From: [redacted]
To: OAIC - Privacy Reform Tranche 1
Subject: Submission - Automated Decision-Making Transparency Obligation (APP 1) Issues Paper
Date: Sunday, 14 June 2026 7:09:18 PM

CAUTION: This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

To the Office of the Australian Information Commissioner,

Thank you for the opportunity to make a submission on the *Automated Decision-Making Transparency Obligation (APP 1) Issues Paper* (18 May 2026).

Dark Horse Governance is an independent governance advisory and research practice focused on AI supply-chain transparency and the board-level governance of AI systems. This submission does not attempt to answer every question in the paper. It responds to **Question 11** (the meaning of "arranged for"), **Question 12** (the extent of disclosure) and **Question 7** (other frameworks the OAIC should have regard to), and raises **one additional matter** the paper does not address: the interaction between the obligation and the s 7B(3) employee-records exemption.

A single proposition runs through all four. The obligation in APP 1.7–1.9 attaches to the entity that *arranges* the program, but the facts needed to comply with it — what personal information the program consumes, what kinds of decisions it makes, whether outputs are human-reviewed, what the provider does with the data — typically sit one or more layers down the supply chain. An entity cannot disclose, accurately and on a current basis, what it has not first mapped. In our advisory work we describe the resulting internal artefact as an "AI bill of materials" (AI-BOM): a maintained, per-system record of each AI deployment's model, provider, processing location, data inputs, decision classes, human-review posture, retention and place in the sub-processor chain. Several of the paper's interpretive questions resolve more cleanly once that artefact is assumed, and we suggest the guidance can usefully point entities toward it.

Question 11 — the meaning of "arranged for"

We support the OAIC's distinction between an entity that has *arranged for* ADM and one that merely *operates* it, and the consequential expectation that entities monitor third-party ADM during and after procurement.

We add one practical observation and identify two scenarios we consider

require guidance.

The practical consequence of the arranger/operator split. Because the obligation rests on the arranging entity while the load-bearing facts sit with the provider, discharging it in a procured deployment is, in substance, an information-supply problem. The arranging entity can only make and maintain an accurate APP 1.8 disclosure if the provider supplies — and keeps current — a description of what the program consumes and the kinds of decisions it makes. We suggest the guidance frame the procurement-and-monitoring expectation concretely: entities should obtain a maintainable description of the ADM system's inputs and decision classes from the provider, and a contractual commitment to notify the entity when models, features or decision classes change. This is the AI-BOM discipline expressed as a procurement control, and it is the mechanism by which the "arranged for" duty is actually met rather than merely asserted.

Scenario 1 — staff-initiated ("shadow") AI. A common and growing pattern is an employee independently adopting a generative-AI tool — a chatbot the entity did not procure — to make, or substantially and directly assist in making, a decision of the kind APP 1.7 captures (for example, drafting an assessment that determines an applicant's progress). Has the entity "arranged for" that program? We submit the better view is that an entity which permits, or fails to take reasonable steps to prevent, such use in a context where it could reasonably foresee triggering decisions has, in effect, arranged for it — and that the guidance should make clear the obligation cannot be avoided by leaving AI adoption informal or unmonitored. A construction under which only formally procured systems are "arranged for" would reward entities for *not* mapping their AI use, which is the opposite of the transparency the obligation seeks. Guidance on this point would also reinforce the OAIC's own emphasis that "computer program" includes generative AI and chatbots.

Scenario 2 — AI features added to existing software after procurement. Where a SaaS provider adds an ADM capability to a product the entity already uses, the entity may be making triggering decisions through a program it never separately procured or assessed. We suggest guidance confirm that the procurement-monitoring expectation is continuing rather than point-in-time: a change in a deployed product's capabilities can bring an entity within APP 1.7 without any new procurement decision, and entities should be expected to keep their inventory and disclosures current as their vendors' products evolve.

Question 12 — the extent of disclosure

We support the five qualities the OAIC identifies. We offer two points.

Meaningful disclosure is downstream of internal mapping. The plain-language, appropriately-tailored, current disclosure the OAIC describes is only achievable where the entity first holds an accurate internal record of which of its systems perform ADM and what they decide. APP 1.3 already requires the privacy policy to be up to date; given the pace at which AI features change, an up-to-date APP 1.8 disclosure is not sustainable without a maintained internal inventory behind it. We suggest the guidance encourage entities to maintain such an inventory (by whatever name) as the precondition for compliant disclosure, and treat the external disclosure as a view onto a living internal artefact rather than a one-off drafting exercise.

