Protecting people from illegal harms online

Annex 5: Assumptions on costs and further analysis on costs and benefits

Annexe

Published 16 December 2024

Contents

A4 Assumptions on costs and further analysis on costs and benefits

A4.1 This annex provides further analysis which has been used to support our conclusions for some of the measures included in our Illegal Content Codes of Practice ("Codes"). We outline:

a) the assumptions we have used to develop quantified cost estimates across a number of the measures. Our analysis of the stakeholder responses to the November 2023 Illegal Harms Consultation ('November 2023 Consultation') related to those assumptions underpinning the costs estimates;
b) the additional analysis of costs and benefits undertaken to support our conclusions about hash-matching measure for child sexual abuse material (CSAM) (discussed in Volume 2, chapter 4: Automated content moderation); and
c) the additional analysis of costs and benefits undertaken to support our conclusions about safety defaults measure for child users (discussed in Volume 2, chapter 8: U2U settings, functionalities, and user support).

Assumptions on costs and stakeholder responses

A4.2 This section describes some common assumptions we have made when estimating the costs of the measures recommended in the Codes, such as the labour costs that we use. When we estimate the costs of particular Codes measures, we combine these common assumptions with other assumptions that are specific to each measure. The assumptions that are specific for each measure are described in the costs section of the relevant chapters of the Statement.

Price level

A4.3 All quantified estimates of costs or benefits are provided in 2023 prices, unless otherwise stated. We have used 2023 prices as that is the year of the most recent Annual Survey of Hours and Earnings ('ASHE'), which we use to develop estimates for the labour cost required to implement some code measures.[1]

A4.4 When source data is not directly available in 2023 prices, we adjust using the UK GDP deflator.[2]

[1] Office for National Statistics ("ONS"), 2023. *Annual Survey of Hours and Earnings ('ASHE'),* Table 14, 2023 provisional estimates. [accessed 04 September 2024]
[2] Using the GDP deflator is the approach recommended in HM Treasury's Green Book. The GDP deflator data from the ONS is available here. [accessed 04 September 2024]

Labour costs

A4.5 To develop estimates for labour costs, we have used ASHE 2023 gross median full-time earnings for the three occupations listed below.[3] These are the occupations we judge services providers are likely to employ to implement the recommended measures.

a) Programme and software development professionals (Standard Occupational Classification 2020 reference (SOC) 2134). We use the salary from ASHE to estimate the cost of 'software engineer' time used when developing our cost estimates.
b) Database administrators and web content technicians (SOC 3133).[4] We use the salary from ASHE to estimate the cost of 'content moderator' time used when developing our cost estimates.
c) Professional Occupations (SOC 2). We use the estimate from ASHE to cover a range of professions that are employed at various online services and might be required to implement code measures. This could be legal advisers, operations, product managers and so forth.

A4.6 ASHE data provides median salary estimates based on data submitted by a wide range of employers in the UK and for occupations at different levels of seniority. Our salary estimates are intended to be averages reflecting the wide range of service providers, the types of occupations involved in implementing our recommended measures, and the different levels of seniority within these occupations.

A4.7 For some service providers, median UK wage rates may be too low. This may be especially the case for larger services based in the US. The costs of hiring some types of staff, such as software engineers with certain specialisms, may be considerably higher in some cases. To take account of this, we refer to the median salaries reported in ASHE data as the 'low' estimate. We have used a simple assumption of doubling this for our 'high' estimate.

A4.8 Some service providers may outsource some work to locations where average pay is lower than the median UK wage rates, which may reduce these costs. To the extent this is the case, our 'low' salary estimates may overstate costs.

A4.9 We assume **a 22% uplift** to the gross wage costs to account for non-wage labour costs, such as employers' National Insurance contributions.[5]

A4.10 Table A5.1 shows the resulting 'low' and 'high' labour cost estimates for different time periods, including the 22% uplift.[6]

---

[3] While ASHE documentation does not explicitly state that gross earnings include bonuses, our understanding is that the gross earnings include bonuses, tips and other payments.

