Re:State

accenture

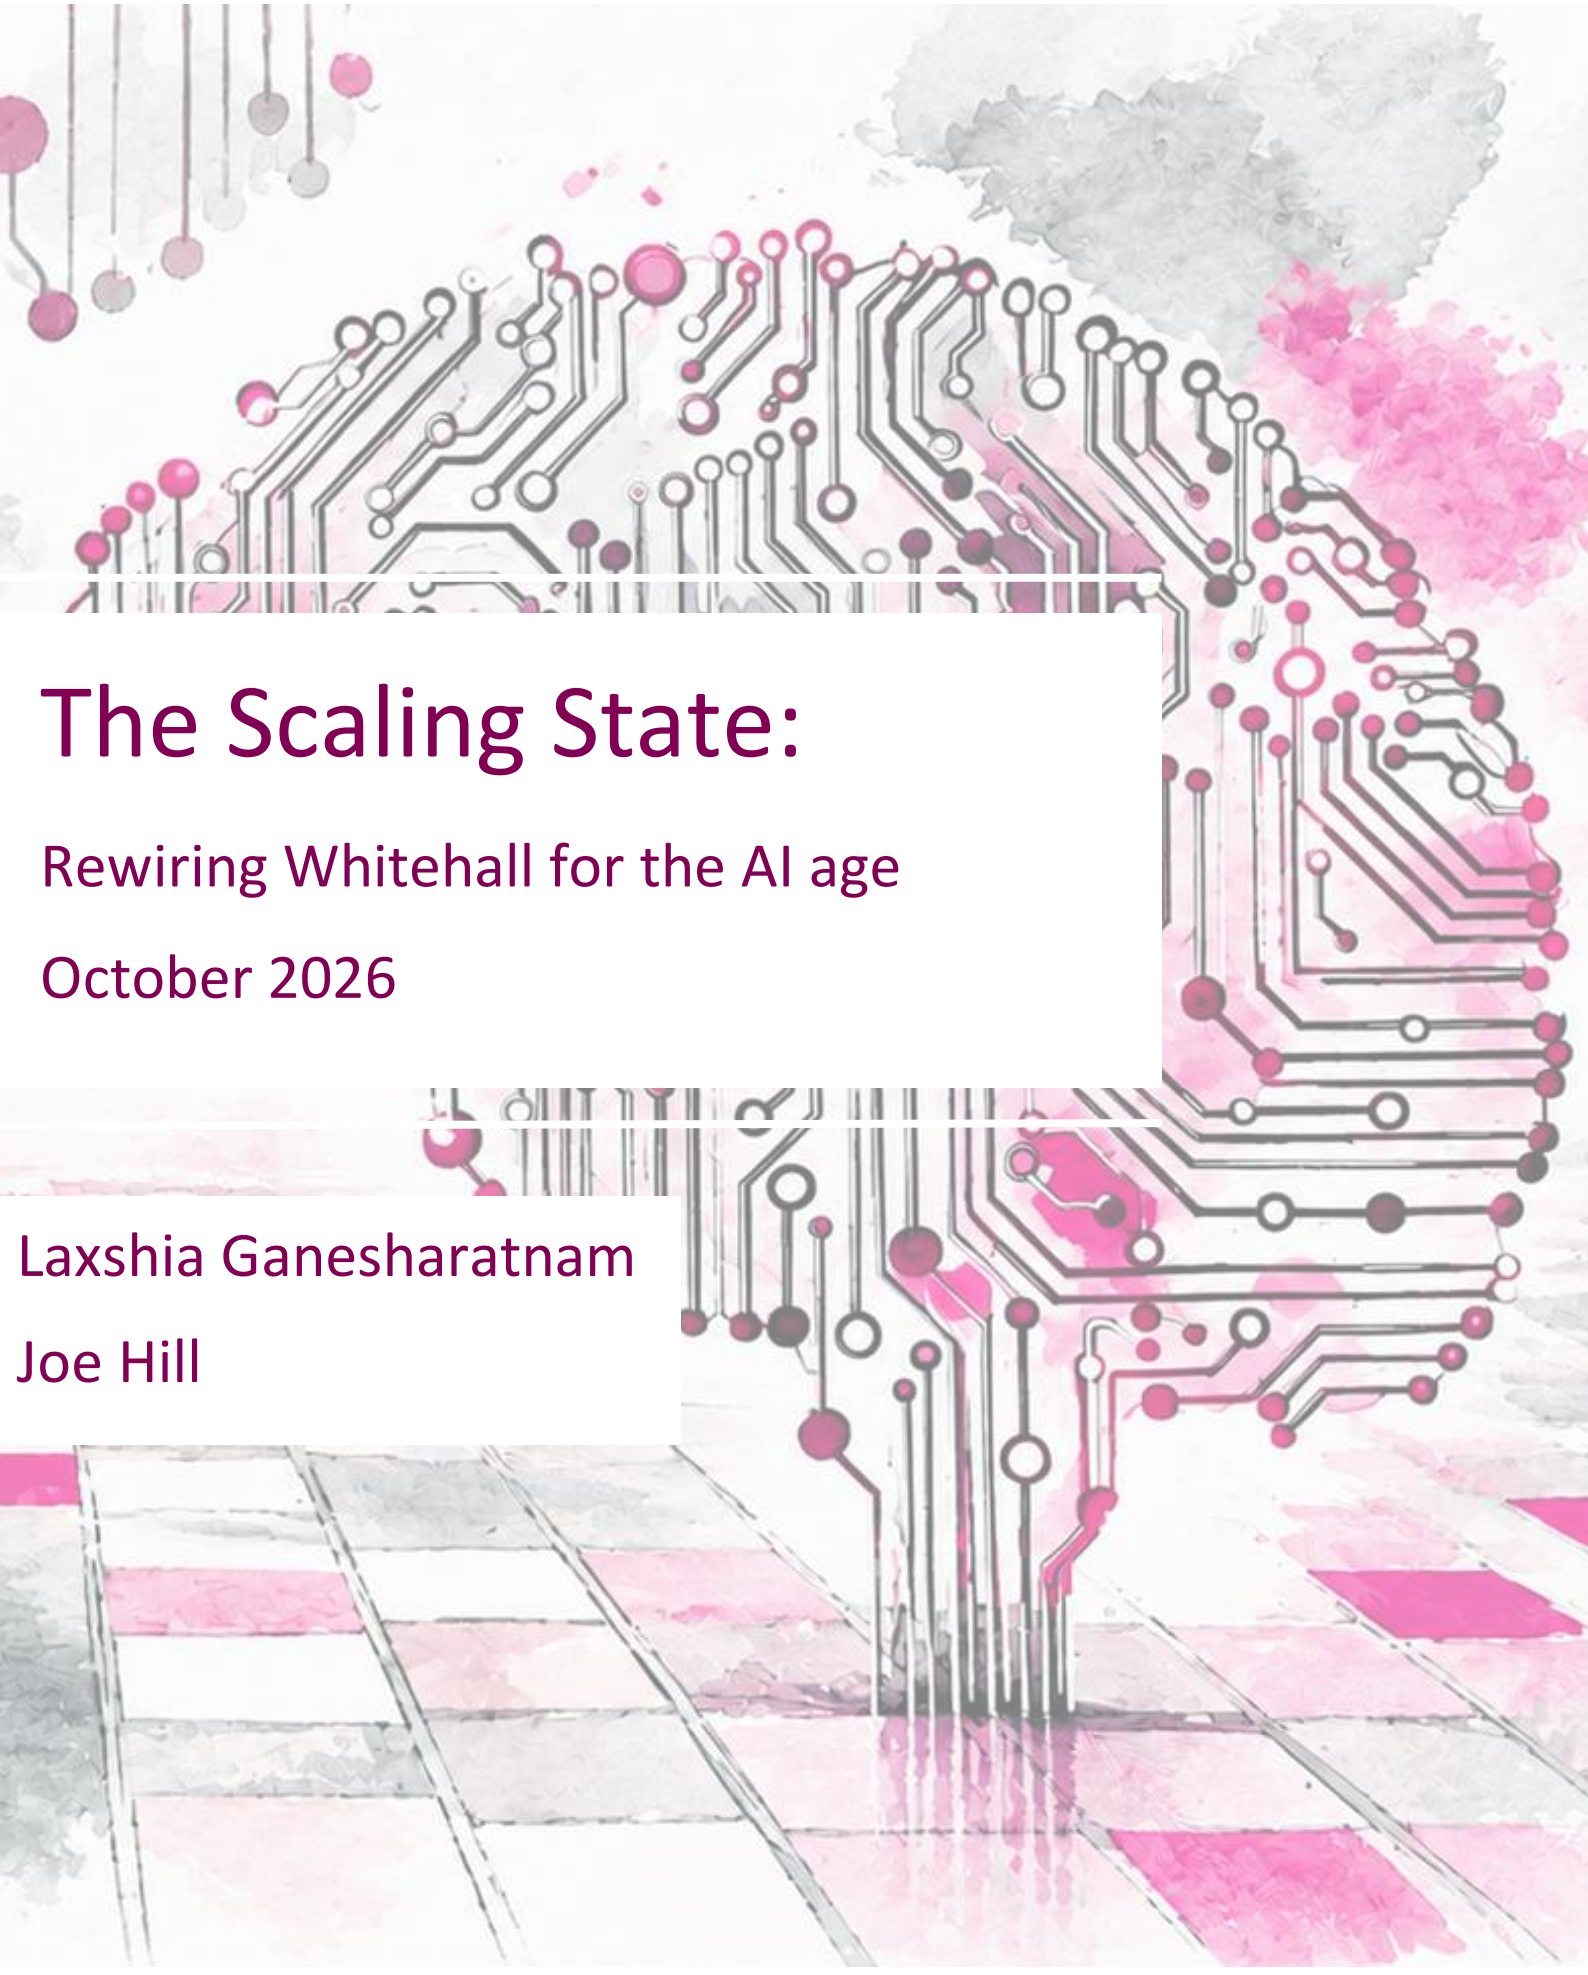

The Scaling State:

Rewiring Whitehall for the AI age

October 2026

Laxshia Ganesharatnam

Joe Hill

ABOUT RE:STATE

Re:State is established as the leading Westminster think tank for public service reform. We believe that the State has a fundamental role to play in enabling individuals, families, and communities to thrive. But our vision is one in which the State delivers only the services that it is best placed to deliver, within sound public finances, and where both decision-making and delivery is devolved to the most appropriate level. We are committed to driving systemic change that will deliver better outcomes for all.

We are determinedly independent and strictly non-party in our approach. This is reflected in our cross-party Advisory Board and our events programme which seeks to convene likeminded reformers from across the political spectrum.

Re:State is a registered charity, the Re:State Trust Ltd, charity no. 1103739.

ABOUT RE:IMAGINING THE STATE

After a decade of disruption, the country faces a moment of national reflection. For too long, Britain has been papering over the cracks in an outdated social and economic model, but while this may bring temporary respite, it doesn't fix the foundations. In 1942 Beveridge stated: "a revolutionary moment in the world's history is a time for revolutions, not for patching." 80 years on, and in the wake of a devastating national crisis, that statement once again rings true. Now is the time to fix Britain's foundations.

Re:State's programme, 'Re:Imagining the State', puts forward a bold new vision for the role and shape of the State. One that can create the conditions for strong, confident communities, dynamic, innovative markets, and transformative, sustainable public services.

ABOUT RE:IMAGINING WHITEHALL

To effectively reimagine the State itself, the institution at its heart needs to radically change – Whitehall. Both the first and second world wars led to a rewiring of the core structures of the State, driven by an imperative to 'build back better' and an understanding that the effectiveness of the machinery of government would make or break that endeavour. After a decade of deep disruption – economic stagnation, pandemic, trade conflict and war in Europe – we face a similar task.

While successive administrations have attempted to modernise the civil service and improve the structure of Whitehall, the same problems recur time and time again. The system is slow to change, and reforms are hard to make last. This paper is part of *Re:State's* 'Re:Imagining Whitehall' workstream. To effectively reimagine the State, major change must occur in the behaviours, processes, and structures of central government.

ACKNOWLEDGEMENTS

Re:State would like to thank Accenture for kindly supporting this research.

Interviewees

We would like to thank all six interviewees for giving their time and candid insights to support this research paper. The list of interviewees is as follows:

- Victoria Bew, Policy and Programs, DeepMind
- Andy Gregory, Deputy Director, Data services and Analytics, Home Office
- Kirsty Innes, Director, Centre for AI-Driven Innovation at Imperial College London
- James Lowe, Member of Technical Staff, Anthropic
- Emily Middleton, Interim Director General, Digital Transformation, Government Digital Service, Department of Digital, Culture, Media and Sport

And one interviewee who wished to remain anonymous.

The arguments and any errors that remain are the authors' and the authors' alone.

