The new AI operating model for government

A study with UK permanent secretaries and AI leaders



OCTOBER 2026

LED BY LORD GUS O'DONNELL

GLOBAL government FORUM

Knowledge Partners

cognizant | aws

Contents

Foreword: Lord Gus O'Donnell 4

Knowledge Partner Foreword 6

Introduction: Beyond the AI foothills – making the choices that matter 8

Participants 10

Methodology 11

1. Set the ambition for AI and demonstrate the impact 12

Efficiency and productivity remain a significant opportunity 12

Bigger opportunities require deliberate choices 12

Refine ambition with priorities 13

Experiment to discover what's possible 14

Measure what matters 14

Implications for a government AI operating model 15

In action: Towards a new operating model for AI 16

2. Establish a more joined-up approach to AI 17

Continue to bolster technical foundations, but keep moving on AI 17

Move beyond ad hoc sharing 18

Increase visibility of AI efforts, successes and capabilities 19

Pool demand for common capabilities 20

Common does not always mean one size fits all 20

Make the centre enabling – but not passive 21

Implications for a government AI operating model 23

In action: Towards a new operating model for AI 24

3. Prepare for how AI will change work **25**

Continue to build technical expertise and AI capability across the civil service 25

Training is not enough 26

Understand how roles will change 26

Career paths will also need to change 27

How AI could reshape the workforce 27

Plan for change and bring people with you 28

Implications for a government AI operating model 29

In action: Towards a new operating model for AI 30

4. Use trust, not hype, to set the pace for AI **31**

The need to balance pace and responsibility 31

Be clear and intentional about risk 32

Transparency is becoming increasingly important 32

Clarify accountability as AI advances 33

Pace and trust need to work together 34

Implications for a government AI operating model 35

In action: Towards a new operating model for AI 36

Conclusion: The next phase of government AI **37**

Next steps **38**

Foreword: Lord Gus O'Donnell

Former cabinet secretary and head of the UK civil service

Our *Rewiring the State* report last year highlighted the need to explore further how AI can be most effectively used in government.

This follow-up report provides that insight. The interviews with permanent secretaries and other senior government figures have helped to identify the priority areas of focus for the next phase of government AI – and indeed the need for a new operating model to achieve them.

As usual, any investigation into government throws up some longstanding challenges. First, the need for clarity about desired outcomes and ways of measuring those outcomes. This led me to thinking about past occasions where we had new tools and were not sure how to use them – for example, the emergence of computers.

Nobel Prize-winning economist Robert Solow famously wrote in a 1987 *New York Times Book Review*: "You can see the computer age everywhere but in the productivity statistics."

This observation became known as the Solow paradox or the productivity paradox. It highlights the mismatch between heavy corporate investments in information technology and the lack of visible acceleration in macroeconomic productivity growth at the time.

Solow made the remark during a period when personal computers and mainframes were entering workplaces rapidly, yet national productivity growth remained sluggish.

Economists eventually realised that new technologies require time to create an impact. Productivity gains did not show up until the 1990s, after businesses redesigned their workflows, reorganised management and retrained workers to use the technology effectively.

Governments need to learn these lessons. It is harder for them because they are designed around an inherently inflexible, siloed system of secretaries of state being accountable for their departments. In business, there are separate divisions but corporate structures emphasise the success of the whole company, as does the market, which assesses the share price. To some extent the bond market also assesses the whole of government by considering how much interest it requires in order to finance state borrowing. To fully embrace the potential of AI, government will need to establish a new operating model that overrides the old silos, pools resources, connects data and, ultimately, delivers efficient, joined-up and bespoke services to the public.

Another problem is that traditional economic statistics like Gross Domestic Product struggle to capture qualitative improvements, such as consumer convenience. This is particularly true in the public sector where the efficiency savings are often in the form of time saved by the public.

This makes it harder to make the business case for departments to invest in AI. Treasuries around the world struggle with assessing so-called 'spend to save' initiatives. They know the spending is certain while the savings are usually much more uncertain.

This is particularly true of AI where savings may end up being gains for customers in the quality of services but with little direct revenue savings.

Having clear priorities is the first and crucial step to being able to measure the success and value of AI.

On implementing AI across government, this report highlights a theme that was emerging from the first study – namely, the need for clarity about what should be done by the centre and what should be devolved to departments. This has not been helped by a number of machinery of government changes but there is now an opportunity to put this right. With an AI champion in the Cabinet, a new set of arrangements that will hopefully be long-lasting, and a new cabinet secretary who sees AI deployment as a priority, solutions are already emerging.

The task now is to prioritise efforts, highlight early successes and learn from failures.

Inevitably there will be setbacks with new technology and it's important that everyone understands what the risk appetite is in their area. When it comes to the use of AI, there are particular problems for government as the level of public trust is worryingly low. This is evident, for example, in citizens' reluctance to have digital IDs in a world where they give away enormous amounts of personal data to the private sector, or in tolerating relatively high levels of human error in road accidents but not in driverless cars. A significant challenge for policymakers is understanding how to build public trust and confidence in the use of AI while recognising that the technology itself will not be infallible.

In this report we have identified some key priorities for the next phase of government AI: setting the ambition for AI and demonstrating impact; establishing a more joined-up approach to scale what works; preparing for how AI changes the way we work to ensure the civil service is fit for the future; and putting the right parameters in place to support both trust and pace. Together, these provide some of the building blocks for a new operating model for government – one that can make the most of AI across departmental boundaries while retaining clear accountability and the flexibility to respond to different needs.

There are undoubtedly risks and we live in a world where messages of doom – whether warranted or not – receive massive attention. At the same time, we should remember that a lot of progress has been made and the opportunities are huge.

This technology is here to stay and is evolving incredibly rapidly. Governments need to embrace it, understand it and recognise how to maximise the opportunities while managing the risks. I would like to thank every colleague who participated in this research. The insights they shared as current leaders grappling with these unprecedented considerations were invaluable in shaping our recommendations. I hope this report can help them, and others at the top of government, in this daunting endeavour.



Lord Gus O'Donnell
Former cabinet secretary and head
of the UK civil service

Knowledge Partner Foreword

Artificial intelligence has rapidly moved from possibility to priority across government. In just a few years, the conversation has shifted from whether public sector organisations should adopt AI to how they can deploy it safely, responsibly and at scale to deliver better outcomes for citizens.

That shift sits at the heart of this study, *The New AI Operating Model for Government*, which builds on last year's *Rewiring the State* report and captures the perspectives of 13 permanent secretaries and AI leaders on how government can move from experimentation to impact. We are particularly grateful to them for the time and insight they contributed to this research, and to Lord Gus O'Donnell for leading the study once again. Their perspectives give a candid view of how government is navigating this next phase of AI adoption.

As long-standing technology partners to UK government, Cognizant and AWS recognise both the opportunity and the responsibility this moment brings. Efficiency gains from early AI adoption are real and valuable, but as this report makes clear, they are not the destination. The bigger prize comes when leaders choose to rewire how a service works, not just speed it up, reorganising around what citizens need rather than departmental processes.

> "The opportunity now is to move with purpose, turning promise into lasting public value."

Realising that bigger prize requires a more joined-up approach to AI across government, with shared standards, reusable platforms and secure data practices that enable departments to innovate without duplicating effort, including clarity on where data resides and how it is protected, so that government retains control and public confidence as it adopts AI at scale. But technology foundations alone are not enough. It also means rethinking how the workforce evolves, not simply how many people need new skills, but how roles, professions and career paths will change as AI reshapes the work itself.

