



European Policy  
Innovation Council



# Age Verification in the Digital Single Market

December 2025



# European Policy Innovation Council

---

This publication was commissioned and funded by the Group of the Progressive Alliance of Socialists & Democrats in the European Parliament (S&D Group) and prepared in the framework of the roundtable discussion on *Age Verification in the Digital Single Market*, hosted at the European Parliament.

The present version is republished by the European Policy Innovation Council (EPIC) strictly for non-commercial, analytical, and documentation purposes. The content has not been altered from the original version.

The S&D Group retains full institutional credit for commissioning and funding the paper, as well as for hosting the related parliamentary discussion. The opinions expressed are those of the author(s) and do not necessarily represent the official position of the S&D Group or the European Parliament.

# Table of Contents

<b>Executive Summary</b>	4
<b>1. Introduction — Setting the Scene</b>	8
<b>2. Overview of Available Age Verification Technologies</b>	9
2.1 Self-Declaration Systems	10
2.2 Document-Based Verification (ID Scans, Passports, National eID)	10
2.3 Device-Based Age Signals	11
2.4 Biometric Age Estimation (Face, Voice, Behavioural Patterns)	12
2.5 Third-Party AV Providers & Privacy-Preserving Tokens	13
2.6 Behavioural Inference (AI-Based Risk Scoring)	14
2.7 On-Device, Real-Time Harm-Prevention Systems	14
<b>3. Comparative Assessment of Age Verification Technologies</b>	15
3.1 Comparative Assessment Matrix	15
3.2 Strengths and Limitations of Each Approach	16
3.3 Key Takeaways	19
<b>4. Key Issues in the EU Debate</b>	19
4.1 Privacy and Data Minimisation Under the GDPR	19
4.2 Enforceability, Evasion, and Practical Feasibility	21
4.3 Complementary Safety-by-Design Measures	21
4.4 Proportionality and Fundamental Rights	22
4.5 Competition, Market Structure, and Innovation	22
4.6 Fragmentation Across Member States	23
4.7 The Role of eIDAS 2.0 and the EUDI Wallet	23
4.8 Governance, Certification, and Trust	24
4.9 Summary: A System Not Yet Ready for One-Size-Fits-All Solutions	26
<b>5. Policy Options for the EU</b>	27
<b>Annex — International Landscape Snapshot</b>	33

## Executive Summary

Children today navigate a digital environment shaped by social media, video-sharing platforms, gaming, generative AI tools, and encrypted communications. The opportunities for learning and creativity are vast, but so are the risks: exposure to harmful content, grooming, sexual extortion, and addictive design patterns continues to affect minors across Europe. As a result, age verification (AV) is emerging as a central issue in the wider debate on online safety and platform accountability.

At the same time, the European regulatory context is evolving. The Digital Services Act (DSA) already requires platforms to assess and mitigate risks to minors, and the European Commission has begun active enforcement through investigations and Article 28 guidelines. Yet several Member States are introducing national AV laws, and policymakers across the European Parliament are calling for additional measures. This raises concerns about fragmentation, proportionality, technological feasibility, and the risk of premature regulatory lock-in.

The EU should reject OS/app-store-centric age-assurance mandates and any model that turns device vendors into de facto identity authorities. Responsibility for age checks must remain with the service offering the regulated activity, with interoperable, standards-based acceptance paths that do not depend on a single commercial ecosystem.

This paper provides a technology-neutral, evidence-based assessment of the current age-verification landscape and outlines policy options for a harmonised and future-proof EU approach. It does not prescribe a single AV method. Instead, it identifies the conditions needed for multiple technologies to be tested, compared, and allowed to compete in ways that protect children, preserve privacy, and support innovation.

## Key Findings at a Glance

The figure below summarises the major issues shaping the EU debate, the technological realities of AV, and the principal policy levers available to legislators and regulators.

Figure 1. Overview of the EU Age Verification Policy Landscape

Dimension	Core Elements	Insights & Constraints	Implications for Policy Design
<b>1. Key Issues in the EU Debate</b>	Privacy and GDPR minimisation; enforceability and evasion; proportionality; competition and gatekeeping; fragmentation; EUDI Wallet readiness; governance gaps.	<ul style="list-style-type: none"> <li>• AV risks unnecessary data collection or de facto ID checks.</li> <li>• Usability and bypassing undermine enforcement.</li> <li>• OS-level centralisation creates competition and sovereignty risks.</li> <li>• Divergent national rules threaten the Single Market.</li> </ul>	<ul style="list-style-type: none"> <li>• Prioritise least-intrusive, privacy-preserving methods.</li> <li>• Keep responsibility/liability with high-risk services; explicitly prohibit OS/app-store gatekeeper mandates.</li> <li>• Promote interoperable EU baselines.</li> </ul>
<b>2. Technology Landscape &amp; Comparative Assessment</b>	Self-declaration; document-based AV; device signals; biometric estimation; third-party tokenisation; behavioural inference; on-device harm-prevention tools.	<ul style="list-style-type: none"> <li>• No single technology meets accuracy, privacy, and proportionality goals.</li> <li>• Hybrid approaches are most realistic.</li> <li>• Biometric and behavioural methods require strong safeguards.</li> <li>• Device-based signals reduce friction but risk gatekeeping.</li> <li>• Safety-by-design tools complement but do not replace AV.</li> </ul>	<ul style="list-style-type: none"> <li>• Maintain technology neutrality; allow multiple AV pathways.</li> <li>• Require transparency, bias testing, and security safeguards.</li> <li>• Avoid embedding a single technical method in law.</li> <li>• Recognise complementary safety tools as part of the ecosystem.</li> </ul>
<b>3. Policy Options for the EU</b>	<ol style="list-style-type: none"> <li>(1) Minimum EU requirements;</li> <li>(2) Privacy-preserving age tokens;</li> <li>(3) AV sandboxes;</li> <li>(4) EUDI Wallet integration;</li> <li>(5) Risk-based tiering;</li> <li>(6) Conditions on biometric AV.</li> </ol>	<ul style="list-style-type: none"> <li>• EU coordination can reduce fragmentation and build trust.</li> <li>• The EUDI Wallet is promising but not ready as a universal solution.</li> <li>• Sandboxes improve evidence and comparability.</li> <li>• Risk-tiering aligns regulatory burden with service risk and supports SMEs.</li> </ul>	<ul style="list-style-type: none"> <li>• Build a flexible, tech-neutral framework.</li> <li>• Focus on governance and open interoperability (OpenID4VP/SD-JWT/ISO 18013-5) – no exclusive or proprietary AV pathways.</li> <li>• Keep Wallet integration optional.</li> <li>• Apply the strongest requirements to the highest-risk services.</li> </ul>

## Conclusions vs Policy Direction

### 1. Technology neutrality must guide EU policy.

No AV technology is mature or universal enough to justify exclusive reliance. The EU should avoid mandating specific tools and instead shape an environment where multiple solutions can evolve safely and competitively.

