

From Playbook to Progress

Delivering AI in the public sector

A practical guide to turning
ideas into impact—securely, at
scale and at pace



Why this guide? And why now?

You've read the playbooks. You've seen the promises. AI is meant to transform public services, ease workloads, and deliver significant results. But here's the reality: according to McKinsey, only 1% of organisations have fully integrated AI into their workflows to genuinely unlock its benefits. The vast majority of projects stall at the pilot phase, never reaching full-scale implementation.

[Read on](#)

Contents

The public sector AI landscape. What's changed?	4
2025 and beyond: moving from experiment to reality	9
The delivery challenge	13
The path to production	15
What good looks like	19
Getting started: 5 practical steps	21
A partnership for delivery	24

Recognising this, the Department for Science, Innovation and Technology (DSIT) launched initiatives such as the AI Playbook and commissioned the AI Opportunities Action Plan—both designed to guide public sector teams toward successful and responsible AI deployment.

More recently, the Cabinet Office published The People Factor: A human-centred approach to scaling AI tools, building on current guidance to advocate an approach that prioritises organisational culture, skills, and behavioural change alongside technology.

Public sector budgets are under pressure, but expectations keep growing. In a 2024 UK trial with 20,000 civil servants, AI tools saved staff 26 minutes a day—adding up to nearly two extra weeks per person each year. That means more capacity, less admin, and better service delivery.

Some departments are already seeing the benefits. Now it's time to scale.

This guide bridges the gap between ambition and reality, helping you move swiftly from promising experiments to meaningful outcomes—delivered securely, responsibly, and at pace.

Let's get into it.



Who is this guide for?

This guide is aimed squarely at leaders and operational teams across government departments and agencies—whether you're heading digital transformation, leading service delivery, overseeing data governance, or managing frontline services, this guide is for you. So you can bridge the gap from promising proofs of concept (PoC) to practical AI solutions that save time, reduce costs, and directly improve citizen services.

Get in touch



The public sector AI landscape. What's changed?

AI in government isn't a fringe experiment anymore. With £45bn to save across government, AI will play a significant part and has now become official strategy. Momentum has surged over the past year and the AI Playbook for Government and the AI Opportunities Action Plan both lay down a clear expectation that AI should be used to improve services, drive productivity, and support better outcomes for citizens.

Read on



The trouble is many public sector organisations are stuck in the pilot phase

The desire and drive for AI has never been stronger and departments and public bodies are being asked to turn policy into working solutions at speed. But many teams are still wrestling with the same questions, such as “Where do we start?” “Is our data good enough?” “How do we scale something safely?” Or often: “What does ‘good’ even look like?”

There’s broad agreement on the potential. But the pressure to deliver has exposed a readiness gap. Many departments are still early in their AI journey—experimenting with pilots, wrestling with legacy infrastructure and siloed data, or trying to build capability on the fly. A few are further ahead, but no one’s cracked it completely.

That’s because the shift from concept to reality isn’t just technical—as highlighted by the Cabinet Office’s People Factor paper and framework. It’s cultural, operational,

and deeply tied to service design. Becoming AI native—ensuring every process can benefit from intelligent automation, every person is equipped to work alongside AI, and every solution is designed with AI in mind—is critical to success.

And with service delivery the top priority, the stakes are high and margin for error small. Get it wrong, and you risk public trust, legal challenge, or unintended harm to the people who rely most on your service.

That’s why practical, grounded guidance matters more than ever.

In the private sector, AI maturity varies—but the commercial incentive to automate, personalise, and optimise has driven faster adoption. Many firms now have embedded AI in customer service, logistics, fraud prevention and product development.



What does it mean to be AI native?

Being AI native means embedding AI at the heart of how your department operates, serves users, and solves problems. It’s about empowering every team with AI capabilities, enhancing processes through automation and intelligence, and designing services and solutions with AI as a core element—not an afterthought. In practice, becoming AI native transforms the way teams work, helping you deliver more efficient, innovative, and responsive public services—all grounded in a responsible approach to adoption.