Above all, it requires trust. Public confidence will shape the pace at which AI can be adopted across government, making proportionate risk-taking, strong governance and genuine transparency essential to ensuring innovation earns and retains that trust.

These findings echo what we see in our own work with government. Real transformation depends on going beyond pilots and personal productivity tools to build AI into the operating models, skills and secure cloud-based infrastructure that provides the scalable foundation on which these capabilities are built and sustained. Matched with strong outcomes, leadership and genuine partnership between government and industry, this is what turns ambition into services that are faster to build, more reliable in operation and simpler for citizens to use.

This report offers an assessment of where government's use of AI stands today and a practical guide to the choices ahead. The opportunity now is to move with purpose, turning promise into lasting public value. Cognizant and AWS remain committed to supporting that journey and helping government translate AI's potential into better outcomes for citizens across the UK.



Yatin Mahandru
Vice president, head of public sector and health UKI, Cognizant



Ash Thankey
Managing director – UK, Germany and international organisations, Amazon Web Services

Introduction: Beyond the AI foothills – making the choices that matter

This report comes at a significant moment for the UK government. AI is becoming increasingly powerful and more widely used across government, with growing implications for how organisations operate, how public services are designed and delivered, and how people work.

Within the last year, the country has seen a new cabinet secretary and prime minister, while recent machinery of government changes have once again moved where responsibility for digital, technology and AI sits. This creates some uncertainty, but also an opportunity to take stock of progress and consider what needs to come next.

This study builds on ***Rewiring the State***, our 2025 report based on interviews with permanent secretaries about the opportunities and barriers to digital transformation in government. That report pinpointed four priority areas for action: elevating digital to the core of government leadership; building transformation capability at every level; unlocking the full power of data and AI; and driving joined-up government.

A year on, progress has been made in some of these areas, while AI, which this report focuses on specifically, makes some of the issues more pressing and raises new ones.

The focus now must be to move beyond encouraging adoption of AI to ensuring it delivers service transformation, value for money and better outcomes.

This is clearly a priority at the top of government. Cabinet secretary Antonia Romeo has described AI as "the defining opportunity of our generation, as well as the defining challenge", and made scaling its use across the civil service a key priority, arguing that it will fundamentally change how the civil service works, operates and delivers services.

Since entering Downing Street, prime minister Andy Burnham has also established an AI taskforce and elevated the AI minister to Cabinet level.

Realising the benefits of AI will mean addressing both longstanding and emerging challenges amid fast-moving technology and citizen expectations.

Discussions of digital transformation in government have historically caused handwringing about stubborn challenges: data, legacy technology, skills, hiring and budgets. While several permanent secretaries interviewed for this study described government as still being in the "foothills" of AI, we also noted a fresh optimism about the benefits AI could bring. Importantly, there was also greater confidence than we have seen in previous research that longstanding barriers can be tackled or managed, rather than having to be fully resolved before government can make progress with AI. We welcome this positive shift.

As one interviewee put it, the task is to move towards "establishing the tooling to scale AI rather than just playing with it".

If the first phase of government AI was about adoption, the next will be about choices. Throughout our interviews, permanent secretaries and other leaders discussed the need to develop a new "operating model" of government to capitalise on AI.

The elements that will make up this new operating model will be explored throughout this study, which is based on interviews with 13 permanent secretaries and AI leaders.

In particular, the leaders point to four areas that will shape the next phase of government AI:

1. **Ambition and impact:** Widespread adoption of AI can deliver significant efficiency and productivity gains, but better services and outcomes will require deliberate choices about where AI can make the greatest difference, alongside experimentation to discover what is possible and measurement to demonstrate impact. Clear, user-centred, cross-government priorities, supported by AI, could drive truly transformative changes in public service delivery. Ministers and senior officials should seize this opportunity to think big, think differently and reimagine the operating model for government.

2. **AI across government:** As AI activity grows, better visibility of what is happening across departments will be needed, alongside a more systematic approach to identifying common needs, sharing capabilities and scaling successful approaches. The challenge is to establish what should be done collectively while retaining freedom to innovate. Many of the leaders we interviewed invited more direction from the centre where a common approach would remove costly duplication but wanted greater flexibility where a bespoke approach is needed.

3. **Workforce:** Building AI skills remains essential, but leaders also need to think ahead about how AI will change tasks, roles and career paths, and what that means for workforce planning for the future. Conversations around workforce and AI have tended to focus on numbers, but there is a pressing need to think about the changing nature of professions in an AI-enabled world and what this means for recruitment and, for example, long standing views on core civil service skills and values.

4. **Trust:** Realising AI's benefits at pace will require clear, proportionate decisions about risk appetite and accountability, alongside transparency that builds public confidence in how the technology is used. As several of our participants stressed, the public will be far less forgiving of AI mistakes than they are of human errors.



Siobhan Benita
Former senior civil servant and advisor to Global Government Forum



Sarah Wray
Research editor, Global Government Forum

Participants

We thank the following interviewees:



Jonathan Brearley: Permanent secretary, Department for Energy Security and Net Zero



Amanda Brooks: Interim permanent secretary, Department for Business and Trade



Nick Dyer: Interim permanent under-secretary, Foreign, Commonwealth and Development Office



Jo Farrar: Permanent secretary, Ministry of Justice



Jerome Glass: Director general, Future Civil Service, Cabinet Office



Dan James: Advisor to the cabinet secretary on AI in the civil service



Samantha Jones: Permanent secretary, Department of Health and Social Care



Cat Little: Permanent secretary of the Cabinet Office/incoming permanent secretary of the Department for Business, Innovation, Science and Trade



John-Paul Marks: Permanent secretary and chief executive, His Majesty's Revenue and Customs



Christophe Prince: Director, data and technology policy and strategy, Home Office



Jo Shanmugalingam: Permanent secretary, Department for Transport



Kalbir Sohi: Government chief AI officer, Government Digital Service, Department for Digital, Culture, Media and Sport



Susannah Storey: Permanent secretary, Department for Digital, Culture, Media and Sport

Job titles and department names reflect those at the time of interview. Quotes from the study's interviews are presented in this report in an anonymous format. Images of participants throughout the pages of this report in no way indicate attribution of individuals to any of the quotes and are for illustration only.

Methodology

For this research, Global Government Forum (GGF) used a semi-structured interview format covering issues such as vision, skills and leadership, culture, technology, and collaboration.

The research was led by Lord Gus O'Donnell, former cabinet secretary and head of the UK civil service, and Siobhan Benita, a former UK civil servant and executive adviser to GGF.

The report was written by Sarah Wray, research editor at GGF.

Cognizant and AWS were GGF's knowledge partners for this project.

Cognizant is a global AI builder and technology services provider, building the bridge between AI investment and enterprise value by building full-stack AI solutions for clients. It is a digital and data partner to UK Government in central government, transport, health and defence, delivering citizen-facing services and outcomes since 2017. In addition, the company helps clients across a wide range of industries in the UK and Ireland to modernise technology, reimagine processes and transform experiences. Its deep industry, process and engineering expertise enables it to build an organisation's unique context into technology systems that amplify human potential, realise tangible returns and keep global enterprises ahead in a fast-changing world.

