INSIGHT

Good practice guide for organisations using AI

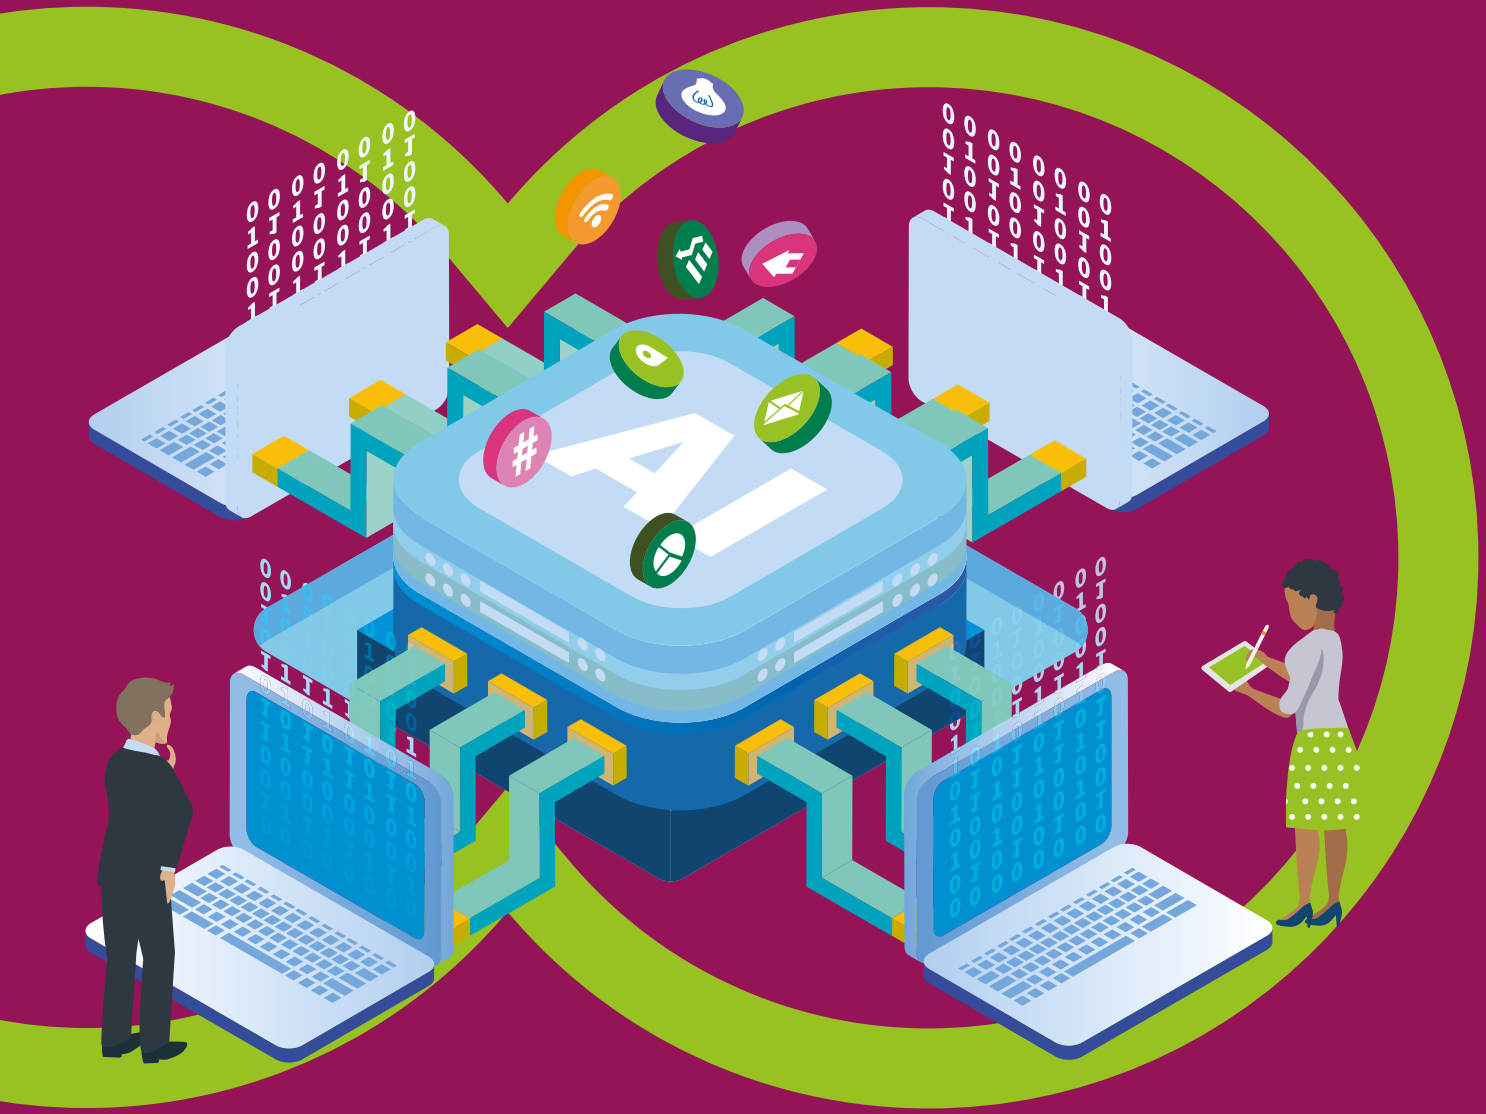

National Audit Office

Good practice guidance
May 2026

This guide highlights key considerations for audit and risk assurance committees when overseeing the planning, deployment and scaling of artificial intelligence within public sector organisations.

We are the UK's independent public spending watchdog

DP Ref: 018501

Insights

Our insights products provide valuable and practical insights on how public services can be improved. We draw these from our extensive work focused on the issues that are a priority for government, where we observe both innovations and recurring issues. Our good practice guides make it easier for others to understand and apply the lessons from our work.

We are the UK's independent public spending watchdog. We support Parliament in holding government to account and we help improve public services through our high-quality audits.

The National Audit Office (NAO) scrutinises public spending for Parliament and is independent of government and the civil service. We help Parliament hold government to account and we use our insights to help people who manage and govern public bodies improve public services. The Comptroller and Auditor General (C&AG), Gareth Davies, is an Officer of the House of Commons and leads the NAO. We audit the financial accounts of departments and other public bodies. We also examine and report on the value for money of how public money has been spent. In 2024, the NAO's work led to a positive financial impact through reduced costs, improved service delivery, or other benefits to citizens, of £5.3 billion. This represents around £53 for every pound of our net expenditure.

Contents

Aim of this guidance

This guide highlights key considerations for audit and risk assurance committees when overseeing the planning, deployment and scaling of artificial intelligence (AI) within public sector organisations.

It draws on NAO findings, the UK Government's AI Playbook, and lessons from digital transformation programmes across government.

This guidance includes:

- where AI is used in government;
- areas that organisations need to consider; and
