



AI Transparency in Practice

An evaluation of Commonwealth entities' compliance with their obligations regarding AI Transparency Statements

Kimberlee Weatherall, José-Miguel Bello y Villarino and Alexandra Sinclair

AI Transparency in Practice

An evaluation of Commonwealth entities' compliance with their obligations regarding AI Transparency Statements

Kimberlee Weatherall, José-Miguel Bello y Villarino and Alexandra Sinclair¹

10 February 2026

¹ The authors thank Shuxuan (Annie) Luo for her valuable research assistance.

Acknowledgement of Country

In the spirit of reconciliation, we acknowledge the Traditional Custodians of country throughout Australia and their connections to land, sea and community. We pay our respect to their elders past and present and extend that respect to all Aboriginal and Torres Strait Islander peoples today.

Suggested citation

Weatherall, K., Bello y Villarino, J-M., Sinclair, A. (2026). AI Transparency in Practice An evaluation of Commonwealth entities' compliance with their obligations regarding AI Transparency Statements.

ARC Centre of Excellence for Automated Decision Making and Society, Melbourne.

DOI: 10.60836/gbx1-8w63

Copyright © 2026 Kimberlee Weatherall, José-Miguel Bello y Villarino, Alexandra Sinclair. The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited.

ARC Acknowledgement

This research was conducted by the ARC Centre of Excellence for Automated Decision-Making and Society (CE200100005), and funded by the Australian Government through the Australian Research Council



Australian Government

Australian Research Council

Contents

Contents.....	4
Executive Summary.....	5
Findings	5
Recommendations.....	6
Context.....	6
The <i>Policy for the responsible use of AI in government</i>	6
AI transparency matters for the Australian public sector	8
Why we did this research.....	8
Question 1: Are AI Transparency Statements Findable?	9
Identifying relevant agencies	9
Looking for AI Transparency Statements	10
Question 2: Are AI Transparency Statements published, as required by the Policy?.....	11
The Digital Transformation Agency list (entities considered bound by the Policy)	11
What we found that was not on the DTA list.....	13
What we could not find	14
Question 3: Are AI Transparency Statements informative?	15
What should AI transparency statements contain?.....	15
Criteria for assessing informativeness	16
Use of AI	16
How are agencies using AI?.....	17
Accountable officers	18
Human oversight.....	18
Use of public generative AI tools	19
Conclusions and Recommendations	19
Findings	19
Recommendations.....	20
Opportunities for better transparency arising from the 2025 update to the Policy for the responsible use of AI in government	21
Bibliography	22

Executive Summary

This report assesses how well Commonwealth government departments and agencies are complying with a key aspect of the government's own *Policy for the responsible use of AI in government*: the requirement to publish an AI transparency statement. This requirement became mandatory for many Commonwealth government departments and agencies from February 2025.

At core, the goal of this policy is to increase trust in government use of AI, through transparency about AI use, and accountability. To achieve this goal, AI transparency statements ought to meet three criteria. They should be:

1. **Published**, as required by the Policy;
2. **Findable** by an ordinary member of the public interested in understanding government use of AI;
3. **Informative** about government AI use.

Six months after the Policy came into effect, we undertook a two stage-exercise to test these three criteria. We sought to **find** AI transparency statements: first through basic internet searches that a member of the public might undertake, then working from lists of agencies supplied by the Office of the Information Commissioner (OAIC) and the Digital Transformation Agency (DTA). We built a dataset of **110 AI transparency statements**.

Next, we reviewed the **content** of the policies found. For this report we have focused on three key elements of the *Transparency Statements Standard* to assess whether the Statement would be informative to an interested member of the public. However, given the multiple dimensions of the standard, further research into the statements is being conducted and will be published at a later date in 2026.

Findings

1. **Findability:** AI transparency statements are **not** as easily found as they ought to be. **Very few** statements we have identified were accessible through 'a link to the statement [...] placed in a global menu' as recommended by the DTA.
2. **Publication:** While the DTA considers 92 agencies bound by the Policy and compliant, we have identified **30 government entities** potentially within the scope of the policy for which we could not find a Statement.
3. **Informativeness:** AI transparency statements vary significantly in the detail provided. Some are useful and informative, but others are less useful: information is missing or Statements literally comply with the policy and standard without providing useful information to the public. For example, across 110 Statements:
 - a. In five cases, we could not determine whether AI is presently being used.
 - b. 11 entities fail to state that accountable officials have been appointed: 8 of those are using AI. A further 7 state that an accountable official has been appointed without providing any further information or contact details.

- c. 19 Statements repeat or paraphrase text from the Standard without actually providing substantive information, making it difficult to understand the impact of AI on the public;
- d. 14 Statements failed to address human oversight at all, 7 from entities using AI.
- e. There is variation across government departments when it comes to allowing the use of publicly available AI systems for official tasks, which deserves further consideration.

Interestingly, we also found that 14 of the 110 Statements state that the relevant government entity is not using AI at all.

Recommendations

Several simple changes that would improve the findability and informativeness of AI transparency statements. We recommend:

1. **Findability:** AI transparency statements should be published by the DTA in one central location: presumably by the DTA. The list should publish both the entities bound, those complying voluntarily and links to all statements.
2. **Publication:** The DTA should consider mechanisms to ensure that the Policy and requirements are complied with, or, at the very least, note when it records a transparency statement that the DTA may consider it not fully compliant with the standard.
3. **Informativeness:** The AI Transparency Standard should be revised by:
 - a. *Separately* requiring agencies to report (1) whether the public interacts with AI used by the entity; (2) whether AI used by the entity has a significant impact on members of the public, and (3) whether any uses of AI give rise to automated decisions without involving human intervention.
 - b. Requiring agencies to list policies and legislation, compliance with which they have considered in designing and deploying AI.