Table of contents

1. INTRODUCTION 8
2. GAP BETWEEN AMBITION AND DELIVERY 11
2.1 False assessment of risk 12
2.2 Governing technology adoption 13
3. AN OPERATING MODEL FOR SCALING AI 17
3.1 Operating model 17
3.1.1 A mismatched departmental State 17
3.2 From pilots to scale 19
3.3 The role of agents 22
4. PROCUREMENT 24
4.1 Buying the building blocks – procuring generative AI 24
4.2 Laying the foundations – procurement for the AI age 27
5. WORKFORCE, SKILLS AND ADOPTION 30
5.1 Capabilities to support adoption. 30
5.2 A workforce transformed by AI 32
6. FUNDING APPROACHES 34
6.1 Systemic financial traps 34
6.2 Funding for an AI-ready State 36
7. CONCLUSION 38

Forewords

FOREWORD BY JOE HILL

Director of Strategy, Re:State, report co-author

Every generation of British government has promised to go digital. Tony Blair's Government set a target for every public service to be available online by 2005. The Coalition Government created the Government Digital Service and declared that services would be 'digital by default'. Each time, the political language moved faster than the policy underneath it, and neither kept pace with the technology itself.

Artificial intelligence raises the stakes for public services. A frontier model can complete a piece of casework for a few pence, but a department that wants to test one will often run a procurement lasting months before anyone is allowed to log in. In the United States, federal agencies have bought access to the leading models for a dollar a year. Meanwhile in the UK, the Government has signed memoranda of understanding with the same companies which explicitly commit neither side to any actual commercial partnership.

So civil servants have gone ahead without permission. *Re:State's* survey of more than 1,200 officials found 30 per cent were using AI in their work, and most of them relied on tools their department had not provided at least some of the time. That is a rational response to a system which does not respond quickly or effectively to a new technology, and which promises to be left behind if it does not catch up.

The barriers to adoption are rarely technological: they sit in the Treasury's budgeting rules, in commercial routes designed for multi-year contracts, in workforce plans that do not account for automation, and in an accountability system where ministers rarely show much interest in digital programmes until they have already gone wrong.

This paper sets out 20 ideas to change that, from an annual Technology Statement scrutinised by a new Office for Technology Assessment to a standing route for buying frontier models based on usage.

AI will be game-changing: just as previous general-purpose technologies have remade the economy, AI too can remake the State for the better. But with great opportunity comes a risk of missing out. The policies ministers set will influence how AI is adopted, but indecisiveness itself is also a choice – in the face of a State standing still, the technology will not pause for it to catch up, and if government falls behind it risks being left behind.

FOREWORD BY HACER EVANS

Managing Director, Health & Public Service Lead, Accenture UKI

This paper comes at a critical moment for government.

Across the public sector, leaders are being asked to deliver better outcomes against a backdrop of rising demand, constrained resources and growing public expectations. At the same time, the increasingly devolved nature of decision-making is placing greater responsibility in the hands of local leaders, creating new opportunities to design services around the needs of places, communities and citizens, rather than around the structures of central government.

The question at the heart of that shift is not whether public services should be more responsive, more productive and more sustainable. They must be. The question is whether the Government is building the foundations that will make that possible. Shared digital infrastructure, common standards and interoperable platforms are not considerations that follow reform, they are conditions for it.

The same logic applies to artificial intelligence. The organisations that will see the greatest benefit will be those that use it to rethink how services are designed and delivered from the ground up, reimagining operating models and ways of working around what these technologies genuinely make possible. As with previous waves of technology-driven change, the real value is only realised when institutions adapt to take advantage of it.

That adaptation is a leadership challenge. The public-sector organisations that navigate it well will be clear about the outcomes they are trying to achieve, disciplined in building the capabilities to deliver them, and deliberate about creating conditions in which innovation can be adopted safely and at scale.

That is what makes this paper both timely and necessary. It moves the conversation beyond technology, beyond pilots and proofs of concept, and addresses the harder question of how government builds the capability, leadership and institutional foundations to deliver lasting public value. These are amongst the most consequential challenges facing government today, and this paper is a serious contribution to meeting them.

20 Ideas for rewiring Whitehall for the artificial intelligence age

Idea 1: The Office for the Prime Minister and Cabinet should mandate a cross-government 'AI first' principle. If policies are not compatible with the adoption of AI, based on technical advice about current best practice, they should be reviewed and amended.

Idea 2: The Government Digital Service should add a new first principle to the Service Standard. 'Use AI first', encouraging service teams to develop and test new products quickly and place them in the hands of users as part of user research and feedback.

Idea 3: The Treasury should require departments to include explicit 'risk of inaction' assessments or evaluations in major digital business cases, adding it to the programme risk register.

Idea 4: The Office for the Prime Minister and Cabinet, alongside HM Treasury, should establish a process for publishing an annual Technology Statement in Parliament, supported by dedicated debates in both Houses of Parliament and joint select committee hearings. The Statement should cover progress against digital transformation objectives, and set out headline measures and policy decisions for the coming years against which the Government's performance should be judged.

Idea 5: The Government should put the publication of the annual Technology Statement on a legislative footing, alongside independent scrutiny of the Statement by a new Office for Technology Assessment (OTA). The OTA should be established as an operationally independent unit with clear terms of reference to scrutinise and publicly report on the Government's programme of digital transformation and use of AI, publishing its assessment to coincide with the Statement.

Idea 6: The Government should create a single digital authority with real delivery powers to oversee digital government transformation across departments, by moving the Government Digital Service back to the Cabinet Office.

Idea 7: Government should appoint a lead Secretary of State and senior official with responsibility for overseeing the Government's wider digital and AI transformation across government.

Idea 8: Government should publish a clear framework for which common digital capabilities will be delivered centrally and what should remain with each department, with strong alignment between ownership of policy, data, technology, and delivery no matter where a service is located.

Idea 9: The Government Digital Service should codify and roll out a 'middle-out' deployment process for AI, led by technical teams funded centrally in each department.

Idea 10: The Government Digital Service should establish a Trusted Agents Service as a key part of cross-government digital policy.

Idea 11: The Government's Digital Commercial Centre of Excellence should create a central route to market for the specific procurement of frontier AI software, available to all contracting authorities and paid for based on actual usage.

Idea 12: For AI-enabled and platform-based services, contracting authorities should be required to design requirements so that more than one supplier can meet them, by separating the model layer from the application layer, and identifying a credible second source before award.

Idea 13: GDS should mandate a standard exit schedule in every central government digital contract above a set value, covering export of data and configuration in documented open formats, no charge for extraction, defined notice, transition and retrieval, and a supplier obligation to assist migration. Departments should be required to rehearse migration of at least one significant workload during the contract term, and report the result to the Government Digital Service.

Idea 14: GDS should launch challenge funds for larger scale-up companies targeted at public service challenges, and based on buying software products rather than solely services.

Idea 15: The Government should launch an 'advance talent commitment' for AI, committing to hire up to 400 staff with machine learning and AI skills on salaries of up to £200,000, across government, with a single point of approval in the Government Digital Service for all salaries. The budget for these roles would come from within hiring departments.

Idea 16: The Government should work with industry partners to launch a Tech Talent Initiative for hiring experienced technology leaders into the private sector.

Idea 17: The Government Digital Service should create an AI Testing Fund to purchase additional capacity from teams across government to test AI in their areas through paid overtime, as part of structured evaluations designed to inform the roll-out of new, more productive ways of working.

Idea 18: Departments should ensure workforce planning incorporates the potential efficiency gains from AI, including implications for staffing levels, skills, organisational design, recruitment and the way external capacity is commissioned.

Idea 19: Shared platform services (identity verification, data exchange, common AI tooling etc) should be funded centrally, where these services create value across multiple departments and should not depend on any individual department's budget allocations.

Idea 20: HM Treasury should create a cross-government Digital Investment Fund. This should be ringfenced for shared infrastructure, reusable and strategic transformation projects. Departments should be able to apply for match funding when a digital or AI investment will benefit multiple parts of government, to encourage co-investment and shared ownership for initiatives with wider benefits.

1. Introduction

The UK does not lack digital strategies. *The Blueprint for Modern Digital Government* offers a long-term vision for public services, while the accompanying *State of Digital Government Review* makes a frank case for change.[1] Together, they recognise that public services need to be joined up, AI should be harnessed for public value, and government must strengthen digital and data infrastructure, leadership, talent, funding, procurement, and transparency.

But strategy is not the same as transformation. The central weakness of the current approach is that digital is still too often treated as a layer of government delivery rather than as an operating model for the State. Government has become better at designing individual digital services and setting standards for them, but it has been less successful at reorganising the underlying systems of funding, ownership, data, accountability, and workforce capability that determine whether those services can be joined up and improved over time.

Eighteen months on from the publication of the *Blueprint*, the evidence is that the change it describes is not happening in full or fast enough.

The Government Digital Service ran the largest AI trial in Whitehall's history in late 2024, with 20,000 officials using Microsoft 365 Copilot and reporting an average of 26 minutes saved a day;[2] but follow-up evaluations at the Department for Business and Trade found no conclusive evidence of a measurable productivity gain,[3] raising questions both about the value of the technology and how it was used, and access to any generative AI tool remains uneven across departments two years later.

HMRC's first Chief AI Officer told a conference in April 2026 that it had been almost impossible to track the AI activity underway in his own department, with separate teams procuring similar tools independently.[4]

Of the £3.25 billion Transformation Fund announced at the 2025 Spending Review, just £42 million went to the three Frontier AI Exemplars intended to show what the technology could do across government.

The result: widespread use of AI by civil servants, much of it on tools their employers have not provided, in ways their employers' policies would not necessarily approve of, but which make the most sense to them as employees.

AI in government can take several different forms. The opportunity is not only to empower individuals with better tools, but to use AI to fundamentally change the nature of government work and public service delivery. The potential for this technology in the public sector can be categorised in at least four ways:

- Individual productivity: using AI tools to support and speed up everyday tasks at the scale of single team members.

[1] Department for Science, Innovation & Technology, Government Digital Service, *A Blueprint for Modern Digital Government* (HM Government, 2025); Department for Science, Innovation & Technology and Government Digital Service, *State of Digital Government Review* (HM Government, 2025).
[2] Government Digital Service, *Microsoft 365 Copilot Experiment: Cross-Government Findings Report* (2025).
[3] Department for Business & Trade, *The Evaluation of the M365 Copilot Pilot in the Department for Business and Trade* (2025).
[4] Christine Horton, *HMRC Rolls out 28,000 Copilot Licences as AI Chief Targets World's 'Most AI-Enabled Tax Authority'* (2026).

- Operational enhancement: integrating AI deeply into operations, including casework and back-office administration – an area with significant potential for rapid efficiency gains.
- New types of public service: using AI's specific abilities to deliver services or forms of support that were previously difficult or impossible to provide at scale, such as highly personalised education using AI tutors.
- Software development: using AI coding tools that make the development and maintenance of government software significantly faster and more accessible.

This pattern stems from Whitehall treating the most important strategic shift in a generation as a technology programme. Generative AI changes the cost of administrative work, which in turn changes what a department should do, how it should be funded, who should be accountable for any given process and what the civil service workforce is for. Every one of those is a question for the Government Digital Service, the Treasury, permanent secretaries and ministers, and the current model has assigned all of them to digital teams in individual departments with no power to answer them. The Government's own plans for a more automated civil service depend on adjustments to spending rules, commercial routes, accountability structures and workforce planning that are not factored into their strategy.

AI safety is also an essential concern for policymakers as this transition accelerates. More capable models create greater opportunities to improve public services, but also create new risks around inaccurate outputs, insecure tool use, data leaks, cyber misuse and systems taking actions that are difficult to explain or reverse. Recent incidents, such as the OpenAI/Hugging Face episode[5], show how developments in model capabilities can challenge existing assumptions about what these systems are able to do. Government cannot treat safety as a one-off approval exercise at the point of procurement. It needs continuing technical evaluation, monitoring, 'red-teaming', and clear accountability throughout the life of an AI-enabled service.

This is particularly important as Government moves from generative AI assistants to agentic systems. The ease and accessibility of creating agents has increased rapidly, lowering the barriers to developing systems that can take actions rather than simply generate information. The challenge will be to maintain safety and consistency as more teams begin engaging with and creating agents. Technical evaluations will be essential to show not only whether an agent can complete a task, but when it fails, how it behaves under pressure, what permissions it uses, whether its outputs and actions can be followed and evaluated, and when it should escalate to a human. This will be important not only for managing risk, but for building the public, parliamentary, ministerial and workforce trust needed to deploy these systems at scale.

This paper sets out what it would require to adjust the Government's digital strategy to account for the systemic barriers to innovation. It argues for a commercial and governance model built for a market that moves faster than a procurement cycle, including a standing route for buying access to frontier models and the foundations for a multi-vendor ecosystem. It makes the case for a central programme to put AI agents into public services on a footing that can scale, with identity, permission and accountability infrastructure. It argues for a 'middle-out' model of transformation, in which the operational units that own a process end to end – an agency, a directorate, a service line – hold the permission and the money to redesign it.

