Age Verification Providers Association
557B Wandsworth Road, London, SW8 3JD
+44 (0) 7811 409769 | iain@avpassociation.com

Position Paper for IAB/W3C/TAG Workshop on Age-Based Restrictions on Content Access

I am the Executive Director of the Age Verification Providers Association (AVPA) which represents companies that deliver privacy-preserving, standards-based, regulator-approved age assurance technologies. I submit this paper *in a personal capacity* with the benefit of that experience, and welcome the opportunity to contribute to the IAB/W3C/TAG workshop.

I am the author of IEEE 2089.1 on online age verification and serve as a UK expert on ISO/IEC 27566-1 and SC27/WG10, helping to shape global consensus on the design of privacy-preserving, trustworthy age assurance frameworks.

This paper offers a contribution to the debate on technical architecture and implementation choices for online age-based access restrictions. Its focus is on the maturity and effectiveness of publisher-level age checks, the emerging standards that underpin them, and the implications of placing responsibility elsewhere in the stack, such as at the OS or app store level.

The workshop scope only includes content but it should be noted that age restrictions Increasingly being applied to data protection and functionality such as talking to strangers online. Concentrating solely on content there is still a wide variety of legislation with for example 24 US states passing laws requiring age checks for adult content each with at least subtle differences in the requirements. I have worked to align these on the basis of international standards but with limited success so far. You will often see similar phrases such as the use of commercially reasonable methods of age verification, methods which rely on transactional data, well the use of government issued identity documents. The UK has taken a more sophisticated approach to defining what it describes as highly effective age assurance and Australia is just completing a trial of technology to provide a scientific basis for its regulations.

Technical requirements are sometimes specified at the level of regulation and guidance but in many cases are left to the interpretation of courts. While there is little case law so far, it is unlikely that those involved in proceedings in US state courts are likely to have technical expertise which may make early cases hard to predict.

The workshop correctly identifies key factors relating to the design of a technical architecture, including, but not limited to, privacy, equity of access, market dynamics (such as centralization), vulnerability to circumvention, cost, accuracy, jurisdiction/geolocation, and censorship. It misses the critical question of *liability* which I will discuss below.

ISO 27566-1 should be the baseline for methods which are used when determining the age of people - it specifies three categories; verification estimation and inference

Identifying content that might need to be restricted is perhaps the least well explored aspect of this problem. It is relatively straightforward for an adult website to know that all its content needs to be restricted but for a social media platform for example to distinguish between a phrase post it in a comment which may be innocent, such as “nice sweater”, and the same two words which maybe the culmination of months of bullying that push the recipient over the edge is a far harder task.

Publisher-Level Age Checks: Proven, Effective and Extensible

Age Verification Providers Association
557B Wandsworth Road, London, SW8 3JD
+44 (0) 7811 409769 | iain@avpassociation.com

Publisher / relying-party-level age assurance is already deployed at scale and has proven highly effective, particularly in regulated contexts such as adult content, online retail, and social platforms. Under the UK Online Safety Act, France's SREN Law, Germany and similar regimes in 24 US States, this approach has been endorsed by regulators, civil society, and courts.

This model places accountability for compliance with legal age restrictions on the digital service. The service itself is the only party in a position to know all the content and functionality on its site, to determine whether laws apply. This structure also enables a clear, contractual relationships between publishers and third-party age assurance providers, rather than relying on signals from other parties in the tech stack with whom there may be no commercial relationship or due diligence. It has also allowed for innovation in privacy-preserving techniques such as zero-knowledge proofs, double-blind exchanges, and verifiable credentials, which are now being codified in international standards, from suppliers which support technical interoperability, data minimisation, auditability, and ultimately, clear accountability.

Other Technical Layers: Complements, Not Replacements

While publisher-level checks are the regulatory norm, age-based access controls can also be implemented further up the stack - at the level of operating systems, app stores, browsers, or network infrastructure. Each of these has potential to reinforce protections.

However, OS- and app store-based models are often voluntary, lack regulatory oversight, and offer little transparency or consistency across jurisdictions. They cannot easily accommodate the diversity of risk profiles and content types that may trigger different age requirements under law. The level of assurance is standardised, not proportionate to risk.

Instead of treating these as alternatives, I see value in a layered approach, where publisher-side controls remain the legal and architectural foundation, supplemented where feasible by indicative platform-level and device-based restrictions and/or parental controls.

Key Considerations in Technical Architecture

I propose that the workshop explore the following principles when comparing and assessing technical architectures for age-based restrictions:

1. **Trust and auditability**: Publisher-level checks offer clear accountability and audit trails via commercial and legal relationships with certified providers.
2. **Privacy by design**: Third-party systems increasingly use cryptographic protocols that deliver strong assurance without revealing identity, location, or behavioural data.
3. **Interoperability**: Verifiable credentials and double-blind exchanges are being standardised to allow age claims to be reused across contexts and jurisdictions.
4. **Layered implementation**: Where upstream controls (e.g. OS, app store, ISP) exist, they may reduce risk by giving an indication of likely age, but publisher enforcement ensures consistency and legal compliance.

Age Verification Providers Association
557B Wandsworth Road, London, SW8 3JD
+44 (0) 7811 409769 | iain@avpassociation.com

5. **Jurisdictional targeting**: Publisher-level controls can adapt content access rules per user geography and risk category without exposing user data.