Since 2006, Amazon Web Services has been the world's most comprehensive and broadly adopted cloud. AWS has been continually expanding its services to support virtually any workload, and it now has more than 240 fully featured services. Millions of customers – including the fastest-growing startups, largest enterprises, and leading government agencies – trust AWS to power their infrastructure, become more agile, and lower costs.

Interviews were carried out virtually. They were conducted under the condition of anonymity of any material quoted to ensure participants could speak openly. This approach was supplemented with desk research.

Example projects outlined in the report use information available in the public domain. Web links referenced were accessed in September 2026.

1. Set the ambition for AI and demonstrate the impact

Global Government Forum's *Rewiring the State* report highlighted the importance of digital transformation being driven by the problem to be solved or the desired outcome, rather than technology. This matters particularly for AI, where expectations are high and technology-led adoption could add cost, complexity or risk without delivering clear benefits.

Interviews for this study suggest that widespread adoption can deliver valuable efficiency gains, but more fundamental changes to services and outcomes will require deliberate choices about where AI can make the greatest difference.

A new operating model will combine top-down vision and clarity about outcomes with experimentation that gradually reveals the art of the possible.

Efficiency and productivity remain a significant opportunity

The first wave of AI adoption has understandably focused on efficiency and productivity. Interviewees pointed to applications that can cut administrative work, speed up processes and help staff manage growing demand as well as supporting financial management.



> "AI is only another enabler; it's a tool. Real transformation is about behaviour change."

With public finances constrained and demand rising across many services, these gains remain important. Leaders pointed out that efficiency gains can be used in different ways, including to handle more work, improve quality or create capacity to improve services.

"It's really a leadership choice," one interviewee said.

Bigger opportunities require deliberate choices

Beyond this, interviewees argued that AI could also enable more fundamental shifts in services, behaviours and outcomes. Interviewees pointed to opportunities such as reducing economic inactivity, improving health outcomes and managing demand on stretched public services. They also envisaged changing how citizens interact with government, organised around people's circumstances and goals rather than departmental processes.

The current conversation in government remains too focused on "process efficiency", said one interviewee, noting that looking at how AI can help achieve the wider outcomes a department is aiming for has the potential to make its use "a lot richer and a lot more transformative".

As one interviewee put it: "AI is only another enabler; it's a tool. Real transformation is about behaviour change."

These bigger opportunities are unlikely to emerge simply because AI becomes more widely available. One interviewee described the current "glide path" as the government deploying personal productivity tools and automation, and using AI largely to optimise existing processes.

A more transformative path that changes how the government interacts with citizens and how services are designed, would, they argued, require not only the right "enabling environment" but also political and leadership appetite.

"I think it's absolutely a choice," they said.

The potential value can also extend well beyond the organisation implementing the technology. One interviewee argued that the benefits for citizens, the exchequer and the economy could "dwarf" the direct efficiency opportunity within the organisation itself. Another described opportunities across the wider system they oversee as "off the scales".

However, widespread adoption within the government is not separate from this more transformational ambition. In some cases, interviewees said, the government remains "significantly further back than it needs to be" in its own journey. Civil servants sometimes recognise the relevance of AI to the wider economy and public services without necessarily seeing its applicability to their own everyday work.



"If we don't get it right, we'll get in the way of it," one said. "We've got some way to go to get people to believe that it is worth investing the time in AI adoption."

To address this, interviewees highlighted the value of colleagues sharing practical examples of how they are already using AI in their work, helping others to see where it could be relevant to them.

Refine ambition with priorities

Interviewees were clear that the question of which AI opportunities to pursue should be guided by what the government and departments are ultimately trying to achieve. Articulating this effectively is likely to require a mix of top-down and bottom-up actions.

As one interviewee put it: "We probably need better shared outcomes in terms of what we are seeking to achieve because if we haven't got the shared outcomes, we will never get the shared underpinnings."

"It would be good if the government could articulate the 10 big behavioural shifts that we are going to try and leverage AI to achieve," one said. "And then take them together: what is that worth to the country?"

"Get each department to [outline] the AI-enabled opportunities to drive behaviour change that maybe previously weren't possible," they suggested. And ask: "What's the next

leap from there? It requires a bit more imagination of AI capability beyond the obvious process reengineering."

The cabinet secretary's new AI Innovation Award will also be an important way to demonstrate what good looks like and direct ambition.

Experiment to discover what's possible

That does not mean waiting for a perfectly defined set of outcomes – exploration is also key. One interviewee argued that government will have to "find our way there" by experimenting and testing different approaches. Without that, they warned, "we will always default to an adoption and productivity mindset".

"There is a huge opportunity… for the mechanics of public service to be quite different," they added. "But I don't think you can find anyone who can tell you what that different looks like. Therefore, I think we need to set ourselves up to find it."

Another said: "It's always better for all sorts of reasons that you have clarity of an outcome" but noted that "government is messy" with various stakeholder groups and trade-offs.

There is a risk of "wishing for a world that doesn't exist, rather than dealing with the world we have", they said. "The role of senior leadership is to make sense of that for your organisation."

This reinforces the finding from *Rewiring the State* that AI cannot be left to technologists – it must be sponsored and driven from the top of the organisation.



> **"There is a huge opportunity… for the mechanics of public service to be quite different. But I don't think you can find anyone who can tell you what that different looks like. We need to set ourselves up to find it."**

Mainstreaming AI across the organisation "is a particular personal focus of mine" said one permanent secretary. "I believe without my personal leadership it won't happen."

Measure what matters

Measurement is also part of how these choices are reinforced. Some benefits, such as reduced time or cost, are relatively easy to quantify, but others, such as improved service quality, reduced risk and long-term outcomes, are considerably harder to capture, and harder still when they cross departmental boundaries.

Measuring time saved is only one aspect. Evaluation also needs to establish what happens to the capacity released and whether it delivers meaningful benefits, such as improved citizen and staff experience, better service outcomes, reduced demand or financial savings. The intended benefits should be defined upfront so they can be measured.

"There's a monetary side, but there's also a quality side that we're trying to measure," one interviewee said.

Another went further, describing investment in impact evaluation or performance measurement as potentially "one of the biggest unlocks" for AI's potential. That also means understanding how existing processes perform to create a baseline against which to judge whether the technology is actually making it better.

The government's guidance on evaluating the impact of AI, published as part of the Magenta Book, emphasises early evaluation planning and assessment of both intended and unintended outcomes.



Cat Little: *Permanent secretary of the Cabinet Office/incoming permanent secretary of the Department for Business, Innovation, Science and Trade*

If the strongest incentives remain around benefits that can be easily expressed as time or departmental savings, more ambitious opportunities may be hard to prioritise.

Implications for a government AI operating model

- **Elevate the conversation on what AI can do in government.** Productivity and efficiency gains can emerge through widespread adoption, but more fundamental changes to services and outcomes are unlikely to happen automatically. Leaders need to articulate their vision and where they want AI to make the greatest difference.
- **Productivity and transformation can run in parallel.** The government can continue to realise productivity gains from AI, while space for experimentation will help teams discover where AI could make a more fundamental difference.
- **AI's impact needs to be captured.** Time and money saved will capture only part of AI's value. Improvements in services and outcomes, including benefits that extend across departmental boundaries, also need to be measured.

In action: Towards a new operating model for AI

Examples are already emerging of the different ways AI can create value, from improving existing processes to supporting more fundamental changes in services and outcomes.

