SUBMISSION TO THE OFFICE OF THE AUSTRALIAN INFORMATION COMMISSIONER

Exposure Draft: Privacy (Children's Online Privacy) Code 2026

Public Consultation — 31 March to 5 June 2026 | Submitted May 2026

Field	Detail
Submitter	[redacted]
Capacity	Individual — attorney and child internet safety infrastructure developer
Initiative	[redacted]
Contact	[redacted]
Submitted to	copc@oaic.gov.au

Executive Summary

This submission is made in an individual capacity by [redacted] J.D., an attorney and child internet safety infrastructure developer based in the United States. The submitter is developing The Reservoir, an additive DNS namespace architecture that creates a structurally isolated, pre-verified internet environment for children — a complement to application-layer regulatory codes such as the one proposed here.

This submission is structured in two parts. Part One engages directly with the specific provisions of the Exposure Draft, offering observations on their strengths and the structural limits of each. Part Two advances a proposal that goes further than the Code can reach: the formation of a multi-stakeholder coalition of child safety organisations, ISPs, and DNS infrastructure providers to apply for an ICANN community generic top-level domain (gTLD) as a structural complement to the Code.

The central argument is this: the Code is necessary and well-designed for its purpose. It sets enforceable obligations on covered entities handling children's data on the open internet. But it cannot solve the underlying structural problem — that the open internet's architecture makes harmful data collection the default, and regulatory compliance the exception. The proposal in Part Two addresses that structural problem directly. It does not replace the Code. It makes the Code more effective by changing the environment in which it operates.

> *The Code governs how covered entities must behave within a harmful architecture. This submission asks the OAIC to consider also encouraging an architecture within which harmful behaviour is structurally precluded.*

PART ONE — OBSERVATIONS ON THE EXPOSURE DRAFT

PART ONE
Observations on the Exposure Draft

Section 8 — Age Assurance

s.8 **Ascertaining the age of end-users of an entity's service**

What the Code does

Section 8 requires entities to take steps that are reasonable in the circumstances to ascertain the age of end-users before collecting personal information. The standard is risk-proportionate: higher-risk data collection requires greater certainty about age. An entity may avoid the requirement altogether by applying child-level protections to all users regardless of age.

This is a sound and practical framework. The risk-proportionate approach avoids over-burdening low-risk services while ensuring high-risk platforms — those most likely to harm children — face the highest verification obligations. The safe harbour for universal child-level treatment is particularly sensible: it incentivises a protective default without mandating expensive age-verification infrastructure across all services.

Where the Code falls short

The Code's age assurance obligation is reactive — it applies at the point of data collection on a service a child has already reached. By the time a covered entity ascertains a user's age, the child has already navigated to the service, been exposed to its interface, and in many cases provided information. Age assurance after contact is protective for data handling purposes but does not prevent the child from reaching a harmful environment in the first place.

The Code also does not address how age assurance information can flow between services and infrastructure layers. A parent who has verified their child's age with one ISP or platform cannot easily port that verification to another service. The child re-encounters the age assurance process at every platform boundary — creating friction without meaningful cumulative protection.

The structural complement

The additive namespace architecture described in Part Two addresses age assurance at the infrastructure layer. A household that has opted into the verified namespace has, by definition, pre-screened every service in that namespace for child-appropriateness. Age assurance is embedded in admission criteria — no service enters the namespace unless its content and data practices have been pre-verified against the Content Metadata Standard. The parent's single opt-in decision replaces hundreds of per-service age assurance encounters.

> *Recommendation: The Code's Explanatory Statement should recognise household-level DNS namespace opt-in as a complementary age assurance mechanism — one that operates prospectively, at the infrastructure layer, before a child reaches a service rather than after.*

Section 9 — Privacy by Default

s.9	**Entities may collect personal information about a child by default only if strictly necessary**

What the Code does

Section 9 is one of the Code's most important provisions. It requires entities to implement technical and organisational measures ensuring that, by default, only personal information strictly necessary to provide the core service is collected, used or disclosed. Targeted advertising, personalised recommendations, and service improvement data collection must not be enabled by default — they require affirmative opt-in.

This provision directly addresses one of the most significant structural harms in the current children's online environment: the design of services to maximise data collection as a default state, with privacy as an opt-in feature buried in settings menus. The strict necessity standard is the right threshold. The requirement that controls be clear, simple and easily accessible is appropriately prescriptive.

