techUK
FOR WHAT COMES NEXT

Industry brief: What it takes to scale the responsible adoption of Agentic AI

July 2026

Seizing the AI Opportunity



About techUK

techUK is the trade association that brings together people, companies and organisations to realise the positive outcomes of what digital technology can achieve. With over 1,100 members across the UK, techUK creates a network for innovation and collaboration across business, government and stakeholders to provide a better future for people, society, the economy and the planet.

Executive Summary

Agentic AI has captured the attention of organisations, policy makers and technologists worldwide.

Unlike more traditional AI tools that produce a single output in response to a prompt or follow a specific sequenced process, agentic AI can plan sequences of actions, break down complex objectives into steps and carry out tasks by interacting with external systems and other agents to pursue goals over longer horizons with limited step-by-step human direction. This capacity for multi-step coordination and action, rather than simply producing an output for review, represents a meaningful shift in how AI can be embedded into organisations and offers significant opportunities across the private and public sectors to improve productivity, efficiency and effectiveness. In February and March 2026, techUK held a four-part workshop series to examine and capture what industry best practice for scaling agentic systems looks like today.[1]

The workshops also explored organisational and regulatory barriers to adoption, workforce preparedness, and the ethical, governance and accountability challenges that the use of agentic AI systems may raise within organisations.

This brief is not a descriptive write-up of the workshop discussions. Instead, it captures key insights from those discussions and offers recommendations to support organisations in overcoming barriers and scaling the use of agentic AI. It also includes recommendations for action by the UK Government and regulators to support increased adoption and use of agentic AI by UK organisations.

Key learnings from our agentic AI workshop series

techUK's workshop series was an opportunity to take stock and assess where organisations are right now in their agentic AI adoption journey and identify challenges and barriers that may be preventing more organisations from adopting and scaling the use of agentic AI. Key findings gathered from the workshop discussions included the following:

- **Most organisations are not yet scaling agentic AI widely.** Deployment remains concentrated in proof of concept and pilot stages, where systems are tested in controlled environments before broader rollout. This is consistent with DSIT's January 2026 AI Adoption Research, which found that 16% of UK businesses currently use AI, and that among those adopters, agentic AI is the least adopted category, with only 7% reporting use.[2]

- **Governance is a prerequisite for scale, not an obstacle.** Trust in data, in model outputs and in governance processes is what allows confident deployment. Governance built in from the outset enables scale. Governance added later, after systems are already deployed, slows it.
- **Workforce preparedness for the agentic era is seen as key.** Short term deployment choices are being made without sufficient consideration of medium term implications for skills pipelines, junior talent and institutional knowledge. We found that many organisations are not yet reflecting a trajectory from task automation to role automation in their operating models.
- **Accountability in multi-agent systems is raising questions.** Questions such as where accountability sits when one organisation's agent interacts with a supplier's agent and at what point accountability transfers are genuinely unresolved and acting as barriers to confident deployment.

Four key ingredients for success

Based on learning and insights gathered from techUK's workshops, we have identified four interdependent areas that organisations need to focus on and have in place to be successful in adopting agentic AI. We believe that progress in any one area will enable progress across the others, while weakness in one area has the potential to constrain the whole. These four ingredients of success are:

1. **technical readiness,**
2. **organisational structures and ways of working,**
3. **workforce preparedness and**
4. **governance.**

This brief examines each area in turn and offers targeted recommendations for organisations to help support their use of agentic AI. These recommendations are grounded in the reality of where deployment stands today and the barriers organisations are facing. Recommendations for action by government and regulators are also included as techUK has identified a number of policy levers that could also be used to help unlock the responsible scaling of agentic AI at pace.

techUK believes the UK is well placed to lead in demonstrating how responsible agentic AI deployment can be achieved. However, realising that potential will require action both from organisations looking to adopt agentic AI and from government and regulators to help overcome existing barriers and address future challenges as these systems continue to develop. This will be essential to ensure UK organisations are able to adopt agentic AI systems with confidence and at scale.

Introduction: What is agentic AI, why does it matter, and where is adoption today?

What do we mean by agentic AI?

Before exploring the four areas techUK has identified that organisations need to focus on to scale the adoption of agentic AI, it is first important to be clear about what is meant by agentic AI. Insights from techUK's workshop series highlighted that right now the phrase agentic AI is often being used to describe a wide range of different systems that are quite different. For example, agents, agentic AI and agentic systems mean different things:

- An **individual agent** is a software component that performs a specific task such as summarising a document or checking a dataset for errors,
- **Agentic AI** describes a class of capabilities that enable autonomy,
- **Agentic systems** are complete deployments that include multiple agents working together, coordination layers that manage how they interact and human oversight mechanisms.

To support organisations looking to adopt agentic AI, the following working definition is offered to promote more consistent terminology across organisations and among policy makers. techUK recognises that the language in this area is still evolving and that this definition may need to develop in parallel with the technology itself.

techUK's definition of agentic AI

Agentic AI has the ability to accelerate AI deployment through its capacity to take actions autonomously to achieve objectives. These systems operate by planning, meaning they can break down goals into sequences of steps; making decisions, meaning they can select between options based on changing conditions; and adapting to new information with minimal human intervention. These systems are capable of learning from interaction, being goal oriented rather than task specific and operating autonomously without continuous user direction.

Why agentic AI matters

Throughout our workshop series, participants identified agentic AI as representing a significant productivity and economic opportunity for UK organisations. Agentic AI can accelerate workflows that currently require constant human intervention, automate complex multi-step processes across functions, and enable new forms of service delivery that were previously too slow or resource-intensive to scale. Participants pointed in particular to productivity gains in high-volume knowledge work, faster decision-making, and the potential to free workers from repetitive tasks to focus on higher-value activities. The potential productivity benefits of scaling agentic AI help explain growing enterprise interest in the technology in the UK. A recent industry report based on insights from 1,100 tech leaders found that 69% of UK organisations say most or all teams have adopted AI agents, running an average of 13 agents each and expecting that number to double within two years.[3]
For UK businesses, this represents a chance to improve competitiveness, reduce operational costs and unlock new sources of value creation.



