Submission to the Office of the Australian Information Commissioner

Automated Decision-Making Transparency Obligation (APP 1) — Issues Paper

18 May 2026 | Closing date: 15 June 2026

Publication consent: This submission may be published in full as submitted. It does not contain confidential information.

ADM transparency lifecycle — the practical model

This submission treats ADM transparency as an eight-stage reasoning governance lifecycle, not a one-off privacy policy disclosure. The value of this submission is not the sum of its individual question responses — it is the integrated architecture connecting them. Reviewers are encouraged to assess this submission as a whole rather than question by question.

Stage	What the entity does	Why it matters
1 Define the terminal decision	Identify the specific decision to be made and specify what a defensible version looks like — what level of accuracy, fairness, and explainability is required	The defensibility standard must be established before any tool or process is designed. Without it there is no basis for testing whether the decision pathway is adequate
2 Map the decision pathway	Working backward from the terminal decision point, identify every information transformation, decision-shaping reasoning activity, and contributory decision point — with or without human involvement — required to reach it	The decision pathway, not the terminal decision point alone, is the unit of governance. Contributory decision points constrain what is possible at the final step and must be identified before they can be governed
3 Test whether the decision pathway is defensible	Before selecting any tool, assess whether the designed decision pathway can withstand scrutiny — is the logic sound, are intermediate conclusions valid, would a reasonable person accept this pathway as adequate?	A defensible terminal decision requires a defensible decision pathway. Tool selection should not precede this test
4 Match tool to contributory decision point	For each step: if structured information and rule-based execution, conventional software is appropriate. If unstructured information or adaptive and strategic reasoning is required, AI may be involved — and the information environment and reasoning type determine governance needed. AI	Tool selection is a function of what reasoning is required and what information is involved — not what technology is available

Stage	What the entity does	Why it matters
	can assist in assuring step design where that assurance function is itself governed	
5 Assess governance risk within each contributory decision point	For each step: what is the likelihood of error given the tool and information environment? What is the impact on the terminal decision? What oversight level is proportionate?	Risk accumulates across the decision pathway. Errors in early steps compound through subsequent steps. Governance must be calibrated at step level, not only at outcome level
6 Implement separated governance functions	Separate compliance checking, risk identification, and access control across the decision pathway. No single system or actor performs all three. These functions need separate authority, separate information access, and independent ability to raise concerns	Collapsed governance functions produce undetectable failure. The institutional design wisdom of separation of powers applies to AI systems performing consequential reasoning
7 Test and iterate until risk is acceptable	Run the governed decision pathway against the defensibility standard from Stage 1. Where risk exceeds the acceptable level, revise the process, tool selection, or governance arrangements and iterate	Governance is not a one-time assessment. ADM systems evolve and new failure modes emerge. The decision pathway must be tested continuously against the standard it was designed to meet
8 Verify audit capability and maintain transparency	For any specific decision, contributory decision points must be reconstructable after the fact — not just explainable in principle but traceable at contributory decision point level. Measure and disclose the governance scorecard. Maintain disclosures against actual practice	Audit capability enables self-correction when errors occur. Without it, errors are discoverable only through external complaint rather than internal monitoring

How to read this submission

Two layers. First, direct responses to Questions 1–12. Second, a cross-cutting reasoning governance architecture connecting those responses through the eight lifecycle stages above.

Architectural value. The value of this submission is the integrated architecture, not the individual parts. Each concept has precedents in existing ADM governance frameworks including the Commonwealth Ombudsman's ADM Better Practice Guide[1] and the Western Australia Privacy and Responsible Information Sharing Act 2024.[2] This submission's contribution is their integration into a single operational lifecycle that prevents the ADM transparency obligation from becoming generic privacy policy boilerplate. Please assess this submission as a whole.

On frameworks referenced. This submission references governance frameworks developed by [redacted], published under CC BY 4.0. All framework concepts appear in plain language in the submission body. Technical vocabulary is in Appendix B only. The frameworks are offered as operational specifications, not as legal authority.

Classification labels. Each question response ends with a classification label indicating what type of action the recommendation requires:

- **Current APP 1 guidance** — the OAIC can include this in guidance now under existing APP 1.7–1.9. No legislative change needed.
- **Good practice** — sensible for entities to do but not strictly required by the current obligation. The guidance could recommend without mandating.
- **High-risk ADM** — appropriate as a stronger requirement for AI systems performing consequential reasoning on complex information, or decisions affecting essential services.
- **Future reform** — requires legislation or broader policy development. Flagged so the OAIC can signal it as necessary next-step legislative development.

Mixed labels (e.g. **Current APP 1 guidance / Good practice**) indicate that the basic version is current guidance and the extended version is good practice — the proportionality table in Q4 calibrates which applies in a given context.

OAIC question	Lifecycle stage(s)	Submission answer — one line	Recommendation
Q1	Stages 2–3	Assess the full decision pathway; apply two additional factors: decision pathway reconstructability (good practice) and reasoning neutrality	Rec 1
Q2–Q3	Stage 4	AI summaries and recommendations substantially assist decisions where the reviewer cannot inspect the decision pathway	Rec 2
Q4–Q6	Stage 5	Significant impact includes practical ability to understand the decision pathway and access available review or complaint pathways	Rec 4
Q7	All stages	Existing frameworks do not address how governance survives machine-speed ADM or AI-generated local rules	Recs 1–7
Q8–Q9	Stage 5	Differential pricing significance depends on challenge pathway, not magnitude	Rec 4

OAIC question	Lifecycle stage(s)	Submission answer — one line	Recommendation
Q10	Stage 1	Non-promotion of a job ad is a decision — the default was never classified as a governed contributory decision point	Rec 3
Q11	Stage 6	Third-party ADM requires consideration of inspection rights and governance arrangements in procurement	Rec 5
Q12	Stages 7–8	Disclosures should be tiered, measured against the governance scorecard (voluntary), and maintained over time	Recs 6–7

Executive summary

This submission proposes that ADM transparency should be treated as an eight-stage reasoning governance lifecycle — not a one-off privacy policy disclosure — and that guidance structured around this lifecycle will prevent the APP 1 obligation from becoming generic boilerplate.

Three connected arguments run through every response, each directly actionable by the OAIC in developing guidance.

First, the "substantially and directly related" test as drafted creates a structural incentive to avoid it. Entities can distribute substantive analytical work across a decision pathway made up of contributory decision points — including decision-shaping reasoning activities such as selecting, filtering, characterising, weighting, and framing information — so that no single point appears to cross the threshold. The governable unit is the decision pathway, not the terminal decision point alone. The terminal decision is the aggregated outcome of all contributory decision points and cannot be meaningfully explained without tracing the pathway that produced it. The guidance should make this explicit and frame the test around two operationalised measures: decision pathway reconstructability (can the entity reproduce the contributory decision points without the system's own output?) and reasoning neutrality (does the system have an undisclosed default position?).

Second, human oversight without access to the decision pathway is not governance — it is ceremony. A human who can technically override an ADM output but cannot independently inspect the contributory decision points that produced it has authority that is formally present and substantively empty. Guidance that treats human oversight as a governance guarantee without examining what the human can actually inspect and independently evaluate will reproduce the accountability failure Robodebt exemplified.[3]

Third, trained-in bias in AI systems sits in a governance gap that no existing framework addresses substantially. It is not discoverable by code audit and not addressable by employment contract. Guidance should name this gap explicitly and signal that future legislative development will need to close it.

This submission also includes a forward-looking section — **Appendix D** — identifying three structural issues that exceed what guidance can deliver but which the OAIC should signal as necessary for future legislative development: a protected right for affected persons to challenge specific contributory decision points in the decision pathway rather than only the terminal decision; mandatory structural separation of governance functions for high-risk ADM; and periodic mandatory reconciliation of disclosed and actual ADM practice.

Evidence from the OAIC's own findings. The October 2025 desktop review found that only 17 per cent of 23 government agencies statutorily authorised to use ADM disclosed that use under existing FOI Information Publication Scheme frameworks.[4] This is the closest existing analogue to the new obligation. The baseline is already failing. The guidance should address why disclosure obligations fail in practice — not only what they should contain in principle.

Seven recommendations:

- **Recommendation 1 —** Assess ADM across the full decision pathway and apply two additional factors: decision pathway reconstructability (good practice) and reasoning neutrality.
- **Recommendation 2 —** Test the quality of human oversight: can the reviewer access, understand, question, override, and record reasons about the ADM output independently?
- **Recommendation 3 —** Require entities to identify hidden ADM use including contributory decision points and upstream decision-shaping reasoning activities not internally classified as decision-making.
- **Recommendation 4 —** Assess significant impact through practical consequences including the affected person's ability to understand the decision pathway and access available review or complaint pathways.
- **Recommendation 5 —** Require third-party ADM transparency in procurement — entities should consider whether contracts include inspection rights and understand the governance arrangements in place.
- **Recommendation 6 —** Design meaningful disclosures tiered to the complexity of the reasoning involved.
- **Recommendation 7 —** Maintain disclosures over time with assigned ownership, review cadence, written update records, and voluntary governance self-assessment.

Question 1

Why guidance should assess the full decision pathway — and two additional factors the OAIC's list does not include

What factors enable an entity to assess whether a computer program substantially facilitates and is directly connected to a human decision-maker's decision?

Supports lifecycle Stages 2–3. Develops Recommendation 1.

Direct answer: A computer program substantially and directly facilitates a decision when the entity cannot reproduce the intermediate contributory decision points that produced the system's output without relying on the system's own description. The guidance should also

require assessment of two additional factors not in the OAIC's current list: decision pathway reconstructability (good practice) and reasoning neutrality.

The decision pathway problem. ADM operates through a chain of contributory decision points — performing decision-shaping reasoning activities: selecting information, filtering what is surfaced, characterising and framing it, weighting factors, generating a recommendation, qualifying that recommendation. These contributory decision points are the components whose aggregate produces the terminal decision. Each step shapes what is cognitively available to the human at the final point. Governance that captures only the final step misses where the substantive reasoning occurs.

