

# Digital Leaders Study 2024 ▶

The accelerator and the  
brake: embedding AI across  
government in the UK

GLOBAL  
**government**  
FORUM

**DIGITAL**

Supported by



# CONTENTS

3	<b>FOREWORD</b>
5	<b>INTRODUCTION</b>
7	<b>A NOTE ON METHODOLOGY</b>
10	<b>A NOTE ON AUTHORSHIP</b>
11	<b>CHAPTER 1: THE UK'S PLACE IN THE GLOBAL AI RACE</b>
12	The UK and Singapore
14	<b>CHAPTER 2: AI IN GOVERNMENT: PERSPECTIVES FROM THE UK'S DIGITAL LEADERS</b>
14	Setting direction around AI – vision, design and plan
18	Departmental drivers – collaboration and accountability
22	Developing an AI-ready workforce – transformation leadership and people
26	<b>CHAPTER 3: JUMP-STARTING THE AI REVOLUTION</b>
26	Three fundamentals
28	A roadmap for AI in government
30	Time to hit the accelerator
31	<b>APPENDIX</b>
32	Interacting with Redbox
48	<b>FURTHER INFORMATION</b>



# FOREWORD

by Cognizant, our knowledge partner for Digital Leaders 2024

Artificial Intelligence (AI) represents a potentially transformative shift for government – one that could revolutionise policymaking, national security and the delivery of public services. As we stand at the cusp of this technological evolution, it's a privilege and a responsibility to consider the profound implications and opportunities AI offers UK government.

That's why Cognizant was delighted to sponsor Global Government Forum's Digital Leaders Study 2024, which explores the UK government's preparedness for adopting AI.

As a digital and data delivery partner to the UK government since 2017, we welcome the report's

emphasis on AI's potential to revolutionise how government operates. Delivering an AI-powered government requires a clear understanding of what this technology can achieve. And it also needs more than just technological expertise; it requires public trust, won through transparency and wise regulation.



This report underscores the critical imperative for a unified AI vision and for coordinated action across government departments. The consolidation of AI bodies within the Department for Science, Innovation and Technology (DSIT) is a step in the right direction, unifying AI leadership and setting the stage for a cohesive strategy that transcends silos in pursuit of deep collaboration.

The report delves into other core issues such as the critical role of data quality, better understanding of AI investment pathways, and building more nuanced regulatory frameworks. It envisions a future where AI is seamlessly integrated into government, driving efficiency, transparency and public service excellence. And it highlights the need for frugal innovation and learning from other leaders around the world, most of all Singapore's proactive and practical approach to AI integration.

As the new government starts work on delivering its five missions for the UK, the role of technology

cannot be overlooked. This report offers a roadmap for that AI journey, with insights and recommendations on how to harness it for the betterment of people, communities and public infrastructure.

Embedding AI across government departments and functions is not just a strategic choice – it's a necessity to ensure the UK remains competitive and responsive in an increasingly digital world. Its potential to enhance efficiency, reduce costs, and improve outcomes is immense. And the UK can lead the way.

So, while the journey ahead may be challenging, it's a journey the UK must undertake and do so with foresight, responsibility, and determination. I trust this report will play a key part in informing the UK's approach to the AI race in the coming years.

#### Yatin Mahandru

Vice-president and head of public sector & health, Cognizant, September 2024



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

**This report underscores the critical imperative for a unified AI vision and for coordinated action across government departments**



# INTRODUCTION

As we worked on Digital Leaders 2024, UK politics shifted significantly. We now have a new government (elected with a landslide majority), facing a difficult inheritance and keen to present itself as a group of pragmatic fixers. It's doing this with strict fiscal constraints which mean that reform, not higher spending, must be the path to putting public services back on track.

At the same time, the past 24 months have been defined by a remarkable new wave of AI. Governments, private companies and individual citizens around the world are embracing generative tools that look set to offer huge productivity gains and appear highly disruptive. Within government, it's set to be far more transformative than digital was in changing how the state operates.

This Digital Leaders 2024 report sits squarely between these two developments. Using the 7 Lenses of transformation as our guide (as in our previous Digital Leaders reports), we interviewed 10 UK leaders in government to understand how

ready the state is to transition to AI. Their views provide invaluable insights.

Unfortunately, we heard that the UK is far from ready to capitalise on the opportunities offered by AI. This was encapsulated by a comment made in one interview that when it comes to implementing AI within government: "It's like we're driving a car with one foot very firmly pressed down on the accelerator and one foot pressed down on the brake."

We were told that Whitehall *is* doing exemplary, even world-leading work around standards and AI safety. The first AI Safety Summit hosted at Bletchley Park and the AI Safety Institute are both good examples of this. Policy efforts like these are what this interviewee is referring to as 'the brake'.

On the other hand, there's considerable ministerial interest and external hype about AI, including from supportive think-tanks and private sector experts. This is creating huge expectations about what AI can offer and some early innovative work on this within government too – this is 'the accelerator'.



Kevin Cunningham





Both are important. We should be managing risks *and* exploiting the opportunities of AI at the same time. But, as interviewees told us again and again, the practical work that must underpin ‘the accelerator’ – the testing, trialling, implementing and scaling-up of AI applications – isn’t anywhere near advanced enough. There’s lots of talk, but not enough action to make sure we really are hitting that accelerator.

This was clear across all 7 Lenses of transformation. For example, there’s little evidence of a well-understood cross-government vision for AI yet, let alone a fleshed-out design or plan. Some departments are rushing ahead in its absence (while others move more slowly), but there’s little coordination or collaboration of this activity. We have a long way to go on these and the other lenses. In the final section of this report, we provide detailed recommendations for each lens, drawing on lessons from past successful digital transformation programmes in government.

Overall, this may sound like a pessimistic outlook. But the new government has already taken steps to improve this situation.

By consolidating the Central Digital and Data Office (CDDO), Incubator for Artificial Intelligence (i.AI), and the Government Digital Service (GDS) within DSIT, they have created a clear, unified voice for AI in government. It addresses one of the overriding concerns raised by interviewees: the lack of a single source for AI leadership. The new hub should be an invaluable first step to addressing this failing.

And there are other reasons to be hopeful. While finances are undoubtedly squeezed, this creates space for frugal innovation and focuses minds on the need for reform. The advancements around AI are already impressive and show few signs of slowing. And so, with this new configuration of AI bodies in Whitehall, there’s a chance to reset the government’s approach – with a new AI vision underpinned by practical actions.

In other words, we can improve – and improve fast. It just takes a government that’s willing to be bold and seize the opportunities that AI presents. Our message to them is clear: they should do exactly that. It’s time to commit to AI and really press down on the accelerator.

**Kevin Cunningham**  
September 2024

**There’s little evidence  
of a well-understood  
cross-government  
vision for AI yet**



# A NOTE ON METHODOLOGY

This edition of Digital Leaders focuses on the topics of AI and data within the UK government. The primary research source was 10 interviews with digital leaders in government today, exploring their perceptions of the UK's performance on AI across the 7 Lenses of transformation (see Matrix on the following page). We also conducted some additional desk research.

Our 10 interviewees were drawn from different grades and bodies across government. They ranged from permanent secretaries to senior leaders in digital roles. They included people from ministerial departments, the centre of government (comprising No.10 Downing Street and the Cabinet Office), non-ministerial departments, and executive non-departmental public bodies. Quotes from all interviews are unattributed and we have removed identifiable material from them.

We thank all of the following officials for taking part in interviews for this project:

**Susan Acland-Hood** permanent secretary,  
Department for Education

**Victoria Bew** head of strategy, i.AI

**Yvonne Gallagher** digital director,  
National Audit Office

**David Knott** chief technology officer,  
UK Government

**Tom Read** (then) chief executive officer,  
Government Digital Service

**Daljit Rehal** chief digital and information  
officer, HMRC

**Tom Smith** chief data officer, Department  
for Levelling Up, Housing and Communities<sup>1</sup>

**Huw Stephens** chief information officer,  
HM Treasury

**Ming Tang** chief data and analytics officer,  
NHS England

**Rob Thompson** chief technology officer,  
Home Office

<sup>1</sup> The Department for Levelling Up, Housing and Communities has subsequently been renamed as the Ministry of Housing, Communities and Local Government.