Where agentic AI is being deployed today

While many organisations are running pilots and proofs of concept, agentic AI deployment remains concentrated in a small number of sectors. Financial services are leading in early adoption, and are beginning to deploy agents across anti-money laundering workflows (AML) to gather and summarise case notes, know-your-customer (KYC) processes to research financial ledgers at scale and credit risk reviews to structure evidence, with humans retaining sign-off authority. A deeper dive into financial services use cases and the agentic AI landscape is explored in techUK's report, Agents of Change: Generative and Agentic AI in Financial Services 2026, which was recently published on 24 June 2026.[4] Beyond financial services, software development has become a cross-sector use case for coding assistance, unit testing and documentation, while pharmaceuticals are also beginning to explore more autonomous systems to support regulatory drafting and patient information workflows, with multiple components checking whether each version aligns with the specific rules of different countries before human review and approval. Looking ahead, agentic commerce, where agents find, compare and purchase on behalf of users, is emerging but remains experimental.

Scaling agentic AI: four ingredients for success

To help make sense of what is needed in practice, this section sets out four areas that workshop participants consistently returned to when discussing how to scale agentic AI responsibly. They are presented separately for clarity, but in practice they are interconnected and organisations are unlikely to scale successfully unless all four are in place.

1. Technical readiness: tooling capabilities for agentic AI

The use of technical tools will be crucial in how organisations move from experimentation with agentic AI to deployment. Without tools and systems for orchestration, process mapping, evaluation, observability and control, organisations may struggle to deploy agents confidently in live environments. This section draws on workshop insights from industry leaders with deployment experience and highlights the technical capabilities and tools that can help organisations scale agentic AI responsibly. These include not only tools for coordinating and monitoring agent behaviour, but also the practical control mechanisms needed to intervene when systems behave unexpectedly.

Orchestration platforms

Orchestration platforms coordinate multiple AI agents working together and are a core part of making an agentic system more than the sum of its single agents. As agentic deployments scale, the number of agents, dependencies and workflow handoffs can increase rapidly. Without effective orchestration, organisations may find that agent interactions become unpredictable, outputs do not reliably pass between stages, and failures are difficult to troubleshoot. Participants also highlighted that within agentic systems complex multi-level agent hierarchies, where agents spawn further agents and pass control across several layers, can create loop problems, unclear ownership of tasks and unpredictable behaviour without corresponding improvements in performance. A clear success factor is

knowing at each step of a workflow which agent controls which task and being able to trace and verify that control structure in operation.

Business process discovery tools

Business process discovery, meaning the systematic mapping of existing workflows before agent deployment begins, was identified as an important enabler of responsible scaling. It helps organisations understand where autonomy can add value, where human intervention must remain, and where dependencies across legacy systems may create operational or compliance risks. Participants noted that this kind of process mapping can allow orchestration layers to be added on top of existing systems without requiring full infrastructure replacement.

Risk-based design and control tooling

Participants emphasised that technical readiness is not only about enabling agents to act, but also about designing control into systems from the outset. Risk-based design, meaning assessing the potential impact of each agent action and building mitigation into the workflow before deployment, was cited as a differentiator between deployments that scaled and those that stalled. This includes tools and approaches for setting permissions, constraining actions, managing fallback rules, defining escalation thresholds, implementing override or shutdown mechanisms where higher-risk tasks require stronger control, and regularly testing these mechanisms in practice so that intervention remains realistic when needed. While these are technical design choices, they also create governance questions about who can intervene, under what conditions and with what evidence.

Evaluation and testing tools

Workshop discussions also pointed to the importance of evaluation and testing tools that go beyond standard model benchmarking.



Organisations need ways to test multi-step workflows, agent-to-agent handoffs and performance under changing conditions before systems are deployed more widely. Participants noted that without structured testing, organisations may struggle to distinguish between impressive pilot performance and genuine production readiness. Workshop participants also argued that organisations need clearer outcome-based measures of value as agentic systems move beyond pilots. They stressed that organisations should assess not only whether systems function technically, but whether they improve outcomes such as decision quality, error reduction, process reliability or service performance in live environments. Without this, organisations may overestimate the value of deployments based on narrow proxies such as time saved or model activity rather than real business impact.

Observability tools

Observability, meaning visibility into what agents are doing throughout workflows, including the chain of decisions and actions from input to output rather than only at the start and end points, was consistently identified by participants as essential. While observability tools exist today, many organisations deploying agentic AI have not yet built adequate observability capabilities into their systems. Monitoring only at endpoints may be insufficient to detect errors, manipulation or undesirable behaviour occurring at intermediate steps, particularly when agents act on each other's outputs at speed. We found that many current deployments would struggle to reconstruct a full decision trail if required by an auditor or regulator because logging and traceability were not designed in from the outset.

2. Organisational structures and ways of working

Alongside technical readiness, organisations also need internal structures and ways of working that support responsible deployment at pace. This section looks at the organisational conditions that can help agentic AI move beyond isolated pilots and into wider use.

Clear internal structures and single ownership of agentic AI

Based on members' experiences shared in the techUK workshops, we found that different functions within organisations may use different terminology for agentic AI, operate on different timelines and hold different views on what constitutes an acceptable pace of deployment. For example, first-line functions such as sales teams, product development and customer-facing operations tend to be capability-focused and experimental in orientation, wanting to move quickly to test what agentic AI can do. However, second-line functions such as legal, compliance and risk management are responsible for governance and risk oversight and may therefore prioritise safety and regulatory compliance over speed of deployment. This can create internal friction, slow decision-making and make it harder to establish clear ownership for deployment decisions. Organisations therefore need clear internal routes for agentic AI deployment decisions, with defined ownership and points of engagement across legal, compliance, risk, security and delivery teams.

3. Workforce preparedness: skills and capabilities needed and the future of work

Workforce preparedness will be just as important as technical capability in determining whether agentic AI can be scaled successfully. This section explores the skills, capabilities and wider labour market questions organisations will need to consider as adoption grows.

AI literacy and fluency

AI literacy and AI fluency have emerged as foundational capabilities for workers to adapt and upskill in the AI era. AI literacy, meaning a basic understanding of what AI systems can and cannot do, how they operate and where they may introduce bias or error, enables workers across functions to engage critically with AI outputs rather than accepting them uncritically. AI fluency, meaning the ability to work effectively with AI tools, validate outputs and know when to intervene, is increasingly relevant beyond technical specialists. This matters not only at organisational level but increasingly in the wider UK policy landscape. The UK Government has already begun to frame AI skills and workforce preparedness as part of a broader economic transition, including through Skills England's work on AI skills and the launch of a Future of Work Unit to examine how AI may affect jobs and labour market patterns over time.[5][6] techUK welcomes this focus by the UK Government and sees it as an important step towards ensuring that the UK's approach to agentic AI adoption is supported by practical thinking on skills, workforce adaptation and longer-term economic resilience.