- **Doing work that would otherwise be impractical.** The Department for Transport is using deep learning image-recognition technology to detect potentially fraudulent or erroneous photographs submitted as evidence for electric vehicle chargepoint grants. The government estimates that carrying out the same checks manually would take more than 100 years of staff time, while the tool has helped recover tens of thousands of pounds in fraudulent claims.
- **'Rewiring' a process, not just speeding it up.** HMRC's Making Tax Digital for Income Tax programme is a wider digital transformation initiative but it illustrates how technology can form part of a more fundamental change in how a service works. The website states that "it's the biggest change since HMRC launched Self Assessment more than 30 years ago", noting that: "We're rewiring the way you do Income Tax". Under the MTD scheme, eligible sole traders and landlords need to keep digital records and provide quarterly updates to HMRC instead of annually. It aims to reduce errors and help reduce uncollected tax. Rather than providing the software itself, HMRC enables a market of compatible products. AI-enabled HMRC Assist can be built into those products, using information submitted to HMRC alongside data it already holds to identify potential errors and provide tailored feedback.
- **Improving human outcomes, not just efficiency.** AI-supported stroke imaging is being used across more than 70 NHS hospitals to help clinicians identify patients who could benefit from time-critical treatment. An evaluation found that, at participating hospitals, thrombectomy rates doubled from 2.3% to 4.6%, while the time taken to transfer patients from primary stroke centres for specialist treatment fell by 64 minutes. Patients whose scans were reviewed using AI were also more likely to have a favourable functional outcome when they left hospital. As of December 2025, 15,000 patients had directly benefitted from having their scans reviewed by the technology, according to an analysis published in The Lancet Digital Health.
- **Managing complexity at scale.** The Grid We Need Now, an independent review led by Lucy Yu, the UK government's AI champion for clean energy, argues that the opportunity from AI goes beyond improving individual tasks such as forecasting and planning. As renewable generation, electric vehicles and flexible demand make the electricity system more complex, it argues that AI could enable a shift towards risk-based operation and planning of the grid, helping operators make decisions with greater precision, speed and confidence. The review estimates the potential value to consumers and the economy to be in the billions of pounds.

2. Establish a more joined-up approach to AI

Rewiring the State found support for a more joined-up approach to government digital transformation but also highlighted some resistance to changes or over-centralisation that could slow delivery.

The interviews for this study confirmed that there is a lot of good work underway on AI across government, but leaders still often struggle to see what is happening beyond their own organisations and networks. Departments are also investing in tools and capabilities to meet needs that are often shared.

As AI activity grows, a new operating model will need to define where government should take a common approach, and where departments and teams need freedom to innovate.

Continue to bolster technical foundations, but keep moving on AI

Many of the technical foundations needed to scale AI are not new. One interviewee summarised the challenge as a combination of "AI tooling, the legacy data, IT systems that can integrate with the AI tooling… [and] security", alongside organisations' ability to operationalise AI.

Legacy technology remains a challenge, particularly as AI moves beyond relatively standalone, "greenfield" applications. As one interviewee put it: "The biggest opportunities are going to be on top of all of these legacy systems, all of these big digital products that are already owned by digital teams inside government."

Data remained another recurring concern. One interviewee described it as "the great enabling thing" for AI but noted that much of the government's valuable information remains "hyper-fragmented", with different organisations holding parts of what is needed without necessarily knowing what exists elsewhere, how it can be brought together or whether it can legally be shared.



Christophe Prince: *Director, data and technology policy and strategy, Home Office*

Another interviewee said: "The data is not in the right place. The regulation is not in the right place… We're not one government, and we're individual departments. We have our own rules and responsibilities, and the governance associated with it."

In one example, a team identified a promising opportunity for automation only to discover that the dataset required to make it work did not exist in a useable form. The limiting factor was not the AI, but "having the right underpinning data".

At the same time, however, interviewees acknowledged significant work already underway to address these longstanding issues, such as **action on legacy technology** and the potential of the forthcoming National Data Library, though most interviewees still had

little detail on the latter and one said it was important but only "a drop in the ocean" in terms of addressing the structural barriers to routine data sharing across government departments.

Notably, though, there was greater optimism than in last year's *Rewiring the State* study that these challenges need not halt progress. Several interviewees felt that AI progress could still be made while the government addresses data and legacy IT, and in some cases, AI could itself help accelerate that work.

Some suggested that AI could make some previously difficult foundational work easier. One interviewee described engineers using AI to understand legacy systems and access their data, arguing that "the opportunity to join data up is massive now" and that it is "so much easier for us to be able to do that using artificial intelligence".

Still, these fundamentals remain a priority. As one put it: "There's some really basic stuff that we just have to get right," citing data interoperability, legacy IT and data-sharing.

"We can't actually use AI at the moment in some of the ways we'd like to because we've not got some of these foundational things right," they said. "If we don't fix some of these things, we're not going to be able to exploit AI as quickly as we would like to."

Move beyond ad hoc sharing

One of the clearest findings of this study is the opportunity to do more to reuse what works and solve common problems.

Two key issues stood out from the interviews: limited visibility of tools, assets and initiatives already being used or developed elsewhere in government, and the opportunity to identify common capability needs that could be addressed across organisations rather than duplicated within each.

Several interviewees felt that cross-government collaboration on digital and AI had been improving recently, particularly through the work of the Department for Science, Innovation and Technology (DSIT), which has now been disbanded. However, they also described collaboration and reuse as still too dependent on individuals knowing who to contact or proactively seeking out information.



> **"If we don't fix some of these things, we're not going to be able to exploit AI as quickly as we would like to."**

Most felt it was too soon to assess the impact of the recent machinery-of-government changes. These include redistributing DSIT's functions to the Department for Digital, Culture, Media and Sport and the Department for Business, Innovation, Science and Trade; establishing the new Office for the Prime Minister and the Cabinet, which will host a new AI Taskforce; and giving the Cabinet Office responsibility for AI strategy and public sector AI

adoption. Several nevertheless noted the disruption such changes can cause, calling for clarity on programme continuity and greater stability from here.

Ultimately, however, the challenge is to build capabilities and models that are resilient to such changes. As one interviewee noted: "There aren't optimal organisational constructs. What we need to be able to do is create things that withstand the organisational constructs and changes to them… a shared underpinning."



Dan James: *Advisor to the cabinet secretary on AI in the civil service*

Work is already underway in parts of government to take successful approaches developed in one service and adapt them for others, but this is far from systematic.

"We're in this world of 'let a thousand flowers bloom'," said one permanent secretary, echoing a term used by more than one interviewee as well as in the previous *Rewiring the State* study.

"I don't think we have the slightest idea across the government as a whole: what are the flowers that we're all tending?"

Another interviewee was concerned that "we're replicating the same idea in different organisations" and therefore "operating sub-optimally".

However, there is an important counterpoint to this, highlighting the balance to be struck. As one leader commented: "I'm not sure it's a very bad thing that you do have a thousand flowers, as long as some of those flowers die and you end up with a garden you want. We've got to have that sandbox mentality and just allow some flowers to keep going."

Without a mechanism for selecting and scaling the strongest approaches, experimentation can result in competing products, duplicated effort and poor use of scarce digital and technical capability.

Increase visibility of AI efforts, successes and capabilities

Visibility is part of the problem. One interviewee said: "I really want to know what other departments are doing." They found it "extraordinary that I don't know how other people do their basic things".

