



Consultation Cover Sheet

Basic details

Consultation title: Guidance for service providers publishing pornographic content

Representing: Free Speech Coalition

Confidentiality

Part of response to be kept confidential: None

Question 1: Do you agree with our proposed guidance on scope? If not, please provide any information or evidence in support of your views, including descriptions of services or content where you consider it is unclear whether they fall within the scope of Part 5.

The section entitled, **What does “published or displayed by the provider on its internet service” mean?** is quite difficult to understand.

Paragraph 3.13 states, “Where an entity or individual has control over which content is published or displayed on an internet service, that entity or individual will be treated as the provider of the internet service...” This seems to contradict the fact that the entity that owns/controls the service (i.e., the public-facing website or application on which content appears) is the only one that can practically implement the guidance, regardless of the degree to which it was actively involved in publishing the content.

In the adult industry, there are a variety of business models for monetizing sexually explicit content. A few examples:

- Subscription “paysites” - websites that sell access to adult content on a subscription basis
- Content retailers - websites that sell access to discrete pieces of adult content (single videos, movies, or collections of videos/photos) for download or streaming. Some offer access on a rental or pay-per-minute streaming basis.
- Tube sites - websites where content producers can upload videos to be accessed for free to entice users to purchase their content on another platform.
- Fan platforms - websites where content creators sell access to content on a subscription basis.
- Clips marketplaces - websites where content creators sell access to discrete pieces of content.
- Live-streaming/cam platforms - websites that content creators use to stream live video content to customers who generally pay on a per-minute or per-performance basis.

In all of these cases, regardless of who is responsible for publishing the content, the website/platform itself is the only entity that can require age assurance. Fictionalised examples of how this can work in practice:

- SimplyFans is a fan platform. Content creators such as Laine Raine have accounts on the SimplyFans website that allow them to upload and publish explicit videos and images. Customers can purchase a subscription to access Laine’s content. Laine decides what

content is published, when, and what a subscription costs. SimplyFans can disallow or take down individual pieces of Laine’s content (or their entire account). Laine has no control over the features of the SimplyFans website, including whether and how it implements age assurance.

- Clips2Own is a clip marketplace. In addition to offering their content as a subscription on SimplyFans, Laine has a storefront on the Clips2Own website for their production company, Swell Wellies. Customers can purchase individual videos from the Swell Wellies storefront on the Clips2Own marketplace. Laine Raine uploads the content, publishes it, and determines the price. Clips2Own can disallow or take down individual pieces of Laine’s content (or the entire Swell Wellies storefront). Laine has no control over the features of the Clips2Own website, including whether and how it implements age assurance.
- ChatStream is a live-streaming/cam platform. Laine Raine engages in live performances on ChatStream’s website where customers can pay a set price per minute to watch them. Laine has some control over the price charged within strict boundaries set by ChatStream. Laine chooses when and how often to schedule live performances, over which ChatStream exercises no influence. ChatStream can end the broadcast of a live performance or ban Laine from using the ChatStream platform. Laine has no control over the features of the ChatStream website, including whether and how it implements age assurance.

In all of these instances, an entity or individual that does not control whether and how a provider implements age assurance does have control over which content is published on an internet service, at least in part. I think that simplifying language where possible and incorporating more examples from business models being used in the adult industry would be helpful for clarity.

Question 2: Do you have any comments on how our proposed guidance applies in respect of pornographic content created by generative-AI services within the scope of Part 5? Please provide any information or evidence in support of your views.

No.

Question 3: Do you have any comments on our proposed guidance in respect of the kinds of age assurance which could be highly effective? If you consider there are other kinds of age assurance which have not been listed that you consider could fulfil the proposed criteria, please identify these with any supporting information or evidence.

As a general matter, the Free Speech Coalition feels that placing the burden of age assurance on individual websites is detrimental to the law's goals for the reasons we outline below.

Our opinion on the specific methods proposed depends heavily on whether Ofcom officially endorses their use (see answer to question 4), the extent to which consumers will trust them (see answer to question 6), and the availability of simple and affordable ways to implement them (see answer to question 10).

Question 4: Do you agree that service providers should use the proposed criteria to determine whether the age assurance they implement which is highly effective at correctly determining whether or not a user is a child? Please provide any information or evidence in support of your views.

Section 82(2)(a) of the Act requires Ofcom to produce “examples of kinds and uses of age verification and age estimation that are, or are not, highly effective at correctly determining whether or not a particular user is a child.” The guidance does not do this.