The test as currently drafted operates as a single-node test on a chain phenomenon. This creates a predictable response from sophisticated entities: distribute the substantive analytical work across multiple contributory decision points so that each appears merely preparatory and no single step clearly crosses the threshold. Entities can insert nominal human review steps, route AI outputs through human checklists, and frame the human as the decision-maker while the substantive reasoning has already occurred upstream. This is gaming the substantial/direct test by distributing reasoning across the chain.

A decision-making process should be understood as a decision pathway made up of multiple decision points. The terminal decision point is the final, legally recognised outcome — the grant or refusal, the approval or rejection, the access provided or withheld. The terminal decision point is the point most readily recognised in APP 1 and administrative law as “the decision.” However, the terminal decision point is the aggregated outcome of contributory decision points that precede it: decision-shaping reasoning activities including selecting information, filtering what is surfaced, characterising and framing evidence, assigning weight to factors, ranking priorities, summarising records, or recommending outcomes.

These contributory decision points are not merely clerical or preparatory where they materially shape the information, analysis, or recommendations available to the decision-maker at the terminal point. For APP 1.7(a) purposes, guidance does not need to resolve in every case whether each contributory decision point is itself the statutory “decision.” It is sufficient that, where a computer program performs a decision-shaping reasoning activity at a contributory decision point and it materially shapes the terminal decision, the program may have done a thing substantially and directly related to making that decision.

The decision pathway — the full sequence of contributory decision points through which information is selected, transformed, evaluated, and acted upon — is therefore the governable unit for traceability, explainability, and transparency purposes, because the terminal decision is the aggregated outcome of all contributory decision points and cannot be meaningfully explained without tracing the pathway that produced it. [9] [11] [10] [1] [12] [13] [14]

The explanation test — the operative test for 'substantially facilitates'. The guidance should include this practical test: can the entity, without relying on the AI system's own output or description, explain the sequence of contributory decision points that produced that output? If the answer is no, the system is substantially facilitating the decision. This test is concrete, falsifiable, and resistant to architectural gaming. If a decision made with ADM assistance is later challenged — by the affected person, by a regulator, by a court — and the entity cannot reproduce the contributory decision without the system's own self-description, it cannot demonstrate the decision was sound. Signing off on unreproducible reasoning is not an accountable decision: it is an authorisation of an unaccountable process.

Two additional factors — decision pathway reconstructability and reasoning neutrality

The OAIC's listed factors assess the impact and integration of ADM. These two additional factors operationalise concepts fundamental to meaningful ADM governance that are absent from the current list.

Additional factor 1 — Decision pathway reconstructability: transparency operationalised. (Good practice) Standard transparency asks whether the entity disclosed that ADM was used. Reconstructability asks whether the entity can show how the decision was actually made, step by step. An entity that discloses the existence of ADM without being able to reproduce the decision pathway has met the form of the transparency obligation but not its substance.

Decision reconstructability is measurable at four levels:

- **Level 0 — Unrecoverable:** the decision pathway cannot be traced for a specific decision after the fact.
- **Level 1 — Category-level:** the entity can identify which system and rule category applied but cannot reproduce the specific reasoning for a specific decision.
- **Level 2 — Step-level:** the entity can reproduce each contributory decision point for a specific decision. Errors can be located in the chain.
- **Level 3 — Independent reconstruction:** a third party with no access to the system's own description can reproduce the decision pathway from primary inputs. The reasoning is independently verifiable, not just self-described.

Entities operating high-risk ADM systems are encouraged to assess and consider disclosing their decision pathway reconstructability level as a governance good practice.

Additional factor 2 — Reasoning neutrality: non-bias operationalised. Both approve-by-default and reject-by-default configurations are bias positions. A system that defaults to recommending approval is not neutral — it is biased toward approval. A system that defaults to rejection is not cautious — it is biased toward rejection. Neutrality requires structural design: separating the function that generates a recommendation from the function that evaluates it, and from the function that acts on it. The guidance should require entities to disclose their ADM system's default position and to assess whether that default has been examined, authorised, and documented. A system's default position is itself an indicator of how substantially it facilitates the decision.

A protected right for affected persons to identify and raise concerns about specific contributory decision points — rather than only the terminal decision — is addressed in Appendix D as a future legislative development.

Rankings of listed factors: (1) ability to explain each contributory decision point independently; (2) degree of reliance on ADM output; (3) integration into workflow; (4) nature of output as advisory versus determinative — least reliable because a system labelled advisory may perform all substantive reasoning; (4) ability and likelihood of human override — least reliable because override without access to the reasoning basis is formally available and substantively empty.

Recommended guidance wording: *Entities should consider whether a computer program performs decision-shaping reasoning activities — selecting, filtering, ranking, scoring, summarising, characterising, or recommending — that materially shape the information or analysis available at the terminal decision point. Where the entity cannot explain those contributory decision points without relying on the system's own output, the program is substantially and directly related to the decision. Entities should also disclose: whether the system has a default position that has been examined and authorised; and what information affected persons can access about the decision pathway and available review or complaint pathways.*

▸ Current APP 1 guidance / Good practice

Question 2

Yes — GPTea is substantially and directly related to making decisions

Would NACIA's use of GPTea be considered substantially and directly related to making decisions?

Direct answer: Yes.

▸ Current APP 1 guidance

Question 3

Why AI summaries and recommendations substantially assist decisions even with human sign-off

Reasons for the above answer.

Supports lifecycle Stage 4. Develops Recommendation 2.

GPTea performs every step of the decision pathway before the human sees anything. Before NACIA staff see anything, GPTea has performed: which fields are processed; what is surfaced versus filtered; how information is described and framed; how factors are weighted; what eligibility recommendation is offered; how confident that recommendation is stated to be. The substantive analytical work is complete before human review begins.

Applying the explanation test. Can NACIA staff, without GPTea's output, explain why this information was surfaced, why other information was filtered, and how competing factors were weighted? GPTea is performing strategic and ethical judgement on complex, ambiguous information — free-text candidate profiles. This is where the gap between what the system did and what the human can independently verify is largest.[5] The answer is almost certainly no. GPTea is therefore substantially facilitating the decision regardless of the formal oversight structure.

Human oversight without epistemic access is ceremony. A human reviewer is not enough if they only see an AI-generated summary and cannot inspect how it was produced. Aviation provides a direct analogy: when navigation was fully automated, pilots who had delegated navigation to the system lost the independent mental model needed when the

automation failed.[6] NACIA staff signing off on GPTea's recommendation places a human signature over an unreconstructable process.

Multiple reviewers evaluating the same AI-generated summary do not constitute meaningful diversity of oversight.[7] Three NACIA staff reviewing the same GPTea output are each evaluating the same pre-constructed frame. Real oversight requires that reviewers can independently access the underlying information.

Recommended guidance wording: *Human involvement is more likely to be meaningful where the reviewer has access to enough information to understand the basis of the ADM output, question it, override it, and record reasons for doing so — not merely to accept or reject an AI-generated recommendation or summary.*

▸ **Current APP 1 guidance**

Question 4

Factors for assessing significant effect — including practical ability to respond

What factors increase the likelihood that a decision could affect an individual's rights or interests?

Supports lifecycle Stage 5. Develops Recommendation 4.

Direct answer: The listed factors measure the impact of a decision. They do not measure the affected person's practical ability to respond to it. Both dimensions are necessary.

ADM type	Minimum expectation	Enhanced expectation
Simple rule-based, low impact	Basic disclosure of rules applied, annual review	—
Scoring and ranking systems	Explanation of factors, weights, and thresholds. Description of reviewer access to that information	More frequent review cadence, documented update triggers
Generative AI or complex recommendation systems	Step-by-step explanation of decision pathway. Reviewer information access disclosure. Vendor inspection rights in procurement	Written disclosure update record. Escalation pathway. Failure response disclosure
High-impact ADM — essential services, employment,	All of the above	Governance scorecard self-assessment. Regular reconciliation of disclosed versus actual practice. Governance function separation

ADM type	Minimum expectation	Enhanced expectation
health, credit, housing, insurance		

Four additional factors — practical ability to respond:

- **Can the affected person see the reasoning?** Inability to see why a decision was made magnifies its effect — the person cannot evaluate whether it was sound.
- Can the affected person access sufficient information about the decision pathway to understand what happened and access available review or complaint pathways?
- **Can the decision be reversed?** An irreversible low-impact decision can be more significant than a reversible high-impact one.
- **Does the effect compound across a population?** Individually modest decisions that accumulate into systemic effects at scale.

Reframing vulnerability. People are functionally vulnerable when they cannot understand that a decision has been made about them, understand the reasoning, evaluate whether it is correct, or take action in response. This applies across every demographic category and is determined by the gap between the person's capacity and the system's opacity — not the person's characteristics alone.

> **Recommended guidance wording:** *In assessing whether a decision could significantly affect rights or interests, entities should consider not only service type, financial outcome, and data sensitivity, but also: whether the affected person can see the reasoning; whether the person can access sufficient information about the decision pathway and available review or complaint pathways; whether the decision is reversible; and whether individual effects compound into population-level outcomes.*

▸ **Current APP 1 guidance**

Question 5

Vulnerability is a relationship between the person and the system, not only a demographic category

What classes of person are vulnerable for the purposes of considering a decision's expected effect?

Direct answer: Standard demographic categories are necessary but not sufficient. Guidance should also define vulnerability by the gap between the affected person's practical ability to understand and respond and the system's opacity.

Standard categories warranting heightened attention: children, elderly persons, persons with cognitive impairment, persons in financial hardship, persons with limited English proficiency, persons in institutional or custodial settings, and persons experiencing family violence.

