AffixIO Technical Paper · WP-024
June 2026
affix-io.com

AFFIXIO WHITE PAPER · WP-024

# Social Media Age Restrictions: Compliant Age Verification Without the Privacy Trade-Off

Laws in Australia, the UK, France, the EU, and the US now require platforms to keep under-16s out. Here is how to do it without turning every sign-up into a data collection exercise.

**AffixIO** | United Kingdom | affix-io.com | June 2026

ABSTRACT

The world's major democracies have, in under two years, moved from debating whether to restrict children's access to social media to passing laws that require it. Australia, France, the UK, and the EU all now require platforms to verify or assure the age of users before granting access to content and features that carry risk for younger users. What none of these laws specifies is exactly how platforms should do it. That is where the compliance challenge lies: age assurance that actually works, at scale, without creating a privacy liability larger than the harm it is preventing. This paper sets out the current global regulatory

> landscape, explains the technical options available to platforms, and describes how privacy-preserving age assurance closes the gap between legal obligation and user protection.

CONTENTS



SECTION 1

## The Global Wave of Age Restrictions

For most of the past decade, the debate about children and social media moved slowly. Researchers documented harms. Governments consulted. Platforms updated community standards. Very little changed. That has now shifted, rapidly and across multiple jurisdictions simultaneously.

In November 2024, Australia passed the Online Safety Amendment (Social Media Minimum Age) Act 2024, the world's most explicit legislative ban on under-16s accessing social media platforms. The Act places the obligation on platforms rather than parents and requires platforms to take reasonable steps

to prevent underage users from creating accounts, with fines for non-compliance. The legislation came into force in stages through 2025, with the eSafety Commissioner developing technical standards through 2026.

France followed a similar path, building on its 2023 law requiring parental consent for under-15s and advancing towards a harder age verification requirement for under-15s that came into effect in 2026. The French approach, shaped by the work of ARCOM (the Autorite de regulation de la communication audiovisuelle et numerique), has emphasised the development of a government-backed age verification infrastructure that platforms can integrate rather than requiring each platform to build its own system from scratch.

In the United Kingdom, Ofcom began enforcing the Online Safety Act 2023's age assurance provisions from July 2025. The Act requires platforms to implement age assurance for content that is harmful to children, with the Protection of Children Codes of Practice setting out what steps platforms in scope must take. Ofcom's guidance makes clear that self-declaration of age does not satisfy the requirement and that platforms must implement technically robust age assurance measures.

In the European Union, the Digital Services Act (DSA) and GDPR Article 8's age of digital consent requirements have long required platforms to obtain verifiable parental consent for users under 16, with member states permitted to lower the threshold to 13. The DSA adds obligations on very large online platforms (VLOPs) to conduct specific risk assessments related to minors and to implement systemic risk mitigation measures, including access restrictions where appropriate.

In the United States, the Kids Online Safety Act (KOSA) and the Children and Teens' Online Privacy Protection Act (COPPA 2.0) passed in 2024, creating federal obligations around the design of platforms used by minors, the collection of their data, and the types of features and content that may be served to younger users. Several US states have enacted their own age verification requirements in advance of, or in addition to, the federal framework, including California's Age-Appropriate Design Code Act, which took effect in 2024.

The picture across OECD countries is one of convergent direction with divergent detail. More than 25 countries now have either enacted or are actively legislating age restrictions on social media access for under-16s or under-18s, or age-gating requirements for specific categories of content. Platforms operating globally cannot treat this as a regional concern. It is a platform-wide compliance obligation that requires a coherent, scalable technical solution.

