Security Considerations for Age Assurance Technologies

Andrew Shaw, Senior Internet Standards Researcher, UK National Cyber Security Centre

This paper presents an overview of security considerations for the use of age assurance technologies online, primarily from an architectural perspective. This paper will also consider current IETF work that could be relevant to this problem space. This paper is written in the context of the UK legal and regulatory environment but does not represent UK government policy. This paper does not make any judgment on the merits of, or issues that may affect, age assurance systems currently in deployment.

Age Assurance Techniques

There are various ways of assessing the age of a user of a web service. Some methods will provide verification of an age (for example verifying the provision of a government photo ID card), while others will provide an estimate of the age or a likelihood of the person being assessed meeting the required age criteria. These definite methods of age verification may not always be suitable or possible, so age assurance technologies, which may only provide a likelihood, offer a complementary alternative in restricting access to content by age.

In the UK, the regulator, Ofcom, lists the following age assurance techniques as capable of being highly effective (Ofcom, 2025):

- Open banking
- Photo-identification (photo-ID) matching
- Facial age estimation
- Mobile-network operator (MNO) age checks
- Credit card checks
- Email-based age estimation
- Digital Identity Services

Some of these techniques rely on the verification of government-issued documentation (photo-ID matching, Digital Identity Services), others rely on use of heuristics based on personal data (email-based age estimation, Open banking), others rely on third parties having already performed age assurance (credit card checks, MNO checks), and others use more involved processing to assess age (facial age estimation). While these are techniques listed by Ofcom in the UK, other jurisdictions may have different assessments of the capability of different techniques.

Each of these techniques, and others that may be used in other jurisdictions, come with their own set of considerations around accuracy, robustness, privacy and security. For those techniques that rely on less precise methods or use of AI, then the accuracy and robustness of the methods used is particularly important to understand. For example, for facial age estimation, understanding how possible it is to trick the software into giving an incorrect answer or whether the system has systematic biases affecting its accuracy. It is not necessary for these methods that do rely on a likelihood to be 100% accurate to be suitable for deployment. In practice, those techniques that do

provide relative certainty are also possible to fool (e.g. a child could use their parent's credit card), so drawing a distinction on that basis when deciding what is suitable for deployment is not particularly helpful. Instead, each technology should be considered on its own merits.

Age Assurance Architectures

There are various different architectures for implementing age assurance. We will consider a few different, high-level, options for architecting these systems.

First party

Providing proof of age to a web service directly provides the simplest architecture and minimises the number of parties involved.

Trust Relationship - Where a user already has a trusted relationship with the service provider this may not require any additional personal information to be shared. For example, if a user has already made purchases with a credit card from a web service, then using the possession of a credit card as a way to indicate the likelihood that the user is over 18[1] provides no extra personal information to the web service. It also minimises the number of parties involved in performing the age assurance. However, if providing information directly to the service, the user needs to trust the service provider to store their information securely, assess age accurately, and not to misuse their personal data. Additionally, if the user has not already trusted the service with this information then this does lead to the service provider obtaining more information on the user.

Implementation - Implementing the necessary infrastructure may be difficult for the service provider, especially if an established provider of age assurance services isn't used. A lack of resources or expertise in implementation could lead to a less reliable age assurance system, and hence assurance failures. It also places even more onus on the service provider to secure the personal data correctly. A compromise of a service provider could yield data that links user's personal data to their activity on the service. Depending on the nature of the service, this could make it an especially attractive target for malicious actors, for example a compromise of information on a site hosting adult content could be used for blackmail by malicious actors.

Trusted third party

A web service could use a mutually trusted third party to provide age assurance services rather than a web service providing them directly. A third-party solution could take the form of an on-device solution or a remote solution. There are various different ways of implementing an on-device solution, and each which will carry different considerations both in terms of parties involved and the efficacy of the solution.

Trust relationship - The primary advantage of using a trusted third party, rather than having users share personal data directly with the web service, is that the user no longer needs to trust the web service to handle correctly any personal data provided as part of age assurance. Instead, they should only need to trust the third-party age assurance service used. For many web services that ask users

[1] While 18 is generally the lower age limit for a credit card in one's own name, provision of credit card details does not prove age – e.g. a child may use a parent's card with or without permission, or an under 18 may have legitimate access to business credit card.

to provide age assurance, it is likely that users will trust a third-party service more than the web service itself.

Execution on-device – For on-device solutions, the user will not only need to trust the age assurance provider with their data, but also to not run malicious or unexpected code on their device. This risk can be mitigated by using a solution provided by an entity already running code on the user's device, for example the Operating System. For code running on-device, methods of assuring the code is trusted and as expected, such as code-signing, will be particularly important.

Partition of data – For maximal privacy, a third-party provider should not be able to learn which service a user is visiting. In practice, this will involve use of a privacy-preserving technique, as discussed below. In the absence of these techniques, the third party needs to be trusted by the user not to find out or store this information. This would ensure that compromise of personal data from a third-party service would not be linked to the web services a user was using.

Centralisation of data - If a trusted third party is storing the personal data of large numbers of users, either deliberately or accidentally, then the age assurance service presents an attractive target for malicious actors. A breach and compromise of any stored data could enable mass exploitation across users from a large number of services.

Centralisation of capability - Centralisation of services provides a greater incentive for malicious actors to trick or compromise their age assurance techniques. This presents a different risk to compromise of personal data, as it could affect the effectiveness of the age assurance itself, as opposed to the privacy of specific users. This could lead to a Denial-of-Service attack, either by preventing age assurance checks from running or by compromising a technique to never return a successful result. Alternatively, a malicious actor could compromise the use of age assurance as a safety measure by ensuring that certain users will incorrectly pass age assurance checks. There are also increased resiliency risks if age assurance checks are consolidated in a small number of providers, with an outage at one provider having the potential to be very wide reaching. However, this same concentration means that passing an age assurance check once could then be sufficient for access to multiple services. This has benefits for useability, as users don't tire of providing age assurance, and means that users only need to share personal data with one party, rather than many.

