The Age Verification Providers Association

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

Ms. Kate Thwaites
Chair, Joint Select Committee on Social Media and Australian Society
PO Box 6100
Parliament House
Canberra ACT 2600
socialmedia.joint@aph.gov.au

24 June 2024

Dear Ms. Thwaites,

Age Verification Providers Association Submission for the Inquiry into Social Media and Australian Society

We are the global trade body representing 30 providers of privacy-preserving online age assurance technology and appreciate the invitation to make this submission to your inquiry. We will address two parts of your terms of reference:

(a) the use of age verification to protect Australian children from social media; and
(e) issues in relation to harmful or illegal content disseminated over social media, including scams, age-restricted content, child sexual abuse and violent extremist material.

Summary

It is relatively straightforward to use age-assurance technology to impose any prescribed minimum age for accessing digital services, or indeed to adjust content and functionality within a service to ensure an age-appropriate experience for children of different age ranges.

This can be achieved:

- **Accessibly**, through offering a wide range of methods of age assurance which mitigate the risk of digital exclusion
- **Privately and securely**, with the user's identity protected through the application of privacy-enhancing technologies (PETs), such as zero-knowledge proof (ZKP), and by applying the principle of data minimization so no new central databases of personally identifiable information (PII) are created
- As **stringently** as required, as defined in the latest international standards, but best applied in proportion to the risk of harm, given there is a trade-off with the user's experience
- **Interoperably**, and thus more conveniently, with one age check used across multiple platforms
- **Cost effectively**, not least as a result of the re-use enabled by interoperability

We will also address common misconceptions that attempts to apply age restrictions online can be easily circumvented through virtual private networks (VPNs), or are undermined by shared devices, and we will consider alternative approaches which are often proposed by opponents of age assurance which we do not believe meet the same policy objectives as effectively.

We would be pleased to give evidence to the Committee in person or remotely if it has further questions.

The Age Verification Providers Association

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

Accessibility

There are a wide range of existing methods of age assurance and others which the Australian government could facilitate. This gives users a choice, so they can exercise their preferences based on their personal priorities and available data sources. This choice will also mitigate the risk of digital exclusion which could arise if, for example, only citizens with a passport were able to prove their age.

Existing methods of age assurance are generally divided between verification and estimation. These terms are defined by the Committee Draft of ISO 27566 A Framework for Age Assurance as follows:

age assurance

process to establish confidence that the output of a system or a method meets the requirements for an age-related eligibility decision

age verification

age assurance method based on calculating the difference computation between the current date and the confirmed date of birth of an individual

age estimation

age assurance method based on using inherent features or behaviours to estimate the age or age range of an individual

The existing methods of age verification include:

Government-issued physical ID – passports, driving licences, military/veteran ID etc. can be read remotely with either optical-character recognition technology or by reading data from chips embedded in the document; this is compared to a selfie image of the user, which itself is checked for liveness to prevent presentation or injection attacks where fake or altered images are used to evade the check. This electronic identification validation technology (eIDVT) is already widely in use around the world for checking identity. Our sector only needs to extract the user's age, and all other personal data is then deleted by the age assurance provider to comply with the principle of data minimisation. Some of our members can perform this process entirely remotely, on the user's own device, so none of their personal data apart from age ever leaves their own control.

Digital ID – increasingly users have access to either government issued or privately provisioned reusable digital ID. The user's age can be selectively disclosed to a third party, such as an age assurance provider or indeed, directly to the digital service that needs to check age, without any other PII being passed on.

Bank Records – through the advent of open banking, such as Australia's ConnectID programme, age can be confirmed by the user logging into their online banking and agreeing to share it with a third party[1].

[1] https://connectid.com.au/

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

Mobile Network Operator Records – subject to an audit of data quality, MNOs may also have age data on their users. For example, in the UK, all new mobile devices are issued with a block on adult content that can only be removed when the user proves they are 18 or older. (We discuss this as a proposed alternative to age assurance below.) Age assurance providers can therefore confirm if a mobile phone user is an adult by contacting their MNO.

Authoritative Databases – this was a popular early method, where a user supplied details such as name, address and date of birth and this was confirmed with a credit rating agency. Because such details are easily known to others, this method is more suitable for when age-restricted products are ordered online, as they will then be delivered to the address supplied in this process, so it is harder to just borrow someone else's information. It is also suitable for wagering, as winnings can be sent only to the person named in the age checking process, making using someone else's name pointless.

Existing methods of age estimation include:

Facial age estimation - Originally, age assurance was developed solely to determine if a user was an adult, to address use-cases such as the purchase of alcohol or access to online wagering and adult-only content.