[4] This four-digit SOC 2020 code (unit group code 3133) includes occupations such as content, chat, web, and website moderators as well as other occupations such as database administrators and web content technicians. ONS, SOC 2020 Volume 2: the coding index and coding rules and conventions. The associated ONS spreadsheet can be found here: SOC 2020 Volume 2: the coding index. [all accessed 04 September 2024]

[5] The 2024 Budget increased the employer national insurance contribution rate from 13.8% to 15%, which might argue for increasing the 22% uplift fractionally. However, we have continued to use 22% as this is the non-wage uplift recommended by the Regulatory Policy Committee ("RPC"). Source: RPC, 2019. *RPC guidance note on 'implementation costs'*. [accessed 04 September 2024] It is also the uplift used by DSIT, 2024. Online Safety Act: Enactment impact assessment [accessed 19 November 2024].

[6] When producing cost estimates for our measures, we have used wage estimates based on different time periods (e.g. hour/day/week/month/year), depending on what is appropriate for the particular measure. The annual wages are derived from the Office for National Statistics ("ONS"), 2023. *Annual Survey of Hours and*

Table A5.1: Low and High Estimates of Labour Costs[7]

Occupation	Low	High
	Annual labour cost estimates	
Software Engineer	£60,305	£120,609
Content Moderator	£38,430	£76,860
Professional Occupations	£52,693	£105,386
	Monthly labour cost estimates	
Software Engineer	£5,025	£10,051
Content Moderator	£3,203	£6,405
Professional Occupations	£4,391	£8,782
	Weekly labour cost estimates	
Software Engineer	£1,177	£2,354
Content Moderator	£736	£1,472
Professional Occupations	£1,047	£2,093
	Daily labour cost estimates	
Software Engineer	£265	£530
Content Moderator	£169	£338
Professional Occupations	£231	£463
	Hourly labour cost estimates	
Software Engineer	£31	£63
Content Moderator	£20	£39
Professional Occupations	£29	£57

Earnings ('ASHE'), Table 14.7a Gross annual pay for full-time employees, 2023 provisional estimates. The monthly wages are derived from dividing the annual wages by the number of months in a year (12). The weekly wages are derived from the Office for National Statistics ("ONS"), 2023. *Annual Survey of Hours and Earnings ('ASHE'),* Table 14.1a Gross weekly pay for full-time employees, 2023 provisional estimates. The daily wages are derived from dividing the annual wages by the number of working days in a year. We assume on average there are 228 working days in a year. This assumes people work 5 days a week and that there are 8 bank holidays and on average people take an additional 25 days leave a year. The hourly wages are derived from the Office for National Statistics ("ONS"), 2023. *Annual Survey of Hours and Earnings ('ASHE'),* Table 14.5a Gross hourly pay for full-time employees, 2023 provisional estimates. [all accessed 04 September 2024]

[7] While the 'high' estimates have been derived from doubling the 'low' estimate, they may not match exactly due to rounding.

A4.11 For a few measures that require input from senior management, we have used salary estimates for additional occupations.[8] These include senior managers or senior leaders with an estimated annual labour cost of £122,000 to £183,000.[9] These also include S&P 500 Independent Director with an estimated annual labour cost of £310,000 per year.[10]

Maintenance costs for system changes

A4.12 Where system or other software changes associated with a measure involve an initial cost, we assume there is also an ongoing annual maintenance cost of 25% of the initial cost.[11] These ongoing costs reflect likely work required to ensure the system continues to operate as intended. We apply this assumption unless we have more specific information about the ongoing maintenance costs related to a specific measure.

Stakeholder feedback and our response

A4.13 In response to our November 2023 Consultation, we received feedback from stakeholders on our cost assumptions. We summarise this feedback and our responses below.

Broad cost estimates and underestimating costs

A4.14 Some stakeholders argued that our cost estimates are generalised and broad.[12] Google additionally highlighted that Ofcom's approach to estimating costs would not be able to generate "reliable cost assumptions in this way, as opposed to making cost estimates on a more individual basis, factoring in the nature of the relevant service provider, its processes, and its resources." Some stakeholders argued that our labour costs, including the high wage estimate, and maintenance costs are underestimated.[13]

A4.15 As explained in paragraphs A5.6 to A5.8, our costs estimates are broad to reflect the wide range of service providers and their characteristics, such as size and existing systems. Where feasible, we have explained the factors that may influence costs of implementing a measure in the costs section relating to specific measures in the relevant chapters. Given the broad scope of the regime, we consider our range of estimates are appropriate, and in some circumstances may overstate the costs (see paragraph A5.8).

---

[8] See Volume 1, chapter 5: Governance and Accountability; Volume 2, chapter 2: Content moderation; Volume 2, chapter 3: Search moderation.