Structural definition. People are functionally vulnerable when they cannot understand that a decision has been made about them, understand the reasoning behind it, evaluate

whether it is correct, or take action in response. The test is the gap between what the person can do and what the system has done. Research on AI use in professional contexts finds that intensive reliance on AI outputs reduces the depth of independent reasoning over time.[6] Where this describes the entity's own human review process, the oversight is nominal regardless of the reviewer's formal authority.

Recommended guidance wording: *In assessing vulnerability, entities should consider not only demographic characteristics but also whether the affected person has a practical ability to understand the decision, identify the reasoning behind it, and seek review — including where the system's decision pathway is not accessible or explainable to the person affected.*

▸ **Current APP 1 guidance**

Question 6

Significant service means high cost of denial, not category label alone

What constitutes a significant service or support?

Direct answer: A service is significant when the practical consequences of denial are high-friction, high-cost, or have no accessible alternative. Categorical lists produce gaps when new service types emerge.

Rankings: access to essential banking and credit services (rank 1); access to essential utilities (rank 1); financial assistance (rank 2); physical assistance and support (rank 2); access to educational services (rank 3); access to telecommunications (rank 3); in-home services (rank 3).

A categorical definition produces gaps whenever new service types emerge. But even where services are already correctly classified as significant, the definition alone does not guarantee ongoing disclosure. Without assigned ownership and review mechanisms, a disclosure made accurately at the point of drafting will drift from actual practice as systems evolve — and services correctly classified as significant will become effectively undisclosed through administrative inaction rather than legal ambiguity.

The OAIC's October 2025 desktop review finding[4] — 17 per cent compliance among agencies providing clearly significant services — demonstrates exactly this failure mode. The problem was not definitional; those agencies were providing services that clearly met any reasonable significance test. The problem was the absence of maintenance mechanisms. Categorical definitions must therefore be accompanied by monitoring, review cadence, and correction obligations.

Recommended guidance wording: *A service or support is likely to be significant where denial or limitation of access would have serious practical consequences for the affected person, including where alternatives are difficult to access, costly, or unavailable.*

▸ **Current APP 1 guidance**

Question 7

Existing frameworks do not address how governance survives machine-speed ADM or AI-generated local rules

Are there other legal frameworks or policies the OAIC should consider?

Direct answer: Yes. Three Australian frameworks are directly relevant. The OAIC should also address two structural gaps no existing framework covers.

Commonwealth Ombudsman ADM Better Practice Guide (March 2025).[1] The most operationally relevant existing Australian framework. It addresses intermediate automated steps, third-party ADM management, and human oversight quality. The OAIC's guidance should build on this.

Western Australia PRIS Act 2024, IPP 10.[2] Commencing 1 July 2026. Establishes materially stronger protections including contestability, right to request information, notification, and human intervention. The most relevant domestic benchmark for future reform.

Robodebt — accountability architecture failure, not transparency failure. The Royal Commission found that Robodebt failed not because the system was unknown but because there was no structure for tracing why specific decisions were made about specific individuals, no independent compliance checking, and no separation of the functions that designed, operated, and reviewed the system.[3] Building AI systems on weak governance infrastructure reproduces failure at scale — it does not remove it.

Operationalisation gap — a consistent absence across all existing frameworks. The frameworks above, and those listed in the OAIC's Issues Paper, share a common characteristic: they identify obligations, principles, and risk categories without providing the operational architecture for meeting them. They define what governed ADM should achieve — transparency, meaningful oversight, defensible decisions — but do not specify how a decision pathway should be designed, how governance functions should be structured within it, or how the lifecycle of rules governing that reasoning should be managed over time. Entities seeking to implement these principles in practice — separating governance functions, designing defensible decision pathways, maintaining disclosures against actual system behaviour — currently lack standardised operational tools that sit beneath the legal framework layer. Appendix B describes frameworks that can assist with this implementation challenge, covering decision pathway design, task-level governance risk assessment, and AI governance architecture.

The AI-generated local rules problem — a gap no existing framework addresses. In practice across organisations deploying AI for record management, information curation, archiving, and customer data processing, AI systems generate their own operational rules to handle novel cases, approve those rules without independent review, execute decisions based on those rules, and report on their own compliance. The same system makes rules, decides on rules, and follows rules. The accumulated institutional design wisdom underlying separation of powers, corporate governance, and professional standards all reached the same conclusion: these functions must be held by structurally separate actors.[8] The guidance should acknowledge this gap.

The disclosure maintenance capacity problem. The APP 1 obligation is well-drafted, but drafting quality at the point of creation is not sufficient. The binding constraint over time is whether the entity's disclosure review and maintenance process can keep pace with ADM system evolution. The guidance should specify how frequently entities should review disclosures, what changes trigger mandatory updates, and how the OAIC itself will monitor whether disclosures remain accurate.

The AI bias visibility gap. None of the frameworks listed by the OAIC addresses bias in AI systems embedded in learned behaviour rather than source code. Not discoverable by code audit and not addressable by employment contract. Guidance should acknowledge this gap.[5]

▸ **Current guidance / Future reform — mixed**

Question 8

Differential pricing: significance depends on challenge pathway, not magnitude

Does the inflation of prices for a consumer product significantly affect the interests of an individual such as Marco?

Supports lifecycle Stage 5. Develops Recommendation 4.

Direct answer: Yes. The harm is not the size of the price differential — it is that Marco has no pathway to discover the differential exists, understand its basis, or challenge the reasoning behind it.

A five per cent differential where the affected person can see the reasoning and contest it is less harmful than a twenty per cent differential where the basis is invisible and unchallengeable. Any magnitude threshold measures the output of the problem rather than its source. It also creates a ceiling that sophisticated operators will design to stay beneath.

The appropriate mechanism is a pathway for the affected person to challenge the differential and receive a substantive response. This is developed in the forward-looking section.

> **Recommended guidance wording:** *In assessing whether differential pricing significantly affects an individual's interests, entities should consider whether the affected person has a practical pathway to discover, understand, and challenge the differential — not only its dollar magnitude.*

▸ **Current APP 1 guidance**

Question 9

Service significance matters more than dollar quantum

How important are: (a) the dollar amount; (b) the significance of the product or service?

Direct answer: Service significance (5/5) matters more than dollar amount (3/5). A small differential on an essential service creates a barrier to access. A large differential on a discretionary purchase does not. Service significance is the more reliable indicator of whether the differential affects the person's substantive position.

▸ **Current APP 1 guidance**

Question 10

Targeted advertising, non-promotion, and hidden ADM effects

Should it be considered a decision that Rafqa did not receive the Manscon job advertisement?

Supports lifecycle Stage 1. Develops Recommendation 3.

Direct answer: Yes. Failing to promote a job advertisement to Rafqa is a failure to act that materially limited her employment opportunities, within the plain meaning of APP 1.9(a).[9]

The default was never classified as a governed contributory decision point. Opaque Window's algorithm does explicitly and at scale what human recruitment processes have always done implicitly: it selects for the existing pattern. The default — male engineers receive the opportunity — was never examined, authorised, or documented as a contributory decision point subject to governance. It was inherited from existing workforce composition and encoded into the optimisation logic. Under the Commonwealth Ombudsman's framework, [1] this is an intermediate automated step that materially shapes who can access an employment opportunity and should be identified, documented, and disclosed.

The algorithm did not create this bias. What ADM has done is make the mechanism detectable for the first time. The obligation is an opportunity to create the visibility that enables challenge.

Perspective matters. From Opaque Window's perspective, the targeting behaviour is a default operational process. From Rafqa's perspective, it is an employment foreclosure. Assessment of whether an ADM contributory decision point constitutes a governed activity should not be made solely from the entity's operational perspective — the effect on the affected person is equally relevant.

Disclosure surfaces the mechanism. Remedy requires a challenge right. Rafqa does not know she did not receive the advertisement. Even if she knew, the current obligation provides no mechanism for her to challenge the targeting reasoning. Disclosure is diagnosis. Remedy requires a protected right to challenge the decision pathway — developed in the forward-looking section.

Recommended guidance wording: *Entities should take reasonable steps to identify ADM activity that affects individuals without being internally recognised as 'decision-making,' including automated targeting, filtering, prioritisation, and non-promotion of information relevant to decisions. The default position of any such system should be examined, authorised, and disclosed.*

▸ Current APP 1 guidance / Good practice

Question 11

Third-party ADM and procurement access

Are there scenarios relevant to the meaning of 'arranged for' that require guidance?

Supports lifecycle Stage 6. Develops Recommendation 5.

Direct answer: Yes. Two scenarios require specific guidance: procurement without inspection capability, and AI systems that generate their own operational rules.

Procurement without inspection capability. In practice, most procurement of AI-powered services does not include contractual rights to inspect the system's decision pathway, training data, or decision logic. The arranging entity carries the disclosure obligation but may have no practical ability to discover what it is obligated to disclose.[1]

Entities should consider whether contracts for third-party ADM include rights to obtain sufficient information to comply with the disclosure obligation. Where such rights are not obtainable, entities should assess whether the arrangement is consistent with their disclosure obligations.

AI bias not visible through code review. Bias in conventional software is in the source code — detectable by audit. Bias in AI systems may be in the learned behaviour — not discoverable by code review and not addressable by employment contract or professional standards.[5]

AI systems generating their own operational rules — a structural governance gap. When AI systems generate operational rules to handle novel cases, approve those rules without independent review, and execute decisions based on them, the functions of rule-making, rule-deciding, and rule-following are collapsed into a single system.

The accumulated institutional design wisdom that produced separation of powers, corporate governance, and professional standards applies: these functions must be held by structurally separate actors.[8]

Guidance should acknowledge this gap and signal that future legislative development will need to address it. See Appendix D for the case for mandatory structural separation of governance functions as a future legislative requirement.

Recommended guidance wording: *Where an entity arranges for third-party ADM to make or assist decisions, the entity should consider whether it can obtain sufficient information about the system's decision pathway, use of personal information, and decision types to support accurate disclosure. Contracts should include inspection rights for this purpose.*

▸ Good practice / High-risk ADM

Question 12

How entities should design and maintain ADM disclosures

What should entities disclose to provide effective and meaningful disclosures about ADM?

