INDUSTRY SUBMISSION · OAIC PHASE 3 CONSULTATION

Children's Online Privacy Code 2026

Submission on the Exposure Draft

Oiva Developments Pty Ltd · ABN 93 693 621 457

Submission to	Office of the Australian Information Commissioner (OAIC) — copc@oaic.gov.au
Consultation	Phase 3 public consultation, 31 March – 5 June 2026
Submitter	Oiva Developments Pty Ltd — ABN 93 693 621 457 · ACN 693 621 457
Authored by	[redacted] — Chief Visionary Officer + Chief Technology Officer
Endorsed by	[redacted] — Chief Executive Officer · legal sign-off
Version	v0.2 — reinforced per red-team analysis · draft pending CEO sign-off
Date	2026-04-18
Classification	Public — will be published by the OAIC as part of the consultation record
Format	Written submission under s.26GC(1) of the Privacy Act 1988

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 support of the Children's Online Privacy Code 2026 as exposed on 31 March 2026, and in particular to confirm that Oiva's platform falls squarely within the Code's intended scope under section 5 Note 1. The examples enumerated in that Note — "applications that track early childhood development, family photo sharing applications, online school management systems that monitor student performance and internet-connected baby monitors" — describe precisely the category of service we operate. We therefore submit as a regulated entity, offering our architectural choices as one existence proof that the Code's intent is technically achievable — not as prescription, and not as a model for the sector to adopt.

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. The architectural patterns described in this submission work at our scale; they may need adaptation at much larger or much smaller operating scales. We support a principles-based Code that sets outcomes rather than prescribing technology, and we are conscious that our submission should not push in the opposite direction.

This submission is structured in two halves. The first half (sections 2 to 4) responds to the Exposure Draft clause by clause on the provisions we consider highest-leverage for the ECEC sector. The second half (sections 5 to 7) proposes concrete architectural patterns that operationalise the Code's requirements, together with one substantive amendment and two worked examples. Our recommendations are summarised below.

We welcome the OAIC's invitation to make our architecture, threat modelling and privacy impact assessments available as public reference material. We believe that children's privacy in the ECEC sector is best served by regulators, providers and platforms pooling technical knowledge of what good looks like, rather than each re-deriving it in isolation.

2. Summary of recommendations

Oiva Developments makes nine recommendations to the OAIC for consideration in the final Code.

No.	Recommendation	Clause(s)
R1	Retain the scope provision in s.5 and the Note 1 examples. They correctly capture ECEC CCMS platforms as regulated entities.	*s.5 · scope*
R2	Clarify in the Explanatory Statement that "primarily concerned with the activities of children" includes platforms whose end-users are educators and parents where the data subject is a child.	*s.5(b)*
R3	Permit s.8 (age ascertainment) to be satisfied by existing verified enrolment records held by the same entity for a pre-existing lawful purpose, subject to (i) no additional collection arising from the reliance and (ii) the reliance being documented in the entity's privacy policy.	*s.8*
R4	Retain the best-interests-of-the-child test in ss.10 and 11 as a mandatory condition. Clarify that architectural controls (fail-closed gates, HITL envelopes) are necessary but not sufficient evidence of operationalising the test — sufficiency requires case-specific documented consideration with an audit trail.	*ss.10, 11*
R5	Strengthen s.28(2)(g) on automated decision-making transparency by requiring a plain-language explanation of model inputs, the confidence threshold at which human review is required, and risk-tiered evidence-chain substantiation (depth calibrated to the risk and volume of each decision type, with the tiering itself documented and auditable).	*s.28(2)(g)*
R6	In s.32 (destruction requests), expressly distinguish between destruction of platform-held data and destruction of records that providers are required to retain under the Family Assistance Law or the National Law. We propose specific drafting.	*s.32*
R7	In s.38 (privacy impact assessments), specify that PIAs must address AI/machine learning	*s.38*

	components separately where present, covering model selection, data flow, training data provenance, and human-in-the-loop escalation.	
R8	Consider a narrowly-drawn companion provision capturing provider-held child records (under the National Law) where those records are processed by third-party platforms. The current Code captures the platform but not the relationship.	*Scope · extension*
R9	We support the inclusion of an explicit independent audit power for the Commissioner in the final Code or Explanatory Statement. Our architecture is designed for audit and we would welcome scrutiny.	*New · Commissioner powers*

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.

