

Submission to the Office of the Australian Information Commissioner

Draft Children's Online Privacy Code

Submitted by: [REDACTED]

1. Introduction

I make this submission in my personal capacity.

I welcome the Office of the Australian Information Commissioner's development of the Children's Online Privacy Code. Children and young people now participate in online spaces that are deeply embedded into social life, education, entertainment, communication, identity formation, gaming, search, media, artificial intelligence services and access to information. These services can provide real value but they also create privacy risks that children are often poorly placed to understand, resist or remedy.

The OAIC consultation page explains that the Code will specify how online services likely to be accessed by children, or primarily concerned with the activities of children, must comply with the Australian Privacy Principles and will set out additional requirements for the handling of children's personal information. The Code will apply to social media services, relevant electronic services and designated internet services as defined by the Online Safety Act 2021. It will be a legislative instrument under the Privacy Act 1988 and must be registered by 10 December 2026 (OAIC, 2026a).

My core position is that the Code should be strong, practical and privacy-preserving. It should protect children through data minimisation, privacy by default, strict limits on profiling and direct marketing, child-friendly transparency, meaningful deletion rights, practical complaint pathways, privacy impact assessment obligations and stronger controls over artificial intelligence systems likely to be accessed by children.

At the same time, the Code should avoid creating incentives for surveillance-by-default, identity-by-default, blanket age verification, biometric age estimation, government identity document uploads, persistent age credentials or excessive parental monitoring of older children and teenagers.

A strong children's privacy regime should not require every child, teenager and adult to prove who they are everywhere online. It should require online services to design their systems so that children are not tracked, profiled, nudged, manipulated, retained, monetised, emotionally exploited or surveilled unnecessarily in the first place.

2. Core position

The final Code should be built around the following principles:

- children's personal information should be collected, used and disclosed by default only where strictly necessary;
- services should be safer and more privacy-protective by design, not merely age-gated;
- applying child-protective defaults to all users should be treated as a preferred pathway where practicable;
- age assurance should be risk-based, privacy-preserving and data-minimising;
- age assurance should not become identity verification by default;
- children should not be forced into unnecessary biometric, document-based or persistent identifier systems;
- anonymity and pseudonymity should be preserved wherever practicable;
- older teenagers should have meaningful privacy and autonomy, especially in relation to support-seeking, health, identity, belief, family conflict and personal development;
- services used through schools and essential institutions should be subject to heightened safeguards;
- default discoverability, contactability and public visibility should be privacy-protective;
- gaming, virtual worlds and avatar-based services should be treated as children's privacy environments;
- children's images, voice, face and media data should receive specific protection;
- privacy notices should be genuinely understandable to children and teenagers;
- consent should be voluntary, specific, current, unambiguous and free from manipulative design;
- direct marketing and profiling of children should be tightly limited;
- children should have practical access, correction, deletion and complaint pathways;
- monitoring and control mechanisms should be transparent to children;
- AI companions, AI assistants and potentially emotionally interactive AI systems should be treated as high-risk where they process children's personal information or are likely to be accessed by children;
- privacy impact assessments should be operationally meaningful, not merely documentary compliance;
- entities should be required to demonstrate how their design choices serve the best interests of the child;
- security, breach response and account recovery should be treated as core children's privacy safeguards.

I broadly support the direction of the Exposure Draft. The strongest parts of the draft are its strict-necessity default, best-interests framing, consent safeguards, child-friendly transparency requirements, destruction rights, monitoring notifications, complaints provisions and privacy impact assessment obligations. My recommendations are directed at ensuring those provisions remain precise, enforceable and resistant to unintended consequences.

3. Key terms and interpretive principles

Several terms used in this submission may be interpreted too broadly or too loosely by some individuals. For clarity, I use the following meanings.

“Artificial intelligence system” should mean a machine-based system that, for explicit or implicit objectives, infers from inputs how to generate outputs such as predictions, content, recommendations or decisions. The term should not be treated as implying human understanding, judgement, care, moral responsibility or therapeutic competence.

“Generative AI” should mean AI systems designed to generate synthetic content, including text, images, audio, video, code or other digital output. Generative AI may be useful but where it is likely to be accessed by children it should be assessed by what it does in practice: what information it collects, what it infers, what it remembers, what it generates, how it responds to distress and whether it encourages dependency or unsafe reliance.

“Large language model” should mean a machine-learning model trained on large amounts of text or multimodal data to predict, generate, transform or respond to language-like input. A large language model should not be described to children as if it understands them, cares about them, knows them or can replace human support under any circumstance.

“Generative Pre-trained Transformer” should mean a family of machine-learning models based on the transformer architecture and pre-trained through self-supervised learning on large datasets. GPT-style systems are a major form of contemporary large language models but they are not the whole of AI. The Code should therefore avoid definitions that are either too narrow by treating all AI as GPT’s or too broad by treating all software automation as equivalent to high-risk AI.

“AI companion” should mean an AI system designed, presented or used to simulate companionship, friendship, intimacy, mentoring, therapy, emotional support, romantic interaction or a persistent personal relationship. An AI companion should be treated as higher risk where it is likely to be accessed by children because it may process sensitive disclosures, infer emotional state, encourage repeated engagement, simulate care or become part of a child’s support-seeking behaviour.

“Age assurance” should mean the broad set of methods used to estimate, infer, declare, verify or otherwise establish that a user falls within a defined age range.

“Age estimation” should mean a reliably probabilistic assessment of a user’s age or age range, including assessment from biometric, behavioural, contextual or other signals.

“Age inference” should mean a probabilistic conclusion about age or age range drawn from existing information, behaviour, metadata, account history, service use or contextual signals.

“Age self-declaration” should mean a user stating their age, date of birth or age range without verification against a meaningfully authoritative source.

“Age verification” should mean establishing age from a verified source such as a government document, digital identity credential or other meaningfully authoritative record.

“Identity verification” should mean establishing who a person is. Identity verification is more intrusive than age assurance and should not be used where the relevant question is only whether a user falls within an age range.

“Sensitive inference” should mean a conclusion or prediction about a child’s sensitive or intimate characteristics derived from behaviour, metadata, content, interactions, images, voice, location, time, search activity, social graph, device signals or other indirect data. Sensitive inferences may concern mental health, distress, sexuality, religion, political belief, family conflict, body image, disability, vulnerability, emotional state, loneliness or self-harm risk.

“Profiling” should mean the use of personal information to evaluate, predict, rank, classify or influence a child’s behaviour, preferences, interests, vulnerabilities, emotional state, likely actions or commercial value. Profiling should not be treated as harmless merely because it is automated or probabilistic.

“Strictly necessary” should mean necessary for the child-facing function reasonably expected by the user, not merely useful for engagement, advertising, analytics, growth, experimentation, personalisation or business optimisation.

“Privacy by default” should mean that the service’s default settings protect children without requiring them to understand complex privacy menus, reject manipulative prompts or repeatedly opt out of collection, profiling, sharing or visibility that is not strictly necessary.

“Surveillance-by-default” should mean a design pattern where users, including children, are routinely monitored, identified, tracked, profiled, inferred about or made visible as the normal condition of participation.

“Identity-by-default” should mean a design pattern where users must routinely prove legal identity, provide government documents, submit biometric information or use persistent age or identity credentials to access ordinary online services.

“Child-protective defaults” should mean service settings and architecture that reduce risk for children by default including less collection, less profiling, less discoverability, less direct marketing, less location exposure, less public visibility, stronger deletion rights, safer recommendations and clearer explanations.

“Dark patterns” should mean interface designs that manipulate or pressure users into choices they would not otherwise make including repeated prompting, misleading button design, unequal accept/refuse options, guilt-inducing language, bundled consent, preselected optional settings or excessive friction to refuse or withdraw consent.

4. Age assurance should be a last-resort design choice, not identity verification by default

Section 8 of the Exposure Draft requires an entity, before collecting personal information about an end-user of the service to take reasonable steps in the circumstances to ascertain the age of the end-user. In determining what steps are reasonable, the entity must have regard to the risk of harm that may arise from collection, use or disclosure of personal information during ordinary provision of the service. The entity may collect personal information before ascertaining age only to the extent necessary to comply with that requirement. Importantly, section 8 does not apply if the entity applies the rest of the Code to the service regardless of the age of end-users (OAIC, 2026b).

This is one of the most important provisions in the Code. If implemented carefully, it can support proportionate protection for children. If implemented poorly, it can become a major privacy risk in itself for all Australian citizens.

The Code should make clear that age assurance is not synonymous with age verification, and that age verification is not synonymous with identity verification. These concepts should remain distinct and must not be confused.