Supports lifecycle Stages 7–8. Develops Recommendations 6–7.

Direct answer: ADM disclosures should be tiered to the complexity of the reasoning involved, measured against a governance scorecard, and maintained as living documents against actual practice.

The OAIC's five proposed criteria[10] — plain language, structured for follow-up, appropriately tailored, logically organised, framed to allow challenge — are the right foundation. We recommend six extensions:

Extension 1 — Treat the disclosure as a maintained document, not a publication

An ADM disclosure is not a static fact — it is a document that must be maintained against actual practice. Entities should disclose: the review cadence; who is responsible for verifying the disclosure reflects current ADM use; and how gaps are identified and corrected. The OAIC's 17 per cent compliance finding[4] shows what happens when disclosure is treated as a one-time publication.

Extension 2 — Keep a written disclosure update record

Entities should maintain a record of what was disclosed, when it was updated, who validated the update, and what triggered the change. The disclosure is the working document. The update record is the audit trail.[7]

Extension 3 — Tier the disclosure to the complexity of the reasoning involved

- **Rule-based execution:** state the rule applied. The logic is fixed and reproducible.
- **Scoring and ranking:** explain factors used, how they were weighted, and what thresholds applied. Describe whether the reviewer has access to this information.
- **Strategic and ethical judgement — recommendation and contextual assessment:** explain what information was selected, how it was characterised, what was filtered, and how competing factors were resolved. Human oversight at this level must be substantive, not nominal.

Extension 4 — Disclose whether governance functions are separated

Entities may consider disclosing whether compliance checking, risk identification, and access control are held by separate actors with separate authority, or collapsed into a single path.

Extension 5 — Disclose what the human reviewer can actually access and independently evaluate

Disclosure of human oversight should include what the reviewer independently accesses, decides, and exercises authority over. Disclosure stating only that humans make the final decision, without specifying what they can independently evaluate, does not enable the affected person to assess whether oversight is substantive.

Extension 6 — Disclose what happens when the ADM system fails or produces erroneous output

Entities should disclose what fallback mechanisms exist, who is notified, how failures are investigated, and how contested decisions are escalated and resolved.

Formal notice versus actual understanding — both required separately. Formally notifying affected persons that ADM is used is not the same as ensuring they can understand and act on the disclosure. The test for meaningful disclosure is whether the affected person can understand and use the information, not only whether it was published.

Governance scorecard — voluntary self-assessment. Entities may wish to document and consider disclosing a governance self-assessment — see Appendix A3 for a voluntary template. This can provide a useful basis for internal governance review.

> **Draft disclosure template — entities may adopt directly:** We use automated systems to assist with [type of decision]. These systems may use [kinds of personal information] to [score / rank / recommend / triage / summarise / escalate]. A human reviewer [does / does not] make the final decision. The reviewer [can / cannot] access [the relevant underlying information / the basis of the automated output / the factors used] and can [record reasons for departing from / override] the automated recommendation. Our decision pathway for this decision type is reconstructable at Level [0 / 1 / 2 / 3]. This disclosure was last reviewed on [date] by [role/team]. You can contact us at [contact point] to ask for more information about how this applies to you, to request an explanation of the reasoning in a specific decision, or to seek review where available.

> **Recommended guidance wording:** *Entities should review ADM disclosures periodically and when there are material changes to ADM systems, decision types, personal information used, or human review arrangements. Entities should assign a responsible person for disclosure accuracy, keep a written record of updates, and may wish to document and consider disclosing a governance self-assessment — see Appendix A3 for a voluntary template.*

▸ **Current APP 1 guidance / Good practice**

Summary of recommended guidance wording

This table consolidates all recommended guidance wording from Questions 1–12 to support efficient extraction of actionable recommendations.

Rec	Stage	Recommended guidance wording
Rec 1	2–3	Entities should consider whether a computer program materially shapes a decision by selecting, filtering, ranking, scoring, summarising, framing, or recommending information a human decision-maker relies on, including steps before the final decision point. Where the entity cannot explain those contributory decision points without relying on the system's own output, the program is substantially and directly related to the decision. Entities should also disclose: whether the system has an undisclosed default position; and what information affected persons can access about the decision pathway and available review or complaint pathways.
Rec 2	4	Human involvement is more likely to be meaningful where the reviewer has access to enough information to understand the basis of the ADM output, question it, override it, and record reasons for doing so — not merely to accept or reject an AI-generated recommendation.
Rec 3	1	Entities should take reasonable steps to identify ADM activity that affects individuals without being internally recognised as decision-making, including automated targeting, filtering, prioritisation, and non-promotion of information relevant to decisions.
Rec 4	5	In assessing significant effect, entities should consider not only service type and financial outcome but also: whether the affected person can see the reasoning; whether the person can access information about the decision pathway and available review or complaint pathways; whether the decision is reversible; and whether individual effects compound into population-level outcomes.
Rec 5	6	Where an entity arranges for third-party ADM, it should consider whether it can obtain sufficient information about the system's decision pathway, use of personal information, and decision types to support accurate disclosure.
Rec 6	7	ADM disclosures should explain: the kinds of decisions involved; personal information used; whether ADM makes or assists the decision; the human reviewer's role and what they can independently access and evaluate; and pathways for further information or review. Disclosures should be tiered to the complexity of the reasoning involved.
Rec 7	8	Entities should review ADM disclosures periodically and when there are material changes to ADM systems, decision types, personal information used, or human review arrangements. Entities should assign a responsible person for disclosure accuracy, keep a written update record, and may wish to document and consider disclosing a governance self-assessment — see Appendix A3 for a voluntary template.

Closing

The ADM transparency obligation is a necessary and welcome first step. The practical lifecycle proposed in this submission — define the terminal decision, map the decision pathway, test defensibility, match tool to contributory decision point, assess governance risk, implement separated governance functions, test and iterate, and verify audit capability — is directly actionable in guidance and directly responsive to the obligation's purpose.

Disclosure without the ability to reproduce the decision pathway is information without verifiability. Disclosure without reasoning contestability is information without recourse. Disclosure not maintained against actual practice is diagnosis of a problem that immediately begins to drift from the reality it describes.

Values without operational architecture do not produce accountability. The obligation creates the values. The guidance must build the structure.

Three issues exceed what guidance can deliver and are addressed in Appendix D: a protected right for affected persons to challenge specific contributory decision points in the decision pathway rather than only the terminal decision; mandatory structural separation of governance functions for high-risk ADM; and periodic mandatory reconciliation of disclosed and actual ADM practice. These require future legislative development and the OAIC is well placed to signal their necessity.

Further engagement:
I would welcome the opportunity to discuss this submission further or provide any clarification that may assist the OAIC.

Preparation note:
This submission was developed with the assistance of AI-based drafting and editing tools. The analysis, judgments, and conclusions are my own, and any errors are mine alone.

Submitted by:

[redacted]
Independent researcher

15 June 2026

Personal research.

Appendix A — ADM Transparency Policy Templates and Governance Tools

This appendix has two parts. Part A0 contains suggested OAIC guidance text — written in the register the OAIC uses when addressing entities, ready to be lifted directly into guidance

documentation. Parts A1–A4 contain entity implementation templates — documents and tools that entities use to comply with the guidance in Part A0.

Part A0 — Suggested OAIC guidance text

The following seven sections provide suggested guidance text for the OAIC's guidance document on the ADM transparency obligation. Each section is written as guidance addressed to APP entities. Cross-references to the relevant question responses in this submission are provided after each example.

A0.1 — Assessing the substantial and direct relationship (Recommendation 1)

Purpose: Operationalise APP 1.7(a) so entities can assess whether a computer program is substantially and directly related to making a decision — covering the full decision pathway, not only the terminal decision point.

Guidance text

Decision pathway assessment. Entities should consider whether a computer program materially shapes a decision by selecting, filtering, ranking, scoring, summarising, or framing information that a human decision-maker relies on — including where those steps occur before the final decision point. The test is not whether the system makes the final decision, but whether the entity can explain those contributory decision points without relying on the system's own output or description of itself. Where the entity cannot independently explain the contributory decision points, the program is substantially and directly related to the decision.

Three additional disclosure properties. Entities should assess and disclose:

- Decision reconstructability: at what level can the contributory decision points be reproduced for a specific decision after it was made? Level 0 = unrecoverable; Level 1 = category-level only; Level 2 = step-level for specific decisions; Level 3 = full independent third-party reconstruction. Entities operating high-risk ADM systems are encouraged to assess their decision pathway reconstructability level and may consider disclosing it as a governance good practice — see Appendix A3 for a voluntary template.
- **Reasoning neutrality:** does the system have a default position — an outcome it tends toward in the absence of information pointing another way? Both approve-by-default and reject-by-default configurations are bias positions, not neutral ones. The default position should be examined, authorised, and documented.

> **Example — Recruitment screening**
> *A recruitment platform scores and ranks candidates before presenting a shortlist to human reviewers. The reviewers can override any recommendation but only see candidates the system selected. The entity should assess whether it can explain, without relying on the platform's own description, why specific candidates were scored above the threshold and others were not. If it cannot, the platform is substantially and directly related to the hiring decision regardless of human sign-off on individual candidates.*
>
> → *See also: Questions 1, 2, and 3*

Proportionality: For simple rule-based systems, explaining the contributory decision points means stating the rules applied. For scoring and ranking systems, entities should explain the factors, weights, and thresholds. For generative AI systems processing

unstructured information, entities should map the full decision pathway and verify that each step is independently explainable. Reasoning contestability — a protected right for affected persons to challenge specific contributory decision points — is addressed as a future legislative development in Appendix D and is not a current APP 1 guidance requirement.