More recently, a need to assess the age of minors has arisen, driven by a desire to enforce not only minimum ages for the use of social media, but also data protection laws where younger children may, for example, require parental consent before agreeing to share personal data. Obviously, not all of the methods described above are suitable, in whole or in part, for children. Fewer children than adults will have a passport, none will have a credit record or a driving licence, only some will have a bank account etc.

Facial age estimation emerged as a more accessible option. The algorithms used to estimate age are improving by the day, and six examples were recently tested by the US Federal Government National Institute of Standards and Technology (NIST), and we would commend their report to the Committee as evidence for consideration in its own right[2]. One of our members, Yoti, publishes regular White Papers disclosing the accuracy of their estimation tool[3].

[2] https://pages.nist.gov/frvt/reports/aev/fate_aev_report.pdf
[3] https://www.yoti.com/wp-content/uploads/2024/04/Yoti-Age-Estimation-White-Paper-December-2023.pdf

The Age Verification Providers Association

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

Mean Absolute Error by age band

Yoti facial age estimation accuracy — Mean estimation error in years split by gender, skin tone and age band

Gender	Female				Male				All
Skintone	Tone 1	Tone 2	Tone 3	All	Tone 1	Tone 2	Tone 3	All	
6-12	1.3	1.4	1.7	1.4	1.2	1.3	1.4	1.3	**1.4**
13-17	1.3	1.5	1.7	1.5	1.0	1.4	1.6	1.3	**1.4**
18-24	2.4	2.3	2.4	2.4	1.9	1.9	2.0	1.9	**2.1**
25-70	2.8	3.2	3.9	3.3	2.6	3.1	3.2	2.8	**3.0**
6-70	2.5	2.7	3.3	2.9	2.2	2.6	3.0	2.6	**2.7**

A degree of error is inevitable, so a policy-decision is required to accept an approximation, knowing some underage users will be false positives. The proportion of false positives can be reduced by testing if a user appears to be a couple of years older than the minimum age – the difference being termed a "buffer", and its width determining with statistical certainty the expected proportion of false positives. NIST specifically studied this in their analysis.

Email address analysis - Email analysis has also proven to be an effective method. Another of our members, VerifyMy, has just published a White Paper which sets out the accuracy that has achieved4. Out of 847 individuals actually under the age of 18 who were tested, they estimated only 19 to be older than they were (2.24% false positive rate of the testing set), and no-one's age was overestimated by more than 2 years.

Estimated minimum age	Actual age 6	7	8	9	10	11	12	13	14	15	16	17	Total
Insufficient data	6	5	8	3	3	2	5	10	17	15	28	23	125
7	1	1	1	1	1	0	0	0	0	0	0	1	6
8	0	0	3	6	5	5	4	6	8	0	3	2	42
9	0	0	1	4	6	6	2	8	4	1	4	3	39
10	0	0	0	0	6	4	7	14	10	1	3	2	47
11	0	0	0	1	1	9	16	16	11	13	3	4	74
12	0	0	0	0	0	0	15	23	16	15	3	0	72
13	0	0	0	0	0	0	1	19	22	13	15	12	82
14	0	0	0	0	0	0	1	2	26	63	23	7	122
15	0	0	0	0	0	0	0	0	2	15	47	18	82
16	0	0	0	0	0	0	0	0	1	1	40	64	106
17	0	0	0	0	0	0	0	0	0	2	3	43	48
18+	0	0	0	0	0	0	0	0	0	0	1	1	2
Total	7	6	13	15	22	26	51	98	117	139	173	180	847

False positive: overestimated age (grey figures)
True positive: didn't overestimate age (yellow figures) - therefore, restricting access to something the user shouldn't be able to access.
Insufficient data: where we do not have enough data to provide a meaningful response or reliably estimate the user's minimum age. This can be due to the email address being newly created, invalid or rarely used, for example.

[4] https://verifymyage.com/email-address-age-estimation

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

Where estimation is used, there is a need to provide alternative methods to correct false negatives, for example when a 14 year-old has been estimated to look under 13 so is initially denied access to social media (assuming 13 is the applicable minimum age). In this case, the user might first be asked if they do have a suitable ID document, such as a passport. Ideally, governments would also facilitate access to the many sources of authoritative data they hold on children's ages – schools, social security benefits, healthcare etc. But there must also be a third level – age assurance of last resort – to ensure that no child is excluded simply because they look too young and lack paperwork – that can be achieved through professional vouching. Social media site Yubo has for several years also addressed this through allowing parents to contact them and after a short interview be relied up to vouch for their child's age as well.