Finally, we note the window in 2026 to **capture the opportunity:** The recently published version 2 of the *Policy for the responsible use of AI in government* creates an opportunity to significantly improve the informativeness of AI transparency statements.

Context

The Policy for the responsible use of AI in government

AI use in the Australian Commonwealth public sector is not subject to a specific legislative framework: it is governed by general legal frameworks, and by a series of policies, frameworks, standards and guidelines issued by various parts of the Australian government, most notably the Digital Transformation Agency.

Arguably the most important, and the **only AI-specific mandatory** document governing public sector use at a Commonwealth level of which we are aware is the *Policy for the responsible use of AI in government* ("Policy for the Responsible Use of AI in Government - Version 1.1" 2024²).³

The 2024 Policy contained **two mandatory requirements** binding on most, but not all Commonwealth government departments and agencies:

1. An obligation to appoint an **accountable official**;
2. An obligation to publish, on the entity's public-facing website, an **AI transparency statement** setting out, at a high level, how the relevant entity uses and governs AI.

These two pillars, transparency and accountability, are the anchor of the regulatory strategy for public sector AI use in Australia.

According to the Policy v1.1, AI transparency statements were required to be on agency websites by February 2025. The expected content of statements is further elaborated in an accompanying *Standard for AI transparency statements* (Digital Transformation Agency 2024).⁴

The stated aim of the Standard is to build public trust by supporting the public service to engage with AI responsibly. The Standard states that transparency is key to public trust and that the public should have confidence that agencies are monitoring how they use AI and protecting the public from negative impacts. To this end, AI transparency statements should explain how agencies use artificial intelligence (AI), in what domains and with what safeguards. The Standard also recommends that "a link to the statement [be] placed in a global menu, aligned to the approach often taken for privacy".

The Policy for the responsible use of AI in government is not mandatory for all Commonwealth government entities. First, corporate Commonwealth entities are only 'encouraged' to apply this policy. Second, the policy does not apply to the use of AI in the defence portfolio and the 'national intelligence community' (NIC) as defined by Section 4 of the *Office of National Intelligence Act 2018* (Cth). Nonetheless, defence and members of the NIC may voluntarily adopt elements of this policy where they are able to do so without compromising national security capabilities or interests.⁵

² Version 1.1 of the Policy has been superseded by its version 2.0, effective from 15 December 2025 (Digital Transformation Agency 2025a). In this report, unless otherwise noted all references and analysis relate to the 2024 version of the Policy.

³ We note that other general legal frameworks may impose additional transparency requirements. For example, the *Privacy and Other Legislation Amendment Act 2024* (Cth) introduced a requirement that entities bound by the Act – including many Commonwealth government agencies – publish information about automated decision-making in their privacy policies: Sch. 1 Pt 15. This requirement comes into effect 10 December 2026. It could also be argued that (at least some) AI uses could fall within the *Freedom of Information Act 1982* (Cth) subs 8(2)(c) and 8(2)(j), requiring agencies to publish details of their functions and decision-making powers that affect the public; and their "operational information," which includes the rules, guidelines, practices, and precedents used to make decisions or recommendations affecting the public

⁴ There is also a new, substantially equivalent Standard for AI transparency statements v 2.0 applicable from 15 December 2025 (Digital Transformation Agency 2025b), but again we refer in this report to the 2024 Standard binding at the time of our assessment.

⁵ This remains the scope of application for the 2.0 version of the policy issued in December 2025 (Digital Transformation Agency 2025a, 25).

AI transparency matters for the Australian public sector

The National AI Plan (Department of Industry, Science and Resources 2025) positions transparency as a fundamental principle in the regulation and governance of artificial intelligence in Australia. In so doing, it follows the path established in 2019 by the first iteration of the *Australian AI Ethics Principles* (updated this year, Department of Industry Science and Resources 2025) and the various frameworks, guidelines, 'standards' and other statements published by the Commonwealth government since then.

According to the National AI Plan, transparency must be embedded in the government's own practices. The Plan sees transparency in the public sector as a tool for the government to build 'public trust in the technology', explicitly 'by leading the way in adopting AI transparently and responsibly', with the final objective of enhancing the capabilities of government agencies with 'appropriate human oversight' to enable them to operate more efficiently (Department of Industry, Science and Resources 2025, 25).

As this makes clear, transparency enables another key pillar of the Australian government's commitments on the use of AI: *accountability* for human decision-makers delegating to the technology or using the technology as part of their decision-making process. The government has a stated objective of ensuring 'that human decision makers remain accountable for key decisions made with the assistance of AI tools'.⁶

If the goal of accountability is to be achieved, and risks and harms arising from AI use are to be identified and mitigated, a reasonable level of transparency will be required (Weatherall et al. 2024, 69–76). This transparency could be seen as the 'availability of information about an actor allowing other actors to monitor the workings or performance of this actor' (Meijer 2014, 511, emphasis omitted). Holding a system to account requires first 'understanding how it works as a system' (Ananny and Crawford 2018, 983). To put it another way, an actor cannot be held accountable if we do not know what they are doing.

