Public Consultation Submission on OAIC Automated Decision-Making Transparency Guidance

[redacted] Move78 International

9 June 2026

To: Office of the Australian Information Commissioner

Consultation: Guidance for Transparency in Automated Decision-Making

Submission date: 9 June 2026

Submitted by: [redacted] Move78 International

Contact: [redacted]

Confidentiality classification: Public. This submission does not contain confidential information and may be published by the OAIC, subject to OAIC publication and redaction processes.

Contents

Section	Topic
-	Executive Summary
-	Source Basis and Scope of Submission
-	Summary of Recommended Responses to OAIC Questions
1	Statutory and Regulatory Context
2	Point A: Functional Definition of a Decision
3	Point B: ADM Trigger Matrix for the “Substantially and Directly Related” Threshold
4	Point C: SaaS and API Shared Responsibility Model for ADM Transparency
5	Point D: Significant Services or Support and Rights or Interests
6	Standardised ADM System Inventory Template
7	Model APP Privacy Policy Disclosure Clauses
8	Recommended OAIC Guidance Text
9	Implementation Timeline for APP Entities
10	Claim and Source Register
11	Conclusion
Appendix A	Question Mapping
Appendix B	Reduced ADM Inventory Template for Smaller APP Entities
Appendix C	Source URLs

Executive Summary

This submission recommends that the OAIC’s final guidance for the automated decision-making transparency obligation under APP 1 should move beyond general principles and adopt auditable, operational instruments that APP entities can implement before the commencement of the new obligation on 10 December 2026.

The core implementation problem is classification. APP entities will need to decide whether a computer program is merely administrative, whether it makes a decision, whether it does a thing that is substantially and directly related to making a decision, whether the decision could reasonably be expected to significantly affect an

individual's rights or interests, and whether the entity has arranged for that ADM system or merely supplies, hosts, or operates a technology component.

If the final guidance leaves these thresholds at a high level, entities will be incentivised to create paper-based human-in-the-loop processes, vague APP Privacy Policy disclosures, and procurement structures that obscure rather than explain responsibility. That would weaken the transparency purpose of APP 1.7 to 1.9.

This submission proposes five regulator-ready instruments:

1. **A Functional Decision Test** for determining when a recommendation, ranking, score, escalation, suppression, price, prioritisation, or default workflow is sufficiently decision-shaping to be treated as a decision or as substantially connected to one.
2. **An ADM Trigger Matrix** that maps Human Override Capability against Model Explainability to distinguish effective human judgement from cosmetic human review.
3. **A Meaningful Human Involvement Test** to prevent sham human-in-the-loop arrangements from being used to avoid disclosure.
4. **A SaaS and API Shared Responsibility Model for ADM Transparency** to clarify responsibility where APP entities use third-party AI systems, APIs, cloud-hosted models, embedded decision engines, or foundation-model services.
5. **A Standardised ADM System Inventory Template** that the OAIC can append to final guidance as a practical compliance artefact.

The submission uses the language and structure of the Privacy Act 1988 (Cth), as amended by the Privacy and Other Legislation Amendment Act 2024 (Cth), and maps the proposed instruments to ISO/IEC 42001:2023 and the NIST AI Risk Management Framework 1.0. The emphasis is on ISO/IEC 42001 Annex A control areas for AI impact assessment, information for interested parties, use of AI systems, and supplier responsibility allocation, and on the NIST AI RMF MEASURE function for transparency, accountability, explainability, validation, and contextual interpretation.

Source Basis and Scope of Submission

This submission is based on the following source materials.

ID	Source	Use in this submission
S1	OAIC, *Consultation on Guidance for Transparency in Automated Decision Making*, published 18 May 2026	Consultation process, submission deadline, email route, publication/confidentiality rules, and OAIC policy objective.
S2	OAIC, *Automated Decision-Making Transparency Obligation (APP 1) Issues Paper*, May 2026	Consultation questions, edge cases, scope concerns, and OAIC's stated points of friction.
S3	*Privacy and Other Legislation Amendment Act 2024* (Cth), Part 15	Statutory text inserting APP 1.7, APP 1.8 and APP 1.9, and related compliance architecture.
S4	OAIC, *Chapter 1: APP 1 Open and transparent management of personal information*, updated 3 October 2025	APP 1.2 governance practices, APP Privacy Policy requirements, and forthcoming APP 1.7 to 1.9 obligations.
S5	ISO, *ISO/IEC 42001:2023 Information technology - Artificial intelligence - Management system*	Management-system basis for impact assessment, responsible use, documentation, traceability, transparency, and supplier controls.
S6	NIST, *Artificial Intelligence Risk Management Framework 1.0* and NIST AIRC Playbook, MEASURE function	Measurement basis for transparency, accountability, audit logs, human oversight, explainability, validation, and contextual interpretation.

This submission does not reproduce the copyrighted text of ISO/IEC 42001:2023. It uses control topics and management-system concepts as implementation anchors. Before final OAIC guidance quotes any ISO control text, the OAIC should verify the wording against the licensed standard.