The final Code and OAIC guidance should state that identity verification should not be used where an age range, age flag, local device signal, service-level default, self-declaration, age estimation or lowest-intrusion method would be sufficient. A service should not ask “who are you?” when the regulatory question is only “are you old enough for this feature?” or “should this user receive child-protective defaults?” In all cases, the method used should be the least intrusive of the selections available to the entity, within reason.

The OAIC’s separate privacy guidance on age assurance technologies already supports this approach. It asks entities to establish whether age assurance is needed, whether the function or activity can be undertaken without age assurance, whether age assurance may cause unjust digital exclusion and whether less privacy-intrusive alternatives are available. It also states that age assurance can affect anonymity or pseudonymity where it unnecessarily documents identity rather than age or age range and that age checks should not become a way to unnecessarily collect identity information (OAIC, 2026d).

The Code should therefore make clear that age assurance is not a shortcut around the Privacy Act. The fact that age assurance is technically available, commercially convenient or useful to enforcement should not establish that it is necessary. The least intrusive effective method should be preferred.

In many cases, the better approach will be to apply child-protective defaults to all users rather than verify or infer the age of every user. If a service can reduce tracking, profiling, recommender manipulation, direct marketing, geolocation collection, behavioural surveillance and unnecessary retention for everyone, it should not need to create a new age-checking layer merely to deliver basic levels of privacy protection.

I recommend that OAIC guidance state that entities should follow a hierarchy of preferred approaches:

1. redesign the service so that the relevant privacy protection applies to all users without age assurance;
2. if age assurance is genuinely needed, use the least intrusive age range or eligibility method available;
3. prefer anonymous, pseudonymous, local, ephemeral or tokenised age signals where practicable;
4. avoid collecting government identity documents, biometric information or persistent identifiers unless strictly necessary;
5. avoid retaining raw age-assurance evidence where a short-lived proof, decision record or yes/no age flag would be sufficient;
6. avoid repeated or always-on age monitoring unless necessary, proportionate and clearly justified;
7. preserve fallback pathways for users who lack government ID, compatible devices, stable internet access, biometric reliability or parental support;
8. ensure age-assurance vendors are subject to strict purpose limitation, deletion, security, audit and no-secondary-use requirements.

Where age assurance is used, the age-assurance method should itself be treated as a high-risk privacy activity. It should not be treated as a neutral compliance step.

Before implementing age assurance, an entity should be expected to assess:

- whether age assurance is genuinely necessary;
- whether child-protective defaults could be applied to all users instead;
- what personal information the age-assurance method collects;
- whether the method requires identity, age range only, biometrics, behavioural inference, device data or government documents;
- whether the method produces false positives or false negatives and the probability of such incidents if practicable;
- whether particular groups of children or adults are more likely to be wrongly excluded;
- whether users have a practical fallback pathway;
- whether users can contest an incorrect age decision;
- whether raw evidence is retained;
- whether age-assurance vendors can reuse, profile, sell, train models on or otherwise process the information for secondary purposes;
- whether the method creates a persistent identifier or cross-service tracking risk;
- whether the method creates identity-theft risks or otherwise underlying cybersecurity and privacy risks;
- whether the method could later be repurposed for broader access control or enforcement.

The purpose of the Code should be to reduce privacy risk for children, not create a new identity infrastructure around children.

5. Blanket age verification is an unsafe substitute for safety-by-design

The Code should explicitly avoid creating incentives for blanket age verification across ordinary online services. Age verification may be appropriate in limited, high-risk contexts but it should not become the default response to online harm. There are four reasons for this;

First, blanket age verification is technically and behaviourally fragile. Children and teenagers may use workarounds, shared accounts, borrowed devices, false declarations, VPNs or proxies, fake identity information, adult credentials or less regulated services. UNICEF has warned that age restrictions may backfire if children and young people continue to access social media through workarounds, shared devices or less regulated platforms, making them harder to protect (UNICEF, 2025a). The Internet Society has similarly warned that privacy concerns can push people towards risky workarounds and that attempts to restrict tools such as VPNs can weaken legitimate security tools used by people and businesses (Internet Society, 2025).

Second, identity-heavy age verification can create serious privacy and security risks. Government ID uploads, facial scans, biometric templates, behavioural age inference, persistent age tokens and repeated verification events create valuable data stores and phishing opportunities. The Australian Human Rights Commission has warned that social media age restrictions may require age assurance across social media accounts, including ID uploads, facial scans or behavioural age signals, and that this risks normalising surveillance as the price of online participation (Australian Human Rights Commission, 2025). The OAIC's age-assurance guidance likewise emphasises privacy by design, data minimisation, anonymity or pseudonymity where practicable and destruction or de-identification once age-assurance inputs are no longer needed (OAIC, 2026d).

Third, blanket access restrictions can remove beneficial online spaces. Social media and online communities are not only sources of harm. They can also provide learning, connection, self-expression, peer support and access to information. UNICEF has described social media as a lifeline for many isolated or marginalised children, providing access to learning, connection, play and self-expression (UNICEF, 2025a). eSafety's research on young LGBTIQ+ people found that LGBTIQ+ teens use the internet to express themselves creatively, access support, connect with others like themselves, learn about the world and communicate with like-minded people while also experiencing higher rates of online hate and harassment (eSafety Commissioner, 2024). The policy conclusion should not be that these young people should be removed from online spaces. It should be that those spaces should be made safer, less exploitative and more accountable.

Fourth, overbroad restrictions can create reactance, distrust and counterproductive behaviour. Psychological research on censorship and reactance has found that restriction can increase desire to access restricted communication and can change attitudes toward the restricted content in some circumstances (Worchel, Arnold and Baker, 1975; Steindl et al., 2015). The Code should avoid design choices that make young people experience privacy protection as exclusion, surveillance or arbitrary control.

Australian psychiatric literature has also warned that the evidence base for broad age restrictions should be treated carefully. Blake et al. argue that there is inadequate evidence at this time to conclude that the rise in youth mental illness is attributable to social media and that restrictive measures affecting all young people should be evidence-based and attentive to unintended consequences (Blake et al., 2025). That does not mean online harms are unreal. It means broad restrictions should not substitute for targeted design, privacy and safety reforms.

For those reasons, the final Code should make clear that the preferred child-safety pathway is not blanket age verification. It is safer service design. This means:

- less collection;
- less profiling;
- less behavioural tracking;
- stricter default privacy settings;
- less direct marketing;
- fewer dark patterns;
- safer recommender systems;
- clear complaint pathways;
- effective content moderation;
- meaningful deletion rights;
- privacy impact assessments;
- transparency about automated decision-making;
- higher standards for services likely to be accessed by children.

A privacy code should make the online environment safer without requiring unnecessary identity disclosure. Children should not be protected from one form of harm by being exposed to another.

6. Anonymity and pseudonymity should be preserved as core internet protections

The final Code should treat anonymity and pseudonymity as central privacy protections, not as loopholes.

APP 2 already provides that individuals must have the option of not identifying themselves or using a pseudonym when dealing with an APP entity unless identification is required or authorised by law or impracticable. The OAIC's age-assurance guidance states that age assurance tools can affect anonymity or pseudonymity when they are overbroad and unnecessarily document identity rather than age or age range. It also states that APP entities should deal with unidentified age-assured individuals to the extent practicable unless identity is required or authorised by law (OAIC, 2026d).

This principle should be carried clearly into the Children's Online Privacy Code.

Anonymity and pseudonymity are important for:

- children seeking information about health, safety, bullying, family conflict or abuse;
- teenagers exploring identity, belief, sexuality, politics, religion, creativity or community;
- young people in regional, isolated or hostile environments;
- children avoiding retaliation, stalking, harassment or doxxing;
- people who need to participate online without tying every action to their legal identity;
- adults who should not be forced to identify themselves merely because children may also use a service.

The Internet Society has stated that age checks must not identify or track people online, should not require government ID or financial accounts, should limit what information is shared to age or age range and should not allow age-check providers to track which online services people are visiting (Internet Society, 2025). Those principles should inform OAIC guidance.

The Code should make clear that, where age assurance is used, the desired end state should be:

- “this user is in an appropriate age range for this feature”;
- not “this service knows who this person is”;
- and not “this vendor can track this person across online services”.

7. Older teenagers should have meaningful privacy and autonomy

The Code should recognise that teenagers, particularly older teenagers have legitimate privacy interests that are not identical to those of younger children.

This does not mean that teenagers should be treated as adults in every context. It does, however, mean that privacy frameworks should account for growing autonomy, capacity, maturity, safety needs and development.

