

ACT | The App Association response to the United Kingdom's Office of Communications (Ofcom) Call for evidence: "Second phase of online safety regulation: Protection of children"

1. To assist us in categorising responses, please provide a description of your organisation, service or interest in protection of children online.

For providers of online services, please provide information about:

- the type of service and functionalities you provide;
- number of users globally, and in the UK (including children and their ages);
- global and UK revenues; and
- your business models and revenue generation.

Please indicate where this information is confidential.

ACT | The App Association is a global trade association for small and medium-sized technology companies. Our members are entrepreneurs, innovators, and independent developers within the global app ecosystem that engage with verticals across every industry. The App Association supports the United Kingdom's leadership in creating a regulatory environment that promotes innovation and job growth. Today, the ecosystem the App Association represents—which we call the app economy—is valued at approximately £736.2 billion globally and is responsible for creating countless jobs across the United Kingdom. Alongside the world's rapid embrace of mobile technology, our members have been developing innovative hardware and software solutions that power the growth of the internet of things (IoT) across modalities and segments of the economy.

Consumer trust is fundamental for competitors in the app economy, especially for smaller firms that may not have substantial name recognition. Strong privacy protections that meet evolving consumer expectations are a key component of developing consumer trust in tech-driven products and services. The App Association helps shape and promote adherence to privacy and other consumer protection laws and best practices in a variety of contexts, including for apps directed to children.

2. Can you identify factors which might indicate that a service is likely to attract child users?

In particular, please provide evidence explaining:

- the types of services which are likely to attract child users;
- any functionalities or other features of a service which are particularly likely to attract child users;
- the type of content that is likely to attract child users;
- if or how the factors you've identified may differ depending on the age of a child user; and
- whether there are any noticeable patterns in the activity of child users.

Where possible, please specify whether this evidence relates to child users in the UK or globally.

Various factors may indicate whether a service is likely to attract child users. Types of services likely to attract children may include games, educational content, or video streaming. Depending on the subject matter, however, the same type of services may be directed to adults. Therefore, the subject matter of an online service is significant when determining whether an online service may attract child users, e.g. an educational app that teaches spelling in a gamified way or a streaming app that hosts children’s cartoons and animated movies are more likely to be interesting to a young child than an adult. Visual content that includes animated characters or child-oriented activities and incentives (e.g. gamification) can further indicate likelihood that an online service is likely to attract children. The use of “celebrities” for children (e.g. a beloved TV or book character) to promote or presence in the online service further increases the likelihood of attracting children to a service. As children age, the visual content they are attracted to and whom they consider “celebrities” may change. Other determining factors can include the music or audio content used, the language of advertising or the website, branding and modes/channels of distribution (e.g. making a service available on a device that children are more likely to use), and using geo-fencing to show targeted advertisements at places children may frequent. Some online services will also state that their intended audience is children.

We believe, however, that language such as “likely to attract” is too vague. Plenty of online services could be reasonably likely to attract a child (which is not explicitly defined in the UK’s age-appropriate design code, but under GDPR covers persons under 16 years of age). A 13-year-old could, for example, use a cooking app to find a recipe to bake cookies with a friend. That app is by the above-mentioned standards not ‘likely’ to attract a child, but it’s perfectly imaginable that children would use it. Therefore, we urge Ofcom to proceed with caution in using this broad wording and carefully consider implications of including it in the guidance Ofcom plans to release. Anticipating Ofcom’s preparation of codes of practice for online safety, we note that the ‘likely to attract’ wording may encompass a larger group of businesses than we believe Ofcom intends to regulate. Much of the consternation regarding children’s online safety stems from the practices of large online community platforms that serve a wide, public base of individual users and rely on surveillance-based advertising for revenue. It is well known that children are a regular by-catch in these environments. Unfortunately, many smaller app developers that do not cater to an audience primarily consisting of children could easily fall in the scope of such codes of practice, which would add further barriers and compliance costs for small businesses that already operate with limited resources.

