Submission to the OAIC

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

Submitted by [redacted] | May 2026

About this submission

I am a data professional with 15 years of experience, including as a Data Governance Specialist in regulatory sensitive environments. [redacted]

My work centres on a practical problem: helping organisations map what systems they have, what data those systems hold, and how that data flows into decisions. I am making this submission from that operational perspective.

I am responding to Questions 1, 6, 9 and 10, where I have direct practitioner experience. I am not addressing the legal-interpretation or vulnerability questions, which are better served by other stakeholders.

Core position

An entity cannot disclose automated decision-making it has not first discovered and mapped. The accuracy of any APP 1.7 privacy-policy disclosure depends entirely on the quality of the entity's underlying system inventory and data-flow mapping. Guidance that acknowledges this, and points entities toward discovery and mapping as the first step, will produce better compliance outcomes than guidance that begins at policy drafting. This position runs through each of my answers below.

Question 1: Substantially and directly related to making a decision

> *In your view, what are the relevant factors to enable an entity to assess whether a computer program substantially facilitates and is directly connected to a human decision-maker's decision?*

The factors listed in the Issues Paper are sound. My suggested ranking, using the 1–5 scale, is as follows:

- **Integration of ADM into decision-making workflow: Strongly agree (5).** This is the most reliable signal. A system embedded in the workflow is difficult to override in practice, even where override is theoretically available.
- **Degree of reliance on the ADM system output: Strongly agree (5).** Where a decision-maker routinely depends on the output, the program is substantially related to the decision regardless of formal labelling.
- **Nature of the output, advisory versus determinative: Agree (4).**
- **Ability and likelihood of human override: Agree (4), with a caveat.** Theoretical override is weak evidence. What matters is whether override occurs in practice.

- **Transparency and explainability of outputs: Neutral to agree (3).** Important for good practice, but a weaker test of whether the program is substantially related to the decision.

I propose one additional factor: traceability of the data input. Whether the entity can identify what personal information the program consumed in order to produce its output. If an entity cannot trace the inputs, it cannot meaningfully assess or disclose the program's role in the decision. This factor connects the Question 1 test directly to the discovery and mapping work that must precede disclosure.

A practitioner note on human override. In my experience this is the factor most likely to be applied loosely. Entities will point to the presence of a human in the loop to argue that a system is not substantially related to a decision. Guidance should make clear that a human who approves outputs without genuine capacity or time to evaluate them does not break the substantial connection. Override must be real to be relevant.

Question 9: Meaning of "arranged for"

> *Are there any scenarios relevant to the meaning of "arranged for" that you consider require guidance? If so, why?*

This is the area I consider most in need of practical guidance, because it is where many entities will fail to recognise an obligation they hold.

The straightforward case, an entity deliberately procuring an AI system to make or assist decisions, is well covered by the examples in the Issues Paper. The harder and more common case is automated decision-making capability embedded inside third-party software that an entity already uses. A customer relationship management system, a case management platform, or a planning tool may contain an automated scoring, ranking, or recommendation feature. That feature may be enabled by the vendor, sometimes through a routine software update, and sometimes after the original procurement decision was made.

In this situation an entity may hold an ADM transparency obligation without being aware of it, because the capability arrived through a product update rather than a deliberate procurement choice. The entity has, in the words of APP 1.7, arranged for the computer program in the sense that it has chosen to use the software, but it has not consciously arranged for the automated decision-making within it.

I recommend that the guidance address this directly, and advise entities to:

- Maintain an ongoing inventory of third-party systems in use and the automated decision-making capabilities within each of them.
- Treat vendor feature updates and configuration changes as a trigger to reassess whether an ADM obligation has arisen.
- Treat third-party ADM oversight as a continuing monitoring activity, not a one-off procurement checkpoint.

A point-in-time procurement assessment is insufficient when automated decision-making features evolve continuously. This is, again, a discovery and monitoring problem. Entities need a living map of their systems and the decisions those systems influence.

Question 10: Extent of disclosure

> *Having regard to the OAIC's considerations, is there anything additional that entities should disclose to provide effective and meaningful disclosures about ADM?*

The five balance criteria in the Issues Paper are sound and I support them. I offer one addition, directed not at what is disclosed publicly but at how the disclosure is produced.

A privacy-policy disclosure is only as accurate as the system inventory and data-flow mapping that sit behind it. If an entity drafts an ADM disclosure without first mapping which of its systems use personal information in decisions, the disclosure is an estimate, not a statement of fact. It will also drift out of date as systems change, unless the entity has a process to keep it current.

I recommend that the guidance encourage entities to treat an ADM disclosure as the output of a documented internal mapping process, rather than as a standalone drafting exercise. In practical terms, an entity should be able to retain internally, and not necessarily publish, a record of how each disclosure was derived: which systems were reviewed, what data flows were identified, and when the review was last refreshed.

This protects consumers, because disclosures remain accurate and current. It also protects entities, because they can demonstrate diligence if the basis of a disclosure is ever questioned. It does not require publishing commercial-in-confidence detail about ADM systems, which the Explanatory Memorandum already excludes from APP 1.7.

Question 6: Other legal frameworks and policies

> *Are there any other legal frameworks or policies that the OAIC should have regard to in considering the meaning of "rights" and "interest", and more broadly in developing the guidance?*

Beyond the question of rights and interests, I suggest the OAIC consider aligning the vocabulary of the ADM transparency guidance with frameworks that many entities are already using to manage AI and data systems. This would help entities reuse existing work rather than treat the ADM obligation as a separate exercise.

Frameworks worth referencing include:

- **NIST AI Risk Management Framework (AI RMF 1.0) and the Generative AI Profile.** Widely used internationally. Its "map" and "govern" functions align closely with the discovery-first approach described in this submission.
- **ISO/IEC 42001, the AI management system standard.** Entities pursuing certification will already maintain an inventory of AI systems and their purposes.
- **Victorian Protective Data Security Standards (VPDSS V2.0) and state equivalents.** Relevant for public sector entities that already maintain information asset registers.
- **APRA's April 2026 letter to regulated entities on artificial intelligence.** It emphasises dedicated AI governance and oversight of AI suppliers, consistent with the third-party point raised under Question 9.

The common thread is that an entity which already maintains an AI or system inventory under one of these frameworks holds most of what the ADM transparency obligation requires. Guidance that cross-references these frameworks will help entities build on existing effort rather than start fresh.

Closing

The ADM transparency obligation will be met most effectively by entities that begin with discovery and mapping, and treat the privacy-policy disclosure as the output of that work. The commencement date of 10 December 2026 gives entities enough time to do this properly, provided they begin the work of identifying and mapping their systems now rather than treating the obligation as a drafting task close to the deadline.

I would be glad to provide further practitioner input if it would assist the OAIC in developing the guidance.

[redacted]