Privacy-preserving solutions

Age assurance techniques all require the user to share personal information with another party in some form – be that documentation, an email address, an image of themselves or other personal information. Privacy-preserving technologies can therefore be a useful tool to allow users to prove their age without unduly disclosing personal data to parties with whom they do not want to share it.

Zero-knowledge proofs (ZKPs) are one such technique used in providing privacy-preserving age assurance (Google, 2025). A ZKP is an advanced cryptographic technique that allows a party to prove the possession of some secret knowledge without revealing it (NCSC, 2025). This can be applied to the age assurance problem in such a way that instead of providing a direct proof of age to a web service, they can provide a zero-knowledge proof that they meet the relevant age restrictions. ZKPs are well-understood, secure cryptography that have been widely applied to anonymous credentials.

Ultimately, to use a privacy-preserving solution the user still needs to provide evidence to a third party of their meeting the relevant age restriction. However, using a solution like ZKPs allows the user to constrain the information they provide to the web service they want to access and constrain the information the third party receives about their request. For example, while they may provide ID

with their full date of birth to the third-party age assurance provider, the web service could just receive a ZKP that the user is over 18, reducing any privacy leak. The protocol can also be constructed such that the third-party age assurance provider learns nothing about the web service the user is visiting or the relationship between the user and that web service – all that's required is that the web service trusts that age assurance provider.

Use of ZKPs, or another advanced cryptographic technique, to provide these privacy-preserving properties is not sufficient for security of personal data. It is still necessary for parties to secure their services, and user personal data, properly. Partitioning information in this way may reduce the impact of a breach, but the considerations for security of third-party providers outlined above all still apply.

Relevant IETF Work

Without addressing the challenge of age assurance directly, the IETF has some strands of work that could be relevant to this problem. In particular, IETF technologies could form part of privacy-preserving solutions for age verification.

An example of an IETF technology that could be used as part of a privacy-preserving age assurance solution is the work of the Privacy Pass Working Group (IETF Privacy Pass WG, 2025). The Privacy Pass architecture is designed to be very flexible and can be deployed in multiple different ways. One of the models it supports is the issuance and redemption of anonymous, unlinkable tokens that allow a client to prove certain properties about themselves without revealing their identity. A pen picture of a model that this architecture could support is the following:

> A client visits a website (Origin) and is challenged for proof of age. They then contact an attesting party (Attester), to whom they prove their age. The attesting party then forwards a client token request to the issuing party (Issuer). This issuing party can then issue a token response which the attester forwards to the client. From this token response, and the challenge given by the website, the client can compute a token. The client can now connect to the website again and provide the token as authorisation that they meet the age restriction that the website has set. No party other than the client is aware of both the detailed personal information and the website visited.

For Privacy Pass, and other similar technologies that rely on the partitioning of information to provide privacy or anonymity, there are additional assumptions around non-collusion between parties that need to be upheld for all of the desirable privacy properties to hold. Given the sensitive nature of personal data involved, understanding the relationships between the entities operating the parties is particularly important in this case.

Other work in IETF and IRTF could also be relevant to age assurance problems, including work on advanced cryptographic algorithms in CFRG which could sit behind privacy-preserving solutions. The work of the Privacy Preserving Measurement (PPM) Working Group (IETF PPM WG, 2025) is not directly applicable to privacy-preserving age assurance, but shows further precedent for the application of advanced cryptography to well-scoped problems within the IETF.

Security impacts of circumvention

In general, it will be possible for a motivated, technically capable, user to circumvent age assurance techniques. This could either be by tricking the age assurance technique itself, or, more likely, by bypassing it to access the content protected by an age assurance check at a different web service

that does not implement robust age assurance. However, this does not render such checks without merit in achieving the goal of preventing inappropriate access to content. This is in part because incremental adoption of these technologies will reduce the ability to bypass these measures over time. Importantly, a measure that is non-trivial to circumvent will still deter or prevent a proportion of users from bypassing it and accessing the age-restricted content.

Controls based on age assurance will only limit access to resources that are generally available, rather than providing any authentication or authorisation of particular users, so circumvention of such measures should not have a direct impact on the security of the users or services involved. There is a risk that adult users may choose to use VPN software with poor security or privacy guarantees in order to circumvent these measures (BBC News, 2025). However, use of VPNs with poor security properties is not a new risk introduced by age assurance.

References

BBC News. (2025). *VPNs top download charts as age verification law kicks in*. Retrieved from BBC News: https://www.bbc.co.uk/news/articles/cn72ydj70g5o

Google. (2025). *An age assurance tool for Europe and beyond*. Retrieved from https://blog.google/around-the-globe/google-europe/age-assurance-europe/

IETF PPM WG. (2025). *Privacy Preserving Measurement WG*. Retrieved from https://datatracker.ietf.org/wg/ppm/about/

IETF Privacy Pass WG. (2025). *Privacy Pass*. Retrieved from https://datatracker.ietf.org/wg/privacypass/about/

NCSC. (2025). *Advanced Cryptography*. Retrieved from https://www.ncsc.gov.uk/whitepaper/advanced-cryptography

Ofcom. (2025). *Guidance on highly effective age assurance for Part 3 Services.* Retrieved from https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/statement-age-assurance-and-childrens-access/part-3-guidance-on-highly-effective-age-assurance.pdf?v=395680