

# RESPONSIBLE AI IN ANZ GOVERNMENT: ETHICAL FRAMEWORKS, GOVERNANCE STRATEGIES, AND CITIZEN TRUST

## AI FOR THE PEOPLE: A ROADMAP FOR SAFE INNOVATION

RESEARCH INNOVATION COUNCIL ANZ BRIEF

### EXECUTIVE SUMMARY

**Artificial Intelligence (AI)** is rapidly transforming global industries, including the public sector in Australia and New Zealand, where it promises to enhance operational efficiency, improve service delivery, and support complex decision-making. However, as AI adoption accelerates, government leaders face significant challenges in navigating the technology's governance, ethical implications, and trust-building requirements. The need for clear regulatory frameworks, workforce education, and responsible AI deployment has emerged as a critical priority across all levels of government.

This whitepaper synthesises insights from seven roundtable discussions hosted by academic partners of the Research Innovation Council—including **Griffith University, RMIT University, QUT Executive Education Centre, Torrens University Australia, Victoria University of Wellington, Curtin University, and University of Technology Sydney**—bringing together senior government officials, academic experts, and industry leaders. These discussions explored the key capabilities required to integrate AI responsibly within the public sector.

A unique perspective within this report comes from Objective Corporation, an Australian based software provider specialising in GovTech applications, through their CTO, Anthony Turco, and VP of Strategy, Cassandra Bisset. Between them, they attended all regional roundtables, allowing them to **observe, compare, and contrast the perspectives of interstate government leaders**. This broad exposure provided them with a nuanced understanding of the shared challenges across jurisdictions, including data sovereignty concerns, public trust barriers, and the practical realities of AI deployment at scale.

While AI's transformative potential is undeniable, its success depends on thoughtful integration into government frameworks — not just as a personal assistant or chatbot, but as a core enabler of public service delivery and operational efficiency. Discussions highlighted that beyond governance and risk management, significant gaps remain in workforce training and technical enablement. Building a viable AI program will require not only upskilling staff in data literacy and AI fluency, but also developing cross-functional capabilities in data engineering, model deployment, and system integration. **Transparent decision-making, robust safeguards, and clear accountability structures** must be embedded at every stage to ensure AI serves the public good.

By **Patrick Joy** | Head of Research & Analysis | [Public Sector Network](#)

# FOREWORD

The conversation around AI in government is often dominated by buzzwords, futuristic promises, and reactionary policies. Yet, when we look beyond the hype, the real challenge isn't just about AI itself—it's about **operationalising** AI in a way that delivers tangible, scalable benefits.

Governments recognise AI's transformative potential, from **reducing administrative burden to accelerating high-impact decision-making**. But too often, AI investments focus on the quick fix—deploying tools like copilots and chat assistants—without establishing the necessary **scaffolding** for sustained, responsible use. AI isn't a magic bullet.

One of the clearest examples of AI's practical value is in **highly information-intensive fields**, where the volume of data exceeds human capacity. For instance, clinicians tasked with policy updates must analyse thousands of pages of legislation, policy and compliance documentation, with changes taking years to enact. With the right AI-driven approach, that timeline can be **reduced to a single year**, dramatically accelerating life-saving decisions. These are the types of high-value applications governments should prioritise, yet too often, investments in AI are made without a clear roadmap for implementation and governance.

A major gap in government AI adoption is the **lack of a centralised framework**. Many agencies operate in silos, deploying AI without consistent data governance, lineage tracking, or interoperability standards. Policy guidance frequently outlines **what not to do** but offers little direction on how to set up AI well. Without structured implementation plans, agencies risk deploying AI in a fragmented, inefficient, and untrustworthy manner. There are varying requirements across organisations, so it's key to consider how data will be collected or curated for these initiatives, data preparation, or how organisations will pass audits and reporting obligations when in production.

"Everyone wants the magic bullet, but agencies aren't putting in the scaffolding to do this—yet they're spending millions on personal assistants without a second thought."

Another challenge is the fast-paced evolution of AI capabilities. The best solution today may be obsolete next week. AI strategies must be flexible, allowing governments to adopt the most efficient, cost-effective, and transparent AI solutions as they emerge. This requires a standardised approach to AI deployment—one that supports **data curation, guardrails, and modular AI services**.