Their suggestion went beyond AI use cases and they pointed to the potential value of "model department functions" and "model department processes", so that organisations could understand how others are doing common things rather than each starting from scratch.

Another commented: "It's so easy for us to be able to coordinate all of the AI activities going on across government, and we just don't do it today."

One leader called for a "mandatory" cross-government tracker of AI activity. This had support elsewhere as long as it was deliberately designed to be useful rather than a "dump of stuff that will be difficult to parse".

Rewiring the State identified the need for stronger central guidance and practical tools for departments adopting AI. That need remains, with one interviewee calling for a 'version

1.0 playbook' to help departments not only implement AI but use it as a "change agent" across their organisations.

Interviewees also saw a much greater role for cross-government professions. One expressed being "deeply frustrated" at the lack of activity from some professions, asking why they were not "actively pushing standardisation of business processes, making use of the latest tooling solutions, and the adaptation and adoption of AI".

Pool demand for common capabilities

Visibility alone is not enough, however. Interviewees pointed to AI capabilities in areas such as casework, compliance, record-keeping, grant delivery and transcription that may apply to different services but represent "really the same competence" applied to different contexts.

And one asked: "How do we make sure… that we are developing things that have multiple organisations buying into them earlier on?"

This could also help establish what genuinely can be common. Testing a capability in different organisational environments may reveal why it transfers successfully in some cases and not others. As one interviewee put it, the question is: "How do we find projects that actually will have resonance beyond a single department?"

There is a further current challenge when a successful departmental product attracts demand elsewhere. The team that developed it must decide whether to divert its own resources towards maintaining and scaling a product for other organisations or leave those organisations to recreate much of the capability themselves.



> **"How do we make sure… that we are developing things that have multiple organisations buying into them earlier on?"**

As one interviewee argued, government could "go further than simply sharing our learnings… our code… because that's what we always do".

This points towards a more active model of reuse: creating visibility of what is being developed and deployed, pooling needs earlier, co-developing or sourcing solutions where requirements are common, and providing a route for successful products to become shared where there is wider demand.

Common does not always mean one size fits all

However, the interviews also highlighted a tension between greater reuse and standardisation and retaining sufficient choice at departmental level.

There was support for the centre of government to provide common tools and capabilities where this makes sense but interviewees also distinguished between common

requirements – "where the organisational value comes from being on the same system" – and more specific applications where departments may need greater freedom to select the technology best suited to the outcome they are trying to achieve.

The opportunity, one argued, is for the centre of government to provide more common infrastructure which can be made "specific to the particular use case of a particular department or domain".



Jerome Glass: *Director general, Future Civil Service, Cabinet Office*

Several interviewees also described benefits from giving staff access to AI tools, but there was an appetite for greater choice. One interviewee argued that departments should have "flexibility and autonomy" to pursue a multi-vendor strategy: "Civil servants should have access to lots of those different tools," they said.

Make the centre enabling – but not passive

These issues, in turn, raise the question of the centre's role. As one interviewee put it: "I think a lot of this comes down to the operating model of how we scale technology and interoperability in government."

That includes greater clarity about the division of responsibilities between the centre of government and departments. One interviewee argued that it is "not always clear" what the government has agreed should be done once collectively, what organisations such as the Government Digital Service and Incubator for AI are expected to deliver, and what departments should work out for themselves. Where significant investment is being made in common capabilities, they argued, priorities should be collectively agreed rather than developed centrally and then presented to departments. This points to a more co-created model: the centre can provide common capabilities and coordinate investment, while departments help shape what is needed and retain freedom where requirements genuinely differ.

One interviewee envisioned the model as an "enabling rather than constraining" centre. This means establishing direction and guardrails, solving genuinely common problems, providing access to shared capabilities, and then allowing organisations to innovate within that.

Enabling may not always mean optional. Some interviewees wanted a stronger and, in defined areas, more directive centre. One said they would welcome a "slightly more dictatorial" approach, where the centre could identify a product, service or platform for wider use. They said they were "happy for more central planning… and standardisation" as AI adoption develops.

Cybersecurity was another example. "We don't need every government department to go off and try and work out… how big is the threat and what to do about it," one interviewee commented. "We need that to be done really once by experts."

The centre of government could be more "bullish and tell departments what's required," they added.

Some interviewees also saw scope to move beyond voluntary sharing where approaches have become sufficiently mature. Rather than either sharing best practice and leaving departments to act on it, or providing a capability centrally, the centre could do more to set common standards, performance measures or frameworks that departments are expected to follow when buying or building. One interviewee argued that the government is "nowhere near assertive enough" in doing this.



Jo Farrar: *Permanent secretary, Ministry of Justice*

While there were several common themes, the interviews revealed a tension between mandate and autonomy. Some of the strongest progress has come where teams have had clear senior backing to pursue AI, but also significant freedom over how they operate – from recruitment and technology choices to how they develop products, including permission to deviate from existing standards. While this approach is not necessarily scalable in itself, an experimental approach can generate lessons on which freedoms are essential to successful AI delivery, and which things genuinely need to be common.

Implications for a government AI operating model

- **Address the foundational technology and data problems that will hinder AI scaling.** The opportunity-cost in government of legacy systems and fragmented data will increase if these long-standing issues are not tackled. The government's acknowledgement of tackling legacy technology as a priority has been a positive step but ongoing focus is needed.
- **Improve cross-government visibility of AI activity.** Interviewees argued that fragmented information-sharing on AI limits opportunities to learn from and reuse work already underway and creates duplication, with some even calling for the creation of a mandatory cross-government tracker of AI activity.

  The government already has frameworks, tools and examples of AI in practice. An evolving AI toolkit could bring together resources on practical aspects such as managing risk, evaluating impact and moving from experimentation to deployment.

- **More AI tools and capabilities can be shared – without assuming one size fits all.** Departments will continue to need applications shaped around their own users and needs, but the underlying capabilities they require are often common. This creates scope to reduce duplication by identifying shared needs earlier and finding ways to develop, source or scale capabilities across organisations, while retaining flexibility for innovation.
- **Strike the right balance between central direction and departmental autonomy.** Interviewees saw value in the centre providing greater direction, while also valuing the freedom to innovate, experiment and choose tools that fit their needs. Involving departments in shaping common priorities and capabilities can help strike that balance.
- **Cross-government professions should play a stronger role in AI rollout.** Because professions span departmental boundaries, they can identify common processes and requirements, support standardisation where it adds value, and help spread effective approaches across government.

In action: Towards a new operating model for AI

Different approaches to developing and scaling technology across the government show some of the choices emerging for AI.

