Response to DSIT 'Growing up in the online world: a national consultation'

Global Partners Digital submission
May 2026

About Global Partners Digital

Global Partners Digital is a social purpose company working to ensure that human rights underpin the development, use and governance of digital technologies.

Through advocacy, partnerships, capacity building, networks, and research, we shape rights-respecting laws and policies and help build a more diverse digital policy ecosystem.

Responses Submitted

Chapter 2: Interventions for safer, more positive experiences

Restricting social media services by age

Question 4: Would you support a legal requirement for social media services to have a minimum age of access?

-yes

-no

-don't know / prefer not to answer

Question 5: To what extent do you agree or disagree with the following statement: "Social media services should have a minimum age of access of at least 16 and should not be accessible to any children under that age"

a. Strongly agree

b. Somewhat agree

c. Neither agree nor disagree

d. Somewhat disagree

e. Strongly disagree
f. Don't know/ Prefer not to answer

Question 6: Would you support a legal requirement for social media services to have a minimum age of access lower than 16? If so, at what age would you set it?

a. Yes – 13
b. Yes – 14
c. Yes – 15
d. No – not lower than 16
e. Other (please specify)
f. Don't know/ Prefer not to answer

We do not support minimum age requirements for social media as a primary policy approach. Our concern is not with any particular age threshold but with the use of access restrictions as the primary mechanism for addressing online harms. As set out in our response to Question 7, any enforceable minimum age requires age verification of all users, creating disproportionate privacy risks at population scale, while evidence from Australia demonstrates that circumvention is widespread and harms are displaced rather than prevented.

Please explain the reasoning behind your answers about minimum age requirements.

GPD favours platform design obligations over access-based restrictions as the primary mechanism for protecting children online. The harms children experience online, including exposure to harmful algorithmic content, manipulative design, and commercial exploitation of their data, originate in how platforms are designed and monetised rather than in the absence of age gates. Any enforceable minimum age requires all users, not only those approaching the threshold, to submit to age verification, creating disproportionate privacy and surveillance risks at population scale that are inconsistent with human rights under the International Covenant on Civil and Political Rights (ICCPR) and children's rights under the UN Convention on the Rights of the Child (UNCRC) Article 16.

Children also hold positive rights under UNCRC Articles 13 and 17 protecting access to information, and Article 16 protecting privacy, that blanket access restrictions directly undermine. GPD's preference is for robust enforcement of existing Online Safety Act obligations on platform design, targeting the algorithmic features (including addictive design), advertisement monetisation and data practices that are the actual sources of harm.

Question 7: What do you think the impacts would be of having a minimum age requirement higher than 13 for social media services?

For example, impacts on the safety and wellbeing of children, or the impact for parents and carers, as well as other users. You could also comment on the impact on all users' privacy and data or on business costs, revenue, and innovation.

Social media services are an integral part of young people's daily lives and have documented benefits. Social media use enhances young people's ability to find and access information, communicate with peers, seek advice or support, and encounter communities to overcome isolation and cultivate safe spaces for interaction. These benefits are particularly key for young people from traditionally marginalised groups or those in vulnerable situations.

Conversely, social media engagement has been also recognised as a source of harm for young people. Tackling those harms is a critical and laudable aim that requires a multilayered strategy. Online spaces can post distinct and compounding dangers for young people, in particular those who may be marginalised, exposed to discrimination or violence. Responding to these harms requires thoughtful policy interventions that tackle the incentives for harm production inside platforms, and transcend the online realm to also target the offline causes of harm. This includes additional educational efforts fostering young people's capacity for critical thinking when consuming information online and emotion regulation skills to encourage their personal development.

A higher minimum age for social media would not eliminate the harms children face online. Those harms originate in how platforms are designed and how they exploit user data to shape user interaction, not in the absence of age gates. Evidence from the Australian eSafety Commissioner on the enforcement of the ban for under-16s from social media in December 2025, found that approximately 70% of affected children were still accessing restricted platforms within three months, with circumvention requiring no technical expertise. Harms were displaced, not prevented. Analysis of the ban found that circumvention followed a tipping point dynamic, with restrictions becoming ineffective once a critical mass of peers had found workarounds, suggesting that enforcement cannot keep pace with social norm change among adolescents.

Raising the minimum age, which would in practice require all users to verify their ages, also carries significant risks for the enjoyment of human rights for children and for adults. The UNCRC and General Comment No. 25 on children's rights in the digital environment affirm children's rights to freedom of expression, access to information, and participation in public life online. For marginalised young people, LGBTQ+ youth, and children from isolated communities, online spaces are often a primary source of support, information, and connection. Restricting access to these online spaces in a blanket manner constitutes a disproportionate restriction to the right to freedom of expression, freedom of association and other rights, limiting children's access to

content that has educational or developmental value. Blanket exclusion from those spaces is not a neutral or harmless policy choice for young people as it might push them to migrate to even less safe spaces or interact with malicious actors to circumvent the restrictions.

