

Policy Research Note

Age Verification for Online Platforms: Policy Considerations for Canada with Reference to the Adult Content Sector

June 2026

Executive Summary

Online age verification is technically achievable, already in use, and will need to form part of any serious Canadian legislative response to online harms affecting children, whether that is a ban on social media or the enforcement of age-appropriate design. This research note examines the current state of age verification technology and practice, drawing on the experience of adult content platforms operating under the United Kingdom's Online Safety Act 2023. The note further outlines a number of policy implications for Canadian policymakers.

Seven distinct methods of age assurance have been recognized as compliant under the UK framework: facial age estimation, photo ID matching, digital identity services, credit card and open banking checks, mobile network operator verification, and email-based estimation. No method is perfect, and policymakers should not expect one to be. Methods that rely on an existing financial or contractual relationship are generally more straightforward to implement but are unsuitable where the relevant threshold falls below contracting age, as it would under proposed under-16 social media restrictions. Facial age estimation is currently the most inclusive option for all age groups, but its accuracy is best evaluated through independent testing rather than vendor claims, and it depends on a robust data-protection regime to address the privacy implications of biometric processing.

The paper identifies a number of genuine trade-offs on the journey to policy adoption. Privacy and anonymity concerns are real, particularly for photo ID and biometric methods, and can only be partially addressed through data minimization and accredited third-party providers. With respect to age verification driving users to less safe services, the evidence is mixed: adult VPN uptake and use has been significant in some jurisdictions, but there is no corresponding evidence of increased VPN use among children. Freedom of expression and children's rights, including the risk of ex-

Sommaire exécutif

La vérification de l'âge en ligne est techniquement réalisable, déjà utilisée à grande échelle et devra faire partie de toute réponse législative canadienne sérieuse visant à réduire les préjudices en ligne touchant les enfants, qu'il s'agisse d'une interdiction des médias sociaux ou de l'adoption d'exigences de conception adaptées à l'âge des utilisateurs. Cette note de recherche examine l'état actuel des technologies et des pratiques de vérification de l'âge en s'appuyant sur l'expérience des plateformes de contenus pour adultes assujetties à la Online Safety Act 2023 du Royaume-Uni. Elle met également en lumière plusieurs implications importantes pour les décideurs publics canadiens.

Le cadre réglementaire britannique reconnaît actuellement sept méthodes distinctes d'assurance de l'âge comme étant conformes : l'estimation de l'âge à partir du visage, la vérification par pièce d'identité avec photo, les services d'identité numérique, les vérifications fondées sur les cartes de crédit et les services bancaires ouverts (open banking), la vérification par les opérateurs de téléphonie mobile ainsi que l'estimation de l'âge à partir d'une adresse courriel. Aucune de ces méthodes n'est parfaite, et les décideurs ne devraient pas s'attendre à ce qu'une solution unique le soit. Les méthodes reposant sur une relation financière ou contractuelle préexistante sont généralement plus simples à mettre en œuvre, mais elles sont peu adaptées lorsque le seuil d'âge visé est inférieur à l'âge légal permettant de conclure un contrat, comme ce serait le cas dans le cadre de restrictions proposées visant les utilisateurs de moins de 16 ans sur les médias sociaux. À l'heure actuelle, l'estimation de l'âge à partir du visage constitue l'option la plus inclusive pour l'ensemble des groupes d'âge. Toutefois, sa précision devrait être évaluée au moyen de tests indépendants plutôt qu'à partir des affirmations des fournisseurs, et son utilisation exige un solide cadre de protection des données afin de répondre aux enjeux de confidentialité associés au traitement de données biométriques.

clusion and the chilling of lawful, sensitive uses, must be addressed explicitly in framework design.

The UK's enforcement record, and Australia's early experience with its social media age ban, together suggest that compliance is achievable but depends heavily on platform cooperation and regulatory capacity. For Canada, age verification requires a broader legislative architecture than is currently in operation — including a privacy regime, an independent regulator, and a Child Rights Impact Assessment — not a standalone technical fix. Age verification is one instrument among several, and how it is designed and governed will matter as much as whether it is adopted.

Ce rapport met en évidence plusieurs compromis réels auxquels les décideurs devront faire face dans l'élaboration des politiques. Les préoccupations liées à la protection de la vie privée et à l'anonymat sont légitimes, particulièrement dans le cas des méthodes basées sur les pièces d'identité avec photo et les données biométriques. Ces risques ne peuvent être atténués qu'en partie grâce à des pratiques de minimisation des données et au recours à des fournisseurs tiers accrédités. Quant à l'argument selon lequel la vérification de l'âge pousserait les utilisateurs vers des services moins sécuritaires, les données probantes demeurent mitigées : dans certaines juridictions, l'adoption et l'utilisation de réseaux privés virtuels (RPV) par les adultes ont été importantes, mais aucune preuve comparable n'indique une hausse de leur utilisation chez les enfants. Les considérations relatives à la liberté d'expression et aux droits de l'enfant — notamment les risques d'exclusion et les effets dissuasifs sur certains usages légitimes et sensibles — doivent être explicitement prises en compte dans la conception des cadres réglementaires.

L'expérience du Royaume-Uni en matière d'application de la loi et les premiers enseignements tirés de l'interdiction d'accès aux médias sociaux fondée sur l'âge en Australie, suggèrent qu'une mise en conformité est possible, mais qu'elle dépend largement de la coopération des plateformes et de la capacité des organismes de réglementation. Pour le Canada, la vérification de l'âge nécessite une architecture législative plus vaste que celle actuellement en place — comprenant notamment un régime robuste de protection de la vie privée, un organisme de réglementation indépendant et la réalisation d'évaluations des répercussions sur les droits de l'enfant — plutôt qu'une simple solution technique autonome. La vérification de l'âge constitue un outil parmi plusieurs autres, et son efficacité dépendra autant de sa conception et de sa gouvernance que de son adoption elle-même.

1. Introduction

Debates about online harms and children's safety increasingly centre on two broad approaches: restricting younger users' access to online platforms and services¹, and reforming the design of those platforms to reduce exposure to harmful content². The two are often combined, and both are under active consideration across a range of jurisdictions, including Canada³.

As of early 2026, Canada is among several countries considering social media restrictions for children under 16, at both the federal and provincial levels⁴. This follows Australia's Online Safety Amendment (Social Media Minimum Age) Act 2024, which imposed an enforceable ban on social media accounts for that age group. Litigation in the United States has added further momentum: courts have found certain platforms negligent⁵ or misleading⁶ on questions of platform safety. The California case is notable for using the law to target platform design rather than the content the platforms carry.

Underlying many of these harms is the advertising-based business model common among major online platforms⁷. That model rewards designs which maximise user engagement, including among younger users. Age-based access restrictions do not address those incentives directly, but they may serve as an interim measure while longer-term reforms to platform design and duty-of-care obligations are developed.

Any age-based access policy, whether a ban or an age-appropriate design standard, depends on a workable age assurance or verification regime. This note examines the current state of that technology and practice, drawing in particular on adult content platforms operating under the United Kingdom's Online Safety Act 2023 (OSA). The sector is specific, but the lessons are not: they apply across the online environment, including to social media. The argument that follows is that age verification is technically achievable and already in use, but that it is one instrument among several, and that how it is designed, governed, and bounded matters as much as whether it is adopted.

2. The Case for Age Verification

Age-based restrictions are well-established offline. In most jurisdictions, adolescents cannot buy alcohol or tobacco, or vote or drive, until they reach an age set by law. These thresholds rest in part on the development of the adolescent brain, and in particular on the prefrontal cortex, which governs impulse control and decision-making and matures only in early adulthood⁸.

The online environment presents analogous, if more complex, challenges. Research has identified associations between high levels of social media use among adolescents and adverse mental health outcomes⁹, reduced attention spans¹⁰, and increased exposure to harmful actors, including those who target children for sexual purposes¹¹. Platform design features, such as infinite scroll, algorithmic recommendation, and variable

reward mechanisms, have been argued to activate neurological pathways associated with addictive behaviour¹². The strength and direction of these effects, however, are debated. The most comprehensive reviews describe the overall association between time spent on social media and adolescent mental health as small and mixed, while more consistent evidence of harm is found in relation to problematic or compulsive use, particularly among adolescent girls¹³. The picture the research supports is not that platforms are harmless, nor that they are uniformly harmful, but that the risk is real and unevenly distributed, and that design and patterns of use matter more than screen time alone.

The difficulty for regulators is the breadth of online access. Unlike buying alcohol or tobacco, using the internet is a near-continuous activity spanning many services, most of them benign, or capable of being used safely if designed in an age-appropriate way. Any regulatory approach must therefore be targeted and proportionate, distinguishing genuinely harmful services from those that pose little risk to younger users.

The smartphone compounds the difficulty. It is a persistent gateway to online services, available throughout the day, and the competitive dynamics of the attention economy push platforms to maximise engagement, often at the expense of users' wellbeing¹⁴.

