



Australian Government
Digital Transformation Agency



Joint Committee of Public Accounts and Audit

Inquiry into the use and governance of artificial intelligence systems by public sector entities

Digital Transformation Agency Submission
October 2024

Contents

Contents	2
Introduction	3
Artificial intelligence across the public service	4
Rapid development of generative and general-purpose AI	4
Impact on the Australian Public Service workforce	5
Government as an exemplar	8
Interim guidance for public generative AI tools.....	8
AI in Government Taskforce	9
National framework for the assurance of AI in government	10
Policy for the responsible use of AI in government	10
<i>AI in government fundamentals training</i>	12
Draft Australian Government AI assurance framework	12
AI technical standards	13
Sovereign capability	14
Trial of Microsoft 365 Copilot	14
DTA Marketplaces and Panels	16
Complementary initiatives	17

Introduction

The Digital Transformation Agency (DTA) welcomes the opportunity to make a submission to the Joint Committee of Public Accounts and Audit (JCPAA) inquiry into the use and governance of artificial intelligence (AI) systems by public sector entities (the Inquiry).

To assist the Committee's inquiry, this submission outlines how the DTA is supporting the safe and responsible use of AI by the Australian Government.

The DTA provides strategic and policy leadership on whole-of-government and shared investments into information and communications technology (ICT) and digital service delivery. It seeks to position the Australian Government as an exemplar in transparency, risk management and governance of not just AI, but all data and digital technologies.

The DTA is focused on driving and securing the government's investment in Australia's digital future and supporting government to respond to the challenges and opportunities presented by emerging technology, including AI. To support these objectives, the DTA:

- leads government's digital transformation strategy through co-delivery of the [Data and Digital Government Strategy](#) and [Implementation Plan](#)
- provides leadership in policy design, strategy, advice and whole-of-government architecture
- oversees and measures the effectiveness of the short, medium and long-term whole-of-government digital and ICT investment portfolio through the Australian Government Digital and ICT [Investment Oversight Framework \(IOF\)](#)
- manages whole-of-government digital and ICT strategic sourcing and contracts.

As part of its role to advise government on emerging technologies, the DTA recently released the [Policy for the responsible use of AI in government](#), which came into effect on 1 September 2024. The policy aims to ensure that government plays a leading role in embracing AI for the benefit of Australians while ensuring its safe, ethical and responsible use, in line with community expectations.

Artificial intelligence across the public service

This section addresses the following terms of reference:

- whether the public sector has the internal capability to effectively adopt and utilise AI into the future
- the purposes for which AI is currently being used by the public sector entity and whether there are planned or likely future uses.

While the use of AI in government is not new, it has traditionally been limited to 'narrow' applications performing specific tasks within defined domains. This includes AI-enabled predictive analytics to identify patterns and relationships in large data sets. The technical expertise and cost involved in developing, deploying, and operating these AI systems created a natural barrier to adoption for many agencies.

These narrow AI applications have been subject to existing general digital and ICT governance mechanisms, which are based on a technology-agnostic approach. This includes mechanisms embedded within internal agency processes and Whole-of-Government mechanisms, such as the IOF. AI is also subject to existing legislation, regulation, and policy related to privacy, online safety, intellectual property, anti-discrimination, and cyber security. These obligations apply regardless of the technology used.

Recent developments in AI technology present new challenges and highlight limitations in the current state. This has led to the development of more targeted measures to support safe and responsible AI adoption.

Rapid development of generative and general-purpose AI

Generative AI is a branch of AI that develops generative models that can learn to generate content such as images, text, and other media with similar properties as their training data. General-purpose AI describes systems that address a broad range of tasks and uses, both intended and unintended by developers.¹ Both share many characteristics with other technologies but stand out due to their speed, scale, adaptability and versatile deployment.

Generative AI has made AI widely accessible and easy to use, exemplified by large language models like ChatGPT. This accessibility has magnified AI's potential value to government, offering opportunities for administrative efficiencies, enhanced productivity, and better public

¹ Department of Industry, Science and Resources (DISR), [Voluntary AI Safety Standard](#), Australian Government 2024

OFFICIAL

outcomes. With global AI investment and development accelerating at unprecedented rates, it is crucial for the government to prepare for this transformation.

