AVPA Response to the UK Government's Growing up in the online world: a national consultation

CHAPTER 2: INTERVENTIONS FOR SAFER, MORE POSITIVE EXPERIENCES

Q15. What do you think the impacts would be of having minimum age requirements higher than 13 for social media services?

If the current minimum age of 13 was effectively implemented and enforced, there would be no substantive marginal cost of raising that minimum age. Age assurance — comprising age verification, age estimation and age inference, according to the ISO/IEC 27566-1 definitions — can be used to apply age restrictions at any of the proposed ages.

In fact, as the minimum age rises above 13, it becomes easier and therefore cheaper to implement, as users will have more potential sources of age data on which age verification methods rely. For example, many more 16 year olds than 13 year olds have a bank account, which is a reliable source of age confirmation. A higher threshold is not inherently more technically demanding. For many estimation systems, accuracy also improves materially as operational thresholds move further above early adolescence.

We are aware that most social media platforms do not currently enforce their minimum age policy, in breach of their long-standing legal obligations under GDPR, and we welcome recent enforcement action by the ICO to address this anomaly, and their joint statement with Ofcom that Highly Effective Age Assurance should be used for this purpose and that this is feasible given the state of the art of current technology. There will be a modest cost to platforms to comply with existing and any new minimum age policies.

While we do not discuss pricing with our members, we are aware of publicly quoted prices. The OpenAge Initiative recently launched a service that facilitates reusable and interoperable age checks for a fee of one third of a US cent per check, once an initial check has been purchased by any participating relying party. So for example, Meta might pay an age assurance provider 11 cents to check a new user, and after that Snapchat, Discord or TikTok could re-use that same check for just one-third of a cent — effectively creating an average cost of 4 cents per user across four platforms (illustrative figures only). Individual age assurance providers discount heavily for high volume clients, and it is today reasonable to assume that even low volume platforms are paying pennies not pounds per check. The Impact Assessment for the Online Safety Act suggested 10 pence per check as a baseline assumption.

Social norming

The existence of clear and consistently enforced age thresholds contributes to social norming. The present environment, in which children routinely access services nominally restricted to older users through simple self-declaration, undermines respect for online safety rules generally. Where children encounter no genuine barrier, the implicit message is that these rules are not serious. Consistently enforced age assurance changes that norm.

Innovation

There are existing clients of our sector who have worked hard to create age-appropriate online experiences for children, informed by comprehensive age assurance, that could fall into some definitions of social media. These sites often not only check that children are not exposed to content or functionality when too young, but also use age assurance to mitigate the risk of bad actors who are adults pretending to be children. Care should be taken to avoid the unintended consequence of restricting access to such sites which have been designed to be safe for children. In theory, a social media platform can be designed to be safe for a child of any age, if design decisions are always taken to deliver that outcome rather than to optimise for engagement, profit or other purposes.

Privacy

Age assurance for people of all ages, including children, can be implemented entirely anonymously, without putting privacy or personal data at risk. Age assurance providers certified under frameworks such as the AVPA's certification scheme are required to demonstrate this, providing independent assurance that privacy protections are genuine rather than merely claimed.

Public debate frequently conflates user friction with privacy intrusion. In practice, many independently certified age assurance systems process substantially less personal data than the advertising and tracking ecosystems already deployed by major platforms as a matter of routine. The choice is not between privacy and age assurance; it is between age assurance conducted to certified privacy standards and the status quo in which children are tracked extensively without meaningful age verification.

Accessibility and inclusivity

Any age assurance requirement must include a duty on platforms to offer multiple methods, so that no child or adult is excluded by virtue of lacking a particular form of ID or biometric. Where lower-cost automated methods are unavailable or unsuitable for a particular user, platforms must be required to fund more expensive backstop methods such as professional attestation services (for example through CitizenCard or PASS-accredited schemes) at no cost to the user.

The burden of providing accessible age assurance must sit with the platform, not the individual.