2.1 The axes of reform: platform versus device, and content versus design

Because operating systems and app stores are so concentrated, a duopoly of Google and Apple in each, some have called for age verification at the device level. Google or Apple would verify a user's age once, and that result would govern every app downloaded or site visited. Neither company has welcomed the idea; both argue that verification is the responsibility of the individual service or website.

A second choice is whether to regulate the content the major platforms carry, or instead to govern the design features that drive continual use.

Children now most often encounter harmful content, such as pornography, self-harm material, and violence, through general-purpose services such as social media rather than through dedicated pornography sites reached by search¹⁵.

Data from the UK's Office of Communications (Ofcom) shows the scale of children's exposure. Ofcom found that roughly eight per cent of children aged 8 to 14 in the United Kingdom visited an online pornography site or app in a given month, including around three per cent of 8 and 9 year-olds, the youngest group studied¹⁶. Boys aged 13 and 14 were the most frequent visitors. The figures probably undercount exposure among under-18s, since older teenagers were not the study's focus. Surveys in the UK and Canada suggest that a large majority of adults, roughly 80 per cent, support age verification for online pornography.

The UK's Online Safety Act (OSA), in force from 2023 with its age verification provisions operative from July 2025, is the most significant legislative development in this area to date. This note draws on the experience of its implementation as a reference point for Canadian policymakers.

3. How Does Age Verification Work?

The UK OSA mandates a “duty of care” for online service providers, particularly those catering to under-18s¹⁷. The duty of care includes protecting children from harmful content, including pornography and violent or self-harm-promoting material. The Act sets out obligations for highly effective age assurance on all sites and apps that display or publish pornographic content. They must ensure that children cannot encounter pornography on their service. Other user-to-user services that allow pornographic content must ensure the same level of safety.

Before the implementation of the OSA, most online platforms relied on self-declaration as their primary mechanism for age verification. Users were asked to confirm that they meet a minimum age threshold, typically by clicking a checkbox or entering an (unverified) date of birth. This approach has proven largely ineffective at preventing access by younger users.

Canadian governments and legislators have attempted to introduce online harms legislation, including attempts with provisions related to age verification, without success. The most recent attempt, Bill S-210, did not complete passage before the session ended¹⁸ and has since been reintroduced as Bill S-209¹⁹, incorporating elements of the UK's legislative model.

3.1 The Main Age Verification Methods

Several jurisdictions, including the United Kingdom and dozens of US states, have already implemented age verification regimes. In the UK, the Information Commissioner's Office has issued guidance on age assurance technologies²⁰ and Ofcom's guidance sets four criteria for highly effective age assurance: accuracy, robustness, reliability, and fairness.

The following seven methods are currently recognised as compliant age assurance approaches under the UK's Online Safety Act:

- Facial age estimation
- Photo ID matching
- Digital identity services
- Credit card age checks
- Mobile network operator (MNO) age checks

- Open banking
- Email-based estimation

3.1.1 Facial age estimation

Facial age estimation has a third-party service compare a live image of the user against a training dataset to estimate an age. It requires no identity document or personal information such as a name or date of birth. Unlike facial recognition²¹, it does not match the image against a database of known identities, and it can operate statelessly, without storing the photo or the biometric data derived from it.

Yoti, a UK provider, began working with Instagram in 2022 to verify users' ages²². According to Yoti, when asked to prove their age, 81 per cent of users chose facial estimation over uploading an identity document. Instagram has also introduced Teen Accounts, which by default give under-18s the strictest privacy settings.

3.1.2 Photo ID matching

The user uploads an image of a government-issued document, such as a driver's licence or passport, alongside a live selfie, and the two are compared to confirm the document belongs to them. Checks for security features, watermarks, and signs of tampering establish that the document is genuine, and software analyses the selfie to confirm a real, three-dimensional person rather than a photograph held to the camera.

3.1.3 Digital identity services

A digital identity provider, which may be a government²³ or a private-sector entity, verifies a person's documents through photo ID matching or open banking, and issues a ‘verified credential’ attesting to facts such as name and age. The person can then share that credential with third parties in future, rather than the original documents.

Reusable digital IDs support data minimisation: a credential can confirm that a person is over 18 without revealing their date of birth or exact age, and the requesting site receives only a yes-or-no answer

3.1.4 Credit card age checks

Users provide their card details, which a payment processor validates²⁴. Because the minimum age to hold a credit card independently is 18 in most jurisdictions, a valid card confirms adult status.

When a site uses a credit card for age verification, they typically:

- Run a small authorization (often £0.00 or £1.00) against the card to confirm it is real and active, not to charge the user but to validate that the card exists.
- Check the card details (number, expiry, CVV) against the card network to confirm validity.
- Sometimes check the billing address via an Address Verification

Service (AVS), which adds a layer of identity confirmation.

The card number itself encodes the issuing bank and card type, but not a user's age directly. The age inference comes from the assumption that the card issuer already checked.

3.1.5 Mobile network operator age checks

Each UK mobile network operator applies, by default, a content restriction filter (CRF) to under-18s under an agreed code of practice, blocking access to age-restricted sites over mobile internet on both pay-as-you-go and contract SIMs. In the United States and Canada, by contrast, parental controls are optional products offered by the major operators.

In CRF markets, an adult can remove the filter by proving their age²⁵, by presenting ID at a retail outlet or uploading documents to the operator's website.

At a later age check, the user enters their mobile number and confirms possession of the device with a one-time password sent by SMS. The provider then draws on SIM registration and carrier metadata to establish whether the SIM is held by an adult²⁶.

In effect, the check asks the network whether the SIM's content restriction filter has been lifted. If it has, the account holder is recorded as an adult, and that yes-or-no result is passed to the relying party²⁷.

3.1.6 Open banking

In the UK, bank-based verification works through Open Banking²⁸, the framework regulated by the Financial Conduct Authority. The user selects their bank, logs in with their normal credentials, and consents to share a single fact: that they are over the required age. The bank returns a yes-or-no answer; no financial data, transaction history, or account details are disclosed, and the process usually takes under 30 seconds.

The process has four steps, and from the user's perspective, it feels like logging into their online banking.

- On the business's website or app, the user is presented with the option to verify their age or identity. They select bank verification from the available methods.
- A list of supported banks appears, and the user picks theirs. In the UK, this covers the major high street banks and building societies that participate in the Open Banking ecosystem, which was established under the Payment Services Regulations 2017 and is regulated by the Financial Conduct Authority.
- The user is then redirected to their bank's secure login page and authenticated using whatever method they normally use for online banking, whether that is a password, biometric login through their banking app, or a one-time code. They are then asked to consent to sharing a specific piece of information, typically just a confirmation of whether they are over 18.

- Finally, the bank sends a response confirming the user's status. The user is redirected back to the original site or app, now verified.

In this scenario, a user does not upload a document, take a photograph, or provide biometric data to a third party. The bank handles the authentication and data confirmation within its own secure infrastructure.

3.1.7 Email-based estimation

Email-based estimation takes a user's email address and analyses other online services where it has been used, such as banks or utility providers, to estimate their age²⁹. It does not scan the contents of the inbox.

The address can be passed by the platform through an API, or entered manually, in which case the user confirms it with a one-time password³⁰. The provider then analyses where the address has been used, on sites such as banks, mortgage lenders, or utilities, to build a digital footprint, and returns an estimated age or age band, for example over 18 or over 13.

3.2 A Comparative Assessment of Methods

No single method of age verification is perfect, and policymakers should not expect one to be. A robust regime could combine the various methods listed above, but any combination risks adding to a user's reluctance to participate. Governments will need to balance robustness with ease of adoption. They will also need established and effective privacy and data-protection regimes in place to reassure users and ensure strict compliance by all third-party age assurance providers.

Methods that rely on an existing contractual relationship, namely credit card verification, MNO checks, and open banking, are generally more straightforward to implement, and Ofcom classifies them as 'highly effective'³¹. These approaches leverage trusted institutional partners with established data protection frameworks, reducing the risk of data misuse.

These methods have clear limits. Under-18s can obtain credit cards or over-18 email addresses, and not all adults have a card, a mobile contract, or a bank account. They also depend on the user being of legal contracting age, which makes them unsuitable wherever the relevant threshold falls below it, as it would under a proposed minimum age of 16 for social media. A 16-year-old is unlikely, and in many jurisdictions unable, to hold a credit card, a phone contract, or a driver's licence.

Photo ID matching and email-based estimation, though also classed as highly effective, carry greater privacy implications. Most internet users have an email account, but not all adults hold a government-issued photo ID, and few children do. And submitting a copy of an identity document to a third party, or letting an external provider map a user's

wider presence online, introduces risks of data breach or identity fraud if the data is mishandled.

A national digital ID, issued to every citizen, is elegant in principle and operates in countries such as Estonia³², but it faces opposition elsewhere, including in the United Kingdom³³, where Prime Minister Keir Starmer recently had to reformulate his government's plans in the face of public objection on privacy grounds.

