SUBMISSION TO THE

Office of the Australian Information Commissioner

Automated Decision-Making Transparency Obligation (APP 1)

Issues Paper

15 June 2026

Information Governance ANZ

Level 26, 1 Bligh Street, Sydney, 2000

www.infogovanz.com

Submission to the Office of the Australian Information Commissioner

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

15 June 2026

About this submission

InfoGovANZ welcomes the opportunity to contribute to the OAIC's development of guidance on the automated decision-making (ADM) transparency obligation introduced into Australian Privacy Principle 1 by the Privacy and Other Legislation Amendment Act 2024, which commences on 10 December 2026.

InfoGovANZ is a multi-disciplinary community of professionals across privacy, cyber and information security, data governance, information governance, FOI, records management, legal, risk, compliance, eDiscovery and AI governance and related disciplines. Our members sit on both sides of the ADM obligation: as practitioners who will design and document compliance within APP entities, and as advisers concerned with the broader information governance ecosystem. We approach this consultation from that practical, governance-centred perspective.

InfoGovANZ supports the introduction of the ADM transparency obligation. It is an important and necessary reform as organisations increasingly use automated systems, artificial intelligence, machine learning, rules-based engines, vendor platforms, generative AI tools and emerging agentic systems to support, influence or make decisions affecting individuals.

In summary, InfoGovANZ submits that the OAIC guidance should:

- Adopt a broad but practical interpretation of ADM, covering rules-based systems, AI, machine learning, generative AI, decision-support tools and embedded vendor functionality;
- Make clear that the presence of "human oversight" does not, by itself, take a use case outside the obligation: what matters is the function the system's output performs in the decision, whether it informs, recommends, prioritises, scores, ranks, escalates or otherwise shapes the outcome, not whether a human formally signs off;
- Recognise that meaningful human oversight requires more than a nominal reviewer. It requires a person with sufficient authority, competence, time, information and independence to interrogate the output and to reject, override or escalate it where appropriate. For higher-impact decisions, this may involve second-person or “four-eyes” review, supervisor or committee sign-off, or escalation to an accountable officer before a decision takes effect;
- Recognise that ADM transparency depends on good information governance – system inventories, data lineage, records of decision logic, vendor due diligence, monitoring and audit trails sufficient to explain how a decision was made, including records of the inputs to generative AI systems (such as prompts and the information supplied to them) where those inputs inform or lead to a decision;

- Encourage layered, proportionate disclosure, with deeper detail for high-impact decisions such as those affecting people more likely to experience vulnerability;
- Ensure individuals are given clear, accessible and workable pathways to obtain further information, seek correction or human review, challenge an outcome, make a complaint, and pursue external review or appeal where applicable; and
- Align the APP 1 guidance with administrative law, FOI, anti-discrimination and consumer protection law, and with responsible-AI frameworks, including the NAIC's Guidance for AI Adoption and the Digital Transformation Agency's Policy for the responsible use of AI in government and Standard for AI transparency statements.

This submission follows the numbering used in the OAIC's "ADM Issues Paper Submission Questions" document. Where the Issues Paper uses slightly different numbering, the substance of each corresponding question is addressed.

Question 1

Substantially and directly related to making a decision

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?

InfoGovANZ considers it would be helpful if the OAIC provided a practical, workflow-based test. The question should not be answered solely by asking whether a human formally makes the final decision; the relevant question is whether the computer program materially contributes to, shapes or constrains the decision-making process.

We rank the OAIC's listed factors, together with a small number of additions, on the 5-point scale (1 = strongly disagree it is relevant; 5 = strongly agree it is relevant). We have deliberately reserved a 5 for the factors that should carry the most weight, so that the ranking conveys a genuine hierarchy:

Factor	Ranking (1–5)
Degree of reliance on the ADM system output in practice	**5**
Integration of the ADM into the decision-making workflow	**5**
Nature of the output (advisory, ranking, scoring, triaging, recommending or determinative)	**5**
Practical likelihood of override, having regard to workload, time, training, authority, system design and organisational culture	**5**
Whether the output is operationally framed as a default or recommended pathway	**4**
Formal ability of a human decision-maker to override the system output	**4**
Whether the system is used at scale, repeatedly or as a standard operating procedure	**3**
Transparency and explainability of outputs	**3**

Comments

The threshold should turn on the weight the output carries in the actual decision-making process, not the technology used or a formal “human in the loop” label. We encourage the OAIC to frame the test around functional reliance rather than nominal oversight.