SECTION 2

## What Each Law Actually Requires

Reading the actual statutory language matters here, because the media discussion of these laws often oversimplifies what is required. Most laws do not require platforms to verify every user's exact date of birth. They require platforms to take reasonable or appropriate steps to ensure that users below a specified age cannot access the platform or specific features. That is a different and, in some ways, more flexible obligation, but it still requires a technical approach that goes well beyond asking users to tick a box.

| JURISDICTION | KEY INSTRUMENT | AGE THRESHOLD | OBLIGATION ON PLATFORMS | ENFORCEMENT BODY |
|---|---|---|---|---|
| Australia | Online Safety Amendment Act 2024 | 16 | Reasonable steps to prevent account creation by under-16s | eSafety Commissioner |
| France | ARCOM age verification framework 2026 | 15 | Verified age check before access; use of approved verification services | ARCOM |
| United Kingdom | Online Safety Act 2023, Children's Codes | 18 (for harmful content); broader age assurance required | Technically robust age assurance; self-declaration not sufficient | Ofcom |
| European Union | DSA, GDPR Article 8 | 16 (member states may lower to 13) | Verifiable parental consent for under-threshold users; VLOP risk mitigation | Digital Services Coordinators; national DPAs |
| United States | KOSA 2024; COPPA 2.0; state laws | 16 (KOSA); 13 (COPPA) | Safe design requirements; data restrictions; age-gating of harmful features | FTC; state attorneys general |
| California (US) | Age-Appropriate Design Code Act 2024 | 18 | Privacy by default for users likely to be under 18; data minimisation; | California Privacy Protection Agency |

| JURISDICTION | KEY INSTRUMENT | AGE THRESHOLD | OBLIGATION ON PLATFORMS | ENFORCEMENT BODY |
|---|---|---|---|---|
| | | | no profiling of minors | |

A key pattern across all of these frameworks is that they distinguish between two types of obligation. The first is a structural design obligation: platforms should not, by default, collect more data from younger users, serve them algorithmically amplified content without restriction, or make features designed to maximise engagement accessible without age-appropriate safeguards. The second is an access control obligation: certain content or platform features should be inaccessible to users below a specified threshold, and the platform must implement a mechanism to enforce that threshold.

It is the second obligation that drives the age verification and age assurance challenge. A structural design obligation can often be addressed through default settings and data governance policies. An access control obligation requires an actual check at the point of account creation or access, and that check must be sufficiently robust that regulators consider it meaningful compliance, not a formality.

SECTION 3

## Why Verifying Age Is Harder Than It Sounds

Most people's instinct about age verification is that it should be simple: ask people how old they are, or ask them to prove it. The practical difficulties are significant, and they explain why no major social media platform has yet solved this cleanly.

### The scale problem

Social media platforms process millions of new account creations per day. Any age verification step that adds meaningful friction to the sign-up process at this scale results in substantial user drop-off. Industry estimates suggest that adding a mandatory document upload step to account creation reduces

completion rates by between 20 and 40 percent, depending on the platform and the user demographic. For a platform that competes for users, this friction is commercially significant, which has historically created strong internal resistance to robust age verification even where it was legally required.

## The device and context problem

A user creating a social media account may be doing so on a shared family device, a school device, a borrowed phone, or a device purchased by a parent for a child. The device's owner is not necessarily the account creator. Document-based verification can only verify the document, not establish that the person submitting the document is actually the one creating the account. A determined under-age user can use a parent's or older sibling's identity document to pass a document check, particularly if the check does not include a biometric liveness test matched to the document.

## The liveness and spoofing problem

Even where platforms implement facial age estimation or facial match to a document, these techniques are subject to spoofing with varying degrees of difficulty. Holding a photograph of an adult in front of a camera defeats many basic liveness detection systems. More sophisticated ISO 30107-3 compliant liveness detection adds cost and latency to the verification flow. The arms race between age verification technology and methods to circumvent it is ongoing and unlikely to be resolved by any single technical approach.

## The global identity infrastructure gap

Reliable age verification using government-issued identity documents presupposes that the user has access to such a document. In many countries, a significant proportion of the population under 18 does not have a valid passport, driving licence, or national identity card. A verification approach that relies entirely on government identity documents will systematically exclude younger users in populations with low formal identity document coverage, which may correlate with demographic groups that already face barriers to digital inclusion.

### The data retention problem

If a platform collects and stores the identity documents or biometric data used for age verification, it becomes a custodian of extremely sensitive personal data for its entire user base. The breach risk this creates is substantial. Age verification data stores are high-value targets precisely because they link verified real identities to the platform accounts of potentially millions of users. This is the central privacy challenge of conventional age verification: the cure creates a new vulnerability.