Commercial-in-confidence should be construed narrowly. We accept that genuinely proprietary system internals — model architecture, weights, source code — may be commercial-in-confidence and outside the disclosure requirement. We urge the guidance to draw the line carefully, because the carve-out is the lever a provider will pull against transparency. The categories that APP 1.8 actually requires — the *kinds* of personal information used, and the *kinds* of decisions made — are not proprietary system internals, and a vendor's confidentiality claim should not be permitted to swallow them. Worked examples distinguishing protectable commercial-in-confidence detail from the disclosable APP 1.8 categories would materially assist entities resisting over-broad confidentiality assertions from their providers.

Question 7 — other frameworks the OAIC should have regard to

We draw the OAIC's attention to the **cross-border disclosure regime in APP 8, read with s 16C**, as the framework most closely entangled with this obligation in practice.

The same AI deployment that triggers APP 1.7 on the *output* side (what the program decides) frequently triggers APP 8 and s 16C on the *input* side (where the personal information goes when the program runs, which for most generative-AI tools is offshore). The facts an entity needs to answer both are the same: model, provider, processing region, retention, and whether the provider uses or trains on the data. Recent analysis of the major US frontier developers' privacy policies found that all six train on user inputs by default and several retain that data indefinitely (King et al., AAAI 2025) — facts that bear directly on both the APP 8 use-versus-disclosure question and the APP 1.8 description.

We suggest the ADM guidance cross-reference APP Guidelines Chapter 8 and encourage entities to satisfy both boundaries from a single inventory, rather

than building two disconnected compliance artefacts. This reduces regulatory burden while improving the quality of both disclosures, and it reflects the reality that the "arranged for" entity under APP 1.7 is usually the same entity that bears the non-delegable s 16C accountability for the cross-border leg.

Additional matter — the s 7B(3) employee-records exemption

We note, with respect, that the Issues Paper does not address the interaction between APP 1.7–1.9 and the **employee-records exemption in s 7B(3)**. We consider this the most significant unaddressed interaction, and one on which entities will need guidance before commencement.

APP 1.7–1.9 sit within the Australian Privacy Principles and inherit the exemptions that apply to the APPs generally. The consequence is that the obligation does not apply to AI-mediated decisions about current or former employees where the act or practice is directly related to the employment relationship and an employee record. The same AI tool that triggers an APP 1.8 disclosure for a decision about a customer or service-recipient is exempt when applied to an employee's performance management, discipline or termination — at precisely the moment that algorithmic management of the workforce is expanding. The OAIC's published guidance on the exemption is dated December 2015, and predates both the ADM amendments and the contemporary AI deployment landscape.

We do not suggest the OAIC can close the policy gap by guidance; its scope is a matter for Parliament. We do recommend that the guidance:

1. **acknowledge the interaction** expressly, so entities do not have to infer the exemption's effect on the new obligation;
2. **clarify the boundary** — in particular that the exemption does *not* extend to prospective employees and unsuccessful job applicants (until an employment relationship is formed), to volunteers, or to contractors and HR-services providers handling another organisation's employee data. Much workplace AI — recruitment screening, candidate ranking, third-party HR and assessment platforms — sits outside the exemption, and entities should not assume blanket coverage; and
3. **note the limitation** for the record, so that the gap between the transparency afforded to consumers and that afforded to workers is a matter of visible policy rather than an unremarked artefact of the APP architecture.

Closing

We would be glad to expand on any of the above, including the practical form an AI-inventory / AI-BOM control takes in procurement. Thank you again for the opportunity to contribute.

Yours faithfully,

[redacted]
Dark Horse Governance
[redacted]

AI-use disclosure: this submission was drafted with the assistance of an AI writing tool for research synthesis and drafting. All legal sources were verified against primary materials, and the author is accountable for every claim and recommendation it contains.

Sources cited

- *Privacy Act 1988* (Cth) — APP 1.3, APP 1.7–1.9 (commencing 10 December 2026); s 7B(3) employee-records exemption and the s 6(1) definition of "employee record"; s 16C and APP 8 (cross-border disclosure).
- *Privacy and Other Legislation Amendment Act 2024* (Cth) — introducing APP 1.7–1.9.
- OAIC, *Automated Decision-Making Transparency Obligation (APP 1) Issues Paper* (18 May 2026) — Questions 7, 11, 12; Explanatory Memorandum paras 335–337 and the addendum (commercial-in-confidence, p 38).
- OAIC, *APP Guidelines Chapter 8 — cross-border disclosure* (v1.3, October 2025).
- OAIC, *Employee records exemption* (web guidance, December 2015).
- King, J., Klyman, K., Capstick, E., Saade, T., & Hsieh, V. (2025). *User Privacy and Large Language Models: An Analysis of Frontier Developers' Privacy Policies.* AAAI 2025; arXiv:2509.05382.