As section 4.12 of the consultation acknowledges, Ofcom does “not have sufficient evidence as to the effectiveness and potential risks of different age assurance methods to recommend specific metrics for assessing whether or not any given age assurance method or process should be considered highly effective.” The guidance instead offers examples of the kinds of age assurance that *could* be highly effective and puts the burden on service providers to prove that they meet the very stringent and technical criteria set forth in sections 4.24 through 4.67. We do not feel that this satisfies Ofcom’s obligation.

The Act specifies in section 81(3) that “[t]he age verification or age estimation must be of such a kind, and used in such a way, that it is highly effective at correctly determining whether or not a particular user is a child.” It does not say that the service provider is responsible for ascertaining this.

If Ofcom lacks sufficient evidence to assess age assurance methods, what reason would it have to expect that pornographic service providers are in a position to generate it? The guidance seems

to assume that providers will be designing *and* implementing age assurance, which will not be the case for the vast majority of those affected by the law. Platforms that publish pornographic content are not experts in age assurance technologies and cannot be reasonably expected to conduct tests to evaluate their technical accuracy, robustness, reliability, and fairness.

Rather than suggesting methods that *could* be effective, Ofcom needs to assess the options and provide a standard, as well as a list of methods that meet that standard, as required by the Act. We appreciate the level of flexibility offered by the guidance – certainly a few platforms will attempt to build their own age assurance technology – but it seems better targeted to the third parties that will provide age verification services, not the overwhelming proportion of platforms that will be implementing their solutions.

Question 5: Do you have any information or evidence on the extent of circumvention risk affecting different age assurance methods and/or on any steps that providers might take to manage different circumvention risks for different methods?

The following is offered as an example of non-compliance: “the service provider has implemented photo-ID matching which easily allows children to verify their age using fake or manipulated ID documents.” Yet there is no standard or benchmark against which to measure.¹

Recent reporting has revealed that fake identification documents capable of fooling commercial identity verification software can be purchased online for as little as \$15.² It’s not unreasonable to expect that generative AI will make measures such as “liveness detection” easily circumventable as well, but in the meantime, hackers are targeting this extremely valuable face scan data.³

Service providers are not in a position to determine the level of risk these technologies present. The guidance should define what constitutes an acceptable or unacceptable level of risk and how to calculate it.

¹ It also seems to exceed the requirement of the law that children are “not normally able to encounter” the content. It seems absurd to suggest that a child who has gone to the effort of procuring fraudulent identification documents capable of fooling age verification software is “normally able to encounter” the content.

² Joseph Cox (5 February 2024) 404 Media, [Inside the Underground Site Where ‘Neural Networks’ Churn Out Fake IDs](#)

³ Taryn Plumb (21 February 2024) Venture Beat, [Face off: Attackers are stealing biometrics to access victims’ bank accounts](#)

Question 6: Do you agree with our proposed guidance that providers should consider accessibility and interoperability when implementing age assurance? Please provide any information or evidence in support of your views.

The burden Ofcom is placing on individual service providers to ensure that the age assurance method they use is highly effective, extremely secure, and accessible by persons with protected characteristics, but not easily circumvented, and doesn't "unduly prevent adult users from accessing legal content" is, simply put, impossible to meet.

As noted in the guidance, "[a]n age assurance method which performs poorly in test conditions will perform worse in a real-world deployment," which makes the observations of a consortium of UK-based companies and organisations called euCONSENT quite concerning. They found that 16% of all adults and 21% of parents they tested were unable to complete the age verification process of two large commercial providers (Yoti and AGEify) under test conditions.⁴ This is unsurprising in light of news reporting that demonstrated the "lengthy, time-consuming process" of age assurance using Yoti – a task that required 52 separate steps to complete.⁵ Ultimately, 42% of the euCONSENT research participants did not rate their experience positively, and 21% rated it negatively.⁶

Australia's eSafety Commissioner has also studied the topic, finding that "stakeholders emphasised that measures which create too much friction have the potential to deter users from accessing compliant sites. Instead, they may follow the path of least resistance toward sites which do not comply with age requirements – and may also contain more extreme and harmful content."⁷

Ofcom's own research demonstrates that consumers are significantly less likely to agree to online AV when accessing pornographic content than when participating in other activities restricted to adults.⁸ In fact, participants "said they would likely go to other sites to access pornographic content if asked to verify their age"⁹ – specifically, "a site where AV is not required."¹⁰

