defend | digital | me

Who is age-gating our life online?

Analysis of tech in the Australian Social Media Minimum Age Ban, emerging proposals in Europe, and the outcomes for privacy and data protection rights



defend|digital|me

Citation: Defend Digital Me. (2026). Who is age-gating our life online? Analysis of the technology used in the Australian Social Media Ban, emerging proposals in Europe, and the outcomes for privacy and data protection rights.

Available at https://defenddigitalme.org/research/

Published under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International licence (CC BY-NC-SA 4.0).

Version 3.0— published October 20, 2026

Authors: Jen Persson and Madeleine Stone

The authors wish to thank contributors including companies, for their time taken to discuss; and every reader and reviewer of drafts for comments and suggestions. We gave the companies mentioned an opportunity for comment and to contribute ahead of publication. This report reflects the publicly available information as of March-October 2026, reported in good faith, represented fairly and accurately to the best of our interpretation. Given the emerging status of laws, policy and technology, we welcome corrections or missing updates identified, and will maintain them online. For all enquiries please contact us via: info@defenddigitalme.org

Generative AI academic usage statement: Claude was used during the preparation of this work to assist with the building of and to consistently format tables, to ensure consistent usage of key terms, to check grammar and spelling, and to identify mistakes in internal numbering and referencing of the appendix structure. Generative AI was not used to develop the core ideas or draft the substance of the report and all responsibility for any mistakes rests with the lead author.

Contents

Executive Summary 4
1. Introduction 7
2. The reality of age checks in Australia 11
3. Liveness checks: what they are, why they are used, and who does them 25
4. The myth of no ID required 31
5. Biometric data and a legal challenge 34
6. The industry: standards, certification and economics 40
7. Is the ban working? 52
8. VPN bans and content blocking 66
9. The European picture 68
10. The "digital age of consent" myth and the age of majority 75
11. Problems with the evidence base 77
12. Four approaches compared: Australia, Canada, the UK and the EU 80
13. Recommendations: beyond age verification 82
14. Conclusion — and what is ultimately at stake 86

References 89

Appendices

Appendix A — What age-check methods are offered (SMMA, Australia)? 104
Appendix B — About the third-party age-check providers 105
Appendix C — Lawful-basis mapping and where Article 8 of the GDPR applies 107
Appendix D — Age-based proposals across Europe 109
Appendix E — Age-based proposals across Rest of World 113
Appendix F — VPN bans and platform-blocking timelines 115
Appendix G – Intentions, predictions and outcomes of the SMMA 117
Appendix H — Age-assurance consent architecture by provider 120
Appendix I — Liveness checks by age-assurance provider 124
Appendix J — Technical standards governing Age Assurance, AV, and identity 126
Appendix K — Proof of identity documents required for age verification 131

Executive Summary

In August 2026, France's Constitutional Council blocked a bill banning social media access for under-15s, saying it infringed upon freedom of expression and privacy[1]. In its decision, the court also said, "By prohibiting minors under the age of fifteen from accessing certain online services, the law inherently requires every person, even an adult, to prove their age before accessing them".

Making children safe online is an outcome. Age assurance is a process. This report shows what happens when the technique is treated as though it were the outcome, and what has been built in the space between those two things in Australia, mapping the reality underpinning these platforms' age-gating measures. This report examines what its requirements mean in practice. Age checks on their own leave the harms children face online largely untouched, while building a permanent identity layer, often tied to a live face scan, between every person and the digital public sphere. The cost and risk of that layer fall first on the children these policies are said to protect.

The UK was among the first governments worldwide[2] in announcing policies to restrict access to social media by age, described as an "Australia-plus" social media ban. Politicians suggested it will go further than the antipodean model and cover more platforms and features. In addition to under-16 restrictions, overnight curfews are to be on by default for 16 and 17-year-olds that can be switched off, and with features that prolong scrolling switched off. Others are calling for a more Canadian-style approach. Full details will only be known once the draft regulations are published.[3]

After Australia became the first country to ban under-16s from social media in December 2025, its government has doubled down on restrictions on platforms "that are not doing enough"[4] to keep them off.

The process behind checking everyone's age is poorly understood by the public and rarely discussed by policymakers. Most methods involve several steps. A person first proves that a claimed age belongs to them, using an identity document, a bank record or another state-accepted source. The restricted online service then checks that proof, itself or through a third-party provider, often with a face scan and a liveness check, and may keep profiling the user's behaviour during use. The more certain the law requires the result to be, the more intrusive the process becomes, and data that leaves the device may pass through sub-processors in several

[1] Décision n° 2026-911 DC du 14 août 2026, Loi visant à protéger les mineurs des risques auxquels les expose l'utilisation des réseaux sociaux The Conseil constitutionnel struck down Article 1 of the law that would have barred under-15s from social media from 1 September. It found the restriction disproportionate to the freedom of expression and communication of the under-15s themselves, in the terms quoted at Section 1. It held separately that because barring under-15s requires every user including every adult to prove age, and because the legislature specified no conditions or limits for how that proof is furnished, the necessary legal safeguards were absent. Privacy under the 1789 Declaration featured as a second ground. The President has asked the Prime Minister to redraft.
[2] Tech Policy Press Global Social Media Age Restriction Tracker https://social-media-age-tracker.onrender.com/
[3] DSIT (2026) Default overnight curfews and addictive features to be switched off automatically to protect 16-17-year-olds https://www.gov.uk/government/news/new-social-media-curfews-and-crackdown-on-addictive-features-to-better-protect-16-and-17-year-olds-online
[4] Prime Minister of Australia (June 2026) Stronger powers double the penalties for world-leading social media law https://www.pm.gov.au/media/stronger-powers-and-double-penalties-world-leading-social-media-law

countries.

To make checks enforceable, states then try to close the routes children use to evade them, through VPN bans and platform blocking. These are tools associated with further harms from censorship and surveillance, and they remove the anonymity that protects whistleblowers, journalists and people escaping domestic violence, as well as children's right to privacy.

These measures engage the rights of adults and children alike: privacy, freedom of expression and access to information, association, play, and children's right to be heard (UNCRC Articles 16, 13, 15, 31 and 12). Children's protection from economic exploitation (Article 32) is also at stake, because age assurance turns children's faces and behaviour into the raw material of a growing market.

The evidence from Australia shows restrictions failing on their own terms. Around 70% of under-16s who previously had accounts have kept access. Of those who lost access, one in two say they now get less news (Notley et al., 2026), and half report no difference to their safety (Molly Rose Foundation, 2026).

Public support for a ban depends on how it works. In June 2026, 69% of British adults supported a law requiring age verification to keep under-16s off social media. Support fell to 50% if everyone in the UK had to upload ID or a credit card, and to 48% if the law required a total ban on VPNs (Ipsos, 2026). The public also trusts parents (51%) and an independent regulator (49%) to decide which platforms suit children, well ahead of technology companies (16%) and government ministers (15%) (IPPR, 2026).

What these measures reliably do is normalise the processing of biometric and bodily data at scale, by third parties placed between platforms and the people who use them, and teach children to perform for a camera on demand. Human rights impact assessments of the UK ban, the EU KIDS Act and the EU age verification app must consider these behavioural changes and their effect on consent, agency and the development of children's identity.

We ask that policymakers rethink the reality of age-gating the Internet and

- treat the normalisation of on-camera coercive compliance as its own child protection risk;
- End the routine processing of children's facial biometric data and bodily data for age checks;
- Commission a full and independent child rights impact assessment of the age-assurance measures or social media age restrictions, for example in the UK social media ban for under 16s, or plans for an EU mini-wallet and of every accredited attestation provider within it, covering all users of all ages, published before any deployment. This must expressly be about behaviour, as well as the technical security and data risks to identity;
- Recognise that with the March 2026 statement of 438 security and privacy researchers from 32 countries, and with the positions of the Council of Europe Commissioner for

Human Rights and the UN High Commissioner for Human Rights, age verification poses a range of risks to children's safety in the digital environment;

- Educate children and parents on the individual and collective risks to data loss and identity fraud using biometric or bodily data or digital ID which may be lasting lifetime risks to the future adult and to the future society population at vast scale worldwide;
- Act on the duty towards children that they be heard under Article 12 of the UNCRC before final policy decisions and drafting of regulations, not after.

Almost all child rights and human rights organisations worldwide whose positions the authors know, oppose blanket bans. Some children's charities take a different view. Barnardo's welcomed the UK ban in June 2026 as one part of a wider package (Barnardo's, 2026), and the NSPCC, which had initially opposed a blanket ban, said in February 2026 that a ban would be "better than nothing" if the government failed to act on technology companies (NSPCC, 2026).

Several decisions due in the coming months will settle how age is checked: Ofcom's assessment of age-assurance methods, due at the end of October 2026; the UK's regulations under section 214A of the Online Safety Act 2023, expected by the end of 2026; the UK bill on device-level protections now being drafted; the EU KIDS Act in the European Parliament and Council; and France's redrafted ban. This report is intended to inform each of them.

Our concern is that the tool becomes the goal and its consequences are far reaching and poorly understood — "a minimum age limit for social media" — presupposes the answer to the question that most needs asking of if, how, and why that keeps children safe online. In racing to age-gate the internet, the practical outcomes of these processes mean policymakers risk steering children directly into modelling the very harms they intend to prevent.

1. Introduction

> “The face you give the world tells the world how to treat you.”
>
> — Gillian Flynn, Sharp Objects

Many advocates of social media bans — policymakers, child protection organisations, parents and children — are unaware of the systems behind the network of global data transfers undertaken in order to check an age online, required for any social media restriction by age.

Platforms make signing up quick and easy. But in doing so, they do not always draw attention to important choices about privacy or how personal information will be passed along a chain of sub-processing. Because of this, users may not fully understand how their data is being used, for how long, or what options they have regarding privacy settings, before providing consent (Krishnan-Barman et al., 2025).

We believe the behavioural demands these intermediaries make of children in particular, are as important as the nuances of debate over whether bodily data falls into the category of biometric data, its security, or what "privacy preserving" may mean. It is changing how children behave.

Most age restriction methods require children to perform for the camera in a facial scan and liveness check, and/or submit selfies to an unknown number of third-party companies with which they have no relationship. Teaching children to trade their bodily data for everyday participation in digital spaces — under manufactured and coercive "consent," choosing the least difficult path to get online even if they don't really want to give up that data — is a terrible way to teach them how to respond when adults instruct them to do something they do not want to do, framed as "required" for their own online safety.

We argue that providers' and politicians' architecture of choice should not be imposed at the expense of children's fundamental rights, nor remove choices a child may not have wanted to make before giving up their digital identity or proof of personhood until their own adulthood — giving up your facial features to a third party is for life. Biometric and bodily data need protection from excessive use and access even by the well-intentioned, and once lost, are lost forever.

These systems process people mechanistically, requiring each of us to perform on demand for a computer in order to prove that we are alive, but in doing so, the obligations can harm our agency, the nature of consent, and determine the everyday experience of what is required to demonstrate proof of being human to a computer for millions of people. What that looks like, on whose terms and conditions, and how we can exercise our rights and or overrule a computer that wrongly says ‘no’ all remains poorly understood and with little democratic debate. What is at stake in authenticating proof-of-liveness is not only privacy, or the technical merits of a security architecture, but who gets to control what it means to be human in the digital environment.

1.1 Scope

We explore what the age verification measures and its choice architecture mean in practice and their implications for human rights under the Online Safety Amendment (Social Media Minimum Age) Act 2024 ('the SMMA Act'), which requires social media platforms to prevent under-16s from accessing their services. Ten platforms are covered as per the Australia ban: Facebook, Instagram, Snapchat, TikTok, Threads, X, YouTube, Reddit, Twitch and Kick. These platforms must take "reasonable steps" to ensure that children under 16 do not hold accounts, or face fines of up to AU$49.5 million, a maximum doubled in September 2026 by the Strengthening Enforcement amendments (Parliament of Australia, 2026).

What was said before the ban was introduced, can now be measured against the March 2026 evidence post-ban introduction. This report therefore compares:
(i) the reality of the SMMA changes in Australia[5] for children's data privacy rights, three months after implementation as found by the eSafety Commissioner ("eSafety"); (eSafety, 2026a)
(ii) what the Office of the Australian Information Commissioner (Australia's data protection authority) said should be done to protect privacy under the ban[6]; and compares that against
(iii) the Australian Child Rights Taskforce warning[7] about the ban as "too blunt an instrument,"
(iv) other proposals discussed in the UK and beyond, to data protection law ("the digital age of consent") and restrictions on Virtual Private Networks (VPNs) (EPRS, 2026).

Finally, we present the global implications. The logic of coercive age-checking everyone does not stop at age. The same escalation that pushes a 15-year-old from an email check, to a face scan connected to verified state ID check pushes the whole system toward a single, state-verifiable, liveness-bound identity, "Know Your Customer" for the entire population, at the point of access to the digital public sphere. We lift a little the curtain that is what a platform user sees, to look at the industry behind it.

The scope of evidence from practice in Australia should inform what happens next in the UK, wider Europe, and the European Commission age-verification "mini-wallet" app. Human rights impact assessment is missing the fundamentally important questions of what these changes mean: the normalisation of these behavioural changes and their effect on the freely given nature of consent, agency, capacity, and the development of personhood and identity in children into adulthood.

In Europe, the EDPB identifies grooming, bullying, stalking and harassment as forms of malicious targeting that the age assurance process must not facilitate, placing them alongside commercial profiling as outcomes barred by purpose limitation and data minimisation (EDPB, 2025, p. 4, para. 16). "*In general, it should not be possible to learn more than necessary about a natural person and their actions by profiling this person based on the information used in the age assurance process*" (EDPB, 2025, p. 4, para. 16). However, we show that in many processes of ID checking for AV, in

[5] Online Safety Amendment (Social Media Minimum Age) Bill 2024 https://www.aph.gov.au/Parliamentary_Business/Bills_Legislation/Bills_Search_Results/Result?bId=r7284
[6] The OAIC Social Media Ban https://www.oaic.gov.au/privacy/your-privacy-rights/social-media-minimum-age
[7] Australian Child Rights Taskforce 2024 Open Letter https://apo.org.au/node/328608

identity document checks, in age estimation and inference techniques, and in liveness checks, the reality is different. The EDPB states outright that a provider processing personal data to check the age of *all* users accessing *all* content, where that content is suitable for all audiences and devoid of risk, would fail the necessity and proportionality tests (EDPB, 2025, p. 4, para. 14). Yet this appears commonplace.

What's next? We suggest that the legal manoeuvring used to argue bodily data lies outside the protection of biometric-data is the same one now being tested, and in places rejected, against another proof-of-personhood technology for know-your-customer online, iris scanning by orbs. If the technical definition of biometric data is weakened to permit workarounds of lawmakers' intent at the time of making, the protection of such bodily data will end irreversibly. Instead of reaching for ever-escalating technical solutions to social problems, we propose an end to the routine use of children's bodily data and methods defined by companies who have lobbied hard to drive demand for their products. We need expert, rational and informed debate about the implications of performative demands of companies on children, on society broadly, and the range of alternatives.

Like hundreds of child rights experts and 438 computer scientists in March 2026, we call for a moratorium on any social media "bans" or restrictions by age, on policy that extracts digital ID and performative liveness checks from children, with effects that last a lifetime.

Risk assessment of the reality of these well-intentioned public policy aims has simply not been carried out or debated.

A note on terms

Age assurance is the umbrella term. Verification, estimation and inference are its three distinct methods as defined in ISO/IEC 27566-1:2025, with differing ranges of certainty and accordingly when a platform may decide to use one over another, often related to the level of obligation that a public policy or legislation demands. This report tries to keep the terms apart and used accurately.

Verifymy describes in a report how methods are evolving. Not only are documentary evidence and in-person physical features or behaviour used to identify or age check a person, but some methods now join up other available evidence about an individual's past digital activity across multiple environments to 'infer' age. "Historically, age assurance has been understood through two core categories: age verification and age estimation. Age verification confirms an individual's exact age using official documents (such as a passport or other form of government-issued ID) to confirm exactly how old someone is by comparing the birthdate document to the current date. Age estimation by contrast, predicts a user's age or estimated minimum age using physical features (facial age estimation) or behavioural characteristics (e.g. email-based age checks) (Verifymy, 2026c.) Part E of Australia's Age Assurance Technology Trial focused specifically on age *inference* – a method of determining an individual's likely age or age range based on verifiable contextual, behavioural, transactional or environmental signals, rather than biometric data or identity documents. Unlike age verification, which relies on a known and validated date of birth or age estimation, which uses biometric characteristics to predict age, age inference draws "reasonable conclusions" about age by analysing facts such as school enrolment, financial transactions, content barring settings, service usage or participation in age-specific activities."

Term	How this report uses it
Age assurance	The umbrella term. Any process used to establish or approximate a person's age or age range.
Age verification	Establishes age from identity information in a document or authoritative record bearing a date of birth.
Age estimation	Approximates age from measured characteristics such as a facial image, a voice sample or an email footprint. Always probabilistic.
Age inference	Deduces age from a proxy: examples can include a credit card or bank-account based details, on-platform behaviour, an account creation date.
Age check	Not an ISO term. We use it only for a single user-facing moment of age assurance, whatever method sits behind it.
Liveness check	Confirms that a real person is at the camera at that moment (usually at the point of access to set up a social media account / platform or every time the user logs on). Says nothing about age itself unless the liveness check also includes facial / bone structure analysis for age assurance / estimation for example.

2. The reality of age checks in Australia

It is key to recognise these are not bans on children accessing social media but obligations on platforms to ensure every user is of a minimum age to participate in a given defined online activity. It places obligations on the providers of services to check the age of every user to tell adults apart from children, or differentiate between children of different ages, and they are doing this through various technologies, both in-house and using a wide range of third-party providers.

While it is third-party age check providers that attract most scrutiny, the more invasive age inference is often being done in-house by the platforms themselves — and receives far less attention because it is framed as a safety feature rather than a verification step. Meta, TikTok and Google all run their own machine-learning models that infer a user's age from on-platform behaviour. Meta continuously scans all Australian Instagram accounts using "age estimation, object detection and keyword detection," drawing on signals from mentions of birthday celebrations or mentions of school grades, and, in selected countries, visual analysis of photos and videos for cues such as height or bone structure (Meta, 2026a).[8] It further uses these data to train its AI models (Meta, 2026b). What Meta calls age estimation here is largely age inference in ISO terms: it reads contextual signals from posts and captions rather than measuring facial characteristics. Only the element drawing on height or bone structure sits at the estimation end. On the other hand, TikTok runs a comparable model over profile, device, published videos and interactions. And Google reserves the right to use Chrome history, search terms, content the user creates or receives, location data, and third-party activity as inputs to its over-18 determination.

Three things make this category meaningfully different from the third-party data flow to perform an age check. First, it is continuous rather than one-off — every action feeds the model. Second, there is no real opt-out or genuine freely given consent; users cannot decline the in-house assessment in the way they can choose between a selfie or an ID upload. Third, the dataset is far broader than anything an outside third-party provider such as Yoti, or an orchestrator such as k-ID, sees, and in Meta's case is openly used to train its other commercial AI products.

The SMMA Act has, in effect, made routine and legitimised the largest platforms running profiling models over Australian user's content under the banner of child safety, with near no public debate. In the United States, Meta's August 2026 consent judgment with state attorneys general requires it, within a year, to apply an age assurance method to every user of Facebook and Instagram in the settling states, with any user not age-assessed within 14 days given teen default protections (Kentucky Office of the Attorney General, 2026, EFF 2026).

Our platform-by-platform summary of methods is at Appendix A; the providers behind them at Appendix B. In this section we set out our analysis of the methods employed by each platform and their processes and subprocesses to meet minimum age restriction obligations.

[8] Meta (2026) New AI-Powered Age Assurance Measures to Place Teens in Age-Appropriate Experiences https://about.fb.com/news/2026/05/ai-age-assurance-teens/

In Appendix K, we set out the ID documents accepted for age verification by platform and providers, typically the establishment of age from identity information contained in an identity document bearing a date of birth, as defined in ISO/IEC 27566-1:2025. It excludes age estimation (facial or voice analysis, email footprint) and age inference (credit card possession, behavioural profiling, account signals), which are addressed at Appendix A and Appendix I.

2.1 The ten social media platforms in the Australian scope

This section examines how the ten social media platforms covered by Australia's Social Media Minimum Age Act actually conduct age checks — what children are asked to do, which companies process the resulting data, and what happens to it afterwards.

Four findings emerge from the platform-by-platform analysis that follows.

- **The platform is rarely the party performing the age-check.**
- **Age assurance is rarely at a single point, but a process or continuous condition of access.**
- **Claims that age verification does not require digital ID are false. ID checks may not happen in the moment of platform use, but the intrusion and obligation are prior to this.**
- **The intrusiveness of the check escalates precisely for the age group the scheme aims at.**

Facial age estimation is a probabilistic guess, and it is least accurate near the boundary it's asked to police, i.e. a 15-year-old and a 17-year-old look far more alike than a 12-year-old and a 30-year-old. So a passive check (one still selfie, no instructions) most often fails to return a confident answer for users close to 16. This means that this is not of high enough confidence to satisfy newly emerging laws that oblige platforms to obtain proof of *knowledge*, not simply '*best guess*' whether the user meets the minimum age requirements. Australia's Age Assurance Technology Trial found false rejection rates at a 16+ age gate of 8.5% for 16-year-olds and 2.6% for 17-year-olds, and identified buffer zones of two to three years around each age gate in which errors cluster (MediaNama, 2025). Yoti, the leading provider, reports a mean absolute error of 1.04 years for 14 to 18-year-olds (Tombs, 2025c), at an age when a single year decides access.

When the passive check isn't confident enough, providers escalate: the user is routed into an active liveness check with in-screen demands such as nodding, turning the head, blinking, speaking aloud, submitting a video rather than photograph, or requiring more ID upload.

The children the minimum age restricts target are those most likely to be pushed into the more demanding, more data-hungry step. Users who appear well outside the threshold — a clear adult, or younger child — get acted on after a lighter check. The burden falls hardest at the margin on the most vulnerable, and children, the audience that such measures are said to be aimed at protecting.

2.1.1 Meta (Facebook, Instagram and Threads)

Meta, owner of three of the platforms covered by the SMMA Act, states that it uses a "menu of options" offering a combination of its own age inference technology and third-party age estimation by Yoti to conduct age checks. A user of any age, wanting to make an account on Facebook, Instagram or Threads is prompted to enter their birthday. If they state they are over 16, they must either submit a video of their face or upload a copy of an ID (Meta, 2022).

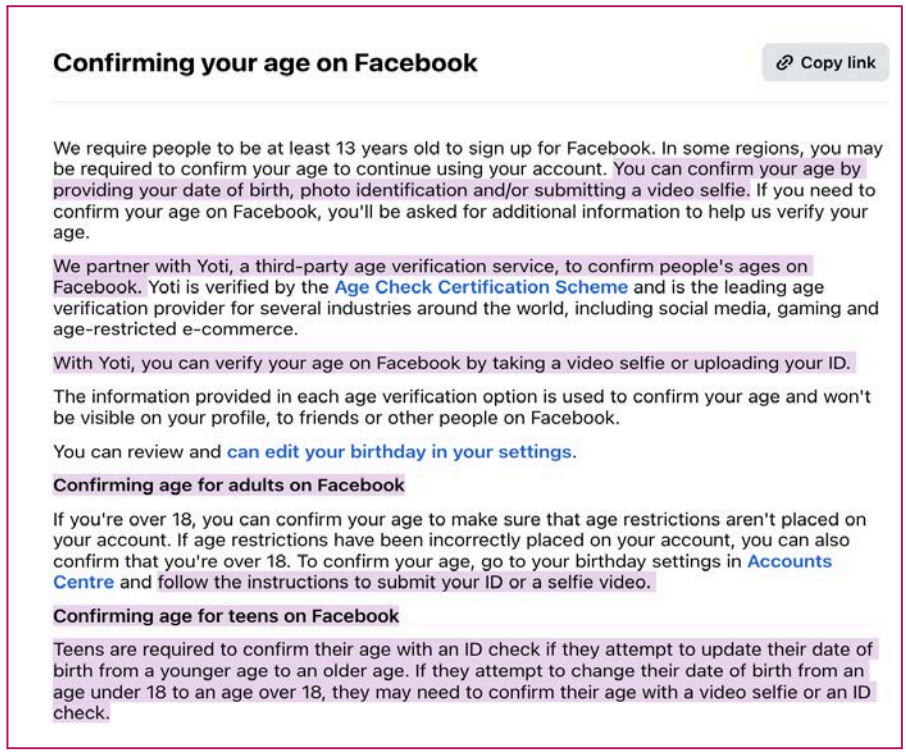

Confirming your age on Facebook

Copy link

We require people to be at least 13 years old to sign up for Facebook. In some regions, you may be required to confirm your age to continue using your account. You can confirm your age by providing your date of birth, photo identification and/or submitting a video selfie. If you need to confirm your age on Facebook, you'll be asked for additional information to help us verify your age.

We partner with Yoti, a third-party age verification service, to confirm people's ages on Facebook. Yoti is verified by the Age Check Certification Scheme and is the leading age verification provider for several industries around the world, including social media, gaming and age-restricted e-commerce.

With Yoti, you can verify your age on Facebook by taking a video selfie or uploading your ID.

The information provided in each age verification option is used to confirm your age and won't be visible on your profile, to friends or other people on Facebook.

You can review and can edit your birthday in your settings.

Confirming age for adults on Facebook

If you're over 18, you can confirm your age to make sure that age restrictions aren't placed on your account. If age restrictions have been incorrectly placed on your account, you can also confirm that you're over 18. To confirm your age, go to your birthday settings in Accounts Centre and follow the instructions to submit your ID or a selfie video.

Confirming age for teens on Facebook

Teens are required to confirm their age with an ID check if they attempt to update their date of birth from a younger age to an older age. If they attempt to change their date of birth from an age under 18 to an age over 18, they may need to confirm their age with a video selfie or an ID check.

Source: Confirming your age on Facebook https://www.facebook.com/help/958848942357089

Meta uses the device's camera to take a video selfie[9]; a facial image is then sent to Yoti behind the scenes, which is used to automatically estimate whether the person is over or under 16 and sends a 'pass/fail' response back to Meta (Meta, 2026d). Alternatively, the person can upload a copy of an ID document — a national ID, passport or driving licence[10] — which is then sent to Yoti for verification. Yoti states it deletes any data relating to the age check, while Meta states it will retain the facial image or ID photo for up to 30 days. Both note that they share personal data with law enforcement or regulatory bodies where required by law (Meta, 2026c; Yoti, 2026a).

[9] https://help.instagram.com/1053588012132894
[10] Types of IDs that Facebook accepts https://www.facebook.com/help/159096464162185

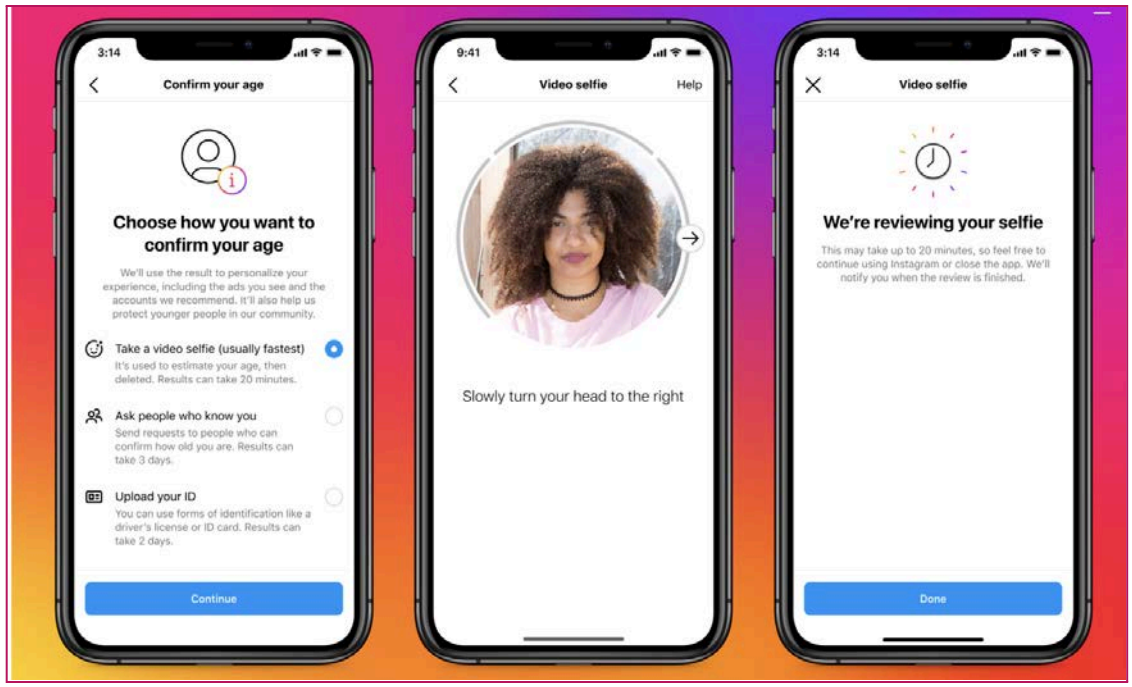

Source: Meta: instagram **New ways to verify your age on Instagram**.
https://about.fb.com/news/2022/06/new-ways-to-verify-age-on-instagram/

How your age is verified

If you choose to take a video selfie, we partner with vendors who use technology to estimate your age. After you've taken a video selfie, we'll share an image from that selfie with the vendor. Our vendors' technology only estimates age. It can't identify specific people.

Once your age has been estimated and shared with us, we will delete the image within 30 days. We don't share any other information about you with the vendor.

How to take the selfie

If you choose to take a video selfie to verify your age, you'll see instructions on your screen for how to take the selfie. If you haven't used the camera on Instagram before, we'll ask you for permission to use the camera so that you can take the selfie.

When you take the selfie, make sure that your whole face is visible and that you have enough light for the camera. Avoid wearing anything that blocks your face, such as a mask or hat. You'll be asked to move your face in different directions, such as to the right and left, and up and down.

Once your age has been estimated and shared with us, we will delete the image within 30 days. Once you have submitted your selfie, you can keep using Instagram or exit the app.

It can take up to 20 minutes to verify your age. You will get a notification with the outcome of the verification.

Source: Meta: **How does video selfie age verification work on Instagram?**
https://help.instagram.com/292521846402382/?helpref=related_article

As an additional measure, Meta scans all Australian Instagram accounts continuously for contextual clues (Meta, 2025) that aim to identify the age of an account holder. Photos, posts, captions, bios and comments are screened for keywords and objects that may indicate a user is under 18, automatically moving the user to a 'Teen Account'. Age estimation, object detection and keyword detection are used across users worldwide, with visual analysis available so far only in selected countries, with Meta stating that everything from mentions of birthdays to the child's height or bone structure will form part of the assessment (Meta, 2026a).

Meta uses all of this data to train its generative AI models. This may be less surprising than first seems, when one considers Meta already collects an extraordinary amount of personal data from its users (and even from people without accounts). Its published generative AI policy says that even children's images who have no Meta account may be part of their training datasets (Meta, 2026b).

> *"Even if you don't use our Products or have an account, we may still process information about you to develop and improve AI at Meta. For example, this could happen if you appear anywhere in an image shared on our Products by someone who does use them or if someone mentions information about you in posts or captions that they share on our Products. This could include information about people who are under 18 years old."*
>
> — Meta, How Meta uses information for generative AI models (2026b)

By partnering with Yoti, Meta outsources part of the age checks and passes sensitive personal data to a third party. Meta's support pages do not always name Yoti, instead referring to "vendors who use technology to estimate your age," (Meta, 2026d) making it challenging for children and parents to understand who is processing their data.

Meta has also announced it said it would begin integrating a new age verification system in Australia and other markets during 2026 — AgeKey, a reusable age credential developed by French company Opale.io, which k-ID acquired in November 2025 and relaunched under the separate OpenAge Initiative, and which Snapchat already offers (Bradshaw, 2025).

2.1.2 Snapchat

Snapchat uses Singapore-based k-ID's age assurance 'middle-man' services, which, according to k-ID, route Australian users to four checks:[11] an ID document check through Veratad, facial age estimation through Privately SA, an age check through ConnectID, in which a participating bank confirms the user is over 16 from the date of birth it already holds or a credit card check through Stripe. Users can also store the result as a reusable AgeKey, issued through the separate OpenAge Initiative, or present one they already hold. Although k-ID manages the process on Snapchat's

[11] https://security.k-id.com/subprocessors

behalf, a range of subcontractors undertake the actual checks. k-ID's privacy policy states that, where it manages the age assurance process, it partners with "age assurance providers" that allow it to offer multiple options (k-ID, 2026c).

If the user uploads an ID, the user is routed to US company Veratad Technologies LLC, which captures the document in its own environment; k-ID says it does not collect or have access to the document and receives the result of the check. While k-ID states it does not store the information provided, Veratad's privacy policy is vague (Veratad, 2026) and says it, "does not use end-user verification data for unrelated secondary commercial purposes," but it is unclear how ID-document data is used to carry out checks or what is done with it afterwards. Veratad states it processes "public and non-public records, images of government-issued IDs, selfies, and other personal information," (Veratad, 2026).

If the user chooses facial age estimation, the check runs on the user's device using software from Swiss company Privately SA, and no image data is passed to k-ID or to Privately SA's servers; once the check is completed, the data is deleted from the user's device. Although Privately SA does not currently use user data to train its AI, its privacy policy (Privately SA, 2026) states "this might change in the future."

For the ConnectID method, a service operated by Australian Payments Plus Ltd, the user chooses a participating bank, signs in to it and consents to the bank sharing an 'over 16' answer. ConnectID, Australian Payments Plus and k-ID do not see the user's bank credentials or underlying identity data (Australian Payments Plus, personal communication, October 2026). ConnectID partners with several Australian banks (Westpac, CommBank, NAB and ANZ Plus), confirming whether a user meets the age requirement and passing a pass/fail response back without accessing identity data. This is assurance at the point of the check. The bank has already verified identity from documents at account opening, so the ID-check sits earlier in the process and before the point-in-time of the age check request by a platform. For example, to set up a CommBank account an under-18 can provide CommBank with either a Medicare (health service) card, or their Australian birth certificate. For adults, CommBank accepts any passport or Australian driving licence (CommBank, 2026). Each time a platform or its intermediary requests an age check through ConnectID, the participating bank creates and retains a "Transaction Record" including the identity of the Service Provider, the date of the transaction, and the types of identity data shared (ANZ Plus, 2026). Westpac and NAB state they will share Transaction Record details with ConnectID for operation, governance, fraud detection and "other purposes." (Westpac, 2026; NAB, 2026). ConnectID says the record it receives is limited to transaction metadata, such as the bank, the platform and the type of attribute shared, and does not include the customer's identity, date of birth or whether they passed the age check (Australian Payments Plus, DDM communication, October 2026). NAB's terms say the "other purposes" are those set out in the ConnectID Privacy Policy (NAB, 2026).

The user can also create an 'AgeKey' — a reusable credential stored as a passkey on the user's device, which presents the result of an earlier check to any service that accepts AgeKeys without a fresh check being run, allowing Snapchat to see only whether the user is over or under 16 and retaining no data about where it was used.

Game platforms (Apple iOS, Google Play, Xbox, Meta Horizon, and k-ID itself) can also provide age data about a player. This data can be passed to k-ID alongside, or instead of, a self-reported age. k-ID uses it to suppress the age gate when appropriate, resolve age conflicts, and, when the signal is considered verified, satisfy age verification requirements without additional verification steps.

Platform age signals work in combination together with further age assurance steps for designated "high-risk features". For example, a verified Apple iOS signal can both skip the age gate and unlock loot box permissions in Brazil without a separate verification step, according to k-ID (K-ID, 2026b).

While UK-based Verifymy is listed by k-ID as an 'available' age assurance method via email check, Verifymy is not a deployed provider for Snapchat or Kick (or any other platform in Australia) via k-ID. Neither platform routes checks to Verifymy via k-ID.

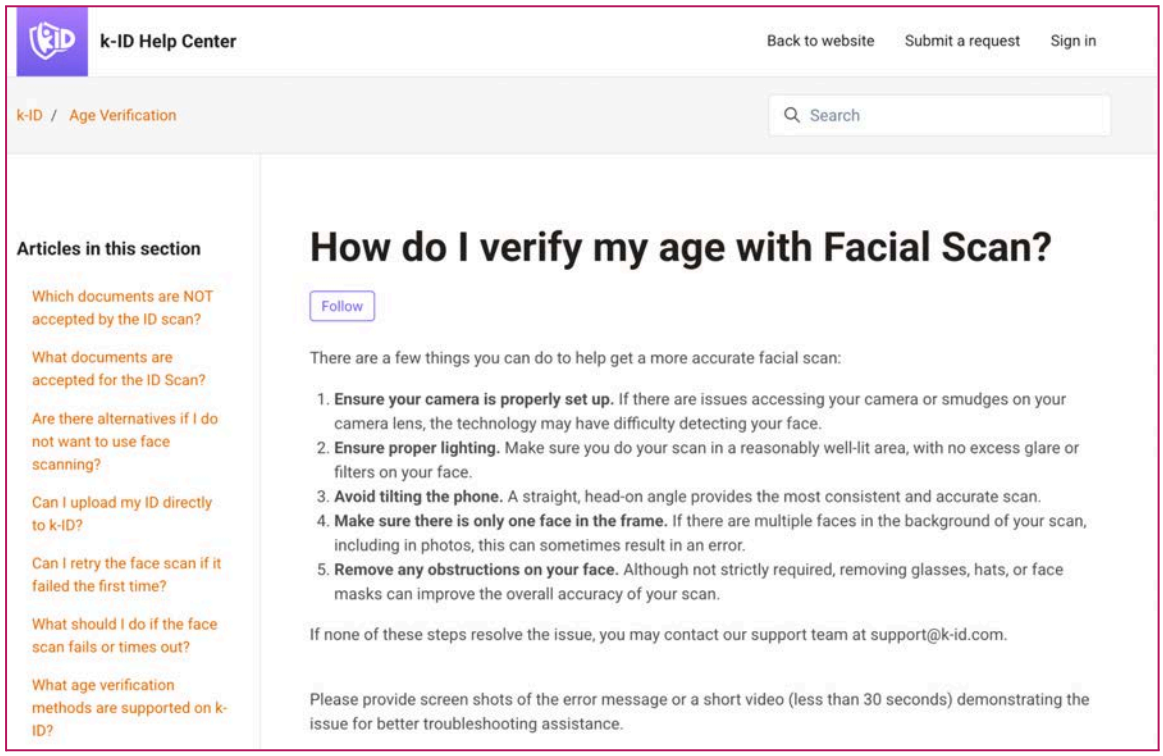

An example of facial age estimation screens shown through k-ID's orchestration service on Snapchat.[12]

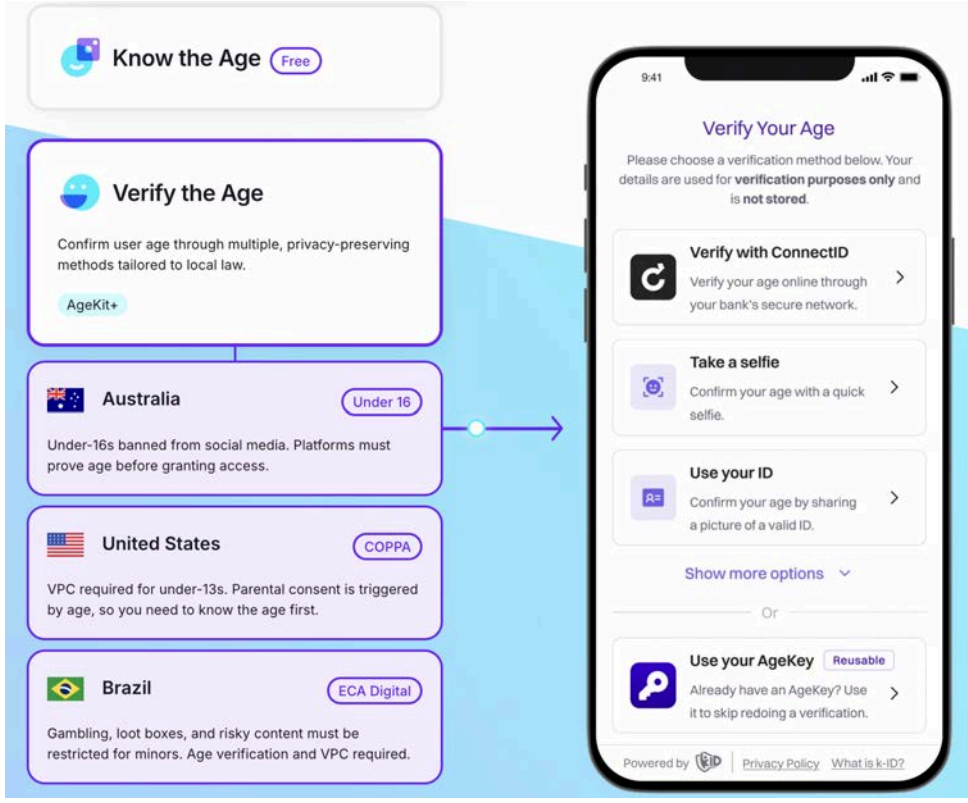

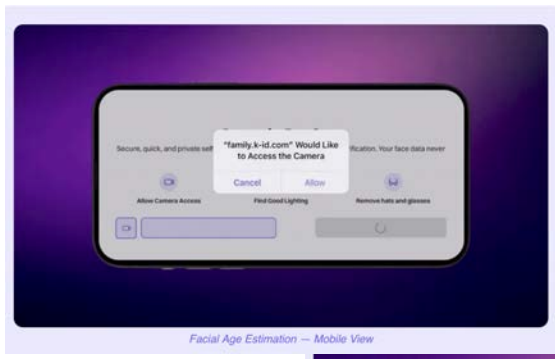

Facial Age Estimation — Mobile View

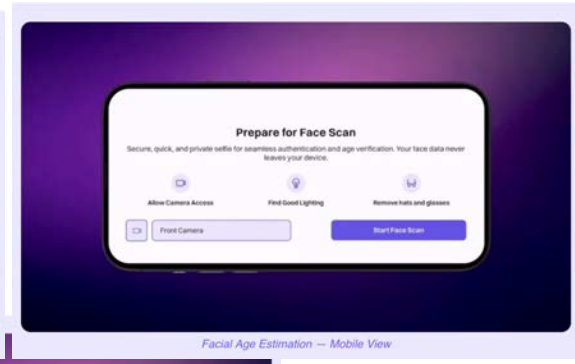

Facial Age Estimation — Mobile View

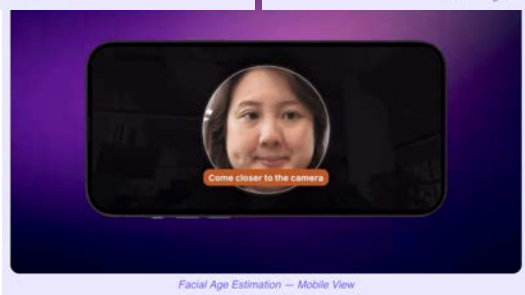

Facial Age Estimation — Mobile View

[12] Snapchat and Kick both ask users to perform a live facial scan for the age check via k-ID https://help.snapchat.com/hc/en-us/articles/42550358600084-Account-Restrictions-for-Australia-Users | Kick https://help.kick.com/en/articles/12947994-age-assurance-faqs#h_ef2a74b4d9

2.1.3 TikTok

TikTok said it uses facial age estimation, credit card authorisation, or ID checks in Australia[13]. Yoti provides the technology for facial age estimation (TikTok, 2025), but TikTok does not identify the third parties it uses for credit card or ID checks, referring only to "our service providers, who help us check if your ID is authentic [and] make a temporary credit card authorization" (TikTok, 2026a). Like Instagram, TikTok also runs an age inference model over the account holder's, drawing on the user's "profile, device information, the videos they published, their interactions with other users… and other on-platform behavior." (TikTok, 2026b). TikTok states it keeps this data "for no longer than necessary." In February 2026, the European Commission preliminarily found TikTok in breach of the Digital Services Act for its addictive design, including infinite scroll, autoplay, push notifications, and "its highly personalised recommender system."