This is a moment of opportunity, with a vast prize in prospect. If the Government can capture even a small fraction of the potential of generative AI, and do so consistently and at scale across government, then it will be able to fundamentally rewire the State for the better. The

[5] OpenAI, *The Hugging Face Incident and the Road Ahead* (2026).

Government's devolution agenda compounds the opportunity, and the challenge. A new generation of public service reform, built on pushing power out of Whitehall into new institutions, creates a moment to rebuild those same institutions in a new image – not replicating an analogue model of government, but instead building public services which were designed for the AI age.

More efficient services, higher staff productivity and whole new categories of public services are only the beginning. In a government where artificial intelligence is as common as human intelligence, then the cost of deploying software to improve services will be increasingly small – ushering in the real potential of a digital state, where bureaucratic and tedious processes are all automated and public servants can spend their time on the activities only they can do best – exercising human judgment and care in service of the public.

2. Gap between ambition and delivery

Government has spent decades announcing its intention to become more digital. The language has changed from ‘e-government’ to ‘digital transformation’ to ‘government as a platform’ and to ‘AI-enabled public services’. But no matter the language used, delivery has rarely lived up to it. The persistent gap between digital ambition and tangible improvements is not a result of a lack of strategy. Rather, it reflects that the core levers of the State such as spending controls, business cases, audit, commercial governance, accountability and departmental budgeting are inhibiting adoption of digital technologies. Often, these processes disincentivise cross-cutting, iterative transformation of the kind software has promised, meaning progress has tended to be slower and harder-won than it should be.

Departments can publish ambitious strategies and make commitments to modernisation. But implementation must then pass through systems designed to control risk, allocate finite resources, and establish clear lines of accountability. These are legitimate functions of a State responsible for ensuring value for money from public spending. However, taken together these governance and policy measures tend to favour programmes that are tightly scoped, predictable, and owned by a single organisation. Digital transformation, though, will require discovery, iteration, shared infrastructure, process redesign, ongoing maintenance, and sustained leadership over a long period of time. This often puts our government system at odds with the necessary conditions for change.

The National Audit Office (NAO) has repeatedly identified this mismatch.[6] Its work finds that standard approval, procurement, funding and assurance mechanisms do not align well with digital programmes, which are often hard to define with the certainty expected in business cases. Teams often need research to understand the scope, the costs, and technical feasibility of a transformation, but they often need an approved business case before they can actually start working or secure the funding that will allow them to do that work. Therefore, the result is a tendency to lock in assumptions before government knows enough to make them reliable.

These challenges were common in previous generations of digital transformation, but are even more acute for the era of AI adoption.

The original *Manifesto for Agile Software Development* described a way of working which fundamentally upended how large enterprise digital transformation projects had previously been led throughout the 1990s, including:

> “Individuals and interactions over processes and tools
>
> Working software over comprehensive documentation
>
> Customer collaboration over contract negotiation
>
> Responding to change over following a plan”.[7]

Today, similar principles are still at the heart of the Government Digital Service’s *Service Standard* and its 14-point plan for how digital teams should build and run public services.[8] AI

[6] *The Challenges in Implementing Digital Change*, No date.; *Digital Transformation in Government: A Guide for Senior Leaders and Audit and Risk Committees*, No date.; *Government’s Approach to Technology Suppliers: Addressing the Challenges*, No date.
[7] Kent Beck and et al., *Manifesto for Agile Software Development* (2001).
[8] ‘Service Standard’, GOV.UK, accessed 8 September 2026, https://www.gov.uk/service-manual/service-standard.

promises to change these further, making software even quicker and easier to deploy and test in a live environment with users. AI is already transforming the software industry that built it: generating more and more digital products, automating software engineering processes, and making them faster to deploy and improve.

If this was an ordinary programme for public service reform, government departments would expect to change the processes of governance slowly – they exist to safeguard important public sector equities, like value for money and proportionality. But AI adoption is not an ordinary programme for public service reform. It is a rapidly maturing technology with the potential to permanently shift the practices of government, and where the advantages of being an early adopter won't come again. With the urgency that demands, the Government should be prepared to radically alter its policies when necessary. This could be reflected by the adoption of an 'AI first' principle across all internal government policies. If policies are not compatible with the adoption of AI in government, based on the advice of the lead technologists in that department or organisation, then they should be reviewed for amendment.

Idea 1: The Office for the Prime Minister and Cabinet should mandate a cross-government 'AI first' principle. If policies are not compatible with the adoption of AI, based on technical advice about current best practice, they should be reviewed and amended.

This same approach should apply to GDS and its own policies about software development. The Service Standard predates AI and is based on different assumptions about how quickly software can be prototyped, tested and improved. Including an 'AI-first' approach as the first principle in the Service Standard would address this and bring it up to date with best practice.

Idea 2: The Government Digital Service should add a new first principle to the Service Standard: 'Use AI first', encouraging service teams to develop and test new products quickly and place them in the hands of users as part of user research and feedback.

2.1 False assessment of risk

Government is right to scrutinise the risks of transformation, however, a false assessment of the risks of change is a significant barrier to digital transformation. The Green Book requires departments to assess costs, benefits, risks, and uncertainty.[9] Its approach is designed to counter the tendency to underestimate the costs or overstate the benefits of a project. This results in a system much better at pricing the risk of doing something than the risk of inaction. Government appears to prefer known inefficiency over reforms that could fail, as current practice is far better at identifying the risks of change than the risks of inaction.

Business cases are built assuming the status quo as a baseline, and that change needs to be decided positively. For example, a business case may catalogue delivery risk, commercial

[9] 'The Green Book (2026)', GOV.UK, accessed 9 September 2026, https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government/the-green-book-2026.

risk, costs and uncertainty. But assumptions around the 'do nothing' cost of the status quo are rarely as comprehensive, such as staff time absorbed by manual workarounds to inefficient processes, growing tech debt, inability to share data, deteriorating cyber resilience, poor user experience, or the much higher eventual accumulated costs or the cost of forced replacement. The cost of legacy systems is a particularly acute example, where *Re:State's* own research has found the ongoing cost of legacy IT is rarely properly budgeted for.[10]

The Public Accounts Committee (PAC) has argued that departments should identify the full costs of the services they operate, track the benefits of transformation, and assess the opportunity cost of not changing[11] – building on the call by many digital practitioners to "fund teams, not projects"[12], and *Re:State's* research which outlined a model for service-by-service budgeting and value-for-money evaluation of spending.[13] PAC has also recommended that business cases show how they will resolve problems caused by legacy systems and data, and how wider redesign will reduce future service costs.[14] While this is the right direction, it needs to become a consistent requirement rather than an occasional recommendation.

Idea 3: The Treasury should require departments to include explicit 'risk of inaction' assessments or evaluations in major digital business cases, adding it to the programme risk register.

2.2 Governing technology adoption

The mismatch between stated government digital adoption plans, and their delivery against them, is fundamentally a problem of collective action. It's easy to make bold claims, but harder to deliver them, and Government lacks a strong system of accountability to force that. This is even more important when that change is perceived as complex or risky, as described in the previous section. One interviewee was particularly frustrated by this, claiming that senior leaders routinely deprioritised digital projects because "Ministers will never ask about them unless they are already going wrong". *Re:State's* research has previously identified how digital spending is often squeezed by other calls for public money.[15]

Within government, there is often abundant reporting on strategy, announcements, expenditure and individual handpicked programmes, especially in recent times, but less clear reporting of whether services are being redesigned, whether technical debt is declining, whether departments are reusing shared capabilities, or whether productivity and user outcomes are changing. This leads to fragmented or lacking accountability across the system. Government can issue guidance, set standards and expectations, while departments retain decisions over systems, budgets and delivery. When delivery stalls, responsibility is

[10] Laxshia Ganesharatnam and Joe Hill, *From Legacy to Leadership: Upgrading the Digital State* (Re:State, 2026).
[11] Cabinet Office, *Digital Transformation in Government: Addressing the Barriers to Efficiency* (National Audit Office, 2023).
[12] Andrew Greenway and Tom Loosemore, *The Radical How* (Public Digital, Nesta, 2024).
[13] Joe Hill, *Byte-Sized Budgeting: Funding Digital Service in Government*, January 2025.
[14] *Digital Transformation in Government: Addressing the Barriers to Efficiency*, n.d.
[15] Hill, *Byte-Sized Budgeting: Funding Digital Service in Government*.

distributed across the centre and departments, with no single mechanism that pinpoints where action is required.

This gains even more importance in light of AI. Without a shared view of readiness, implementation and barriers, Government risks repeating the pattern of from-scratch experimentation in each department without system-wide learning. A department may have a promising pilot, but no route to scale it, no common infrastructure on which to build, and no incentives to share resulting capabilities with other departments or across government. This challenge never sees the light of day outside government, unless forced through a large public failure or a select committee inquiry, and those processes are insufficient to drive the right approach across a digital footprint as large and complex as Whitehall's.

Case study 1: Bank of England's modernisation of the Real Time Gross Settlement System

The Real Time Gross Settlement System (RTGS) is a piece of critical national infrastructure, underpinning the UK's payments system and playing a significant role in maintaining the resilience and stability of the UK economy. The original system went live in 1996, and the new system was launched in April 2025 following a major renewal programme costing £431 million.

Its failure would pose a serious risk to the UK's financial and monetary stability. The Bank needed to replace an ageing system while maintaining uninterrupted, highly resilient payment settlement. It invested in planning, governance, procurement, testing, and risk management alongside technical delivery, and adjusted its delivery plans as technical dependencies and implementation risks became clearer.

The renewed platform gives the Bank a more modern and adaptable foundation for the UK payments system. It is designed to support future innovation as well as current settlement needs, and the Bank has planned further improvements over the following two to three years, rather than treating launch as the end goal.

RTGS shows that public institutions can renew highly complex legacy infrastructure where they treat it as a strategic national capability rather than a conventional, time-limited IT project. The Bank combined clear ownership, sustained leadership, substantial upfront planning, and low tolerance for operational failure with a commitment to continuous improvement after going live.

Source: Committee of Public Accounts, *The Bank of England's Real-time Gross Settlement Renewal Programme*, Seventy-Eighth Report of Session 2024-26, April 2026.

Similar challenges in government are solved through a combination of governance and external scrutiny processes. The most powerful is the system of annual Budgets and regular multi-year Spending Reviews, 'fiscal events' designed to draw together the different policy commitments the Government is making on spending, taxation and borrowing and present their plans comprehensively for external scrutiny. Combined with the Office for Budget Responsibility, which independently assesses progress against the Government's own fiscal plans (in the form of the fiscal rules) it is a powerful (if intensive) system for holding the Government to account for the economy and public finances.

There is no equivalent process which demands clarity on progress against digital transformation, or framework for tracking changes and the performance of different projects.

Given the significance of digital transformation, particularly through AI adoption, this is untenable. The Government should establish a similar framework for external scrutiny of digital transformation, recognising its fundamental importance to the future of the State, based on an annual Technology Statement produced by the Government, and analysis of its plans by an Office for Technology Assessment, inspired by the Office for Budget Responsibility.

Idea 4: The Office for the Prime Minister and Cabinet, alongside HM Treasury, should establish a process for publishing an annual Technology Statement in Parliament, supported by dedicated debates in both Houses of Parliament and joint select committee hearings. The Statement should cover progress against digital transformation objectives, and set out headline measures and policy decisions for the coming years against which the Government's performance should be judged.

Idea 5: The Government should put the publication of the annual Technology Statement on a legislative footing, alongside independent scrutiny of the Statement by a new Office for Technology Assessment (OTA). The OTA should be established as an operationally independent unit with clear terms of reference to scrutinise and publicly report on the Government's programme of digital transformation and use of AI, publishing its assessment to coincide with the Statement.

To provide the necessary technical authority at the centre of government to lead delivery of the plans in the Statement, the Government should reconsider the recent change to move GDS into the Department for Digital, Culture, Media and Sport (DCMS). The 'digital centre of government', including GDS and other units like i.AI, was moved to DSIT in a machinery of government change following the last general election. Its move to DCMS in July is another disruption which will take significant time to work through, and leaders will be hesitant to make further changes.

But getting the structures at the heart of digital government right is crucial – because the current status of these organisations doesn't provide sufficient capacity to drive central change. The proper place for the Government's digital capability is at the centre of government in the Cabinet Office, along with other core enabling functions for the State, and the convening power of a body that can act across Whitehall.[16]

The Cabinet Office should then take on the responsibility of working with departments to drive digital transformation and AI adoption across the wider public sector. While this should not mean centralising delivery or taking responsibility for departmental services away from departments, it should mean providing the coordination, technical authority and support needed to turn cross-government priorities into practical changes within departments. The Cabinet Office should work with departments to identify opportunities for shared capabilities and common infrastructure, support adoption of common standards and approaches, and help resolve barriers when transformation requires action across organisational boundaries. This would give the centre a clearer mandate to drive adoption while enabling departments to remain accountable for delivery in their own areas.