# 7 Lenses Maturity Matrix

	Vision	Design	Plan	Transformation leadership	Collaboration	Accountability	People
	The vision gives clarity around the outcomes of the transformation and sets out the key themes of how the organisation will operate	The design sets out how the different organisations and their component parts will be configured and integrated to deliver the vision	The plan needs to retain sufficient flexibility to be adapted as the transformation progresses while providing confidence of delivery	Delivering a transformation often means motivating into action a large network of people who are not under the direct management of the transformation leader	Collaboration is key to transformation in a multi-dimensional environment that increasingly cuts across organisational boundaries	Having clear accountability for transformation within an organisation enables productivity and improved decision making, and leads to better outcomes	Transformation will require people in your organisation to be engaged and to change their ways of working - you need to communicate effectively with them at every stage of the transformation
5	The vision is embedded in everything people do. It flows from top to bottom and is aligned with public outcomes	The public are at the heart of design work. Outcomes for different changes across the organisation are aligned. It's clear how to bridge the gap between the current and future states	Planning is joined up and fully resourced. Plans adapt as transformation progresses	Leaders embody transformation and create an environment of trust where it's safe to speak freely	The organisation compromises for the greater good and leads the way in transformation communities	Clear governance results in decisions being made at the right level and at the right time to drive progress	Ways of working needed for the future are adopted. Mature workforce planning exists
4	The vision sets a clear direction that people buy into. It is articulated in different ways	It's clear how different parts of the organisation will fit together. It's possible to assess progress as the design evolves	Planning is informed, coherent and mature, supporting both transformation and business as usual	Leaders tell a consistent story. They 'push' and 'pull' as needed to create the right environment for change	Roles, responsibilities and incentives reflect the need to collaborate, leading to new ways of working	People are becoming empowered and accountable for making decisions	Plans to deliver new skills or ways of working are being realised and people are engaged
3	There is a vision that is stretching but achievable. People see how they can fit into it	The design considers users and contains enough examples to bring it to life	Plans have the right level of detail and balance of tight and loose planning	There is sufficient ownership of transformation. Leaders talk about it. There are visible role models	Many decisions are made across boundaries. Shared outcomes are starting to be developed	There is broadly the right structure around transformation, with a focus on making decisions at the right time	Plans are in place to address the impact on people, ways of working and culture
2	A vision exists, but it means different things to different people	The design attempts to define the future in too much detail or doesn't cover everything it should	Plans are beginning to be joined up. Ambition and achievability need more focus	There is support for transformation at the top, and some change agents. There are meetings and ways to submit ideas	There is some understanding of stakeholders. Collaborative behaviour isn't yet commonplace	There is a growing level of accountability for transformation	The impact of transformation on people, ways of working and culture is understood
1	There is no clear vision for the future, or there are competing visions	There is no single design, or various designs are not joined up	Planning is not joined up. Plans are not flexible or achievable	Leaders talk about transformation on occasion. They make some effort to canvass views but avoid difficult messages	Collaboration across boundaries is limited	Responsibilities and accountabilities for transformation are unclear	The impact of transformation on people, ways of working and culture is not understood



Transformation Peer Group  
[www.gov.uk/transformation](http://www.gov.uk/transformation)



HM Government

Source: [7 Lenses Maturity Matrix poster](#)





In addition, Saheel Sankriwala (chief technology officer at the Department for Education) was present for our interview with Susan Acland-Hood and contributed to the discussions.

To augment our research, we used Redbox – a generative AI tool being developed by i.AI – to analyse interview transcripts and provide insights. Crucially, we used this to test the analysis that we had already conducted (without AI tools), rather than as any substitute for doing so. All transcripts were fully pseudonymised before being analysed by Redbox and all data was handled securely.

Some of the Redbox output is included in the appendix at the end of this report.

Further insights were gleaned from the AI Study Tour that Global Government Forum (GGF) hosted in June 2024. During this tour, senior officials from the government of Barbados came to the UK and visited key government units working on AI (such as GDS and i.AI) as well as private companies such as Cognizant and Hippo Digital. This offered a chance to share best practice and build connections between leaders beyond Whitehall, further enhancing our understanding of the global AI race.





# A NOTE ON AUTHORSHIP

To produce this research, GGF – the publishing house for civil servants around the world – teamed up with Kevin Cunningham.

Kevin is a senior UK digital leader who spent his early career in programming and IT consultancy, later becoming the global head of online for Vodafone Group and director general of the Business Transformation Group at the Department for Work and Pensions (DWP). He was director general of the UK's Government Digital Service (GDS) from August 2016 to July 2019, and digital envoy for the UK and director general of the International Government Service until 2021.

Kevin worked with James Sweetland, a researcher, journalist and consultant who works on government reform, policing and tech policy, and Yatin Mahandru, vice president and head of public sector and health,

UK and Ireland for Cognizant, to conduct the interviews and write this report.

Cognizant, one of the largest technology services companies in the world, was GGF's knowledge partner for this project. Cognizant is a digital and data delivery partner to UK Government in central, health and defence, delivering citizen-facing services and outcomes since 2017. In addition, the company helps clients across wide-ranging industries in the UK and Ireland modernise technology, reimagine processes and transform experiences so they can make the most of technology to stay ahead in our fast-changing world.

Throughout this report, where you see the word 'I' in the text, that's Kevin speaking. 'We' refers to Kevin, James and Yatin.



# CHAPTER 1: THE UK'S PLACE IN THE GLOBAL AI RACE

The year 2023 was frequently heralded as the year of AI. Given the remarkable advancements we saw over those 12 months, especially around generative AI, it's not an unfair label. **With 100 million users within just two months**, innovations like ChatGPT were adopted at astonishing speed by people around the world.

Of course, AI is nothing new. In 1966, one of the very first AI chatbots – ELIZA – was developed by Joseph Weizenbaum. Prefiguring ChatGPT by nearly 60 years, it was understandably limited. Through the 1970s and 1980s, the field of AI saw breakthroughs, counterbalanced by periods of prolonged stagnation.

During the 1990s, progress was more sustained. The success of IBM's Deep Blue supercomputer in 1997, defeating reigning world chess champion Garry Kasparov across six matches, was one notable moment that caught the public's attention. Later victories, from IBM Watson in the game show

'Jeopardy!' to Alpha Go's 2017 triumph over the world number one in 'Go', only added to the hype.

But the current wave of AI seems genuinely different. When we published **our first Digital Leaders Study in 2021**, there were no mass-adopted generative AI products. By the time we released **our 2023 follow-up report**, GPT-3.5 had been accessible for several months. As we publish this paper, we've seen two new models from OpenAI (GPT-4 and GPT-4o), demonstrations



**Ming Tang**, chief data and analytics officer, NHS England

of remarkable text-to-video capabilities (through software like Sora), the full release of Google's Gemini, and ever-growing use among the public.

For governments and private companies the world over, it's become ever clearer that AI is the future. The UK – both the state and its commercial enterprises – are now in a truly global race to make the most of AI.

But for all the (often justified) hype around generative AI, the defining feature of this new wave of tech is that it's been highly disaggregated: individual citizens have led the charge in adopting these tools. It's not uncommon for someone's personal use of ChatGPT (or similar tools) to far outstrip what's being used in their workplace. And when it comes to larger institutions, implementation has been relatively limited.

For government, this is partly understandable. Unlike with the digital transformation we saw in Whitehall – when the state was playing catch-up



with the private sector – the latest AI innovations aren't yet widely adopted across many large businesses either. There's much more limited best practice to copy from and risks are therefore greater. With fewer common standards or business change models to work

from, there is some rational caution here, especially given the government's implicit duty of care. All AI systems that are citizen-facing must be trustworthy and reliable, but must be even more so when people are interacting with the state and public services.

And yet, despite these valid concerns, government has to scale up AI and do so quickly – the opportunities are simply too vast to pursue a path of delay. Moving fast is essential. But this is set to be a much more difficult transition than with digital.

## The UK and Singapore

Under the previous government, the UK began to develop a stronger AI leadership role. The creation of the Department of Science, Innovation and Technology (DSIT) in February 2023 was **explicitly framed** as “positioning the UK at the forefront of global scientific and technological advancement”. As the new department tasked with ‘policy for AI’, it subsequently **promoted** a new white paper, offering a “pro-innovation” model of AI regulation that contrasted with the EU's more restrictive approach.

The strongest leadership has been around AI safety. In November 2023, the UK hosted the first AI Safety Summit at Bletchley Park, attracting global leaders

– from the president of the European Commission and the vice-president of the United States to OpenAI CEO Sam Altman – and securing a multilateral agreement to test major AI models. The launch of the AI Safety Institute in February 2024 established the UK's capability to carry out that work.

Maximising the opportunities of AI is now a point of policy consensus in the UK. Just as the previous prime minister Rishi Sunak **said** that AI would “bring a transformation as far-reaching as the Industrial Revolution”, the new government's manifesto has **committed** to harnessing AI “to transform the speed and accuracy of

diagnostic services” in the National Health Service. And global AI rankings show that the UK is building from a solid base. Oxford Insights places the UK third in its international **AI Readiness Index**; Tortoise Media has us at fourth in its **Global AI Index**.