3. What information do services have about the age of users on different platforms (including children)?

In particular, please provide evidence explaining:

- the methods used to gather any information that can assist in estimating or assuring a user’s age, either at the point a user first accesses the service or subsequently;
- what, if any, mechanisms are available to enable services to identify children in different age groups (for example children below age 13, aged 13-15, or aged 15-17); and
- how approaches to assessing the age of users are evolving.

Platforms or websites use various methods to gather information to estimate or assure a user’s age. Such mechanisms include:

- self-declaration (sometimes called age-gating), in which a website or platform asks users to state their own age or check a box to certify they are above a certain age (e.g. over 16, 18, or 21 years old);
- questionnaires through which the user provides information that allows the platform or website to identify the user’s age (e.g. by cross-referencing the user provided information against a database that contains authenticated age information);
- hard identifiers such as document submission to the website or platform by the user to show their age, such as a government-issued form of identification, which is then reviewed by the online service provider;
- automated estimation uses facial recognition software to estimate a user’s age or categorise them as a minor or an adult;
- profiling/inference;
- cross-account verification; and
- third party age assurance, using e.g. digital identity or age tokens by a trusted online provider

An online service provider can implement age assurance tools at the point of access to a service or use these tools sequentially, e.g. to age restrict certain parts of content or certain functionalities of the service. Approaches to age assurance are evolving to a certain extent as technological capacities change, e.g. using facial recognition was not a possibility a few years ago. Please see question 5 for benefits and concerns associated with each age assurance method.

4. How can services ensure that children cannot access a service, or a part of it? In particular, please provide evidence explaining:

- how age assurance policies have been developed to date and what age group(s) they are intended to protect;
- if the service is tailored to meet age-appropriate needs (for example, by restricting specific content to specific users or part of a service), how this currently works or could work;
- how the efficacy of age assurance policies is or could be monitored; and
- how services can identify users that do not meet any relevant age limits and how is this appropriately addressed?

Age assurance policies should correspond to the nature and risk a product or service presents in relation to a child's age. The less risky a service is, the less intrusive and the lower the bar of age assurance should be. If a service is entirely appropriate for users of all ages, it may not need to implement any age assurance mechanism at all. Tailoring a service to meet age-appropriate needs by restricting specific content or features to specific users could be accomplished by implementing an age assurance mechanism at the time a user tries to access those features. Another way to accomplish it could be to make certain content or features available only to a certain age group and ask the user to verify their age when they sign in/up. Their stated age follows them across the platform and the platform only shows appropriate content according to age.

In terms of efficacy, any age assurance mechanism must be balanced against its potential impact on the child. For example, restricting a child's access to extreme violent or sexual content or protecting them from grooming or contact with anonymous adults requires highly accurate and efficient age assurance methods. That said, monitoring efficacy is difficult because it's impossible to measure the incidents that did not occur due to age assurance. Perhaps a decrease in negative online incidents that involve children could be an indicator. We also note that the phrasing of the question uses the word 'ensure', which means to have certainty, may be too high of a standard. Depending on the content and the possible harm to children, there may be a range between discouraging use by children to gaining near certainty/ensuring that children cannot access a certain service.

5. What age assurance and age verification or related technologies are currently available to platforms to protect children from harmful content, and what is the impact and cost of using them?

In particular, please provide evidence explaining:

- how these technologies can be assessed for effectiveness or impact on users' safety;
- how accurate these technologies are in verifying the age of users, whether accuracy varies based on any user characteristics, and how effective they are at preventing children from accessing harmful content;
- any potential unintended consequences of implementing age assurance (such as risk of bias or exclusion), and how these can be mitigated;
- the safeguards necessary to ensure users' privacy and access to information is protected, and over restriction is avoided;
- which methods of age assurance users prefer, when offered a number of ways to verify their age;
- the cost of implementing and operating such technologies;
- how age assurance and age verification or related technologies may be circumvented; and
- what mitigations exist to reduce circumvention among users.

