for advice. Support is therefore most effective when delivered through trusted intermediaries that can translate AI adoption into practical, sector-specific guidance.⁴⁴ Targeted programmes such as Made Smarter Adoption, which provides manufacturing SMEs with advice, training, and technology support, show how this can work in practice.⁴⁵

“Are we creating a system in which only the large UK-based companies and multinationals have access to this? What about smaller and mid-sized companies?”

— Taskforce member

For the employed workforce as a whole, the case for broader public intervention is more limited, even while some firms note challenges around accessing affordable training offers that meet their needs. Needs are heterogeneous and larger businesses are already investing at pace. The highest-value role for government is therefore to extend that momentum to smaller firms, rather than to substitute for private investment.

“One of our big concerns is the cost of using AI tools as the inference prices have suddenly rocketed, combined with the ongoing training costs”

— CBI AI and Data Protection working group member

The sharpest risks of market failure sit beyond the reach of firm-led training altogether.

Businesses will invest where skills are needed for their own operations, particularly for employees already in post. That leaves a weaker offer for people before they enter work, for those moving between roles, employers, or sectors, and for later-career workers whose expertise could be retained but may not fit traditional training routes.

44 As highlighted in the prior CBI AI Skills Report. Source: London School of Economics (2024), [LSE-CBI survey on AI adoption in UK firms](#); CBI submission to the Business and Trade Committee inquiry into the Small Business Strategy (2025).

45 For further detail refer to: Made Smarter Adoption [website](#).

AI is reshaping early-career work and the routes into future capability

AI is beginning to reshape the entry-level work through which future capability is built.

As routine analytical, administrative, and drafting tasks become easier to automate or augment, businesses are reassessing the role junior employees can play. Some expect this to shift organisational structures from a pyramid towards a diamond, with a larger middle and fewer entry-level roles. Others expect the pyramid to remain, but with entry-level work redesigned around AI-enabled ways of working. Wider evidence reflects the same pattern, with some companies reducing junior roles typically while shifting their workforce towards mid-level roles.⁴⁶

“You definitely don’t hire graduates to do the same thing as five years ago, but we do need them”

— Taskforce member

Businesses recognise the importance of entry-level roles, both for individuals’ first step on the ladder and for organisations in developing future capability.

There is an imperative for businesses and sectors to redesign and reimagine entry-level roles and career pathways.⁴⁷ If routine junior tasks are automated before new learning pathways are built, the pipeline into mid-level and senior roles could weaken over the next decade.

⁴⁶ Oliver Wyman Forum x NYSE Survey 2025 and 2026. Refer to [The CEO Agenda 2026](#).

⁴⁷ CBI (2026), [AI Skills Report: Building an AI Ready Workforce](#).

Some businesses are already responding by redesigning entry routes rather than simply reducing them, rebuilding graduate and apprenticeship programmes around AI-enabled ways of working, running short, structured work-experience programmes jointly with other firms in some sectors, and selecting for judgement, domain curiosity, and the ability to direct and verify AI outputs. Investing in the progression of young talent will remain necessary for organisations that want to preserve a sustainable pipeline of future leaders.

“We now risk a gap where entry-level roles are perceived to have disappeared because a lot of that repetitive work is being handled by AI models. We’ve talked a lot about how you change the way you interact with early-career workers”

— Taskforce member

“There isn’t, in principle, any reason why AI should replace people entirely rather than augment them”

— Taskforce member

Government has begun to strengthen support at this stage. Alongside the wider Youth Guarantee, the package announced at the AI Adoption Summit in June 2026 funds AI bootcamps leading to guaranteed AI apprenticeships and a jobs pilot in the North East AI Growth Zone. It also introduces new foundation apprenticeships and apprenticeship units aligned to Industrial Strategy priorities, including AI.⁴⁸

These are welcome steps, but the central challenge is to ensure that early-career routes evolve quickly enough as the content of entry-level work changes across the economy.

⁴⁸ For further detail refer to: HM Government (2026), [Youth Guarantee](#); DSIT and DfE (2026), [Entry-level AI skills package announced at the AI Adoption Summit](#), including AI bootcamps leading to guaranteed AI apprenticeships and a jobs pilot in the North East AI Growth Zone; new foundation apprenticeships and apprenticeship units aligned to Industrial Strategy priorities including AI.

Vocational and professional routes are adapting unevenly to AI adoption

The UK's AI talent pipeline is strongest at the academic and specialist end, supported by world-class universities and doctoral training. It is more uneven across the technical, vocational, and professional routes through which most workers build, apply, and evidence new skills. Several professional sectors are moving quickly to embed AI capability into professional practice. However, technical colleges, apprenticeships, sub-degree provision, and professional accreditation will be increasingly important as AI changes roles across the economy, but current provision remains fragmented and small relative to demand.

Structural problems hold these pathways back:

Many reskilling routes still require people to step out of work, while businesses need flexible, work-based training that enables employees to build and apply new skills directly in the workplace.

Professional accreditation for AI-relevant skills remains limited outside a few sectors. Sector-led models such as the Financial Services Skills Commission show how employers can define emerging skills needs and coordinate provision. Greater flexibility in how levy funding can support shorter, AI-relevant training is widely cited as a practical priority, as is the treatment of levy-raised funding across the nation, where its availability for skills is a longstanding issue.