But for the vast majority of people who are by definition more than +/- 2 years of any given minimum age, age estimation is a quick, convenient, privacy-preserving and effective method of age assurance.

Privacy and security

When age verification was first considered for adult sites, they quickly realised their users would be reluctant to share personal data with them directly and so the concept of using a third party to check age and then confirm simply yes or no as to whether a user was old enough to access a site became the foundation of the age assurance industry's approach to privacy. Obviously, if a provider is compliant with relevant data protection laws, this should be sufficient to protect the privacy of the user. However, given the sensitivity, regulators, such as the French CNIL have become proponents of double-blind solutions, which technically guarantee that it is impossible for the adult site to discover the identity of the user and also that the user's online behaviour cannot be tracked by the age verification provider. The age assurance industry has been considering how this can be applied, although the theoretical model in its purest form makes it impossible to sustain a commercial market, simply because the age assurance providers would not be able to know who had used their services and could not therefore charge their clients.

We have, through the nonprofit euCONSENT project, developed an industry wide ecosystem which will allow for billing, but which maintains the guarantee of anonymity through the use of privacy enhancing technologies. This innovative approach is the one we are advocating that the Australian pilot of age assurance technologies should include within its scope.

Stringency

Not all age checks are created equal. Their liability can vary based not only on the level of accuracy but also on how well a system prevents circumvention by deepfakes, or how often it authenticates that a user is still the same user who completed a previous check. We have developed in partnership with international standards bodies such as IEEE and ISO, detailed descriptions of different levels of age assurance, which provide a common understanding and which we hope lawmakers and regulators will adopt. Securing alignment around the world will make it a lot easier to deploy global solutions for global platforms and will ultimately keep costs to both those services and their users as low as possible.

In general, the more accurate and reliable a check is compared to the more expensive it will be to perform and when accuracy and reliability are improved, this usually comes at the cost of increased

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

friction, leading to a poorer user experience, so regulators must as ever weigh the balance between cost and effectiveness when they set the rules.

We do not argue for or against any particular approach – we suggest instead that policymakers, perhaps best articulated through detailed regulations, specify the level of accuracy they require for different use cases. For example, we recently recommended to the UK regulator, Ofcom, that it defines "Highly Effective Age Assurance" which is reserved for the most significant forms of harm such as pornography and content relating to suicide and self-harm, as:

"Highly effective age assurance systems must demonstrate that their certified expected outcomes are such that more than 95% of children under 18 are prevented from accessing primary priority content, and more than 99% of children under 16 are prevented."

This approach not only limits the proportion of false positives, but it also indicates the acceptable standard deviation of the results – a system which allows some 17-year-olds to pass as 18 is more likely to be considered sufficient than one which allows the same number of 7-year-olds to be falsely accepted at 18+. A similarly constructed regulation could apply to social media, but set at less exacting levels of expected outcome accuracy.

Interoperability

Historically, each digital service required its own age check. That was tolerable for the occasional purchase of alcohol from a new online supplier, or to open an account with an adult website. But as the requirement to prove your age online has grown, doing so afresh each time is not a tenable proposition.

Interoperability greatly enhances user convenience by streamlining the age verification process. With a single age check, users can verify their age once through one age assurance provider, and then this verification can be reused on various platforms by providers recognising previous age checks. This reduces friction, saves users' time, shares the cost, and makes the overall experience much smoother and more efficient.

The age assurance industry has been developing interoperability for over five years. During 2021-2022, intense effort was facilitated by a grant from the European Commission, instigated by the European Parliament. A consortium of age assurance suppliers, academics, auditors and ourselves delivered a project called euCONSENT[5]. This mirrored the existing European digital identity solution at the time, eIDAS. It allowed age assurance providers to offer users a token to keep on their device when they completed an age check, so that in future, other providers would be able to see that a check had already been processed, and could ask the provider of that check to re-confirm age eligibility, rather than involving the user again. The project ran a largescale pilot with some 1600 adults and children across five countries, and was considered a successful proof of concept.

Since then, new requirements have emerged from EU regulators such as the French CNIL's emphasis on "double-blind" cryptographically enabled checks, and the Spanish AEPD's preference for checks to be processed on a user's own device. This has driven an upgrade to the euCONSENT infrastructure, re-christened AgeAware®, that delivers a double-blind, device-based,

[5] www.euCONSENT.eu

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

signed tokenized approach. The industry's design also facilitates, through a Tallying Service, a sustainable commercial business model for suppliers of age assurance services which was not obviously possible with the proposals from the two aforementioned data protection authorities.