TikTok LIVE allows viewers and creators to interact in real time, and at the time of writing, to use that service TikTok requires users to confirm their age using a selfie with ID through Incode, which requires users to submit photo(s) of your government-issued ID, and capture a selfie of your face taken in real time (TikTok, 2026c).

Incode expanded its architecture in 2026 with the acquisition of Israeli-founded Identiq.[14] The acquisition is Incode's third in two years. The company bought identity verification rival MetaMap Inc[15] in 2024 (founded in 2017 and headquartered in San Francisco) and AuthenticID in 2025.

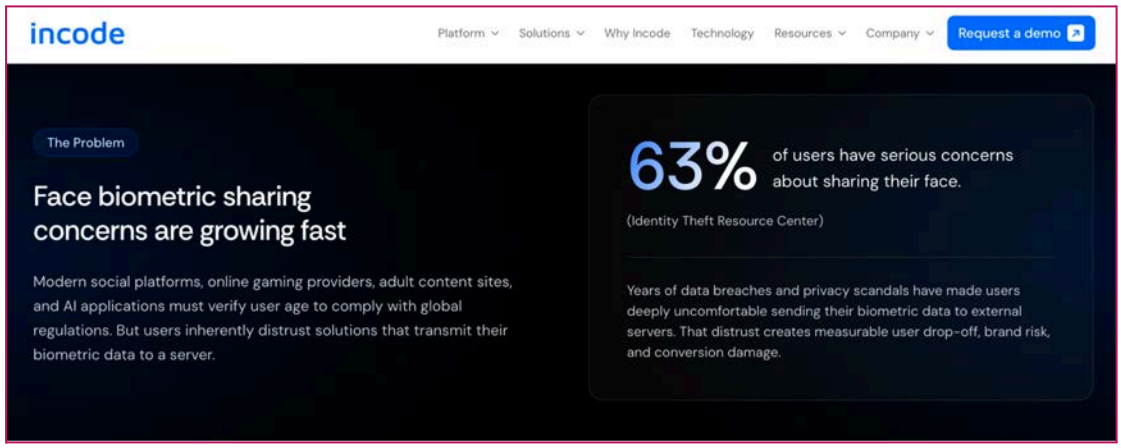

2.1.4 Reddit

Reddit uses US company Persona, stating that users must use "ID, facial selfie, [or] open bank account checks." (Reddit, 2026). Persona's own site states (Persona, 2026b) that Australian Reddit users can use only government ID or facial age estimation. To undertake a check, Persona collects facial images and biometric data, or ID-document data, much like other providers. However,

[13] https://newsroom.tiktok.com/how-tiktok-is-complying-with-a-new-government-law-in-australia
[14] https://www.incode.com/blog/incode-identiq-acquisition
[15] https://siliconangle.com/2024/06/11/incode-acquires-metamap-enhance-identity-verification-services/

Persona states it can also collect device information (including IP address, device type, OS, browser, cookie and device identifiers), account information, geolocation data, usage data (including “copy and paste detection”), and wireless device data.

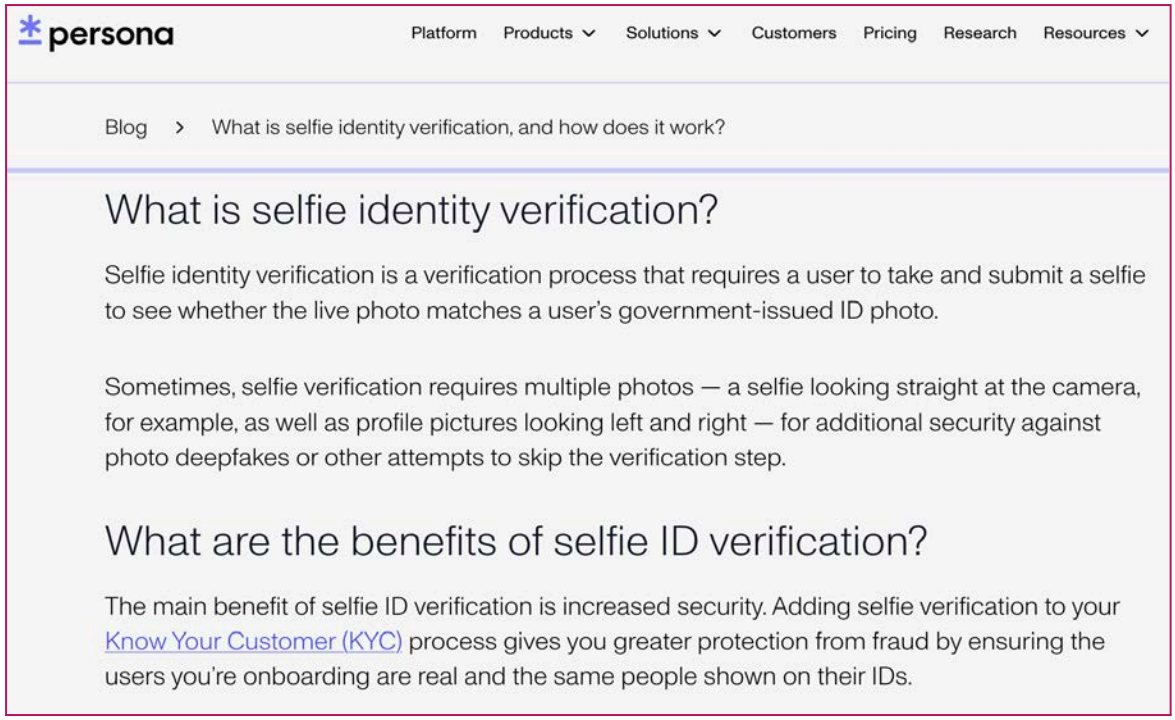

Source: Persona: What is Selfie identity verification?

A profile of Persona in Forbes noted that its “machine learning models can also pick up risk signals such as the network a person’s using, how far they are from the location on their ID, the way they’re interacting with a device and their online presence.” (Shrivastava, 2025).

Persona also states it can cross-reference this data with a “network of trusted third-party data sources” including “publicly available sources (such as open government databases), government and national ID registries, consumer credit bureaus, utility companies, mobile network providers and postal address databases.” Persona grants itself a wide margin to collect considerable data on each adult or child, states it deletes Reddit users’ data after three days, and notes it “may engage third parties” — without naming them. Persona says that what it collects depends on the age check method and on how the platform buying its service has configured it, and that facial age estimation needs no identity document (Persona, n.d.). A child or parent has no way to see which of these settings their platform has chosen. Persona describes seven separate age check groupings, of which the platform picks one or more. Two of them draw on data held elsewhere. Email-based age inference checks “public and proprietary data sources” for how long an address has existed, its digital footprint and whether its domain is linked to a particular demographic. Database verification cross-references what a user provides against credit bureaus and government records; Persona says these queries check existing records and add nothing to them

(Persona, communication to DDM October 2026). Some people and groups are concerned about the large-scale collection of government issued IDs and facial data that companies can collect over time especially where there are connections with the defence industry, for example, Persona is financially backed by Founders Fund, Peter Thiel's VC fund (Castro, 2026). In February 2026, researchers examining Discord's age checks found a publicly exposed Persona frontend. Its code showed the software could run 269 distinct verification checks, including facial recognition against watchlists and lists of politically exposed persons and screening across 14 categories of adverse media, and could retain IP addresses, device fingerprints, government ID numbers, names and faces for up to three years. Persona said the exposed environment was a non-production system holding no customer data, and that no customer uses all 269 checks. Roblox and OpenAI are also reported to use Persona for age checks (Arntz, 2026).

2.1.5 X

X uses a range of methods. For existing accounts, it uses inference signals: self-attested age, verified status and account creation date, alongside email-address age estimation. X states it uses US companies Stripe, that currently supports five types of verification checks: document, selfie, ID number, address, and phone; and Persona for ID checks for identity verification, although this is a separate, voluntary process for X Premium users. If a user cannot be verified by those methods, X's policy states it will undertake facial age estimation or ID verification. (X, 2026b). X does not explicitly state who carries out these checks for Australian users, though it is likely Persona, given their role in ID verification for X Premium. X provides no further information about how much data is shared, stating only that "third party providers will also be bound by strict data protection standards." (X, 2026a).

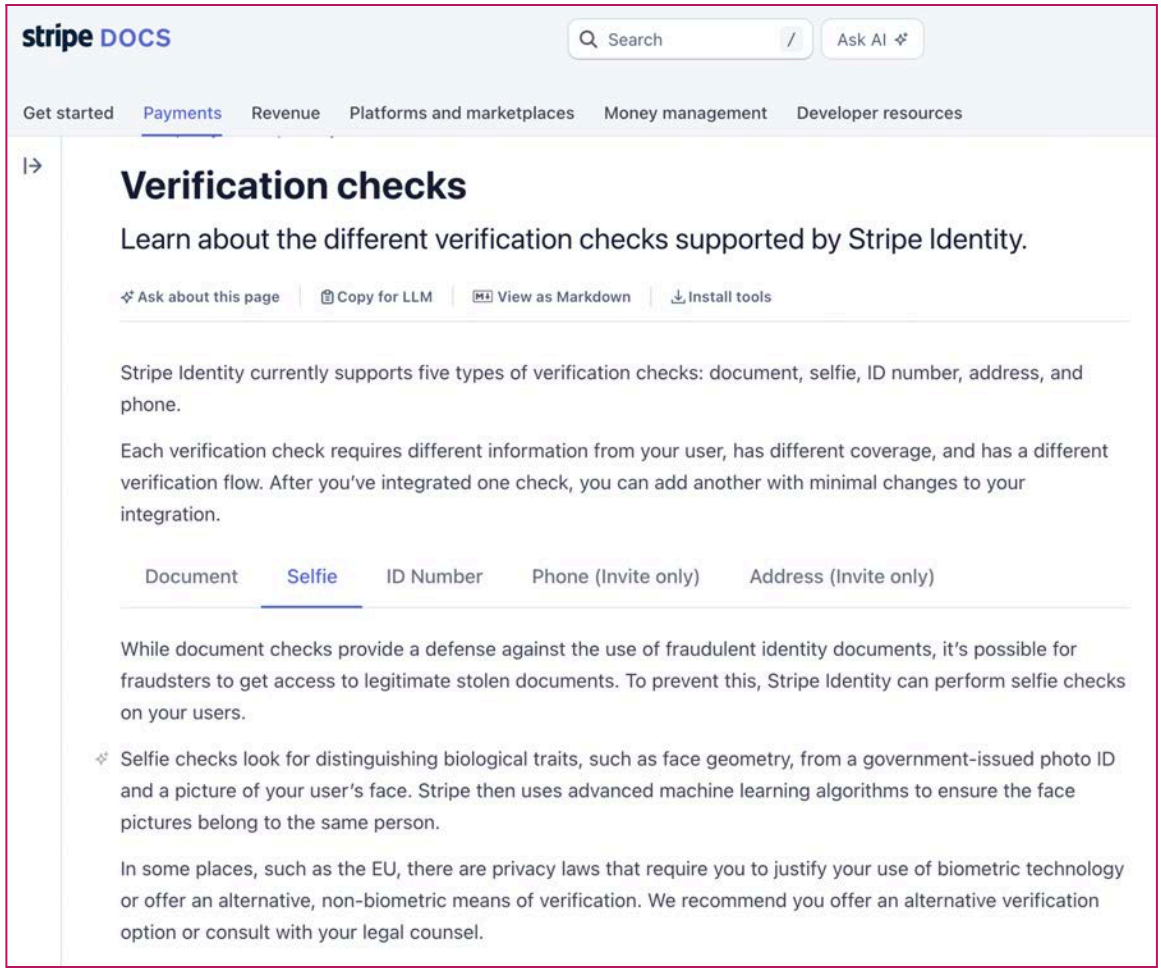

Source: Stripe
https://docs.stripe.com/identity/verification-checks

2.1.6 YouTube

YouTube, owned by Google, gives users four options: uploading a government ID, facial age estimation, email-address age estimation, or a credit card check (Google, 2026a). If a user uploads an ID, Google undertakes the check itself and states the data can be used to "improve our verification services for Google products and protect against fraud and abuse." (Google, 2026b). Google says ID data will be "stored securely" and "won't be made public," without commitment to deletion.

Google uses US company Private Identity LLC ('PrivateID') for facial age estimation. The selfie and biometric data are stored on-device and not uploaded; the results are "scrambled with strong encryption," sent to PrivateID's server, which returns a 'pass/fail', and stored for 90 days. (PrivateID, 2023).

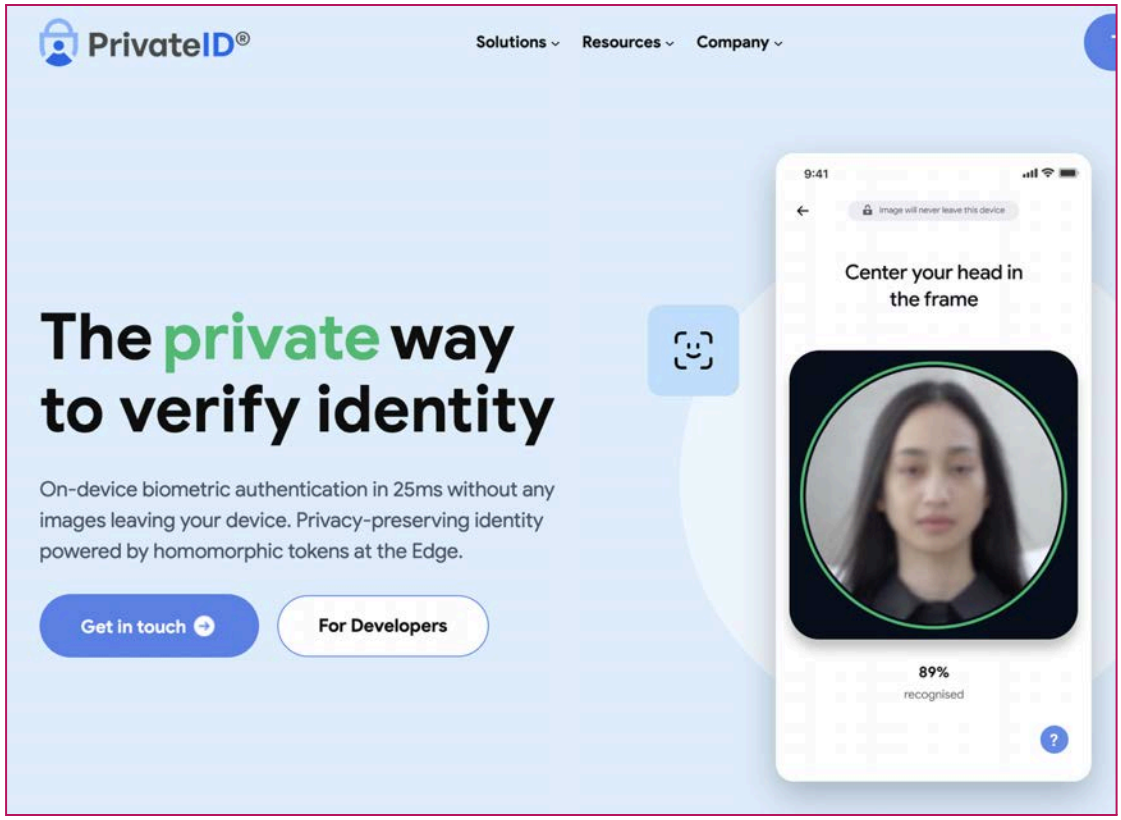

Source: PrivateID https://privateid.com/

Google uses another partner, Verifymy, for email-based age checks, which are age inference in ISO terms (Verifymy, 2026d). The user submits their email to Google, which passes it to Verifymy, where "analysis is performed on a user's digital footprint" (Verifymy, 2026b) reviewing the websites and apps where the email has previously been used, "such as a financial institution, mortgage lender or utility provider." (Verifymy, 2025). Verifymy's privacy policy for Google checks says it shares the email address with unnamed "trusted third parties, such as data aggregators" to run those queries, deletes the data once the result has been shared with Google, and refers to the method as its "age inference model" (Verifymy, 2026d).

For a credit card check, Google makes a pending payment authorisation without charging the account, and states it "will retain this data as necessary to meet legal and regulatory requirements."

In addition, Google's privacy policy reserves the right to use machine learning over a wide range of activity — content the user creates or receives, synced Chrome browsing history, search terms, location, and activity on third-party sites — to infer whether a user is likely over 18. This is age inference in ISO terms, and on the broadest signal set of any provider in scope. Without a single dedicated privacy policy for age assurance, parents and children are given only broad brushstrokes across multiple pages, making it harder to follow.

2.1.7 Twitch

Twitch uses facial age estimation provided through k-ID's orchestration service. As above, the actual facial scan is conducted by a third party, Privately SA. They say that, "Privately does not store, collect, or transmit this media data. Where voice is used in a configured workflow, that data follows the same handling principles. After the age estimation process is completed, "captured image data is immediately destroyed on your device". (Twitch, 2026).

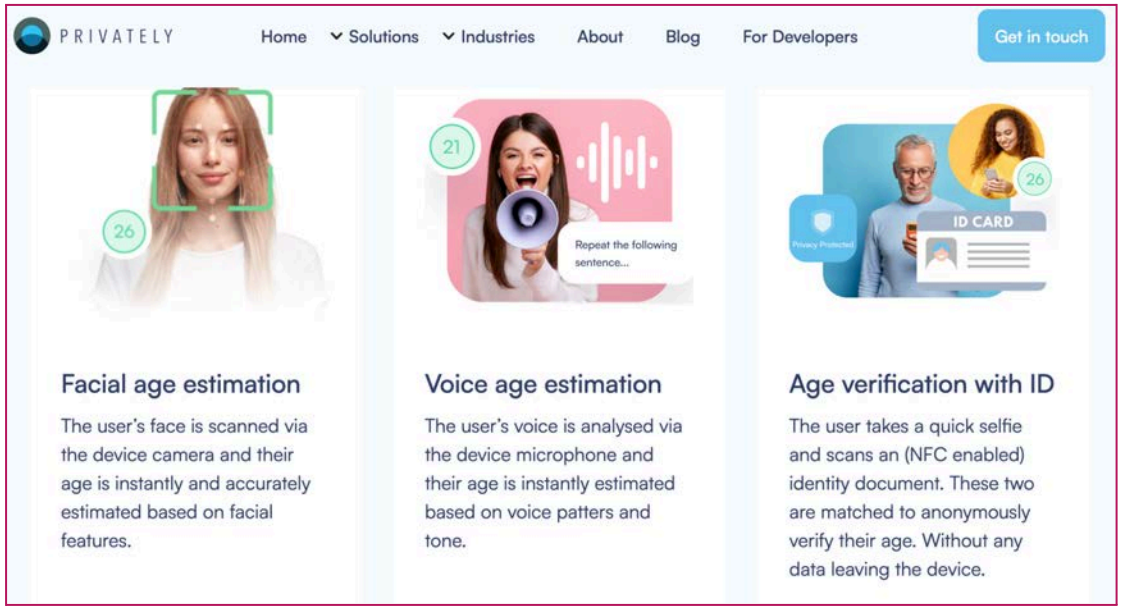

Source: Privately SA
https://www.privately.eu/solutions/online-age-verification

2.1.8 Kick

Kick also outsources to the range of options in k-ID's system,[16] offering facial age estimation, ID checks, bank-based age checks through ConnectID, credit card checks and the creation of an 'AgeKey', according to k-ID. Kick's own Help Centre documents photo ID, facial age estimation, credit card and AgeKey. Kick's Help Centre states that "not all verification methods are available to all viewers," so parents and children may not necessarily know in advance what personal data they will be expected to provide. k-ID says the methods available vary by jurisdiction, according to each country's rules and identity infrastructure.

2.2 Same provider, multiple models: Yoti app vs Yoti in-browser checks

A reusable Yoti ID account and one-off check are two different products: the Yoti Digital ID consumer wallet app (the user has a pre-verified, reusable Yoti ID account, and shares an 'over X age' attribute) versus Yoti's in-browser one-off check (no Yoti account; a single-use webcam selfie or ID upload).

[16] https://help.kick.com/en/articles/12947994-age-assurance-faqs

Meta currently uses the second of these Yoti products to support age assurance on Facebook, Instagram and Threads. Two flows are technically available, but Meta does not offer both, with materially different implications for what personal data is captured and retained.

Flow A — Yoti Digital ID app (mobile, pre-existing wallet)

A user who already has the Yoti Digital ID app, with their identity verified once in advance by Yoti against a validated ID document, can use that wallet to confirm their age when signing up to a platform. The user opens the app, consents with a PIN or biometric unlock, and the app sends only a "yes, over 16" attribute to the platform / website. No new biometric data is captured for the platform / website's check itself. At the time of writing Meta does not offer this Yoti-model.

Flow B — Yoti in-browser one-off check

A user signing up on a desktop browser without a pre-existing wallet is offered two options: submit a webcam selfie or a photograph of an identity document. At the time of writing Meta offers this method. Meta captures the image, transmits it to Yoti for processing, and receives a pass or fail. Yoti deletes its copy after the check; Meta retains its own copy for up to 30 days.

Differences between the two flows

The flows differ in three respects that matter for child-rights and data-protection analysis: where biometric data is captured, what is shared with Meta, and what is retained. In the Digital ID flow, the identity document and biometric selfie are processed once during setup, on the user's own device, and are not re-transmitted; each verification reuses that credential and shares only the "over 16 or not" attribute. In the in-browser flow, a fresh selfie or document photograph is captured every time the user creates an account; the image is uploaded to Yoti's servers, Yoti deletes its copy after the check, but Meta retains its copy for up to 30 days — and Meta's continuous AI profiling proceeds in parallel regardless of which is used. Yoti also markets a standalone age estimation product that performs no document check at any stage

Meta's support documentation describes the in-browser/in-app capture pattern (Flow B) explicitly. Yoti markets the Digital ID wallet (Flow A) as available, but Meta has not publicly offered it as a verification option for the Australian SMMA compliance. Even under the Digital ID flow, the user's age is verified against a biometrically-anchored identity held by a third-party provider in another jurisdiction. Data minimisation reduces what crosses to Meta, but it does not change the underlying architecture.

In a January 2025 discussion paper (Dawson, 2025), Yoti argues age checks should stay close to the point of service access, cautioning against shifting verification to operating-system or device level, citing privacy, competition, parental-burden and legacy-device risks.

3. Liveness checks: what they are, why they are used, and who does them

A liveness check is the step in a face-based age or identity check that tries to confirm the face in front of the camera belongs to a real, physically present human at that moment in time— rather than using a photograph taken prior, a video replayed on a second screen, a mask, or an

AI-generated image injected into the camera feed. It is distinct from the age estimation or identity match that follows it: liveness does not ask *how old is this face* or *whose face is this*, only *is this a live human, here, now*. (In the language of the standards, it is Presentation Attack Detection (PAD), tested against ISO/IEC 30107-3.

The digital proof of age cannot merely be genuine. There needs to be adequate assurance that it belongs to the person presenting it.

The distinction matters for the argument of this report because liveness is the point at which a child is most often required to *perform* — to place their face before a camera and, in some systems, to move it on command. It is the bodily act that the rest of the age-assurance chain is built upon: without a face offered to the camera, there is no template to analyse, no age to estimate, and nothing to retain.

Readers unfamiliar with this subject area need to pay attention to precision in the language used to discuss it. There are different meanings given to some of the same words which can cause confusion. The word **"passive"** is one to watch, because it is used to mean opposite things depending on its use, and the difference matters for anyone judging how intrusive a check is.

When a **regulator** uses it, "passive" usually means *no check happens to the user at all*. Ofcom describes behavioural age-inference models as passive "because they analyse users' activity on the service to infer their age, rather than users actively undergoing an age check (such as via facial age estimation or photo ID matching)" (Ofcom, 2026b). In that sense, "passive" is in fact the *most* data-hungry approach — it profiles a user's whole pattern of activity — Ofcom regards it as not highly effective.

Ofcom describes behavioural age inference models as “passive” in their July 2026 report “because they analyse users’ activity on the service to infer their age, rather than users actively undergoing an age check (such as via facial age estimation or photo ID matching).” And speaks about, “Social media services taking a passive age inference approach.”

Often when an **age-verification provider** uses it, "passive" means almost the reverse: a live facial image *is* captured, but the user is not asked to perform on-screen directed behavioural movements during an on-screen liveness check. Yoti for example also talks about a “passive” check that requires an action ‘in-the-moment’ of a live facial image to be taken, but no action-based response.

> “*Passive liveness requires users just to capture a selfie. In contrast, active liveness requires users to act on one or more random prompts, such as blinking, smiling, looking in a specific direction or moving their face closer to or further from the camera*.”
>
> (Yoti, November 2025).[17].

[17] https://www.yoti.com/blog/thoughts-from-our-ceo-november-2025/

Persona uses the term the same way — "it's 'passive' because the individual doesn't need to do anything more than take a selfie" (Persona, n.d.) — as do Regula (Regula, n.d.), Ondato (Ondato, n.d.) and Ping Identity (Ping Identity, n.d.). Here "passive" describes only the *absence of a directed gesture* — a biometric face scan is still taken every time.

A third, more loose usage treats the term "passive" as **silent or invisible**, which a reader can easily mistake for "private." Ondato calls it a check that "operates silently and invisibly … no action is required" and, “In essence, **liveness detection is the silent bouncer at the door of digital identity systems** checking not just who someone is, but whether they’re really there” (Ondato, n.d.) and iProov markets a "passive, effortless facial biometric ceremony" (iProov, 2020). But invisibility to the user is not the same as data minimisation; if anything, a check the user barely notices raises sharper questions about awareness and consent. So the same word spans a method that takes *no image but tracks everything*, and one that takes a *biometric scan the user is scarcely aware of* — precision in language matters when a term appears in policy debate. This may be particularly important **where technology emerges that a user cannot clearly identify as processing bodily data in the interaction.** For example, as used in the technology from the company used in the GOV.UK One Login fraud detection capability (GOV.UK One Login, 2025):

> “iProov's Dynamic Liveness projects subtle light patterns onto the user’s face to confirm live presence” (iProov, n.d)

In its use with the NHS, the iProov app also asks users to upload their date of birth, postcode, phone number and a photo of either their passport or driving licence during the sign-up process. A photo clipped from the ID document is supplied to IProov (The Guardian, 2021).

Ondato explains that many vendors opt for **hybrid liveness detection** such as:

- **“3D depth sensing**: Uses stereo cameras or infrared sensors to verify the three-dimensionality of the face, which 2D forgeries cannot replicate.
- **Infrared imaging**: Helps detect heat signatures or differentiate live skin from synthetic materials.
- **Contextual data analysis**: Examines behavioral signals (like device motion), network conditions, or even keystroke patterns to add contextual verification.
- **Challenge-response AI**: Dynamically adapts the level of required liveness testing based on perceived risk—for example, tightening checks if anomalies are detected.” (Ondato, n.d.)

3.1 The reasons why liveness checks are done

Providers give four reasons why their systems demand proof of “liveness”, and each is worth explaining because each also carries a cost that falls on the person undergoing the check.

First, to stop *presentation attacks* — a minor holding up a printed photo of an older sibling or parent, or replaying a video of an adult's face, to fool the computer into believing it sees an adult.

Second, to stop *injection attacks*, where software bypasses the physical camera and feeds a pre-recorded or synthetic video directly into the verification stream; providers such as Yoti, market dedicated secure image capture technology (Yoti's "SICAP"[18]) for this.

Third, to defeat *deepfakes*, by analysing signals that generative models reproduce inconsistently — the way light from the screen scatters across the curved geometry of real skin, or sub-surface indicators of a living face. Fourth, *regulatory accountability*: because the eSafety Commissioner rejects simple self-declaration as offering no integrity, and because the SMMA Act discourages platforms from mandating government-ID upload as the sole route, anonymous facial age estimation — and the liveness check that gives it integrity — has become the default alternative.

Liveness checks exist, say the vendors, to make the *rest* of the age check trustworthy. By its own logic, the more convincing the spoofing threat, the stronger the liveness demand, and the more bodily data the person must surrender to satisfy it. As AI technology improves over time, this need to prove your humanness to a machine may only escalate if this direction of travel is built into legislation. What does that mean in practice?

> *" No policy maker has asked, what does this performing-on-demand do to a person, especially a child, what are the lifetime implications for their development, shaping their understanding of interactions with the digital world and the companies behind them, and the security of their identity for life?*
>
> **— Jen Persson, Defend Digital Me (2026)**

3.2 Active and passive: the two types, and why the difference matters

There are typically two groupings of liveness check, and the difference determines what the child/user is asked to *do*. Both are in-the-moment and actively involve the user.

Passive liveness works through taking a single still selfie in real-time, on demand. The user simply looks at the camera; the system analyses skin texture, light reflection, depth and 3D geometry in the background, typically in one to five seconds, with no instruction to follow. It is the industry default, favoured because it is fast, accessible, and — providers argue — harder for deepfakes to defeat than a scripted movement.

Active liveness requires an additional one or more actions performed by the user following on-screen instructions: nodding, turning the head left and right, smiling, blinking, or reading words aloud. It takes longer (commonly ten to twenty seconds), needs the user to understand and follow the instruction, and captures a video rather than a still photograph.

[18] Yoti Secure Image Capture https://developers.yoti.com/age-estimation/secure-image-capture

Some AV providers themselves increasingly argue *against* active liveness, and the reasons they give are precisely the child-rights concerns this report raises. Yoti's own MyFace liveness paper (Yoti, 2026e) sets out three: active challenges create friction and can fail users reading instructions in a second language; they exclude users with motor or physical disabilities who cannot reliably move on command; and modern generative AI can now simulate an instructed head-turn or smile in real time, making the movement *less* secure defeating the aim of a liveness check.

Above all, friction in the process often means people can abandon the transaction which is not perceived as a threat to the stated goal of online safety, but a generic threat to the commercial success of AV providers:

"26% of US adults abandon an online purchase because they're forced through a cumbersome login process. Authentication friction isn't an annoyance; it's a revenue leak."

Source: https://realeyes.ai/blog/authentication-friction-user-experience/ January 2026

When the passive check is not confident enough, active checks are kept and triggered as a *fallback*. And the passive check is least confident exactly where the SMMA Act targets, near the age-16 threshold, where facial age estimation is known to err. The cohort the scheme most wants to catch is therefore the cohort most likely to be routed onward into a harder, more intrusive check — including, for at least one provider used on children's platforms, an *active* gesture. This 16-18 age threshold is a whole new territory for Ofcom to deal with.

We are yet to see regulators ask the question, is what performing on demand does to a person in particular the development and normalisation of children's online behaviour.

3.3 Which platforms use which type, and who actually performs the check

No platform covered by the SMMA Act performs its own face-based liveness check in isolation. Every platform either runs an in-house behavioural model, hands the face-based check to a third-party sub-processor, or both — and it is the sub-processor layer, not the platform, that determines what the child is asked to do. Email-based age checks involve no face, so liveness does not apply. Verifymy says they instead confirm that the user controls the email address with a time-limited one-time passcode, which it describes as serving a similar integrity purpose. Appendix I sets out the full detail:

In-house, profiling and always-on but non-facial: Meta, TikTok and Google each run their own continuous age-inference models over on-platform behaviour (posts, videos, interactions, and — in Google's case — Chrome history and search terms). These are not liveness checks and involve no camera moment; they are always-on profiling. They matter here only because they run *in parallel* with whatever face-based check the platform also uses, and because they are the least visible part of the system.

Face-based, via Yoti (passive): Meta's Facebook, Instagram and Threads, and TikTok, route their facial age estimation to Yoti, which uses passive liveness — a single still selfie, no head-turn, no speech. Crucially, Yoti runs this liveness at different points depending on the product: in the reusable Digital ID wallet, liveness happens **once, at issuance**, and later sign-ins reuse a token; in the in-browser Age Estimation flow — the one Meta actually uses — a fresh passive selfie is captured **at each check**.

Face-based, via Privately SA (both passive and active): Snapchat, Twitch and Kick route facial age estimation through k-ID to Privately SA, whose CEO has confirmed it uses *both* passive and active liveness, with the active path — nodding, or speaking aloud — triggered when the passive evidence is insufficient. This is the clearest instance of children on mainstream platforms being pushed into an active gesture precisely at the threshold the scheme targets. k-ID argues that an active check takes a few seconds, happens on the user's device with no image transmitted, and is less intrusive than uploading a passport. Our concern is with what repeated performance on command asks of a child, and its behavioural normalisation, which the volume of data collected does not measure.

Face-based, via Persona (active by default): Reddit, and X for its Premium ID checks, use Persona, whose default selfie configuration is an *active* head-turn — look centre, then left, then right — configurable to add smiling, blinking one eye, or holding up a hand or a number of fingers. Persona's platform also offers reverification on password recovery, multi-factor reset and before high-risk actions. Persona describes this as an account protection feature that its customers choose whether to switch on, separate from age assurance. Where a platform switches it on, the demand to perform for the camera recurs across the life of the account; we found no public statement on whether Reddit or X have done so (Persona, n.d.).

Face-based, via PrivateID (passive, with active step-up)*:* YouTube uses Private Identity for on-device facial age estimation — passive by default, with active as an optional step-up for high-risk cases — and runs liveness at both enrolment *and* each authentication.

Face-based, via Stripe (passive): X's Premium ID route uses Stripe Identity, a passive selfie matched to an ID document, with no documented gesture challenge.

No camera-based liveness: *Snapchat's and Kick's bank-based route (ConnectID) needs no camera and no face. The user instead signs in to their bank, for example with a PIN or Face ID in the CommBank app, and consents to the bank sharing an over-16 answer (CommBank, n.d.).*

The European comparison: the EU age-verification "mini-wallet" presents only a zero-knowledge "over-16/18" token to platforms, with no liveness at the point of presentation. But that token is only as clean as its *issuance* point. Where the Proof-of-Age issuer is a face-based provider rather than a national eID, the child still performs that provider's liveness check — passive selfie, or an active gesture on step-up — once, at issuance. The mini-wallet does not remove the biometric moment; it relocates it upstream and out of view.

The picture, read across the table, is that the platforms parents recognise are not the entities performing the check. The child's face is captured, and the instruction to move it is issued, by a layer of third-party providers with which the child has no relationship — and whether that instruction is a passive one-second glance or an active, filmed sequence of head-turns depends on which provider a given platform has chosen, and on how close the child is to the age the check is trying to enforce.

Full detail — provider by provider, including exact gestures, where in the process each check occurs, and how often it repeats — is described in Appendix I.

3.4 One-off issuance versus repeated at each access

The child-rights significance rests not only on liveness type but on how *often* a child is required to perform it. The evidence suggests a three-way split:

- **Liveness once, at credential issuance, then reused:** the Yoti Digital ID wallet (Flow A) — face captured once at set-up, later sign-ins share only a token. Note, however, that Meta does not enable this option although using Yoti among its meta AV choices.
- **Liveness repeated at every access / sign-up:** the Yoti in-browser flow that Meta uses (a fresh selfie each account creation), and the facial paths of Privately SA and Stripe.
- **Liveness repeated at access and, where the platform configures it, re-triggered on account events (most intrusive): PrivateID (liveness at enrolment and every authentication) and Persona (liveness at the initial check, with reverification on password recovery, MFA reset and high-risk actions available for platforms to switch on).**

Normalising children taught to “perform for the camera on demand” is therefore strongest with PrivateID, where liveness is a recurring rather than one-off demand, with Persona wherever a platform switches on reverification, and Privately SA, whose fallback forces the child into an *active* gesture (nodding, or speaking aloud) precisely when the passive check fails — which, near the age-16 threshold, is the targeted cohort.

4. The myth of “no ID required” in age-gating social media

The SMMA Act does not set out a method by which the ten platforms should prevent under-16s from holding accounts. Instead, platforms are free to choose from a range of approaches including age estimation, age verification, and age inference, ID checks, bank-based and email checks. However, every person wishing to create an account on any of the ten platforms must now submit some of their most personal data to swap for an age check. Measures that are deemed robust enough for over 18+ content may not need to be as stringent as for setting up a teenage appropriate game at 13.

For verification to be as the UK regulator Ofcom describes, “highly effective,” they nearly always require evidentiary proof of age at some point in the process, whether the age-verifying body hands those data over to the requesting platform or not, or keeps any of the data or not, or hands

over only a yes/no to the question of whether the user meets the required age threshold in a “zero knowledge proof” at the point of using the social media platform, the only way for an age verifying organisation to be able to *verify* is to have a document to verify *against* at an earlier point in the process.

4.1 What the document route shows

Of the 10 named platforms in the Australia Minimum Age for Access measures, only three of the ten platforms publish a specific list of accepted documents. Snapchat is the most transparent. TikTok, X, Reddit and Kick say only "government-issued ID", and TikTok names no sub-processor at all: a user cannot establish who will receive their passport before handing it over. Twitch offers no ID route at the point of the check at all — only an appeal after the face scan has failed.

Kick, Snapchat and Twitch each contract with k-ID, but k-ID performs no check itself. It routes the user to one of several sub-processors (its published list has ten entries). k-ID receives the result of the check, which for an ID check can include the user’s date of birth or an age range, and passes only a pass or fail to the platform. A child who has read Snapchat's help page has been told the name "k-ID"; but the entity that actually receives their passport is Veratad Technologies LLC in New Jersey, which appears nowhere in Snapchat's user-facing text. k-ID states it stores no images and its policy links to a sub-processor list on a separate Vanta-hosted Trust Center page, which loads only in a browser and does not appear in archived copies.

Several routes commonly presented as document-free are not. A bank-based check returns a threshold assertion only because the bank has already verified the account holder's identity from documents at account opening. The bank did that for its own purposes, including anti-money-laundering rules, and ConnectID enables the customer to reuse it without showing a document again (Australian Payments Plus, personal communication, October 2026). The age answer still rests on that earlier document check. The same is true of the EU mini-wallet and of Yoti's Digital ID wallet: each is a reusable assertion whose issuance rested on an ID document, ‘seen’ at that time. An AgeKey created from an ID scan or bank check carries the same upstream document step; one created from facial age estimation carries none, and the platform relying on it need not know which. What these products remove is repetition of showing the documentation step, not the document step itself at the start.

Only Privately SA's on-device facial estimation and Verifymy's email check have no document ID step at any point in the chain, doing estimation or inference, not verification.

Second, every document route boils down to the similar three artefacts: a passport, a driver's licence, or a state or national identity card. The only addition anywhere is YouTube's reference to a state Proof of Age card. A person under 16 does not, as a matter of course, hold any of the three in their own name, and Australia has no accessible fourth. The document route is therefore not merely harder for the targeted cohort, it is substantially unavailable to them, and the routing to a face scan follows from the process design rather than from any choice the child makes. This compounds the exclusion for children in out-of-home care, refugee and asylum-seeker children, Indigenous children, and children whose documents do not match their name or appearance.

Third, the verification, estimation and inference routes are not alternatives of equal kind, and the platforms do not distinguish them. A "menu of options" that presents a passport scan alongside a selfie, an email address and a credit card is not offering a choice between comparable things; it is offering one route that establishes age, several that estimate it from the body, and several that infer it from a proxy, under a single label, with the guessing routes carrying the lower friction. That asymmetry is what makes the least-friction path the chosen path, and it is why the consent relied upon is manufactured rather than freely given.

Education about what these systems require has not been a mandatory part of their adoption, and leaves an enormous hole in understanding the impact of the implementation in Australia that other countries should avoid repeating. Given the myriad of schemes across different platforms, it will likely prove near impossible for parents and children to understand how their data is shared and processed. They may not be aware that it is not the social media platform itself undertaking the age check, but rather a third-party sub-processor. Some platforms are vague about where data is sent, referring to “vendors” or “partners” and use more than one. Many offer a confusing variety of provider choice while mandating use of age assurance. Australia’s Privacy Commissioner, Carly Kind, has noted that, “age assurance solutions are in many cases fragmented across multiple providers” (OAIC, 2026).[19]

Personal data may be passed on multiple times across these processes to different degrees, to different companies in different parts of the world. While many of these companies stress their commitment to privacy, every time sensitive personal data is transferred to a new processor it creates a new risk of hacking or exploitation, and when it comes to biometric data, that means high-risk for a lifetime that, in our view, cannot be mitigated to an acceptable level for children.

[19] OAIC guidance on age assurance technologies (March 17, 2026) https://www.oaic.gov.au/news/media-centre/privacy-commissioner-publishes-new-guidance-to-ensure-proportionate-age-assurance-as-a-gateway-to-access-online-experiences

4.2 What the age verification process can require step-by-step

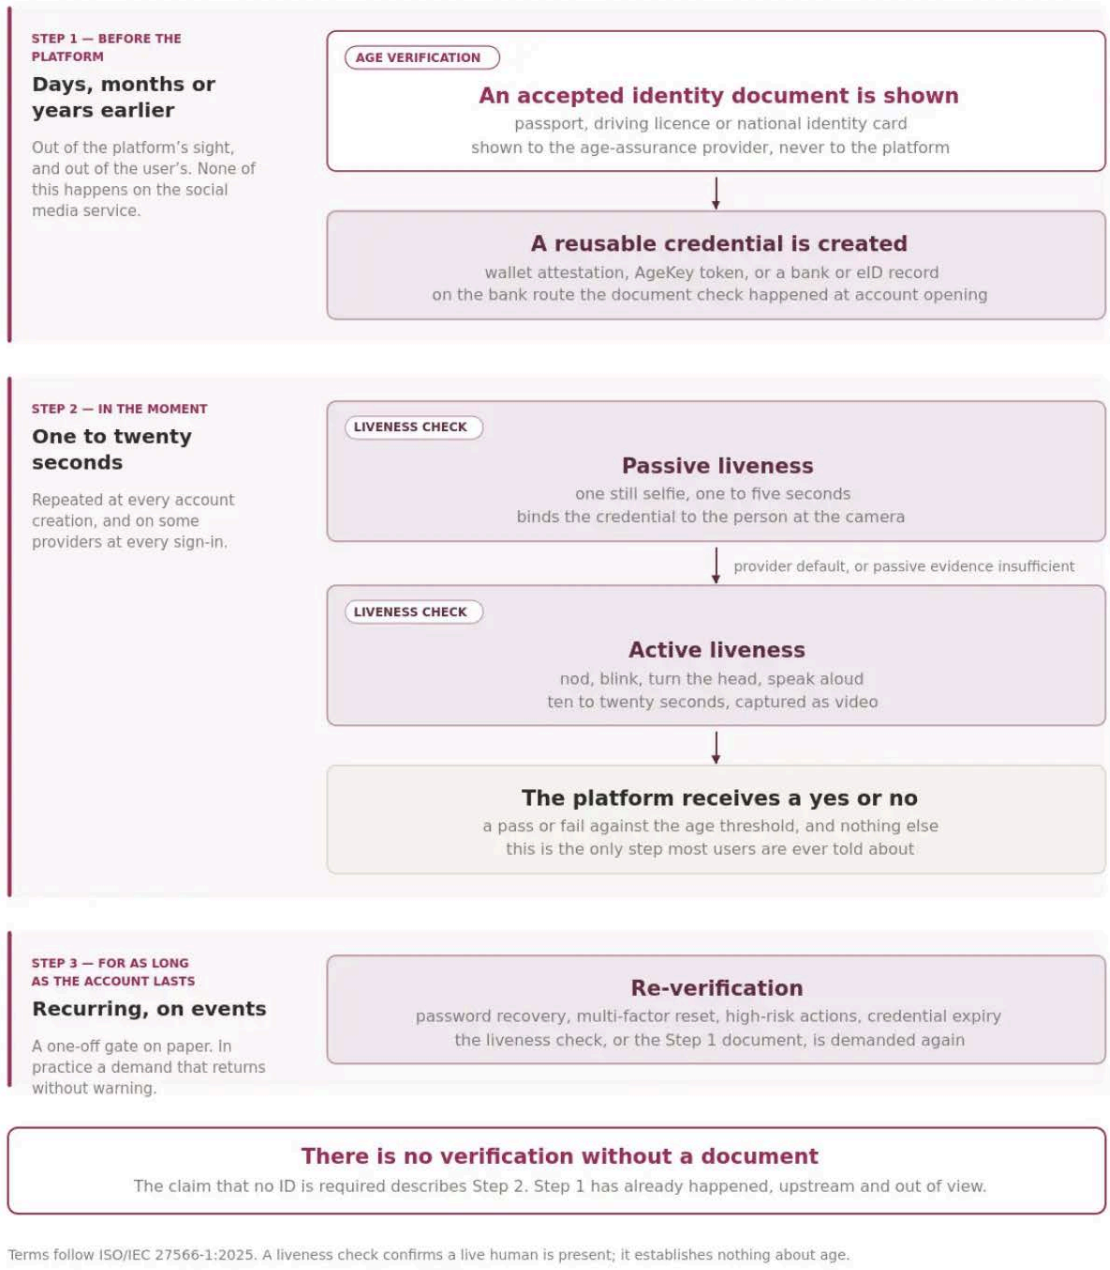

5. Biometric data and a legal challenge

To a child or parent uploading a video or photo selfie, the proposition is straightforward, a company behind age check is scanning their face. They are unlikely to read the result through the technical lens of whether the system "uniquely identifies" them, performs estimation rather than verification, or stores a template rather than an image. People recognise a face scan using their body to do so, and expect it to be treated as sensitive. The distinction that allows a provider to assert its product is "not biometric" for one regulatory purpose or at a point in time, in only part

of the process, has no obvious distinction to a user— and to many child-rights advocates, like us, it is splitting hairs to obtain an exemption from the more stringent laws around processing of biometric data.