- Degree of reliance is the central factor. Where a human decision-maker routinely adopts the program's output without independent assessment, the program substantially facilitates the decision even where a human formally signs off. “Rubber-stamp” oversight should not take an entity out of scope.
- The nature of the output and degree of workflow integration, together with reliance, are the factors most likely to determine the assessment in practice. An output that is determinative, that ranks or scores individuals, or that sets a default a human must affirmatively displace is substantially related; an output that is one of several independent inputs a human genuinely weighs is less likely to be.
- Evidence of practical and appropriate override decisions matters more than the theoretical ability to override, absent such examples. Robodebt shows how a nominal capacity for human intervention can coexist with de facto automation. Relevant matters include workload, time pressure, system design, default settings, authority to depart from the recommendation, training and organisational culture.
- Explainability is ranked lower as a factor for scope. It affects how an entity should describe a system, but we caution against allowing low explainability to be argued as a reason a system falls outside the obligation. Complexity should not be used as an excuse for lack of scrutiny.

To assist entities, the guidance could confirm that the following uses will ordinarily be substantially and directly related to decision-making: scoring or ranking individuals; triaging applications or complaints; recommending eligibility, entitlement, priority or risk level; generating draft reasons or assessments; filtering candidates, customers or applicants; flagging individuals for investigation, escalation or exclusion; generating prompts or recommendations for a decision-maker; determining the evidence or options presented to the decision-maker; and automatically initiating workflow steps that materially affect the decision path.

Distinct from the question of scope, we encourage the guidance to promote proportionate human-control measures for higher-impact ADM – four-eyes or second-person review, supervisor or committee escalation, documented override rights, and recorded reasons for accepting or departing from the output. These are best treated as good-practice expectations rather than as elements of the scope test itself.

Question 2

Substantially and directly related to making a decision – Edge Case: Generative AI with human oversight

Would you consider NACIA's described uses of GPTea to be substantially and directly related to making decisions?

Yes.

Reasons

On the facts, GPTea performs two functions each substantially and directly related to the eligibility decision: it summarises candidate profiles (shaping what information the decision-maker sees and how it is framed), and it produces eligibility recommendations. The recommendation is, by description, the thing the human acts on. That is a key factor facilitating the decision and is directly connected to it. The context – eligibility for significant aged-care support for older persons – heightens, rather than lessens, the case for transparency.

The existence of human oversight and a final human sign-off does not, of itself, take the use case out of scope. The statutory test is whether the program does a thing substantially and directly related to making the decision and not whether the program makes the decision alone. APP 1.8(c) expressly contemplates disclosure of decisions for which such a thing is done, as distinct from decisions made solely by the program under APP 1.8(b). NACIA's use squarely engages APP 1.8(c).

Generative AI summarisation is not neutral: variation in what information is included, omitted or emphasised can materially steer the human assessor, so the summarisation step is substantially related, even setting the recommendation aside. We suggest the guidance state explicitly that genuine, substantive human review can be relevant to how a decision is described, but that the presence of a human in the loop does not by itself remove an entity from the obligation where the program's output is relied upon as a key input.

Generative AI also introduces transparency-relevant risks the guidance could usefully acknowledge – inaccurate or incomplete summaries, hallucinated or unsupported conclusions, over-reliance on apparently authoritative outputs, inconsistent recommendations for similar cases, difficulty explaining how a recommendation was generated, and inadequate records of prompts, outputs and human review. In the generative AI context in particular, we encourage the guidance to make clear that adequate record-keeping extends to the inputs as well as the outputs: the prompts entered, the personal information and source material supplied to the model, and any system or configuration instructions, where those inputs inform or lead to a decision. For high-impact contexts, entities should retain records showing the prompts and information supplied, the AI output, the evidence considered, the human assessment, any departure from the recommendation, and the reasons for the decision.