GDS should then aid the operation of national digital infrastructure, set standards for interoperability, coordinate departmental transformation agreements, provide technical

[16] Joe Hill and Sean Eke, *Getting the Machine Learning: Scaling AI in Public Services* (2024).

support, and report publicly on adoption and progress through the Government Technology Statement.

Idea 6: The Government should create a single digital authority with real delivery powers to oversee digital government transformation across departments, by moving the Government Digital Service back to the Cabinet Office.

One reason digital transformation repeatedly loses momentum is that responsibility is diffused across ministers, departments, central units, and delivery teams, while no significant political or administrative role is accountable for turning cross-government commitments into results.

The annual Technology Statement and the OTA would create visibility and scrutiny, but the system still lacks ownership. There should therefore be a clear political and administrative owner for delivering the Government's digital transformation agenda. A lead Secretary of State and senior official should therefore be responsible for overseeing implementation of the Technology Statement and the Government's wider digital and AI transformation programme, ensuring that cross-government commitments are followed through and that barriers to delivery are addressed.

Idea 7: Government should appoint a lead Secretary of State and senior official with responsibility for overseeing the Government's wider digital and AI transformation across government.

3. An operating model for scaling AI

Even with greater governance and accountability for digital transformation, and an appraisal of risk and reward from change which is much more permissive of radical changes, the Government's vision for an intelligent State powered by AI misses a crucial enabler: an operating model for digital government which can realise the potential of AI.

Whitehall treats digital as a functional capability which fits around department structures, whereas an AI-first State would organise around end-to-end user journeys and shared infrastructure that can cut across institutional boundaries.

One key operating model for government digital services has been the development of shared products like GOV.UK, Notify, Pay and One Login which can replace individual solutions used in many different government departments. GDS's focus in the *Blueprint* is building more of these services – GOV.UK App, Wallet and Chat.[17] There are good reasons to build these single user-facing services, which show clear value to the public.

But many of the benefits from AI adoption in government will not come from enterprise-wide services like these, and will instead come from the iterative improvement of overly siloed systems which sit in each department's own operations and budgets. Casework, customer service, logistics and business planning are all services with digital components which have huge promise from automation,[18] but where the technological leadership does not come from GDS, it comes from departments where local technology choices, separate budgets, and fragmented data still determine how services are delivered.

The Blueprint leaves unresolved fundamental questions: who has authority, capability and accountability to drive ambitious changes that necessarily cut across departmental structures as we know them? And what theory of change will they use for the adoption of AI?

3.1 Operating model

3.1.1 A mismatched departmental State

User experiences go beyond a departmental State[19] as citizens experience life events such as having a child, starting a business, bereavement, moving home, losing work or needing support later in life. These journeys often require contact with several public bodies, each holding a different part of the relevant data and each operating a separate process as government remains organised around departments, agencies, arm's length bodies, local authorities, and professional silos. This creates a persistent mismatch between the way services are owned and the way they are used. The result is not just inconvenience but duplicated processes, repeated requests for the same information, inconsistent service quality, and barriers to data-driven redesign.

The *State of Digital Government Review* illustrates the problem clearly. It notes that a 2021 review identified 19 different account setups and 44 unique sign-in methods across

[17] Department for Science, Innovation & Technology, Government Digital Service, *A Blueprint for Modern Digital Government*.
[18] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.
[19] Department for Science, Innovation & Technology and Government Digital Service, *State of Digital Government Review*.

government, a fragmentation that GOV.UK One Login is intended to reduce.[20] This reflects a deeper problem: without a clear expectation that foundational infrastructure should be shared, each organisation develops its own solution and the State accumulates duplication.

Some functions have a particularly strong case for central ownership to both avoid duplication and because their value increases when they are consistent, interoperable and available across organisations.[21] These include identity and authentication, payments, notifications, data exchange, core cloud and security capabilities, common AI tooling, and standards for interoperability and evaluation. These are not merely departmental IT products, but enabling infrastructure in a truly digital government, without which departments cannot easily design services around the needs of users or make reliable use of data and AI. In these areas, central provision can reduce duplication, strengthen security, create economies of scale, and ensure that services and systems can work together.[22]

Other functions should continue to be owned by departments: policy objectives, frontline judgement, service-specific discretion, and responsibility for outcomes in their sectors. A more centralised digital operating model should not mean centralising all decisions. It should mean creating a clearer division of labour: the centre owns the capabilities that need to work consistently across government whereas departments own how those capabilities are used to meet policy and service needs.

Case study 2: Estonia's digital government reforms

Estonia's digitisation journey began back in 1994 with the signing of the first draft policy on digitisation. From the mid-1990s, the State invested in three foundational layers: a mandatory digital identity for all residents, a secure national data exchange layer otherwise known as X-Road, and clear legal and technical rules for how data could be shared across public services.

By 2002, digital ID was mandatory and the means to access basic public services such as a driver's licence or passport. It also became the single way citizens and businesses authenticate themselves when accessing public services, so there is no need for separate logins for different parts of the State. X-Road then allowed different IT systems to talk to each other securely, so that a school, a hospital, and a tax authority can exchange data without building custom point-to-point integrations. The State adopted a "once only" principle whereby citizens should not have to provide the same information twice, because the Government can retrieve it from existing registers via X-Road.

This architecture means that new services can be built by composing existing capabilities rather than recreating them. A citizen can file taxes, register a business, or access health records using the same digital identity and the same underlying data infrastructure. Officials can see relevant information from other parts of the State without needing to ask citizens to resubmit documents. Audit trails show who accessed which data and when, which supports both security and trust.

[20] Department for Science, Innovation & Technology and Government Digital Service, *State of Digital Government Review*.
[21] *Update on Government Shared Services*, No date.; 'Digital Infrastructure', IMF, accessed 9 September 2026, https://www.imf.org/en/publications/fandd/issues/2026/03/digital-infrastructure-diane-coyle.
[22] Department for Science, Innovation & Technology and Government Digital Service, *State of Digital Government Review*.

Estonia's experience shows that digital government becomes transformational, and marginal transformation much easier, only when interoperability is designed into the system from the start. The UK can learn from this by treating identity, authentication, and data infrastructure as core infrastructure that must be centralised and standardised, rather than leaving each department to build its own solutions. The case supports the argument that an AI-ready State depends on shared data, common standards, and reusable platforms, not just departmental pilots.

Source: URBACT, *E-Estonia: Digital Society & Open Data*, July 2024

The Government Digital Service (GDS), established in 2011 as the UK government's central digital transformation unit, is one of the clearest examples of what central leadership can achieve in government. It brought together specialist digital capability at the centre of government to lead the transformation of public services. GOV.UK consolidated over 2,000 government websites into a single front door and created a shared design system for public services.[23] Its core principles, such as user-centred design, multidisciplinary teams, iterative delivery, open standards, and reuse of common capabilities, remain the right foundations for digital government.[24]

But AI is not a single digital service, nor a specialist capability that can be held in one team like GDS. Rather, AI depends on shared data, secure access to models and tools, common evaluation and assurance standards, interoperable systems, and workflows redesigned around the technology. In other words, it requires the State to behave more like a platform than a federation of isolated technology estates.

Government should not aim to recreate GDS exactly as it was. Instead, it should focus on applying the principles that made GDS successful to building and operating national digital infrastructure, deploying AI at scale, and giving the centre sufficient authority to improve policy without stripping departments of their legitimate policy and operational responsibilities.

Idea 8: Government should publish a clear framework for which common digital capabilities will be delivered centrally and what should remain with each department, with strong alignment between ownership of policy, data, technology, and delivery no matter where a service is located.

3.2 From pilots to scale

As businesses and governments around the world experiment with AI adoption, there is no single clear model for best practice on how to scale AI throughout organisations. As the technology evolves, and more organisations experiment with different tools and products, it is

[23] Department for Science, Innovation & Technology and Government Digital Service, *State of Digital Government Review*.
[24] 'Use GOV.UK Domains and the GOV.UK Design System', GOV.UK, accessed 7 September 2026, https://www.gov.uk/guidance/use-govuk-domains-and-the-govuk-design-system; Design community, *Designing Good Government Services: An Introduction* (GOV.UK, 2024); 'Government Design Principles', GOV.UK, 2 April 2025, https://www.gov.uk/guidance/government-design-principles; 'GOV.UK - 1 Year Old and Improving All the Time', GOV.UK, accessed 7 September 2026, https://www.gov.uk/government/news/govuk-1-year-old-and-improving-all-the-time.

likely that a variety of better practices will emerge which the Government can begin to implement. But at present, little evaluation is going into whether current approaches are working well.

The previous section has already discussed the areas where centrally delivered AI services like GOV.UK Chat will be appropriate, and the areas where departments will need to continue to lead. But within departments, technology leaders still need a practical approach to turn experimentation into results. Despite broad experimentation, it isn't clear that departments are learning what works. Each organisation is running its own pilots, buying its own tools, building its own use cases and learning in isolation. Without shared platforms, patterns, data access, assurance, and deployment support, each department must solve similar problems independently.

Interviewees discussed the merits and drawbacks of two different approaches to identifying good practice in the use of AI:

1. Top-down direction. A central team, such as GDS, selecting crucial use cases where there are likely to be efficiency savings or productivity improvements, making business cases to invest in order to unlock those benefits, and running agile projects to deliver a new solution underpinned by AI. Many teams have identified key use cases for public services to adopt, such as *Re:State's* six priority areas from *Getting the machine learning* in 2024.[25] But as technologies evolve, a top-down approach may not be responsive enough to keep pace.
2. Bottom-up discovery. Providing staff with access to AI tools and allowing them to experiment, finding new ways to improve their work at a personal level. Often supplemented by support from 'forward-deployed engineers' (FDE), embedded within teams and building software to solve specific local pain points, such as many of the use cases pursued by the No10 Innovation Fellows. This could lead to a proliferation of approaches that are ultimately incompatible with each other.

In practice, both approaches are being tried across central Government, with neither emerging as a dominant approach. Top-down, many departments are piloting systems, either led by dedicated AI units which operate cross-government (such as the Incubator for AI or the Customer First team, both in the Government Digital Service), or by departmental teams like the Justice AI Unit. And bottom-up, many departments have had Microsoft Copilot available as part of their Office 365 licences since 2024, evaluating the time saved by civil servants on routine tasks.

But as practiced, both routes have clear drawbacks, and neither is likely to reflect the full potential of AI across public services.

Top-down approaches have often led to a kind of 'pilotitis': a growing number of pilots that do not translate into fundamental changes to how government operates.[26] Backed by insufficient investment to scale, and lacking the confidence of leaders that putting AI at the heart of frontline government operations will be operationally feasible, it can leave teams in Whitehall stuck in a cycle of experimentation without the means or mandate to scale what works.

[25] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.
[26] ibid

While bottom-up approaches may improve staff productivity, it seems unlikely that this will ever unlock savings in the biggest areas of cost, such as frontline caseworking operations, where a structured approach to adoption is key to retaining public confidence and embedding sustainable changes. Many of the benefits of adoption will also be harder to capture for departments themselves: for example, from the productivity benefits flowing to individuals and enabling them to deliver the same level of output while working much less, or much higher burdens on managers from quality-assuring work which is delivered mostly by AI. Individual officials may become more productive, but the organisation may not capture those gains because services, processes, staffing, and performance measures remain unchanged.

A more effective model would combine the strengths of central coordination with greater autonomy for teams closer to delivery and introduce an additional approach:

3. Middle-out deployment – Providing technical AI teams with the resources and the powers to identify good practice in key areas and scale it up, drawing down on central resources and budgets as they demonstrate progress. Unlike a top-down approach, where a central management team selects use cases and commissions projects, the teams closest to delivery would identify where AI can have the greatest impact, test solutions in practice, and build the case for scaling them. They could draw on central resources and budgets as they demonstrate progress, in return taking on the key oversight and approvals for new systems from other parts of the DDaT function that currently manage these.

Idea 9: the Government Digital Service should codify and roll out a 'middle-out' deployment process for AI, led by technical teams funded centrally in each department.

To support this middle-out deployment, and building on the No.10 Innovation Fellowship model, Government should recruit more multidisciplinary teams of engineers, data scientists, designers, product managers, and transformation specialists. The No.10 Innovation Fellowship demonstrates the value of a coordinated central approach to hiring and deploying this kind of talent across government – a route which is already expanding and scaling up to provide more technologists with a route into public service.

Case study 3: The No 10 Innovation Fellowship

The No. 10 Innovation Fellowship programme was created by No.10 and brings world-class talent into government for short, high-impact tours of duty. The programme recruits AI engineers, data scientists, software developers, and product specialists on 6-to-12-month stints across government focusing on priority projects and working closely with senior officials and delivery teams. They bring skills from industry that are often scarce in government, such as machine learning, software engineering, and product management. The model is deliberately centrally managed: fellows are recruited and managed at the centre but deployed into departments to work on live problems.