Some leading companies maintain that data captured through their checks are not used to train their models in a way that would engage biometric protections, because subsequent processing is characterised as "research and development" or internal product-improvement rather than a new biometric processing activity. If a child does not offer their face to the camera, the product training does not happen. Their bodily data is processed in order to create the training data. A regime in which a child must submit bodily data simply to access mainstream online platforms — with no meaningful alternative — cannot easily be reconciled with the standard set by the AEPD, the GDPR, Convention 108+, or General Comment No. 25.

The fact that there is a variety of choices of which face-tool to use, even including those that do not require biometric data processing or could be replaced by a visit to a Post Office for example, does not mean the user *chooses* to use the tool. If age checks are mandatory and the face-processing is the process with the least friction, and the incentive is to do so or lose access, the "choice" is manufactured and not freely-given. It is not based on consent.

On July 11, 2026, Jonathan Haidt, a world leading champion of minimum age restrictions to social media, tweeted his own views on this saying, "Authoritarian countries like China can mandate identity verification via government-issued ID. But democracies like Australia must "start more softly" even if "the first step did not bring them to the final destination."[20]

We and others have warned that this intent has no democratic mandate. Age verification that requires identity verification via government issued ID is the only logical natural conclusion of a model that will not tolerate a method of age assurance that social media users of any age can more easily work around than evidentiary proof, and age verification providers need some sort of state-accepted ID from which they, the AV provider, obtained the proof that they go on to offer as a yes/no credential or "zero-knowledge proof". Despite not sharing details of the ID itself onwards to the platform, still requires the ID to be shown earlier in the process shared with the AV provider.

It is increasingly clear that it will become mandatory to use government-issued or state approved digital ID over time, and that starting with a range of options is a soft start to make it more acceptable. Jurisdictions without national [children's] digital ID, will seek to move towards one as a result. The UK 2026 consultation on national digital ID already proposed this use-case for social media (Cabinet Office, 2026). This is, as Haidt describes, the model of "authoritarian countries."

There is a tipping point where the volume of use makes little difference whether it began as optional or mandatory, as the outcomes become the same. The CEO of Yoti wrote in March 2026, that, "in around 3 years time, when more than 20 million UK citizens own a certified digital ID, it

[20] https://defenddigitalme.org/2026/07/12/this-was-the-plan-all-along-social-media-minimum-age-restrictions-socialmediaban/

will become much harder for fraudsters to use real ID details to impersonate an intended victim. Businesses will be able to check the interaction with the real owner of those ID details."[21] The architecture being normalised today, marketed (only) for children's social media access, is the same architecture that will become normalised and mandatory, or at least *expected* to use, for a much wider range of daily activities for their foreseeable lifetime.

5.1 What counts as biometric data

The face-based methods documented above involve processing biometric data — facial images, liveness signals, and in some products voice samples — from every user who undergoes a face-based check, including children. Some providers, platforms and policymakers argue that on-device processing, immediate deletion of source images, and binary "pass/fail" outputs make these systems privacy-preserving and place them outside the scope of special-category data protections. In the EU and the UK, Article 9(1) GDPR establishes a general prohibition on processing special-category data, which explicitly includes "biometric data for the purpose of uniquely identifying a natural person"; processing is permitted only where an Article 9(2) exemption applies, such as explicit consent.

5.2 One company, two regulators: the UK ICO and Spain's AEPD

In its 2022 Regulatory Sandbox report, the UK ICO looked at Yoti's facial age *estimation* — the tool that estimates an age band from a face without matching it to any identity. The ICO assumed, without reaching a formal 'decision', that this involves biometric data, but concluded it is not special-category data under Article 9, because the tool categorises a person by age rather than uniquely identifying them (ICO, 2022). The finding was however contested by some among the data protection community, who believe that a more accurate reading of data protection law does not require identification in the sense of being able to name an individual, but rather to single them out from others.

The Spanish AEPD assessed a different Yoti product. Its 2025 resolution (EXP202317887) concerned *the Yoti Digital ID app*, in which a selfie becomes a stored biometric template later compared 1:1 when a user changes a PIN or recovers an account. This was not the same age-estimation product as assessed by the ICO in the UK.

The AEPD held this is special-category data and fined Yoti €500,000 for processing it without a valid Article 9 basis, rejecting Yoti's argument that 1:1 authentication is not identification: comparing a live scan against a stored template is unique identification whether the match is 1:1 or 1:N. €200,000 for invalid consent (a click-through privacy policy and a pre-ticked opt-in to biometric use for research) and €250,000 for excessive retention (biometric templates kept for the life of the account plus three years), took the total to €950,000 (≈US$1.1m) (Burt, 2026). Aggravating factors included the involvement of minors, international data transfers (Yoti operates

[21] Tombs, R. (2026). Thoughts from our CEO: Spanish regulator AEPD fining Yoti https://web.archive.org/web/20260711134928/https://www.yoti.com/blog/thoughts-from-our-ceo-march-2026/

a manual-verification Security Centre in India[22], where “trained document checkers who have demonstrated an above-average ability to recognise people conduct the online identification process. The Security Operations Centre personnel check the images of the user-submitted ID documents against a global database of ID documents”),[23] and the scale of processing.

Part 1 — the pre-ticked R&D consent box (Article 7 finding), Section VI "Obligación incumplida. Condiciones del consentimiento":

> The AEPD found that consent for processing users' data for internal research and development ("investigación interna"/I+D) was obtained through a pre-ticked box, which cannot constitute the clear affirmative act required by Article 4(11) GDPR. As the resolution states, "los consentimientos realizados mediante casilla declarativas activadas por defecto no son el resultado de un acto afirmativo y claro que refleje una manifestación de voluntad libre, específica, informada e inequívoca" (AEPD, 2026, Exp. EXP202317887, Fundamento de Derecho VI).

Part 2 — that this default-on setting applied to children’s biometric data, Antecedentes plus the Article 76.2(f) aggravating factor:

> The same default-on configuration applied to children’s biometric data: the resolution records that R&D processing covered "los datos de menores de 14 años, ya que para tratar datos biométricos de menores solo hay que marcar una declaración de ser mayor de edad" — a minor needed only to tick a declaration of being over the threshold age (AEPD, 2026, Exp. EXP202317887, Antecedentes). The AEPD applied the effect on minors' rights as an aggravating factor under Article 76.2(f) LOPDGDD to each of the three sanctions (AEPD, 2026, Fundamentos de Derecho V, VIII, XI).

The reuse of data is bundled with the consent to use the app and set to default in, from which users can opt out only later in the process. Yoti accepts that in its app, the face scan is “special category biometric data” and rests its defence on explicit consent and user choice — the freedom to use the app or “a different service” presented as consent— and Yoti is appealing to the Spanish High Court (Yoti, 2026c) and has withdrawn that product from the Spanish market pending the decision, saying in a blogpost, “We hope to republish the Yoti ID app in the app stores following the conclusion of the appeals process with the Spanish High Court, after which we hope app updates will become available again to existing users (Yoti, 2026d).” The app was removed from Spanish app stores on September 10, 2026. Yoti described biometric authentication as the cornerstone of the app’s protections and withdrew the app rather than change it to align with the AEPD’s decision (Yoti, 2026d). In July 2026 the AEPD argued that alternatives to the EUDI Wallet and to its biometric authentication must exist, so that the wallet does not become a source of digital exclusion (Biometric Update, 2026b). The Age Verification Providers Association warned

[22] https://committees.parliament.uk/writtenevidence/90067/pdf/
[23] https://committees.parliament.uk/writtenevidence/96702/html/

that the AEPD's position could undermine the EUDI Wallet scheme and asked the EDPB for guidance on how Article 9 applies to biometric authentication in digital identity and age verification apps (Age Verification Providers Association, 2026) which may seem an extraordinary position to object to rather than meet the standards of legal compliance expected by the Spanish regulator.

Privacy International had flagged the underlying problems years earlier. Its 2019 concern was not the core age check but the secondary use of customers' data (Privacy International, 2019). Yoti had fed users' selfies, plus the date of birth and gender from their uploaded ID documents — more than 72,000 people by mid-2019, with skin tone tagged by researchers — into the dataset training its Age Scan model. PI argued this used data for a purpose users would not reasonably expect, that the opt-out was buried instead of an opt-in consent process, and that once data has trained a model it cannot be extracted. The AEPD decision found the same, and it appears Yoti would rather withdraw its product than redesign to meet the obligations the regulator enforced upon.

5.3 The same argument, scaled up: Worldcoin and "proof of personhood"

Arguments about whether AV companies process special-category biometric data matter for the governance of all biometric data going forward. The claim is often made that scanning a person's face escapes Article 9 as long as the system is not, at that moment, able to identify them. The dispute looks technical, but the future of bodily-data protection rests on it, because the same argument underpins the most far-reaching biometric identity project yet attempted: Worldcoin, run by Tools for Humanity and co-founded by OpenAI's Sam Altman.

Its "Orb" device scans people's irises to issue a "World ID" — a credential meant to prove the holder is a unique living human ("proof of personhood") without, the company says, knowing who they are. Its legal defence is, in substance, Yoti's: because the iris data is hashed and "anonymised once it leaves the Orb," and because no name is attached when the credential is checked, the company argues it is not processing special-category biometric data and may not even be a data controller (Tools for Humanity, 2026). So far, regulators have refused that claim as firmly as the AEPD refused it for Yoti. Germany's Bavarian authority ordered it to delete its iris data and grant users an unrestricted right to erasure (BayLDA, 2024); Spain's AEPD issued a warning before a planned Barcelona relaunch; Colombia ordered an immediate and permanent shutdown and deletion; the Philippines' regulator issued a cease-and-desist. The thread running through all of them is the same: "we don't identify you at each scan" is no answer to the fact that a biometric was captured and processed (WP29, 2007).

The BayLDA found that Worldcoin obtained consent only for collecting the iris code and the *active* comparison at registration, but not for the permanent storage and *passive* comparison that follows — the larger, more intrusive, and longer-lasting processing. As the decision states, the consent given "covers only the case of 'to calculate derivatives of the Image data (like the Iris Code) and actively compare it against our database'," so "the storage of the iris code after registration and the passive comparison is not covered by this" (BayLDA, 2024, Worldcoin

Foundation decision, paras. 507–508). With no valid Article 9(2)(a) consent and no overriding legitimate interest under Article 6(1)(f), the processing was held unlawful under Articles 9(1) and 6(1) GDPR, triggering an erasure obligation under Article 17(1)(d).

This is why the Yoti decisions matter far beyond age checks. Both cases ask the same legal question, packaged differently: when does processing a person's biological features fall outside special-category protection simply because the operator says it is or is not at a particular point in time of a process or the nuance of the data processing? If the answer becomes that "estimation, not identification" or "hashed and anonymised" is accepted as a route that avoids Article 9 obligations, it will not merely exempt one product or provider but opens up the processing of highly sensitive data to a range of providers that say they need not obtain an individuals' consent.

We must ask whether this is what the drafters of data protection and privacy law intended and whether it meets the reasonable expectations of the public about what protections are afforded to bodily data and why, and whether those limitations are robust under these emerging definitions.

5.4 Children's biometric data and its lawful basis for processing in Europe

The split matters most for children. Under the GDPR (and the UK GDPR), Australia's Privacy Act amendments, and Council of Europe Convention 108+, biometric data used to identify a person is a special category requiring heightened protection. Children's biometric data is afforded particular emphasis in Convention 108+ guidance and in the UN Committee on the Rights of the Child's General Comment No. 25, which calls on states to ensure that biometric processing of children is necessary, proportionate, and accompanied by strong safeguards (UNCRC, 2021).

Defenders of age assurance as consensual point to choice. They claim a child can usually pick from several verification methods, so no one is 'forced' into a face scan. But where age checks are mandatory and the facial route is the path of least friction — regardless of choices on offer 'how' to age verify— there is no freely-given consensual choice whether to verify. Any decision is an agreement built on coercive compliance, not consent.

5.5 Why this is high risk and irreversible

Biometrics cannot be reset. A breached password can be replaced; a scanned iris or face cannot. As the Bavarian regulator put it, once a company holds your iris pattern you can never be anonymous to it again. The harm is therefore permanent for everyone who has already handed their data over, and the exemption being argued for is both global and forward-looking — arriving precisely as the infrastructure to demand a biometric at the door of the digital world is being built. There is no path back that restores the privacy of anyone who has already given up their biometric data.

A March 2026 written question to the European Commission (E-001251/2026)[24] noting Yoti is used by Meta, TikTok and pornography sites, recites the AEPD's three findings, and asks how the

[24] https://www.europarl.europa.eu/doceo/document/E-10-2026-001251_EN.html

Commission will enforce GDPR compliance on age-verification providers, whether it will bar them from transferring data outside the EU, and what can be done about the irreversibility of biometric identity theft (Griset, 2026). A further question (E-001659/2026) followed in April.[25]

6. The industry: standards, certification and economics

6.1 The certification picture

A consistent line in provider marketing is some claim of regulatory approval or certification. The reality is that the standards landscape is still very new, and largely voluntary (Open Rights Group, 2026). The most important recent development is considered to be ISO/IEC 27566-1:2025, the first global standard to try to lock down definitions for terms like "age assurance", "age verification" and "age estimation", which providers and regulators have until now used interchangeably. Veratad has been among the first to be certified to it; broader take-up is still being worked through. The standard sits alongside two longer-running schemes operated by the Age Check Certification Scheme (ACCS): ACCS 1:2020 for age estimation and ACCS 2:2021 for data protection and privacy, which issue Level 1 (Basic), 2 (Effective) and 3 (Highly Effective) certifications (Allen, 2025). Privately SA and Private Identity LLC are both certified at Level 3 for facial age estimation. Yoti's facial age estimation has been certified for the UK's Challenge 25 retail policy at a stated accuracy of "at least 98.89%."

Separately, the US National Institute of Standards and Technology (NIST) runs a regular benchmarking programme for facial age estimation; Yoti, Privately and Private Identity have all submitted to it. Regulator-level recognitions also exist: the German youth-protection body KJM, the German self-regulatory body FSM, the UK's Ofcom, and (historically) the BBFC have all accepted particular methods. That said, certification is not an audit of any specific deployment. The ACCS expressly states that inclusion in the Age Assurance Technology Trial documents "does not constitute endorsement, certification or approval," and that not all claims have been independently verified. None of these schemes is mandatory under the SMMA Act: a platform can meet the "reasonable steps" threshold using an uncertified provider, provided it can otherwise justify its choice. The certification picture is therefore more a marketing differentiator than a regulatory floor. Since September 2026, eSafety can require information and documents from third parties such as age-assurance and app-store providers, and can summon people to give evidence, to test what platforms claim about their compliance (Wells, 2026).

Open Rights Group has sounded the alarm about the incentives the UK Online Safety Act creates for adopting cheap technology providers at the expense of users' privacy and security, and providers whose terms allow them to repurpose collected data for marketing. Its March 2026 blog gives case studies on the Discord data breach and on Persona. The Discord breach, in which around 70,000 users' government ID images were exposed, took place at 5CA, a third-party customer service provider handling Discord's age-related appeals; Persona was not involved (Discord, 2025).

[25] ibid

Separately however, researchers found Persona's frontend code publicly exposed on a non-production subdomain. The code showed the platform could run up to "269 individual verification checks", including screening "against 14 categories of adverse media from terrorism to espionage" (Open Rights Group, 2026; Arntz, 2026). Persona says the exposure involved no customer data or backend systems, that these checks belong to separate identity verification and financial crime products for regulated customers rather than to its facial age estimation, and that no customer uses all 269 of them. It has published a post-incident review (Persona, 2026c). Persona's policy states that it does not sell personal data or use it for marketing or model training (Persona, n.d.). The distinction Persona draws between age assurance and identity verification matters but how many are applied for age and how many for other purposes? All checks run on the same platform, and which checks apply is decided by the company buying the service or by Persona by default, but not by the person using it for an age check in front of the screen (section 6.2).

6.2 Consolidation, lobbying and the KYC convergence ("faces, please")

The Australian rollout has coincided with, and accelerated, a sharp consolidation of the global age and identity assurance market, projected to reach US$18.6 billion by 2026 (LexisNexis Risk Solutions, 2025). Large players have been buying their way in: in February 2025, LexisNexis Risk Solutions (part of RELX) closed its acquisition of Australian-founded IDVerse; Entrust acquired Onfido in 2024 (Entrust, 2024); and k-ID acquired Opale.io, the French developer of the AgeKey credential, in November 2025 and relaunched it as the separate OpenAge Initiative. Persona was valued at around US$2 billion in its April 2025 Series D, and Yoti reported revenue growth of 62% in 2025 (to £29 million) on the back of 56% the year before. These are venture-capital and private-equity-backed businesses whose growth is closely linked to the spread of statutory age and identity checks — a connection their own leadership acknowledges, with Yoti's CEO attributing anticipated UK revenue growth (Tombs, 2025a) to companies deploying facial age estimation to comply with the Online Safety Act (Yoti, 2026a).

The commercial and regulatory interests of these companies are closely aligned. The same companies that supply the technology actively participate in the policy discussions that mandate it: Yoti, Privately and k-ID all took part in Australia's Age Assurance Technology Trial and have been vocal in UK, EU and Australian consultations, while the Age Verification Providers Association lobbies on the industry's behalf (k-ID is not a member, though the OpenAge Initiative it launched is). Several are vertically integrated with platforms or banks. ConnectID, for example, is operated by Australian Payments Plus, whose 25 shareholders include the major Australian banks that also act as its identity providers — so each check uses a bank's data through a service the banks part-own (Australian Payments Plus, DDM communication, October 2026).

The OAIC's own worked examples include a "creator" flow in which "an ID and selfie are captured for AML/CTF onboarding," and "the entity also needs to know that the creator is over 16 years to comply with the SMMA obligation" (OAIC, 2025)— one biometric-and-document capture discharging anti-money-laundering identity verification and age assurance at the same moment.

That is the direction of travel. Once that fusion is normalised, "prove you are over 16" and "prove who you are to a state-verifiable standard" become the same act — and the companies consolidating this market (Persona, Yoti, k-ID, Stripe, ConnectID) are positioned to be the gatekeepers of both. ConnectID says platforms receive only an over-16 answer through its service, and none of the bank's anti-money-laundering data (Australian Payments Plus, personal communication, October 2026).

The old surveillance fear had a shorthand: "papers, please." That image now undersells what is being built. The paper age feared papers, please; the AI age brings faces, please — the document dissolves and the face itself becomes the credential. Your face already sits in the chip of your passport; a digital state ID simply strips away the paper around it, leaving the biometric, held by the state and verifiable on demand. A single capture — asserted, of course, to be "not biometric data" — becomes the key to every transaction and entitlement. As China's street-level identification shows, infrastructure that grants access can equally withdraw it, score it, or simply watch. A face check to open a social media app is not a contained measure but the first and most normalised instance of a credential that, once accepted, follows a person everywhere, for life.

6.3 Making the market: the underlying demands for age verification and KYC

If the movement toward social media bans has not produced child safety, it is clear that it has produced a market. Jonathan Haidt tweeted as much in July 2026, in his post welcoming the spread of age verification now that Australia had created it.[26]

That outcome is largely absent from public debate. Creating a market and delivering online safety are different aims, but they have been packaged and presented as exclusively the latter in recent policy discussion.

The Philippines shows where this leads. Senate Bill 2424, the Child Online Safety and Protection Act, passed second reading in the Senate at the end of September 2026. It would bar under-16s from social media and require platforms to use independent third-party age verification providers, which writes the age-check industry into statute. It would fine those providers, along with platforms and government agencies, between P1 million and P20 million per offence, and it would fine parents who help a child get round the restrictions between P5,000 and P50,000 (Daily Tribune, 2026). Australia's SMMA Act places no penalty on parents. The House of Representatives is moving in parallel and goes further on gaming. Its Committee on the Welfare of Children has approved a substitute Children's Social Media and Online Gaming Safety Act, consolidating 41 bills, which would bar under-13s from the social and interactive features of online gaming platforms as well as from social media accounts (Journal News, 2026). Lawmakers have also proposed a classification board for online games, and a rule that social media and gaming platforms must establish a presence in the Philippines before they can serve Filipino users (GMA News, 2026). The two chambers have yet to agree on a minimum age.

The UK ambition was always broader than policing social media. Age gating was proposed for

[26] https://defenddigitalme.org/2026/07/12/this-was-the-plan-all-along-social-media-minimum-age-restrictions-socialmediaban/

gaming as early as 2019, and industry has since talked up its use on the high street and in vending machines (Yoti, 2021; Hadley, 2022; FinGo, 2022). What platforms, the Home Office, and GCHQ share in their verification of Children Online ("VoCO") project, is a common interest in knowing the customer — and its outcome is the end of anonymity online as it currently exists.

GCHQ and the Home Office provided the government with the groundwork in 2019 (DSIT, 2020) for building a lobbying message, and expanding future testing to assist in making age assurance a reality. DCMS sought, "to boost innovation and promote technology solutions in the age assurance space. Insights and research resulting from VoCO will inform ongoing work in this area."

Collective international efforts to shape norms were seen by the UK as vital to helping to prevent cyberspace from being used as a global platform for serious crimes, including fraud and sexual abuse of children (HM Government, 2021). In June 2021, the G7 leaders' Carbis Bay communiqué endorsed the Digital Ministers' Internet Safety Principles, which set out common approaches to improving online safety (G7, 2021, para. 34; G7 Digital and Technology Ministers, 2021).

In October 2020, the "Five Eyes" (UK, US, Australia, New Zealand, Canada) plus India and Japan issued a joint statement calling on technology companies to design encrypted services so that law enforcement can obtain readable content under lawful authorisation, framing end-to-end encryption as a barrier to detecting child sexual abuse material (US Department of Justice, 2020).

They completely missed the fact that their vision and technical models ignore any detrimental or developmental behavioural effects on children. Their VoCO 2030 vision of newly normalised behaviour that their age-gating the Internet brings about, creates a world in which children have been programmed to perform on camera, on demand, in much the same way as predators demand of them. It creates new and additional risks very much like those same risks that the Home Office and GCHQ described as the VoCO aims they are protecting children from.

Making age-gating mandatory for social media means children (*and adults* to tell them apart from a child) need ID, clearly speeding up the normalisation of adoption by all of society. Requiring ID to go online, means the end of anonymity online. And there are tipping points, the world is on the way towards (Tombs, 2025b), after which it is perceived the majority of society has accepted a new practice, after which the minority need no longer be catered for, or accounted for.

The *Digital identity sectoral analysis report 2026 (*DSIT, 2026a) from the UK DSIT Office for Digital Identities and Attributes, points to some of the shared market growth. *"In October 2025, GBG, a UK-headquartered global identity firm, acquired DataTools Pty Ltd, a leading provider of address validation and data quality solutions in Australia and New Zealand, for AUD $16 million (£7.9 million), extending its identity data capabilities in the region*." If that name rings a bell with readers, it might be a reminder of when in 2020 gambling clients of GBG had been caught accessing information from the UK Department for Education national Learner Records Service of 28 million users personal details, enabled unlawfully via Trustopia (Defend Digital Me, 2020).
GBG wrote then on its website,

> "W*ithout credit reference data to draw on, businesses are having a hard time verifying millennials' identities and turning away potentially good customers.*" "*Millennials are taking longer to reach the 'traditional' markers of adulthood, such as buying their first house or car. Some are opting out altogether.* ***As a result, conventional identity verification methods are falling short."*** (GBG, 2020)

Social media bans meaning the market gets to start the youth earlier and more reliably than otherwise possible, might therefore appear very fortunate timing for such age verifying companies. In 2020, GBG delighted in reusing school children's records as an ID database to be able to increase their 'uplift' in getting new younger gamblers signed up. Many might consider that justification for expanding age verification data use and access, the very opposite of enhancing children's online safety.

Incode in its acquisition announcement of AuthenticID in 2025, bringing together two leaders in digital identity said it intended *"to form a global AI powerhouse built for the next era of secure, scalable IDV—a market projected to grow to $116 billion by 2027." (Incode, 2025)*

The EU Special Panel on Child Online Safety chairs handed over its report to President von der Leyen on July 13th (Fegert and Melchior, 2026) and on 17 September 2026 the Commission published its proposal for an EU KIDS Act (see section 9.1). While outside the EU, UK policymakers are no doubt paying attention to this and it has wider global implications too.

Many countries have centralised, digitally accessible registers of births and deaths. The UK does not, but its March 2026 consultation on national digital ID proposed that one could underpin age-gated access to social media (Cabinet Office, 2026). In July 2026 the new Prime Minister announced that the digital ID scheme would be scrapped (Government Technology, 2026). Absent a national ID, the UK is instead likely to end up with commercial age-verification systems enrolling children by exact date of birth, on proprietary terms. Either route ought to carry open requirements as a minimum: a published specification, an auditable enrolment process, and public disclosure of what is recorded, for how long, and by whom — so that the register of children's ages is not a closed commercial asset whose contents and accuracy cannot be independently checked.

Systems like this, can become non-tariff barriers to new competitors who wish to enter a market and existing market players can try to keep competitors out using various means.

Safecast, in 2021 evidence to the UK Parliament, said that, *"For example, if Facebook or YouTube were to establish an age verification enrolment system based upon the exact age of a Facebook or YouTube user, this could be used for the commercial benefit of Facebook or YouTube respectively and their "walled gardens" of commercial services." (Safecast, 2021).*

The providers and businesses that are required to age-gate either have to pay to run their own system in-house, build and maintain it, or they outsource to a third party that charges them for every new user and/or per every age-check each time.

There are clearly a range of per-check-prices on the market but most companies do not publish these transparently and instead websites tend to read, 'come and talk to us about pricing'. In July 2026, Ofcom reported they found the average per-check cost incurred by the service providers does not appear high. "It generally sits towards the lower end of the range of £0.05 to £0.30 we provided in our April 2025 Statement, where we estimated per-check cost figures in the range of £0.05 to £0.30." (Ofcom, 2026b). But found that social media providers said photo ID matching per-check costs were above the upper end of that range. Ofcom's figures also show the scale of the market. More than 69 million age checks were completed in six months across the 32 services it examined, a 23-fold increase, and it says the true UK total is likely to be "materially higher". The median cost per check was about £0.06, while some providers running their own systems reported global costs in the tens to low hundreds of millions of pounds a year. Reusable digital ID, often presented as the privacy-preserving answer, made up only 2% of completed checks on the pornography services Ofcom examined, and one provider dropped it after a week because users abandoned the process (Ofcom, 2026b, paras. 3.2, 4.12, 7.14, 7.20).

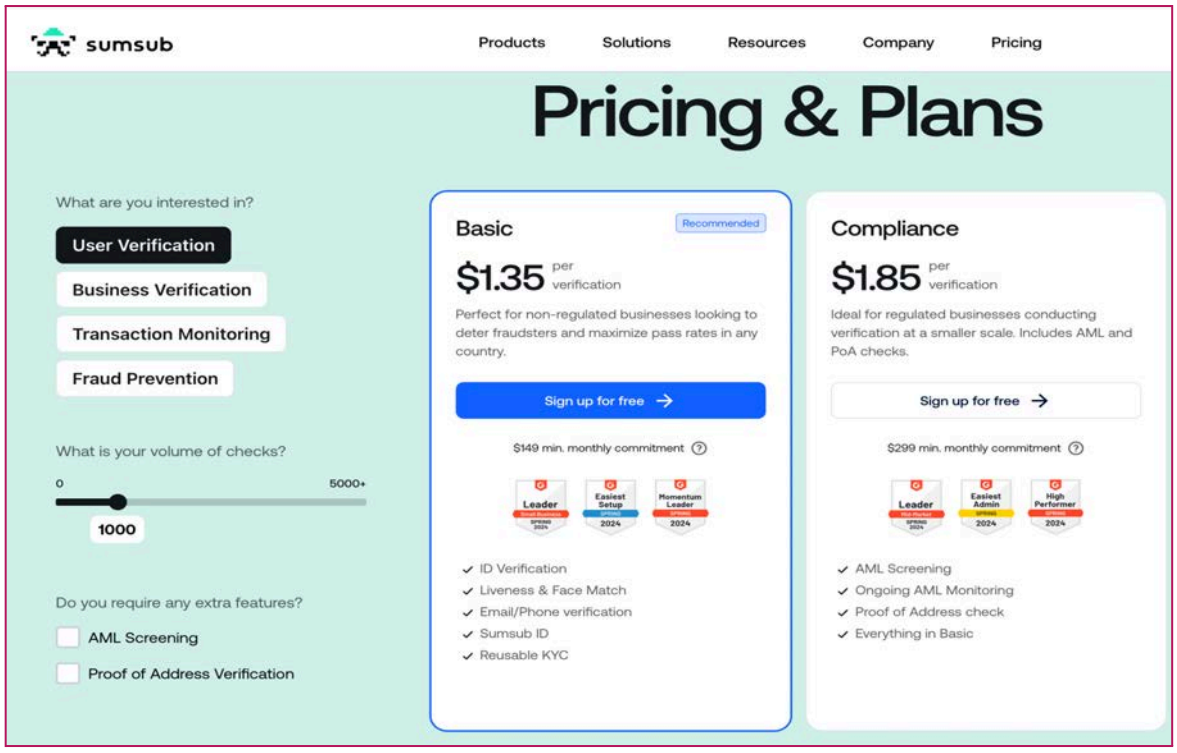

Source: Sumsub pricing https://sumsub.com/pricing/

The example from the company Sumsub above, even suggests as much as $1 above that, per check. Ofcom went on to assess these costs as manageable. "*Even where per-check costs could be a concern, commercial arrangements such as fixed-fee or subscription-based models may reduce*

the extent to which costs increase directly with each additional check. This could mitigate cost pressures for some service providers, where those arrangements are available." (ibid, 7.28, p.51).

The changing economic models of online business are also recognised as a driver creating new demands for age verification behind "know your customer" KYC models. "Buy now pay later", a relatively recent commercial model, demands reliable fraud detection. Reed Taussig, pioneer in AV, in a 2022 interview, said that, "*50% of the Australian population has tried buy now pay later… 30% of the UK population.*" He went on to articulate why it matters to industry in particular when customers are young.

> *"A lot of the people that are going to use buy now pay later are what are called thin-file customers. People like my daughter who don't really have a big credit file, don't own a credit card never has bought anything on credit. Those are called thin file customers and trying to figure out who those people are, is a big challenge."* (Intellicheck, 2022).

In summary, age assurance as a cost to the platform (monetary) or cost to the user (the privacy risk) is usually discussed as a single event when a document is uploaded, or a face scanned. That is assumed to stop at the moment of capture, and misses what the age signal does to everything else a platform already holds about you.

An age check does not only answer "over 16 or not." It hands the platform a verified age for an account it was previously only guessing at. Sometimes this is not only about the individual (a child) but their relations to others (the adult who approves their account, or with whom they are gaming together). Watch time, interaction graph, location, scroll behaviour, viewing habits, attention signals, pause and re-watch — all of it becomes more valuable once it is connected to a known age, because everything else can be read with more confidence. The biometric might be deleted the same day; but the commercial gain of that check does not disappear with it. The information that is retained is not the face but the certainty. Minimising personal data at the point of capture does nothing about that transfer of retained knowledge.

"Zero-knowledge-proof" is therefore a misnomer and misleading. The receiving platform might not know which given name you have in human terms, but their machine-learned knowledge of your behaviour expands.

This is why the in-house profiling at section 2.1 matters more than its low profile suggests. Meta, TikTok and Google run their age inference not alongside the third-party check but with the benefit of it, and Meta says openly that the same material trains its generative AI models. A duty to know a user's age has given the largest platforms a mandated reason to build richer profiles of everyone, under the banner of safety. Age assurance is not a compliance cost sitting outside the surveillance business model. It feeds it.

6.4 Zero-knowledge-proofs – a model in search of a market?

A common flippant reply to privacy concerns is that the problem has already been solved: *"Of course age checks are privacy preserving. Have you never heard of zero-knowledge proofs?*" The reply mistakes one step for the whole process. A zero-knowledge proof (ZKP) lets one party prove that a statement is true, such as "this person is over 16", without revealing anything beyond the fact that it is true. The EDPB recommends that approaches such as zero-knowledge proofs be made available where age assurance may involve high risks to privacy (EDPB, 2025, p. 9, para. 34). *But. What a ZKP protects is only in the moment of presentation.* What the platform learns when the user shows a credential in the moment of social media use for example, is less than the age verification provider needed to know to create the credential in the first place.

Before that moment, something has to establish the user's age so that a credential can be issued. In the systems examined in this report, that step involves showing a state-accepted identity document, a bank record, a face scan with a liveness check, or a database lookup to an age-verification provider or credential issuer (sections 4 and 9.2). That provider sees the full document, the face or the account details, and how it handles, retains and shares them is outside anything the proof itself guarantees. The same model is already sold for know-your-customer, anti-money-laundering and sanctions compliance, where a firm uses a zero-knowledge proof to show that it meets its obligations without holding the client's identity documents itself, while the verification company that checked those documents has seen them in full. After presentation, the platform still holds everything it already knew about the account, now linked to a confirmed age (section 6.3). Around the proof sit the device, the wallet app, the network and the issuer, each of which can generate metadata about when and where a credential was used.

Privacy therefore has to be assessed end to end, at every stage where data is collected, processed, retained or linked: enrolment, issuance, storage, presentation, revocation, recovery and re-verification. A process that is anonymous at the final step can still be identifying at the first. The French Conseil constitutionnel reached the same point in law when it found that a ban on minors requires every person to prove their age, and that a statute silent on how that proof is collected and governed had not provided the necessary guarantees (section 7.7). Zero-knowledge proofs address one of those stages, and the expert assessments below examine how much of the rest they leave untouched.

The cryptographic community's assessment of ZKPs for age verification is more mixed and complex than vendor and platform marketing suggests.

Olivier Blazy (École Polytechnique), who developed an open-source double-anonymous ZKP age-verification proof of concept with the CNIL and PEReN, has nonetheless argued that age verification is a poor policy even when ZKPs are used, underscoring that a sound cryptographic primitive does not cure the socio-technical harms of a mandate (Blazy, 2023). This distinction between the mathematics and the deployment recurs throughout expert commentary.

As a technology, zero-knowledge proof can limit what a single verifier / platform learns. However for privacy and security, a complete framework is only as safe as its weakest link, and as such ZKP

is not a magic fix. Getting privacy wrong at the initialisation of the process, will create a deep mistrust from the general public, and just democratise even further the use of avoidance mechanisms. VPN being the easiest, readily available but not the only one.

Steven Bellovin, a computer scientist who has worked on privacy-preserving credential schemes, reaches the same conclusion from inside the field. He argues that it is technically straightforward to build a privacy-preserving, credential-based age verification system for the web, using anonymous credentials that cannot be linked to each other or to the person who holds them, but that the legal, economic and social obstacles are formidable and possibly insurmountable (Bellovin, 2025). Someone has to check documents such as birth certificates or passports before issuing a credential, which is hardest in countries without a national ID card, such as the UK, and difficult for people who are poor, elderly, rural or homeless. Shared family computers and public library terminals raise further privacy problems. Regulating the issuers implies audit logs that are likely to include identifying personal data such as names, addresses, ages and the documents presented. And because a credential used only for age checks is easily shared, Bellovin argues it would need to be used for other purposes to deter sharing, which creates pressure for scope creep and a single point at which a government could revoke a person's access to many services at once. He notes that these human and procedural obstacles apply to any credential-based age verification scheme, privacy-preserving or not.

Sofía Celi, Kyle Den Hartog, and Hamed Haddadi, writing for Brave Research, caution that even semantically "private" range proofs may leak information depending on how they are constructed or interpreted, that many proposals remain fragile with respect to soundness and composability, and that ZKPs alone cannot resolve the broader problems of content categorisation, revocation, and centralised credentialing. Their core argument is that ZKPs cannot simply be "dropped in" to solve the socio-technical problems of online trust and access control; rather, they must be built into wider architectures that protect user autonomy, limit central control points, and reckon with the realities of exclusion, censorship, and profiling (Celi et al., 2025).

The Electronic Frontier Foundation similarly notes that ZKPs do nothing to constrain verifier abuse, over-asking, or repeatedly demanding age over time, and that they cannot undo the data-broker profiles that already exist (Electronic Frontier Foundation, 2025b) or that may be expanded as a result of the ID checks required further up the chain, before the point of the ZKP 'ask and respond'.

Empirical scrutiny reinforces the caution. A large-scale measurement study reverse-engineered Yoti, the dominant US provider, and reported that its age-verification flow collects high-entropy device metadata and shares data with less-visible "fourth parties," concluding that deployed systems can fall short of their privacy claims (Minocha et al., 2026). However, Yoti disputes the study's characterisation, and following the universities' investigation the specific allegation that Yoti shares users' facial images with third parties was removed from the associated press article; Yoti maintains that facial-image data is processed solely to estimate age and deleted immediately (Tombs, 2026).

At a design level, the trade-offs are structural rather than incidental: even privacy-preserving credential schemes based on ZK-SNARKs carry a trusted-setup requirement and witness-generation and revocation burdens that constrain real-world deployment (Kumara et al., 2023).

The broader position advanced by researchers including in March 2026 from Georgia Tech, is that identity and verification infrastructures are themselves security and privacy liabilities, attractive targets for breach, and vectors for scope and function creep and state demands, independent of the elegance of the cryptography layered on top (Minocha et al., 2026).

The resulting consensus is not that ZKPs are worthless to individuals but that they are a narrow tool wrongly presented as a systemic remedy for "privacy preserving" age verification. They can limit what a single verifier learns in one transaction and what the platform requesting the check receives, but does not solve the questions of coercion, exclusion, centralisation, revocation, and mandate scope that determine whether age assurance protects children, or simply expands surveillance of everyone in the digital environment into which age-gating is deployed.

6.5 Pricing the check: what age assurance costs, and what users will accept

Age-assurance providers rarely publish their prices. A 2026 rate card from one seen by the authors, shows how the service is sold. Compliance is priced in layers. An organisation-wide tier, from free up to US$25,000 a month, sets the markets and regulatory regimes covered. Each app, game or website then needs its own subscription, from about US$500 to about US$50,000 a month depending on audience size, with a quota of checks included. Further checks cost US$0.10 each on pay-as-you-go, falling to US$0.01 each at 100 million checks a year. A low price per check does not make the whole system affordable. Integration, appeals, security, support and maintenance all add to the cost, and the Commission's own analysis of the EU KIDS Act relies on third-party services and the EU age verification app to keep costs down while acknowledging that the evidence for some of its integration estimates is limited (Gamers Voice, 2026).

The price list also ranks the methods. Server-side facial age estimation is included by default. On-device estimation, which keeps the face image on the user's own device, is included only for the gaming sector and is otherwise charged at US$0.16 a check. Email-based checks cost US$0.32 and credit card checks US$0.35. National identity-number and digital-ID checks are priced country by country, from US$0.15 to US$0.99, and the Australian bank-account route costs US$0.85, more than eight times the standard list price. Going on these prices, the cheapest route for a platform is a server-side scan of the user's face.

Two further lines show where the market is heading. Human review of failed checks is sold as a premium service, from about US$9,000 a month for up to 6,000 appeals, so whether a child or adult wrongly refused by an automated check can reach a person depends on what the platform has chosen to buy and is an expensive "optional extra" in the service. Ofcom has found that many services have no appeals process at all, and only six of the 32 it examined could say how many

appeals they had upheld (Ofcom, 2026b, para. 3.17). The same infrastructure is also sold per user, from US$0.48 a year falling to US$0.11 at scale, to marketplaces opening paid services to young users, with parents approving purchases. Here, age assurance is priced as the gateway to children's spending as well as a barrier to their access. Large platforms also gain the most. Games sold through major storefronts can rely on the platform's own age checks and account-level guardian systems, which makes independent distribution harder, and the games publisher Epic now even sells its own age-verification service, KWS (Gamers Voice, 2026). Age checks therefore displace the commercialisation of children's data to new companies: verification providers, orchestrators, banks and device makers now see children's faces, documents and sign-ups.

Evidence on how users respond to these methods points the same way. In a deceptive experiment with 1,635 US adults who believed they had to verify their age to view R-rated film clips, 99% completed a checkbox self-declaration and 86% an email-based estimate, while 51% completed AI facial age estimation and between 18% and 27% a government-ID upload; ID with an active liveness check had the lowest completion of all, at 18% (Lin et al., 2026). Statements that data would be used only for age verification and deleted afterwards made no significant difference to completion. Between 53% and 79% of those who did verify by a method other than the checkbox said they were uncomfortable doing so, only 0.86% opened the privacy policy, and 45% believed the website running the study, and not only the verifier, would receive their data. The authors treat these figures as an upper bound, because participants trusted a university study more than they would a commercial provider. Ofcom's survey of adults who had stopped using pornography sites rather than prove their age found the same distrust: 76% did not want to share their personal information, 75% were worried it would not be protected, and 71% did not trust the third-party providers that run the checks (Ofcom, 2026b, para. 7.60).

The price list and the user evidence together suggest what drives the choice of method in practice. Providers price by volume, platforms pay for completed checks, and users respond to friction and trust. No part of that chain rewards the method that collects the least data. Privacy promises, on which most regulation relies, did not change behaviour in the one controlled test.