Summary of Recommended Responses to OAIC Questions

OAIC topic/question area	Recommended OAIC position	Proposed operational instrument
Meaning of computer program	Interpret broadly, but distinguish ordinary utility software from decision-shaping automation by asking whether the program materially changes the analytical path, ranking, priority, eligibility, price, risk classification, escalation pathway, or practical access pathway.	Functional Decision Test.
Q1: substantially and directly related	Treat a program as substantially and directly related where its output is a key factor that directly shapes, constrains, ranks, filters, prioritises, recommends, scores, gates, prices, escalates, suppresses, or defaults the outcome pathway.	ADM Trigger Matrix; Meaningful Human Involvement Test.
Q2-Q3: generative AI with human oversight	Treat the NACIA/GPTea scenario as substantially and directly related where the chatbot summarises candidate profiles and provides recommendations about eligibility for aged-care support. Human oversight does not remove disclosure where the output materially guides the decision.	Trigger Matrix classification: likely H2-H3 depending on override evidence; X1-X3 depending on explainability.
Q4-Q7: rights, interests, significant services and vulnerability	Use a structured impact rubric including legal entitlement, contractual rights, financial effect, access to essential or significant services, employment, housing, education, healthcare, insurance, essential utilities, telecommunications, public support, vulnerability, sensitivity of information, recourse difficulty, and scale.	Downstream Impact Score and significant service schedule.
Q8: differential pricing	Avoid a fixed percentage safe harbour. A 5% price difference in an ordinary low-value consumer product may often be outside scope, but differential pricing should be in scope where it restricts access to significant goods or services, materially affects vulnerable persons, causes financial exclusion, or is combined with sensitive or inferred personal attributes.	Differential Pricing Materiality Rule.
Q10: targeted job advertisement	Treat the non-receipt of a targeted job advertisement as a decision, or as a thing substantially and directly related to a decision, where the ad-delivery system materially affects access to employment opportunities.	Functional Decision Test.
Q11: arranged for	Treat an APP entity as having arranged for ADM where it procures, configures, directs, permits, integrates, or relies on a third-party computer program for decisions about individuals. Merely hosting, operating, or developing a tool for another entity should not automatically make the supplier responsible for the customer's APP Privacy Policy disclosure.	SaaS and API Shared Responsibility Model.
Extent of disclosure	Require concise, decision-domain-level disclosure in the APP Privacy Policy, supported by internal inventories and evidence records. Do not require disclosure of model weights, source code, trade secrets, or commercial-in-confidence material where not needed for meaningful transparency.	Standardised ADM System Inventory Template and model APP Privacy Policy clauses.

1 Statutory and Regulatory Context

APP 1.7 to 1.9 create a privacy policy transparency obligation, not a general licensing regime for ADM systems and not a complete individual contestability right. The guidance should therefore be proportionate. However, it must be concrete enough to prevent avoidable misclassification.

APP 1.7 requires an APP Privacy Policy to contain the information covered by APP 1.8 where three conditions are met:

1. the entity has arranged for a computer program to make, or do a thing substantially and directly related to making, a decision;
2. the decision could reasonably be expected to significantly affect the rights or interests of an individual; and
3. personal information about the individual is used in the operation of the computer program.

APP 1.8 requires disclosure of the kinds of personal information used in the operation of such computer programs, the kinds of such decisions made solely by the operation of such computer programs, and the kinds of such decisions for which a thing substantially and directly related to making the decision is done by such computer programs.

APP 1.9 clarifies that making a decision includes refusing or failing to make a decision; doing a thing includes refusing or failing to do a thing; a decision may affect rights or interests whether adversely or beneficially; and

examples include decisions to grant or refuse benefits under legislation, decisions affecting contractual rights, and decisions affecting access to a significant service or support.

The OAIC should make three points explicit:

1. **Disclosure classification must be evidence-based.** Assertions that a system is advisory, human-reviewed, or out of scope should be supported by technical and workflow evidence.
2. **Human involvement is not a safe harbour.** A nominal approval click, bulk review, or workflow handoff should not displace APP 1.7 where the program output remains a key factor in the decision.
3. **Third-party procurement does not dissolve responsibility.** An APP entity that configures or relies on a third-party system for decisions about individuals remains responsible for its own APP Privacy Policy disclosure, even where the supplier has separate technical documentation obligations.

2 Point A: Functional Definition of a Decision

2.1 Recommended OAIC Guidance Text

For the purposes of APP 1.7 to 1.9, the OAIC should define a decision functionally rather than by job title, software label, or internal workflow label.

A decision should be understood as any determination, conclusion, classification, ranking, prioritisation, exclusion, grant, refusal, escalation, suppression, price, eligibility assessment, access pathway, contractual outcome, or failure to act that materially affects how an individual is treated by an APP entity.

A computer program may be substantially and directly related to making a decision even where a person performs a later workflow step, if the program output materially shapes, constrains, prioritises, ranks, filters, scores, gates, defaults, recommends, escalates, suppresses, or prices the outcome pathway.

2.2 Functional Decision Test

The OAIC should recommend that entities apply the following five-factor test.