- **Doing common things once.** GOV.UK One Login provides a common way for people to sign in and prove their identity when using government services, rather than individual departments developing separate solutions. More than 250 government services can now be accessed through One Login, with HMRC and Companies House among those moving services onto it.
- **Moving from experimentation to scale.** Justice Transcribe uses AI to transcribe and summarise probation supervision sessions and other case interactions. It was developed by the Ministry of Justice's Justice AI Unit, which has deliberately adopted a different way of working. Its technical staff have **considerable freedom** over the problems they pursue and the tools they use, while 'forward deployed engineers' work directly alongside staff in prisons, courts and probation services to develop and test products. Justice Transcribe has progressed from frontline pilots to wider deployment, with **Justice AI reporting** a 50% reduction in note-taking time across pilot and early-scale sites and a staff rating of 4.7 out of 5. The approach raises a wider question about which aspects of this model could be replicated as AI activity scales across the government.
- **Tackling common problems.** Ruby, also developed by the Ministry of Justice's Justice AI Unit, brings private-office work including briefings, commissions and box work into a single digital system, replacing fragmented spreadsheet-based processes. Early pilots show reduced reliance on manual tracking, faster coordination of submissions, briefings and diary activity, and improved visibility and continuity across private-office workflows. Ruby is now being piloted in additional departments. As it has expanded, the team has adapted Ruby for different private offices while putting common security and governance arrangements in place. The example shows the potential to identify problems shared across departments and develop capabilities that can be reused, rather than each department solving them separately.

3. Prepare for how AI will change work

The interviews made clear that people, culture and organisational readiness are just as much if not more of a challenge for AI than the issue of technology.

According to one interviewee, the civil service needs to move beyond "come and have a go with AI" towards much more routine use of the technology. "Everybody should be using it every day" until that becomes normal, they said.

A new operating model will require continuous investment in staff skills, alongside planning for how AI will change work: what people will do, how roles will evolve, and what skills and capabilities the government will need as a result.

Continue to build technical expertise and AI capability across the civil service

Technical talent remains important to getting the most from AI. Some of the strongest examples of AI delivery in government have relied on skilled engineers working closely with frontline teams. Interviewees also stressed that the government needs the right technical expertise to work effectively with suppliers and "hold vendors to account".

Although progress has been made, attracting and retaining that expertise still remains challenging. Interviewees highlighted pay and career structures as particular barriers. One argued that the government needs to recognise that "a software engineer can be as impactful as a director general", while another called for career frameworks that allow highly skilled technical staff to progress as individual contributors rather than having to move into conventional management roles.

Technical specialists are only one part of the capability government will need for this new way of working.



Kalbir Sohi:
Government chief AI officer, Government Digital Service, Department for Digital, Culture, Media and Sport

A major finding of the *Rewiring the State* study was the need to upskill all civil servants in data and AI, from the frontline to the top of the organisation. A cross-government learning offer has since been introduced, and interviewees described further efforts underway within departments.

Several interviewees stressed the need to continue the focus on this and go further. One described upskilling and investment in people's skills and confidence as the "number one thing", adding: "I don't think we can do it fast enough". They pointed to an important role for the National School for Government in building that capability.

They also argued that this needs to start at the top. Leaders "set the tone" through their expectations, their own use of AI and as a role model they said, calling for government to "double down" on training and learning for senior leaders. This should extend to political leaders too, so that they have "the best possible understanding" of AI and its potential.

Several highlighted the value of practical AI masterclasses for senior leaders, including sessions in which they designed and built applications for themselves.

One interviewee described current training overall as "moderately good" but argued that the range of tools available to civil servants "places a ceiling on what value people can get out of those tools", including the scope for creativity and experimentation.

Training is not enough

The interviews also suggest that training people to use AI is only part of the challenge. The civil service also needs to think about how roles will change, what that means for the skills, shape and potentially size of the workforce, and how to manage the transition.



Jo Shanmugalingam: *Permanent secretary, Department for Transport*

The June 2025 Spending Review, under the previous prime minister, committed to publishing the first-ever civil service strategic workforce plan. As this develops, interviewees argued that this plan should take a broader view than effectively assuming that the civil service will have roughly the same number of people in five years doing "broadly the same sort of things".

Instead, they argued that the government needs to ask "what a job's going to look like in five years' time" and then shape recruitment, training and the workforce around that vision.

Understand how roles will change

As AI can increasingly summarise information, surface relevant material and recommend next steps, for example, one interviewee described how employees will increasingly be operating in an "AI-augmented world".

"Teaching and coaching our colleagues to operate in an AI-augmented world with confidence" will be increasingly important, they said.

Another spoke of the opportunity to "elevate the human aspects of the job, whether it's autonomy, critical thinking, using judgement, having more relational aspects of a role that may end up resulting in far better citizen outcomes than simply just doing what we're doing today, which often isn't good enough".

The implications could extend beyond individual tasks and roles to professions themselves.

One interviewee anticipated "an enormous amount of upheaval" in the policy profession, for instance. Some have recently argued that Whitehall needs to move away from the 'policy generalist', with a broad range of knowledge across different areas, towards deeper domain expertise.

"I think the opposite is going to be true for the next 5 or 10 years," they said. "I think what we're looking for actually, what we're going to develop, is multidisciplinary people."

With this in mind, some called for the professions to be more proactive in informing the discussion on how AI changes work in their area.

Career paths will also need to change

With many expecting the government to increasingly recruit people who are "AI natives", this also raises questions about how people develop foundational capabilities early in their careers.

If AI increasingly performs some of the research, evidence gathering, synthesis and initial analysis traditionally undertaken by junior staff, organisations will need to consider how these people develop judgement and decision-making.

"To ask the AI the right question, you really need to understand it. To understand whether you're getting junk out, you really need to understand it," one interviewee said.

"Our model needs to change," they added. "And certainly, we don't yet have an answer to that expertise and model question."

More than one leader described today's workforce as a "pyramid", with large numbers of more junior staff progressing towards a smaller senior cohort. They suggested it could become more like a "diamond" if AI reduces some of the work currently done at junior levels.

"You're probably going to have to find ways to bring entry-level people up to a level of breadth of responsibility that normally you would be looking at five or ten years in," they said. "And how do you do that? How do you do that quickly? Because the risk is, if you just cut it all off at the bottom, you haven't got your pipeline."

"I can see the end state vaguely," they added. "I cannot see the pathway, and I think that is our challenge."

How AI could reshape the workforce

Changes to tasks and roles will clearly have implications for the overall shape – and potentially size – of the civil service workforce as the government adapts its operating model for AI.



John-Paul Marks: *Permanent secretary and chief executive, His Majesty's Revenue and Customs*

While some anticipated a smaller civil service, interviewees were cautious about simplistic predictions of large workforce reductions, with the impacts likely to vary substantially between different types of work. In some areas AI may help address existing workforce shortages, while elsewhere it may reduce the need for particular tasks or roles.

"It's not a threat to the workforce because we don't have enough [people]. It enhances their jobs," one said.

Some warned against a "jobmageddon" outlook, suggesting that many predictions about large-scale job losses are based on the experience of the private sector and software industries, which are different in nature to public service.

Interviewees also saw potential for AI to "raise the workforce up", creating "more valuable jobs" and "a better career path".

Examples of private sector organisations redesigning roles as AI takes on existing work were referenced in the interviews. **IKEA, for example,** retrained thousands of call-centre workers as interior-design advisers as an AI assistant took on a growing share of routine customer enquiries. The company made a decision to use the capacity released by AI to move people into different work and develop a new source of revenue.

Plan for change and bring people with you

The precise workforce implications are therefore uncertain and unlikely to be uniform across the civil service. However, leaders agree that a 'wait and see' approach will not suffice.

One interviewee argued that AI change needs to be "accompanied with a workforce plan".



Jonathan Brearley:
Permanent secretary, Department for Energy Security and Net Zero

"It's not about getting rid of your job. It's about helping you to develop and grow," they said.

They also warned that "fear of change and people worrying about losing their job is the thing that gets in the way of change".