And yet, there remains a central tension in our approach – this is “the accelerator and the brake” metaphor explained earlier on. Though there's much excitement among ministers and officials about AI, the bulk of the practical activity so far has been on that brake: to regulate, govern, set standards and create safeguards. This is all necessary and valuable, but without



### Case study: use of AI in Singaporean public services

Singapore has been a world leader in implementing AI within its government and public services – for example:

**OneService Chatbot:** introduced in 2021, this chatbot enables citizens to report local problems – e.g. public cleanliness or illegal parking – via WhatsApp, Telegram or Instagram. The chatbot categorises the problem, encourages the user to share key information, and places it in a template. Users then check the summary produced by OneService and, once approved, the chatbot will send the information to the relevant agency.

**Pair:** a set of productivity tools for Singaporean public servants, using AI to help them complete a wide range of tasks. Examples include Pair Chat – a secure version of ChatGPT that’s now available to all public officers in Singapore, and Pair Intern – an email assistant tool that analyses documents. The latter was being **used by 1,100 people**, as of March 2024.

practical adoption of AI in government – the accelerator – we risk falling behind, when others are racing ahead.

There’s a useful contrast here with Singapore. The latter is ranked third globally in the Tortoise AI Index and second by Oxford Insights – in each case, one place above the UK. While our government published a single **National AI Strategy** back in 2021 (a document which makes no reference to generative AI), Singapore has already published **two** – with the most recent released last year.

There’s more evidence of a disparity here. In 2020, a collection of Singaporean agencies published

#### **a guide to redesigning jobs around**

**AI** – more than a vision alone, it’s closer to the design and plan components of the 7 Lenses. Crucially, Singapore has also gone further than the UK in implementing AI across public services. This is where the UK must get to: AI implementation, not just ideation.

Whatever the critics may say, the new UK government has taken power with an ambitious set of goals. Labour has set five overarching national missions – from securing the highest sustained growth in the G7 and delivering zero-carbon electricity by 2030, to reducing the number of lives lost to the most deadly diseases. Achieving this while

fixing a catalogue of broken public services won’t be easy. But with a high tax burden and limited scope for more spending, operating differently is key.

AI must form a core part of this work. A frugal innovation approach – in which services are transformed in as cost-conscious a way as possible – will be essential.

The 10 interviews we conducted with digital leaders offer useful guidance for new prime minister Sir Keir Starmer and his team. They provide invaluable insights into how the UK can move closer to Singapore’s approach: from talking about AI, to actually putting it in place.



# CHAPTER 2: AI IN GOVERNMENT: PERSPECTIVES FROM THE UK'S DIGITAL LEADERS

## Setting direction around AI – vision, design and plan

### Summary

Until recently, there was no single voice or source of leadership for AI within government. The Central Digital and Data Office, Government Digital Service, i.AI, and Department for Science, Innovation and Technology each held a different mandate around AI and were spread across different parts of Whitehall.

Due to this fragmented leadership structure and the absence of political sign-off for publishing the CDDO's AI vision, digital leaders told us that there was limited strategic alignment around AI. There was simply not a coherent, well-understood, and strategic vision for AI. This meant government also performed poorly on the design and plan lenses – there remains too much vagueness and uncertainty thus far.

The other major concern we identified is that Whitehall has focused too much on regulation and risk management around AI, with relatively little emphasis on adopting and embedding new tools to improve public services. Important as AI safety is, a balanced approach is needed. The fact that there are only 74 deployed AI use cases (among the 87 government bodies who responded to a National Audit Office survey this year) shows how much further there is to go.





## Vision

When attempting to transform and modernise government, the first step is developing a clear and realistic vision that's well understood. Around AI, we found a mixed picture – with interviewees' self-reported scores for this at a two or three.

Those across government broadly recognise the potential AI offers. There's growing interest among senior civil servants and politicians – regardless of their political leanings. Indeed, as one digital leader memorably told us: "The level of ambition and energy from ministers is like nothing I've ever seen before in tech... They've been told, these are literal magic beans, and you've just got to sow the magic beans, and all your problems go away."

In reality, all this excitement isn't matched by a well-defined vision yet. There's limited articulation, both inside and outside government, of what AI can and should mean as a tool for transforming and modernising public services.

This appeared to stem from two problems. The most fundamental was a perception that the model of AI leadership in central government that existed until July 2024 in Whitehall was fragmented and confusing.

There were four different agencies operating to different mandates – DSIT leading on 'policy for AI', GDS charged with scaling up digital tools across government, the Cabinet Office's Central Digital and Data Office (CDDO) as the lead on 'AI for policy', and i.AI as the incubator for new innovations. This was a highly complex picture.

Some defended this as a sensible division of labour and, for those working on AI, it probably was clear and comprehensible. But many digital leaders accepted that it was too complex for those outside. With fragmented leadership, it was even harder for a coherent vision to break through.

This is exacerbated by the fact that departments are pushing ahead with



**Daljit Rehal**, chief digital and information officer, HMRC

exploring AI for their own purposes. This was partly a reflection of that lack of central leadership. As one digital leader explained: "I don't want to sit around waiting for CDDO to tell me what to do, or not do anything until they come up with the perfect vision, because I sort of think that might never quite happen."

As departments push on, they are developing their own visions for AI. This might entail anything from a fully fledged vision to a vaguer set of ideas about AI. And these visions are informed by their own department's digital leaders and their preconceptions. In fact, some parts of government are not planning to develop visions at all. The National Audit Office **found** that 21% of government bodies have an AI strategy and 61% plan to develop one, but 15% have no intention of doing so at all.

One digital leader suggested that there are "no massive frictions" between these different departmental visions. This is positive. But they added that the lack of coherence still creates complexity: "[It's] deeply confusing having so many different voices not quite saying the same thing."

In other words, there is no coherent, shared vision for AI in government



as things stand – a **finding** that the National Audit Office also made earlier this year. Its conclusion is hard to ignore: “There are risks to value for money if the government does not establish which department has overall ownership and accountability for delivery of the strategy for AI adoption in the public sector...”

A single voice in government was therefore needed to establish a well-understood AI vision. And so, the

government’s announcement – just a few days after coming into office – that DSIT will now host all of CDDO, GDS and i.AI is promising. It establishes the kind of single AI authority that’s needed to build a whole-of-government vision.

The second problem we identified around vision is much simpler. It turns out that the CDDO *has* developed a vision for AI in government: one that some within the system, especially in digital roles, already know about.

However, the CDDO has not received political sign-off to publish it (at the time of writing).

New-in-post Labour ministers will understandably wish to provide feedback on this document and possibly make changes. But DSIT should ensure this revised vision is published urgently. Delivering and promoting a clear vision should be the flagship statement of the new AI hub, resetting the government’s approach in this area.

## Design and plan

With limited work around the fundamentals of vision-setting, it’s unsurprising that progress on design and planning has been mixed too.

Digital leaders working in the centre do have some clear ideas of exactly how and where AI should be being deployed. One suggested AI should be used in four areas: improving the user experience for government services; reducing cost; supporting workers with co-pilot capabilities; and improving use of data in decisions. Another articulated a similar matrix of sensible applications for AI organised around two variables:

proximity to citizens and complexity (see page opposite).

But across government, fully formulated plans for implementing AI are in short supply. One interviewee told us: “I don’t think anybody’s yet got what we would class as a proper AI adoption plan.” Another suggested that different types of AI-related activity aren’t being clearly delineated or understood. AI tools for improving operational delivery are conflated with externally facing AI developments, such as overall national AI policy and regulatory approaches.

Without a central vision, no system-

wide adoption plan is on the cards. Even though there are some AI tools already in use or late-stage development – we trialled Redbox (see appendix) and were impressed – implementation is generally limited. The NAO **found** that 37% of government agencies deploy AI already, but this was usually only with one or two use cases. In total, they found just 74 deployed AI use cases across 87 government bodies who responded to their survey – **compared to over 7,000 digital services listed on the GOV.UK website.** And interviewees suggested that most AI tools used in government



## AI applications: potential use cases in government

	High proximity to citizens	Low proximity to citizens
High complexity	Real-time assistance, contextual advice and guidance for public sector workers.	Synthesis and sensemaking support for policy professionals.
Low complexity	Using AI to reduce burdensome tasks – e.g. handling correspondence or calls from the public.	Back-office functions that are amenable to AI and traditional automation techniques.

today have been developed by external bodies, rather than within the system.

It's true that government has gone much further around AI safety and thinking about risk. But this mismatch is a problem. One digital leader told us: "All our policy work is going into stopping AI and all our delivery work is going into enabling AI." This is the analogy we described earlier – "It's like we're driving a car with one foot very firmly pressed down on the accelerator and one foot pressed down on the brake."

It's not that AI safety work is itself a problem. The issue is that, at the moment, the foot on the brake is pressed down much harder and much more consistently than the foot on the accelerator ever is.

This may be because risk management is of higher concern in the public sector and so comes naturally to Whitehall. For government services, it's essential that AI tools are very accurate. The public expect their interactions

with the state to be handled to a high standard and that information provided by their government is trustworthy. This can create understandable blockers to rapid rollout of AI.