The Fellowship has significantly helped accelerate specific AI and data projects, brought modern technical practices into government teams, and created a cohort of people who understand both the public sector and advanced technology. It has also shown that the centre

can attract and deploy technical talent quickly, rather than waiting for permanent recruitment or long procurement cycles.

For example, Fellows have worked directly with frontline justice teams to apply AI to operational problems. William Tarr, an Innovation Fellow, developed AI tools to automate bureaucratic tasks and support staff-prisoner engagement. This illustrates how the Fellowship can put highly technical capability directly into government teams, allowing promising ideas to be tested and developed in an operational setting rather than waiting for lengthy recruitment or procurement processes.

The scheme shows that it is possible to import exceptional talent into the centre rather than relying on limited departmental capacity to tackle every challenge. The case supports the argument that government can make progress with a stronger central technical function but also highlights that short-term fellowships are only a partial solution. For an AI-ready State, the UK may need to scale this into a more permanent, embedded capacity rather than a small, elite programme.

Source: Lucy Fisher, *'The "innovators and disrupters" hired to bring AI to the UK public sector'*. Financial Times, 2026

3.3 The role of agents

One area the Government seems far from a route to scale is the use of AI 'agents' within government. There is no clear path to taking a small number of internal pilots to scale, and without clear governance and rules it is unlikely that one will ever emerge. Interviewees were near-unanimous in believing that agentic uses of AI were the single biggest opportunity for the public sector, but also the furthest from reach because of governance, policy and sometimes legislative uncertainty on the precise role that autonomous software can have in public sector organisations.

There are some examples of early experimentation. The Department for Science, Innovation and Technology has been working on agentic AI for GOV.UK since March 2025, when it published a preliminary market engagement notice inviting frontier AI companies into a "National AI Tender".[27] The programme follows the scan, pilot, scale framework from the AI Opportunities Action Plan, with go/no-go gates between stages. Anthropic announced in January 2026 that DSIT had selected it to help build and pilot the assistant, starting with employment services.[28]

Departments have begun local agent deployment as well. HMRC's first Chief AI Officer told a conference in April 2026 that the department had rolled out 28,000 Copilot licences and was days away from switching on the agents functionality in Copilot Chat,[29] with a specific agent under development to summarise customer complaints and guide the handlers who deal with them. And DEFRA built its new digital service manual using agentic design and engineering, with AI acting as an autonomous software agent that took high-level user needs, chose

[27] Department for Science, Innovation & Technology, *National AI Tender for Frontier AI Companies to Support Agentic AI Powered GOV.UK Chat - Pre Market Engagement* (Find a Tender, 2025).
[28] Anthropic, *Anthropic Partners with the UK Government to Bring AI Assistance to GOV.UK Services* (2026).
[29] Horton, *HMRC Rolls out 28,000 Copilot Licences as AI Chief Targets World's 'Most AI-Enabled Tax Authority'*.

technical approaches and iteratively refined the code while human designers held editorial control.[30]

The leadership of GDS has identified agentic AI as the area "where we've got the biggest work to do – but [it offers] the most potential".[31] But there is no roadmap from these isolated use cases to deployment of agents at scale.

The questions agents raise are the ones that government departments are not yet prepared to manage. When an agent may take an action with legal or financial effect on a member of the public, as distinct from drafting something a human decision-maker approves, it becomes a question about administrative law and ministerial accountability. Government departments are likely to sidestep these tricky use cases in favour of more 'back office' implementations of agents. But the real value will come from the support that much more autonomous agents can give when operating in frontline services. Here the Algorithmic Transparency Recording Standard, used for transparency on automated decision making in government, was built for algorithmic tools that inform a decision, and doesn't fit agents that execute a multi-step workflow across the boundaries of systems and decisions.

Departments may try to fill the gap themselves, and bottom-up experimentation to work out how agents are best deployed should also be welcome. But progress is likely to be much slower than if different technical teams have central cover and a consistent policy on how to procure agentic software, orchestrate agents operating in government environments, monitor and evaluate their work, manage any security classification and data protection issues, and comply with legislation on automated decision-making. Instead, the Government Digital Service should launch the Trusted Agents Service – a central programme building the right policy and assurance mechanisms for AI agents in Whitehall. As one interviewee suggested, it should stress test and evaluate the use of agents in sandbox environments, kite-mark particular products along with publishing notes on how to integrate them into government workflows, and publish minimum standards on oversight and assurance for departments to comply with.

Idea 10: The Government Digital Service should establish a Trusted Agents Service as a key part of cross-government digital policy.

[30] Mark Chillingworth, *Defra Creates New Service Manual with AI* (2026).
[31] Sam Trendall, *GDS Leader Picks AI Agents as Area with 'Biggest Work to Do – but Most Potential'* (2026).

4. Procurement

For AI to transform the State, government will have to build a fundamentally different relationship with the companies provisioning the models and infrastructure which combine to create this new technology. It is unlikely that the Government will build much of the core AI 'stack' itself – frontier models continue to develop at incredible speeds, driven by a 'scaling law' which relies on massive amounts of private investment on a scale that the public finances cannot support. Instead, AI will mostly be bought from a fast-growing economy of companies selling different solutions to integrate the use of generative models into large and complex organisations.

While a small number of companies can build language models themselves, more provide the infrastructure and applications which will combine to let organisations use AI at scale. Government can provide some parts of this stack itself, particularly in the 'application' layer which provides the public interface with services like 'Gov Chat', but overall it will require large and ongoing procurement of software from a market which is developing faster than most Whitehall has dealt with.[32]

Most public procurement exercises are set up to deliver certainty and stability for suppliers and contracting authorities via long-term, multi-year contracts which give both parties clear economies of scale. But a fast-developing market requires much more flexibility to adapt to new products than these processes are designed for, and to buy these from markets which are often much less competitive than other parts of the digital economy. Specifications created up front for a procurement exercise are not likely to be relevant benchmarks for AI performance in six to twelve months, because the capabilities will have grown significantly in that time – Stanford's 2026 AI Index identified how frontier models had improved 30 percentage points in a single year against 'Humanity's Last Exam', one of the key benchmarks, and already regularly outperform humans.[33]

If Government buys AI the way it bought the last generation of technology, it will get the same results: becoming locked into outdated capabilities, at prices driven by incumbents and their market power, and with 28 per cent of digital systems classified as legacy.[34]

4.1 Buying the building blocks – procuring generative AI

One of the most striking differences between government's stated policy on AI adoption and its practice of it is the lack of availability of many of the basic building blocks of AI – in particular, access to language models. Chapter 5 of this paper highlights the widespread use of AI models in government, but also that much of the use is of externally-available models, often in ways which are not approved of in government policy and raise challenging questions about oversight and accountability.

It would be preferable for civil servants using AI, as they are encouraged to do so by the Government's own plans for a more automated civil service workforce, to be using tools provided by their employers. But interviewees routinely fed back that government access to language models was incredibly challenging, particularly outside of specialist technical teams, and also identified that uncertainty around the right procurement approach was the primary

[32] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.
[33] Stanford Institute for Human-Centred AI, *AI Index 2026* (No date.).
[34] Department for Science, Innovation and Technology, *State of Digital Government Review*, January 2025.

barrier. They listed the following concerns which procurement teams raised when it came to procuring AI:

1. That commercial teams, often organised 'vertically' around business areas rather than 'horizontally' around categories like AI, were unsure how to buy the best tools.
2. That procurement regulations (defined mostly under the Procurement Act 2023) required long competitive exercises to identify a preferred supplier to demonstrate value for money, and evaluating quality would be challenging when comparing similar enterprise tools.
3. That procurement of generative AI should be done at a different level to an individual team – for example, at a departmental level or a cross-government level.
4. That there was a lack of clarity on whether to buy generative AI tools in isolation from partners, through developer environments or contracts which offered system integration support, and whether choosing to use an AI tool through another supplier required an entire new procurement.
5. That there was a lack of clarity on how to avoid supplier lock-in and address data sovereignty which are more relevant to cloud-native technologies like generative AI than other categories.

These are all realistic concerns. While there is now an AI Dynamic Purchasing system that provides access to AI technologies from Frontier model providers like OpenAI and Anthropic, these providers are still not consistently and directly listed as suppliers in the other main procurement frameworks for accessing digital technology: G-Cloud 15, Technology Services 4 or Digital Outcomes and Specialists 7.

Case study 4: Memorandums of Understanding with frontier AI companies

The UK Government signed Memorandums of Understanding (MoUs) with several frontier AI companies: Anthropic (February 2025), OpenAI (July 2025) and Google DeepMind (December 2025). All commit the parties to work together on public sector AI adoption, but they are all explicitly not legally binding, and contain the commitment to "not prejudice against future procurement decisions". In practice, these commitments have no impact on the accessibility of the tools provided by each frontier AI company to government departments.

The contrast with the United States' approach at a federal level is striking. In the same year the UK Government signed these terms, the US General Services Administration (GSA) announced OneGov agreements to access ChatGPT Enterprise for nominal fees – ChatGPT Enterprise (at the value of $1 per agency per year), Claude at the same value, and Gemini for Government at 47 cents.

Source: Department for Science Innovation and Technology, *Memorandum of Understanding between UK and OpenAI on AI opportunities*, July 2025; Department for Science Innovation and Technology, *Memorandum of Understanding between the UK and Anthropic on AI opportunities*, February 2025; Department for Science Innovation and Technology, *Memorandum of Understanding between the UK and Google DeepMind on AI opportunities and security*, December 2025; Miranda Nazzaro, Anthropic, Google and OpenAI land GSA contract for governmentwide use, August 2025.

This lack of clarity on the right approach to procuring access to generative AI models is unsustainable, and institutional paralysis on identifying the right approach is standing in the way of making any significant progress.

The Digital Commercial Centre of Excellence in GDS has an opportunity to clarify many of these issues for contracting authorities in central government, giving them a route to market for generative AI which they currently lack.

Rather than adapting existing procurement routes, a new route for procuring AI should be designed from first principles. Based on discussions with interviews and the practice covered in this paper, the key principles of this route should include:

- **Specificity** – government bodies should aim to procure access to generative AI tools either directly from providers (such as Anthropic, OpenAI or Google), cloud supplier contracts (like Amazon Bedrock, Azure OpenAI, Vertex) rather than through larger contracts which combine access to those tools with service integration support. These capabilities require very different commercial models and approaches, and contracts which do both are unlikely to be appropriate to either.
- **Proportionality** – though generative AI spending could become a big category for government, individual tasks and token requirements are comparatively cheap. Frontier AI companies don't require long-term commitments from buyers, because they offer a relatively standardised product. Government should optimise for a proportionate approach based on experimentation with smaller commercial commitments to individual suppliers for short periods of time with no ongoing commitment, rather than creating large, bureaucratic and long-term contracts which are disproportionate to the risk involved.
- **Competition** – though there are a relatively small number of providers of generative AI models, the market is still more competitive than the market for many government contracts (where one study found that 20 per cent of large contracts which are openly competed still only received one bid).[35] Unlike many technologies, generative AI is very easy to access without extensive integration (meaning buyers don't need suppliers to 'bid' and then integrate, they can just pick up tools and use them) and are easily interoperable – buyers can swap out between models relatively easily, including when it comes to integrating those models into government digital services. Concerns about dependency on any one supplier are not realistic, provided the Government does not award a single, large 'winner takes all' contract to any one supplier for multiple years.
- **Quality** – point-in-time assessments of the quality of individual AI models at particular tasks are not a good measure of quality. Models are constantly improving in different ways against multiple benchmarks. Rather than using a single point-in-time assessment of quality to award large contracts, contracting authorities should award multiple companies contracts, paying them based on usage, and let technical teams evaluate multiple tools at once and decide which to spend budget with as they go. This applies the principle of 'consumer choice', where users in government are the consumer, to identify which product is the best quality for a given set of tasks.
- **Cost** – token costs are likely to change significantly over coming years, and government should aim to secure good value for money on AI contracts without locking itself to paying for any one individual provider. Ongoing competition between companies will be essential to keeping costs down, reinforcing the need to use multiple suppliers at once rather than awarding contracts to any one supplier.
- **Sovereignty** – the Government should set an overall policy on digital sovereignty,[36] but in the absence of one, Government departments should follow a proportionate approach to evaluating data sovereignty concerns – applying the same tests they

[35] Public Accounts Committee, *Competition in Public Procurement* (2023).
[36] Matthew Feeney, *Sovereignty by Design: Security, Resilience and Growth* (2026).

would for deciding whether to use cloud storage of information at different levels of sensitivity (personal or non-personal data, and security classification).

- **Scope** – use of generative AI in government comes with many operational considerations, including operating risk, reputational risk, and financial management of usage costs. Although these are all important, they should not all be solved through a procurement approach – appropriate oversight of these tools into the business with separate policies and processes is more appropriate. Procurement processes for AI should be focused on getting officials the best performing products at the best price, not compensating for all the different ways those products might be used in government.