Age assurance and age verification that are currently available to platforms include self-declaration, hard identifiers (accessing existing databases of previously established identification data), biometric estimation, profiling based on user behaviour, capacity testing, cross-account authentication, third-party age assurance provider (digital identity, age tokens, B2B), account holder confirmation, and device/operating system controls. Assessing these technologies for effectiveness and impact on users' safety should consider the following factors: privacy preservation, proportionality, ease of use, impact on user experience, security, accessibility, transparency/accountability, and the level of friction.

Self-declaration requires a user to enter their date of birth or check a box to certify they meet the minimum age required to use the online service. Self-declaration is easy to use and implement but offers a comparatively low level of assurance and puts responsibility directly on the child to report their age accurately. It is not a reliable tool because users can misreport their age, and children may not understand the consequences on their user experience of pretending to be older. While easy for children to use and for businesses to implement, self-declaration seems only suitable for low risk and non-intrusive online services.

Relying on hard identifiers means users must provide verified sources of their age, such as a copy of their ID or passport or other identifying information that can be cross-checked against official databases, e.g. in the UK this could be a national insurance number. While such hard identifiers provide a high level of age assurance (the documents in question have already been verified), they often contain more information than just a user's age such as their name and address or sensitive data like race and gender, making this a highly privacy-invasive mechanism. Additionally, most children will not have access to such documentation and most services require an additional verification by the parent to match the identification to the child. Additionally, this measure may be exclusionary to those who don't have government-issued or other official age documentation.

Biometric scanning or facial recognition-based age assurance has become more accurate in recent years, but it continues to fail to accurately recognize facial characteristics of both young children and people with darker skin tones. The level of age assurance thus varies from low to high confidence. While facial recognition can be implemented in privacy preserving ways, e.g. by discarding the user's image immediately after age has been estimated, most users do not understand the type of data that this tool collects, how it is used and how it may be further shared and stored. Such automated estimation often creates serious privacy risks because a person's face is highly sensitive personal information and if it is 'digitally stolen' it can impact a person's life without any good fixes.

Inferring age by profiling creates all kinds of privacy risks and excessive data collection issues. Processing data to estimate a user's age consists of information a user chose to share about themselves as well as information the online service provider infers or collects

from the user's engagement with the service. Such information can include the time spent on the website or app, times of day the service is accessed, where a user is located, what interests are, whom they interact with, and more. Building detailed profiles of users, especially children, is highly invasive and heavily restricted by GDPR. Age inference based on profiling not only interferes with a child's right to privacy, but it also offers a low level of assurance if the data quality is poor or wholly inaccurate. Profiling also usually violates the data protection principle of data minimisation, and it is likely that an online service provider would collect more than it needs to estimate age.

Capacity testing means that a user's age is assessed based on an aptitude or capacity test, such as completing a puzzle or another task that would indicate their age or age range. While these tests are privacy- and child-friendly and easy to implement for service providers, an adult could easily complete them on behalf of a child. Children's capacity does not equal age, and children develop at different speeds, so capacity testing is not suitable for situations where the user's exact age is necessary because they may only suggest whether a child is above or below a certain age range. It also may be an exclusionary tool for children with lower aptitudes or developmental disabilities.

When using cross-account authentication, a child can use an existing account to access a new service or product or feature (e.g. sign in with Google or sign in with Apple). When the child enters the correct username and password for their existing account, the online service provider allows access for the new service or product to the child's user data via an API. Often it is unclear what user data is being shared between the two providers, e.g. whether it is just the child's age or if name, location, and data are also being shared. While this method is convenient for children, the level of age assurance is unclear as the original authenticating provider determines the method and therefore the level of age assurance in this scenario. The opacity around which data is shared between providers further risks violating children's privacy rights.

There are companies that provide identify confirmation and/or age assurance services. Such third parties can help online services providers by offering tokenised age checking, API solutions, or background checks or to users directly by providing digital IDs.

Digital identities or credentials offer a high level of age assurance, and they can minimise personal data sharing, offering users more control over their identity. With a digital identity, a user does not repeatedly need to submit documents for information on their age, as they have already done so once to the third-party digital identity provider. Once the digital identity is established, users can store it in a digital wallet and use it to identify themselves when signing up for other online services. Privacy risks exist nonetheless, as holding large amounts of personal information in a centralised database like a digital wallet can increase danger of fraud or commercial misuse.