3.2 Why we are in scope

Our platform is a designated internet service under the Online Safety Act 2021. It is primarily concerned with the activities of children in the sense that every record it handles is, at its origin, about a child enrolled in an ECEC service. Our core capabilities include: enrolment management; educator observation and learning documentation (the EDIE capability, grounded against the Early Years Learning Framework v2.0 and My Time Our Place v2.0); child safety and incident management (the GuardianIQ capability, including handling of Reportable Conduct Scheme notifications); and reporting to Services Australia under the CCS. All are regulated interactions involving personal information about children.

Note 1 of s.5 of the Exposure Draft enumerates examples of services primarily concerned with the activities of children. Two of the four examples — "applications that track early childhood development" and "online school management systems that monitor student performance" — describe precisely our platform. We welcome this clarity and respectfully encourage the OAIC to retain these examples in the final Code.

3.3 Why our perspective may be useful

The ECEC sector handles child personal information at an exceptional density per end-user. A child in long day care typically has several hundred observations, photographs, communications, incident

records and assessment notes created about them per year, stored for a minimum of seven years under the Family Assistance Law. These records are created by educators (workers), consumed by parents (persons with parental responsibility), and reported to regulators (Services Australia, ACECQA). The Code's application to this sector is therefore both high-stakes and architecturally distinctive, and we offer this submission to help the OAIC understand how its provisions interact with the technical and regulatory environment we operate in.

4. Clause-by-clause responses

4.1 Section 5 — Additional entities bound by this Code

We support the scope provision as drafted. The three-element test (provider of a social media, relevant electronic or designated internet service; primarily concerned with the activities of children; not providing a health service) is workable, and the exclusion in s.6 of carriage service providers is appropriate.

Recommendation (R2). We recommend clarifying in the Explanatory Statement that the phrase "primarily concerned with the activities of children" captures platforms whose immediate end-users are adults (educators, parents) but whose data subjects are children. ECEC platforms fall into this category: an educator operates the platform, but the records produced are about a child. Without this clarification a strict reading could exclude platforms where children do not directly interact with the user interface, which we believe would not be the OAIC's intent given the examples in Note 1.

4.2 Section 8 — Ascertaining the age of end-users

We support the default-minimal collection principle embodied by s.8. In the ECEC context, however, the age of every child in the platform is known with high confidence at the point of enrolment — it is a mandatory field under the National Law and is verified against Medicare for CCS eligibility. Re-ascertaining the age of the child via the platform would be duplicative and would introduce, rather than reduce, privacy risk.

Recommendation (R3). Permit s.8 to be satisfied by reliance on existing verified enrolment records, subject to three conditions that preserve the privacy benefit the section intends: (i) the records are held by the same entity for a pre-existing lawful purpose; (ii) no additional personal information is collected as a consequence of the reliance; and (iii) the reliance is documented in the entity's privacy policy so that a child or person with parental responsibility can see how age has been established. Suggested supplementary drafting: "An entity satisfies subsection (1) if the entity relies on age information derived from enrolment records verified under the A New Tax System (Family Assistance) (Administration) Act 1999 or equivalent, the entity holds those records for a pre-existing lawful purpose, no additional personal information is collected for the purpose of this subsection, and the reliance is disclosed in the APP privacy policy."

4.3 Sections 10 and 11 — Best interests of the child

We strongly support the mandatory best-interests test on collection, use and disclosure of children's personal information. We note that this test can be operationalised architecturally as well as by policy. In our platform:

- Collection is default-minimal and contextual. Observations, photographs and learning records are only captured when an educator initiates the act and the system has verified the device is service-authorised under National Quality Framework s.175F and s.175I.
- Use is bounded by a structural exclusion list. Financial and subsidy data is explicitly excluded from any AI inference context; this is enforced by Lambda assertions and by a minimum of fourteen continuous-integration contract tests that fail the build if the exclusion is bypassed.
- Disclosure requires both express consent (per s.13) and human-in-the-loop approval via an AWS Step Functions waitForTaskToken pattern. No AI-generated content affecting a child is auto-published; every artefact crosses a human approval point.