Why we did this research

The research reported here has been conducted as part of *The Regulatory Project* at the Australian Research Council Centre of Excellence for Automated Decision-Making and Society (ADM+S). The mission of ADM+S is to connect 'academic research, industry, government and the community to develop responsible, ethical and inclusive automated decision-making systems, for the benefit of all Australians'. We proposed this work to the Office of the Australian Information Commission (OAIC), in order to support the mandate of OAIC in monitoring agency administration, and its role in advising the public, organisations and agencies about the handling of information by government entities within its remit (*Australian Information Commissioner Act 2000* (Cth), s 7). We also hoped to provide an evidence-based analysis to the outcomes of the important work of the Digital Transformation Agency (DTA) promoting responsible adoption of AI across the Commonwealth government.

We have focused on AI transparency statements, to better understand the effectiveness of the *Policy for the responsible use of AI in government*, and the type and quality of information

⁶ It is also consistent with other transparency commitments of the government, such as those stated in the objects clause in the *Freedom of Information Act 1982* (Cth) s 3. That clause states that by requiring agencies to publish information, the 'Parliament intends ... to promote Australia's representative democracy by contributing towards the following: (a) increasing public participation in government processes, with a view to promoting better-informed decision-making; [and] (b) increasing scrutiny, discussion, comment and review of the Government's activities.'

accessible to the public. We have restricted the research to information publicly available on the departments and agencies' websites, to reflect the access of the general public to these statements.

Our findings are also relevant in assessing the Commonwealth government's current approach of governing AI use in the public sector through frameworks, guidelines and policies only, rather than enacting a legal framework. Neither the Policy, nor the Standard creates consequences for government entities which fail to comply. Nor are they likely to be enforceable using existing legal frameworks: administrative law doctrine is traditionally reluctant to find that a breach of a policy (as opposed to legislation) can provide a legal remedy in a claim for judicial review.⁷ In other jurisdictions, transparency for AI is guaranteed through binding regulatory requirements.⁸

This reality makes it all the more important to assess – and promote – the effectiveness of the government's approach. If light touch, high-level guidance does not create the conditions for accountability for public sector uses of AI, more may be needed.

Question 1: Are AI Transparency Statements Findable?

Our first goal was to assess whether, six months after the Policy came into effect, AI transparency statements were **findable** for all Commonwealth government agencies bound by the Policy.

We faced an immediate challenge. At the time we commenced our research in August 2025, we found:

1. There was no central repository of published transparency statements – or indeed any publicly available information about the levels of compliance with the policy;
2. There was no list that we could find of agencies required to publish AI transparency statements.

This was surprising, because the Standard directs agencies to 'send the DTA a link to the statement when it is published or updated by emailing' the dedicated address within the DTA (Digital Transformation Agency 2025b, 5). The DTA, at least, ought to know, and might perhaps be expected to publish, links to all extant policies.

Identifying relevant agencies

To identify relevant Commonwealth government agencies, we used a list of public federal entities supplied to us by the OAIC. The list, incorporated in the dataset supporting this report,⁹ is structured by portfolios and identifies a total of 224 portfolios, departments, agencies and other entities in the federal government. Of these, 221 seemed to be unique entries.¹⁰

⁷ See for example *Kioa v West* [1985] HCA 81; (1985) 159 CLR 550 *Minister for Arts, Heritage and Environment v Peko-Wallsend Ltd* (1987) 15 FCR 274.

⁸ See a broad review of its governance in Krook et al. (2025).

⁹ The complete dataset will be published in 2026. If earlier access is needed, please email the authors.

¹⁰ Three headings did not seem to correspond to singular entities at the time of compiling the dataset.

We matched that list to the websites of each entity as listed in the official government directory (“Commonwealth Entities and Companies - Directory” 2025). During this process it became apparent that some entities in the list were not envisaged by the Policy and so likely fell outside of the scope of our research. For example, the Attorney-General (the role itself) was listed as an entity, but has no separate website from that of the Attorney-General’s Department.¹¹

In total, we matched 198 of the 221 entities to existing websites in the official Commonwealth Directory. Not all would be bound by the Policy. We nevertheless consider this ‘expanded list’ as the relevant scope of our research, because, as noted above, the Policy explicitly encourages those not bound to publish statements where possible and consistent with their mandate.

Looking for AI Transparency Statements

To determine the **findability** of AI transparency statements six months after the deadline for the agencies and departments to comply with the policy (end of August 2025), we ran two separate search processes.

The first step was a search of the kind that might be used by a technically-savvy individual, who does not have access to the list of entities shared by OAIC, but is still interested in accessing all the relevant documents. With the help of a research assistant, we wrote and ran code to scrape from the internet all the AI transparency statements within the domain ‘.gov.au’ that we could identify. Essentially, the code did a series of iterative Google searches and extracted the data.¹² This process generated an initial dataset with the text of 29 separate transparency statements. After further refining of the process to focus on distinct links to webpages instead of text, we generated a new list with 110 hyperlink results, although the text of some of the statements could not be extracted.¹³

We conducted a manual review of the data to assess its reliability. Notably, **very few** of the statements we had identified were accessible through “a link to the statement [...] placed in a global menu, aligned to the approach often taken for privacy”, as recommended in the Standard.