Business-to-business age assurance (B2B) minimises user engagement in age assurance, but the process lacks transparency and oversight. Users are often unaware a third party is part of the assurance process, making it difficult to obtain valid consent, especially from children. Adding a third party into the process also increases personal data sharing, exposing users to heightened privacy risks.

Age tokens contain information exclusively related to a user’s age, allowing the online service provider to confirm whether a user meets age requirements without having to collect any other personal information. The attribute provider that generates the age token determines the initial method of age assurance, so age tokens minimise data sharing, but the level of assurance depends on the method the provider chooses. The technology to generate age tokens is not yet widely available or taken up, but age assurance could evolve as age token innovation progresses.

Account holder confirmation requires a service provider to get confirmation of a child’s age or age range from a person that the provider knows to be an adult (e.g. a parent or caregiver). The adult can then either set up shared accounts or set up a child-specific account. All aspects of such a child-specific account, including design and content filtering, should provide for age-appropriate experiences. This method seems appropriate for younger children, but for older children, it could pose risks to children’s privacy (including privacy from parents). It could also exclude children who struggle to obtain confirmation from a parent or caregiver as well as those children with parents who may not have access to hard identifiers like a government-issued ID.

Parents already go through the process of providing confirmation of a child’s age to the device manufacturer when setting up a new device for their child. To ease the burden on parents, it may be useful to allow device manufacturers to retain and provide that verification to covered platforms available for the device, should the developer choose to do so. This method would be easy to use and privacy preserving, but the full picture is complicated by the fact that many families have either a lot of devices children may use or shared devices that people of different ages use. Using the device as source of truth is therefore also difficult.

As the above demonstrates, there is currently not an ideal way to conduct age assurance in a way that is both accurate and privacy preserving, and while each may have benefits, there are also drawbacks that require implementation of further safeguards to protect a child’s best interests.

8. How do services currently assess the risk of harm to children in the UK from content that is harmful to them?

In particular, please provide evidence explaining:

- how risks from harmful content are identified (including any relevant internal processes, policies and documents); and

- in considering the potential risk that children may encounter harmful content, the extent to which services factor in evidence on users' behaviour and age.

Generally, App Association members' services assess the nature of the app/site/service, whether it is likely to attract children, and whether the content poses any risk to children. For user-contributed content, services may implement a process that detects bad actors, the intensity of which is usually proportionate to the threat level. For example, a recipe app will probably not have many people incorporating inappropriate content but may have special functionality to restrict the discussion of alcohol.

9. What are the exacerbating risk factors services do or should consider which may have an impact on the risk of harm to children in the UK?

In particular, please provide evidence of:

- how the user base of the service may have an impact on the risk of harm to children;
- how the business model of the service may have an impact on the risk of harm to children;
- the functionalities or features of services which may have an impact on the risk of harm to children; and
- what mitigations exist for these risk factors

It seems extremely difficult to create a general list of exacerbating risk factors for impact on children since many situations are specific and unique to different apps, services, or sites. So rather than creating specific hard rules, it may be better to focus on processes where issues can be reported and addressed effectively.

10. What are the governance, accountability and decision-making structures for child user and platform safety?

As part of your answer, please outline how different teams may consider child user safety risks across different business functions such as product development, management, engineering, public policy, safety, legal, business development and marketing. Please consider:

- how staff are/should be trained in a service to understand how their own roles and responsibilities can create risks to child user safety;
- how services can ensure consistency in consideration of child user safety across teams;
- examples of best practice or innovative approaches to governance, accountability and decision making; and
- possible costs associated with assessing risks of harm to child users, including specific reference to costs associated with ensuring services have governance and decision-making processes for child user safety where applicable.