The private sector has contributed to AI uncertainty by **overhyping capabilities** and failing to communicate practical implementation strategies. But AI governance doesn't have to be complex. With controlled data models, thoughtful temperature settings, and purpose-built AI agents, risks can be mitigated. Instead of simply adopting AI tools in response to industry trends, governments should focus on building structured AI frameworks that ensure reliability, security, and adaptability.

This whitepaper delves into these critical issues—governance, scalability, responsible adoption, and practical AI deployment. By focusing on structured, operationalised AI approaches rather than hype-driven experimentation, governments can move beyond pilot programs and towards **real, measurable impact** to embed within business practices.

Cassandra Bisset | VP Strategy & Anthony Turco | Chief Technology Officer | [Objective](#)

# INTRODUCTION TO ARTIFICIAL INTELLIGENCE IN GOVERNMENT

The potential of Artificial Intelligence (AI) to transform the public sector is immense, though its disruptive impact is not new. Long before the rise of Generative AI (Gen AI), advanced AI and machine learning (ML) capabilities have been driving innovation and automation, with roots tracing back within the field of computer science. These technologies have been deployed to enhance **citizen services**, redefine internal **government processes**, and make **data-driven decision-making** more efficient. However, the latest wave of democratised large language models (LLMs)—embodied by tools like ChatGPT—has quickly become nearly synonymous with AI in the minds of many. This has overshadowed the broader, more diverse landscape of AI technologies that span from automation systems handling routine tasks to machine learning algorithms that power complex decision-making. Misunderstandings about the nature of AI, particularly regarding its **intelligence and agency**, have profound implications—not just for how we view the future of work, but for society at large.

As AI technologies proliferate, sectors must come to grips with what these advancements mean for their operations. For citizen-driven industries like government, the **democratisation of AI tools** presents both tremendous opportunities to improve services despite significant risks. The expansion of AI-driven smart services offers a chance to enhance citizen experiences, yet it also brings the challenge of managing evolving dangers such as scams and the vulnerability of customers—concerns that are trending far more rapidly than workforces can adapt.

As the Australian government seeks to avoid falling **behind the curve of AI adoption**, there remains palpable hesitation within agencies, fuelled by concerns about ethical governance, transparency, and data integrity at the critical juncture of AI integration into public sector operations.

“Generative AI is having a big impact on workforces and the way we're thinking about the future of work...it's actually going to be much more of a cultural and leadership challenge than it is a technological one.”

## GENERATIVE AI AND PUBLIC PERCEPTION

Many public sector leaders currently view AI through a narrow lens, focusing primarily on generative tools like ChatGPT, which generate responses based on prompts. While these tools offer significant potential to improve user productivity, such as content generation, summarisation, and customer chat, the public sector must expand their scope to support more mission-critical business activities and unlock AI's transformative potential. For example, **AI-driven agents** that autonomously handle citizen enquiries, **intelligent virtual assistants** in service delivery, and **decision-making systems** for policy recommendations are already being tested and proven effective in leading-edge domains.

These capabilities are further enhanced by the vast rise in processing power, as demonstrated by competing advancements releasing in the industry, which enable AI models to learn, process and analyse far larger datasets at faster speeds. The increase in computational power is pushing the boundaries of what AI can achieve, unlocking opportunity to address more complex problems, automate more tasks, and generate more precise, context-aware outputs.

## A PRIMER: GOVERNMENT'S RESPONSIBLE USE OF AI IN ANZ



**Policy Aim:** The responsible adoption of AI is a priority for both Australia (AU) and New Zealand (NZ), ensuring that AI benefits society and the economy while aligning with community expectations. Both policies emphasise the principles of safety, ethics, and transparency in AI governance, though the NZ Public Service AI Framework places an additional focus on lawful AI use that aligns with public service values.




**Embrace the Benefits:** AU's Policy for the Responsible Use of AI in Government establishes a governance baseline across the Australian Public Service (APS) to foster responsible AI adoption and improve public service efficiency. Similarly, NZ's framework outlines an AI Work Programme under the Government Chief Digital Officer (GCDO) to support agencies in their AI initiatives. Both approaches seek to remove adoption barriers and ensure AI enhances government decision-making.