Of the methods on offer, facial age estimation is the most inclusive: everyone has a face and it does not depend on contracting age or possession of identity documents, making it applicable across age groups. And while some studies have shown that age estimation struggles with certain skin tones³⁴, accuracy rates are reported by the leading providers themselves to be high across all skin tones and ages, with 'true positive rates' of approximately 99.3 per cent³⁵. The technology can still be circumvented, for example through 'injection attacks' that substitute a different person's image into the verification stream, and providers must keep developing countermeasures as workarounds evolve. This figure is best read as a vendor's claim about its own product rather than a settled benchmark. Independent evaluation gives a more textured picture. The United States National Institute of Standards and Technology, in its most authoritative public benchmark, reports mean absolute errors of roughly two to five years for young adults, with accuracy that varies by sex, age, image quality, and region of birth rather than holding uniformly across groups. Peer-reviewed studies document the same demographic bias, with one analysis finding it more pronounced in AI systems than in human judgement³⁶. In practice, providers manage this imprecision by setting a conservative 'challenge age' well above the legal threshold, as this note describes for Yoti, which makes the method workable but routes more borderline users into secondary checks. The fair conclusion is that facial age estimation is useful and improving, but that its accuracy is best judged through independent testing, such as the NIST evaluations, rather than vendor figures.

Working alongside regulators, facial age estimation providers in the UK are taking steps to mitigate against broader privacy concerns³⁷. In general, age estimation providers do not require a user's personal details from the site seeking verification, and do not keep a database of the images sent to them or require a comparison with a government-issued ID (only a training sample of images), mitigating some of the most serious privacy criticisms (more on which in section 4.0 below).

Even if a facial image is not retained after processing, it must be captured and analysed. That processing involves biometric data, which is among the most sensitive categories of personal information, and it creates risks during transmission even where storage is minimised. The European Data Protection Board's 2025 statement on age assurance³⁸ noted that even transient processing of biometric data engages serious rights including non-discrimination and bodily integrity. Countries like Canada with less developed online privacy frameworks might

struggle to reassure citizens of the safety of facial age estimation technologies.

In summary, effective age estimation and verification depend on system design, data minimisation, third-party auditing, and the proportionality of the deployment, with robust technical and regulatory standards needed throughout.

3.3 Case Study: Australia's Social Media Age Ban

The existence of 'highly effective' methods does not guarantee effective implementation. Age verification adds friction to the user journey and can cut against commercial incentives³⁹, and even regulators operating under Europe's most advanced data-protection regimes are struggling to secure compliance⁴⁰.

Australia is the clearest case. Its Online Safety Amendment (Social Media Minimum Age) Act 2024⁴¹ took effect in late 2025, making it the first jurisdiction to enforce a ban on social media accounts for under-16s. The Act penalises platforms that fail to take 'reasonable steps' to enforce the ban, including an age verification process, but does not prescribe the methods⁴². eSafety's guidance instead encourages a layered, or 'waterfall', approach to age assurance across the user journey.

Despite nearly five million account terminations, teenagers and parents reported that many children kept access to their accounts⁴³. This example shows how heavily regulatory outcomes depend on robust enforcement and on the active cooperation of platforms.

An assessment by Australia's eSafety Commissioner⁴⁴, three months into the compliance period, identified several deficiencies across the ten platforms covered by the ban (Facebook, Instagram, Kick, Reddit, Snapchat, Threads, TikTok, Twitch, X, and YouTube), including:

- Permitting children who had declared themselves to be under 16 prior to the ban to subsequently claim they were over 16 without adequate verification.
- Allowing under-16s to repeatedly attempt the same age assurance method.
- Insufficient safeguards to prevent new accounts being created by under-16s.
- The absence of effective reporting mechanisms for parents and others to flag non-compliant accounts.

eSafety is actively investigating potential non-compliance in relation to Facebook, Instagram, Snapchat, TikTok and YouTube. A decision on any enforcement action is due by the middle of 2026.

As part of the review, the Australian Minister for Communications, Anika Wells, registered a new rule that ties the definition of an age-restricted social media platform more closely to the features that pose

the greatest risk to children⁴⁵: recommender systems, infinite scroll, feedback features such as likes and comments, and time-limited features such as disappearing stories.

eSafety is due to review its guidance for the social media platforms in mid-2026.

4. Concerns and Trade-offs

Age verification involves a set of genuine trade-offs, and opponents raise several substantive concerns. A credible regime has to weigh each of them honestly rather than treat them as obstacles to be cleared away. This section sets out the main concerns and what a careful design would do about them.

4.1 Privacy, Anonymity, and Security

Several verification methods require users to share personal data, including biometric information or identity documents, with third-party providers. That creates risks of data breach, identity fraud, or misuse of information about a person's online behaviour. The risks are most pronounced for photo ID matching and facial age estimation, where users hand sensitive data to entities that may lack the data-protection frameworks of established institutions such as the major banks and mobile network operators.

While valid, these concerns must be weighed against the fact that users in most cases have already shared much of the same information with the devices and social media platforms legislators are seeking to regulate for harm. This includes face scans to unlock devices and political opinions⁴⁶. Adding a third-party into age verification introduces a degree of additional risk, and that risk can be only partly contained, and only where a robust data privacy regime and strong enforcement by a fit-for-purpose regulator, i.e. one dedicated to the task of online safety as in the Australian model, are already in place. Privacy-preserving designs can reduce this risk substantially. The open-banking checks and reusable digital credentials described above can confirm that a user meets an age threshold without disclosing identity to the service, and that data minimisation is a real mitigation, not a cosmetic one. But it does not remove the risk, and it depends on a data-protection and enforcement regime that Canada does not yet have. The honest framing, then, is that the privacy implications of third-party verification are a genuine trade-off, to be managed by careful design and regulation rather than treated as already resolved.

Where verification is required to reach adult content, there is a further concern: that a malicious actor could coerce, embarrass, or expose someone for visiting a site to consume lawful content. As noted above, verification can be designed so that identity is not linked to the check, which limits the risk of blackmail or 'honeypotting'.

A closely-related concern is anonymity. The ability to use online services without disclosing one's identity is not merely a convenience. It protects people seeking sensitive health information, survivors of abuse, LGBTQ youth, and others for whom identification carries real risk. Verification methods that return only a yes-or-no answer limit how much identity is exposed to any single service, and that is a meaningful safeguard. Even so, mandatory verification narrows the space for the anonymous and pseudonymous use that the open internet has historically allowed. The blackmail and honeypotting risk described above, in the adult-content context, is one instance of a broader chilling effect on lawful, sensitive, and protected uses, and it should be named as such.

A third concern, broader still, is systemic security. Requiring identity or biometric checks across a wide range of services does not create a single risk; it creates many. Each verification point, and each third-party provider, becomes a store of sensitive data, and therefore a target. The European Data Protection Board has cautioned that even the transient processing of biometric data engages serious rights. Data minimisation, short retention periods, and the use of accredited providers can reduce the size of each target, though they do not remove it. Related to this is the problem of function creep: verification infrastructure built for one purpose tends, over time, to be extended to others, and an architecture of identity checkpoints, once built, is difficult to confine. None of these concerns is, on its own, a decisive argument against age verification. But a measured assessment has to weigh them, and the strongest case for a regime is one designed with them in view rather than around them.

4.2 Effectiveness and Enforcement

Critics note that age verification may displace rather than prevent access to harmful content. Where mandatory verification has been introduced, in the UK and several US states, the use of Virtual Private Networks (VPNs) has risen substantially⁴⁷. Following the introduction of age verification requirements in Florida, VPN usage increased by over 1,000 per cent⁴⁸. UK regulator Ofcom similarly recorded a doubling of VPN use among adults after the OSA's provisions came into force⁴⁹.

Several of the US states that have brought in age verification measures are now seeking ways to counter increased VPN use. Utah's Senate Bill 73⁵⁰, for example, deems any user physically located in the state to be accessing a site from there regardless of any VPN use, and prohibits commercial entities hosting material harmful to minors from facilitating or encouraging its use to circumvent age checks.

Aggressive measures to police VPN use have drawn criticism, on the ground that they restrict the legitimate use of a technology that provides real privacy and security protections⁵¹.

The evidence on displacement is genuinely mixed. Research by the not-for-profit organisation Internet Matters found no corresponding rise in VPN use among children in the UK following the introduction

of age verification⁵². Whether displacement undermines a regime depends on what the regime is for. If the aim is to restrict adult access, the VPN surges suggest real leakage. If the aim is the narrower one of reducing children's exposure, the same evidence is less damaging, and the Internet Matters finding is mildly encouraging. Either way, these regimes are too new for a confident assessment, and displacement remains a real and unresolved concern.

4.3 Scope and Proportionality

Some critics argue that mandatory age verification should be narrowly targeted, applied to clearly harmful services such as pornography or gambling sites, rather than extended to general-purpose services such as social media. Broad application, they contend, would impose disproportionate compliance burdens and create real enforcement challenges, particularly since many of the relevant platforms are both free to use and based in other jurisdictions⁵³.