Factor	Indicator of in-scope decision or decision-related thing	Indicator of ordinary utility tool
Outcome pathway	The program changes the pathway, priority, eligibility, price, risk class, access, escalation, or likelihood of human action.	The program only stores, formats, transmits, displays, or performs neutral arithmetic.
Consequence	The output affects legal status, contractual rights, financial position, employment opportunity, access to service, safety, support, or meaningful opportunity.	The output has no practical effect on the individual's treatment.
Degree of reliance	Human actors ordinarily accept, transmit, or act on the output.	Human actors use the tool only as one neutral information source and make a separately reasoned assessment.
Counterfactual effect	If the program output changed, the likely outcome or workflow priority would also change.	If the program output changed, the final treatment would be unaffected.
Evidence trail	Logs, prompts, thresholds, rank order, scores, recommendations, or default actions show that the program shaped the decision path.	Logs show only recordkeeping, retrieval, arithmetic, formatting, or administrative routing.

The counterfactual limb can be expressed as follows, where D is the practical decision outcome and C is the computer program output:

$$\frac{\partial D}{\partial C} \neq 0$$

This is not proposed as a literal mathematical test for all legal purposes. It is a concise expression of the operational proposition that the computer program is decision-shaping where changing the program output would change the probability, pathway, timing, ranking, or content of the individual's practical outcome.

2.3 Boundary Examples

A spreadsheet that calculates a person's age from a date of birth should generally be treated as an ordinary utility tool. A spreadsheet or rules engine that produces a risk score used to triage domestic violence crisis calls should be treated as doing a thing substantially and directly related to making a decision.

A generative AI tool used only to improve grammar in a document should generally be outside scope. A generative AI tool used to summarise an applicant profile and recommend whether to grant aged-care support, employment progression, refund approval, credit eligibility, insurance acceptance, or complaint escalation should be treated as in scope where personal information is used and the decision may significantly affect rights or interests.

3 Point B: ADM Trigger Matrix for the "Substantially and Directly Related" Threshold

3.1 Problem Statement

The most foreseeable compliance failure is the creation of sham human-in-the-loop arrangements. These will appear in three forms:

1. a human reviewer sees only a score or recommendation and clicks approve;
2. a human reviewer technically can override but lacks the data, explanation, time, authority, or training to do so meaningfully; and
3. a human review step occurs only after the computer program has ranked, filtered, suppressed, priced, escalated, or defaulted the pathway in a way that practically determines the individual's treatment.

The OAIC should prevent these designs from becoming de facto safe harbours.

3.2 Proposed Operative Principle

A computer program should be treated as doing a thing that is substantially and directly related to making a decision where its output is a material input into the decision pathway and the entity cannot demonstrate effective human override capability and sufficient explainability for the human decision-maker to make an independent determination.

This principle aligns with ISO/IEC 42001 control areas for AI system impact assessment and use of AI systems, because an entity cannot assess impact or govern use if it cannot identify how the system output affects the decision path. It also aligns with NIST AI RMF MEASURE 2.8 because transparency and accountability require meaningful visibility across pipelines, workflows, and organisational responsibilities.

3.3 Human Override Capability Scale

Code	Classification	Minimum operational meaning
H0	No human override	The system executes the decision or material workflow step without any human approval gate.
H1	Formal-only human step	A person clicks, confirms, transmits, batch-approves, or reviews only a final result without practical ability to alter the outcome.
H2	Limited human review	A person can approve, reject, or escalate, but lacks adequate data access, explanation, authority, time, training, or non-retaliatory discretion.
H3	Effective human review	A person can inspect the material inputs, explanation, model or rules version, confidence or uncertainty signals, and case file before the outcome takes effect, and can vary or reject the output.
H4	Determinative human decision	The system is advisory only. The human must record an independent determination, and the workflow technically permits departure from the output in every case.

3.4 Model Explainability Scale

Code	Classification	Minimum operational meaning
X0	Opaque	The entity cannot produce reliable documentation, case-level reasons, version traceability, or reproducible evidence for the output.
X1	Minimal	High-level product documentation exists, but the entity cannot explain a specific output in the operational context.
X2	Limited	Some reason codes, scores, feature importance, or prompt records exist, but the explanation is unstable, incomplete, or insufficient for meaningful review.
X3	Operational	Global and case-level explanation, model or system documentation, threshold information, performance limitations, and input-output traceability are available to authorised reviewers.
X4	Reproducible	The entity can reproduce or reconstruct the relevant output using preserved model/rules version, parameters, thresholds, input set, and decision pathway records, and can provide counterfactual or equivalent review support where appropriate.

3.5 ADM Trigger Matrix

Human Override Capability	X0 Opaque	X1 Minimal	X2 Limited	X3 Operational	X4 Reproducible
H0 No human override	Trigger	Trigger	Trigger	Trigger	Trigger
H1 Formal-only human step	Trigger	Trigger	Trigger	Trigger	Trigger
H2 Limited human review	Trigger	Trigger	Trigger	Trigger	Trigger
H3 Effective human review	Trigger	Trigger	Trigger	Trigger unless rebutted	Trigger unless rebutted
H4 Determinative human decision	Trigger	Trigger	Trigger	Trigger unless rebutted	Non-trigger only where advisory use and independent human determination are audit-proven

Trigger means the entity should disclose the relevant kinds of personal information and kinds of decisions under APP 1.8 unless the entity can rebut the classification with documented technical and workflow evidence.

The only presumptive non-trigger cell is H4 plus X4. Even then, the non-trigger classification should be available only where the entity can show that the system is genuinely advisory and that removal of the system would not materially alter the individual's practical pathway to the outcome.