Any enforceable minimum age would impose age verification on all users, not only those approaching the threshold, creating disproportionate privacy risks at population scale. The surveillance infrastructure required would affect the entire user base to address a minority of accounts, and the data collected would far exceed what is necessary for the stated purpose. Age verification tools would provide a widespread tracking infrastructure ready to be instrumentalised by governments or malicious actors for surveillance or exploitative purposes.

These risks are not uniform: different age verification methods pose distinct and compounding harms. Some risk exclusion by design, relying on methods of identification that are not widely available across populations (such as ID, passport or other legal credentials, or the possession of smart devices and the skills to interact with them). Others are insufficiently accurate, relying on age estimation technologies that have produced discriminatory outcomes across demographic groups and cannot reliably distinguish users near the age threshold. Age verification methods that rely on intermediaries such as device providers or trusted authenticators risk creating additional centralization, further concentrating market power and producing associated security risks. This latter risk has been already experienced, with evidence of serious breaches of UK users' data, including the 2025 Discord incident, where government ID photos submitted for age verification were exposed.

GPD's position is that policy should instead target the actual sources of harm through platform design obligations: prohibiting behavioural profiling of children and targeted advertising, restricting addictive algorithmic features and enforcing existing Online Safety Act requirements relating to transparency and strengthening user control through content filters and reporting mechanisms. This is also the emerging multilateral consensus, as reflected in a joint civil society statement signed by GPD and others, as well as in the Council of Europe Commissioner for Human Rights February 2026 statement, which explicitly cautioned against blanket social media bans and called on governments to regulate platform design rather than restrict children's access, and a statement by UNICEF noting age restrictions alone will not keep children safe and called on platforms to proactively redesign their products.

Restricting access to services based on features and functionalities

Question 15: What do you think the impacts would be if some online services were required to introduce age restrictions on specific features and functionalities?

For example, impacts on the safety and wellbeing of children, or the impact for parents and carers, as well as other users. You could also comment on the impact on all users' privacy and data or on business costs, revenue, and innovation.

GPD is concerned about the structural risks posed by the widespread introduction of age-based restrictions across platforms. The enforcement of such restrictions through a regulatory mandate would require the introduction of user-wide identity authentication and tracking infrastructure, obligating the collection of the sensitive data of all users and not just those around a certain age threshold. These impacts are discussed in more detail in our response to Question 33.

GPD takes the view that these considerable limitations and risks reinforce the case for preferencing platform design obligations rather than age assurance wherever possible. Such an approach is consistent with General Comment No.25 of the Committee on the Rights of the Child, which emphasises the requirement to integrate privacy-by-design into digital products and services that affect children, and the narrow use of robust age verification systems in the case of products and services that are illegal for them to own and use (paragraph 70; paragraph 114). It also aligns with the Council of Europe's Recommendation CM/Rec(2026)4, adopted in April 2026 across 46 member states, which positions safe-by-design obligations as the primary regulatory tool, with age assurance as a secondary and narrowly scoped measure.

GPD is concerned that restricting access to privacy-preserving functionalities poses significant risks for the enjoyment of human rights, failing to meet the principle of proportionality, and would negatively impact the privacy and security of all users, including young people. Privacy-preserving solutions like Virtual Private Networks (VPNs) and encryption have been designed by technologists to protect the privacy of individuals' digital communications, shielding them from the surveillance, extraction and breach of their personal data.

VPNs can be essential privacy and security tools for all users. As noted by this consultation, the main reason they are used is "to access the privacy and data protection benefits that VPNs offer". Children and young people rely on VPNs for legitimate purposes. For instance, they can offer enhanced data protection and shield them from monitoring and exploitation by third parties when connecting to unsecured public WiFi, such as public or school networks. VPNs are an easy-to-use and accessible tool, helping children and young people to assert their agency and protect their personal data when navigating the online world.

End-to-end encryption (E2EE) likewise provides critical privacy and security benefits. E2EE mathematically scrambles sensitive data to make sure that only the sender and the intended recipients can access it, preventing third parties from interfering with sensitive financial information, identity, or other sensitive data. When E2EE is not adopted, this provides the technical means for businesses and others to intercept

children's data, allowing for the harvesting, monetisation or exploitation of children's data. Additionally, research by cybersecurity experts and by the European Parliament Research Service has shown that there are currently no feasible means of allowing access to content on systems that are E2EE without compromising the security of the system as a whole. General Comment No. 25 of the Committee on the Rights of the Child explicitly identifies encryption as a tool for protecting children's right to privacy in the digital environment. Previous cases have shown that measures to undermine encryption have led providers to withdraw privacy-preservising features from the UK market.