The proportionality concern is real. Narrow targeting reduces compliance burdens and enforcement difficulty, but it leaves untouched the social media services through which most children first encounter harmful content. One intermediate option, noted earlier, is to concentrate verification at the level of the primary app stores, so that a single age check governs access to multiple downstream services. That has administrative appeal, but it places compliance with a small number of dominant platform operators and does not reach harms that arise from content or design choices at the service level. Proportionality here is not a single setting but a design choice with costs on either side.

4.4 Freedom of Expression

Opponents of online age verification argue it curtails a user's ability to navigate the open web and access public-interest information⁵⁴. Legal challenges have been brought in Australia and the United States arguing that social media bans restrict children's rights to information and public discourse⁵⁵. While these arguments are not fully resolved, they represent a legitimate consideration that any Canadian legislative framework should address explicitly. The most significant recent development is in the United States, where in June 2025 the Supreme Court, in *Free Speech Coalition v. Paxton*, upheld a Texas age-verification law for adult content. The Court applied intermediate rather than strict scrutiny, on the reasoning that the law only incidentally burdened adults' access to protected speech⁵⁶. That reasoning matters for the Canadian debate because it frames the tension precisely. Age verification for harmful content also conditions adults' access to lawful material, and a Canadian framework will need to address that burden directly rather than treat freedom of expression as a question of children's access alone.

Social media is increasingly a primary route through which young people find and share information, and restricting access carries a

real expressive cost. That cost is mitigated, though not erased, by the range of other sources that remain open, including books, libraries, and reference sites such as Wikipedia. It is also why an age-appropriate design code is worth weighing alongside access restrictions. A design standard can reduce harm while preserving more of the informational and participatory value that social media provides.

4.5 Age Verification and Children's Rights

A further consideration, less often raised in the Canadian debate, is that age verification engages the rights of children themselves. Discussion of online safety tends to treat children only as subjects to be protected. They are also rights-holders. The UN Convention on the Rights of the Child, and the UN Committee on the Rights of the Child's General Comment No. 25 (2021) on children's rights in relation to the digital environment, set out a framework in which protection is one of three registers, alongside provision, meaning access to information, education and play, and participation⁵⁷.

General Comment No. 25 speaks directly to age assurance. It accepts that age verification can be a legitimate tool, but it cautions that the measures used to establish age must themselves respect children's rights. Age assurance should not result in the excessive collection of children's data, should not be used to profile or track them, and should not exclude children from services, information, or civic participation to which they are entitled. An age-verification regime that is careless about data minimisation, or that locks children out of beneficial services because they cannot establish their age, can therefore infringe the very rights it is intended to serve.

This has two practical implications for a Canadian framework. First, the rights of children should be assessed explicitly, rather than assumed to be satisfied by the goal of protection. The emerging instrument for this is a Child Rights Impact Assessment, which evaluates a proposed measure against the full range of children's rights before it is adopted⁵⁸. Second, the design of any regime should be informed by the views of young people themselves. Public and parental support for age verification is consistently high, but research with the teenagers who are themselves subject to these measures finds more ambivalence, including considered concerns about biometric and identity-based methods and what they mean for privacy⁵⁹. Engaging children as rights-holders, and not only as objects of protection, would strengthen both the legitimacy and the design of a Canadian regime.

4.6 Exclusion

A final concern is exclusion. Every verification method assumes that a user can produce something: a document, a payment card, a bank or mobile account, or a face that the system reads accurately. People who cannot are at risk of being shut out of lawful online services. This falls hardest on groups that are already marginalised, including people

without government identification, newcomers, low-income users without credit or banking relationships, and people whose appearance or documents do not match what a system expects. Offering several verification routes, and a non-digital fallback, narrows this problem, and a well-designed regime can do so. But a regime that looks proportionate in the aggregate can still be exclusionary at its edges, and that residual burden is not evenly distributed.

5. Lessons from the Adult Content Industry

The adult content sector is a useful reference point, for two reasons. First, adult sites were among the earliest platforms made subject to mandatory age verification, which gives an empirical basis for assessing different approaches. Second, the main route by which children encounter pornography is not adult sites themselves but social media.

The lessons drawn from the adult content sector are therefore relevant not only to pornography regulation, but to the broader regulatory environment, including social media.

5.1 The UK Regulatory Context

The UK's OSA distinguishes between two categories of online pornography service⁶⁰:

- Part 5 Services: Platforms that publish their own pornographic content (e.g. studios, subscription pay sites), where operators control the available material.
- Part 3 Services: Platforms hosting user-generated content, including tube sites, cam sites, and fan platforms.

In late 2025, Ofcom produced a report addressing the successes and challenges of implementing the Online Safety Act during the first year of its enforcement across both categories of service⁶¹. On age verification in the pornography sector, it said the following:

“...age assurance measures were implemented by major platforms including the top 10 pornography sites by the 25 July deadline, resulting in a **one-third reduction in UK visitors to pornography sites**. Ofcom has launched **enforcement investigations into over 80 pornography websites** for non-compliance and imposed penalties including £1m on a provider of multiple adult sites and £50,000 on a nudification app. The regulator **reviewed 104 risk assessment records** from service providers, **challenging 11 providers** to revise their assessments and taking enforcement action against two providers for failing to respond to information requests.”

Overall, Ofcom found that over half of the 100 most popular adult services in the UK had deployed age checks, and that those sites accounted for more than three quarters of daily visitors. As noted earlier,

there was no increase in children using VPN services to bypass age verification measures for these sites.

On children completing age verification checks, Ofcom said:

“...research also found a significant rise among children aged 8-17 who completed a highly effective age verification method, encountering an age check when trying to access age restricted content on an online service such as social media (from 30% to 47%) after the 25 July deadline. In the first month of the new age-check rules being in force, visits to third-party age-check sites grew significantly, with 7.5 million UK visits to the top five age assurance providers, compared to under a million in June.”

Ofcom also laid out specific improvements for the coming year (i.e. 2026), which included companies having to name a specific ‘senior responsible person’ in order to deliver organisational accountability to the regulator for all enforcement efforts.

As evidenced by Ofcom's enforcement activity, individual sites chose different ways to respond to the Online Safety Act. The following examples provide a good cross-section of those responses.

5.2 Adult Content Site Responses: A Range of Approaches

5.2.1 Pornhub

Canada-based Pornhub (a subsidiary of Aylo, which is owned by Ethical Capital Partners) is the world's most visited pornography site and has adopted various approaches to age verification across different jurisdictions.

In US states that passed age verification laws, Pornhub blocked its service entirely rather than comply⁶². In the UK, it initially complied, requiring an account and age verification through digital ID, facial estimation, or payment⁶³. Estimates of the resulting traffic decline varied widely. Ofcom put it at 31 per cent⁶⁴, the platform itself at 77 per cent, and third-party analytics at 47 per cent⁶⁵.

In January 2026, Aylo announced it would no longer participate in the UK system⁶⁶, arguing that the OSA had not protected minors but had instead driven them toward less regulated sites. It blocked all new UK users while allowing those already verified to continue⁶⁷.

In May 2026, it shifted again after Apple built an age verification measure into its latest iOS update, which restored access for any user running that version⁶⁸.

5.2.2 xHamster and xVideos

xHamster, the second most visited pornography site in the UK, chose to comply, using facial age estimation provided by Yoti⁶⁹. Third-party data indicated a 39 per cent decline in daily visits in the two weeks

after implementation⁷⁰.

xVideos adopted similar measures⁷¹ with third-party data indicating a 47 per cent traffic decline⁷². These figures are disputed. Some analyses suggest more modest decreases, in part because adult users reach the sites through VPNs⁷³.

5.2.3 4chan

The imageboard 4chan declined to comply, arguing that Ofcom's notices placed no binding obligation on a US-based entity and that the UK's approach conflicted with First Amendment principles⁷⁴. It has not paid the resulting Ofcom fine, and has appealed to the Trump administration to act against "extraterritorial censorship mandates"⁷⁵. Ofcom later issued a provisional notice of contravention for 4chan's failure to respond to its information requests⁷⁶.

4chan is not the only offender. Adult content company AVS Group Ltd received a £1 million fine from Ofcom for non-compliance with the OSA⁷⁷.

In summary, the implementation of the UK Online Safety Act has produced compliance amongst a significant majority of the most popular adult content websites. How effective that compliance is at keeping people away from harmful or illegal content remains an open question, with more study required. What the UK example demonstrates is that a regime can be designed and brought into force and that a regulator, in this case Ofcom, can enforce a framework even as it evolves.

That said, these examples all involve free services that earn revenue by advertising to their users. And they cover only pornography sites, not the social media services through which most children first encounter pornography. Those services will also need to be brought into compliance if youth access to pornography is to fall significantly.

One platform, however, shows some of the characteristics of a social media service, and was classed as one for the purposes of Australia's ban⁷⁸, while also carrying significant adult content.