The experience of Australia is instructive. The central problem was not the technology — age assurance technology was available, proven and affordable — but the failure of social media platforms to deploy it in the first place. AVPA published a lessons learned report drawing on this experience which we commend to the consultation.

Q17. What risks or burdens may be associated with raising the minimum age of digital consent?

If the existing minimum age of digital consent was being implemented and enforced, there is no substantive marginal cost to amending it. If it is raised, it is likely that costs would fall as reliable age data becomes more available.

Q18. What should be considered to make raising the digital age of consent effective and workable?

A number of our members have extensive experience over the past 26 years implementing the USA's Children's Online Privacy Protection Act (COPPA) for their clients. This requires verifiable parental consent before the personal data of children under 13 is processed, when platforms are either directed at children or gain actual knowledge that a user is under 13.

This experience highlights the challenges of any policy requiring parental consent. First, there are no easily accessible records of which adults are legally responsible for which children. COPPA is written to be limited to what is possible within the state of the art, and in practice the primary check — often the only check — is that the person giving consent is an adult.

There are also important issues in families where parents are separated and may disagree on whether to give consent: a parent with care may refuse consent only to find a non-resident parent then gives it. Children in care also need special consideration.

We are working with the IEEE to define a new standard for parental consent in the 2089 series, which already includes 2089.1 on age verification itself. This will look at what evidence might be used, alone and in combination, to build confidence in the relationship being the correct legal basis for consent, and will standardise what parents and guardians are asked to consent to.

On balance, we would warn that any policy relying on parental consent will be legally and technically hard to implement. It also leads to disparities between children and could undermine the achievement of any new social norming around the appropriate age for general access to unrestricted social media.

Q19. There is a case for changing the digital age of consent for some online services but not others — strongly agree.

Where a service has been designed from the ground up for a particular age group and has deployed robust, independently certified age assurance to verify that its users fall within that group, there is a strong case for treating it differently from a general-purpose social media platform. The key principle is that the age of consent applied to a given service should reflect both the nature of the service and the quality of the age assurance deployed. Services that invest in higher-quality age assurance should be able to demonstrate compliance with a higher degree of confidence and should be recognised for doing so.

Q23. What do you think the impacts would be if some online services were required to introduce age restrictions on specific features and functionalities?

We refer to our previous comments on privacy, data and costs. Restricting functionalities by age is only effective if the underlying age assurance is robust. Functionality restrictions built on top of self-declaration or easily circumvented device-based age signals provide very limited genuine protection. The quality of the age assurance underpinning any functionality restriction is therefore as important as the restriction itself.

Q28. What do you think the impacts would be if online platforms were required to restrict specific features or functionalities, or to introduce time limits?

Restricting a single user's total time on social media or any other category may quickly lead to a requirement for proof of identity, not just age, for affected users. In effect, each underage citizen would be required to be given a virtual voucher for the time permitted per day to spend across all platforms in scope. To prevent users obtaining more than one voucher, it would need to be tied to their unique identity. A less rigorous approach could limit use by device, which can be circumvented with extra devices. In practice, this would be a per-platform, per-account limit without significant shared infrastructure.

Q37. Should AI chatbots have minimum age restrictions? Yes — both minimum age requirements and restrictions on access to certain features and functionalities.

While we generally take a policy neutral approach, we make two exceptions — one for adult content and one, as a recent addition, for AI chatbots. It is our collective view as technologists that current general-purpose AI chatbot systems do not yet have sufficiently mature and consistently reliable safeguards to be deployed to children without robust age-based restrictions and graduated protections.

