

Insights from
building PolicyNova

Fast tools, high stakes: AI and the policy professional

At a glance

- From boosting public sector efficiency to building AI scenarios, governments worldwide are advancing AI initiatives for policy research and design.
- AI adoption among policy professionals is growing, but remains uneven, with 56% using it only weekly or not at all. This reflects deliberate caution, shaped by the high stakes of policymaking.
- Most off-the-shelf AI tools don't explain how they arrived at an answer, which makes it hard for policy professionals to introduce AI into a process that demands substantiation.
- Launching in October 2026, PolicyNova keeps policy professionals in the loop. The tool grounds its outputs in traceable, trusted evidence and helps policymakers assess whether interventions that worked elsewhere might fit their own context.

Across most of government, AI adoption and readiness are surging. In policymaking, it's split, with 44% of policy professionals using AI daily, while one in ten don't use it at all.

Part of the split might be attributed to nobody having a clear picture of what good looks like. How does the AI-enabled policy professional work? Do they approach policy design differently? Are they confident in how AI is being used to improve outcomes for the public? These questions are top-of-mind for policy professionals navigating the introduction of AI into their work, and into policy itself.

Apolitical's research into the problems and solutions public servants run into day-to-day offers some answers, while simultaneously accentuating the need for a global community of policymakers grappling with the thorny new challenges AI brings. In this briefing, Apolitical shares the data collected while developing PolicyNova, a purpose-built AI tool for policy research, launching in October 2026.

Apolitical began developing PolicyNova by asking, if we look at the problems policy professionals are running into day-to-day, which ones could AI help resolve, and which ones is AI not suited to support?

The following are the conditions in which AI is being introduced or adopted:

- There is a contradiction at the core of policy work, where it is relentless yet painstakingly slow.
- Individuals work under time pressure and resource constraints, but often without a clear, straightforward path to implementation.
- Political priorities frequently override evidence-based recommendations.
- Siloed information blocks effective communication and collaboration.
- There's too much information to parse through.

In interrogating the nature of these problems, Apolitical learned a lot about why policy professionals have difficulties with incorporating AI into their function's work, despite it seeming to have core tasks that lend themselves to support from AI.

In a domain where recommendations directly shape policy outcomes, the risk of hallucination or factual error is too high to accept. Simultaneously, policy professionals told us that off-the-shelf AI models are trained on data representative of only a small slice of the world, and they're uncertain about applying them to local policy problems. They also question whether AI can produce a truly novel solution, and not just a well-constructed version of what already exists.

At the same time, policy professionals also recognise that AI has the potential to fundamentally transform their work. AI tools can enable a continuous loop of feedback and evidence, supporting research and analysis in ways that can reshape the menu of solutions available to any government.

Join the discussion

Do you have an insight or point of view to share on AI in policy?
Head to Apolitical's Policymaking Community to join the conversation.



Some **questions** to get you started:

- How are you using AI in your policy work?
- How is AI helping you?
- Have you come across any interesting or creative uses of AI?

AI in policy, today

Current applications of AI in policy are aimed at reducing the large pile of paper endemic to work in the profession. There is more of it than any single individual can process, let alone absorb, including evidence reviews, briefing notes, consultation responses, hours of committee hearings and council decisions, to name a few. AI's first intervention has been to make that high volume of documentation readable. Through [Government AI Navigator](#), a living, growing library of government AI projects worldwide, Apolitical tracks what AI use in government looks like. The examples below are drawn from that library.

NAME OF PROJECT	COUNTRY	DESCRIPTION	INNOVATION IN POLICY /SOME KIND OF FUTURE GLANCE	SCALE OF IMPACT
<u>Parlex</u>	UK	Parlex uses AI to analyse years of parliamentary debate transcripts from MPs and Peers, forecasting how Parliament might react to a new policy if it were debated tomorrow.	AI-powered analysis turns years of otherwise-wasted parliamentary debate data into predictive stakeholder positioning.	Pilot/small team
<u>YouTube-to-docs</u>	Maryland, USA	A free, in-house tool that turns YouTube recordings of long public meetings into transcripts, structured summaries and other written content.	A hearing that would take hours to re-watch can be turned into a summary and searchable transcript that takes minutes to review.	Regionally or nationally
<u>AI Assessment Framework (AIAF)</u>	New South Wales, Australia	The NSW Government developed the AI Assessment Framework (AIAF), a mandatory risk assessment tool that every state government agency must use whenever it designs, develops, deploys, procures or uses any system containing AI.	The framework's innovation lies in mandating independent review for the highest-risk AI systems and embedding AI governance within the state's existing oversight structures.	Regionally or nationally
<u>GENAI</u>	Japan	Japan's Digital Agency built a generative in-house AI platform for civil servants called GENAI, offering a wide range of general tools for everyday work, such as chat, drafting, proofreading, summarising and translation, plus applications built for administrative practice.	In-house AI platform built to accelerate and simplify the workflow of civil servants across multiple working dimensions.	Pilot/small team
<u>OpenFisca</u>	France	OpenFisca is a policy simulation framework that models the impact of social and fiscal laws and regulations. It is accessible via open APIs, making it usable by developers, data scientists and researchers.	By using AI to model the impact of policy interventions, policymakers could adjust their approach before piloting changes in the real world.	Regionally or nationally

A **thematic analysis** of the policy-adjacent case studies catalogued in Government AI Navigator surfaced **four key themes**.