'Addiction', compulsive design and displacement

Question 20: What do you think the impacts would be if online platforms were required to restrict specific features or functionalities, or to introduce time limits? *For example, impacts on the safety and wellbeing of children, or the impact for parents and carers, as well as other users. You could also comment on the impact on all users' privacy and data or on business costs, revenue, and innovation.*

The impacts of requiring platforms to restrict specific features depends entirely on which features are concerned. Restricting harmful design choices, including infinite scrolling, push notifications, behaviour profiling and recommender systems optimised for engagement over wellbeing, could help reduce the exploitation of children's attention and data and would be broadly consistent with children's rights under UNCRC. On the other hand, if restrictions extend to privacy-preserving features such as VPNs or end-to-end encryption, the impact would be severely detrimental as these tools protect children, and other vulnerable users from surveillance, data exploitation, and abuse.

Tackling exploitative algorithms and data practices should be prioritised by combining robust enforcement of the UK's existing platform design obligations under the UK Online Safety Act with the timely development of new, rights-respecting policies. Proposals could usefully target design features such as infinite scrolling and alerts and push notifications, operating at the design level to limit exploitative data collection.

GPD supports the recommendation on algorithmic transparency contained in the House of Commons Science, Innovation and Technology Committee's report "Social Media, Misinformation and Harmful Algorithms". We also support the principle that platforms should demonstrate how they are addressing all risks identified in their own risks assessments, and not only those covered by Ofcom's Codes of Practice. However, we would caution against framing this as a hard legal requirement, which could incentivise over-removal of content and have unintended consequences for freedom of expression. We also urge caution on the Committee's other recommendations, which carry similar risks.

As we have previously proposed, we encourage policies that require service providers transparency and strong user tools to enhance children's and parents agency over the types and amounts of content that they see and interact with. This also includes community-led moderation in groups and forums, allowing young users to set their own terms of engagement. Platforms should also ensure that children are clearly informed of how to opt out of seeing particular content types that they do not want to see, and of how to report harmful content easily to the platform, and be redirected towards alternative content, such as helplines or resources when a harmful interaction is reported.

What type of services should restrictions apply to

Question 21: What factors are important when determining which apps, sites or services to apply minimum age of access restrictions to?

For example, user-to-user interaction, the ability to post material, persuasive design features, risky functionalities, the ability to generate non-text mediums such as video or images, the target age group, the size of the service.

As set out in our responses to Questions 4, 7, and 15, GPD favours platform design obligations over access-based restrictions as the primary mechanism for protecting children online. Where the government proceeds with access-based restrictions, the following factors should guide their application.

The primary factor should be whether the harm is specifically attributable to children's developmental vulnerabilities in ways that cannot be adequately addressed through design obligations applicable to all users. Where a service's harmful features can be made safe for all users through design, restricting children's access is disproportionate. Restriction should be reserved for services whose fundamental purpose exploits children's developmental vulnerabilities in ways that cannot be mitigated without restricting access, or for services already illegal for children.

A second factor is proportionality: whether the privacy and data collection costs imposed on all users by enforcing age restrictions are proportionate to the specific harm being addressed. As the Knight-Georgetown Institute's January 2026 technical assessment found, no currently available age assurance technology can verify user age without creating structural risks for all users. These costs must be weighed seriously against the harm any given restriction aims to prevent.

A third factor should be children's evolving capacities under UNCRC Article 5 and General Comment No. 25. Rather than binary access restrictions, policy should consider whether graduated or tiered approaches calibrated to age and maturity can deliver protective outcomes while preserving children's rights to access information and participate in digital life. Harmful design features that target all users, not only children, should be addressed through design obligations rather than age restrictions.

Question 25: Some services are already exempt from the Online Safety Act. Examples include internal business services, services with limited functionalities and services provided by persons providing education or childcare. Are there additional types of service which you think would be appropriate to exempt from age restrictions?

There are several categories of service for which age restrictions would be disproportionate or counterproductive, and which GPD considers should be explicitly exempt.

Privacy-preserving services, including VPN providers and end-to-end encrypted messaging applications, should be exempt from age-based access restrictions. Their core purpose is to enable users to communicate and browse without their identity being linked to their activity. As set out in our responses to Questions 15, 32, and 39, applying age assurance requirements to these services requires providers to collect identity data that defeats the purpose of the service and undermines the protections it exists to provide, including for the children and vulnerable users who rely on it most: those in unsafe home situations, LGBTQ+ youth, and those for whom online privacy is a safety necessity rather than a preference. This exemption does not extend to obligations to address clearly illegal content such as CSAM; GPD's position is that such obligations should be met through methods that do not require compromising encryption or collecting identity data from all users.

Services providing access to educational content, health information, essential services (such as banking and government services) and general information resources should also be exempt, consistent with children's rights under UNCRC Articles 13, 17, 28, and 29. Restricting children's access to information they have a right to seek would be disproportionate regardless of the platform hosting it.