A0.2 — Meaningful human oversight (Recommendation 2)

Purpose: Operationalise the human oversight question under APP 1.7 so entities can assess whether human involvement constitutes genuine governance rather than nominal sign-off.

Guidance text

The meaningful oversight test. Human involvement is more likely to be meaningful where the reviewer has access to enough information to understand the basis of the ADM output, question it, override it, and record reasons for doing so. An entity should not treat human oversight as meaningful merely because a human makes the final decision. The relevant questions are: what information can the reviewer independently access; what can they independently evaluate without reference to the ADM output; and what authority can they independently exercise?

Disclosure of human reviewer role. Entities should disclose the nature of the human reviewer's role, including: what information the reviewer can independently access; what they can independently evaluate without reference to the ADM output; what authority they can independently exercise; and how they record reasons for departing from the ADM recommendation. Where the reviewer's evaluation depends entirely on the ADM system's own characterisation of the information, entities should disclose this limitation.

Multiple reviewers. Multiple human reviewers evaluating the same AI-generated output do not constitute meaningful diversity of oversight. Real oversight diversity requires that reviewers can independently access the underlying information, not just react to the AI's presentation of it.

Example — Aged-care eligibility assessment (GPTea scenario)

> *An aged-care eligibility officer who reviews a generative AI summary of a candidate profile and approves or rejects the AI's recommendation is not exercising meaningful oversight if they cannot independently access the underlying profile and identify what information the AI selected, filtered, or weighted. The entity should disclose this limitation rather than characterising the process as human-reviewed. If NACIA staff cannot independently verify GPTea's reasoning, the oversight does not make the decision explainable — it places a human signature over an unreconstructable process.*

→ *See also: Questions 2 and 3*

Proportionality: *For simple rule-based systems, human oversight may be meaningful at a category level — the reviewer understands and can verify the rules applied. For systems performing judgement-level reasoning on unstructured information, the reviewer must have access to the underlying information, not just the AI's output, to exercise meaningful oversight.*

A0.3 — Identifying hidden ADM use (Recommendation 3)

Purpose: Help entities identify ADM activity that affects individuals without being internally classified as decision-making, consistent with APP 1.9(a) which provides that making a decision includes refusing or failing to make a decision.

Guidance text

ADM inventory. Entities should take reasonable steps to identify all systems that make, recommend, rank, score, triage, filter, summarise, escalate, or withhold information relevant to decisions that could significantly affect individual rights or interests. This includes upstream systems that operate before the formal decision point, systems that determine who receives information rather than what decision is made, and systems that determine priority or sequencing of human review.

Inventory maintenance. Entities should maintain an ADM inventory that is reviewed periodically and updated when new systems are deployed or existing systems are materially changed. The inventory should include systems not internally classified as decision-making systems but whose outputs materially shape decisions.

Default position disclosure. The default position of any ADM system — the outcome the system produces in the absence of information pointing another way — is an important governance fact. Entities should consider examining, authorising, and documenting the system's default position as part of the inventory process. A system whose default position has never been examined may be systematically biasing decisions without the entity's awareness.

> **Example — Targeted job advertising (Rafqa scenario)**
> *A job advertising platform that uses an algorithm to select which graduate profiles receive a particular job advertisement is performing an ADM function that affects employment opportunities, even though the platform does not characterise itself as making a decision. The entity arranging for the platform should identify this as a governed ADM activity. The algorithm's default — promoting to candidates who match the existing workforce profile — was never examined, authorised, or documented as a contributory decision point subject to governance. It was inherited from existing workforce composition and encoded into the optimisation logic.*
>
> → *See also: Question 10*
>
> **Proportionality:** *All entities using ADM should maintain an inventory. The depth of documentation proportionate to the significance of the decisions affected: simple inventory listing for low-impact systems; full decision pathway documentation for high-impact systems.*

A0.4 — Assessing significant effect on rights or interests (Recommendation 4)

Purpose: Help entities apply the significant effect test under APP 1.7(b), including factors beyond service category and financial outcome.

Guidance text

Factors for assessment. In assessing whether a decision could reasonably be expected to significantly affect an individual's rights or interests, entities should consider:

- **Service importance:** a service or support is likely significant where the practical consequences of denial are high-friction, high-cost, or where alternatives are difficult to access or unavailable.

- **Financial effect:** the magnitude of financial impact on the affected person, including differential pricing that creates barriers to access.
- **Reversibility:** whether the decision can be reversed and at what cost to the affected person in time, money, and practical resources.
- **Vulnerability:** vulnerability should be assessed not only by demographic characteristics but also by whether the affected person has a practical ability to understand the decision and seek review — including where the system's reasoning is not accessible or explainable.
- Practical ability to understand and challenge: whether the person has access to available review or complaint pathways and sufficient information about the decision pathway to understand what happened.
- **Population-scale compounding:** where individually modest ADM decisions accumulate into significant effects across a population, the cumulative population-level effect should be considered even where the individual-level effect appears minor.

Proportionality of disclosure and governance. The level of disclosure and governance required should be proportionate to the combination of how explainable the AI's reasoning is given the information it processes, and how significant the decision's effect is. Higher risk on either dimension requires stronger governance.

> **Example — Differential pricing (Daintree scenario)**
> *An e-commerce platform that charges different prices to different customers based on postcode data may significantly affect the interests of customers in lower-income areas. The harm is not only the price differential itself — it is that the affected person has no pathway to discover the differential exists, understand its basis, or challenge the reasoning behind it. A five per cent differential where the affected person can see the pricing rationale and contest it is less harmful than a twenty per cent differential where the basis is invisible and unchallengeable.*
>
> → *See also: Questions 8 and 9*
>
> **Proportionality:** *High-impact ADM affecting essential services, employment, health, credit, housing, or insurance should apply all factors with enhanced disclosure requirements. Lower-impact rule-based ADM may apply a simplified assessment covering service importance and financial effect.*

A0.5 — Third-party ADM and the meaning of 'arranged for' (Recommendation 5)

Purpose: Clarify what 'arranged for' means in APP 1.7(a) and what entities must do to discharge their disclosure obligations where a third party operates the ADM system.

Guidance text

Disclosure obligation follows the arranging entity. An entity that arranges for a third-party system to make or assist decisions is subject to the APP 1 disclosure obligation in relation to that system. The obligation is not discharged by the fact that the entity does not itself operate the system.

Procurement requirements. Entities should consider whether contracts for third-party ADM include rights to obtain sufficient information about the system's use of personal information, decision types, and decision pathway to support accurate disclosure. This includes rights to

obtain information about the contributory decision points the system performs, not only the outputs it produces.

Where inspection rights are unavailable. Where such inspection rights cannot be obtained from the vendor and the system will make or assist decisions that significantly affect individual rights or interests, entities should assess whether the arrangement is consistent with their disclosure obligations.

AI-generated local rules. Entities should assess whether any third-party AI system generates its own operational rules to handle novel cases. Where an AI system is generating, applying, and reviewing its own operational rules without human authorisation for each new rule, entities should consider what governance arrangements are in place for this characteristic of the system — see Appendix D for the case for mandatory structural separation as a future legislative requirement.

> **Example — Third-party AI for document processing**
> *An entity procures an AI system for record management that is configured to generate category rules for new document types it encounters. The entity has arranged for a system that will make decisions about document handling using rules the system itself created — without human authorisation of those rules. The entity cannot discharge its disclosure obligation without inspection rights that reveal what rules the system has generated, what decisions those rules have affected, and whether those rules align with the entity's original instructions.*
>
> → *See also: Question 11*
>
> **Proportionality:** *For simple third-party software implementing fixed rules, inspection of the rules applied is generally sufficient. For AI systems that adapt, learn, or generate operational rules, entities should require vendor access to decision pathway documentation as a condition of procurement.*

A0.6 — Designing meaningful ADM disclosures (Recommendation 6)

Purpose: Specify what an ADM privacy policy disclosure should contain to be effective and meaningful under APP 1.8.

Guidance text

Required disclosure content. An entity's privacy policy should, in relation to each ADM system subject to APP 1.7, disclose:

- The kinds of decisions the system makes or assists.
- The kinds of personal information used in the system's operation.
- Whether the system makes the decision or assists a human decision-maker.
- **Where a human decision-maker is involved:** what information the reviewer can independently access and evaluate; what authority they can independently exercise; and how they record reasons for departure from the ADM recommendation.
- Available pathways for the affected person to obtain further information or seek review where another framework provides for it.
- What happens when the system produces an erroneous output or fails — what fallback mechanisms exist, who is notified, and how the failure is investigated.

Entities operating high-risk ADM systems may also wish to disclose their self-assessed decision pathway reconstructability level as an indicator of how traceable the decision pathway is for a specific decision.

Tiered to reasoning complexity. The level of detail in the disclosure should be proportionate to the complexity of the reasoning involved:

- **Rule-based systems:** state the rules applied. The logic is fixed and reproducible from the rule statement.
- **Scoring and ranking systems:** explain what factors are used, how they are weighted, and what thresholds apply. Describe whether the reviewer has access to this information.
- **Judgement and recommendation systems:** explain what information was selected, how it was characterised, what was filtered, and how competing factors were resolved. Disclosure of the reviewer's independent access to underlying information is required.

Plain language. Disclosures should be written in plain language and structured around what the affected person needs to understand and respond to the decision. Technical descriptions of system architecture are not a substitute for plain-language explanation of how the decision was made.

> **Example — Insurance claim triage**
> *An insurer uses a model to classify claims as low, medium, or high risk, which affects how quickly the claim is reviewed and what evidence is requested. The disclosure should explain: that ADM is used to assist claim triage; what kinds of information are used in the triage assessment; what factors and thresholds determine the risk category; whether the claims officer can see those factors and override the triage category; and how a claimant can request review of the triage classification. Stating only that 'automated systems are used to assist claim handling' does not meet this standard.*
>
> *→ See also: Questions 1 and 12*

Proportionality: *Simple rule-based disclosures can be short and specific. Disclosures for complex AI recommendation systems should describe the decision pathway at a level sufficient for the affected person to identify which step they wish to contest.*

A0.7 — Maintaining ADM disclosures over time (Recommendation 7)

Purpose: Establish what entities must do to keep ADM disclosures accurate over time, addressing the compliance gap identified in the OAIC's October 2025 desktop review.