Technical skills and capability

Workshop participants highlighted that organisations looking to deploy agentic AI at scale will need a broader technical capability base than many currently have in place. This includes skills in prompt and workflow design, orchestration, tool integration, API connectivity, system monitoring, observability, evaluation and testing, and security controls for agent behaviour. Participants also noted the need for technical teams that can design reliable handoffs between agents and other systems, set clear operating parameters and fallback rules, and diagnose failures across multi-step workflows. Taken together, these capabilities are essential for moving from experimentation to robust, production-grade deployment.

Managing agentic AI fleets

According to workshop participants, this workforce generation may be among the last that primarily manages only humans. As AI agents become more common in the workplace, a key skill will be knowing how to work alongside AI agents and coordinate multiple agents working together. Ensuring employees have the right skills to do this will be critical. Organisations therefore need to treat workforce readiness as a core part of deployment planning rather than something that follows technical implementation.

Some of these skills are specific to agentic AI. Workers will need to know when to step in and oversee agent decisions versus when to let agents work autonomously. They will need to understand what agents can and cannot reliably do, so they can set appropriate expectations and catch problems early. Validating agent outputs



is essential. Workers must be able to assess whether agent decisions are sound rather than accepting them automatically. This includes being able to define clear problem boundaries and parameters for what agents should achieve, as well as thinking ahead about what could go wrong and planning accordingly.

However, participants also emphasised that broader capabilities remain essential, including critical thinking, risk assessment, and the ability to question and challenge system outputs. The balance between technical skills like prompt design and system oversight, and these broader judgement-based capabilities, will likely vary by role and context. But the ability to validate what agents are doing and exercise human judgement appears to be emerging as a core capability across many different types of roles.

Impact and implications for graduate and entry-level work

A recurring concern in our workshops was the disruption of agentic AI for graduate roles and entry-level work. Participants discussed whether repetitive foundational work historically performed by junior staff has genuine developmental value or whether it is simply accepted tradition. Participants did not reach a settled view, but a shared concern emerged: if this work is automated before organisations develop alternative pathways for junior talent to build judgement and contextual knowledge, the availability of mid-level capability could be weakened in the years to come. This is not just an organisational risk but potentially a macroeconomic one for the UK. Organisations that reduce graduate hiring or heavily automate entry-level roles today may find they lack mid-level talent with the experience and judgement to manage complex work in the future.

These concerns are also beginning to be recognised amongst policy makers. We believe that institutions and initiatives such as the proposed AI Economics Institute and the Future of Work Unit all have a role in helping monitor how agentic AI affects hiring, progression and skills formation over time, and in identifying where targeted intervention or retraining support may be needed.[7]

4. Governance: accountability, oversight and assurance

Governance frameworks play a key role in supporting organisations to scale agentic AI adoption responsibly. As these systems take on greater autonomy and operate across organisational boundaries, ensuring clear accountability, effective oversight and appropriate regulatory compliance become increasingly important for organisations. This section examines the governance questions and challenges that organisations are currently facing when deploying agentic AI at scale and offers advice and support for how these issues can be overcome. This includes the role that AI assurance tools and mechanisms can play in building confidence in autonomous systems, and how cyber security considerations can and must be addressed to enable safe deployment.

Accountability

Accountability was identified as one of the most pressing challenges holding back confident scaling of agentic AI. When agents make decisions autonomously, and particularly when multiple agents from different systems work together, existing accountability models designed for human decision-making or single-model AI struggle to apply clearly according to participants.

Questions from participants included: where does accountability sit when an agent from one organisation interacts with an agent from a supplier and their combined output causes harm? At what point in an agent chain does accountability transfer and under what evidential threshold? How should responsibility be allocated when a harmful outcome results from an update to one agent that creates problems downstream in other agents?

While it is recognised that issues of shared accountability can and are being addressed through commercial contractual agreements today, workshop participants raised that guidance from government and regulators on how accountability should be allocated in multi-agent deployments could further help organisations to deploy with greater confidence. For example, this could include offering clarity on issues such as how existing data protection concepts such as data controller apply when agents operate autonomously and providing sector-specific guidance for regulated industries such as healthcare and financial services where professional accountability structures already exist and could be adapted for agentic deployment. Participants also noted that some of these questions may require further engagement between government, regulators and industry as deployment evolves. As agents increasingly operate across different legal jurisdictions, questions also arise about how oversight should work when systems from different jurisdictions interact. This is an area where international coordination may also become increasingly important.