### 2. AV should be proportionate and risk-based.

Stricter verification makes sense only for high-risk services such as adult content or gambling. Imposing heavy AV requirements on low-risk or general-purpose services risks exclusion, friction, and chilling effects without improving safety.

### 3. Privacy-preserving architectures offer the most promising path forward.

Selective disclosure, age tokens, on-device processing, and zero-knowledge proofs can reduce data exposure while enabling compliance. By contrast, OS-level centralisation or identity-heavy models pose significant privacy, competition, and governance risks.

### 4. Governance and evaluation capacity are prerequisites for effective regulation.

Before imposing new obligations, the EU should establish certification frameworks, performance benchmarks, bias-testing regimes, and regulatory sandboxes. This enables policymakers to understand what works, where, and under what conditions.

The EU should focus on **enabling conditions** — and explicitly rule out gatekeeper-centric AV architectures that centralise age, identity, or parental relationships at OS/app-store level. This includes:

- Minimum EU-wide requirements for transparency, security, and data minimisation;
- Privacy-preserving verification pathways, including tokenisation and selective disclosure;
- Optional integration with the EUDI Wallet as it matures;
- Risk-based tiering, focusing the strongest requirements on the highest-risk services;
- And regulatory sandboxes to generate evidence before legislating.

- Prohibit tying AV to proprietary OS/app-store identity layers; require non-exclusivity and open, standards-based interfaces for any age-attribute signals.
- Keep liability with high-risk services, not with upstream device or platform intermediaries that do not control the regulated activity.

A flexible, technology-neutral framework will allow innovation to continue, protect children more effectively, and uphold fundamental rights across the Digital Single Market.

# 1. Introduction — Setting the Scene

## Why Age Verification Matters Now

Children today navigate an online environment profoundly different from that of just a few years ago. Digital platforms — from social media and streaming services to gaming, adult content, and generative AI tools — have become central to communication, learning, and entertainment. At the same time, risks have intensified: exposure to harmful content, grooming, sexual extortion, addictive design, and algorithmic amplification all disproportionately affect minors.

Policymakers across Europe increasingly recognise that **a baseline of age-appropriate digital experiences must be ensured**. Platforms are expected to understand whether a user is a child and to apply appropriate safeguards. As a result, political momentum behind age verification (AV) is accelerating.

With parliamentary momentum behind hard age-verification proposals across several files, including minors' protection, AI systems, and media regulation, the need for a proportionate and evidence-based framework is increasingly urgent.

## Regulatory Drivers in the EU

Several EU legislative and policy developments converge on the question of how platforms can reliably determine age while respecting fundamental rights:

- **Digital Services Act (DSA)** — Requires platforms to assess risks to minors, adapt recommender systems, and apply “appropriate and proportionate” age verification where necessary.
- **CSA Regulation (ongoing negotiations)** — Raises questions about the role of AV in preventing access to harmful content and in broader risk mitigation.
- **AVMSD (Audiovisual Media Services Directive)** — Obligates video-sharing platforms to protect minors from harmful content, prompting national AV requirements.
- **National legislation** — France, Germany, Ireland, Spain, and others have introduced or proposed AV laws, leading to emerging fragmentation in the Digital Single Market.
- **eIDAS 2.0 and the EU Digital Identity Wallet (EUDI Wallet)** — Will eventually allow selective disclosure of age attributes without revealing identity.

## The Core Policy Dilemma

A structural tension lies at the heart of the debate:

**How can the EU protect children online without creating new risks of surveillance, privacy intrusion, discrimination, or exclusion?**

Robust AV can reduce exposure to harm — but if poorly designed, it may:

- require excessive data collection
- normalise identity checks for everyday online activities
- restrict access for marginalised or uncredentialed communities
- expose children and adults to data breaches or misuse of biometric or identity data

## Purpose and Scope of This Paper

This paper offers a **technology-neutral assessment** of the main AV approaches, evaluates their implications for rights, feasibility, and innovation, and maps policy options for a harmonised and future-proof European approach.

**It does not recommend a single AV technology.** Instead, it identifies conditions under which multiple technologies can be safely tested, compared, and allowed to compete, recognising that the ecosystem is not yet mature enough for a single mandated solution.

Finally, policymakers should avoid introducing new legislative mandates before the impact of ongoing DSA and GDPR enforcement is fully assessed, particularly in high-risk environments where investigations are already under way.

## 2. Overview of Available Age Verification Technologies

Age verification technologies span a wide range of methods, each with distinct implications for privacy, accuracy, usability, and proportionality. No single technology is currently capable of serving all use cases across the Digital Single Market. This section maps the main categories and clarifies what each can — and cannot — deliver.

## 2.1 Self-Declaration Systems

**Examples:** “Enter your date of birth,” click-through age gates, parental confirmation pop-ups.

**How it works:** Users manually state their age or confirm they meet the age threshold.

**Strengths:**

- Zero friction; universally deployable
- No additional infrastructure or cost
- Minimal data processing

**Limitations:**

- Trivial to bypass
- No meaningful assurance
- Ineffective in medium- and high-risk contexts

**Conclusion:** Appropriate only for low-risk contexts; not considered a credible AV measure elsewhere.

## 2.2 Document-Based Verification (ID Scans, Passports, National eID)

**How it works:** Users upload or scan a government-issued ID, or authenticate through a national digital identity scheme.