It is also the service the UK Prime Minister, Keir Starmer, recently held up as a standard for social media companies on age verification⁷⁹: OnlyFans.

5.3 Case Study: OnlyFans and Age Verification

Disclosure: The author's firm, Trafalgar Strategy, has advised OnlyFans' corporate leadership since January 2023.

OnlyFans is a distinct case within the adult content sector. Unlike most sites hosting adult content, it runs on a subscription model and carries no third-party advertising, and, unlike algorithmic social media platforms, does not use a content-discovery algorithm to surface posts to users⁸⁰. It also hosts a range of non-adult creators, including comedians, athletes, and coaches.

Because its subscription model does not depend on maximising a user's time on site, OnlyFans lacks the structural incentive to facilitate permissive access that advertising-based social media platforms have. The company has also chosen to operate as a strictly adult platform, barring both creators and users under 18, whatever content they intend to produce or consume⁸¹. That choice, together with the sensitive nature of the content, makes an effective age verification regime essential.

5.3.1. Verification Requirements for Users and Creators

OnlyFans' 'fans' must provide payment details and other documentation to view content on the platform. The company applies distinct verification processes to each side of its platform, creator and 'fan':

Requirements for Users ('fans')

In the UK, all users must provide payment information and complete an age assurance process through the external provider Yoti before an account is activated⁸². Yoti's facial age estimation tool is set to a threshold of 21; users assessed below it are offered the option to submit photo ID, matched by a live selfie, to confirm they are 18 or older.

Requirements for Creators

In all markets, OnlyFans content creators must provide a minimum of nine pieces of identifying information before posting content. These include:

- Full name
- Verified email address
- Postal address
- Date of birth
- Valid government-issued photo ID
- Standalone selfie
- Selfie held alongside photo ID (validated by Yoti)
- Social media account handles
- Bank account information
- In the United States: W-9 form and Social Security Number

Creator identity checks are renewed every 30 days to confirm continuity between the account holder at sign-up and the active creator.⁸³

5.3.2. OnlyFans' Privacy and Technical Standards

Despite leaving the European Union, the UK maintains its own General Data Protection Regulation (GDPR) regime, modelled on the EU original. The EU has granted the UK an 'adequacy decision', confirming that its data protection rules are close enough to the EU's to allow information to flow freely between the two⁸⁴.

OnlyFans is subject to GDPR and uses Yoti for its age estimation checks. It chose Yoti to “provide the right balance between accurately assessing users’ ages and respecting their privacy”⁸⁵.

Yoti’s facial age estimation system is built to the GDPR principle of ‘privacy by design’ and reports a ‘true positive rate’ of 99.3 per cent for correctly estimating those aged 13 to 17 as being under 21⁸⁶, with no discernable bias across gender or skin tones. The system does not process ‘special category data’ as defined under GDPR (including biometric data used for unique identification)⁸⁷, as the model categorises faces by age group rather than identifying specific individuals. Images captured during the estimation process are immediately and permanently deleted and are not shared, stored, or used for model training. Yoti’s approach is externally assessed under the SOC2⁸⁸ and PAS1296⁸⁹ standards⁹⁰.

OnlyFans operates as the data controller in this arrangement, with Yoti acting as the data processor⁹¹. Under UK GDPR, it is the controllers who make decisions about processing activities; they exercise overall control of the personal data being processed and are ultimately in charge of and responsible for the processing. A processor may make its own day-to-day operational decisions, but guidelines dictate that it should only process personal data in line with a controller’s instructions, unless it is required to do otherwise by law. This separation means that clients using external services like Yoti to process data must have an independent legal basis for processing under their own jurisdiction.

OnlyFans’ regime has not been without incident. The platform received a fine from Ofcom in 2025 for an age verification breach. The challenge age for Yoti’s estimation had been set at 20 rather than the declared 23, which Ofcom found could have let some users under 18 gain access, given the margins of error inherent in the technology⁹². The configuration was later set to 21, where it remains.

In summary, payment information combined with age estimation allows OnlyFans to deliver what the regulator considers a ‘highly effective’ age verification regime, one backed by enforceable penalties set out in primary legislation.

6. Policy Considerations for Canada

The evidence reviewed in this note suggests several considerations relevant to Canadian policymakers as they develop online safety legislation.

6.1 Age Verification Is Technically Feasible at Scale

The adult content sector in the United Kingdom shows that age verification can be implemented at scale, and that a range of methods can meet the criteria for ‘highly effective’ age assurance while still protecting

user privacy. The OnlyFans case shows that a fairly comprehensive verification regime can be maintained without the government acting as the verifier or the holder of user data. A model in which the government sets and enforces standards, and private companies manage verification through accredited third parties, is technically viable. Feasibility in this operational sense, however, is not the same as a settled finding that the regime achieves its goal. The UK’s enforcement record shows that compliance can be secured and that traffic to non-compliant sites can fall sharply, which is real evidence that the model can work. What it does not yet show, because the regime is new, is how durably that reduction holds once circumvention adapts, or whether harm to children falls as a result. The fair conclusion is that large-scale age verification is administratively feasible, and that its longer-term effectiveness is still being established.

The May 2026 guidance from the Office of the Privacy Commissioner of Canada acknowledged as much, recognising that age verification, age estimation, or age inference could each be implemented in a way that is highly effective⁹³.

6.2 Age Verification Requires a Broader Regulatory Architecture

The UK’s regime works because primary legislation, the Online Safety Act, is in place, with supporting regulations behind it. The Act provides clear direction and gives the regulator, Ofcom, the tools to enforce it. Other bodies, most notably the Information Commissioner’s Office, set guidance and standards for third parties like Yoti that deliver private-sector age verification services. All of this rests on the UK GDPR privacy regime. Together, this architecture helps build user trust in the verification process. For Canada, that architecture should also include one element the UK model does not yet make explicit. Because age verification engages the rights of children directly, as set out in section 4.5, any Canadian framework should require a Child Rights Impact Assessment of the proposed regime before it is enacted, and should build the data-minimisation and non-exclusion safeguards of General Comment No. 25 into the standards that regulators enforce. This treats children’s rights as a design input rather than an afterthought.

6.3 An Existing Customer Relationship Aids Verification

A major obstacle to age verification on social media and adult content sites is that most of their users have no direct or pre-existing financial relationship with the company providing the service. Paying for a service creates a bond between customer and business, and an expectation of a certain standard. It also gives the provider one means of age-verifying its customer base.

The government is unlikely to mandate prices for these sites. But age verification methods such as open banking, mobile network operator checks, and email-based estimation can draw on the trust users have already built with their banks, mobile operators, and email providers.

6.4 Enforcement Depends on Platform Cooperation

The Australian experience shows how much enforcement depends on active cooperation from platforms. Major social media companies have shown limited enthusiasm for age verification, in part because their advertising-based revenue models benefit from permissive access. Any Canadian framework will need robust enforcement mechanisms, clear penalties for non-compliance, and, following the UK's 2026 approach, a requirement that named senior individuals within platforms be accountable for compliance.

6.5 Verification Methods Must Cover Users Under Contracting Age

Where the relevant policy threshold is set below 18, as it is with proposed social media bans for under-16s, methods that rely on contracting age, such as credit cards and mobile phone contracts, will be insufficient as primary tools. Facial age estimation is currently the most viable technology for verifying users below the age of contracting, though it depends on the architecture set out in section 6.2 and on a sustained commitment to address evolving circumvention techniques.

6.6 VPN Displacement Is a Real but Partial Concern

The available evidence suggests that adult users, not children, are the primary adopters of VPNs in response to age verification on dedicated pornography sites. That does not eliminate the concern, but it indicates that the core protective aim of verification, limiting children's access, may still be achievable.

Whether the same holds for social media is less clear. A social media age limit applies to services that young people use constantly and value highly, so the incentive for under-age users to seek workarounds, including VPNs, may be stronger than it is for pornography sites. Whether VPN use rises in Australia after its ban, as it has in other jurisdictions, remains an open empirical question. Some jurisdictions have begun to consider VPN-side responses: Utah's legislation, also noted above, targets the use of VPNs to evade age checks, and more far-reaching ideas, such as restricting the availability of VPNs at the operating-system or app-store level, are sometimes raised. But VPNs are mainstream privacy and security tools, and constraining them is a serious step that should be weighed openly, as a policy question in its own right, rather than adopted as an incidental technical fix.

6.7 Age Verification Is a Necessary but Insufficient Response

Age verification and access restrictions address one dimension of online harms to children, but they do not substitute for broader reforms. The advertising-based business models that underpin the major social media platforms create structural incentives for harmful design choices, and age gating alone does not alter those incentives. A comprehensive approach to children's online safety would also address content moderation, algorithmic accountability, and platform design, not only access controls.

Subscription-based platforms, such as OnlyFans, run on a different model, one in which the revenue relationship is directly with the user rather than with advertisers. That alignment of incentives is one consideration among others as policymakers design requirements for platforms that serve younger audiences. What the evidence in this note makes clear is that the central questions now are ones of regulatory design and enforcement, not of technical possibility.