Where the Code falls short

Privacy-by-default obligations apply after a service is designed, deployed, and already collecting data. Enforcement is necessarily retrospective: the OAIC investigates non-compliance after harm has occurred. The Code does not create a mechanism for prospective verification that a service is actually built to the privacy-by-default standard before it reaches children.

The Code also leaves the definition of 'strictly necessary' to each entity's own judgment, subject to regulatory review. This creates compliance uncertainty for well-intentioned entities and gaming opportunity for bad actors. Without a shared, machine-readable standard for what counts as strictly necessary in the context of children's services, the privacy-by-default obligation is difficult to operationalise consistently across covered entities.

The structural complement

The Content Metadata Standard developed for the proposed additive namespace requires each registrant to declare its data collection practices across a structured taxonomy — no collection, opt-in only, anonymised, or identifiable with parental consent. Admission to the namespace is conditioned on that declaration being accurate and auditable. A service that cannot truthfully declare one of the first two categories is ineligible for admission. Privacy by default becomes a condition of structural access, verified prospectively, not a compliance obligation applied retrospectively.

> *Recommendation: The OAIC should encourage development of a national structured metadata standard for covered entities, enabling machine-readable declaration of data collection practices that can be independently verified at the point of service admission to a verified namespace — and enforced in real time by household DNS resolvers.*

Sections 10–11 — Best Interests of the Child

s.10–11	**Collection, use and disclosure must be consistent with the best interests of the child**

What the Code does

Sections 10 and 11 introduce a best interests of the child standard as the overarching test for all collection, use and disclosure decisions. This is the Code's most significant structural innovation — it imports a child welfare principle from family law and child protection frameworks into privacy regulation, requiring entities to ask not merely whether data handling is lawful but whether it serves the child's interests.

This is the correct framing. Children are not small adults with reduced privacy interests. Their developmental stage, their vulnerability to manipulation, and the long-term consequences of early datafication require a protective standard that goes beyond adult consent models. The best interests test provides that standard.

Where the Code falls short

The best interests test, as drafted, is applied by the entity whose commercial interest is in maximising data collection. The entity is simultaneously the data collector and the primary best-interests assessor — a structural conflict of interest that regulatory oversight addresses retrospectively but cannot eliminate prospectively. Well-resourced entities will apply the test carefully. Entities with weaker compliance cultures will apply it loosely. The standard is only as strong as the weakest compliant entity.

The Code also does not define the best interests standard with sufficient specificity to enable consistent application across the enormous range of covered entities — from global social media platforms to small educational technology developers. Without guidance on how the test applies to specific data collection scenarios, the obligation risks becoming a box-ticking exercise.

The structural complement

The additive namespace model converts the best interests test from a self-assessed compliance obligation into a structural admission condition enforced by a multi-stakeholder governance body. A registry governance board — comprising child welfare organisations, privacy experts, and parent representatives — applies the best interests standard prospectively at the point of namespace admission. Services that fail the test are not admitted. The entity is not the assessor of its own compliance. This is not a replacement for

the Code's best interests obligation — it is a prospective enforcement mechanism that makes the Code's retrospective enforcement more effective by reducing the population of non-compliant services that reach children.

> *Recommendation: The OAIC should recognise multi-stakeholder infrastructure governance bodies that apply a child best-interests standard as a condition of registry admission as a complementary enforcement mechanism — and consider whether services admitted to a verified namespace should receive a rebuttable presumption of best-interests compliance under the Code.*

Section 13 — Parental Consent for Under-15s

s.13	Consent given by child or person with parental responsibility

What the Code does

Section 13 establishes 15 as the age below which parental consent is required for the collection, use or disclosure of a child's personal information. Below that threshold, the entity must obtain consent from a person with parental responsibility and take reasonable steps to verify that the consenting person actually has that responsibility. Where parental consent is obtained, the entity must provide the child with an age-appropriate notice of the consent and its terms.

The age-15 threshold is well-calibrated. It aligns with the OAIC's existing guidance on children's presumed age of consent and reflects research on child brain development and adolescent decision-making capacity. The requirement to verify parental responsibility — not merely accept a checkbox — is appropriately rigorous. The child notification requirement is a meaningful innovation: it respects the child's developing agency while maintaining the parental consent structure.

