INDUSTRY SUBMISSION · OAIC ADM TRANSPARENCY ISSUES PAPER

Automated Decision-Making Transparency Obligation — APP 1.7–1.9

Submission on the Issues Paper

Oiva Developments Pty Ltd · ABN 93 693 621 457

Field	Value
Submission to	Office of the Australian Information Commissioner (OAIC) — privacyreformtranche1@oaic.gov.au
Consultation	ADM Transparency Issues Paper, 18 May 2026
Closing date	Monday 15 June 2026
Submitter	Oiva Developments Pty Ltd — ABN 93 693 621 457 · ACN 693 621 457
Authorised by	[redacted] Chief Visionary Officer + Chief Technology Officer
Endorsed by	[redacted] Chief Executive Officer · legal sign-off
Version	v0.1 — draft for CEO sign-off; for lodgement on or before 12 June 2026
Date	2026-05-23
Classification	Public — will be published by the OAIC as part of the consultation record
Prior OAIC engagement	Oiva Developments lodged a submission on the Exposure Draft of the Children's Online Privacy Code 2026 (COPC Phase 3 consultation, copc@oaic.gov.au, acknowledged by the OAIC Privacy Reform Implementation and Social Media Taskforce on 21 April 2026). References in this submission to 'our COPC submission' refer to that lodged document.
Correspondence	Correspondence regarding this submission should be directed to the CEO [redacted].

1. Opening statement

Oiva Developments Pty Ltd is an Australian-owned developer of a sovereign Child Care Management System (CCMS) for the Early Childhood Education and Care (ECEC) sector. The platform is currently in private pilot with Early Learning Management Pty Ltd (ELM), an approved provider operating approximately 74 long day care and outside school hours care services nationally.

We write in connection with our prior engagement on the Children's Online Privacy Code 2026. In that submission, we observed that s.28(2)(g) of the Exposure Draft — which includes automated decision-making transparency as an element of a child's right to request information about their data handling — would be strengthened by requiring three specific elements: (i) a plain-language explanation of the categories of input data; (ii) the confidence threshold below which a human review is required; and (iii) risk-tiered evidence-chain substantiation calibrated to the risk and volume of each decision category. The ADM transparency obligation under APP 1.7–1.9 is the Privacy Act equivalent of that s.28(2)(g) requirement, and this submission develops those proposals in that broader context.

We submit as a regulated entity. Our platform falls within the ADM obligation's scope because it uses personal information about children, families and educators in computer programs that make, or substantially and directly support, decisions that could reasonably be expected to significantly affect their rights or interests. The decisions in our domain include CCS subsidy entitlement (made under the A New Tax System (Family Assistance) Act 1999), enrolment approval (affecting access to a significant service), debt determination, and fee setting. We offer our perspective not as an observer but as an entity that has already designed for the obligation and that wishes to assist the OAIC to develop guidance that is workable at the engineering level.

We are explicit about our scale. Oiva operates a single pilot tenant with approximately 74 services, and our production experience is measured in months, not years. We do not propose our approach as the only way or the preferred way. We present it as one existence proof that the obligation can be implemented in a technically rigorous and commercially proportionate manner.

We confirm that this submission may be published in full on the OAIC's consultation record.

2. Summary of recommendations

Oiva Developments makes five recommendations for the OAIC's consideration in developing its Guidance.

No.	Recommendation	Issues Paper questions
R1	Adopt a two-branch test for 'substantially and directly related': a computer program is substantially and directly related to making a decision when it either (a) operates the decision fully without human involvement, or (b) provides the primary evaluative data or recommendation on which a human decision-maker substantially relies. Degree of reliance and ability to override are the dominant	Q1, Q2