Australian public systems already recognise meaningful privacy and control for young people around age 14 in important contexts. Services Australia states that when a child turns 14, parents no longer have online access to that child’s Medicare claims history and parents also no longer have access to the child’s immunisation history statement. The child can view their immunisation history statement using their own Medicare online account through myGov or the myGov app (Services Australia, 2025). My Health Record guidance states that authorised representatives are automatically removed from a child’s record when the child turns 14 and individuals aged 14 or older who have capacity are responsible for their own My Health Record (Australian Digital Health Agency, 2026). The Department of Health states that young people aged 14 to 17 can register and provide consent for MyMedicare without a parent or guardian (Department of Health, Disability and Ageing, 2026).

These examples do not settle every online privacy question but they show that Australian law and public administration already recognise that a 14-year-old may have privacy interests and decision-making capacity that differ from those of a younger child.

The Exposure Draft itself recognises some of this complexity. Section 13 provides that a child may give consent to collection, use or disclosure only if at least 15 years old but it also recognises that a child under 15 may consent where seeking legal or health-related information or support in connection with a person with parental responsibility (OAIC, 2026b). That is important and should be preserved.

The final Code and guidance should recognise that teenagers may need privacy from parents, peers, schools, platforms, advertisers, data brokers and even other users in specific contexts. This is especially important where the teenager is:

- seeking health information;
- seeking legal or family violence support;
- dealing with bullying or harassment;
- exploring identity or belief;
- communicating with support services;
- in family conflict;
- at risk of outing, retaliation or punishment;
- seeking mental health information;
- asking for help in a sensitive context.

Parental involvement can be protective but parental monitoring is not automatically safe or proportionate. For older teenagers, meaningful privacy can itself be a safety protection.

8. Services used through schools or essential institutions require heightened safeguards

The Code should address services that children use because a school, government body, employer, sporting organisation, community organisation or other institution effectively requires them to do so.

In those contexts, children may not have a real choice. A child may be told to use a learning platform, school communication app, assessment tool, online classroom, AI tutoring system, sports registration platform or welfare-related service as a condition of participation. In that environment, consent is weaker, refusal may be impractical and the child may not be able to avoid the service without social, educational or otherwise practical consequences.

For mandatory or institutionally required services, the Code should require stricter safeguards:

- no behavioural advertising;
- no sale or brokerage of children's personal information;
- no unrelated profiling;
- no use of children's information for unrelated product development;
- no use of children's information for AI training unless strictly necessary, clearly disclosed and legally justified;
- strict role-based access for school staff, administrators and vendors;
- clear retention and deletion periods;
- deletion or transfer pathways when a child leaves the school or organisation;
- privacy impact assessments before procurement or deployment;
- child-readable explanations of what the institution and vendor can see;
- clear limits on monitoring, analytics and behavioural dashboards.

Services used in education or essential participation contexts should be held to a higher standard because children cannot meaningfully protect themselves by refusing to use them.

9. Strict necessity should be operationalised as data minimisation by design

Section 9 of the Exposure Draft requires entities to implement technical and organisational measures that, by default, ensure the entity only collects, uses or discloses personal information about a child that is strictly necessary to provide the service. Those measures must also allow a child to control whether non-strictly-necessary information is collected, used or disclosed, and must be clear, simple and easily accessible (OAIC, 2026b).

This is one of the strongest parts of the draft. The Explanatory Statement states that section 9 establishes privacy protections by default and is intended to minimise the amount of personal information in circulation online. It also states that minimising collection and retention reduces breach risk and impact, improves quality and accuracy of personal information and lowers storage and management cost (OAIC, 2026c).

The final Code should preserve this requirement. It should also provide clear examples of what is and is not strictly necessary.

Strictly necessary information should be limited to information required for the core function reasonably expected by the child. It should not include information collected for advertising, profiling, engagement optimisation, behavioural prediction, cross-service tracking, data brokerage, product analytics beyond essential service operation or experimentation unless that handling is demonstrably necessary for the specific service function being used.

The Code should make clear that the following should not be treated as strictly necessary merely because they are useful to the provider's business model:

- targeted advertising;
- behavioural profiling;
- personalised recommendations where not essential to the service;
- cross-context tracking;
- location tracking where not required for the service feature;
- contact list access;
- device fingerprinting;
- analytics beyond security, reliability and essential service operation;
- retention of personal information after a feature is disabled;
- data sharing with unrelated third parties;
- dark-pattern testing or engagement optimisation.

Where a child turns off a non-essential feature, any personal information collected solely for that feature should be deleted or de-identified where APP 11 permits. The Explanatory Statement already recognises this logic in relation to personalised content and APP 11 (OAIC, 2026c). That should become a clear operational expectation.

10. “Best interests of the child” should be specific enough to be clearly enforceable

Sections 10 and 11 of the Exposure Draft require collection, use and disclosure of children’s personal information to be consistent with the best interests of the child. The Explanatory Statement explains that this reflects Article 3 of the United Nations Convention on the Rights of the Child and that children merit specific protection because they may be less aware of the risks, consequences and rights relating to the handling of their personal information (OAIC, 2026c).

I support this framing. However, “best interests of the child” must not become an abstract phrase that entities can satisfy through general statements. The Code should require entities to record a genuine assessment of how the handling of personal information benefits or protects children, rather than only how it supports business objectives.

The final Code or OAIC guidance should require entities to consider, at minimum:

- whether the handling is necessary for the child-facing function;
- whether the same purpose can be achieved with less personal information;
- whether the handling exposes the child to profiling, manipulation, discrimination or loss of autonomy;
- whether the handling increases the risk or severity of future data breaches;
- whether the handling may affect the child’s dignity, safety, development, mental health or social participation;
- whether the child can reasonably understand the handling;
- whether the child can refuse, withdraw consent, delete data or correct information;
- whether the service creates different risks for children of different ages;
- whether parents, schools, peers, advertisers, data brokers or other third parties may gain inappropriate access to the child’s information.

A best-interests assessment should not be treated as a substitute for necessity. A provider should not be able to say that a privacy-invasive feature is in a child’s best interests merely because it increases engagement, personalisation or convenience.

11. Consent requirements should be preserved and strengthened

The Exposure Draft contains important consent safeguards. It provides that consent may be given by a child only if the child is at least 15 years old, with parental consent required for children under 15 unless an exception applies. It also requires consent to be voluntary, informed, current, specific and unambiguous, and states that consent is not voluntary if obtained through manipulative, deceptive or misleading practices, bundled consent requests or denial of service access where the information is not strictly necessary. It also states that consent is not unambiguous if it arises from omission, with examples including pre-ticked boxes or assumed continued use (OAIC, 2026b).

I support these requirements. They address a real problem in online services; formal consent often exists in theory but not in substance. Children and young people are especially vulnerable to consent fatigue, confusing prompts, nudges, false urgency, confirmshaming, repeated pop-ups, bundled requests and take-it-or-leave-it design.

The Explanatory Statement properly identifies nudge techniques and repeated prompts as practices that may frustrate users and incentivise them to agree to more handling than they otherwise would. It also identifies confirmshaming as a manipulative design practice (OAIC, 2026c).

The final Code should preserve these requirements and clarify that consent is not voluntary where the interface uses:

- repeated prompting after refusal;
- guilt-inducing language;
- exaggerated benefits or consequences;
- intentionally confusing button labels;
- preselected optional disclosures;
- unequal prominence between accept and refuse options;
- bundled consent for unrelated purposes;
- forced consent for non-essential data handling;
- excessive friction when withdrawing consent;
- persistent reminders designed to wear down refusal.

Withdrawal of consent should be as easy as giving consent. A child should not need to navigate multiple screens, contact support, use legal language or understand complex account settings to withdraw consent for non-essential data handling.

12. Assent for children under 15 should not be treated as a loophole

Section 20 of the Exposure Draft requires the assent of children under 15 in certain circumstances, including where the child enables collection of sensitive information, use or disclosure for another purpose, or direct marketing. The Explanatory Statement explains that children consulted on the Code expressed a desire to be more involved and aware of how their personal information is handled, and that assent gives children autonomy and empowerment in a manner appropriate to their age and maturity (OAIC, 2026c).

I support this approach. Children under 15 should not be treated as if their privacy interests belong only to adults. However, assent should not become a mechanism for shifting responsibility onto children.

The final Code should make clear that assent is an additional safeguard, not a substitute for strict necessity, best interests, fairness, parental consent where required or privacy-by-default design. A child's assent should not validate a harmful or unnecessary data practice.

Guidance should also address situations where a child's privacy interest may differ from the interest of a parent, guardian, school, platform or other adult. This is particularly relevant for older children, children experiencing family conflict, children seeking health or wellbeing support, children exploring identity or belief and children subject to excessive monitoring.