Where the Code falls short

Section 13's consent architecture is per-service. A parent must navigate a consent flow for every platform their child encounters. In practice, a child who uses ten services requires ten separate parental consent processes — each with its own notice format, consent period, and withdrawal mechanism. No parent operates at that level of granularity. The practical result is that parental consent becomes a compliance formality rather than a genuine protective mechanism.

The per-service consent model also creates an asymmetry of burden: the friction of consent falls on parents, while the commercial incentive to obtain consent — and to design consent flows that minimise friction for the entity while technically satisfying the requirement — falls on covered entities. Dark patterns in consent design are specifically prohibited under section 21, but detection and enforcement of dark pattern violations is difficult at scale.

The structural complement

The additive namespace model operationalises parental consent at the infrastructure layer. A parent who opts their household into the verified namespace has made a single, informed, affirmative choice about the class of services their child can access. That single decision — made once, at account level, covering all devices — is structurally more protective than hundreds of per-service consent flows. The namespace admission criteria substitute for the per-service consent negotiation: every service in the namespace has already been vetted against standards that, by definition, serve the child's best interests.

> *Recommendation: The Code's Explanatory Statement should explicitly recognise household-level DNS namespace opt-in as a valid mechanism for giving effect to the protective intent of section 13's parental consent requirement — acknowledging that infrastructure-layer parental choice is more durable than per-service consent flows for the cohort of children under 15.*

Sections 22–24 — Transparency and Privacy Notices

s.22–24	Transparency, privacy policy requirements, and notification of collection

What the Code does

Division 3 requires entities to produce child-directed privacy policies that are clear, concise, age-appropriate, and accessible. Policies must avoid legal and technical language, use non-text material where appropriate, and distinguish between necessary and optional data collection. Collection notifications must meet the same standards. Privacy practices must be reviewed and updated at least annually.

These requirements address a genuine harm: the current state of children's privacy notices is characterised by verbose, legalistic documents that no child — and few parents — read or understand. The age-appropriate standard (calibrated to 10–12 year olds where no specific age range is targeted) is a practical and sensible default.

Where the Code falls short

Transparency obligations, however well-designed, are documentary in nature. They require covered entities to produce accurate documents and deliver them to children and parents. They do not ensure that children and parents read, understand, or act on those documents. Research consistently shows that privacy policies — even simplified ones — do not meaningfully inform user behaviour. The transparency obligation is necessary for accountability but insufficient as a protective mechanism.

The per-service nature of transparency obligations also creates information overload. A child who uses ten services encounters ten privacy policies. Even if each is age-appropriate and well-designed, the cumulative cognitive burden on parents is unmanageable. No household can meaningfully exercise informed consent across the full range of digital services a child uses.

The structural complement

The Content Metadata Standard creates machine-readable transparency that does not depend on parents reading documents. Structured metadata — declared by the service, verified at registry admission, and enforced in real time by the household DNS resolver — provides operational transparency rather than documentary transparency. The parent's resolver knows what data a service collects before the child visits it. Parental configuration of resolver preferences — set once, enforced continuously — replaces per-service notice review.

> *Recommendation: The OAIC should encourage development of a national structured metadata standard for covered entities that creates machine-readable, resolver-enforceable transparency — complementing the Code's document-based transparency requirements with operational protections that do not depend on parental reading behaviour.*

Section 32 — Right to Destruction

s.32 **Request to destroy personal information about a child**

What the Code does

Section 32 creates a right for children, and persons with parental responsibility, to request destruction of a child's personal information. This is one of the Code's most significant innovations — and one of the most difficult to enforce. The right to destruction addresses the problem of data persistence: information collected about a child at age 8 remaining in a data ecosystem at age 28, with no mechanism for the now-adult to remove what was collected without meaningful consent.

Where the Code falls short

Destruction rights are retrospective — they apply to information already collected. They do not address the structural incentive to collect as much data as possible before a destruction request is made. An entity that collects aggressively and destroys on request has complied with the Code. An entity that never collects unnecessary data in the first place has done something qualitatively different — and the Code does not adequately distinguish between these two models.