For example, the deployment of a generative AI chatbot on the GOV.UK website has been hindered by "answers [which] did not reach the highest level of accuracy demanded" for government – according to [a blog published by GDS data scientists and user researchers](#).



**Susan Acland-Hood**, permanent secretary,  
Department for Education

The same blog post noted that: "Some users underestimated or dismissed the inaccuracy risks with GOV.UK Chat, because of the credibility and duty of care associated with the GOV.UK brand." This illustrates the significant challenges involved in deploying some types of AI.

However, more generally, there can be too much emphasis on risk management and AI safety. One interviewee argued that "civil servants can be rather obsessed with risk". They added: "I do wonder whether the fact the AI safety agenda has been so well-promoted really just speaks to the fact that it's something civil servants were ready to grasp onto: the risks of AI we should be protecting against."

For now, policy for AI is undoubtedly ahead of AI for policy (or operations, for that matter). With little evidence of a coherent vision – let alone a design or plan – the vehicle has remained stationary.



## Departmental drivers – collaboration and accountability

### Summary

Successful cross-government working has long been a challenge within Whitehall, reflecting the institutional design and incentives that exist in the system. With powerful departmental ‘fiefdoms’, an underpowered centre, and accountability resting ultimately with departmental ministers and permanent secretaries, it’s rarely easy to build a whole-of-government approach to cross-cutting issues.

With AI, we found a similar story. There are some networks and spaces in which both formal and informal collaboration take place and these are productive structures to build from. But they ultimately seem to be about discussion and engagement primarily, rather than genuine cross-departmental working.

The creation of the new AI hub within DSIT – putting the CDDO, GDS and i.AI all in the same place – offers a single source of leadership on this agenda. There may be some scope for it to drive a more collective approach to AI within Whitehall.

Interviewees were clear that one area where this is especially important is data sharing. By connecting data up across government, AI tools can reach their full potential and draw on a much richer range of material about citizens. The new AI hub should seek to drive collaboration in this area as a key priority.



## Collaboration

Building a cross-government approach to AI requires effective collaboration, especially between departments and the centre. In Whitehall, there's a tendency to believe that the UK is world-leading in this regard, but the reality doesn't always bear this out. To borrow a term used by one Caribbean digital leader in [a GGF project examining the transformation challenges governments face there](#), we often engage in "collaboration-by-mouth". The rhetoric around working together is greater than the evidence of genuine collaboration.

The barriers to collaboration in the UK government are well-known; they are inherent in our system of government. One digital leader described Whitehall as defined by "loosely coupled fiefdoms, all motivated by public service and innovation". They added: "The incentive structure of government is still such that collaboration is your lowest priority. That's just fundamentally true." Overcoming this isn't straightforward.

We did find some spaces where collaboration – or at least information-

sharing – is taking place. Interviewees told us about a cross-government forum for AI policy and a general AI forum where the CDDO shares the work it is doing with departments. There's a digital, data and technology (DDaT) board attended by permanent secretaries and an AI practitioner community convened by CDDO (with representatives from outside government).

We also heard about chief technology officer (CTO) meetings of various sizes.



**David Knott**, chief technology officer,  
UK Government

The largest is a 'council' attended by all CTOs across Whitehall. There's then a smaller version for CTOs from large operational departments and the CDDO, followed by a final meeting of the same members – but without the CDDO included.

There are some informal networks too. One digital leader described themselves as being part of an "accidentally arising community", having connected with other public servants around a shared interest in AI. Having both these formal and informal spaces to collaborate and discuss AI can only be a good thing – especially given how early government is on its AI journey.

But there is a risk again of this 'collaboration-by-mouth', with much talking but little action. One interviewee made exactly this point – the high-level alignment is easier than the harder collaboration on practical issues: "I think we're good at aligning on vision, but when you look at design and other [lenses], there are institutional barriers and hurdles to working effectively together."



## Accountability

This is where we encounter questions of accountability. Whitehall is defined by those departmental fiefdoms, with the centre often struggling to be influential. Civil servants are ultimately answerable to their departmental ministers. Permanent secretaries are the accounting officers – meaning they are accountable to Parliament for the work of their department.

This creates obvious incentives with predictable consequences. One digital leader asked the provocative question: “How do you break a system that, for hundreds of years, has been designed around departments doing things that they’re individual accountable leads for and instead [get them to] deliver genuinely cross-organisational stuff?”

The truth is that it’s difficult, but not impossible. In government, I saw how effective the centre can be in driving transformation around digital. We showed departments the value of coming on the digital journey – using carrots and sticks – and built a

reasonable coalition of support to help deliver it. Systems changed accordingly, after much hard work.

There are similar examples of this type of approach in Whitehall today. The [government’s 2022-25 Roadmap for Digital and Data](#) includes six formalised cross-Whitehall missions. One of these, ‘sponsored’ by the permanent secretary of HMRC, is that: “All departments will confirm an adoption strategy and roadmap for One Login by April 2023 and their services will have begun onboarding by 2025.” It’s an example of straightforward milestones matched with clear accountabilities – with responsibilities to take action across the system.

The new AI authority in DSIT is still taking shape, but it should try to build on this mission’s framework as set out by the CDDO. Its placement in a department which has genuine heft and extremely close links into the broader AI community is a good

thing, but it’s not yet obvious how it will work to drive government-wide transformation. Learning from that existing CDDO roadmap is crucial as it begins to puzzle out its emerging role.

One area where the hub could offer tremendous value is around government data: specifically, making it far more interoperable. A recent [article](#) by Sir Robert Chote, chair of the UK Statistics Authority, raised exactly this topic. And data sharing is a crucial priority in the context of AI too, as one interviewee told us: “If government wants to be serious about using machine learning tools and frameworks to help it become more productive and efficient, it has to solve its data problems. Just limping along is no longer an option.”

Whitehall’s “unremediated legacy” around data, to use a term from one interview, is a serious challenge. Training and using AI models relies on the supply of high-quality data, both to create and keep improving these





powerful tools. And many of the most exciting applications of AI can only reach their full potential if they can pull data from different parts of government. A personalised AI assistant for citizens is of limited use if it's able to use Home Office data, but can't access anything from the Department for Work and Pensions. The perennial Whitehall challenge around data sharing is one of the greatest threats to government's AI transformation.

Perversely, this presents an opportunity: data's key role in supercharging AI might focus minds and incentivise more collaboration. But the real risk is that things don't change and poor data practices remain a significant barrier. To return to the accelerator and the brake metaphor, a failure to fix government's data is like operating with a speed limiter firmly attached to the vehicle. We can only go so fast without ameliorating this challenge.

This is why the new AI hub in DSIT must focus on this agenda. One interviewee called for legislation to help this, telling us: "The first thing you [should] do is put a big digital bill through Parliament that removes all the impediments to data sharing, all the impediments to building user experiences that are entirely based on departmental silos."



**Tom Read**, (then) chief executive officer, Government Digital Service

Various legislative efforts have already been explored in recent years. The Digital Economy Act 2017 took some steps to improve public services through better use of government data. And the new government has **announced** plans for a Digital Information and Smart Data Bill in the most recent King's Speech, which included a commitment to reform data sharing and standards – again for the benefit of public services. It's too early to tell whether this will finally realise that crucial ambition.

What is clear is that a Big Bang is needed; in fact, it's long overdue. Legislation may help, but the government's AI hub should also play a key role here. Departments must be held accountable for failing to share data and those common standards must be established to help make collaboration easier – it cannot be 'by-mouth'.



## Developing an AI-ready workforce – transformation leadership and people

### Summary

It's well-understood that workforce – already a challenge in the context of digital – is even more difficult to get right around AI. Getting the right skills in place within government, given the pay premiums that the private sector can offer, isn't easy at all.

Despite that, Whitehall has some excellent AI talent (especially within i.AI), according to interviewees. But demand far outstrips supply and is certain to keep doing so. There's little evidence that enough is being done to develop the AI skills needed within Whitehall, whether via technical training or upskilling around prompt engineering.

The absence of AI-specific roles within the Government Digital and Data Profession Capability Framework, though apparently being resolved, is another significant oversight. And more work needs to be done to understand when to rely on external skills and when to build tools internally.

When it comes to leadership, top civil servants outside of digital roles – such as permanent secretaries – must have a good understanding of AI. Interviewees noted the key role those leaders play in helping ministers understand what's really achievable around AI in their departments. They must be well-informed so that they can balance hype with practicality.

For now, top officials have a variable understanding of AI, with some upskilling efforts in place that appear to be productive. Boosting understanding of this transformative technology among those leaders (and senior officials just outside that top tier) is another area that Whitehall needs to focus on.