Question 3

Meaning of significantly affect rights or interests

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

We treat these as factors that raise the likelihood and the gravity of an effect, not as gateways that must each be met. We have spread the rankings to indicate relative weight rather than mark every factor as paramount:

Factor	Ranking (1–5)
Financial outcomes (income, credit, pricing, employment, essential goods and services)	**5**
Impact on access to essential services or supports (health, aged care, disability, housing, social services)	**5**
Intrusive practices, including surveillance, monitoring, biometric or identity verification, behavioural tracking, location tracking or automated risk flagging	**5**
Effect on a person more likely to experience vulnerability	**5**
Effect on legal status, entitlement, licence, benefit or right of review	**5**
Exclusion from opportunities (employment, education, insurance, credit, housing)	**5**
Use of sensitive information, or proxies for protected attributes	**4**
Profiling, segmentation or use of inferred, behavioural or third-party data	**4**
Difficulty for the individual to detect, challenge or correct the decision	**4**
Irreversibility or lasting consequences of the decision	**4**
Scale, repeatability and cumulative effect of the decision	**3**

Comments

A decision can significantly affect rights or interests without involving sensitive information – for example, an automated eligibility or pricing decision drawn from ordinary transactional data. As the Issues Paper and the EU Guidelines note, an effect that is modest for the general population may be significant for particular cohorts, so the likelihood assessment should be made with the affected population in mind, not an abstract “average” individual.

Beyond the listed examples, we encourage the OAIC to recognise: irreversibility or persistence of an effect (a decision hard to undo or that follows a person across services); cumulative or aggregate effects, where an individually modest automated decision becomes significant when combined with others or repeated over time; opacity to the individual, where the person cannot readily detect that automation was involved (as in targeted advertising); and effects flowing from the attributes or actions of others, such as group- or postcode-based inferences.

Question 4

Meaning of significantly affect rights or interests

What classes of person do you consider to be vulnerable for the purposes of considering a decision's expected effect on a person's rights or interests, and why?

We recommend a contextual rather than a closed-list approach. Vulnerability is situational: it arises from the interaction between a person's circumstances and the particular decision, not solely from category membership. Robust guidance would identify cohorts commonly more likely to experience vulnerability while making clear the list is non-exhaustive and that vulnerability can be temporary, contextual or intersectional.

Cohorts commonly more likely to experience vulnerability include children and young people; older persons; people with disability, including cognitive or psychosocial disability; people experiencing family violence; people experiencing financial hardship, homelessness or housing insecurity; people with limited digital access, digital literacy or English proficiency; Aboriginal and Torres Strait Islander peoples, particularly where decisions concern government services, health, welfare, justice or child protection; migrants, refugees and people with uncertain visa status; people receiving social security, aged-care, disability, health or housing services; people in regional, rural or remote areas; people affected by mental ill-health, trauma or crisis; people in custody or under supervision; job applicants and employees subject to algorithmic recruitment, performance management or work allocation; and individuals who lack meaningful alternatives to essential services.

Because an entity may not always know whether a person is vulnerable, the guidance should direct entities to consider the foreseeable class of affected individuals and the potential consequences of the decision, and – in contexts involving vulnerable persons – to design review pathways that are accessible in practice (clear language, non-digital alternatives where appropriate, reasonable assistance, prompt escalation of urgent matters, and review by a person with authority to change the outcome).

Question 5

Meaning of significantly affect rights or interests

What do you consider to be a significant service or support?

In our view, a "significant service or support" is one whose grant, refusal, conditioning or pricing has a material effect on a person's health, safety, financial security, livelihood, housing, education or ability to participate in everyday economic and social life. We rank the most clearly essential highest, but emphasise that the assessment should be functional rather than category-driven – most of the following can be significant in context:

Factor	Ranking (1–5)
Health, aged-care and disability services	**5**
Financial assistance and essential banking and credit services	**5**
Housing and accommodation	**5**
Essential utilities	**5**
Insurance (health, life, income protection, home, general)	**5**
Employment, recruitment and work allocation	**5**
Government benefits, licences, entitlements and identity services	**5**
Physical and in-home assistance / support	**5**
Telecommunications and connectivity	**5**
Digital identity, authentication and account access (as a gateway to other services)	**5**
Educational services on admission, exclusion, scholarship, assessment integrity or material support	**5**
Transport and mobility services	**3**

Comments

We make three further points.

First, we have ranked education at 5, however, other education-related decisions may be less significant. In our view, the guidance should treat education decisions as capable of being significant where they affect access, progression, opportunity or support.

Second, significance can arise from access being restricted or from the terms on which access is offered (including price), not only from outright refusal.

Third, telecommunications and connectivity have become functionally essential to reach most other services; their ranking depends on the decision's effect in context, and we caution against treating them as inherently lower-tier. Digital identity and account access warrant attention because they increasingly operate as gateways to other essential services.