**Strengths:**

- High accuracy and clear age attribute
- Legally recognised documents
- Suitable for high-risk environments where strong assurance is needed

**Limitations:**

- Requires processing highly sensitive data

- Potential exclusion of minors without documents or parental support
- Security risks if documents or identity data are mismanaged
- High friction for users; may deter participation
- Overly burdensome for low-risk services

This method is most appropriate for specific high-risk use cases, but its proportionality is questionable in broad, everyday online environments.

## 2.3 Device-Based Age Signals

**Examples:** OS-level settings, family accounts, parental controls, device profiles.

**How it works:** The device provides an age attribute to the platform.

### Strengths:

- Privacy-preserving; minimal personal data
- Low friction and seamless user experience
- Useful when implemented consistently at system level

### Limitations:

- No guarantee that the device owner is the user
- Heavy dependency on major OS providers (Apple, Google, Microsoft)
- Potential gatekeeping and competition concerns
- Many families do not configure parental/child profiles consistently
- Limited interoperability across EU services

Most effective as *one layer* within a multi-factor AV approach, rather than a standalone solution. Although emerging app-store/OS age-signal APIs can reduce friction, they centralise control in gatekeeper ecosystems and risk cross-context tracking. The EU should not endorse models that make OS vendors the authoritative source of age/parental status. Any OS-level signal must be optional, standards-based, and non-exclusive, with clear bans on persistent identifiers and secondary uses.

## 2.4 Biometric Age Estimation (Face, Voice, Behavioural Patterns)

**Face Estimation:** AI models estimate age from a short selfie or video.

**Voice Estimation:** Models infer age from vocal characteristics.

**Behavioural Biometrics:** Analyses user interaction patterns (keystrokes, mouse movements, reading or scrolling speed).

### Strengths:

- Does not require identification
- Can be fast, low friction, and privacy-preserving if processed on-device
- Suitable for a simple “under/over 18” distinction

### Risks and Limitations:

Accuracy varies across demographics (age, gender, skin tone, disability)

- Requires strict safeguards to avoid storage or misuse of biometric data
- Not yet standardised; quality varies significantly between providers
- Sensitive under GDPR even when not used for identification

Any deployment should:

- prohibit template storage,
- publish accuracy and bias metrics (e.g., confusion matrices),
- enable independent audits.

Biometric AV is promising but remains politically and technically sensitive, requiring cautious and well-governed implementation.

## 2.5 Third-Party AV Providers & Privacy-Preserving Tokens

How it works: **A trusted third party verifies age and returns a privacy-preserving token or “yes/no” attribute to platforms. Platforms never see identity data.**

**Strengths:**

- Strong alignment with GDPR data minimisation
- Clear separation between verification and service access
- High scalability and low friction for users
- Supports interoperability across services
- Reduces the burden on platforms, especially SMEs

**Limitations:**

- Requires certification, audits, and clear liability frameworks
- Risk of market concentration without proper governance
- Token-sharing needs mitigation (e.g., device/session binding, revocation lists)
- Dependent on user trust in intermediaries

This model supports a competitive and privacy-preserving ecosystem without favouring specific technologies. Privacy-preserving age attributes can be conveyed through verifiable-credential presentations using open standards (e.g., OpenID4VP, SD-JWT, or ISO 18013-5-compatible formats), enabling selective disclosure such as “over 16” without revealing identity. Zero-knowledge proofs (ZKPs) may further minimise data by proving an age threshold without sharing the underlying date of birth.

## 2.6 Behavioural Inference (AI-Based Risk Scoring)

**How it works:** Algorithms infer whether a user is likely a minor based on metadata or behavioural patterns (e.g., browsing or interaction habits).

**Strengths:**

- Invisible to users; near-zero friction
- Useful as part of safety-by-design systems
- Helps flag suspicious or underage usage patterns

**Limitations:**

- Low reliability for individual users

- High risk of discrimination or opaque outcomes
- Cannot be used as a primary AV method
- May introduce bias or false positives without transparency

AI age estimation is increasingly used in the EU for broad age categorisation (e.g., under/over 18) and, when subject to strict safeguards, can provide effective support for age assurance without identifying the user.

## 2.7 On-Device, Real-Time Harm-Prevention Systems

A further category of tools is emerging that uses **on-device artificial intelligence** to detect and block sexually explicit or exploitative content in real time. Operating at the operating-system or device level, these systems analyse images, videos, and live camera feeds locally – including within end-to-end encrypted environments – to prevent the viewing, recording, sharing, or live-streaming of harmful material. Because all processing occurs on the device, these tools do not require identity checks or the transmission of content to external servers.

While these systems **do not perform age verification**, they represent an important complementary layer in a multi-layered ecosystem approach to children’s online safety. They illustrate how safety-by-design solutions can mitigate risks such as coercion, sextortion, and exposure to explicit content, even where age verification is not present or is bypassed.

# 3. Comparative Assessment of Age Verification Technologies

Age-verification methods differ significantly in accuracy, privacy impact, cost, usability, and interoperability. No single technology is mature enough to serve all use cases in the Digital Single Market. This section provides a structured comparison to help policymakers understand the trade-offs and identify gaps in the current ecosystem.

### 3.1 Comparative Assessment Matrix

Technology Type	Accuracy	Privacy / Data Impact	Cost & Complexity	User Experience	Interoperability	Notes
Self-declaration (age gates)	•	•	•	•••	•••	Minimal assurance; suitable only for low-risk contexts.
Document-based verification (ID scans, eID)	•••	•••	•••	•	••	Highly accurate; privacy-intensive; may exclude some minors.
Device-based age signals	••	•	••	•••	•	Dependent on OS ecosystems; useful as one layer only.
Biometric age estimation (face, voice, behavioural)	••	•••	••	•••	•	Sensitive; requires strict safeguards; variable accuracy across groups.
Third-party AV with privacy-preserving tokens	••• - •••*	•	••	•••	•••	Accuracy depends on upstream method; strong privacy profile; needs certification + governance.
Behavioural inference (AI risk scoring)	•	•	•	•••	•	Supplementary only; not reliable as standalone AV.

\* Token accuracy reflects the upstream method (ID-based verification vs. estimation-based verification); platforms receive only the age attribute, not identity data.