Recommendation (R4). Retain ss.10 and 11 as drafted. We propose supplementary guidance in the Explanatory Statement that clarifies the relationship between architectural controls and the best-interests test: architectural controls — fail-closed gates, default-minimal collection, human-in-the-loop approval — are a necessary condition for operationalising the test, but they are not sufficient. They evidence the absence of certain risks. Sufficiency requires case-specific, documented consideration of the child's interests for each category of collection, use or disclosure, with an audit trail accessible to the Commissioner and, on request, to a person with parental responsibility. Architecture enables the audit trail; it does not substitute for the consideration.

4.4 Section 23 — APP privacy policy requirements

We support the requirement for a child-directed APP privacy policy including non-text material, and will comply. We note that in the ECEC context the more common pattern is that parents (persons with parental responsibility) will read the policy on behalf of a child under 15. We encourage the OAIC to recognise this in guidance: both a child-directed and a parent-with-parental-responsibility-directed version of the policy should be produced, with clear signposting between them.

4.5 Section 28 — Right to request information about handling

We welcome s.28(2)(g)'s explicit inclusion of automated decision-making transparency, particularly given the $50 million civil penalty for APP 1.7–1.9 breaches commencing 10 December 2026 under the Privacy Act amendments. We believe s.28(2)(g) can be strengthened in three specific ways.

Recommendation (R5). We propose s.28(2)(g) should require a plain-language explanation of: (i) the categories of input data used in the automated decision; (ii) any confidence threshold below which the automated decision is escalated to a human; and (iii) risk-tiered evidence-chain substantiation, with depth calibrated to the risk and volume of the specific automated decision category. Routine pedagogical tagging (low risk, high volume) should not require the same evidence depth as an automated child-safety escalation (high risk, low volume). The Code should permit tiered substantiation provided the tiering is itself documented, justified and auditable. This calibration respects the practical limits of audit storage at scale while preserving traceability where it matters most.

The third element is architecturally achievable. Our platform maintains a three-domain audit graph (user activity, communications, AI inference events) with tamper-evident storage via AWS S3 Object Lock in COMPLIANCE mode, aligning with Family Assistance Law record-keeping requirements. We describe this further in section 5 as one approach, not as a required approach.

4.6 Section 32 — Request to destroy personal information

Section 32 requires careful drafting to interact cleanly with the Family Assistance Law and the National Law, which impose minimum retention periods (seven years for attendance and session records; variable for safety records). The Exposure Draft partially addresses this via s.32(3)(d) (retention required by Australian law) but the interaction in ECEC is more nuanced than a binary retain-or-destroy.

A worked example illustrates the issue. Suppose a parent requests destruction of all learning story narratives about their child held by an ECEC platform. Some narratives have been exported to the child's portfolio (which the parent has received); some are retained as evidence for the provider's Quality Improvement Plan (QIP) under NQS Element 1.3.2; some underpin CCS session decisions. The parent's right to destruction should be respected, but the platform cannot lawfully delete records that substantiate regulatory obligations without exposing the provider to compliance risk.

Recommendation (R6). We propose amending s.32 to distinguish between (a) destruction of the record itself, where retention is not required, and (b) destruction of the link between the record and the child's identity, where the underlying record is required to be retained. The latter preserves the substance of the parent's request (the child is no longer identifiable from the retained record) while respecting the retention obligation. We would suggest supplementary drafting along the lines of: "Where destruction of the information is not permitted under subsection (3), the entity must destroy any link between the information and the child's identity to the extent practicable, and describe this in the written notice under subsection (4)."

4.7 Section 33 — Notification of monitoring mechanisms

We support s.33. In our context, we note that the NQF Child Safety Reforms effective 27 February 2026 under NQF s.175F and s.175I make it a criminal offence for staff to capture images of children on personal devices; only service-authorised devices are permitted. Our platform embeds a service-authorised device pathway as an architectural default, and we would encourage the OAIC to reference this parallel regulatory regime in the Explanatory Statement as an example of how sector-specific law and the Code reinforce each other.