No.	Recommendation	Issues Paper questions
	factors; nature of the output and integration into the workflow are secondary indicators.	
R2	Recognise children in ECEC services as a category of vulnerable person for the purposes of assessing the significance of a decision; and recognise access to an approved ECEC service as a significant service under APP 1.9. Guidance should include ECEC-specific examples to reduce interpretive uncertainty for the sector.	Q3, Q4, Q5
R3	Accept live-generated privacy policy fragments backed by runtime telemetry as satisfying the 'up to date' requirement in APP 1.3, read with APP 1.7–1.9, provided the fragment is version-controlled, approved by a named responsible officer before each publication, and its generation mechanism is disclosed in the policy itself.	Q10
R4	Require entities to disclose, for each class of ADM decision, whether human review is available and, if so, the conditions under which it applies. Where a human review threshold exists, the threshold itself (calibrated to the risk and volume of the decision class) should be described in accessible terms. This requirement, consistent with COPC R5 in our prior submission, is feasible for entities that implement a formal tier classification.	Q10
R5	Recognise the proportionality of enforcement-edge patterns for platform-type services. Where a regulated platform routes all AI-mediated calls through a single sovereign compliance layer, the resulting audit trail and disclosure telemetry constitute evidence of systematic compliance that the OAIC should recognise as at least equivalent to per-capability bespoke documentation.	Q10

3. Submitter context and scope relevance

3.1 Who we are

Oiva Developments is a small Australian software company building an AI-native Child Care Management System for the ECEC sector. Our platform integrates with Services Australia's Child Care Subsidy (CCS) through the PRODA B2B gateway, supports the National Quality Framework (NQF) and National Quality Standard (NQS), and is hosted entirely within the ap-southeast-2 (Sydney) region of Amazon Web Services. No child or family personal information leaves Australian jurisdictional boundaries in operation.

We have thirty-plus AI-mediated capabilities currently in development or live, spanning enrolment notice generation, fee suggestion, workforce forecasting, CCS reconciliation intelligence, and educator support. Each capability uses personal information — enrolment data, session data, educator credentials, family details — as inputs to computer programs that make or substantially support decisions affecting families, educators, and providers.

3.2 Why we are in scope

Our platform is squarely within the ADM obligation. We use personal information about individuals in computer programs that make decisions or substantially support decisions that could reasonably be expected to significantly affect their rights or interests. Specific decision categories and why we consider them in scope:

Decision category	APP 1.9 branch	Personal information used
CCS subsidy entitlement determination	Decision made under an Act (A New Tax System (Family Assistance) Act 1999)	Child CRN, date of birth, enrolment details, session hours, immunisation status
Enrolment approval or refusal by provider	Decision affecting access to a significant service (childcare)	Family details, child details, immunisation status, CCS eligibility flag
Debt determination (CCS debt notices)	Decision made under an Act (Family Assistance Law)	Payment history, session records, subsidy calculation inputs
Fee schedule suggestion for a family	Decision affecting rights under a contract or arrangement (the fee agreement)	Enrolment details, session pattern, family circumstances
Child safety incident escalation classification	Decision significantly affecting a child's rights under the NQF reportable conduct framework	Incident narrative, educator records, child records
Educator workforce rostering (AI-assisted)	Decision affecting an educator's rights under the Children's Services Award	Educator qualifications, availability, hours worked

We do not suggest all of the above are unambiguously in scope on current guidance. Some will require legal confirmation specific to each capability's implementation. We raise them to

illustrate the range of decisions that an ECEC platform must navigate and to signal that guidance with ECEC-specific examples would materially reduce uncertainty for our sector.

4. Responses to Issues Paper questions

Q1: Substantially and directly related to making a decision — ranking of factors

We rank the five listed factors as follows, on a 1–5 scale where 5 is most important:

Factor	Rating (1–5)	Rationale
Degree of reliance on the ADM system output	5	This is the most operationally determinative factor. Where a human decision-maker adopts the computer program's output with minimal independent assessment, the program substantially facilitates the decision. Where the decision-maker exercises independent professional judgement alongside the output, the degree of reliance is lower. We propose this as the primary diagnostic.
Ability and likelihood of human override over an ADM decision	4	Where human override is available but rarely exercised (because the output is trusted), this compounds the reliance factor. Where override is systematically exercised for a meaningful proportion of cases, the human contribution is more substantive. We note that architectural override mechanisms (a review button exists) and actual override rates can diverge significantly; guidance should direct attention to the latter.
Nature of the output (advisory vs determinative)	4	A determinative output leaves the decision-maker no discretion; an advisory output still requires independent evaluation. This correlates with reliance but is analytically distinct. A nominally advisory output that is in practice determinative (because no one has time to check it) should be treated as determinative.
Integration of ADM into the decision-making workflow	3	Deeper integration increases the risk that the output becomes default-adopted. However, integration per se is not dispositive — a well-integrated tool with trained reviewers can have lower effective reliance than a loosely integrated one that is rubber-stamped.
Transparency and explainability of outputs	2	Transparency is an important condition for meaningful human review but does not itself determine whether the program is substantially facilitating the decision. An opaque program might still be below threshold if its output is genuinely treated as one input among many.
Additional factor — Calibrated confidence tier (our proposal)	—	We propose a sixth factor: whether the entity has documented a formal human review threshold (a confidence level, risk tier, or HITL classification) below which human review is required. An entity that has formally classified its AI-mediated decisions into review tiers and enforces those tiers architecturally — as we have — has made the substantial-and-direct