[9] This is based on simple assumptions we have made of £100,000 salary for a senior manager and £150,000 salary for a senior leader, which are then uplifted by the 22% uplift for non-wage labour costs.

[10] The annual compensation for an S&P Independent Director is based on Spencer Stuart, 2023 S&P 500 Compensation Snapshot, converted to GBP, with additional assumption that such persons spend approximately 250 hours per year in their role based on PwC's 2022 Annual Corporate Directors Survey. [all accessed 04 September 2024]. The resulting annual wage is £254,000 to which 22% uplift is applied for non-wage labour costs.

[11] The estimate draws on our own expertise, stakeholder engagement, and engagement with external technical and research specialists.

[12] Google response to November 2023 Consultation, pp.33-34; Meta response to November 2023 Consultation, p.21; [✂]; Reddit response to November 2023 Consultation, p.9. We note that [✂].

[13] [✂]; Meta response to November 2023 Consultation, p.21; [✂]; Name Withheld 3 response to November 2023 Consultation, p.9; Skyscanner response to November 2023 Consultation, p.16.

Further analysis on CSAM hash matching measure

A4.16 In this section, we present our analysis undertaken to quantify the costs and benefits associated with the CSAM hash matching measure, as set out in Volume 2, Chapter 4: Automated content moderation (for U2U services). The costs and benefits of the measure will vary by service due to many factors. However, for the purpose of our analysis we assume that costs and benefits vary based on two factors:

- a service's number of users; and
- a service's risk of image-based CSAM.

A4.17 In our November 2023 Consultation, we recommended that the CSAM hash matching measure should apply to the following services: large services which are at medium or high risk of image-based CSAM; services which are at high risk of image-based CSAM and have more than 700,000 monthly UK users; and services which are at high risk of image-based CSAM and are file-storage and file-sharing services that have more than 70,000 monthly UK users.

A4.18 Although stakeholders expressed general support for the scope of services captured by our measure, there were some that suggested the scope of the measure should be widened, while others suggested it had been applied too broadly.

A4.19 In light of this feedback, we reviewed our evidence base and decided to apply the measure to all file-storage and file-sharing services which are at high risk of image-based CSAM (removing the user threshold proposed in the November 2023 Consultation). Overall, therefore, we have decided the measure should apply to:

i) large services which are at medium or high risk of image-based CSAM;
ii) services which are at high risk of image-based CSAM and have more than 700,000 monthly active United Kingdom users ("monthly UK users")[14]; and
iii) services which are at high risk of image-based CSAM and are file-storage and file-sharing services.

A4.20 We first present our analysis, slightly updated based on feedback received from responses to the November 2023 Consultation, for three hypothetical services that would be in scope of the measure, defined in terms of their number of users and risk of image-based CSAM:

i) a medium-risk service that reaches 7 million monthly UK users;
ii) a high-risk service that reaches 700,000 monthly UK users; and
iii) a file-storage and file-sharing service that reaches 70,000 monthly UK users.

A4.21 These hypothetical services represent: the smallest service with a medium risk for CSAM that we will be recommending the measure to (7 million monthly UK users); the smallest service (700,000 monthly UK users) with a high risk for CSAM that we will be recommending the measure to (excluding file-storage and file-sharing services); and a much smaller file-storage and file-sharing service (70,000 monthly UK users) with a high risk for image-based CSAM. However, we recognise that there will be services in scope of the measure that will have fewer than 70,000 monthly UK users and services that will have more than 7 million monthly UK users.

[14] As calculated in accordance with the methodology set out in the Codes of Practice. See the Our approach to developing Codes measures chapter for more information.

A4.22 We then go on to consider the impact of the changes made to the scope of the measure, on our costs and benefits analysis.

A4.23 For our analysis, we have drawn on market prices, industry experts, stakeholder feedback, and our own expertise. Our quantitative analysis of benefits is limited by a lack of evidence on the monetary value associated with removing CSAM from internet services. Monetised estimates are only available for the social cost of contact child sexual abuse (CSA), so we only estimate the benefit of removing CSAM from a service in as much as this results in a reduction in contact CSA. Due to data limitations, we do not quantify any other benefits associated with removing CSAM from a service, such as reduced re-victimisation and fewer people inadvertently viewing CSAM. This means our estimates will underestimate the benefits of the measure. Even when limiting our quantitative analysis to focus on the reduction in contact CSA, we find that this sole benefit could outweigh the direct costs of the measure for many services.