3.6 Evidence Required to Rebut a Trigger

The OAIC should state that a non-trigger or "advisory only" classification should be supported by evidence including:

1. workflow diagrams showing where the program output enters the decision path;
2. role-based access records showing who can override and when;
3. interface screenshots or equivalent records showing the information presented to reviewers;
4. model, rules, prompt, or system version records;
5. input and output logs sufficient to reconstruct sampled decisions;
6. override, escalation, and disagreement statistics;
7. reviewer training and competency records;
8. complaint, review, correction, and recourse records;
9. vendor system cards or equivalent documentation where a third-party system is used; and
10. a documented impact assessment aligned to the entity's AI governance process.

3.7 Sham Human-in-the-Loop Red Flags

The OAIC should include the following rebuttable indicators of ineffective human review.

Red flag	Regulatory implication
Override rate below 0.5% across more than 1,000 decisions in a rolling 90-day period	Indicates possible rubber-stamping unless justified by independent quality evidence.
Batch approval or one-click acceptance above 95%	Indicates the human step may be administrative rather than substantive.
Reviewer sees only a binary recommendation, risk score, or colour flag	Indicates insufficient explanation for independent assessment.
System auto-executes after timeout or default acceptance	Indicates automated control over the practical outcome.
Reviewer lacks access to material source data	Indicates the reviewer cannot verify whether the output is justified.
Reviewer cannot override without manager, vendor, administrator, or engineering intervention	Indicates the override is formal rather than practical.
Reviewer incentives penalise override, delay, or escalation	Indicates the workflow suppresses independent judgement.
Entity cannot produce logs, model version, rule version, prompt record, or case explanation	Indicates insufficient auditability and likely X0 or X1 explainability.

These are audit triggers, not absolute liability thresholds. They should require review, documentation, and, where appropriate, disclosure.

3.8 Minimum Test Procedures

The OAIC should recommend three test procedures for medium and high-impact ADM systems.

Test	Purpose	Minimum procedure
Override Effectiveness Test	Determine whether the human reviewer can practically change the outcome.	Sample decisions and require the reviewer to vary, reject, or escalate the output without code changes, vendor intervention, or administrator override.
Blind Review Test	Determine whether the human decision is independent.	Have reviewers record a provisional decision without the model output, then compare against final decisions where the output is visible.
Explanation Sufficiency Test	Determine whether the explanation supports meaningful review.	Verify whether reviewers can access material reasons, source data, threshold conditions, confidence or uncertainty indicators, and version identifiers.

3.9 Response to NACIA/GPTea Edge Case

The NACIA example should be treated as substantially and directly related to making eligibility decisions unless NACIA can demonstrate H4 plus X4 conditions. The chatbot summarises profiles and provides recommendations about eligibility for aged-care support. That output materially guides a decision about access to significant support. The presence of human oversight does not remove the APP 1.7 trigger where the AI output is a key factor in facilitating the human decision.

4 Point C: SaaS and API Shared Responsibility Model for ADM Transparency

4.1 Problem Statement

Modern ADM systems often use third-party APIs, model-as-a-service platforms, SaaS decision engines, cloud AI products, embedded scoring systems, retrieval-augmented generation pipelines, or foundation-model services. The entity that decides the business purpose and acts on the output may not be the same entity that hosts or operates the model.

APP 1.7 uses the phrase “has arranged for”. The OAIC should interpret this phrase in a way that reflects operational control rather than physical hosting.

4.2 Foundational Principle: Responsibility Follows Operational Control

An APP entity should ordinarily be treated as having arranged for a computer program to make or assist a decision where it procures, configures, integrates, authorises, directs, permits, or materially relies on a third-party computer program for a decision about an individual.

A supplier that merely develops, hosts, or operates a system for a customer should not automatically be responsible for the customer’s APP Privacy Policy disclosure unless the supplier determines the decision purpose, controls the customer-facing decision, or itself uses the personal information to make or assist the decision.

This model avoids an accountability gap. The customer cannot avoid APP 1.7 by outsourcing the model. The supplier cannot be expected to disclose the customer’s decision categories where the supplier does not determine how the model is used. The supplier must, however, expose sufficient information for the customer to classify, disclose, and audit the ADM use.

4.3 Three-Tier Responsibility Architecture

Tier	Actor	Core responsibility
Tier 1	API, SaaS, model, or infrastructure provider	Expose technical documentation, model/system limitations, versioning, logging, data-handling terms, and material change information.
Tier 2	Integration layer, systems integrator, reseller, workflow orchestration provider, or internal platform team	Transmit and preserve documentation; map prompts, configurations, thresholds, routing rules, and data flows; maintain integration evidence.
Tier 3	APP entity using the system for decisions about individuals	Determine APP 1.7 classification; maintain inventory and impact assessment; update APP Privacy Policy; manage human review, complaints, recourse, and vendor oversight.

4.4 Vendor “Must Expose” Package

The OAIC should recommend that APP entities contractually require the following information before using third-party AI or ADM services in decisions about individuals.