Factor	Rating (1–5)	Rationale
		assessment concrete and auditable. We submit this should be weighted highly in the guidance.

Q2: Fictional edge case — NACIA / GPTea generative AI with human oversight

Would we consider NACIA's use of GPTea to be substantially and directly related to making decisions? Yes.

Applying our ranking: NACIA staff receive GPTea's summaries and eligibility recommendations before making decisions. If staff routinely adopt GPTea's recommendation — which the edge case does not preclude — the degree of reliance is high. The output is advisory in form but effectively determinative if staff lack the time or means to independently re-evaluate each candidate profile. Integration is deep (GPTea is used systematically, not occasionally). Even if transparency is limited (generative AI outputs are often difficult to audit), the substantiality test is met where reliance is high.

The presence of human oversight does not of itself remove the obligation. In our ECEC context, an educator who reviews an AI-generated enrolment notice draft before submitting it to Services Australia is exercising human oversight — but the AI is still substantially and directly related to the decision that is lodged. The OAIC's guidance should state explicitly that human oversight does not break the chain; it is a HITL Tier 2 arrangement (human-assisted but AI-primary-evaluative-data), not a Tier 3 arrangement (AI merely assistive, human independently leads).

> **Recommendation (R1).** Adopt a two-branch test: fully automated decisions (Tier 1) and AI-primary-evaluative-data decisions with human sign-off (Tier 2) are both in scope. Human oversight does not break the 'substantially and directly related' chain unless the human decision-maker independently evaluates the decision without material reliance on the ADM output. Degree of reliance and actual override rates are the primary diagnostic factors.

Q3: Factors increasing likelihood of significantly affecting rights or interests

We rank the listed factors as follows:

Factor	Rating	ECEC context
Vulnerable persons	5	Children in ECEC services are among the most vulnerable persons the Privacy Act's ADM obligation should protect. They cannot self-advocate; their privacy interests are mediated through parents, educators, and regulated providers. Guidance should explicitly recognise children as a category of vulnerable person.
Financial outcomes	5	CCS subsidy decisions directly determine the financial cost of childcare to families. Errors or unexplained automated decisions in CCS contexts have material, ongoing financial impact.

Factor	Rating	ECEC context
Sensitive information	4	ECEC records include sensitive information — child health, immunisation, NQF child safety incidents (Reportable Conduct Scheme), developmental observations. All are in scope for the sensitive information uplift.
Intrusive practices	3	Some AI-mediated decisions in ECEC (e.g. learning story generation from observation notes) are relatively non-intrusive. Others (incident escalation classification) are highly intrusive. The factor is relevant but must be assessed per decision class, not per platform.

Q4: Vulnerable persons for the purposes of significant effect

Children in ECEC services are a category of vulnerable person that the OAIC's guidance should explicitly recognise. The reasons are structural, not circumstantial: children cannot self-advocate; their personal information is managed by adults on their behalf under a regulated framework; the decisions made about them — enrolment, subsidy, learning assessment, child safety escalation — are consequential and cumulative over years.

Persons with parental responsibility who are themselves in vulnerable circumstances (experiencing family violence, financial hardship, or language barriers) are a secondary category. ECEC platforms serve many such families in the normal course of operations.

> **Recommendation (R2).** Guidance should recognise children in ECEC services as a category of vulnerable person. Decisions made using personal information about a child — or decisions affecting a child's access to an approved service — should carry a strong presumption that they could reasonably be expected to significantly affect rights or interests, absent countervailing factors specific to the capability.