- areas of focus and suggested questions to ask.

This guide will evolve as AI itself continues to evolve.

Introduction: Context and responsibilities

Context for government use of artificial intelligence

The government is keen to increase the use of artificial intelligence (AI) beyond optional experimentation to more deliberate, well-managed adoption in departments and other public sector bodies. It sees AI as a means for improving public services, productivity and economic growth, while also helping the government cope with rising demand and capacity constraints.

AI is a broad term for the ability of computer systems to perform tasks that normally require human intelligence, such as understanding language, learning from data and making decisions. The use of AI can deliver huge benefits to organisations. It can increase efficiency, automate routine tasks, analyse large volumes of data quickly and improve service quality by identifying patterns and insights to support faster and more informed decision-making.

However, AI also comes with significant risks and challenges. These include protecting personal data, ensuring decisions are fair and explainable, and avoiding inaccuracy, bias or discrimination. Poorly designed or managed AI systems can increase workloads, cause delays, undermine public trust and lead to harm to citizens if used irresponsibly.

The government has produced an Artificial Intelligence Playbook. Its aim is to provide practical guidance to help public sector organisations use artificial intelligence safely, responsibly and effectively and in a way that delivers public value while managing risk and maintaining public trust. This will be a valuable resource for those responsible for implementation.

The key questions for audit and risk assurance committees to consider are whether the organisation is clear on why it is using AI, what risks it must manage and how it ensures responsible adoption.

Central leadership

The **Department for Science, Innovation and Technology (DSIT)** provides AI policy leadership and strategy. It supports AI in the wider economy by setting national AI strategy, funding research and innovation, with the aim of ensuring AI growth is safe, trusted and competitive.