Estimating costs

A4.24 We have estimated the following costs incurred by a service provider that implements the measure:

a) one-off cost of building a hash-matching system ('build cost');
b) ongoing cost of maintaining a hash-matching system ('maintenance cost');
c) ongoing cost of software, hardware, and data ('tech cost'); and
d) ongoing cost of reviewing matches, moderating content, and reporting CSAM ('moderation cost').

A4.25 We understand that these costs are the most material costs associated with a service provider implementing the measure.[15] This is also consistent with the stakeholder feedback we received from respondents to our November 2023 Consultation.[16] Below, we explain how we have estimated these four costs. We recognise that there may be other costs to the service provider as well as to other organisations and individuals.

A4.26 To calculate costs, we have made a series of assumptions, which means that our estimates are imprecise. To reflect this uncertainty, we have calculated costs in terms of a range. We would expect the costs for most services to lie within the range, but we recognise that there will be exceptions to this.

The one-off build cost

A4.27 We estimate the one-off build cost to range from £17,000 to £339,000.[17] This assumes that the system takes around two to 18 months of full-time work for a software engineer to build. We have also included equivalent time for other professional occupation staff such as product managers, analysts, and lawyers.[18] The lower bound assumes service providers pay

[15] Based on our own expertise and stakeholder engagement leading up to the November 2023 Consultation.

[16] Name withheld 5 response to November 2023 Consultation, p.11; INVIVIA response to November 2023 Consultation, p.14; and UK Interactive Entertainment (UKie) response to November 2023 Consultation, p.20.

[17] This is the cost of building the hash-matching system, including integrating the system into the service. We also acknowledge there may be other activities involved for service providers. For example, one stakeholder, [✂], mentioned that they build [✂]. [✂].

[18] The time range is based on our own expertise and engagement with industry experts and stakeholders. This time range is conservative relative to the assumption made by the EU, that it takes 120 hours per year to build a system to detect known CSAM (see p.220 of the EU's impact assessment). Source: European Commission,

median wages (see Table A5.1 above) and build their system in around two months; this amount of time is consistent with the use of a third-party API-based solution.[19] The upper bound assumes service providers pay double the median wage and take the equivalent of 18 months for engineers and other professionals to build the system. This is consistent with stakeholder feedback we received from [✂].[20]

The annual maintenance cost

A4.28 We estimate the annual maintenance cost to range from £4,000 to £85,000. Maintaining the system involves activities such as applying updates, reviewing the technology's performance and adjusting parameters (at least every six months), ingesting new hash lists, and integrating with new functionalities. As explained in paragraph A5.12 above, we assume that annual maintenance costs are 25% of the initial cost required to build the system, which could include input from engineers and other professionals.

The annual tech cost

A4.29 We assume that the annual tech cost will vary by the size of service. In response to our November 2023 Consultation, a few stakeholders (INVIVIA, Microsoft, and [✂]) highlighted that there can be costs for accessing hash matching databases and services from third-parties.[21] Protection Group International commented that our cost estimates "do not appear to show 'membership' fees to organisations that can supply the hash data".[22] We did take such costs into account in our estimate of the annual tech cost. Our estimate draws on our own expertise, engagement with industry experts, and publicly available data from third parties such as fees required to access hash lists and prices of all-in-one software solutions.[23] In light of further information since the November 2023 Consultation, we have also updated our estimate for the minimum annual tech cost likely to be incurred by services. We have assumed a minimum annual cost of £1,000[24] and a theoretical maximum annual cost of £1 million for a service that reaches the entire UK population. We are aware that there may be other low cost or free options available in the market. We have assumed that this cost will scale linearly with a service's user base. We expect that larger services will pay more for software, hardware, and data, as they are more likely to: build in-house hash-matching systems that ingest multiple hash lists and involve bespoke software solutions; pay more to access hash lists and solutions from third-party organisations that tend to base their price on a service's capacity to pay for the product; and overall have more content on their service to be hashed.[25]

2022. Proposal for a regulation laying down the rules to prevent and combat child sexual abuse: Impact assessment. [accessed 17 September 2024].

[19] Safety technology organisations offer API-based solutions which automatically check whether material on a service matches with any material on a CSAM hash list.

[20] [✂].