13. Transparency must explain actual data practices, not merely legal categories

The Exposure Draft requires child-friendly privacy policies, age-appropriate explanations and clear distinctions between strictly necessary and non-strictly-necessary data handling. It also requires explanations of anonymous or pseudonymous use and how pseudonyms are managed (OAIC, 2026b).

The Explanatory Statement gives an important example: if an entity tells a child that it has collected personal information to show videos they might like, without explaining that this involves tracking activity to build a behavioural profile, this may misrepresent the consequences of the handling (OAIC, 2026c). This example should be treated as a central principle. Transparency should explain what actually happens.

A child-friendly privacy explanation should not merely say:

- “we use your information to improve your experience”;
- “we personalise content”;
- “we may share data with trusted partners”;
- “we use cookies and similar technologies”;
- “we use analytics”.

Instead, it should explain in age-appropriate language:

- what information is collected;
- whether behaviour is tracked;
- whether a profile is built;
- whether the profile affects what the child sees;
- whether the child’s information is used for advertising;
- whether the child’s location is collected;
- who receives the information;
- whether the information is sent overseas;
- how long the information is kept;
- whether automated decisions or recommendations are made;
- whether AI systems process the information;
- whether the information is used to train, fine-tune or evaluate AI systems;
- how the child can turn the feature off;
- what happens to previously collected data if the feature is turned off.

The Code should encourage layered explanations; a short child-readable summary first, with deeper detail available for older children, parents, advocates and regulators.

14. Sensitive inferences about children should be treated as high-risk personal information

The Code should explicitly address inferred information about children. Modern online services may not only collect information that a child provides directly. They may infer sensitive or intimate information from behaviour, searches, clicks, watch time, messages, location, device use, friend networks, images, voice, purchases, pauses, emotional expression or interaction patterns.

Inferred information may include:

- mental health or distress;
- loneliness;
- self-harm risk;
- sexuality or gender identity;
- religion or political belief;
- family conflict;
- body image concerns;
- financial vulnerability;
- disability or neurodivergence;
- emotional state;
- peer relationships;
- risk-taking behaviour;
- susceptibility to persuasion.

The Code should make clear that privacy risk does not depend only on whether a child typed a sensitive fact into a form. A platform that infers sensitive facts about a child may create equal or, in some cases, greater privacy risk.

Entities should not infer, store, rank, segment, target, monetise or disclose sensitive inferences about children unless doing so is strictly necessary, in the child's best interests, clearly explained and subject to strong safeguards. Sensitive inferences should not be used for advertising, engagement optimisation, manipulation, recommender tuning, AI companion personalisation or unrelated product development.

15. Profiling and automated decision-making require stronger treatment

Section 28 of the Exposure Draft allows children or persons with parental responsibility to request information about the handling of personal information, including automated decision-making and profiling. The Explanatory Statement says entities are expected to provide simple, easy-to-understand and age-appropriate information about the context of automated decisions, including the logic involved and consequences, without requiring technical knowledge (OAIC, 2026c).

This is important but the final Code should go further. Profiling and automated decision-making are among the most important risks in children's online privacy because they can affect attention, behaviour, self-image, content exposure, advertising, recommendations, risk scoring and access to opportunities.

The Code should require entities to treat the following as high-risk handling likely to require a privacy impact assessment:

- behavioural profiling;
- recommender systems based on personal information;
- ad targeting or marketing segmentation;
- engagement optimisation;
- inferred interests, vulnerabilities or emotional states;
- geolocation-based profiling;

- automated risk scoring;
- personalisation that materially affects content exposure;
- profiling for safety, moderation or enforcement where it may affect account access;
- use of AI systems that respond to, infer or simulate emotional states.

Children and parents should be able to ask not only whether automated decision-making exists but what meaningful categories of data are used, what the decision affects, whether the child can opt out and whether a human review pathway exists for significant consequences.

The Code should also make clear that “meaningful information about logic” does not mean disclosing trade secrets or source code. It means explaining the operational effect in terms a child or parent can understand.

16. AI companions and AI assistants should be treated as high-risk children’s privacy environments

The Code should explicitly address artificial intelligence systems likely to be accessed by children especially AI companions, emotionally interactive chatbots, AI assistants embedded into platforms, AI tutoring systems, AI search interfaces and AI systems that simulate friendship, intimacy, counselling, mentoring or emotional support.

These systems can create serious privacy and wellbeing risks because children may disclose highly sensitive information to them, including, health concerns, family conflict, sexual questions, identity issues, loneliness, bullying, self-harm thoughts, suicidal ideation, secrets, photos, voice recordings, location information and relationship information.

eSafety has warned that popular AI companion chatbots are failing to protect Australian children from exposure to sexually explicit content and are not doing enough to prevent users generating child sexual exploitation and abuse material. Its March 2026 report also stated that most of the AI companion services it examined failed to refer users engaged in suicide or self-harm chats to appropriate support services. eSafety reported that 79 per cent of surveyed Australian children aged 10 to 17 had used either an AI companion or AI assistant and estimated that around 200,000 Australian children had used an AI companion (eSafety Commissioner, 2026a).

eSafety has also warned that AI companions can share harmful content, distort reality, give dangerous advice, encourage ongoing interaction, feel addictive, contribute to overuse and dependency and reduce time spent on genuine social interactions in ways that may contribute to loneliness, low self-esteem and further social withdrawal (eSafety Commissioner, 2025).

UNICEF has identified novel risks for children from AI systems including AI-generated disinformation, emotional dependency on companion chatbots, AI-generated explicit deepfakes and AI-generated child sexual abuse material (UNICEF, 2025b). UNICEF has also warned that AI companions can pose serious privacy risks for children by collecting sensitive information such as photos, voice notes, location, health information, sexual behaviour information and mental health information often without proper safeguards (UNICEF, 2025c).

Common Sense Media and Stanford Medicine's Brainstorm Lab have found that major AI chatbots are fundamentally unsafe for teen mental health support because they miss warning signs, can create dangerous trust through perceived competence, are designed for engagement rather than safety and show degraded safety in extended conversations that mirror real-world teen usage (Common Sense Media, 2025).

The risk is not only that AI systems may give obviously harmful answers. A subtler risk is dependency. If a child or teenager begins to rely on an AI system for emotional validation, decision-making, companionship, motivation, conflict advice or mental-health support, the system may displace human support while also reflecting the child's own inputs back to them in a more persuasive form. This can create echo-chamber effects, reinforce distorted beliefs, intensify isolation or provide seemingly confident but hallucinated advice. The Code should treat this as a privacy and design problem because the system's ability to personalise, remember, infer and respond depends on the processing of highly sensitive personal information.

The Code should therefore treat AI companions and emotionally interactive AI systems as high-risk services where they are likely to be accessed by children. Relevant entities should be required to address:

- whether the AI system is likely to be accessed by children;
- whether it simulates friendship, romance, therapy, counselling, mentoring or emotional attachment;
- whether it encourages ongoing engagement or dependency;
- whether it asks for or receives sensitive information;
- whether conversations are used for profiling, advertising, model training or product development;
- whether users can delete conversation history and derived profiles;
- whether crisis or self-harm signals trigger appropriate human or professional support information;
- whether children are warned that the system is not human, not a friend and not a mental health professional;
- whether the system has age-appropriate boundaries;
- whether the system can hallucinate harmful advice;
- whether the system reinforces false beliefs, delusions or self-destructive thinking;
- whether the system creates echo chambers by agreeing too readily with the child;
- whether the system undermines motivation, learning, critical thinking or real-world relationships through over-reliance.

The Code should not try to solve all AI safety questions. However, it should make clear that child privacy includes protection from exploitative AI data practices, emotional profiling, dependency loops, intimate data extraction, unsafe mental-health interactions and undisclosed use of children's conversations for training or commercial purposes.

17. Default discoverability, contactability and visibility should be private by design

The Code should address not only what information is collected but how visible and reachable a child becomes inside a service.

Many privacy harms arise because children are made discoverable, searchable, contactable or publicly visible by default. This can occur through public profiles, searchable usernames, friend suggestions, contact syncing, location features, public follower lists, open direct messages, public comments, group invitations, voice chat, livestreaming, avatar worlds, gaming lobbies or recommendation systems that connect children with strangers.

The final Code should make clear that services likely to be accessed by children should default to privacy-protective settings for both visibility and contactability.

In particular, services should not make children:

- publicly searchable by default;
- discoverable through phone number, email address or contact upload by default;
- visible in location-based discovery by default;
- exposed to public follower lists or friend graphs by default;
- reachable through open direct messages by default;
- visible in recommender systems designed to increase social engagement by default;
- automatically included in people-suggestion, friend-suggestion or “nearby user” systems by default.

Where contact, discovery or public visibility features exist, they should be off by default or limited to age-appropriate, clearly explained, child-protective settings. Children and teenagers should be able to understand who can find them, contact them, see them, follow them, message them, invite them or view their profile clearly.