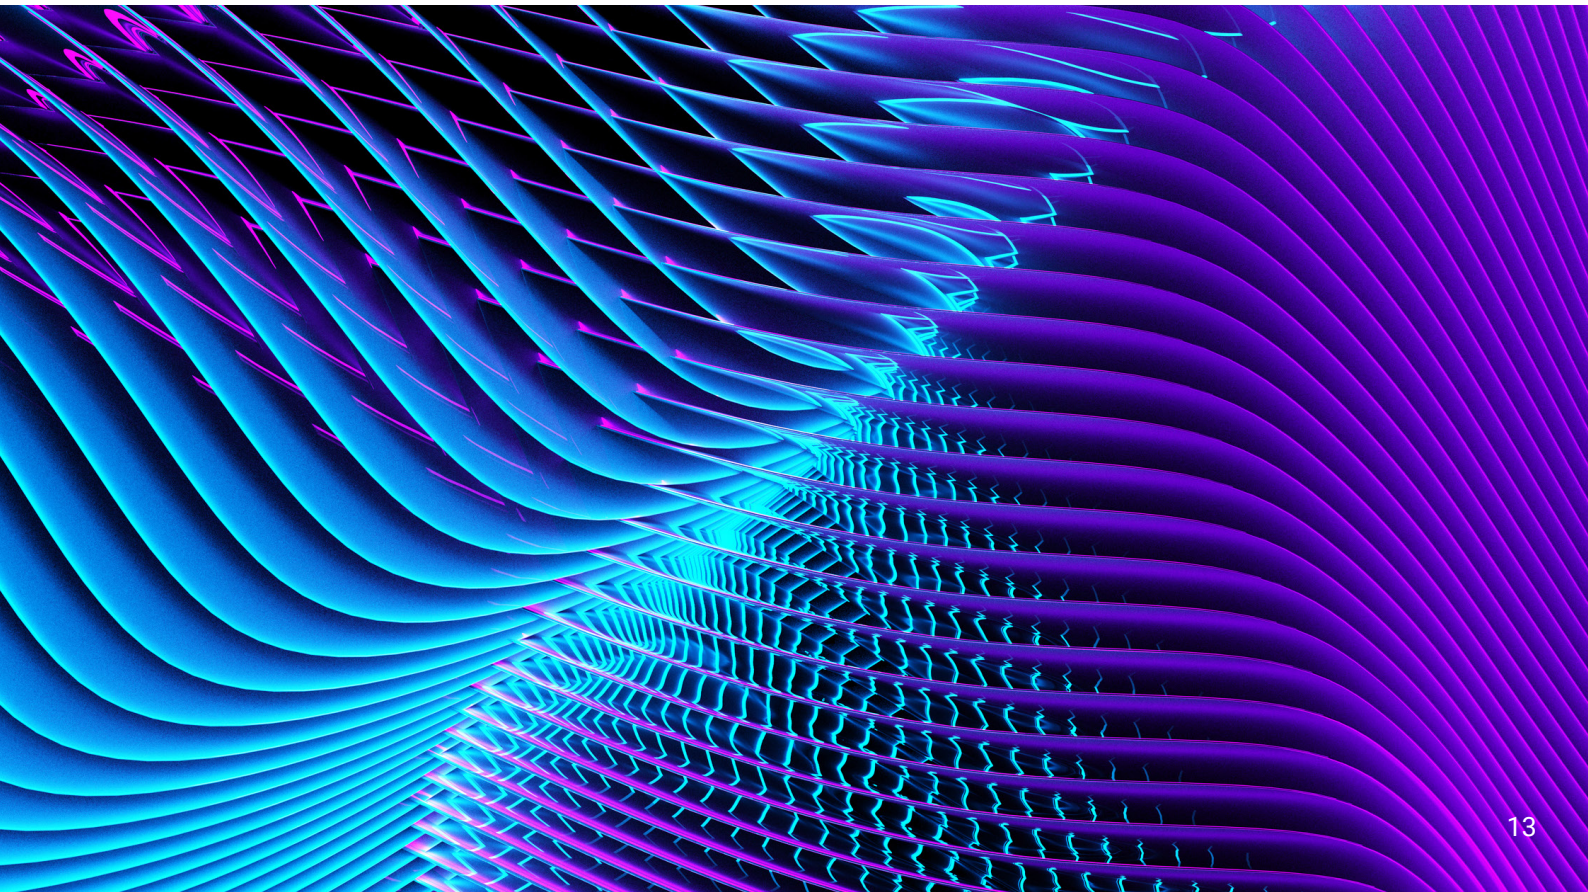
Data governance

Participants stressed that data governance is the foundational layer on which everything else depends for successful agentic AI deployments. This includes getting the data basics right, that is to say clean, reliable real-time data, as a prerequisite for confident deployment. As agentic systems become more capable and take on greater autonomy, the quality of the data they

access becomes even more critical. Agents acting on poor quality, incomplete, or outdated data can make decisions and take actions that compound errors rapidly before human oversight can intervene. When agents operate at speed and coordinate with other agents, data quality issues that might be manageable in supervised systems can cascade into significant failures in autonomous workflows. Organisations that have not invested in the underlying data estate including data quality processes, master data management and real-time integration capabilities are unlikely to deploy agentic systems responsibly at scale.

Meaningful human oversight

Meaningful human oversight was identified by workshop participants as an essential condition for responsible agentic AI deployment. In practice, this means including genuine veto power, clear override and escalation routes, and enough time and information for a person to exercise judgement rather than simply rubber-stamp agent decisions. Without these features, oversight risks becoming procedural rather than effective. Furthermore, autonomy is changing what oversight means for AI systems. We have heard from members that as systems operate at greater speed and across more complex workflows, the traditional “human in the loop” model, where a person checks every action, will need to evolve to remain meaningful and effective. This can mean, in some contexts, moving to a “human on the loop” approach, where a person supervises and intervenes when needed. At scale, meaningful oversight should not require a person to review every individual action. Members pointed instead towards a more targeted model of oversight. Rather than a person checking every output,



a set of safeguards can flag the situations that need a human to step in, so that attention goes to the decisions that carry real consequences, rather than to routine activity that is working as intended. These safeguards include anomaly detection, which watches for behaviour that departs from the expected pattern and raises an alert when something looks unusual; escalation thresholds, which are pre-defined triggers for human review, for example when an agent attempts a higher-risk action, behaves unexpectedly, exceeds agreed confidence or risk limits, or could create consequences that are difficult to reverse; and, where appropriate, automated circuit breakers, which pause or stop a system automatically when one of those limits is crossed, in the same way a fuse cuts the power before a fault can cause damage. Moreover, the appropriate model, whether a human in the loop, on the loop or out of the loop, should be determined by a risk-based assessment of the specific use case rather than applied uniformly. We believe that an approach with tighter oversight for novel, high-stakes or irreversible actions and lighter-touch monitoring for well-understood, low-risk and reversible tasks, is more practical than defaulting to maximum oversight everywhere. Organisations should therefore develop tiered oversight models calibrated to the actual risk of each deployment.

Oversight and intervention responsibilities

Participants stressed that even where control mechanisms are designed, whether to monitor outputs and actions, to set and manage permissions or to review and manage a systems overall performance, governance still needs to determine who is responsible for operating them, under what conditions intervention should happen and what evidence is needed to support those decisions. In practice, this means organisations need clear responsibilities for oversight, documented escalation routes and confidence that human intervention remains meaningful rather than nominal. Participants argued that these organisational questions become more important as agentic systems operate across longer workflows, involve multiple agents and affect higher-risk decisions.

We recognise that the term oversight is often used to cover different things, and clarity on what it means is important. For example, it can mean keeping watch over what an agent does, by monitoring its actions and the outputs it produces. It can also mean limiting what an agent can do in the first place, by setting clear permissions and boundaries when the system is designed. Finally, it can cover overseeing how the system is behaving, including its overall performance, whether its behaviour changes over time, and whether agents begin to interact in unintended ways.

