DIGI Submission on the Automated Decision-Making Transparency Obligation (APP 1) Issues Paper, 2026

To: Office of the Australian Information Commissioner
By email: privacyreformtranche1@oaic.gov.au

29 June 2026

Dear Commissioner,

The Digital Industry Group Inc. (DIGI) thanks the Office of the Australian Information Commissioner (OAIC) for the opportunity to contribute to its guidance on the automated decision-making transparency obligation under Australian Privacy Principle 1.

DIGI is a non-profit industry association representing leading technology companies in Australia, including Apple, Discord, eBay, Google, Meta, Microsoft, Pinterest, Snap, Spotify, TikTok, Twitch and Yahoo. We work closely with the Australian Government to support a thriving digital economy in which innovation, privacy and online safety are protected.

We are pleased to provide this submission in response to the OAIC's Issues Paper on guidance for the automated decision-making (ADM) transparency obligation. DIGI remains a strong supporter of meaningful transparency where automated systems materially shape decisions that directly affect individuals' rights or interests (e.g. their employment, liberty, health, or shelter) in accordance with Parliament's intention; we believe that clear, accessible disclosures are essential to fostering trust and understanding in our digital economy.

Our submission aims to ensure transparency guidance is practical, proportionate, and avoids the implementation complexity that arises from duplicative regulatory mapping. We focus our comments on four areas where we believe industry clarity is most critical:

1. clarifying "significantly affect rights or interests" to defined consequential decision domains, rather than ordinary operational processes
2. clarifying that the "arranged for" obligation relates to the entity that deploys a consequential use, rather than every upstream technology used
3. applying a meaningful threshold so that AI or automated tools used for routine purposes with human oversight are not inadvertently captured
4. adopting a plain language disclosure standard, rather than requiring technical system detail.

We have proposed recommendations that protect the integrity of the transparency obligation. We believe these insights will help regulated entities deliver information that is useful for individuals without transforming privacy policies into overly complex technical inventories.

DIGI appreciates the ongoing dialogue on this important issue. We hope these insights are valuable as you finalise the guidance, and we would welcome the opportunity to discuss our recommendations further or assist in the development of practical worked examples.

Best regards,

[redacted] DIGI

DIGI

Table of contents

1. Introduction 2
2. Define the consequential decisions within scope 3
3. Confine “arranged for” to responsible use 6
4. Require disclosure when technology materially shapes a consequential decision 8
5. Enable meaningful disclosure without technical overreach 9
6. Promote regulatory coherence 11
7. Conclusion 12

1. Introduction

1.1. DIGI supports transparent use of ADM where automated systems materially shape outcomes significantly affecting individual rights or interests. To be effective, OAIC guidance must address three core questions: the scope of consequential decisions, appropriate due diligence for disclosures, and the threshold for material technological involvement. Without clear answers, the obligation risks becoming a broad inventory list of software and supply-chain actions rather than a targeted transparency measure that supports users’ appreciation of ADM.

1.2. DIGI therefore recommends that the Australian Information Commissioner (OAIC) guidance on transparency in privacy policies resolve three threshold positions.

- First, the guidance should define the consequential decisions within scope, by narrowing “significantly affect rights or interests” to a defined list of consequential decision domains that are materially related to individuals' rights or interests such as their employment, liberty, health, or shelter, and to a binary decision in accordance with Parliament’s intent as outlined in the Explanatory Memorandum, rather than an accumulation of low-level interactions.
- Second, it should confine “arranged for” to the responsible use, and clarify that the obligation carries a due diligence protection available to the entity that determines and deploys the relevant consequential use, not for every upstream infrastructure or software provider.
- Third, it should apply meaningful thresholds to the technology and its role. The relevant question is whether the technology materially shapes a decision that materially affects an individual's rights or interests, so that use of artificial intelligence (AI), automated tools or other software for ordinary, common, or operational purposes and processes that have human review or judgement, as well as commercially sensitive information and information relevant to the safety, integrity and security of services, are not inadvertently captured.