"If you give them some certainty and you help them see a future, then actually that takes away a lot of the anxiety and therefore a lot of the resistance," they said.

With leaders across departments facing many of the same challenges, there is significant scope for a shared and collaborative approach. One questioned: "How far and firmly ahead do we plan to realise those [AI] benefits without causing a spectre of something that actually we don't really know exactly what it means?"

The need to navigate this uncertainty and actively shape the future of work reinforces the finding from *Rewiring the State* on the importance of digital and AI being a leadership issue, not just a technical one.

Implications for a government AI operating model

- **AI capability will need continued investment at all levels.** Attracting and retaining technical talent will mean addressing barriers such as pay and career structures. Ongoing AI skills and capability development will also be needed across the wider civil service, from frontline staff to senior and political leaders.
- **Workforce planning needs to take account of how the work is changing.** Starting from today's jobs risks missing more fundamental changes in what people will actually be needed to do, with implications for recruitment, skills and career development.
- **Planning will have to happen despite uncertainty.** The interviews do not point to one clear outcome for the future workforce, and the impact of AI is likely to vary significantly between different types of work. Workforce plans will need to respond to what is already changing while retaining flexibility as the impact of AI becomes clearer.
- **Bring people with you.** Leaders will need to engage staff openly about how AI will impact how work is done, even where the longer-term implications remain uncertain. Helping people understand the opportunities and develop their skills will be important to build confidence and support adoption.

In action: Towards a new operating model for AI

Early applications of AI are beginning to show different ways the technology could impact work.

Freeing up time. The Department for Business, Innovation, Science and Trade **uses generative AI** to help staff find answers to routine questions from departmental information. During testing and piloting, users reported getting answers more quickly and needing to ask colleagues fewer questions, while some said it freed up time for more valuable work. BIST's **Business Finder** uses machine learning and government data to help export advisers identify and prioritise businesses likely to benefit from support, allowing advisers to spend less time identifying and verifying leads and more time supporting businesses.

AI assistance for policy work: The Department for Health and Social Care's Data Science & Applied AI Team has introduced SIFT, an AI-assisted tool developed by the Department for Science, Innovation and Technology that helps process consultation, call for evidence and survey responses submitted in unstructured formats such as emails and PDFs. The tool uses AI to identify relevant consultation questions, extract the corresponding content and present the results for human review and approval. This reduces the time analysts spend manually reading and restructuring responses, allowing them to focus on analysis, insight generation and policy advice. Early evidence suggests a 50-75% reduction in processing time per response, while retaining human oversight throughout.

Potential knock-on impacts. NHS England says that ambient voice technologies, which capture clinician–patient conversations and use AI to generate real-time transcriptions and clinical summaries, could save clinicians up to two or three minutes for each patient consultation, freeing up more time for them to see other patients.

However, at a Global Government Forum roundtable, **NHS leaders warned** that this time could simply be converted into greater throughput, potentially leaving clinicians "always on" and affecting wellbeing and resilience. They also noted that apparently 'inefficient' moments in a working day can have value, allowing clinicians to connect with patients and recalibrate. This underlines that the value of time saved depends partly on what it is then used for.

4. Use trust, not hype, to set the pace for AI

Trust was top of many interviewees' minds. They were concerned that, if not handled well, AI could damage citizen trust and make further adoption harder. The government, therefore, faces a difficult balance: moving quickly enough to capture the benefits of AI while giving leaders and the public confidence that it is being used responsibly.

A new operating model will require clarity on risk appetite and accountability, as well as transparency with the public.

The need to balance pace and responsibility

There is considerable pressure on the civil service to move quickly with AI. The capabilities of the technology are developing rapidly and departments are pressed to do more with constrained resources.



"I think it's important that there is encouragement, but I think it's also important that that is matched with a tone about responsibility. It mustn't be scale at the risk of being reckless."

One interviewee noted that AI, alongside people's experiences of consumer technology and services, is dramatically changing what citizens expect from government and public services.

"The burning platform that we already feel will feel even more burning," they said.

However, while AI is sometimes characterised as an 'arms race' between organisations, sectors and governments, interviewees also cautioned that pace is not the only important factor.

"I think it's important that there is encouragement, but I think it's also important that that is matched with a tone about responsibility. It mustn't be scale at the risk of being reckless."

Leaders noted that a significant failure could undermine confidence in AI more widely and make subsequent adoption harder.

"You only need your first AI disaster, and that will set us back years," one warned.

However, they also warned against over-caution, instead arguing for active, pragmatic risk management.

Be clear and intentional about risk

Different applications carry different levels of risk. Using AI to support relatively contained internal processes is different from using it with sensitive personal data or in decisions that have significant consequences for individuals.

One interviewee argued for being more "intentional about our risk appetite" and starting with targeted applications where risks can be more easily understood and managed. This allows organisations to build experience, develop controls and establish confidence before moving into more complex areas.



Nick Dyer: *Interim permanent under-secretary, Foreign, Commonwealth and Development Office*

Another described a similar approach of starting in "the smallest possible context" to enable evidence and confidence to build. Rather than committing to a predetermined solution, they argued for starting with a hypothesis, exploring it and being prepared to "throw things away when they don't work", while extending successful approaches as they develop.

Another added that "risk proportionality is super important in this" and that AI operating across "hundreds of thousands of citizens in life-or-death scenarios" requires very different assurance from "an agent to improve inbox management or to do a recording of a meeting".

"Those judgments about risk assurance in terms of scaling, I don't think should be left to individuals," they said, suggesting that there should be clear, common criteria.

Transparency is becoming increasingly important

Public trust was a recurring concern. As AI becomes more embedded in government and public services, interviewees called for more focus on transparency.

One interviewee argued that citizens should not have "a moment when they discover AI is being used" without previously knowing about it.

The government already requires departments to publish information about certain AI tools through the Algorithmic Transparency Recording Standard. One interviewee made the case for going further, calling for departments to publish AI action plans that provide a wider picture of what they are doing with AI and why, the direction of future travel, and how they are acting safely.

These openly available plans, they said, should be "very crunchy, very precise", including where human oversight sits.

The Justice AI Unit's **AI Action Plan Explorer** offers an emerging example of greater public visibility, bringing together information on AI initiatives and their stages of delivery.

At the same time, some were concerned that the government sometimes "over-indexes" on concerns about data and underestimates citizens' appetite for services that make better use of the information government already holds.

This suggests a need to understand more clearly where public confidence is lower and where over-caution may be preventing better services. It also requires a more mature conversation with the public about AI.

"I feel we're still in the cheerleader phase," one leader said. "Whereas I think we need to be able to explain that better for people. We have to make this about people, not about technology. And at the moment, there's a real risk that we're still too heavily focused on the technology rather than the people."

The recent example of digital identity illustrates the importance of communication and how quickly trust can be lost. Following the Starmer government's initial digital ID announcement and later amendments to the proposed approach after a public backlash, the Home Affairs Committee **concluded** that the plans had been poorly explained and had "spooked" the public, warning that rebuilding trust would not be a "quick-fix solution".

One leader commented: "We know how to do public engagement, and we are going to have to get a bit more uncomfortable in that transparency process."

This could mean increasingly needing to "expose some of the detail of our policy and analytical thinking. So, I think this is soluble but it's a choice."

There is a balance here too, with some interviewees also flagging the need to give teams space to test ideas before committing publicly to a particular solution or outcome.