Setting limits at the design stage can often be an efficient way to manage risk at scale, because if an agent is never permitted to take a particular action, that risk does not arise and does not need to be watched for. Constraints of this kind are best understood as prevention rather than oversight: they can be the first line of defence, narrowing the range of behaviour that oversight then has to cover. However, this approach does not remove the need for other forms of oversight. Ongoing monitoring, regular review of how the wider system is performing, and meaningful human oversight remain essential for the risks that design constraints cannot fully rule out.

It is also worth recognising that safety controls can sit at more than one level of the value



chain. Some, such as input and output filtering, rate limiting, content moderation and monitoring, are increasingly built into the foundation models and platforms that organisations deploy on top of, while others are implemented by the deploying organisations itself. Both the developers and deployers therefore have a part to play in assessing whether these controls are adequate for a given use case and context. Where these controls technically sit, however, does not determine where accountability lies. As set out above, how responsibility should be allocated across developers, providers and deployers remains an open question on which this brief does not take a position.

Cyber security risks and how they can be addressed

As agentic AI systems take on more complex tasks and interact across tools, data sources and other agents, organisations will also need to think carefully about how cyber security risks are identified and managed in practice.

Workshop participants supported exploring whether adversarial testing and red-teaming, practices already established in cyber security contexts, should become standard governance requirements for high-risk agentic deployments. However, identifying risks is only part of the solution, as organisations also need practical guidance on how to address them.

Workshop participants also pointed to the AI Security Institute's work on evaluation, testing and AI security as an important source of emerging methods and lessons for agentic AI deployment. However, this work is pitched at frontier model evaluation rather than at enterprise deployment teams, meaning that organisations deploying agentic systems in production do not yet have sufficiently practical guidance on how to manage risks such as prompt injection (when hidden or malicious instructions are slipped into the content an agent reads so that it carries out actions its operator never intended), tool misuse, agent-to-agent manipulation and cascading failures across multi-step workflows. This highlights a wider gap between existing

frontier-focused assurance work and the more operational guidance organisations need to deploy agentic AI confidently at scale. It is worth noting that some practical mitigations for risks such as prompt injection, and tool misuse are already being developed and shared within the open-source community to be used by organisations to mitigate against any potential risks. Fully addressing this gap however will require further and future work to translate relevant methods and insights that could support organisations into more practical guidance for enterprise deployment.

Participants also noted that this sits within a wider UK effort to build a more credible AI assurance ecosystem, including work on the Trusted Third Party AI Assurance Roadmap intended to improve access to independent assurance, audit and validation as adoption scales.[7] Workshop discussions also highlighted a more practical deployment challenge: many organisations do not yet have easy access to usable assurance methods, clear expectations about evidence, or credible routes to independent validation for more autonomous and multi-step systems. These barriers are likely to be felt most acutely by smaller firms and by organisations moving from proof of concept to production. The roadmap therefore creates a practical policy hook for supporting adoption. techUK believes government should use the AI Assurance Innovation Fund to back practical pilots that test assurance methods, metrics, logging and audit approaches for real agentic AI deployments. This would help ensure the fund supports not only frontier capability assurance, but also the more operational and enterprise-facing assurance methods organisations need in order to deploy agentic AI with confidence at scale.

Assurance mechanisms: building confidence through independent validation

Workshop participants also pointed to the role that impact assessments have to play as a practical assurance tool for organisations to use to support higher-risk agentic AI deployments. Participants highlighted both workforce impact assessments and algorithmic impact assessments as useful ways to document risks, potential biases and safeguards. It was felt this was one of the clearer areas where government could take practical action by encouraging their use, and where appropriate requiring them, in higher-risk or public sector deployments. Participants also pointed to existing governance frameworks and standards as a useful baseline for agentic AI deployment. However, the discussion suggested that the challenge is not only that additional deployment-specific guidance may be needed, but that key frameworks such as ISO/IEC 42001 were not designed with multi-agent coordination and autonomous operational workflows in mind. Participants argued that the wider standards and assurance ecosystem around these frameworks is not yet mature enough for agentic deployment, particularly when compared with the more established ecosystem that has developed around standards such as ISO/IEC 27001, the globally recognised international standard for Information Security Management Systems (ISMS).

The role of UK regulators

techUK believes UK sector regulators have an important role to play in supporting organisations to deploy and scale agentic AI responsibly by offering guidance, advice and practical support. For example, the Information Commissioner's Office (ICO) set out its early thinking in its

January 2026 Tech Futures: Agentic AI paper, highlighting both the opportunities of agentic AI and the implications for accountability, risk management and privacy by design in agentic systems.[8] In March 2026, the Competition and Markets Authority (CMA) published guidance for consumer-facing organisations clarifying how consumer protection law applies to AI agents and emphasising that businesses remain responsible for agent behaviour.[9] The Digital Regulation Cooperation Forum (DRCF), which brings together the CMA, ICO, Ofcom and the FCA, recently published its foresight paper on The Future of Agentic AI.[10] This scoped what cross-regulator work on agentic AI should look like. The next step is to operationalise this. Alongside this, the DRCF, has used its Thematic Innovation Hub to first focus on agentic AI adoption to surface practical issues including assurance at machine speed, cross-regulatory coherence and delegation.[11] While this has provided a useful forum for continued and sustained engagement with industry stronger coordination going forward will be needed to address the barriers organisations may face when moving from pilots into live deployment. While techUK welcomes the collaborative approach being taken, we also encourage clear mechanisms for regulatory coherence in practice that avoid duplicative reporting and ensure consistency when multiple duties apply to a single AI deployment. As recent joint actions have shown, a single AI application may need to engage and comply with ICO, CMA, Ofcom and FCA guidance and codes simultaneously. For this reason, coordination between regulators is essential to make the system workable.



Recommendations

The recommendations below set out practical actions for organisations, government and regulators to support the responsible scaling of agentic AI.

Actions for organisations looking to adopt and scale agentic AI

The following actions should be considered by organisations looking to move their use of agentic AI beyond proof-of-concept and pilot stages.

