Transparency In Automated Decision-Making: Third-Party Systems And The Meaning Of “Arranged For”

[redacted] | Independent Executive & Practitioner
[redacted]

Submission to the Office of the Australian Information Commissioner consultation on the Automated Decision-Making Transparency Obligation (APP 1)

02 June 2026

Question Addressed

This submission responds to Question 11 of the OAIC Issues Paper.

It makes one narrow point.

> **An APP entity should not avoid the automated decision-making transparency obligation merely because the computer program that makes, or does a thing that is substantially and directly related to making, a decision is supplied, hosted, operated or embedded by a third party.**

The relevant question should be functional:

Did the APP entity arrange for the program to enter the decision pathway, and to do a thing substantially and directly related to making a decision, through procurement, configuration, authorisation, integration or operational reliance?

If yes, the APP entity should assess whether the ADM transparency obligation is engaged where the other APP 1.7 conditions are met.

1. Legal Frame

APP 1.7 applies where an APP entity has arranged for a computer program to make, or do a thing that is substantially and directly related to making, a decision, where the decision could reasonably be expected to significantly affect the rights or interests of an individual and personal information is used in the operation of the program.

The critical phrase for this submission is **“arranged for.”**

The OAIC Issues Paper recognises that a computer program may be operated by one entity while another entity is responsible for arranging for that program to make or assist a decision.

That distinction matters.

If “arranged for” is interpreted too narrowly, APP entities may be able to avoid transparency by placing technical operation, hosting or system provision with a vendor, platform or outsourced service provider while still relying on the program’s contribution to decisions about individuals.

If it is interpreted functionally, the obligation follows the entity that caused the program to enter the decision pathway.

This submission does not propose that every vendor, host, software maintainer or infrastructure provider should be treated as arranging ADM. It addresses the APP entity that selects, authorises, configures, integrates or relies on the program for decisions affecting individuals.

2. Response To Question 11

Question 11 asks whether there are scenarios relevant to the meaning of “arranged for” that require guidance.

Yes.

Guidance should clarify that “arranged for” includes situations where an APP entity, through its own choices, causes a computer program to make, or do a thing that is substantially and directly related to making, a decision using personal information.

Paragraph 337 of the Explanatory Memorandum is important here. It explains that the phrase “do a thing that is substantially and directly related to making a decision” was included because a computer program may recommend a decision to a human decision-maker, or guide a human decision-maker. That is the relevant statutory context for third-party systems, embedded tools and authorised staff use of generative AI.

Guidance should outline at least the following scenarios.

Procurement

An APP entity procures a third-party AI or software system to screen, rank, recommend, triage, approve, decline, score or otherwise guide decisions about individuals.

In that case, the entity has not merely acquired software. It has arranged for the program to enter the decision pathway.

Configuration and integration

An APP entity configures a system, selects decision thresholds, chooses settings, defines workflow rules, sets escalation triggers or integrates the system into a case-management, claims, recruitment, credit, benefits, healthcare or service-access process.

The entity should be treated as having arranged for the program’s decision role even where the underlying software is supplied or hosted by a third party.

Authorised staff use of generative AI

An APP entity permits, directs, enables or, through policy, practice or institutional expectation, normalises staff use of an AI tool to draft, summarise, score, recommend or assess material that guides or influences decisions about individuals.

Where staff are authorised to use an AI chat tool to generate eligibility summaries, performance assessments, complaint responses, risk assessments or recommendations, the entity should consider whether it has arranged for that program contribution where the output is substantially and directly related to making the decision.

Outsourced decision support

An APP entity contracts with a third-party service provider to carry out decision support on its behalf, including automated approval, rejection, ranking, fraud flagging, eligibility assessment or service triage.

The transparency analysis should not turn solely on the location of the server, ownership of the model or identity of the technical operator. The relevant question is whether the APP entity arranged for the program to enter the decision pathway.

Embedded platform functionality

An APP entity deploys a platform, workflow tool, customer relationship management system, recruitment product, claims system, case-management system or fraud-management tool that contains embedded ADM functionality.

Where that embedded functionality provides recommendations, scores, summaries, rankings, risk flags or eligibility signals that are substantially and directly related to making decisions about individuals, the entity should consider whether it has arranged for ADM through adoption and operational use of that platform.