1.3. Further, the central implementation question is what level of detail will satisfy Australian Privacy Principles (APPs) 1.7–1.9. DIGI advocates for a privacy-policy plain language standard, not a technical list standard: individuals should be able to understand that ADM is used, the kinds of decisions affected, and the broad categories of personal information

involved, without disclosure of model architecture, training data or proprietary system design.

1.4. DIGI would welcome consulting with the OAIC to develop concrete examples to help anchor the OAIC's guidance in a practical, proportionate direction.

1.5. Sections 2, 3 and 4 of this submission address the three threshold positions set out above. Sections 5 and 6 address two further implementation matters that flow from those positions: the standards for meaningful disclosure in line with the APP 1 and the Explanatory Memorandum to the Privacy and Other Legislation Amendment Bill 2024 (Cth) (Explanatory Memorandum), and the case for interoperability with international ADM frameworks.

2. Define the consequential decisions within scope

2.1. The guidance should adopt the ordinary meaning of "decision" as a discrete, operative determination directed at an individual, rather than treating every automated act, process, prediction or intermediate step as a decision. The relevant question is whether a system determines an outcome that materially affects an individual's rights or interests, rather than whether AI or automated tools were used somewhere in the process. This approach is consistent with the ordinary meaning of "decision" and the Administrative Review Tribunal Act 2024, as foreshadowed in the Issues Paper. A failure, delay, suppression or other non-action may amount to a decision where it has the practical effect of determining an outcome, but should not be captured merely because automation plays some role in the surrounding process.

2.2. The OAIC should include both positive and negative worked examples. Negative examples are particularly important because they show entities what may properly be excluded from inventories and privacy policy disclosures. Suggestions for what would be helpful to cover include, but are not limited to, fraud detection, recruitment screening and other eligibility decisions in consequential domains (e.g., employment, credit, insurance, healthcare, housing and education).

2.3. For example, a fraud, safety, or moderation signal should not be treated as a 'decision' merely because it identifies higher-risk activity. If the system flags Activity A over Activity B for triage, prioritisation, or human review, it is assisting a process rather than deciding the outcome. But if that signal automatically denies, grants, or otherwise determines access to an essential service, statutory benefit, or contractual right (e.g., loan approval, healthcare access, visa grant) without human input, review, or judgment, then the system may be materially determining a consequential decision.

2.4. This reading is consistent with existing domestic legal frameworks. The meaning of "making a decision" in APP 1.7 should carry its ordinary meaning, informed by its use across administrative law, corporations law, and the Administrative Review Tribunal Act 2024. In those contexts, a "decision" is understood as a discrete, operative determination directed at an individual that produces a definitive outcome, such as granting or refusing a benefit, imposing a condition, or issuing or revoking a licence.

2.5. The Explanatory Memorandum at paragraph 343 confirms the intended scope of decisions that significantly affect the "rights or interests". The examples included therein – admission to a country, eligibility for housing benefits, contracts for life insurance, and access to healthcare – share the feature of direct, material consequences for legal, economic, or similarly significant positions. Similarly, in our view, the Issues Paper appropriately anchors this to APP 1.9(d) and the Explanatory Memorandum's examples of legislative benefits or entitlements, contractual rights, and access to essential services. This suggests that the phrase "rights or interests" should be interpreted by reference to decisions affecting matters of genuine significance, such as liberty, health, shelter, financial security, employment, or access to essential services, rather than remote, incidental, or ordinary commercial effects.

2.6. The application of those "rights or interests" should therefore be subject to a meaningful materiality threshold, having regard to APP 1.9 and the Explanatory Memorandum's examples (e.g., government entitlements, healthcare access), and should not be diluted by indirect or speculative impacts or the accumulation of micro-interactions..

2.7. DIGI is concerned that categorising routine or ordinary functions, such as a score or personalised product ranking, as consequential decisions would be an over-inclusion, as they do not typically determine access to an essential service, entitlement, or opportunity as Parliament intended. Treating them as such would compel entities to report on numerous low-impact systems, yielding dense privacy disclosures that would not serve a proportionate privacy benefit, and would obscure, rather than illuminate, the ADM applications that genuinely affect individuals. As personalised product ranking differs fundamentally from consequential decisions involving credit, educational opportunities, or statutory entitlements, the guidance must draw that line clearly.