6.6 Proposed alternatives: credit reference checks and operating-system age signals

Two alternatives are often put forward as less intrusive than document uploads or face scans: checking age against credit reference data, and moving the check into the device's operating system. Both change where age data is collected and who holds it, and each needs the end-to-end assessment set out in section 6.4.

Credit reference checks compare the details a user enters with third-party databases, most commonly those of credit reference agencies, along with the electoral roll and other records (ICO, 2024). They are long established in online gambling and other 18+ services, and the credit reference agency Equifax has urged Ofcom to add them to its list of highly effective age assurance methods (Equifax, n.d.). Ofcom's January 2025 guidance does not list them among the methods it considers capable of being highly effective (Lin et al., 2026).

Their limitations stem from what the data is. A credit file can show only that someone is an adult, so in a scheme with a threshold of 15 or 16 a credit check cannot confirm the age of any child the scheme is meant to cover. It can only admit adults who have a financial history, leaving young adults, people new to the country and others with thin or no credit files to fall back on another method. The check relies on financial profiles built for lending decisions, and it leaves a soft search on the person's credit file, a record of which company looked them up (Equifax, n.d.). The industry's interest in reaching young people earlier is already documented in section 6.3 that describes an identity company's mis-use of a state education data set to verify young adults with little credit history.

Operating-system age signals are the other, moving the check from each app onto the device. California's Digital Age Assurance Act, which takes effect on 1 January 2027, requires operating system providers to ask for a user's age or birth date at account set-up and to pass an age-bracket signal (under 13, 13 to 15, 16 to 17, or 18 and over) to any app developer that requests it through a real-time interface (Alston & Bird, 2025). Developers must treat the signal as the primary indicator of a user's age and may not share it with third parties for unrelated purposes. Apple, Google and Meta all supported the bill (Reclaim The Net, 2025). The EU KIDS Act proposal would likewise require operating systems that hold a compliant age signal to share it with in-scope services, with the user's consent (European Commission, 2026c, Art. 29(6)). In Australia, k-ID already accepts age signals from Apple, Google Play, Xbox and Meta Horizon alongside, or in place of, a self-declared age (section 2.1.2).

The UK has gone further, combining device-level age checks with controls on content. On 8 June 2026 the government said it wants Apple and Google to block nudity across the whole device by default, so that the block can be lifted only through age assurance, which means every adult who wants it lifted must verify their age. The government said it gave the companies three months to act before legislating (The Stack, 2026). Apple had already introduced age checks for all UK iPhone and iPad users in March 2026, asking new users to verify with a credit card or by scanning an identity document. Debit cards are not accepted, and anyone who cannot or will not prove they are 18 has Apple's web content filter and Communication Safety switched on automatically (Computing, 2026; TechRepublic, 2026). On 8 September 2026, Lisa Nandy committed the government in parliamentary debate to primary legislation requiring major tech platforms to build in device-level protections for children (Nandy, 2026), and the new Prime Minister, Andy Burnham, has said he wants protections "built into the device and switched on by default" (ISPreview, 2026). Child-protection organisations that welcomed the plan included the NSPCC, and the Internet Watch Foundation said device-level protections could help reduce the creation of new child sexual abuse material (Family Law, 2026). The National Crime Agency said such protections were needed to stop harmful content being created and shared in the first place (The Stack, 2026).

By contrast, Big Brother Watch called it a "backdoor digital ID requirement", and Signal described it as "a dystopian combination of age verification and content scanning". Google said it was working with UK partners to find "effective, privacy-preserving solutions". Apple had itself dropped

plans to scan devices for child abuse images, warning in 2023 that such scanning could create a "dystopian dragnet" (The Stack, 2026). In response to the September deadline, Apple also said it had shared with UK officials plans to strengthen Communication Safety further (AppleWorld.Today, 2026).

The UK model differs from California's in two respects. It rests on verified adult identity, including a credit card route, and blocking images before they are sent implies inspecting content on the device itself (The Next Web, 2026). It therefore joins the two alternatives in this section and adds content scanning to both. In the United States, a federal Digital Age Assurance Act introduced in July 2026 would create device-level age signals while prohibiting age verification by government-issued ID, biometric data or facial age estimation (Free Speech Coalition, 2026).

The approach has attractions for some. One check per device could replace repeated checks by every service, and the EDPB favours architectures based on user-held data and secure local processing on the device (EDPB, 2025, p. 9, para. 34). But its limitations may be as significant. The California model rests on a declared age, so the signal is an assertion; any stronger check at device level would mean the operating system provider verifying the identity of everyone who sets up a device. Age data for every app would be concentrated in a handful of operating system providers. A standing age signal available to any app on request has commercial value to advertisers and data brokers as well as to child-safety compliance.

But most simply, such a device-level based approach may be impractical. Devices are often shared within families, and a child using a parent's phone or a school device inherits the signal of whoever set it up. And age-gating moves below the platform into the device itself, so that the device vouches for its user's age to every app it runs (Reclaim The Net, 2025). Yoti, which sells checks at the point of service, has argued *against* moving verification to the device (Dawson, 2025), which would undermine its business model, so the debate between these approaches is as much a commercial as political one. Meta argues the reverse, describing checks at app store and operating system level as, "*centralized, consistent, and privacy-preserving*" (Meta, 2026a).

The UK public is not persuaded that the location of a check changes much. Ipsos found support for an under-16 ban fell to about half if everyone had to upload ID or a credit card, whether the data would be held by platforms (50%) or by device makers such as Apple or Google (51%) (Ipsos, 2026).

Neither alternative answers the questions we ask in this report. Credit checks admit adults by reference to their financial lives and cannot identify a child at all. Operating-system signals spread one declared or verified age across every app on a device and concentrate it with the largest technology companies. Device-level signals do not work reliably on shared devices, and checks tied to a user's apparent location can be worked around with a virtual private network.

We must therefore ask the question whether such technological measures can ever "work" to provide children's online safety at all? Have we even a consensus on how to define what online safety means and how it is measured?

7. Is the ban working?

Evidence from Australia found that facial "*age estimation is a popular and effective form of age assurance*" (eSafety, 2026a) in particular, for confirming that a person is considerably younger or older than a threshold. However, it "*is known to have higher error rates for children near the age threshold of 16 years*" (eSafety, 2026a). Platforms that pushed self-declared 14- and 15-year-olds to facial age estimation "*would have been aware that many of them would likely receive a false 16+ outcome*."

eSafety's recommended "layered/waterfall" approach pushes more processing in a second check such as a bank- or ID-based check, plus age-inference (behavioural profiling) models to flag suspected under-16s. This level of uncertainty means more intrusive checks for those children.
And in the UK, Ofcom reported in July 2026,[27] "*Serious doubts whether age inference methods currently used by some popular social media sites and apps are fit for purpose, with investigation launched into TikTok.*" Ofcom told social media services that rely on age inference to switch to methods listed in its guidance as highly effective without delay, and said it would give an update on the TikTok investigation in October 2026, alongside a rapid assessment for Parliament of what highly effective checks look like for over-16s (Ofcom, 2026c).

Ofcom's own statutory report on age assurance, published the same month, shows how little the checks have changed for children. Before the duties came into force, 70% of 11 to 17-year-olds recalled seeing at least one kind of harmful content in the previous four weeks. Four months after, the figure was 73%. Only 11% of children said they could no longer use something they had used before, and 16% said they saw less upsetting content. Ofcom concludes that age assurance, together with services' other safety measures, is "not yet delivering the intended outcome" of protecting children from harmful content (Ofcom, 2026b, paras. 5.19, 5.21, 5.28).

Checks also push children elsewhere. Almost half (47%) of the pornography services visited by children aged 8 to 14 had no age checks at all. Estimated daily visitors to adult services fell by about a third, while some services without checks gained traffic, in one case rising from nearly 10,000 to around 25,000 visitors a day. On dating apps, 12% of online 15 to 17-year-olds were still visiting the three most popular in December 2025, the same share as before the duties began (Ofcom, 2026b, paras. 4.37, 4.40, 6.10, 7.32).

27

https://www.ofcom.org.uk/online-safety/protecting-children/age-checks-helping-make-online-experiences-safer-for-uk-children-but-job-not-done-and-tech-industry-must-act-to-strengthen-protections

Ofcom says circumvention "appears low", but notes that children may not admit it to researchers. Its own figures show that 16% of 13 to 17-year-olds know how to get round age checks and 5% say they have done so in the past year (Ofcom, 2026b, paras. 7.72, 7.76). Low self-reported circumvention tells us little about safety when exposure to harm has not fallen and children are moving to services with no checks at all.

Ofcom opens its report by welcoming "a generational shift towards a more age-aware internet" (Ofcom, 2026b, p. 3). That shift is the normalisation this report warns about. Ofcom also found that services showed "strong awareness of privacy requirements", but that rests on the services' own answers to its information notices, with every service's name redacted (Ofcom, 2026b, p. 48). Our review of the providers' published policies in Appendices B, H and I gives a weaker picture. And Ofcom did not assess age checks used to enforce minimum ages of 13 or 16, so nothing in its report yet speaks to whether an under-16 ban can work (Ofcom, 2026b, paras. 2.17, 2.19).

7.1 The evidence from Australia, six months in

Making children safe online is an outcome. Age assurance is a process. This report shows what happens when the technique is treated as though it were the outcome, and what has been built in the space between those two things in Australia.

In March 2026 eSafety confirmed that, "*Five investigations are currently underway: eSafety is actively investigating potential non-compliance in relation to Facebook, Instagram, Snapchat, TikTok and YouTube.*"

Prior to the SMMA Act implementation in December 2025, the Australian government commissioned the Age Check Certification Scheme to assess whether age assurance of social media users was possible. The trial found it would be technically feasible but noted: "We did not find a single ubiquitous solution that would suit all use cases, nor did we find solutions that were guaranteed to be effective in all deployments." Some methods, such as facial age estimation, were found less accurate for children.

And its findings were contested. Electronic Frontiers Australia (EFA) issued a statement claiming that the findings of the Australian Age Assurance Technology Trial were "strong on hype, rhetoric and difficult to reconcile with the evidence." EFA served on the Stakeholder Advisory Board for the trial, the only representative on the 20-person board from a digital rights organisation. Tim Levy of children's safety tech provider Qoria had previously resigned from the board, claiming a disconnect between the early results and findings.

Australia may be seen as a test-case for how effective age assurance techniques are at keeping them off social media. Surveys of young Australians have found that over 60% of under-16s who previously had accounts have retained access (Molly Rose Foundation, 2026; Bursztyn, 2026). The Molly Rose Foundation found that "children still using restricted sites say that it was 'easy' to circumvent the ban" and that around a quarter had been "successfully able to get around an age check", while others had family members open accounts for them (Molly Rose Foundation, 2026).

Some under-16s told the BBC they set up fake accounts before the deadline (Livingstone, H., 2026). The Foundation draws a parallel with Korea's 2011 Youth Protection Act — a midnight gaming "shutdown" ("Cinderella") law for under-16s, discontinued in 2021–22, which research (Lee et al., 2017) found "increased juvenile Internet hours, instead of reducing them."

In contrast with the Australian AAT trial, 438 scientists and researchers from 32 countries have publicly called[28] for a pause on age-assurance policies, saying that there is no proven benefit to children's wellbeing, online checks far exceed offline equivalents, they are trivially circumvented, and there is a flawed threat model. They say that age estimation/inference is unreliable and invasive, there is no global trusted infrastructure and that among other unintended consequences, age verification reduces rather than increases safety, erodes privacy and entrenches discrimination, while centralising power and threatening the open internet design. In a joint statement, security and privacy experts warned that "age-assurance checks are easy to bypass, as evidenced by current deployments being circumvented using VPNs, bought or borrowed credentials, or props or AI-based tools (e.g., deepfakes or AI-generated profiles)" (da Costa, 2026) and that such checks require Internet-wide trust infrastructures that do not yet exist. In the UK, children are already circumventing checks using fake IDs, makeup, and AI-generated selfies (Williams, 2025). In one widely copied method in July 2025, Discord users fooled the facial age estimation offered through k-ID with an image of a video-game character; k-ID says the workaround was patched within hours.

In March 2026, the Australian Children's eSafety Commissioner ("eSafety") reported platforms removed or restricted about 4.7 million age-restricted accounts by mid-December 2025 and a further 300,000+ in Q1 2026, and that parent-reported account ownership across all 8–15s fell from 49.7% to 31.3% (eSafety, 2026a). That whole-population drop, however, sits alongside a high retention rate among children who already had accounts. Of those who held accounts before the cut-off, around seven in ten still held them on Facebook, Instagram, Snapchat and TikTok (about half on YouTube), and the single most common reason for retention (66.8%) was that the platform had not yet asked the child to verify their age (eSafety, 2026a, p.14). Under-16s continued to report harm with no discernible drop. Five platforms are under investigation (Facebook, Instagram, Snapchat, TikTok and YouTube) (eSafety, 2026a, p.7). eSafety's first longitudinal evaluation report, published on 31 July 2026, found account ownership among 10 to 15-year-olds fell from 52.4% to 42.1% three months after the ban began. Most children who held accounts before the ban had kept them or opened new ones, and platforms' failure to implement effective age assurance was the main reason given. Parents' awareness of their children's social media use also fell, particularly for girls and for children aged 10 to 12 (eSafety, 2026b). (We have not found a published outcome to the five investigations at the time of writing).

Pushing people, especially young people away from visible services in fact sees a rise in usage of "alternative" platforms with no care for the law, or user protection, creating other risks for them. In December 2025 before the SMMA came into effect in Australia, teenagers were "flocking to

[28] https://csa-scientist-open-letter.org/ageverif-Feb2026

alternative social media platforms not caught by Australia's new minimum age laws, propelling three (Lemon8 a sister company to TikTok also owned by ByteDance, Yope and Coverstar) to the top of Apple's App Store on the first day of the ban" (AFR, 2025).

It also appears that practice has not met in spirit nor on paper what Australian parliamentarians were told would be a privacy-preserving feature in the adoption of legislation. The SMMA consent must be freely given, voluntary, informed, current, specific and unambiguous – this is an elevated requirement that precludes platforms from seeking consent *through preselected settings or opt-outs*. The third-party age verification data processing operates on an opt out basis for retention for their product development.

7.2 Were the 2024 Child Rights Taskforce fears realised?

In October 2024 the Australian Child Rights Taskforce open letter — signed by over 140 academics, experts and civil-society organisations — argued a ban was "too blunt an instrument" (Australian Child Rights Taskforce, 2024). Measured against the March 2026 evidence, most warnings have been substantially borne out within the first three months.

It concluded, *"that to date, social media companies have failed to enforce their minimum age requirements properly. Together with our report findings, this shows that social media companies will need to take a more effective approach to ensure that any social media restrictions for under 16s work in practice."*

The pre-social media ban warning	Did it happen?	Evidence seen by March 2026
There are not yet effective techniques for age assurance.	Largely yes	Facial age estimation is unreliable near 16; under-16s obtained false 16+ outcomes, re-tried until passing, or were never asked; ~7 in 10 retained accounts on major platforms.
Privacy concerns remain.	Yes, escalating	Practice is drifting toward ID/bank verification and behavioural inference; the OAIC had to issue a detailed safeguards regime; the trial found over-retention of biometric/document data by some providers.
Parental consent does not fix unsafe products; parents are not the experts but are still asked to decide.	Partly	SMMA dropped parental consent, but parents reporting under-16 accounts were obstructed (e.g. asked for a lawyer's letter) and platforms often took no action.
A ban creates more risks for children who still use platforms.	Yes / too early	Substantial retention and continued harm reporting with no discernible drop; a circumvention culture. eSafety claims no significant migration to fringe services, which other studies dispute — unresolved.
Bans affect children's rights to access and participation.	Realised	Millions of accounts removed or restricted; access curtailed for those who comply.
Systemic safety/privacy standards on platforms would be the better route.	Partly vindicated	The 25 March 2026 rule pivots the definition toward harmful design features (recommender and endless-feed/feedback/time-limited 'logged-in' features) — a more systemic, design-based lens, though still inside a ban.

Table 7.1 — The 2024 Australian Child Rights Taskforce warnings against the March 2026 evidence (Australian Child Rights Taskforce, 2024; eSafety, 2026a).

7.3 Guidance vs practice: the minimise-or-escalate ratchet

The Australian case is the closest real-world test of the model the UK is consulting on, and it carries two lessons. First, the gap between a privacy regulator's minimisation ideal and a safety regulator's 'reasonable steps' escalation logic does not resolve in favour of minimisation: in practice it ratchets toward biometrics, ID/bank verification and behavioural profiling, even where a 'choice' of method is formally offered. Second, the specific privacy and efficacy warnings made in 2024 were largely borne out within three months.

A fuller side-by-side of the four framings — what the SMMA Act and its Explanatory Memorandum set out to do, what the Australian Child Rights Taskforce warned in 2024, what the OAIC's privacy guidance asks of platforms, and what the early outcome evidence from both the Molly Rose Foundation poll and eSafety's own compliance update actually shows — is set out in Appendix G (Table G). Read across each dimension, the pattern is consistent: the legislation's protective intent and the regulator's minimisation guidance both pull against an escalation logic in practice, while the 2024 efficacy and privacy warnings have largely been borne out within the scheme's first three months.

Issue	OAIC privacy guidance (minimise)	eSafety 'reasonable steps' + observed practice (escalate)
Front-line method	Least-intrusive method that works; prefer binary "16+ yes/no"; biometrics are sensitive.	Facial age estimation (biometric) is the popular default; error-prone near 16.
When a check is 'not enough'	Escalate to more intrusive handling only where genuinely necessary and justified.	Escalate to a "more robust" bank- or ID-based check when a user is flagged or low-confidence.
Profiling / inference	Most intrusive route; prefer non-sensitive, point-in-time signals; avoid always-on monitoring and long-lived profiles.	Age-inference models encouraged to flag likely under-16 accounts on an ongoing basis.
Data retention	Destroy inputs immediately; keep only a minimal ring-fenced artefact; no reuse for analytics/advertising.	Trial found a minority of providers over-retaining full biometric/document data 'just in case'.
Burden on parents	Transparency and accessible processes; minimise data demands.	Some platforms demanded a lawyer's letter to accept a parent's report of an under-16 account. Some accept parental authorisation "consent" and their own ID-check to age-verify their child.[29]

Table 7.2 — Guidance versus practice, at a glance.

Three forces push the 13–17 cohort toward biometrics and then toward ID/bank verification and behavioural profiling. Estimation is unreliable exactly where it matters: near the 16 threshold, facial age estimation errs, so the cohort the scheme targets is precisely the cohort routed onward to a "more robust" method. "Resistant to circumvention" drives escalation: eSafety treats repeated low-confidence passes and easy re-tries as poor practice and expects a second, harder check once a user is suspected of being under 16. And inference fills the gaps: to catch accounts

[29] https://k-id.zendesk.com/hc/en-us/articles/44154566493339-How-do-I-verify-my-age-with-the-help-of-a-Parent-or-Guardian

that pass at sign-up, platforms are pushed toward age-inference models — the route the OAIC itself flags as most privacy-intrusive. The net effect is a ratchet: in practice the equilibrium drifts toward more biometrics, more ID/bank verification, and more behavioural inference — the opposite of data minimisation — even though a 'choice' of method is formally preserved.

Ofcom's July 2026 report confirms the more intrusive direction of travel in the UK. It tells services to add liveness detection to every facial age estimation and photo ID check, to use a challenge age (most often set at 21), and to consider repeating age checks after the first one when triggered by "suspicious behaviour or new log-ins". Services that rely on age inference should replace or supplement it with methods Ofcom considers highly effective, which in practice means face scans and ID. Ofcom also wants search engines, operating systems and app stores to play a part, and will publish a report on app stores by January 2027 (Ofcom, 2026b, pp. 4–5, 65–66; paras. 4.47, A2.7). Each step asks more of the child or adult in front of the camera.

7.4 The risks to children

Although well-intentioned, age verification checks place children at risk online. Social media companies already have a predatory approach to collecting children's personal data, and in reality provide scant detail for children and parents about where their data will end up if they attempt an age check. Hacks, leaks, poor security practices and differing standards across jurisdictions all place children's data at risk. Age and ID verification requirements have only been mandated in recent years, yet there have already been multiple high-profile incidents:

- Discord reported that 70,000 users had their government IDs exposed on the internet when a third-party customer service provider handling its age-related appeals was breached (Discord, 2025).
- In a TechCrunch investigation, journalists linked users of social media platform 'TeaOnHer' to their driving licences within 10 minutes (Whittaker, 2025).
- 404 Media found that one of TikTok's third-party age verification partners, AU10TIX, left the photos and driving licences of those who had used its service exposed for almost a year (Cox, 2024).

Anyone — child or adult — attempting to verify their age now faces a dizzying range of systems, each demanding various forms of sensitive personal data. Tracing the data streams involves reading complex privacy policies from multiple companies in different jurisdictions. Although the SMMA Act seeks to make children safe online, it instead risks exposing them to serious privacy violations and normalises regular, intrusive identity checks to access the internet.

The Meta/Instagram incident reported in June 2026 (Neowin, 2026) illustrated the mechanism how. Attackers manipulated Instagram's AI-powered account-recovery assistant — a support tool with write access to email-binding and password-reset functions — into issuing password-reset links to attacker-controlled addresses, simply by sending a natural-language instruction naming a target username (Cybersecurity News, 2026). The exploit required no credential theft and no

breach of the server. Meta confirmed its backend was not compromised, and accounts were hijacked by abusing the recovery logic rather than defeating two-factor authentication directly. The flaw was a "confused deputy" — a long-recognised class of security failure in which an intermediary that legitimately holds elevated privileges is tricked by someone who lacks them into exercising those privileges on their behalf. Here the deputy was a probabilistic language model placed in front of irreversible account actions, with no deterministic verification gate behind it. High-value "OG" handles were taken over and listed for resale on Telegram within minutes, including the dormant Obama White House account (SQ Magazine, 2026).

The significance for age assurance is structural rather than a claim about any particular provider. The same platforms now legally required to hold children's biometric data are simultaneously wiring AI agents into their account and support infrastructure. As we (section 2.1) document, Meta, TikTok and Google already run in-house machine-learning models with privileged access to user data and account state, and the Instagram case shows what happens when such an agent can execute state changes without a hard checkpoint. Age verification then concentrates precisely the data that makes any such failure irreversible. A compromised password can be reset, a leaked face scan, iris template or identity document cannot (section 5.5). The breaches already on record show this is not hypothetical: Discord users' government IDs were exposed when a third-party customer service provider handling its age-related appeals was breached; one of TikTok's verification partners, AU10TIX, left users' photos and driving licences exposed for almost a year; journalists linked TeaOnHer users to their driving licences within ten minutes; and researchers who examined exposed Persona code found the platform could run up to 269 verification checks, including adverse-media screening, although Persona says these are not part of its age estimation (section 6.1). Taken together, systems whose safety case rests on age verification are adding both a high-value, permanent biometric honeypot and an increasingly AI-mediated access surface of exactly the type that has now demonstrably failed in production — raising the stakes of a breach to lifelong, unfixable identity loss while introducing new ways for that breach to occur.

7.5 HIdden harm: Teaching children to perform identity on camera, on command

Adolescence is the developmental window in which identity is least fixed and most in need of experimentation — Erikson's (1968) "identity versus role confusion," confirmed empirically by Branje et al. (2021) as an iterative, provisional process rather than a settled state. This is precisely what compelled facial performance corrodes. Butler (1990) argues that identity is not a possession expressed by a pre-existing subject but is *constituted* through reiterated, citational acts — yet her account reserves freedom for agentic resignification, the capacity to repeat norms differently and thereby subvert them. This is the distinction the age-assurance apparatus collapses: performativity in Butler's sense is not the theatrical performance of a role a subject knowingly chooses, and the liveness check extracts the reiteration while stripping out the improvisation that makes it generative — not self-authorship but externally-scripted iteration, turn left, blink, hold still, performed to a machine on command as the price of participation. In Foucault's (1977) terms this is a disciplinary ritual of the "examination," binding subjection and objectification by rendering the subject visible, individualised and recorded, and thereby governable. It also destroys what

Goffman (1959) called audience segregation — the ordinarily protective ability of young people to present a different self to parents, peers and strangers; a single verified identity that follows the user everywhere produces permanent context collapse (Marwick and boyd, 2011), leaving nowhere to be an unfinished, experimental self.

We believe this is going to be disastrous not only for individuals, but collectively for global society.

Against a system premised on one authoritative identity, Castells (2010) and Hall (1996) insist the self is plural, fragmented and multiply constructed across intersecting discourses — and, as Mukasa et al. (2022) warn, centralised identity infrastructures become instruments of exclusion and control, a danger sharpest in a culture-wars climate where trans, gender-diverse and minority young people lose most when the right to be more than one thing, or not yet anything settled, is engineered out of the system for life.

7.6 Children's rights implications of social media minimum age restrictions and AV

Social media can be places where young people access their rights — to self-expression, to play, and access to information. Human rights law holds that all rights are indivisible, interdependent and interrelated. For children, a properly conducted "best interests" assessment helps determine which rights prevail where they conflict, by weighing the risks that consider both their "agency" and "protection" rights (Livingstone et al., 2024). The balancing test is often seen as a zero-sum game of where one right is protected and it is acceptable to infringe on another. Protection from a risk of harm is likely to outweigh other considerations, a key question is whether an assessment can ever be accurate and more than procedural, if it is remote rather than a substantive individual check. But how intrusive would that have to become and would anyone want it?

The right of the child to be heard under Article 12 of the Convention on the Rights of the Child (OHCHR, 1989) is engaged directly in this policy area, because such measures plainly concern children and restrict their access to online spaces. Australia had no consultation on age social media minimum age restrictions, where the UK did. What children say is mixed. In the UK government's 2026 consultation, which drew more than 116,000 responses from parents, children and experts, two thirds agreed that under-16s should not be allowed to use at least some social media platforms (DSIT, 2026b). The 2026 UK Youth Poll found that 62% of young respondents supported an under-16 ban, down from 67% in 2025, though it did not survey 13 to 15-year-olds, the group most affected (University of Glasgow, 2026). In Australia, teenagers have said they want better protection online but described current age assurance proposals as cumbersome and in some cases overly restrictive (Woodley et al., 2025). The European Commission says the EU KIDS Act proposal drew on dedicated consultations with children, parents and educators, and it is seeking feedback on the proposal, including from children, until 26 November 2026 (European Commission, 2026d).

But we must ask whether a consultation was or is meaningful or genuine, if the children or adult respondents were not informed of or cannot understand what the implications of any processes required as a result of the ban will be?

- Which company has the platform contracted to carry out the age check?
- Where the third-party is based and how they process your image or video?
- How long these data are kept, and how they are used, including for machine training?

The SMMA Act is framed as a child-protection measure, but the age-assurance infrastructure it has produced engages a wide range of children's rights — many adversely. Under the UNCRC and General Comment No. 25, children have a right to privacy (Article 16), freedom of expression and access to information (Article 13), freedom of association (Article 15), non-discrimination (Article 2), to have their best interests taken as a primary consideration (Article 3), and to be heard (Article 12). Routing the biometric data, identity documents and behavioural profiles of every Australian child through a fragmented network of foreign processors — and, in Meta's case, into the training data for commercial generative AI models — is a privacy intrusion that would not be considered proportionate if applied to adults.

Children's rights apply equally and fully in digital environments as they do everywhere else.

Under the UN Guiding Principles on Business and Human Rights and the Children's Rights and Business Principles (illustrated in General Comment no.16), companies bear the responsibility to respect children's rights through human rights due diligence, including child rights impact assessments, and to ensure safety by design. Governments bear the primary duty to protect children's rights, which requires effective legislation, regulation, oversight of digital platforms, and provision of effective remedies.

While accepting that more must be done online, the Australian Child Rights Taskforce makes four warnings that bear on this report. On efficacy, it doubts a blanket under-16 ban would make platforms safer and warns it shifts the burden onto young people rather than the companies that should make platforms safe by design. On displacement, it cautions that a ban could push young people toward less-regulated platforms. Most distinctively, it identifies a "cliff edge" at 16, where a young person is abruptly exposed to harmful content without the gradual support or media literacy to navigate it — reframing a hard threshold as deferred risk and strengthening the case for digital-literacy investment regardless of any minimum age. And on rights, it warns a ban may infringe young people's ability to connect, learn and seek support online. It closes by calling for evidence-based decisions and investment in youth work as a positive alternative, alongside stronger platform accountability rather than access restriction alone.

The discrimination dimension is more substantial than the public debate so far has acknowledged. Facial age estimation is, on the trial's own data, less accurate for children than for adults — meaning under-16s are more likely than anyone else to be wrongly flagged and pushed into a second, more invasive check. Bias by skin tone and gender persists in many of the underlying models (NIST, 2014). Children without identity documents — those in out-of-home care, refugee and asylum-seeker children, Indigenous children, children whose names or appearance do not match their documents (including trans and gender-diverse young people) — face structural exclusion from any check that depends on an ID, as do children without a smartphone, a stable

address, or reliable connectivity. Where verification fails, the published flows offer no meaningful appeal route.

In the UK, Article 12 binds the UK in international law throughout England, Scotland, Wales and Northern Ireland, but its domestic force varies. In Scotland, where the Convention is incorporated within devolved competence, a devolved measure affecting children's online access could be challenged for incompatibility. There the Children and Young People's Commissioner Scotland, Nicola Killean, has said that a ban for under-16s would not currently be a proportionate, effective or enforceable way to protect children's rights, based on a children's rights impact assessment informed by young advisers aged 12 to 17 (BBC, 2026b).

In Wales, the 2011 Measure requires Ministers to have due regard to Article 12, so a failure to consider children's views is reviewable as a process defect. Its Children's Commissioner also opposes a ban.

New legislation handed ministers the power to restrict under-16s' digital access to more than social media platforms and adoption appeared a foregone conclusion after a junior minister said even before public consultation closed in May 2026 (Hooker and Cress, 2026), it was "clear that under any outcome" they would impose "some form of age or functionality [restriction] on children under 16," (House of Commons, 2026). The ban will be made by regulations under section 214A of the Online Safety Act 2023, a draft of which must be approved by both Houses; the government aims to lay them by the end of 2026, with the ban in force from spring 2027 (Woodhouse, 2026).

In England and Northern Ireland, where Article 12 is neither incorporated nor subject to a comparable duty, a measure that excludes children without hearing them cannot generally be struck down for breaching Article 12 alone; the right operates instead as an international obligation and an interpretive principle. Across all four nations, the international obligation persists: a regime built without children's input is vulnerable to the charge that it skipped a step the Convention treats as a precondition of any valid best-interests determination.

In general, these measures treat children only as objects of protection rather than rights-holders. A regime that prevents under-16s from accessing the spaces where their peers communicate also restricts their rights to expression, association, information, cultural participation, and access to support, including for mental-health, LGBTQ+, disability and minority community resources that many young people find through these platforms and nowhere else. Those costs of non-participation sit on one side of the scale; the evidence that the ban actually reduces harm (with around 70% of under-16s who previously had accounts reportedly retaining access, and the Korean precedent showing such bans can increase total online time) sits on the other.

According to the EDPB, age assurance poses specific risks to data protection with the potential to adversely impact not only natural persons' right to the protection of their personal data, but also other rights and freedoms such as the right to non-discrimination, the right to the integrity of the

person, the right to liberty and security, and the right to free expression and information (EDPB, 2025, p. 2, para. 4).

A child-rights-consistent approach would require a proportionality analysis that weighs both sides — considering safeguards for the children most likely to be excluded, harmed or surveilled by the verification infrastructure itself, and asking whether the costs are proportionate at all. The SMMA Act, as currently implemented, and therefore similar measures emerging in Europe, do neither.

7.7 What the French ruling found

On 14 August 2026 the French Constitutional Council struck down the central provision of the loi visant à protéger les mineurs des risques auxquels les expose l'utilisation des réseaux sociaux. Article 1 would have barred under-15s from social media services from 1 September, with existing accounts closed from 1 January 2027. The law had been adopted by 279 votes to 81 on 21 July, three weeks before the Council received two referrals from opposition deputies. In *Décision n° 2026-911 DC*, the Council held Article 1er contrary to the Constitution on two grounds. The first was freedom of expression. The second, and most relevant for age assurance, was privacy.

The Council did not reject the aim. It accepted that protecting the best interests of the child, together with the constitutional objective of preventing threats to public order, could justify limiting minors' access to these services. What it rejected was the instrument. At paragraph 17 it found that the legislator could not impose a prohibition of general scope, depriving minors of access to numerous online communication services, without regard to the minor's own situation or to the risks specific to each service. The ban made no distinction between platforms by function, content, or the protections they already offered. It allowed no assessment of a child's age or degree of maturity, and it left parents no means of lifting or adjusting it for their own children. At paragraph 18 the Council concluded that the interference with the freedom protected by Article 11 of the 1789 Declaration was "n'est pas adaptée, nécessaire et proportionnée à l'objectif poursuivi".

The privacy finding is the more significant of the two for anyone designing age assurance. Having recalled at paragraph 19 that the liberty proclaimed by Article 2 of the 1789 Declaration entails the right to respect for private life, the Council turned at paragraph 20 to what a ban actually requires in practice. By prohibiting access for every minor under fifteen, it found, the contested provisions imply by themselves that every person, including every adult, must prove their age before access. The applicants showed that an under-15 prohibition necessarily entails technical processes to verify the identity of social media users, and the law contained nothing governing the processing of that sensitive data, nothing naming who would be responsible for it, and nothing on the technical arrangements by which verification would occur. The Council agreed. Having failed to determine the conditions and the limits within which age would have to be proved, the legislator had not provided the legal guarantees necessary to secure these constitutional requirements. That omission was a distinct and sufficient ground for censure.

This is a constitutional court finding, based on its own reasoning rather than merely accepting the applicants' framing, that a measure aimed at children builds identity infrastructure for the entire adult population, and that a legislature which does not confront that consequence in the text of the law has not done its job. The finding was made without any evidence before the Council about breach rates, retention periods, biometric processing, or the commercial chains through which verification data actually travels.

This should be instructive for the UK and for the Irish and wider EU and global debates. Within hours of the decision the Élysée announced that the Prime Minister had been tasked with producing a legally robust rewrite, taking account of the ruling and of the European framework, with the President's determination to deliver by spring 2027 undiminished. On 14 September 2026 the government notified a revised text to the European Commission, which reportedly allows parents to lift the restriction for 13 to 15-year-olds (Franceinfo, 2026; 20 Minutes, 2026). The Conseil d'État had already indicated in its avis of 8 January 2026 the shape such a text would take: targeting by service rather than a blanket prohibition, and a mechanism of parental authorisation. Neither of those adjustments addresses the privacy ground. A ban that applies to fewer services still requires everyone to prove their age at the door of the services it does cover, and a parental authorisation gateway adds a second layer of identity obligations rather than removing the first. Unless the replacement text specifies what may be collected, by whom, for how long, and on what lawful basis, the defects the Council identified in August 2026 will not be resolved.

7.8 What others say about social media bans for children

A wide range of organisations have warned against the access-restriction model.

- **5Rights Foundation:** restriction of access through blanket bans is "misguided." (5Rights, 2026).
- **Amnesty International:** a ban is "an ineffective quick fix that's out of step with the realities of a generation that lives both on and offline." (Amnesty, 2025).
- **Article 19:** the UK approach "could be copied by States around the world setting a dangerous precedent for human rights online."
- **Back Youth Alliance:** "It remains to be seen whether a blanket ban for under-16s from using social media would make platforms any safer. There is a risk this approach could place an additional burden on young people rather than address the responsibility of technology companies to make their platforms safer by design. It could also have unintended negative consequences of pushing young people towards less known and more unregulated online platforms." (Back Youth Alliance, 2026).
- **Baroness Kidron** "it is likely that at least some of the increased use of VPNs by adults is the result of a legitimate concern that the introduction of age checks has not been accompanied by a rigorous focus on user privacy. Until privacy is central to Ofcom's concern, it will always meet resistance to age checking. How many services have been

referred by Ofcom to the ICO for failing to uphold users' privacy rights while performing age checks?" (VPN Debate, House of Lords, September 2025).

- **Big Brother Watch:** third-party verifiers will hold deeply personal data linked to browsing, and "that can be very dangerous." (Big Brother Watch, 2023).
- **Council of Europe (Commissioner for Human Rights, Michael O'Flaherty):** banning access "shifts the responsibility for safety from the platforms that create the environment to the children who navigate it." (O'Flaherty, 2026).
- **Electronic Frontier Foundation:** AV mandates are "wrapped in the legitimate concern about children's safety [but] operate as tools of censorship." (EFF, 2025a).
- **European Digital Rights:** these tools "systematically leak data in ways that threaten the privacy and data protection rights of children and adults alike." (EDRi, 2025)
- **FIRE:** age verification "raises serious privacy concerns, intrudes on anonymous comment and access to information." (FIRE, n.d.)
- **Index on Censorship:** the provisions risk "creating a chilling effect on legitimate speech" and fostering "an environment of self-censorship." (Ball, 2025).
- **Molly Rose Foundation:** Australia's ban delivers "a false sense of safety that quickly unravels." (Molly Rose Foundation, 2026).
- **NetChoice:** "no one should have to hand over a government ID to speak in digital spaces." (NetChoice, 2025)
- **438 security and privacy scientists:** it is "dangerous and socially unacceptable" to deploy these systems before scientific consensus. (Joint statement, 2026)
- **Open Rights Group:** the Online Safety Act "creates incentives for poor security and privacy," risking "a free-for-all and potential privacy debacles." (Open Rights Group, 2025)
- **United Nations High Commissioner for Human Rights**[30] "The risks posed by social media to children's safety stem from platform design and practices that prioritize profit at all costs. Addressing them through regulation and human rights can take us further than blanket bans." (OHCHR, 2026).
- **UNICEF:** "social media bans come with their own risks, and they may even backfire." (UNICEF, 2025b)
- **Wikimedia Foundation:** "We will not be age-gating Wikipedia under any circumstances." (Rodgers, 2025).

On June 15, 2026 the UK government announced "Social media platforms to be blocked from offering services to under-16s, marking a line in the sand and setting a new normal for future generations." (DSIT, 2026b)

[30] https://www.ohchr.org/sites/default/files/documents/issues/children/getting-childrens-safety-online-right.pdf



Source: Social media to be banned for under-16s in landmark government move (gov.uk,) DSIT, 2026b

8. VPN bans and content blocking

At the time of the UK announcement that social media platforms would be blocked from offering services to under-16s from 2027, it was reportedly still under consideration whether or not to also put age restrictions or a ban on VPNs. The Minister for AI and Online Safety wrote to civil society advocates saying, "*we are not banning or age gating VPNs. However, we must also ensure that these tools are not misused to undermine child safety protections. We will continue to keep this area under review*."

The UK's own experience with pornography sites shows both effects. After highly effective age checks became mandatory on 25 July 2025, visits to the ten most-visited pornography services fell and settled at a lower level, and daily VPN use rose sharply (Freevacy, 2025). Ofcom's latest figures put UK daily VPN users at around 2.2 million, compared with 1.2 million before 25 July 2025, and show the higher level holding since the initial spike; Ofcom says an earlier data supplier had understated it (Ofcom, 2026b, para. 7.67). Its data does not show how much of the increase came from children and how much from adults seeking privacy.

Two 2026 surveys have reviewed use of VPNs by children. A YouGov poll of 2,167 GB children aged 11–17, in March 2026 for Ofcom,[31] explored the extent to which children use VPNs, including the extent to which they are used for certain activities associated with VPN use, and how frequently they use a VPN. Of the respondents, 3.2% had used a VPN to access age-inappropriate content in the last 6 months. Among children who accessed content meant for older people, only 11% used a VPN— versus 54% who simply said they were older than they are, 22% who had an adult account, and 17% who used someone else's device or account. VPN use is more extensive in "watched

[31]https://www.ofcom.org.uk/online-safety/protecting-children/experiences-of-age-assurance-use-of-vpns-and-attitudes-to-the-protection-of-children-measures

sports/TV/movies not available in the UK" (43% of VPN users vs 15% of all children) and gaming content (23% vs 9%). But the most-used-for purpose remains privacy. 46% of VPN users have taken steps to hide identity/location and 56% to keep personal information private, which is 21% and 29% of all children asked. Under a third of those children asked, had ever used a VPN at all. (Ofcom, 2026a). Ofcom's July 2026 report gives a figure of 5% from the same study, because it adds children who used a VPN to reach features with age restrictions. It also found that a quarter of 11 to 17-year-olds had used a VPN in the previous six months, and that 77% of 16 and 17-year-olds have ever used one (Ofcom, 2026b, paras. 7.68–7.72).

DSIT conducted further research, conducted by BMG between 18 and 29 May 2026 with 2,299 UK children (DSIT, 2026c), found 26% lifetime use and 22% in the last three months. **Most children who use a VPN feel safer for doing so, and value their privacy highly.** Eight in 10 VPN users (81%) say it makes them feel safer online. Privacy matters to 86% of all children, rising to 90% among lifetime VPN users and falling to 84% among non-VPN users.

DSIT treats the Ofcom figures as comparison rather than corroboration, noting only that its findings sit broadly in line with other recent research, and setting them alongside Childnet's November 2025 figure of 21% for the wider 8–17 age range (Childnet and Nominet, 2025). Both surveys also probe children's motivations, covering much the same ground — privacy, geo-blocked sport and video, gaming, and access to age-restricted content — but they frame the question in opposite directions. Ofcom asks, for each activity a child undertook, which method they used to do it, basing the answer on children who did that activity; DSIT asks VPN users to select their reasons, basing the answer on all lifetime VPN users. The resulting percentages therefore share a subject but not a denominator: Ofcom finds that 17% of children who accessed age-restricted features did so using a VPN, while DSIT finds that 22% of VPN users cite gaining age-restricted access as a reason, from which it derives its headline figure of 7% of all children use a VPN for this reason.

8.1 The pattern: tools of censorship and surveillance

To make age checks enforceable, states try to close the routes children use to evade them. The natural target is the VPN, and platform blocking more broadly. But the historical record is stark. VPN bans and platform shutdowns are overwhelmingly the tools of authoritarian states. China permits only government-licensed VPNs and detects others through the Great Firewall; Iran, Russia, Myanmar, Turkmenistan, North Korea, Belarus, Oman, and others variously criminalise, fine or imprison for unauthorised use. Platform blocking follows the same map — long-term bans of Facebook, X, YouTube, Instagram and TikTok clustered around elections, protests and "national security" episodes in dozens of states. Turkey's 2026 under-15 ban is notable for pairing age verification with explicit anti-VPN measures in the same law. Gabon shows how closely the two can be tied. Its media regulator suspended access to major social media platforms from February to April 2026 during anti-government protests, and a new regulation requires social media users to provide verified names, addresses and ID numbers, with fines and prison terms for

non-compliance (Egbejule, 2026). (Full timelines of VPN restrictions and platform-blocking incidents are at Appendix F.)

9. The European picture

9.1 National proposals across the Council of Europe

Other countries are considering similar provisions, including Denmark (Bryant, 2025), Spain, Norway and the UK (Wertheimer, 2026). European Commission President Ursula von der Leyen (Vohra, 2026) announced a European age verification app on 15 April 2026, with France, Denmark, Greece, Italy, Spain, Cyprus and Ireland "piloting the development of the app" and planning to integrate it into national eID wallets. Approaches vary widely: outright bans (France, Spain, Norway, Greece, Slovenia, Turkey), parental-consent restrictions (Portugal, Romania, Italy), states relying only on the EU DSA baseline, and outliers such as Estonia, which openly opposes bans and prefers a digital-literacy approach. Eleven Member States signed a joint letter advocating for mandatory age verification for social media to be included in the minor-protection guidelines of the EU's Digital Services Act (DSA). A full country-by-country breakdown, across 46 Council of Europe states and observers (including the 27 EU member states) — is at Appendix D.