Question 6

Meaning of significantly affect rights or interests

Are there any other legal frameworks or policies that the OAIC should have regard to in considering the meaning of 'rights' and 'interest'?

Yes. We support the OAIC's approach of interpreting 'rights' and 'interests' by reference to adjacent frameworks, and we agree the sources the Issues Paper already canvasses including administrative law, anti-discrimination law, the *Freedom of Information Act 1982* and Information Publication Scheme, the General Insurance Code of Practice, the *Privacy and Responsible Information Sharing Act 2024* (WA), and the EU GDPR and EU AI Act, are the right starting points.

On the frameworks already identified, we would emphasise the FOI Act connection in particular: the OAIC's own October 2025 review found only 17% of agencies authorised to use ADM disclosed that use in their IPS, underscoring the need for the APP 1 guidance and IPS expectations to reinforce one another.

We suggest the OAIC also have regard to records management and archives obligations (including duties to create and retain records sufficient to explain decision logic, the inputs and evidence relied on, including prompts and information supplied to generative AI systems, human review, reasons and system changes), which are not addressed as a distinct framework in the Issues Paper but are essential to making any right to understand or contest a decision practically effective.

Our broader point is interoperability: where the OAIC aligns its concepts with these frameworks, entities can build a single defensible governance process rather than reconciling competing definitions.

Question 7

Meaning of significantly affect rights or interests – Edge Case: Differential pricing

Would you consider the inflation of prices for a consumer product to have significantly affected the interests of Marco, and if so at what amount?

Yes. On the facts, the practice should be treated as capable of significantly affecting Marco's interests. The increase from $35 to $100 is substantial, approximately an 186% increase, and has the practical effect of preventing Marco from purchasing the book. The use of postcode as a proxy for wealth also illustrates the risk of inaccurate or unfair inferences: a person living in a higher-income postcode may not personally have high disposable income.

InfoGovANZ cautions against a fixed percentage or dollar threshold. The better question is whether the pricing practice could reasonably be expected to significantly affect the interests of affected individuals, having regard to the nature of the product or service, the amount of the differential, the recurrence of the cost, the affected cohort, and the practical effect on access.

As an indicative guide only, a 20% differential should generally be a strong signal of potential significance for non-essential consumer products. A 5% differential may also be significant where the product or service is essential, recurring, high-value, linked to financial hardship, or affects a person more likely to experience vulnerability. Larger differentials, such as the increase in Marco's case, are more clearly capable of significantly affecting interests. These figures should be treated as illustrative signals, not safe harbours.

A systematic differential-pricing practice that uses personal information or inferred characteristics to set prices is materially different from an ordinary one-off discount or general dynamic pricing. Such a practice can shut out access where prices become prohibitive, entrench disadvantage where it correlates with protected or sensitive attributes, and operate opaquely so individuals cannot detect or contest it. The significance lies not only in the dollar amount of a single transaction, but also in the practice, its cumulative reach and its practical effect on access. We recommend that the guidance distinguish ordinary dynamic pricing from ADM-driven differential pricing based on personal information, inferred characteristics or proxy variables.

Importance of factors (1 = not important, 5 = very important):

Factor	Importance (1–5)
The dollar amount of differential pricing	**3**
The significance of the product or service being purchased	**5**

We rank the significance of the product or service, the practical effect on access, and the affected cohort above the raw dollar amount. Differential pricing on essential goods and services – credit, insurance, utilities, healthcare-adjacent products, education, telecommunications or other significant services – is far more likely to significantly affect interests than the same differential on a discretionary purchase. The dollar amount and percentage difference are relevant but should be assessed in context.

Question 8

Meaning of making a decision – Edge Case: Discriminatory targeted job ad

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

Yes – though the harder question is whose obligation it is, addressed below.

Reasons

APP 1.9(a) provides that making a decision includes refusing or failing to make a decision, and APP 1.9(b) that doing a thing includes refusing or failing to do a thing. Opaque Window's algorithm determines which candidates to promote the role to, and the act of not surfacing the ad to Rafqa is the output of that determination – a decision (or at least a thing substantially and directly related to a decision about ad distribution) made by the operation of a computer program using personal information, here her gender.