2.8. This interpretation also gives meaningful and predictable force to the language of APP 1.7, which requires that the decision could reasonably be expected to "significantly affect the rights or interests of an individual". That threshold should attach to a specific determination affecting rights or interests of the kind discussed above, rather than extending to every automated act, process, prediction or intermediate step. It would also avoid creating novel Privacy Act meanings that do not have established interpretive history, and would improve cross-statutory certainty, preserve existing protections for commercial-in-confidence and security-sensitive information, and support worked examples grounded in practical outcomes rather than technical process.

2.9. A clear threshold is also necessary to avoid unintended operational outcomes. Where the definition of "significant effect" is assessed by reference to an individual's undisclosed circumstances such as undisclosed vulnerability, income level, or other personal characteristics, entities face a structural dilemma: they have no reliable means of knowing whether their automated output crosses that threshold without first collecting and analysing additional, and potentially sensitive, personal information about users. This is directly contrary to the principle of data minimisation. The OAIC should accordingly confirm that the "significantly affects" test is assessed against objective criteria, and that incidental processing of information in the inputs of a decision tool does not, of itself, bring that decision within the scope of the ADM transparency obligation.

2.10. This is particularly important when applying APP 1.9(b), which provides that "doing a thing includes refusing or failing to do a thing". The relevant question is whether an automated process results in a substantive refusal, exclusion, or denial of access, not merely whether it affects the order, timing, or prominence of information presented within a digital service. For example, a system that ranks Item A above Item B has not, by that fact alone, "refused" to show Item B, particularly where the system is continuously organising competing items within a finite display space rather than determining whether an individual may access a service, entitlement, or opportunity.

2.11. Accordingly, an act or omission should fall within scope only where it materially shapes the determination of whether an individual can access an essential service, entitlement, or opportunity. Reading "refusing or failing to do a thing" as capturing every item not displayed to a user at a particular point in time would make the concept of "decision" effectively unlimited, sweep in ordinary personalisation across digital services, and create an unworkable disclosure obligation. The focus should remain on whether the automated system substantively determines an outcome that significantly affects an individual's rights or interests, not merely whether automation plays some role in the surrounding process. This reading also applies to the pricing context considered below.

Pricing transparency should match the consequence

2.12. Guidance should confirm that differential pricing is only a consequential ADM decision where the price determination materially affects access to an essential product, service, or opportunity, rather than every instance of dynamic or personalised pricing. Differential pricing should be distinguished from ordinary dynamic pricing driven by demand, supply, timing, or inventory, and from low-impact retail personalisation.

2.13. Comparable frameworks such as GDPR Article 22, as referenced in the Issues Paper, reinforce that transparency should be calibrated to the risk of harm. In this way, higher-risk automated decision-making is typically limited to decisions producing legal or similarly significant effects within the text of APP 1.

Define the domains that matter

2.14. To give entities and individuals a workable basis for identifying in-scope decisions, the guidance should identify non-exhaustive examples aligned with the Explanatory Memorandum (e.g., access to healthcare, access to housing, access to government entitlements), rather than leaving entities to assess significance across every personalised interaction. Clear categories of consequential domains would give the "interests" threshold practical meaning. DIGI thus recommends adopting the narrower formulation. Read together, DIGI's comments in Section 2 define both the relevant outcomes and the domains in which they matter. This would give entities a clearer basis for identifying in-scope decisions for both individuals and services, consistent with Parliament's intent.

Safe harbour

2.15. To the extent that the definition of decisions within scope can be interpreted differently across different products and services, we recommend a show cause notice or similar to allow the opportunity to explain and/or remedy breaches, ahead of any issuance of penalty.

Summary of recommendations in Section 2

A. Define "decision" as a discrete, operative determination directed at an individual that produces a definitive outcome, rather than a mere automated act, prediction, process, or intermediate step.

B. Confirm that failures, delays, suppressions, or non-actions are only captured where they have the practical effect of determining such an outcome.

C. Apply the "significantly affect rights or interests" threshold by reference to decisions affecting material consequences for legal, economic, or similar significance such as liberty, health, shelter, financial security, employment, and access to essential services.

D. Clarify that APP 1.9(b) captures "refusing or failing to do a thing" only where automation substantively refuses, excludes, or denies access to a service, entitlement, or opportunity.