#### Complementary Safety Technologies (Not Age Verification)

Technology Type	Accuracy Relevance	Privacy / Data Impact	Cost & Complexity	User Experience	Interoperability	Notes
On-device, real-time harm-prevention systems	Not an AV method	•	••	•••	•••	Local AI analyses screen or camera activity to block harmful sexual content in real time; complements but does not replace AV.

Legend: Low = • | Medium = •• | High = •••

(High privacy/data impact = higher sensitivity or risk; High accuracy = better performance.)

## 3.2 Strengths and Limitations of Each Approach

### Self-Declaration Systems

The simplest and least reliable AV method.

#### **Strengths:**

- Zero friction and universally accessible
- Cost-free and easy to deploy

#### **Limitations:**

- Trivially easy to bypass
- Provides no meaningful assurance
- Inadequate for any risk category beyond low

**Conclusion:** Acceptable only in low-risk environments where minimal signalling is sufficient.

### Document-Based Verification

Includes ID upload, live capture comparisons, and authentication via national eID systems.

#### **Strengths:**

- Highly accurate; directly confirms age
- Legally recognised credentials

#### **Limitations:**

- Significant privacy and security risks
- Exclusion risks for minors lacking documents or parental support
- Creates friction and may discourage usage
- Disproportionate for many everyday online services

Best suited to *specific high-risk contexts*; not proportionate for general-purpose platforms.

### **Device-Based Age Signals**

Age inferred from device settings, parental controls, or OS-level profiles.

#### **Strengths:**

- Privacy-preserving with minimal data transfer
- Seamless user experience
- Works well when parental ecosystems are properly configured

#### **Limitations:**

- No guarantee that device user = device owner
- Strong dependency on major OS vendors
- Low interoperability across platforms
- Many households do not use child profiles consistently

Effective as a supporting layer within broader AV systems.

### **Biometric Age Estimation**

Uses AI models to estimate age from facial images, voice recordings, or behavioural signals.

#### **Strengths:**

- Can be frictionless and non-identifying
- Suitable for broad distinctions (e.g., over/under 18)
- Potential for on-device processing

#### **Limitations:**

- Accuracy varies across demographic groups

- Requires strict safeguards (no template storage, auditability)
- Sensitive under GDPR even without identification
- Public trust concerns and risk of misinterpretation

Deployment requires transparent governance, bias testing, and independent auditing.

### **Third-Party AV Providers & Privacy-Preserving Tokens**

A trusted intermediary verifies age and issues a reusable “yes/no” attribute.

#### **Strengths:**

- Strong alignment with GDPR data minimisation
- Platforms avoid processing identity data
- High interoperability, especially in multi-platform ecosystems
- Scalable for SMEs
- Supports competition among AV providers

#### **Limitations:**

- Requires certification frameworks and liability allocation
- Potential centralisation of market power
- Token-sharing must be mitigated (device/session binding, expiry, revocation lists)

Well suited to a tech-neutral, competitive ecosystem if governance is defined.

### **Behavioural Inference (AI Risk Scoring)**

Algorithms infer age likelihood based on metadata or user interaction patterns.

#### **Strengths:**

- Invisible to users; near-zero friction
- Useful for flagging suspicious patterns or guiding safeguarding measures
- Can complement other AV methods

**Limitations:**

- Insufficient accuracy for verification
- Risk of discrimination or opaque decision-making
- Should not be used as a standalone AV mechanism

Appropriate only as a contextual signal within a multi-layered approach.

### 3.3 Key Takeaways

- **No single technology meets all requirements** for accuracy, privacy, usability, and proportionality.
- **Hybrid approaches** are the most realistic path, combining different signals depending on context and risk.
- **Privacy-preserving architectures** — such as tokenisation or selective disclosure — align strongly with EU principles but require further development and governance.
- **Biometric and behavioural inference methods** show promise but require caution, transparency, and robust safeguards.
- **SMEs face disproportionate burdens** if solutions are too complex or costly, underscoring the need for scalable approaches.
- **Innovation is ongoing**, and the ecosystem is too immature for a single mandated EU-wide technology.

These findings support **regulatory sandboxes**, comparative evaluation environments, and a long-term commitment to **technology-neutral competition**, rather than prematurely choosing winners.

## 4. Key Issues in the EU Debate

As EU policymakers consider whether and how to introduce more structured age-verification (AV) requirements, several unresolved issues shape the debate. These debates sit at the intersection of fundamental rights, enforcement, competitiveness, and the functioning of the Single Market. Understanding these tensions is essential before moving toward any specific regulatory direction.

## 4.1 Privacy and Data Minimisation Under the GDPR

Age verification intersects directly with children’s rights and broader fundamental-rights obligations. Many AV technologies require processing personal – and sometimes sensitive – data.

Privacy-enhancing cryptography, including selective disclosure and zero-knowledge proofs (ZKPs), can reduce data flows by allowing verifiers to check an age threshold without accessing identity data or full birth dates. Hard-verification models relying solely on electronic IDs risk excluding significant parts of the population and are proportionate only in the highest-risk contexts.

### Key considerations include:

- **Data minimisation:** AV must use the least intrusive method capable of achieving the objective.
- **Purpose limitation:** AV data cannot be repurposed for advertising, profiling, security scoring, or unrelated platform functions.
- **Security risks:** Identity documents and biometric data require exceptionally strong protection.
- **Biometric sensitivity:** Even biometric estimation may raise GDPR concerns due to public perception and technical risks.
- **Identity creep:** Risk that users are required to prove identity-like attributes for ordinary online activities.

The EDPB and EDPS have repeatedly stressed that AV must not lead to de facto universal identity checks. This is why privacy-preserving architectures – such as selective disclosure, privacy-preserving tokens, and on-device processing – are attracting increasing attention.

A further consideration relates to proposals in which age, date of birth, and parent-child relationships are collected and stored directly at the operating-system or app-store level. While such models may simplify downstream implementation for applications, they create a highly centralised repository of sensitive information, increasing systemic risk and challenging GDPR requirements for data minimisation and purpose limitation. Concentrating persistent age attributes and family-relationship data within the device infrastructure introduces the possibility of identity creep, whereby a child’s personal data is reused across contexts in ways unnecessary for most online services. From a privacy perspective, any solution that embeds age assurance directly into the operating system must be assessed with heightened caution.