[21] INVIVIA response to November 2023 Consultation, p.14; Microsoft response to November 2023 Consultation, p.13; and [✂].

[22] Protection Group International response to November 2023 Consultation, p.5.

[23] IWF publish their membership fees, and AWS publish price points for Thorn's Safer.

[24] In the November 2023 Consultation, we were aware that IWF's membership fees, which vary based on industry sector and company size, can start at £1,000 per year. However, we are now aware that these fees would include access to both IWF's hash list, as well as hash-matching technology such as Microsoft's PhotoDNA. IWF, 2024. Membership fees. [accessed 4 November 2024]; and IWF, 2024. Image hash list. [accessed 4 November 2024].

[25] For example, IWF and Thorn charge services based on factors such as industry sector, company size, number of API queries, etc.

The annual moderation costs

A4.30 We assume that the annual moderation cost will depend on the number of matches that are detected on the service by the hash-matching technology. We estimate this number based on the historical number of reports made to NCMEC by service providers that already use hash-matching technology. To allow us to generalise this number to services of any size, we assume that there is a relationship between a service's number of reports and number of users.

A4.31 Taking as our example a service that is likely to identify as high risk for image-based CSAM and that already does hash matching, we calculate that this service made 0.0004 reports per user on average between 2022 and 2023.[26] This ratio allows us to estimate the expected number of reports for similarly risky services of different sizes. For example, a service that has a reach of 700,000 UK users and makes 0.0004 reports per user would be expected to make 311 UK-based reports per year.[27] We use the example of another internet service, [✂], to benchmark the cost of dealing with this much CSAM, including with false positives and content that does not relate to the UK. The service provider told us that they employ [✂] FTE content moderators just for CSAM.[28,29] In the same year, the service made [✂] reports of CSAM that pertained to UK offenders or victims, implying that they employed around one moderator per 500 such reports (or 0.002 moderators per such report).[30] Applying this ratio, we would expect a high-risk service that makes, for example, 311 UK-based reports, to employ the full-time equivalent of 0.7 moderators. Together with our assumptions about the salary of a content moderator, this implies that a high-risk service that reaches 700,000 monthly UK users would be expected to spend £18,000 to £55,000 on labour to review, moderate, and report CSAM per year. This calculation can be generalised to estimate the moderation costs for services with any number of users and different levels of risk.[31]

A4.32 In response to our November 2023 Consultation, a few stakeholders ([✂], Microsoft, [✂])[32] highlighted the importance of human review in determining whether the content detected through hash matching is CSAM. As part of our measure, we do recommend that human moderators should be used to review an appropriate proportion of content detected as CSAM.[33]

---

[26] This is based on the number of reports made by the service with CSAM pertaining to UK offenders or victims, relative to the number of UK users. The National Crime Agency (NCA) provided us with data on reports passed to them from NCMEC. CSAM reports sourced from: [✂]. Adult user numbers data sourced from: Ipsos, *Ipsos iris online audience measurement service*, January 2023, age: 18+, UK.

[27] We note that a report may include multiple images or videos.

[28] Ofcom / [✂] meeting, 2 November 2022.

[29] This is also consistent with stakeholder feedback we received from [✂] as part of our further engagement with them. They indicated that [✂] services would typically employ [✂] FTE content moderators for CSAM reviews, but that in reality this would vary according to the service provider's size and other characteristics. ☯].

[30] We add or subtract 20% from the central FTE estimate in order to account for the likely variation in the number of content moderators needed across services. This means that we obtain two different estimates for our FTE moderators per report. However, both estimates approximately round to 0.002.

[31] We have modelled services with different risk levels by using the ratio of reports to users from three different internet services that could be considered reflective of high-risk, medium-risk, and high-risk file-storage and file-sharing services.

[32] Name witheld 5 response to November 2023 Consultation, p.10; Microsoft response to November 2023 Consultation, p.13; and [✂].

[33] We are also aware that some service providers review all CSAM hash matches. [✂]; and [✂].