The approaches in use across publicly available large language model products currently fall into two categories. The first is conversational inference, where the model analyses a user's language and interaction patterns to form a probabilistic estimate of whether they may be a

minor. The second is self-attestation at sign-up, where a date of birth or age confirmation checkbox is recorded without verification. Neither approach constitutes highly effective age assurance. Both are systematically vulnerable to low-effort circumvention — a child need only type convincingly or enter a false date of birth. Furthermore, AI systems can readily be induced to move from restricted to unrestricted behaviour through conversational manipulation: an inability to directly answer a question about dangerous content does not prevent a system from answering it when reframed as fiction or creative writing. Robust, independently certified age assurance is required, not inference or self-attestation.

Age assurance should not be treated as a one-time event in this context. Higher-risk services including AI chatbots should consider age assurance across the full account lifecycle, including registration, behavioural changes suggesting age inconsistency, access to newly restricted features and periodic re-assurance over time.

CHAPTER 3: ENFORCEMENT AND COMPLIANCE

Q40. What should be considered to make minimum age restrictions effective and workable?

Our mission is to create an Age Aware Internet; we believe all digital services should be able to know the age range of their users with proportionate certainty to prevent harm to children.

The appropriate regulatory standard is not perfect prevention, but reasonably effective age assurance that materially reduces underage access compared with the status quo. No safeguarding system in either the online or offline world is entirely immune from circumvention. Policymakers should resist arguments that cast any possibility of circumvention as a reason not to require robust age assurance; the relevant question is whether systems materially reduce harm at scale.

The single most important factor in making minimum age restrictions effective is the requirement for independent, third-party audited and certified age assurance. The experience of Australia demonstrates this clearly. Australia's minimum age of 16 for social media has been in force, yet platforms were not required to use independently certified age assurance and were permitted to rely on self-assessed compliance. As AVPA's lessons learned report documents, the central failure was not a lack of available technology but a failure by platforms to deploy age assurance at all.

YouGov data commissioned for 7News Australia illustrates the consequence. Conducted after the law came into force, the survey sampled 1,298 children aged 13-15 — every one of whom was below the legal age of access. It found that these underage users were still actively using the major platforms in large numbers:

Social media platforms still currently used by 13-15 year olds in Australia (YouGov / 7News,

n=1,298)
YouTube — 73%
TikTok — 55%
Instagram — 49%
WhatsApp — 45%
Snapchat — 39%
Facebook — 20%
Discord — 10%
Reddit — 3%
Other — 3%
None of these — 1%

The technology to prevent this existed. The legal obligation existed. What was absent was any requirement to deploy independently certified age assurance, without which self-assessed compliance is meaningless.

Regulators should also be alert to the risk of performative or "malicious" compliance, where platforms deploy age assurance measures in ways that are technically weak, inconsistently applied or deliberately friction-heavy in order to create the impression that age assurance itself is ineffective, rather than that their particular implementation is inadequate. This dynamic has been observed in previous waves of online safety regulation and will recur unless certification and independent audit requirements are sufficiently robust.

AVPA published a lessons learned report on the Australian implementation which we commend to this consultation. The central lesson is that the question is not whether the technology exists — it does — but whether platforms are required to deploy it and to have that deployment independently verified.

Certification frameworks already exist. The Age Check Certification Scheme already operates a certification scheme for age assurance providers. The IEEE and other accredited conformity assessment bodies can audit compliance. Any statutory requirement should specify that age assurance must be provided by a supplier certified under a recognised scheme, and that platforms must demonstrate their compliance to an independent auditor rather than self-certify.

Platforms should retain sufficient auditable evidence of age assurance decisions and assurance levels to enable meaningful regulatory oversight and independent review, without retaining any personal data.

The proximity principle

Age checks should happen as close as possible to the moment a user attempts to access age-restricted content or services — what AVPA calls the proximity principle. A useful analogy is a construction site: the health and safety check happens at the gate, not at the point where

someone decided to enter the building industry. Checking age at the app store is the equivalent of checking at the point of career choice rather than at the site entrance. The principle is already well established in regulated sectors such as gambling and adult content, where verification occurs at the point of sign-up or content access.