Despite the significant benefits of AI adoption, there is a clear need to ensure its safe and responsible use to meet community expectations. The long-term insights briefing on AI produced by the Department of the Prime Minister and Cabinet (PM&C) specifically highlighted the potential risks of using AI in the public service.² These include risks around bias, privacy, unintended consequences, dependency, inaccuracies, opaque processes, and job displacement.

The rapid evolution of AI also requires continuous monitoring and refinement of AI models to prevent inaccuracies and ensure their relevance over time. This is to ensure models reliably operate in accordance with intended purposes and enable the identification of unintended impacts upon people, society or the environment. This may include biases which could exclude or harm vulnerable individuals, groups and communities. In unlocking AI's potential, governments should adopt a measured and pragmatic approach to potential AI risks and harms.

Impact on the Australian Public Service workforce

The adoption of AI within the Australian Public Service (APS) has the potential to enhance productivity and efficiency. AI technologies can streamline service delivery, administrative processes, reduce manual workloads, and improve the accuracy of tasks such as fraud detection and case management.

Based on information collected by the DTA and the AI in Government Taskforce in late 2023 there are four broad categories AI in use within Australian Government agencies:

- **Category 1:** Predictive analytics; identifying patterns in data to then make predictions – often used for large business tasks or statistics
- **Category 2:** Natural language processing, which powers services like Google Translate
- **Category 3:** Computer vision and audio recognition, which lets your phone recognise the people in your photos, or use dictation to type
- **Category 4:** Generative AI, such as the capabilities of tools like Microsoft Copilot, Google Gemini and ChatGPT.

² Department of the Prime Minister and Cabinet, [Long Term Insights Briefing - How might artificial intelligence affect the trustworthiness of public service delivery](#), Australian Government 2023

OFFICIAL

In service delivery, AI-enabled tools can help identify patterns in, and triage, contact centre calls and provide staff with relevant information quickly, improving service delivery times and the accuracy of information provided to the public. Public facing AI-enabled digital assistants can also reduce demands on contact centres by making it easier and faster for the public to find the right information for themselves.

The use of generative AI tools in an APS office environment can also assist staff draft and edit documents and emails, generate visually engaging presentations, summarise long and complex documents, conduct preliminary research, brainstorm ideas and develop bespoke spreadsheets for data analysis. By simplifying and speeding up these tasks staff will have more time to focus on more complex and important work.

To maximise these benefits, while mitigating potential risks, it is crucial to ensure that APS staff are adequately trained in the responsible use of AI. The APS must therefore invest in training programs to equip staff with the necessary skills to use AI technologies and apply human judgment where it makes sense to.

Our [evaluation of a trial into generative AI solutions](#) found that 75% of surveyed users felt more confident about their use of a generative AI assistant when they received 3 or more forms of training. This was 28% higher than those who only received one form of training.

The [Policy for the responsible use of AI in government](#) strongly recommends that agencies implement AI fundamentals training for all staff, regardless of their role. To support agencies in implementing this recommendation, the DTA has released [Guidance for staff training on AI](#) and [AI in government fundamentals](#) training modules.

Jobs requiring specialist AI skills are growing significantly. While AI digital assistant tools will be able to help APS staff by filling other skills gaps, such as writing code or providing specialist content, the skills required to safely build, implement and support these specialised AI tools are in short supply. While there are existing data science skillsets in specialised areas of the APS, the current number of APS employees with both data science and AI skills, or even specialty AI skills alone, appears to be low. To realise the opportunities of AI and meet demand, the APS will need to grow its AI skills supply.

Relying solely on university graduates or entry-level talent is not going to meet these rapidly changing workforce demands. Instead, targeted upskilling initiatives for mid-career professionals will be essential in developing a pipeline of senior talent with AI knowledge. There is also a significant opportunity to develop retraining programs that enable diverse tech workers to re-enter the APS workforce equipped with new AI skills.

OFFICIAL

It is also important that current APS leadership possess the knowledge and situational awareness to enable the safe and responsible use of AI. Senior executives should consider community expectations, risks, opportunities and value of adopting AI technologies in the contexts of their agencies and portfolios.

Between August and October 2024, the DTA hosted 4 webinars for senior leaders on how AI might impact government workforces, accountability, effectiveness and the way it interacts with an AI-enabled society. Each webinar included leading subject matter experts from within and outside the service, presenting challenging perspectives. The series was designed to broaden the high-level, strategic understanding for senior leaders and encourage continuing conversation within their agencies.