A4.33 A few stakeholders (Integrity Institute, INVIVIA, [✂]) also commented that the costs associated with reviewing, moderating, and reporting CSAM could be particularly onerous for small and medium sized service providers. For example, Integrity Institute commented that "smaller services will struggle when dealing with reporting and legal aspects".[34] INVIVIA commented that the human review required to manage false positives and negatives can "strain limited resources".[35] Another stakeholder, [✂], also raised concerns that [✂].[36] We sent follow-up questions to these respondents, as well as to other small service providers already known to undertake hash matching, to better understand the costs associated with human moderation.[37] Two stakeholders ([✂] and [✂]) submitted responses to our follow-up questions. Both indicated difficulties in specifically quantifying the proportion of moderation costs dedicated towards CSAM hash matching. We found that their estimated annual costs of human moderation were broadly consistent with our own cost estimates, once adjusting our model to more closely reflect the number of users and assumed risk level of these stakeholders (our modelling is based on these two factors, see paragraph A5.16).[38] Therefore, we remain of the view that our annual moderation cost assumptions and estimates are reasonable.[39]

A4.34 We generally expect the costs of human moderation to scale with the amount of content on a service and to be higher for service providers that have more known CSAM.[40] This is based on information from services. For example, Match Group indicated that across the services it operates, which are services it considers as relatively low risk for image-based CSAM, hash matches for CSAM are negligible and they allocate staffing levels for content moderation proportionate to this level of risk.[41] Therefore, we expect that the costs of human moderation, including those incurred by small service providers, will correlate with the extent of harm (CSAM) that exists on a service, and therefore are likely to be proportional to the benefits of adopting our measure.

A4.35 Using the assumptions set out above, we can calculate the overall one-off and ongoing costs for the services outlined in paragraph A5.20. Table A5.2 sets out these cost estimates. We have calculated lower and upper estimates for costs as bounds to reflect the uncertainty in our assumptions. We have modelled costs for services with a low-cost profile and large user base, as well as for services with a high-cost profile and small user base, even though we consider both scenarios unlikely.

---

[34] Integrity Institute response to November 2023 Consultation, p.12.

[35] INVIVIA response to November 2023 Consultation, p.14.

[36] [✂].

[37] We sent follow-up questions via email to six stakeholders.

[38] [✂]; and [✂].

[39] As mentioned above, we have generally sought to quantify the most material costs involved for service providers. However, we recognise that there may be other costs to the service provider in this area. For example, two stakeholders, Name withheld 5 and [✂] indicated that there are costs associated with supporting and maintaining the wellbeing of human content moderators. Another stakeholder, [✂], moderators may also need to spend time on other activities such as developing documentations and updating policies, for which there is a cost. Name withheld 5 response to November 2023 Consultation, p.11; [✂]; and [✂]. [40] We are also aware that the costs of human moderation may also be higher for services with more false positives. However, we expect that service providers will be able to account for such false positives in their configuration of their hash-matching technology.

[41] Ofcom / Match Group meeting, 15 July 2024, subsequently confirmed by Match Group by email on 5 September 2024.

Table A5.2. Estimated costs of three hypothetical services implementing the measure

	File-storage and file-sharing service with 70,000 monthly UK users	High risk service with 700,000 UK monthly users	Medium-risk service with 7 million monthly UK users
Build costs (one-off)	£17,000 - £339,000	£17,000 - £339,000	£17,000 - £339,000
Ongoing costs[42] (annual)	£8,000 - £91,000	£33,000 - £150,000	£118,000 - £224,000

Source: Ofcom analysis, various sources

A4.36 In response to our November 2023 Consultation, a few stakeholders also commented that some of the one-off and ongoing costs (beyond moderation costs) associated with our measure could be particularly costly for smaller service providers. For example, INVIVIA said that the initial set-up, including integration of hash-matching technology and accessing of hash lists could be "costly and technically demanding for smaller platforms",[43] however they did not provide us with any alternative cost estimates. The Scottish Government commented that "potential challenges around the affordability" of the measure for "small community/voluntary organisations" should be considered.[44] We are aware that some hash database and software solution providers charge different services different prices and that it is not uncommon for some NGOs (e.g. C3P)[45] to work with smaller service providers to overcome barriers to implementing hash-matching, by providing low-cost options. We nonetheless acknowledge that some of the one-off and ongoing costs associated with our measure could represent a more significant proportion of smaller service providers' income.