On 17 September 2026 the European Commission proposed the EU KIDS Act. Under-13s would be barred from social media. Children aged 13 and 14 could use only a mini account managed by a parent or guardian, with limited features and a limit of one hour a day, and those aged 15 and over could hold their own accounts. Platforms would have to verify age whenever a new account is opened, and the Commission states that its age verification app does not retain identity documents or biometric data (European Commission, 2026b). The proposal now passes to the European Parliament and the Council. If adopted, it would make age verification a legal condition of opening any social media account in the EU, whatever the user's age.

Article	Applies to	What the proposal requires
⊘ **Access limits and minimum ages (Chapter II)**		
Delayed accounts Art. 6(1)	☺ ▶	No account for under-15s where a service has any listed risky feature: livestreaming, contact with strangers, profiling recommender systems, unsolicited suggestions, endless feeds or notifications that prompt return.
Parent-held accounts, ages 13–14 Arts. 6(2)–(3)	☺ ▶	A guardian may open a limited account in the guardian's own name, with guardian tools always on, a daily cap of up to one hour and pre-approved contacts. Parental responsibility and the child's age of 13 must be verified.
Existing accounts Art. 6(4)	☺ ▶	Within six months of application, platforms must check existing accounts and disable those of under-15s or of users whose age cannot be established.
GDPR Article 8 Art. 6(5)	☺ ▶	The age limits apply in parallel with national digital ages of consent where processing rests on consent (see section 10).
Under-13s Art. 7	▶	Access only through a guardian's own account, on video-sharing services designed for under-13s, for children aged 3 and over, for up to one hour a day, with personalisation and search off.
⛨ **Safety by design (Chapter III)**		

Article	Applies to	What the proposal requires
Protected by default Art. 8	☺ ▶ ♘ ♡ ❝ ▦	Protective design applies to every user, with or without an account, until age assurance shows the user is an adult.
Addictive design and usage limits Art. 9	☺ ▶ Partial: ♘ ♡ ❝	No design that encourages compulsive use, such as autoplay, notifications not triggered by the minor or streak penalties. Time limits and interruptions must protect school time and core sleep hours.
Recommender systems Art. 10	☺ ▶	Explicit preferences take priority; engagement-based recommendations are off by default; no data from outside the service; a non-profiling option must be offered.
Safe default settings Art. 11	☺ ▶ Partial: ♘ ♡ ❝	Geolocation, microphone, camera, contact sync, account suggestions and push notifications are off by default. Features that drive social comparison or embellish a child's image are unavailable.
Contact and visibility Art. 12	☺ ▶ Partial: ♘	No contact from strangers without the minor's pre-approval; profiles hidden from unaccepted users; no screenshots or downloads of minors' content; livestream hosting off by default.
Economic transactions Art. 13	☺ ▶ Partial: ♡ ❝	Purchases labelled in real time, virtual currency shown in national currency, and no variable reward systems such as loot boxes.
AI companions and chatbots Art. 14	♡ ❝ When embedded: ☺ ▶ ♘	No features that create emotional dependency; memory of past interactions off by default; under-13s only through guardian tools; testing before launch and monitoring after. Embedded chatbots must not switch on automatically and must be easy to opt out of.
Online games Art. 15	♘	Addictive-design, safe-settings and contact rules apply; guardian tools are mandatory; under-13s only through guardian tools; games must not entice minors into contact on riskier services.
App stores Art. 16	▦	Every app must carry an age rating. Stores must assess each user's age, block age-inappropriate apps and offer the EU age verification app.
Agency, reporting and guardian tools Arts. 18–20	☺ ▶ ♘ ♡ ❝	Child-friendly controls and reporting, priority handling of minors' reports, and guardian tools that minors are told about. Very large platforms' guardian tools must work with third-party tools.
⚷ Parental responsibility (Chapter IV)		
Verifying parental responsibility Arts. 26 and 31(1)	☺ ▶ ♘ ♡ ❝	Providers may use signals from official databases (recital 49 names education records and birth registries), past use of the service, or a guardian's self-declaration with reasonable checks. Member States must offer a free electronic attestation.
✓ Age assurance and verification (Chapter V)		
Definitions Art. 3(5)(h)–(i)	*All in scope*	Age verification rests on identification documents or other verified sources of identity. Self-declaration is not age assurance.
Principles and data protection Arts. 27–28	*All in scope*	Accurate, non-intrusive and non-discriminatory; must not identify, track or profile users; must use zero-knowledge proof. An age signal may be stored at account level.
Certified verification at sign-up Arts. 29(1)–(3)	☺ ▶	Only a certified EU age verification solution using a third-party EU proof-of-age attestation; certified EUDI Wallets qualify.
Alternative age assurance Art. 29(4)	☺ ▶ ♘ ♡ ❝ ▦	For safety-by-design duties and app stores, other methods are allowed if they meet Articles 27–28.
Operating system age signals Art. 29(6)	⚙	Operating systems holding a compliant age signal must share it with in-scope services, with the user's consent.
Proof of age for all Art. 31	*Member States*	States must provide ways to obtain proof-of-age attestations (recital 60 lists biometric ID cards and passports, digital ID, apps and in-person checks) and a free EU verification solution.

Article	Applies to	What the proposal requires
Checking existing accounts Art. 32	☺ ▶ Partial: ▦	Verification is required unless the provider has a high degree of confidence, such as account age or a credit card on file (recital 61). Very large platforms must file a plan.
⚖ **Enforcement, timing and review (Chapters VII and IX)**		
Supervision and fines Art. 34	*All in scope*	DSA and AI Act structures apply. Data protection authorities supervise age assurance processing with GDPR-level fines; AI providers face fines of up to 6% of turnover.
Expedited procedure Art. 35	*Very large platforms; Commission-supervised AI*	Preliminary findings within 30 working days and a final decision within 90.
Timing and review Arts. 42–43	*All in scope*	Applies six months after entry into force. Commission review due by 31 August 2030.

Table 9.1 — The EU KIDS Act proposal: what applies to whom, as proposed on 17 September 2026 (European Commission, 2026c).

Key: ☺ social networks ▶ video-sharing platforms ♘ online games ♡ AI companions ❝ general chatbots ▦ app stores ⚙ operating systems. "Partial" means only some of the article's duties apply, through Articles 14 and 15. Not-for-profit encyclopaedias, educational, research and public-authority services, and open-source software development and sharing platforms are out of scope (Art. 2(4)).

The proposal also asks a great deal of parents. Table 9.2 sets out what parents would have to do or provide under the articles above, and what follows for children and families (or "step in" State).

Table 9.2 — What the EU KIDS Act proposal asks of parents, as proposed on 17 September 2026 (European Commission, 2026c; EDPB, 2025).

Requirement	What the parent must do or provide	Implications for children and families
Parent-held accounts, ages 13–14 Arts. 6(2)–(3)	Open and hold the account in their own name; pass age assurance themselves; prove parental responsibility and that the child is 13.	Adults are brought into the age-check system by design. The child's online life sits inside the parent's account, under the parent's identity. (There are extraordinary implications for children who have no parents to become in essence "digital wards of the State"too).
Under-13s on video-sharing services Art. 7	Give access only through their own account, on services designed for under-13s, for up to one hour a day.	A very young child's viewing is linked to a verified adult account, joining the parent's identity data to the child's activity on the platform.
Parental / Guardian tools Arts. 6(2) and 20	Use parental / guardian controls.	Monitoring of 13 and 14-year-olds becomes the default. This engages the child's right to privacy (UNCRC Art. 16) and their evolving capacities (Art. 5), and General Comment 25's call for oversight proportionate to the child's age. In online games, guardian tools are a condition of access for under-13s even in low-risk titles rated PEGI 3, and deciding when guardian-controlled access is needed may put pressure on providers to collect age information about every player (Gamers Voice, 2026).
Verifying parental responsibility Art. 26; recital 49	Be matched to the child through official databases such as education records or birth registries, through past use of the service, or through a self-declaration with reasonable checks.	The family relationship becomes a verified data point that platforms check. Children's education and birth records are reused for a commercial access decision.
State attestation of	Obtain an electronic attestation of	A new state-issued credential linking a parent to a child, which most

Requirement	What the parent must do or provide	Implications for children and families
parental responsibility Art. 31(1)	parental responsibility from the Member State.	countries do not have today. Separated parents, kinship and foster carers, children in care, refugee and unaccompanied children, and parents without ID may not fit a database match; the EDPB says alternatives should be offered to people without suitable documents (EDPB, 2025, p. 6, para. 22(a)).
GDPR Article 8 in parallel Art. 6(5)	Where processing rests on consent, authorise it for a child below the national digital age of consent.	Two parental gates stack: the KIDS Act account limits and Article 8 consent (see section 10). The burden moves further onto parents, as the Australian Child Rights Taskforce warned (Table 7.1).

> *"The reality is that these tools systematically leak data in ways that threaten the privacy and data protection rights of children and adults alike – whether they use government-issued identity cards, AI analysis to predict a person's age or other methods. They create ways for people's digital footprint to be constantly tracked, can have a chilling effect on everyone's ability to access legitimate content, and frequently perpetuate discrimination and exclusion. Although this step from the European Commission is purportedly to age gate access to porn platforms, the plans that have been laid out show a clear intention to control access to a much wider range of platforms and services."*
>
> *"specifications will require people to have access to modern smart phones and government issued IDs."*
>
> — EDRi, 2025

9.2 The EU Commission's age-verification app and its unresolved contradictions

The EU Special Panel on Online Safety, in its report of 13 July 2026, advocated stepped age-gating of features, safety by design, digital literacy tailored to age, and greater emphasis on children's rights — though no child rights impact assessment of the age-gating plans has yet been carried out. The European Commission has developed an age verification framework built around a free, open-source app, often called the "mini-wallet," intended to let users prove they meet an age threshold without revealing their identity. The ecosystem separates four roles: issuers (Proof of Age Attestation Providers) verify a user's age from a strong source such as a national eID or, failing that, an ID card, passport, counter check, or accredited third-party app (European Commission, 2025a); the Age Verification Solution Provider is the app itself, standalone or integrated into the EU Digital Identity Wallet (EUDIW); verifiers (relying parties) are the online services requesting a proof-of-age attestation (King and Spalding, 2026); and trust list providers — the Member States — operate the list of accredited issuers. Built on the same technical specifications as the EUDI Wallets due by the end of 2026, the app is designed for compatibility and future integration into those wallets (European Commission, 2026a).

The framework is intended to comply with the DSA (Article 28, protection of minors) and to align with the amended eIDAS Regulation (eIDAS 2.0) and the EUDI Wallet reference framework. Under eIDAS 2.0, every Member State must offer at least one wallet by 6 December 2026, and very large online platforms must accept it. The Commission's DSA Article 28 guidelines identify scenarios requiring full age verification rather than softer estimation for access to pornography, gambling, and alcohol and tobacco advertising, and recommend the EU app and EUDI Wallet as the standard. The first blueprint was published on 14 July 2025; von der Leyen announced the app was "technically ready" on 15 April 2026; Birgit Sippel, a prominent German centre-left lawmaker, called the app a "half-baked app solution that doesn't live up to [the EU's] own standards," after claims the app had been hacked in under 2 minutes, and that within hours of the EU's app release, a security consultant found it would store sensitive data on a user's phone and leave it unprotected. Another found it was possible to bypass the app's biometric authentication features, meaning someone would be able to forgo entering a PIN code or using Touch ID to access the app (Politico, 2026).

The non-binding Recommendation establishing a common framework is dated 29 April 2026; and the Commission announced deployment readiness on 11 May 2026 (Inside Privacy, 2026). Distribution in pilot-country app stores was expected by summer 2026, though most Member States are unlikely to meet the end-of-2026 deadline (Biometric Update, 2026a). In July 2026 the Commission's DG CONNECT stated, in reply to an access-to-documents request, that it holds no privacy impact assessment written specifically for the app and that responsibility for one lies with the future publisher (PPC Land, 2026). The KIDS Act proposal itself rests on a Staff Working Document rather than a full impact assessment, drawing on a study the Commission procured on the impact of age assurance on fundamental rights (European Commission, 2026c). Notably, the blueprint[32] largely skips what happens between step 1 ("download an age verification app") and step 2 (proof of age issued from national eIDs, passports, ID cards, or banking data).

[32]https://digital-strategy.ec.europa.eu/en/factpages/blueprint-age-verification-solution-help-protect-minors-online

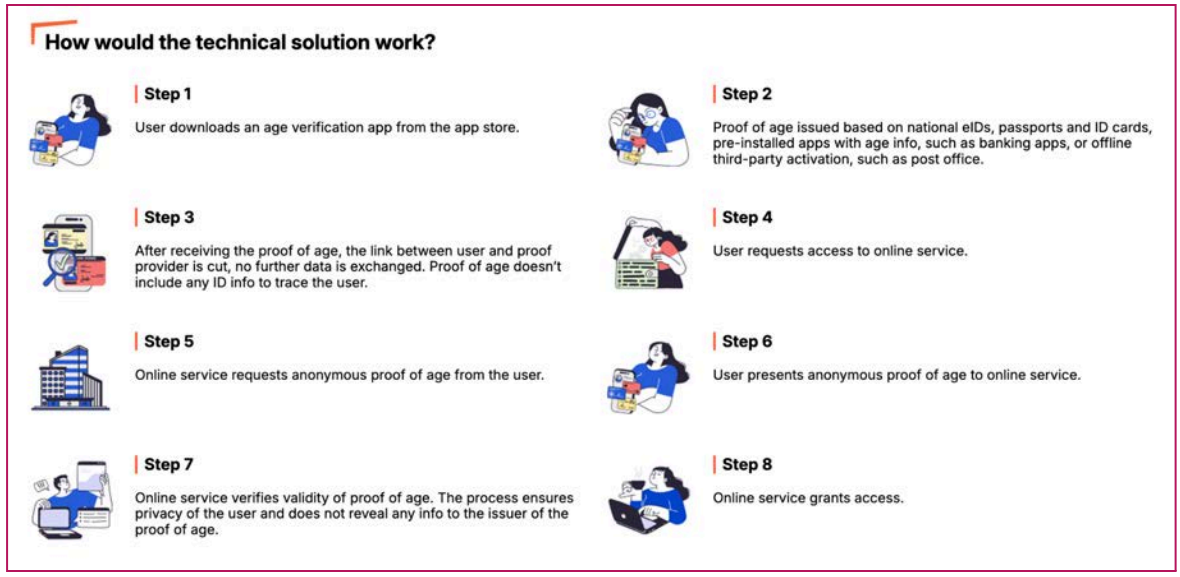

Fig.1 EU Commission Blueprint for an age verification solution (European Commission, 2025a).

The central tension is that a system marketed as cryptographically "privacy preserving" at the point of verification still depends on a state-accepted identity document being shown to the AV provider at the earlier credential-issuing stage. Providers typically claim that once the ID is seen, identifiable data is hashed: the system no longer knows that Person A is named A, and so cannot pass that name to a third party. But hashing still lets the provider single out Person A from Person B, recognise them again, and tell users apart. This is not resolved by the fact that, in the moment of verification, the online service receives only a token confirming the user meets the age threshold and the proof provider is not told where the proof is used (the mini-wallet does not learn it was used with Meta, and Meta receives no more than "yes, this user is over 18"). The solution relies on Zero-Knowledge Proof cryptography, which lets the app prove a single fact without revealing other attributes and prevents different relying parties from correlating responses (European Commission, 2026a). In the Commission's technical manual, presentation over the W3C Digital Credentials API and OpenID4VP is supported, ZKP support is listed as upcoming, and the data model is based on ISO/IEC 18013-5 (mobile driving licence) and ISO/IEC 23220-2. Under the EU KIDS Act proposal, self-declared age would not count as age assurance at all, and platforms would have to rely exclusively on a certified EU age verification solution using a proof-of-age attestation provided by a third party; the Regulation would apply six months after entry into force (European Commission, 2026c). The Commission states that its app does not retain identity documents or biometric data, yet the proposal expects proof-of-age attestations to be issued from biometric identity cards and passports, digital identity schemes, third-party applications or in-person checks (European Commission, 2026c, recital 60). That is the earlier contradictory reality of the document-based stage of the process this report describes.

As with the G7 Common Set of Principles defining a safer and more secure digital space for minors, adopted by G7 digital ministers on 29 May 2026 and welcomed by G7 leaders at Évian on

17 June 2026 (G7 Digital and Technology Ministers, 2026; G7, 2026), policy frames protective measures as needing to be least-invasive, privacy-preserving, and balanced against free expression, but does not resolve the tensions these demands create — leaving the design architecture to industry and the resulting adoption problems to implementers with new risk passed on to the users themselves.

Further inconsistencies remain in the EU approach. Under Article 91(1) DSA, the Commission has evaluated Article 33 and the DSA's interaction with other acts (European Commission, 2025b). The DSA exempts micro and small enterprises from the Article 28 duty, yet comparable rules reach small video-sharing platforms via the Audiovisual Media Services Directive, so the obligation to protect minors is applied inconsistently. As a full-harmonisation instrument, the DSA means Member States should not maintain additional national requirements within its scope; national age-restriction measures in France, Spain, Greece and elsewhere may therefore be tested against full harmonisation and the country-of-origin principle, risking parallel proceedings. The EU KIDS Act proposal of September 2026 would set minimum ages at EU level, with the stated aim of avoiding national fragmentation (European Commission, 2026b). Crucially, Article 28(3) DSA "*shall not oblige providers of online platforms to process additional personal data in order to assess whether the recipient of the service is a minor*" — it creates no legal basis to process additional data, and the GDPR's prohibitions on special-category data (including biometrics) and its data-minimisation duty (Art. 25.2) continue to apply. These questions are also relevant for the forthcoming Digital Fairness Act, the AVMSD review, and the full DSA evaluation due by 17 November 2027.

Because these systems are designed to operate around processing the human body, they should in our view be treated as handling bodily data throughout the process end-to-end, with the protection afforded to biometric data extended to all bodily data, regardless of any identification.

9.3 The consent failure built into choice-architecture

The SMMA Act requires consent to age assurance that is voluntary, informed, current, specific and unambiguous — an elevated standard that precludes reliance on preselected settings or opt-outs. Of the thirteen provider flows we mapped at Appendix H, none shows a documented, specific, standalone opt-in for retention period or reuse purpose. Only two (Persona and PrivateID) show any documented, discrete opt-in for the act of undergoing the check itself, and even there, retention and reuse are not separately consented to—either bundled into that same single action, or governed by the social media platform’s own instruction rather than the user’s choice during the age check part of the process.

In our own view, to the child or parent presenting their face, the distinction between biometrics or not is invisible. Where age checks are mandatory the “choice” relied on is not a real one. Consent is neither informed nor freely given. Consent is manufactured and invalid.

Companies behind age verification tools rely on consent in data protection law in Europe, as a lawful basis for processing biometric data. They claim this is valid because children are given a

"choice" of more than one age verification method to use to get into an online space. However, the route with the least friction is most likely the one “chosen” by children who are very keen on taking part, and in fact, is no choice at all. Participation is incentivised and power imbalanced with the provider and the platforms.

By requiring state-accepted ID at the start of any age verification process, against which a variety of third-party age verification providers must verify to create the credential, the idea that these systems are "privacy preserving" is only a claim about further on in the process. The widespread adoption of age verification enables a previously unseen scale of ID distribution and access by these commercial providers.

Worse, this normalises that children should do something online that they may be uncomfortable with doing, because they are incentivised to do and an adult tells them to.

Frequently, schools told children in COVID remote school learning that it was OK to keep cameras off. The IWF in the UK has warned schools to remove children's photographs from school websites and school social media platforms as schools are being threatened with AI generated CSAM created from pupils' photos online, posted by the school.

Yet we are being told to encourage children to give away their facial image freely, just to get online.

On the *60th World Day of Social Communications*, Pope Leo XIV warned that we must protect faces and voices to safeguard human qualities in an AI world. (Holy See, 2026).

Instead, children are learning to perform for the camera in a facial scan and turn left and right on demand, for "liveness checks", just to seek permission to participate in conversations with friends, play a game or join in a local community forum, solely because age verification duties demand it. Politicians are unwittingly demanding that children trade their bodily data for fundamental participation online.

We suggest this is an abuse of consent from which children must be protected. And it is the opposite of the message that politicians should want young people to learn about what consent means in today's digital environment.

10. The “digital age of consent” myth and the age of majority

What is meant when policymakers and the media discuss a “digital age of consent” to discuss restriction of online content access on the basis of age? The challenge of making content restriction or access based on age, is that it masks the consequences of the methodology that is not designed for the purpose intended. A focus on age arguably incentivises skipping the most important part of the policy process, problem structuring and definition (Veale, 2019). There is little legal definition on the ends of harm across jurisdictions, much less how it is achieved through age-gating. As Veale wrote on AI in 2019, the focus on online safety should be, “*in taking off the tech-goggles to identify problems, challenges and needs and not to be afraid to discover that other policy options are superior to a technology-based solution*.”

10.1 What Article 8 GDPR does and does not do, in European law

Article 8 of the GDPR is often wrongly described as a "digital age of consent." It is neither a minimum age at which a child's data may be processed, nor a lowest age-limit for using social media. It comes into play only where two conditions both hold: (1) the lawful basis for the processing is consent from the data subject (Art 6(1)(a)); and (2) the processing is for an information society service offered directly to a child. In that narrow case, Article 8 sets the age at which a child can give consent for themselves, rather than needing it given or authorised by the holder of parental responsibility — 13 in the UK, and between 13 and 16 depending on the Member State across the EU.

We frequently see critical claims that "most social media platforms do not currently enforce their minimum age policy, in breach of their long-standing legal obligations under GDPR," when in fact if they process no data under the basis of consent, they need not apply Article 8 nor ask for consent, nor age, at all. There is no GDPR-based floor below which a child's data cannot be processed. Where a platform relies on contract or legitimate interests instead of consent, it can process a younger child's data with no minimum age and no parental intervention required by data protection law. The "you must be 13" rule seen across UK social media is therefore generally a contractual / terms-of-service floor (reinforced in the US by COPPA), not a GDPR processing threshold.

For example, the minimum age for YouTube, Instagram, Facebook, Snapchat and TikTok is 13 years of age. The minimum age for Roblox is 5 years of age. Under 13s may use YouTube through a supervised account (with parental consent). YouTube also offers YouTube Kids, for which there is no minimum age.[33]

Changing Article 8 of the GDPR, is the wrong fit for constraining social media use.

This is why in the UK that the ICO, as an enforcer of data protection law has to date been the regulator involved with age appropriate design, not Ofcom in the UK whose remit to date in the Online Safety Act 2023, was limited to the restriction of adult content to 18+.

10.2 Why raising or lowering the age of the GDPR Article 8

Article 8 simply marks the line above which a parent no longer has to authorise — and supply their own personal data in addition to that of the child to verify and approve consent-based data processing. It does not grant a child generic "permission" to use social media, nor does it validate any other lawful basis such as legitimate interests with no opt-out, or processing on the basis of contract. Lowering or raising it would therefore be unlikely to make any material difference to children's social media use, because most of the processing involved does not rely on consent, and Article 8 never comes into play.

[33] Ofcom (2026). Keep underage children off your platforms, Ofcom tells tech firms https://www.ofcom.org.uk/online-safety/protecting-children/keep-underage-children-off-your-platforms-ofcom-tells-tech-firms

Because platforms use several different lawful bases across the many stages of a single user journey, Article 8 is engaged only at the precise points where a platform has chosen consent as its basis — and nowhere else. For any consideration of a change to the wrongly labelled "digital age of consent" in the UK GDPR Article 8 as a result of the Children's Wellbeing and Schools Act 2026, policymakers should first be fully aware of what it will not affect and why. In the EU, the proposed KIDS Act would add a second parental gate that runs alongside Article 8 rather than replacing it, through parent-held accounts and verified parental responsibility (see Table 9.2).

Unless the processing is consent-based, so that it places some restriction on (or duty for) parental intervention, a change to the age will have no effect at all for social media processing. Understanding under which lawful bases the various services process children's personal data is therefore essential before making changes to it in law. (A full lawful-basis mapping, and an account of exactly where Article 8 parental authorisation is triggered, is set out at Appendix C.)

10.3 Contradictory age thresholds across jurisdictions

Across the Council of Europe's member states and Australia, the law fixes a child's transition to adulthood not at a single moment but through a patchwork of separate age thresholds: the age of majority, of full contractual capacity, the end of compulsory education, the minimum working age, the ages at which a young person may buy alcohol and tobacco, gamble, drive, consent to sex, consent to medical treatment, marry, vote, be held criminally responsible, and consent to the processing of personal data under Article 8 GDPR. Far from aligning on a common boundary of childhood, these ages differ markedly from one country to the next and are fragmented even within a single jurisdiction — so the same person may be treated as an adult for one purpose and a child for another, with no consistent threshold.

11. Problems with the evidence base

Inaccurate evidence presented to policymakers and in the media frequently misstates the problems around children's online activity, likely leading to misunderstanding and misplaced political solutions, from why the term 'screen time' is unhelpful (Etchells, 2026), to definitions of harm and online safety.

Whether about social media use and mental health and wellbeing associated with social media use, online gaming, and mobile phone use by children, these errors and misrepresentations of still-emerging academic research, leave parliamentarians vulnerable to overestimating the scale of mental health harms, gaming "addiction" (European Parliament, 2019) and problematic phone use by children. A misinformed view of the scale of any problem misinforms the choice of solution. A common feature is an under-assessment of the positive impacts, and their scale, and how these should be considered alongside the negative consequences of restrictive laws, or harms that affect the whole population or even harms across an age group, such as treating all teenagers as a one-size-fits-all policy. Three recent examples stand out.

11.1 Claims of causal evidence given to the Australian parliament

The Explanatory Memorandum (EM) that accompanied the Australian Bill in 2024 claimed, that, "a United Kingdom (UK) study published in 2022, which examined longitudinal data from more than 17,400 participants, found that adolescent social media use is predictive of a subsequent decrease in life satisfaction for certain developmental stages including for girls aged 11 to 13 years old and boys 14 to 15 years old." Whose research goes unstated in the EM, but it is likely that of Orben et al. published in 2022. In fact their 17,400 participants were part of a far larger dataset and the research is far more nuanced than suggested.

The problem with this citation is that the depth of the study is lost, it is the only one cited and it supports the policy intention without suggesting there is a range of available evidence, and there is no balance presented from any other findings, or analysis which is found most compelling and why.

For example, a separate UK study of 3,228 participants in a similar time period who were 10- to 15-year-olds from Understanding Society (2009-2019), a UK longitudinal household survey, found, "There was little evidence to suggest that more time spent on social media was associated with later mental health problems in UK adolescents." (Plackett et al., 2023).

A further global study of 2,305,058 users across all age groups, who provided data between October 2020 and August 2023 found, that *of those who believed it affected wellbeing [at all]*, it was marginally negative, playing online games was positive, but there were also ranges of those who felt social media use increased their wellbeing of up to 20% as well as lowered it by 15%. In other words, enormous ranges of feeling across age groups and within them, and that,

> "the perceived impact of social media is small, negative, and varies across populations, gender, and age groups. To serve a global population, recommendations and regulations will need to be informed by local data, taking into account the highly individualized experiences of social media users." (Dienlin et al., 2025)

11.2 The EU: "one in four minors" and problematic gaming

The evidence adopted by EU parliamentarians before their key 2025 vote — deciding that children should be at least 16 to access social media — is a prime example of misleading findings given as fact to support a policy view, from far more nuanced research. MEPs adopted the report by 483 votes to 92, expressing grave concern that "one in four minors display 'problematic' or 'dysfunctional' smartphone use, i.e. behavioural patterns mirroring addiction." (European Parliament, 2025). This is misleading of the underlying findings.

The authors of the key supporting paper state clearly in its caveats that "the majority of the papers analysed addressed problematic gaming." (European Parliament, 2019) In other words, the research did not examine *all* internet and social media activity and then find that a quarter of it was problematic; it examined a subset — only *already-problematic* online activity — and found

some of that involved children. Presented the other way round, it gives a false impression of scale, and leads policymakers in only one direction, how to respond.

11.3 The UK: the 91% CSAM claim

The third example is from the former UK Minister for Violence against Women and Girls, who repeatedly used a wrong statistic about children creating "self generated imagery" or abuse and nudes on their own phones in her resignation letter and subsequent interviews (BBC, 2026a). It suggests a substantial misunderstanding that really matters if the government is designing policy change based on it. She emphasised on Sky News, "91% as I said in my letter, of the child abuse imagery that was found by the IWF and reported to police, last year [in 2025] 91%, so *nearly all CSAM* imagery was self-made by children on their own phones" (SkyNews, 2026). In fact, the IWF reports 27% not 91% nearly third not nearly all of all CSAM was, as they label it, "self-generated" in 2025 which is not defined as with being made "on their own phones" either (IWF, 2026a).

After the AI and Online Safety Minister's statement made in the House of Commons repeated the same mistake in June, DSIT later published a correction to the parliamentary record (Hansard, 2026) at our request. But, instead of a correction to the figure from 91% to 27% the subject of the percentage was *changed to fit the number*, without any contextual explanation.

Worse, on September 8th 2026 the new Minister repeated the exact same mistake in parliamentary debate when announcing legislation for device-level protections (Nandy, 2026) and it is still online in a published ministerial statement.[34] We requested a second Ministerial correction same day which remains outstanding 6 weeks later at the time of publication. The Minister for VAWG had resigned over what she claimed was the government inaction saying, "technology companies have the technology to do it" and talking about "operating systems' leaders" being made to take action and "if they haven't made progress in 3 months we will legislate towards it." Towards what is not specified (SkyNews, 2026). "We could make this possible on every phone and device in the country" and "the announcement was meant to be in March". What announcement went unsaid (BBC, 2026a).

The most likely source of the confusion is an IWF figure that 91% of reported URLs contained some self-generated content but, that is *the share of reported pages* that contain any self-generated image at all, and that could be all of the same single image, it is **not** the same as share of all imagery. As the IWF explains, each URL "could contain one, tens, hundreds or even thousands" of images, but perhaps a single self-generated image among a thousand that were not – targeting policy towards the former would misdirect effort from the latter.

Any amount of this material is appalling, but accuracy matters precisely so that (i) solutions target the real problems and best help children, and families and any victims and survivors of harm, as well as (ii) seeking the best technical solutions for everyone when considering the necessity and proportionality test in balancing public policy interventions at scale. This statistic very likely misled

[34]https://web.archive.org/web/20261016143326/https://www.gov.uk/government/speeches/lisa-nandys-statement-on-protecting-children-online

UK MPs and the public by inflating the prevalence of CSAM of children with no adult visibly involved, from just over a quarter to "nearly all." If British policymakers are proposing "on-device" intervention or law to address the problem, it matters that the problem is understood, and its evidence base is robust and accurate.

11.4 The case for restriction, and the harms behind it

Problems with how evidence has been presented do not make the concerns behind restrictions unfounded, and the case made by supporters of bans is understandable. Its first strand is collective action that makes it easier for parents to regulate what their children do. Parents report that limits set by one family fail when every other child is online, and supporters of bans argue that only a rule applying to everyone can change the norm. Jonathan Haidt for example, has been its most prominent advocate. The second is public demand. In June 2026, 44% of the UK public preferred a ban and 39% tighter regulation, and parents of under-16s favoured a ban over regulation by 54% to 36% (IPPR, 2026). Its third is distrust of platforms, which have failed for years to enforce their own minimum ages. Bereaved parents have been among the most powerful voices for change.

The harms ministers now cite most are sexual. This adds emotive and moral dimensions to debate. The Internet Watch Foundation took action on 397 reports of sexually coerced extortion, or sextortion, in 2025. Since 2023, 98% of such reports have involved teenagers aged 14 to 17, and 92% have involved boys (IWF, 2026b). The National Crime Agency says sextortion is typically carried out by organised crime groups overseas who deliberately target boys aged 14 to 18 (Police Professional, 2025). The UK government's device-level proposals are framed around stopping children taking, sharing or viewing nude images (section 6.6).

These figures and volume of activity matter for the design of any response. Most victims are 14 to 17, so an under-16 ban leaves 16 and 17-year-olds outside its scope, and perpetrators may reach children through messaging, gaming and other services that social media age limits will not cover.

The measures that would have most bearing directly on these harms are design measures: private-by-default accounts, no contact from strangers without a child's prior approval, on-device nudity detection, fast reporting and takedown through services and enforcement of the illegal content duties that already apply under the Online Safety Act 2023. Several appear in the EU KIDS Act's safety-by-design articles (Table 9.1), and most do not require anyone to prove their age.

12. Four approaches compared: Australia, Canada, the UK and the EU

Governments restricting children's access to social media differ in what they restrict, who verifies age and by what method, and the role they give to parents. Those choices decide how much personal and biometric data each system collects, and from whom. This section compares four approaches as they stand in October 2026.

Australia restricts specified platforms. Since December 2025, age-restricted social media platforms have had to take "reasonable steps" to prevent under-16s from creating or keeping an account (Woodhouse, 2026). Parents cannot provide "consent" to permit a child back in, and the

Australian Act does not prescribe a verification method. Each platform has therefore chosen its own combination of facial age estimation, ID upload, bank checks and behavioural inference, run through chains of sub-processors (section 2 and Appendices A and B). Enforcement rests on fines, doubled in September 2026, and on the eSafety Commissioner's information powers (Parliament of Australia, 2026).

Canada's Safe Social Media Act, introduced in June 2026, would bar under-16s from social media platforms unless those platforms meet safety criteria set by the government (CEPA, 2026). The ban therefore works as leverage instead. A platform that redesigns its service to the government's standard keeps its younger users. Of the four approaches, it gives design duties the most weight, while keeping exclusion as the sanction for platforms that do not comply. The bill must still pass the Canadian House of Commons and the Senate.

The UK government says its ban will follow the Australian model but go further, with obligations to block functions such as livestreaming and contact with strangers for under-16s. For 16 and 17-year-olds, overnight curfews will be on by default and features that prolong scrolling will be off by default (Woodhouse, 2026). The ban will be made by regulations under section 214A of the Online Safety Act 2023, which both Houses must approve in draft, and sits on top of Ofcom's "highly effective age assurance" standard. The method(s) of checking age remains open. In June 2026 the government asked Ofcom for a rapid assessment of methods for establishing whether a user is over 16, due by the end of October 2026 (Woodhouse, 2026).

The EU KIDS Act proposal grades access by age (European Commission, 2026c). Under-13s are barred, 13 and 14-year-olds may hold only a parent-controlled account, and under-15s are kept off services with listed risky features. Account sign-up must use a certified EU age-verification solution backed by a third-party proof-of-age attestation, which in practice anchors the check in state-recognised (digital) identity obligations. Safety-by-design duties apply to every user until age assurance shows the user is an adult, and parental responsibility becomes a verified attestation (Tables 9.1 and 9.2). As a Regulation, it would displace emerging national bans in member states.

The United States shows where litigation leads. In June 2025 the Supreme Court held in Free Speech Coalition v. Paxton that Texas's age verification law for pornography sites was subject only to intermediate scrutiny and upheld it, and by January 2026, 25 US states had enacted age verification laws, most for pornography and some for social media (Lin et al., 2026). California's Digital Age Assurance Act moves the check into the operating system from January 2027 (section 6.6).

Table 12.1 summarises the differences.

	Australia	Canada	United Kingdom	European Union (proposal)
What is restricted	Named age-restricted platforms	Social media platforms, unless they meet government safety criteria	Platforms, plus risky features such as livestreaming and contact with strangers	Accounts graded by age; services with listed risky features closed to under-15s

	Australia	Canada	United Kingdom	European Union (proposal)
Age thresholds	16	16	16; default curfews and limits on addictive features for 16–17	13 and 15, with a parent-held tier at 13–14
Role of parents	None; parents cannot consent a child in	Not a central feature	No consent route announced	Parent-held accounts, verified parental responsibility and a state attestation
Verification	Platform's choice of method	Not yet set out	Ofcom's "highly effective" standard; method under assessment	Certified EU solution backed by a third-party proof-of-age attestation
Design duties	Separate from the ban	The route out of the ban	Feature-level restrictions added	Apply to all users until shown to be adults
Legal form and status	Act in force since 10 December 2025	Bill introduced June 2026	Regulations under s.214A Online Safety Act 2023; expected by end 2026, in force from spring 2027	Proposed Regulation (17 September 2026); would apply six months after entry into force

Table 12.1 — Four approaches to restricting children's access to social media, October 2026 (CEPA, 2026; European Commission, 2026c; Woodhouse, 2026).

Read together, the four approaches show the trade-offs involved. Leaving the choice of method to platforms, as Australia does, hands the decision to platforms and their sub-processors, and the Australian evidence shows the result: widespread biometric processing, layered checks and behavioural profiling (section 7). Mandating one certified route, as the EU proposes, removes that variation by anchoring every account in state-approved digital identity obligations. That is the outcome the French Conseil constitutionnel identified when it struck down France's under-15 ban in August 2026: a measure aimed at children that requires every person, including every adult, to prove their age, with no legal limits on how that proof is collected or used (section 7.7). Canada's proposal makes platform design the condition of access, which comes closest to the systemic approach most child-rights organisations support, though exclusion remains the sanction. The UK has borrowed from each without yet deciding how age will be verified, so the privacy consequences of the UK ban may rest on Ofcom's assessment and the government's response to it.

Each approach can also be measured against the EDPB's Statement 1/2025 on age assurance. Age assurance must be shown to be necessary and proportionate to an identified risk, and checking the age of every user for content that is suitable for all audiences would fail that test (EDPB, 2025, p. 4, para. 14). It must not give providers or third parties additional means to identify, locate, profile or track people (EDPB, 2025, p. 4, para. 15), and viable alternatives should be offered to people who cannot or do not wish to use a particular method (EDPB, 2025, pp. 4–5, para. 17). A ban that requires every account holder to be checked meets the first of these conditions only if the risk is shown to justify checking everyone. The French Conseil constitutionnel held that a law with that effect must at least set out how proof of age is collected and governed (section 7.7).

13. Recommendations: beyond age verification

There is substantial common ground between child protection organisations and tech-civil society. Both agree age-based restrictions do not solve the problems of online harm, that platform redesign and systemic regulation should be the priority, and that poorly designed age checks risk harming the children they aim to protect (UNICEF, 2025a; EDRi, 2023, 2025), as no current method meets the necessity, proportionality and GDPR thresholds required to be adequately rights-respecting. The recommendations below consolidate consensus of positions. Public support for restriction is also conditional on how it is done. Support for an under-16 ban fell from 69% to 50% when respondents were told everyone would have to upload ID or a credit card (Ipsos, 2026).

13.1 What to do instead

Policy should start from the harms children face and the design choices that cause them. Online harm has systemic and societal causes, and a focus on age-gating can obscure more effective solutions. Regulation that changes how services are designed has a track record. Between 2017 and 2024, Meta, Google, TikTok and Snap made 128 changes to children's privacy and safety, with a peak of 42 in 2021, the year the UK's Age Appropriate Design Code took effect, and the changes with most effect were protections built in by default (Wood, 2024). A follow-up study of 70 platforms between 2024 and 2026 found regulation still driving change. It also found age assurance in use on 46 of those platforms, applied inconsistently and largely unaudited, and called for a product safety model with child safety testing before launch (Wood, 2026).

Five measures that might do more for children than checking everyone's age and its consequences.

- **Safe by default for every user.**

  This remains to be defined. But campaigners who use these terms tend to suggest private accounts; no content-sharing or contact from strangers without prior approval; no profiling-based recommendations or amplification of harmful content by default; no features built to compel return; user controls over what content people see, such as keyword filters a child can set to block a nickname used to target them; and reporting routes with effective redress. These may protect children whether or not their age is known, and protect adults too. Many are already in the EU KIDS Act's safety-by-design articles (Table 9.1). Below these, limit the collection and use of personal data for targeting content and features, for all users, since this drives the engagement features families object to (Luria, Bhatia and boyd, 2026).

- **Testing before launch.**

  Make human rights assessment [including child safety and data protection impact assessments] a condition of launching high-risk features, with independent testing and published results, and require qualified child rights due diligence (Wood, 2026).

- **Enforcement of the law that exists.**

  The illegal content and children's safety duties of the Online Safety Act 2023, data protection law including the Age Appropriate Design Code, and Article 28 of the Digital Services Act already apply. Regulators should use them in full before any new access restriction is introduced. Age-related laws must never stand in for companies' duty to invest in safer design and effective moderation.

- **Action on the specific harms.**

  For sextortion, that means blocking unsolicited contact from adults, funding reporting and takedown through services without ID obligations, disrupting the organised groups responsible, and education aimed at boys aged 14 to 17, who are targeted most (section 11.4).

- **Digital literacy for children, parents and educators.**

  Teach digital literacy, critical thinking and online safety from a young age, whatever minimum age is set, including what age and identity checks involve and the rights people have. Choices about how to prove age cannot be informed, and consent cannot be valid, without it. Child protection online is not appropriate to package as "stranger danger"

13.2 Before any age restriction is introduced

Where governments still choose to restrict access by age, they should first meet these conditions.

- **Assess the impact.** Carry out a human rights impact assessment, with a dedicated child rights impact assessment, before regulations are laid. It should cover every age-verification provider likely to sit in the chain, users of every age, and the effects on behaviour as well as the data and security risks. The European Commission should do the same for its age verification app and every provider approved within it.
- **Meet the EDPB test.** Show necessity and proportionality against an identified risk (para. 14); carry out a data protection impact assessment before processing (para. 13); give no additional means to identify, locate, profile or track users (para. 15); offer viable alternatives to people who cannot or will not use a given method (para. 17); and favour user-held, on-device and unlinkable approaches (para. 34) (EDPB, 2025).
- **Meet the French benchmark.** A measure that requires every person, adults included, to prove their age must set out in law what may be collected, by whom, for how long and on what lawful basis, and must be justified for everyone it reaches (section 7.7).
- **Define the scope.** Specify what counts as social media, whether any threshold applies to whole platforms or to specific features, and what terms such as "passive" mean.
- **Protect children's rights and anticipate the harms.** Account for children's evolving capacities and circumstances, and protect their rights to freedom of expression, access to information and free exploration of identity. Access to information is essential to how young people develop and form their identity (section 7.5; Luria, Bhatia and boyd, 2026).

Plan for the children who will be cut off from education, health and community resources, the rural, disadvantaged and marginalised children who will be harmed most, and those pushed into less-regulated spaces.

- **Learn from other countries** (section 12). Leaving the method to platforms, as in Australia, has produced widespread biometric processing; mandating one state-anchored route, as the EU proposes, builds identity infrastructure for every user. Where governments persist with access restrictions, conditions should fall on platform design, as in Canada's proposal, rather than on users' identity.
- **Respect family judgement.** Respect families' judgement about their own children, instead of overriding it by default. Many parents weigh access to peers, information and safety skills against the risks, and decide for themselves (Luria, Bhatia and boyd, 2026).
- **Hear children first.** Involve children before decisions are made and before regulations are drafted, as Article 12 of the UNCRC requires, and draw on the experience of young people, parents and caregivers who call for participation, privacy *and* protection and upholding their full range of human rights.

13.3 Reducing harm where checks are already imposed

Defend Digital Me opposes age verification for access to social media and the wider internet. Where checks are already in force, the measures below would reduce the harm they do in the meantime. Meeting them would not make age-gating acceptable however, there are also existing harms from such practices that we can reduce.