Vendor exposure category	Minimum information required
System identity	Product/service name, model family, model or rules version, deployment mode, release notes, and deprecation schedule.
Intended and prohibited uses	Intended use cases, prohibited uses, high-risk use disclaimers, known limitations, and any use cases for which the vendor does not warrant fitness.
Data handling	Whether customer inputs, prompts, files, outputs, metadata, logs, or feedback are used for service improvement, fine-tuning, abuse detection, human review, or model training.
Data location and recipients	Hosting regions, subprocessors or sub-services, routing arrangements, and whether overseas disclosure may occur.
Input/output schema	Expected input fields, output fields, scores, labels, reason codes, confidence or uncertainty semantics, and API response structures.
Evaluation evidence	Relevant accuracy, calibration, reliability, robustness, bias, safety, security, and limitation information, proportionate to the use case.
Explainability support	System card, model card, reason-code schema, explanation API, feature-attribution method, prompt logging, retrieval context logging, or equivalent artefact.
Logging and audit support	Request identifiers, timestamps, model version identifiers, configuration metadata, retention settings, export functionality, and audit assistance.
Operational behaviour	Latency, timeouts, retries, fallback behaviour, safety-filter behaviour, rate limiting, and escalation behaviour that may affect outcomes.

Vendor exposure category	Minimum information required
Change management	Prior or prompt notice of material model changes, endpoint changes, safety-filter changes, threshold changes, subprocessors, hosting geography, or evaluation methodology.
Incident notification	Notice of personal information incidents, unexplained output degradation, model regressions, security events, or service events that could affect decision integrity.
Contractual assistance	Regulatory inquiry assistance, audit cooperation, complaint investigation support, and record preservation commitments.

Where a vendor cannot or will not provide this package, the APP entity should classify the system as X0 or X1 by default under the ADM Trigger Matrix and should not rely on a non-trigger classification without senior legal and privacy approval.

4.5 Enterprise Customer Disclosure Duties

Where APP 1.7 is engaged, the APP entity's APP Privacy Policy should disclose at least:

1. the decision domain or business function;
2. the kinds of personal information used;
3. the kinds of decisions made solely by a computer program;
4. the kinds of decisions where a computer program does a thing substantially and directly related to the decision;
5. whether the system is in-house, third-party, or hybrid;
6. whether the output is determinative, advisory, ranking, filtering, scoring, triaging, pricing, escalation-related, or default-setting;
7. whether human review occurs, and the broad nature of that review;
8. complaint and review channels;
9. likely overseas disclosures where required by APP 1.4(f) and APP 1.4(g), if practicable; and
10. the date of last policy update.

The APP Privacy Policy should not be required to disclose source code, model weights, proprietary training data, security-sensitive details, or commercial-in-confidence material unless such information is necessary to satisfy another legal requirement. Meaningful transparency can be achieved at the decision-domain and data-category level.

4.6 Pro Forma Contract Clauses

Transparency package. The Supplier must provide and keep current documentation sufficient to enable the Customer to classify the Service under applicable automated decision-making transparency obligations, including current system documentation, known limitations, intended uses, prohibited uses, version identifiers, deployment regions, and performance information relevant to the agreed use case.

Material change notice. The Supplier must give the Customer prior written notice of any material change to the model, rules engine, serving stack, safety filters, thresholds, evaluation methodology, subprocessor chain, logging functionality, or hosting geography that may reasonably affect output behaviour, explainability, privacy risk, or legal compliance. For emergency security or safety changes, notice must be given as soon as practicable after implementation.

Logging and evidentiary support. The Supplier must maintain and make available technical records reasonably necessary for audit and incident investigation, including request identifiers, model or rules version identifiers, timestamps, safety-filter outcomes, and configuration metadata, subject to lawful confidentiality and security constraints.

Incident notification. The Supplier must notify the Customer without undue delay of any personal information incident, material model malfunction, unexplained output degradation, or service event that may reasonably affect individuals, decision integrity, or the Customer's APP compliance obligations.

Audit assistance. The Supplier must provide reasonable cooperation, documentary assistance, and suitably scoped technical access to enable the Customer to respond to regulatory inquiries, internal audit, independent assurance, and complaint investigation relating to the Customer's use of the Service for decisions about individuals.

5 Point D: Significant Services or Support and Rights or Interests

5.1 Proposed Interpretive Rule

The OAIC should avoid an exhaustive list. The final guidance should include an indicative schedule and a materiality safeguard.

A significant service or support should include a service, support, opportunity, entitlement, facility, or access pathway that materially affects an individual's legal position, physical or mental health, financial resilience, housing security, employment or education opportunity, access to public or private support, safety, mobility, communications, or ability to participate in ordinary economic or civic life.

5.2 Indicative Schedule of Significant Services or Support

Domain	Examples of in-scope decisions
Government benefits and statutory entitlements	Eligibility, prioritisation, suspension, recovery, investigation, sanction, or failure to grant a benefit, licence, visa, permit, subsidy, or support.
Healthcare and aged care	Triage, diagnostic prioritisation, clinical risk scoring, aged-care eligibility, allocation of support hours, mental health escalation, or health insurance acceptance.
Housing and accommodation	Rental application ranking, social housing eligibility, tenancy risk scoring, mortgage hardship decisions, or shortlisting for housing support.
Credit, banking, insurance, and essential finance	Credit approval, credit-limit changes, insurance acceptance, premium setting, claim approval, hardship, fraud blocking, or account restriction.
Employment and livelihood	Job ad targeting, resume screening, candidate ranking, performance scoring, promotion, disciplinary escalation, shift allocation, or termination recommendation.
Education and training	Admissions ranking, scholarship eligibility, academic risk flagging, exclusion, support allocation, or special accommodation decisions.
Essential utilities and telecommunications	Electricity, gas, water, broadband, mobile access, hardship plans, disconnection prioritisation, fraud blocking, or service restoration.
Safety, policing, migration, and border-related decisions	Risk profiling, watchlist matching, visa triage, enforcement priority, bail support, family safety intervention, or public safety escalation.
Children and vulnerable persons	Any ADM that materially affects child welfare, school support, disability support, domestic violence support, aged-care access, or persons experiencing vulnerability.