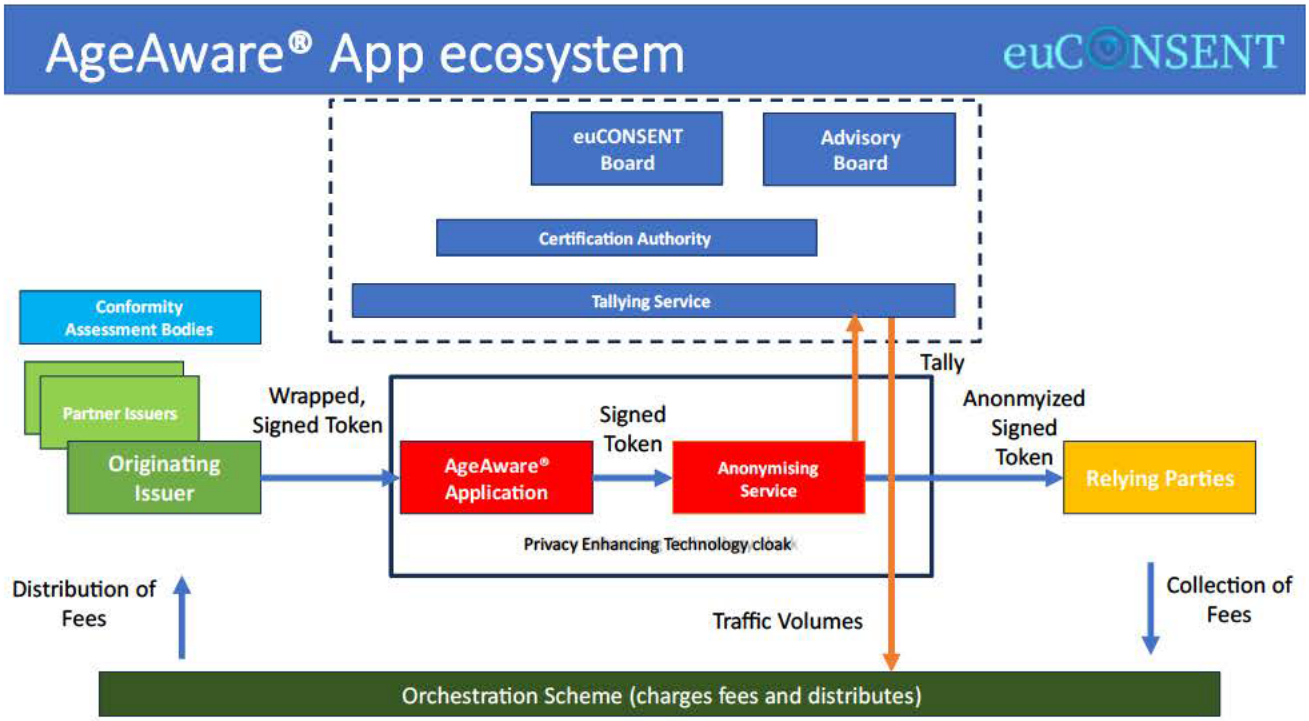

Cost effectiveness

The earliest age verification services were essentially the same as full identification checks, and pricing five years ago may have been $1 or more per check. Competition and technological advances have had a dramatic impact on that. As a trade body, we avoid any discussion of pricing with our members, but we can quote the UK government's impact assessment for the Online Safety Act, published in 2022, which, having conducted extensive surveys of our members, estimated a cost per check of 12 pence, and predicted the price would fall further, not least as a result of the re-use enabled by interoperability[6].

Industry practice is to charge for an initial age check, but it is rare to then re-charge for subsequent checks where a user is merely re-authenticating themselves – just proving they are the same user who completed a previous check. So, costs can generally be considered, prudently, as an annual fee per user. Obviously, if a user is incognito, and does not wish to open an account with a digital service, then a fresh check is required each time. But the tokenized solution the industry has agreed to adopt will now mitigate this otherwise recurrent cost.

[6] https://assets.publishing.service.gov.uk/media/6231dc9be90e070ed8233a60/Online_Safety_Bill_impact_assessment.pdf

The Age Verification Providers Association

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

Misconceptions

We will also address the common misconceptions that attempts to apply age restrictions online, that they can be easily circumvented through virtual private networks, or are undermined by shared devices:

Shared devices - a frequent concern is that if a user proves their age on a device which is then shared with another user who may be underage, then the controls will fail. This can be completely avoided but only by requiring a user to prove that they are the same user who completed the age check very frequently. That is clearly going to be inconvenient and so pragmatically a balance must be found, but we would suggest regular authentication as this process is known, should be required with the regularity determined in proportion to the risk of harm. For example, it may be necessary to do a new age check or re authenticate an existing check each and every time you buy a hunting knife online but perhaps this should only be required once a week for accessing adult content.