The **Government Digital Service (GDS)**, as part of DSIT, publishes cross-government frameworks and guidance. These include:

- the AI Playbook for the UK Government that sets the principles for lawful, ethical, secure and responsible AI use across departments, and
- the Algorithmic Transparency Reporting Standard for documenting and publishing algorithmic tools used in decision-making processes, which is mandatory for central government and in-scope arm's length bodies.

Departmental responsibilities

Individual departments and government bodies are responsible for developing their own AI strategy and solutions and applying any guidance locally, including security, data protection and assurance within their own organisations.

Where AI is used in government

Use of AI

AI is being deployed throughout the public sector to help improve productivity, decision-making and service delivery. Typical internal applications of AI range from the automation of routine tasks to advanced analytics, either integrated into existing workflows or as standalone tools.

Figure 1 sets out examples where we have seen AI being used or contemplated in government. It is important for committees to understand the characteristics of the specific type of AI being proposed in their organisations in order to evaluate, manage and mitigate its associated risks.

Figure 1
Examples of AI in a government context

Type of AI	What it does	Examples	Uses personal data?
Fraud and error detection	Identifies potential anomalies within datasets	Analysing claims or invoices to flag potential fraud	Yes
Imaging	Identifies and classifies physical objects	Medical imaging to assist diagnosis	Possibly
		Analysing CCTV images to monitor road traffic	
		Facial recognition in passport e-gates	
Document processing	Extracts information from formatted documents	Extracting information from an application form and copying it into case management systems	Yes
		Comparing documents to identify discrepancies	
Managing operations	Monitoring or automating business processes	Automating routine checks as part of an application process (e.g. passport, driving licences)	Yes
		Assessing cases to determine whether further checks are required (e.g. health, taxation, social security)	
		Triaging correspondence received	
Research and monitoring	Seeing trends/patterns in data and predicting outcomes over time	Monitoring markets to identify trends and issues that could lead to consumer harm (e.g. regulators)	Possibly
		Identifying taxation fraud and illegal practice	
Text generation	Produces written statements and documents	Analysing large volumes of text to summarise key points (e.g. ChatGPT)	Possibly
		Transcribing and summarising of meetings	
Virtual assistants	Provides information based on user prompts	Drafting emails, presentations, documents (e.g. Copilot)	Possibly
Coding assistants	Writing computer programming code	Translating obsolete code into supported versions as when addressing legacy systems (e.g. GitHub)	No
		Coding tools to improve developer productivity (e.g. Claude Code)	

Source: NAO

Areas that organisations need to consider

AI is increasingly present across government bodies, whether developed to meet a specific need, embedded in productivity tools, offered as cloud services or built into third-party systems.

Embedded AI

Modern office productivity suites, such as Microsoft 365 and Google Workspace, include embedded AI tools by default (Copilot and Gemini respectively). These support everyday tasks such as drafting emails and documents, summarising reports and meetings, answering questions, analysing data and generating ideas or code. AI is now also routinely embedded in other mainstream software products as background functionality rather than as a separate feature, enhancing how systems work without users always noticing.

In UK central government, Copilot is the main approved AI tool, having been formally trialled and rolled out across many departments as a secure, enterprise service. ChatGPT Enterprise is approved in a small number of departments under strict criteria. Individual departments and other public bodies remain responsible for security, data protection and assurance decisions within their own organisations.

Cloud services

Infrastructure cloud service providers offer ready-to-use AI features that organisations can add to their software, such as tools that understand text and speech, recognise images or make predictions from data. They also supply the powerful computing needed to run AI and handle tasks like scaling, security and updates. This means organisations can use AI to improve services and decisions without having to build or manage complex technology themselves.

Using pre-trained models available as cloud services differs fundamentally from developing bespoke models. It is important to recognise that these require different specialist skills. Some organisations may wish to consider an intermediate option of tailoring pre-trained models to their own needs without training a model from scratch, for example using approaches like retrieval-augmented generation to draw from a trusted set of internal data. This still requires specialist skills.

Continued
Areas that organisations need to consider

Productivity and efficiency gains

External evidence suggests that promises of greater productivity translating into substantial organisational financial savings are not yet borne out by real world experience, either in government or in commercial organisations. AI can make individual activities faster. However, speeding up discrete steps and activities in isolation does not automatically improve overall throughput, especially when work is subject to approvals or governance processes, or human judgment. Saving a few minutes across many tasks rarely frees up a meaningful block of productive time during the working day. Scattered gains, even though they feel helpful at an individual level, do not readily extrapolate cleanly into large, organisation-wide savings.