Estimating benefits

A4.37 Hash matching identifies known CSAM on a service which can then be removed. Removing CSAM has many benefits (see Volume 2, Chapter 4), such as reduced victim re-traumatisation and fewer people inadvertently viewing CSAM.[46] Identifying known CSAM can also prevent contact child sexual abuse. Evidence suggests that people who view CSAM are more likely to go on to commit other sexual offences against children, including contact offences. By reducing the availability of CSAM, hash-matching technology can prevent potential perpetrators from viewing CSAM and going down an abusive pathway.[47] Investigations of users storing known CSAM may discover unknown CSAM stored in the

[42] This consists of maintenance costs, tech costs and moderation costs.
[43] INVIVIA response to November 2023 Consultation, p.14.
[44] Scottish Government response to November 2023 Consultation, p.7.
[45] C3P's response to November 2023 Consultation, p.19.
[46] A survey commissioned by Ofcom found that 3% of adults and 5% of children said they had encountered CSAM online in the previous year. Source: Ofcom, 2020. Internet users' experience of potential online harms: summary of survey. [accessed 18 September 2024].
[47] Protect Children (Insoll T., Ovaska A., Vaaranen-Valkonen N.), 2021. CSAM Users in the Dark Web: Protecting Children Through Prevention. [accessed 18 September 2024].

same or a nearby location. In this way, hash-matching technology leads services to identify CSAM that was not previously known to law enforcement agencies and can help identify children and/or abusers; this means the children can be protected and abusers arrested.[48]

A4.38 Below, we seek to estimate the number of children that could be safeguarded or protected from contact child sexual abuse (referred henceforth as the number of 'safeguards'), if a service implements hash matching. Drawing on existing estimates of the social cost of contact child sexual abuse (CSA), we then value the expected benefit from a reduction in contact CSA. No estimates exist of other social costs of CSAM being present on a service (e.g., the potential costs associated with victim re-traumatisation). Therefore, we do not attempt to quantify the other benefits associated with removing CSAM from a service, although we expect that these benefits will be significant.

A4.39 In paragraph A5.31, we explained how we estimated the expected number of CSAM reports that service providers would make if they implemented hash matching on their service. We now estimate the expected number of safeguards that would result from the reports. What matters for the number of safeguards is the number of *actionable* reports a service provider makes.[49] We can estimate the number of actionable reports a service provider would make if it implements hash matching, based on the historical number of actionable reports made by service providers that already implement hash matching. Taking the same potentially high risk service modelled above as our example, we estimate that this service provider made around 0.0001 actionable reports per user on average between 2022 and 2023.[50] As the Act constrains Ofcom from including measures in Codes of Practice that recommend the use of proactive technology to analyse user-generated content communicated privately, we assume that only 40% of these actionable reports stem from content communicated publicly, which would be hash-matched by a service provider implementing our measure on CSAM hash matching.[51] This implies that, for a similarly high-risk service, we would expect around 0.00005 actionable reports per user, when excluding reports based on private communications.

A4.40 To estimate the additional number of safeguards based on the number of actionable reports, we calculate the frequency with which actionable reports lead to safeguards. Between 2022 and 2023, the NCA received on average [✂] actionable referrals from NCMEC. Over the same time period, UK law enforcement authorities protected or safeguarded on average [✂] children against CSA,[52, 53] of which we assume that 40%

[48] Industry experts told us that this is a common channel through which services identify unknown CSAM.
[49] Reports are actionable based on a set of criteria used by NCMEC analysts to identify referrals where the reporting company has provided sufficient information to evidence that a crime has been committed, and requires review and assessment by law enforcement agencies.
[50] This is based on the number of actionable CSAM reports submitted by the high-risk service modelled to NCMEC, that were subsequently passed on to the NCA between 2022 and 2023, relative to the service's number of UK users. CSAM reports sourced from: [✂]. Adult user numbers sourced from: Ipsos, *Ipsos iris online audience measurement service*, January 2023, age: 18+, UK.
[51] This assumption is based on NCMEC's estimate that "more than half of its CyberTipline reports will vanish" with end-to-end encryption. For these purposes it is assumed that content communicated privately would be the content that is end-to-end encrypted. Source: NCMEC. End-to-end encryption. [accessed 2 December 2024].
[52] NCA provided us with data on the number of reports passed to them from NCMEC, and the number of children safeguarded or protected from CSA in the UK, between 2022 and 2023. Source: [✂].
[53] In response to the November 2023 Consultation, Protection Group International commented that "the figures for actionable content from NCMEC referrals isn't clear to make a full judgement of costs about

depend on actionable reports received as a result of hash matching of content communicated publicly.[54] This implies that around 0.1 children are protected or safeguarded against contact CSA for every actionable report received by the NCA.[55] If there are 0.1 protections/safeguards per actionable report, and if (as in the above example) there are around 0.00005 actionable reports per user, then this implies that a child would be protected/safeguarded as a consequence of hash matching for every 0.000005 users. This rate will vary by a service's risk level.