E. Confirm that dynamic pricing, product ranking, and personalisation of content or advertising do not, in themselves, constitute 'decisions' within the meaning of the provision, because they do not produce discrete determinations about individuals that significantly affect their rights or interests in clear consequential decision domains such as employment, healthcare, housing, or access to essential government services.

F. Develop positive and negative worked examples, including for fraud detection, recruitment screening, age estimation, content moderation, safety detection, recommendation systems, and general-purpose AI used with genuine human input, review, or judgment.

G. Implement a safe harbour by providing a show-cause notice or similar mechanism to allow entities to explain or remedy potential breaches before penalties are issued.

3. Confine "arranged for" to responsible use

3.1. To ensure regulatory clarity, the guidance should clarify the meaning of "arranged for" in APP 1.7(a), which the Explanatory Memorandum describes as the entity responsible for arranging the computer program to make or assist a decision. In practice, this should focus on the entity with control over, and visibility of, the relevant decisional purpose: for example, the entity that selects, commissions, configures or deploys the system for a particular decision-making objective, and is best placed to understand whether the resulting decision may significantly affect an individual's rights or interests.

3.2. This distinction is important in multi-party digital supply chains. Infrastructure providers, software vendors, model developers, hosting providers, and other technology suppliers including processors providing back-end, hosting, software or infrastructure functions on another entity's instructions, should not be treated as having "arranged for" a decision merely because their technology is used in the decision-making process. A provider may understand the technical characteristics of a tool but lack visibility over the downstream use case, decision-making purpose, affected cohort, safeguards, or consequential impact on individuals. The OAIC's own examples support this distinction: procuring an AI system to screen job applications, directing employees to use AI for performance assessments, and contracting software to approve or decline refunds all involve the commissioning entity, not the software vendor. Treating such providers as *arrangers* would risk producing redundant, speculative or inconsistent privacy disclosures, rather than meaningful information for individuals. It would also create disproportionate compliance burdens for service providers who are neither closest to the decisional objective nor best placed to explain the relevant automated decision-making practice.

3.3. For that entity, DIGI's position is that a proportionate due diligence protection should be available to the entity that determines and deploys the relevant consequential use: an entity that incorporates a third-party ADM tool into its consequential decision-making process should not be considered to have breached the obligation if it has taken reasonable steps to obtain the information required for a functional privacy policy disclosure. It would enhance regulatory certainty if the guidance identifies safe-harbour steps that ordinarily satisfy this expectation, such as:

- contractual terms requiring the vendor to describe the decisions the tool is designed to make or assist and the categories of personal information it uses;
- a proportionate assessment at procurement or before deployment;
- reasonable reliance on information supplied by the provider unless the entity knows, or ought reasonably to know, that it is materially inaccurate.

3.4. These steps would support accountability without turning a privacy policy obligation into a guarantee of a third-party system's architecture or performance.

Summary of recommendations in Section 3

H. Clarify that "arranged for" in APP 1.7(a) attaches to the entity with control over, and visibility of, the relevant decisional purpose, including the entity that selects, commissions, configures or deploys the system for that purpose.

I. Confirm that infrastructure providers, software vendors, model developers, hosting providers, and other technology suppliers are not "arrangers" merely because their technology is used in a downstream decision-making process.

J. Recognise that upstream providers may understand a tool's technical characteristics but lack visibility over the downstream use case, affected cohort, safeguards, and consequential impact on individuals.

K. Establish proportionate due diligence steps for entities using third-party ADM tools, including contractual information requests, pre-deployment assessment and reasonable reliance on provider information unless known to be materially inaccurate.

L. Confirm that APP 1 disclosure does not require entities to guarantee a third-party system's architecture, performance, or all possible downstream uses.

4. Require disclosure when technology materially shapes a consequential decision

4.1. Technology type alone should not trigger the disclosure obligation. The transparency obligation centres on 'computer programs', which in our reading encompasses rule-based systems, AI, and generative AI, but their presence or use in the technical stack is not sufficient to trigger disclosure. The relevant question is whether the technology materially shapes a decision that materially affects an individual's rights or interests, rather than the mere use of AI, automated tools, or other software.