This would not require new legislation. The Procurement Act 2023 created dynamic markets, which suppliers can join at any point rather than at a single tendering moment, and open frameworks, which have to be reopened to new entrants during their life rather than closed at the beginning. Both were designed for markets where the list of suppliers and products evolves fast, such as frontier AI. The Government should create a worked example of how these provisions can be applied to frontier AI models: a standard commercial route, published centrally, which a commercial team in a medium-sized department can pick up and use without commissioning fresh legal advice.

That route should look very little like a normal government technology framework. Qualification should be an initial test of security, data handling and contractual terms, open continuously to any provider who wants to take it, and every provider who passes should be awarded. It should include standard schedules for pricing agent-based, transaction-based, consumption-based and outcome-based work alongside labour-based rate cards, giving contracting authorities and suppliers more flexibility to keep pace with a changing market.

Departments should then draw down on that framework based on actual usage, with no minimum spend and no exclusivity, and technical teams should be free to move budget between providers month to month rather than only after a two-year term and extensive evaluation. Existing open frameworks could also be amended when they reopen to apply learnings from this new commercial route more widely.

Idea 11: the Government's Digital Commercial Centre of Excellence should create a central route to market for the specific procurement of frontier AI software, available to all contracting authorities and paid for based on actual usage.

4.2 Laying the foundations – procurement for the AI age

But the technology stack which supports an AI-native State is more complex than simply getting departments access to the foundational models themselves. Although that is the most notable absence from the current procurement strategy, the government will always need to rely on a range of different products and services to deliver the digital capabilities it cannot feasibly build in-house.

Buying for a multi-vendor market

The cost of dominance of markets by a small number of suppliers is well-documented in Whitehall. Procurement markets are often highly uncompetitive. But it often persists because the incentives are for procurement teams to encourage reliance on single suppliers. The commercial officer who consolidates six requirements into one supplier gets a simpler contract, a single account manager and needs to develop fewer business cases. But the cost of that resulting dependence often falls three or four years later. This risk is just as true in the age of AI – the model layer, the infrastructure layer and the application layer are separable in a way that the last generation of enterprise software was not – and a contract that bundles all three back together throws away this structural advantage.

Idea 12: for AI-enabled and platform-based services, contracting authorities should be required to design requirements so that more than one supplier can meet them, by separating the model layer from the application layer, and identifying a credible second source before award.

Open standards and interoperability

Designing for multiple vendors only works if leaving one is technically possible in practice. The Government does not need legislation to give itself these protections, because it is the buyer and can simply require them. Guidance already asks departments to think about open standards and exit planning. But this requires departments to know what to ask for: the Government Digital Service should break that dependency, by creating standard interoperability terms which can be included in every contract.

Idea 13: GDS should mandate a standard exit schedule in every central government digital contract above a set value, covering export of data and configuration in documented open formats, no charge for extraction, defined notice, transition and retrieval, and a supplier obligation to assist migration. Departments should be required to rehearse migration of at least one significant workload during the contract term, and report the result to the Government Digital Service.

Crowding in new suppliers

A typical challenge for the UK's procurement policy is that it makes it incredibly hard for new companies to bring new products to market in the public sector.[37] Direct awards to individual companies are frowned on, and only permissible in a small number of cases under the Procurement Act 2023. This means that most procurements have to start with a government-initiated 'requirement' – making it almost impossible for them to buy things they don't know they need.

This challenge of partnering with government is compounded in difficulty when it comes to AI because most AI services will be offered through enterprise software business models – i.e. on paying recurring fees to integrate and use a 'product'.[38] But this is a challenging fit for

[37] Joe Hill and Sean Eke, *Procure and Simple: A Proportionate Approach to Public Procurement* (2025).
[38] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.

government procurement approaches, which are still mostly designed around 'Time and Materials' contracts for buying time-limited services from a company – a model which rules out a whole host of companies from working with government.

In September 2026 the UK's Sovereign AI Fund ("Sov/AI") launched a "Sovereign AI R&D Procurement Scheme", which provides contracts for companies wanting to support the government in developing sovereign AI capabilities in four challenge areas – including the National Health Service and defence. The scheme promises "partnership with public sector users, establishing a direct supplier-customer relationship"[39]. Challenge-based approaches are a welcome way of crowding in new suppliers who may have a different offer in mind to the government departments who would ordinarily contract for something else,[40] but the approach will only work sustainably if it provides a route to scale for those companies. The Scheme is currently only limited to R&D procurement – focused on "pre-commercial research and testing". The Government should expand this challenge-fund based approach to technology at a higher level of readiness which is ready for wider use.

Idea 14: GDS should launch challenge funds for larger scale-up companies targeted at public service challenges, and based on buying software products rather than solely services.

[39] Sovereign AI, *Competition Guidance: Government as Your Customer* (2026).
[40] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.

5. Workforce, skills and adoption

The adoption of generative AI will have profound implications for the civil service workforce. While the fundamental building blocks of AI services will have to be procured by Whitehall, the rate-limiting step to deploying them will likely be its own staff – their skills, capabilities and how their teams are structured to best drive adoption. And while the route to adoption will rely disproportionately on the skills of technologists, successful adoption will have widespread implications for the whole workforce – hugely increasing staff's own productivity, and opening up routes to widespread automation. Done well, this should make central government much more productive, but done badly the returns to higher productivity won't be felt in public budgets, and will instead be absorbed by individuals being able to achieve the same amount while working less.

Both challenges require a strategic approach to the workforce which is largely absent from government thinking, and insofar as workforce AI policies do exist, these are insufficient to the scale of the challenge.

5.1 Capabilities to support adoption

Integrating, testing, monitoring and improving on AI-driven software in the public sector at any scale will require a significantly better-resourced government digital, data and technology (DDaT) profession, along with a broader base of knowledge and skills in the rest of the civil service. The National Audit Office found that 70 per cent of the 87 government bodies they surveyed in 2024 were struggling to recruit and retain staff with AI skills,[41] and since then multiple interviewees for this paper confirmed there had been a widespread departure from the government of many specialists with AI backgrounds.

This is partly underpinned by lower salaries in the public sector, where technical specialists can earn much more in private companies,[42] but also by ways of working. As one interviewee for a previous *Re:State* paper said of their staff (at a private company) "you'd have to pay [their own staff] even more to work in government",[43] based on their sense that the public sector often misuses the skills of technical specialists, who don't want to work in heavily compliance-driven areas where they can't use their skills as freely to develop and implement software.

Clearly the factors which prevent government accessing technical AI talent are complex, but there should be no doubt that workforce capability is essential to deploying AI, along with adaptation to the unique working patterns and different circumstances in each organisation. Consider the published evaluations of Microsoft's Copilot tool – the Government Digital Service's cross-government evaluation found that civil servants saved an average of 26 minutes a day,[44] the Department for Work and Pensions found an average saving of 19 minutes a day,[45] and the Department for Business and Trade found no evidence that any time savings actually improved productivity.[46] The exact same software in different settings, with different staff, shows varied results.

[41] National Audit Office, *Use of Artificial Intelligence in Government* (2024).
[42] Department for Science, Innovation and Technology, *State of Digital Government Review*.
[43] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.
[44] Government Digital Service, *Microsoft 365 Copilot Experiment: Cross-Government Findings Report*.
[45] Department for Work and Pensions, *An Evaluation of DWP's Microsoft Copilot 365 Trial* (2026).
[46] Department for Business & Trade, *The Evaluation of the M365 Copilot Pilot in the Department for Business and Trade*.

Interventions to upskill the workforce which are targeted at the whole civil service, such as the One Big Thing campaign which reached 212,000 civil servants in 2023,[47] are unlikely to do more than marginally raise awareness, including in many roles which are not crucial to AI adoption. AI seems to be a focus for the new National School of Government, although it's unclear how this (mostly externally-delivered) training will be applied.[48] Rather than whole-civil service interventions, there are a small number of key areas where the absence of internal AI capability is the main bottleneck to deployment and needs targeted investment.

Interviewees were clear that one of the biggest bottlenecks was senior technical talent, paid to a level which was commensurate with their skills. One advocated for hiring staff at more senior levels to get around grade-level pay restrictions, even if technical talent hired that way in practice ended up reporting to more junior colleagues in the grade structure. Other workarounds exist – previous *Re:State* research has suggested centralising the set of allowances which staff can attract for technical skills into a single approval point, and making these portable between internal roles,[49] but take-up has been low because changing the system for existing allowances across departments is hard, and interviewees fed back that the prospect of higher salaries after applying for allowances was less attractive to external candidates who felt they could go through the whole process only to be offered the basic salary at that pay point.

Idea 15: the Government should launch an 'advance talent commitment' for AI, committing to hire up to 400 staff with machine learning and AI skills on salaries of up to £200,000, across government, with a single point of approval in the Government Digital Service for all salaries. The budget for these roles would come from within hiring departments.

Alongside removing the friction to external entrants into government digital roles, the Government should find a new structure to explicitly identify and attract technical talent into Whitehall. External job adverts are frequently opaque, particularly to technical talent,[50] and as one interviewee explained there is limited support with navigating the system or identifying key hires and encouraging them to apply. One interviewee mentioned the Tech Talent Project in the United States, which was founded in 2017 by a cross-party group of officials to bring better technical leaders into government. Today they bolster a network of 10,000 professionals, and have placed over 200 leaders into government roles.[51] An initiative like this built from within the Government would likely fail to find the right networks, but if an industry partnership brought it together then that would provide a better catalyst for interested talent to engage with. The Government should support private sector partners in developing an equivalent project in the UK, for attracting and placing candidates into central government roles.

Idea 16: the Government should work with industry partners to launch a Tech Talent Initiative for hiring experienced technology leaders into the private sector.

[47] Department for Science, Innovation and Technology, *State of Digital Government Review*.
[48] Cabinet Office, *Centralised Civil Service Learning and Training Services* (2026; Find a Tender, No date.).
[49] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.
[50] Sean Eke et al., *Making the Grade: Prioritising Performance in Whitehall* (Re:State, 2024).
[51] New Venture Fund, *Tech Talent Project*.

5.2 A workforce transformed by AI

The scale of central government, employing over 549,000 staff,[52] provides an unprecedented arena for productivity and efficiency improvements from AI adoption. If it were a private company, the civil service would be one of the largest in the world, and although its operations are very different to those of any individual company, many of its functions perform similar tasks – customer service, caseworking, payment processing, advisory services and many more functional services which cut across all departments (such as human resources).

Quantifying the scale of potential workforce benefits from AI transformation is challenging, as the technology continues to advance and there is no clear single model for reckoning with the benefits or making sure the civil service captures the advantages of any deployment. The *State of Digital Government Review* described more than £45 billion of potential annual savings and productivity benefits, though little information on the methodology behind this number was published. The Government Digital Service's internal analysis of civil service roles specifically found £6.3 billion of savings, made up of £1.1 billion of cost reductions and £5.2 billion of productivity gains (i.e. non-cashable estimates of improved performance).[53]

Other organisations have attempted to estimate the 'size of the prize' using top-down methods – Public First put the total at £17 billion by 2035,[54] based on a similar methodology to the GDS approach: classifying combinations of tasks in civil service roles depending on how automatable they could be. And the Tony Blair Institute estimated up to £40 billion a year after discounting the investment required to deliver those savings.[55]

The range and spread of these estimates shows how hard it is to have confidence on the scale of the productivity benefits from AI. But whatever the scale, one concern discussed by several interviewees was the extent to which the benefits will be captured by staff themselves, through working fewer hours, rather than being translated into a wider workforce and commercial strategy to maximise productivity, efficiency and service quality.

Despite official guidance to carefully deploy AI, use is already widespread, as shown by research from the Alan Turing Institute in 2024.[56] *Re:State's Alternative People Survey* in 2025, which collected responses from 1,265 civil servants across Whitehall, found 30 per cent were using AI in their work, with 59 per cent of those using externally-available AI tools at least some of the time.

But one barrier to AI adoption throughout central government that was frequently raised by interviewees was capacity in teams to test and evaluate tools and workflows. Often, it is the same teams where AI would have the most benefit that are also the most strained – not just operational and frontline teams, but also policymakers and programme teams. While Whitehall may have excess staff in other areas, often they won't have the right skills and specialist knowledge to support when there are shortages in these areas.

There is an acute risk that the productivity benefits of AI aren't generated where they are most needed because of this capacity constraint, and that paradoxically the Government gets the most testing and evaluation of AI in the areas it least needs it. GDS should be prepared to pay

[52] *Statistical Bulletin - Civil Service Statistics: 2025* (No date.).
[53] Government Digital Service, *Digital and Data Benefits Framework* (2026).
[54] Microsoft and Public First, *Unlocking the UK's AI Potential* (2024).
[55] Alexander Iosad et al., *Governing in the Age of AI: A New Model to Transform the State* (Tony Blair Institute for Global Change, 2024).
[56] Alan Turing Institute, *Generative AI Is Already Widespread in the Public Sector* (2024).

overtime to frontline staff in return for dedicated testing of AI products, to a high standard which can be used to inform the middle-out adoption priorities described in Chapter 3.