- **People choose the method.** People, not platforms or age-verification providers, should choose how they prove their age, from options that regulation requires to be available, including options that need no biometrics. Where no method meets privacy and security thresholds, providers must not be obliged to deploy one.
- **Necessary and proportionate checks only.** Limit age assurance to specific high-risk content or features, process data on the device, prefer less invasive options such as self-declaration backed by safety by design, and never require everyone to present ID. Reject document-based verification and age estimation where they fail the tests of necessity and proportionality or exclude people without official documents. Do not shift the burden onto parents and caregivers. As Australia's Privacy Commissioner, Carly Kind, put it, age assurance "is not a blank cheque to use personal or sensitive information in all circumstances and must not erode Australians' privacy rights" (OAIC, 2026).
- **Bodily data treated as biometric data.** Treat bodily data as biometric data throughout the process, closing the estimation-versus-identification workaround and raising the level of protection for children before the question is settled in case law.
- **No active liveness checks.** Active checks ask a person, most often a child, to perform gestures on camera. In the only controlled test of how users respond, government ID with an active liveness check had the lowest completion rate of any method, at 18%, and was rated among the least convenient (Lin et al., 2026). Where a liveness step is retained, it should be passive and should not be repeated at every access (section 3.4).

- **No reliance on privacy notices.** Statements that data would be used only for age verification and deleted afterwards made no significant difference to whether people verified; fewer than 1% opened the privacy policy, and 45% misunderstood who would receive their data (Lin et al., 2026). Require data minimisation and deletion by design, subject to independent audit, and a clear statement at the point of the check of exactly who receives what.
- **One mandatory certification standard.** A single enforceable open standard for age assurance should cover data minimisation, retention limits, a ban on secondary uses including model training, and interoperability, so that people can choose their own tool and are not required to perform with their face. Providers should list each risk and harm and say how they address it.
- **No VPN restrictions or platform blocking.** These are the tools of censorship regimes and would not survive a proportionality test in Europe (section 8).
- **No escalation in response to circumvention.** Do not respond to circumvention by escalating checks. Workarounds show where a policy fails to meet children's needs, and closing them means collecting more data from everyone (Luria, Bhatia and boyd, 2026).

13.4 Measure, review and be ready to stop

- **Measure outcomes**. Publish measures of harm to children in place of counts of accounts removed, with companies reporting material child-rights impacts in their public filings.
- **Build in an off-ramp.** Include review clauses and sunset provisions in every measure, so it can be ended if it does not do what it promises or once the unintended consequences are made apparent and can be weighed against the aims of the measures taken.
- **Consult on the method**. The UK government should publish Ofcom's assessment of age-assurance methods, and its own choice of method, for consultation in time for scrutiny and expert input before the section 214A regulations are laid.

> *"As several European countries consider introducing a minimum age to access social media platforms, I urge caution in imposing sweeping bans. The focus on restricting access should not deflect attention from ensuring that platforms respect human rights through clear legal duties, independent oversight, and effective accountability."*
>
> — Council of Europe Commissioner for Human Rights

14. Conclusion — and what is ultimately at stake

The growing trend of countries implementing strict minimum age laws will harm online anonymity, creating an internet that is less private and less free (Tuta, 2026). Australia's ban is being watched

as a template, and Europe is moving in the same direction. But the direction of travel does not make the course either immovable or correct. The Council of Europe's Human Rights Commissioner urged caution, echoing much of civil society: the harm lies in platform design and engagement-driven business models, not in children's access.

Two European authorities have already set limits that the current wave of proposals has not met. In August 2026 the French Conseil constitutionnel struck down France's under-15 ban for infringing freedom of expression and privacy, finding that a ban on minors requires every person, including every adult, to prove their age, and that the law set no conditions or limits on how that would be done (section 7.7). The EDPB has stated that age assurance must be necessary and proportionate to an identified risk, that checking the age of every user for content suitable for all audiences fails that test, and that age assurance must not become a new means to identify, profile or track people (EDPB, 2025, paras. 14–16). The proposals that followed, including the EU KIDS Act and the UK's planned regulations, still rest on checking the age of everyone who opens an account and do not meet that baseline obligation towards necessity and proportionality of the infringements.

The Australian evidence is clear. Age restrictions and age-based curfews do not create online safety, and there is no clear drop in the harms these measures are meant to address. What the ban has reliably produced instead, is commercial infrastructure that lasts a lifetime for the person involved. The face-based methods, which are the default route in practice, force children to hand over bodily data — biometric data that attracts the highest protection under the GDPR, Convention 108+ and the UNCRC, yet which providers reclassify as "not biometric" to sidestep those safeguards. Where age checks are mandatory, there is no meaningful choice: consent is manufactured, not freely given. In some cases that data is then fragmented across foreign processors and fed into commercial AI training. Verification of this kind does not loosen the grip of the largest platforms — it deepens dependence on them, and on a consolidating industry of commercial age-assurance companies, as the new gatekeepers of the digital public sphere.
Neither Australia nor the UK has conducted a holistic risk assessment of the measures introduced in safety's name. Routine age-gating ends privacy not to promote online safety but to promote online commerce, trading bodily data for everyday transactions in an environment where providers anticipate life online and offline being entirely automated. The next stage of that "readiness" is an internet without humans, where agents act on our behalf. Fewer humans in the process may look like less friction, but the convenience is bought at the cost of identity risk for life.

Mandating age verification asks children to deposit permanent biometric data into platforms that are simultaneously wiring AI agents into their account infrastructure — agents that, as the 2026 Instagram account-takeover incident showed, can be talked into irreversible actions they should never have authorised (Neowin, 2026; SQ Magazine, 2026). The security case for concentrating this data is correspondingly weak.

The contradiction at the heart of the policy was exposed in May 2026, when the UK's National Crime Agency, the Internet Watch Foundation and the safeguarding minister urged schools to strip

pupil photographs from their websites, after a school was blackmailed with AI-generated abuse images of its own pupils.[35] In July, the IWF's chief technology officer said he was “very uncomfortable” telling parents not to put pictures of their children on public display, but felt there was no other option, because “*there is no protection*” (Milmo, 2026).[36] Our own research in 2023-24 (Defend Digital Me, 2026) found pupil photographs lifted from school websites and local newspapers sitting in databases used to train AI and generate simulated images and video. Having concluded that a child's face in the digital environment is so high-risk that schools and parents should withhold it, policymakers might be expected to follow suit. Instead, age checks mandates require people to hand over exactly that data — a selfie, a live facial scan, video of an on-screen facial performance — to commercial third parties, often abroad. What is judged too hazardous for a school website will be made compulsory for all.

Open Rights Group argues the Online Safety Act left a critical technical-standards gap. By setting no mandatory minimum for age assurance and relying solely on existing data-protection law, it lets commercial drivers —convenience, cost and compliance— dictate how providers handle highly sensitive data. ORG calls for the Act to be amended to mandate a single, enforceable certification standard for all age-assurance tools: high data security, data and metadata minimisation, minimal retention, the least possible disclosure to platforms, a ban on secondary uses like profiling, and interoperability so users can pick their own trusted tool — favouring zero-knowledge techniques and cryptographic protection. In the UK, age-gating nonetheless appears a foregone conclusion.

The consultation closed on 26 May 2026 with no "due regard" or child rights impact assessment, and a junior minister had already said in April: "We are consulting on the mechanism, which is the right thing to do, but we are clear that under any outcome we will impose some form of age or functionality [restriction] on children under 16."

What is ultimately at stake is larger than children's access to social media or the nuances of "privacy preserving" technology. The global adoption of age assurance is the end of anonymous participation in the digital environment. Participation is becoming conditional on a verifiable, liveness-bound identity, checkable against a single authoritative, state-approved record. The infrastructure is owned by a consolidating industry of private providers; nation states supply its lifeblood in the identity data to verify against; and digital life is built around "know your customer".

European data protection lawfulness rests on the claim that age checks are privacy preserving, but this ignores the earlier part of the process that requires a digital ID. Others claim that processing a face or an iris is "not biometric data" because no name is attached at the instant of the scan, and because what is retained is a mathematical hash rather than the "actual" face print. The argument some AV companies make about age estimation is the argument others make about their iris-scanning orbs.

[35] https://www.theguardian.com/technology/2026/may/08/uk-schools-remove-pupils-photos-online-ai-blackmail-threat-grows
[36] https://www.nationalcrimeagency.gov.uk/news/new-guidance-for-parents-and-carers-as-ai-manipulated-images-of-children-become-a-growing-concern

In Europe, if regulators hold the line the AEPD and the Bavarian authority have drawn (Biometric Update, 2024), bodily data may be pushed to the limits of normalisation but keeps its status in law. If they do not, the special protection its drafters intended ends by definition, for everyone, and at rapid global scale if AV becomes a de facto norm for participation in digital life. Bodily data cannot be reset or easily replaced. If proof-of-liveness becomes required in an AI world, personal privacy in any meaningful sense is over. The Netherlands Digitalisation Minister Eddie van Marum has also expressed reservations, calling age verification an "intrusive" measure (Gkritsi, 2025). Estonia points to a better path. "*Estonia believes in an information society and including young people in the information society*," Liisa-Ly Pakosta, Minister of Justice and Digital Affairs, told POLITICO.

Excluding children from the digital environment, built on mandatory identity checks and routine biometric processing, does not create "safety." It creates a lifelong loss of identity control and of access to news and information, and it will damage youth democratic participation and trust in both public authorities and technology. This global rush to impose minimum ages in the name of children's online safety is the opposite of genuine child protection. It is already spreading beyond children and social media to gaming and to ever more situations offline. Unless policymakers stop the current push, it will affect a wide range of everyday activity for everyone and reach a tipping point beyond which control has been handed to these companies for good. Everyone online, including families and young people, should be fully in control of their online experiences and able to navigate the web according to their preferences and needs. There is a significant opportunity to empower users by design, with easy, effective opt-out rights and mandated granular user controls, without penalties or downgraded versions of services, or ID becoming a universal obligation.

Performative coercive “consent” on camera, and processing children's bodily data in ways beyond their understanding or lifelong control, must not become acceptable, normalised or routine in the name of safety. We encourage policymakers to promote enforcement action against unlawful platform design in line with the DSA and its data minimisation principles, to make full use of existing European data-protection laws, and to empower individuals and communities by investing in digital literacy to defend protection, participation *and* privacy, for life online, for everyone.

References

20 Minutes. (2026). France : nouvelle loi pour restreindre l'accès des -15 ans aux réseaux sociaux. https://www.20min.ch/fr/story/reseaux-sociaux-la-france-soumet-une-nouvelle-loi-pour-les-moins-de-15-ans-103632849

5Rights. (2026). Access limitations must be part of age-appropriate design, and effectively restrict companies from exploiting children https://5rightsfoundation.com/access-limitations-must-be-part-of-age-appropriate-design-and-effectively-restrict-companies-from-exploiting-children/

Access Now. (2026, February). #KeepItOn: Authorities must reverse social media shutdown order and restore access in Gabon. https://www.accessnow.org/press-release/keepiton-social-media-restore-access-in-gabon/

AFR. (2025). *Teens pile into Lemon8, Yope, Coverstar as social ban kicks in*. Australian Financial Review. https://www.afr.com/companies/media-and-marketing/teens-pile-into-lemon8-yope-coverstar-as-social-ban-kicks-in-20251209-p5nm88

Age Verification Providers Association. (2026). Biometrics, digital identity and GDPR: Why the AEPD decision matters for every European. https://avpassociation.com/thought-leadership/biometrics-digital-identity-and-gdpr-why-the-aepd-decision-matters-for-every-european/

Agencia Española de Protección de Datos. AEPD. (2026). *Resolución de procedimiento sancionador PS-00164-2025, Expediente EXP202317887* [Yoti Ltd]. https://www.aepd.es/documento/ps-00164-2025.pdf

Al Jazeera. (2021). FinGoPay used at Costcutter, Brunel University. Vein technology from Starla. https://www.youtube.com/watch?v=xjmfWWMoP5E

Allen, T. (2025). Age Assurance Technology Trial — Part C: Age Verification. Australian Dept of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts. https://www.ageassurance.com.au/report

Alston & Bird. (2025). California enacts digital age verification law. Privacy, Cyber & Data Strategy Blog. https://www.alstonprivacy.com/california-enacts-digital-age-verification-law/

Amnesty. (2025). Australia: Social Media Ban for Children and Young People an “Ineffective Quick Fix”. Amnesty International https://www.amnestyusa.org/press-releases/australia-social-media-ban-for-children-and-young-people-an-ineffective-quick-fix/

Amnesty International. (2026). Iran: Internet shutdown hides violations in escalating deadly crackdown on protesters. https://www.amnesty.org/en/latest/news/2026/01/internet-shutdown-in-iran-hides-violations-in-escalating-protests/

AppleWorld.Today. (2026). The UK plans to crack down on nude images regarding kids on Apple and Google platforms. https://appleworld.today/2026/09/the-uk-plans-to-crack-down-on-nude-images-regarding-kids-on-apple-and-google-platforms/

Arntz, P. (2026). Age verification vendor Persona left frontend exposed, researchers say. Malwarebytes. https://www.malwarebytes.com/blog/news/2026/02/age-verification-vendor-persona-left-frontend-exposed

Australian Child Rights Taskforce. (2024). Proposed social media bans for children under 16 [Letter]. https://apo.org.au/sites/default/files/resource-files/2024-10/apo-nid328608.pdf

Australian Payments Plus. (n.d.). ConnectID privacy policy. https://www.auspayplus.com.au/connectid-privacy-policy

ANZ Plus. (2026). ConnectID Terms and Conditions. https://www.anz.com.au/plus/terms-conditions/connectid-terms-conditions/

Back Youth Alliance. (2026). *Back Youth Alliance statement on the Government Consultation on online safety for young people* [Statement]. UK Youth. https://www.ukyouth.org/2026/01/back-youth-alliance-statement-on-the-government-consultation-on-online-safety-for-young-people/

Ball, J. (2025). *Free expression concerns over Online Safety Act's age verification requirements.* https://www.indexoncensorship.org/2025/07/free-expression-concerns-over-online-safety-acts-age-verification-requirements/

Barnardo's. (2026). Barnardo's responds to social media ban for under-16s. https://www.barnardos.org.uk/news/barnardos-responds-social-media-ban-under-16s

Bavarian Data Protection Authority (BayLDA). (2024). *Decision in the matter of the Worldcoin Foundation* (Decision 1594). https://www.edpb.europa.eu/system/files/2025-02/decision1594_0.pdf

BBC. (2026a). Read all the ministers' resignation letters in full (Jess Phillips, resigning as Minister for VAWG) https://www.bbc.co.uk/news/articles/cglpz89265jo

BBC. (2026b). Social media ban 'won't keep children safe', commissioner warns. BBC News. https://www.bbc.co.uk/news/articles/cvgj168m546o

Bellovin, S. M. (2025). Privacy-preserving age verification—and its limitations. Paper for the IAB/W3C Workshop on Age-Based Restrictions on Content Access. https://www.cs.columbia.edu/~smb/papers/age-verify.pdf

Big Brother Watch. (2023). *The threat to privacy posed by online age verification laws.* https://bigbrotherwatch.org.uk/press-coverage/context-by-reuters-the-threat-to-privacy-posed-by-online-age-verification-laws/

Biometric Update. (2024). *World does not comply with GDPR, says German regulator*. https://www.biometricupdate.com/202412/world-does-not-comply-with-gdpr-says-german-regulator

Biometric Update. (2026a). EU recommends white label age verification app, but member states are wary. https://www.biometricupdate.com/202604/eu-recommends-white-label-age-verification-app-but-member-states-are-wary

Biometric Update. (2026b). AEPD makes case for alternatives to EUDI Wallet, biometric authentication. https://www.biometricupdate.com/202607/aepd-makes-case-for-alternatives-to-eudi-wallet-biometric-authentication

Blazy, O. (2023). *Using PETs for age verification (and more)* [Conference presentation]. Brussels, Belgium. https://blazy.eu/PocBxl.pdf

Bradshaw, T. (2025). Meta adopts new age-check system to meet global child safety laws. Financial Times. https://www.ft.com/content/8164c36c-f224-4006-bc1f-561952c119b6

Branje, S., de Moor, E. L., Spitzer, J., and Becht, A. I. (2021). Dynamics of identity development in adolescence: A decade in review. *Journal of Research on Adolescence, 31*(4), 908–927. https://doi.org/10.1111/jora.12678

Bryant, M. (2025). Denmark plans social media ban for under-15s. The Guardian. https://www.theguardian.com/world/2025/oct/07/danish-pm-plans-to-ban-social-media-for-under-15s-warning-it-is-stealing-childhood

Bursztyn, L. et al. (2026). Why Bans Fail: Tipping Points and Australia's Social Media Ban. Becker Friedman Institute. https://bfi.uchicago.edu/wp-content/uploads/2026/04/BFI_WP_2026-57.pdf

Burt, C. (2026). Spain's AEPD fines Yoti $1.1M for biometric data handling violations. Biometric Update. https://www.biometricupdate.com/202603/spains-aepd-fines-yoti-1-1m-for-biometric-data-handling-violations

Butler, J. (1990). *Gender trouble: Feminism and the subversion of identity.* Routledge.

Cabinet Office. (2026). *Making public services work for you with your digital identity*. GOV.UK Consultation. https://www.gov.uk/government/consultations/making-public-services-work-for-you-with-your-digital-identity/making-public-services-work-for-you-with-your-digital-identity

Castells, M. (2010). *The power of identity* (2nd ed.). Wiley-Blackwell.

Castro, C. (2026). *The price to pay: This is why 400+ scientists are calling for a halt to mandatory age verification*. TechRadar. https://www.techradar.com/vpn/vpn-privacy-security/the-price-to-pay-this-is-why-400-scientists-are-calling-for-a-halt-to-mandatory-age-verification

Celi, S., Den Hartog, K., and Haddadi, H. (2025). *The limits of zero-knowledge for age-verification*. Brave Research. Brave. https://brave.com/blog/zkp-age-verification-limits/

CEPA. (2026). Mapping the spread of child safety rules (updated 25 September 2026). Center for European Policy Analysis. https://cepa.org/comprehensive-reports/mapping-the-spread-of-child-safety-rules/

Childnet and Nominet. (2025). Young people's use of VPNs. UK Safer Internet Centre. https://www.childnet.com/wp-content/uploads/2025/12/Young-peoples-use-of-VPNs.pdf

CommBank. (2026). *Secure Online Identity Check (to set up a bank account, later used as age assurance check)* https://www.commbank.com.au/banking/online-id-check.html

CommBank. (n.d.). What is ConnectID? https://commbank.com.au/digital-banking/connect-id.html

Computing. (2026). Apple brings age checks to UK devices. https://www.computing.co.uk/news/2026/legislation-regulation/apple-brings-age-checks-to-uk-devices

Cox, J. (2024). *ID Verification Service for TikTok, Uber, X Exposed Driver Licenses*. 404 Media. https://www.404media.co/id-verification-service-for-tiktok-uber-x-exposed-driver-licenses-au10tix/

CRIN. (2026). *The Big Debates* (to ban or not to ban; defining social media; VPNs for children). Child Rights International Network. https://home.crin.org/the-big-debates

Cybersecurity News. (2026). *Instagram Meta AI vulnerability allegedly enables password reset for accounts*. https://cybersecuritynews.com/instagram-meta-ai-vulnerability/

da Costa, C. (2026). *Kids Are Outsmarting Age Verification With Fake Mustaches*. Gadget Review. https://www.gadgetreview.com/kids-are-outsmarting-age-verification-with-fake-mustaches

Daily Tribune. (2026). Senate advances under-16 social media ban. (Philippines). https://tribune.net.ph/2026/10/02/senate-advances-under-16-social-media-ban

Data (Use and Access) Act 2025, c. 18. https://www.legislation.gov.uk/ukpga/2025/18/contents

Dawson, J. (2025). *Where in the tech stack should age assurance checks occur?* Yoti. https://www.yoti.com/blog/discussion-paper-where-in-the-tech-stack-should-age-assurance-sit-and-how-should-it-be-done/

Defend Digital Me. (2020). *The Learner Records Service data breach and ICO audit: A connected chronology*. https://defenddigitalme.org/the-learner-records-service-data-breach-and-ico-audit-a-connected-chronology/

Defend Digital Me. (2026). News: UK schools should remove pictures of pupils' faces from websites https://defenddigitalme.org/2026/05/09/news-uk-schools-should-remove-pictures-of-pupils-faces-from-websites/

Dienlin, T., Smahel, D., and Orben, A. (2025). Global analysis of perceived social media effects on well-being. Cyberpsychology, Behavior, and Social Networking. https://doi.org/10.1177/21522715251359204

Discord. (2025). *Update on a Security Incident Involving Third-Party Customer Service*. https://discord.com/press-releases/update-on-security-incident-involving-third-party-customer-service

DSIT. (2020). *VoCO (Verification of Children Online) Phase 2 report*. DSIT, GCHQ and the Home Office. https://assets.publishing.service.gov.uk/media/5faa9cffd3bf7f03a841cfc2/November_VoCO_report_V4__pdf.pdf

DSIT. (2026a). *Digital identity sectoral analysis report 2026*. https://www.gov.uk/government/publications/digital-identity-sectoral-analysis-report-2026/digital-identity-sectoral-analysis-report-2026

DSIT. (2026b) *Social media to be banned for under-16s in landmark government move to give kids their childhood back* [Press release]. GOV.UK. https://www.gov.uk/government/news/social-media-to-be-banned-for-under-16s-in-landmark-government-move-to-givekids-their-childhood-back

DSIT. (2026c). Children's circumvention behaviours online. Research conducted by BMG Research. https://www.gov.uk/government/publications/childrens-circumvention-behaviours-online/childrens-circumvention-behaviours-online archived at https://www.gov.uk/government/publications/childrens-circumvention-behaviours-online/childrens-circumvention-behaviours-online

EDPB. (2025). European Data Protection Board Statement 1/2025 on Age Assurance, adopted on 11 February 2025 (with minor changes to paragraphs 16 and 22(b) on 1 April 2025). https://www.edpb.europa.eu/system/files/documents/2025-04/edpb_statement_20250211ageassurance_v1-2_en.pdf

EFF. (2025a). *EFF Launches Age Verification Hub as Resource Against Misguided Laws.* https://www.eff.org/press/releases/eff-launches-age-verification-hub-resource-against-misguided-laws

EFF. (2025b). *Zero-knowledge proofs are not a solution to age-gating mandates*. https://www.eff.org/pages/zero-knowledge-proofs-are-not-solution-age-gating-mandates

EFF (2026). Meta's $17 Billion Settlement is a Bad Deal for Teens and All Social Media Users https://www.eff.org/deeplinks/2026/09/metas-17-billion-settlement-bad-deal-teens-and-all-social-media-users

Egbejule, E. (2026). 'Blatant disregard for rights': concern grows over Gabon's social media clampdown. https://www.theguardian.com/world/2026/may/13/concern-gabon-social-media-clampdown-human-rights

Entrust. (2024). *Entrust completes acquisition of Onfido* [Press release]. https://www.entrust.com/company/newsroom/entrust-completes-acquisition-of-onfido-creating-a-new-era-of-identity-centric-security

Equifax. (n.d.). Response to Ofcom consultation on guidance for service providers publishing pornographic content. Ofcom. https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/272586-consultation-guidance-for-service-providers-publishing-pornographic-content/responses/equifax.pdf

Erikson, E. H. (1968). *Identity: Youth and crisis*. Norton.

EPRS | European Parliamentary Research Service. (2026). *Virtual private networks and the protection of children online* (EPRS_ATA(2026)782618). European Parliament Think Tank. https://www.europarl.europa.eu/thinktank/en/document/EPRS_ATA(2026)782618 VPNs.

eSafety. (2026a). *Australia SMMA Compliance update: March 2026*. https://www.esafety.gov.au/sites/default/files/2026-03/SocialMediaMinimumAgeComplianceUpdateMarch2026.pdf

eSafety. (2026b). Early days, early insights: Understanding experiences of social media age restrictions at the three-month follow-up. https://www.esafety.gov.au/research/social-media-age-restrictions-evaluation/early-days-early-insights-three-months-report

Etchells, P. (2026). *Why is the term 'screen time' unhelpful?* – Teachappy interview with Professor Peter Etchells [Video]. YouTube. https://www.youtube.com/watch?v=ZXIy9HBu1y4

European Commission. (2025a). *The Commission makes available an age-verification blueprint*. https://digital-strategy.ec.europa.eu/en/news/commission-makes-available-age-verification-blueprint

European Commission. (2025b). Report on the application of Article 33 of Regulation (EU) 2022/2065 (COM(2025) 708 final). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2025:708:FIN

European Commission. (2026a). The Age Verification Manual — EU Digital Identity Wallet; EU age verification blueprint technical portal. https://ageverification.dev/

European Commission. (2026b). EU KIDS Act: Helping children navigate a safer online world. https://commission.europa.eu/news-and-media/news/eu-kids-act-helping-children-navigate-safer-online-world-2026-09-17_en

European Commission. (2026c). Proposal for a Regulation of the European Parliament and of the Council: EU KIDS Act, 'EU Keeping Internet Digital Spaces Accountable and Trustworthy' (COM(2026) 681 final, 2026/0286(COD)). https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52026PC0681

European Commission. (2026d). Commission seeks feedback on EU KIDS Act [Press release]. https://digital-strategy.ec.europa.eu/en/news/commission-seeks-feedback-eu-kids-act

European Digital Rights (EDRi). (2023–2025). Position papers on age verification and children's rights. https://edri.org/our-work/policy-paper-age-verification-cant-childproof-the-internet/

European Digital Rights (EDRi). (2025). Showing your ID to get online might become a reality: a closer look at the EU's new age verification app. https://edri.org/our-work/showing-your-id-to-get-online-might-become-a-reality-a-closer-look-at-the-eus-new-age-verification-app/

European Parliament. (2019). Harmful internet use — Part I: Internet addiction and problematic use (Lopez-Fernandez and Kuss; PE 624.249). STOA. https://doi.org/10.2861/315951

European Parliament. (2025). Children should be at least 16 to access social media, say MEPs [Press release]. https://www.europarl.europa.eu/news/en/press-room/20251120IPR31496/children-should-be-at-least-16-to-access-social-media-say-meps

Facebook. (2026). How using your ID to verify your age works on Facebook. Help Centre. https://www.facebook.com/help/671565971895323

Family Law. (2026). Government plans device-level restrictions on nude images for children. https://www.familylaw.co.uk/news_and_comment/government-plans-device-level-restrictions-on-nude-images-for-children

Fegert, J. M., and Melchior, M. (2026). *Child safety online: Protecting and empowering minors in a digital world. Report by the Co-Chairs of the Special Panel on Child Safety Online*. European Commission.

FinGo. (2022). CCO Simon Binns introducing FinGoVend at the recent Vendex Show https://www.youtube.com/watch?v=vdm1J0rx7AM

FIRE (n.d). FIRE statement on age-based restrictions on social media access https://www.fire.org/research-learn/fire-statement-age-based-restrictions-social-media-access

Foucault, M. (1977). *Discipline and punish: The birth of the prison* (A. Sheridan, Trans.). Pantheon Books.

Franceinfo. (2026). « Ne rien lâcher » : le gouvernement soumet une nouvelle version de l'interdiction des réseaux sociaux aux moins de 15 ans à la Commission européenne. https://www.franceinfo.fr/internet/reseaux-sociaux/le-gouvernement-soumet-une-nouvelle-version-de-l-interdiction-des-reseaux-sociaux-aux-moins-de-15-ans-a-la-commission-europeenne-apres-la-censure-d-un-precedent-texte-par-le-conseil-constitutionnel_8191493.html

Free Speech Coalition. (2026). Digital Age Assurance Act of 2026 [Bill tracker]. From the trade association of the adult entertainment industry based in the United States. https://action.freespeechcoalition.com/?p=4071

Freevacy. (2025). Age verification measures reduce UK porn traffic, VPN use spiked but now falling [Summary of The Guardian reporting on Ofcom's Online Nation report]. https://www.freevacy.com/news/the-guardian/age-verification-measures-reduce-uk-porn-traffic-vpn-use-spiked-but-now-falling/7000

G7. (2021). Carbis Bay G7 Summit communiqué: Our shared agenda for global action to build back better. GOV.UK. https://www.gov.uk/government/publications/carbis-bay-g7-summit-communique/carbis-bay-g7-summit-communique

G7. (2026, June 17). Leaders' call on a safer digital space for minors. Élysée. https://www.elysee.fr/en/G7evian/2026/06/17/leaders-call-on-a-safer-digital-space-for-minors

G7 Digital and Technology Ministers. (2021, April 28). G7 Internet Safety Principles (Annex 3 to the G7 Digital and Technology Ministerial Declaration). https://www.g7.utoronto.ca/ict/2021-Annex_3__Internet_Safety_Principles.pdf

G7 Digital and Technology Ministers. (2026, May 29). G7 Common Set of Principles defining a safer and more secure digital space for minors. Ministry of Internal Affairs and Communications, Japan. https://www.soumu.go.jp/main_sosiki/joho_tsusin/eng/pressrelease/2026/pdf/20260601_1_3.pdf

Galissaire, J., Gomez, S., and Böswald, L.-M. (2026). *AgeGate tracker* [Data set]. (Referenced in Appendix D). Retrieved July 9, from https://www.interface-eu.org/publications/agegate-tracker

Gamers Voice. (2026, October 4). Briefing: How the proposed EU KIDS Act could impact video games and the Stop Killing Games campaign. Gamers Voice Ltd. Licensed under CC BY 4.0.

Guardian, The (2021). NHS app storing facial verification data via contract with firm linked to Tory donors

https://www.theguardian.com/society/2021/sep/15/nhs-app-storing-facial-verification-data-via-contract-with-firm-linked-to-tory-donors

GB Group. (2020). *The GBG UK Education Data Set: The exclusive data that boosts millennials' onboarding rates*. https://web.archive.org/web/20191226143014/https://www.gbgplc.com/newsroom/the-exclusive-data-that-boosts-millennials-onboarding-rates/

Girlguiding.(2026). Social Media companies must do more to protect young people. https://www.girlguiding.org.uk/about-us/press-releases/social-media-companies-need-to-do-more-to-protect-young-people/

Google. (2026a). Access age-restricted content & features. Google Account Help. https://support.google.com/accounts/answer/10071085

Google. (2026b). Privacy Policy. https://policies.google.com/privacy?hl=en

Government Technology. (2026). No Tech Minister in Burnham's new government. https://governmenttechnology.co.uk/node/9313

Griset, P. (2026). Written question E-001251/2026 to the Commission on the AEPD ruling against Yoti. European Parliament. https://www.europarl.europa.eu/doceo/document/E-10-2026-001251_EN.html

Gkritsi, E. (2025). *Estonia speaks out against EU social media ban for kids*. POLITICO. https://www.politico.eu/article/estonia-liisa-ly-pakosta-breaks-the-ranks-on-minors-protection/

GMA News. (2026). Romulo eyes social media ban for kids 13 years and below; regulation for 14-17 teens. https://www.gmanetwork.com/news/topstories/nation/999539/total-social-media-ban-eyed-for-13-years-and-below-regulation-for-14-17-years-old/story/

Goffman, E. (1959). *The presentation of self in everyday life*. Doubleday.

GOV.UK One Login. (2025).DSIT: GOV.UK One Login - Liveness & Likeness Checks. https://www.gov.uk/algorithmic-transparency-records/dsit-gov-dot-uk-one-login-liveness-and-likeness-checks

Hadley, L. (2022). *For first time consumers can use digital IDs to buy vaping goods from a vending machine thanks to 1account*. UK Vaping Industry Association. https://ukvia.co.uk/for-first-time-consumers-can-use-digital-ids-to-buy-vaping-goods-from-a-vending-machine-thanks-to-1account/

Hall, S. (1996). Introduction: Who needs 'identity'? In S. Hall & P. du Gay (Eds.), *Questions of cultural identity* (pp. 1–17). SAGE.

Hansard. (2026). Science, Innovation and Technology Volume 789: debated on Wednesday 15 July 2026 https://hansard.parliament.uk/commons/2026-07-15/debates/9E099359-6B4F-4888-AD04-B21AE0236811/ScienceInnovationAndTechnology

HM Government. (2021). *Global Britain in a competitive age: The integrated review of security, defence, development and foreign policy* (CP 403). https://assets.publishing.service.gov.uk/media/60644e4bd3bf7f0c91eababd/Global_Britain_in_a_Competitive_Age-_the_Integrated_Review_of_Security__Defence__Development_and_Foreign_Policy.pdf

Holy See. (2026). *Message of His Holiness Pope Leo XIV for the 60th world day of social communications* https://www.vaticannews.va/en/pope/news/2026-01/pope-leo-xiv-messsage-world-day-social-communications-ai-human.html

Hooker, L. and Cress, L. (2026). UK social media ban for under-16s consultation begins. BBC News. https://www.bbc.co.uk/news/articles/cvg3vjkx9d7o

House of Commons. (2026). Children's Wellbeing and Schools Bill [Debate]. Hansard, HC Deb vol 784, col 699 (27 Apr 2026). https://hansard.parliament.uk/Commons/2026-04-27/debates/09FEE6E7-FF12-45A0-A358-1E18BE4A44DF/

Incode. (n.d.-a). Biometric data policy and notice. https://incode.com/?p=90921

Incode. (n.d.-b). Incode privacy notice: End users of Incode services. https://www.incode.com/privacy-policy

Incode. (2025, August 19). Incode acquires AuthenticID [Blog post]. https://www.incode.com/blog/incode-acquires-authenticid

Information Commissioner's Office (ICO). (2022). Age assurance and biometric data — Regulatory Sandbox report on Yoti. https://ico.org.uk/for-organisations/advice-and-services/regulatory-sandbox/

Information Commissioner's Office (2024). Age assurance report: Technical annex. (IFF Research) https://ico.org.uk/media2/migrated/4030925/20240704-ico-age-assurance-report-technical-annex.pdf

Inside Privacy. (2026). EU sets the clock on age verification: Rollout urged by end-2026. Covington and Burling.https://www.insideprivacy.com/minors/eu-sets-the-clock-on-age-verification-rollout-urged-by-end-2026/

Intellicheck. (2022). Q&A with Digital Identity Pioneer Reed Taussig. https://www.intellicheck.com/resource-library/qa-with-reed-taussig

Internet Society Pulse. (2026). Social media blocked in Gabon, February 2026. https://pulse.internetsociety.org/en/shutdowns/social-media-blocked-in-gabon-february-2026/

IPPR. (2026). Vast majority support clampdown on social media to protect children, IPPR finds [Press release]. Institute for Public Policy Research. Polling by YouGov, fieldwork 10–11 June 2026. https://www.ippr.org/media-office/vast-majority-support-clampdown-on-social-media-to-protect-children-ippr-finds

iProov. (n.d). https://www.iproov.com/liveness-detection archived at https://web.archive.org/save/https://www.gov.uk/algorithmic-transparency-records/dsit-gov-dot-uk-one-login-liveness-and-likeness-checks

iProov. (2020). *iProov introduces flexible biometric customer authentication for global enterprises* [Press release]. Business Wire. https://businesswire.com/news/home/20201006005635/en/iProov-Introduces-Flexible-Biometric-Customer-Authentication-for-Global-Enterprises

Ipsos. (2026). 2 in 3 support social media ban for under-16s, though support softens if it requires nationwide ID verification. Online survey of 1,082 British adults aged 18–75, conducted 19–22 June 2026.

https://www.ipsos.com/en-uk/2-3-support-social-media-ban-under-16s-though-support-softens-if-it-requires-nationwide-id

ISO/IEC 27566-1:2025. (2025). ISO. https://www.iso.org/standard/88143.html

ISPreview. (2026). Government look to impose mandatory age verification on UK smartphones. https://www.ispreview.co.uk/index.php/2026/09/government-look-to-impose-mandatory-age-verification-on-uk-smartphones.html

IWF. (2026a). 'Self-generated' imagery

https://www.iwf.org.uk/annual-data-insights-report-2025/hashing-insights/self-generated-imagery/

IWF. (2026b). Sexually coerced extortion. Annual Data and Insights Report 2025. Internet Watch Foundation. https://www.iwf.org.uk/annual-data-insights-report-2025/emerging-and-persistent-harms/sexually-coerced-extortion

Joint statement (2026). Joint statement of security and privacy scientists and researchers on Age Assurance. https://csa-scientist-open-letter.org/ageverif-Feb2026

Journal News. (2026). House panel advances stricter safeguards for children on social media, online gaming. https://journalnews.com.ph/?p=399014

k-ID. (2025). Newly launched OpenAge Initiative introduces AgeKey [Press release].

k-ID (2025b) Age Assurance for Mobile Games https://product.k-id.com/p/age-assurance-for-mobile-games

k-ID (2026a) Security Center for k-ID and the OpenAge Initiative: List of sub-processors https://security.k-id.com/subprocessors

k-ID (2026b) *Age assurance for High Risk Features* https://docs.k-id.com/cdk/age-assurance/

k-ID. (2026c). Privacy Policy. Last updated 10 July 2026. https://k-id.com/privacy-policy

k-ID (n.d) Age Verification Methods https://docs.k-id.com/concepts/verification-methods/

Kentucky Office of the Attorney General. (2026). The Meta consent judgment: Briefing prepared for the Kentucky General Assembly. https://apps.legislature.ky.gov/CommitteeDocuments/366/44471/09%2015%202026%206A.%20Duncan%20-%20Meta%20Settlement%20Legislative%20Briefing.pdf

Kick. (2026). Age Assurance FAQs. https://help.kick.com/en/articles/12947994-age-assurance-faqs

King and Spalding. (2026). When regulators build the age verification app. https://www.kslaw.com/news-and-insights/when-regulators-build-the-age-verification-app-it-is-not-necessarily-for-the-greater-good

Krishnan-Barman, S., Kaldor, I., Fenn, L., Woodhouse, L., Persian, R., and Kolker, E. (2025). *A behavioural audit of online services in the UK: Thematic report*. Ofcom. https://www.ofcom.org.uk/siteassets/resources/documents/online-safety/research-statistics-and-data/online-services-research/a-behavioural-audit-of-online-services-in-the-uk_thematic-report.pdf

Kumara, B., Hooper, M., Maple, C., Hobson, T., and Crowcroft, J. (2023). *Redactable signature schemes and zero-knowledge proofs: A comparative examination for applications in decentralized digital identity systems*. arXiv. https://arxiv.org/abs/2310.15934

Lee, C., Kim, H., and Hong, A. (2017). Ex-post evaluation of illegalizing juvenile online games after midnight: the South Korea shutdown policy. Telematics and Informatics, 34(8). https://doi.org/10.1016/j.tele.2017.07.006

LexisNexis Risk Solutions. (2025). LexisNexis Risk Solutions completes IDVerse acquisition https://risk.lexisnexis.com/about-us/press-room/press-release/20250219-idverse-acquisition

Lin, Y., Zhang, C., Xiao, M., Cranor, L. F., and Scheffler, S. (2026). What adults will (and won't) do to prove their age: Empirical evidence from a deceptive web experiment. In Proceedings of the 35th USENIX Security

Symposium (pp. 5573–5592). USENIX Association. https://www.usenix.org/conference/usenixsecurity26/presentation/lin-yanzi

Livingstone, S. Stoilova, M. and Nandagiri, R. (2019) Children's data and privacy online: Growing up in a digital age. An evidence Review. London: London School of Economics and Political Science. https://www.lse.ac.uk/media-and-communications/assets/documents/research/projects/childrens-privacy-online/Evidence-review-Executive-summary.pdf

Livingstone, S. et al. (2024). The best interests of the child in the digital environment. LSE. https://researchonline.lse.ac.uk/id/eprint/122492/3/Best_Interests_of_the_Child_FINAL.pdf

Livingstone, H. (2026). How does Australia's under-16 social media ban work? BBC News. https://www.bbc.com/news/articles/cwyp9d3ddqyo

Luria, M., Bhatia, A., and boyd, d. (2026, October 16). Kids are outsmarting social media bans. That's not a problem we need to solve [Opinion]. MIT Technology Review. https://www.technologyreview.com/2026/10/16/1146834/kids-are-outsmarting-social-media-bans-thats-not-a-problem-we-need-to-solve/

Marwick, A. E., and boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society, 13*(1), 114–133. https://doi.org/10.1177/1461444810365313

MediaNama. (2025). Limits of age assurance tools [Report on Australia's Age Assurance Technology Trial]. https://www.medianama.com/2025/09/223-limits-age-assurance-tools/

Mediazona. (2026). Russia's internet censorship in 2026: VPN crackdowns, mobile shutdowns, Telegram blocks and the state messenger Max. https://en.zona.media/article/2026/04/07/russian_internet_censorship_2026

Meta. (2022). *Introducing new ways to verify age on Instagram*. Meta Newsroom. https://about.fb.com/news/2022/06/new-ways-to-verify-age-on-instagram/

Meta. (2025). Working With Parents to Enroll Teens Into Teen Accounts. https://about.fb.com/news/2025/04/meta-parents-new-technology-enroll-teens-teen-accounts

Meta. (2026a, May 5). New AI-powered age assurance measures to place teens in age-appropriate experiences. Meta Newsroom. Archived 6 May 2026: https://web.archive.org/web/20260506000615/https://about.fb.com/news/2026/05/ai-age-assurance-teens/ (the reference to height and bone structure has since been removed from the live page, https://about.fb.com/news/2026/05/ai-age-assurance-teens/)

Meta. (2026b). *How Meta uses information for generative AI models*. https://www.facebook.com/privacy/genai/

Meta. (2026c). Meta Privacy Policy. https://mbasic.facebook.com/privacy/policy/printable/version/23954169707588482/

Meta. (2026d). *How does video selfie age verification work on Instagram?* https://help.instagram.com/292521846402382

Milmo, D. (2026). UK parents warned over posting images of children amid AI sexual abuse fears. The Guardian.https://www.theguardian.com/society/2026/jul/03/ai-prey-watchdogs-telling-parents-protect-children-nudification-apps

Minocha, S., Sheridan, I., Oppenheimer, H., Pearce, P., and Specter, M. A. (2026). Papers, please: A first look at age verification on the web. In *Proceedings of the 47th IEEE Symposium on Security and Privacy*. IEEE. https://doi.org/10.5281/zenodo.19515393

Molly Rose Foundation. (2026). *Australia's social media ban — is it working?* https://mollyrosefoundation.org/wp-content/uploads/2026/04/MRF_Australia-Social-Media-Ban-Research_Briefing-April-26.pdf

Molly Rose Foundation. (2026b). *More than 60% of Australian children still using social media despite ban for under-16s, research shows*

https://mollyrosefoundation.org/more-than-60-of-australian-children-still-using-social-media-despite-ban-for-under-16s-research-shows/

Mukasa, D., Garcia, L., and Hosein, G. (2022). Modern ID systems: Watching a global identity crisis unfold. *Al Jazeera.* https://www.aljazeera.com/opinions/2022/9/16/modern-id-systems-watching-a-global-identity-crisis-unfold

NAB. (2026). ConnectID customer terms and conditions. https://www.nab.com.au/help-support/connectid-your-digital-identity/terms-conditions

Nandy, L. (2026). Secretary of State statement on protecting children online [Oral statement to Parliament]. GOV.UK. https://www.gov.uk/government/speeches/lisa-nandys-statement-on-protecting-children-online

Neowin. (2026). *People are using prompt injection to trick Meta's AI into handing over Instagram accounts.* https://www.neowin.net/news/people-are-using-prompt-injection-to-trick-metas-ai-into-handing-over-instagram-accounts/

NetChoice. (2025). Court halts Georgia's ID-for-speech law for NetChoice [Press release]. https://netchoice.org/court-halts-georgias-id-for-speech-law-for-netchoice/

Niestadt, M. (2026). *Debate on setting a minimum age for social media* (PE 789.316). European Parliamentary Research Service. https://www.europarl.europa.eu/RegData/etudes/BRIE/2026/789316/EPRS_BRI(2026)789316_EN.pdf

NIST. (2014). Face Recognition Vendor Test (FRVT): Performance of automated age estimation algorithms (NIST IR 7995). https://doi.org/10.6028/NIST.IR.7995

Notley, T., Chambers, S., Dezuanni, M., Park, S., and Lee, J. Y. (2026). *Social media ban: The impact on young people's news engagement.* Western Sydney University, Queensland University of Technology, and University of Canberra. https://apo.org.au/node/334363

NSPCC. (2026). Government must stand up to big tech or under 16s social media ban is coming, warns NSPCC CEO. https://www.nspcc.org.uk/about-us/news-opinion/2026/social-media-ban-three-actions/

OAIC. (2025–2026). Privacy Guidance on Part 4A (Social Media Minimum Age) of the Online Safety Act 2021; new proportionate age-assurance guidance. https://www.oaic.gov.au/privacy/privacy-legislation/related-legislation/social-media-minimum-age

OAIC. (2026). Privacy Commissioner publishes new guidance to ensure proportionate age assurance as a gateway to access online experiences [Media release]. Office of the Australian Information Commissioner. https://www.oaic.gov.au/news/media-centre/privacy-commissioner-publishes-new-guidance-to-ensure-proportionate-age-assurance-as-a-gateway-to-access-online-experiences

O'Flaherty, M. (2026). Regulate online platforms, not children [Statement]. Council of Europe Commissioner for Human Rights. https://www.coe.int/en/web/portal/-/regulate-online-platforms-not-children

Ofcom. (2026a). *Experiences of age assurance, use of VPNs, and attitudes to the protection of children measures*. https://www.ofcom.org.uk/online-safety/protecting-children/experiences-of-age-assurance-use-of-vpns-and-attitudes-to-the-protection-of-children-measures

Ofcom. (2026b). *Report on the use of age assurance: Report prepared under s157 of the Online Safety Act 2023* https://assets.publishing.service.gov.uk/media/6a5770ae6d6d5c7e2845c9d0/Web-_accessible_Report_on_the_use_of_age_assurance_1.pdf Also available from Ofcom, archived at: https://web.archive.org/web/20260804033534/https://www.ofcom.org.uk/siteassets/resources/documents/online-safety/information-for-industry/age-assurance-report-2026/report-on-the-use-of-age-assurance.pdf?v=423162

Ofcom. (2026c). Age checks helping make online experiences safer for UK children, but job not done and tech industry must act to strengthen protections. https://www.ofcom.org.uk/online-safety/protecting-children/age-checks-helping-make-online-experiences-safer-for-uk-children-but-job-not-done-and-tech-industry-must-act-to-strengthen-protections

OHCHR. (1989). Convention on the Rights of the Child. Adopted by UN General Assembly resolution 44/25 of 20 November 1989. ("UNCRC"). https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-child

OHCHR. (2026a). Protection of children online needs to be done right – UN Human Rights Office issues guidelines for child safety online https://www.ohchr.org/en/press-releases/2026/05/protection-children-online-needs-be-done-right-un-human-rights-office-issues

OHCHR. (2026b). Office of the United Nations High Commissioner for Human Rights. (2026). *Getting children's safety online right*. https://www.ohchr.org/sites/default/files/documents/issues/children/getting-childrens-safety-online-right.pdf

Ondato (n.d) https://ondato.com/blog/liveness-detection/

Open Rights Group. (2025). Regulating age verification: Keeping users safe online. https://www.openrightsgroup.org/publications/regulating-age-verification/

Open Rights Group. (2026). Regulating age verification; Break privacy to make privacy? https://www.openrightsgroup.org/blog/break-privacy-to-make-privacy-age-verification-isnt-the-answer/

Orben, A., Przybylski, A. K., Blakemore, S. J., and Kievit, R. A. (2022). Windows of developmental sensitivity to social media. Nature communications, 13(1), 1649. https://doi.org/10.1038/s41467-022-29296-3

Parliament of Australia. (2024). *Online Safety Amendment (Social Media Minimum Age) Bill 2024: Supplementary explanatory memorandum*. Parliament of the Commonwealth of Australia. https://parlinfo.aph.gov.au/parlInfo/download/legislation/ems/r7284_ems_ca8c5dba-cc80-4846-92f5-bea56885dcdf/upload_pdf/Supplementary%20EM_Online%20Safety%20Amdt%20(Social%20Media%20Minimum%20Age)%20Bill%202024.pdf

Parliament of Australia. (2026). Online Safety Amendment (Strengthening Enforcement for the Social Media Minimum Age) Bill 2026 (Act No. 83, 2026). https://www.aph.gov.au/Parliamentary_Business/Bills_Legislation/Bills_Search_Results/Result?bId=r7512

Persona. (n.d). Processor Privacy Policy Last updated: April 9, 2026 https://withpersona.com/legal/privacy-policy#age-assurance

Persona. (2025). What is liveness detection? Persona https://withpersona.com/blog/what-is-facial-liveness-detection/

Persona. (2026). What is selfie identity verification, and how does it work? https://withpersona.com/blog/what-is-selfie-identity-verification-and-how-does-it-work

Persona. (2026b). FAQ: For Reddit Users Using the Persona Service. https://help.withpersona.com/articles/7F6BaF9h8Fxf0XWkwQscXN/

Persona. (2026c). Post-incident review: source map exposure on non-production subdomain. https://withpersona.com/blog/post-incident-review-source-map-exposure-non-production-subdomain

Pillitteri, P. (2026). EU age verification app: How the anonymous mini-wallet works. https://pasqualepillitteri.it/en/news/949/eu-age-verification-app-anonymous-mini-wallet-2026-guide

Ping Identity. (n.d.). PingOne Verify and face verification [Product documentation]. https://docs.pingidentity.com/pingone/identity_verification_using_pingone_verify/p1_verify_face_verification.html

Plackett R., Sheringham J., and Dykxhoorn J. (2023). The Longitudinal Impact of Social Media Use on UK Adolescents' Mental Health: Longitudinal Observational Study. J Med Internet Res. 2023 Mar 24;25:e43213. doi: 10.2196/43213. Erratum in: J Med Internet Res. 2023 Mar 31;25:e47678. doi: 10.2196/47678. PMID: 36961482; PMCID: PMC10132039.