Without a stronger vocational and professional middle layer, the UK risks having world-class specialist talent at one end and broad AI awareness at the other, but too few practical routes for the majority of workers to build, update, and evidence AI capability as their jobs change. The priority is therefore not to duplicate the academic pipeline, but to make AI-relevant learning more modular, employer-led, and portable across roles and sectors.

AI literacy is a core foundation for workforce readiness

Beyond vocational and professional routes, the UK also needs a stronger baseline of AI literacy across the population. Access to skills and talent is one of the principal barriers to AI adoption identified by the government's own analysis,⁴⁹ and demand for AI capability will extend well beyond the businesses currently investing. Jobs involving core AI activities could rise to around 3.9 million by 2035, roughly 12% of the UK workforce.⁵⁰ The current base is not aligned to that demand. Only 17% of adults say they can explain AI in detail, 28% feel confident using AI in daily life, 21% of people in work feel confident using AI tools in the workplace, and 84% of workers report no AI-related training in the past year.⁵¹

Government has set an ambitious agenda, but delivery remains complex. The AI Opportunities Action Plan⁵² sets the overall direction, while provisions span population-level literacy through the AI Skills Boost, which aims to upskill 10 million workers by 2030, technical support through conversion courses, scholarships, and doctoral training.⁵³ Following the recent machinery-of-government changes, responsibility for AI skills and workforce readiness is now distributed across several departments and delivery bodies, spanning digital inclusion and skills, AI strategy and adoption, education, employment support, and workforce transition. This creates an opportunity to embed AI more directly across relevant economic and social policy portfolios, but increases the importance of clear ownership, coordination, and routes for delivery across the system. Skills England's 2026 Annual Skills Report reinforces this concern, identifying accelerated AI adoption as one of the key challenges facing the skills system.⁵⁴ This creates a risk that AI literacy is treated as a standalone initiative, rather than as part of the wider skills, digital inclusion, and workforce transition agenda. Delivery at that scale will also depend on provision beyond the public system, and while technology firms, professional bodies, and industry academies already deliver free and low-cost AI training to large numbers of UK learners, it is not consistently mapped to national frameworks, so it sits alongside government programmes rather than counting towards them.

49 DSIT, DBT and HMT (2025), [Technology Adoption Review 2025](#).

50 DSIT (2026), [AI Skills for Life and Work](#).

51 Ibid. Based on 2024 survey evidence.

52 DSIT (2025), [AI Opportunities Action Plan](#).

53 DSIT (2026), [AI Skills Boost: explainer](#).

54 Skills England (2026), [Annual Skills Report](#).

Provision also starts too late. As AI literacy is not yet consistently embedded across schools, technical education, and apprenticeships, new cohorts often enter the workforce needing the foundational training that adult programmes are designed to provide. Adult provision is left permanently catching up, closing for one cohort a gap that the education system recreates for the next. Peer economies are already moving upstream. China has integrated AI literacy into compulsory schooling, with Beijing mandating a minimum of eight hours of AI teaching a year, while Japan has embedded AI and data-science literacy as a core component of national education policy (see exhibit 6).

For the workforce already in post, the challenge is different again. Much current AI training is pitched at degree or master's level, yet most near-term adoption depends on people already in their roles, who need practical, role-specific training rather than formal qualifications. Businesses report that training that is genuinely engaging and applicable to specific jobs is hard to find, expensive, and, because tools change constantly, not a one-off cost but a recurring one. For many firms, particularly SMEs, the combination of high training costs, uncertain returns, and the pace of change make it difficult to justify sustained investment. Provision therefore needs to shift towards affordable, modular, role-specific training aimed at the existing workforce, not only at new entrants or specialists. The same gap reaches the top of the organisation, where board and executive capability is among the least consistently built. Some firms are already investing on their own initiative, but with no common baseline or established training route for AI, as there is in cybersecurity,⁵⁵ each is doing so in isolation.

“I want every executive committee member to be able to say, proudly, that they are certified in AI, and to move towards an industry certification”

— Taskforce member

Moreover, AI literacy cannot be built in isolation from wider basic and digital skills.

Seventeen percent of UK labour force adults do not have essential digital skills for work, and 15% have not attained the foundation level for essential digital skills.⁵⁶ For a significant part of the population, AI training will only be effective if it is delivered alongside support for the literacy, numeracy, and digital confidence needed to use new technologies safely and productively.

“Thirteen million adults haven't got basic digital skills, never mind fancy AI stuff. You need some of those foundation stones to be able to adapt to how things change”

— Research panel participant

⁵⁵ For example, refer to: NCSC's Cyber Security Toolkit for Boards.

⁵⁶ Lloyds Banking Group and Ipsos (2025), [Lloyds Essential Digital Skills 2025](#). Commissioned by Lloyds Banking Group.

Workforce transition is a system-wide challenge

Workforce transition arising from AI operates at three levels: the redesign of roles within firms, the movement of people between firms, and the reallocation of people between sectors as some categories of work contract while others grow. Larger businesses are better placed to manage the first through workforce planning, retraining, and redeployment, though the availability of suitable training remains a constraint even here. The harder challenge is movement beyond the boundary of a single employer, which depends on national labour market institutions and where UK provision is least developed.

Within firms, role redesign is already under way. Some businesses are redeploying experienced staff into areas where their knowledge of the organisation remains valuable, even where their previous technical role is changing.