**Strengthen Public Trust:** Transparency and accountability are key components in both national policies. AU mandates public disclosure of AI use, while NZ emphasises the principle of explainability and risk mitigation. Agencies in both countries must conduct AI risk assessments to enhance accountability and maintain public confidence.



**Adapt Over Time:** Given the rapid evolution of AI, both frameworks adopt an iterative, adaptive approach. AU's policy stresses the importance of ongoing evaluation and revision of AI practices, while NZ's framework introduces six key pillars to support AI governance, including ongoing risk monitoring and policy refinement.



**Implementation and Application:** AU's policy applies to all Non-corporate Commonwealth entities from September 2024, complementing existing governance frameworks on data management and cybersecurity. NZ's framework provides agency-specific AI guidance under the GCDO's oversight, reinforcing AI's alignment with national digital transformation efforts.



**Principles and Requirements:** AU's 'Enable, Engage, Evolve' model focuses on responsible AI deployment through training, accountability, and risk assessments. In contrast, NZ's five key principles prioritise lawfulness, ethical AI design, explainability, human oversight, and continuous evaluation, offering a more structured ethical approach.



**Risk Assessment:** Both policies mandate a structured AI risk assessment framework. While AU promotes risk matrices to evaluate fairness, privacy, and security, NZ requires agencies to assess AI systems against public service values and ensure alignment with legal and ethical obligations.

## PROBLEM STATEMENT

The Australian public sector is facing a series of challenges as it looks to adopt AI technologies. These challenges include a lack of understanding and readiness among leadership, ethical and regulatory concerns, data quality issues, and the growing risks to public trust in AI systems.

### Leadership and Workforce Readiness

While some government agencies are embracing AI, many are still in the early stages of adoption, largely due to a lack of AI literacy at senior levels. Leaders are often unsure of how to implement AI responsibly and effectively, which has led to a fragmented approach across agencies. In particular, public sector leaders and staff often fail to scale AI's potential to automate routine tasks or to enhance decision-making through predictive analytics.

One particularly insightful point raised in the roundtable discussions was the realisation that the public sector is often behind the private sector in terms of AI integration, not because of a lack of resources, but because of organisational inertia and reluctance to embrace technological change.

As one participant pointed out,

“AI tools are not intelligent...from a social sciences perspective, we start to pull that apart because the minute we tell the public this is how this tool works, they start to believe that the tool can do what people can do.”

This challenge is further compounded by the lack of consistent governance and policy frameworks across agencies.

The Australian Policy for the Responsible Use of AI in Government addresses this by requiring **accountable officials** within agencies who are responsible for implementing the policy's provisions and ensuring that AI systems are deployed in alignment with ethical standards and transparency. This ensures that leadership readiness and responsibility are central to AI adoption across the public sector.

### Data Quality and Management

Many government agencies struggle with poor-quality data, which affects the effectiveness of AI systems. AI models are only as good as the data they are trained on, and public sector data is often messy, incomplete, or inconsistent. This lack of reliable data has led to fears that AI systems could perpetuate existing biases or produce inaccurate results, undermining public trust.

“Data is messy, and AI will heighten all of the existing problems you have behind the scenes.”

This concern was particularly acute in public-facing decision-making roles. If AI tools are used without proper data governance, they could make decisions that are discriminatory or incorrect, leading to significant harm. Therefore, addressing data management challenges is crucial for the success of AI deployment in the public sector.

Conversely, data science experts raised that AI can also offer solutions by making sense of messy datasets. For instance, advanced models can identify patterns, clean data, and reduce systemic bias, transforming disorganised inputs into actionable insights.

The AI risk assessment framework outlined in the Australian Government Responsible AI Policy (Page 17) encourages agencies to proactively assess and mitigate the risks associated with data quality, including issues such as bias and privacy concerns. This aligns with this paper's emphasis on ensuring that data management practices are robust and designed to eliminate inaccuracies that could harm the effectiveness or fairness of AI systems.

## STRATEGIES & SOLUTIONS

### Building AI Governance Frameworks

Establishing clear and robust AI governance frameworks is essential for ensuring that AI is used ethically and responsibly in the public sector. Governance frameworks should include policies for transparency, accountability, and fairness, ensuring that AI tools are deployed in a way that benefits citizens without compromising public trust.

“The challenge is, how do we reclaim our own agency over these tools? How do we ensure that they're doing what we would like them to do?”