By addressing the challenges associated with newer forms of AI and workforce impacts, the APS can harness the full potential of AI while safeguarding the interests and wellbeing of its workforce.

Government as an exemplar

This section addresses the following terms of reference:

- the existing legislative, regulatory and policy frameworks that are relevant to the use of AI and whether they are fit for purpose
- whether the internal governance structures that currently exist for AI will ensure its ethical and responsible use by public sector entities
- the internal framework/policies or additional controls used for assessing the risks associated with the use and possible misuse of AI, including the areas of security, privacy, ethics, bias, discrimination, transparency and accountability
- whether there is an adequate line of sight to the output of AI, and the decisions made through its use

The Australian Government is committed to embracing AI while acting as an exemplar in transparency, risk management and governance in AI adoption. This commitment to exemplary AI adoption aligns with the government's response to its consultation on safe and responsible AI in Australia.³

The Australian Government has a unique position in the community due to the datasets it holds, the breadth of its interactions with individuals and businesses, and its ability to deliver positive societal impacts through embracing technology. Its roles and responsibilities, such as creating laws, implementing policies, delivering services and funding grants, can amplify both the positive and the negative impacts of using AI. For these reasons it is critical government takes a leading role in the use of AI aligned to community expectations, setting a standard that other areas of the economy can follow.

As the AI landscape rapidly evolves, government must also be flexible and able to quickly adapt. AI policy settings will need to respond to change and embed best practices as they emerge to further reduce the risk of negative consequences while maximising the opportunities.

Interim guidance for public generative AI tools

Responding to the immediate need for guidance to safely using publicly available generative AI tools, the DTA and Department of Industry, Science and Resources (DISR) released the [Interim guidance on government use of public generative AI tools](#) in July 2023.

³ DISR, [Safe and responsible AI in Australia Consultation Australian Government's interim response](#), Australian Government, 2024.

OFFICIAL

The guidance outlines principles and considerations for APS agencies and staff when using these tools. The principles focus on accountability, transparency and explainability, privacy protection and security, fairness and human-centred values, and human, societal, and environmental wellbeing.

These principles in practice expand on Australia's 8 AI Ethics Principles, a voluntary framework to reduce the risk of negative AI impacts by maintaining safe, reliable and fair outcomes for Australians.⁴

AI in Government Taskforce

In August 2023, the Australian Government established the AI in Government Taskforce (the Taskforce) in response to continued interest and rapid developments in AI. Jointly led by the DTA and DISR, the Taskforce was comprised of a team of multidisciplinary secondees from 11 agencies across the APS.

The Taskforce, which concluded on 30 June 2024, focused on supporting government to harness the opportunities of AI technologies in a safe and responsible way. The Taskforce experienced high and consistent demand for guidance and support from across government. Its AI Community of Practice was attended by over 300 participants across 82 agencies wanting to share best practice in AI adoption. It also established an advisory group comprised of over 70 experts across 18 agencies to support the development of consistent approaches for responsible AI use.

The Taskforce's achievements include:

- revising the guidance on the use of publicly available generative AI tools
- establishing an extensive network to coordinate AI work
- supporting the development of the [National framework for the assurance of artificial intelligence in government](#) with state and territory governments under the auspices of the Data and Digital Ministers Meeting (DDMM) process.
- supporting the rollout and evaluation of the Australian Government's 6-month trial of a Microsoft 365 Copilot.

The Taskforce made important progress in supporting the government's safe and responsible AI objectives. The DTA continues to build on the work of the Taskforce while also maintaining and growing the extensive network of professionals across government.

⁴ DISR, [Australia's AI Ethics Principles](#), Australian Government 2022

National framework for the assurance of AI in government

On 21 June 2024, the DDMM endorsed and released the [National framework for the assurance of artificial intelligence in government](#). The national framework sets foundations for a nationally consistent approach to the assurance of AI use across the Australian Government and all state and territory governments, informed by Australia's AI Ethics Principles.