General-purpose AI tools and chatbots used for educational, research, or creative purposes should be exempt from blanket service-level age restrictions. As set out in our responses to Questions 26 and 28, restrictions on AI tools should be targeted at specific high-risk features rather than applied at the service level, and implemented through provider design obligations rather than user-facing age verification requirements.

Chatbots and AI

Question 26: What are the benefits to children of using AI chatbots?
For example, this might include as a search function, for educational purposes, for creativity.

Children hold positive rights that AI chatbots can directly support. UNCRC Articles 13 and 17 guarantee the right to seek and receive information from diverse sources; Articles 28 and 29 protect the right to education; and General Comment No. 25 confirms these rights apply in digital environments. AI chatbots extend these rights in

practical terms: they provide personalised learning support, assist children with learning differences, and make expert knowledge accessible to children without well-resourced schools or tutoring. For children in lower-income households or geographically isolated communities, this is not a marginal benefit but a meaningful equalisation of access to rights the UNCRC already guarantees.

UNCRC Article 16 guarantees children the right to privacy, including the ability to seek information without exposure or surveillance. AI chatbots can provide a space for children to ask about health, identity, sexuality, or family situations that they may not feel safe raising with adults, particularly for LGBTQ+ young people, children in abusive home situations, and those facing discrimination. Accessing this information privately is not incidental; it is itself a rights-bearing activity. Overly broad restrictions on AI chatbot access risk severing this for the most vulnerable children. The Electronic Frontier Foundation has warned that wide-ranging AI age-gating proposals would in practice block access to general-purpose educational and information tools, not just companion chatbots, disproportionately harming children who rely on them most. Policy must be targeted at harmful features rather than the tools themselves.

Question 27: Which AI chatbot features are most risky for children?

(Please select all that apply)

-the realism of interactions, including realism of content generated

-the personalisation of interactions

-how they mimic relationships (friendship)

-how they mimic relationships (romantic)

-how they mimic empathy

-flattering language

-features to encourage more questions/requests (e.g. asking questions back)

-the ability to recall interactions across sessions

-the type of content generated - A) video, B) text, C) audio, D) image

-allowing children to have accounts

-hallucination or false, misleading responses

-ability to engage in and generate mature content (e.g., sexual / romantic roleplay)

-don't know / prefer not to answer

-none of the above/ai chatbot features are not risky for children

-other (please specify)

We have selected the above options as illustrative examples of the features that most clearly implicate children's rights. The full reasoning behind which features warrant restriction, and the human rights framework informing that assessment, is set out in our response to Question 28.

Question 28: Which functionalities of AI chatbots should minimum age restrictions apply to?

AI chatbots are tools for accessing information, as such any restrictions must satisfy the principles of legality, legitimate aim, and be necessary and proportionate, meaning they are no broader than required. UNCRC Article 3 requires children's best interests to be a primary consideration, and Article 5 recognises their evolving capacities, meaning restrictions should be graduated and targeted rather than blanket. The features identified in our response to Question 27 share a common characteristic: they are engineered to exploit children's emotional and cognitive vulnerabilities to manipulate them, foster dependency, erode their agency or expose them to harm. UNCRC Article 36 requires protection from this form of exploitation regardless of the specific mechanism. This is the test: not whether a feature appears on a list, but whether it is designed to exploit children's developmental vulnerabilities in ways that undermine their rights under Articles 16 (privacy), 36 (protection from exploitation), and 17 (access to accurate information).

On this basis, the functionalities warranting restriction are companion and relationship-simulation features designed to foster emotional attachment, cross-session persistent memory building intimate user profiles, romantic or sexual roleplay, features simulating therapeutic relationships without clinical oversight or crisis referral, and the ability to generate sexually explicit content. The American Psychological Association's Health Advisory on AI and Adolescent Well-being found that adolescents are particularly susceptible to these features given their heightened social sensitivity and developing impulse control. Legal cases, including those reported on by NPR, have confirmed the severity of the resulting harms, including the deaths by suicide of 14-year-old Sewell Setzer III (Character.AI) and 16-year-old Adam Raine (ChatGPT) (*Garcia v. Character Technologies*, US District Court Florida, 2024.

General-purpose conversational, educational, creative, and information features must remain freely accessible. Restricting them would violate children's rights under Articles 13, 17, 28, and 29, and would fall hardest on the most vulnerable children who rely on AI tools for access to information they cannot safely seek elsewhere. Critically, restrictions must be implemented through design obligations on providers rather than age verification requirements imposed on users. Requiring all users to verify their identity to access a chatbot service creates disproportionate privacy risks for the entire user base and is inconsistent with General Comment No. 25's emphasis on privacy-by-design as the appropriate mechanism for protecting children online. Providers should be required to demonstrate through their product architecture that high-risk features are inaccessible below the relevant threshold, without any user-side identity verification.