External impacts of AI on an organisation

External impacts of AI on an organisation are becoming a significant and growing risk area that audit and risk committees need to be aware of. Organisations may be disrupted, degraded or see significant increases in workload arising from AI-driven external activity (such as automated tools generating surges in online submissions or sharp increases in traffic to public-facing websites). This could manifest itself in the following ways:

- an increase in legitimate demand as barriers and friction in accessing public services are reduced,
- more repeat, low-value or poor-quality submissions, which can overwhelm existing capacity, and
- an elevated risk of fraud, including more sophisticated and harder-to-detect attempts.

External uses of AI may also result in targeted cyber-attacks or attempts to extract sensitive data, as it lowers the barriers and saves time and effort for criminals. Audit committees should ensure their organisations can anticipate, monitor and mitigate these risks if they are to provide effective oversight of the full impact of AI.

Areas of focus and questions to ask

Delivering the transformational benefits of AI in the public sector depends on establishing the foundational infrastructure and digital enablers. These can be constrained by the issues we have highlighted in our previous work, such as outdated legacy systems and fragmented, poor-quality data. These issues can make getting benefit from AI difficult and costly. Audit committees should focus on whether an AI strategy exists and that appropriate foundations are in place for successful implementation.

The evidence informing this guide reflects a rapidly evolving AI landscape and variable maturity across departments. Observations may not apply uniformly across all bodies, and examples illustrate themes rather than represent comprehensive coverage. Nevertheless, organisations that are contemplating deploying or extending their use of AI should be aware of the factors set out below.

Innovation

There are different types of innovation, and how organisations approach AI reflects their thinking about innovation more broadly.

One approach is to reimagine what could be possible, using 'blue-sky' thinking to explore what AI might enable and to experiment with new ideas. Another approach starts with today's most pressing business problems and asks how AI could help to solve them. Both approaches are valid but require different mindsets.

Blue-sky innovation calls for a different way of thinking about value. Traditional value for money assessments are often too narrow to judge innovation properly. Investing in innovation is fundamentally different from buying established technology. Innovation involves longer time horizons, higher uncertainty and a higher risk of failure, but the potential rewards are much greater.

Applying the same value-for-money lens to innovation and experimentation as used for mature hardware, software and business-as-usual services is a major flaw, and it often sets both the public and private sectors up to fail. To innovate successfully, there is a need for senior leaders to have digital fluency, data fluency and AI fluency. All three are important, in that order. Addressing digital fluency helps with addressing data fluency, namely understanding how to connect, link and govern data to support the government's digital ambitions.

On using AI to solve business problems, organisations must identify and understand the business need before they determine the best solution for the problem. Without careful consideration at the outset of the complexities and interdependencies involved, the risk of programme failure increases.

Questions that audit and risk assurance committees could ask

- Is there a clear, well-articulated opportunity?
- Does it align with departmental objectives and represent value for money?
- Is there a technically feasible solution available within today's AI technology?
- Does enough accessible data exist, and is it of sufficient completeness and quality?
- Is there an implementable outcome that aligns with ethical requirements?

Continued
Areas of focus and questions to ask

AI strategy

While AI has existed since the 1950s, its use across organisations represents relatively new thinking.

Leadership capability for enterprise-wide AI is even less developed and, in many cases, is not yet clearly established.

A strong digital and AI strategy should be clearly articulated, business-led and aligned with other organisational strategies. It needs a realistic level of ambition backed by visible leadership support and clear ownership. The strategy should set out clear goals focused on the value that AI can deliver, such as improving services or reducing costs. It should be clear on how to achieve these goals, whether by undertaking research and experimentation, piloting solutions or scaling up successful pilots. It should also define measurable objectives with a credible roadmap, underpinned by committed funding and resources aligned to the wider business agenda.

In complex cross-government transformation programmes of the kind required for AI adoption at scale, a strategy with clear accountabilities and supporting governance arrangements is essential for success.

In our 2024 report *Use of artificial intelligence in government*, we found that departments were at an early stage in developing their own AI strategies and supporting governance arrangements.