Idea 17: The Government Digital Service should create an AI Testing Fund to purchase additional capacity from teams across government to test AI in their areas through paid overtime, as part of structured evaluations designed to inform the roll-out of new, more productive ways of working.

The civil service has a poor track record at reducing its own headcount. Previous research has found that even during recent periods of recruitment freezes, the civil service still grew in size.[57] Recent headcount reductions pursued by the civil service through voluntary exit schemes have not been sufficiently targeted, removing staff from many key areas and enabling them to remain in areas which are not priorities for the current government.[58]

To prevent a repeat of these failures, the Government Digital Service will need to take a strategic approach to automation which includes clear plans for headcount reductions. Departments should use the results of AI testing to identify how automation will affect future staffing levels, skills, team structures and recruitment requirements, and incorporate these changes into their workforce plans.

To be clear, headcount reductions are not the only source of benefits from AI adoption – productivity improvements for staff, able to deliver more for the public, and similar improvements in public-facing services which make them faster and cheaper are also valuable. But given Whitehall's poor track record on strategic workforce planning, it will be important to incorporate the impact of AI into future strategic workforce planning processes. Workforce planning should therefore distinguish between cashable savings, productivity improvements and service enhancements, so that the Government can identify where benefits are being generated and take steps to capture them.

This should also be linked to commercial reform. Changes in the organisational structure of work as a result of AI should inform decisions about what departments need to buy from external suppliers, as well as the skills and capacity they retain in-house. Where work can be defined clearly and performance measured consistently, departments should look for opportunities to commission suppliers against the work delivered or service provided, rather than simply purchasing staff time. This could include arrangements based on defined volumes of work, particular processes or agreed service standards. Where requirements are less predictable, or delivery depends heavily on specialist expertise, labour-based approaches may remain more appropriate. Aligning these commercial decisions with workforce planning would help ensure that the benefits of AI are reflected in how government manages external capacity, rather than simply allowing additional resource to be deployed to meet demand.

Idea 18: Departments should ensure workforce planning incorporates the potential efficiency gains from AI, including implications for staffing levels, skills, organisational design, recruitment and the way external capacity is commissioned.

[57] Jack Worlidge et al., *Whitehall Monitor 2025* (No date.).
[58] Joe Martin, *Voluntary Madness: Whitehall's Obsession with Exit Schemes Is Hollowing out the Wrong People While Dodging Real Reform* (Civil Service World, 2025).

6. Funding approaches

Funding is one of the clearest examples of mismatch between government's digital ambitions and its current operating model. AI-enabled platforms, shared data infrastructure and end-to-end service redesign require sustained investment over time. Yet current funding practices mean that budgets work on a departmental, project-by-project basis, which makes it difficult to invest in shared infrastructure like data platforms or AI tooling that delivers value across multiple services. Often, government departments try to fit digital projects into an outdated project-by-project funding model which assumes large peaks of investment followed by years of none – this is unrealistic, as many researchers have pointed out.[59] This model is poorly suited to shared digital capability.

The *State of Digital Government Review 2025* illustrates the size of the problem. It found that in 2023, the public sector spent around £26 billion on digital and data – around 30 per cent below peer benchmarks – and digital and data spend per full-time employee was estimated to be 78 per cent below peer benchmarks.[60] The Review also emphasises that current budgets make it difficult to maintain the existing technology estate, respond to legislative change, and keep pace with cyber threats.

At the same time, *A Blueprint for Modern Digital Government* notes that spending is biased towards new programmes, with insufficient priority given to maintaining and improving existing systems, particularly legacy systems.[61] This reflects another key pattern in government spending whereby funding is often made available for a crisis or for a politically visible new commitment, but less reliably for continuous improvement and preventative forward-looking maintenance of legacy systems.[62]

6.1 Systemic financial traps

Spending practices can also reinforce departmental boundaries and act as a barrier to cross-departmental transformation and whole system change. This means that there is appetite for policy change from ministers and senior leaders, but the delivery system is not designed to support it and instead acts as a barrier to funding the kinds of digital transformation that government has committed to.

The fragmentation in funding and departmental split creates systemic barriers to funding and a lack of accountability.

Examples of systemic barriers include:

- Departments are accountable for their own budgets, but the benefits of investment and transformation often stretch beyond their department. The different models outlined in *Managing public money* for departments to pool resources are rarely used in practice, and when investment in shared capabilities is made available centrally by HM Treasury, departments rarely pursue it.[63]

[59] Jennifer Pahlka, *Project vs Product Funding: Modernisation Shouldn't Even Really Be a Thing at All* (Eating Policy, 2024).
[60] Department for Science, Innovation and Technology, *State of Digital Government Review*.
[61] Department for Science, Innovation & Technology, Government Digital Service, *A Blueprint for Modern Digital Government*.
[62] Ganesharatnam and Hill, *From Legacy to Leadership: Upgrading the Digital State*.
[63] Public Accounts Committee, *Cross-Government Working*, HC 75 (The Stationery Office, 2024).

- Business cases tend to be structured around a specific time-limited programme rather than an ongoing capability.
- Audit and commercial systems seek clarity, scope, and ownership to be established early, while digital transformation requires adaptation as teams learn.
- Senior officials tend to be held more accountable for visible programme failure than for gradual service deterioration.
- The transition of digital spending from mostly capital to resource budgets, in line with the shift from building digital through time-limited projects to build and then run a capability to a model based on iterative development of products, is technically correct but has squeezed change budgets available for AI integration.

These systemic barriers mean that, as one interviewee described, the State maintains continuity of existing services by default. The Treasury's budgeting model was designed both for a different time and a different economy. It is still better suited to acquiring physical assets or delivering defined capital projects than to operating subscription-based platforms and iteratively-developed software services built largely by in-house teams. Technology spending has shifted from a capital-intensive model to a revenue-intensive one, underpinned by licences and credit or token-level spending year-to-year, yet central government's digital and data funding remains weighted towards capital expenditure.

This creates a structural mismatch between how modern digital services are built and how government budgets for them: teams can secure funding to build or procure a system, but face greater difficulty securing the predictable, multi-year resource needed to run, maintain and continuously improve it. This can encourage short-term projects over sustained service development, make it harder to retain technical capability in-house, and create uncertainty around recurring technology costs.

The *State of Digital Government Review* illustrates this problem, with only one in five survey respondents feeling that the current funding model enables effective investment in and running of digital services.[64] Other countries have taken different approaches. The Singaporean Open Government Products Team built significantly more central capabilities to avoid these kinds of systemic barriers to change. Funded through a block grant based on a single business case, and developing products internally on behalf of other government teams, they share all usage, cost and benefit information on a publicly available dashboard.

6.2 Funding for an AI-ready State

The funding approach for digital technology should reflect the different roles digital investment plays in a modern State. Some capabilities, such as digital identity and common AI tooling are national foundations: they enable multiple public services and should be planned, funded and governed accordingly. Other investments may be specific to a department or service, but may still create reusable components, shared learning or wider public value. On top of these, government must fund the ongoing work of operating, securing and improving live services. To break open the systemic barriers to transformation contained in the current budgeting

[64] Department for Science, Innovation & Technology and Government Digital Service, *State of Digital Government Review*.

framework, the Government needs a new approach. This should distinguish between three different kinds of digital spending:

- National infrastructure – which should be centrally funded because it supports multiple services and departments.
- Departmental transformation – which should remain locally owned but receive central support where it creates reusable capability or cross-government value.
- Continuous service operation and improvement – which should be treated as a core revenue cost rather than an optional programme allocation.

Without this distinction, government will continue to mix infrastructure, transformation, maintenance and business-as-usual activity into the same short-term funding competition. The predictable outcome will be that the most visible new programme wins while the foundations on which every service depends will continue to deteriorate.

Idea 19: Shared platform services (identity verification, data exchange, common AI tooling etc) should be funded centrally, where these services create value across multiple departments and should not depend on any individual department's budget allocations.

These funding challenges point to a clear conclusion: government cannot make serious progress on AI and digital transformation while treating shared capability as a series of isolated departmental decisions. The current model pushes departments to prioritise their own immediate pressures, even where the most effective investment would be one that benefits the whole system. That is why a different mechanism is needed.

A cross-government Digital Investment Fund would give HM Treasury a way to back infrastructure and transformation on the basis of cross-government value, rather than departmental convenience. It would help fund capabilities that are too important to leave to individual budgets, while encouraging departments to co-invest where they will directly benefit. In doing so, it would create a more realistic route for building reusable, strategic digital capacity across government.

Match funding would retain departmental ownership and ensure that proposals respond to operational needs, while reducing the disincentive to invest where costs are concentrated but benefits are distributed across government. The Fund should take a multi-year, portfolio approach, supporting adoption and reuse rather than treating each initiative as an isolated project.

Idea 20: HM Treasury should create a cross-government Digital Investment Fund. This should be ringfenced for shared infrastructure, reusable and strategic transformation projects. Departments should be able to apply for match funding when a digital or AI investment will benefit multiple parts of government, to encourage co-investment and shared ownership for initiatives with wider benefits.

7. Conclusion

Through AI the State has been offered a chance to radically remake the services it provides to the public, and it is still processing the implications of that offer while the rest of the world rushes to take it. The effectiveness of AI challenges the very processes the State uses to evaluate and deploy tools, because it changes fundamental assumptions about how software should be bought, funded and staffed across the public sector. A frontier model can complete a piece of casework for a few pence, but a department that wants to buy access to even test it feels the need to run competitions lasting months to establish whether this is good value for money. The Treasury scores the savings from automation before a single process has changed, then removes them from the budget of the team that would have to deliver those same savings. Analysts argue over the performance of AI-supported work in isolation, rather than comparing it to the performance of existing staff.

These are all decisions that leaders in Whitehall can and should take differently, but also decisions which the leaders of new devolved public services will need to take for the first time. In both cases, past choices and precedent will be a poor guide to how to make the best of a technology which is changing the world as quickly and radically as AI is.

The changes this paper proposes are within the gift of the current Government and require no new technological developments or legislation to be passed. It is not technically challenging to embrace AI in government, but failure to do so could prove to be administratively existential.

Adoption can proceed in any number of ways, influenced by Government policy, but it will proceed regardless. Civil servants are already using AI in their thousands – and many are using tools their employers did not provide and do not monitor. That is what happens when the centre declines to lead, and it is the same situation the next *State of Digital Government Review* will describe if the machinery of adoption stays as it is.

This report sets out some aspects of an alternative path – inverting many of the fundamental ideas about how government should work – should the Government decide to lead, rather than follow, the adoption of AI throughout the economy.

Bibliography

Alan Turing Institute. *Generative AI Is Already Widespread in the Public Sector*. 2024.

Anthropic. *Anthropic Partners with the UK Government to Bring AI Assistance to GOV.UK Services*. 2026.

Beck, Kent, and et al. *Manifesto for Agile Software Development*. 2001.

Cabinet Office. *Centralised Civil Service Learning and Training Services*. 2026; Find a Tender, No date.

Cabinet Office. *Digital Transformation in Government: Addressing the Barriers to Efficiency*. National Audit Office, 2023.

Chillingworth, Mark. *Defra Creates New Service Manual with AI*. 2026.

Coyle, Diane, David Eaves, and Beatriz Vasconcellos. 'Digital Infrastructure'. *IMF*, March 2026.

Department for Business & Trade. *The Evaluation of the M365 Copilot Pilot in the Department for Business and Trade*. 2025.

Department for Science, Innovation & Technology. *National AI Tender for Frontier AI Companies to Support Agentic AI Powered GOV.UK Chat - Pre Market Engagement*. Find a Tender, 2025.

Department for Science, Innovation & Technology, Government Digital Service. *A Blueprint for Modern Digital Government*. H M Government, 2025.

Department for Science, Innovation & Technology and Government Digital Service. *State of Digital Government Review*. H M Government, 2025.

Department for Science, Innovation and Technology. *State of Digital Government Review*. January 2025.

Department for Work and Pensions. *An Evaluation of DWP's Microsoft Copilot 365 Trial*. 2026.

Design community. *Designing Good Government Services: An Introduction*. GOV.UK, 2024.

Digital Transformation in Government: A Guide for Senior Leaders and Audit and Risk Committees. No date.

Digital Transformation in Government: Addressing the Barriers to Efficiency. No date.

Eke, Sean, Joe Hill, and Charlotte Pickles. *Making the Grade: Prioritising Performance in Whitehall*. Re:State, 2024.

Feeney, Matthew. *Sovereignty by Design: Security, Resilience and Growth*. 2026.

Ganesharatnam, Laxshia, and Joe Hill. *From Legacy to Leadership: Upgrading the Digital State*. Re:State, 2026.