Question 29: Should AI chatbots have minimum age restrictions?

-Yes - minimum age requirements for AI chatbots

-Yes - restrict access to certain features and functionalities

-Yes - both minimum age requirements and restricting access to certain features

-No

-Don't know / prefer not to answer

Please explain the reasoning behind your answer:

A blanket minimum age for AI chatbots would be disproportionate and incompatible with children's rights under UNCRC Articles 13 and 17. The category spans general-purpose educational and information tools at one end and companion applications designed to simulate intimate relationships at the other. Treating them as a single regulatory category is not consistent with a rights-based or proportionate approach. Restrictions must be narrowly targeted at the high-risk features identified in Question 28 and implemented through design obligations rather than age gates that impose identity verification on all users. As noted in our response to Question 28, the test is not which features appear on a list but whether a feature exploits children's developmental vulnerabilities in ways that undermine their rights under the UNCRC.

GPD's primary concern is that AI chatbot regulation must not become a pretext for compelling surveillance of private communications. UNCRC Article 16, ICCPR Article 17, and ECHR Article 8 all protect the confidentiality of communications, and General Comment No. 25 emphasises the requirement to integrate privacy-by-design into digital products and services that affect children, and the narrow use of robust age verification systems in the case of products and services that are illegal for them to own and use (paragraph 70; paragraph 114). Conversations with AI chatbots are private communications and must be protected by strong end-to-end encryption in transit and at rest. As the Internet society maintains, any regulation that mandates on-device scanning of conversations, client-side content monitoring, or other mechanisms that circumvent encryption directly violates these rights. The Online Safety Act already contains powers capable of compelling such scanning. AI chatbot regulation must explicitly exclude their activation. Access Now has consistently argued that human rights-centric AI governance requires that the right to private communication be preserved in AI design, not compromised by it.

Question 30: What do you think the impact would be of introducing age restrictions on AI chatbots or certain features and functions?

For example, impacts on the safety and wellbeing of children, or the impact for parents and carers, as well as other users. You could also comment on the impact on all users' privacy and data or on business costs, revenue, and innovation.

Age restrictions requiring all users to verify their identity before accessing AI chatbot services would have an immediate chilling effect on the exercise of rights under UNCRC Articles 13 and 17 and ICCPR Article 19. Children seeking health information, information about their identity, or educational support would be deterred from

accessing tools they have a right to use. This chilling effect falls most heavily on those with the greatest privacy needs, including children in unsafe home situations, those exploring identity, and those from communities facing discrimination, precisely the groups the UNCRC is designed to protect.

Blanket restrictions would also produce displacement: children excluded from regulated AI tools would turn to unregulated alternatives where no safety obligations apply, producing worse outcomes for the rights and safety the restrictions were intended to protect. Feature-level design obligations, grounded in UNCRC Article 3 (best interests) and proportionality, could address the specific harms identified without restricting children's access to the informational and educational benefits of AI tools that Articles 13, 17, 28, and 29 protect.

The most serious impact of poorly designed AI chatbot regulation would be the creation of a legal basis for surveilling private communications at scale. Any requirement for providers to scan, log, monitor, or report conversation content would violate UNCRC Article 16, ICCPR Article 17, and ECHR Article 8. It would build surveillance infrastructure into AI services that would, once established, be available for purposes far beyond the child safety context that justified it. Freedom of expression under ICCPR Article 19 and UNCRC Article 13 requires that individuals be able to seek information without fear of monitoring; mandatory conversation scanning would destroy that assurance. GPD, alongside the broader civil society coalition that signed the May 2026 joint statement, urges the Government to adopt a rights-based approach grounded in platform design obligations rather than creating new surveillance obligations that would outlast any particular regulatory moment.

Chapter 3: Enforcement and compliance

Improving age assurance

Question 31: To what extent do you agree or disagree with the following statement: "Adults should complete age checks more often, if it means children are safer online"

-strongly agree
-somewhat agree
-neither agree nor disagree
-somewhat disagree
-strongly disagree
-don't know/prefer not to answer

Question 32: What should be considered to make minimum age restrictions effective and workable?

This could mean either age restrictions for access to whole services, or for specific risky or 'addictive' features or functionalities.

Any age assurance system must be built around privacy by design. Verification should confirm only that a user meets an age threshold, establishing nothing further about their identity, location, or online activity. Technologies such as zero-knowledge proofs can achieve this cryptographically, confirming a fact without disclosing the underlying data. However, as the Internet Society's policy brief on age restrictions observes, this technology is still under development and not widely deployed. Systems that prevent any linkability between a user's identity and the services they access must be the baseline, not an aspiration.