Questions that audit and risk assurance committees could ask

- Is there a clear strategic use for AI in the organisation that aligns with departmental priorities, service outcomes and wider government strategy, following technology-led experimentation?
- Does it distinguish between AI for innovation and solving existing business problems within the organisation's current priorities?
- Is there a senior leader on the organisation's executive committee or equivalent who understands the wider business priorities and how technology can be harnessed?
- Is there proper coordination and monitoring to ensure that AI development is being shaped in a safe, secure and organisationally coherent way?

Leadership and skills

Leadership capability and digital fluency at senior levels are essential to exploit the opportunities for AI.

While specialists may understand AI and digital risks, strategic decisions about funding, scope and procurement are typically made by non-specialist leaders and this will require a rethink of the culture of the organisation. A lack of alignment between the existing organisational ways of working and the new AI requirements can hinder adoption.

Our reports and insights guides have highlighted persistent shortages of specialist skills. Acquiring new AI skills is a further challenge. Organisations should gain the skills needed to use, design, build and maintain AI solutions, bearing in mind that developing bespoke AI solutions and training in-house models require different specialist skills from using pre-trained models (such as ChatGPT). There is also an asymmetry between the providers of technology (who are often US-based 'big tech' corporations) and government digital leaders. It is important for the government to have someone with the skills to engage effectively with suppliers and cut through the 'hype'.

Questions that audit and risk assurance committees could ask

- Do senior leaders have sufficient understanding of AI to make informed decisions?
- Do they ask the right questions to provide effective, well-informed challenge that balances innovation with risk?
- Is there clear accountability for safe, ethical and effective AI use, including where AI is embedded in third party products and services?
- Is there adequate in-house capability and a plan to build and retain critical skills, rather than treating AI as a purely technical or outsourced issue?

Continued
Areas of focus and questions to ask

Data

Data quality, accessibility and governance are foundational risks.

Government bodies often hold large volumes of data, but these are frequently incomplete, inconsistent, siloed and difficult to share, with manual workarounds still common. Large volumes of good-quality data are important to train, test and deploy AI models. Weaknesses directly affect model performance, bias, explainability and reliability.

Data governance, metadata and remediation of legacy data issues are as important as the data itself, and overestimating data readiness is a recurring cause of failure.

If data sovereignty is a concern, organisations should also clarify where and by whom it will be processed.

Questions that audit and risk assurance committees could ask

- Does management understand the condition of underlying data, how it was collected and where it is stored?
- Is there assurance over data provenance, quality, licensing and the legal basis for use (for example, if legislation restricts this to a specific purpose)?
- Is data assessed for fitness for purpose before being used to train or deploy AI?

Security

AI can increase exposure to operational and security risks, particularly where ageing legacy systems, complex integrations and cloud-based services coexist and systems become more interconnected and complex.

For AI, this amplifies concerns around data breaches, model manipulation, supply chain risk and resilience of critical services.

'Secure by Design' is a set of ten mandatory principles for central government that requires security to be built into new digital services from the outset, with clear relevance for developing AI solutions.

Questions that audit and risk assurance committees could ask

- Is security by design treated as an integral part of AI development rather than a technical afterthought?
- Is accountability clear for protecting sensitive data used by AI?
- Have AI-specific threats been assessed?
- Are interfaces between legacy and modern systems secured?
- Is there active management of new risks introduced by cloud and third-party providers, such as reduced control over data and supply chain exposure?

Continued
Areas of focus and questions to ask

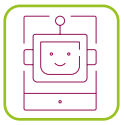

Pilots

Pilots and technology-led experimentation can be valuable for showing the art of the possible.

By testing ideas on a small scale, pilots allow teams to explore potential uses, benefits and limitations in a controlled way, building practical insight before wider adoption. They can also help challenge assumptions, surface risks early and inform better decisions about whether, how and where AI could be deployed at scale.

Best practice for piloting AI is to start with small, clearly defined, low-risk use cases, with strong senior ownership and multidisciplinary oversight, and to treat pilots as learning exercises rather than routes to rapid automation. Benefits include building capability, testing value at low cost and improving services or decision support before scaling. The main pitfalls include poor data quality, unmanaged bias or errors, and pilots drifting into live use without proper controls.

Government guidance encourages departments to build evaluation into their activities so they can learn from failure and drive continuous improvement. In piloting AI solutions for the public sector, the government should expect failures, and mechanisms are needed to ensure that the insights from pilots are disseminated and used to improve future pilots.