[Get in touch](#)


Key points from the guidance documents

The AI Playbook for Government sets out practical guardrails to help public sector teams use AI safely and effectively. It emphasises principles like meaningful human oversight, transparency, explainability, and the need to design for inclusion. Rather than treating AI as a black box, teams are urged to understand how systems make decisions and to be able to explain those decisions to the public. The playbook also stresses the importance of robust data governance, bias mitigation, and clear accountability, especially when AI is used in high-impact areas like policing, welfare, or healthcare.

The AI Opportunities Action Plan shifts the focus from ambition to implementation, highlighting practical steps for laying strong foundations and achieving citizen impact. It specifically calls out key steps including:



Investing in data foundations. Quickly identifying and unlocking high-impact public datasets, ensuring they're high-quality, accessible, and responsibly managed to accelerate innovation



Scaling what works. Implementing a structured "Scan, Pilot, Scale" framework, allowing departments to swiftly identify viable AI opportunities, rapidly prototype them, and then efficiently scale solutions that demonstrate clear value



Building skills and adoption. Creating dedicated pathways and programmes to rapidly develop the necessary AI skills across the public sector—ensuring AI literacy isn't limited to technical roles but integrated into operational teams



Most importantly, the plan strongly advocates a citizen-first approach, prioritising AI use cases that directly improve public services and experiences for both citizens and employees. Real-world examples include AI assistants reducing administrative burdens by up to 20% for teachers, AI-driven report drafting cutting document creation times significantly in healthcare and professional services, and automated anomaly detection helping police forces rapidly identify threats. These practical, grounded applications not only streamline processes but free staff to focus more of their time on frontline, high-value tasks, ultimately delivering faster, better, and more responsive services to citizens.

Complementing these practical steps, the Cabinet Office's People Factor specifically addresses the human and organisational dimensions of AI adoption. It highlights the critical importance of user-focused strategies, structured into three key stages:



Adopt: encouraging initial uptake through targeted support, clear communication, and practical training.



Sustain: embedding AI into daily workflows by helping users develop consistent habits and routines around the tools.



Optimise: continuously refining how AI tools are used in practice, proactively managing risks, and enhancing user skills and confidence.

This structured, people-centric approach ensures AI solutions genuinely integrate into everyday operations, driving lasting impact.

Each of these documents clearly signal that the era of isolated pilots is ending. More importantly, they represent a fundamental shift in how government views and uses AI—moving away from technology-driven experiments towards solutions deeply embedded in real-world service delivery. This matters, because lasting change won't come from flashy one-off demonstrations. It demands sustained focus, clear foundations, and genuine alignment with citizens' and employees' everyday needs. What's needed now is scale, momentum, and joined-up thinking to transform AI from abstract strategy into sustained, citizen-centric impact.

Why it's important to act now

The pressure to modernise public services isn't new, of course, but AI adds real opportunity for departments to do an awful lot more for less. Rising demand, tighter budgets, and growing citizen expectations mean that departments simply can't afford to wait. At the same time, the available tools are evolving fast, from automation and generative AI to agentic AI. If departments don't act, they risk falling behind, duplicating effort and wasting time and resources, or investing in solutions that don't scale.

But it's not just about seizing opportunity—it's about managing risk. Without clear direction and shared standards, AI rollouts have the potential to erode trust, make inequality worse, or simply fail to deliver. Getting ahead of that risk means investing in skills, governance, and delivery capability today—not tomorrow.

[Get in touch](#)

2025 and beyond: moving from experiment to reality

Every team is at a different stage of delivering AI at scale, but most have now moved beyond experimentation. Last year was about exploring possibilities; this year is about practical implementation, operational excellence, and scaling what works. Here are four common scenarios we're seeing across government—alongside critical considerations to keep front of mind.

[Read on](#)



01



**“We’ve experimented
—now we need to scale”**

You’ve completed pilots and proven the tech works. Now it’s time to scale and operationalise. At this stage, the focus shifts from proving concepts to robust, sustainable delivery. It’s about standardising infrastructure, embedding consistent data governance, and rolling out adoption plans that secure real buy-in from users.