3. Boundary: Technical Operation Alone Should Not Be Enough

Guidance should also distinguish arranging ADM from merely operating, hosting or maintaining a system for another entity.

An entity that only hosts infrastructure, maintains software or provides technical support does not arrange for ADM if it has no role in selecting the decision use case, configuring the decision pathway, authorising staff reliance or incorporating the output into decisions about individuals.

This boundary matters.

The purpose of APP 1.7 is not to impose transparency duties on every technical infrastructure actor in a supply chain. It is to ensure that the APP entity responsible for bringing a computer program into a consequential decision pathway discloses that use in its privacy policy.

A functional interpretation should therefore ask:

Who caused the program to be used in the decision pathway?

Who selected, configured, authorised, integrated or relied on the program for decisions affecting individuals?

Who incorporated the program's output into the decision process?

Those questions help distinguish arranging ADM from merely operating technology.

4. Third-Party Arrangements And Transparency

Guidance should make clear that third-party arrangements require active attention by APP entities.

Where an APP entity procures or uses third-party software or services, it should understand enough about the decision role of the system to assess whether the product uses personal information to make, or do a thing that is substantially and directly related to making, decisions about individuals.

This does not require public disclosure of proprietary model logic, commercial-in-confidence information, source code or case-level review records. It requires the entity to know enough about the decision role of the system to make accurate privacy-policy disclosures under APP 1.7 and APP 1.8.

The relevant disclosure is not merely:

“A third party provides software.”

The relevant disclosure is:

“A computer program is used to make, or do a thing that is substantially and directly related to making, these kinds of decisions, using these kinds of personal information.”

Where the APP entity has arranged for that use, outsourcing should not sever the transparency obligation.

These disclosures can remain category-level and plain-language. The point is not to expose a vendor’s proprietary system design. It is to make the role of the computer program in the APP entity’s decision pathway intelligible to the individual.

5. Suggested Wording For OAIC Guidance

The OAIC could adopt wording along the following lines:

> *In assessing whether an APP entity has “arranged for” a computer program to make, or do a thing that is substantially and directly related to making, a decision, entities should consider the substance of the operational arrangement rather than the identity of the technical operator. An APP entity may have arranged for ADM where it procures, configures, authorises, directs, integrates or operationally relies on a computer program to make, or do a thing that is substantially and directly related to making, decisions about individuals, including where the program is supplied, hosted or operated by a third party. The decisive question is whether the APP entity caused the program to enter the decision pathway, not whether it owns or operates the underlying software.*

6. Provenance And Material Considered

OAIC Issues Paper, Automated Decision-Making Transparency Obligation (APP 1), 18 May 2026
Privacy Act 1988 (Cth), Schedule 1, APP 1.7 to 1.9

Explanatory Memorandum, Privacy and Other Legislation Amendment Bill 2024 (Cth), paragraph 337

This submission draws on the author’s published legal-evidential work on provider and user roles, reliance-chain evidence and institutional proof in AI-assisted decision-making.

[redacted] *Deployment Entanglement: Frontier AI Providers, Operational DNA Transfer And Liability In The Reliance Chain*, https://doi.org/10.5281/zenodo.20214194

[redacted] *Permission Is Not Proof: The Evidential Burden Of AI-Assisted Execution Under V = 0*, https://doi.org/10.5281/zenodo.20155810

Transparency In Automated Decision-Making: Assessing Human Involvement By Function, Not Workflow Presence

[redacted] | Independent Executive & Practitioner
[redacted]

Submission to the Office of the Australian Information Commissioner consultation on the Automated Decision-Making Transparency Obligation (APP 1)

02 June 2026

Questions Addressed

This submission responds to Questions 1, 2, 3 and 12 of the OAIC Issues Paper.

It makes one narrow point.

> **Where a computer program is a key factor in facilitating a decision that could reasonably be expected to significantly affect rights or interests, nominal human involvement should not be treated as displacing the program from the decision pathway. Human involvement should be assessed by function, not appearance.**

For the purposes of this submission, verification capacity means the practical capacity of a human decision-maker to understand, test and depart from a program's output.

The relevant question is therefore simple:

Could the human understand the program's contribution, test its output and reach a different decision?

If not, and where the program remains a key factor in facilitating the decision, the program remains substantially and directly related to making the decision where the other statutory conditions are met.