This is both a privacy issue and a safety issue. A child’s personal information is not protected merely because the service has a privacy policy. It is protected when the service’s default design does not expose the child to unnecessary discovery, profiling, contact or attention.

18. Gaming, virtual worlds and social features should be treated as children’s privacy environments

The Code should clearly apply to online games, virtual worlds, avatar-based services, social gaming platforms and other interactive environments likely to be accessed by children where personal information is collected, inferred, disclosed or used.

These services are not only “games” in a narrow sense. Many include persistent identities, friend graphs, voice chat, text chat, private messages, public lobbies, virtual items, behavioural analytics, in-game purchases, creator economies, location or device data, moderation systems, advertising, recommendation systems and AI-assisted interactions.

The Code should therefore require entities operating child-accessible gaming or virtual-world services to consider:

- whether children can be contacted by strangers;
- whether adults can contact children;
- whether voice, chat or behavioural data is recorded or analysed;
- whether children's behaviour is used for profiling or monetisation;
- whether in-game purchases or offers are personalised using behavioural data;
- whether friend graphs or social graphs are visible;
- whether location, device or session data is collected;
- whether moderation systems process children's communications;
- whether virtual-world interactions create sensitive inferences;
- whether safety tooling is used for unrelated advertising, analytics or engagement optimisation.

A child's privacy can be affected just as seriously in a game, virtual world or avatar-based service as on a conventional social media platform. The Code should avoid narrow assumptions about what children's online services look like.

19. Images, voice, face and media data should receive specific protection

The Code should give specific attention to children's images, videos, voice recordings, face data and other media-based personal information.

Children increasingly interact with online services through images, video, voice, livestreams, avatars, filters, AI tools, classroom platforms, gaming platforms and social apps. These media types can reveal far more than ordinary account information. They can reveal appearance, age, location, school, home environment, family members, disability, emotional state, health information, accent, ethnicity, religion, gender expression, social relationships and routines.

Media data can also be used to create or support biometric templates, face recognition, voice recognition, deepfakes, synthetic images, sexualised misuse, impersonation, harassment, bullying, doxxing, advertising profiles or AI training datasets.

The final Code should therefore state that entities should not collect, use, disclose or retain children's images, voice, face or other media data unless necessary for the relevant service feature and consistent with the best interests of the child.

In particular, entities should avoid:

- using children's face or voice data for unrelated analytics;
- retaining raw images, videos or voice recordings longer than necessary;
- using children's media data for AI training without clear legal basis and strong safeguards;
- using children's images or voice to infer sensitive information;
- creating biometric templates unless strictly necessary;
- enabling public reuse, scraping or download of children's media by default;
- exposing children's media to search indexing or unauthorised third-party access.

Where media data is collected, children and parents should receive clear, age-appropriate explanations of what is collected, why, whether it is retained, whether it is analysed, whether AI systems process it, whether it can be deleted and whether it is shared with third parties.

20. Direct marketing to children should be tightly limited

The Exposure Draft requires consent for direct marketing involving children's personal information and requires the use or disclosure to be consistent with the best interests of the child (OAIC, 2026b).

This is a necessary safeguard. The final Code should take a particularly cautious approach to direct marketing to children because children are less able to understand persuasion, profiling, sponsorship, influencer dynamics, algorithmic nudging and commercial manipulation.

The Code should clarify that direct marketing to children using personal information is unlikely to be in the best interests of the child where it involves:

- sensitive information;
- inferred vulnerabilities;
- emotional state;
- body image or health concerns;
- financial pressure;
- gambling-like mechanics;
- repeated prompting;
- targeted pressure based on engagement data;
- cross-platform behavioural tracking;
- data obtained from third parties;
- profiling of a child's friends or contacts;
- AI-generated persuasive messaging;
- AI companions or chatbots that steer children toward products or paid features.

The opt-out process should be immediate, child-friendly and no more difficult than the opt-in process. The Explanatory Statement already notes that design practices that make it hard for a child to opt out of direct marketing such as misleading buttons or multiple screens should be precluded (OAIC, 2026c). That expectation should remain strong.

21. The destruction right is one of the most important protections

Section 32 of the Exposure Draft creates a request process for destruction of personal information about a child. The Explanatory Statement states that this is intended to provide stronger privacy protection than APP 11.2 by generally requiring destruction upon request and that de-identification is not sufficient because of re-identification risk (OAIC, 2026c).

I strongly support this approach.

Children should not have to carry a permanent data trail from childhood into adulthood merely because they used ordinary online services. Information collected when a child was younger may become embarrassing, sensitive, inaccurate, misleading or harmful later. Childhood data can also be used to create long-term behavioural profiles or feed future analytics systems.

The final Code should preserve the destruction right and provide practical guidance on:

- backup systems;
- derived data;
- behavioural profiles;
- inferred attributes;
- advertising segments;
- recommendation profiles;
- analytics datasets;
- AI training data;
- AI memory or personalisation systems;
- chatbot conversation histories;
- third-party disclosures;
- cloud storage;
- data processors;
- audit records;
- account deletion versus specific data deletion.

A destruction right will be weak if entities delete the visible account record but retain derived profiles, segment memberships, event logs, advertising identifiers, device identifiers, AI memory, recommendation profiles, third-party copies or model-training datasets linked to the child. The Code should make clear that a destruction request should extend to personal information derived from or linked to the child where reasonably practicable and not subject to a valid exception.

Entities should also be required to explain where complete destruction is not possible, why not, what will be retained, for how long and under what legal basis.

22. The Code should be a step toward stronger erasure and data broker protections for children

The Exposure Draft's destruction right is valuable but it should be understood as part of a wider privacy problem; children's data can spread across online services, advertising technology, analytics vendors, data brokers, AI vendors, third-party SDKs and platform ecosystems in ways that are difficult for children or parents to see, understand or reverse.

International frameworks show stronger approaches. The European Commission explains that under the GDPR, individuals have rights including access, rectification, erasure, restriction of processing, portability, objection and rights in relation to automated decision-making and profiling. It states that personal data provided when a person was a child can be deleted at any time and that this right also applies online and is often referred to as the "right to be forgotten" (European Commission, 2026).

California's Delete Act and DROP system provide another example of a stronger data broker deletion mechanism. California's privacy regulator explains that DROP allows consumers to request deletion of data held by over 500 data brokers in one request and that data brokers must register, process deletion requests, report the types of information they collect and share and undergo regular audits (California Privacy Protection Agency, 2026).

Australia has considered broader reforms. The Attorney-General's Department's Privacy Act Review proposed reforms aimed at strengthening control individuals have over their information, including improving control over personal information through a right to seek erasure and giving individuals more transparency and control over direct marketing, targeting and sale of personal information (Attorney-General's Department, 2023).

This consultation is not the place to redesign the entire Privacy Act or create a national data broker regime. However, the Code should recognise the child-specific version of the problem; once children's data enters advertising, analytics, profiling, AI or broker ecosystems, deletion becomes practically difficult.

The Code should therefore require entities to:

- identify third parties that receive children's personal information;
- identify whether children's personal information is disclosed for advertising, analytics, profiling, model training, data enrichment or brokerage purposes;
- ensure third parties delete or de-identify children's data when the entity must do so;
- prohibit onward sale or unrelated secondary use of children's personal information;
- require deletion of derived profiles and inferences where practicable;
- provide children and parents with meaningful information about whether children's information has been disclosed beyond the primary service;
- make destruction requests operationally effective across processors and vendors.

The Code should not allow a situation where a child can delete an account but cannot meaningfully delete the data ecosystem created around that account.

23. Monitoring and control mechanisms should protect the child's privacy, not only parental visibility

Section 33 of the Exposure Draft requires notification to a child where a mechanism allows a person with parental responsibility to control or monitor service use or geolocation or where another end-user monitors the child's geolocation. Notification must be age-appropriate, given as soon as practicable and provided in a way that reasonably ensures the child is aware of the monitoring (OAIC, 2026b).

I support this requirement. The Explanatory Statement states that the provision is intended to ensure children know when their service use or location may be observed, increasing transparency and promoting children's privacy (OAIC, 2026c).

The final Code should preserve this. It should also recognise that monitoring is not automatically benign simply because it is parental or framed as safety. Monitoring may protect children in some situations, but it may also create risks for older children's autonomy, dignity, help-seeking, identity exploration, friendships, political or religious thought and safety in abusive or high-conflict family environments.

Guidance should distinguish between:

- safety-related parental controls;
- spending controls;
- content controls;
- location sharing;
- communication monitoring;
- message scanning;
- friend/contact monitoring;
- behavioural analytics;
- school or institutional monitoring;
- AI-generated risk scoring or emotional monitoring.

