

The Hon. Dr Sarah Kaine, MLC
Committee Chair
Standing Committee on Social Issues
Parliament of New South Wales,
6 Macquarie St,
Sydney NSW 2000,
Australia

31 May 2025

Dear Madam Chair and Committee Members,

Age Verification Providers Association Follow up: Submission for the Inquiry into Impacts of harmful pornography on mental, emotional, and physical health

It was a pleasure to give evidence to your committee. Thank you for your follow-up questions which we address below, as well as an answer to the question I took on notice which was to provide a summary of the two main international standards for age assurance.

(1) If certain persons wished to self-exclude from pornography sites, could age-verification software be used to achieve this?

Yes – provided the age verification system uses identity-linked verification, self-exclusion can be supported, and in a way that still preserves user anonymity at the point of access to adult content.

A useful precedent is the UK's **Gamstop** scheme for online gambling (being replicated by Australia's forthcoming **BetStop** system). When a user self-excludes, they supply personal details including their name, date of birth, postcode, email address and phone number. Gambling operators must then check new or existing customers against this secure register. A match typically requires an exact name and date of birth, plus a matching contact detail such as postcode, email or phone number. Operators receive only a "match/no match" response – they do not see the exclusion list or the user's identity – we call this a "one-way blind check".

A similar model could be adapted for adult content but only if a user verifies their age using a government-issued ID or similar credential. The age verification provider (acting as a trusted intermediary) could check a hashed version of that ID against a self-exclusion list. If a match is found, access would be denied – but the adult site itself would not see the ID or learn the user's identity, and because the query is 'hashed', nor would the operator of the list know the identity of those being checked against it. ('Hashing' is the process of converting data into a fixed-length code using a mathematical algorithm, so it can be stored or compared without revealing the original information.)

By contrast, anonymous or non-ID-based age assurance methods – such as facial age estimation or behavioural inference – do not obtain any identity information so cannot support self-exclusion, since there is, deliberately, no way to recognise the same individual across sessions or platforms.

In short: self-exclusion is technically feasible in ID-linked, privacy-preserving age verification systems, but not in estimation or some inference based models.

(2) If age-verification software could not be used in this way [for self-exclusion], are you aware of any other software which could be used to exclude certain users from accessing pornography sites?

Where the initial age verification process is not identity-linked, it cannot facilitate self-exclusion. However, there are alternative tools that individuals can use to limit their own access:

- **Parental control software** and **device-level content filters** can block access to adult sites. These are often used by parents but can also be installed voluntarily by adults on their own devices.
- **Network-level DNS filtering services** allow users to block categories of content – including pornography – across all devices connected to a home network.
- Some **browser extensions** or third-party firewalls offer self-managed restrictions or blacklists for adult content.

These tools operate independently of any identity verification process. They may be suitable for individuals seeking to avoid adult content on their own devices, but they are not enforced at the platform level and are easily circumvented. They also lack the oversight, consistency and user accountability that a trusted, regulated self-exclusion mechanism – such as one linked to identities – could provide.

(3) Can age-verification software be used to block ads for certain sites being directed towards those under a specified age?

Yes – in principle, age verification technology can support more responsible advertising by limiting the targeting of age-restricted marketing content to users who have proven they are above a specified age threshold.

This would typically work by providing the advertising platform with an **age token** – a cryptographic proof that a user is over (or under) a certain age, without revealing their identity. The ad platform could then decide which ads to serve based on this information. This enables compliance with age-appropriate advertising standards without collecting or sharing personal data.

In effect, this creates a separation between the advertiser, the publisher, and the age verification provider – with only the latter knowing the user's age. This is consistent with emerging privacy-first regulatory models, such as those being implemented in the EU and France.

The challenge lies in integrating such systems across the ad tech ecosystem, which has historically relied on inferred or self-declared ages rather than verified ones. Nonetheless, this is technically feasible and aligns with best practice in online safety, facilitated by the latest innovative interoperable age assurance ecosystems, such as euCONSENT's AgeAware.

(4) If age-verification software could not be used in this way, are you aware of any other software which could be used to block ads for certain sites being directed towards particular users or particular classes of users?

Yes – there are other tools that can help limit exposure to inappropriate advertising, although most are less reliable than verified age-based controls:

- **Ad platforms** already offer settings to target or exclude users based on age, geography or behaviour, but this relies heavily on self-declared or inferred data, which is often inaccurate and easily manipulated.
- **Parental controls** and **safe browsing modes** (e.g. Google SafeSearch, YouTube Kids) can reduce exposure to adult-oriented ads but offer no guarantees.
- **Device-level ad blockers** and **content filtering software** can be used to block entire ad categories or known domains associated with adult content.

However, none of these approaches offer the same assurance as using verified age information. Inferred profiles can be wrong, especially for shared or child-accessible devices. By contrast, a trusted age verification layer can provide a consistent, privacy-preserving method of classifying users by age group – and could help raise standards across the online advertising sector.

As the Commonwealth Government moves forward with the implementation of largescale age assurance, following the completion of the current technology trial, we hope it will encourage two important foundational steps:

1. Requiring that platforms deploy age assurance that is audited and certified against international standards to deliver the level of accuracy required by Australian regulators for each use-case e.g. adult content, social media, data sharing. Those standards also guarantee privacy, data security and human rights.
2. The adoption by the age assurance industry as it is rolled out across Australia of interoperable solutions, so users need only prove their age once to access multiple services, often without any further action required for a defined period of time, determined by regulators in proportion to the level of risk.