Agencies should begin by defining the role of AI within their organisational structure. This includes determining which processes can be automated, which require human oversight, and where AI can assist in decision-making.

Effective governance also requires setting clear boundaries around the ethical use of AI, particularly with regard to sensitive data, decision-making processes, and citizen interactions.

### Data Management and Customisation

Effective AI implementation hinges on the quality and integrity of the data used to train AI models. Agencies should prioritise data management and ensure that data is clean, structured, and free from bias. This may involve revising data collection processes, standardising data formats, and investing in data cleaning tools to ensure that the AI systems can work with reliable and accurate information.

“We've recently announced that we will be moving the Chief Data Officer role from IT to the business...because data is our business. It's not an IT function.”

Customising AI systems to align with the unique data sets of each agency will enhance their performance and ensure they deliver actionable insights. By focusing on data governance and grounding AI models in accurate, unbiased data, public sector leaders can mitigate risks associated with AI deployment and improve the effectiveness of their AI-driven applications.

The Policy for the Responsible Use of AI in Government mandates that agencies adopt a clear, transparent approach to their AI use, including AI transparency statements that outline how AI is being implemented and monitored within each agency.

The policy's focus on data governance (Page 11) complements this approach by urging agencies to integrate AI considerations into existing frameworks such as privacy, cyber security, and data governance. This ensures that the data integrity concerns highlighted in this paper are addressed with a structured, policy-driven approach.

The AI literacy and training requirements outlined in the policy (Page 11) reinforce a call for widespread education. Agencies are encouraged to implement AI training for all staff and tailor it to specific roles, which aligns directly with this paper's call for workforce readiness and AI literacy across all levels.

## Education & Training

The research roundtables emphasised that building AI literacy across the public sector is critical to successful adoption. Training should not be limited to technical teams but should extend to all levels of government employees, including policymakers, managers, and frontline staff. This training should focus on understanding AI's potential and limitations, ethical considerations, and how to use AI tools responsibly.

“We need to build public confidence...we still had a human do a long read on 300 submissions...so although we used AI, we still had a long read i.e. a human long read every submission.”

Training should also include hands-on experience with AI tools and real-world applications, allowing staff to understand how these technologies function and how to integrate them into their workflows effectively. The roundtable discussions also highlighted the importance of teaching critical thinking skills so that public servants can assess AI outputs with a discerning eye.



## Safe AI Experimentation

To mitigate the risks associated with AI adoption, public sector agencies should focus on pilot programs and controlled experimentation. By starting small and testing AI in non-public-facing, low-risk environments, agencies can gain valuable insights into the technology's capabilities and limitations.

**“In many cases, there are steps before you actually get to launch the data sets...the challenge is, if people trust the system, and they're not critical about it, it will simply be bad data in, bad response out.”**

A phased approach allows agencies to monitor AI's performance and make adjustments before scaling up. This will help build confidence among both public servants and citizens, ensuring that AI tools are used responsibly and effectively.

This recommendation is further supported by the policy's emphasis on adaptability (Page 7), which encourages agencies to adopt an evolving approach to AI governance, ensuring that AI systems are tested, monitored, and refined over time as technologies and practices evolve.

Given one of the most significant challenges identified in the discussions is maintaining public trust, Safe AI is paramount to the highly sensitive environment within which agencies operate, where even small errors in decision-making can have wide-reaching consequences.

The roundtable participants underscored the importance of ensuring that AI systems are explainable and auditable, and that human oversight is maintained in decisions that impact citizens directly.

# AI IN THE PUBLIC SECTOR

## INTERNAL USE CASE IN RECORDS MANAGEMENT

One of the most successful examples of generative AI adoption in the public sector is the use of AI for records classification. Several agencies have successfully implemented AI tools to automate the process of categorising records, reducing manual workloads and improving accuracy. This internal use of AI allows agencies to demonstrate its value without risking citizen-facing decisions.

**“We are starting where we don’t have people using it to do the work of the workforce...using AI for the things that we don’t already have people doing”**

This example highlights the potential for AI to improve internal efficiency and operational effectiveness, especially in areas where human labour is constrained or where there is a need for large-scale data processing.