The stronger the monitoring, the stronger the transparency and safeguards should be. Location monitoring, communication monitoring and emotional or behavioural monitoring should receive particular attention.

Children should be given age-appropriate information about:

- who can see their information;
- what information is visible;
- when monitoring is active;
- whether monitoring is continuous or event-based;
- how long monitoring records are retained;
- whether the child can ask questions or seek help.

24. Child-friendly complaints and inquiries should be operationally real

Sections 35 and 36 of the Exposure Draft require clear, concise, transparent and age-appropriate information about children's privacy rights, and require child-friendly inquiry and complaint processes. The Explanatory Statement says entities must embed processes into the service and that design elements should not make it difficult for end-users to locate and understand how to make an inquiry or complaint (OAIC, 2026c).

I support these requirements. A privacy right is weak if a child cannot use it.

The Code should require complaint and inquiry pathways that are:

- visible in the service;
- written in plain language;
- accessible without legal terminology;
- usable by children;
- usable by parents or guardians where appropriate;
- capable of anonymous or pseudonymous inquiry where practicable;
- able to handle deletion, correction, access and explanation requests;
- responsive within clear timeframes;
- designed not to punish or shame the child for asking;
- connected to escalation pathways, including the OAIC.

Entities should be encouraged to include a “What happens next?” explanation so children know whether a complaint will be read by a person, whether their account will be affected, whether their parent will be contacted, and, approximately when they should expect a response.

25. Privacy impact assessments should be published in meaningful summary form

Section 38 of the Exposure Draft requires privacy impact assessments for new services or activities likely to be accessed by children or new or changed handling likely to have a significant impact on children's privacy. The assessment must include the nature, scope, context, flow and purposes of handling, why collection is strictly necessary, how collection is lawful and fair, whether handling is consistent with the best interests of the child, compliance information and risk of harm or potential impact on children (OAIC, 2026b).

I support this requirement. It is one of the strongest governance provisions in the draft.

The final Code should require entities to maintain and, where appropriate, publish meaningful PIA summaries for high-risk activities. These summaries should not expose security-sensitive details or trade secrets but should be more useful than a title and date.

A useful public PIA summary could include:

- the service or feature assessed;
- the categories of children affected;
- the categories of personal information handled;
- whether sensitive information, geolocation, profiling or direct marketing is involved;
- whether automated decision-making is involved;
- whether AI systems process children's personal information;
- whether conversation logs, prompts, voice, images or behavioural data are processed;
- whether data is disclosed overseas;
- the main privacy risks identified;
- the main mitigations adopted;
- whether children or parents were consulted;
- whether the design changed because of the PIA.

Privacy impact assessments should not be purely internal compliance artefacts. Public summaries would improve accountability and trust.

26. Annual review obligations should be measurable

Section 25 of the Exposure Draft requires entities to review and update privacy practices, procedures and systems at least annually, keep records of the reviews and provide them to the Commissioner if requested (OAIC, 2026b).

I support this requirement. However, an annual review should not be a mere box-ticking exercise. Guidance should state that annual reviews should include:

- testing whether default privacy settings remain strict;
- reviewing whether data fields remain strictly necessary;
- reviewing whether non-essential data handling has expanded;
- reviewing whether consent flows remain compliant;
- reviewing whether withdrawal of consent remains easy;
- reviewing deletion request handling;
- reviewing access and correction request handling;
- reviewing complaints and inquiry data;
- reviewing third-party data sharing;
- reviewing profiling and automated decision-making;
- reviewing AI systems likely to be accessed by children;
- reviewing security incidents and near misses;
- reviewing whether age-assurance methods remain proportionate;
- reviewing whether children can still understand privacy notices.

Entities should be encouraged to maintain evidence of control effectiveness, not merely evidence that a policy was reviewed.

27. The Code should address third-party SDKs, advertising technology and analytics

Many children's online services rely on third-party SDKs, analytics tools, advertising technology, crash reporting, cloud services, payment providers, age-assurance vendors, moderation tools and AI systems.

The Code should make clear that entities remain responsible for children's personal information handled through these systems. A service should not be able to avoid responsibility by saying that profiling, analytics, advertising, telemetry or age assurance was performed by a vendor.

Guidance should require entities to assess:

- what third-party code is embedded in the service;
- what personal information third parties receive;
- whether third parties can track children across services;
- whether third parties use children's information for their own purposes;
- whether SDKs collect device identifiers or behavioural telemetry;
- where the data is stored;
- how long it is retained;
- whether third parties may disclose it further;
- whether the entity can delete it on request;
- whether the third party has suffered breaches or security issues;
- whether the third party uses the data for AI training, model evaluation, advertising or profiling.

This is particularly important because children and parents cannot realistically inspect SDK-level data flows.

28. Cross-border disclosure notices should be practical and meaningful

Section 26 of the Exposure Draft sets out requirements for expressly informing a child about cross-border disclosure matters under APP 8. The provision recognises that children should understand when privacy protections may differ if information is disclosed overseas (OAIC, 2026b).

I support this requirement. However, cross-border disclosure notices should not become formal warnings that children cannot understand.

The Code should encourage explanations that identify:

- what information may be sent overseas;
- what country or region it may be sent to, where reasonably known;
- whether the recipient is a service provider, parent company, analytics provider, advertising partner, cloud provider, AI provider or support provider;
- what the recipient may do with the information;
- whether the child can avoid the disclosure by turning off a feature;
- whether the service still works without the disclosure;
- how long the overseas recipient may retain the information.

For high-risk disclosures involving sensitive information, geolocation, profiling, persistent identifiers, conversation logs, voice, images, AI training data or mental-health-related disclosures, the entity should be expected to provide a clearer and more specific explanation than for low-risk, operational disclosures.

29. The Code should avoid over-reliance on parental consent

The draft Code uses parental consent for children under 15 while also requiring assent in some circumstances. This is understandable. However, the final Code should avoid treating parental consent as complete protection.

Parental consent may not always reflect the child's privacy interests. Parents may not understand data handling practices. Parents may be pressured by children. Children may be pressured by parents. Some children may be in unsafe family environments. Some older children may have legitimate privacy interests in health, wellbeing, identity, belief, association or help-seeking.

The Code should therefore maintain the child's own rights and interests as central. Parental consent should not override strict necessity, best interests, fairness, security, deletion rights, transparency or privacy-by-default requirements.

30. The Code should explicitly discourage surveillance-by-default and identity-by-default

A major risk with children's online safety and privacy regulation is that services may respond by collecting more information about children in order to demonstrate compliance. This would undermine the Code's purpose.

The final Code should expressly discourage surveillance-by-default and identity-by-default approaches. In particular, it should discourage:

- blanket identity verification;
- routine government identity document uploads;
- unnecessary biometric age estimation;
- unnecessary facial scans;
- persistent age credentials that can be reused to track users;
- centralised databases of children's age or identity status;
- broad behavioural monitoring for age inference;
- default geolocation collection;
- default profiling;
- default direct marketing;
- excessive parental monitoring;
- retention of raw age-assurance evidence after a decision has been made;
- secondary use of age-assurance data for advertising, analytics, profiling, fraud scoring, AI training or unrelated compliance purposes.

The OAIC's own age-assurance guidance says privacy risks arise where age assurance paves the way for a service to collect identity when that is not required and that APP entities should deal with unidentified age-assured individuals to the extent practicable unless identity is required or authorised by law (OAIC, 2026d). The Code should carry that principle into the children's online privacy context.

The preferred model should be "safe by default", not "identify first". A child should not need to surrender more personal information to obtain privacy protection. Nor should adults be forced to identify themselves merely to access lawful online services that children may also use.

This matters because systems built for child protection can become general-purpose identity and access-control infrastructure. Once deployed, they may be expanded to new categories of content, new kinds of speech, new forms of platform access, new enforcement purposes or new regulatory schemes. The Code should therefore be drafted to prevent function creep.

Where age assurance is used, it should be:

- necessary for a clearly identified risk;
- proportionate to that risk;
- the least intrusive effective method available;
- limited to age or age range rather than identity wherever possible;
- separated from advertising and profiling systems;
- subject to strict vendor controls;
- protected by privacy impact assessment;
- accompanied by deletion and de-identification obligations;
- supported by fallback pathways;
- transparent to children and parents;
- reviewable where a user is wrongly excluded.

The Code should make clear that children's privacy is not served by creating larger stores of children's identity data. It is served by reducing how much data services collect, retain, infer, share and monetise overall.

31. Online safety obligations should not become a privacy loophole

The Code should make clear that child safety, online safety and privacy obligations must be read together. An entity should not be able to justify unnecessary collection, profiling, surveillance, identity verification or biometric processing merely by pointing generally to online safety.