Guidance text

Disclosure as a maintained document. ADM disclosures are not one-off publications — they are governance documents that must be maintained against actual practice. The gap between what an ADM disclosure states and what a system actually does compounds silently as systems are updated, models are retrained, and processes change. An entity that published an accurate disclosure at the time of first compliance but has not reviewed it since may be significantly non-compliant without being aware of it.

Maintenance requirements. Entities should:

- Assign a named person or role responsible for ADM disclosure accuracy.
- Establish a review cadence specifying how frequently disclosures are reviewed against actual ADM system behaviour.
- Identify the changes to ADM systems that trigger mandatory disclosure updates.

- Maintain a written record of disclosure updates — what changed, when, who validated the change, and what triggered it.

Update triggers. Disclosures should be updated promptly when there is a material change to: the ADM system or its underlying model; the kinds of decisions it makes or assists; the kinds of personal information it uses; the role of any human reviewer; or available review pathways.

Periodic reconciliation. Entities should conduct a periodic reconciliation comparing the disclosure against actual system behaviour at a point in time, and document the result. Where a gap is identified, it should be corrected and the correction recorded. Entities may wish to document their governance self-assessment and consider including a summary in their ADM privacy policy disclosure — see Appendix A3 for a voluntary template.

> **Example — Evolving credit scoring model**
> *A bank discloses that its credit scoring system uses income, employment history, and credit history as factors. Over the following 18 months, the model is retrained twice on updated data, and the weighting of employment history is substantially reduced while geographic location indicators are introduced. The disclosure continues to state the original three factors. The gap between the disclosure and the actual system has grown materially — but because no review was triggered and no owner was assigned, neither the bank nor the affected person is aware of the discrepancy.*
> *Entities should treat model retraining and parameter changes as potential triggers for disclosure review, not only changes to the system's source code or stated purpose.*

→ See also: Question 12

Proportionality: *All entities subject to APP 1.7 should implement at minimum an annual review cadence and assign a named disclosure owner. Entities using AI systems that adapt or retrain continuously should implement more frequent review cycles and automated change-detection mechanisms.*

Parts A1–A4 — Entity implementation templates

Parts A1–A4 from the main submission document (ADM-Submission-v6.docx) follow here in the final assembled document. They contain: A1 — ADM Transparency Policy (eight-stage entity policy statement); A2 — ADM Disclosure Template; A3 — Governance Scorecard Template; A4 — ADM Inventory and Governance Checklist.

These templates are addressed to entities and use first-person or fill-in-bracket language. They are the implementation tools entities use to comply with the guidance text in Part A0 above.

A1 — ADM Transparency Policy

This is a voluntary governance framework for entities. It covers the full ADM transparency lifecycle and exceeds the minimum requirements of APP 1 in several respects. Entities operating high-risk ADM systems are encouraged to adopt this framework in full. Other entities may adopt the elements most relevant to their circumstances and risk profile.

Policy statement. [Organisation name] is committed to transparent use of automated decision-making systems. This policy establishes how we identify, assess, disclose, and

maintain information about automated systems that make or assist decisions that may significantly affect individuals.

Scope. This policy applies to all systems that score, rank, filter, recommend, summarise, triage, escalate, or frame information used in decisions that could significantly affect individual rights or interests — including systems operated by third-party vendors on our behalf.

Stage 1 — Define the terminal decision and defensibility standard

Before deploying any automated system, we document: the specific decision to be made; what a defensible version of that decision looks like; and what level of accuracy, fairness, and explainability is required. This standard is reviewed when the decision type or system changes materially.

Stage 2 — Map the decision pathway

For each ADM system, we document the full chain of information transformations and contributory decision points required to reach the terminal decision — including which steps are automated, which involve human judgment, and how intermediate conclusions influence subsequent steps.

Stage 3 — Test whether the decision pathway is defensible

Before deploying any tool, we assess whether the designed decision pathway can withstand scrutiny. We document this assessment and update it when the process changes materially.

Stage 4 — Match tool to contributory decision point

For each step in the decision pathway, we assess: the nature of the information involved (structured versus unstructured); the type of reasoning required (rule-based execution, adaptive process reasoning, or strategic and ethical judgement); and what tool is appropriate. We document our tool selection rationale and review it when systems change.

Stage 5 — Assess governance risk within each contributory decision point

For each step, we assess: the likelihood of error given the tool and information environment; the impact of that error on the terminal decision; and the level of oversight proportionate to that risk. We maintain this risk assessment and update it when systems change.

Stage 6 — Implement separated governance functions

Where ADM systems perform consequential contributory decision points, we separate the functions of compliance checking, risk identification, and access control. These functions are held by separate actors with separate authority and separate information access. We document our governance function structure and review it when systems change.

Stage 7 — Test and iterate

We test our governed decision pathway against the defensibility standard from Stage 1 before deployment and when systems change. Where risk exceeds the acceptable level, we revise the process, tool selection, or governance arrangements and retest.

Stage 8 — Verify audit capability and maintain transparency

We verify that intermediate contributory decision points for any specific decision are reconstructable after the fact. We disclose our governance scorecard self-assessment. We maintain our ADM disclosures against actual practice and keep a written update record.

A2 — ADM Disclosure Template

> How we use automated decision-making We use automated systems to assist with [type of decision, e.g. assessing eligibility for support services / ranking job applications / approving credit applications]. These systems may use [kinds of personal information, e.g. your financial history, employment records, application responses] to [score / rank / recommend / triage / summarise / escalate]. A human reviewer [does / does not] make the final decision. The reviewer [can / cannot] access [the relevant underlying information / the basis of the automated output / the factors used in scoring] and can [record reasons for departing from / override] the automated recommendation. Further information about the decision pathway — including the decision-shaping reasoning activities involved and the role of any automated system — is available on request at [contact point]. This disclosure was last reviewed on [date] by [role/team]. It will next be reviewed by [date] or when our systems change materially. If you would like more information about how this applies to you, to request an explanation of the reasoning used in a specific decision, or to seek review where available, please contact us at [contact point].

A3 — Governance Scorecard Template

This is a voluntary self-assessment tool. It is not a compliance requirement under APP 1. Entities operating complex or high-risk ADM systems are encouraged to complete this assessment and consider disclosing a summary alongside their ADM privacy policy disclosure.

Entities may complete this scorecard for each ADM system and consider disclosing a summary as part of their ADM privacy policy disclosure.

Governance measure	What it measures	Level (0–3)	Required level for this decision type	Gap and action
Decision reconstructability	Can the entity reproduce each intermediate contributory decision point for a specific decision without relying on the system's own output? Level 0 = unrecoverable; Level 1 = category-level only; Level 2 = step-level for			

Governance measure	What it measures	Level (0–3)	Required level for this decision type	Gap and action
	specific decisions; Level 3 = full independent third-party reconstruction			
Oversight effectiveness	Is the human reviewer performing independent evaluation — accessing the basis of the ADM output, questioning it, and recording reasons for departure — or rubber-stamping an AI-generated recommendation?			
Reasoning integrity	Is the AI system operating within its designed parameters, or has it drifted from what was assessed and approved? Measured by: consistency of outputs for equivalent inputs; detection rate for out-of-distribution inputs; frequency of outputs outside expected range			
Governance separation	Are compliance checking, risk identification, and access control held by separate actors with separate authority — or collapsed into a single path? Can the compliance function be disabled by the system being checked?			
Self-correction capacity	Can the system identify and correct its own errors before downstream harm compounds? Is the correction mechanism independent of the system that made the			

Governance measure	What it measures	Level (0–3)	Required level for this decision type	Gap and action
	error? Measured by: time between error and detection; proportion detected internally versus reported externally			
Disclosure accuracy gap	What is the measurable difference between what the ADM disclosure states and what the system actually does at this point in time? Assessed by periodic reconciliation between disclosed practice and observed system behaviour			
Audit trail completeness	For any specific decision, what proportion of intermediate contributory decision points are recoverable from the audit record? Scale: 0% = nothing recoverable; 50% = recommendation and major inputs; 100% = full chain reconstructable independently			

A4 — ADM Inventory and Governance Checklist

Stage	Check	Question	Status	Action
1	1	Have we defined the terminal decision and its defensibility standard?		
2	2	Have we mapped the full decision pathway including all contributory decision points?		
2	3	Have we identified all upstream automated steps		

Stage	Check	Question	Status	Action
		not internally classified as decision-making?		
3	4	Have we tested whether the decision pathway is defensible before tool selection?		
3	5	Have we applied the explanation test: can we explain contributory decision points without relying on the system's own output?		
3	6	Have we identified and documented each system's default position?		
4	7	Have we matched each tool to the contributory decision point based on information environment and reasoning type?		
5	8	Have we assessed governance risk at each contributory decision point?		
5	9	Have we considered whether affected persons have access to sufficient information about the decision pathway to understand what happened and access available review pathways? (Recommended good practice / High-risk ADM)		
6	10	For high-risk ADM: have we considered separating compliance checking, risk identification, and access control functions? (High-risk ADM best practice / Future reform)		
6	11	For third-party ADM: have we considered whether contracts include rights to obtain sufficient information about the system's decision-shaping reasoning activities to support accurate disclosure? (Good practice)		

Stage	Check	Question	Status	Action
7	12	Have we tested the governed decision pathway against the defensibility standard?		
8	13	Is there an assigned owner for disclosure accuracy?		
8	14	Is there a written disclosure update record?		
8	15	Is the governance scorecard completed and, if used, disclosed as a voluntary self-assessment?		