Explicit red line: system-level accumulation of children’s age, DOB, and

family relationships by OS/app-store vendors is incompatible with GDPR minimisation and purpose-limitation at EU scale. The EU should preclude such centralised repositories and require selective disclosure patterns (e.g., verifiable presentations, ZKPs) that avoid persistent, cross-service identity scaffolding.

## 4.2 Enforceability, Evasion, and Practical Feasibility

A key challenge is whether AV technologies can be reliably deployed across diverse services and user populations.

**Practical issues include:**

- **Evasion:** Minors can circumvent weak systems using VPNs, borrowed devices, false accounts, or generated images.
- **Accuracy variation:** Technologies such as biometric estimation show demographic variability in performance.
- **Platform diversity:** Requirements suitable for major platforms may be disproportionate for SMEs or low-risk services.
- **Fraud and misuse:** Widespread AV may generate incentives for attackers to compromise third-party providers or token systems.
- **User experience:** Excessively high friction drives users toward circumvention or unregulated services.
- **Platform-level age estimation already in use:** Some services deploy AI age-estimation to apply default protections for likely minors, with optional higher-assurance verification paths where appropriate.

Enforceability therefore depends not only on technology, but also on usability, adoption incentives, system integration, and proportionality.

## 4.3 Complementary Safety-by-Design Measures

In parallel to age verification, on-device, real-time harm-prevention systems are emerging as complementary tools that help mitigate risks such as sextortion, coerced self-generated imagery, and exposure to sexual content. These technologies use local AI models to analyse screen or camera activity directly on the device — including in end-to-end encrypted environments — to prevent the viewing, recording, or sharing of harmful material without transmitting data externally. While not a substitute for age verification, they demonstrate that AV alone cannot address the full range of online harms faced by minors and that a broader, multi-layered ecosystem approach is required.

## 4.4 Proportionality and Fundamental Rights

AV must be proportionate to the risks posed by a service — a principle reflected in the DSA and broader EU jurisprudence.

**Key concerns include:**

- **Privacy and anonymity:** Users should not have to reveal identity information to access general-purpose platforms.
- **Chilling effects:** Intrusive verification may deter socially beneficial or legitimate online interactions.
- **Digital inclusion:** Not all minors have compatible devices, documents, or stable parental support.
- **Mass-surveillance concerns:** Overreliance on biometric or behavioural inference at scale raises systemic risks.

Balancing child protection with fundamental rights remains one of the most complex tensions in the debate.

## 4.5 Competition, Market Structure, and Innovation

AV obligations may unintentionally affect market dynamics.

**Key risks include:**

- **Incumbency advantage:** Large platforms can absorb AV development costs; SMEs may struggle.
- **Gatekeeper effects:** Device-based AV could increase dependence on major OS providers.
- **Market concentration:** Without governance, a handful of AV providers may dominate the market.
- **Innovation slowdown:** Premature standardisation may freeze out emerging solutions.

A **technology-neutral, competitive framework** is essential to avoid locking in early technologies or reinforcing dominant positions.

**Gatekeeping risk:** placing OS/app-store providers at the centre of age assurance would consolidate market power, create developer dependency, and distort competition in AV markets, contrary to the spirit of the DMA. The EU should

avoid AV architectures that depend on a single commercial device ecosystem, require non-exclusivity and open standard interfaces, and assign liability to the service providing the regulated activity (e.g., adult sites, gambling services) rather than to upstream intermediaries.

## 4.6 Fragmentation Across Member States

Diverging national AV laws are already emerging in France, Germany, Ireland, Spain, and other Member States.

### Risks include:

- **Breakdown of the Digital Single Market:** Platforms may need to implement different systems in each national market.
- **Disproportionate burden on SMEs:** Fragmentation raises compliance costs and operational complexity.
- **Uneven protection:** Children in different Member States may receive unequal levels of safety.
- **Regulatory arbitrage:** Services may situate operations in jurisdictions with lighter requirements.

This fragmentation strengthens the case for **EU-level coordination** and minimum requirements.

Device-centred or operating-system-level proposals also raise questions of interoperability. System-level age and parental-consent mechanisms may function well within a single device ecosystem, but they often do not generalise across different operating systems, browsers, or device types. In multi-device or multi-OS households — as well as in contexts involving shared devices, public computers, or cross-border mobility — the assumptions underlying these models break down. Without strict interoperability safeguards, OS-based approaches could introduce uneven protection and complex implementation barriers across the Single Market.

## 4.7 The Role of eIDAS 2.0 and the EUDI Wallet

The EUDI Wallet enables selective disclosure of attributes — such as “over 13,” “over 16,” or “over 18” — without sharing identity information.

### Potential advantages:

- Strong privacy by design; no need to reveal full date of birth or identity.
- Supports cross-border interoperability.
- High-quality assurance when backed by government-issued credentials.

### Challenges:

- Adoption will take years, especially for minors.
- Not all minors will have access to Wallet credentials; parental involvement introduces friction.
- Integration standards and certification processes are still evolving.
- Should be an **optional acceptance path**, not a prerequisite for access.

The Wallet is a promising future component but **cannot yet serve as the core system** for AV across the EU.

In addition, proposals that rely on proprietary operating-system infrastructures to collect and distribute age attributes risk creating private, parallel identity layers that sit outside the EU's emerging framework for digital identity. As the EUDI Wallet develops, the EU's strategic objective is to ensure that digital identity attributes — including age — are governed through interoperable, standards-based, and sovereign mechanisms. System-level AV solutions controlled by commercial device manufacturers could undermine this objective by establishing de facto identity authorities without the transparency, auditability, or public oversight required under EU law.

Public over proprietary identity. As EUDI matures, age attributes should flow through EU-governed, standards-based wallets — not private, opaque OS identity layers. Proprietary system-level AV risks creating parallel identity infrastructures without public oversight.

## 4.8 Governance, Certification, and Trust

Technical performance alone is insufficient; AV requires trustworthy governance frameworks.

### Key questions:

- Who sets evaluation criteria and performance benchmarks?