4.2. Left broadly, this definition risks capturing routine operational tools, such as cybersecurity, customer relationship management (CRM), analytics, and general-purpose productivity functions (e.g., coding, translation, or summarisation) and the use of AI, automated tools or other software for safety, integrity, and security purposes.

4.3. "Substantially and directly related" should be assessed by reference to the system's actual role in the decisional workflow. For example, relevant factors include whether the output is factual versus functionally determinative, whether humans have genuine discretion and access to underlying or other relevant materials to form an independent view, and whether they are practically able to depart from the system output. Intermediate technical steps (e.g., data preparation, routing) should not be treated as decision-making in themselves.This ensures proportionality without creating arbitrary distinctions based on technical architecture.

Summary of recommendations in Section 4

M. Confirm that technology type alone does not trigger disclosure; the obligation applies only where technology materially shapes a consequential decision.

N. Define "substantially related" as technology that effectively determines an outcome, replaces human oversight, or operates where human reviewers lack the authority or information to override the output.

O. Define “substantially and directly related” as the function that produces the substantive recommendation or outcome, not intermediate steps such as data preparation, sorting, routing or formatting.

P. Exclude routine operational and productivity tools, including cybersecurity, CRM, analytics, coding, translation and summarisation tools, where they do not materially determine a consequential decision.

Q. Confirm that safety, integrity, and security tools should not be captured merely because they assist detection, triage, prioritisation or review, unless they materially determine the consequential outcome.

5. Enable meaningful disclosure without technical overreach

5.1. This section addresses the standard for what meaningful disclosure looks like in practice, flowing from section 2, 3, and 4: DIGI supports a plain-language, outcomes-focused standard that gives individuals a clear picture of how ADM affects them without requiring regulated entities to publish technical inventories or commercially sensitive system details. DIGI supports plain-language, functional disclosure. APP 1.7 requires privacy policies to describe the kinds of personal information used and the kinds of decisions made or materially assisted by computer programs, while APP 1.8 excludes commercial-in-confidence information. DIGI considers this overall framework sensible and workable in principle.

5.2. The guidance should also clarify the meaning of "tailored" disclosures (Issues Paper page 24). DIGI supports the principle that disclosures should be appropriate and accessible to users, but we caution against any interpretation that would require organisations to tailor notices to specific individuals or sub-groups of individuals. "Tailored" should be understood to mean that organisations have the flexibility to present their disclosures in a way that best reflects the nature of their specific ADM use and the character of their services, rather than that each user must receive a personalised disclosure. A requirement for individualised ADM notices would risk producing disclosures that are inconsistent, confusing, and ultimately less useful to users than a well-designed entity-level disclosure in a privacy policy.

5.3. Three principles should govern the scope and granularity of disclosure. These are a focus on category-level descriptions of ADM functions rather than granular technical logic, clear protection for commercial-in-confidence information, and an express exclusion for information that could compromise safety, integrity, or security systems. We recommend the following principles to ensure transparency remains useful for users:

- **Prioritise category-level descriptions:** Disclosure should focus on category-level practices of sensitive personal information used and the

practical role of the automated output. Excessive granularity would force frequent policy updates and overwhelm individuals with technical detail, ultimately reducing comprehension.

- **Maintain clear commercial-in-confidence boundaries:** The line should be drawn between explaining an individual-facing function and disclosing the technical means by which that function is implemented. Exclusions should expressly cover source code, model architecture and weights, training data composition, proprietary algorithms, feature weightings, training datasets, system prompts, confidence and detection thresholds, weighting of signals, and information that could compromise fraud, safety, or cybersecurity controls. This is particularly important for safety and integrity systems, where detailed disclosure enables adversarial gaming.

- **Utilise worked examples:** We encourage the OAIC to provide positive and negative compliant examples that demonstrate how to meaningfully inform readers without turning privacy policies into technical inventory lists.

5.4. We also recommend removing reference to challenge and contest. The Issues Paper suggests that disclosure should be "framed in a way that allows the information and the decision to be challenged or contested".[1] This changes the complexion of the transparency requirement from transparency (enhancing transparency so that individuals understand how their personal data is being handled in the context of ADM) toward redress. In DIGI's view, redress is a separate concept and we recommend removing this language to align with the legislative text.