7. Conclusion

The direction of policy travel across comparable jurisdictions is clear. Some form of age assurance or verification for online platforms is likely to be part of any serious legislative response to online harms affecting children. The experience of the adult content sector, and the implementation of the UK's Online Safety Act in particular, gives Canadian policymakers a useful evidence base as they design their own framework.

Key lessons include the technical feasibility of large-scale age verification; the importance of robust enforcement and platform accountability; the limitations of methods that depend on contracting age when the relevant threshold is below 18; and the need for ongoing adaptation as technology and user behaviour evolve.

Age verification is one element of a broader policy framework, not a standalone solution. Protecting children online will ultimately depend as much on the structural features of platform design and business models as on the technical mechanisms used to control access. How verification is designed, and what it is paired with, will matter as much as whether it is adopted.

Contributor

Andrew MacDougall

Senior Policy Fellow at the Centre for Media, Technology and Democracy, Partner at Trafalgar Strategy, and former Director of Communications to Prime Minister Stephen Harper.

Endnotes

1 eSafety Commissioner (Australia), "Social Media Age Restrictions," eSafety Commissioner, n.d., <https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions> (accessed 29 May 2026).

2 UK Department for Science, Innovation and Technology, "Principles of Safer Online Platform Design," GOV.UK, 29 June 2021 (updated 14 May 2026), <https://www.gov.uk/guidance/principles-of-safer-online-platform-design> (accessed 29 May 2026).

3 CBC News, "Ottawa 'very seriously' considering age restrictions for social media, AI chatbots," CBC News, 2026, <https://www.cbc.ca/news/politics/social-media-ai-minimum-age-9.7164902> (accessed 29 May 2026).

4 Sean Previl, "Manitoba could fine social media companies billions if youth ban not enforced: Kinew," Global News, 28 April 2026, <https://globalnews.ca/news/11820904/manitoba-social-media-ban-tech-giant-fines/> (accessed 29 May 2026).

5 Bobby Allyn, "Jury finds Meta and Google negligent in social media harms trial," NPR, 25 March 2026, <https://www.npr.org/2026/03/25/nx-s1-5746125/meta-youtube-social-media-trial-verdict> (accessed 29 May 2026).

6 BBC News, "Meta told to pay \$375m for misleading users over child safety," BBC News, 2026, <https://www.bbc.co.uk/news/articles/cql75dn07n2o> (accessed 29 May 2026).

7 Competition and Markets Authority, Online Platforms and Digital Advertising: Market Study Final Report, CMA, 1 July 2020, https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final_report_1_July_2020_.pdf (accessed 29 May 2026).

8 Jay N. Giedd, "Risky Teen Behavior Is Driven by an Imbalance in Brain Development," *Scientific American* 312, no. 6, 1 June 2015, <https://www.scientificamerican.com/article/risky-teen-behavior-is-driven-by-an-imbalance-in-brain-development/> (accessed 29 May 2026).

9 Office of the U.S. Surgeon General, Social Media and Youth Mental Health: The U.S. Surgeon General's Advisory, U.S. Department of Health and Human Services, 2023, <https://www.hhs.gov/sites/default/files/sg-youth-mental-health-social-media-advisory.pdf> (accessed 29 May 2026).

10 Teun Siebers, Ine Beyens, J. Loes Pouwels and Patti M. Valkenburg, "Social Media and Distraction: An Experience Sampling Study among Adolescents," *Media Psychology* 25, no. 3 (2022): 343–366, <https://doi.org/10.1080/15213269.2021.1959350> (accessed 15 June 2026).

11 National Academies of Sciences, Engineering, and Medicine, "Online

Harassment," in *Social Media and Adolescent Health* (Washington, DC: National Academies Press, 25 March 2024), <https://www.ncbi.nlm.nih.gov/books/NBK603433/> (accessed 29 May 2026).

12 Lauren E. Sherman et al., "The Power of the Like in Adolescence: Effects of Peer Influence on Neural and Behavioral Responses to Social Media," *Psychological Science* 27, no. 7 (2016), <https://pmc.ncbi.nlm.nih.gov/articles/PMC5387999/> (accessed 29 May 2026).

13 Patti M. Valkenburg, Adrian Meier and Ine Beyens, "Social Media Use and Its Impact on Adolescent Mental Health: An Umbrella Review of the Evidence," *Current Opinion in Psychology* 44 (2022): 58–68, <https://doi.org/10.1016/j.copsyc.2021.08.017> (accessed 29 May 2026).

14 Tim Wu, *The Attention Merchants: The Epic Scramble to Get Inside Our Heads* (New York: Alfred A. Knopf, 2016).

15 Ofcom, "Proving Age to Access Online Pornography: Survey Technical Report (2025, Wave 1)," Ofcom, 2025, <https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/online-research/keeping-children-safe-online/proving-age-to-access-online-pornography/survey-proving-age-to-access-online-pornography-survey-technical-report-2025-wave-1.pdf> (accessed 29 May 2026).

16 Ofcom, "Age Checks for Online Safety: What You Need to Know as a User," Ofcom, 2025, <https://www.ofcom.org.uk/online-safety/protecting-children/age-checks-for-online-safety--what-you-need-to-know-as-a-user> (accessed 29 May 2026).

17 UK Department for Science, Innovation and Technology, "Online Safety Act: Explainer," GOV.UK, updated 24 April 2025, <https://www.gov.uk/government/publications/online-safety-act-explainer/online-safety-act-explainer> (accessed 29 May 2026).

18 Bill S-210, An Act to restrict young persons' online access to sexually explicit material, 44th Parliament, 1st Session, LEGISinfo, Parliament of Canada, <https://www.parl.ca/LegisInfo/en/bill/44-1/s-210> (accessed 29 May 2026).

19 Bill S-209, An Act to restrict young persons' online access to pornographic material, 45th Parliament, 1st Session, LEGISinfo, Parliament of Canada, <https://www.parl.ca/LegisInfo/en/bill/45-1/s-209> (accessed 29 May 2026).

20 Information Commissioner's Office, "Age Assurance" (Age Appropriate Design toolkit), ICO, n.d., <https://ico.org.uk/for-organisations/advice-and-services/audits/data-protection-audit-framework/toolkits/age-appropriate-design/age-assurance/> (accessed 29 May 2026).

21 Kayee Hanaoka et al., *Face Analysis Technology Evaluation: Age Estimation and Verification*, NIST IR 8525, National Institute of Standards and Technology, May 2024 (updated 29 April 2026), https://pages.nist.gov/frvt/reports/aev/fate_

aev_report.pdf (accessed 29 May 2026).

22 Meta, “Introducing New Ways to Verify Age on Instagram,” Meta Newsroom, 23 June 2022 (updated 2 March 2023), <https://about.fb.com/news/2022/06/new-ways-to-verify-age-on-instagram/> (accessed 29 May 2026). The 81 per cent figure is Yoti’s own: Julie Dawson, “AI Adoption Case Study: Learn about Yoti’s Facial Age Estimation Tool,” techUK, 14 October 2024, <https://www.techuk.org/resource/ai-adoption-case-study-learn-about-yoti-s-facial-age-estimation-tool.html> (accessed 29 May 2026).

23 Department for Science, Innovation and Technology, “Digital ID Scheme: Explainer,” GOV.UK, updated 18 March 2026, <https://www.gov.uk/government/publications/digital-id-scheme-explainer/digital-id-scheme-explainer> (accessed 29 May 2026).

24 Zarema Plaksij, “Credit Card Age Verification: Limitations, Risks, and Alternatives,” Ondato, 10 September 2024 (updated 24 February 2026), <https://ondato.com/blog/credit-card-age-verification/> (accessed 29 May 2026).

25 Lina Ghazal, “The Online Safety Act,” Verifymy, updated 14 January 2026, <https://verifymy.io/resources/regulation/uk/the-online-safety-act/> (accessed 29 May 2026).

26 All Pass Trust, “Summary of Age Assurance Written Record” (UK Online Safety Act), All Pass Trust, updated 5 August 2025, <https://www.allpasstrust.com/UKOnlineSafetyAct> (accessed 29 May 2026).

27 OneID, “Mobile Number Age Check: User Guide,” OneID, n.d., <https://oneid.uk/knowledge/user-guide/mno-age-check> (accessed 29 May 2026).

28 Financial Conduct Authority, “Open Banking and the FCA,” FCA, 16 December 2025 (updated 31 December 2025), <https://www.fca.org.uk/firms/open-banking-fca> (accessed 29 May 2026).

29 Ofcom, “Age Checks for Online Safety: What You Need to Know as a User,” Ofcom, n.d., <https://www.ofcom.org.uk/online-safety/protecting-children/age-checks-for-online-safety--what-you-need-to-know-as-a-user> (accessed 29 May 2026).