Summary of IEEE 2089.1-2024

Title: *IEEE Standard for Age Appropriate Digital Services Framework Based on the 5Rights Principles for Children*

Purpose and scope:

IEEE 2089.1-2024 provides a **design framework** to help digital service providers implement age-appropriate experiences, particularly for children. The standard aligns with the principles of the 5Rights Foundation, aiming to ensure that digital products and services proactively respect children's needs, rights and capabilities.

It includes technical, design and organisational controls, offering guidance on implementing these principles through a structured and risk-based approach.

Core Phases of Age Assurance in IEEE 2089.1

The standard defines a structured lifecycle for implementing age assurance, consisting of **five sequential process phases** supported by **three continuous cross-cutting concerns**:

1. Determination

- Assess the context of the product or service.
- Identify whether age assurance is required based on legal, ethical or contractual obligations.
- Define risks to children and document potential harms or exposure.

2. Selection

- Choose the most appropriate age assurance method(s) for the context.
- Balance effectiveness, user experience, proportionality and privacy.
- Select from methods including age verification, estimation or inference.

3. Assurance

- Execute the age assurance process for individual users.
- Apply the selected method to establish the user's age or age range.
- Capture input data and produce an assurance result.

4. Categorization

- Assign a **level of confidence** to the assurance result (Basic, Standard, Strict, Enhanced).
- Factor in the accuracy, source quality, binding strength, and resistance to spoofing.
- Document how the result informs eligibility or access decisions.

5. Exchange (Interoperability)

- Securely and privately share assurance results across services or domains where appropriate.
- Enable interoperability between relying parties and age assurance providers.
- Support reuse through cryptographic tokens or trusted frameworks without re-collecting identity data.

Supporting Phases (Continuous)

These three continuous concerns apply **throughout** the above lifecycle:

- **Privacy** – All processes must minimise data collection, avoid unnecessary identity disclosure and align with data protection laws (e.g. GDPR, CCPA).
- **Data Security** – Systems must secure inputs, results and processes against unauthorised access, misuse, and breaches.
- **Interoperability** – Systems should allow age assurance results to be transferred or reused across services in a secure and privacy-preserving way.

Annex A – Levels of Age Assurance (with minimum accuracy thresholds)

IEEE 2089.1 defines four levels of age assurance, each with an expected minimum accuracy threshold, allowing service providers to align their safeguards with the risk of harm and context of use:

Level	Use Case Example	Minimum Accuracy
Basic	Low-risk or general access sites	≥ 90%
Standard	Medium-risk content or features	≥ 99%
Strict	Access to restricted content, chat functions	≥ 99.5%
Enhanced	Access to high-risk content (e.g. pornography)	≥ 99.9%

By adopting common levels of age assurance, interoperability can be offered to consumers because the tokens that demonstrate a user has recently completed an age check will also record the level of age assurance it achieved. So if you proved your age for a low risk use-case, you may need to do a fresh age check for a higher risk use-case, but the same check would be re-usable in situations with a similar level of potential harm.

The standard emphasises that higher accuracy does not necessarily require identity disclosure, as long as the method can securely prove age. These thresholds support risk-based, privacy-preserving design, allowing flexibility for innovation and regulatory alignment.

Summary of ISO/IEC 27566-1

Title: *Information Technology – Age Assurance Systems – Part 1: Framework*

Purpose and scope:

This ISO/IEC standard sets out a **comprehensive framework for the development, evaluation and use of age assurance systems**, covering both verification and estimation techniques. It is intended to provide **global consistency** in the way such systems are designed, deployed and assessed – especially where regulatory or safety objectives apply.

Key Concepts:

1. Types of Age Assurance

- **Age verification:** Confirmation of age against an authoritative source, such as a government-issued ID or certified account.
- **Age estimation:** Use of biometric, behavioural or contextual signals to infer a likely age range.
- **Age inference:** Use of indirect indicators (e.g. owning a payment card) to draw probabilistic conclusions about age.

2. Successive Validation

- The standard encourages **using multiple methods in succession** (e.g. estimation followed by ID check) to improve accuracy and reduce false positives/negatives.

3. Performance Measures

- Accuracy, confidence intervals, bias detection and statistical robustness are required metrics.
- Error rates (e.g. under-age passed or over-age blocked) must be measurable and regularly reported.

Key System Characteristics:

- **Binding** – Age assurance must be linked to the user so it cannot be reused by someone else (e.g. sibling or friend).
- **Portability** – Systems should allow users to carry their verified age between platforms.
- **Transparency** – End users and regulators must be able to understand what method was used and how reliable it is.
- **Context sensitivity** – Different levels of assurance are acceptable depending on the sensitivity of the service or feature.

Governance and Compliance Expectations:

- **Practice Statements** – All participants in the age assurance process (e.g. providers, platforms, relying parties) should publish clear statements explaining how their systems meet the standard's expectations.
- **Auditability** – Systems should be independently auditable, with results available to regulators.
- **Privacy and Data Minimisation** – The standard requires that age assurance methods avoid collecting identity data unless strictly necessary, and that systems do not retain data longer than needed.

Recommended Deployment Practices:

- Age assurance should be applied as close as possible to the point of risk, such as when accessing a harmful feature, not just when registering for a service.
- Systems must be inclusive and non-discriminatory, ensuring that users without access to certain forms of ID are not excluded unfairly.
- The framework is technology-neutral, allowing for methods to evolve over time as new solutions emerge.

Do please let us know if we can be of any further assistance to the Committee.

Yours sincerely

Iain Corby

Executive Director

The Age Verification Providers Association