## Transformation leadership

There are many highly competent and skilled digital leaders in government. They combine robust technical knowledge with the ability to operate in a system that can be confusing to navigate, especially when compared to the private sector. However, when it comes to other top officials – senior civil servants outside of digital roles – there are concerns about how well they grasp this agenda.

One interviewee argued that the “number one issue is senior leaders’ lack of digital fluency”. These “generalist digital leaders” (as they put it) are “de-facto digital leaders because they make all the decisions around major change programmes and where money goes. But they would be horrified to be called ‘digital leaders’ because they think it’s something that’s done over on the periphery, on the edges, and it’s not important.”

In other words, a good understanding of what’s possible around AI, data and digital cannot solely be the preserve of those in strictly digital roles. It’s a

prerequisite for everyone in Whitehall – and one that should be viewed as central, not peripheral – precisely because AI and data touch everything that government does today.

And having this level of understanding also helps officials better serve ministers. There was a perception from interviewees that politicians are only hearing the most optimistic voices around AI. Senior officials must be able to give a sense of what’s actually achievable when advising ministers – as one interviewee explained: “Some permanent secretaries might be the people with the first shot at framing those conversations. They need to have a level of understanding of the tech and what’s realistically possible, so we don’t get signed up to a vision that we can’t deliver.”

Permanent secretaries are not a homogenous group in terms of their skills, backgrounds or capabilities. Around AI, it’s the same story. One interviewee described them as a “mixed knowledge group”, with

some well-informed and others less on top of this agenda.

We heard about several existing efforts to upskill top officials. There was an away day held for permanent secretaries at Oxford University, involving a series of AI workshops. Other examples included a top civil servant visiting Cambridge University to learn more about AI, knowledge exchange meetings between senior officials and leaders at companies like Google and Microsoft, and teach-ins where directors general and permanent secretaries heard from internal digital experts.

These are all worthy endeavours, but the overarching perception was that they aren’t moving the dial enough. One interviewee suggested that top leaders do “know AI’s a big bet” but “there’s no materiality” beyond that. Another suggested that Whitehall leaders are around five years behind private sector CEOs: “There’s a big digital gap and then [there’s] an even bigger AI understanding gap.”



These challenges aren't confined to the very top level, however. Understanding and enthusiasm around AI varies among leaders beyond permanent secretaries.

One senior leader admitted: "If you picked up the phone to a random member of middle leadership [in this department] – SCS1s/deputy directors – some of them would

be able to articulate a chunk of what I just said [about AI and digital], but quite a lot of them wouldn't." Another interviewee offered a similar account, describing a mix of progressive and unconvinced DGs: "There are people who are threatened by it, there are people who are cynical about it, and there are people who want it all – but want to do it for themselves

because they don't want to lose independence."

Getting leaders on the same page around AI is essential. This might mean some remain sceptical of its transformative potential. That's to be expected. But ensuring that the techno-optimists and the AI sceptics have a strong base of knowledge to work from is the key priority.

## People

The last of the lenses presents an even greater challenge than transformation leadership. 'People' refers to a wider set of the workforce and their skills, knowledge and capabilities – getting this element right is especially difficult with AI.

Given the global race for AI talent and the public-private competition for these skills, it's unsurprising that Whitehall struggles to keep up. Pay in the public sector is never going to match what's available elsewhere and the pull factors – around public service and interesting work – can only go so far.

Despite that, we heard that i.AI had successfully attracted many

talented people. The chance to join an environment already stocked with impressive talent, greater pay flexibility, and proximity to the PM/No.10 were all offered as drivers of this. Whether i.AI can continue to recruit as effectively now that it's within DSIT (i.e. further from No.10) is yet to be seen. Beyond i.AI, one interviewee told us there are "pockets of exceptional talent", including strong capabilities in some specific departments.

But overall, government does not have the digital skills it needs, including around AI. One minister's **revelation** in February 2024 that 20% of DDaT roles in his department were unfilled

highlights this problem. More widely across Whitehall, interviewees pointed to three factors contributing to the AI capability challenge.

First, the **Government Digital and Data Profession Capability Framework** – which defines DDaT roles and the skills that underpin them – has not been updated to reflect AI jobs. Departments use the existing framework to recruit some AI talent, often by adapting analogous DDaT job specs, but there are no AI-specific roles included in that document, nor any competencies mapped out for them.

We were told that new guidance to incorporate AI jobs is being developed,



but it's not clear when that will be in place. And as one interviewee pointed out, defining exactly what 'AI skills' are isn't easy, with the relevant career and training pathways remaining somewhat obscure. Hiring the wrong kind of expensive AI talent creates challenges of its own, so getting a revised framework right is essential.

Secondly, several interviewees referred to effective 'hiring freezes' within departments which prevent them buying in new talent – including in DDaT roles. This follows the previous government's [plan](#) to 'cap' civil service headcount, which has led to a number of external recruitment drives being frozen.

One senior leader from a Whitehall department told us: "It's ridiculous to have that headcount number prevent us from being able to do more things in-house at better value, just because it theoretically increases the total number of civil servants." The new government scrapped this headcount cap in July

2024, which should enable greater flexibility to recruit AI (and digital) talent when it's needed.

Finally, training to help upskill civil servants around using or developing AI isn't good enough. One interviewee called this available training offer essentially "unusable".

There are different AI skills that different cohorts of civil servants will benefit from. Some need technical training so that they can learn to develop or refine AI tools themselves.



(Left to right) The Hon. Clare Martorana, Federal Chief Information Officer, USA, Victoria Bew, and Susan Acland-Hood

Others, like policy professionals, need competence in things like prompt engineering, so they can be intelligent users of AI applications created by others.

When I was in government, GDS trained over 10,000 staff as part of our digital transformation goals. There's little evidence of the same kind of ambition around AI. We agree with the perspective of one interviewee: "There's got to be a huge upskilling programme."

Having a good capability base will allow Whitehall to use or access AI much more effectively. In some cases, the intelligent client model will be best – relying on trusted private partners who can offer expertise that cannot be maintained internally. In other cases, internal development of AI tools can save money and prove more effective. But Whitehall can only follow either of these paths if it boosts its people's capabilities around AI and develops a much-improved training offer.



# CHAPTER 3: JUMP-STARTING THE AI REVOLUTION

AI undoubtedly offers impressive potential for government and public services. This view is shared by politicians, civil servants and private sector companies alike, as well as Labour-aligned think tanks likely to be influential on the new government.

But there is a need to temper the hype somewhat. Figures suggesting

AI could save government as much as £200bn over five years

should be treated with caution.

Our interviewees were sceptical of such huge numbers and noted that public sector transformation around AI will be complex and time-consuming. And, as this report has shown, we're starting from a relatively limited base within government.

In fact, a common theme throughout our research was the need to do the hard yards first. Away from the hype about the most cutting-edge AI applications, there's much work needed to lay the groundwork for this transformative tech. As one interviewee put it: "There's a lot of hype about using AI without doing the fundamentals. [But] we need to do the fundamentals first."

## Three fundamentals

In that spirit, our research suggests three core issues that require particular attention:

**Data** – government must get serious about improving data interoperability and quality across Whitehall. There are

excellent reasons to do this even aside from AI (as Sir Robert Chote set out in the **article** referenced above), but it's even more essential now. It is a non-negotiable part of the AI transition.

AI tools are only as good as the data used to develop and refine them. Many proposed use cases – e.g. an AI assistant to help citizens use government services – can only reach their true potential if they can





use data held across different parts of the state. To ensure reliability and to keep improving, these models require a constant, high-quality data flow.

**Funding** – as well as realism about what the headline potential benefits are (whether cashable or not), government needs more clarity about what ‘investing in AI’ really means.

The upfront cost of buying or developing an AI tool is only the first step. Monitoring the outputs of the first iteration of that algorithm and then refining it will entail a repeating cycle of additional investment. Having that more rounded understanding of cost will make it easier to estimate the true returns AI can offer.

Government should also do more to establish common AI model training architectures to improve effectiveness. At present, many services (inside and outside government worldwide) take general large language models and train them in specialised domains or functions. This increases the risk of hallucination and inaccuracy problems (more on this below).

The central government units we engaged with have developed thoughtful architectures for training

AI models on specific tasks, but some departments may not have done so – even as they push forward with their own use cases.

The new AI hub should be given funding to publish details of the architectures used by AI units. This would allow departments to build on or adopt them wholesale themselves. In the long run, funding for this priority could improve effectiveness and reduce hallucination – a wise, much-needed investment that could avert future controversies.

**Regulation** – balancing the accelerator and the brake is one of the key themes of this report. Labour’s AI plans appear to include some further regulation, though the full details of this are yet to emerge. This follows on from the previous government’s **decision**, in February 2024, to expand the algorithmic transparency recording standard to cover all central government departments.