⁴ euCONSENT (May 2022) [Pilot Execution Report – third euCONSENT](#), page 11

⁵ Samantha Cole (May 2023) Vice, [Accessing Porn In Utah Is Now a Complicated Process That Requires a Picture of Your Face](#)

⁶ euCONSENT (May 2022) [Pilot Execution Report – third euCONSENT](#), page 49

⁷ eSafety Commissioner (March 2023) [Roadmap for age verification and complementary measures to prevent and mitigate harms to children from online pornography](#)

⁸ Ofcom and Yonder Consulting (2022) [Adult Users' Attitudes to Age Verification on Adult Sites](#), page 5

⁹ *ibid*, page 6

¹⁰ *ibid*, page 11

This data has been borne out by the experiences of adult content service providers that have implemented age assurance procedures. In May of 2023, Pornhub shared that traffic to its website from the state of Louisiana “dropped 80 percent since the introduction of its age-verification law.”¹¹ Another platform reported that only one quarter of users presented with the age assurance check “even clicked the link to verify their age and only 9 percent of those users completed the process” – just 2% of all users.¹²

This is rational behaviour on the part of users. The mere transmission of age assurance data exposes them to the risk of data breaches, extortion, and identity theft. According to the government's data, we live in an online world where “79% of businesses have faced phishing attacks” in the last 12 months and “information and communications businesses, and professional, scientific and technical businesses are more likely than average” to have been targeted.¹³

And it’s not just the age assurance providers who criminals are targeting. Phishing websites that appear to be legitimate age assurance vendors collect government identification, biometrics, and even webcam video footage from victims. The information is then sold on the dark web, used by hackers to commit fraud, or leveraged in “sextortion” schemes. People desperately hoping to avoid having their browsing history, identities, and intimate images shared are an attractive and lucrative target. And the danger continues to grow – researchers “detected a 178% increase in sextortion emails between the first half of 2022 and the same period” in 2023.¹⁴

These systems are obviously *not* easy to use and appear to be overwhelmingly untrusted by users. These are extremely serious problems that service providers have no way to solve.

Question 7: Do you have comments on the illustrative case study we have set out in the guidance? Do you have any supporting information or evidence relating to additional examples of how the criteria and principles might apply to different age assurance processes?

In step 3 of the case study, the explanatory text makes clear that even if a service provider is utilising a third-party age assurance vendor, they are responsible for testing and evaluating the performance of that vendor’s tool:

¹¹ Anna Iovine (May 2023) Mashable, [Pornhub blocks Utah because of age verification law](#)

¹² Makena Kelly (August 2023) The Verge, [Child safety bills are reshaping the internet for everyone](#)

¹³ Department for Science, Innovation & Technology (19 April 2023) [Official Statistics: Cyber Security Breaches Survey 2023](#), Chapter 4

¹⁴ Infosecurity Magazine (24 August 2023) [Sextortion Scams Surge 178% in a Year](#)

- “The service provider assesses the results of performance testing of the age estimation method and determines that the level of technical accuracy is not high enough.”
- “The service provider carries out real-world testing to ensure that the estimation method performs to a suitable level in practice.”
- “The service provider ensured that during development of the solution, steps were taken to train the model on a diverse dataset.”

Beyond the impracticality of expecting businesses and individuals with no expertise in this scientific and highly technical research to carry it out, Ofcom does not put forward any standard by which to assess whether the “technical accuracy is not high enough” or whether “the estimation method performs to a suitable level in practice.” And they certainly have no control over how diverse the dataset used to train a vendor’s model is.

This guidance places 100% of the liability on the party least capable of mitigating it and exceeding the actual requirements of the law, which obligate regulated providers to ensure that children aren’t normally able to encounter sexually explicit content by using age verification or estimation that is highly effective. It does not say that providers must be the ones to verify and take responsibility for the technical accuracy, robustness, reliability, and fairness of tools that they did not create.

Question 8: Do you agree with our proposed guidance on the record-keeping duties? Please provide any information or evidence in support of your views.

It would be helpful if this section linked to a template demonstrating compliance with Part 5 of the Act (similar to [the one provided for the DPIA](#)).