Destruction requests also face significant practical challenges: data held by third-party processors, data that has been aggregated or de-identified, and data that has been shared across jurisdictions may not be practically destroyable even where the covered entity acts in good faith.

The structural complement

The additive namespace model addresses data minimisation at the point of architecture rather than at the point of destruction request. Services admitted to the namespace must

declare that they collect only data strictly necessary to provide their service. Data that is never collected cannot be the subject of a destruction request — and does not carry the risks of persistence, re-identification, or cross-border disclosure that section 32 is designed to address. The namespace is a data minimisation architecture, not a data remediation architecture.

> *Recommendation: The Code's guidance should explicitly distinguish between entities that collect minimally by architecture (as required by the namespace model) and entities that collect aggressively and destroy on request — and should develop regulatory incentives that reward the former model as the higher standard of compliance.*

PART TWO
Going Further — A Structural Proposal

The Structural Problem the Code Cannot Solve

The observations in Part One share a common thread. The Code is an application-layer instrument. It regulates how covered entities behave within an internet architecture that was not designed for children and is not structurally safe for children. Every provision of the Code — age assurance, privacy by default, best interests, parental consent, transparency, destruction rights — is a regulatory correction applied to a harmful default.

The harmful default is this: on the open internet, data collection is the architecture, and privacy is the exception. Services are designed to collect as much data as possible, to monetise that data through advertising and profiling, and to make opting out as difficult as possible. Children arrive in this environment without meaningful protection because the environment was not designed with them in mind. The Code makes covered entities behave better within that environment. It cannot change the environment itself.

> *Every existing child internet safety mechanism — DNS filters, parental controls, content moderation, age-gating, regulatory codes — is a correction applied to a harmful default. None of them changes the default. This proposal does.*

The Additive Namespace Architecture

The Reservoir is a proposed ICANN community generic top-level domain (gTLD) designed to create a structurally isolated, pre-verified internet namespace for children. It inverts the fundamental architecture of child online safety.

Subtractive vs. Additive — The Core Distinction

Every existing child safety tool is subtractive: it begins with the open internet — an environment of unlimited, largely harmful-by-default content and data practices — and

attempts to remove or block what is harmful. This approach is reactive, incomplete, and technically circumventable.

The Reservoir is additive. It begins with an empty namespace and admits only content and services that have been pre-verified against a structured Content Metadata Standard. No filtering is required because harmful content and data architectures are structurally absent — they were never admitted. A VPN cannot reveal content that does not exist in the namespace. Circumvention is architecturally precluded.

How the Architecture Works

Layer	Actor	Mechanism	Effect
Registry (gTLD)	ICANN-authorised registry operator	Admits only content providers meeting the Content Metadata Standard as a condition of domain registration	Harmful data architectures excluded at point of admission — before they reach any child
ISP	Internet service provider (voluntary)	Offers the verified namespace as an optional household subscriber service; configures DNS resolver at account level	All devices on the household account resolve only within the verified namespace when the service is activated
Household	Parent or carer (voluntary opt-in)	Activates the namespace for the household; configures metadata filters to household preferences	Single parental decision replaces hundreds of per-service consent flows across all household devices
Content Provider	Publisher (voluntary)	Applies for a registry domain; accepts the Content Metadata Standard as a binding condition of registration	Data minimisation, best interests, and privacy-by-default are conditions of structural access — not retrospective compliance obligations

The Content Metadata Standard

The Content Metadata Standard is the technical heart of the architecture. It is not a content filter. It is a structured declaration system — analogous to nutritional labelling on food products — that requires registrants to accurately describe their data practices in machine-readable form. The registry is a standards body, not a gatekeeper. It defines how services must describe themselves. The parent decides what their child sees.

Metadata Layer	Required Declarations	Regulatory Relevance
Inclusion Criteria (Floor)	No targeted advertising directed at minors; no persistent tracking without parental consent; no dark patterns; verified age-appropriateness of content	Directly operationalises s.9 (privacy by default), s.10-11 (best interests), s.21 (no coercion)

Descriptive Metadata	Age band; content type; interaction type (passive/interactive/social); commercial presence; privacy level (none / opt-in only / anonymised / identifiable with parental consent)	Enables machine-readable transparency consistent with Division 3; creates resolver-enforceable parental controls consistent with s.13
Parental Control Integration	Metadata tags read by household DNS resolver in real time; parental preferences configured once at account level; automatically enforced across all devices and all services	Operationalises s.13 parental consent at infrastructure layer; reduces per-service consent burden to a single household opt-in decision

The Coalition Proposal

The submitter recommends that the OAIC use its convening authority to encourage the formation of a multi-stakeholder coalition of organisations with expertise in child online safety — including child welfare organisations, ISPs, DNS infrastructure providers, academic researchers, and international regulatory counterparts — to develop and apply for an ICANN community gTLD as a structural complement to the Code.