“A 53-year-old former network engineer turns out to be an excellent customer service agent: they understand how the technology works and can empathise with senior buyers”

— Taskforce member

Within-firm redeployment has limits. Members report that workforce planning built around growing headcount is being reshaped, and even organisations that redeploy effectively cannot absorb every affected role internally. The larger challenge arises when role change is broad-based across the economy at the same time. Evidence points to simultaneous contraction in some entry-level and routine roles alongside growth in others, a pattern that agentic AI could accelerate as systems take on larger bundles of tasks.

The challenge is no longer only whether firms can retrain their own employees.

It is whether the UK has a labour market system that can help people move from roles where demand is falling into roles where demand is growing. At present, support is still largely delivered through general re-employment services, rather than through provision linked to AI-related reskilling, sector demand, and live vacancies. Provision is weakest of all for later-career workers and those approaching retirement, despite domain expertise that could be retained through flexible, advisory, or transition roles at a time of tight labour supply and rising demand for experienced judgement.

This raises a question of responsibility. Businesses will fund the skills and redeployment their own operations require, but support that helps people move into a different employer, occupation, or sector creates benefits across the wider labour market. No single firm has a strong incentive to fund that transition on its own. This points to a shared role for government and industry in supporting transition between firms and between sectors.

Large-scale economic transitions have historically delivered their benefits only where displaced workers could move into new and growing parts of the economy. Where that reallocation was not supported — as during the contraction of heavy industry and manufacturing in the 1980s and 1990s, the period often described as deindustrialisation — the consequences for affected regions were long-lasting.⁵⁷ The occupations exposed to AI are different, spanning service and knowledge roles across all regions, but the same principle applies: systems to support reallocation are most effective when established ahead of large-scale change rather than assembled in response to it.

Other economies offer useful design principles. Sweden's job security councils support workers before redundancy takes effect.⁵⁸ Austria's labour foundations operate at company, regional, and sectoral level, helping move people out of contracting sectors and into areas facing labour shortages.⁵⁹ Singapore's train-and-place programmes connect mid-career workers to hiring sectors, with financial support during retraining.⁶⁰ The common features are early intervention, income support during retraining, provision linked to actual hiring demand, documentation that evidences courses or trainings completed, and mechanisms that move people between employers and sectors, not only within them.

57 Beatty, C., Fothergill, S. and Gore, T. (2019), *The State of the Coalfields 2019*, Sheffield Hallam University.

58 OECD (2025), *OECD Economic Surveys: Sweden 2025*, OECD Publishing, Paris, <https://doi.org/10.1787/75e94b2f-en>.

59 OECD (2025), *Promoting Better Career Mobility for Longer Working Lives in Austria*, Ageing and Employment Policies, OECD Publishing, Paris, <https://doi.org/10.1787/db85473f-en>.

60 For further detail, see here: [SkillsFuture Singapore](#), [SkillsFuture Career Transition Programme](#), and [SkillsFuture Mid-Career Training Allowance](#).

Exhibit 6: UK positioning across population segments and examples of key practices in peer countries

Population segment	Select initiatives and UK interventions	Assessment of UK position and examples of peer practices
Pre-employment groups: Primary and secondary education	- National Centre for Computing Education support for computing, machine learning and AI teaching; wider TechFirst activity to build digital and AI awareness among young people - Department for Education (DfE) guidance and support materials for safe and effective AI use in education settings, alongside pilots to test AI-enabled education tools - Former DSIT's Early Career Jobs Alliance, launched June 2026⁶¹	Moderate — Positive early activity but an opportunity for government to be more systematic in incorporating AI literacy across the secondary education curriculum supported by national targets to signal and codify an ambition to embed AI learning throughout education systems China AI literacy integrated into compulsory education, with Beijing providing a minimum eight hours per year of AI education in primary and secondary schools⁶² Japan AI and data science literacy embedded within national education policy as a core component of its AI Strategy,⁶³ supported by government guidance for teachers on the safe use of generative AI⁶⁴ EU Joint development between European Commission and OECD of an AI literacy framework for primary and secondary education, alongside guidance on the ethical use of AI for teachers⁶⁵

61 DSIT and DWP (2026), [Entry-level jobs support, AI bootcamps and tech training as government supports young people into the jobs of the future](#).

62 For further detail, refer to: PRC State Council (2026), [China aims to build an AI literacy system](#).

63 For further detail, refer to: Government of Japan, [AI Strategy 2022](#).

64 For further detail, refer to: Ministry of Education, Culture, Sports, Science and Technology, [Guideline for the Use of Generative AI in Primary and Secondary Education](#).

65 For further detail, refer to: OECD / European Union (2026). Empowering learners for the age of AI: An AI literacy framework for primary and secondary education. OECD Publishing, Paris. <https://doi.org/10.1787/65cd27d4-en>.