Q5: Significant service or support

We rank the listed examples as follows:

Service or support	Rating	ECEC relevance
Access to educational services	5	Access to an approved ECEC service is access to an educational service. Under the National Quality Framework, approved services are regulated educational environments. Denying or limiting access — through automated enrolment rejection or fee-setting that makes attendance unaffordable — significantly affects the child's and family's rights.
Financial assistance	5	CCS subsidy is financial assistance under the Family Assistance Law. Automated decisions in the CCS pathway (eligibility, entitlement, debt) directly affect the financial assistance a family receives.

Service or support	Rating	ECEC relevance
Physical assistance/support	3	Not directly applicable in standard ECEC.
In home services	2	Not directly applicable in standard ECEC.
Access to essential banking and credit	1	Not applicable in ECEC.
Access to telecommunications	1	Not applicable in ECEC.
Access to essential utilities	1	Not applicable in ECEC.

Q6: Other legal frameworks relevant to 'rights and interests'

Two legal frameworks in the ECEC context are directly relevant and should be reflected in the OAIC's guidance:

- **Family Assistance Law (FAL):** decisions made under the A New Tax System (Family Assistance) Act 1999 and related instruments directly confer and affect rights — CCS entitlement, ACCS eligibility, debt. These are statutory rights with defined review pathways. The ADM obligation's disclosure requirement should extend to these pathways in guidance.
- **National Law / National Quality Framework:** the National Law (effective in all states and territories) imposes obligations on approved providers and regulates access to approved services. NQF child safety obligations under ss.175F and 175I (personal device ban, effective 27 February 2026) are a current example of sector-specific law that interacts with privacy and ADM obligations in the same way we identified in our COPC submission regarding the device-authorisation pathway.
- **National Employment Standards and Children's Services Award:** AI-mediated workforce decisions (rostering, leave management, qualification compliance) affect educators' rights under employment law. These decisions may be in scope where the AI program substantially facilitates a significant employment decision.

Q7: Fictional edge case — Daintree differential pricing

Would we consider the inflation of prices for Marco to have significantly affected his interests? Yes, on the facts as stated — not primarily because of the dollar amount, but because of the combination of factors: the pricing was personalised, automated, and opaque; it affected access to a service Marco intended to purchase; and the price differential (roughly 186 per cent uplift) is substantial in any ordinary sense.

On the amount question: we do not think a fixed dollar or percentage threshold is workable. A 5 per cent uplift on a $5 item is trivial; a 5 per cent uplift on a $100,000 professional services contract is not. The significant-service factor matters more: differential pricing on a children's educational service or a health service is more serious than differential pricing on a discretionary consumer product, because access to the service has a higher rights-bearing weight. We recommend guidance adopt a principles-based approach: does the pricing decision,

taken together with the nature of the service and the personal circumstances of the affected person, reasonably represent a significant impact on their interests?

On the importance of individual factors: dollar amount (4/5), significance of the product or service (5/5). The significance of what is being priced matters more than the absolute dollar delta.

Q8: Fictional edge case — Manscon / Opaque Window discriminatory targeted job ad

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

APP 1.9 makes clear that a decision includes refusing or failing to make a decision or do a thing. The failure to show Rafqa an employment opportunity is a 'doing a thing' — or more precisely, a failing to do a thing — that is substantially and directly related to a decision about her access to an employment opportunity. Employment opportunities are significant services in any sensible reading of the threshold. The discriminatory criterion (gender) makes this a decision that affects rights or interests adversely. The fact that no human at Opaque Window consciously decided not to show Rafqa the ad is the precise scenario the obligation is designed to capture: the computer program made the decision, and the entity 'arranged for' it to do so.

Guidance should state clearly that omissive ADM — failing to act, failing to serve, filtering out — is in scope as a 'doing a thing substantially and directly related to making a decision' where the filtering mechanism uses personal information and the outcome significantly affects rights or interests.

Q9: 'Arranged for' — scenarios requiring guidance

Two scenarios in our operational context require guidance:

- **Third-party integration:** our platform connects to external services (Services Australia PRODA, payment gateways, HR platforms). Where an external party's computer program makes or substantially supports a decision, and we have 'arranged for' its use by integrating it into our service — is the obligation ours, the external party's, or shared? Guidance should address the position of a regulated platform that orchestrates third-party ADM. Our current approach: we treat the obligation as ours for decisions we have integrated and in which our platform's data flows are used, regardless of where the model resides.
- **Multi-stage pipelines:** a CCS session report flowing from our platform to Services Australia, processed by a Services Australia decision system, results in a CCS payment decision. Our platform's computer programs make intermediate decisions (session eligibility flags, exception classifications) that substantially and directly feed the Services Australia decision. The chain of 'arranged for' across two regulated entities requires guidance. We submit the obligation should attach at every node where personal information enters a computer program in a way that substantially contributes to the ultimate decision — not only at the final decision point.