Safety can require information handling in some cases. However, safety-related processing should still be necessary, proportionate, purpose-limited, transparent, secure, retention-limited and subject to review.

In practice, this should mean:

- safety systems should not collect more personal information than needed;
- safety classifiers should not become general behavioural profiling systems;
- moderation tooling should not be used for unrelated advertising or engagement optimisation;
- age-assurance data should not be reused for marketing, fraud scoring, analytics or AI training;
- safety logs should not be retained indefinitely;
- children should receive clear explanations where safety systems materially affect them;
- children should have review pathways where automated safety decisions restrict access or incorrectly classify them.

The Code should prevent a false trade-off where children are asked to give up privacy in order to receive safety. The better standard is privacy-preserving safety by design.

32. Children's personal information should receive stronger security, breach and recovery safeguards

The Code should address security and breach response more explicitly. Children's personal information can have long-term consequences if misused, breached, exposed or linked across services. A child cannot easily replace their identity history, family context, childhood images, school information, health-related disclosures, location history, behavioural profile or sensitive inferences.

APP 11 requires APP entities to take active measures to protect personal information from misuse, interference, loss, unauthorised access, modification or disclosure and to actively consider whether personal information may be retained. The OAIC also states that APP 11 compliance is key to minimising data breach risk (OAIC, 2025a; OAIC, 2024).

For services likely to be accessed by children, reasonable security should include:

- strong access controls for children's personal information;
- encryption in transit and at rest where appropriate;
- strict controls over administrative access;
- logging and review of access to high-risk children's data;
- protection against scraping and bulk export;
- secure account recovery processes;

- controls against account takeover;
- breach response plans that consider child-specific harm;
- prompt notification where a breach creates serious risk;
- deletion of information that is no longer necessary;
- vendor and processor security requirements;
- testing of deletion and account recovery workflows.

Account recovery deserves particular attention. Weak recovery systems can expose children to account takeover, impersonation, blackmail, grooming, harassment or loss of access to important social or educational accounts. Overly strict recovery systems can also lock children out of support, education or community spaces. Entities should design recovery pathways that are secure, age-appropriate, privacy-preserving and resistant to social engineering.

The Code should make clear that protecting children's privacy requires not only limiting collection but also securing what is collected and deleting what is no longer required.

33. The Code should draw from international child privacy models, without importing their weaknesses

The OAIC should consider the UK Information Commissioner's Office Age Appropriate Design Code as a useful comparative model. The ICO describes that Code as a set of flexible standards that provide built-in protection so children can explore, learn and play online with the best interests of the child as a primary consideration. Its standards include high default privacy, data minimisation, limits on profiling, controls around geolocation, transparency about parental monitoring and restrictions on nudge techniques (Information Commissioner's Office, n.d.).

Australia should adopt the strongest lesson from that model; child privacy should be built into service design, not left to children to manage through complex settings.

However, Australia should avoid treating child protection as a reason to normalise identity checks for all users. The Australian Code should be privacy-preserving by design. Its direction should be safer defaults, less data, less profiling and less manipulation, not more identity infrastructure.

34. Small and medium providers need practical implementation guidance

The Code should be enforceable and should not be weakened merely because compliance may require change. However, practical guidance will be needed for smaller providers, open-source projects, educational services, hobby communities and small Australian businesses.

Guidance should include examples of:

- privacy-by-default settings;
- acceptable and unacceptable consent flows;
- age-appropriate notices;
- deletion workflows;
- minimal age-assurance patterns;
- privacy impact assessment templates;

- child-friendly complaint forms;
- direct marketing opt-out designs;
- third-party SDK assessment checklists;
- default settings for geolocation and profiling;
- examples of strictly necessary versus optional data handling;
- examples of AI companion or AI assistant risk assessments;
- examples of how to avoid unnecessary identity verification;
- school and educational technology procurement checklists;
- sensitive inference risk assessments;
- online safety processing boundaries;
- child-specific breach response and account recovery guidance;
- privacy-protective defaults for games, virtual worlds and social discovery features;
- media-data handling guidance for images, video, voice, face and livestreaming.

Good guidance would improve compliance and reduce the temptation for providers to adopt overly intrusive age checks as the simplest apparent solution.

35. Recommendations

I recommend that the OAIC consider the following recommendations in finalising the Children's Online Privacy Code.

1. Preserve the strict-necessity default in section 9.
2. Treat application of child-protective defaults to all users as a preferred privacy-preserving pathway where appropriate.
3. Make clear that age assurance should be a last-resort or risk-based measure where safer design cannot reasonably address the relevant risk.
4. Distinguish clearly between age assurance, age estimation, age inference, age self-declaration, parental attestation, age verification and identity verification.
5. State that identity verification should not be used where an age range, age flag, local device signal, service-level default or lower-intrusion method would be sufficient.
6. State that age assurance should not require government identity documents, biometric information, facial scans, persistent identifiers or reusable age credentials unless strictly necessary and proportionate to a clearly identified risk.
7. Require entities to minimise retention of age-assurance evidence.
8. Require age-assurance methods themselves to be assessed through privacy impact assessment where they create material privacy, exclusion, biometric, identity, vendor or tracking risks.
9. Require practical fallback and contestability pathways where age assurance wrongly excludes a user or requires a user to disclose more identity information than necessary.
10. Clarify that targeted advertising, behavioural profiling, cross-service tracking and engagement optimisation are not strictly necessary merely because they support a business model.
11. Preserve strong consent requirements, including prohibitions on bundled consent, manipulative design, pre-ticked boxes and consent by omission.
12. Clarify that repeated prompting, confirmshaming, unequal button prominence and excessive friction to refuse or withdraw consent may undermine voluntary consent.
13. Treat child assent as an additional safeguard, not a substitute for best interests, strict necessity or parental consent where required.

14. Recognise that older teenagers may have meaningful privacy interests distinct from parents or guardians, especially around health, wellbeing, identity, belief, support-seeking and family conflict.
15. Require services used through schools or other essential institutions to meet heightened privacy safeguards because consent is weaker and refusal may be impractical.
16. Require transparency notices to explain actual data practices, including behavioural profiling, recommendation systems, advertising, overseas disclosure, retention and AI processing.
17. Treat sensitive inferences about children as high-risk personal information, especially where they relate to mental health, sexuality, religion, belief, vulnerability, emotional state, family conflict or self-harm risk.
18. Strengthen treatment of profiling and automated decision-making, including meaningful explanation and opt-out or review pathways where appropriate.
19. Treat AI companions, emotionally interactive AI systems and AI assistants likely to be accessed by children as high-risk where they process children's personal information or simulate emotional support.
20. Require AI systems likely to be accessed by children to disclose whether conversations, prompts, images, voice, emotional signals or behavioural data are retained, profiled, shared or used for training.
21. Require clear safeguards against AI systems creating dependency loops, reinforcing harmful thoughts, hallucinating dangerous advice, simulating therapy, or using sensitive disclosures for profiling or commercial purposes.
22. Require services likely to be accessed by children to use privacy-protective defaults for discoverability, contactability, public visibility, location visibility, friend suggestions, direct messages and profile search.
23. Clarify that gaming, virtual worlds, avatar-based services, social gaming platforms and other interactive environments likely to be accessed by children must assess privacy risks from chat, voice, friend graphs, behavioural analytics, in-game monetisation, moderation, AI systems and social discovery features.
24. Require specific safeguards for children's images, video, voice, face data and other media-based personal information, including limits on retention, biometric processing, AI training, scraping, public visibility and third-party reuse.
25. Maintain strong destruction rights and clarify that deletion should extend, where practicable, to derived profiles, advertising segments, event logs, AI memories, recommendation profiles and third-party copies.
26. Preserve child notification requirements for monitoring, control and geolocation mechanisms.
27. Provide guidance on risks associated with parental monitoring and other-user location monitoring.
28. Require child-friendly complaint and inquiry processes that are genuinely easy to find and use.
29. Require meaningful privacy impact assessments for high-risk services and features.
30. Encourage public summaries of privacy impact assessments where security and commercial sensitivity allow.
31. Require annual reviews to test actual control effectiveness, not merely policy currency.
32. Require entities to assess third-party SDKs, analytics, advertising technology, age-assurance vendors, AI providers and cloud providers.