## AI IN DECISION-MAKING

While AI adoption in internal processes is progressing, its use in public-facing decisions is still met with caution. Agencies are wary of allowing AI to make decisions that directly affect citizens, particularly in areas such as welfare, law enforcement, and public health. The roundtable discussions emphasised the need for human oversight and transparency in any AI-driven decision-making process.

**“In terms of the health context, AI can be a numbers game...but when AI starts interacting with citizens...if it goes wrong, you lose trust so quickly.”**

This caution is not unfounded, as errors in AI decision-making, especially in public services, can have significant consequences. Ensuring that AI decisions are explainable and traceable is crucial for maintaining public trust.

## CONCLUSION

As AI becomes an integral part of the Australian public sector, the need for responsible and ethical governance is paramount. The Australian government must approach AI adoption carefully, ensuring that it is used to enhance service delivery, improve efficiency, and make better decisions, while safeguarding against potential risks.

By focusing on AI governance, data management, training, and safe experimentation, the public sector can unlock AI's full potential while maintaining public trust and confidence. With careful planning and a commitment to responsible practices, AI can transform the way the Australian government serves its citizens, driving innovation and improving the quality of services provided to the public.

# ACTIONABLE STEPS: EXECUTION & IMPACT

## 1. DEVELOP CLEAR AI GOVERNANCE FRAMEWORKS

- **Action:** Establish AI governance frameworks that align with AU's 'Enable, Engage, Evolve' model and NZ's AI guiding principles (lawfulness, ethical use, explainability, human oversight, and continuous evaluation). Assign clear accountability for AI implementation at the senior leadership level in both jurisdictions.
- **Impact:** Enhances public trust by ensuring AI deployment is ethical, transparent, and accountable. Encourages interoperability between ANZ governance structures to support regional alignment on responsible AI.

## 2. INVEST IN WORKFORCE EDUCATION & TRAINING

- **Action:** Roll out AI literacy programs for all levels of government employees, including policymakers, technical teams, and frontline staff. Training should cover foundational AI skills and specialised role-based competencies, including ethical AI, compliance, and risk.
- **Impact:** Builds an AI-aware workforce capable of integrating AI into daily operations responsibly. Ensures consistency in AI literacy across ANZ, reducing misapplication risks and boosting adoption confidence.

## 3. IMPLEMENT PHASED, CONTROLLED AI EXPERIMENTATION

- **Action:** Establish AI sandbox environments and controlled pilot projects to address complex, high-value information challenges (e.g., records management, workflow automation). AI prototypes can be particularly valuable for processing vast datasets and large volumes of complex documentation—such as maintenance manuals, overlapping legislation, and extensive policies—to extract actionable insights. NZ's AI Work Programme and AU's phased AI adoption principles guide these approaches.
- **Impact:** Mitigates risks while enabling agencies to test AI solutions in real-world, high-value scenarios, improving confidence in AI reliability. Engaging with industry and vendors is key to leveraging best practices, accessing proven methodologies, and ensuring alignment with emerging AI capabilities. By focusing on critical information challenges, organizations can unlock transformative efficiencies while integrating lessons from the broader AI ecosystem and aligning with cross-government AI safety priorities.

## 4. ESTABLISH DATA GOVERNANCE PROTOCOLS

- **Action:** Standardise AI data management practices to ensure data accuracy, completeness, security, and privacy. Agencies must conduct Privacy Impact Assessments (PIAs), integrate privacy-by-design principles, and comply with ANZ data protection laws. AI systems should be continuously audited to prevent data leaks, unauthorised access, and cyber risks. Additionally, agencies can consider private cloud or private GPU infrastructure as an alternative to external SaaS providers, enabling a fully closed-loop AI deployment.
- **Impact:** Minimises data exposure, reduces costs, and provides sensitive departments with greater control over security and compliance. Ultimately, citizen data is protected while maintaining public trust and AI accountability. Compliance with AU's Privacy Act 1988 and NZ's Privacy Act 2020 is also ensured, strengthening AI resilience against data security threats.

## 5. LAUNCH AI TRANSPARENCY STATEMENTS

- **Action:** Within six months of AU's policy implementation, each agency should issue a public AI transparency statement. These statements should outline AI use cases, risk management strategies, safeguards against bias, and auditability measures.
- **Impact:** Fosters public confidence by making AI decision-making processes transparent, explainable, and accountable. Aligns with NZ's emphasis on AI explainability and AU's commitment to public disclosure.