4.8 Sections 38 and 39 — Privacy impact assessments

We strongly support the mandatory PIA regime under s.38 and the public register requirement under s.39. We note that s.38(2)(e) requires "specific information about how the entity complies with this Code and the Act". Where AI and machine learning are involved, this specific information should address elements that are specific to AI: model selection and version; training data provenance and whether the model was fine-tuned on personal information; inference data flow and cross-border implications; confidence thresholds; and human-in-the-loop escalation paths.

Recommendation (R7). We propose s.38(2) be supplemented with a new paragraph (g): "Where the service involves artificial intelligence or machine learning, an assessment that separately addresses model selection, training data provenance, inference data flow, and the conditions under which automated decisions are escalated to human review." This aligns the Code with the Australian AI Ethics Framework and with emerging international practice (EU AI Act Article 27 risk assessment).

4.9 On independent audit of entities bound by this Code

We wish to make a clear and unprompted recommendation on a matter the Exposure Draft does not directly address. Self-reported compliance, including self-authored PIAs under ss.38 and 39, can only carry the Code so far. The Commissioner's ability to examine entities' compliance independently — whether through the existing powers under the Privacy Act or through a new power specific to this Code — is foundational to public confidence in the regime.

Recommendation (R9). We support, and would welcome, the inclusion of an explicit independent audit power for the Commissioner in either the final Code or the Explanatory Statement. Our platform architecture is designed with external audit in mind: tamper-evident logs, complete evidence chains, natural-language audit query for regulator access, and signed PDF export. We are pleased to be audited, and we think the sector as a whole is well-served by the prospect of independent verification rather than unqualified self-assertion.

5. Architectural patterns — one existence proof

This section describes six technical patterns that we have implemented to operationalise the Code's requirements at our current scale. **We do not propose these patterns as the only way, the preferred way, or a reference implementation for the sector.** We present them as one existence proof that the Code's intent is technically achievable with currently available technology in a sovereign Australian context, and to help the Commissioner calibrate what a set of workable controls can look like in practice. Other entities will reach the same objectives through different technical choices appropriate to their scale and context, and the Code should remain principles-based to permit this.

5.1 Three-layer tenant isolation

Personal information about children is protected by three independent layers. First, Amazon Cognito issues a JSON Web Token on authentication containing the tenant identifier. Second, every server-side function verifies the tenant binding against the requested resource. Third, the Aurora PostgreSQL database applies row-level security via FORCE RLS, setting the tenant context for every database session and rejecting any cross-tenant read or write attempt at the storage layer. A failure at any single layer is contained by the others.

5.2 Three-domain audit graph

Every action that touches a child's personal information is logged into one of three correlated log domains: the User Activity Log (authenticated user actions), the Communications Log (inbound and outbound messages and government API calls), and the AI Inference Log (every AI-influenced decision, including the prompt, the response, the confidence score, and the human approval step). Logs are stored in Amazon S3 Object Lock in COMPLIANCE mode with a seven-year retention period to align with Family Assistance Law record-keeping. The audit graph is queryable by natural language, producing cited evidence packs that can substantiate any child's data handling history to a parent, a regulator, or a court.

5.3 Fail-closed PII gate before AI inference

Every text payload passed to a large language model is first scanned by Amazon Comprehend's PII detection service. If PII is detected above configured thresholds, the gate fails closed: no inference

request is made, and the originating user receives a machine-readable reason. This pattern operationalises the default-minimal principle at the AI boundary, where data leakage risk is highest.

5.4 Structural exclusion of financial data from AI

Financial and subsidy data are structurally excluded from all AI inference contexts by design. This is enforced not by policy but by code: the Lambda functions that construct AI prompts are unit-tested against a minimum fourteen contract tests that verify no financial fields reach the prompt. The continuous-integration pipeline fails the build if the exclusion is broken. The result is a hard architectural boundary that cannot be accidentally breached in production.