Questions that audit and risk assurance committees could ask

- Is there organisation-wide visibility of what type of AI is being piloted and why?
- Is the pilot managing risks around data quality, security, bias and accountability?
- Is there a clear understanding of how decisions will be made to stop, scale or redesign, based on evidence?
- Is there clarity on how success criteria are scored and measured?

Scaling

Once an AI pilot successfully concludes, the shift from pilot to production needs careful consideration.

Scaling new, leading-edge technologies such as generative AI presents significant, multifaceted risks and often reveals deep-seated technical issues. These risks can extend beyond traditional business challenges.

Scale and complexity increase risk exponentially. Large, multi-year digital programmes have a much higher scope for error, and this is especially relevant for enterprise AI deployments that cut across multiple services and datasets.

Complexity and integration are increased, as typically new technologies (particularly AI) may not work well with existing legacy systems. This can cause bottlenecks, disruptions and unexpected compatibility problems. Testing AI solutions at scale must be extremely rigorous and demands time and resource.

Questions that audit and risk assurance committees could ask

- Are ambitions for AI at scale realistic, given the starting point, legacy and data constraints and organisational maturity?
- Are the risks of scaling AI understood and managed at the organisational level, not just within a pilot?
- Is there sufficient assurance over the ability of existing systems, data and business processes to integrate AI at scale?
- Are the impacts on surrounding work processes that may not be directly involved in the pilot also considered?
- Is there a robust, planned and resourced testing regime?

Continued
Areas of focus and questions to ask

Guardrails and guidelines

Guardrails are the rules, controls and safeguards put in place to ensure AI systems are used safely and ethically.

Guardrails are needed because AI can produce harmful, biased or inaccurate results, expose sensitive data or be used in ways that create ethical, reputational and legal risks. Guardrails can include:

- policies on acceptable use,
- data protection controls,
- bias testing,
- human oversight of automated decisions, and
- clear accountability for AI outcomes.

Disciplined governance, early scrutiny and adherence to agreed standards are essential to avoid optimism bias and ensure that digital and AI initiatives deliver sustainable value rather than new long-term risks.

Questions that audit and risk assurance committees could ask

- Are there documented principles on risk appetite, acceptable use, ethics, transparency and traceability?
- Are there measures in place to prevent harmful outputs (bias, misinformation, toxicity) and ensure that AI systems remain safe and effective over time?
- Do senior leaders actively assure themselves that data protection, ethics and legal obligations are being met in practice?

Workforce and culture

AI adoption plans that are ambitious enough to deliver transformational benefits will require more than new technology alone.

They will also demand significant changes to business processes and the workforce that supports them.

AI is likely to reshape organisations by introducing new skills requirements, redesigning roles and increasing the need for digital and AI fluency at all levels. Workforce planning should anticipate changes to job design, including potential risks to entry-level learning, early career development and the need for substantial upskilling across the organisation.

As routine and rules-based tasks are increasingly automated and redirect more routine tasks away from human staff, demand will shift towards roles that emphasise judgement, oversight and the ability to deal with the more complex cases. The organisation will also need new specialist capabilities such as data science, model assurance, cyber security and AI ethics.

At the same time, a much broader group of staff will need basic AI literacy to use AI-enabled tools safely and effectively in their day-to-day work. Organisations will therefore need to invest in reskilling and continuous learning, update professional standards and job roles, and strengthen leadership capability to manage AI-enabled teams. Workforce planning will also need to address emerging risks, including overreliance on automated systems, loss of institutional knowledge and increased dependence on external suppliers for scarce technical skills.

In our 2024 report, we said that the implications for the overall composition of the workforce and the skills required are not yet being considered in detail. Audit committees should therefore ascertain whether the organisation is addressing this as part of its workforce planning.

Questions that audit and risk assurance committees could ask

- Is the organisation thinking about how roles will be redesigned in the short, medium and longer term?
- Are there credible plans for upskilling staff from basic AI literacy to specialist skills?
- Are risks to entry-level learning and long-term capability being managed?
- Are the risks of overreliance on automation and loss of institutional knowledge and ‘corporate memory’ being actively identified and mitigated?
- Is the organisation monitoring and assessing the risk that unforeseen external uses of AI may dramatically increase the workload in unexpected ways and require more resources?