Data minimisation is equally essential. Age assurance must not produce centralised databases of identity documents or biometric data. Such databases represent serious security threats and a potential infrastructure for surveillance. The 2025 Discord breach, in which government ID photos of approximately 70,000 users submitted for age-related appeals were exposed, illustrates the real-world consequences of collecting sensitive identity data at scale. No personally identifiable information should be retained beyond the moment of verification. A January 2026 technical assessment by the Knight-Georgetown Institute (KGI) found that most widely deployed age assurance architectures require users to trust an age verification provider (AVP) with no technical mechanism preventing that provider from colluding with the service to match the user's identity to their activity. Users have no way of verifying this is not happening. This is a structural risk inherent to centralised AVP models, not a question of good intentions or contractual commitments.

Age assurance must also be accessible, accurate and non-discriminatory across the full population. Requirements based on government-issued ID, bank accounts, or biometric cameras will exclude care-experienced young people, refugees, disabled users, elder not digitally skilled and those from lower-income households, undermining the right to equality, right of access to information and digital participation for those who most need protection. Facial estimation systems carry documented performance disparities across demographic groups, raising direct discrimination concerns. Research of specific age estimate and age inference technologies have also shown them to be highly inaccurate close to age thresholds, calling into question their reliability and whether users trust in these systems. Finally, systems must be interoperable across jurisdictions to prevent fragmentation of the open internet, and subject to independent audit so that claims of privacy compliance can be verified. As CDT and the Internet Society have both documented, there are severe limitations and risks, which reinforce the case for preferencing platform design obligations rather than age assurance wherever possible.

A further constraint is that age assurance must be carefully designed to preserve the full protective value of end-to-end encryption and not only its technical integrity. The

working document of the Internet Engineering Task Force (IETF) on age verification architecture explicitly states that "when security tools are considered services that need age-gating such as in proposals to not allow youth to use end-to-end encryption this puts them at great risk and would never be supported by security considerations. Nor would age-gating of encryption be possible without some kind of intervention akin to backdooring encryption." For users who rely on these services precisely because no one should be able to link their identity to their communications, this is a serious harm to the value and purpose of encryption.

A more acute concern arises where verification extends to content-level checking: on 8 May 2026, Ofcom published its Statement on Technology Notice powers setting out minimum standards of accuracy for technologies that could be required under section 121 of the Online Safety Act to scan communications within encrypted services for CSEA and terrorism content, confirming that these powers are now being actively operationalised. Any age assurance framework must explicitly exclude their activation and must permit encrypted services to satisfy age assurance obligations through device-level or credential-based approaches requiring no access to communication content.

Question 33: What do you think the impacts might be from requiring age assurance across a greater number of online platforms?

For example, impacts on the safety and wellbeing of children, or the impact for parents and carers, as well as other users. You could also comment on the impact on all users' privacy and data or on business costs, revenue, and innovation.

The primary risk of expanding age assurance across more platforms is the construction of population-level identity and tracking infrastructure. Every service that mandates verification becomes a point at which sensitive data is collected about all users, not only minors. That data, whether government IDs, biometric profiles, or behavioural signals, can be breached, repurposed, or shared. The risk that age assurance infrastructure becomes a general mechanism for online identification is not hypothetical; it must be addressed through binding technical constraints, not contractual commitments.

Wider mandates will also have a chilling effect on the exercise of rights. Users who fear that accessing a service requires submitting personal data will self-censor or disengage entirely. This particularly affects those for whom anonymity or privacy is not a preference but a necessity: survivors of domestic abuse, political dissidents, LGBTQ+ individuals in hostile environments, journalists, human rights defenders, whistleblowers, and others. As the Internet Society's policy brief on age restrictions observes, even the perception of a privacy intrusion is sufficient to deter users from engaging with age-gated services, noting specifically that teens and parents are skeptical of facial age estimation requiring a webcam scan. When that deterrence falls

on the most vulnerable users, access to information and free expression are directly compromised.

At the international level, UK mandates set a precedent. Governments that wish to impose identity requirements on internet users as a condition of access will point to UK policy as justification. The UK's historic support for an open, interoperable, and globally accessible internet is a genuine asset; it should not be traded for age assurance measures that could be replicated in far more restrictive forms elsewhere. The Council of Europe's Recommendation CM/Rec(2026)4, adopted in April 2026 across 46 member states, takes precisely the opposite approach: it positions safe-by-design obligations as the primary regulatory tool, with age assurance as a secondary and narrowly scoped measure. The UK's consultation choices will be assessed against this emerging international consensus.

These risks are compounded where mandates extend to privacy-preserving services such as VPNs and encrypted messaging applications. As observed by the working document of the Internet Engineering Task Force (IETF) on age verification architecture, applying age assurance technologies to these services "puts them at great risk, and would never be supported by security considerations". Privacy-preserving services should be explicitly exempt from expanded age assurance mandates, or when exceptionally implemented adopt device-level approaches requiring no collection of identity data.

Question 34: How, if at all, could age assurance be made more effective?