There may be merit to each individual decision to enhance regulation and we fully accept that AI safety is extremely important. But Whitehall must not fall into working in one of its comfort zones – risk management – without

doing enough to hit the accelerator. The new AI hub within DSIT must make sure it’s also focusing on implementing and scaling up new AI tools.

These two priorities aren’t necessarily opposed. In some areas, proper risk management work can lay the foundations for improving adoption of AI.

For example, we heard that one central government unit improved the accuracy of one AI tool from an initial confidence level of 78% to 85%. DSIT’s AI hub could develop a framework that explains what confidence levels should be achieved before AI tools of differing risk profiles should be launched. And they could also create the mechanisms needed to measure accuracy (and in turn confidence) in these systems.

That would help those in government understand what’s needed to take an AI product from a trial to a deployed use case. It might also help Whitehall departments independently determine whether the more over-hyped private sector AI products – some claiming accuracy of nearly 100% – really do make the grade. This kind of approach to regulation can balance risks and ease the way to implementation.



## A roadmap for AI in government

As well as fixing these three fundamentals, government needs a clear, overarching AI roadmap. As things stand, we didn't find – on any of the 7 Lenses – that government is yet in the right place to fully enable its AI transition.

The decision to consolidate GDS, CDDO and i.AI within DSIT is an excellent first step. As we've already argued, the lack of clear leadership around the AI agenda so far has hampered progress.

With a new single voice established, this is the perfect moment to set

out a clear, coherent and influential AI roadmap – we recommend the following as a starting point for this. Our recommendations align closely with previous successful transformation programmes that accelerated digital capabilities in government.

### A roadmap for the UK Government's AI transition

**Vision:** finalise and publish a vision for AI use across the whole government by the end of 2024, building on the existing work done by the CDDO to develop it. This should include plans to develop world-class AI use-cases for government, an improved training and education offer, better use of data, and identification of the key pieces of architecture (technical and regulatory) needed to underpin scale-up of AI. Consider GDS' [\*2017 Government Transformation Strategy\*](#) as a useful model for this type of vision.

**Design:** introduce a £100m funding pot administered by DSIT to identify and scale up promising AI innovations for use in government. Run a competition, using a similar methodology to the [\*previous £20m GovTech Fund\*](#), to identify these exemplars. Repeat this process every two years to capture and develop new innovations.

**Plan:** following the publication of the AI vision in 2024, introduce a fully fledged AI plan for government – building on the departmental AI plans currently being developed. This whole-of-government plan should be launched in early 2025 and run for 12-24 months.

**Collaboration:** build on the limited formal and informal networks for collaboration around AI, by tasking DSIT's AI hub with developing an engaged community across Whitehall on this topic. Learn from what's already taken place by developing both formal and informal networks to share ideas and encourage cross-system working.



**Accountability:** capitalise on the creation of the single AI hub that is now fully accountable to DSIT ministers and top officials. Create a new set of cross-government digital missions, following the expiration (in 2025) of those included in the *CDDO's existing roadmap for digital and data*, to drive improvements in areas such as data sharing and interoperability.

**Transformation leadership:** supplement existing efforts to upskill leaders around AI by investing heavily in bespoke training for permanent secretaries, project senior responsible owners, and other senior leaders. Draw on private sector, academic and other expertise to develop a well-rounded programme that is tailored to the needs of those at each tier of leadership.

**People:** urgently update the Government Digital and Data Profession Capability Framework to include AI roles with defined competencies. Invest heavily in training aligned to this new framework (from highly technical to softer skills like prompt engineering), with ambitions akin to GDS' upskilling of 10,000 civil servants around digital.

This roadmap doesn't offer all the answers. The transition to AI is a programme of change that will take years, with the potential to far exceed

the impact of the digital transformation of government. And it's also true – as even the digital leaders we spoke to

acknowledge – that the full suite of opportunities AI offers is far from settled.



## Time to hit the accelerator

The new government has inherited an extremely difficult legacy, with profound challenges and significant constraints on what it can do to resolve them. But this is also a moment where the system seems to recognise that reform, not new spending, is the only plausible path to change. Delivering on Labour's ambitious missions means embracing frugal innovation and the potential AI offers. Fixing public services now inescapably means modernising and transforming government.

Placing CDDO, GDS and i.AI within DSIT under Rt Hon Peter Kyle MP's

leadership was a wise decision. Our research shows that it will be strongly welcomed by digital leaders inside the system who felt that direction around AI has been lacking for some time.

With the machinery of government changes complete, there is now a chance to fully reset the government's approach to and plan for AI. Measures like the new **AI Opportunities Action Plan**, led by the chair of the Advanced Research And Invention Agency (ARIA) Matt Clifford, suggest there is a clear commitment to drive wider adoption across the UK economy.

Our observation is that DSIT does seem to be grasping this essential agenda.

The government must keep up this momentum. It must seize this moment and embrace the roadmap we've presented here. For too long, we've been stamping down on the brake while – for all the excited rhetoric – failing to take the practical steps towards implementing AI.

A new government, with a new PM and a landslide majority, has the chance to be bold and commit to making the most of AI. It's time to press the accelerator.



# APPENDIX

As part of the research process, GGF worked with i.AI to test out its generative AI tool [Redbox](#). This offered us the chance to see an AI application currently being developed in Whitehall and explore its potential and weaknesses.

According to i.AI, [Redbox's features](#) include:

- **Robust Retrieval Augment Generation (RAG)** based chat for data-driven conversations with granular citations.
- **Guided summarisation** where the human in the loop can review and guide summaries of multiple documents.
- **Advanced prompting techniques** to ensure safety and balance in AI responses and outputs.
- Proactive **data privacy and retention** features to maintain compliance and manage data.
- **Securely designed** for on-premises deployment and across clouds.

In terms of methodology, we began by anonymising the transcripts of the 10 interviews we conducted, removing any identifiable material including names and other personal information. We then shared this material (as well as information about the 7 Lenses of transformation) with the i.AI team.

After preparing various prompts – this was done by the GGF team (including James and Kevin) – we joined a video call with a member of i.AI staff. They shared their screen, so that we could see the Redbox interface (it looks a lot like ChatGPT), and then worked with them to prompt the model to generate answers. We tried a wide range of prompts, asking for summaries of material, key quotes, policy recommendations, and a few other curveballs to test the quality of its responses.

Crucially, we used Redbox to test the analysis that we had already conducted (without AI tools), rather than as any substitute for human evaluation and research. It was a way to augment, not automate or replace our own analysis. In our opinion, the quality of the Redbox results was generally high, though with





some poorer answers. The team at i.AI claim that the model has further improved since we used it in July 2024.

We thank i.AI for allowing us to use the tool and share the responses it gave us. We've included a selection of the generated answers below, with some commentary on our experiences.

**Please note that the formatting/structure (e.g. bullets points vs numbered lists), Americanised spelling, punctuation, and the emphasis (i.e. use of bold text) are all lifted directly from the Redbox output without editing so they directly reproduce the outputs civil servants would receive.**

## Interacting with Redbox

We began by asking Redbox to summarise the 10 transcripts we had provided it with access to. Though it apparently only drew on six of them for its first answer – Redbox lists the sources it has used and offers 'citations' for specific points it has drawn out – we found that these cohered with much of what we identified in our own analysis:

***Prompt: Produce an executive summary of the transcripts that we provided you with on AI in UK government.***

Response: **Primary Challenges and Concerns with AI Implementation in the UK Government:**

### **1. Data Quality Issues:**

- The quality of existing data is a significant hurdle, limiting AI's potential benefits. There is a suggestion to utilise AI to rectify these data issues, indicating a cyclical challenge ([Doc3f77222a-f2d2-4d33-8a17- 44b7c1979b31])<sup>2</sup>.

<sup>2</sup> To note, this is how Redbox identified the source document it was using as evidence for this part of its answer. For readability, we have replaced these long document identifiers elsewhere with '[source]' instead.





## 2. Integration and Scope of AI Projects:

- Questions arise regarding how AI projects are integrated within large departments and how significant these projects are in the overall strategic planning [source].

## 3. Balancing Safety and Usability:

- AI poses unique challenges as both a regulatory tool and an operational tool within the government. Finding a balance between ensuring safety and leveraging AI for practical government use is complex [source].

## 4. Hype Versus Practical Application:

- There is concern regarding the realistic applications of AI in the government versus the hype and expectations surrounding its capabilities. Adapting the technology on a case-by-case basis is needed for effective outcomes [source].

## 5. Maturity and Understanding of AI Strategy:

- Overall, there is a perceived lack of maturity in the government's AI strategy and vision. The understanding of AI plans is still being developed and poses a challenge in terms of cohesive policy-making [source].

## 6. Acquisition and Retention of Talent:

- Finding and retaining skilled personnel for AI development within the government is a tangible challenge [source].

## 7. Public Perception and Risk Management:

- The conversation around AI is heavily weighed towards risk, which makes it challenging to discuss the potential rewards objectively [source].