App-store or device-level age signals happen much earlier — at the point of installation or device setup — long before a user engages with any specific content. This timing failure has several practical consequences. First, it cannot prevent underage access to age-inappropriate material within apps that contain mixed-age content, such as social media platforms or online gaming services, without banning the entire app — a step that raises legitimate concerns about children's rights to knowledge, networks and support. Second, even a correctly age-gated app may contain an in-app browser, providing a portal to the wider web that bypasses any installation-level restriction entirely. Third, any framework that relies primarily on app-store controls will fail to address direct browser access, progressive web apps, sideloaded applications and shared-device usage.

App-store-level verification also concentrates sensitive age data in a small number of powerful intermediaries, creating centralised databases that become targets for data breach and enabling operating system providers to surveil or profile users based on verified age status — risks that do not arise with independently certified age assurance deployed at the service level.

The Internet Matters report published in May 2026 provides relevant evidence on bypass methods in practice. The most common method children reported was entering a false date of birth, cited by 13% of children. This is not a failure of age assurance technology; it is the absence of it. The next most common methods were using someone else's login (9%) and someone else's device (8%) — the online equivalent of an adult purchasing age-restricted goods for a minor, a risk no regulatory system eliminates entirely but which regular re-authentication and behavioural monitoring can substantially reduce. VPN use was reported by only 7% of children, consistent with wider sector data and a direct counterpoint to claims that age checks would be easily sidestepped.

Age assurance must therefore be applied at the point of access to regulated content or functionality, not upstream at device or operating system level. App-store and device signals may form a useful supplementary layer within a broader ecosystem, but cannot and should not be treated as sufficient for compliance in their own right.

The Trust Platform Model and secure enclave concern

We would further caution against any regulatory framework that treats device-level or app-store age signals, including those delivered through Trusted Execution Environments or secure enclaves, as sufficient in themselves for compliance. Such approaches do not generally provide the relying service with independently auditable assurance of: how the age was established; whether the signal remains accurate over time; whether the current user of the device is the

same individual originally age-checked; or the applicable level of assurance.

A recent example illustrates this point. When Pornhub reopened access to UK iPhone users following Apple's iOS age attestation mechanism, some commentators treated this as a compliance milestone. In AVPA's view, it does not clearly satisfy the underlying objective of highly effective age assurance. Pornhub does not check the user itself, cannot know whether the adult who verified with Apple is the person actually using the device, and receives no independently auditable confirmation of how the age was established. The duty to implement highly effective age assurance sits with the service provider, not with the device manufacturer. Device and operating system level signals may form a useful part of a layered ecosystem, but should not displace interoperable, independently certified age assurance at the point of access to regulated services.

Proportionality and graduated assurance

Different risks require different levels of assurance. A proportionate framework should permit lower-friction methods for lower-risk use cases while requiring stronger forms of assurance where the risks to children are higher. The framework should specify minimum levels of age assurance — see IEEE 2089.1 Appendix A — by risk category rather than prescribing specific methods, so as to preserve space for innovation and accessibility.

Interoperability and competition

The government should avoid approaches that inadvertently create centralised technical choke-points or distort competition by favouring a single operating system, wallet provider or platform ecosystem. In particular, government should avoid creating de facto mandatory reliance on any single state-issued wallet, operating system provider or identity infrastructure, as this risks reducing competition, innovation and long-term resilience in the market. Interoperability and standards-based portability between age assurance providers are essential. Systems such as Yoti Age Connect, OpenAge and AgeAware allow a single age check performed for one platform to be reused by multiple other platforms, with the user's consent, substantially reducing friction for users and cost for platforms.