5.5 Human-in-the-loop by default

No AI-generated content that affects a child is auto-published. Every AI-drafted learning story, every proposed child-safety notification, every classification that would alter a child's record enters an AWS Step Functions state machine with a waitForTaskToken step. The state machine cannot progress past the human review point without an explicit approval, amendment or rejection. This is the architectural expression of s.28(2)(g) transparency.

5.6 Sovereignty invariant

The platform's deployment manifest enforces ap-southeast-2 (Sydney) region as an invariant at infrastructure-as-code synthesis time. No operator, developer or system administrator can override this constraint without rewriting the deployment manifest and triggering a new build. Personal information about children never leaves Australian jurisdictional boundaries as a consequence of this design choice, which aligns with the Code's s.26 cross-border provisions.

6. Responding to what children told the OAIC

The Phase 1 consultation report identified six things that 337 children, young people and parents wanted organisations to do better. We find each of the six maps directly to an architectural or interaction pattern. We summarise the mapping here as evidence that children's stated preferences can be met technically, not only through policy exhortation.

Children told the OAIC	How the architecture responds
Ask for permission!	Every collection, use or disclosure passes through an explicit consent moment (ss.13–21). Our human-in-the-loop envelope means that consent is enforced architecturally, not merely requested in a policy.
Give more explanations!	The AI-assist layer presents plain-language explanations with cited evidence whenever an action is taken that affects a child's record. Every automated decision points back to the inputs that produced it.
Make privacy policies short and simple!	Section 23 of the draft Code mandates this, and we support it without qualification. We will publish a child-directed policy with non-text material alongside a parent-directed version.
Make important information and buttons easier to find!	Opt-out from direct marketing (s.29), withdrawal of consent (s.17), and destruction requests (s.32) are each surfaced as first-class actions with persistent visibility — not buried in settings menus.
Help children when they have questions!	Section 36's inquiry and complaints process is implemented as a visible channel reachable from every screen, with pseudonymous and anonymous options where lawful.
Make it fair!	The fairness question is partly answered by the structural exclusion of financial information from AI contexts, and by the Code's s.11(3) prohibition on bundled direct marketing consent. We support both.

7. A scope extension the OAIC may wish to consider

We have confined recommendations R1 through R7 to the current scope of the Code. This section proposes a narrowly-drawn extension (R8) that we believe would materially improve the Code's effectiveness in the ECEC sector without over-reaching.

The gap. The Code binds the platform but not the relationship between the platform and the approved provider. In ECEC, the approved provider is the legal controller of the child's records under the National Law. The platform processes those records on the provider's behalf. The Code's requirements on the platform are clear, but if a provider misuses the platform or configures it improperly, the child's privacy can still be compromised despite the platform itself being Code-compliant. The Code does not currently extend to provider-side governance of platform use.

A narrow drafting option. Consider a new provision (for example, a s.40A) requiring entities bound by the Code to publish, as part of their s.39 register, a model Data Processing Agreement (DPA) or equivalent that any controller using the platform must sign. The DPA would bind the controller to use

the platform's privacy features as designed (for example, not to attempt to bypass the human-in-the-loop envelope). This reaches the provider-side governance gap without expanding the Code's personal scope.

We flag this as a suggestion, not a firm recommendation, because we recognise it may exceed the OAIC's current statutory authority under s.26GC. If the OAIC judges the extension appropriate, we would be pleased to contribute drafting.

8. Closing

Oiva Developments supports the Children's Online Privacy Code. It is substantively the right framework for the moment, and it arrives at a time when the ECEC sector in particular needs a clear floor on children's privacy protections. Our recommendations in this submission are offered in the spirit of helping the final Code operate as intended in one of the sectors most visibly affected.

We welcome further engagement with the OAIC on any aspect of this submission. In particular, we would be pleased to: (a) make our architecture documentation available as public reference material under the Code; (b) participate in any follow-up industry roundtable the OAIC convenes during the Phase 3 consultation; (c) contribute technical drafting on ss.28, 32, 38 and on the DPA extension at s.40A as outlined above; and (d) support the design of any independent audit power that the Commissioner chooses to introduce.

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

Authorisation

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

[redacted]

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