5.3 Materiality Safeguard

A decision should not be treated as significant merely because it involves personalisation or ordinary commercial recommendation. The effect must be more than trivial or nominal. However, a low-value decision can become significant where it is repeated, cumulative, discriminatory, exclusionary, difficult to contest, based on sensitive information, or targeted at a vulnerable person.

5.4 Differential Pricing Rule

The OAIC should not adopt a fixed percentage threshold as a safe harbour. Differential pricing should be treated as potentially significant where one or more of the following applies:

1. the product or service is essential or significant;
2. the price change materially limits access for the affected individual or cohort;
3. the pricing uses sensitive information or proxies for protected characteristics;
4. the pricing targets individuals experiencing vulnerability;
5. the pricing is repeated, cumulative, or difficult to detect; or
6. the individual lacks realistic alternatives or recourse.

A 5% differential for an ordinary discretionary consumer product may often be outside scope. A 20% differential may be significant in a higher-value or repeated transaction. Any differential that effectively denies access to a significant service or support should be treated as significant, regardless of percentage.

6 Standardized ADM System Inventory Template

6.1 Purpose

The OAIC should append a standardised inventory template to the final guidance. The template should be framed as a practical mechanism for APP 1.2 governance and for ensuring APP Privacy Policies remain accurate and up to date.

Entities cannot disclose what they have not catalogued. A privacy policy update without a supporting system inventory is likely to be incomplete, inconsistent, and difficult to audit.

6.2 Five-Phase Compliance Workflow

Phase	Action	Output
Phase 1: Scope and discovery	Identify every system, model, rules engine, scoring engine, workflow automation, AI assistant, or API call that uses personal information in a process capable of influencing a decision about an individual.	Candidate ADM list.
Phase 2: ADM eligibility assessment	Apply APP 1.7 criteria, the Functional Decision Test, the ADM Trigger Matrix, and the Downstream Impact Score.	In-scope, pending, or out-of-scope classification.
Phase 3: Vendor and data mapping	Record data lineage, providers, model family, versions, subprocessors, regions, logs, retention, and human review design.	Vendor and data evidence pack.
Phase 4: Disclosure drafting	Draft APP Privacy Policy disclosure using kinds of personal information and kinds of decisions rather than excessive technical detail.	APP Privacy Policy update.
Phase 5: Governance and monitoring	Review material changes, model updates, new decision domains, complaint data, override rates, and incidents.	Current inventory and audit trail.

6.3 Standardised ADM System Inventory

Field	Required entry	Guidance
System ID	Persistent identifier, e.g., ADM-001	Use the same ID across privacy, legal, engineering, risk, audit, and procurement records.
System name	Business-facing name	Use the name compliance teams and business owners recognise.
Business owner	Role and team	Identify accountable business owner and second-line reviewer.
Decision domain	Employment, credit, refund, claims, healthcare, housing, support, pricing, etc.	Match the public disclosure category.
Model family	Rules engine, scorecard, decision tree, gradient boosting, LLM, embedding model, RAG, ensemble, SaaS API, hybrid	Avoid generic descriptions such as “AI tool”.
Model/version	Version, endpoint, rules release, prompt template, or deployment date	Required for reproducibility and change control.

Field	Required entry	Guidance
Provider	In-house, third-party, hybrid; vendor name if applicable	Supports the “arranged for” analysis.
Data lineage	Source systems, collection channel, personal information categories, transformations, prompts/features, storage, retention, deletion rule, overseas recipients	Supports APP 1.8(a), APP 1.4, APP 5, APP 6, and APP 8 analysis.
Sensitive data flag	Yes/No; describe category	Record health, biometric, children, financial hardship, racial or ethnic origin, criminal record, disability, or other sensitive/proxy data.
Output type	Decision, recommendation, score, rank, triage, price, flag, escalation, suppression, summary, default workflow	Determines whether APP 1.8(b) or APP 1.8(c) language is required.
Degree of HITL	H0, H1, H2, H3, or H4	Use the Human Override Capability scale.
Explainability level	X0, X1, X2, X3, or X4	Use the Model Explainability scale.
Downstream Impact Score	0-10	Use rubric below. Systems scoring 6 or above require second-line approval before any out-of-scope classification.
APP 1.7 classification	Trigger, trigger unless rebutted, non-trigger, pending	Link to evidence and reviewer approval.
Disclosure status	Required, disclosed, not required, pending, remediation required	Include APP Privacy Policy URL, section heading, and last update date.
Vendor evidence	System card, model card, API docs, contract clauses, change notices, subprocessors, security docs	Required for third-party systems.
Human review evidence	Workflow diagram, interface screenshot, reviewer guidance, training record, override logs, review-time data	Required for H3/H4 claims.
Complaint and recourse path	Internal complaint channel, review channel, external dispute body if applicable	Align with APP Privacy Policy complaint requirements.
Monitoring cadence	Monthly, quarterly, six-monthly, annual, material change	High-impact systems should not be reviewed only annually.
Risk mitigations	Thresholds, fallback, escalation, data minimisation, access control, testing, monitoring	Must be practical and evidence-backed.
Audit evidence location	Link or repository path	Avoid reliance on individual email inboxes or informal notes.