1. Legal Frame

APP 1.7, read with APP 1.8 and APP 1.9 of Schedule 1 to the Privacy Act 1988 (Cth), creates a transparency obligation where an APP entity has arranged for a computer program to make, or do a thing substantially and directly related to making, a decision, where the decision could reasonably be expected to significantly affect the rights or interests of an individual and personal information is used in the operation of the computer program.

The critical phrase for this submission is "substantially and directly related to making a decision."

The Explanatory Memorandum states that a computer program may be used to recommend a decision to a human decision-maker, or guide a human decision-maker. In that context, "substantially" means the program is a key factor in facilitating the human's decision-making, and "directly" means the program has a direct connection with making the decision.

That raises the practical question at the centre of this submission:

When does human involvement displace the program as the key factor in facilitating the decision?

This submission proposes the following answer:

Only where the human involvement is substantive enough to understand, test and depart from the program's output.

Where a person merely reviews, approves or records a recommendation without that capacity, the human step should not sever the connection between the program and the decision. The program remains substantially and directly related to making the decision.

This is an interpretive point about the scope of APP 1.7. It is not a proposal for a new decision-governance regime, case-level audit duty or human-oversight framework.

2. Response To Question 1

Question 1 asks what factors enable an entity to assess whether a computer program substantially facilitates and is directly connected to a human decision-maker's decision.

The OAIC identifies relevant factors including degree of reliance on the ADM system output, ability and likelihood of human override, advisory versus determinative output, transparency and explainability of outputs and integration into the decision-making workflow.

Those factors should be read through one organising lens: **Did the human have practical capacity to understand, test and depart from the program's output?**

That question helps determine whether the program remains a key factor under the Explanatory Memorandum's test.

Understanding

Did the human decision-maker understand what the program contributed, what personal information it used and how the output affected the decision pathway?

If the person sees only a recommendation, score, summary or eligibility signal without understanding its role, the human involvement is likely to be nominal.

Access

Did the person have access to the relevant inputs, context, limitations, reasons or uncertainty needed to test the output?

Access to the final recommendation alone is not enough. A person cannot meaningfully depart from an output they cannot examine.

Authority And Practical Freedom

Could the person reject, modify, escalate or override the recommendation in practice, not merely in theory?

A formal override right has limited significance if the workflow, targets, organisational design or default pathway strongly expects approval of the program's output.

Effect

Could the person's involvement actually change the outcome?

If the decision process normally follows the program's output, or if human review operates only as confirmation, the program remains a key factor in facilitating the decision.

Using the OAIC's 1 to 5 scale, I would rank the listed factors as follows:

Degree of reliance on the ADM system output: 5
This is central to whether the program is a key factor.

Ability and likelihood of human override: 5
This is central to whether human involvement can alter the decision pathway.

Transparency and explainability of outputs: 5
Without explainability, meaningful review and departure from the output are weak.

Integration of ADM into decision-making workflow: 4
This is strong evidence of direct connection, though not always decisive without reliance.

Nature of the output: 4
This is important, but labels such as advisory or determinative should not control the analysis.

Practical verification capacity should not be treated as just another factor in the list. It is the organising lens through which the listed factors should be assessed: whether the human could understand, test and depart from the program's output.

3. Boundary: Substantive Human Involvement

This submission does not suggest that every use of a computer program in a decision pathway is necessarily captured by APP 1.7. The statutory question remains whether the program is a key factor in facilitating the decision and has a direct connection with making it.

Human involvement may be substantive enough to displace the program as the key factor where the human decision-maker has sufficient information, functional understanding, authority and practical freedom to test the program's output, depart from it and make the decision on an independently reasoned basis.

By contrast, where the human step consists mainly of receiving, reviewing, approving or recording the program's output, and where that involvement cannot in practice change the outcome, the program should remain substantially and directly related to making the decision.

4. Response To Questions 2 And 3

Question 2 asks whether NACIA's described use of GPTea is substantially and directly related to making decisions.

Yes.

Question 3 asks for reasons.

On the facts provided, GPTea summarises candidate profiles and produces eligibility recommendations for aged-care support. NACIA staff then make the final eligibility decision based on the recommendation.

GPTea's role is both substantial and direct.