SECTION 4

## Current Approaches and Their Limitations

Platforms have deployed a range of age verification and age assurance approaches, each with distinct strengths and weaknesses. Understanding these helps clarify why more sophisticated approaches are gaining traction.

### Self-declaration

The simplest approach: users type in their date of birth, and the platform accepts whatever they type. Self-declaration is inexpensive and creates no friction, but it provides no meaningful assurance. A child who knows they need to claim to be over 16 to access a platform will do so. Regulators in Australia, the UK, and France have explicitly stated that self-declaration does not satisfy their age assurance requirements. Its continued use as the primary mechanism on most major platforms is a compliance gap that enforcement action is expected to close.

### Parental consent flows

Some platforms have implemented parental consent mechanisms for younger users, requiring a parent or guardian to verify their own age and provide consent before a child's account is activated. These approaches are required under COPPA in the United States for under-13s. The challenge is that the verification of the parent's age faces the same difficulties as the verification

of the child's age: if the parent's age check is a self-declaration, a child can complete it themselves. Robust parental consent flows require the parent's identity to be verified, which adds cost and complexity.

### Document-based verification

Requiring users to upload a government-issued identity document is reliable where the document is genuine and the person submitting it is the document holder. It creates the data retention problem described above, has high friction and drop-off, excludes users without qualifying documents, and cannot prevent a determined under-age user from submitting an older person's document. Document-based verification is appropriate for high-stakes access scenarios (opening a bank account, accessing gambling services) but is widely considered disproportionate for general social media access.

### Facial age estimation

AI-powered facial age estimation analyses a selfie or live video feed and estimates the user's age range. It is friction-light and document-free, but it has accuracy limitations at the margins: a 15-year-old who looks 17, or a 17-year-old who looks 14, may be misclassified. Accuracy is also affected by lighting, camera quality, and the composition of the training data underlying the model. Facial age estimation is most effective as a supplementary signal or as a triage mechanism rather than as a standalone compliance solution. It also raises its own privacy concerns around biometric data processing.

### Credit card and financial proxy checks

Using credit card ownership as a proxy for adult status works for some demographics but excludes adults without credit cards (a significant population in many markets) and can be bypassed by children with access to a parent's payment details. Financial proxy checks also raise GDPR concerns around the appropriateness of the data processing.

### Mobile network operator (MNO) age signals

Mobile network operators hold age data for their contract customers and can provide age confirmation signals without sharing the underlying date of birth. MNO-based age verification is privacy-friendly and covers a large proportion of adult users in markets with high mobile penetration. It does not cover users on prepaid SIMs, Wi-Fi-only devices, or those who create accounts on devices other than their primary mobile. It is a useful component of a layered age assurance approach rather than a complete solution on its own.

SECTION 5

## The Privacy Paradox of Age Verification

There is a genuine tension at the heart of social media age verification that no amount of good intentions can dissolve through conventional technical means. The protection of children's safety and privacy requires platforms to implement age checks. But conventional age checks require collecting sensitive personal data, which creates new privacy and security risks, including risks that may fall disproportionately on the children the checks are intended to protect.

A child whose parent provides a copy of their passport to a social media platform's age verification system is in a situation where their most sensitive identity document is now stored by that platform, or by a third-party verification provider the platform has contracted. If that data is breached or sold, the consequences for the child include exposure of their real name and identity alongside their social media activity, which is precisely the kind of linkage that privacy-conscious parents and child safety advocates are most concerned about.

The Open Rights Group, the Electronic Frontier Foundation, and similar civil liberties organisations have pointed to this paradox in their responses to age verification consultations in the UK, the US, and elsewhere. Their argument is not that children should not be protected online. It is that an age verification system that requires children and their parents to hand over more sensitive data to more parties creates new vulnerabilities that offset or exceed the benefits of the age check itself.

Regulators have begun to acknowledge this tension. Ofcom's guidance on age assurance explicitly states that platforms should prefer approaches that minimise the personal data collected in the course of the age check. The ICO's Children's Code (the Age Appropriate Design Code) embeds data minimisation as a core principle. GDPR Article 25 requires privacy by design and default, which means that the default age verification approach should be the one that collects the least data consistent with providing sufficient assurance.