Population segment	Select initiatives and UK interventions	Assessment of UK position and examples of peer practices
Pre-employment groups: Tertiary education and training (including universities, apprenticeships, and vocational training)	- University studies are well supported by government programmes such as TechGrad undergraduate and Spärck AI scholarship, and Office for Students and DSIT's conversion courses and graduate scholarships for those with STEM degrees⁶⁶ - Alternative training pathways emerging through apprenticeships, for example, Skills England's AI and automation practitioner⁶⁷	Mixed — strong higher education base and specialist AI pathways but weaker coverage across technical colleges, apprenticeships, and sub-degree routes Japan Explicit target to produce ~2,000 advanced AI specialists annually⁶⁸ with a further ambition for all university graduates to acquire baseline AI and data-science literacy (~500,000 graduates annually)⁶⁹ China Five-Year Plan (2026-2030)⁷⁰ references construction of 'research-orientated universities', refinements to education curricula to integrate emerging technologies and specialisations needed by AI, and a proposal to establish a 'high-tech talent immigration system to cultivate world-class outstanding talent.'⁷¹ This is supported by its AI+ Education Plan that plans to incorporate AI skills into traditional industry training programmes to ensure talent is capable of adapting to the AI transition⁷² Sweden AI Strategy outlines the need for higher vocational education institutions and universities to offer short courses aligned with labour market needs.⁷³ A university network called AI Competence for Sweden was launched in 2018 to encourage academic collaboration on AI courses and professional development⁷⁴

66 For further detail, refer to: GOV.UK, [TechFirst](#) and [Spärck AI](#), OFS, [Postgraduate conversion courses in data science and artificial intelligence](#).

67 For further detail, refer to: Skills England, [Artificial intelligence \(AI\) and automation practitioner](#).

68 For further detail, refer to: Government of Japan, [AI Strategy 2022](#). This reflects continuity from the AI Strategy 2019.

69 Ibid.

70 International Institute for Strategic Studies, [China's 15th Five-Year Plan](#).

71 For further detail, refer to: [China's 15th 5-Year Plan](#).

72 For further detail, refer to: People's Republic of China (2026), [China aims to build an AI literacy system](#).

73 For further detail, refer to: Government Offices of Sweden, [Sweden's AI Strategy in five minutes](#).

74 For further detail, refer to: AI-Competence, [A National Cooperation Initiative for Education in AI](#).

Population segment	Select initiatives and UK interventions	Assessment of UK position and examples of peer practices
Pre-employment: Doctoral studies and specialist R&D	UKRI has announced £1.6 billion in funding for AI research until 2030. Research and training are supported through programmes such as: • UKRI's AI Centres for Doctoral Training • Turing AI Fellowship • Expansion of Doctoral Focal Awards to cover AI research in biosciences • National AI Research Resource	Strong — UK has extensive research leadership in AI-related doctoral topics, supported by significant government investment • Sweden Wallenberg AI, Autonomous Systems and Software Program (WASP) is the country's largest individual research programme, with six partner universities and 500+ active PhD students as of August 2026, designed to expand the pipeline of advanced AI talent through coordinated doctoral training, academic-industry collaboration, and long-term research funding⁷⁵ • US Several large-scale, federally supported doctoral programmes, including the National Science Foundation Graduate Research Fellowship Program that has supported research-based master's and doctoral studies in AI,⁷⁶ and the National AI Research Institutes, a collection of 29 institutes that provides ~\$20 million of funding each over five years towards AI R&D⁷⁷ • EU AI research funded through a variety of programmes, including Horizon Europe, Digital Europe, the AI Innovation Package, and the European Networks of Excellence in AI to spearhead AI research and develop a world-leading excellence in frontier AI "made in Europe".⁷⁸ The EU's Marie Skłodowska-Curie Action programme acts as a key funding resource for doctoral education and post-doctoral training, with a €6.6 billion programme budget for 2021-2027, supporting AI-related research and training projects since 2014⁷⁹

75 For further detail, refer to: [WASP](#).

76 For further detail, refer to: US National Science Foundation, [Artificial Intelligence](#).

77 For further detail, refer to: US National Science Foundation, [National AI Research Institutes](#).

78 For further detail, refer to: [European Commission, European research development and deployment of AI](#).

79 For further detail, refer to: [European Commission, EU's Marie Skłodowska-Curie Actions Research in AI](#).

Population segment	Select initiatives and UK interventions	Assessment of UK position and examples of peer practices
Displaced workers: Unemployed populations	Digital-first jobs and careers service, personalised coaching, and wider skills support delivered through the Department for Work and Pensions (DWP) in England, and similar programmes in Scotland and Wales ⁸⁰	Moderate — Broadly in line with other peers as nascent policy area. Opportunity to adapt mainstream support given likely scale of AI transition by embedding AI literacy and practical reskilling into existing jobseeker support • US Department of Labor explicitly allows Workforce Innovation and Opportunity funding to support AI literacy, training, and reskilling for dislocated workers to support a return to employment⁸¹ • Sweden Arbetsförmedlingen, Sweden's public employment service, provides practical guidance on how jobseekers can use AI tools throughout the job application and interview process⁸²
Late-career workers and retirees	Digital Inclusion Action Plan highlights the importance of digital inclusion as a priority within government with a focus on older people and the need to tailor training modules appropriately to support a possible re-entry into work, especially for technical professionals ⁸³	Weak — Limited AI-specific support for retirees and workers in transition is an underexploited policy area. Government opportunity to incorporate AI-specific advice and guidance into existing digital inclusion and literacy initiatives supported by the private sector and charities. This includes experienced technical professionals who might consider re-entering the workforce via advisory, non-executive director, and board roles but require foundational AI literacy training • Sweden AI Strategy references digital inclusion for older people, ensuring provision of services are accessible, useful and understandable.⁸⁴ Supported by the previous DigiSen project designed to improve the digital lives of the elderly through education.⁸⁵ • Singapore “Train-and-place” scheme is the SkillsFuture Career Transition Programme. Participants train before securing a new job; Government agencies oversee funding, provider approval, labour market alignment, and outcomes. Courses are designed using employer demand, vacancy data, and national Skills Frameworks. Universities, polytechnics, the Institute of Technical Education, and approved private providers deliver training and career support. Documentation is spread across programme rules, course listings, provider materials, and quality reports

⁸⁰ For further detail, refer to: GOV.UK, [Launching a modern jobs and careers service to support people into work](#).