The same concern applies to proposals that would make government-issued digital identity the primary or default mechanism for age assurance. A government digital ID scheme crosses a conceptual line that age assurance policy has deliberately avoided: the requirement for age confirmation becomes, in public perception and potentially in practice, a requirement to present identity to access online services. This conflation is corrosive to public trust and to the privacy principles underpinning the regulatory framework. It is also operationally problematic: government digital identity schemes are not generally designed for use by minors, and the EU's own experience illustrates this clearly. The European Digital Identity Wallet was not built with provision for under-18s, requiring the EU to develop a separate age verification app to deliver anonymous checks for minors. France has made the same distinction deliberately, launching a

standalone age verification app precisely so that it is not conflated with the France Identité EUDI wallet in the public mind. The UK government should follow the same principle: age confirmation and identity verification are distinct functions that should be delivered by distinct infrastructure, with age assurance remaining in the hands of a competitive, independently certified private sector rather than concentrated in a single state-issued credential.

Measurable outcomes

Platforms should be required not merely to claim compliance, but to demonstrate measurable effectiveness through independent audit, testing and reporting against standardised metrics. The current absence of mandatory outcome reporting is a significant gap.

Q41. What do you think the impacts might be from requiring age assurance across a greater number of online platforms?

The growth of interoperable age assurance infrastructure — such as Yoti Age Connect, OpenAge and AgeAware — means that requiring age assurance across more platforms will generate network effects that reduce the marginal cost of each additional check, making compliance progressively cheaper as adoption increases. The privacy argument against wider age assurance is also weaker than is sometimes suggested: independently certified age assurance systems process substantially less — sometimes zero beyond the user's own device — personal data than the advertising and tracking ecosystems already deployed by major platforms as a routine matter.

Q42. How, if at all, could age assurance be made more effective?

Age assurance can be made more effective through six principal means.

First, the requirement for independent certification. Platforms should only be permitted to satisfy their age assurance obligations using suppliers certified under a recognised independent scheme. Self-certification, device-based signals without independent verification and parental attestation without supporting evidence should not constitute compliance. Two decades of largely ineffective parental controls demonstrate that platform self-assessment does not protect children.

Second, application of the proximity principle. Age checks must happen at the point of access to restricted content or functionality, not upstream at device or app-store level. Browser-based access to restricted content cannot be controlled by app-store signals at all, and any compliance framework that does not address this gap will be routinely circumvented.

Third, a requirement to offer multiple methods and for platforms to fund backstop options. No single age assurance method is suitable for all users. Platforms must be required to offer a range of certified methods, and where automated methods fail or are unavailable, they must

fund access to professional attestation services such as CitizenCard or PASS-accredited schemes at no cost to the user.

Fourth, interoperability. The government should support and where appropriate encourage adoption of interoperable age assurance infrastructure, so that a verified check performed once can be reused across services.

Fifth, account lifecycle assurance. Platforms should consider age assurance across the full account lifecycle, including registration, behavioural changes suggesting age inconsistency, access to newly restricted features and periodic re-assurance over time. Age assurance should not be treated as a one-time event at registration.

Sixth, VPN-aware compliance frameworks. Age assurance requirements must be accompanied by guidance on how platforms should handle traffic that arrives via VPN, so that VPN use does not become a straightforward route to circumvention, using behavioural and technical signals to flag accounts that will not disable the VPN for age or geolocation checks.

Q43. What should be considered when assessing the effectiveness of age-verification and age-assurance technologies?

Independently tested accuracy rates, including both false positive rates (adults incorrectly blocked) and false negative rates (children incorrectly passed), should be the primary measure. Accuracy must be assessed by independent testing bodies, not claimed by the vendor, and assessed across demographic groups including age, gender and ethnicity.

The level of assurance provided should be assessed against the ISO/IEC 27566-1 framework, which distinguishes between verification, estimation and inference. Regulators and assessors should resist treating these as equivalent simply because all three are described as "age assurance".

Privacy preservation should be assessed as a technical property of the system, not a policy commitment by the vendor. Systems that achieve age verification without transmitting or retaining personal data beyond what is strictly necessary should be preferred.