Q10: Extent of disclosure — what should be disclosed and how

This is the question where our operational experience is most directly relevant, and where we have the most developed views. We address the OAIC's five listed considerations in turn, then add three positions that our design experience has surfaced.

4.10.1 On the OAIC's five considerations

- **Clearly articulated in plain language:** we agree, without qualification. Our disclosure helper (an internal AI assistant) drafts the privacy-policy fragment from aggregated telemetry and targets a Flesch-Kincaid reading level of Year 9 or lower. The Director reviews and approves each draft. We propose this as one existence proof of a generation mechanism; we do not propose it as the only acceptable approach.
- **Structured to enable consumers to request further information:** we support a layered disclosure model. The public privacy policy fragment discloses categories and tiers. A family or educator who wants more should be directed to a data-subject-access channel. We are building this channel; we note it is architecturally non-trivial for a multi-tenant platform to respond to data-subject-access requests while maintaining tenant isolation and confidentiality obligations to providers.
- **Appropriately tailored — specific but not overwhelming:** guidance should distinguish between the granularity of 'kinds of decisions' (should be category-level, not per-capability) and the granularity of 'kinds of personal information' (should list information categories, not individual fields). A CCMS platform may have thirty-plus AI capabilities; forcing thirty-plus separate disclosures in the privacy policy would overwhelm readers. We propose that capabilities be grouped by decision category and information category, with the underlying capability-level detail held in an accessible register.
- **Organised so similar information is grouped:** we support decision-category grouping. In our implementation, the public disclosure groups by: (a) CCS and subsidy decisions; (b) enrolment and access decisions; (c) fee and financial decisions; (d) learning and developmental decisions; (e) safety and regulatory decisions; (f) workforce decisions.
- **Framed to allow information and the decision to be challenged:** we support this. Contestability is a design goal in our platform — every audit envelope preserves the evidence a family would need to challenge a decision, including the input data used and the human reviewer's identity for Tier 2 decisions.

4.10.2 Our three additional positions

> **Recommendation (R3).** Accept live-generated privacy policy fragments backed by runtime telemetry as satisfying APP 1.3 / APP 1.7 currency requirements. A live-generated fragment is updated automatically as the platform's AI use evolves, cannot drift from operational reality, and is version-controlled with a named approver. Guidance should provide a safe harbour for this approach, provided the entity discloses the generation mechanism in the fragment itself.

Our implementation: the Capability Register holds the authoritative classification of every AI-mediated capability (its HITL tier, decision category, and personal information categories used). The Transparency Service aggregates disclosure events from all live capabilities and generates a plain-language fragment for Director approval. The fragment is refreshed weekly during the

stabilisation period and monthly thereafter. Every published version is time-stamped and version-controlled. This means that from 10 December 2026, Oiva's privacy policy will contain a fragment that reflects its actual AI use as at the prior weekly refresh, not a static document drafted months earlier.

> **Recommendation (R4).** Require disclosure of human review availability and threshold. For each class of ADM decision, entities should be required to disclose: (a) whether human review is available before the decision takes effect; and (b) if so, the conditions under which it applies. Where a confidence threshold or risk tier determines when human review is triggered, that threshold should be disclosed in accessible terms. We first proposed this in COPC R5 in relation to s.28(2)(g). It applies equally under APP 1.7–1.9.

In our platform, Tier 2 decisions are those where the AI provides primary evaluative data and a named human approves before the decision takes effect. Tier 1 decisions are fully automated. The disclosure fragment will say, for each decision category, which tier applies. This gives a family or educator a meaningful signal: 'for fee suggestions, a human educator reviews and approves every AI-generated recommendation before it is presented to your family'; or 'for routine session data submission, the process is fully automated with tamper-evident logging that your provider can access on request.'