6. **Market fairness**: Publisher integration allows SMEs to comply using modular solutions without relying on dominant OS or app ecosystem vendors. Some industries are reliant on discriminating between adults and children and should not become dependent on dominant tech companies to operate legally.

A Note on Liability and Control

Placing responsibility at the publisher level aligns liability with content control. The publisher defines the access rules, integrates the verification system and faces enforcement if non-compliant. Shifting enforcement to the OS or app store level could centralise control in a handful of private actors, weaken transparency, and undermine regulatory reach.

Moreover, the publisher-relying party model offers regulatory resilience: it can be mandated, audited and independently certified in ways that upstream layers, governed by voluntary ecosystem policies, cannot.

Shaping Global Practice

I have been a key contributor, through my role at the AVPA, to the development of age assurance policy and technology worldwide. It co-developed a cross-border pilot system under my leadership, AgeAware by euCONSENT, and have advised the UK, French, Australian, multiple US state and Irish governments, as well as the European Commission and Ofcom, on how to implement age assurance responsibly and effectively.

The way forward lies in tokenised, double-blind solutions that decouple identity from age verification and prevent both content providers and verification services from learning more than necessary. Systems like **euCONSENT's AgeAware** model demonstrate how cryptographic tokens and two-factor reusability can be implemented to provide seamless, privacy-preserving proof of age across multiple services. These approaches allow users to verify their age once and reuse that credential without revealing their identity or being tracked across sites. This aligns with the principles of data minimisation, technical neutrality and user control, and reflects the direction of emerging standards such as IEEE 2089.1 and ISO/IEC 27566-1. Future technical architectures should encourage this model of decentralised, standards-based interoperability rather than reinforcing siloed or centralised solutions.

Age Verification Providers Association
557B Wandsworth Road, London, SW8 3JD
+44 (0) 7811 409769 | iain@avpassociation.com

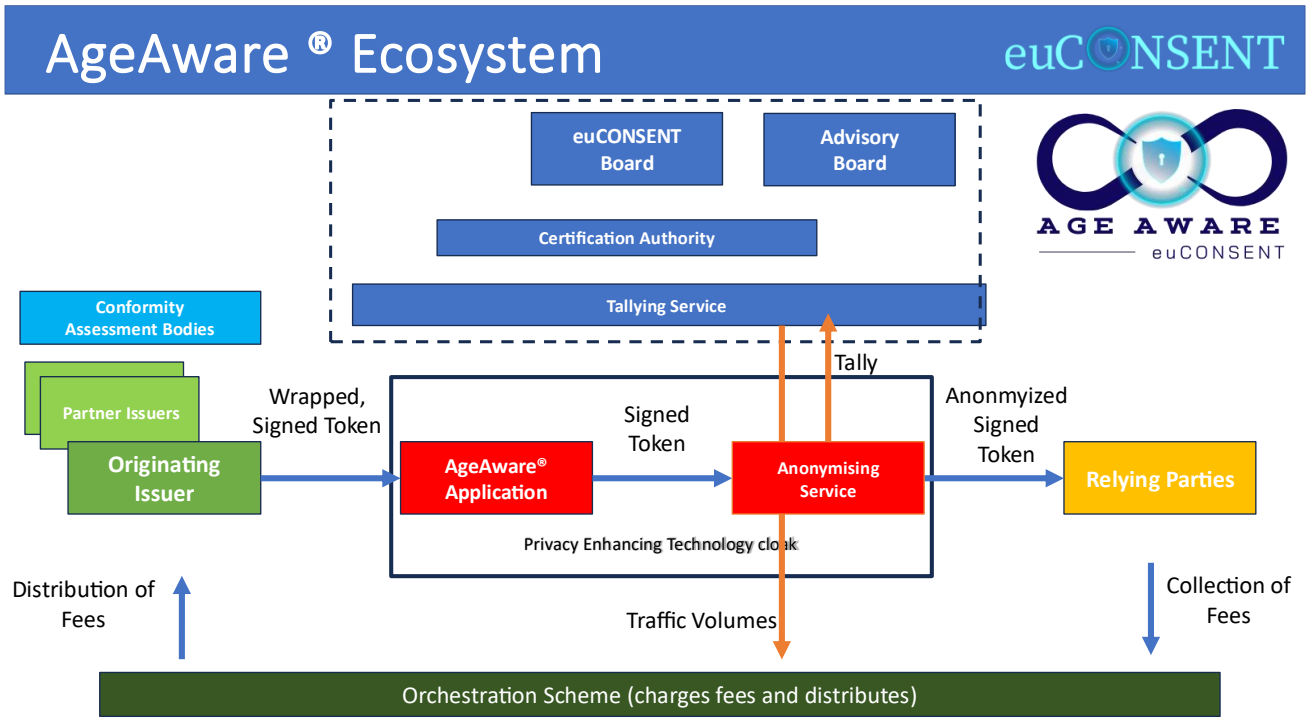

I would welcome the opportunity to join this workshop, to share these insights and help ensure that future architectural principles support privacy, innovation, and child protection across the Web.

Submitted by:
Iain Corby
iain@avpassociation.com
Author, IEEE 2089.1
UK Expert, ISO/IEC 27566-1 & SC27/WG10