The most effective age assurance approaches are those that do not require service providers to collect or store sensitive personal data at all. Client-side and device-level tools, including operating system age signals and parental control mechanisms, allow age-appropriate filtering without any server-side data collection. Credential-based systems, where a verified age attribute is held by the user and presented on a need-to-know basis, mean that services learn only whether a threshold is met, nothing more.

Zero-knowledge proof technologies and device-based enforcement are the most rights-preserving options currently available, and the only architectures the Knight-Georgetown Institute's January 2026 technical assessment (referenced in Q32) identifies as ensuring that neither the age verification provider nor the service provider learns the user's identity at all. ZK proofs allow a user to prove cryptographically that they satisfy an age requirement without disclosing their actual age, identity, or the source of their credential. Device-based enforcement checks age on the device itself, with no server-side data transmission. Crucially, both approaches can operate without any access to the content of communications, making them fully compatible with end-to-end encryption. The KGI assessment also found that no single age signal is sufficient on its own, and that allowing users to select among multiple privacy-preserving signals protects privacy more than requiring a predetermined

sequence, a principle that should be embedded in any UK framework. The UK Government could consider standardising ZK-based and device-level approaches instance than centralised AVP models, which the KGI assessment found to be structurally incapable of preventing collusion between providers and services.

However, as the Internet Society's policy brief on age restrictions observes, this technology is still under development and not widely deployed. GPD supports the recommendation by CDT and others that more work and multistakeholder dialogue is needed among developers, policymakers, civil society and academics to address privacy, security, and human rights issues prior to wide deployment. The OECD's "Towards Digital Safety by Design for Children" framework reinforces design obligations over access bans as the evidence-based approach. A final element to consider is that, as a 2026 Joint statement of security and privacy scientists and researchers on Age Assurance points out, these privacy preserving approaches can bolster centralization by pushing users towards mainstream phone manufacturers that amass more power on the market and create single points of failure for security.

Finally, effectiveness also depends on platform accountability. Evidence from Australia shows that the main reason under-age users access restricted services is not circumvention through technical tools; it is that platforms fail to implement meaningful checks or allow users to re-verify at low confidence. Age assurance cannot substitute for enforcing design obligations on the platforms where harm actually originates.

Question 35: What should be considered when assessing the effectiveness of age-verification and age-assurance technologies?

Assessments of age assurance technologies must be grounded in rights as well as accuracy. A system that reliably identifies age but requires invasive data collection, performs poorly for marginalised groups, or enables tracking is not an effective or acceptable solution. GPD recommends assessments consider the following.

Privacy and security: Does the system collect only what is necessary to confirm the age threshold, with no retention of personally identifiable data after that point? Are there technical safeguards preventing the data from being used for tracking, profiling, or secondary purposes? Are security protections built into the protocol itself, and has the system been independently tested for vulnerabilities? The Knight-Georgetown Institute's January 2026 technical assessment provides a ready-made framework for this evaluation, assessing age assurance systems across four criteria: baseline accuracy, circumvention resistance, availability, and privacy. Its finding that most deployed systems require users to trust an AVP with no technical mechanism preventing collusion should be treated as disqualifying for any system that cannot demonstrate a technical rather than contractual solution to this risk.

Encryption compatibility: Does the technology require, or create incentives for, access to the content of encrypted communications? Does its deployment depend on

client-side scanning, backdoors, or any mechanism that would compromise end-to-end encryption? Any system that cannot be deployed without undermining E2EE must be disqualified from use in regulated contexts. This criterion is not technical pedantry; it reflects the fact that encrypted services are used by the most vulnerable people online, and that age assurance cannot be permitted to become a justification for dismantling the security architecture that protects them.

Accessibility and non-discrimination: Can the technology be used without a government ID, bank account, smartphone, or working camera? Has its accuracy been tested across demographic groups, including different skin tones, genders, disabilities, and ages? Systems that systematically misidentify or exclude particular groups violate equality obligations and restrict access to information for those least able to contest the outcome. The KGI assessment specifically found that facial age estimation cannot reliably distinguish whether a user is just above or just below the age threshold, and must therefore reject users who are not clearly older than the threshold, meaning it systematically excludes users who are entitled to access rather than only those who are not. It also found that government-ID-based systems, while accurate, are not available to many users, and are particularly unavailable to minors themselves, making them unsuitable as a primary mechanism for any threshold below 18.

Transparency, user rights, and interoperability: Do users know what data is collected and have meaningful recourse when errors occur? Is the system interoperable, allowing a credential to function across services and jurisdictions without repeated data exposure? All assessments must be conducted by bodies independent of government and the technology provider, and results published in full.