⁸¹ For further detail, refer to: US Department of Labor, [Training and Employment Guidance Letter No. 03-25](#).

⁸² For further detail, refer to: Arbetsförmedlingen, [This is how AI can help you in your job search](#).

⁸³ For further detail, refer to: GOV.UK, [Digital Inclusion Action Plan](#).

⁸⁴ For further detail, refer to: Government Offices of Sweden, [Sweden's AI Strategy](#).

⁸⁵ For further detail, refer to: European Commission, [DigiSen collected insights about how to make digital services more accessible](#).

Recommendations

For government

Business investment in AI skills needs to be matched by public investment in the skills base of the wider population. Government should ensure that AI literacy, practical AI use, and responsible digital behaviour are embedded across education, workforce transition, employment support, entry-level jobs, and later-life learning. This should support growth, reduce labour market disruption, and help citizens use AI-enabled services safely.

RECOMMENDATION

Systematically infuse AI knowledge into education, employment support, entry-level jobs, workforce transition, and later-life training

Embed AI literacy more formally across secondary education and teacher support, broadly modelled on Japan and the EU

The DfE should build on its generative AI guidance and the 2026 support materials for schools and colleges by embedding age-appropriate AI literacy more formally across the curriculum. Young people should learn how AI systems work at a basic level, how to use them safely, how to evaluate outputs, how to protect data, how AI affects work and society, and how human creativity and judgement remain essential. Teacher training and school leadership guidance should be expanded so that educators can use AI to reduce workload and improve learning while managing safeguarding, privacy, and academic integrity risks.

Provide foundational AI skills support to transitioning workers, unemployed people, and entry-level workers, drawing on design principles from Sweden's job security councils, Austria's labour foundations, and Singapore's train-and-place programmes

Government, led by the Minister for AI and working across DCMS, BIST, DWP, DfE, and Skills England should align AI Skills Boost, AI bootcamps, Jobcentre Plus support, apprenticeships, the Youth Guarantee, and the new partnership on entry-level jobs exposed to AI. Support should focus on practical employability: using AI tools safely, applying AI to job search and workplace tasks, understanding rights and responsibilities, and building pathways into roles in data, cyber, AI assurance, software, operations, and AI-enabled sectors. Programmes should prioritise groups at higher risk of disruption and should aim to minimise time out of work by linking training to real vacancies and employer demand. Programmes should develop consent-based means for individuals to evidence transferable capability acquired — through a clearly defined skills taxonomy or framework — visible and searchable for future employers.

Deliver meaningful Growth and Skills Levy reform so businesses can invest in AI skills, building on findings in the previous CBI AI Skills report

Building on recent reforms, businesses should be able to access more of the funding they contribute and invest it in a wider range of modular, stackable, and accredited training alongside apprenticeships. Government should also establish a sustainable approach to funding skills provision for non-levy-paying employers, ensuring SMEs can access the training needed to adopt and deploy AI effectively.

Incorporate AI-specific advice into digital inclusion initiatives for older people and retirees, similar to Sweden's previous DigiSen Project

AI literacy should not be limited to students and current workers. Older people and retirees need support to use AI-enabled public services, financial services, health information, and consumer tools safely and confidently. Government should work with local authorities, libraries, community organisations, further education providers, and charities to provide age-friendly guidance on AI, scams, privacy, misinformation, accessibility, and practical use. This will support inclusion and increase confidence in AI-enabled services, while creating demand for trustworthy AI products.

Develop a national AI Literacy Standard, building on findings in the previous CBI AI Skills report and the existing skills frameworks set out earlier in this chapter

DCMS should work with BSI, Skills England, DfE, employers, and training providers to develop a national AI literacy standard that defines the knowledge, skills, and behaviours required for safe and effective AI use. This could even include modules around ethics and critical thinking. The standard should align with existing AI skills frameworks, covering both technical and complementary AI skills, and should support modular credentials that employers, colleges, and public bodies can recognise, including a board-level credential. It should also provide a route to recognise quality-assured training delivered by employers, technology firms, and professional bodies, so that industry provision counts towards a common baseline rather than sitting outside it. A common standard would reduce fragmentation, improve training quality, and give businesses greater confidence that employees and jobseekers have a shared baseline capability for using AI productively, safely, and ethically.

Build specialist pathways for the AI stack and responsible adoption to support a national approach focused on capturing value from parts of the AI stack where the UK has credible strengths, as outlined in Chapter 2

Foundational literacy must be complemented by specialist pathways in AI assurance, data engineering, cybersecurity, compute operations, AI hardware, model evaluation, product management, and responsible AI governance. Skills England, DfE, UKRI, universities, colleges, and employers should align apprenticeships, bootcamps, higher technical qualifications, and professional accreditation to these needs. This will support both business adoption and the UK's ambition to develop a sustainable AI stack, because infrastructure, assurance, and deployment all depend on a pipeline of specialist skills.

For business

Businesses should continue investing in AI capabilities and workforce skills to maximise the value of adoption. Training should therefore evolve beyond generic awareness sessions into role-specific learning paths linked to actual use cases and business priorities. There must also be recognition that AI skills needs will change as tools mature. Some skills will become easier to access through vendors and automated platforms, while others — such as judgement, problem framing, systems thinking, collaboration, communication, user-centred design skills, domain expertise, process redesign, data stewardship, assurance, and change leadership — will become more valuable.

