

UNPACKING AGE ASSURANCE: TECHNOLOGIES AND TRADEOFFS

Age Assurance is the umbrella term for existing and emerging methods used to establish an individual's age or age range. In a rapidly evolving regulatory landscape, there is no 'one-size-fits-all' solution; instead, policymakers and organizations should focus on proportionate, risk-based solutions tailored to specific use cases. By adopting a layered approach and leveraging emerging standards like ISO/IEC 27566, platforms can ensure that age checks prioritize functionality, performance, privacy, security, and acceptability.

AGE ASSURANCE QUESTIONS

WHAT ARE THE GOALS?

- Place an individual within an age threshold or age band. (e.g. 13-15)
- Limit access to an age-restricted service, provide age-appropriate experience or facilitate parental consent.

WHAT IS THE APPROPRIATE ASSURANCE METHOD?

Choose a method or methods that provides a level of age assurance (accuracy) proportional to the goals and risks of the service, keeping in mind that legal obligations may dictate a specific method.



DOES THE AGE ASSURANCE METHOD....

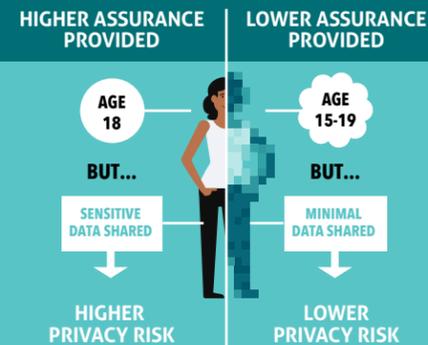
- Provide a result that relates to the individual.
- Maintain low errors over different demographics.
- Facilitate access based on true positive (TP) and prohibit for true negative (TN) results.

WHO IS RESPONSIBLE FOR AGE ASSURANCE?

Answering this question requires policy and technical choices in a complex and evolving environment where age assurance functions may be split across the technology stack.

IS ASSURANCE BALANCED WITH PRIVACY RISKS?

After considering privacy risks and mitigations, confirm that the assurance goal warrants the level of privacy risks and other impacts associated with the chosen age assurance method.



DECLARATION

A user self-asserts their birthdate without providing supporting evidence.

AGE-GATING

This common method is most appropriate in low risk situations, as children and teens can bypass by providing a false birthdate. Privacy risk is low, especially if birthdates are not retained or matched with a name or other indirect identifier.

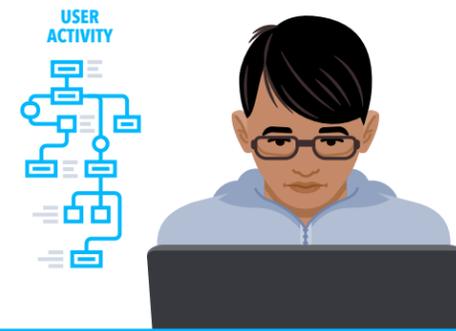


INFERENCE

Drawing reasonable conclusions about an age range based on user activity, financial transactions, email address, existing accounts, etc..

BEHAVIORAL / ACCOUNT

Email linked to workplace apps, mortgage lender, 401k provider. Login patterns show weekday activity during business hours. Usage profile—professional services, financial products—infer adult user.



ESTIMATION

Using AI and machine learning to deduce likely age based on biological or behavioral traits e.g. face, voice, iris, typing patterns.

TRAIT CHARACTERIZATION

Estimates age using a facial image, but the individual is not uniquely identified. Best used to place users in age bands, or signal that a user meets an age threshold, such as under 13 or 21+. Estimation is less effective for discerning age in a narrow range like 17 vs 18.

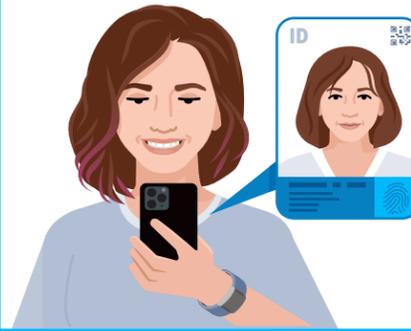


VERIFICATION

Determining age by referencing an authoritative, verified date of birth (DOB) from documents or databases.

ID + BIOMETRIC

Matches a scan of a government-issued ID and live photo or video using facial recognition. This method is more appropriate for regulated or higher risk age restricted services.



PARENTAL VOUCHING: A parent establishes an account with an above method and confirms the child or teen's age.

EMERGING AGE ASSURANCE CONCEPTS: Age Assurance methods are rapidly evolving — below are some of the key emerging technologies.

AGE SIGNALS AND AGE TOKENS

Age signals and tokens work together to replace the sharing of raw identity documents with verified attributes. An age signal is the real-time communication of a user's age status (e.g., "Over 18"), while an age token is the secure, reusable digital credential that stores this status on a user's device or browser. Functioning like a "digital ticket," this combination allows platforms to instantly confirm age without accessing sensitive source data.

HOW USER-BINDING WORKS

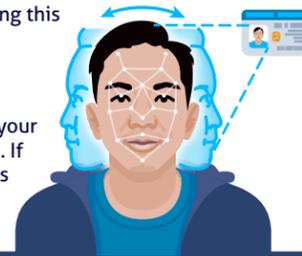
User-binding ensures the age assurance result is accurately linked to the correct individual. This mitigates the risk of shared devices and can be done using biometrics, multi-factor authentication (MFA), or passkeys.