The effect is significant because it limits Rafqa's access to an employment opportunity, a consequence the Explanatory Memorandum expressly identifies. That the mechanism is ad-targeting rather than a formal adjudication does not change the analysis: the Issues Paper recognises that targeting individuals with content can significantly affect them, including by limiting access to employment. The opacity of the practice, that is, Rafqa cannot see the ad she did not receive, heightens, rather than diminishes, the case for transparency.

The harder question is allocation. Opaque Window operates and designed the targeting and is the more obvious in-scope actor; Manscon's position depends on what it arranged for. Guidance should make clear that an advertiser who arranges for a third-party platform to target distribution can fall within the obligation in respect of that arranged use, even though the platform operates the underlying program, and that more than one entity may have obligations in respect of the same decision. The discriminatory use of gender also engages anti-discrimination law, reinforcing the synergy the Issues Paper describes.

Question 9

Meaning of 'arranged for'

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

Yes. “Arranged for” is, in our members' experience, the single area most likely to cause uncertainty and inconsistent practice. Many organisations will not build ADM themselves; they will procure, configure, enable, rely on or permit the use of systems supplied by third parties. The OAIC should make clear that “arranged for” includes situations where an entity selects, procures, configures, authorises, directs, enables or relies on ADM functionality, even if the technical operation is performed by a third party.

The guidance should identify practical indicators that an entity has “arranged for” ADM, including whether the entity determines the decision-making purpose, selects or configures the tool, sets eligibility or ranking criteria, supplies or authorises use of the personal information, receives and acts on the output, incorporates the output into a business workflow, relies on the output to affect an individual, or has contractual rights to require information, correction, audit or change.

Worked guidance across common sectors would assist, including: recruitment platforms (screening, ranking or targeting candidates); insurance (claims triage, fraud detection, risk scoring, premium setting); banking and credit (identity verification, credit and affordability assessment, hardship triage); telecommunications and utilities (service access, disconnection, hardship assessment); health, aged-care and disability services (eligibility, prioritisation, care-plan generation); education (admissions, scholarships, student-risk identification); employment (performance assessment, rostering, work allocation, termination risk); government service delivery (eligibility, compliance, risk scoring); generative AI in decision workflows (summarising evidence, drafting reasons, recommending outcomes); and case-management or CRM systems that automatically escalate, classify or prioritise.

We also encourage the guidance to address agentic AI and autonomous workflow tools – systems that can take steps, call APIs, update records, generate communications or escalate cases with limited human intervention. These strain the “arranged for” concept because an entity may authorise an autonomous workflow without controlling each step it takes; guidance on how responsibility attaches in that setting would be valuable as such tools are adopted ahead of commencement.

Beyond sectors, we recommend worked treatment of recurring structural scenarios:

- Procured and SaaS/AI tools: an entity buys or licenses a product (including general-purpose generative AI or embedded vendor features) and applies it to in-scope decisions. Guidance should confirm that adopting and directing a tool to a decision is “arranging for” it, even though the vendor built and hosts the model.
- General-purpose tools repurposed for decisions: staff use everyday software (spreadsheets, office tools, chatbots) to assist decisions. The distinction between operating and arranging for should turn on whether the entity directs the tool to the decision-making purpose, not on whether the tool was designed for it.
- Layered supply chains: an arranging entity uses a platform that itself relies on sub-processors or third-party models (as in the Question 8 scenario). Guidance should help allocate responsibility along the chain.

- Default or shadow use: employees adopt ADM tools without formal approval. Guidance should address when permitting or failing to prevent such use constitutes “arranging for”, and encourage governance controls (acceptable-use policies, tooling registers).
- Group and outsourced structures: one group entity arranges a tool used by another, or decision-making is outsourced. Clear allocation of the disclosure obligation would assist.

The Issues Paper's expectation that entities identify, assess and maintain oversight of third-party ADM during and after procurement is sound but demanding, particularly where the vendor (which may not be an APP entity) holds the information needed to comply. We recommend the guidance pair that expectation with practical due-diligence and contracting prompts: vendor disclosure of where and how ADM is used, what personal information and records are involved, how outputs are explained and corrected, and contractual rights to the information needed to populate the privacy policy and to support correction, review and audit. We acknowledge that an entity can only disclose what it can reasonably ascertain from its vendor.