Technical readiness: tooling capabilities for agentic AI

- **Determine the right tools for tasks and design orchestration platforms effectively.** Organisations should use robotic process automation for routine tasks that follow fixed rules, traditional machine learning where pattern recognition and prediction are needed, and more advanced AI models where tasks genuinely require reasoning, planning or adaptation. In agentic deployments, they should also design orchestration layers deliberately so that handoffs, permissions, sequencing and escalation routes across agents remain clear and controllable. This is particularly important where multiple agents or systems interact across a workflow, as not every step needs the same level of autonomy. A layered approach can reduce cost, improve control and make it easier to identify where failures occur when systems act across multiple steps.
- **Build observability in from day one.** For agentic systems, monitoring inputs and outputs alone is unlikely to be sufficient. Organisations should design in logging, trace-ability and visibility across the full workflow, including agent-to-agent handoffs, tool use and decision sequences. This is necessary not only for governance and security but also for practical deployment, as organisations may otherwise struggle to reconstruct why a system acted as it did when something goes wrong.
- **Map processes before deploying agents.** Agentic AI is most likely to scale successfully where organisations understand the process being automated before introducing autonomy into it. Business process discovery can help identify where agents should and should not act, where human intervention remains necessary and where dependencies across systems could create operational or compliance risk. This is particularly important where multiple agents or external systems may be interacting in sequence.
- **Measure whether agents are improving real outcomes.** Organisations should track whether agentic deployment is improving business outcomes such as decision quality, lower error rates,

faster case resolution or better customer outcomes, rather than relying only on measures such as time saved or token usage. This matters because agentic systems can create new forms of risk or rework that may be missed if deployment is judged only through narrow productivity measures.

Organisational structures and ways of working

- **Create clear internal structures and single ownership of agentic AI.** Organisations should establish clear internal routes for agentic AI deployment decisions, with defined points of engagement for legal, compliance, risk, security and delivery teams. This is important because agentic deployments often cut across multiple business functions and can stall where ownership is fragmented or approval pathways are unclear. A single ownership point can help align first-line capability teams with second-line oversight functions while still allowing proportionate review.

Workforce preparedness: skills and capabilities needed and the future of work

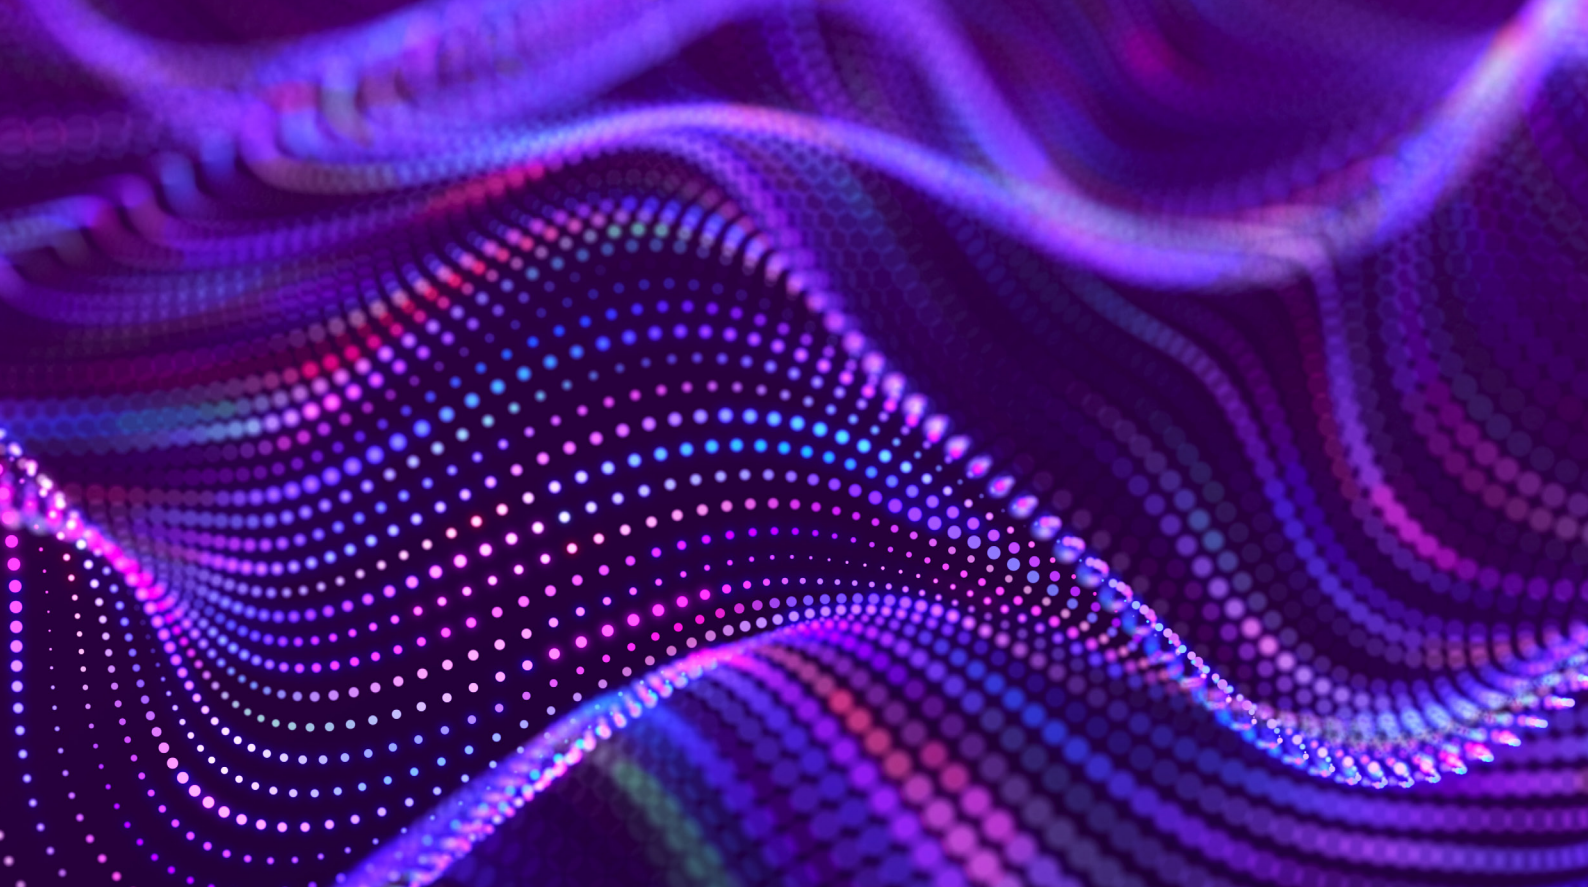
- **Build AI literacy and fluency across the workforce.** Organisations should not treat AI capability as relevant only to technical teams. A broader base of AI literacy and fluency will be needed so that workers across functions understand what agentic systems can and cannot do, can interrogate outputs appropriately and know when to intervene rather than relying on systems uncritically.