Police Professional. (2025). NCA launches online campaign to tackle 'sextortion' among teenage boys. https://policeprofessional.com/?p=141976

Politico (2026). Brussels launched an age checking app. Hackers say it takes 2 minutes to break it. https://www.politico.eu/article/eu-brussels-launched-age-checking-app-hackers-say-took-them-2-minutes-break-it/

PPC Land. (2026). EU age-check app: 9 files released, no privacy review, claims unproven. https://ppc.land/eu-age-check-app-9-files-released-no-privacy-review-claims-unproven/ based on https://www.linkedin.com/posts/morten-larsen-80542b335_ease-2026-2179-disclosure-and-confirmatory-ugcPost-7484217865672380416-43FU

Privacy International. (2019). Yoti — the trouble with age estimation and biometrics. https://privacyinternational.org/

PrivateID. (2023). Secure Age Estimation privacy policy. https://g.privateid.com/privacy

Privately SA. (2026). Privacy Policy. https://www.privately.eu/privacy-policy-en

Public First (January 2026). Polling https://publicfirsttech.substack.com/p/to-ban-or-not-to-ban

Reclaim The Net. (2025). California Digital Age Assurance Act (AB 1043) [Bill tracker]. https://reclaimthenet.org/bill/california-ab-1043

Reddit. (2026). How does age verification work in Australia? Reddit Help. https://support.reddithelp.com/hc/en-us/articles/47163443918228

Regula. (n.d.). Liveness check on web service [Regula Face SDK documentation]. https://docs.regulaforensics.com/develop/face-sdk/web-service/development/usage/liveness/

Rodgers, S. (2025). Wikipedia's Jimmy Wales warns of "political showdown" with the UK government over Online Safety Act. *PoliticsHome*. https://www.politicshome.com/news/article/wikipedia-v-uk-government-jimmy-wales-warns-political-showdown-online-safety-act

Safecast. (2021). Evidence to parliamentary committee. https://committees.parliament.uk/writtenevidence/38543/pdf/

Shrivastava, R. (2025). AI Is Making The Internet's Bot Problem Worse. Forbes. https://www.forbes.com/sites/rashishrivastava/2025/04/30/

SkyNews (2026). June 8th, Jess Phillips Interview https://news.sky.com/share/13551922

Snap Inc. (n.d.). Snap and ads privacy: European contextual advertising. https://values.snap.com/privacy/ads-privacy (archived at https://web.archive.org/web/20261015120937/https://values.snap.com/privacy/ads-privacy)

Snapchat. (2026). Account Restrictions for Australia Users. https://help.snapchat.com/hc/en-us/articles/42550358600084

SQ Magazine. (2026). *Meta fixes Instagram AI flaw used in account takeovers.* https://sqmagazine.co.uk/meta-fixes-instagram-ai-flaw-account-takeovers/

Stripe (2026) Learn about the different verification checks supported by Stripe Identity https://docs.stripe.com/identity/verification-checks?type=selfie

SkyNews. (2026). Ex-minister Jess Phillips gives first interview since resigning from Starmer's government https://www.youtube.com/shorts/S0evMGuT6q0

Tech Policy Press Global Social Media Age Restriction Tracker. (n.d). Tracking legislation to restrict or ban teen social media access worldwide https://social-media-age-tracker.onrender.com/

TechRepublic. (2026). Millions of UK iPhone users will need to verify their age — here's why. https://www.techrepublic.com/de/article/news-apple-iphone-age-verification-uk/

The Next Web. (2026). UK will legislate to make Apple and Google block nude images for children. https://thenextweb.com/news/uk-law-apple-google-block-nude-images-children

The Stack. (2026, June 8). UK calls for “device-based” nude controls. https://www.thestack.technology/uk-device-based-nude-controls-child-safety/

TikTok. (2025–2026). How TikTok is complying with a new government law in Australia; Confirm/Technology to determine your age. https://newsroom.tiktok.com/how-tiktok-is-complying-with-a-new-government-law-in-australia

TikTok. (2026a). Confirm your age on TikTok. https://www.tiktok.com/support/faq_detail?id=7543604789788039686

TikTok. (2026b). Technology to determine your age on TikTok. https://www.tiktok.com/support/faq_detail?id=7543604782246959621

TikTok Live (2026c) Age requirements for TikTok Live https://www.tiktok.com/support/faq_detail?id=7581820708762753548&category=web_tiktok_live

Tombs, R. (2025a). ‘Yoti reaches EBITDA profitability, expects further facial age estimation growth.’ Biometric Update. https://www.biometricupdate.com/202504/yoti-reaches-ebitda-profitability-expects-further-facial-age-estimation-growth

Tombs, R. (2025b) *Thoughts from our CEO: Yoti reaches 20 million global downloads*. https://www.yoti.com/blog/thoughts-from-our-ceo-october-2025/

Tombs, R. (2025c). Thoughts from our CEO: Despite opposition, most Brits support online age protections. Yoti. https://www.yoti.com/blog/thoughts-from-our-ceo-september-2025/

Tombs, R. (2026, May 22). *An open letter to Georgia Institute of Technology and University of California, Irvine requesting retraction and correction of false statements*. Yoti. https://www.yoti.com/blog/open-letter-to-georgia-institute-of-technology-university-of-california-irvine-requesting-retraction-correction-false-statements/

Tools for Humanity. (2026). *Orb: Proof of human hardware and software*. Retrieved June 1, 2026, from https://www.toolsforhumanity.com/orb

Tuta. (2026). Age verification: What countries require you to perform ID checks? https://tuta.com/blog/age-verification-kills-anonymity

Twitch. (2026). Age Verification on Twitch. https://help.twitch.tv/s/article/age-verification

UNICEF. (2025a). Drawing a line in digital spaces: Policy note on age-based restrictions for social media. https://www.unicef.org/media/170666/file/Policy%20note_age%20restrictions%20social%20media-new.pdf

UNICEF. (2025b). Age restrictions alone won’t keep children safe online [Statement, 9 December 2025]. https://www.unicef.org/press-releases/age-restrictions-alone-wont-keep-children-safe-online

UNICEF. (2026). Drawing a Line in Digital Spaces: Age-based restriction of social media, Rapid Analysis https://www.unicef.org/media/179716/file/Drawing-line-digital-spaces-Rapid-analysis-2026.pdf

UNCRC. (2013). General comment No. 14. Best Interests of the Child (Article 3). https://www2.ohchr.org/english/bodies/crc/docs/gc/crc_c_gc_14_eng.pdf

UNCRC. (2021) General comment No. 25 on children’s rights in the digital environment. https://www.unicef.org/bulgaria/en/media/10596/file

University of Glasgow. (2026). Is a ban on social media the answer to tackling online harms? What do young people really think? [Student blog]. https://www.gla.ac.uk/schools/socialpolitical/study/student-assessed-blogs/headline_1277930_en.html

US Department of Justice. (2020). *International statement: End-to-end encryption and public safety* [Press release]. https://www.justice.gov/opa/pr/international-statement-end-end-encryption-and-public-safety

van der Hof, S., Hannema, T. S. P., and Liefaard, T. (2020). The child's right to protection against economic exploitation in the digital world. Int'l Journal of Children's Rights, 28(4). https://doi.org/10.1163/15718182-28040003

Veale, M. (2019). A Critical Take on the Policy Recommendations of the EU High-Level Expert Group on Artificial Intelligence. *European Journal of Risk Regulation*. 2020;11(1):e1. doi:10.1017/err.2019.65

Veratad. (2026). Privacy Policy. https://veratad.com/privacy-policy

Verifymy. (2025). Email address age estimation FAQs. https://verifymy.io/age-verification-and-estimation/email-address-age-estimation-faqs/

Verifymy. (2026a). Age assurance Privacy Policy. https://verifymy.io/age-verification-and-estimation/age-assurance-privacy-policy/

Verifymy. (2026b). Email address age estimation. https://verifymy.io/age-verification-and-estimation/email-address-age-estimation/

Verifymy. (2026c). *Email-based age checks: The benchmark for frictionless age assurance.* Published April 2026. https://verifymy.io/wp-content/uploads/2026/06/Whitepaper_Email-based-age-checks_Innovative-age-assurance_Jun26.pdf

Verifymy. (2026d). Verifymy Privacy Policy for age assurance for Google https://verifymy.io/age-verification-and-estimation/google-age-assurance-privacy-policy/

Vohra, A. (2026). European nations mull social media ban for children. DW. https://www.dw.com/en/european-nations-mull-social-media-ban-for-children/a-75841184

Wells, A. (2026). Australia doubles down on social media law [Media release]. Minister for Communications. https://minister.infrastructure.gov.au/wells/media-release/australia-doubles-down-social-media-law

Wertheimer, T. (2026). Spain announces plans to ban social media for under-16s. BBC News. https://www.bbc.co.uk/news/articles/c5y2nddvmryo

Westpac. (2026). Westpac ConnectID Customer Terms and Conditions. https://www.westpac.com.au/about-westpac/innovation/connectID-your-digital-identity/connectID-customer-terms-and-conditions/

Whittaker, Z. (2025). How we found TeaOnHer spilling users' driver's licenses in under 10 minutes. TechCrunch. https://techcrunch.com/2025/08/13/how-we-found-teaonher-spilling-users-drivers-licenses-in-less-than-10-minutes/

Williams, T. (2025). 'Unbelievably easy': Online age checks tricked in UK. Information Age. https://ia.acs.org.au/article/2025/unbelievably-easy-online-age-checks-tricked-in-uk.html

Wood, S. (2024). Impact of regulation on children's digital lives. Digital Futures for Children centre, LSE and 5Rights Foundation. https://www.lse.ac.uk/digital-futures-for-children/digitalfutures-assets/digitalfutures-documents/Impact-of-regulation-on-children-DFC-Research-report-May-2024.pdf

Wood, S. (2026). Impact of regulation on children's digital lives: Phase 2. Digital Futures for Children centre, LSE. https://www.lse.ac.uk/news/more-work-needed-for-regulation-to-drive-systemic-change-in-child-online-safety-and-privacy-protections

Woodhouse, J. (2026). Proposals to ban social media for children (Research Briefing CBP-10468). House of Commons Library. https://commonslibrary.parliament.uk/research-briefings/cbp-10468/

Woodley, G., See, H. W., O'Neill, B., Green, L., Staksrud, E., and Haskell-Dowland, P. (2025). Australian teen voices on age verification and age assurance measures. Policy & Internet, 17(4), e70019.

WP29. (2007). Article 29 Data Protection Working Party. *Opinion 4/2007 on the concept of personal data* (01248/07/EN, WP 136). European Commission. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2007/wp136_en.pdf

X. (2026a). Age Assurance on X. Help Center. https://help.x.com/en/rules-and-policies/age-assurance

X. (2026b). Australia Online Safety Information. https://help.x.com/en/rules-and-policies/australia-resources

Yoti. (2021). Written evidence to the Online Safety Bill Committee. https://committees.parliament.uk/writtenevidence/38770/html/

Yoti. (2025). Thoughts from our CEO: Yoti reaches 20 million global downloads https://www.yoti.com/blog/thoughts-from-our-ceo-october-2025/

Yoti. (2026a). Age Verification Privacy Policy; Thoughts from our CEO (Jan 2026). https://www.yoti.com/blog/thoughts-from-our-ceo-january-2026/

Yoti. (2026b). Age Verification Privacy Policy. https://www.yoti.com/privacy/age-verification/

Yoti. (2026c). Yoti's response to AEPD sanctions and fine. https://www.yoti.com/blog/yoti-response-to-aepd-sanctions-and-fine/

Yoti. (2026d). *Withdrawing the Yoti ID app from Spain* archived at https://web.archive.org/web/20260905165441/https://www.yoti.com/blog/withdrawing-yoti-id-app-spain/

Yoti. (2026e). My Liveness Checks https://defenddigitalme.org/wp-content/uploads/2026/07/Yoti-MyFace-Liveness-Paper-Jan-2026-public.pdf

Appendix A — What age-check methods are offered (SMMA, Australia)?

Platform	Owner	Methods offered	Third parties and their sub-processors	Notes
Facebook	Meta	Video selfie age estimation; ID upload (passport/driving licence)	Yoti (facial estimation + ID verification)	Meta retains facial/ID image up to 30 days; Yoti deletes after check. AgeKey (OpenAge Initiative) announced for 2026.
Instagram	Meta	Video selfie age estimation; ID upload; continuous AI scanning of photos, posts, captions, bios and comments for age signals	Yoti (facial estimation + ID verification)	All Australian accounts continuously scanned; data used to train Meta's generative AI models. AgeKey (OpenAge Initiative) announced for 2026.
Threads	Meta	Video selfie age estimation; ID upload	Yoti (facial estimation + ID verification)	Same Meta/Yoti pipeline as Facebook/Instagram. AgeKey (OpenAge Initiative) announced for 2026.
TikTok	ByteDance	Facial age estimation; credit card authorisation; ID check; ML behavioural assessment (profile, device, videos, interactions)	Yoti (face); Incode (ID + real-time selfie for TikTok LIVE); unnamed 'service providers' for ID and credit card checks	TikTok does not name third parties for ID or card checks.
Snapchat	Snap Inc.	ID upload; facial age estimation; bank-based age check (over-16 answer from the user's bank); credit card check; AgeKey token	k-ID (manager); range of sub-processors >> Veratad (ID); Privately SA (face, on-device); Australian Payments Plus / ConnectID (bank); Stripe (credit card); OpenAge Initiative (AgeKey)	Veratad's privacy policy is vague on data use. Privately SA does not currently train its AI on user data but says this 'might change in the future'.
Reddit	Reddit Inc.	ID upload; facial selfie age estimation (AU users: ID + face only)	Persona (plus undisclosed sub-processors)	Persona may also collect device, geolocation, usage and wireless data and cross-reference government/credit/utility/telco databases.
X	X Corp.	Account signals (self-attested age, verified status, creation date); email-based age checks (estimation/inference); facial age estimation; ID check	Stripe and Persona (ID, for X Premium); other providers undisclosed	X does not disclose its email-based age checks (estimation/inference) provider or confirm who handles age-assurance ID/face checks.
YouTube	Google (Alphabet)	Government ID upload; facial age estimation; email-based age checks (estimation/inference); credit card authorisation; ML assessment of account activity	Google (ID + card); Private Identity LLC / PrivateID (face, on-device); Verifymy (email)	Google makes no clear commitment to delete ID data; broad privacy policy covers a wide range of activity data.
Twitch	Amazon	Facial age estimation	k-ID (manager) → Privately SA (on-device facial scan)	No ID route at the point of the check; users who dispute the result can appeal to Twitch Support with unspecified documentation; daily attempt limits apply.

Platform	Owner	Methods offered	Third parties and their sub-processors	Notes
Kick	Stake.com (Easygo)	Facial age estimation; ID check; bank-based age check; credit card check; AgeKey token	k-ID (manager) → range of sub-processors Privately SA, or Veratad, or Australian Payments Plus, or Stripe; OpenAge Initiative (AgeKey)	'Not all verification methods are available to all Viewers' — users don't know in advance what data they'll need to provide.

Table A — Types of age check tools offered by platforms under Australia's Social Media Minimum Age Act.

Appendix B — About the third-party age-check providers

Provider (legal entity)	HQ / background	Uses biometric data	Methods offered	Platforms using it
Australian Payments Plus (ConnectID) — Australian Payments Plus Ltd	Sydney, Australia. Member-owned entity created in 2021 by the merger of BPAY, eftpos and NPP Australia.	No	Bank-based age verification: partners with Australian banks (CommBank, NAB, ANZ Plus, Westpac) to confirm an age threshold from existing bank records. Returns pass/fail without sharing ID documents. The identity document step sits upstream, at account opening.	Snapchat (via k-ID); Kick (via k-ID)
AU10TIX — AU10TIX Ltd	Parent company ICTS International N.V. in Amsterdam; operating HQ and R&D in Hod Hasharon, Israel; offices in NY, London, Singapore. Founded 2002.	Yes (server-side document authentication, facial match and liveness)	ID document verification, facial match, liveness detection, organised-fraud detection. 404 Media found AU10TIX left TikTok users' photos and driving licences exposed for almost a year (2024).	Variable by client (incl. TikTok partners)
Incode — Incode Technologies, Inc.	San Francisco, USA. Founded 2015. Acquired MetaMap (2024), AuthenticID (2025) and Identiq (completed 25 June 2026).	Yes (ID document photo matched to a real-time selfie; liveness detection)	ID document verification; real-time selfie matching; liveness and deepfake detection; age assurance products.	TikTok (TikTok LIVE)
k-ID — Kidentify Pte. Ltd.	Singapore. Series A led by Andreessen Horowitz and Lightspeed; other investors include Konvoy, TIRTA, Okta, Z Venture Capital and EDBI.	No (orchestrator only — biometric processing handled by sub-processors; k-ID receives the result, which for ID checks can include date of birth or an age range; platforms receive pass/fail)	Orchestration platform routing users to facial age estimation, ID upload, bank check, credit card check and country-specific routes (Social Security Number in the US, Singpass in Singapore, KISA-designated methods in South Korea, CPF in Brazil), or an on-device 'AgeKey' credential. Underlying checks run by sub-processors (Privately SA, Veratad, AP+/ConnectID, Stripe, Singpass, KG INICIS, Privy). In Singapore and South Korea the provider returns the user's date of birth to k-ID.	Snapchat; Twitch; Kick; AgeKey (OpenAge Initiative) announced for Meta platforms in 2026

Provider (legal entity)	HQ / background	Uses biometric data	Methods offered	Platforms using it
OpenAge Initiative — OpenAge Global Initiative Pte Ltd	Singapore. Launched by k-ID in November 2025 after it acquired Opale.io (France).	No (stores the result of a prior check as a passkey on the user's device)	Reusable 'AgeKey' credential recording age range, method used and date created; relying services check it against their own threshold without a fresh check.	Snapchat; Kick; announced for Meta platforms in 2026
Persona — Persona Identities, Inc.	San Francisco, USA. Founded 2018. Valued at ~US$2B in its April 2025 Series D.	Yes (collects facial images and biometric data server-side for facial age estimation and ID-to-selfie matching)	Seven methods, chosen by the platform buying the service: selfie age estimation (with checks for masks, faces shown on screens and deepfakes), email-based age inference, mobile network operator check, credit card check, government ID (optionally matched to a selfie), database check against credit bureaus and government records, and digital ID (Persona, communication to DDM, October 2026). May also collect device data, IP/geolocation and usage signals (incl. copy/paste detection), and cross-reference government/credit/utility/telco databases, depending on the method and the customer's configuration.	Reddit (age assurance); X (ID checks for X Premium)
Private Identity (PrivateID) — Private Identity LLC	Concord, North Carolina, USA. Founded 2017.	Yes (on-device; selfie and biometric data stay on device; encrypted age result transmitted to servers)	On-device facial age estimation using fully homomorphic encryption; only a scrambled/encrypted age result is sent to PrivateID's server, which returns pass/fail. Scrambled age data stored 90 days.	YouTube
Stripe — Stripe, Inc.	Dual HQ San Francisco / Dublin. Founded 2010.	Yes (Stripe Identity captures selfies, ID images and runs biometric face match)	Stripe Identity: ID document verification, biometric face match, liveness detection. Used primarily for payments-adjacent KYC rather than as a dedicated age-assurance product.	X (ID verification for X Premium); Snapchat and Kick (credit card checks, via k-ID)
Privately SA	Lausanne, Switzerland (EPFL Innovation Park); UK branch registered. Founded 2013/14.	Yes (on-device; facial images and biometric data processed locally and deleted after check)	On-device facial age estimation (AgeAI offline / FaceAssure online); only pass/fail returned. Also offers voice-based age estimation.	Snapchat, Twitch, Kick (all via k-ID)
Veratad Technologies, LLC	Teaneck, New Jersey, USA. Founded 2006.	Yes (collects selfies; biometrics / liveness + facial match in its identity flows)	ID document verification; knowledge-based authentication; data-based age verification against public/private records; reusable identity credentials (Veratad VX / IDMax).	Snapchat, Kick (both via k-ID)
Verifymy Limited	London, UK. Founded 2019.	No (for the email-based age checks (estimation/inference) product used by YouTube; biometric products exist but the email path does not process biometrics).	Email-based age checks (age inference, in ISO terms): analyses the digital footprint of the user's email (which sites/services it has registered with). For Google checks, shares the email address with unnamed 'trusted third parties, such as data aggregators' and deletes the data once the result is shared.	YouTube (direct). Listed by k-ID as an available method but not deployed for any platform in Australia.
Yoti — Yoti Ltd	London, UK. Incorporated 2014.	Yes (facial images processed for age estimation and ID-to-selfie	Facial age estimation (selfie + liveness); ID document verification; reusable Digital ID app (Yoti ID wallet); reusable age tokens / 'Yoti	Facebook; Instagram; Threads; TikTok

Provider (legal entity)	HQ / background	Uses biometric data	Methods offered	Platforms using it
		matching; underlying processing built on biometrics)	Keys'. Also email, credit card, mobile-provider and database age checks. The Yoti ID app was withdrawn from Spanish app stores on 10 September 2026 after the AEPD fine.	

Table B — Key age-check providers used under the Australian social media ban.

Appendix C — Lawful-basis mapping and where Article 8 of the GDPR applies

UK GDPR largely mirrors EU GDPR on lawful bases, but since 5 February 2026 the Data (Use and Access) Act 2025 has added a UK-only 'recognised legitimate interests' basis that needs no balancing test or to offer an opt out to the data subjects (the child), and a duty on services likely to be accessed by children to consider 'children's higher protection matters' (Data (Use and Access) Act 2025). This makes it far easier for companies to avoid prior data protection obligations to respect what users want. Article 8 GDPR is NOT a minimum age to process a child's data: it applies ONLY where two conditions are met: (a) the basis is consent and (b) the service is an information society service offered directly to a child. There is no GDPR age-related floor for contract- or legitimate-interest based processing; the '13' age gate on these platforms is a Terms and Conditions based floor, not GDPR data processing threshold.

Processing activity	Article	Lawful basis (across platforms)	Notes
Core service (account, feed, posting, messaging, friend/follow graph)	Art 6(1)(b)	Contract (all platforms)	Consent not used.
Customer support / queries, claims, disputes	Art 6(1)(b)/(f)	Contract / Legitimate interests	—
Safety, security, integrity; fraud, spam, abuse prevention; moderation	Art 6(1)(f)	Legitimate interests (all)	LI is problematic to rely on for children (Recital 38).
Legally-required moderation (illegal content, DSA/OSA duties)	Art 6(1)(c)	Legal obligation (all)	—
Service improvement, analytics, R&D, debugging	Art 6(1)(f)	Legitimate interests (all)	—
Contextual / non-personalised advertising	Art 6(1)(f)/(b)	Legitimate interests (all)	—

Processing activity	Article	Lawful basis (across platforms)	Notes
Personalised / behavioural advertising (on-platform)	Art 6(1)(a)/(f)	Consent (Meta, YouTube); Legitimate interests (X, Reddit, Twitch, Kick); TikTok and Snapchat: LI for 18+ only	Below the Art 8 age a parent must authorise. TikTok and Snapchat avoid this by not personalising ads for under-18s. In the EU, the UK, Norway and Switzerland, Snapchat shows teen accounts (13–17) only contextual ads, based on surrounding content, time of day, age, device language, city-level location and device information (Snap Inc., n.d.). In the EU, Article 28(2) of the DSA already bars online platforms from presenting profiling-based ads to users they know with reasonable certainty to be minors.
Off-platform / cross-site and cross-app data for ads	Art 6(1)(a)	Consent (all)	Parental authorisation required below the Art 8 age.
Non-essential cookies / tracking pixels / SDKs	PECR/ePrivacy + 6(1)(a)	Consent (all)	Applies where the consenting user is a child below the Art 8 age. In the UK, the Data (Use and Access) Act 2025 exempts some low-risk cookies, such as certain analytics, from the consent requirement.
Special-category data (including biometrics)	Art 9(2)(a)	Explicit consent (all)	Higher 'explicit' standard; parental authorisation also required below Art 8 age.
Optional features: precise geolocation, contacts upload, camera/microphone	Art 6(1)(a)	Consent (all)	Per-feature parental authorisation below Art 8 age.
Direct marketing (email / push)	Art 6(1)(a)/(f)	Consent / Legitimate interests	Consent-based marketing to a child below Art 8 age needs parental authorisation.
Legal compliance (law-enforcement, retention)	Art 6(1)(c)	Legal obligation (all)	Consent not used.
Protecting life / emergencies and self-harm	Art 6(1)(d)	Vital interests (all)	Consent not used.

Table C — Data-processing lawful basis by activity for platforms affected by the Australia under-16 ban, and the points at which Article 8 parental authorisation is engaged

Further related recent legislative changes of note: Section 72 of the Children's Wellbeing and Schools Act 2026 lets the Secretary of State change the UK digital age of consent (currently 13) to any age from 13 to 16 by regulations (Woodhouse, 2026). In the EU, the proposed KIDS Act would apply its account age limits in parallel with Article 8 (European Commission, 2026c, Art. 6(5)).

Appendix D — Age-based proposals across Europe (October 2026, ongoing changes)

Summary of age-check plans and legislation emerging across the 46 member states of the Council of Europe and selected observers. Where not otherwise specified, we have not been able to establish from publicly available sources at the time of writing that specific legislation is planned. Member states in the EU must comply with EU regulation and the DSA baseline. States with no bill introduced are omitted below for space. The EU KIDS Act proposal of September 2026 would, if adopted, set EU-wide minimum ages for social media accounts, which would affect many of the national approaches below (European Commission, 2026b).

Country	Approach	Threshold	Status (Oct 2026)	Mechanism	VPN ban
Albania	No specific legislation	—	No bill; aligning with EU as accession candidate; TikTok ban lifted February 2026	—	No
Armenia	No specific legislation	—	No bill introduced	—	No
Azerbaijan	No specific legislation	—	No bill introduced	—	No
Austria	Ban	<14	Draft law published September 2026: under-14s barred from platforms with personalised feeds, infinite scroll or autoplay; signed 11-state letter	Verification method undecided	No
Belgium	Considering	TBD	Superior Health Council report Dec 2025; draft proposal expected June 2026; Wallonia-Brussels Federation signed 11-state letter; Flanders introduced an under-13 restriction on harmful platforms (April 2026)	TBD	No
Bulgaria	No national bill	—	Will comply with EU regulation; Proposed: GERB amendments to Child Protection Act tabled 27 May 2026; <13 ban, 13–15 parental consent	EU DSA baseline only	No
Croatia	National bill rejected	—	Domestic bill was rejected in March 2026 (Možemo! *Prijedlog zakona o digitalnoj zaštiti djece*), parental consent <15. Signed 11-state letter.	EU DSA baseline; supports EU-wide rules	No
Cyprus	Ban proposed + piloting EU app	<15	Under-15 ban proposed April 2026, with verification through the Digital Citizen app; action plan approved by the Council of Ministers September 2026; signed 11-state letter; one of 7 frontrunner pilot states	EU age verification mini-wallet pilot	No
Czechia	Considering	—	No bill; PM Babiš and Deputy PM Havlíček back an under-15 ban; Jutland signatory.	—	No

Country	Approach	Threshold	Status (Oct 2026)	Mechanism	VPN ban
Denmark	Banned for under-13s (parental approval 13-15)	Under discussion (not yet agreed).	Political agreement Nov 2025; legislation drafting; signed 11-state letter; piloting EU app. Under discussion (not agreed); <13 strict ban + 13–15 consent; AltID launched 4 June 2026	National eID + EU app; DSA fines up to 6% turnover	No
Estonia	Opposes bans	n/a	Education Minister opposes; one of few EU states not to sign the digital age of majority declaration	Prefers digital literacy	No
Finland	No ban proposed (debating)	—	PM supports under-15 ban; health ministry working group examining the issue since February 2026; no bill yet	—	No
France	Ban	<15	Law adopted 21 Jul 2026; ban struck down by Conseil constitutionnel 14 Aug 2026 (Décision n° 2026-911 DC); revised text notified to the Commission 14 Sept 2026; piloting EU app	Refuse/delete under-15 accounts; phone ban in lycées	Discussed (not in bill)
Germany	Divided	<13 recommended by expert commission; tiered access 13–16 and 16–18	Expert commission (June 2026) recommended an under-13 ban with tiered access to 18; family minister supports; no bill yet; supports EU mandatory verification; federal/state complications	—	No
Greece	Ban	<15	TRIS notification 2026/0233/GR, 8 May 2026; tiered 15–18 teen accounts. target Jan 2027; signed 11-state letter; piloting EU app. Withdrew notification to Com on drafting legislation. PM announced on 8 April 2026 that platforms must re-verify all users' ages twice a year.	Children's digital wallet + EU app	No
Hungary	Restrict (draft)	<13 for high-risk platforms; restrictions 13–16	Draft amendment to the Consumer Protection Act submitted May 2026; Jutland signatory.	EU DSA baseline only	No
Iceland	Considering	<15	Ministry of Education and Children proposed an under-15 ban (July 2026); not yet debated in parliament; signed digital age of majority declaration	TBD	No
Ireland	Restrict (no age ban yet)	Online Safety Code in force since Oct 2024	Digital wallet using PPS number in development; signed 11-state letter; piloting EU app; Minister welcomed the EU KIDS Act; Taoiseach says Ireland will follow EU positions	Online Safety Code; considering Australia-style ban	No

Country	Approach	Threshold	Status (Oct 2026)	Mechanism	VPN ban
Italy	Restrict / debating	<15 (govt bill); 2018 law requires parental consent <14	Bill DDL 1136: <15 ban + 15–16 parental consent for under-15 restriction incl. child influencers; piloting EU app; signed 11-state letter. PM Meloni said in June 2026 that the government will not introduce its own age-restriction bill.	Existing parental-consent rule	No (adult sites banned from promoting VPN workarounds)
Lithuania	Restrict (parental consent) proposed	<16	MP's proposal (April 2026) to require parental consent under 16; parliament evaluating	—	No
Luxembourg	Considering	TBD	Parliament debated an age limit (January 2026); PM says national legislation will follow if EU-wide rules do not materialise by end of 2026	—	No
Malta	Restrict (announced)	<13	PM announced restrictions for under-13s modelled on Australia's law; may later extend to older ages	—	No
Montenegro	Ban proposed	<13; parental consent 13–16	Bill introduced by two MPs; also bans sexualised and deepfake content featuring children	—	No
Netherlands	Restrict → ban being prepared	<15	New coalition intends ban; supports EU mandatory verification	Privacy-friendly age verification — TBC	No
Norway	Ban	<16 (raised from <15)	Bill announced Apr 2026, expected in parliament late 2026	Platform-side verification (BankID); GDPR consent age raised to 15; fines up to NOK 20m	No
Poland	Ban	<15	Ruling party drafting (Feb 2026); tool launch end 2026; phone ban in primary schools	Platform verification; fines up to 6% revenue	No
Portugal	Restrict (parental consent) + ban under 13	<13 outright; 13–16 needs verified parental consent	Approved in Parliament February 2026; final vote pending	Digital Mobile Key; fines up to 2% global turnover	No
Romania	Restrict (parental consent)	<16 (digital adulthood)	Digital Adulthood Law approved by Senate; additional VLOP bill (fines up to 3% global turnover)	Platform verification + parental controls + monthly reports	No
Slovakia	Ban (government bill)	<16	Government approved under-16 ban August 2026, pending parliament; signed 11-state letter	EU DSA baseline + push for mandatory verification	No
Slovenia	Ban	<15	Bill drafted; experts consulted; signed 11-state letter	—	No
Spain	Ban	<16	Announced Feb 2026 (PM Sánchez); in parliamentary process; signed 11-state letter; piloting EU app. <16 ban; Cartera Digital Beta by AEPD	Mandatory verification; criminalisation of algorithm	No

Country	Approach	Threshold	Status (Oct 2026)	Mechanism	VPN ban
				manipulation; fines + 'Hate and Polarization Footprint'	
Sweden	Not proposed a ban (debating); pushing for EU-wide ban	<15 (recommended)	Government-appointed commission recommended barring under-15s (June 2026); signed digital age of majority declaration	—	No
Switzerland	Considering	TBD	Draft 'Communication Platforms Act' (transparency, not age bans); Federal Councillor open to access bans; Federal Commission for Children and Youth Affairs opposes a ban	KomPG draft; existing Film and Video Games Act 2022	No
Turkey	Ban	<15 (parental controls 15–17)	Passed by parliament 22 Apr 2026; effect 6 months after Gazette publication	Identity verification + explicit anti-VPN measures in same law	Yes — anti-circumvention in the same law
Ukraine	Considering (under-15 ban proposed)	<15	Draft bill in Verkhovna Rada; follows French model	TBD — platform-side	No
United Kingdom	Age assurance enforced + ministerial power to restrict for under-16s by regulation	Adult content: 18; possible <16 restriction by regulation	Online Safety Act enforcing since 25 Jul 2025; Children's Wellbeing and Schools Act 2026; consultation closed 26 May 2026; first regulations expected to be laid by end of 2026, in force from spring 2027.	Ofcom enforcement (£18m / 10% turnover); new SoS regulation power; smartphones-in-schools ban	Power exists — Lords' under-18 VPN amendment rejected, but Act gives ministers regulation-making power

Table D — Age-check approaches across Council of Europe member states and observers (selected). Sources include national legislation trackers, CEPA, Biometric Update and national press, May 2026, updated October 2026 from CEPA (2026) and Age Gate Tracker https://www.interface-eu.org/publications/agegate-tracker (Galissaire et al)

Appendix E — Age-based proposals across Rest of World (Oct. 2026, subject to ongoing change)

Selected non-EU jurisdictions where a statutory minimum age for social media access has been proposed, adopted, or is under discussion.

Country	Stage	Proposed minimum age
Australia	In force since 10 December 2025; maximum penalties doubled September 2026	16
Bahrain	Proposed; Shura Council reviewed draft amendment, government opposes a blanket ban	15
Brazil	In force since March 2026 (Digital Statute of the Child and Adolescent); sets no social media age limit	Under-16 accounts must be linked to the account of a parent or guardian
Canada	Safe Social Media Act introduced June 2026; platforms meeting government safety criteria exempt	16
Costa Rica	Proposed	14 (with parental consent) / 18
Ecuador	Proposed (bill introduced February 2026)	15
Egypt	Under discussion	Not indicated
Fiji	Under discussion	16
Gabon	Adopted; ordinance published 8 April 2026; also requires identity verification for all users	16
India	National tiered restrictions for under-18s under debate; several states discussing or proposing restrictions	13–16 (states); tiered to 18 (national)
Indonesia	In force since 28 March 2026; phased account deactivation ongoing	16
Japan	No ban; draft guidelines on age verification and platform risk assessments (June 2026)	Not indicated
Kazakhstan	Proposed amendment; messaging apps exempt	16
Malaysia	In force since 1 June 2026 (Online Safety Act 2025)	16
Mongolia	Draft law introduced July 2026; feature restrictions for ages 13–16	13
New Zealand	Government Online Safety Bill introduced 24 August 2026	16

Country	Stage	Proposed minimum age
Nigeria	Under discussion	Not indicated
Norway	Proposed; bill due in parliament by end of 2026	16
Pakistan	Proposed 2025; withdrawn, no new bill	Not indicated
Philippines	Senate Bill 2424 (under-16 ban; third-party age verification for 16–17s) passed second reading, Sept/Oct 2026; House committee substitute bill (under-13 ban, restrictions to 18) approved Sept 2026	16 (Senate) / 13 (House)
Republic of Korea	At least seven proposals under review (August 2026)	14
Rwanda	Government drafting legislation (announced April 2026)	16
Singapore	Legislation planned for 2027 (time limits, feature restrictions); platforms with weak safeguards may face a minimum age of 18	Not indicated (18 if safeguards not met)
South Africa	Under discussion; no formal proposal	Not indicated
Switzerland	Under discussion; Federal Commission for Children and Youth Affairs opposes a ban	Not indicated
Tajikistan	Proposed by an MP (late 2025)	14
Trinidad and Tobago	Under discussion	12
Türkiye	Adopted April 2026	15
United Arab Emirates	Adopted June 2026; one year to comply; restrictions for ages 15–16	15
US	No federal law; House passed KIDS Act June 2026; state laws vary (California, September 2026) and several are blocked in court	Varies by state; 13 in some federal proposals
Vietnam	Draft decree July 2026 bars under-16s from posting, commenting or reacting; accounts registered to a parent	16 (posting)
Zimbabwe	Under discussion	18

Source: UNICEF March 2026 (pp. 9–12), with reference to the Tech Policy Press Global Social Media Age Restriction Tracker; updated from CEPA, Mapping the Spread of Child Safety Rules (25 September 2026).

Appendix F — VPN bans and platform-blocking timelines

Year	Country	Action	Outcome
2009 onward	China	Only government-licensed VPNs allowed; Great Firewall detects and blocks others	Unlicensed providers blocked; account termination
2010	Oman	Long-standing restriction; only licensed VPNs permitted	Fines for unauthorised use
2011 onward	Syria	Repeated shutdowns; censorship targets VPN protocols	Among the world's most censored states
Since 2013	Iran	Only government-approved VPNs permitted	Fines or up to ~1 year imprisonment for unapproved use; in March 2026 police warned that connecting to the internet by VPN would be treated as a crime
Ongoing	North Korea	Total ban; citizens have only the state intranet (Kwangmyong)	VPN use heavily punishable
Ongoing	Turkmenistan	Effective total ban; single state-owned ISP	Enforcement reaches individual users
2015	Belarus	Outlawed "anonymisers" including VPNs and Tor	Banned without government approval
2019–2026	Russia	Bans non-compliant providers; ~100 VPN apps ordered removed (2024); fines for advertising/using VPNs to reach "extremist" content; more than 400 VPN services blocked by January 2026, and VPN traffic over mobile internet heavily restricted	Targeted fines; using a VPN not itself an offence per officials; mobile internet limited to whitelisted services during routine shutdowns
2022	Venezuela	State ISP blocked Tor; VPN sites blocked	Used to limit access to news and bypass sanctions
2025	Myanmar	Cybersecurity restrictions; unauthorised VPN install criminalised	Up to 6 months imprisonment

VPN restrictions by country. The underlying list is drawn from a VPN-vendor summary (nordvpn.com/blog/are-vpns-legal) and subject to change.