- Should the EU maintain reference datasets for accuracy and bias testing?
- What certification processes should providers undergo?
- How should liability be allocated across platforms, AV providers, and intermediaries?
- How will auditing, recertification, and oversight be conducted?
- Should certifications be mutually recognised across Member States?

Without governance, even privacy-preserving designs can be undermined by poor implementation.

Centralising age assurance at operating-system level also raises governance concerns. Proposals of this type typically lack clear frameworks for certification, auditing, or liability allocation, despite relying on persistent and privileged system-level access. Without explicit oversight mechanisms, transparency obligations, and enforceable safeguards, it would be difficult to ensure that system-level age-assurance functions operate in a rights-respecting and non-discriminatory manner. Governance gaps of this kind would pose significant challenges in an EU regulatory environment grounded in accountability and high-assurance privacy protections.

Certification boundaries. OS/app-store AV functions should not be treated as a default “trusted authority.” Certification should be provider-agnostic, with open test suites, demographic bias reporting, and revocation/accountability that any compliant AV provider can meet – not privileges granted by device control.

Figure 2. Summary of Key Issues in the EU Debate on Age Verification

Issue	Description	Why It Matters	Policy Tension	Practical Implications
<b>Privacy &amp; Data Minimisation (GDPR)</b>	AV often requires processing personal or sensitive data.	Fundamental rights, public trust, GDPR compliance.	Safety vs. identity creep and over-collection.	Risk of excessive data flows; potential for de facto ID checks; need for selective-disclosure and minimisation.
<b>Enforceability &amp; Evasion</b>	Systems face bypassing, demographic accuracy gaps, and UX friction.	Ineffective AV undermines child safety and compliance.	Rigour vs. usability and adoption.	VPNs and borrowed devices; friction-driven circumvention; AI estimation already used for defaults in some services.
<b>Complementary Safety-by-Design Measures</b>	On-device, real-time harm-prevention that operates independently of AV.	AV alone cannot address grooming, sextortion, CSAM creation/exposure.	AV scope vs. holistic safety design.	Functions in E2EE; reduces harm even when AV is bypassed; illustrates need for multi-layered safety frameworks.

Issue	Description	Why It Matters	Policy Tension	Practical Implications
<b>Proportionality &amp; Fundamental Rights</b>	AV strength must match service risk.	Preventing overreach, chilling effects, and exclusion.	Protection vs. anonymity and free expression.	Heavy AV on low-risk services is disproportionate; risk of deterring legitimate use.
<b>Competition &amp; Gatekeeping Risks</b>	OS/app-store-centred AV may concentrate power.	Impacts innovation, SME viability, and DMA objectives.	Open market vs. gatekeeper-centric identity/AV.	Developer lock-in; foreclosure of innovative AV; <b>single-ecosystem dependency</b> ; harder SME entry; risk of <b>cross-context tracking</b> via persistent system IDs.
<b>Fragmentation Across Member States</b>	Divergent national AV laws and technical requirements.	Threatens the Digital Single Market and legal certainty.	National flexibility vs. EU coherence.	Multiple per-market builds; uneven protection; regulatory arbitrage.
<b>Role of the EUDI Wallet</b>	Selective age-attribution disclosure via EU digital identity infrastructure.	Long-term, privacy-by-design pathway.	Readiness vs. ambition.	Slow rollout; minors may lack credentials; parental involvement adds friction; cannot be sole AV solution in near term.
<b>Governance &amp; Certification</b>	Need for common standards, audits, and liability rules.	Ensures accountability and comparability across AV providers.	Innovation pace vs. oversight.	Difficult to assess provider claims; lack of EU benchmarks; OS-level models lack clear certification pathways.
<b>Lack of One-Size-Fits-All Solution</b>	AV ecosystem still maturing.	Prevents premature lock-in and rights risks.	Urgency vs. evidence.	Premature mandates risk high cost with little safety gain; supports need for sandboxes and evaluation environments.

## 4.9 Summary: A System Not Yet Ready for One-Size-Fits-All Solutions

Across all issues, several themes emerge:

- The AV ecosystem remains **technologically immature**.
- **Key components — governance, standards, interoperability — are not fully developed.**
- A single mandated EU-wide AV technology would be premature and likely disproportionate.
- The EU should instead enable **controlled experimentation, innovation, and rigorous evaluation**.

This supports a strategic approach built on:

- **technology neutrality**
- **competition and innovation**
- **transparent assessment environments**
- **regulatory sandboxes**
- **evidence-based policymaking**

## 5. Policy Options for the EU

The EU faces a dual challenge: ensuring strong online protection for minors while safeguarding privacy, proportionality, innovation, and the integrity of the Digital Single Market. Because no single age-verification (AV) technology is mature or universal enough to meet all requirements, the EU's role is not to choose a technology but **to create conditions for safe experimentation, competition, and evidence-based decisions.**

The following options offer a structured, non-prescriptive menu for policymakers. They can be combined or sequenced over time.

### Option 1 — Establish Minimum EU-Wide Technical and Governance Requirements

The EU could define **baseline expectations** for AV performance, security, and governance without mandating specific technologies. This would reduce fragmentation while maintaining strict technology neutrality.

#### Possible elements:

- Context-dependent performance ranges (not fixed thresholds)
- Transparency obligations for providers
- Data minimisation and security standards
- Requirements for demographic bias testing and reporting
- Clear retention and deletion rules

- Strict prohibition of identity disclosure unless strictly necessary
- Ban exclusive dependence on proprietary OS/app-store age signals; require open, standardised acceptance paths.

**Benefits:**

- Reduces fragmentation across Member States
- Protects minors while upholding GDPR principles
- Allows technological competition and innovation

**Limitations:**

- Difficult to set performance ranges without stifling innovation
- May still burden SMEs without guidance or shared tools

## Option 2 — Promote Privacy-Preserving Age Tokens and Verified Attributes

This model separates **verification** from **access**. A trusted intermediary verifies age and issues a reusable “yes/no” token (e.g., “over 13,” “over 16,” “over 18”). Platforms receive only the age attribute — never the identity. Acceptance should be mandatory for standards-conformant presentations (OpenID4VP, SD-JWT, ISO 18013-5), ensuring no single commercial ecosystem controls AV acceptance.