6.4 Downstream Impact Score Rubric

The Downstream Impact Score (DIS) should use a 0-10 structure. It can be expressed as:

$$DIS = L + F + R + S + V + A + O$$

where each component is scored as follows.

Component	Score range
L: legal entitlement, statutory benefit, visa, licence, penalty, investigation, or formal status effect	0-2
F: contractual, financial, credit, insurance, employment, housing, healthcare, education, or essential service effect	0-2
R: severity, duration, reversibility, and practical difficulty of recourse	0-2
S: sensitive information, inferred sensitive traits, or proxy variables	0-1
V: vulnerable person, child, disability, hardship, family violence, aged care, or comparable heightened impact	0-1
A: scale, repeatability, automation speed, or systemic cohort effect	0-1

Component	Score range
O: opacity, weak explanation, low override, or third-party black-box dependency	0-1
Total	0-10

Recommended interpretation:

Score	Assurance level
0-2	Low impact. Record rationale; annual review may be sufficient.
3-5	Moderate impact. Complete inventory and legal/privacy review before non-disclosure classification.
6-8	High impact. Presume disclosure unless rebutted; require second-line approval and evidence pack.
9-10	Very high impact. Require senior accountable owner approval, impact assessment, vendor evidence, monitoring, and APP Privacy Policy disclosure unless a documented legal basis for non-disclosure exists.

7 Model APP Privacy Policy Disclosure Clauses

7.1 General Clause

We use computer-based systems, including third-party software services in some cases, to assist with [screening / ranking / fraud detection / eligibility assessment / pricing / claims handling / customer support escalation]. These systems may use the following kinds of personal information: [insert categories]. Depending on the use case, the system may make a decision automatically, generate a recommendation, rank or prioritise matters for human review, or trigger a workflow step that is substantially and directly related to a decision. Where human review occurs, the reviewer may approve, vary, reject, or escalate the output before the decision takes effect.

7.2 Third-Party API Clause

We may use third-party software or API services to help process information and generate outputs used in [insert decision domain]. We remain responsible for how those outputs are used in our decision-making workflow. We require relevant suppliers to provide documentation about system limitations, version changes, data handling, logging, and audit support where needed for our privacy and governance obligations.

7.3 Complaint and Review Clause

Individuals may contact [insert contact point] to ask questions or make a complaint about how personal information is used in automated decision-making processes. Where a decision has affected access to a significant service, support, entitlement, contractual right, or similar interest, we will explain the available review or complaint pathway.

These clauses are starting points. They should be adapted to the entity's actual system inventory, decision domains, data categories, and complaint-handling arrangements.

8 Recommended OAIC Guidance Text

The OAIC could insert the following language into final guidance.

> An APP entity should not treat a human step as removing a computer program from APP 1.7 merely because a human remains involved somewhere in the workflow. Where a computer program materially shapes, ranks, filters, prioritises, recommends, scores, gates, escalates, suppresses, prices, or defaults the pathway to an outcome, the entity should assess whether that program does a thing that is substantially and directly related to making the decision.
>
> A human review step will be relevant to this assessment only where it is meaningful in practice. In assessing whether human involvement is meaningful, the entity should consider whether the reviewer

has sufficient authority, time, training, information, explanation, and technical means to make an independent determination before the outcome takes effect.

A nominal approval step, bulk confirmation, default acceptance, or escalation-only process should not be treated as meaningful human involvement where the reviewer ordinarily follows the program output or cannot practically interrogate and vary the output.

Where an APP entity procures, configures, integrates, permits, or relies on a third-party AI system, API, SaaS service, model, or rules engine for decisions about individuals, the entity should ordinarily treat itself as having arranged for the computer program for the purposes of APP 1.7 unless another entity alone determines the decision purpose, decision logic, and operational use of the system in relation to the individual.

Entities should maintain a current ADM system inventory that records the relevant system, model family, provider, data lineage, human review design, downstream impact assessment, explainability level, disclosure classification, and evidence supporting the classification.

9 Implementation Timeline for APP Entities

The OAIC should recommend a staged implementation plan for entities preparing for 10 December 2026. The dates below are readiness milestones, not statutory deadlines.

Period	Action
June-July 2026	Assign accountable owner, define ADM inventory scope, and begin discovery of systems using personal information in decision pathways.
By 30 September 2026	Complete initial system inventory, vendor mapping, and APP 1.7 triage for known high-impact systems.
By 31 October 2026	Complete impact assessment, trigger matrix classification, vendor evidence collection, and draft APP Privacy Policy disclosure language.
By 15 November 2026	Conduct legal, privacy, procurement, security, and business-owner review. Validate complaint and review pathways.
By 30 November 2026	Approve final APP Privacy Policy updates and internal evidence records.
From 10 December 2026	Maintain inventory, monitor material changes, review new systems before deployment, and update APP Privacy Policies when decision categories or personal information categories materially change.