Why ICANN's 2026 New gTLD Program

ICANN has opened its 2026 New gTLD Program application window (30 April to 12 August 2026). A community gTLD application — submitted by a coalition of child welfare organisations — would qualify for Community Priority Evaluation, which grants priority consideration to applications made by and for a defined community with a legitimate purpose. A child safety coalition is precisely the kind of community for which this evaluation pathway was designed.

A successful application would establish a verified, ICANN-recognised namespace for child-appropriate content and services — giving the architecture described in this submission a global DNS infrastructure foundation that no individual government or regulator could create unilaterally. Australia's participation in such a coalition — through OAIC, eSafety, or Australian child welfare organisations — would position Australia as a founder of a global child safety infrastructure initiative at a moment when that leadership is achievable.

What the OAIC Could Do

- Encourage Australian child welfare organisations, ISPs, and academic institutions to participate in a multi-stakeholder coalition exploring a community gTLD application.
- Engage with international regulatory counterparts — particularly the UK ICO, the EU's supervisory authorities, and the US FTC — to signal regulatory support for the infrastructure model as a complement to application-layer codes.
- Include in the Code's Explanatory Statement a recognition that DNS infrastructure operators, ISPs, and registry operators are relevant stakeholders in the child online safety ecosystem — and that their participation in multi-stakeholder governance bodies is consistent with the Code's objectives.

- Reference the additive namespace model in the Code's guidance as a recognised complement to per-service compliance obligations — acknowledging that services admitted to a verified namespace are operating within an architecture designed to serve the Code's protective intent.

Why This Is Not Sufficient Without Both

To be direct: the additive namespace alone is not sufficient, and the Code alone is not sufficient. They address different layers of the same problem and are most effective in combination.

	Code Alone	Namespace Alone	Code + Namespace
Covers all services on open internet	Yes	No — only opt-in participants	Yes
Prevents harmful data collection by architecture	No — governs behaviour within harmful architecture	Yes — excludes harmful architectures at admission	Yes
Enforceable against non-compliant entities	Yes — regulatory investigation and penalty	Yes — registry agreement revocation	Yes — dual enforcement
Provides household-level parental control	No — per-service consent only	Yes — single household opt-in	Yes
VPN-immune	No — applies to entity behaviour regardless	Yes — content is structurally absent	Yes
Addresses pre-existing data harms	Yes — destruction rights, audit obligations	No — prospective only	Yes
International scalability	Limited — jurisdiction-specific	Yes — ICANN is a global infrastructure authority	Yes

Conclusion

The submitter commends the OAIC for producing a Code that is substantive, technically informed, and grounded in genuine engagement with children, parents, and industry stakeholders. The Exposure Draft represents the most serious effort yet made in the Australian context to translate child welfare principles into enforceable privacy obligations.

The observations in Part One are offered in that spirit — not as objections to the Code's approach but as identification of the structural limits that any application-layer instrument faces when operating within an internet architecture that is not designed for children. The proposal in Part Two is offered as a path beyond those limits: a structural complement that does not require amending the Code, does not expand OAIC's regulatory mandate, and does not depend on government action — but that would make everything the Code is trying to achieve more effective.

The ICANN application window closes 12 August 2026. The moment to act on this proposal is now. The submitter welcomes the opportunity to present this framework at a Virtual Roundtable and to engage further with the OAIC and relevant Australian stakeholders on the coalition formation process.

> *Australia helped design the world's first social media age restriction. It can now lead the formation of the world's first structurally safe internet namespace for children. The two instruments together — a regulatory code and a verified infrastructure — would represent a genuine step change in child online safety. Either alone is insufficient. Both together is transformative.*

Submitted by:
[redacted]