Accessibility should be assessed by testing the system against a representative range of users, including those without standard forms of ID, those with disabilities affecting biometric capture, and those using older or lower-specification devices.

Circumvention resistance should be assessed, including resistance to VPNs, borrowed or stolen credentials and AI-generated synthetic identity documents.

Proximity should be assessed: whether the check occurs at the point of access to the restricted content or service, or upstream at device or app-store level. Upstream checks are substantially

less effective.

Persistence should be assessed: whether the system provides for re-authentication and ongoing lifecycle assurance or relies solely on a one-time check at registration.

Interoperability with other certified systems should be considered, as systems that support interoperability reduce the overall burden on users and the market as a whole.

Auditability should be assessed: whether the platform retains sufficient evidence of age assurance decisions and assurance levels to enable meaningful regulatory oversight and independent review.

Suitability for the full user population should be assessed, including whether a system is designed and provisioned for use by minors. A compliance mechanism that cannot be used by under-18s cannot satisfy an age assurance obligation that applies to them.

Q44. What methods to circumvent online safety rules do you think children in the UK use, beyond VPNs or similar technologies?

The most widely used method of circumvention is simple: lying about age when self-declaration is the only check required. This requires no technical knowledge and has no friction. Where no genuine barrier exists, many children will not even perceive this as circumvention — it is simply how they access services. The persistence of self-declaration as the primary age check on most major platforms is therefore the single largest systemic failure in online child safety, as the Australian experience demonstrates. It is a failure of platform compliance and of the absence of any social norm that treats age limits as real. Consistently enforced age assurance would address both.

Beyond self-declaration, the principal circumvention methods and the technical responses to each are as follows.

Borrowed, shared or purchased credentials. Children use accounts belonging to older siblings, parents or other adults, or in some cases purchase access credentials. This cannot be addressed by a one-time registration check alone. It requires ongoing re-authentication beyond password, and continuous behavioural monitoring to flag users whose activity pattern is inconsistent with the verified age of the account holder.

Facial age estimation spoofing. Anecdotal reports of children using avatars, photographs or drawn disguises to defeat facial age estimation have received disproportionate attention. Certified facial age estimation systems are specifically tested for these simplistic presentation attacks as part of the audit process; a system defeated by a drawn moustache or a still image would not pass certification. Robust liveness detection, live image capture rather than file upload, and limits on retry attempts address the meaningful attack vectors. As Project DefAI —

AVPA's joint UK-Swiss research programme into presentation and injection attacks on age assurance systems — demonstrates, the more serious emerging threats are deepfake and video injection attacks using AI-generated synthetic imagery. Project DefAI is developing the test infrastructure and standards needed to assess and certify resistance to these attacks, which represents the genuine frontier of circumvention risk rather than the low-sophistication methods reported anecdotally.

Fake or borrowed identity documents. Document-based verification is vulnerable to the use of older siblings' or parents' documents, or to AI-generated synthetic documents. Live image capture at the point of verification, rather than acceptance of pre-existing files, substantially reduces this risk, as does liveness detection and document authenticity checking.

Device-based circumvention. Children circumvent device-level parental controls by using a second device, a friend's device, or by accessing services via a browser rather than an app — a route that app-store controls do not address at all.

VPNs. Recent evidence from both Childnet and Internet Matters showed no increase in VPN use among children when age assurance was required by adult sites. The Internet Matters report found only 7% of children reported using a VPN to circumvent age checks, consistent with wider sector data showing VPN use among children aged 9-17 has remained stable at around 8% since Online Safety Act protections came into force. VPNs are in any case more commonly used by older teenagers with greater technical confidence, and platforms should monitor VPN traffic for behavioural signals consistent with underage use.