Businesses can draw on existing skills frameworks but they should translate those frameworks into practical role expectations, training time, incentives, and career pathways.

RECOMMENDATION

Sustain investment in AI skills across every level of the organisation, from workforce literacy to role-specific capability, specialist knowledge, and leadership at board and executive level, to support business process change and implementation

Strengthen baseline AI literacy and practical usage across the workforce

Employees need confidence to use AI tools safely and effectively in day-to-day work. Baseline literacy should include what AI can and cannot do, how to evaluate outputs, how to avoid over-reliance, how to protect confidential data, when to challenge a tool, and how to use approved systems within company policy. Training should cover practical issues such as hallucinations, bias, copyright, data privacy, cybersecurity, record keeping, and the limits of automation. It should be tailored by role and function, because safe use in finance, human resources, customer service, legal, engineering, or sales will differ. Businesses should also provide employees with time to practise on approved tools, because confidence and adoption are unlikely to improve through classroom training alone. Training curricula delivered by employers should clearly link back to the foundational AI skills support to be developed by the government (see recommendation above on foundational AI literacy). This should ensure that employees can evidence the skills and training they acquire over time to facilitate workforce planning (see recommendation below on workforce planning) and support transition to new or different employment.

Build skills around business process redesign and data engineering to support scaled adoption

The constraint on scaling is more often delivery capability than technology, with firms limited by their capacity to redesign and run processes and services rather than by access to AI tools. Firms should therefore develop capability in process mapping, process mining, workflow and service redesign, data engineering, data governance, integration, delivery and programme management, product management, user research, change management, and benefits measurement. These skills depend in turn on the complementary human capabilities that make redesign work: creative problem-solving, systems thinking, collaboration, and communication. Together they translate AI tools into operating model change. For example, an organisation deploying AI in customer service needs people who can redesign triage, update knowledge bases, integrate AI into case-management systems, define escalation routes, monitor service quality, and train supervisors.

These capabilities are especially important for SMEs, which may not need large in-house AI teams but do need enough internal knowledge to specify requirements and adopt external tools well.

Develop specialist AI, risk, and assurance capability where needed

Only a minority of roles will require deep technical AI expertise, but organisations will increasingly need access to specialist skills to design, tailor, evaluate, monitor, and govern AI systems. This includes data science, machine learning engineering, model evaluation, cyber resilience, privacy engineering, procurement, legal, intellectual property, supplier risk, audit, and assurance. Firms deploying high-impact AI should be able to assess model performance, bias, robustness, security, data provenance, documentation, and failure modes. Larger firms may build internal assurance teams; smaller firms may rely on external providers, sector bodies, or shared services. The key is that someone in the organisation must be capable of asking the right questions of suppliers and understanding whether assurance evidence is meaningful.

Equip leaders to manage AI-enabled transformation

AI adoption is more likely to succeed where there is visible sponsorship from the top. Boards and senior leaders need to understand AI's strategic relevance, the organisation's risk appetite, the investment case, the workforce implications, and the governance responsibilities. Leaders should be able to decide where AI should be used, where it should not be used, how benefits will be shared, and how employees will be involved in redesigning work. Middle managers also need support, because they translate strategic ambition into changes in tasks, performance management, and team culture. Leadership development should therefore include practical modules on AI value creation, responsible adoption, workforce engagement, organisational design, and communication.

Link AI skills investment to workforce planning and role redesign

Businesses should treat AI skills as part of strategic workforce planning rather than a one-off training offer. Firms should map where AI is likely to change tasks, identify roles that will be augmented or redesigned, and create pathways for employees to move into higher-value work. This may include apprenticeships, internal academies, secondments into AI implementation teams, professional accreditation, and partnerships with colleges, universities, or vendors. Businesses should also support workers whose tasks are most exposed to automation by offering retraining early, not only after disruption occurs.

A credible workforce plan will help build trust, reduce resistance, and ensure that productivity gains are not undermined by unmanaged labour market or employee-relations risks.

Embed training into implementation rather than treating it as a separate human resources activity

Employees are most likely to build durable skills when they learn on real workflows, with managers reinforcing new behaviours and with incentives that reward adoption, quality, and responsible use. Firms should give employees enough time to practise, provide approved tools rather than leaving staff to improvise, and collect feedback on where training, guidance or tooling is falling short.

05

From
recommendations
to delivery

Three immediate actions to deliver on the recommendations in this report

To support coordinated execution and begin making progress at pace, we propose three immediate courses of action:

1. Establish a time-limited Business AI Adoption Delivery Group, a government and industry partnership feeding into the Prime Minister's AI Taskforce
2. For business, pursue immediate "no-regret" actions as a first step
3. For government, pursue immediate interventions that will support businesses in acting on adoption

ACTION 1

Establish a time-limited Business AI Adoption Delivery Group, a government and industry partnership feeding into the Prime Minister's AI Taskforce

The strategy, institutions, and ambition are in place. With AI now represented at Cabinet through the Minister for AI, and a dedicated cross-government AI Taskforce chaired by Lord Vallance, there is a natural home for driving delivery. But given the breadth of government AI activity, spanning public sector adoption, transforming public services, and supporting the growth of UK AI businesses, we see a need for focused attention on the actions government can take to accelerate private sector AI adoption as part of those wider ambitions.