## 6. PROMOTE CROSS-AGENCY COLLABORATION ON AI

- **Action:** Establish ANZ AI governance working groups and/or a cooperative research centre to foster regular dialogue between ANZ public sector agencies. These groups should share best practices, lessons learned, and case studies to create a cohesive regional approach to AI adoption.
- **Impact:** Encourages a unified approach to AI adoption across government, reducing fragmentation and enabling shared learning. Helps agencies avoid duplication of effort and leverage successful initiatives from other sectors.

## 7. ENGAGE IN ONGOING MONITORING AND RISK ASSESSMENT

- **Action:** Implement real-time monitoring and continuous evaluation of AI systems, using adaptive risk frameworks that align with both NZ's structured AI risk assessment requirements and AU's risk-based AI approach.
- **Impact:** Ensures AI remains safe, effective, and aligned with evolving public needs, reinforcing public trust in AI-driven services.

## 8. FOCUS ON HUMAN-AI COLLABORATION IN DECISION-MAKING

- **Action:** Establish strict human oversight policies in AI-driven decision-making, particularly in sensitive areas such as public health, social services, welfare, and law enforcement. AI should remain assistive rather than autonomous, ensuring human judgment prevails in critical decisions.
- **Impact:** Protects citizens from harmful or opaque AI decisions, reinforcing NZ's commitment to human oversight and AU's focus on AI accountability.

## 9. INTEGRATE AI ETHICS INTO POLICY DEVELOPMENT

- **Action:** Mandate bias detection and fairness assessments in AI applications to ensure equitable outcomes across all demographic groups. Agencies should implement regular audits of AI models to detect and mitigate discriminatory biases and ensure alignment with ANZ anti-discrimination laws.
- **Impact:** Reduces ethical risks and improves AI accountability, encouraging fair, transparent, and non-discriminatory AI systems to ensure AI benefits all citizens equitably.

## 10. DEVELOP CUSTOMISED AI TOOLS FOR SPECIFIC AGENCY NEEDS

- **Action:** Design AI tools customised for each agency's needs while ensuring interoperability across public sector AI applications, effectively managing growing risks of vendor lock-in.
- **Impact:** Improves AI efficiency, adaptability, and effectiveness, enabling agencies to tailor AI solutions while ensuring compliance with both national governance models.


# ABOUT THE RESEARCH INNOVATION COUNCIL

**Connecting government, industry and academia:** we've been able to engineer a new program antithetical to the classical red tape, administration, and risk-aversion that impedes innovation.

Our program offers senior executives a unique opportunity to gather monthly for PSN-facilitated roundtable sessions with colleagues, academics, and technologists. We run the program at no cost to government executives or academics, and provide attendees a private online group on our global insights exchange platform, access to research briefs, and expert engagement with PSN analysts and facilitators.



## Greatest challenge partnering with Academia or Industry



Source: PSN Research Innovation Council Survey 2024. Total Sample Size: 188 ANZ/ 83 USA/ 80 CA Gov Executives

## ABOUT PUBLIC SECTOR NETWORK

Public Sector Network is a research company that represents public sector professionals across Australia, Canada, New Zealand, and the USA. It develops roundtables, seminars, and conferences to suit current areas of interest to government agencies and their suppliers.

PSN's growing community spans across federal, state, and local government departments, healthcare, and education, allowing members to share information, access the latest in government innovation, and engage with other like-minded individuals on a secure and closed-door network.

### AUSTRALIA / NEW ZEALAND

P +61 2 9057 9070

E [INFO@PUBLICSECTORNETWORK.COM.AU](mailto:INFO@PUBLICSECTORNETWORK.COM.AU)

### USA

P +1 (647) 969 4509

E [HELLO@PUBLICSECTORNETWORK.COM](mailto:HELLO@PUBLICSECTORNETWORK.COM)

### CANADA

P +1 (647) 459 8904

E [CONTACT@PUBLICSECTORNETWORK.CO](mailto:CONTACT@PUBLICSECTORNETWORK.CO)

### Public Sector Network (Australia) Pty Ltd

ABN - 46 617 870 872 20-40

Meagher Street, Chippendale, Sydney NSW 2008, Australia