> **The core tension:** Robust age verification built on conventional document checks creates a centralised store of sensitive identity data linked to children's platform activity. Privacy-preserving alternatives exist that provide equivalent assurance without creating this store. The regulatory frameworks actively encourage their use.

SECTION 6

## Age Assurance vs Age Verification

The distinction between age verification and age assurance matters because it opens up a broader range of technically valid compliance approaches. Understanding the difference is important for platforms building their compliance architecture.

**Age verification**, in its strict sense, means confirming a specific date of birth using an authoritative identity document or record. It produces a high-confidence result but requires processing a significant amount of personal data, specifically at minimum the date of birth and some form of identity credential, in order to produce that result.

**Age assurance** is the broader category. It covers any method that provides sufficient confidence that a user meets an age threshold for a particular purpose. Age assurance does not necessarily require establishing an exact date of birth. It can be satisfied by a range of approaches that together provide proportionate confidence, calibrated to the risk level of the content or features being protected.

Ofcom's approach in the UK explicitly adopts the age assurance framing. Its guidance describes a spectrum of age assurance methods ranging from self-declaration (low reliability, low risk to use as a component of a layered approach) through to hard identity checks (high reliability, high data processing burden) and states that the appropriate method or combination of methods depends on the risk level of the services and content being protected.

This risk-calibrated approach means that a platform protecting content that poses moderate harm to younger users may be able to use a combination of age estimation signals and MNO data confirmation, while a platform protecting content that poses serious harm may be required to implement a more robust check including document verification or cryptographic age proof.

The advantage of framing compliance as age assurance rather than age verification is that it creates space for privacy-preserving approaches that provide high-confidence eligibility signals without requiring the collection and storage of dates of birth or identity documents. Zero-knowledge age proofs fall squarely within the age assurance category and can, depending on the trust architecture behind them, satisfy even the higher-confidence requirements.

SECTION 7

## How Privacy-Preserving Age Assurance Works

The technical foundation of privacy-preserving age assurance is the separation of two questions that conventional age verification conflates: "Is this person the age they claim to be?" and "How old are they specifically?" For access control purposes, the second question is almost never necessary. What a platform needs to know is whether a user is above a threshold, not their exact date of birth.

The approach works in three steps.

### Step one: age verification with a trusted issuer

The user's age is verified once by an entity that the user already has a trust relationship with and that is authorised to process identity data for this purpose. This might be a bank (which already holds verified KYC data including date of birth), a government identity service (such as a national digital identity scheme), a mobile network operator (which holds age data for contract customers), or a specialist age verification provider (AVP) operating under the relevant national framework.

The issuer verifies the user's date of birth against the threshold relevant to the service and records the outcome: this user is, or is not, above the threshold. The issuer may also record whether the verification was performed to a sufficient level of rigour for the intended use case. The issuer does not share the date of birth with the platform.

### Step two: converting the outcome to a cryptographic proof

The binary outcome from the issuer is converted into a zero-knowledge eligibility proof. The proof demonstrates that the user's age was verified by a trusted issuer and that the verification outcome satisfies the relevant threshold. The proof contains no information about the user's actual age, date of birth, or identity. It is a mathematical attestation of the eligibility outcome, nothing more.

The proof may be time-limited, so that it cannot be used indefinitely without refreshing the underlying verification. It may be bound to a specific platform context, so that it cannot be replayed across platforms without the user's explicit involvement in the presentation. And it may include a freshness timestamp that lets the platform confirm the verification was performed recently, without learning anything else about the user.

### Step three: platform verification

The platform receives the proof and verifies it using a public verification key. The verification confirms that the proof is valid, was issued by a trusted issuer, and attests to the relevant threshold outcome. The platform does not receive the user's date of birth. The platform does not need to store the proof itself

beyond the point of access control confirmation. The data minimisation obligation is satisfied structurally: there is no date of birth in the platform's system to protect, breach, or misuse.