Year	Country	Action	Outcome
2009 onward	China	Twitter/Facebook blocked (2009); YouTube, X, Facebook, Instagram banned via the Great Firewall	Long-term; users moved to WeChat, Weibo
2009; 2022 onward	Iran	Facebook blocked (2009); Instagram, X, YouTube, Telegram restricted (Jul 2022) during protests	Most Western platforms remain blocked
2017–2020	Turkey	Wikipedia blocked on "national security" grounds	~3 years; Constitutional Court ruled it violated free expression
2020; 2025	India	Banned TikTok and dozens of Chinese apps (2020); 7-day regional all-social-media ban (2025)	TikTok ban permanent; regional blocks temporary

Year	Country	Action	Outcome
2021 onward	Uganda	Cut internet and banned Facebook around the 2021 election; near-total blackout at 2026 election	Facebook blocked
2021	Nigeria	X (Twitter) banned June 2021 over a deleted presidential post	Lifted January 2022 (~7 months)
2022 onward	Russia	Facebook, Instagram and X restricted amid Ukraine-war information controls; WhatsApp blocked and Telegram restricted from February 2026	Ongoing; Meta designated “extremist”
Aug–Oct 2024	Brazil	Supreme Court ordered ISPs to block X; Starlink assets frozen; user fines for VPN bypass	X complied; block lifted 8 Oct 2024; fined $5.2m
Mar 2025	Albania	Nationwide TikTok ban for one year after a fatal stabbing linked to platform content	Lifted February 2026
Mar 2025	Turkey	42-hour social media shutdown during protests over the arrest of Istanbul mayor Ekrem İmamoğlu	Draft law would let the telecoms regulator block platforms on national security grounds without a court order
Sept 2025	Afghanistan	Instagram, Facebook, Snapchat banned with regional shutdowns; TikTok restricted since 2021–22	—
Sept 2025	Nepal	26 platforms, including Facebook, X and YouTube, blocked for failing to register with the government	Lifted 9 September after protests in which 19 people were killed
Jan 2026	Uganda	Near-total internet blackout ordered ahead of the general election	Coincided with vote and exam period
Jan–May 2026	Iran	Nationwide internet blackout from 8 January during protests, continuing through the 2026 war	Longest nationwide shutdown on record
Feb–Apr 2026	Gabon	Media regulator suspended social media ‘until further notice’ amid strikes and protests	Restored April 2026

Selected platform blocking, throttling and shutdown incidents. The pattern is overwhelmingly associated with elections, protests and “national security” episodes in authoritarian or crisis contexts.

Appendix G – Intentions, predictions and outcomes of the SMMA

Dimension	SMMA Act 2024 + Explanatory Memorandum (intended)	Australian Child Rights Taskforce, Oct 2024 (warned)	OAIC privacy guidance, Mar 2026 (how to comply)	Molly Rose Foundation, Apr 2026 (outcome — polling)	eSafety compliance update, Mar 2026 (outcome — regulator)
Age-assurance feasibility	Mandates assurance without prescribing a method; bans requiring govt ID / Digital ID as the only option (s 63DB).	"Not yet effective techniques for age assurance nor to verify parental consent."	Treats assurance as workable if layered (waterfall); facial age estimation noted as the most popular method offered.	70% of children still on restricted sites found it "easy" to circumvent.	Confirms the warning: facial age estimation errs near the 16 threshold; some platforms invited self-declared 14–15s to "have a go" and many got false 16+ results.
Children's data protection	s 63F limits use of age-check data to age-checking and requires destruction after use; consent must be voluntary, informed, current, specific and unambiguous.	Warns that age-gating cannot be delivered without processing sensitive data across the population, including children; "privacy concerns remain."	Full APP framework applies: minimise collection (APP 3), accuracy (APP 10), security, ring-fencing, destruction and vendor due diligence (APP 11); biometric templates are sensitive data needing extra safeguards.	–	A new 25 Mar 2026 rule ties the definition to high-risk features (recommender / logged-in feeds); five platform investigations are due to conclude by mid-2026 (no published outcome found by September 2026). The trial found some providers over-retaining children's biometric / document data.
Children's right to privacy	Framed as protective; Age Discrimination Act amended to shield the scheme. Does not engage the child's privacy right directly beyond the s 63F use/destruction limits.	Central concern: bans and the checking they require infringe children's privacy and their rights to access / participation (cites UN CRC).	Uphold anonymity / pseudonymity where practicable (APP 2); warns against "unjust digital exclusion"; privacy intrusion should be proportionate to a legitimate aim.	–	–
Core position	Give under-16s "more time to develop digital literacy and resilience"; platforms must take "reasonable steps" to prevent accounts (ss 63B, 63D).	A ban is "too blunt an instrument"; systemic product-safety regulation is the better route.	Neutral how-to: if age assurance is used, make it necessary and proportionate, privacy-protective, and redesign services safer by default.	The ban is in force but porous; "significant questions about effectiveness."	Platforms have taken "some steps"; world-leading but "gaps remain." Whether steps are "reasonable" is contextual, not a one-size test.

Dimension	SMMA Act 2024 + Explanatory Memorandum (intended)	Australian Child Rights Taskforce, Oct 2024 (warned)	OAIC privacy guidance, Mar 2026 (how to comply)	Molly Rose Foundation, Apr 2026 (outcome — polling)	eSafety compliance update, Mar 2026 (outcome — regulator)
Effectiveness / circumvention	Civil penalties (up to AU$49.5m for bodies corporate) assumed to drive compliance. Maximum penalty doubled by the Strengthening Enforcement Act 2026, passed 10 Sept 2026.	Children would "still use platforms," creating more risk for those slipping through.	–	61% of prior account-holders retain ≥1 active account; >half still on TikTok, YouTube, Instagram.	~7 in 10 parents say their child kept Facebook / Instagram / Snapchat / TikTok accounts; some platforms allowed double-digit re-tries (vs recommended ≤5).
Privacy (general)	Privacy Division (s 63F): information collected for age-checking may be used only for that purpose, on a strict consent standard (voluntary, informed, current, specific, unambiguous), with OAIC oversight added by amendment. The age-gating duty itself, however, requires checking every user, not only children.	"Privacy concerns remain" with assurance — the warning is that age-gating cannot be done without building intrusive checking infrastructure for the whole population.	Establish whether age assurance is needed at all (APP 1); uphold anonymity / pseudonymity where practicable (APP 2); collect the minimum data by the least-intrusive means and prefer a binary "16+ yes/no" over identity (APP 3); just-in-time notices on what is collected, why, by whom and for how long (APP 5); conduct a Privacy Impact Assessment per method.	–	Did not assess privacy of the methods directly — explicitly referred privacy to the OAIC — while its "reasonable steps" / waterfall logic pushes platforms toward biometrics, ID/bank checks and behavioural inference for the whole user base.
Privacy: biometric data	– (the Act does not single out biometric data; it bans requiring government ID / Digital ID as the only method (s 63DB) but says nothing specific about facial-scan templates).	– (the 2024 letter predates the deployed methods; raises privacy in general, not biometrics specifically).	Biometric information used for verification/identification, and biometric templates created by facial age estimation, are sensitive information attracting higher protection (APP 11); entities should apply additional safeguards and extra technical/organisational measures to ensure destruction or de-identification of templates.	–	Notes facial age estimation is biometric and error-prone near 16, but defers the privacy assessment of that processing to the OAIC rather than evaluating it.
Privacy: data retention	Mandatory destruction: s 63F(3) requires the entity to destroy the information after using it	–	Destroy inputs (selfies, ID images, liveness video, biometric templates)	–	The Government's own trial found a minority of providers over-retaining

Dimension	SMMA Act 2024 + Explanatory Memorandum (intended)	Australian Child Rights Taskforce, Oct 2024 (warned)	OAIC privacy guidance, Mar 2026 (how to comply)	Molly Rose Foundation, Apr 2026 (outcome — polling)	eSafety compliance update, Mar 2026 (outcome — regulator)
	for the purpose for which it was collected; failure to destroy is itself an interference with privacy.		immediately once age is established; retain only a minimal, ring-fenced artefact (e.g. a 16+ flag, method, timestamp, non-linkable token) with a short time-to-live; wipe caches and transient storage; avoid purpose-"padding"; no reuse for analytics, advertising or model training without consent (APP 11).		full biometric / document data "just in case" — the opposite of the OAIC instruction.
Reporting / parental pathways	No parental-consent exemption; the ban applies regardless of parental approval.	Parents aren't "the experts" and shouldn't carry the burden unaided.	Complaint mechanisms must be easy to navigate (APP 1).	–	Reporting pathways "generally not accessible or effective" — some platforms demanded a lawyer's letter to prove parentage; reports often left accounts active.
Safety outcomes	Object is to "reduce risk of harm" (s 63B); 2-yr independent review built in (s 239B).	Predicted no improvement to product safety; platforms disincentivised from protecting under-16s who remain.	Focus is privacy harm from the checks, not safety outcomes.	51% say no difference to safety; 14% feel less safe; mixed wellbeing.	No discernible drop in cyberbullying / image-based-abuse complaints from under-16s vs the same period in 2025; educators report mixed effects.
Scale of action taken	No targets specified; outcome-based "reasonable steps."	–	–	Measures retained access, not removals.	~4.7m accounts removed / restricted by mid-Dec 2025; +300,000 since — but accounts, not users, and not determinative of compliance.

Table G — Stated legislative intent, the 2024 expert warning, the OAIC compliance guidance, and the early outcome evidence from both an independent poll and the regulator (Online Safety Amendment (SMMA) Act 2024 and Supplementary Explanatory Memorandum; Australian Child Rights Taskforce, 2024; OAIC, 2026; Molly Rose Foundation, 2026; eSafety, 2026a).

Appendix H — Age-assurance consent architecture by provider

This appendix sets out, for each in-house and third-party age-assurance provider identified across the platforms covered by the Australian Online Safety Amendment (Social Media Minimum Age) Act 2024, whether publicly available policy documentation evidences a discrete opt-in, an opt-out default, or an undocumented consent step, at three separate points in the process: (i) undergoing the check itself, (ii) the retention period applied to data collected, and (iii) reuse of that data for any other purposes. *Where ('not documented'), this reflects the absence of any identifiable discrete consent in the provider's own publicly available policy at the time of writing.*

Provider (type)	Platform(s)	Privacy policy	Consent to undergo the check	Consent to retention period	Consent to reuse for product dev / other purposes
Meta in-house	Facebook, Instagram, Threads	about.fb.com/news/2026/05/ai-age-assurance-teens/	Mandatory step in account creation — not a discrete opt-in, required to proceed sign-up	Opt-out / none — 30-day retention stated as default, no user choice mechanism	Opt-out / none — continuous scanning of photos/posts/captions/bios runs with no disclosed opt-out; feeds Meta's generative AI training
Yoti	Facebook, Instagram, Threads, TikTok	yoti.com/privacy/age-verification/	Mandatory step to proceed with account creation	Opt-out — AEPD found biometric templates retained "for the life of the account plus three years," fined €250,000 for excessive retention. Yoti said it would appeal; it withdrew the Yoti ID app from Spain on 10 September 2026.	Opt-out — AEPD Fundamento de Derecho VI: R&D consent obtained via a box "activadas por defecto" (pre-ticked); Antecedentes records a minor "solo hay que marcar una declaración de ser mayor de edad"
k-ID	Snapchat, Twitch, Kick; AgeKey	k-id.com/privacy-policy	Not documented	Retention periods documented for Trusted Adult data only (privacy policy updated 10 July 2026); none specific to age-check results	Not documented
Veratad	Snapchat, Kick (via k-ID)	veratad.com/privacy-policy	Not documented	Not documented	Ambiguous / effective opt-out — states data not used for "unrelated secondary commercial

Provider (type)	Platform(s)	Privacy policy	Consent to undergo the check	Consent to retention period	Consent to reuse for product dev / other purposes
					purposes," implying related use is permitted with no further consent gate
Privately SA	Snapchat, Twitch, Kick (via k-ID)	privately.eu/privacy-policy-en	Not separately documented — capture is part of platform sign-up	N/A — on-device, deleted after check	Opt-out / reserved — policy states future AI-training use "might change," no consent gate described for if/when it activates
Verifymy	YouTube (direct)	verifymy.io/age-verification-and-estimation/google-age-assurance-privacy-policy/	Consent requested by Google before the check; documented in Verifymy's Google-specific policy	Data deleted once the result is shared with Google	Email shared with unnamed "trusted third parties, such as data aggregators" to run the check; no discrete opt-in identified
Persona	Reddit; X (Premium ID checks)	withpersona.com/legal/privacy-policy/	Opt-in — policy states biometric data collected "only with your express consent"	Bundled, not separately elected — default is automatic deletion once an outcome is determined, but the customer (Reddit/X) "may direct us to retain" data longer, "as disclosed to you at the time you provide consent" — the extension rides on the same single consent action and is controlled by the platform, not chosen separately by the user	Opt-in / stated absence — Persona states it does not use personal or biometric data for AI/model training and does not sell or share it
PrivateID	YouTube	g.privateid.com/privacy	Opt-in — Google's support page documents a genuine per-attempt screen: "Review the information from Private ID. Decide if you want to continue… tap No thanks"	Not separately elected — the same screen only asks whether to proceed with the check; the 90-day retention of the scrambled result is not itself presented as a choice, it follows automatically once the user proceeds	Not documented as a reuse clause

Provider (type)	Platform(s)	Privacy policy	Consent to undergo the check	Consent to retention period	Consent to reuse for product dev / other purposes
AU10TIX	Sub-processor for various clients; not confirmed as used in Australian SMMA specifically	au10tix.com/…/Au10tix-Biometric-Data-Policy-November-2025.pdf	No consumer-facing consent step documented — the Biometric Data Policy addresses the "Referring Business," not the end user. The only consent language found anywhere is bundled into an unrelated third-party account ToS (Interactive Brokers): "By agreeing to the terms, you consent to [name] and AU10TIX collecting and storing your personal and biometric data" — one click covering an entire account agreement	Not documented for end users	Not documented for end users
Stripe Identity	X (Premium)	stripe.com/privacy	Opt-in — "Stripe will ask for your consent before collecting and using your biometric information," stated as occurring "before starting the verification process"	Not documented as a separate step	Not documented as a separate step — general privacy policy covers fraud-prevention/AI training broadly, not disclosed as a discrete consent for age-check use specifically
TikTok in-house	TikTok	tiktok.com/support/faq_detail?id=7543604782246959621	Mandatory step to proceed with account creation	Opt-out / none — runs automatically over "profile, device information, videos… interactions… other on-platform behavior"; retention stated only as "no longer than necessary," undefined	Opt-out / none — same undefined-purpose language, no separate reuse clause
Google in-house	YouTube	support.google.com/accounts/answer/10071085	Mandatory step to proceed with account creation	Opt-out / none — Google "makes no commitment to deleting" ID data	Opt-out / none — reserves right to use Chrome history, search terms, location, third-party

Provider (type)	Platform(s)	Privacy policy	Consent to undergo the check	Consent to retention period	Consent to reuse for product dev / other purposes
					activity to infer age, folded into "improve our verification services… and protect against fraud and abuse"
ConnectID / member banks	Snapchat, Kick (via k-ID)	anz.com.au/plus/terms-conditions/connectid-terms-conditions/ · nab.com.au/help-support/connectid-your-digital-identity/terms-conditions · auspayplus.com.au/connectid-privacy-policy	Discrete consent per transaction: the user chooses ConnectID, picks a bank, reviews what will be shared and consents or cancels in the bank's app	Each bank keeps its own "Transaction Record" of every sharing event, which the customer can view in their banking app; no separate consent to this record. ConnectID says it sees only transaction metadata (bank, platform, attribute type), with no personal data or age outcome	Westpac/NAB state Transaction Record details are shared with ConnectID "for operation, governance, fraud detection and other purposes"; NAB says the other purposes are set out in the ConnectID Privacy Policy. ConnectID says these details are metadata only. No discrete consent point
Incode	TikTok (TikTok LIVE)	incode.com/privacy-policy Biometric Data Policy and Notice	Set by the platform: where Incode acts as processor, "the Customer's privacy notice governs" (here TikTok's). Incode's US biometric policy requires written consent before collection	Set by the platform: Incode refers users to the customer's privacy notice for retention periods, and leaves retention instructions for minors' data to the customer	Opt-in or "other lawful legal basis" — Incode may use personal data "with your consent or other lawful legal basis, to improve our products or services"; states it does not "sell or share" minors' personal data

Table H1 — Consent architecture by age-assurance provider: undergoing the check, retention, and reuse for product development or other purposes.

Appendix I — Liveness checks by age-assurance provider: type, action and frequency

This appendix isolates the liveness component of the age checks mapped in Appendices A and B: whether a provider runs a liveness check, whether it is passive (a single still selfie, no action required) or active (the user must perform a named gesture), the exact action demanded of the user, at what point in the process it occurs, and how often it is repeated. The final row extends the same analysis to the EU age-verification "mini-wallet", whose token-based presentation hides a liveness step at the issuer stage. Sourced from each provider's own documentation or stated where none was found available.

Provider (platforms served under SMMA)	Liveness type	Exact user action (gesture)	Where in the process	Frequency
Yoti (dual processes) *Facebook, Instagram, Threads, TikTok*	Passive (MyFace). Yoti argues against active liveness.	Keep still — single selfie, head-on to camera. No nod, turn, blink or speech. Auto-capture or one button press. 3 attempts allowed.	In-browser one-off check (see Flow B, used by Meta): fresh selfie at each account creation. For their Digital ID wallet it is Flow A: once at credential set-up, then reused.	Per access / sign-up (one-off flow); one-off at issuance (wallet flow).
Privately SA *Snapchat, Twitch, Kick (via k-ID)*	Both active + passive (confirmed by CEO, Dec 2025).	Passive path: keep still. Active path: nodding head / saying certain words aloud.	At the facial-check step each time k-ID routes the user to facial age estimation; processed on-device.	Per check; active triggered only when passive evidence is insufficient (i.e. near the threshold).
Persona *Reddit; X (Premium ID)*	Both — active challenge + passive analysis.	Default gesture set: look centre, then left, then right (head-turn). Configurable to add a smile, blink one eye, turn head, hold up a hand or a number of fingers. Multilingual prompts.	ID capture → then selfie + liveness. Reverification on password recovery, MFA reset and before high-risk actions is offered marketed as a customer-configured account protection feature.	Per check; repeated at reverification events (recovery / step-up) where the platform switches this on. Configurable retry limit.
Private Identity / PrivateID *YouTube*	Passive by default; active as optional step-up for high-risk.	Passive: keep still, single selfie, on-device (~20ms). Active step-up gestures not individually named in provider docs.	Runs at both enrolment and each authentication ("during enrolment and authentication").	Repeated at each authentication — not only at issuance; passive unless stepped up.

Provider (platforms served under SMMA)	Liveness type	Exact user action (gesture)	Where in the process	Frequency
Stripe Identity *X (Premium ID)*	Passive — "a live selfie"; no gesture challenge documented in Stripe's own flow.	Single selfie matched to ID photo; detects digital layering. No named head-turn / blink / speech action.	At the selfie check step, paired with ID-document capture (*require_matching_selfie*). Age-16+ gate on the selfie check itself.	Per verification session; also on liveness re-checks (onboarding / some jurisdictions).
Veratad *Snapchat, Kick (via k-ID, ID path)*	Both (passive / active) capable; certified liveness on face-based check.	Specific gestures not named in Veratad's own materials. Confirm by direct enquiry.	Primarily ID-document verification via k-ID; liveness on the face-based path.	Not specified in source — flag for provider enquiry.
Verifymy *YouTube (email)*	Not applicable: email-based check with no face. Ownership of the email address is confirmed with a time-limited one-time passcode instead.	Not applicable (no camera, no selfie).	Email-address analysis step.	Not applicable.
ConnectID / Australian Payments Plus *Snapchat, Kick (bank path)*	No camera-based liveness — bank sign-in and consent process required instead.	Not applicable (no camera).	Bank sign-in and consent step.	Not applicable.
Incode **TikTok (TikTok LIVE)**	Liveness detection on a real-time selfie; passive or active not documented for TikTok.	Capture a real-time selfie alongside photos of a government-issued ID.	At the TikTok LIVE age check, paired with ID-document capture.	Not documented.
OpenAge Initiative (AgeKey) **Snapchat, Kick; announced for Meta**	No liveness at presentation; inherits whatever liveness the issuing check used.	Unlock the passkey with Face ID, fingerprint or device passcode.	At each presentation of the AgeKey to a participating service.	Liveness, if any, once at the original issuing check; the credential is reused without a fresh scan.
EU age-verification "mini-wallet"	No liveness at proof-of-age presentation (token only). But liveness is inherited	At presentation: none — a zero-knowledge "over-16/18" token. At issuance, if the Proof-of-Age issuer	Two-step: an *issuer* (national eID, ID card, passport, face-to-face counter, or accredited third-party app) establishes age	Liveness once at issuance (if a face-based issuer is used); token reused

Provider (platforms served under SMMA)	Liveness type	Exact user action (gesture)	Where in the process	Frequency
European proposal (pilot); not an SMMA provider	from whichever *issuer* verifies age at onboarding.	is a face-based provider (not a national eID), the user performs that provider's liveness — e.g. passive selfie, or an active gesture on step-up.	once; the *wallet app* then presents a token to relying parties. Liveness, if any, sits at the issuer step.	thereafter without re-scan — unless re-issuance is required.

Appendix J — Technical standards governing Age assurance, AV and identity

(16 July 2026). Status reflects the published stage codes of each responsible committee at the date of compilation. Standards under development are subject to change; committee draft scopes are cited from the ISO project record rather than from the draft text, which is not publicly available. Reference numbers prefixed BS denote the UK national adoption of the international standard.

Reference	Title	Body / committee	Status (July 2026)	Scope and relevance
1. Age assurance systems — the ISO/IEC 27566 series				
ISO/IEC 27566-1:2025 (BS ISO/IEC 27566-1:2025)	Information security, cybersecurity and privacy protection — Age assurance systems — Part 1: Framework	ISO/IEC JTC 1/SC 27; UK mirror committee BSI IST/33/5	Published December 2025. Edition 1. 29 pp. ICS 35.030. ISO project 88143.	Establishes a framework for age assurance systems and describes their core characteristics, including privacy and security, for enabling age-related eligibility decisions. Covers verification, estimation and inference as three distinct methods, and a layered "step-up" approach. Explicitly framed to support age-related eligibility decisions without mandating full identity verification in all cases. Made available free of charge (CHF 0) — unusual for ISO and indicative of intended regulatory uptake. Silent on age thresholds and on behavioural harms arising from the act of age checking.
ISO/IEC 27566-2	Age assurance systems — Part 2: Technical approaches and guidance for implementation	ISO/IEC JTC 1/SC 27	Under development.	Describes different technical approaches suitable in different ecosystems, and guidance for their implementation. Bridges the foundational concepts of Part 1 with the analytical approaches of Part 3; addresses how to ascertain objectives, parameters and the required level of assurance.

Reference	Title	Body / committee	Status (July 2026)	Scope and relevance
ISO/IEC CD 27566-3.2	Age assurance systems — Part 3: Approaches to analysis or comparison	ISO/IEC JTC 1/SC 27	Committee Draft, stage 30.60. ISO project 88147. Close of voting / comment period 30 May 2026.	Establishes considerations for analysing, comparing or differentiating the characteristics of age assurance systems or components. Includes metrics, elements and indicators of effectiveness — i.e. the benchmarking layer on which regulatory claims of "effective" age assurance would rest.
2. Age-appropriate design and online age verification — the IEEE 2089 family				
IEEE 2089-2021	Standard for an Age Appropriate Digital Services Framework Based on the 5Rights Principles for Children	IEEE Standards Association	Active.	Foundational standard of the family. Establishes the overarching age-appropriate digital services framework grounded in the 5Rights principles; the ethical blueprint from which 2089.1 derives.
IEEE 2089.1-2024	Standard for Online Age Verification	IEEE SA; Online Age Verification Working Group	Active. Published May 2024. 55 pp. Corrigendum 1 in progress.	Framework for the design, specification, evaluation and deployment of online age verification systems. Sets out key terms and the roles and responsibilities of actors; requirements for four levels of confidence (asserted, standard, enhanced, strict); and requirements for privacy protection, data security and information systems management specific to the age assurance process. More prescriptive than 27566-1 on what each actor should do, with an explicit requirement to consider the unique contexts of children. Corrigendum 1 removes incorrect annex references and revises figures in Annex 1.
IEEE P2089.2	Standard for Formatting Terms and Conditions for Children's Online Engagement	IEEE SA	Active project.	Relevant to consent-flow analysis and to comprehension of terms presented at the point of an age check.
IEEE P2089.3	Standard for Managing and Verifying Online Parental Consent	IEEE SA	Active project.	Directly engaged by parental-consent architectures in the Australian SMMA model and in EU proposals.
3. UK national standards and certification schemes				

Reference	Title	Body / committee	Status (July 2026)	Scope and relevance
PAS 1296:2018 (BS PAS 1296)	Online age checking — Provision and use of online age check services — Code of practice	BSI, with the Digital Policy Alliance	WITHDRAWN. Superseded in guidance by BS ISO/IEC 27566-1:2025.	Published 31 March 2018. Gave recommendations for a framework for the provision and use of online age check services: age-restricted merchandise, online content, online services, and age-gated material. Normatively referenced BS ISO/IEC 29115, ISO/IEC 19790 and BS ISO/IEC 9594-8. Business Companion guidance has removed reference to it and substituted BS ISO/IEC 27566-1:2025. Note for citation: still referenced by some certification marketing.
BSI IST/33/5	National committee on age assurance	BSI	Active.	Coordinated UK input to ISO/IEC 27566-1; the UK provided expert leadership in the editorial process. 5Rights Foundation contributed to the standard's development through this committee. Relevant to provenance analysis of civil society footprint in the standard.
ACCS-4-2020	Technical Requirements for Age Check Systems	Age Check Certification Scheme (ACCS)	Active — private certification scheme, not a formal standard.	ACCS audits and certifies against PAS 1296:2018 and now offers ISO 27566-1 certification, marketed as the "international gold standard". Conformity-assessment layer rather than normative standard; treat accordingly in any evidence base.
4. Associated identity, credential and biometric standards				
ISO/IEC 29115	Information technology — Security techniques — Entity authentication assurance framework	ISO/IEC JTC 1/SC 27	Published (BS version 2013).	Levels of assurance framework underpinning PAS 1296. Historic anchor for "assurance level" language now recast in 27566-1 and 2089.1.
ISO/IEC 29003	Information technology — Security techniques — Identity proofing	ISO/IEC JTC 1/SC 27	Published.	Identity proofing process — the step age verification (as opposed to estimation or inference) necessarily invokes.
ISO/IEC 18013-5	Personal identification — ISO-compliant driving licence — Part 5: Mobile driving licence (mDL) application	ISO/IEC JTC 1/SC 17	Published.	Data structures and protocols for presenting a licence held on a mobile device to a verifier. The technical basis for privacy-preserving proof-of-age attestations. The EU Digital Identity Wallet age verification specification requires mDoc support and follows the 18013-5 recommendation of setting hh, mm and ss to identical values on each Proof of Age Attestation to limit linkability.

Reference	Title	Body / committee	Status (July 2026)	Scope and relevance
ISO/IEC 18013-7	Mobile driving licence — Add-on functions	ISO/IEC JTC 1/SC 17	Published / under revision.	Extends mDL to unattended and online presentation — the pathway by which mDL becomes an online age gate.
ISO/IEC 23220 series	Cards and security devices for personal identification — Building blocks for identity management via mobile devices	ISO/IEC JTC 1/SC 17	Multi-part; parts published and in development.	Underlying architecture for wallet-based attribute attestation, including age attributes.
ISO/IEC 19795-1:2021	Information technology — Biometric performance testing and reporting — Part 1: Principles and framework	ISO/IEC JTC 1/SC 37	Published 2021.	The performance-testing baseline against which any facial age estimation accuracy claim must be read.
ISO/IEC 19795-10:2024	Biometric performance testing and reporting — Part 10: Quantifying biometric system performance variation across demographic groups	ISO/IEC JTC 1/SC 37	Published 2024.	Demographic differentials. Central to any equalities or discrimination argument on facial age estimation. Required by NIST SP 800-63A for biometric performance testing, including demographics testing.
ISO/IEC 30107-1:2023	Information technology — Biometric presentation attack detection — Part 1: Framework	ISO/IEC JTC 1/SC 37	Published, Edition 2, 2023. Publicly available.	Framework for automated detection of presentation attacks (PAD). The standards layer beneath "liveness check".
ISO/IEC 30107-3:2023	Biometric presentation attack detection — Part 3: Testing and reporting	ISO/IEC JTC 1/SC 37	Published, Edition 2, 2023 (cancels and replaces 30107-3:2017).	Principles and methods for performance assessment of PAD mechanisms and reporting of results. The operative citation for active vs passive liveness claims by providers. Directly relevant to DDM's argument that liveness checks constitute a behavioural harm — teaching children to perform on camera on demand — since the standard governs only spoof-resistance, not the effect on the subject.
ISO/IEC 30107-4	Biometric presentation attack detection — Part 4: Profile for testing of mobile devices	ISO/IEC JTC 1/SC 37	Published.	A 30107-3 profile specific to mobile devices, reflecting accelerated product development timelines and a single biometric subsystem integrated across many device types.

Reference	Title	Body / committee	Status (July 2026)	Scope and relevance
ISO/IEC 2382-37	Information technology — Vocabulary — Part 37: Biometrics	ISO/IEC JTC 1/SC 37	Published.	Normative vocabulary referenced by the 19795 and 30107 series.
ISO/IEC 29100 / 29134 / 27701 / 27018	Privacy framework; Guidelines for privacy impact assessment; PIMS extension to 27001 and 27002; Protection of PII in public clouds	ISO/IEC JTC 1/SC 27	Published.	The privacy management layer against which age assurance DPIAs and vendor claims are assessed.
5. Adjacent non-ISO / non-IEEE instruments				
NIST SP 800-63-4 (incl. 800-63A)	Digital Identity Guidelines; Identity proofing requirements	NIST, US Department of Commerce	Published.	Requires PAD conformant to ISO/IEC 30107-3:2023; an impostor attack presentation accept rate (IAPAR) below 0.07 for remote biometric collection; and biometric performance testing conformant to ISO/IEC 19795-1:2021 and 19795-10:2024, including demographics testing, with results made publicly available. The most concrete numeric threshold available in this space.
ETSI TS 119 461	Policy and security requirements for trust service components providing identity proofing	ETSI	Published.	European identity proofing requirements; the eIDAS-facing counterpart to ISO/IEC 29003.
EUDI Wallet AV specification	Age Verification — Architecture and technical specifications (white-label solution)	European Commission / EU Digital Identity Wallet consortium	In development; published on GitHub.	Requires mDoc support; batch issuance of Proof of Age Attestations (recommended 30 per batch); re-identification of the user at least every three months; no revocation mechanism, on grounds that it would significantly increase complexity for providers and relying parties. Claims are statically calculated from the enrolment document, so a new attestation is required on reaching the threshold age.

Notes

No standard in this table addresses the effect of an age check on the child subjected to it. The 30107 series governs resistance to spoofing; the 19795 series governs accuracy and demographic differentials; 27566-1 and 2089.1 govern process, privacy and assurance levels. The behavioural consequence of requiring a child to present their face to a camera on demand falls outside the scope of all of them.

Appendix K — Proof of identity documents required for age verification

This appendix is confined to age verification: the establishment of age from identity information contained in an identity document bearing a date of birth, as defined in ISO/IEC 27566-1:2025. It excludes age estimation (facial or voice analysis, email footprint) and age inference (credit card possession, behavioural profiling, account signals).

Where a cell records that no list is published, this reflects the absence of any accepted-document list in the platform's own publicly available documentation at the time of writing; it is not asserted as proof that no such list exists internally or in an undocumented in-app screen.

K1 — Document-based age verification under the Australian SMMA Act 2024

Table K1 — The document route only. Facial estimation, email estimation, credit card and token routes are excluded as they are not age verification but estimation or inference. The bank route through ConnectID is included because it relies on a document check the bank made at the time of account opening, therefore relying on those ID documents required at that time.

Platform	Owner	Chain of processing for the ID check	Document types accepted	Selfie match to document required?	Notes and gaps
Facebook	Meta	Meta → Yoti Ltd (UK). Yoti sub-processes to Yoti India (manual fallback document authentication and face match, where the client requests manual review) and stores with AWS EMEA.	Government-issued photo ID bearing a face photo and date of birth. Meta names passport, driving licence and national ID card. No Australia-specific list published.	Yes, where the client configures it. Yoti captures multiple images and face-matches the clearest against the document photo.	Meta retains the ID image up to 30 days; Yoti deletes on completion, or at 28 days where manual review is configured. Yoti's own notice records the India transfer for manual document checks.
Instagram	Meta	As Facebook: Meta → Yoti Ltd → Yoti India (manual fallback) → AWS EMEA.	As Facebook: government-issued photo ID with face photo and date of birth (passport, driving licence, national ID card).	As Facebook.	As Facebook. Meta retains the ID image up to 30 days.
Kick	Stake.com (Easygo)	Kick → k-ID (Kidentify Pte Ltd, Singapore) → Veratad Technologies LLC (US) performs the document check. k-ID receives the result, which can include date of birth or an age range; the platform receives pass/fail.	Photo ID, front and back, all required information readable. No document-type list published by Kick; via Veratad these are passports, driver's licences and national ID cards.	Same k-ID conditional face-match provision as Snapchat: a face scan "might also" be required to match the document photo.	"Not all verification methods are available to all Viewers." Manual review available via Kick Support — an undocumented human-in-the-loop route.

Platform	Owner	Chain of processing for the ID check	Document types accepted	Selfie match to document required?	Notes and gaps
Reddit	Reddit Inc.	Reddit → Persona Identities Inc (US). Persona "may engage third parties" without naming them.	Government-issued ID with date of birth. No per-country accepted-document list published.	Yes — Persona's documented flow is ID capture followed by selfie and liveness, matched to the document photo.	Reddit's help page and Persona's own page contradict each other on whether an ID route is available at all: Reddit lists ID, facial selfie and bank; Persona states Australian users may use only government ID or facial estimation.
Snapchat	Snap Inc.	Snap → k-ID (Kidentify Pte Ltd, Singapore) → Veratad Technologies LLC (US) performs the document check. k-ID receives the result, which can include date of birth or an age range; the platform receives pass/fail.	Passport, driver's licence, or any other state-issued ID. Birth certificates expressly not supported. Front and back scan.	"Enhanced security implementations might also require a face scan to confirm the user's face matches the photo in the ID document" (k-ID). Not stated whether Snap enables this.	The only platform publishing an explicit exclusion. Parental approval is not available in Australia.
Threads	Meta	As Facebook: Meta → Yoti Ltd → Yoti India (manual fallback) → AWS EMEA.	As Facebook: government-issued photo ID with face photo and date of birth.	As Facebook.	eSafety confirms Threads uses the same Meta/Yoti selfie-or-ID route.
TikTok	ByteDance	TikTok → unnamed "service providers, who help us check if your ID is authentic". For TikTok LIVE: TikTok → Incode (which has acquired MetaMap 2024, AuthenticID 2025, Identiq 2026).	Government-issued photo ID. No accepted-document list published for Australia.	Not stated for the main flow. Yes for TikTok LIVE: Incode requires ID photograph plus a real-time selfie.	TikTok is the only platform naming no ID sub-processor at all. A user cannot establish who receives their document before submitting it.
Twitch	Amazon	No ID route in the primary flow. Facial age estimation only, via k-ID → Privately SA. Appeals are handled by Twitch Support directly.	None offered in the check itself. Twitch's help page states that a user who believes their estimated age is incorrect "can appeal through Twitch Support with proper age documentation" — the document type is unspecified.	Not applicable.	The only platform of the ten offering no ID route at the point of the check: the choice is a face scan or an appeal after failure. Verification attempts are capped daily; exceeding the limit means a 24-hour wait.

Platform	Owner	Chain of processing for the ID check	Document types accepted	Selfie match to document required?	Notes and gaps
X	X Corp.	X → Stripe Inc. (US) and/or Persona (US) for X Premium ID verification. X does not state who performs the check for Australian users.	Government-issued ID, used to estimate age by AI. No list published.	Yes, for the Stripe Identity route: passive selfie matched to the document. Stripe advises clients in the EU to offer a non-biometric alternative.	X's ID route is tied to the Premium subscription flow; free-tier availability is inconsistent and users report repeat prompting after completion.
YouTube	Google (Alphabet)	Google performs the ID check in-house. No third-party ID processor.	Valid government ID. Google names driver's licence, passport and national ID card generally; for the Australian flow specifically driver's licence, Proof of Age card, or passport.	Not required for the ID route (ID upload and facial estimation are separate routes).	Google gives no deletion commitment: ID data "stored securely", "won't be made public", and may be used to "improve our verification services for Google products".

Table K1 — The document route only. Facial estimation, email estimation, credit card and token routes are excluded as they are not age verification. The bank route through ConnectID is included because it rests on a document check the bank made at account opening.

K2 — The k-ID chain: one orchestrator, five downstream sub-processors in Australia

Kick, Snapchat and Twitch each contract with k-ID, but k-ID performs no check itself. It routes the user to one of the sub-processors below. k-ID receives the result of the check, which for an ID check can include date of birth or an age range, and passes only a pass or fail to the platform. Sources: k-ID Privacy Policy (10 July 2026); k-ID Developer Hub, Verification methods; k-ID Trust Center sub-processor list.

Sub-processor	Legal entity and location	What it does in the k-ID chain	Is this age verification?	Platforms reached
Veratad	Veratad Technologies LLC, Teaneck, New Jersey, USA	Age checks based on ID documents, and server-based facial age estimation. Receives the document; returns the result to k-ID, which can include date of birth or an age range.	Yes — the only true verification route in the k-ID chain.	Kick, Snapchat
Privately SA	Privately SA, Lausanne, Switzerland	On-device facial age estimation. Images do not leave the device.	No — estimation.	Kick, Snapchat, Twitch
Stripe	Stripe Inc., USA / Ireland	Credit or debit card age check. Small authorisation placed and refunded.	No — inference from card possession.	Kick, Snapchat (adult routes)

Sub-processor	Legal entity and location	What it does in the k-ID chain	Is this age verification?	Platforms reached
Australian Payments Plus (ConnectID)	Australian Payments Plus Ltd, Sydney, Australia	Bank-based age check. The bank has already verified identity from documents at account opening; ConnectID returns a threshold assertion.	Verification, but the document step sits upstream at the bank, outside the age check.	Kick, Snapchat
OpenAge Initiative (AgeKey)	OpenAge Global Initiative Pte Ltd, Singapore (launched by k-ID after acquiring Opale.io, France, November 2025)	Reusable on-device token recording age range, method used, and date created.	No — a record of a prior check, whatever its type.	Kick, Snapchat; announced for Meta

Sources — which identity documents are required or accepted

1. Facebook Help Centre. Types of IDs that Facebook accepts. https://www.facebook.com/help/159096464162185
2. Facebook Help Centre. How using your ID to verify your age works on Facebook. https://www.facebook.com/help/671565971895323
3. Facebook Help Centre. Confirming your age on Facebook. https://www.facebook.com/help/958848942357089/
4. Meta Help Centre. Types of ID that Meta supports for ID verification. https://www.meta.com/help/policies/804481810668573/
5. Yoti Ltd. Age Verification Privacy Notice (Identity Document Verification). Last updated 23 February 2026. https://www.yoti.com/privacy/age-verification/
6. Yoti Ltd. Supported documents. https://developers.yoti.com/digital-id/supported-documents
7. eSafety Commissioner. Information from age-restricted platforms. 2026. https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions/information-from-age-restricted-platforms
8. Kick. Age assurance FAQs. https://help.kick.com/en/articles/12947994-age-assurance-faqs
9. Reddit Help. How does age verification work in Australia? https://support.reddithelp.com/hc/en-us/articles/47163443918228-How-does-age-verification-work-in-Australia
10. Persona. FAQ: For Reddit users using the Persona service. https://help.withpersona.com/articles/7F6BaF9h8Fxf0XWkwQscXN/
11. Persona. Age assurance: how Persona can help determine an individual's age. https://help.withpersona.com/articles/6sBm3U81qDvDMyQ4bxTKVB/
12. Persona. What is selfie identity verification, and how does it work? https://withpersona.com/blog/what-is-selfie-identity-verification-and-how-does-it-work
13. Snap Inc. Newsroom. Implementing Australia's Social Media Minimum Age Law. https://newsroom.snap.com/australia-social-media-minimum-age-law
14. Snapchat Support. Account Restrictions for Australia Users. https://help.snapchat.com/hc/en-us/articles/42550358600084-Account-Restrictions-for-Australia-Users
15. Snapchat Support. Age Verification to Continue Using Snapchat. https://help.snapchat.com/hc/en-us/articles/35461114902804-Age-Verification-to-Continue-Using-Snapchat
16. k-ID Developer Hub. Verification methods (ID scan verification). https://docs.k-id.com/concepts/verification-methods
17. k-ID (Kidentify Pte Ltd). Privacy Policy. Last updated 10 July 2026. https://k-id.com/privacy-policy
18. k-ID Trust Center. Sub-processors. https://security.k-id.com/subprocessors
19. Veratad Technologies LLC. Privacy Policy. https://veratad.com/privacy-policy
20. TikTok Newsroom. How TikTok is complying with a new government law in Australia. https://newsroom.tiktok.com/how-tiktok-is-complying-with-a-new-government-law-in-australia

21. TikTok. Confirm your age on TikTok. https://www.tiktok.com/support/faq_detail?id=7543604789788039686
22. TikTok. Age requirements for TikTok LIVE. https://www.tiktok.com/support/faq_detail?id=7581820708762753548
23. Twitch. Age Verification on Twitch. https://help.twitch.tv/s/article/age-verification
24. X Help Centre. Age Assurance on X. https://help.x.com/en/rules-and-policies/age-assurance
25. X Help Centre. Australia Online Safety Information. https://help.x.com/en/rules-and-policies/australia-resources
26. Stripe. Verification checks (Stripe Identity). https://docs.stripe.com/identity/verification-checks
27. Google Account Help. Access age-restricted content and features. https://support.google.com/accounts/answer/10071085
28. YouTube Help. Watch age-restricted videos. https://support.google.com/youtube/answer/10070779

Sources accessed 16 July 2026

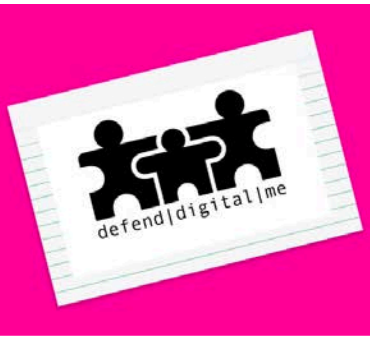

DEFENDDIGITALME.ORG | (CC BY-NC-SA 4.0)