30 Verifymy, “Innovative Age Assurance: Email Address as the New Benchmark for Frictionless Age Estimation” (consultation response/white paper), hosted by Ofcom, 2024, <https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/284469-consultation-protecting-children-from-harms-online/responses/verifymy---white-paper.pdf> (accessed 29 May 2026).

31 Ofcom, Guidance on Highly Effective Age Assurance (Part 3), in “Statement: Age Assurance and Children’s Access,” Ofcom, 16 January 2025 (updated 24 April 2025), <https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/statement-age-assurance-and-childrens-access/part-3->

[guidance-on-highly-effective-age-assurance.pdf](https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/statement-age-assurance-and-childrens-access/part-3-guidance-on-highly-effective-age-assurance.pdf) (accessed 29 May 2026).

32 e-Estonia (Estonian Business and Innovation Agency), “ID-Card / e-Identity,” e-Estonia, 10 June 2024 (updated 13 February 2025), <https://e-estonia.com/solutions/estonian-e-identity/id-card/> (accessed 29 May 2026).

33 “Keir Starmer Denies Roll-back of Digital IDs Is Yet Another U-turn,” The Guardian, 14 January 2026, <https://www.theguardian.com/politics/2026/jan/14/keir-starmer-denies-roll-back-of-digital-ids-is-yet-another-u-turn> (accessed 29 May 2026).

34 “How Accurate Are Age Checks for Australia’s Under-16s Social Media Ban? What Trial Data Reveals,” The Guardian, 19 September 2025, <https://www.theguardian.com/news/2025/sep/19/how-accurate-are-age-checks-for-australias-under-16s-social-media-ban-what-trial-data-reveals> (accessed 29 May 2026).

35 Yoti, “How Accurate Is Facial Age Estimation?,” Yoti, n.d., <https://www.yoti.com/blog/how-accurate-is-facial-age-estimation/> (accessed 29 May 2026).

36 Tzvi Ganel et al., “Biases in Human Perception of Facial Age Are Present and More Exaggerated in Current AI Technology,” Scientific Reports 12 (2022): article 22519, <https://doi.org/10.1038/s41598-022-27009-w> (accessed 29 May 2026). On independent accuracy benchmarks see Kayee Hanaoka et al., Face Analysis Technology Evaluation: Age Estimation and Verification, NIST IR 8525, National Institute of Standards and Technology, May 2024, <https://doi.org/10.6028/NIST.IR.8525> (accessed 29 May 2026).

37 Julie Dawson, “AI Adoption Case Study: Learn about Yoti’s Facial Age Estimation Tool,” techUK, 14 October 2024, <https://www.techuk.org/resource/ai-adoption-case-study-learn-about-yoti-s-facial-age-estimation-tool.html> (accessed 29 May 2026).

38 European Data Protection Board, Statement 1/2025 on Age Assurance, EDPB, adopted 11 February 2025, https://www.edpb.europa.eu/our-work-tools/our-documents/statements/statement-12025-age-assurance_en (accessed 29 May 2026).

39 Zarema Plaksij, “Age Verification: How It Works and How Businesses Approach It Today,” Ondato, 7 August 2025 (updated 23 February 2026), <https://ondato.com/blog/what-is-age-verification/> (accessed 29 May 2026).

40 “Meta Found in Breach of EU Law for Failing to Keep Children off Platforms,” The Guardian, 29 April 2026, <https://www.theguardian.com/technology/2026/apr/29/meta-found-in-breach-of-eu-law-for-failing-to-keep-children-off-platforms> (accessed 29 May 2026).

41 Parliament of Australia, Online Safety Amendment (Social Media Minimum Age) Bill 2024 (Act No. 127 of 2024, assented 10 December 2024), https://www.aph.gov.au/Parliamentary_Business/Bills_Legislation/Bills_Search_

Results/Result?bId=r7284 (accessed 29 May 2026).

42 eSafety Commissioner, “Social Media Minimum Age: Regulatory Guidance,” eSafety Commissioner (Australia), September 2025, <https://www.esafety.gov.au/industry/regulatory-guidance#social-media-minimum-age> (accessed 29 May 2026).

43 eSafety Commissioner, Social Media Minimum Age: Compliance Update (March 2026), eSafety Commissioner (Australia), March 2026, <https://www.esafety.gov.au/sites/default/files/2026-03/SocialMediaMinimumAgeComplianceUpdateMarch2026.pdf> (accessed 29 May 2026).

44 eSafety Commissioner, Social Media Minimum Age: Compliance Update (March 2026), eSafety Commissioner (Australia), March 2026, <https://www.esafety.gov.au/sites/default/files/2026-03/SocialMediaMinimumAgeComplianceUpdateMarch2026.pdf> (accessed 29 May 2026).

45 Anika Wells (Minister for Communications), “Social Media Minimum Age Update Targets Harmful Tools,” Minister for Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, 2026, <https://minister.infrastructure.gov.au/wells/media-release/social-media-minimum-age-update-targets-harmful-tools> (accessed 29 May 2026).

46 “Every scary thing Meta knows about me — and you,” The Times, n.d., <https://www.thetimes.com/uk/healthcare/article/meta-user-data-collection-facebook-instagram-whatsapp-9klqhzlq9> (accessed 29 May 2026).

47 George S. Ford, “Unintended Consequences of Age Verification Laws: Evidence from Search Behavior,” Phoenix Center Perspectives No. 25-06, 12 November 2025, <https://www.phoenix-center.org/perspectives/Perspective25-06Final.pdf> (accessed 29 May 2026).

48 “After Pornhub Left Florida, VPN Demand Surged by More than 1,000%,” Business Insider, January 2025, <https://www.businessinsider.com/pornhub-exited-florida-vpn-demand-surged-by-over-1000-percent-2025-1> (accessed 29 May 2026).

49 Ofcom, Online Safety in 2025: A Summary of the Technology Sector’s Response to Our Rules, Ofcom, 2025, <https://www.ofcom.org.uk/siteassets/resources/documents/online-safety/research-statistics-and-data/os-standards/online-safety-in-2025-summary-of-the-technology-sectors-response-to-our-rules.pdf> (accessed 29 May 2026).

50 Utah State Legislature, S.B. 73, Online Age Verification Amendments, 2026 General Session (introduced version), <https://le.utah.gov/Session/2026/bills/introduced/SB0073.pdf> (accessed 29 May 2026).

51 Rindala Alajaji, “Utah’s New Law Targeting VPNs Goes into Effect May 6th,”

Electronic Frontier Foundation, 30 April 2026, <https://www EFF.org/deeplinks/2026/04/utahs-new-law-regulating-vpns-goes-effect-next-week> (accessed 29 May 2026).

52 Katie Freeman-Taylor, “New Data Shows No Rise in Children’s VPN Use after the Introduction of Online Age Checks,” Internet Matters, 4 December 2025, <https://www.internetmatters.org/hub/research/data-shows-no-rise-childrens-vpn-use-amid-online-age-checks/> (accessed 29 May 2026).

53 Simeon de Brouwer, “Why Age Verification Misses the Mark and Puts Everyone at Risk,” European Digital Rights (EDRI), 25 November 2025, <https://edri.org/our-work/why-age-verification-misses-the-mark-and-puts-everyone-at-risk/> (accessed 29 May 2026).

54 “Australia’s Disproportionate Response to Online Harms,” Index on Censorship, 10 December 2025, <https://www.indexoncensorship.org/2025/12/australias-disproportionate-response-to-online-harms/> (accessed 29 May 2026).

55 BBC News, “Teens Launch High Court Challenge to Australia’s Social Media Ban,” BBC News, 2025, <https://www.bbc.co.uk/news/articles/c9d9l7w1g8yo> (accessed 29 May 2026). In the United States, see *NetChoice, LLC v. Yost*, No. 2:24-cv-00047 (S.D. Ohio, 16 April 2025), holding that the First Amendment protects minors’ rights to speak and access information online and striking down Ohio’s Social Media Parental Notification Act, <https://netchoice.org/netchoice-v-yost/> (accessed 15 June 2026).

56 *Free Speech Coalition, Inc. v. Paxton*, 606 U.S. 461 (2025) (No. 23-1122), United States Supreme Court, decided 27 June 2025, <https://supreme.justia.com/cases/federal/us/606/23-1122/> (accessed 29 May 2026).

57 UN Committee on the Rights of the Child, General Comment No. 25 (2021) on children’s rights in relation to the digital environment, CRC/C/GC/25, OHCHR, 2 March 2021, <https://www.ohchr.org/en/documents/general-comments-and-recommendations/general-comment-no-25-2021-childrens-rights-relation> (accessed 29 May 2026).

58 Sonia Livingstone and Kruakae Pothong, “Child Rights Impact Assessment: A Policy Tool for a Rights-Respecting Digital Environment,” *Policy & Internet* 17, no. 3 (2025): e70008, <https://doi.org/10.1002/poi3.70008> (accessed 29 May 2026).