Appendix B — Governance frameworks: stage mapping and descriptions

Three open-source governance frameworks are referenced in this submission. This appendix first maps each framework to the lifecycle stages where it is operationally relevant, then provides a short description of each with key concepts in ADM-friendly language. All are published under CC BY 4.0. For attribution and enquiries: https://au.linkedin.com/in/kfkchau.

B1 — Framework mapping by lifecycle stage

The table below maps each lifecycle stage to the governance frameworks that are operationally relevant at that stage, with the key concepts and how they apply. Full framework descriptions are in B2–B4.

Stage	Framework	How it applies
Stage 1 — Define terminal decision and defensibility standard	Open Meta-Governance Standard (OMGS)	The terminal decision is a meta-rule — a rule about what outcome is acceptable. Requires explicit human authorisation before reasoning begins. The defensibility standard corresponds to the analysis stage of the OMGS proposal cycle, which must complete before implementation begins.
Stage 2 — Map the decision pathway	OMGS; Task Boundary Framework (TBF)	OMGS separates who governs from who is governed and helps identify ungoverned contributory decision points that affect

		people without being disclosed. TBF decomposes the decision pathway into single-step transformations — one input, one reasoning operation, one output — so explainability can be tested contributory decision point by contributory decision point.
Stage 3 — Test whether the decision pathway is defensible	OMGS	The proposal cycle — monitoring, analysis, and reporting — must complete before implementation begins. Defensibility testing therefore occurs before tool deployment.
Stage 4 — Match tool to contributory decision point	TBF	TBF matches tools to contributory decision points by reasoning tier, information structurality, and risk. What-tier reasoning can usually be handled by conventional software; How-tier reasoning may allow AI assistance; Why-tier strategic and ethical judgement requires human reasoning. Structured information supports more auditable AI use; unstructured information requires stronger human oversight.
Stage 5 — Assess governance risk within each contributory decision point	TBF	Risk scales with both reasoning tier and information complexity. Oversight must increase as those dimensions increase. Highest-tier reasoning on unstructured information produces the highest governance risk and therefore requires the strongest oversight.
Stage 6 — Implement separated governance functions	OMGS; Project White Collar (PJWC)	OMGS supports structural separation where AI generates and follows its own operational rules. PJWC separates compliance checking, risk identification, and access control as distinct governance functions. In

		complex ADM settings these map into an Accountability Committee structure.
Stage 7 — Test and iterate	PJWC; OMGS	PJWC emphasises fail-safe orchestration designed for sustained governed performance under component failure, not just speed under normal conditions. OMGS treats monitoring and analysis as continuous functions, with enforcement outputs feeding back into monitoring inputs.
Stage 8 — Verify audit capability and maintain transparency	OMGS; PJWC; all	OMGS distinguishes formal notice from actual understanding: disclosures must be both published and accessible/actionable. PJWC distinguishes immutable logs from editable knowledge bodies, supporting both audit trail integrity and maintained working disclosures. Together these support governance self-assessment alongside privacy policy disclosure.

B2 — Open Meta-Governance Standard (OMGS)

GitHub and pitch deck | © [redacted] 2025, CC BY 4.0

What it is. OMGS is a rules-of-rules standard — a framework for governing decision pathwayes and decisions that affect groups of people. It provides a portable, domain-agnostic lifecycle for how rules are proposed, adopted, implemented, communicated, enforced, revised, and retired. It applies wherever rule users, rule makers, and decision makers exist, across any sector, jurisdiction, or institutional setting.

Short context. Every information transformation in a decision process is itself a decision governed by rules. Those rules — whether explicit policy rules or implicit optimisation objectives in an AI system — need their own governance lifecycle. When AI systems are involved in making and following operational rules within a decision process, OMGS identifies the minimum structural requirement: three roles must be structurally separate.

Minimum role structure for ADM systems generating local rules

Role	Function in ADM context	Why separation matters
Rule maker (Governance authority)	The entity or function that makes the rules governing each contributory decision point — including what the AI system is authorised to do, what decisions it can make, and what it must escalate to a human	If the AI system makes its own rules without human authorisation, there is no external check on what rules govern the decision. The rules drift from human intent silently and without detection
Decision maker (Authorising body)	The entity or function that decides whether a rule has been correctly applied to a specific case — whether the contributory decision point followed its designed logic in this instance	If the same system that applies rules also decides whether they were correctly applied, there is no independent verification. Errors are self-certified and systematically invisible
Rule user (Subject of rules)	The AI system, automated process, or human actor that executes the contributory decision point according to the rules made and decided upon by the other two roles	The rule user must be distinct from both rule maker and decision maker so that compliance is assessed independently — not by the system being assessed

B3 — Task Boundary Framework (TBF)

GitHub and pitch deck | © [redacted] 2025, CC BY 4.0

What it is. TBF is a framework for identifying the human-machine reasoning boundary, breaking down a decision pathway into its component steps, assessing what type of reasoning each step requires, and surfacing the governance risk at each step.

Short context. Every information transformation is a decision, and every decision involves reasoning. This means every contributory decision point in an ADM process is a governable decision requiring assessment — not just the terminal decision. TBF provides the tools to make this assessment systematic and auditable.

Reasoning tiers — the human-AI reasoning boundary

Tier	Type of reasoning	Information environment	Human-AI boundary	Governance implication
What-tier	Rule-based execution — applies fixed rules to known inputs	Structured: fixed fields, standardised data, predictable format	AI and conventional automation appropriate where rules are explicit and information is standardised	Step-by-step explainability is inherent. State the rule applied. Periodic audit sufficient
How-tier	Adaptive process	Semi-structured:	AI can assist but human oversight	Explain factors, weights, and

Tier	Type of reasoning	Information environment	Human-AI boundary	Governance implication
	reasoning — optimises process given constraints	partial patterns, mixed signals, some ambiguity	required for consequential process decisions	thresholds. Human reviewer must access this information
Why-tier	Strategic and ethical judgement — contextual evaluation, purpose-setting, ethical reasoning	Unstructured: free text, qualitative assessments, narrative records, implicit context	Human reasoning required. AI can inform but must not substitute for judgement at this level	Full decision pathway explanation required. Human reviewer must independently verify — not just sign off on AI output

Information structurality — the limits of AI and conventional automation tools

The more unstructured the information environment, the harder AI reasoning is to explain step by step, and the stronger the human oversight requirement. This is not a risk preference — it is a structural property of unstructured information environments. Governance must scale inversely with information structurality and upward with reasoning tier.

How to use TBF in ADM governance

- Step 1 — Decompose the decision pathway into atomic steps: one input, one reasoning operation, one output.
- Step 2 — Classify each step by reasoning tier (What / How / Why) and information environment (structured / semi-structured / unstructured).
- Step 3 — Assess governance risk for each step: likelihood of AI reasoning error given the combination, multiplied by impact of that error on the terminal decision.
- Step 4 — Match tool and oversight requirement to each step based on the risk assessment.
- Step 5 — Verify that step-by-step explainability is achievable for the governance standard required. Where it is not, escalate to human reasoning for that step.

B4 — Project White Collar Framework (PJWC)

GitHub and pitch deck | © [redacted] 2019, CC BY 4.0

What it is. PJWC is an AI governance architecture designed for interoperability with human institutional governance. It provides a modular structure for governing AI execution in complex decision pathwayes.

Short context. Three principles that human governance arrived at through centuries of institutional design are directly applicable to AI systems performing consequential reasoning: separation of powers, diversity of perspective, and collective decision making.

Three principles applied to ADM

Separation of powers. Compliance checking, risk identification, and access control must be held by separate functions with separate authority and separate information access. No single system or actor should perform all three. This is the same structural requirement underlying constitutional separation of powers and corporate governance — applied to AI systems performing consequential reasoning.

Diversity of perspective. Real governance diversity requires distinct reasoning functions with distinct information access — not the same function performed by different people on the same AI output. Multiple human reviewers evaluating the same AI-generated summary do not constitute diverse oversight. Each governance function must be able to identify problems the other functions cannot see, because they access different information and perform different reasoning tasks.

Collective decision making. Consequential governance decisions — overrides, escalations, access changes — should require structured agreement among multiple governance functions rather than unilateral action by any single function. This prevents any single governance actor from both identifying and resolving a problem without independent check.

The Accountability Committee — governance structure for complex ADM setups

In complex ADM decision pathwayes involving multiple AI systems, the three governance functions map to a minimum viable governance architecture:

Governance function	Role in ADM context	Information access	Authority
Compliance checking function	Verifies that each AI contributory decision point is operating within its designed parameters and following authorised rules. Flags deviations from designed behaviour	Full visibility of AI reasoning outputs and audit logs. Cannot modify systems or grant access	Can halt a contributory decision point for review. Can activate independent review of the full decision pathway
Risk identification function	Detects risks beyond routine compliance — including ethical concerns, unexpected outputs, and risks to affected persons that fall outside the designed rule set	Full visibility of AI outputs. No visibility of compliance checking logs — this structural blindfolding prevents capture of the risk function by compliance concerns	Can escalate to owner or senior governance. Can pause operation for ethical review
Access control function	Controls AI system access to data, tools, and resources.	No visibility of AI reasoning logs or compliance	Can grant or revoke system access to data and tools. Cannot see

Governance function	Role in ADM context	Information access	Authority
	Executes access decisions mechanically without content discretion — it cannot see what it is controlling access to	checking logs — structurally blinded to content to prevent content-based interference with access decisions	the reasoning that triggered access requests

The structural blindfolding of the access control function — its designed inability to see the reasoning content it controls access to — implements the judicial impartiality principle: the function that controls access to resources must not be the function that evaluates whether the reasoning is sound. Impartial authority requires structural separation from content evaluation.

Appendix C — Plain language glossary

The following terms appear in this submission. Where a term has a technical equivalent in the referenced frameworks, this is noted.