It is substantial because the summary and eligibility recommendation appear to be a key factor in facilitating the human staff member's decision-making. It is direct because the output is connected to the eligibility decision itself, not merely to a remote administrative step.

The fact that a human makes the final decision does not, by itself, change that conclusion. The example states that the final decision is made based on the recommendation. That describes reliance on the program. It does not establish that the human decision-maker independently tested the program's contribution.

The relevant question is whether the staff member had sufficient information, functional understanding, authority and practical freedom to understand GPTea's contribution, test its summary or recommendation and depart from it where appropriate.

Nothing in the example indicates that human oversight displaced the program as the key factor. On the facts provided, the use of GPTea should therefore be treated as substantially and directly related to making the decision.

The decision also concerns access to aged-care support. That is a decision that could reasonably be expected to significantly affect rights or interests, including access to a significant service or support.

5. Response To Question 12

Question 12 asks whether entities should disclose anything additional to provide effective and meaningful disclosures about ADM.

Yes.

Where a computer program produces a recommendation, score, summary, ranking, eligibility signal or triage output that is substantially and directly related to making a decision, the privacy policy should describe that role in the decision pathway.

It should not be enough to say:

"Decisions are reviewed by staff."

That phrase may be accurate at a formal level but misleading at a functional level. It can imply meaningful human judgement where the person in the workflow could not understand, test or depart from the program's output.

Effective disclosure should therefore identify, at category level:

The kind of program contribution
For example: recommendation, eligibility signal, summary, triage score, ranking or risk flag.

The type of decision affected
For example: eligibility, access, pricing, benefits, contract rights, service access or support.

The role of human involvement
Where human review is relied on, the policy should describe whether the program assists a human decision-maker, whether the human may depart from the output and, at a high level, how a person can seek further information or contest the decision.

These category-level disclosures should be written in plain language, avoiding dense technical or legalistic descriptions while still making the program's role in the decision pathway intelligible.

This does not require disclosure of proprietary model logic, commercial-in-confidence information or case-level review records. It requires plain-language transparency about how personal information is used in the decision pathway.

These disclosures sit within the existing APP privacy-policy framework, including APP 1.4, and would help ensure that the ADM-specific information required by APP 1.7 and APP 1.8 is presented in a way that is meaningful rather than merely formal.

That is consistent with the purpose of APP 1: open and transparent management of personal information. It also supports individuals who may need to request further information, exercise review rights or challenge a decision affecting them.

That kind of disclosure also helps individuals understand when to seek further information, request review or challenge a decision where other legal or sectoral frameworks make those routes available.

6. Suggested Wording For OAIC Guidance

The OAIC could adopt wording along the following lines:

> *In assessing whether a computer program does a thing that is substantially and directly related to making a decision, an APP entity should consider the substance of any human involvement rather than its presence in the workflow. A program that recommends, scores, summarises or guides remains a key factor in facilitating a decision unless the human involved had sufficient information, functional understanding, authority and practical freedom to understand the program's contribution, test its output and reach a different decision. Human involvement should not be treated as taking a decision outside APP 1.7 merely because a person reviews, approves or records the program's output where that involvement could not in practice change the outcome.*

7. Comparative Context

The OAIC Issues Paper itself considers comparable frameworks, including GDPR Article 22, UK ICO material on explaining AI decisions and the EU AI Act's human oversight provisions. Those comparators support a functional approach to human involvement: the presence of a person in a workflow is not enough if that person cannot understand the automated output, act on it and alter the decision pathway.

This submission remains grounded in APP 1.7 to APP 1.9. The comparative point is confirmatory only.

8. Provenance And Material Considered

OAIC Issues Paper, Automated Decision-Making Transparency Obligation (APP 1), 18 May 2026
Privacy Act 1988 (Cth), Schedule 1, APP 1.7 to 1.9

Explanatory Memorandum, Privacy and Other Legislation Amendment Bill 2024 (Cth), paragraph 337

This submission draws on the author's published legal-evidential work on human oversight, verification capacity and institutional proof in AI-assisted decision-making.

[redacted] *From J = 0 To J = 1: Human Presence Is Not Proof Of Human Verification*, https://doi.org/10.5281/zenodo.20305912

[redacted] *Human Oversight Is Not A Liability Firewall* (KEV Legal-Evidential Note No. 3), https://doi.org/10.5281/zenodo.20305986