Government Digital Service. *Digital and Data Benefits Framework*. 2026.

Government Digital Service. *Microsoft 365 Copilot Experiment: Cross-Government Findings Report*. 2025.

Government's Approach to Technology Suppliers: Addressing the Challenges. 2025.

GOV.UK. 'Government Design Principles'. 2 April 2025. https://www.gov.uk/guidance/government-design-principles.

GOV.UK. 'GOV.UK - 1 Year Old and Improving All the Time'. Accessed 7 September 2026. https://www.gov.uk/government/news/govuk-1-year-old-and-improving-all-the-time.

GOV.UK. 'Service Standard'. Accessed 8 September 2026. https://www.gov.uk/service-manual/service-standard.

GOV.UK. 'The Green Book (2026)'. Accessed 9 September 2026. https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government/the-green-book-2026.

GOV.UK. 'Use GOV.UK Domains and the GOV.UK Design System'. Accessed 7 September 2026. https://www.gov.uk/guidance/use-govuk-domains-and-the-govuk-design-system.

Greenway, Andrew, and Tom Loosemore. *The Radical How*. Public Digital, Nesta, 2024.

Hill, Joe. *Byte-Sized Budgeting: Funding Digital Service in Government*. January 2025.

Hill, Joe, and Sean Eke. *Getting the Machine Learning: Scaling AI in Public Services*. 2024.
Hill, Joe, and Sean Eke. *Procure and Simple: A Proportionate Approach to Public Procurement*. 2025.
Horton, Christine. *HMRC Rolls out 28,000 Copilot Licences as AI Chief Targets World's 'Most AI-Enabled Tax Authority'*. 2026.
Iosad, Alexander, David Railton, and Tom Westgarth. *Governing in the Age of AI: A New Model to Transform the State*. Tony Blair Institute for Global Change, 2024.
Martin, Joe. *Voluntary Madness: Whitehall's Obsession with Exit Schemes Is Hollowing out the Wrong People While Dodging Real Reform*. Civil Service World, 2025.
Microsoft, and Public First. *Unlocking the UK's AI Potential*. 2024.
National Audit Office. *Use of Artificial Intelligence in Government*. 2024.
New Venture Fund. *Tech Talent Project*. No date.
OpenAI. *The Hugging Face Incident and the Road Ahead*. 2026.
Pahlka, Jennifer. *Project vs Product Funding: Modernisation Shouldn't Even Really Be a Thing at All*. Eating Policy, 2024.
Public Accounts Committee. *Competition in Public Procurement*. 2023.
Public Accounts Committee. *Cross-Government Working*. HC 75. The Stationery Office, 2024.
Sovereign AI. *Competition Guidance: Government as Your Customer*. 2026.
Stanford Institute for Human-Centred AI. *AI Index 2026*. No date.
Statistical Bulletin - Civil Service Statistics: 2025.
The Challenges in Implementing Digital Change. No date.
Trendall, Sam. *GDS Leader Picks AI Agents as Area with 'Biggest Work to Do – but Most Potential'*. 2026.
Update on Government Shared Services. No date.
Worlidge, Jack, Teodor Grama, Jordan Urban, et al. *Whitehall Monitor 2025*. No date.

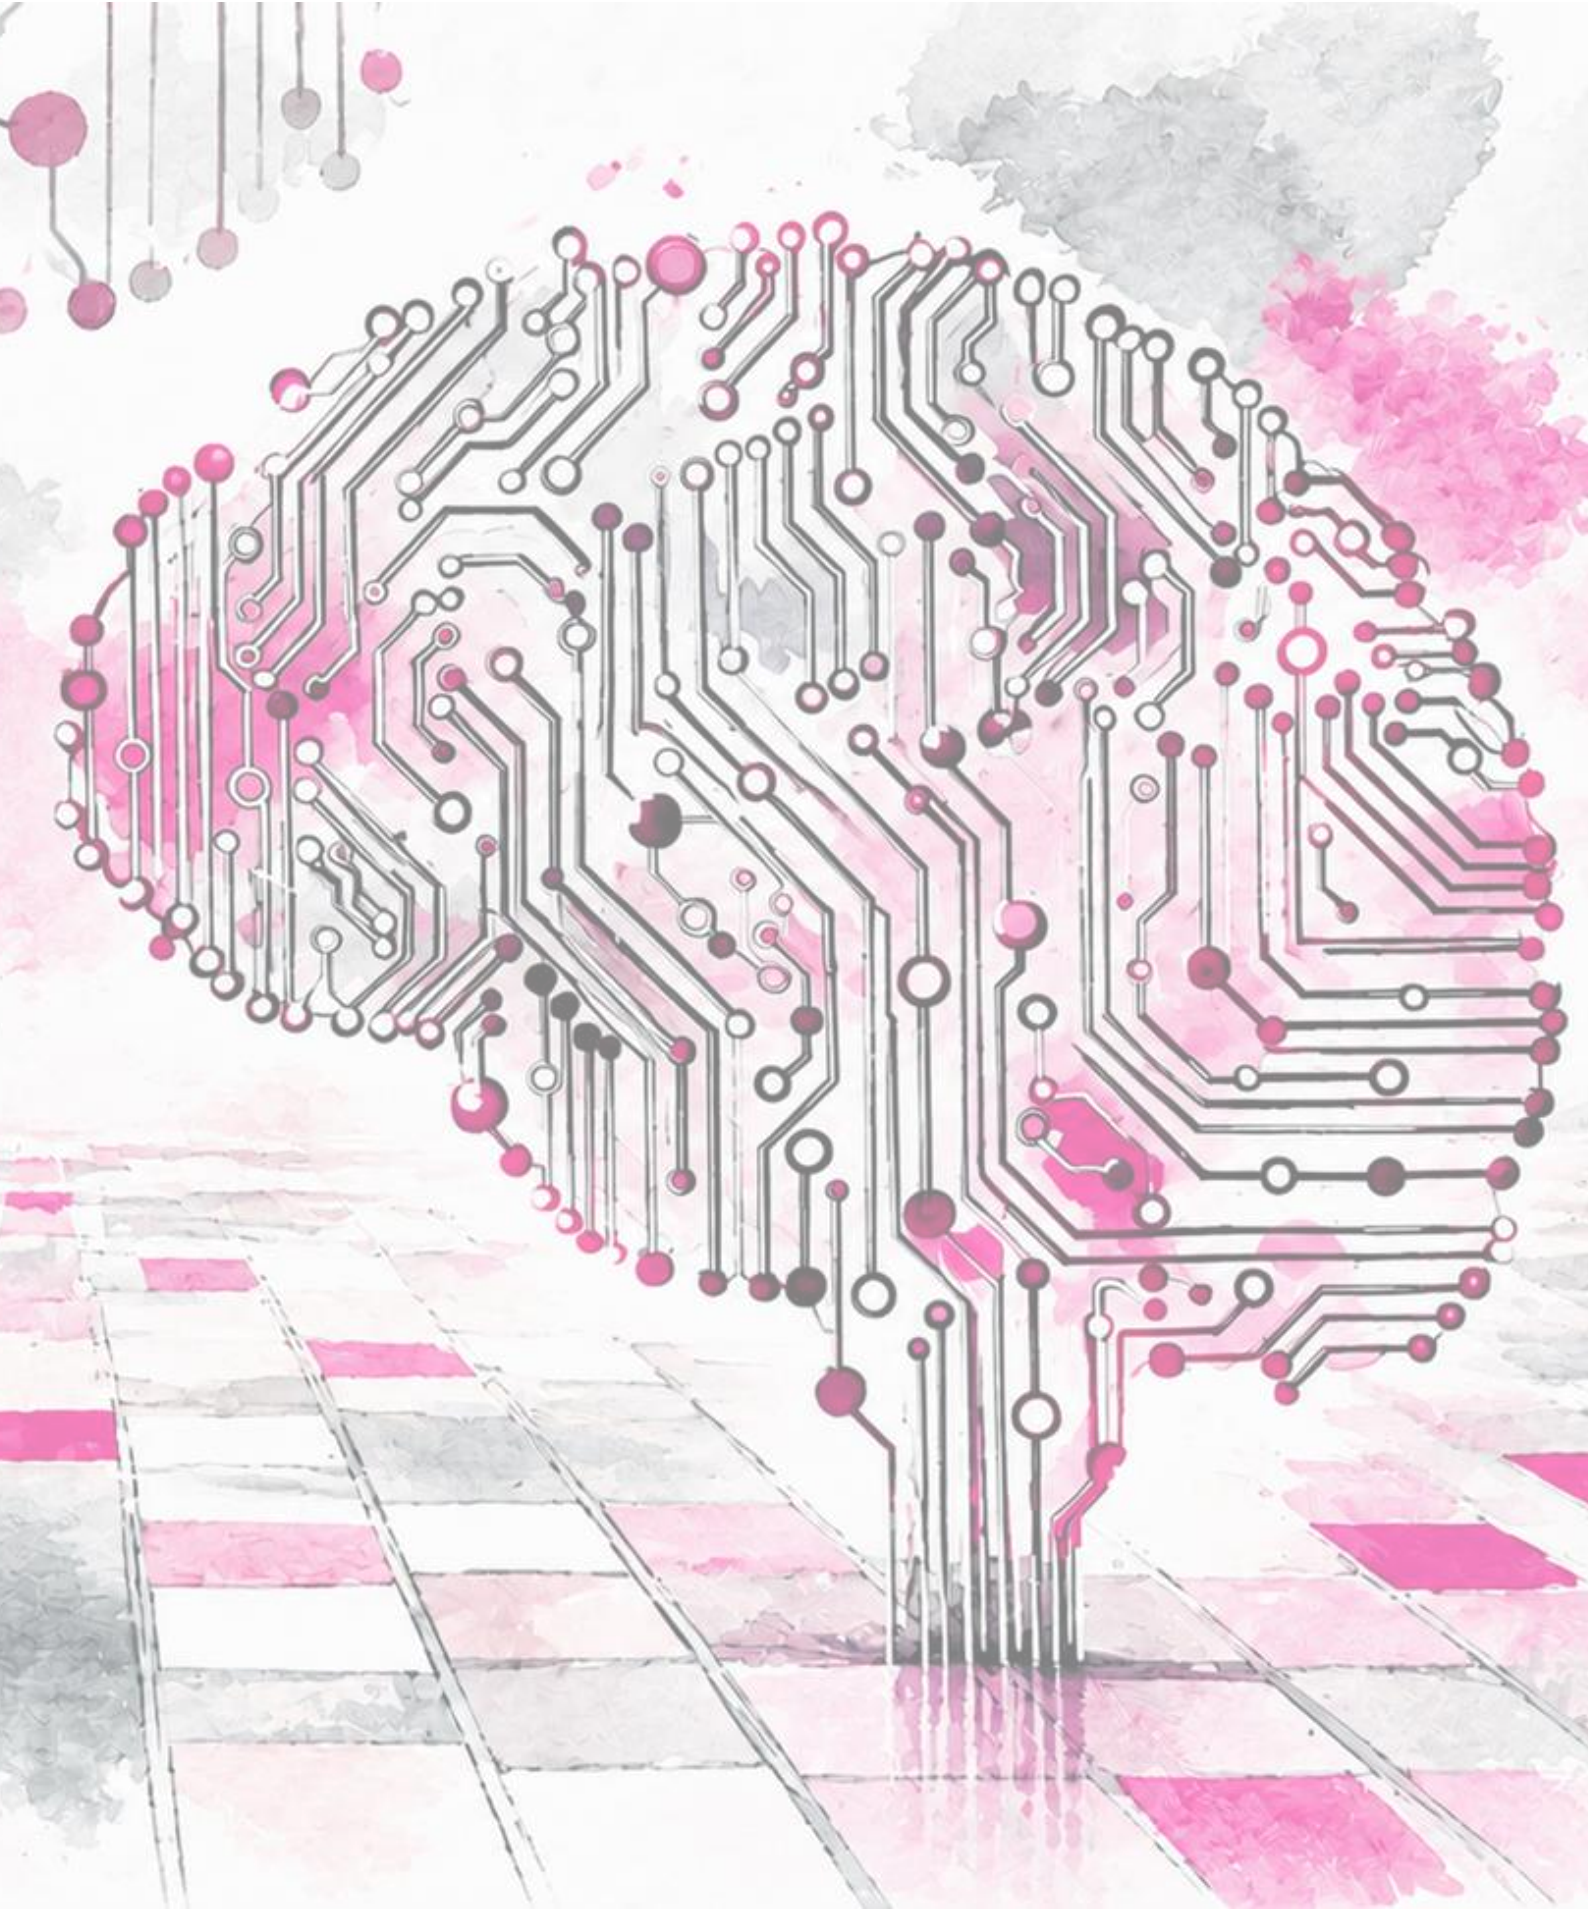

Re:State

ISBN: 978-1-918941-13-5

@restate_thinks

@re-state.bsky.social

www.re-state.co.uk