01

Governments are investing in building in-house AI platforms.

These are intentionally designed to meet security requirements and fit within existing organisational workflows. While early government AI experiments favoured general chat assistants, the newer generation, from [Japan's GENAI](#) to [Berlin's BärGPT](#), has moved toward multi-purpose platforms that support different dimensions of public servants' everyday work.

02

A rising share of initiatives focus on turning previously unusable institutional archives into evidence that can be acted on.

Data inputs come from a variety of sources, including parliamentary debate transcripts ([Parlex](#)), municipal code and council decisions ([STARA](#)), nightly legislative text ([Legi-Assist](#)) and video recordings of public hearings ([Maryland's YouTube-to-docs](#)). AI is increasingly being explored as a way to standardise how information already held by governments can be accessed, understood and applied.

03

National governments are developing AI governance and frameworks in parallel with deployment.

Initiatives such as [NSW's mandatory AI risk assessment](#), [Singapore's agentic AI framework](#) and [Vietnam's AI law](#) are building oversight infrastructure alongside adoption.

04

Human validation remains central across initiatives, with expert sign-off explicitly retained at every stage.

For instance, [Canada's agricultural synthesis project](#) intentionally uses a human-in-the-loop design, catching inaccuracies before outputs are acted on.

Contribute to Government AI Navigator

Do you have an AI project in policy design or research to share?
Publish it directly through Government AI Navigator's AI-powered assistant.

EXPLORE




GOVERNMENT AI
NAVIGATOR

Caution is a feature of policy, not a bug

While Government AI Navigator demonstrates that AI is being used by policy professionals, the kind of the AI projects they're developing today are cautious ones (the risk of using AI to process and extract insight from large volumes of publicly available documentation is low). This restraint, however, is not stalling adoption, with 44% of policy professionals globally using AI daily. And these two findings are not in tension. Policy professionals use AI a lot, but mostly on tasks where an error is unlikely to prove costly.

FIGURE 1

Frequency of AI use among policy professionals



SOURCE

Apolitical (2026), AI Adoption Survey (n=88). Based on responses to "How often do you use AI at work?" Only responses from policy professionals are included.

This caution and risk aversion that constrain AI use in policy are driven by **three concerns**:

➤ Accuracy

Policy professionals remain sceptical of the validity of AI-generated evidence, having encountered cases where AI tools produced content that sounded credible but was inaccurate or entirely fabricated. While AI tools are recognised as having the potential to accelerate processes, widespread concern remains that improvements in speed could come at the expense of the robustness of the evidence and the recommendations derived from it.

➤ Loss of critical thinking


Another concern is that over-reliance on AI could erode critical thinking and independent analysis. Policy analysts and advisors are reluctant to give up the process of working through issues themselves, valuing the opportunity to apply their own expertise and context-specific judgement. To guard against over-reliance, they often use AI as a starting point or for idea generation, while retaining full responsibility for drafting final recommendations and conducting analysis.

➤ Data security

Policy professionals are cautious about inputting confidential or sensitive information into AI tools, worrying that such information could be leaked or exposed. As a result, 4 in 5 rely exclusively on approved tools, such as Microsoft Copilot, which they trust to meet security and data protection standards¹.

FIGURE 2

Top factors that would improve AI use among policy professionals



SOURCE

Apolitical (2026), Poll results from How to Write the Perfect Policy Brief event (n=241). Based on responses to "Imagine you could improve one thing about using AI in your policy briefs. What would you pick?"


1. Apolitical, AI Adoption Survey (n=88) Based on responses to "Which AI tools do you use at work?" Only responses from policy professionals are included.

While these concerns are shared across government functions, they are more prevalent among policy professionals, where 80% say the possibility of AI making a mistake deters them from using it. This compares with 69% in communications and 68% in delivery and project management. The OECD's Digital Government Outlook 2026 contains a similar conclusion of AI use being more widespread in internal processes and in design and delivery than in policymaking².