**Possible mechanisms:**

- Certified third-party AV providers
- Attribute issuance through the EUDI Wallet
- One-time verification with reusable tokens
- Mechanisms to bind tokens to sessions/devices
- OpenID4VP/SD-JWT/ISO 18013-5 presentations for age attributes, with guidance on revocation, replay resistance, and verifier governance.

**Benefits:**

- Strong GDPR alignment; data minimisation by design
- Scalable for SMEs

- Enhances user privacy and anonymity
- Supports cross-platform interoperability
- Reduces the need for platforms to process sensitive data

**Limitations:**

- Governance required (certification, auditing, liability)
- Token-sharing risks require mitigation (expiry, revocation lists, device binding)
- Not all users will have digital credentials in the short term

## Option 3 — Create EU-Level Age-Verification Sandboxes

A sandbox model would allow controlled, regulator-supervised testing of AV technologies, generating real-world evidence before large-scale deployment.

**Potential priorities:**

- Comparative testing of biometric vs. non-biometric AV
- Evaluation of privacy-preserving methods (e.g., tokens, selective disclosure)
- UX testing to minimise friction
- Measurement of demographic accuracy and bias
- Fraud and evasion analysis
- Interoperability trials (platform ↔ token providers ↔ identity systems)

**Benefits:**

- Supports innovation without premature standardisation
- Generates evidence for future policymaking
- Identifies unintended consequences early
- Helps SMEs test solutions at low cost

**Limitations:**

- Requires specialised technical and regulatory expertise
- Not a substitute for enforcement in the short term

## Option 4 — Encourage Integration with the EUDI Wallet

The EUDI Wallet could, over time, provide a trusted infrastructure for selective disclosure of age attributes.

### User experience examples:

- “Age attribute only” shared without identity
- Single verification reused across platforms
- No transfer of underlying documents

### Benefits:

- Strong privacy by design
- High assurance when credentials are verified
- Potential cross-border consistency

### Limitations:

- Slow rollout; adoption will take time
- Minors may lack credentials or parental support
- Integration and certification processes are still evolving
- Should be an **optional acceptance path**, not a universal requirement

The Wallet offers a **long-term pathway**, but interim solutions remain necessary.

## Option 5 — Apply a Risk-Based, Tiered Approach Across Online Services

Blanket restrictions on general-purpose or low-risk digital tools, including educational or AI-driven services, risk undermining digital inclusion and European competitiveness without meaningfully improving child safety. Rather than imposing uniform AV requirements, the EU could tailor AV expectations to the risk level of the service.

**Illustrative tiers:**

- **Low-risk services:** Self-declaration + contextual signals (forums, informational sites)
- **Medium-risk services:** Device-based signals, behavioural indicators, or tokens (social platforms with messaging or UGC)
- **High-risk services:** High-assurance AV (adult content, regulated goods), using ID-based or certified biometric estimation where lawful and proportionate

**Benefits:**

- Aligns with the DSA's proportionality principle
- Avoids overburdening services with minimal risk
- Helps SMEs by focusing obligations where they matter most

**Limitations:**

- Requires clear EU-level criteria for risk classification
- Increases enforcement complexity if not harmonised
- Risk of divergence if Member States interpret tiers differently

Tiering applies to service-side obligations; it does not justify outsourcing AV liability to OS/app-store intermediaries.

## Option 6 — Restrict or Condition the Use of Biometric Age Verification

Given high sensitivity and legal complexity, the EU could impose **strict conditions** on biometric AV or prohibit certain forms unless safeguards are met.

**Possible restrictions:**

- On-device processing only
- No storage of biometric templates
- Mandatory demographic accuracy/bias reporting
- Independent audits and recertification
- Explicit user transparency

- Use strictly limited to age estimation — never identity verification or law-enforcement matching

**Benefits:**

- Addresses public concerns about surveillance
- Encourages development of privacy-preserving biometric solutions
- Provides regulatory clarity

**Limitations:**

- Overly restrictive rules could slow innovation
- Requires ongoing oversight as models evolve

This option balances caution with room for responsible innovation.

Figure 3. Policy Options for the European Union

Policy Option	Description	Intended Benefits	Key Limitations / Risks
<b>Option 1 — Minimum EU-Wide Technical &amp; Governance Requirements</b>	Establish baseline expectations for AV performance, security, transparency, data minimisation, bias testing, and governance — without prescribing specific technologies.	Reduces fragmentation; increases legal certainty; protects minors while upholding GDPR; maintains tech neutrality.	Premature thresholds could limit innovation; SMEs may face compliance challenges without guidance or shared tools.
<b>Option 2 — Privacy-Preserving Tokens &amp; Verified Attributes</b>	Separation of verification and access through reusable, privacy-preserving age tokens using open standards (OpenID4VP, SD-JWT, ISO 18013-5).	Strong GDPR alignment; minimal data flows; high interoperability; scalable for SMEs; reduces platform liability exposure.	Requires certification, auditing, liability clarity; token-sharing risks must be mitigated; uneven access to digital credentials in short term.
<b>Option 3 — EU-Level AV Sandboxes</b>	Controlled environments for comparative testing of AV technologies, privacy techniques, UX, bias, fraud, and interoperability.	Generates evidence; derisks innovation; informs future regulation; supports SMEs; avoids premature mandates.	Requires specialised regulators and infrastructure; not a short-term enforcement tool.
<b>Option 4 — Integration with the EUDI Wallet</b>	Use of the EU Digital Identity Wallet for selective age-attribute disclosure and high-assurance verification when available.	Strong privacy-by-design foundation; cross-border consistency; reusable verification; high trust.	Slow rollout; minors may lack credentials; parental involvement adds friction; not suitable as a sole AV method in the near term.