The second search process focused on each specific entity’s website. We did this initially through an automated search, similar to the one described above, and then through a manual check, navigating the website. Through this process we could identify 101 hyperlinks to statements. We consider this second approach more reliable for several reasons, but, mainly, because:

1. It restricted the search to entities that had individual websites and so were capable of complying with the Policy.

¹¹ This is likely a function of the fact that the list came from the OAIC and reflects their mandate: perhaps the Attorney-General has a role as a decision-maker in some administrative procedures.

¹² Google search: “allintitle: Artificial OR Intelligence OR OR OR AI "transparency statement" site:.gov.au”.

¹³ There are several possible explanations for such a limited output of texts in the initial search, but mainly they seemed to be linked to the scraping techniques in combination with the format of the statements; or because there were crossed links to one single statement. Among others, it was possible that: the content of the pages were designed to be loaded dynamically, rather than serving complete HTML in one go; the link may have been to a ‘.pdf’ file, which could not be extracted; and there could be loading delays until users scroll to that section, as the statements were not in separate pages but embedded into a subsection of other documents, which required a ‘click’.

2. It excluded entities beyond the scope of the policy but with a '.gov.au' address: such as Parliament-related entities or some courts over which the DTA would have no competence.
3. It allowed for a direct matching of entities and transparency statements.

On 27 October 2025 we published the initial results on transparency statement findability in an article on *The Conversation* (Bello y Villarino et al. 2025). We noted our difficulties in finding transparency statements, and that there were still some entities we understood to be bound by the policy, for which we could not find the relevant transparency statement.

Our preliminary results, published in October 2025, remain valid: regardless of the approach (manual or automated) as at the 6-month mark, many AI transparency statements were not readily findable by an ordinary member of the public.

Recommendation: AI transparency statements should be readily findable by the public. The Digital Transformation Agency (or some other entity the Government may consider appropriately resourced for this task) should publish the list of AI transparency statements (with links) in one centralised location.

Question 2: Are AI Transparency Statements published, as required by the Policy?

While the Policy only recommends that AI transparency statements are published in a manner readily findable by members of the public, the Policy *requires* that each bound entity publishes a statement. A first question is: which entities are bound?

The Digital Transformation Agency list (entities considered bound by the Policy)

After publishing our preliminary findings, we were contacted by the DTA through OAIC. After a subsequent meeting, on 6 November 2025 the DTA generously shared their list of all entities with published statements so we could check our dataset. We received a list of **92 entities** from the DTA, which we understand to be a **full list of every entity the DTA considers bound by the Policy**. We also understood that the DTA considered there had been 100% compliance with the Policy by these entities. The DTA list did not share links to, or copies of, the AI transparency statements of those 92 entities.

Comparing our dataset to the DTA list of published AI transparency statements, there are **11 entities** that were not part of our initial dataset. Table 1 lists the entities, why they were not initially found in our research, our judgment whether the entities are in (or out of) scope of the Policy, and the date of any AI transparency statement.

Table 1: Additional Commonwealth entities identified by the DTA

Agency	Reason	Scope	Date of noted publication
Bureau of Meteorology	Not found automatically or manually	In	21-Feb-2025
Australian National Audit Office	Not in OAIC list	In	21-Jul-2025
Australian National Preventive Health Agency	Not in OAIC list + Not identified as having a website	In	Not found
Department of Parliamentary Services	Not in OAIC list	Out (Parliament)	21-Feb-25
Department of the House of Representatives	Not in OAIC list	Out (Parliament)	2-Jul-25
Department of the Prime Minister and Cabinet	Not in OAIC list	In	28-Feb-25
Department of the Senate	Not in OAIC list	Out (Parliament)	21-Feb-25
Office of the Inspector-General of Intelligence and Security	Not in OAIC list	In (Intelligence, but not as NIC)	22-Feb-25
Parliamentary Budget Office	Not in OAIC list	Out (Parliament)	28-Feb-25
Parliamentary Workplace Support Service	Not in OAIC list	Out (Parliament)	26-Feb-25
Royal Australian Mint (the Mint)	Not in OAIC list	In (non-corporate but chartered)	6-Mar-25

As the table shows, ten of these entities were not in the OAIC list. Half of these cases (5) are Parliament-related entities which we would consider completely out of scope of the Policy, even on a “recommended” basis, as they are not part of the executive. The Mint, as a chartered entity, may have a complex status. For the sake of completeness, we have added these entities to our dataset. This may distort slightly other findings.

We agree with the DTA that the Australian National Audit Office and the Department of the Prime Minister and Cabinet are within scope and bound by the policy. It is surprising that they were not in the OAIC list. Both have been added to our dataset and analysed.

We could not find a statement in October for the Bureau of Meteorology (BOM) via automated or manual means, but checking again in November we found a statement. Our conjecture is that this is an artefact of the update of the BOM website (effective on 22 October 2025) (The Bureau of Meteorology 2025). We conducted our searches on bom.gov.au, whereas it is likely that the statement was at the time on the new website, which was running then as beta.bom.gov.au. This is speculation, because the DTA did not provide a link to the statement notified to them. The current BOM statement is in our dataset.

The DTA list indicates that the Australian National Preventive Health Agency has an AI transparency statement. At the time of writing, we have still not found either a website, or a transparency statement for this entity. The Commonwealth government directory notes that the Agency ceased operations in 2014 (“Australian National Preventive Health Agency in Search - Directory” 2025).