This should be a government and industry partnership, convened by the CBI with other partners and sponsored by the Minister for Artificial Intelligence. As the minister leading the AI Taskforce, and with a joint appointment spanning the centre of government and the business department, the Minister is the natural link between the taskforce and business, and the group should feed directly into it.

The group should be time-limited, focused on setting shared outcomes, priorities, measures, and accountability, which can then be taken forward by standing groups and functions. The recommendations in this report represent a view from industry on where targeted action could accelerate enterprise adoption, and offer a starting point for that work.

DEEP DIVE

Developing a shared scorecard for the next five years

The first recommended task for the delivery group should be to validate success measures, establish authoritative baselines, test the level of ambition against credible data, and ensure progress is reported consistently over time.

A starting point for the 2031 ambition is summarised in the table below:

Headline KPI ⁸⁶	How it should be measured	2026 starting point	Provisional 2031 ambition ⁸⁷
Target outcome 1: Achieve productivity benefits through AI adoption			
AI-enabled productivity	Official estimate of AI's annual contribution to labour productivity and GDP, with sector and firm-size analysis	No official baseline; methodology under development	A robust official series and a sustained, measurable contribution; provisional ambition of about 0.5 percentage points a year, subject to validation
Scaled business AI adoption	Share of firms deploying AI in at least one core workflow with an accountable owner and benefits tracked; segmented by size, sector, and region	Baseline to be established using evidence from AI adoption research published by the former DSIT, BICS, and AIEI work	Fastest adoption in the G7, with a materially narrower SME and regional adoption gap
Innovation-active enterprises	OECD share and ranking of innovation-active businesses, used as a wider indicator of absorptive capacity	22 nd in the OECD, based on years 2020–22, per OECD Business Innovation Statistics Database ⁸⁸	Top three OECD performance

⁸⁶ Measurement note: AIEI and ONS should coordinate definitions and quality assurance, while lead departments, regulators, and industry provide administrative data, surveys, and appropriately aggregated telemetry. Targets should be confirmed only after baselines are validated.

⁸⁷ Ambition subject to baseline validation and should be treated as indicative only.

⁸⁸ OECD Business Innovation Statistics Database, April 2026, <http://oe.cd/inno-stats>.

Headline KPI ⁸⁶	How it should be measured	2026 starting point	Provisional 2031 ambition ⁸⁷
Target outcome 2: Scale adoption and build a competitive, resilient UK AI stack			
Advanced compute access	Public advanced-compute capacity, utilisation, median time to access, and user mix, including SMEs and researchers	Baseline to include 2025–2026 AIRR capacity and access performance	Top-five OECD capability with predictable, mission-focused access and high utilisation
Capital for AI infrastructure	Annual investment into UK AI infrastructure plus the value of AI infrastructure financing, insurance and risk transfer arranged from the UK	Not measured consistently	Top-five OECD destination for annual investment and a recognised global hub for financing and insuring AI infrastructure
AI-ready data access	Number and use of high-value datasets and trusted research environments, plus median time to access and user satisfaction	National Data Library discovery and kick-starter phase	Top-tier G7/OECD performance, with published service standards and predictable access for research and innovation
Sustainable and resilient infrastructure	Share of AI Growth Zone and supported capacity covered by low-carbon power, resource-efficiency, cyber and service-continuity plans; energy and emissions per workload	Baseline to be established across zones and major supported facilities	All supported capacity covered by published, tested plans, with declining energy and emissions intensity per workload

Headline KPI ⁸⁶	How it should be measured	2026 starting point	Provisional 2031 ambition ⁸⁷
Target outcome 3: Enable trusted, responsible, and human-centric AI adoption			
Responsible assurance uptake	Share of high-impact deployments with documented risk assessment, human accountability, monitoring, incident response, and independent assurance where appropriate.	Emerging baseline; assurance institutions and roadmap in development	A clear majority of high-impact deployments demonstrably assured; the UK recognised as a leading assurance market
Regulatory navigation performance	Median time to an informed first response, share of cross-regulatory cases with a named lead, resolution time and business satisfaction	Not currently measured	A no-wrong-door service operating to published standards, with consistently high user satisfaction and fewer unresolved cross-regulatory barriers
Target outcome 4: Build an AI-ready workforce and society			
AI literacy and confidence	Share of the population and workforce meeting a common baseline for safe, practical AI use; segmented by age, region, and employment status	Baselines emerging through skills surveys and AI Skills Boost	High baseline literacy across the population; 10 million workers with key AI skills by 2030; major participation gaps narrowed
Role-specific capability and workforce transition	Share of AI-using firms providing role-specific training; quality of early-career and technical pathways; employment and progression outcomes for supported transitions	Not measured nationally and provision remains fragmented	Most AI-using firms provide role-specific capability building, with strong transition and progression outcomes across ages, sectors, and regions

ACTION 2

For business, pursue immediate “no-regrets” actions as a first step

Complementary to the delivery group, there are no-regret actions that link directly to the recommendations in the report. Businesses should act now to:

Self-assess against the seven components of adoption best practice

Boards and executive teams should test whether leadership is aligned, priority workflows are being redesigned, governance is proportionate, central guardrails and local ownership are balanced, core data and technology enablers are in place, experimentation is disciplined, and value is being measured.