Policy Option	Description	Intended Benefits	Key Limitations / Risks
<b>Option 5 — Risk-Based, Tiered Approach Across Services</b>	Tailor AV expectations to service-level risk categories (low/medium/high). Avoid blanket restrictions for general-purpose or educational/AI tools.	Aligns with DSA proportionality; reduces unnecessary burden; focuses strongest requirements on high-risk services; supports innovation and inclusion.	Requires clear EU criteria for risk tiers; risk of divergent national interpretations; enforcement complexity.
<b>Option 6 — Conditions or Restrictions on Biometric AV</b>	Allow biometric age estimation only under strict safeguards (on-device processing, no template storage, audits, bias reporting).	Protects rights; addresses public concerns; encourages development of privacy-preserving biometric methods; provides legal clarity.	Overly restrictive rules may slow innovation; requires ongoing oversight as models evolve.

## Overall Guidance Emerging from the Options

Across all pathways, several consistent principles emerge:

- **Do not choose a single technology.**
- **Avoid premature EU-wide mandates** before sufficient evidence exists.
- **Enable competition and innovation** by keeping the regulatory framework open and tech-neutral.
- **Prioritise privacy, proportionality, and interoperability.**
- **Focus on governance, certification, and evaluation standards** rather than fixed technical prescriptions.

The overarching goal is to create an ecosystem where **multiple AV solutions can develop, be tested, and compete**, guided by safety objectives and fundamental-rights protections.

## Annex — International Landscape Snapshot

Age verification (AV) has become a global policy priority, yet governments are approaching it through diverse regulatory pathways shaped by different legal traditions, risk perceptions, and technological capabilities. Around the world, authorities are experimenting with varying combinations of standards, pilots, and enforcement strategies, and no jurisdiction has yet converged on a definitive AV model.

Although these approaches cannot be directly imported into the EU's fundamental-rights framework, they provide valuable comparative insights into market maturity, regulatory risks, and emerging best practices. This annex summarises developments in four influential jurisdictions — the United Kingdom, the United States, Australia, and Canada — each of which illustrates distinct regulatory assumptions and policy trajectories in the international AV landscape.

Figure A1. International Approaches to Age Verification — Comparative Snapshot

Jurisdiction	Regulatory Approach	Key Features	Relevance for the EU
<b>United Kingdom</b>	Online Safety Act (OSA) with Ofcom-led codes of practice	<ul style="list-style-type: none"> <li>• “Highly effective” AV standard without prescribing technologies</li> <li>• Conformity assessments and audits central to compliance</li> <li>• Emphasis on privacy-by-design and minimisation</li> <li>• Codes effectively shape permitted AV practices</li> </ul>	Shows how regulator-defined standards and audits guide AV without empowering device gatekeepers or naming a technology.
<b>United States (State-Level)</b>	Patchwork of state laws targeting social media & adult content; no federal framework	<ul style="list-style-type: none"> <li>• Divergent requirements across states (ID-based AV, third-party AV)</li> <li>• Widespread litigation on First Amendment grounds</li> <li>• Geofencing required to comply with different rules               <ul style="list-style-type: none"> <li>• Strong focus on parental consent</li> </ul> </li> </ul>	Fragmentation drives up costs and legal risk — a warning against multiple proprietary, OS-tied AV paths inside the Single Market.
<b>Australia</b>	eSafety Commissioner’s Age Verification Roadmap; phased and exploratory	<ul style="list-style-type: none"> <li>• Strong focus on harm reduction for adult content</li> <li>• Recognition that AV ecosystem not ready for mandatory rollout</li> <li>• Supports pilots, trials, industry codes               <ul style="list-style-type: none"> <li>• Evaluates biometric and non-biometric AV in controlled settings</li> </ul> </li> </ul>	Supports the EU argument for sandboxes, incrementalism, and evidence gathering before imposing AV mandates.

Jurisdiction	Regulatory Approach	Key Features	Relevance for the EU
Canada	AV embedded in broader online harms agenda; model still evolving	<ul style="list-style-type: none"> <li>• Debate on linking AV to national digital identity systems</li> <li>• High emphasis on privacy, anonymity, anti-surveillance</li> <li>• Active exploration of tokenisation and selective disclosure</li> <li>• No federal AV mandate</li> </ul>	Mirrors EU tensions around privacy, digital identity, and proportionality; reinforces the need for rights-compatible design in AV solutions.

## Insights from International Trends

Across jurisdictions, several consistent patterns emerge in how governments are approaching age verification. Despite differing legal traditions and policy cultures, six clear themes can be observed:

- 1. No country has adopted a definitive or fully mature AV solution.** All rely on experimentation, pilots, phased implementation, or transitional arrangements. The global landscape is exploratory rather than settled.
- 2. Technology neutrality is the prevailing principle.** Regulators avoid mandating specific AV technologies and instead focus on performance criteria, safeguards, and governance structures.
- 3. Privacy and proportionality concerns are universal.** Even in strong regulatory environments, biometric and ID-based AV are treated cautiously due to sensitivity, public trust concerns, and proportionality requirements.
- 4. Fragmentation creates compliance burdens.** The United States illustrates how patchwork obligations quickly become unmanageable for industry and inconsistent for users — a critical warning for the EU's Single Market.
- 5. Sandboxes and staged development are gaining traction.** Australia and the UK emphasise pilot programmes, controlled trials, and iterative regulation before imposing binding mandates.
- 6. Governance matters as much as technology.** Certification schemes, audits, liability frameworks, and trust mechanisms are emerging as the central preconditions for effective and safe AV systems.



# European Policy Innovation Council

## About EPIC

EPIC is a thinking movement that combines the research capacity of traditional think tanks with the community engagement and mobilisation of political movements. By blending classic initiatives such as policy papers and conferences with innovative digital formats and on-site tours, EPIC is building a large community of policymakers, industry leaders, experts, activists, and influencers in Europe. Together, we aim to become Europe's first public policy influencer, shaping the future of competitiveness, sustainability, and democratic resilience.

## Transparency and Legal Notice

EPIC is a nonprofit registered under Belgian law as an ASBL (BE.1016.989.669) and the EU Transparency Register (REG 683253994641-35). All opinions and publications represent the views of the respective authors and do not necessarily reflect the positions of EPIC or any affiliated organizations or individuals.

## Contact Information

European Policy Innovation Council (EPIC) asbl  
Rue du Chatelain, 8, Box 4  
1000 Brussels, Belgium

2026 © European Policy Innovation Council

[www.thinkepic.eu](http://www.thinkepic.eu)

---