Clarify accountability as AI advances

All of these issues become more significant with the emergence of agentic AI and as AI potentially moves beyond supporting individual tasks towards undertaking more complex activities and playing a greater role in decision-making.

A future model could involve humans monitoring the overall accuracy and effectiveness of AI systems rather than reviewing every individual case, for instance – "human across the loop" rather than "human in the loop".

One interviewee warned that the risks of errors can increase significantly in an agentic environment because "problems are multiplied a thousandfold because you're taking them much faster".



Susannah Storey: *Permanent secretary, Department for Digital, Culture, Media and Sport*

That raises questions that go beyond the accuracy of individual systems. Existing approaches to public administration, accountability and civil service values largely assume that a person is ultimately undertaking the work or making the decision.

One interviewee questioned whether that remains sufficient as the government potentially moves towards an environment in which activity and decisions may increasingly be shared between "humans and machines". A wide-ranging review of the civil service led by cabinet secretary Antonia Romeo will include refreshing the Civil Service Code, and the interviewee suggested that AI and decision-making could potentially be considered as part of this.

Interviewees were agreed on the importance of clarity about how decisions are made and verified, who is accountable when AI performs more of the work, and what meaningful human oversight looks like.

While the permanent secretary is ultimately responsible for what happens in a department, "technology has already made that feel a bit uncomfortable" said one interviewee, noting that AI takes this further and simply assuming that "the buck stops" at the top does not create the right culture of accountability.

One department runs a programme for senior leaders focused on making it clear that they are responsible for the processing of data and for an algorithm throughout its lifecycle.

"It's a team sport," they added. "And you need to engage with the expert individuals who can provide you with the understanding and the information that is needed for you to say, 'yes, I am content that this is operating within an appropriate set of performance criteria'."

It is not acceptable, they said, for something to go wrong and for everybody then to start to "point the finger" over where responsibility lies.

Pace and trust need to work together

Leaders did not see these questions as reasons to hold back adoption, but as highlighting "sensible things you need to do" to enable more ambitious use of AI.

Indeed, some interviewees were concerned that excessive caution could itself prevent government from capturing the benefits of AI.

The interviews point to several things that can help build that confidence: clarity on risk appetite and accountability, transparency and appropriate safeguards. Recognition that departments understand the risk profile of their own areas and use cases can also give them greater confidence to move at an appropriate pace.



Samantha Jones:
Permanent secretary, Department of Health and Social Care

"I think we need to allow for people to move at different pace in an organisation of this scale," said one interviewee.

Balancing trust and pace is not an insurmountable challenge. "Good risk management is what we should be experts in in the civil service," one interviewee said, pointing to the need to articulate risk appetite, understand different types of risk and manage them effectively. They highlighted the Orange Book, the government's existing framework for risk management and risk appetite: "It gives you all of the tools in a very practical way." An aligned **AI Risk Management Toolkit** was recently released. The key is to make these frameworks a systematic part of how AI is approved and managed.

Implications for a government AI operating model

- **Risk appetite needs to be clear and proportionate.** Interviewees wanted teams to have space and backing to experiment where risks are manageable, with firm boundaries around more sensitive or consequential uses.
- **Transparency about how government uses AI will become increasingly important.** Existing transparency requirements provide a starting point, but interviewees see a need for people to be able to understand what government is using AI for, why, where its use is heading and how humans remain involved. This needs to be more nuanced and detailed than "cheerleader" messages about AI.
- **Accountability needs to remain clear as AI advances.** As AI takes on more of the work and potentially plays a greater role in decisions, questions need to be addressed on who remains accountable, how decisions are checked and what meaningful human oversight looks like.
- **A more nuanced understanding of public trust in AI will be needed.** Attitudes are likely to differ considerably depending on where and how AI is used. Understanding those differences and the nuance will become more important as government decides where it can move quickly and where greater explanation, involvement or safeguards are needed.

In action: Towards a new operating model for AI

The government needs to balance pace and trust. Emerging examples show how different parts of the public sector are beginning to navigate this.

Setting safety standards for the market. The Department for Education has **developed** product safety standards for generative AI in education, setting out the capabilities and features products should meet to be considered safe in schools and colleges. Aimed primarily at edtech developers and suppliers, the standards cover areas including filtering, security, privacy, testing, governance and children's development.

Understanding public acceptance of risk. The Automated Vehicles Act requires the safety principles for self-driving vehicles to aim for a level of safety equivalent to, or higher than, that of "careful and competent human drivers". In developing the safety standard, the government **has sought views** from the public and industry on where the threshold should sit, including the implications for public acceptance, safety and the pace at which the technology can be deployed.

Taking a proportionate approach to risk. The NHS is increasingly using AI-enabled ambient voice technology to turn clinical conversations into documentation such as notes and letters. Current **guidance** distinguishes between functions such as transcription and summarisation and those used to support diagnosis, treatment or automated clinical actions, which are subject to medical-device regulation.

The National Commission into the Regulation of AI in Healthcare has **proposed** a new framework intended to allow patients to benefit from AI more quickly while maintaining safety and public trust. Its recommendations include staged authorisation of new AI models, continuous monitoring once they are in use and greater public access to information about their safety. The Commission's engagement with more than 12,000 people found broad support for AI in healthcare, but with clear expectations around safety, human oversight and transparency.

Conclusion: The next phase of government AI

A year on from Global Government Forum's *Rewiring the State* report, many of the longstanding challenges around data, technology, skills and capability remain but the conversation has also moved on. There is greater optimism that government can make progress on AI while continuing to tackle these underlying barriers. The focus is now moving beyond adoption towards how AI can deliver more fundamental improvements in services and outcomes.

This study points to four key areas that will support this and form the foundations of a new operating model linking AI more clearly to the outcomes government wants to achieve and demonstrating its impact; greater visibility of AI activity across government and a more systematic approach to scaling and reusing what works; continuing to build the skills to use AI effectively while preparing for how it could change tasks, roles and professions; and ensuring clear parameters around risk, accountability and transparency.

The opportunity is to move beyond the 'glide path' of doing existing things more efficiently and towards using AI to deliver better outcomes for citizens and the economy.

Next steps

The findings of this and other Global Government Forum research will provide the basis for further exploration of the issues identified through our ongoing programme of content, including events, research, training, newsletters and more.

Stay connected:

- Share your feedback on the study and your insights on how you are addressing the challenges and opportunities identified. Please email Sarah Wray, research editor at Global Government Forum, at **sarah.wray@pendragonim.com**
- **Sign up to our newsletter** to receive our digital transformation content and opportunities for further collaboration.
- **Attend our in-person events** for unrivalled learning and networking.
- If your organisation would like to partner with Global Government Forum on our research programme as a Knowledge Partner, please contact Ed Tillotson at **ed.tillotson@pendragonim.com**

Legal Notice

© Pendragon International Media Ltd 2026.

All rights reserved

The information contained within the document is given in good faith and is believed to be accurate, appropriate and reliable at the time it is given, but is provided without any warranty of accuracy, appropriateness or reliability.

The author does not accept any liability or responsibility for any loss suffered from the reader's use of the advice, recommendation, information, assistance or service, to the extent available by law.

Global Government Forum is part of Pendragon International Media Ltd

Registered address: Hamilton House, Mabledon Place, London, WC1H 9BB, United Kingdom

www.globalgovernmentforum.com

www.pendragonim.com