Finally, the DTA list included the Office of the Inspector-General of Intelligence and Security as an entity bound by the policy. We believe this is correct. Although it is intelligence-related (generally exempted), it is not registered in the list of explicit exclusions for the National Intelligence Community. This Office was not in the OAIC list, but we have now included it in the dataset and analysis.

After these additions, our final list consists of **110 transparency statements** – some from agencies bound by the Policy, others apparently voluntarily published.

What we found that was not on the DTA list

The DTA listed 92 entities as bound by the Policy and having published AI transparency statements. We found (and included in our dataset) a number of transparency statements of entities not identified in the list shared by the DTA. The complete list is below in Table 2.

The main question raised by Table 2 is which (if any) of these entities should be considered to be bound (and therefore should have been included in the DTA list), and which are fulfilling the policy as voluntary compliance (i.e., justified not to be on a list tracking binding policy implementation).

*Table 2: Entities with AI transparency statements **not** listed by the DTA as bound by the Policy*

<ul style="list-style-type: none"> • Administrative Review Council • Food Standards Australia New Zealand • Australian Charities and Not-for-Profits Commission • Independent Health and Aged Care Pricing Authority • Australian Digital Health Agency • Inspector-General of Water Compliance • Australian Institute of Health and Welfare • Murray-Darling Basin Authority • Australian Maritime Safety Authority • National Library of Australia • Australian National Maritime Museum • National Native Title Tribunal 	<ul style="list-style-type: none"> • Australian Reinsurance Pool Corporation • National Offshore Petroleum Safety & Environmental Management Auth. • Australian Skills Quality Authority • National Transport Commission • Comcare • Office of the eSafety Commissioner • Defence Housing Australia • Screen Australia • Department of Veterans' Affairs • Tax Practitioners Board • Family Law Council
---	--

There are clearcut cases, such as the Department of Veterans' Affairs or Defence Housing Australia. They are squarely within Defence and exempt. These entities probably published AI transparency statements on the basis that they were able to do so without compromising national security capabilities or interests. Others such as the Administrative Review Council

could be justified as they may not be bound by a policy for the executive branch, but see benefits in transparency about AI use.¹⁴

However, for the majority of these agencies, it is difficult to see why they would *not* be bound, meaning they should have been tracked by the DTA.

What we could not find

Of our final list of 208 entities (198 from the OAIC list plus 10 from the DTA list), we could not find, at the time of collecting the data transparency statements for 100 entities. Among those, we filtered out those that are clearcut cases of corporations, the national intelligence community, defence, international or part of the judiciary. To be conservative, we also filtered out some that could be covered by another statement, such as the Australian Energy Regulator, which may be covered by the one from the Clean Energy Regulator.

Following this process, we still found **30 entities that could be potentially within the scope of the policy for which we could not find a statement**. The complete list is below in Table 3.

Table 3: Entities potentially bound by the Policy for which no AI transparency statement was found

<ul style="list-style-type: none"> • Australian Pesticides and Veterinary Medicines Authority • Australian Anti-Slavery Commissioner • Australian Human Rights Commission • Australian Institute of Marine Science • Australian Renewable Energy Agency • Repatriation Medical Authority • Australian Curriculum, Assessment and Reporting Authority • Safety, Rehabilitation and Compensation Commission 	<ul style="list-style-type: none"> • Australian Sports Commission • Civil Aviation Safety Authority • Creative Australia • National Film and Sound Archive of Australia • National Gallery of Australia • Australian Institute of Aboriginal and Torres Strait Islander Studies • Executive Director of Township Leasing • Indigenous Business Australia • Merit Protection Commissioner • Gene Technology Regulator 	<ul style="list-style-type: none"> • Torres Strait Regional Authority • National Disability Insurance Agency • Australian Competition Tribunal • Companies Auditors Disciplinary Board • Financial Reporting Council • Housing Australia • Reserve Bank of Australia • Takeovers Panel • Net Zero Economy Agency • Office of the Registrar of Indigenous Corporations • Remuneration Tribunal • Australian Film, Television and Radio School
---	--	--

¹⁴ It is also possible, that the publication of a transparency statement may be triggered by the ‘portfolio agencies approach’ in those cases. That is, a Secretariate supports council members and the Secretariate is provided by and reports to an entity bound by the policy, so it would make sense for the Secretariate to publish a statement. For the Administrative Review Council, the relevant portfolio seems to be the Attorney-General’s Department.

It is not for us to consider why these entities are not part of the DTA list. In our view, at least some of them, such as the Australian Renewable Energy Agency, should be bound by the Policy and, hence, may be considered non-compliant. For some others, such as the Australian Human Rights Commission, even if one can see why they should not be bound by such a Policy to guarantee their independence, not having a discoverable AI transparency statement seems, at least, slightly ironic, recalling that transparency figured highly in their Human Rights and Technology Final Report (Australian Human Rights Commission 2021).

Question 3: Are AI Transparency Statements informative?

What should AI transparency statements contain?

The Digital Transformation Agency's (DTA) *Standard for AI Transparency Statements* notes that AI transparency statements should publicly disclose:

- how AI is used and managed by the agency
- a commitment to safe and responsible use; and
- compliance with the policy.