This approach can be deployed as part of a layered age assurance architecture, where lower-risk access decisions use lighter-weight signals (age estimation, MNO data) and higher-risk access decisions require a cryptographic age proof from a trusted issuer. The layering allows platforms to calibrate friction to risk, reducing drop-off at the lower end of the risk spectrum while maintaining robust protection for the highest-risk content.

SECTION 8

## Reusable Age Credentials and the One-Check Vision

One of the most significant emerging concepts in age assurance is the reusable age credential, sometimes described as a verified age token. The core idea is straightforward: a user should be able to verify their age once, with a trusted issuer, and carry a credential that they can present to multiple platforms without repeating the underlying identity check each time.

This matters for several reasons. First, it reduces user friction significantly. Under a one-check model, the burdensome step of identity document submission or biometric capture happens once, at a trusted issuer, and the resulting credential can be used across dozens of platforms. The friction of subsequent age checks becomes minimal: the user presents their credential and access is granted.

Second, it concentrates the sensitive data processing at the point of highest trust and lowest risk. The organisation that processes identity documents for age verification is the one that has the regulatory authorisation to do so, the technical infrastructure to protect the data, and the legal accountability for its use. That organisation is not the social media platform, which may have no background in identity document handling and whose core systems are not optimised for this kind of sensitive data custody.

Third, it aligns well with the EU Digital Identity Wallet framework (eIDAS 2.0), which is establishing the infrastructure for portable, reusable identity credentials across the EU. The EU Digital Identity Wallet will support selective disclosure of identity attributes, including age-over proofs, allowing users to present just the relevant eligibility confirmation without exposing their full identity record. Age assurance is one of the primary intended use cases for this infrastructure.

The challenge for reusable age credentials is interoperability. For a credential issued by a bank in Australia to be accepted by a platform operating in the EU, both parties need to recognise the issuer as trustworthy and the credential format as valid. This is an active area of standards development, with GSMA, the W3C Verifiable Credentials specification, and national digital identity frameworks all working towards compatible approaches. Progress is real but the interoperability layer is not yet mature in all markets.

In markets where national digital identity infrastructure is more developed (the Nordic countries, Estonia, and increasingly Germany and France with their eID schemes), reusable age credentials built on government identity infrastructure are closer to deployable today. In markets with less centralised identity infrastructure (the US, much of Asia-Pacific), the credential issuer role is more likely to be played by banks, MNOs, or established AVPs.

SECTION 9

## Where AffixIO Fits in the Age Assurance Stack

AffixIO's role in age assurance is in the verification and attestation layer, sitting between age verification providers and the platforms that need to act on an age eligibility outcome. We are not an age verification provider. We do not process identity documents, capture biometrics, or make the primary determination of a user's age. That work belongs with entities that have the appropriate regulatory authorisation and established trust relationships with users.

What AffixIO provides is the cryptographic proof layer that converts an age eligibility determination into a tamper-resistant, privacy-preserving attestation that a platform can verify, act on, and retain as a compliance record.

## The eligibility proof

When an age verification provider confirms that a user meets an age threshold, AffixIO's API converts that confirmation into a zero-knowledge eligibility proof. The proof attests that a trusted issuer confirmed the threshold outcome. The proof does not contain the user's date of birth, identity document details, or any other personal data. It is the compliance record for the age check: evidence that the check was done, by whom, when, and with what outcome, without containing the data that was checked.

## The audit trail

Every age eligibility proof is anchored in AffixIO's Merkle-based audit infrastructure, signed with post-quantum ML-DSA-65 signatures. A platform that receives an age eligibility proof has a verifiable record it can present to a regulator demonstrating that age assurance was performed for a given user session. The eSafety Commissioner, Ofcom, or any other enforcement body examining the platform's compliance can verify the existence and integrity of the age check records without the platform needing to expose user personal data to do so.

## Reusable credentials

AffixIO's approach supports the reusable credential model. A user who has been age-verified by a trusted issuer and whose eligibility proof has been issued can present that proof to multiple platforms that use AffixIO's verification endpoint. The user does not need to repeat the identity check for each platform. The platform does not need to independently contact the issuer. The proof is self-contained and verifiable, with a validity period that reflects the currency of the underlying verification.