59 Giselle Woodley et al., “Australian Teen Voices on Age Verification and Age Assurance Measures,” *Policy & Internet* (2025): e70019, <https://doi.org/10.1002/poi3.70019> (accessed 29 May 2026).

60 Online Safety Act 2023, c. 50, [legislation.gov.uk](https://www.legislation.gov.uk/ukpga/2023/50), <https://www.legislation.gov.uk/ukpga/2023/50> (accessed 29 May 2026).

61 Ofcom, Online Safety in 2025: A Summary of the Technology Sector’s Response to Our Rules, Ofcom, 2025, <https://www.ofcom.org.uk/siteassets/resources/documents/online-safety/research-statistics-and-data/os-standards/>

online-safety-in-2025-summary-of-the-technology-sectors-response-to-our-rules.pdf (accessed 29 May 2026).

62 Anna Iovine, "All the States Pornhub Is Blocked In," Mashable, updated April 2026, <https://mashable.com/article/pornhub-blocked-states-2025> (accessed 29 May 2026).

63 "UK Starts Porn Age Verification: How Do the Internet Age Checks Work and Will Users' Data Be Safe?," Euronews, 25 July 2025, <https://www.euronews.com/next/2025/07/25/uk-starts-porn-age-verification-how-do-the-internet-age-checks-work-and-will-users-data-be> (accessed 15 June 2026).

64 Ofcom, Online Nation 2025, Ofcom, 2025, <https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/online-research/online-nation/2025/online-nations-report-2025.pdf> (accessed 29 May 2026).

65 "Pornhub to Stop New UK Users Accessing Site from Next Week," The Guardian, 27 January 2026, <https://www.theguardian.com/technology/2026/jan/27/pornhub-to-stop-new-uk-users-accessing-site-from-next-week> (accessed 29 May 2026).

66 Tony Allen / Age Check Certification Scheme, "Statement: Aylo Withdrawal from UK Market," ACCS, 27 January 2026, <https://accscheme.com/statement-aylo-withdrawal-from-uk-market/> (accessed 29 May 2026).

67 BBC News, "Pornhub to Restrict Access for UK Users from Next Week," BBC News, 2026, <https://www.bbc.co.uk/news/articles/czr428rxg57o> (accessed 29 May 2026).

68 BBC News, "Pornhub to become accessible again for some UK users," BBC News, May 2026, <https://www.bbc.co.uk/news/articles/cwy27q05gj2o> (accessed 29 May 2026).

69 Semrush, "Most Visited Adult Websites in the United Kingdom" (Traffic Analytics), <https://www.semrush.com/trending-websites/gb/adult> (accessed 15 June 2026); Bill Bowkett, "Pornography Websites Flouting New Age Verification Rules — Despite Threat of Huge Fines," The Evening Standard (via Yahoo News UK), 25 July 2025, <https://uk.news.yahoo.com/pornography-websites-flout-age-verification-173412481.html> (accessed 29 May 2026).

70 Danielle Desouza, "Dip in Traffic to Popular Porn Sites amid UK Age Verification Rules Being Enforced," LBC, 13 August 2025, https://www.lbc.co.uk/article/traffic-porn-sites-dip-uk-age-verification-rules-5Hjd8sn_2/ (accessed 29 May 2026).

71 XVIDEOS, "How to Pass Age Verification on Xvideos" (FAQ), info.xvideos.net, n.d., https://info.xvideos.net/faq/question/139-everyone-how_to_pass_age_verification_on_xvideos (accessed 29 May 2026).

72 BBC News, "UK Traffic to Adult Websites Drops Sharply as Age Verification

Measures Take Effect," BBC News, August 2025, <https://www.bbc.co.uk/news/articles/c17n9k54qz2o> (accessed 29 May 2026).

73 Age Verification Providers Association, "No, UK Porn Use Was Not Halved by Age Verification," AVPA, 8 September 2025, <https://avpassociation.com/thought-leadership/no-uk-porn-use-was-not-halved-by-age-verification/> (accessed 29 May 2026).

74 BBC News, "4chan will refuse to pay daily online safety fines, lawyer tells BBC," BBC News, 2026, <https://www.bbc.co.uk/news/articles/cq68j5g2nr1o> (accessed 29 May 2026).

75 BBC News, "4chan will refuse to pay daily online safety fines, lawyer tells BBC," BBC News, 2026, <https://www.bbc.co.uk/news/articles/cq68j5g2nr1o> (accessed 29 May 2026).

76 Ofcom, Confirmation Decision: Investigation into 4chan Community Support LLC, Ofcom, 2025, <https://www.ofcom.org.uk/siteassets/resources/documents/enforcement/2026/4chan-confirmation-decision.pdf> (accessed 29 May 2026).

77 Ofcom, "Ofcom Fines Porn Company £1 Million for Not Having Robust Age Checks," Ofcom, December 2025, <https://www.ofcom.org.uk/online-safety/protecting-children/ofcom-fines-porn-company-1million-for-not-having-robust-age-checks> (accessed 29 May 2026).

78 eSafety Commissioner, "OnlyFans," The eSafety Guide, eSafety Commissioner (Australia), n.d., <https://www.esafety.gov.au/key-topics/esafety-guide/onlyfans> (accessed 29 May 2026).

79 Financial Times, "OnlyFans-style age checks should be social media standard, says Starmer," Financial Times, April 2026, <https://www.ft.com/content/482a3067-602e-4111-931c-eea8b1edc8dc> (accessed 29 May 2026).

80 OnlyFans, "Terms of Service," OnlyFans, n.d., <https://onlyfans.com/terms> (accessed 15 June 2026), on the platform's subscription model; and, on the absence of an internal content-discovery algorithm (which obliges creators to draw their own audiences from other services), Margherita Di Cicco and Davide Beraldo, "Sex Work as Cross-Platform Self-Branding: Challenges and Strategies of OnlyFans' Content Creators in a Precarious Ecosystem," *Platforms & Society* (2025), <https://doi.org/10.1177/29768624251359797> (accessed 15 June 2026).

81 OnlyFans, "Terms of Service," OnlyFans, n.d., <https://onlyfans.com/terms> (accessed 29 May 2026).

82 OnlyFans, "Verification," Transparency Center, OnlyFans, n.d., <https://onlyfans.com/transparency-center/verification> (accessed 29 May 2026).

83 OnlyFans, "Verification," Transparency Center, OnlyFans, n.d., <https://onlyfans.com/transparency-center/verification> (accessed 29 May 2026).

84 European Commission, “Commission Renews Decisions to Allow for the Free and Safe Flow of Personal Data with the UK” (IP/25/3059), European Commission Press Corner, 19 December 2025, https://ec.europa.eu/commission/presscorner/detail/en/ip_25_3059 (accessed 29 May 2026).

85 Yoti, “How OnlyFans Became the First UK Subscription-Based Platform to Protect Children and Create Age-Appropriate Experiences” (OnlyFans case study), Yoti, 12 June 2023, https://cdn.aws.yoti.com/wp-content/uploads/Onlyfans_Case-Study_120623.pdf (accessed 29 May 2026).

86 Yoti, Yoti Facial Age Estimation White Paper (full version, July 2025), Yoti, July 2025, <https://cdn.aws.yoti.com/wp-content/uploads/2026/01/Yoti-Age-Estimation-White-Paper-July-2025-PUBLIC-v1.pdf> (accessed 29 May 2026).

87 Information Commissioner’s Office, “Special Category Data” (UK GDPR guidance), ICO, n.d., <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/a-guide-to-lawful-basis/special-category-data/> (accessed 29 May 2026).

88 Secureframe, “What Is SOC 2? A Beginner’s Guide to Compliance,” Secureframe, n.d., <https://secureframe.com/hub/soc-2/what-is-soc-2> (accessed 29 May 2026).

89 Age Check Certification Scheme, “The Evolution of Age Assurance Standards: PAS 1296 to IEEE 2089.1 to ISO/IEC 27566-1,” ACCS, updated 13 October 2025, <https://accscheme.com/iso-comparison-guide/> (accessed 29 May 2026).

90 Yoti, “Our Approach to Security and Privacy,” Yoti, n.d., <https://www.yoti.com/blog/our-approach-to-security-and-privacy/> (accessed 29 May 2026).

91 Information Commissioner’s Office, “What Are ‘Controllers’ and ‘Processors’?” (UK GDPR guidance), ICO, n.d., <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/controllers-and-processors/controllers-and-processors/what-are-controllers-and-processors/> (accessed 29 May 2026).

92 Ofcom, “Ofcom Fines Provider of OnlyFans £1.05 Million,” Ofcom, 27 March 2025, <https://www.ofcom.org.uk/online-safety/protecting-children/ofcom-fines-provider-of-onlyfans-1.05-million> (accessed 29 May 2026).

93 Office of the Privacy Commissioner of Canada, “Assessing Whether and How to Use Age Assurance: Guidance for Websites and Online Services,” OPC, 4 May 2026, <https://www.priv.gc.ca/en/privacy-topics/age-assurance/aa-gd-web/> (accessed 29 May 2026).