Critical considerations.

Clearly document outcomes, secure senior sponsorship, and build scalable infrastructure and processes.

02



**“We’ve proven value, but we’re
stuck operationalising it”**

Your pilots showed clear promise, but integration into day-to-day operations is proving challenging. Here the blockers typically aren’t about the tech itself—they’re about cultural resistance, data integration hurdles, or gaps in skills and processes needed for sustainable operational delivery.

Critical considerations.

Involve operational teams early to identify barriers and proactively plan for change management and skills development.

03



“We have multiple solutions running, but they’re fragmented”

You have several AI solutions deployed, but they aren’t yet joined up or delivering coherent departmental impact. Fragmented solutions often duplicate efforts, waste resources, and create confusion. Now’s the time to step back, consolidate, standardise your approach, and embed unified governance and oversight.

Critical considerations:

Unify governance, adopt interoperable standards, and design modular infrastructure to support cohesive delivery.

04



“We’re delivering, but need to demonstrate impact clearly”

Your AI projects are operational, but now you face pressure to clearly demonstrate value to senior stakeholders, funders, and citizens. At this advanced stage, measurement, transparency, and ongoing refinement become crucial. Effective monitoring and evaluation are no longer optional—they’re critical.

Critical considerations:

Implement clear performance frameworks and consistently measure, refine, and communicate outcomes linked to strategic priorities.



If none of these sound like you, try mapping your work against the question:

“What problem are we trying to solve ...and why now?”

The answer is often a better compass than any strategy document.

When moving AI from experiment to reality makes sense (and when it doesn't) A visual guide for public sector teams

When it makes sense



Proven, measurable benefits: Clear outcomes from pilots showing tangible value



Repeatable, high-volume processes: Frequent, routine tasks that clearly benefit from automation



Well-structured data environments: Accessible, clean, and consistent data available



Clear citizen or staff impact: Directly improves user experience or reduces staff burden



Sustainable infrastructure: AI solutions built to scale and integrate seamlessly with existing platforms

When it doesn't



Unclear value or evidence: Benefits haven't been consistently proven or documented yet



Persistent operational barriers: Adoption, cultural issues, or skill gaps remain unresolved



Data readiness issues: Data quality, accessibility, or consistency hasn't been sufficiently addressed



Weak governance or oversight: Lacking clear accountability structures or ethical guardrails



Strategically misaligned solutions: Not clearly linked to departmental priorities or citizen needs

Get in touch

The delivery challenge

As highlighted in the previous section, most teams now face the challenge of moving beyond successful pilots into scalable, operational solutions. But let's be honest - getting AI into production is hard and even great ideas can stall. You can have a great use case, a capable team, and support from the top...and still get stuck.

The numbers tell the story: according to a recent study by McKinsey, only 6% of GenAI PoCs actually make it to production. That's a lot of effort with very little impact. So, what's going wrong?

Read on





Data that's not ready:

Poor quality, patchy, or siloed data is undoubtedly the biggest barrier to getting out of PoC purgatory. Without clean, consistent information to learn from and work with, even the smartest model will underdeliver



User resistance:

Even the best tools will fail if people don't use them. If your frontline teams aren't involved early or don't see how the tool helps them, then adoption inevitably suffers



Unclear business cases:

If the value's not obvious or the costs are hard to justify, it's tough to get buy-in for your AI project beyond the pilot phase. That's especially true when budget holders want fast, measurable results




Ethics and compliance concerns:

Public sector teams tread carefully (and rightly so). But concerns around explainability, oversight, or bias can become dealbreakers if they're not addressed head-on




Lack of sponsorship: Projects often stall without senior champions who can unblock issues, defend funding, and build cross-department support

[Get in touch](#)

The path to production

Getting AI from pilot to everyday use isn't about having the shiniest tech stack. It's about solving real problems for real people and doing it in a way that sticks. The difference between projects that scale and those that gather dust often comes down to a few fundamentals that have nothing to do with algorithms.