Finally, policymakers should note that circumvention is not exclusively a child-to-adult problem. Adults also attempt to present themselves as children, for example to gain access to spaces intended for minors for the purposes of grooming or sextortion. Systems such as device-based parental controls that lack independent age assurance make this trivially easy. Age assurance that verifies age in both directions — confirming both that adults are adults and that those presenting as children are genuinely children — is therefore important for safeguarding as well as for access restriction. This is a further reason why robust, independently certified age assurance at the point of service is preferable to device-level or self-declared approaches.

Q45. Which of the options below do you think the government should prioritise to reduce circumvention of online safety rules in the UK?

More education for children, as a complement to — not a substitute for — mandatory age assurance. Education helps children understand why age limits exist and builds the social norms that make compliance the default expectation.

Q46. Everyone should go through age checks to access a VPN if it would prevent children using them
strongly disagree

Q47. What do you think the impacts would be if VPNs were age-restricted?

We oppose limiting access to VPNs for children. VPNs offer much-needed additional security, particularly when using public wifi connections. Restricting children's access to VPNs would expose them to greater risk in precisely the contexts where they are most vulnerable, such as on school, library or public transport networks.

There is no need to restrict children from using VPNs in order to implement age assurance effectively. The presence of VPN traffic is a signal that platforms can and should use to inform their age assurance approach, not a reason to prohibit VPN use altogether. VPNs are increasingly built into browsers, making any practical ban on child use both technically difficult to enforce and potentially counterproductive if children simply switch to less secure circumvention tools.

The correct policy response to VPN-related circumvention is to require platforms to implement VPN-aware age assurance rather than to restrict VPN access. Australia's eSafety Commissioner has published Reasonable Steps guidance that sets out the technical and behavioural indicators available to platforms to identify when traffic arriving via VPN is potentially from an underage user. These indicators include browser settings such as timezone and currency, content choices, usage patterns and first-party user data associated with geographic advertising targeting. Commercial VPN detection technology independently tested by Kingsmead Security in September 2025 achieves a 99.1% detection rate with a zero percent false positive rate, and is already deployed by the BBC, Netflix and Amazon Prime Video for geo-licensing enforcement and ubiquitously by regulated online gambling operators for state-by-state licensing compliance in the United States. The claim that VPN detection is technically impractical is therefore not credible.

Where a platform detects that a user may be connecting via VPN, the appropriate response is to offer that user a choice: to proceed through an age assurance check; to disable their VPN so that their IP address can confirm they are accessing from a permitted jurisdiction; or to use device-based geolocation tools of the kind deployed extensively in the regulated gambling market. This layered approach targets only the small fraction of users whose location cannot be reliably confirmed by IP address alone, and does not require treating all VPN traffic as high-risk.

Q48. What should be considered to make age-restricting VPNs effective and workable?

As set out above, we do not support age-restricting VPNs and believe that VPN-aware age assurance is both more proportionate and more effective. The government should instead consider publishing guidance analogous to Australia's Reasonable Steps guidance, setting out the technical and behavioural signals available to platforms to identify potentially underage VPN users and the compliance steps platforms should take in response. Such guidance would give platforms the clarity they need to implement compliant age assurance without either prohibiting

VPN use or treating the presence of a VPN as automatically rendering age verification impossible.

CHAPTER 5: SUPPORTING FAMILIES

Q60. Parents should have control over the online experiences of their children — somewhat disagree.

There are practical and principled reasons to be cautious about parental controls.

It is very demanding on parents to monitor what platforms each of their children is using, to learn what controls do and how to use them, and then to maintain those controls and adjust them with age.

Some parental controls are invasive of children's privacy and may conflict with the UN Convention on the Rights of the Child and General Comment 25.

More fundamentally, the evidence of the past two decades demonstrates that parental controls are not an effective substitute for platform-level age assurance. Parental controls place the burden of child protection on individual families rather than on the platforms that profit from children's engagement. They are inconsistently deployed, inconsistently effective and easily circumvented. The record of parental controls as a primary child safety mechanism is poor, and any policy framework that relies primarily on parental controls rather than mandatory platform-level age assurance will replicate that failure.