As stated in the Standard, statements 'are intended to provide a high-level overview of agency AI use and management'. Specifically, the Standard requires the following:

At a minimum, agencies must provide the following information regarding their use of AI in their transparency statement:

- *the intentions behind why the agency uses AI or is considering its adoption*
- *classification of AI use according to usage patterns and domains, as listed at Attachment A*
- *classification of use where the public may directly interact with, or be significantly impacted by, AI or its outputs without human review*
- *measures to monitor the effectiveness of deployed AI systems and protect the public against negative impacts*
- *overview of compliance with the requirements under the Policy for responsible use of AI in government*
- *compliance with applicable legislation and regulation*
- *when the statement was most recently updated.*

The criteria are mandatory, not illustrative.

A member of our team read all 110 AI transparency statements in our dataset and coded them as described in the next section to address the points required for this report. We are still

conducting a deeper review of the published statements against the broader stated aims and requirements of all the criteria: in this report we focus on 3 high level criteria for informativeness. Reviewing the transparency statements has provided a number of insights into how government departments are engaging with both AI itself and with the Standard.

Criteria for assessing informativeness

To evaluate the informativeness of AI transparency statements, we first set out to assess whether each AI transparency statement **met the criteria** quoted above: an objective question of compliance. It quickly became evident, however, that focusing only on compliance would be insufficient to draw conclusions about **informativeness**. Reading and coding the Statements, we found it was possible for an AI transparency statement to *formally* meet each requirement but still fail to be informative to members of the public seeking to understand government use of AI.

We therefore coded each Statement on the following three criteria, which we consider core elements of the standard:

1. Does the Statement record that the relevant entity is using AI? If not, further detail in the Statement would not be expected.
2. Does the Statement identify, and provide contact details for, any accountable official?
3. If the entity is using AI, how did the entity explain how their AI use interacted with the public?

As a general comment, we observe that the AI transparency statements vary widely in the level of detail provided. Some Statements exhibit best practice, going beyond what is required by the Standard: for example, by detailing examples of specific AI use cases.¹⁵ Some provide descriptions of internal AI governance processes and working groups,¹⁶ and give the name and contact details for the relevant accountable AI official so that the public can contact them.¹⁷

We also observed cases of **'literal compliance'**: where the AI transparency statement literally responds to a criterion in the Standard, but in such a way that little information can be gleaned about how a department is actually using AI.

Finally, there are several instances of outright non-compliance with the Standard. Obviously, when agencies indicated that were not using AI, not setting out all the criteria in the Standard was understandable. However, for those using AI, the statements that did not cover all points failed to comply with the Standard, as the criteria are not optional.

Use of AI

One of the Standard's key purposes is to inform the public as to whether a department is using AI. **14 of the 110 government entities said they were not using AI at all.**

¹⁵ See for example the detailed use case provided by IP Australia < <https://www.ipaustralia.gov.au/about-us/accountability-and-reporting>>.

¹⁶ The Australian Law Reform Commission set out in detail the duties of the accountable AI official, the official's name and their internal reporting structures. See < <https://www.alrc.gov.au/about/corporate-information/ai-transparency-statement/>>.

¹⁷ For example the Australian Research Council provides the name, phone number and email address of their accountable AI official see < <https://www.arc.gov.au/about-arc/arc-strategies-and-policies/ai-transparency>>.

But there were another **5** statements from which **we could not determine if the agency is currently using AI**. These statements are vague or are framed in the future tense, expressing a potential interest in piloting or developing AI capabilities without being clear whether AI is currently used. For example, the Statement of the Office of the E-Safety Commissioner states:

'We may employ AI across various corporate and enabling functions, including software engineering, data analytics and workplace productivity'.

It is not evident from this statement whether the use is current practice, an intermittent use or something the Office may do in the future.

How are agencies using AI?

Many departments expressly state that they are not using AI to carry out decision-making. Very few Statements set out monitoring or evaluation processes for their AI use. Overwhelmingly government entities using AI cited **workplace productivity** as the most common use.

Notably, **many Statements do not provide specific AI use cases**. Although use cases are not required explicitly by the Standard, in practice the failure to clarify this point often means that the information provided is so general that it is difficult to have any sense of how the public's data may be employed or what specific AI applications are used. For example, the Australian Federal Police Statement says:

The AFP's current application of AI has been focused on exploring opportunities to enrich data to enhance human efficiency when processing large amounts of data.

It is difficult ascertain from this statement how (or even if) AI is being used: 'enriching data' could be, for example, removing noise from audio or visual data – or it could involve making predictions about potential offenders. 'Exploring' could be testing options with providers.

Other agencies equally offer very limited insight into AI uses, through a completely different approach, providing instead a long list of very general uses. For example, the Department of Agriculture, Fisheries and Forestry (DAFF) notes that:

The department is using AI to:

- *Automate routine tasks, manage workflows, create content, and facilitate communication*
- *Summarise high volumes of documents or information*
- *Identify, produce or understand insights from data*
- *Understand patterns and trends in large data sets*
- *Categorise documents for storage and retention.*

Finally, many government entities note that they are only piloting or trialling AI and will provide a detailed transparency statement once the pilot is over.¹⁸ We note that the DTA Standard makes no exception for trial or pilot uses, which should therefore in theory be accorded the same levels of transparency.

¹⁸ See for example the National Anti-Corruption Commission <<https://www.nacc.gov.au/ai-transparency-statement>>.