Question 9: Do you have any comments on our proposed approach to assessing compliance with the duties on service providers who publish or display pornographic content, including on the proposed examples of non-compliance? Please provide any information or evidence in support of your views.

I appreciate that Ofcom acknowledges that service providers – especially small or inexperienced ones – “may need time to understand the new regime, assess the risks their services pose to users

and make the necessary adaptations to their systems and processes” and that those challenges will be taken into account.

It’s unclear to me what is meant by, “This will be balanced against the need to take swift action against serious breaches...” What type of breaches are being referred to here?

Further and more detailed examples of circumstances in which Ofcom is likely to consider a provider has not complied would also be helpful.

Question 10: Do you have any comments on the impact assessment set out in Annex 1? Please provide any information or evidence in support of your views

The Impact Assessment grossly underestimates the burden this guidance will create on service providers, especially small and micro businesses and adult users.

Ofcom considers the risk of negatively impacting service providers' revenue low because the guidance instructs providers to use multiple equitable and easy-to-use age assurance methods that adults trust enough to share their most private information with. But, all evidence indicates that those methods do not exist (as discussed in our response to question 6).

Additionally, in section A1.13, Ofcom asserts: “We do not consider that our proposed guidance will materially increase the costs to adult users.” But age assurance is not free, nor is the requirement for its use limited to paying customers.

To create a very rough estimate of the costs a service provider will incur to perform age assurance on every unique UK visitor, I compiled a list of 30 service providers with estimated annual revenues ranging from £5M to £150M, calculated their monthly unique visitors for 2023 using semrush.com, and grouped them into three size-based segments. I also assigned them an estimated monthly revenue (which will, of course, vary widely on an individual basis) derived from my experience working for a variety of adult content platforms.

According to an impact assessment conducted by the UK Department for Digital, Culture, Media and Sport, the cost to check the age of one user ranges “from less than 1p to more than £1”¹⁵ I estimated the cost per unique visitor for each segment based on the discounts available at scale and the assumption that a service provider would follow Ofcom’s advice to integrate multiple age assurance vendors.

¹⁵ UK Department for Digital, Culture, Media and Sport (January 2022) [Online Safety Bill: Impact Assessment](#), page 44

Segment	Count	Monthly Unique Visitors	Expected Cost per Verification	Estimated Monthly Revenue	Average Monthly Cost
Small	9	< 100,000	£0.50	£250k	£25,320
Medium	14	100k – 600k	£0.25	£1M	£60,818
Large	7	1M – 50M	£0.10	£25M	£1,805,700

Using this very rough calculation, I would expect that performing age assurance checks will cost a provider between 6% and 10% of their total revenue, with smaller providers paying a higher percentage of earnings.

Most age assurance vendors also charge implementation or set-up fees to integrate their solutions with a provider’s website. These fees can range from hundreds to thousands of dollars. In addition, there are substantial costs to integrating a new technology on the provider’s side. Based on conversations I’ve had with FSC members, the integration takes between 1 and 3 days of software engineering work to complete.

According to the job posting website Indeed, the average base salary for a software engineer in the UK is £49,516.¹⁶ If it takes a single staff engineer 2 days to integrate one age assurance provider, the cost to the provider will be approximately £272 (which doesn’t account for opportunity costs). Hiring a software engineer to perform this task could easily cost over £1,000 for small providers who cannot employ them on staff. These costs obviously multiply if service providers follow Ofcom’s guidance to integrate several age assurance options.

Services can't absorb all of these expenses without materially increasing the costs to adult users. Especially providers who are small and micro businesses. Ofcom’s assertion that “some costs could be somewhat lower” for these businesses misunderstands the market. They are charged the highest price per verification because they represent the least value to age assurance vendors. They pay more to integrate age assurance technologies because they lack the staff resources to do it themselves.

For these reasons, the assessment that this guidance will not unduly affect competition is wholly incorrect. Compliance will be so costly that, in direct contradiction with Ofcom’s duty to promote investment and innovation, small and independent providers will be driven out of

¹⁶ Indeed (updated 15 February 2024) [Software engineer salary in United Kingdom](#)

businesses, entrepreneurs and startups will be unable to enter the market, and well-resourced corporations may become further entrenched if they have the means to comply.

The real winners, from what we've seen in the United States, are almost certain to be websites that disregard the laws, are unreachable by the authorities, and present the greatest danger to users. That is, pirate and dark web sites that also have no incentive to police illegal content such as child sex abuse material (CSAM).