- **Assess workforce impacts as deployment scales, including on graduate and entry-level roles.** Organisations should consider using workforce impact assessments where they add practical value, particularly when agentic systems may affect graduate and entry-level work, skills pipelines or the future availability of mid-level capability. This can help organisations make more informed deployment decisions and avoid creating longer-term capability gaps through short-term automation choices.

- **Invest in workforce readiness alongside technical deployment, including managing agentic AI fleets.** Deploying agentic systems without preparing the workforce to validate, challenge and govern outputs introduces avoidable risk. Organisations should equip staff not only to use these systems, but to supervise, coordinate and intervene in multi-agent workflows in practice, including knowing when to step in, how to manage handoffs across agents and how to respond when systems behave unexpectedly. They should also consider the medium-term implications of current deployment choices for skills pipelines, junior talent development and knowledge retention, particularly where agentic systems are automating work that has historically played a developmental role.

Governance

- **Treat governance, assurance, safety and meaningful human oversight as enablers of scale.** As discussed in the governance section above, organisations are more likely to scale agentic AI confidently where they have clear red lines, defined escalation routes, risk mitigation efforts and



agreed conditions for when agents must seek human authorisation. Meaningful human oversight should distinguish between the different functions oversight is used to describe. For example, monitoring what an agent does, setting strict design-stage permissions and tracking systemic behaviour. Strict permissions set at the design stage can often be an most efficient first line of defence, since an action an agent cannot take is not one that needs to be watched for, but these built-in limits are a form of prevention rather than a substitute for oversight. Meaningful human oversight and ongoing monitoring remain essential for the risks that design limits cannot fully rule out. However, oversight should also be designed to work at scale. Organisations should use technical safeguards such as guardrails, circuit breakers, anomaly detection, escalation thresholds and audit mechanisms to focus human attention on exceptions, higher-risk actions and system-level performance rather than every individual output. In this context, governance and assurance should not be seen as something added after deployment but as an iterative part of the operating model that makes broader rollout possible from end-to-end. Overall governance should be proportionate and dependent on the use case in question. For example, controls applied to an agent that drafts email summaries may not need to match those applied to one that executes financial transactions.

- **Test shutdown, fallback and override mechanisms regularly.** These mechanisms can act as an important safety backstop that complements meaningful human oversight, rather than replacing it. For example, in a situation where a failure unfolds faster than a person is able to respond, systems are able to contain them automatically without waiting for human intervention, based on thresholds set and reviewed regularly by organisations. However, it is not enough for override mechanisms to simply exist in principle. Organisations should test them regularly in practice and under real conditions, to confirm that intervention remains operationally realistic and users remain confident about when and how to step in. This should be treated as part of standard governance and security practice rather than as an emergency-only control.

- **Use existing AI governance standards as a foundation, but complement them with more deployment-specific controls.** Frameworks such as ISO/IEC 42001 and the NIST AI Risk Management Framework can provide a useful foundation for AI governance and risk management. They can help organisations build more consistent internal processes. However, organisations deploying agentic AI should recognise that these frameworks will need to be backed up with more specific controls for observability, multi-agent coordination, escalation and oversight.

Government actions to support agentic AI adoption

Responsible scaling will depend not only on organisational readiness and action, but also on practical guidance, regulatory clarity and policy support. Workshop participants identified issues and areas where action from government and regulators could help scale adoption. These actions are outlined below.

These recommendations are intended to support the practical and responsible scaling of agentic AI, not to create a wholly new regulatory regime. techUK is not calling for new regulatory regimes, sector-specific mandatory testing requirements for agentic AI, or central government definitions of concepts such as fairness or transparency for agentic systems. Nor are we calling for new public-sector programmes or new institutions.

Frameworks and guidance

- **Future-proof existing work on AI accountability for multi-agent contexts:** Questions around accountability, particularly where more autonomous systems operate across longer chains of decision-making or involve multiple actors, were identified as a barrier to confident deployment. techUK believes that existing government work on accountability remains relevant and provides a sound foundation for organisations adopting and deploying AI. However, as deployment evolves from single-model systems toward more autonomous and multi-agent contexts, new questions are likely to arise. For example, key questions remain around liability and the extent to which existing accountability principles and practices remain applicable. These areas may warrant closer examination by government. However, given the pace at which AI technologies, are advancing, it is essential that any work on accountability and liability is technology neutral so that it can be built to last. This calls for an agile, flexible approach by government to work in this area that prioritises ongoing industry engagement and focuses on the issues and questions being raised by how all AI technologies, not just agentic AI, are being developed, deployed and used.

- **Provide clearer cross-regulator guidance on agentic AI deployment.** The DRCF's Future of Agentic AI paper and the CMA's recent guidance on AI agents in consumer-facing contexts provide a useful starting point, but the next step is to translate that work into more practical cross-regulator coordination and clearer best practice for deployers. techUK would like to see guidance that sets out a more consistent approach to accountability across regulators, shared

expectations for autonomous and multi-agent operation, and a clear route for deployers to seek joined-up advice where issues fall across more than one regulator's remit. This should build on regulators' existing work to create practical guidance that helps organisations with evaluation, documentation, logging and testing in real deployments.

- **Provide more specific operational guidance from individual regulators in their own remits.** Alongside joined-up cross-regulator work, organisations would benefit from clearer sector specific guidance from a wider range of UK regulators. For example, similar guidance from individual sector regulators that mirror what has been seen from others such as guidance from ICO, FCA and CMA. More specific guidance of this kind covering other and industries would support the wider adoption of Agentic AI in other sectors.

Standards, assurance and risk management

- **Support a more mature assurance and standards ecosystem for agentic AI deployment.** Workshop participants saw value in standards such as ISO/IEC 42001, ISO/IEC 27001 and the NIST AI Risk Management Framework as reference points, but also stressed that the ecosystem around them remains immature for agentic deployment, particularly when it comes to auditing, certification, product-level assurance and practical usability for SMEs. Rather than relying only on



future international updates, government should work with UK standards and assurance bodies to help develop more usable reference frameworks, audit approaches and supporting guidance that make responsible adoption easier in practice and help the UK lead in building a more mature market for AI assurance.