The national framework recommends practices and cornerstones of AI assurance, an essential part of broader AI governance in the public sector, and provides case studies as follows:

- **Practices.** Demonstrates how governments can practically apply Australia's AI Ethics Principles. From 'maintain reliable data and information assets' to 'comply with anti-discrimination obligations', each practice is aligned to one of the principles.
- **Cornerstones.** Outlines 5 mechanisms (governance, data governance, standards, procurement, and a risk-based approach) that governments often already use that have been identified as valuable enablers for practicing the assurance of AI.
- **Case studies and resources.** Includes several examples of different governments' previous experience, such as the intersection of AI and recordkeeping with transparency and explainability.

The national framework reflects the Australian and state and territory governments' commitment to a lawful, ethical approach that places the rights, wellbeing and interests of people first. Jurisdictions are to develop specific policies and guidance aligned with their own legislative, policy and operational contexts to implement that national framework.

Policy for the responsible use of AI in government

In effect from 1 September 2024, the [Policy for the responsible use of AI in government](#) applies to all non-corporate Commonwealth entities, with some exceptions for Defence and national intelligence community agencies.

The policy aims to create a coordinated approach to the use of AI in government, complementing and strengthening existing frameworks within the APS. It is underpinned by the 'enable, engage, evolve' framework which outlines policy principles and requirements that aim to ensure the government's use of AI is forward leaning, engages with risk and recognises the need for ongoing evolution.

OFFICIAL

As part of this ongoing evolution and in considering next steps, the DTA will examine the outcomes of the pilot of the [draft Australian Government AI assurance framework](#) and developments in the broader regulatory landscape, including the whole-of-economy safe and responsible AI agenda, led by DISR.

The policy principles under each phase of the framework set clear expectations about how AI should be used by government. This includes that APS officers need to be able to explain, justify, and take ownership of advice and decisions when utilising AI. The principles also state that agencies should have clear accountabilities for the adoption of AI and understand its use.

The policy outlines mandatory actions for agencies that build on the principles, including identifying an accountable official responsible for implementing the policy and publishing a transparency statement regarding AI use.

To support the implementation of these mandatory actions the DTA has published a [Standard for accountable officials](#) and a [Standard for AI transparency statements](#).

Accountable officials, who must be appointed by 30 November 2024, have 5 key responsibilities. They must:

- be accountable for implementation of the policy within their agencies
- notify the DTA where their agency has identified a new high-risk use case
- be a contact point for whole-of-government AI coordination
- engage in whole-of-government AI forums and processes
- keep up to date with changing requirements as they evolve over time.

Transparency statements, which must be published on agency public facing websites by 28 February 2025, will provide a high-level overview of AI use and efforts to protect the public from any negative impacts. These statements must include:

- why the agency uses or is considering using AI
- AI uses by classification, including where the public may directly interact with, or be significantly impacted by, AI without a human intermediary or intervention
- measures to monitor the effectiveness of deployed AI systems
- compliance with applicable legislation and regulation
- efforts to identify and protect the public against negative impacts

The policy also strongly recommends agencies implement AI fundamentals training for all staff, with additional training as required for specific roles.

OFFICIAL

The policy positions the Australian Government as an exemplar under its broader safe and responsible AI agenda, ensuring that AI use meets community expectations and mitigates potential risks.

AI in government fundamentals training

As noted above, the policy strongly recommends agencies implement AI fundamentals training for all staff, with additional training as required for specific roles. To support this recommendation, the DTA has released [Guidance for staff training on AI](#) and [AI in government fundamentals](#) training modules which all agencies can access via APSLearn. The training is designed to provide all staff, regardless of their experience, with foundational knowledge of AI, the principles of safe and responsible use and how to apply these in the APS.

The course introduces key AI concepts, explains responsible AI use and equips staff to be:

- informed on the applications and risks of AI
- capable of applying relevant advice
- confident about when it is suitable to use generative AI
- empowered to apply their own judgement with AI.

The training leverages existing resources and frameworks that APS staff need to apply, including their own agency's policies, the APS Values and APS Code of Conduct, Australia's AI Ethics Principles and the interim guidance on government use of public generative AI tools.

Under the policy, agencies are encouraged to consider additional training for staff in consideration of their roles and responsibilities, such as those responsible for the procurement, development, training and deployment of AI systems.