[Read on](#)





Process beats tech every time

The specific AI tool you choose today might be outdated in eighteen months. That's fine - the technology landscape moves fast, and your approach needs to flex with it. What matters more is how you work: how you gather requirements, how you test, how you iterate, and how you learn. Build ways of working that can adapt to new tools rather than getting locked into today's technology choices. Think platforms and processes, not point solutions.



Create clear governance that enables, not blocks


Set up decision-making structures that can move at the speed AI projects need. Establish an AI steering group with clear escalation paths and decision rights. Define what constitutes an acceptable risk versus a show-stopper. Create lightweight approval processes for common changes such as model updates, feature additions, or performance tweaks. Good AI governance feels like guardrails, not roadblocks.



Communicate like you mean it

Don't just announce when you're going live. Talk about what you're building, why you're building it, and what it means for people's daily work. Regular updates, honest about both progress and challenges, build trust and reduce resistance. Create space for questions, feedback and even pushback. If people understand the problem you're solving and how it affects them, they're more likely to engage constructively rather than resist change.





Invest in capability, not just technology

Training isn't a tick-box exercise at the end. It's ongoing capability building that starts before you build anything. Help people understand not just how to use new tools, but why they work the way they do and how to spot when something's going wrong. Build internal expertise so you're not dependent on vendors for every adjustment. The goal isn't just to use AI tools, it's to become an organisation that can work intelligently with AI.



Design for tomorrow's changes

Assume your needs will evolve and your tools will change. Build on open standards where possible. Document your decisions and reasoning so future teams can understand why choices were made. Create feedback loops that help you spot when something's not working or when needs have shifted. Plan regular review cycles - not just "is it working?" but "is it still solving the right problem?" The most successful projects aren't the ones that solve today's problem perfectly—they're the ones that create the foundation for solving tomorrow's problems too.



Remember that delivering your successful AI project isn't a straight line.

Things will go wrong.

But recognising the traps early (and designing to avoid them) is what separates stalled pilots from sustainable success. That's what delivery means and that's what this next section is all about.

[Get in touch](#)

What good looks like

AI doesn't have to change everything to make a difference.

In fact, many of the most effective and valuable projects are focused, low-risk, and quietly transformative - directly linked to improving how services are delivered. They automate routine tasks, help teams with complex decisions, or make public services easier to access. The public sector doesn't need sci-fi use cases, it needs grounded, proven applications that ease pressure, cut delays, and support staff.

[Read on](#)

The real-world examples below were chosen from our own work and government pilots because they show real impact, delivered at pace, without over-engineering. These are what good looks like, showing where AI is already making a difference on the ground and helping teams work smarter, not harder.



Personalised communication at Cafcass

Client: Cafcass - the organisation representing children in England's family court system.

Solutions delivered:

AI-powered tool embedded in their case management system to personalise and simplify communication.

Key benefits:

Clearer letters, better accessibility, improved efficiency, and stronger engagement with children and families.

Challenge:

Cafcass needed to improve engagement with children and families during court cases. Family Court Advisors (FCAs) were spending too much time editing templates in an outdated system— ChildFirst—that

made it hard to use existing case data. With 1,500 FCAs generating over 80,000 letters a month, the process was inefficient and failed to meet the needs of children, especially those with different learning needs or language barriers.

Solution:

Partnering with Version 1, Cafcass developed Scribe, an AI-powered feature added directly into ChildFirst.

Using Microsoft Azure AI services, Scribe helps FCAs generate personalised letters quickly and easily. It pulls in relevant data, adjusts language complexity to suit a child's age and comprehension, and includes tools like translation and audio versions for accessibility.

Scribe improves how Cafcass communicates by making sure each letter is clear, relevant, and easy to understand. It cuts down on admin work and gives FCAs more time to support families. Built-in security and compliance features also ensure data is handled safely and responsibly.