> **Recommendation (R5).** Recognise enforcement-edge patterns as a proportionate compliance approach for platform-type services. A sovereign platform that routes all AI-mediated calls through a single compliance layer — enforcing tier classification, PII redaction, rate limits, disclosure events, and audit envelope creation at one architectural boundary — generates systematic, auditable evidence of APP 1.7–1.9 compliance across all capabilities simultaneously. The OAIC's guidance should recognise that this pattern is at least equivalent in evidentiary weight to per-capability bespoke compliance documentation, and should not inadvertently favour approaches that are prohibitive for smaller regulated entities.

5. Architectural pattern — one existence proof

This section describes the approach we have taken to operationalise the ADM transparency obligation at our current scale. We do not propose this pattern as the only way, the preferred way, or a reference implementation for the sector. We present it as one existence proof that the obligation can be technically achieved in a sovereign Australian context with current technology, to help the OAIC calibrate what a workable compliance posture can look like.

Capability Register. An authoritative database records, for every AI-mediated capability: its HITL tier (Tier 1 fully automated, Tier 2 AI-primary-evaluative-data with human approval, Tier 3 AI-assistive below threshold), the kinds of decisions it makes or supports, and the kinds of personal information it uses. Changes to the register require joint approval by the Chief Visionary Officer and the CEO.

Policy enforcement edge. Every AI-mediated call routes through a synchronous enforcement boundary before any external invocation. The boundary: validates the declared tier against the register (preventing tier drift); runs a fail-closed personally identifiable information gate before any content reaches a large language model; emits a disclosure event (non-identifying) to an immutable store; and creates a tamper-evident audit envelope capturing the call's inputs, outputs, and human review outcome.

ADM transparency service. Disclosure events are aggregated by decision category and personal information category. A generation assistant drafts the plain-language privacy policy fragment from these aggregates. The CEO reviews and approves each draft before it is published. The published fragment is version-controlled and time-stamped. The public endpoint serves the live, approved fragment continuously.

Audit graph. Every action that touches personal information — user activity, government API calls, AI inference events — is logged into one of three correlated audit domains with tamper-evident storage (AWS S3 Object Lock in COMPLIANCE mode, seven-year retention, aligned with Family Assistance Law record-keeping). The audit log is queryable by natural language, producing cited evidence packs that can substantiate any AI-mediated decision to a parent, a provider board, or the OAIC.

Architecture enables the audit trail; it does not substitute for the consideration. Consistent with our COPC submission's R4, we note that architectural controls — fail-closed gates, default-minimal collection, HITL enforcement — are necessary conditions for operationalising the transparency obligation, but they are not sufficient. Sufficiency requires case-specific, documented consideration of what the AI is doing and why, with an audit trail accessible to the OAIC and, on request, to the individual. We offer our architecture as one means to generate and preserve that trail.

6. Closing statement

Oiva Developments supports the ADM transparency obligation as an important protection for individuals whose personal information is used in consequential AI-mediated decisions. In the ECEC sector in particular — where the affected individuals are children and their families, the decisions are regulated, and the volume of AI use is growing rapidly — clear guidance will reduce both interpretive uncertainty and compliance cost.

Our five recommendations are offered in the spirit of helping the Guidance be practically workable for entities that have designed for the obligation in good faith, while being robust enough to address entities that have not.

We welcome further engagement with the OAIC on any aspect of this submission. In particular, we would be pleased to:

- make our Capability Register schema, disclosure event specification, and audit envelope design available as reference material if the OAIC is developing worked examples for the Guidance;
- participate in any follow-up industry roundtable the OAIC convenes prior to the release of guidance (anticipated September 2026);
- contribute technical drafting on the live-generated fragment safe harbour (R3) and the enforcement-edge recognition (R5) if the OAIC wishes to develop specific guidance text; and
- support any independent audit exercise the OAIC undertakes, consistent with our view expressed in our COPC submission that Oiva's architecture is designed with external audit in mind.

We confirm that this submission may be published in full on the OAIC's consultation record.

Authorisation

Role	Name	Signature · Date
Chief Visionary Officer · Chief Technology Officer	[redacted]	
Chief Executive Officer	[redacted]	

Contact. [redacted] · Oiva Developments Pty Ltd · ABN 93 693 621 457 · ACN 693 621 457.
Correspondence regarding this submission should be directed to the CEO [redacted]