We note that only apps that are targeted towards children should have such structures for child safety. A determination of whether an app or website targets children should be based on its subject matter (visual content, use of animated characters or child-oriented activities and incentives, music or other audio content, age of models, presence of child celebrities or celebrities who appeal to children, language or other characteristics of the website or online service, as well as whether advertising promoting or appearing on the website or online service is directed to children), and such a designation should require actual knowledge by the developer/operator that the app or website is collecting personal information directly from users of another website or online service directed to children (and, an app or website should not be considered to be directed towards children if (1) it does not collect personal information from any visitor prior to collecting age information; and (2) it prevents the collection, use, or disclosure of personal information from visitors who identify themselves as under age 13 without first complying with the notice and parental consent provisions of this part). Regular, general-use apps or services should not be required to train staff about child-specific issues and should not be deemed directed to children solely because they refer or link to a commercial website or online service directed to children by using information location tools, including a directory, index, reference, pointer, or hypertext link. In cases of services directed at children, the relevant staff should be trained to understand that they should act in the ‘best interest of the child’ when designing, developing, marketing, and operating an online service likely to be accessed by a child. Children should enjoy special privacy protections and the online service should provide features for parental control and restrictions.

11. What can providers of online services do to enhance the clarity and accessibility of terms of service and public policy statements for children (including children of different ages)?

Please submit evidence about what approaches make terms or policies clear and accessible to children.

Terms of service and public policy statements towards children must be easily accessible, easy to understand, concise, and in intelligible language. It may also make sense to use standardised, machine-readable icons in policy statements directed at children.

15. What actions do or should services take in response to reports or complaints about online content harmful to children (including complaints from children)?

Please provide relevant evidence explaining your response to this question.

We believe the actions services should take in response to reports or complaints are highly situation/context-specific, but it is advisable for service providers to have a process for these situations in place and take good faith actions consistent with these policies based on actual knowledge.

19. With reference to content that is harmful to children, how can a service mitigate any risks to children posed by the design of algorithms that support the function of the service (e.g. search engines, or social and content recommender systems)?

In particular, please provide evidence explaining:

- if different from the risk assessment process outlined in response to Q8 how services assess the risk to children from algorithms central to the function of the service;
- what safeguards services have in place to mitigate the risks posed by algorithms (e.g. testing them before they are put into use, and monitoring their performance in real world settings);
- what safeguards services have in place to mitigate the risks posed using recommender systems in particular (e.g. providing users with controls over what they are shown, such as through keyword filters);
- additional requirements for safety in algorithms (e.g. accurate content categorisation);
- the costs involved in implementing these safeguards. In the absence of specific costs, please provide indication of the key cost drivers;
- how services can measure the effectiveness of these safeguards, in terms of reducing harm to users;
- what information services can provide to demonstrate the effectiveness of such safeguards; and
- how services can assess the impact of these safeguards on users' privacy and minimise the risk of over restriction.

While the types of data items analysed by AI and other technologies are not new, AI-driven analyses will provide greater potential utility of those data items, including in the context of preventing harm to children. There are many new uses for, and ways to analyse, data collected through apps and websites. While this raises privacy issues and questions surrounding consent to use data in a particular way (e.g. research, commercial product/service development), it also offers the potential for more powerful and granular access controls for consumers. While it is important for Ofcom to address AI-related privacy, consent, and modern technological capabilities in this consultation, requirements and obligations for online safety of children should be scalable and assure data is properly protected while also allowing the flow of information and responsible evolution of AI. Further, we generally encourage frameworks impacting AI to, consistent with our views above, also promote data access, including open access to appropriate machine-readable public data, development of a culture of securely sharing data with external partners, and explicit communication of allowable use paired with informed consent.

28. Other than those covered above in this document, are you aware of other measures available for mitigating the risk and impact of harm from content that is harmful to children?

We would be interested in any evidence you can provide on their efficacy, in terms of reducing harm to child users, cost, and impact on user rights and user experience

Another aspect Ofcom could consider is producing resources and educational materials for parents/caregivers and/or teachers/educators to help teach children appropriate cyber hygiene, how to use reporting tools, how to identify harmful content, etc. Such materials should be developed through a public-private partnership model that is inclusive of SMEs.