Draft Australian Government AI assurance framework

In September 2024, the DTA began a pilot of a [draft AI assurance framework](#) that explores mechanisms for Australian Government agencies to implement the practices within the national framework. The draft framework, published on 21 October 2024, is designed to address the specific legislative, policy and operational context of Australian Government agencies.

The framework guides agencies through a practical impact assessment of their AI use cases against Australia's AI Ethics Principles. The impact assessment process helps agencies identify and mitigate AI risks at each stage of the AI lifecycle, and document the steps they

OFFICIAL

are taking to ensure their AI use is responsible. This approach to AI assurance positions the Australian Government to ensure human oversight of AI, proportionate to risks, and to prioritise the rights, wellbeing and interests of people.

By sharing their feedback on the assurance framework impact assessment process, participating agencies are playing a key role in shaping the government's approach to AI assurance. Evidence gathered through the pilot will inform further updates to the framework and the DTA's recommendations to government on future Australian Government AI assurance settings, as part of next steps for the [Policy for the responsible use of AI in government](#).

AI technical standards

To ensure government can confidently experiment with and develop AI use cases, the DTA is leading the development of technical standards for government's use of AI. These standards will deliver technical guidance to ensure that responsible use of AI is embedded at a technical level, supporting the policy and Australian Government AI assurance framework.

These standards will also enable greater consistency, interoperability and reuse across the APS and contribute to the expanding role the DTA is playing in positioning government as an exemplar in AI use.

Technical standards for AI in government are key to ensuring government has the capability to responsibly use and innovate with AI. These standards will cover the AI use case lifecycle, ensuring that responsible use is built into every stage from design and development through to deployment and retirement.

To deliver these standards, the DTA has convened a multidisciplinary team of secondees from across the APS, with representatives from the Australian Taxation Office, the Department of Employment and Workplace Relations, the Department of Infrastructure, Transport, Regional Development, Communication and the Arts and Services Australia.

The team is developing these standards in collaboration with a working group of practitioners and experts from across government. The working group will be invaluable in testing the coverage of the standards to their agencies' contexts.

Sovereign capability

This section addresses the following terms of reference:

- whether there are sovereign capability issues to consider given that most AI tools currently used in Australia are sourced from overseas

The DTA has previously submitted a response to the [Inquiry into supporting the development of sovereign capability in the Australian tech sector](#), which outlines how the DTA's actions support the development of sovereign capability in the Australian tech sector from policy direction to implementation.

While Australia currently lacks a local large language model, the DTA-managed marketplaces offer products and services that could enhance international models. These marketplaces are regularly refreshed, allowing suppliers to re-apply, which could introduce new and innovative technologies and services in the future.

Trial of Microsoft 365 Copilot

The DTA led the Australian Government's 6-month trial of Microsoft 365 Copilot (Copilot), making it one of the first governments in the world to deploy generative AI tools across government. Almost 7,700 staff across over 60 agencies participated in the trial.

Copilot integrates generative AI capabilities into applications like Microsoft Word, Excel, Outlook, and Teams. It helps process and generate information for users during their work activities but does not make decisions or finalise outputs without user involvement. It leverages access to organisational data for context and accuracy, based on permissions set by each agency for their own environments. Users interact with Copilot through chat-based interfaces and dynamic, automated suggestions within M365 applications.

The DTA administered the trial and provided whole-of-government safeguards and minimum training requirements for trial participation. Participating agencies were responsible for undertaking security risk and other assessments for their circumstances and requirements.

The DTA's evaluation of the trial included insights on the specific use of Copilot and broader product-agnostic learnings that could inform government implementation of other generative AI and AI assistant tools. The [evaluation report](#) was published on 23 October 2024, followed by a public briefing on the findings on 25 October 2024. The DTA is progressing actions against the evaluation report's 8 recommendations, outlined below:

- Agencies should consider which generative AI solutions are most appropriate for their operating environment and specific use cases.