Summary of recommendations in Section 5

R. Confirm that APP 1.7 permits plain-language, outcome-focused disclosure that explains the kinds of ADM decisions and personal information involved without requiring technical inventories.

S. "Tailored" should be understood to mean that organisations have the flexibility to present their disclosures in a way that best reflects the nature of their specific ADM use and the character of their services.

T. Allow category-level descriptions, layered information, and links to more detailed material where this improves user comprehension.

U. Avoid requiring model-by-model, version-by-version, or granular technical disclosure that would overwhelm users, require constant updates, and reduce the usefulness of privacy policies.

[1] Office of the Australian Information Commissioner, *Automated Decision-Making Issues Paper* (Issues Paper, May 2026) 23.

V. Formally exclude commercial-in-confidence and security-sensitive information, including source code, model weights, proprietary algorithms, feature weightings, training data, prompts, thresholds and signals that could compromise fraud, safety, or cybersecurity controls.

W. Provide positive and negative disclosure examples that show how to inform individuals about consequential ADM without exposing proprietary or adversarially sensitive information.

X. Remove language suggesting disclosure must enable a decision to be "challenged or contested", as redress is distinct from the APP 1 transparency obligation.

6. Promote regulatory coherence

6.1. A second implementation concern is regulatory coherence: a uniquely expansive Australian interpretation of 'decision,' 'arranged for,' and similar concepts would create duplicative mapping, disclosure, and governance obligations without improving transparency, and would increase reporting complexity and risk compliance success for entities already operating across multiple privacy and AI regimes.

6.2. Entities increasingly operate across multiple privacy and AI regimes. A uniquely expansive Australian interpretation of 'decision' etc. is not needed under the current legislative framework and would create duplicative mapping, disclosure and governance processes without necessarily improving transparency, while leading to significant implementation complexity.

6.3. We agree with the Issues Paper that the OAIC should expressly draw on established and shared international concepts, to avoid duplicative governance and support coherent cross-jurisdictional compliance. We recommend particular consideration be given to associated concepts concerning decisions producing legal or similarly significant effects, the EU AI Act's distinctions between providers and deployers and outcomes-based approaches, and US state consumer privacy and AI laws, including Colorado's treatment of consequential decisions and developer/deployer functions, and the Administrative Review Tribunal Act 2024 and the GDPR noted above.

6.4. Consistency does not require copying another regime wholesale. An effective ADM transparency regime could be achieved using established concepts where they address the same problem, establishing meritorious Australian deviations and enabling organisations to build coherent governance and transparency processes across jurisdictions.

Summary of recommendations in Section 6

Y. Draw on established international concepts where they address the same problem, including the GDPR, the EU AI Act's provider/deployer distinction and US state-law approaches to consequential decisions.

Z. Avoid uniquely expansive Australian interpretations of "decision", "arranged for" or related concepts since this would create duplicative governance without improving transparency.

AA. Use international interoperability as a starting point while developing Australian guidance that is fit for the Privacy Act, APP 1, and the OAIC's transparency function.

7. Conclusion

7.1. DIGI strongly supports a clear and workable transparency obligation for consequential automated decision-making, and the OAIC's Issues Paper provides a sound foundation on which final guidance can be built. To be effective, that guidance must do three things: it must confine the obligation to the decisions Parliament actually intended (discrete determinations with direct and material consequences for individuals' rights or interests, made or arranged for by the entity with the relevant decisional purpose), it must set a disclosure standard that is plain-language and outcomes-focused, protecting commercial-in-confidence and security-sensitive information rather than requiring technical inventories, and it must align with established international frameworks so that regulated entities are not required to maintain duplicative governance processes across jurisdictions.

7.2. Guidance designed on these three principles would focus transparency where it is most valuable to individuals, give industry the regulatory certainty it needs before the obligation commences, and strengthen Australia's position as a coherent and credible jurisdiction for responsible AI governance. DIGI looks forward to continued engagement with the OAIC as the guidance is finalised and would welcome the opportunity to assist with practical worked examples or further implementation consultation.