## Integration

AffixIO's age assurance layer integrates with existing age verification providers via a standard adapter interface. Platforms that have already contracted with an AVP for identity-based age verification can add the AffixIO proof layer without replacing their existing verification flow. The AVP continues to handle the document and biometric processing; AffixIO converts the outcome into the cryptographic proof that the platform stores and presents to regulators. Platforms building age assurance from scratch can integrate AffixIO's API as the proof and audit layer from the outset, connecting it to whichever AVP or digital identity infrastructure they choose.

SECTION 10

## GDPR, the Children's Code, and Data Minimisation

Age verification done badly creates GDPR problems as well as solving a compliance problem. Platforms that collect dates of birth, document images, or biometric data during the age verification process are processing special categories of personal data (biometrics) and data about children, both of which attract heightened obligations under GDPR. Getting age verification wrong from a data protection perspective can create a GDPR liability that is, in financial terms, considerably larger than the underlying Online Safety Act or DSA fine for non-compliant age assurance.

The ICO has been explicit about this. Its 2023 guidance on age verification noted that platforms must consider the data protection implications of any age verification approach before implementation and must be able to demonstrate that the approach chosen is proportionate, minimises data collection, and is consistent with GDPR's purpose limitation and data minimisation principles.

GDPR Article 25, privacy by design and default, requires that the default approach to age verification be the one that processes the least personal data consistent with providing sufficient assurance. This is a direct endorsement of privacy-preserving age assurance: where a cryptographic

eligibility proof can provide the same assurance as a date of birth without requiring the platform to store the date of birth, the GDPR default is the proof-based approach.

The UK's Age Appropriate Design Code (the Children's Code) extends this further. It requires platforms that are likely to be accessed by children to apply high privacy standards by default, collect only data necessary for the service, and not use personal data in ways that are detrimental to children's wellbeing. A platform that uses age verification data for advertising targeting or profiling violates the Children's Code even if the age verification itself was conducted properly.

The Children's Code applies to any online service that is likely to be accessed by under-18s, regardless of whether the service is primarily intended for children. A general social media platform that 17-year-olds can access is in scope, and its privacy defaults must reflect that. This creates a structural incentive to minimise data collection at the age verification stage: the less data collected, the smaller the surface area for Children's Code compliance issues.

Privacy-preserving age assurance, where the platform receives an eligibility proof rather than a date of birth, naturally satisfies the Children's Code's data minimisation requirement. The platform has no date of birth to use for profiling because the platform never received one. The compliance record is the proof, and the proof contains no personal data. This is data minimisation by architecture, not by policy, which is precisely what both GDPR Article 25 and the Children's Code are designed to encourage.

SECTION 11

## What Does Not Work

It is worth being direct about the approaches that are likely to fail either technically or legally, because many platforms are under pressure to implement something quickly and may be tempted by approaches that offer low friction without providing genuine compliance.

**Self-declaration alone does not work.** Regulators in Australia, the UK, and France have all stated explicitly that a date of birth tick-box does not satisfy their age assurance requirements. Platforms that continue to rely solely on self-declaration after enforcement deadlines risk significant fines and, in Australia's case, injunctions requiring the platform to implement genuine age verification within a specified timeframe. The risk is not hypothetical: enforcement actions have begun.

**Device-based age signals without corroboration do not work for high-risk content.** Inferring a user's age from their device settings, app store age declaration, or browsing history provides a probabilistic signal rather than a verified outcome. These signals may be appropriate as a component of a layered, lower-risk age assurance approach, but they cannot stand alone for content that poses serious harm to younger users. They are also susceptible to manipulation by users who want to access age-restricted content.

**Age verification that creates a centralised identity document store without robust security does not work legally.** A platform that implements document-based age verification and stores the collected documents in a system that is subsequently breached will face simultaneous regulatory action under the Online Safety Act (for a deficient age assurance implementation) and under GDPR (for inadequate security of special category personal data). The combination of these liabilities is likely to be more severe than the original age verification non-compliance would have been.

**Blocking entire jurisdictions does not work commercially or reputationally.** Some platforms have considered geo-blocking users from jurisdictions with strict age verification requirements rather than implementing compliant systems. This approach has proven commercially damaging and has not prevented regulatory action: platforms can still be fined for failing to protect users in those jurisdictions where they previously operated.