Term	Plain English meaning	Technical equivalent	Lifecycle stage
Terminal decision point	The final, legally recognised outcome at the end of the decision pathway	—	Stage 1
Defensibility standard	The explicit specification of what it means for a decision pathway to be good enough — what level of accuracy, fairness, and explainability is required for the decision to withstand scrutiny	—	Stage 1
Decision pathway	The full sequence of contributory decision points and decision-shaping reasoning activities that produce a terminal decision, whether performed by humans, AI systems, or both	Decision chain	Stage 2
Information environment complexity	How structured or unstructured the information being processed is. Structured: fixed fields,	Data structurality	Stages 2, 4

Term	Plain English meaning	Technical equivalent	Lifecycle stage
	standardised data. Unstructured: free text, qualitative assessments, narrative records. More unstructured means harder to explain step by step and stronger governance required		
What / How / Why reasoning tiers	Three levels of reasoning depth. What-tier: rule-based execution. How-tier: adaptive process reasoning. Why-tier: strategic and ethical judgement. Higher tiers require stronger human governance involvement and cannot be substituted by AI without human-level oversight	Cognitive tiering	Stage 4
Decision pathway reconstructability	The ability to reproduce each contributory decision point for a specific decision after the fact, without relying on the AI system's own description. Measured at four levels: 0 (unrecoverable) to 3 (full independent third-party reconstruction). Assessed as good practice for high-risk ADM systems. Not a mandatory disclosure requirement under current APP 1. See Appendix A3 for a voluntary governance scorecard template.	—	Stage 8
Reasoning neutrality	Whether the AI system has an undisclosed default position — an outcome it tends toward in the absence of information pointing another way. Both approve-by-default and reject-by-default are bias positions. Neutrality requires structural design, not intent alone	—	Stages 3, 5
Reasoning contestability	A proposed future legislative protection enabling affected persons to identify and raise concerns about a specific contributory decision point in the decision pathway. Not a	—	Stage 8

Term	Plain English meaning	Technical equivalent	Lifecycle stage
	current APP 1 requirement — see Appendix D, FL1.		
Governance function separation	Holding compliance checking, risk identification, and access control as separate functions with separate actors, separate authority, and separate information access so that no single system or person performs all three	Functional separation; Accountability Committee	Stage 6
Gaming the substantial/direct test	Distributing substantive analytical work across multiple contributory decision points so each appears merely preparatory and no single point clearly crosses the threshold — avoiding the obligation through architectural design rather than by genuinely not using ADM	Forum-shopping by architecture	Stages 2–3
Disclosure accuracy gap	The growing difference between what an entity's ADM disclosure states and what its ADM systems actually do, which compounds silently as systems evolve without disclosure updates	Disclosure debt; governance debt	Stage 8
Formal notice vs actual understanding	Two distinct obligations: formally telling affected persons that ADM is used (formal notice) and ensuring they can actually understand what was disclosed and act on it (actual understanding). Both are required and neither substitutes for the other	Announcement vs communication	Stage 8
Governance scorecard	A set of seven measurable indicators showing whether the governance arrangements around an ADM system are actually working: decision reconstructability; oversight effectiveness; reasoning integrity; governance separation; self-correction	—	Stage 8

Term	Plain English meaning	Technical equivalent	Lifecycle stage
	capacity; disclosure accuracy gap; audit trail completeness		
AI-generated local rules	Operational rules that an AI system generates itself to handle novel cases — performing the functions of rule-making, rule-deciding, and rule-following simultaneously without human authorisation for each new rule	Rule generation without governance cycle	Stages 6, FL2
Human-only reasoning space	Reasoning steps that require substantive human judgement and cannot be adequately governed by AI alone — specifically Why-tier reasoning on unstructured information, where step-by-step explainability cannot be achieved without a human holding the judgement	Human-only space	Stage 4
Contributory decision point	Any point in the decision pathway where a decision-shaping reasoning activity materially shapes the information, analysis, or recommendations available at the terminal decision point	Captured within APP 1.7(a)'s 'do a thing substantially and directly related to making a decision.'	Stages 2–3
Decision-shaping reasoning activity	A specific action at a contributory decision point: selecting, filtering, ranking, scoring, characterising, summarising, recommending, or prioritising information	The legally preferred term for what AI systems do at intermediate points in the decision pathway, consistent with APP 1.7(a)'s second limb	Stages 2–4

Appendix D — Future legislative development

The following issues exceed what APP 1 guidance can deliver. They are offered as contributions to future legislative development. The Western Australia model2 demonstrates more comprehensive protections are achievable domestically.

FL1 — Protected right to challenge specific contributory decision points in the decision pathway

The disclosure obligation creates the information condition for a challenge right. It does not create the right. Current ADM legislation primarily prohibits certain practices and constrains how they occur. The enabling mechanisms — actively securing positive outcomes and empowering affected persons to access challenge pathways — are almost entirely absent.

A challenge right should operate at two levels. At the first level — already addressed in WA IPP 10 — the affected person can challenge the terminal decision and receive a substantive response. At the second level, which is not yet achieved in any Australian jurisdiction, the affected person can identify a specific contributory decision point in the decision pathway where they believe an error occurred and challenge that specific point. This is reasoning contestability as a protected right. It does not require challenging the entire decision — it enables targeted correction of the specific decision-shaping reasoning activity that produced the error. Because the terminal decision is the aggregated outcome of all contributory decision points, an error in any contributory decision point can invalidate the terminal outcome even where other points were sound.

The disclosure obligation as currently drafted creates the information condition for this right — affected persons who can access the decision pathway can identify where they believe an error occurred. The right itself requires legislation. Guidance can frame the disclosure obligation as the mechanism a contestability right is designed to serve. The WA model demonstrates this is domestically achievable. The Commonwealth framework should follow.

- Future reform

FL2 — Structural separation of governance functions for high-risk ADM

AI systems generating their own operational rules — already occurring in practice — have crossed from tools into rule-generating institutions without the institutional design safeguards that human societies developed to govern that transition. Future legislation should mandate structural separation of rule-making, rule-deciding, rule-executing, and compliance-oversight functions as a precondition for deploying ADM in consequential decisions.
This mandatory structural separation was flagged in Question 11 as a gap that guidance alone cannot close — FL2 is the forward-looking legislative response to that gap.

- Future reform / High-risk ADM best practice

FL3 — Three further reform considerations

Mandatory disclosure reconciliation. Periodic verification that disclosed ADM practice matches actual system behaviour.

Human role redesign ahead of AI absorption. As AI absorbs rule-based execution and process-level tasks, human roles must be redesigned toward strategic and ethical judgement ahead of time. Future legislation should govern the quality of human reasoning functions within ADM systems, not merely their presence.

Governance language alignment across layers. Vertical alignment between what legislation requires, what policy specifies, what procedures mandate, and what technical systems implement is a precondition for disclosure producing governed outcomes rather than documented compliance.

- Future reform

Appendix E — References

All sources cited in this submission are listed below in citation order.

Footnotes

1. Commonwealth Ombudsman, 'Automated Decision-Making Better Practice Guide' (March 2025). https://www.ombudsman.gov.au/__data/assets/pdf_file/0025/317437/Automated-Decision-Making-Better-Practice-Guide-March-2025.pdf

2. Privacy and Responsible Information Sharing Act 2024 (WA), Information Privacy Principle 10. https://www.legislation.wa.gov.au/legislation/prod/filestore.nsf/FileURL/mrdoc_47994.htm/$FILE/Privacy%20and%20Responsible%20Information%20Sharing%20Act%202024%20-%20%5B00-00-00%5D.html?OpenElement

3. Royal Commission into the Robodebt Scheme, Final Report (Commissioner Catherine Holmes AC SC, 7 July 2023). https://robodebt.royalcommission.gov.au/publications/report

4. OAIC, 'Automated decision-making and public reporting under the Freedom of Information Act' (Report, January 2026). https://www.oaic.gov.au/freedom-of-information/information-commissioner-decisions-and-reports/foi-reports/Automated-decision-making-and-public-reporting-under-the-Freedom-of-Information-Act

5. [redacted], Task Boundary Framework (TBF) Pitch Deck (CC BY 4.0, 2025). https://github.com/kfkchau/project-ahi-fin/blob/main/subfw-ahibom/pitch-deck.pdf

6. Air Vice-Marshal John Blackburn AO (RAAF Retd), 'A Pandemic of the Mind?' (Substack, 8 November 2025). https://johnblackburn1.substack.com/p/a-pandemic-of-the-mind

7. [redacted], Project White Collar Framework (PJWC) Pitch Deck (CC BY 4.0, 2019). https://github.com/kfkchau/project-white-collar/blob/main/pitch-deck.pdf

8. [redacted], Open Meta-Governance Standard (OMGS) Pitch Deck (CC BY 4.0, 2025). https://github.com/kfkchau/Open-Meta-Governance-Standard/blob/main/pitch-deck.pdf

9. Privacy and Other Legislation Amendment Act 2024 (Cth). https://www.legislation.gov.au/C2024A00128/asmade/text

10. OAIC, 'Automated Decision-Making Transparency Obligation (APP 1) — Issues Paper' (18 May 2026). https://www.oaic.gov.au/__data/assets/pdf_file/0027/263925/ADM-Issues-Paper.pdf

11. Explanatory Memorandum, Privacy and Other Legislation Amendment Bill 2024 (Cth). https://www.aph.gov.au/Parliamentary_Business/Bills_LEGislation/Bills_Search_Results/Result?bId=r7148

12. Information Commissioner's Office and The Alan Turing Institute, Explaining Decisions Made with Artificial Intelligence (2020). https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-artificial-intelligence/

13. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (EU AI Act), arts 11 - 14 and 14. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

14. NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (2023). https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf

Attribution

This submission draws on governance frameworks developed by [redacted], published under Creative Commons Attribution 4.0 International (CC BY 4.0).

[redacted] Open Governance Standard, [redacted] CC BY 4.0

For enquiries: [redacted].

Personal research.