To some extent, this is no different from the real-world risk that an adult magazine is left on the coffee table for a child to discover. Parents need to take some responsibility to keep harmful materials away from their children and that applies online as well as offline. To achieve this online, adults should remember to log out of age restricted websites and applications if they know the device may be used by a child. This is not an inherent flaw with the technology, as we know we could, in theory, prevent it altogether, but it is a matter for policymakers to consider when determining an appropriate balance of effectiveness versus convenience.

Virtual Private Networks - we often hear the claim that children can easily circumvent age restrictions by downloading a virtual private network or other location spoofing technology which allows them to pretend they are in a jurisdiction where age restrictions are not required. That is, of course, true, but we have not yet seen any laws passed anywhere around the world which provide a get out of jail free card to age restricted sites if they allow underage users to access their services, however those users connect to the Internet. It is possible for services to detect when a user is deploying a virtual private network, so such traffic can either be blocked or age checks can be applied to all traffic from these sources. Alternatively, users can be asked to prove that they are in a jurisdiction that does not require age assurance, which is the approach taken when a person wishes to place a bet and needs to demonstrate that they are in a state which permits online gambling. It is, in effect, easier to prove where you are, than where you are not.

Overseas enforcement - As has been illustrated recently with the controversy about compliance with existing online safety laws in Australia, some international services may feel exempt from Australian jurisdiction. This problem has been tackled directly in the UK by giving the regulator powers under the online Safety Act to not only block access to particular sites but also to require critical business support services such as advertising, hosting and payments to be withheld from non-compliant services. We saw when MasterCard introduced new rules requiring adult websites to check the age and consent of performers, known as a AN5196, this had a dramatic impact on the operations of sites around the world as they feared the loss of major revenue streams.

Alternatives

A number of alternative approaches are proposed, often by opponents of age assurance, but it is important to recognise that some do offer additional safeguards, and it may be wise to adopt a

The Age Verification Providers Association

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

layered approach, deploying a range of protective measures for children when they are online. Having considered them all carefully, we will still put forward the case to the Committee that online age assurance offers the most comprehensive policy response.

Device Based Filtering - All the leading operating systems offer parental controls which can limit screen time and access to particular sites and applications. The use of these is generally at the discretion of the parent, so the first point to note is that this is a different policy measure from age assurance which makes it less straightforward for a parent to enable their child to access platforms and websites under a legal minimum age.

Secondly, filtering is not an exact science. It either relies on blacklisted sites which will either self-identify, for example through the use of the restricted to adult (RTA) label or need to be identified by third party surveillance services. Or it relies on a specific list of approved sites and services. These may be under inclusive or excessively restrictive. There are AI based options for spotting nudity, for example but if the aim is to address the full range of potential harms to children, the level of sophistication required in automatically screening content and functionality is well beyond anything which is available today.

This approach also applies to the device not to the user and so once the device has had the controls removed anybody borrowing, sharing or inheriting the device will not benefit from any protection unless the controls are re-applied.

Internet Service Provider Filtering - this is very similar to device based filtering but is not within the control of the parent so is arguably a more comprehensive approach, but suffers from the same challenge of knowing which content should be permitted and which should be prevented. The ISP will also not know which users are adults and which are children so this will tend to apply to all users in a household which may be considered overly restrictive and could lead to parents removing the controls in the interests of their own freedom with the obvious impact on their children.

App Store restrictions - it is possible to apply age assurance at the point when a user accesses the App Store on their device so that only age appropriate apps can be downloaded. However, this is a binary choice that may suit apps which are clearly intended for adults only, or indeed those we wish to restrict for use only by children, but many apps offer access to a wide variety of content and functionality some of which may be entirely suitable for children and other aspects of which may be harmful. It is a blunt instrument to prevent access to an app altogether rather than seeking to selectively restrict access based on the particular content or functionality within the app.

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

Conclusion

Society has imposed age-restrictions for the past 100 years. As our lives move increasingly online, it is not surprising that there is a widespread desire to apply those same controls on the Internet. Our industry was created to enable this, mirroring the real world including the way a cashier might be satisfied a customer looks well over 18 before selling alcohol without requiring proof, and how bar staff only concern themselves with a drinker's age, not their full identity.

We were fortunate to accelerate the development of our technology through the foresight and sponsorship of the European Union. Australia, through its proposed pilot of age assurance, now has the opportunity to take up the mantle and deploy the next generation of this technology, which will incorporate important concepts such as device-based, double-blind design.

Yours sincerely

Iain Corby
Executive Director
The Age Verification Providers Association