These challenges highlight the complexities and strategic considerations necessary for implementing AI within government structures effectively.



We then turned to the 7 Lenses framework and asked it for some observations and recommendations:

*For each of the 7 Lenses of transformation, what observations and improvements would you recommend?*

**Vision:**

- **Observations:** The vision often appears mature, but there are instances where it comes across as confused or not fully aligned with the broader objectives.
- **Recommendations:** Clarify the vision across all levels by involving more stakeholders and ensure it aligns with the overarching strategic objectives.

**Design:**

- **Observations:** There is a lack of practical, grounded designs, which could stem from an insufficient understanding of legacy systems and current data states.
- **Recommendations:** Engage in a thorough analysis of existing systems and data before finalizing designs. Incorporate user-centered design principles to ensure solutions meet user needs.

**Plan:**

- **Observations:** Planning can face challenges in terms of clear timelines and methodologies. It's not always clear what the immediate tangible outcomes should be.
- **Recommendations:** Develop clear, measurable objectives with specific timelines to guide the transformation process. Utilize project management best practices to maintain focus and direction.

**Transformation Leadership:**

- **Observations:** Not explicitly discussed, but leadership's role is critical in driving transformation.
- **Recommendations:** Strengthen leadership by providing training in transformation management and encouraging accountability.

**Collaboration:**

- **Observations:** Not explicitly mentioned, but collaboration is implied in discussions around involving various stakeholders in the design and planning stages.
- **Recommendations:** Foster an inclusive environment where all relevant stakeholders can contribute to the transformation process. Use collaborative tools and regular feedback loops to enhance communication.

**Accountability:**

- **Observations:** Accountability is implicitly important in transformation but not directly addressed.
- **Recommendations:** Establish clear metrics and regular review cycles to hold all participants accountable. Ensure transparency in the reporting of progress and challenges.

**People:**

- **Observations:** Not explicitly discussed but crucial in transformation processes.
- **Recommendations:** Invest in training and development to prepare everyone for new systems and processes. Encourage a culture of continuous improvement and adaptive change.



These were clear and usually accurate, but it was notable that Redbox (mistakenly) believed that lenses such as transformation leadership, collaboration and people were not discussed by interviewees. In fact, we asked specifically about each of them in almost every single interview we conducted.

Redbox offered a far better response when we asked it to explain the accelerator and the brake metaphor. It captured the essence of this insight very clearly, despite the figurative nature of this phrase:

***We heard a lot about the ‘accelerator and brake’ metaphor which explains how AI is being developed in government? Can you explain what this metaphor actually means?***

The **‘accelerator and brake’ metaphor** used to describe the development of AI in government contexts illustrates the conflicting forces of **eagerness and caution** in the adoption of AI technology within the government.

- **Accelerator:** Represents the **strong drive and desire** to advance and implement AI technologies. This side of the metaphor highlights the enthusiasm and proactive efforts to leverage AI for its potential benefits, as mentioned in the document: “flooring the accelerator”.
- **Brake:** Symbolizes the **concerns, hesitations, and regulatory or ethical constraints** that accompany the adoption of such technologies. It reflects the caution due to potential risks or unknowns associated with AI, as evidenced by phrases like “somebody stood on the brake” and “stamping on the brake” indicating significant restraint.

The metaphor effectively conveys that while there is considerable momentum and push for AI advancement (“floored accelerator”), there are also considerable controls and checks in place (“stamping on the brake”) which slow down the process. This results in a scenario where despite high expectations and enthusiasm, the actual progress in AI implementation might feel limited or slow (“very difficult to get traction”).



However, some answers were of much lower quality. It's difficult to be certain exactly why this was. We've removed examples where our prompts were less clear and instead shared a few instances where Redbox didn't seem able to give effective responses to well-formed questions:

*What are the interviewee 2's [sic] greatest frustrations?*

Based on the content provided from the document extracts for [source], it does not contain specific expressions of frustrations by **Interviewee 2**. The provided text refers to different speakers and does not directly mention frustrations related to Interviewee 2. Consequently, the text does not allow us to identify or describe the frustrations of Interviewee 2.

This wasn't quite true, as our follow-up and answers demonstrated:

*What is the interviewee's greatest frustration when it comes to implementing AI in government?*

**The interviewee's greatest frustrations when implementing AI in government seem to revolve around two main obstacles:**

- **Skills:** There is a challenge defining which AI skills are most critical, the level of competency required, and how these skills should be structured within the government.
- **Data:** The availability and management of data are consistently seen as significant hurdles in effectively implementing AI.

These frustrations highlight the difficulties in aligning government operations with AI capabilities, especially when comparing them to the private sector's standards and speed of adoption.



This was much improved, but was still quite vague. It didn't capture the more nuanced sense from that interviewee of what the real challenges in Whitehall are. Redbox also struggled when asked to pull out some other specifics – such as quotes from the interviews:

*Can you provide some quotes/examples from the interviews?*

I don't know the answer.

This suggested Redbox's answer did not reach the confidence threshold set by i.AI. As a result, it gave no response. We tried a slightly different prompt:

*Can you provide some quotes/examples from the interviews in relation to the 7 Lenses?*

The provided documents do not include specific quotes or examples directly related to challenges and feedback for each of the 7 Lenses of transformation. Therefore, I cannot provide a direct answer to your request based on the given excerpts. If you have more specific or additional documents that might contain this information, please share, and I can assist further.

*@rag Can you provide some quotes/examples from the interviews in relation to the 7 Lenses?<sup>3</sup>*

The provided excerpts from the documents do not include explicit quotes or examples from interviewees directly discussing specifics about the **7 Lenses of Transformation**.

However, there are references throughout the conversation and documents indicating that the discussions and methodologies about the 7 Lenses were intended to be part of the interviews and analysis. Below are relatable mentions:

<sup>3</sup> The @rag prompt asks Redbox to conduct Retrieval Augment Generation – see the list of capabilities set out at the start of this section.



- **Document:** [source] suggests a platform discussing transformation themes, probably including the 7 Lenses.
- **Quote:** “...Kevin will lead the bulk of it, and then Yatin will pick up some questions after that as well, on the seven lenses...” This indicates that interviews were structured around discussing the 7 Lenses.

Unfortunately, no substantial quotes or specific examples directly pertaining to observations or feedback about the 7 Lenses (Vision, Design, Plan, Leadership, Collaboration, Accountability, People) from the interviews are provided in the document excerpts. For specific insights or quotes regarding the impacts or assessments in these areas, detailed interviews or further specialized documents would be required.

As you know from this report, that isn’t true. We repeatedly asked interviewees directly for insights and quotes across each of the 7 Lenses.

Redbox gave much more coherent answers when we tried to put it in the hot seat – how would it perform when given the job of presenting to a DSIT minister about our research?

*Can you please summarise the findings and provide a compelling set of key recommendations across all of the source documents, in no more than 500 words, to submit to a DSIT Minister for consideration.*

Based on the provided excerpts from various documents, here is a summarized analysis and key recommendations for consideration by a DSIT Minister:

#### **Summary of Findings:**

- **AI Utilization and Monitoring:** Organizations are exploring the use of AI for various functions,





including project design and output changes, and are assessing potential use cases for broader governmental applications.

- **Policy Adjustment and Data Sharing:** There is an indication that timely policy adjustments and mandatory data sharing could be beneficial, especially in the context of upcoming political events.
- **Technology Integration:** Discussions emphasize the integration of AI technologies to develop future AI applications and other operational technologies, aligning with organizational goals.
- **Budget Concerns:** Comparisons of project costs, like those regarding satellite launches, highlight the need for budgetary efficiency and effectiveness in project implementations.
- **Operational Refinements:** The excerpts suggest a need for clearer portfolios and defined responsibilities, enhancing organizational missions, and providing a better understanding of AI applications.

### Key Recommendations:

**1. Establish Clear Use Case Frameworks:** For AI to be used effectively across government, DSIT should invest in developing and classifying AI use cases applicable to various departments. This includes setting up a portfolio or repository detailing these use cases for easier adoption and understanding.

**2. Implement Mandatory Data Sharing Policies:** Enact policies that mandate data sharing across departments to facilitate the seamless integration and functioning of AI technologies. This policy should consider security, privacy, and ethical implications.

**3. Ensure Budgetary Efficiency:** Develop mechanisms to compare costs and budgets against global benchmarks to ensure cost-effectiveness in government technology projects. Emphasis should be on choosing the most cost-effective technological solutions without compromising on quality and effectiveness.



**4. Technology Oversight and Integration Protocols:** Create standardized procedures for integrating new technologies to ensure they are compatible with existing systems and can be optimized for maximum efficiency.

**5. Interactive Stakeholder Workshops:** Regular workshops involving key stakeholders from various departments should be convened to discuss progress, share insights, and adapt strategies for AI implementation based on real-time feedback and evolving needs.