GPD wishes to be explicit that even a well-designed age assurance framework that satisfies all of the above criteria would not eliminate all risks. Some residual harms cannot be addressed through technical safeguards alone. The structural risk that verification infrastructure, once built, is later repurposed for broader surveillance or identification cannot be fully mitigated by privacy-by-design commitments at the point of deployment; it requires ongoing legislative vigilance, enforceable use-limitation rules, and genuinely independent oversight. Age assurance of encrypted services remains incompatible with E2EE regardless of how the verification mechanism itself is designed. And the chilling effect on access to information for users who distrust any data-collection requirement may persist even where the technology is demonstrably privacy-preserving. GPD acknowledges these limits, and submits that they reinforce rather than undermine the case for preferring platform design obligations over age assurance wherever possible.

Circumvention of age limits

Question 37: Which of the options below do you think the government should prioritise to reduce circumvention of online safety rules in the UK?
(Please select the most important one to you)
-more education for children
-restricting children's access to VPNs
-none of the above
-don't know/prefer not to answer
-other (please specify):

Research indicates that VPNs are not predominantly used as a circumvention tool. For instance, evidence from the Australian eSafety Commissioner shows that less than 7% of those surveyed said that they believed their child still held an account across age-restricted social media platforms because of VPN use. Instead, it notes that the most common reason children still had their social media accounts was that they had not been asked by the platform to verify their age.

In the UK, research by Internet Matters found that only 8% of children have used a VPN in the past 12 months. This indicates that the primary motivation for the use of VPNs is not circumvention, but more likely legitimate concerns around user privacy and data security. This consultation also refers to early evidence that the spike in VPN use following the enforcement of the UK Online Safety Act's age assurance requirements for primary priority content "does not suggest the peak in VPN use was driven by children seeking to bypass age assurance processes."

Question 38: To what extent do you agree or disagree with the following statement: "Everyone should go through age checks to access a VPN if it would prevent children using them"
-strongly agree
-somewhat agree
-neither agree nor disagree
-somewhat disagree
-strongly disagree
-don't know/prefer not to answer

Question 39: What do you think the impacts would be if VPNs were age-restricted?
For example, impacts on the safety and wellbeing of children, or the impact for parents and carers, as well as other users. You could also comment on the impact on all users' privacy and data or on business costs, revenue, and innovation.

As discussed in the response to Question 15, VPNs are essential privacy and security tools for all users. Children and young people rely on VPNs for legitimate purposes, helping them to assert their agency by protecting their personal data and shielding themselves from third-party intrusions while using unsecured public or school networks. Restricting VPN use by age would prevent children from benefiting from the critical privacy and security benefits that VPNs provide. This conflicts with the aims of improving digital security online and promoting the use of digital skills. It also conflicts with children's rights under UNCRC Articles 13 and 16, which protect their right to access information and privacy online, and with General Comment No. 25's recognition of privacy-preserving tools as mechanisms that support rather than undermine children's safety.

Additionally, as our response in Question 37 indicates, research suggests that restricting VPNs would not be an effective measure to reduce circumvention of the UK's online safety rules. Imposing through a regulatory mandate a blanket age-restriction on VPNs would represent a disproportionate restriction of the rights of young people. In our view, arbitrary restrictions on technical, privacy-preserving tools compromise the principle of respect for the evolving autonomy, capacity and privacy of the child, reflected in General Comment No. 25.

Another key concern is the structural risks posed by the imposition of age-restrictions on the use of VPNs. The enforcement of such restrictions would require all VPN users to verify their age, obligating them to provide sensitive personal data to be able to access a privacy-preserving technical tool. This undermines the inherent value of strong VPNs, and exposes a number of at-risk professions and groups who rely on the privacy protections that VPNs afford. This includes journalists protecting their confidential sources, political dissidents, public servants requiring high levels of security, whistleblowers or domestic abuse survivors and providers requiring confidentiality and privacy to protect themselves or their users from harm.

Question 40: What should be considered to make age-restricting VPNs effective and workable?

For example, public trust and engagement with increased age assurance requirements, accessibility of age assurance methods and variations of age assurance approaches across services, interaction with legitimate uses of VPNs.

As earlier discussed, GPD does not consider age-restricting VPNs to be an effective or proportionate policy approach. Age-restricting VPNs by imposing identity verification on all users would significantly undermine the positive privacy protections afforded by the use of VPNs, impacting both children and adult users. This would most damagingly prevent at-risk professions and groups from benefitting from the anonymity and privacy guarantees afforded by high-quality VPN providers.

Furthermore, any age assurance system applied to VPNs would face the same structural limitations identified in our responses to Questions 32 and 35: centralised identity verification creates databases of sensitive data vulnerable to breach, facial age estimation performs poorly across demographic groups, and no currently available technology can verify VPN users' ages without collecting identity data that defeats the purpose of using a VPN in the first place.

If the government decides to pursue the age-restriction of VPNs, we refer to our response in Question 32, which reflects the principles of a privacy-by-design approach to the adoption of an age assurance system.