AI assurance

- **Consider using the AI Assurance Innovation Fund, established through the Trusted Third Party AI Assurance Roadmap, to support responsible agentic AI deployment and the growth of assurance applications for agentic systems.** The bottlenecks organisations face when deploying agentic systems include access to assurance capacity, understanding necessary professional skills for assurers, shared methods for evaluating more autonomous systems, clearer expectations about evidence, and routes for smaller firms to obtain credible external validation. For agentic AI in particular, this matters because organisations may need independent third-party evidence that controls are working across multi-step or multi-agent workflows that are difficult to supervise manually. Directing additional funding for the AI Assurance Innovation Fund towards developing novel assurance metrics and methodologies to address these four bottlenecks would continue to support the broader assurance ecosystem while tackling the specific risks posed by agentic systems.

Cyber security

- **Make security testing insights more usable for organisations deploying agentic systems in production.** The AI Security Institute's work on evaluation, testing and AI security is valuable, but it is primarily pitched at frontier model evaluation rather than the needs of enterprise deployment teams. Government should therefore consider how selected methods and findings can be translated into more practical guidance for organisations managing prompt injection, tool misuse, agent-to-agent manipulation and cascading failures in production environments.

Skills and the future of the workforce

- **Make agentic AI deployment a defined research priority for the AI Economics Institute.** Individual organisational decisions to automate entry-level or developmental work may be rational in isolation but could create wider skills risks over time if they weaken the pipeline of future mid-level capability. Government should therefore ensure that the AI Economics Institute's programme of work includes the medium-term workforce effects of agentic AI deployment, including impacts on entry-level work, progression routes and future skills formation.

Conclusion

Agentic AI is not a distant idea. Adoption remains gradual, but deployments are underway and early use cases are already demonstrating measurable value through pilot results across leading sectors. Organisations that invest now in the technical, organisational, workforce and governance foundations will be better positioned as the technology continues to mature. The question workshop participants returned to across all four sessions was not whether agentic AI will be adopted at scale but how and whether the conditions for responsible scaling are in place before deployment outpaces the frameworks needed to govern it.

techUK's view, informed by the workshops, is that the UK is well placed to lead on responsible deployment of agentic AI but that this position is not guaranteed. It will require coordinated action across the four areas surfaced in this brief: technical readiness, organisational design, workforce preparedness and governance. It will require continuous dialogue between organisations across sectors adopting AI, government and regulators grounded in practical experience and responsive to emerging evidence. techUK will continue to develop this work throughout 2026 engaging members and government on the practical and policy dimensions of responsible agentic AI adoption. We welcome input from members, regulators and policy makers as this conversation evolves.



References

1. techUK (January 2026), Scaling Responsible Adoption of Agentic AI - Workshop Series. https://www.techuk.org/resource/scaling-responsible-adoption-of-agentic-ai.html. Workshop 1: Industry best practices and routes for scaling agentic systems (18 February 2026). Workshop 2: Organisational barriers and regulatory levers affecting agentic AI adoption (2 March 2026). Workshop 3: Workforce preparedness, agentic AI and the future of work (18 March 2026). Workshop 4: Ethics, governance and accountability in agentic systems (25 March 2026).

2. DSIT (January 2026), AI Adoption Research. Reports that AI adoption across UK businesses remains modest, with around 16% of UK businesses currently using any AI technology. Of those, only 7% have adopted agentic AI — making it the least-adopted AI technology category, which the report attributes to its relative newness compared with other technologies. The same research finds that businesses face the most significant implementation barriers with agentic AI (32%) compared with any other AI category. https://www.gov.uk/government/publications/ai-adoption-research

3. MuleSoft (2026), 2026 Connectivity Benchmark Report. Reports that 69% of UK organisations say most or all teams have adopted AI agents, are running an average of 13 agents, and expect that number to double within two years. https://www.mulesoft.com/lp/reports/connectivity-benchmark

4. techUK (24 June 2026), Agents of Change: Generative and Agentic AI in Financial Services 2026. Explores financial services use cases and the agentic AI landscape. https://www.techuk.org/resource/techuk-report-agents-of-change-generative-and-agentic-ai-in-financial-services-2026.htm

5. Skills England (April 2026), AI leadership - AI strategy and opportunity - Apprenticeship unit. https://skillsengland.education.gov.uk/apprenticeship-units/AU0009

6. HMT and DSIT (17 March 2026), UK will win AI race as Chancellor sets out economic 'big choices', and AI Opportunities Action Plan 2026 progress updates. https://www.gov.uk/government/news/uk-will-win-ai-race-as-chancellor-sets-out-economic-big-choices and https://delivery.ai.gov.uk/19/

7. DSIT (3 September 2025), Trusted third-party AI assurance roadmap. Policy paper setting out government ambitions for the UK third-party AI assurance market, including actions to support access to independent assurance, audit and validation. https://www.gov.uk/government/publications/trusted-third-party-ai-assurance-roadmap/trusted-third-party-ai-assurance-roadmap

8. ICO (January 2026), ICO Tech Futures: Agentic AI. Sets out the ICO's early thinking on the opportunities and data protection implications of agentic AI, including accountability, risk management and privacy by design. https://ico.org.uk/about-the-ico/research-reports-impact-and-evaluation/research-and-reports/technology-and-innovation/tech-horizons-and-ico-tech-futures/ico-tech-futures-agentic-ai/

9. CMA (9 March 2026), Complying with consumer law when using AI agents. Clarifies how consumer protection law applies to businesses deploying AI agents and emphasises that businesses remain responsible for agent behaviour. https://www.gov.uk/government/publications/complying-with-consumer-law-when-using-ai-agents

10. DRCF (31 March 2026), The Future of Agentic AI. Foresight paper exploring agentic AI and how UK regulatory frameworks can help realise the opportunities of this technology in a responsible and safe way. https://www.drcf.org.uk/publications/papers/thefutureofagenticai

11. DRCF (10 October 2025; 10 April 2026), DRCF launches Thematic Innovation Hub for innovators with first focus on agentic AI and Key Takeaways from the DRCF Thematic Innovation Hub first series: Agentic AI. https://www.drcf.org.uk/news-and-events/news/drcf-launches-thematic-innovation-hub-for-innovators-with-first-focus-on-agentic-ai and the DRCF news page

techUK
FOR WHAT COMES NEXT



linkedin.com/company/techuk



youtube.com/user/techUKViews



@techuk.bsky.social



info@techuk.org

Image credit | iStock by Getty Images