33. Ensure cross-border disclosure explanations are specific enough to be meaningful to children and parents.
34. Avoid over-reliance on parental consent as the sole privacy safeguard.
35. Explicitly discourage surveillance-by-default, identity-by-default, blanket age verification and unnecessary biometric age estimation.
36. Require age-assurance vendors to meet strict purpose limitation, data minimisation, deletion, security, audit and no-secondary-use obligations.
37. Keep the Code focused on reducing children's data exposure, not expanding age, identity or biometric collection.
38. Require entities to prevent children's personal information being sold, brokered, enriched, profiled or reused for unrelated secondary purposes.
39. Require deletion and destruction workflows to propagate to processors and third parties where practicable.
40. Make clear that online safety obligations do not justify unnecessary identity verification, biometric processing, profiling, surveillance, indefinite retention or secondary use.
41. Require entities likely to be accessed by children to apply child-specific security, breach response and account recovery safeguards for personal information.
42. Require services to protect children against scraping, account takeover, bulk export, unauthorised access, insecure recovery flows and excessive retention.
43. Require breach response plans to consider child-specific harm, including identity misuse, exposure of sensitive media, location history, family context, school information, health-related disclosures and behavioural profiles.
44. Draw from international child privacy models, including the UK Age Appropriate Design Code, while preserving Australia's own emphasis on anonymity, pseudonymity and data minimisation.
45. Provide practical implementation guidance for smaller providers and services.
46. Include clear definitions or guidance distinguishing age assurance, age verification, identity verification, profiling, sensitive inference, generative AI, AI companions, surveillance-by-default, identity-by-default, privacy by default and strictly necessary processing.

36. Closing

The Children's Online Privacy Code is an important opportunity to improve the privacy position of children and young people in Australia.

The Code should recognise that children should not have to understand complex data ecosystems in order to be protected from them. Services likely to be accessed by children should be designed so that privacy protection is the default, not the result of perfect user behaviour.

The strongest version of the Code will reduce unnecessary collection, profiling, retention, disclosure and manipulation of children's personal information. It will make privacy settings clear, consent meaningful, deletion practical, complaints usable, age assurance proportionate and privacy impact assessments operationally real.

The Code should protect children by reducing harmful data practices and unsafe design. Age assurance should be available where necessary and proportionate but it should be implemented in the least intrusive way possible. The Code should not normalise blanket identity verification, biometric checks, persistent age credentials or surveillance-by-default as the price of ordinary online participation.

Children's online safety should be achieved by making services safer, less exploitative and less data-hungry, not by building identity infrastructure around every user.

Australia should protect children online by requiring services to collect less, retain less, profile less, sell less, infer less, manipulate less, expose less, secure more, and explain more clearly.

REFERENCES

Attorney-General's Department (2023) Privacy Act Review Report. Available at: <https://www.ag.gov.au/rights-and-protections/publications/privacy-act-review-report>

Australian Digital Health Agency (2026) Authorised representatives. Available at: <https://www.digitalhealth.gov.au/initiatives-and-programs/my-health-record/getting-started/authorised-representatives>

Australian Human Rights Commission (2025) Keeping kids safe shouldn't mean a loss of privacy for everyone. Available at: <https://humanrights.gov.au/about-us/media-centre/opinion-pieces/rights-and-freedoms/keeping-kids-safe-shouldnt-mean-a-loss-of-privacy-for-everyone>

Blake, J.A., Sourander, A., Kato, A. and Scott, J.G. (2025) 'Will restricting the age of access to social media reduce mental illness in Australian youth?', Australian & New Zealand Journal of Psychiatry, 59(3). Available at: <https://pubmed.ncbi.nlm.nih.gov/39968662/>

California Privacy Protection Agency (2026) About DROP and the Delete Act. Available at: <https://privacy.ca.gov/drop/about-drop-and-the-delete-act/>

Common Sense Media (2025) Common Sense Media finds major AI chatbots unsafe for teen mental health support. Available at: <https://www.common sense media.org/press-releases/common-sense-media-finds-major-ai-chatbots-unsafe-for-teen-mental-health-support>

Department of Health, Disability and Ageing (2026) Information for MyMedicare patients. Available at: <https://www.health.gov.au/our-work/mymedicare/patients>

eSafety Commissioner (2024) The digital lives of young LGBTIQ+ people. Available at: <https://www.esafety.gov.au/research/the-digital-lives-of-young-lgbtq-people>

eSafety Commissioner (2025) AI chatbots and companions: risks to children and young people. Available at: <https://www.esafety.gov.au/newsroom/blogs/ai-chatbots-and-companions-risks-to-children-and-young-people>

eSafety Commissioner (2026a) eSafety report shows AI companions are putting children at risk. Available at: <https://www.esafety.gov.au/newsroom/media-releases/esafety-report-shows-ai-companions-are-putting-children-at-risk>

European Commission (2026) Information for individuals: data protection rights. Available at: https://commission.europa.eu/law/law-topic/data-protection/information-individuals_en

Information Commissioner's Office (n.d.) Age appropriate design: a code of practice for online services. Available at: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/childrens-information/childrens-code-guidance-and-resources/age-appropriate-design-a-code-of-practice-for-online-services/>

Information Commissioner's Office (2020) Age appropriate design code of practice. Available at: <https://www.privacylaws.com/media/3093/final-age-appropriate-design-code-of-practice.pdf>

Internet Society (2025) Policy Brief: Age Restrictions and Online Safety. Available at: <https://www.internetsociety.org/resources/policybriefs/2025/age-restrictions-and-online-safety/>

NIST (2023) Artificial Intelligence Risk Management Framework (AI RMF 1.0). Available at: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>

NIST (2025) Generative pre-trained transformer, Computer Security Resource Center Glossary. Available at: https://csrc.nist.gov/glossary/term/generative_pre_trained_transformer

OECD (2024) Explanatory memorandum on the updated OECD definition of an AI system. Available at: https://www.oecd.org/en/publications/explanatory-memorandum-on-the-updated-oecd-definition-of-an-ai-system_623da898-en.html

Office of the Australian Information Commissioner (2024) Part 1: Data breaches and the Australian Privacy Act. Available at: <https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/preventing-preparing-for-and-responding-to-data-breaches/data-breach-preparation-and-response/part-1-data-breaches-and-the-australian-privacy-act>

Office of the Australian Information Commissioner (2025) Chapter 11: APP 11 Security of personal information. Available at: <https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-11-app-11-security-of-personal-information>

Office of the Australian Information Commissioner (2026a) Draft Children's Online Privacy Code (consultation for industry, civil society, academia). Available at: <https://www.oaic.gov.au/engage-with-us/consultations/draft-childrens-online-privacy-code-consultation-for-industry%2C-civil-society%2C-academia>

Office of the Australian Information Commissioner (2026b) Exposure Draft: Privacy (Children's Online Privacy) Code 2026. Available at: https://www.oaic.gov.au/_data/assets/pdf_file/0020/262631/Exposure-Draft-Childrens-Online-Privacy-Code.pdf

Office of the Australian Information Commissioner (2026c) Draft Explanatory Statement: Exposure Draft Children's Online Privacy Code. Available at: https://www.oaic.gov.au/_data/assets/pdf_file/0019/262630/Draft-Explanatory-Statement-Exposure-Draft-Childrens-Online-Privacy-Code.pdf

Office of the Australian Information Commissioner (2026d) Privacy guidance on age assurance technologies. Available at: https://www.oaic.gov.au/_data/assets/pdf_file/0017/262043/OAIC-privacy-guidance-on-age-assurance-technologies.pdf

Services Australia (2025) What happens when your child turns 14. Available at: <https://www.servicesaustralia.gov.au/what-happens-when-your-child-turns-14>

Steindl, C., Jonas, E., Sittenthaler, S., Traut-Mattausch, E. and Greenberg, J. (2015) 'Understanding psychological reactance: New developments and findings', *Zeitschrift für Psychologie*, 223(4), pp. 205–214. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4675534/>

UNICEF (2025a) Age restrictions alone won't keep children safe online. Available at: <https://www.unicef.org/press-releases/age-restrictions-alone-wont-keep-children-safe-online>

UNICEF (2025b) Guidance on AI and Children 3.0. Available at: <https://www.unicef.org/innocenti/media/11991/file/UNICEF-Innocenti-Guidance-on-AI-and-Children-3-2025.pdf>

UNICEF (2025c) The risky new world of tech's friendliest bots. Available at: <https://www.unicef.org/innocenti/stories/risky-new-world-techs-friendliest-bots>

Vaswani, A. et al. (2017) 'Attention Is All You Need', *Advances in Neural Information Processing Systems*. Available at: https://papers.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

Worchel, S., Arnold, S.E. and Baker, M. (1975) 'The effects of censorship on attitude change: The influence of censor and communication characteristics', *Journal of Applied Social Psychology*, 5(3), pp. 227–239. Available at: https://www.uni-muenster.de/imperia/md/content/psyifp/aeechterhoff/vorlesungskommunikation/worchelarnoldetal_censorshattchan_jasp1975.pdf