With Scribe, communication is no longer one-size-fits-all—it's personalised, inclusive, and effective.

Outcomes:

Personalised communication:

Letters tailored to age, reading level, and preferences.

Improved accessibility:

Includes translation, audio, and neurodiverse-friendly options.

Better engagement:

Clearer messages that children and families can understand.

Increased efficiency:

Reduces manual edits and admin workload.

Better use of data:

Pulls accurate info directly from ChildFirst.

Customer satisfaction:

Builds trust with clearer, more relevant content.

Compliance and security:

Ensures privacy and meets AI governance standards.



Scan me to find out more

Getting started: 5 practical steps

You don't need perfect conditions or a ten-year plan to start delivering with AI. But you do need a first beyond PoCs. Start by aiming to embed an AI native mindset— where people have the right skills, processes are intelligently automated, and AI is integral to delivering outcomes. Success depends on clear departmental alignment, robust leadership, and a structured approach to embedding AI across your teams.

Read on

Here are five clear steps to get you moving:

01



Set a clear vision and mission


AI initiatives thrive with clarity. Start by defining precisely what AI success looks like for your department and ensure everyone understands how AI aligns to your strategic priorities. Clearly articulating your AI vision and mission—like Version 1's ambition to become an exemplar AI native organisation—establishes a common direction, creates alignment, and accelerates practical delivery.

02




Secure visible senior sponsorship

Gaining executive-level support isn't optional—it's essential. Establish strong sponsorship from senior leaders who actively champion AI initiatives, clear barriers, secure resources, and visibly drive adoption across all departmental functions. Effective senior sponsorship dramatically increases the likelihood of moving beyond experimentation to tangible, scalable success.



03



Adopt structured change management

Use structured frameworks—like the ADKAR model (Awareness, Desire, Knowledge, Ability, Reinforcement)—to systematically guide your teams through AI-driven change. Communicate clearly, explain benefits explicitly, and provide structured, continuous support to overcome resistance and build excitement

04



Engage, educate, experiment

AI skills development is key. Implement tailored, practical training programmes that encourage your people to engage, educate themselves through role-specific learning, and immediately experiment to embed their new knowledge. This hands-on, practical approach builds lasting capability rather than superficial familiarity

05



Embrace a human-in-the-loop approach

People—not technology—remain central. Clearly define roles such as AI trainers, validators, orchestrators, collaborators, and decision-makers to ensure human oversight and ethical AI use. This “human-in-the-loop” model builds trust, reduces fear, and enhances productivity by creating clear boundaries and responsibilities

Following these steps won't just help you move from experimentation to execution—they'll enable you to scale your AI initiatives and deliver lasting impact for your teams, your citizens, and your department as a whole.

[Get in touch](#)



A partnership for delivery

Starting matters. But delivery is what counts.

That's where we support public sector teams—not just with ideas, but with the experience to turn them into working services. Our approach is delivery-first, people-centred, and grounded in the realities of government—and helping you take the right steps to becoming that truly efficient and productive, AI native department. We know that success isn't about flashy pilots. It's about solving the right problems, building trust, and making progress in the open.

[Read on](#)

We've supported AI, automation and data transformation projects across departments and agencies—always with a focus on outcomes, not hype. From virtual assistants to risk analytics and intelligent triage, we've helped teams move from exploration to live services that save time, improve access, and support public value.


Where it makes sense, we bring accelerators and delivery frameworks—like our AI Labs, automation tooling, or virtual assistant starter kits—all designed with ethical use and inclusion in mind. And through strategic partnerships with Microsoft and AWS, we help teams tap into secure infrastructure, innovation support, and trusted platforms.

But most of all, we work in partnership. That means listening closely, sharing what works, and helping teams deliver AI in a way that's practical and real.

If you're stuck, unsure where to begin,
or ready to scale, we're here to help.

Get in touch to see how we can support your
next step—from sandbox to service delivery.

ali.bovis@version1.com





Talk to us today

E: ali.bovis@version1.com

version1.com