Move a small number of priority workflows from pilot to scale

Each should have a business owner, a pre-deployment baseline, a benefit case, explicit scale-or-stop criteria, and a plan for process, operating model, and workforce change.

Put trust and resilience into the design

Firms should risk-tier use cases, keep people accountable for high-impact decisions, strengthen supplier and data controls, test continuity arrangements, and monitor performance after deployment.

Build skills in the flow of implementation

Role-specific learning, manager capability, and workforce engagement should accompany each deployment rather than sit as a separate generic training programme.

Share what can be shared in forums that will develop and evolve over time

Businesses should contribute anonymised implementation patterns, governance templates, and evidence to sector and regional networks, while protecting intellectual property, security, and commercial confidentiality.

ACTION 3

For government, pursue immediate interventions that will support businesses in acting on adoption

Government and public bodies should clear the path for adoption nationwide.

Immediate actions include:

Use government adoption and procurement as anchor demand

Move successful public-sector exemplars into reusable components, publish implementation lessons, and make procurement proportionate, interoperable, and accessible to innovative UK suppliers and SMEs.

Pilot a “no-wrong-door” regulatory approach

The DRCF, RIO, and relevant sector regulators can establish an informed first-response protocol, a practical regulatory map, and named coordination for complex cross-regulatory cases without creating a new approval authority.

Mobilise trusted local diffusion channels

AI Champions, universities, colleges, professional bodies, trade associations, combined authorities, and AI Growth Zones should combine diagnostics, implementation support, skills, finance signposting, and peer learning for SMEs.

Keep removing infrastructure and data bottlenecks

Public compute, AI-ready data, planning, grid connections, power, mobile network connectivity, cyber resilience, and place-based supplier development can progress in parallel with the measurement work.

Start the baseline now

The AI Economics Institute, ONS, and BIST should align survey instruments and administrative data so the first national scorecard measures adoption depth and outcomes, not simply awareness or tool usage.

CBI commitment

The CBI also recognises that it can play a role in diffusing AI adoption nationwide. Key commitments following this report include:

-
- 01 Provide the private sector voice and participate in the proposed Business AI Adoption Delivery Group (proposed in Action 1).
 - 02 Support knowledge sharing across firms through relevant knowledge-sharing forums (proposed in Action 2) with a priority deliverable of developing a clear “employer view” skills framework and taxonomy to advise the government as it develops and evolves foundational AI skills support programmes (for instance AI Skills Boost).
 - 03 Provide ongoing feedback on adoption behaviours and challenges from firms as findings emerge, supporting government to prioritise the right interventions for business (proposed in Action 3).
 - 04 Work with an independent and credible convener to bring together business, workers, civil society, and government to explore shared principles for responsible AI adoption (as set out in Chapter 3).

Oliver Wyman and CBI project team

Oliver Wyman

Lisa Quest

CEO, Marsh UK
Head of UK and Ireland, Oliver Wyman

Davide Taliente

Partner, Oliver Wyman

Rupal Kantaria

Partner, Oliver Wyman

Deborah O'Neill

Partner, Oliver Wyman

Anthony Charrie

Partner, Oliver Wyman

Elizabeth Hoyler

Principal, Oliver Wyman

Trupti Rao

Senior Consultant, Oliver Wyman

CBI

Rain Newton-Smith

CEO

John Foster

Chief Policy and Campaigns Officer

Naomi Weir

Director, Innovation

Shivani Gupta

Senior Policy Advisor

Helena Coe

Policy Manager

About Oliver Wyman

Oliver Wyman, a Marsh (NYSE: MRSB) business, is a management consulting firm driven by deep industry insight, bold innovation, and a collaborative approach that cuts through complexity to help organizations navigate their most defining transformative moments. Marsh is a global leader in risk, reinsurance and capital, people and investments, and management consulting, advising clients in 130 countries. With annual revenue of \$27 billion and more than 95,000 colleagues, Marsh helps build the confidence to thrive through the power of perspective.

For more information, visit oliverwyman.com, or follow us on [LinkedIn](#) and [X](#).

Americas
+1 212 541 8100

Europe
+44 20 7333 8333

Asia Pacific
+65 6510 9700

India, Middle East & Africa
+971 (0) 4 425 7000

About CBI

The CBI represents 850 members who themselves comprise 1,200 separate registered companies and over 150,000 trade association members. Founded 60 years ago and representing some of the biggest names in business, household brands and globally traded corporations that employ people in all sectors and across every region and nation of the UK, the CBI is the voice of business.

Copyright © 2026

This report has been prepared by Oliver Wyman and the Confederation of British Industry (the "Authors"), with input from broader industry, as well as AI experts and practitioners. The Authors shall not have any liability to any third party in respect of this report nor any actions taken or decisions made as a consequence of the results, advice, or recommendations set forth herein. This report does not represent legal advice. The opinions expressed herein are valid only for the purpose stated herein and as of the date hereof. Information furnished by others, upon which all or portions of this report are based, is believed to be reliable but has not been verified. No warranty is given as to the accuracy of such information. Public information and industry and statistical data are from sources the Authors deem to be reliable; however, the Authors make no representation as to the accuracy or completeness of such information and have accepted the information without further verification. No responsibility is taken for changes in market conditions or laws or regulations and no obligation is assumed to revise this report to reflect changes, events, or conditions, which occur subsequent to the date hereof. All rights reserved. This report may not be reproduced or redistributed, in whole or in part, without the written permission of the Authors.