10 Claim and Source Register

Claim or proposition	**Status**	**Source basis**
OAIC consultation was published on 18 May 2026 and closes on 15 June 2026.	Verified	S1, S2
OAIC requests submissions by email to privacyreformtranche1@oaic.gov.au and asks submissions to include corresponding question numbers.	Verified	S1
Submissions may be made available to the public on the OAIC website, and confidential information should be identified with reasons.	Verified	S1
APP entities using personal information in relevant ADM must update APP Privacy Policies from 10 December 2026.	Verified	S1, S2, S3, S4
APP 1.7 to 1.9 cover decisions made by a computer program and things substantially and directly related to making decisions.	Verified	S3
APP 1.8 requires disclosure of kinds of personal information used and kinds of decisions.	Verified	S3
APP 1.9 includes refusing/failing to decide and beneficial or adverse effects on rights or interests.	Verified	S3
The OAIC seeks feedback on factors such as degree of reliance, human override, output nature, explainability, and workflow integration.	Verified	S2
NIST AI RMF MEASURE 2.8 concerns transparency and accountability risks and supports	Verified	S6

Claim or proposition	Status	Source basis
documentation, audit logs, explanations, and measuring human oversight.		
NIST AI RMF MEASURE 2.9 concerns explanation, validation, documentation, and contextual interpretation of AI system output.	Verified	S6
ISO/IEC 42001 is an AI management system standard for organisations providing or using AI-based products or services.	Verified	S5
Specific ISO/IEC 42001 Annex A control wording is not reproduced.	Conservative copyright control	S5; licensed standard required for verbatim use.
Proposed trigger thresholds, HOC/X scales, override-rate red flags, and Downstream Impact Score are recommended guidance instruments, not current law.	Proposal	This submission.

11 Conclusion

The OAIC's final guidance should not require APP entities to publish source code, model weights, or excessive technical detail. It should require them to know and disclose, at a meaningful level, where personal information is used in automated decision pathways that can significantly affect individuals.

The most practical way to achieve this is to attach operational instruments to the guidance: an ADM Trigger Matrix, a Meaningful Human Involvement Test, a SaaS and API Shared Responsibility Model, and a Standardised ADM System Inventory Template.

These instruments would give APP entities a clear path to compliance, give auditors and regulators objective evidence to inspect, and reduce the risk that nominal human review or third-party outsourcing is used to avoid the transparency purpose of APP 1.7 to 1.9.

12 Appendix A: Question Mapping

OAIC question	Submission section
Q1: Relevant factors for assessing substantial facilitation and direct connection	Point B, especially “ADM Trigger Matrix”, “Evidence Required to Rebut a Trigger”, “Sham Human-in-the-Loop Red Flags”, and “Minimum Test Procedures”.
Q2-Q3: NACIA/GPTea edge case	Point B, “Response to NACIA/GPTea Edge Case”.
Q4-Q5: Rights, interests, vulnerability	Point D and “Downstream Impact Score Rubric”.
Q6-Q7: Significant services or support and relevant frameworks	Point D and “Standardised ADM System Inventory Template”.
Q8: Differential pricing edge case	Point D, “Differential Pricing Rule”.
Q10: Targeted job advertisement as a decision	Point A, especially “Functional Decision Test” and “Boundary Examples”.
Q11: Arranged for scenarios	Point C.
Extent of disclosure	“Enterprise Customer Disclosure Duties”, “Model APP Privacy Policy Disclosure Clauses”, and “Standardised ADM System Inventory Template”.

13 Appendix B: Reduced ADM Inventory Template for Smaller APP Entities

Smaller APP entities should still maintain an inventory, but the OAIC could permit a reduced template where the system is low complexity and low impact.

Field	Entry
System ID	ADM-[number]
System name	[insert]
Business owner	[insert]
Decision domain	[insert]
Personal information used	[insert]
Output type	[insert]
Third-party provider	[insert]
Human review level	H0/H1/H2/H3/H4
Explainability level	X0/X1/X2/X3/X4
Downstream Impact Score	0-10
APP 1.7 status	Trigger / Not trigger / Pending
APP Privacy Policy disclosure location	[insert]
Evidence link	[insert]
Last reviewed	[insert date]

14 Appendix C: Source URLs

S1. OAIC, *Consultation on Guidance for Transparency in Automated Decision Making*: https://www.oaic.gov.au/engage-with-us/consultations/consultation-on-guidance-for-transparency-in-automated-decision-making

S2. OAIC, *Automated Decision-Making Transparency Obligation (APP 1) Issues Paper*: https://www.oaic.gov.au/__data/assets/pdf_file/0027/263925/ADM-Issues-Paper.pdf

S3. *Privacy and Other Legislation Amendment Act 2024* (Cth): https://www.legislation.gov.au/C2024A00128/asmade/2024-12-10/text

S4. OAIC, *Chapter 1: APP 1 Open and transparent management of personal information*: https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-1-app-1-open-and-transparent-management-of-personal-information

S5. ISO, *ISO/IEC 42001:2023 Information technology - Artificial intelligence - Management system*: https://www.iso.org/standard/42001

S6. NIST AIRC, *AI RMF Playbook - Measure*: https://airc.nist.gov/airmf-resources/playbook/measure/

S7. NIST, *AI Risk Management Framework*: https://www.nist.gov/itl/ai-risk-management-framework