Accountable officers

11 agencies fail to state that accountable officials have been appointed. Only 3 of those 11 were not using AI.¹⁹ **That is, 8 departments have a published AI transparency statement which notes that they are using AI and but fail to record whether they have appointed an accountable AI official.** The Standard requires the DTA to be notified of the accountable official. It is possible that the non-compliant agencies may have notified DTA of their official but have not recorded it on their transparency statement.

An additional 7 departments stated they had appointed an accountable official but provided no further information about that person such as their role in the department or their contact details. If no information is provided about the official, it is hard for the public to contact them to find out about the department's use of AI.

Human oversight

In relation to how AI interacts with the public, the Standard states that:

At a minimum, agencies must provide the following information regarding their use of AI in their transparency statement:....classification of use where the public may directly interact with, or be significantly impacted by, AI or its outputs without human review.

Some departments simply repeat this statement verbatim. For example, the Department of Treasury states:

The Treasury does not use AI in ways that the public may directly interact with or be significantly impacted by without human intervention.

The problem with this Statement is that it cannot be determined whether the public:

- is not *directly* interacting with AI or;
- may be impacted by AI, but is not *significantly* impacted; or
- *is* directly interacting or significantly impacted by AI, but there is human intervention.

This is caused by the drafting of the Standard, which elides three distinct ways that departments may use AI. Hence, it cannot be determined how the department is using AI at all. **19 of the 110 statements repeat or paraphrase this wording of the Standard**, saying that there is no use where the public direct interacts with AI or is significantly impacted by AI without human intervention.

This is a problem, not with compliance, but with the Standard itself: and, again, links to this idea of **literal compliance** discussed above. Departments are complying with the requirements of the Standard but in such a way that we cannot glean useful information about interaction with the public or any human oversight process.

These statements contrast with those providing very detailed information. Many stated that they used a 'human in the loop' to review all AI-generated outputs. IP Australia explained they had a solely automated tool for checking trade marks that directly interacted with the public.

¹⁹ Or in the case of Defence Housing Australia it was unclear.

Finally, and probably most problematically in this context, **14 of the statements failed to address the question of human oversight at all**, even though this is a requirement of the Standard. **7 of the 14** agencies that did not address human oversight were **using AI but did not explain whether the public could directly interact with or be significantly impacted by AI**.

Use of public generative AI tools

A final interesting observation relates to generative AI tools. Publicly available generative AI tools' outputs may (significantly or not) impact the public, but there is likely to be a human intervention in the interaction with the tool. That is, the generation of text is unlikely to happen automatically. At the time of our snapshot, government entities varied in their approach. Some departments specify that staff can only use Generative AI tools offered as part of their internal Enterprise system. The Australian National Audit Office for example says: 'The policy limits the use of publicly available generative AI tools' and the Trade and Investment Commission state that only Copilot is approved.

The Department of Finance allows staff to use publicly available tools as well as enterprise versions. They stated:

We allow our staff to use AI in their work with the objective of enhancing productivity and service delivery. This includes enterprise AI deployed in our closed internal ICT environment like Microsoft 365 Copilot, as well as publicly available AI like ChatGPT, Claude and Gemini that is that is not deployed in our closed internal ICT environment.

The Parliament of Australia we would consider outside the scope of the policy, but nevertheless have a transparency statement which notes use of public tools is allowed, but there is a splash screen warning user of the risks. Similarly, the Department of Agriculture Fisheries and Forestry (DAFF) also permits use of public generative AI tools.

Conclusions and Recommendations

Findings

Our review shows that six months after the Commonwealth government's *Policy for the responsible use of AI in government*:

- Many (but, very likely, not all) Commonwealth government agencies bound by this Policy had published an AI transparency statement online;
- Few had made those Statements findable by regular members of the public as recommended by the DTA; and
- While some of the Statements are useful and informative, many offer less transparency than the public is entitled to expect, or need in order to trust in government use of AI.

These findings suggest that the requirements of version 1 of the *Policy for the responsible use of AI in government* and *Standard for AI transparency statements* are not achieving their aspiration of providing the kind of transparency that is likely to increase public trust or enable accountability for AI use in government.

Recommendations

There are several simple changes that would improve the findability and informativeness of AI transparency statements. We recommend that:

1. **Findability:** AI transparency statements should be published in one central location: presumably by the DTA (although it could be another Commonwealth government entity). This could be by publication of Statements, or links to Statements. This would make it 'easier [for the public] to understand and compare how government agencies adopt AI' which is a stated goal (Digital Transformation Agency 2024, 4). This would also make it easier to determine which departments have or have not complied with their obligations to prepare one.
2. **Publication:** The DTA should reconsider the entities bound and explore mechanisms to ensure that the Policy and requirements are complied with, including by considering what consequences flow from non-compliance.
3. **Informativeness:** The AI Transparency Standard should be revised by:
 - a. Removing elision of issues by **separately** requiring agencies to report (1) whether the public interacts with AI used by the entity; (2) whether AI used by the entity has a significant impact on members of the public, and (3) whether which, if any uses of AI give rise to automated decisions, that is, do not involve human intervention. These are critical issues currently elided in text from the Standard which is parroted in numerous AI transparency statements.
 - b. Requiring the DTA or other entity to conduct a minimal assessment of substantive compliance with the Standard, focusing on: envisaging a way to contact any accountable official; that statements make clear whether AI is being used; and considering whether the descriptions of use are minimally informative.

Opportunities for better transparency arising from the 2025 update to the Policy for the responsible use of AI in government

This is an evolving space, and the Commonwealth government is making considerable investments in upskilling the Australian Public Service. With these actions the government is also increasing the mandatory requirements for agencies in relation to their responsible AI use.