1 Initial Verification: You perform a one-time age assurance (e.g., scanning a government ID or using AI-based facial age estimation).

2 The "Binding" Step: The system captures a biometric template (such as a facial map). Rather than storing an image of your face, it utilizes a "fuzzy extractor"—a privacy-preserving method that converts unique biometric features into a stable, irreversible cryptographic key.

3 Encryption: A digital "Age Assurance Certificate" (e.g., "Verified 18+") is encrypted and locked using this unique biometric key.

4 Re-Authentication: Whenever you access the restricted service, the system performs a brief "liveness check." This regenerates the key from your biometrics in real-time to decrypt the certificate. If an unauthorized user—such as a child—attempts to use the device, their features will not match the key, leaving the certificate inaccessible.



ONE-TIME VS. REUSABLE AGE CREDENTIALS

One-time checks require repetitive identity uploads, increasing friction and data exposure. Reusable credentials store a single verification as a secure token, enabling instant age confirmation across platforms without re-submitting sensitive data. However, interoperability is currently limited to specific trust frameworks as universal standards for cross-platform recognition are still maturing.

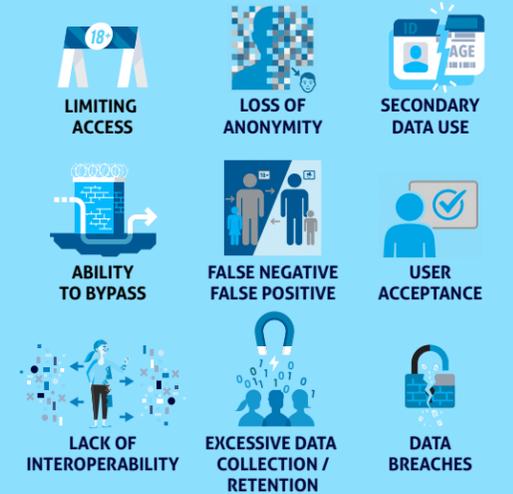
ZERO KNOWLEDGE PROOF (ZKP)

A Zero-Knowledge Proof (ZKP) is a cryptographic method that confirms a user meets a specific requirement (e.g., "18+") without exposing any underlying personal data. The verifier learns only whether the claim is true—not the user's birthdate, identity, or other details.

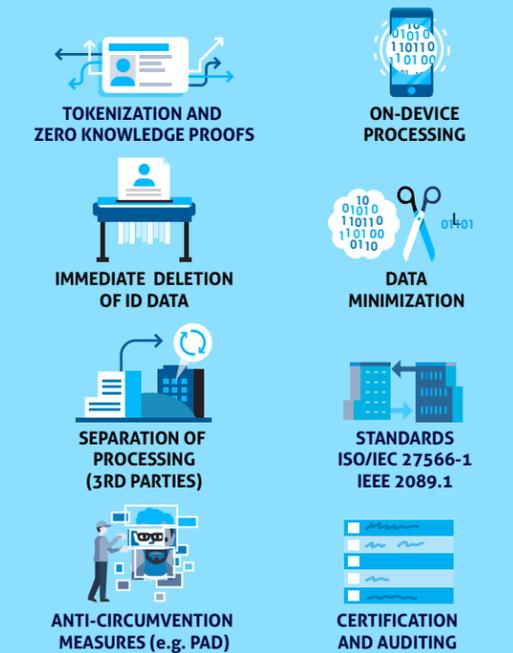
DOUBLE-BLIND ARCHITECTURE

Separates identity from activity: an external service verifies your age without knowing which site you visit, while the website confirms you are of age without ever learning your identity. This prevents linking user data to online behavior.

RISKS AND CHALLENGES OF AGE ASSURANCE



RISK MANAGEMENT TOOLS



EXAMPLE USE CASE: AGE ASSURANCE FOR ONLINE GAMING

In this scenario, Miles, a 16 year old, is accessing an online gaming service that is designed for teens and older. It has optional age-restricted features.

INITIAL EXPERIENCE DECLARATION

Miles starts by providing a self-asserted birthdate (Age Declaration). This low-assurance method allows entry to the main game.



SECONDARY FEATURE ESTIMATION

When Miles attempts to enable 16+ features, the system performs Age Estimation via a "live selfie." By applying a 3-year buffer to the 16-year-old threshold, the system establishes a grey zone range of 13–19. Because Miles' estimate falls within this buffer, he is moved to the next step for stronger verification.



GREATER ASSURANCE NEEDED: VERIFICATION, INFERENCE AND VOUCHING

Miles does not have a driver's license, instead of blocking access, the system offers *Inference* or *Parental Vouching* as fallbacks. **Inference:** Miles connects a bank account or school record that confirms he is in the 16+ age bracket. **Parental Vouching:** Miles' parent (who has been verified as an Adult) confirms that Miles is 16.



AGE SIGNAL / TOKEN

Once verified via one of these fallbacks, the system generates an Age Signal (e.g., "Verified 16+") and all data and metadata used in the verification process is deleted. The "signal" is stored in the browser as an "age token" (cookie).



BINDING

The age credential is cryptographically bound to Miles' device passkey. This ensures that if Miles shares his phone with James, a 15-year-old friend, James cannot access 16+ features. The age signal is only released when a PIN, pattern, or local biometric is successfully entered.