We note, however, that contractual mechanisms place the burden on each deploying entity to negotiate for information the vendor is best placed to provide. In other domains, Australia has begun to address comparable supply-chain information asymmetries through supplier-side obligations, for example, the *Cyber Security Act 2024* (Cth) and its Security Standards for Smart Devices Rules require manufacturers and suppliers of relevant connectable products to ensure those products are accompanied by a statement of compliance. Without suggesting that the OAIC can or should impose obligations on vendors through APP 1 guidance, we encourage the OAIC to flag, for the attention of government, that the long-term effectiveness of the ADM transparency obligation may benefit from clearer expectations that ADM vendors make standardised disclosure information available to deploying entities.

For high-impact ADM, contracting should not displace internal accountability: an accountable system owner, adequate oversight, escalation arrangements, records and review pathways should be retained. This is a strong example of where good information governance operationalises the obligation.

Question 10

Extent of disclosure

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

InfoGovANZ strongly supports the OAIC's five stated principles, particularly plain language, appropriate tailoring, and framing that allows individuals to understand, challenge or contest relevant information and decisions. We suggest that ADM disclosures address:

1. The kinds of decisions affected by ADM (for example, eligibility, prioritisation, approval or refusal, ranking, targeting, triage, escalation, pricing, risk scoring, investigation, identity verification or service allocation);
2. The role of the computer program in the workflow – whether it makes the decision, recommends an outcome, scores or ranks, summarises information, flags risk, triages or drafts reasons;
3. The kinds of personal information used, including whether sensitive, inferred, behavioural, location, third-party or proxy data, or derived scores, are involved;
4. Whether third-party or vendor systems are involved, at a meaningful level, without requiring disclosure of commercial-in-confidence information, proprietary model details or security-sensitive information;
5. The potential effect on individuals (approval, refusal, delay, prioritisation, price difference, additional scrutiny or exclusion);
6. Whether and how a human is involved, avoiding any misleading impression that a human has substantively reviewed a matter where that review is limited or procedural;
7. Any material limitations of the ADM system that are relevant to individuals, including whether the system is used for recommendation, triage, risk scoring or summarisation, and whether outputs require human assessment before being acted upon;
8. How to request further information, including a clear contact point;
9. How to seek correction or human review, challenge or contest an outcome, obtain reasons and make a complaint, including external review or appeal pathways where available, which should be practical, timely, accessible and supported by reasonable assistance for vulnerable persons; and
10. When the ADM disclosure was last updated, supporting the APP 1.3 requirement that the policy be kept current.

For significant decisions, entities should also be encouraged to maintain records sufficient to explain the role of ADM, the inputs relied on (including, for generative AI, the prompts and information supplied to the system), the personal information used, the human assessment and the reasons for the outcome.

How to disclose effectively

- Layered disclosure: a concise, accessible summary at the top layer with the ability to drill down. This reconciles being “sufficiently specific to be meaningful” with “avoiding overwhelming levels of detail”.
- Proportionate depth: greater detail where the decision is high-impact, involves sensitive information, or affects people more likely to experience vulnerability; simpler presentation (and child-appropriate language where relevant) where the audience requires it.
- Accessibility and inclusion: plain language consistent with the Australian Government Style Manual, translations or accessible formats where warranted, and testing that real users understand the disclosure (echoing the UK ICO's 2024 recruitment-tools findings).
- Consistency across transparency artefacts: an entity's ADM disclosures should align with its other public-facing statements about AI and information handling, including any AI transparency statements, IPS information (for agencies) and related privacy notices, so that individuals receive a single, clear and consistent explanation rather than divergent or contradictory information across documents.

Finally, we encourage the OAIC to provide model wording and worked examples spanning common scenarios (a determinative automated decision, an AI-assisted human decision, and a procured third-party tool). Practical templates materially improve consistency and reduce the compliance burden, particularly for smaller entities, and would help realise the obligation's purpose of building public trust through clear, comparable disclosure.

Conclusion

InfoGovANZ commends the OAIC for consulting ahead of the December 2026 commencement and for the practical, scenario-based structure of the Issues Paper. The new ADM transparency obligation in APP 1 is an important step towards improving transparency, accountability and public trust in automated and AI-assisted decision-making.

To be effective, the guidance should be practical, risk-based and implementation-focused. It should make clear that transparency requires organisations to understand and govern their ADM systems, maintain adequate records, oversee third-party tools, ensure meaningful human control, and provide individuals with accessible pathways to obtain information and reasons, seek correction, challenge outcomes and request independent review where available.

Contact

[redacted]
Executive Director
E: [redacted]
T: [redacted]
www.infogovanz.com