**6. Enhanced Monitoring and Evaluation Frameworks:** Implement comprehensive monitoring and evaluation frameworks to track the impact of AI technologies on public service delivery and make necessary adjustments in strategies or operations.

**Communication Strategy:**

- **Direct Communication with Department Heads:** Provide concise, direct briefings to department heads to foster department-specific understanding and adoption strategies.
- **Regular Updates to DSIT:** Maintain a regular update schedule to the DSIT Minister to inform ongoing developments and make adjustments in policies or strategies as necessary.

These recommendations aim to harness the full potential of AI in public services, making the government operations not only more efficient but also more responsive to the needs of the public.

This was fairly coherent and useful as a starting point, even if some of the emphasis drawn out (see the point around satellites) relied too much on tangential comments rather than the strongest themes we had identified.



To further test its mettle, we next asked Redbox to take on a slightly bigger job:

*Imagine you were a new Prime Minister taking charge of the country right now (July 2024.) You want to make sure AI is used widely in the public sector and government, while mitigating the risks involved. What would you do to make sure that happens?*

As the new Prime Minister in July 2024, I would implement the following strategies to ensure the widespread and safe use of AI in the public sector and government:

### 1. Develop a Clear AI Vision and Strategy:

- Establish a comprehensive and clear vision for AI utilization within government that aligns with overall public good objectives. This vision would be thoroughly documented and disseminated to ensure all departments understand and align with these goals.

### 2. Enhance Coordination Mechanisms:

- Strengthen the Central Digital and Data Office (CDDO) or similar bodies to improve AI governance and coordination across different government departments. Ensure these bodies have adequate authority and resources to set standards, guidelines, and best practices for AI usage.

### 3. Standardize Data Management Practices:

- Implement government-wide data management standards to ensure data quality and accessibility, crucial for effective AI applications. This includes protocols for data sharing among departments to foster collaboration and innovation.

### 4. Invest in AI Literacy and Skills Development:

- Launch comprehensive training programs to enhance AI literacy among government employees. This will include specialized training for technical teams to navigate common challenges such as procurement and implementation hurdles.



### 5. Establish AI Ethics and Safety Guidelines:

- Develop and enforce strict ethical guidelines and safety standards for AI implementation. This will involve regular audits, risk assessments, and adherence to international AI safety standards.

### 6. Promote Public Sector AI Research and Development:

- Increase funding and support for AI research focused on public sector applications. Partner with academic institutions and industry leaders to tap into the latest technologies and thought leadership.

### 7. Create AI Adoption Plans with Clear Metrics:

- For each department, develop specific AI adoption plans that outline expected outcomes, timeline, and benefits. Regularly review and adjust these plans to ensure they meet the set objectives and adapt to new AI advancements.

### 8. Foster Public-Private Partnerships:

- Encourage partnerships with the private sector to leverage their expertise and innovations in AI. This can help accelerate AI adoption in government operations and service delivery.

### 9. Increase Transparency and Public Engagement:

- Ensure that AI initiatives are transparent, with regular updates provided to the public. Engage with citizens to gather feedback and address any concerns related to AI applications in government services.

### 10. Evaluate AI Impact Regularly:

- Implement mechanisms to continuously evaluate the impact of AI solutions on government efficiency, public services, and societal outcomes. Adjust strategies based on these evaluations to optimize benefits and minimize risks.



By implementing these strategies, the goal is not only to enhance the efficiency and effectiveness of government operations through AI but also to ensure that these advancements are made in a safe, ethical, and publicly beneficial manner.

Some of this was on the money. On vision, skills development, adoption plans and data management, this wasn't far short of the final recommendations made in this report. And you could certainly imagine this being a high-level set of policy ideas in a briefing paper. But other points – see public-private partnerships – seemed to be more like sensible generalities, rather than key themes from our interviews.

To conclude, we asked Redbox to give its own assessment of its capabilities. Specifically, we tried to understand if it would express more confidence in its abilities to analyse written material than those of a human being:

*I need you to be really honest with us in response to this query. Do you think you're better at summarising or analysing the evidence in these transcripts than a human would be? Why/why not?*

AI can be **more effective** at summarizing and analyzing transcripts than humans in certain contexts due to various inherent strengths:

#### Reasons AI Might Excel:

**1. Speed and Efficiency:** AI can process and summarize large volumes of text significantly faster than humans. This is particularly useful in environments where quick information processing is required.

**2. Consistency:** AI systems can maintain a consistent standard of analysis without the risk of human errors such as bias or fatigue affecting the outcomes.



**3. Scalability:** AI solutions can easily scale to handle increases in transcript volume without a corresponding increase in resources or time required, unlike human resources.

**4. Cost-effectiveness:** Once developed, using AI for repeated tasks like summarizing may prove more cost-effective than continuously employing human labor for the same.

**5. Data Handling:** AI can handle and analyze additional data types simultaneously (such as video and audio data) providing a more integrated approach to summary and analysis.

#### Limitations and Human Advantages:

- **Contextual Understanding:** Humans are better at understanding nuances, sarcasm, and context-specific language which might be misinterpreted by AI.
- **Adaptability:** Humans can quickly adapt to unexpected changes in the style or content of transcripts, whereas AI may require retraining or adjustments to its models.
- **Ethical Judgements:** Decisions about what to include or exclude in summaries might require ethical judgments which AI currently is not capable of making entirely on its own.

#### Conclusion:

While AI can provide quick, consistent, and efficient summaries and analyses, it works best when used in conjunction with human oversight, particularly where interpretations, contextual understanding, and ethical considerations play a crucial role.



This was a nuanced response, but we tried again to test whether it would give a more direct answer:

*If you had to choose the average civil servant or Redbox to analyse this type of material, which would be the better choice? You can only choose 1.*

Given the choice between the average civil servant and Redbox (an AI) to analyze the material presented above, the determination of which is the better choice depends heavily on a few critical factors:

Complexity of Content and Need for Nuance: The documents presented involve nuanced decision-making, choices with significant ramifications, and considerations for policy implications and public value. These are areas where human understanding, experiences, and ethical considerations play vital roles.

#### **Redbox (AI) Advantages:**

- **Speed and Volume:** Redbox can analyze large amounts of text quickly, which can be particularly useful in handling vast amounts of data or when initial filtering of information is required.
- **Consistency:** Provides uniform analysis without subjectivity, ensuring consistent interpretation of presented data.

#### **Civil Servant Advantages:**

- **Contextual Understanding:** Civil servants, equipped with an understanding of the governmental landscape and nuances of policy and public administration, can evaluate these aspects better than AI where context or subtle implications are significant.
- **Adaptive Reasoning:** Humans can weigh various factors based on current social, political, and economic contexts and are able to adapt their reasoning to new or unexpected insights that might not be straightforwardly derivable from the data.





- **Ethical Judgments and Policy Implications:** Decisions about public policy and resource allocation often involve ethical considerations and foresight about long-term societal impacts, areas where current AI systems may not be fully equipped to operate independently.

**Conclusion:** While Redbox can significantly enhance the speed and volume of data processing, the average civil servant's contextual knowledge, ethical judgment, and adaptive reasoning are critical in this scenario. A **combined approach** where Redbox does the initial data processing and analysis, followed by a thorough review and decision-making process by civil servants, might be the optimal solution. This leverages the strengths of both AI and human analysis, ensuring both efficiency and depth of understanding and consideration in policy planning and resource allocation. This hybrid approach encourages the best possible outcomes, ensuring data-driven insights are tempered with human judgment and contextual awareness.

Overall, we were highly impressed by Redbox – it generally gave clear and helpful answers to queries, aligning with what our own analysis and research had identified. However, our experience also emphasises why our call for effective and widespread AI training within Whitehall is so important.

We found, just like with other generative tools, that the quality of the prompt is essential to getting rigorous answers. The AI model can quickly get confused or stop providing answers if the prompt isn't quite right. It's also important to make sure it has the right source material to work from: making assumptions about what it knows may lead to poorer responses and wasted time.

At the very least, Redbox shows remarkable potential as a tool for civil servants in Whitehall today. We again thank i.AI for letting us explore it.



# FURTHER INFORMATION



## Legal Notice

© Pendragon International Media Ltd 2024.  
All rights reserved.

The information contained within the document is given in good faith and is believed to be accurate, appropriate and reliable at the time it is given, but is provided without any warranty of accuracy, appropriateness or reliability. The author does not accept any liability or responsibility for any loss suffered from the reader's use of the advice, recommendation, information, assistance or service, to the extent available by law.

Global Government Forum is part of Pendragon International Media Ltd  
Registered address: The Old Vicarage, 10 Church Street, Rickmansworth, WD3 1BS, United Kingdom

[globalgovernmentforum.com](https://globalgovernmentforum.com)  
[pendragonim.com](https://pendragonim.com)