**Age estimation alone does not work for compliance.** Facial age estimation is a useful supplementary signal but has accuracy limitations that regulators have noted. A system that relies solely on facial age estimation to distinguish 15-year-olds from 17-year-olds will produce a meaningful error rate in both

directions. It is most useful as a triage mechanism, routing edge cases to a more robust verification step, rather than as a standalone compliance solution.

SECTION 12

## Conclusion

The global wave of social media age restriction legislation is not a trend that will reverse. More than two dozen countries have moved from debate to legislation in under two years. The direction of travel is towards more stringent requirements, not fewer, as regulators respond to continued evidence of harm to younger users and to the evident inadequacy of self-regulatory approaches by the platforms themselves.

For platforms, this creates a genuine compliance obligation that requires investment in age assurance infrastructure. The question is not whether to implement age assurance but how to implement it in a way that provides genuine protection, satisfies regulators, respects users' privacy, and does not create a new data liability in the process of solving the existing compliance gap.

The answer that best satisfies all of these constraints is privacy-preserving age assurance: a layered approach that uses trusted issuers to make the primary age determination, converts the outcome to a cryptographic eligibility proof, and allows platforms to verify that proof without receiving or storing users' personal data. This approach satisfies the regulatory requirement for technically robust age assurance, aligns with GDPR's data minimisation and privacy by design principles, and eliminates the centralised identity document store that makes conventional age verification a security liability.

The infrastructure to support this approach is available today, is evolving rapidly through standards bodies and national digital identity programmes, and is increasingly aligned across jurisdictions as the EU Digital Identity Wallet, GSMA's mobile identity standards, and national AVP frameworks

converge. Platforms building their age assurance architecture now have the opportunity to do so in a way that is both compliant today and compatible with where the global standards are heading.

AffixIO's role in that architecture is the proof and audit layer: converting trusted age eligibility determinations into cryptographic attestations that platforms can act on, retain as compliance records, and present to regulators without exposing any user personal data in the process. The underlying age determination belongs with the entities best placed to make it. The audit trail, the tamper-resistance, and the privacy-preserving attestation layer are what we contribute.

## Related reading

- WP-009: Privacy-Preserving Age Verification: Zero-Knowledge Proof of Age Threshold for the Online Safety Act and DSA
- WP-006: A PII-Free KYC Schema by Design: Structural Data Minimisation via Zero-Knowledge Identity Circuits
- WP-017: ZK Selective Disclosure for eIDAS 2.0 and the EUDI Wallet
- WP-008: Zero-Knowledge Proofs as GDPR Article 25 Infrastructure

## Frequently asked questions

### What is the difference between age verification and age assurance?

Age verification typically means confirming a specific date of birth using an identity document. Age assurance is the broader category, covering any method that provides sufficient confidence that a user meets an age threshold. Regulators such as Ofcom use the age assurance framing deliberately to allow for privacy-preserving approaches that do not require capturing an exact date of birth.

### Does Australia's social media ban require age verification?

Australia's Online Safety Amendment (Social Media Minimum Age) Act 2024 requires platforms to take reasonable steps to prevent under-16s from creating accounts. It does not mandate a specific verification method. The eSafety Commissioner is developing technical standards that will define what counts as reasonable steps, and the direction is towards technically robust age assurance rather than mere self-declaration.

### How can platforms verify age without storing date of birth?

Zero-knowledge proof age assurance allows a trusted third party to confirm that a user meets an age threshold and issue a cryptographic proof of that fact. The platform receives the proof, not the date of birth. The proof confirms the eligibility outcome without containing any underlying personal data, satisfying both the age assurance requirement and GDPR's data minimisation principle.

### What is a reusable age credential?

A reusable age credential, sometimes called a verified age token, allows a user to verify their age once with a trusted issuer and carry a credential they can present to multiple platforms without repeating the identity check each time. This reduces friction, concentrates the sensitive data processing at the most trusted point in the chain, and is the direction that national digital identity frameworks including the EU Digital Identity Wallet are moving towards.