In the second half of 2025, significant additional rules, guidance, and upskilling offered by the Commonwealth government for the Australian public sector have been announced. In particular:

- In July 2025, the DTA published a detailed new *AI Technical Standard* for the public sector, outlining best practices for end-to-end design, development, deployment and use of AI systems.
- In November 2025, an *AI Plan for the Australian Public Service* was published. The plan includes updates to policies, the appointment of Chief AI Officers and an AI Review Committee, and work to provide tools, training and skills to members of the public service.
- In December 2025, Version 2 of the *Policy for the responsible use of AI in government* was published. Version 2 sets out to strengthen the government's approach to safe and responsible AI through requirements for agencies to:
 - Develop a strategic approach to adopting AI;
 - Establish an approach to operationalise the responsible use of AI, including, for example, pathways for staff and members of the public to report AI safety concerns, as well as clear processes to address AI incidents.
 - Assess all new (and over time, all existing) AI use cases, and, if they are within scope of the policy, conduct an annual AI use case impact assessment. Agencies must create a register of in-scope use cases, and designate accountable use case owners.

As yet, **transparency** is not an area where Commonwealth policy has yet shifted. None of these new requirements are required to be described or outlined in agency AI transparency statements: not the strategy, the operationalisation of responsible AI, the AI use case register, and not the impact assessments.

We suggest that there is an opportunity build on version 2 of the Policy, to make AI transparency statements better. In particular, we suggest that an agency's AI use cases and governance arrangements should be described in its AI transparency statement.

It is worth noting that in our reading of statements published in 2025, the most useful information was where departments provided specific use cases such as IP Australia explaining how their solely automated trade mark checker interacts with the public. Only with the disclosure of specific use cases are the public able to meaningfully withdraw consent in certain circumstances or seek further information. Ideally, the government should already be working towards a version 3 in which the repository of transparency statements also includes a list of use cases or systems, following models in other jurisdictions.

Bibliography

- Ananny, Mike, and Kate Crawford. 2018. "Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability." *New Media & Society* 20 (3): 973–89. <https://doi.org/10.1177/1461444816676645>.
- Australian Human Rights Commission. 2021. *Human Rights and Technology Final Report*. https://tech.humanrights.gov.au/sites/default/files/2021-05/AHRC_RightsTech_2021_Final_Report.pdf.
- "Australian National Preventive Health Agency in Search - Directory." 2025. December. <https://www.directory.gov.au/search?keywords=Australian+National+Preventive+Health+Agency+&op=Search>.
- Bello y Villarino, José-Miguel, Kimberlee Weatherall, and Alexandra Sinclair. 2025. "Most Australian Government Agencies Aren't Transparent about How They Use AI." *The Conversation*, October 26. <https://doi.org/10.64628/AA.qw4rcwnrc>.
- "Commonwealth Entities and Companies - Directory." 2025. August. <https://www.directory.gov.au/commonwealth-entities-and-companies>.
- Department of Industry Science and Resources. 2025. "Australia's AI Ethics Principles." November 28. <https://www.industry.gov.au/publications/australias-ai-ethics-principles>.
- Department of Industry, Science and Resources. 2025. "National AI Plan." December. <https://www.industry.gov.au/sites/default/files/2025-12/national-ai-plan.pdf>.
- Digital Transformation Agency. 2024. "Standard for AI Transparency Statements." September. <https://www.digital.gov.au/ai/ai-in-government-policy/standard-ai-transparency-statements>.
- Digital Transformation Agency. 2025a. "Policy for the Responsible Use of AI in Government - Version 2.0." December 1. https://www.digital.gov.au/sites/default/files/documents/2025-12/Policy%20for%20the%20responsible%20use%20of%20AI%20in%20Government%202.0_0.pdf.
- Digital Transformation Agency. 2025b. "Standard for AI Transparency Statements - Version 2.0." December 1. https://www.digital.gov.au/sites/default/files/documents/2025-12/Standard%20for%20AI%20transparency%20statements%202.0_0.pdf.
- Krook, Joshua, Peter Winter, John Downer, and Jan Blockx. 2025. "A Systematic Literature Review of Artificial Intelligence (AI) Transparency Laws in the European Union (EU) and United Kingdom (UK): A Socio-Legal Approach to AI Transparency Governance." *AI and Ethics* 5 (4): 4069–90. <https://doi.org/10.1007/s43681-025-00674-z>.
- Meijer, Albert. 2014. "Transparency." In *The Oxford Handbook of Public Accountability*. Oxford University Press.
- "Policy for the Responsible Use of AI in Government - Version 1.1." 2024. September 1. <https://www.digital.gov.au/ai/ai-in-government-policy>.
- The Bureau of Meteorology. 2025. "New Website Launches 22 October 2025." October 15. <https://www.bom.gov.au/news-and-media/new-website-launches-22-october-2025>.



Weatherall, Kimberlee, Paul Henman, Jose-Miguel Bello y Villarino, Rita Matulionyte, Lyndal Sleep, and Melanie Trezise. 2024. *Automated Decision-Making in New South Wales: Mapping and Analysis of the Use of ADM Systems by State and Local Governments*. Report. ARC Centre of Excellence for Automated Decision-Making and Society. New South Wales. <https://apo.org.au/node/325901>.