For the policy professionals in Apolitical's network, however, none of this is surprising. There is very little publicly available research on what it means to bring AI into the policy process, so policy analysts, researchers and advisors are all making these judgements without an evidence-base to lean on. What doesn't change, no matter the level of AI adoption, is that policy professionals are responsible for the outcome of a policy intervention, and they must be able to explain how it came about, and the decisions they made along the way.

Off-the-shelf AI models aren't helping. They don't disclose their reasoning, where an interpretation comes from or why something has been expressed a particular way. They are frictionless, and great policy work necessitates friction, with assumptions challenged, drafts abandoned, deliberation, argument and more.

FIGURE 3 % of OECD countries using AI in the public sector by policy area, 2023 and 2025



SOURCE Source: OECD (2026), Digital Government Outlook 2026.

2. OECD. (2026). Digital Government Outlook 2026.

FIGURE 4

Perceived barriers to greater AI use by government function

BARRIER	Policy	Communications	Delivery and PM	Digital and data	Operational delivery
AI can make mistakes	80%	69%	68%	73%	71%
I don't want to become overly reliant on AI	68%	59%	67%	51%	58%
My organisation does not help me use AI	27%	25%	17%	20%	27%
AI may mishandle sensitive data	81%	81%	74%	68%	70%
AI is not good at doing certain tasks	76%	63%	65%	63%	53%
I do not have time to learn about AI	33%	34%	45%	22%	30%
I do not know what else I can use AI for	32%	44%	38%	20%	43%
AI overwhelms me and I do not know where to start	16%	19%	26%	17%	23%
AI can harm the environment	63%	72%	62%	55%	51%

SOURCE

Apolitical (2026), AI Usage Survey (n=313). Based on responses to "To what extent do each of the following factors stop you using AI more than you already do?" Figures combine "Strongly deters me from using AI" and "Slightly deters me from using AI" responses. Only government functions with more than 30 responses per category are included.

Building PolicyNova

The decisions and trade-offs behind a purpose-built AI tool for policy research.

The research presented in this briefing note laid the foundation upon which Apolitical built PolicyNova, an AI-powered research tool designed around one core principle: delivering evidence that policy professionals can interrogate and use to develop what they end up proposing to ministers.

That principle is reflected in how PolicyNova approaches evidence. Rather than relying on a broad or generalised evidence base, it draws only on material policymakers can trust, including academic research, legislative text and regulatory documents. Every citation is traceable, so it's possible to see exactly where each piece of information comes from, dig deeper and fact-check it.

PolicyNova does not give recommendations. Instead, it provides evidence and visual tools (such as cross-jurisdictional comparison tables) to reason with and draw conclusions from. In doing so, PolicyNova protects the judgement policymakers have built over years of practice, and keeps them accountable in designing a policy intervention.

It also helps policy professionals see, at a glance, how ready their own jurisdiction is to take on a policy that is already in place elsewhere. This feature compares one jurisdiction with another across a set of dimensions, laying the two over each other so that areas of alignment are easy to spot. Rather than giving a single verdict on whether a policy will transfer, the feature shows the conditions underneath, highlighting where the enabling factors are already in place, where they are missing and where more work would be needed before an intervention could realistically be adopted.

Together, these features keep policymakers in the loop, while simultaneously introducing AI into a process it supports safely, namely research and evidence synthesis.



See what other governments have already tried.

Purpose-built AI for policy research. PolicyNova compares the evidence across jurisdictions and cites the source for each claim.

HOW POLICYNova WORKS

Where do you work, and what are you trying to solve?

I'm a public health official in Los Angeles preparing for the 2028 Olympics. How have other host cities managed the pressure on emergency services?



GEMINI

BEGIN RESEARCH

The future of AI in policy

Beyond research, analysis and synthesis.

Today, AI in policy sits at one end of the process, helping to find evidence, analyse it and pull it together. A handful of initiatives already point to further applications, including using AI to stress-test proposed legislation before it's written and surfacing how a policy might perform under conditions that haven't occurred yet. AI could be used to build tools for exploring futures that haven't yet been imagined, testing not just how a known policy option would perform, but surfacing options that wouldn't have occurred to a policy team working alone.

In the future, AI won't just be used to speed up the work, but to improve what was considered part of it.

Which leaves a question this report cannot answer: Will AI use produce better policy?

It's harder than it sounds. Adoption is measurable. Speed is measurable. Outcomes for citizens, however, are often hard to see, and are shaped by far more than the quality of the analysis behind them.

Nobody yet has a credible way to attribute a better result to a better-informed official. Until someone does, the case for AI in policy rests on inputs rather than results.

What are some exciting potential applications of AI in policy, that go beyond speeding up the work?



Save the date!

PolicyNova launches in **October 2026**, available for all public servants and government officials.

SIGN UP



FOR EARLY ACCESS