OFFICIAL

- The DTA will develop procurement guidance that will help agencies source AI solutions that best meet their needs.
- Agencies need to ensure their information systems, permissions and processes are appropriately configured to safely accommodate generative AI products.
 - The Australian Government AI assurance framework, which the DTA is currently piloting, will enable agencies to identify and mitigate risks for higher risk AI use cases.
- Agencies should provide specialised training that reflects agency-specific use cases, as well as training that develops general AI capability.
 - The DTA will offer an AI fundamentals training module to agencies that they can adapt to suit their needs. The [Policy for responsible use of AI in government](#) strongly recommends agencies implement additional training for staff as required.
- Appropriate organisational change management interventions are needed to drive adoption, including identifying ‘generative AI champions’.
 - Accountable officials under the policy (as they are designated by agencies) will help to drive the safe and effective adoption of AI within each agency.
- The APS needs to provide clear guidance regarding the use of generative AI, including when consent and disclaimers are needed, and clear articulation of accountabilities.
 - The DTA will update the Interim guidance on government use of public generative AI tools to ensure it provide clear advice on leading practices in using generative AI tools.
- Agencies should conduct detailed analyses of workflows across job families and classifications to identify further use cases that could improve generative AI adoption.
 - The policy provides a clear mandate for agencies to explore the use of AI in their own context, which includes understanding their own unique requirements and use cases and embrace the technology in a responsible way.
- Agencies should share use cases in appropriate whole-of-government forums to facilitate the adoption of generative AI across the APS.
 - The AI and Machine Learning community on the Australian Public Service Commission’s Data and Digital Profession platform and the DTA’s AI Community of Practice provide opportunities for agencies to engage across government on the adoption of AI.

OFFICIAL

- The APS should proactively monitor the impacts of generative AI, including on its workforce, to manage current and emerging risks.
 - The Framework includes impact assessments and will help agencies to manage the risks of AI adoption. The review of the policy and the evolving AI landscape over time will consider emerging issues and risks within the APS.

In line with the DTA's neutrality and role in managing whole-of-government digital procurement arrangements, the report does not endorse future adoption of Copilot specifically. Consistent with the trial participation conditions, agencies will need to undertake their own assessments as part of any future procurements to ensure the product meets their needs and provides value for money for their agency.

Insights from the trial provide a knowledge base that informs the DTA's ongoing work to support future procurement of AI products and services. Following the trial and the evaluation, the DTA assessed that Copilot met our productivity requirements, and have procured licences to continue using the service.

DTA Marketplaces and Panels

The DTA has not used AI in any of its panels or tender evaluation processes to date. The DTA is exploring whether AI can improve consistency and efficiency of future tender evaluations, particularly when assessing applicants' technical case studies against set criteria.

The DTA is currently undertaking a trial using approved case data from the current Digital Marketplace Panel 2 tender to test AI functionality and potential use in future evaluations. Presently, each applicant's technical capabilities are assessed by two evaluators (a 'pair'), with a consensus process when the technical assessments (ratings) differ substantially between the two evaluators. If the AI trial is successful, the DTA may use generative AI alongside a human evaluator in each pair to assess applicant case studies. If the AI's rating significantly differs from the human evaluator's rating, the application will be reviewed by an additional independent human to reach a consensus score and commentary. The DTA's evaluation team could cover a larger number of evaluations, and shorten the time to evaluate applicants when it refreshes its marketplaces.

Complementary initiatives

This section addresses the following terms of reference:

- any other related matters

Learning from the Taskforce's successful multiagency engagement approach, the DTA is leveraging the significant AI related work underway across government to maintain broad alignment and consistency. This includes DISR's work to formulate risk-based regulatory settings for safe and responsible AI across the economy, including the [Voluntary AI Safety Standard](#) and the development of [mandatory guardrails for AI use in high-risk settings](#).

This work has been informed by extensive public and whole-of-government consultation. These initiatives also align to the cross-jurisdictional work of the DDMM including the [national framework for the assurance of AI in government](#).

As AI use is subject to existing legislation, regulation and policy, the DTA have also focused on aligning our work to existing frameworks. This includes working with the agencies responsible for these frameworks to ensure ongoing alignment with potential reforms.

As an example, the DTA is helping to inform the Attorney General's Department's (AGD) work to develop a whole-of-government framework to support automated decision-making systems for better delivery of government services by Australian Government agencies, which may include systems run by AI. The DTA is also engaging with AGD to align our work with reforms associated with the [Privacy Act Review](#).

The DTA is also engaging with international counterparts and organisations to ensure continued alignment with global AI policy developments. This includes engaging with the Organisation for Economic Co-operation and Development to share insights and influence work to establish global best practice approaches.

The DTA will continue to work with agencies across government to respond to the opportunities and challenges presented by AI.