A4.41 The Home Office estimates the economic and social cost of contact CSA to be £89,000 per victim (in England and Wales in the year ending 31 March 2019). Accounting for inflation, this figure would be £107,731 in 2023 prices.[56] This value allows us to quantify the benefit of reducing contact CSA. For example, if a service has 700,000 monthly UK users and if (continuing with the above example) the service's risk level is such that it would make 0.00005 actionable reports per user, then we would estimate the service to make 32 actionable referrals per year which, assuming 0.1 safeguards per actionable report, we would expect to result in 3 children being safeguarded against contact CSA each year as a result of the service implementing the measure. This is equivalent to an annual benefit of £314,000.[57] This is shown in Table A5.3 below, along with the equivalent estimates of the annual benefit for the other two hypothetical services.

safeguarding". They also commented that our data (on the number of actionable referrals from NCMEC) was from 2022 and did not account for the latest year available. Protection Group International response to November 2023 Consultation, p.5. We note that we have not sought to use the figures, on the number of actionable referrals from NCMEC, to make a judgement about the costs of safeguarding but to instead arrive at a ratio for the number of children safeguarded against contact CSA for every actionable report, as part of our benefit calculations. As part of our post-consultation updates, we have also now incorporated data (on the number of actionable referrals from NCMEC) received by [✂] for 2023.

[54] An industry expert at a [✂], told us that 30-50% of national arrests and safeguards depend on actionable reports received as a result of hash-matching.

[55] However, we are aware that a child safeguarded or protected from CSA in response to an actionable report may have already incurred some harm.

[56] This estimate accounts for costs to the wider economy (e.g. healthcare costs, police costs), as well as costs of the physical and emotional harm caused to the victim. However, this is a conservative estimate of the economic and social cost of contact CSA, as it does not fully account for the long-term mental health impacts on victims and does not reflect the harm where CSA has led to death. Source: Home Office (Radakin, F., Scholes, A., Soloman, K., Thomas-Lacroix, C., Davies, A.), 2021. The economic and social cost of contact child sexual abuse. [accessed 18 September 2024]. The figure is broadly consistent with another study which estimated that the average lifetime cost for victims in the US of non-fatal CSA is $282,734 and $74,691 for female and male victims, respectively. Source: Letourneau E., Brown D., Fang X., Hassan A., Mercy J., 2018. The economic burden of child sexual abuse in the United States, *Child Abuse Neglect,* 79. [accessed 18 September 2024].

[57] We note that this represents a slight increase from the equivalent annual benefit presented in the November 2023 Consultation, driven by an increase in the number of actionable reports made to NCMEC by the modelled service. In contrast, our estimated annual benefit (as depicted in Table A5.3) has fallen for the other two modelled services: a file-storage and file-sharing service that reaches 70,000 monthly UK users; and a medium-risk service that reaches 7,000,000 monthly UK users. This is due to a decrease in the number of total reports and/or actionable reports made by the relevant modelled services to NCMEC in 2023.

Table A5.3. Estimated annual benefit of three hypothetical services implementing the measure

	File-storage and file-sharing service that reaches 70,000 monthly UK users	High risk service that reaches 700,000 monthly UK users	Medium risk service that reaches 7 million monthly UK users
Value of expected reduction in contact CSA in the UK (annual)	£72,000	£314,000	£692,000
Other benefits	See Volume 2, Chapter 4 for a list of the non-monetised benefits.		

Source: Ofcom analysis, various sources

Additional supporting analysis

A4.42 Based on the assumptions set out above, we can estimate the benefits and costs of a service implementing hash matching for CSAM given a service's user base and risk level. Above, we presented the annual estimated benefits and costs for the three hypothetical services. In the table below, we compare costs and benefits after projecting them over a 10-year appraisal period and discounting by 3.5%.[58]

Table A5.4. Estimated costs and benefits of three hypothetical services implementing the measure over ten years

		File-storage and file-sharing service that reaches 70,000 monthly UK users	High-risk service that reaches 700,000 monthly UK users	Medium-risk service that reaches 7 million monthly UK users
Costs (present value)	Low-cost profile	£73,000	£262,000	£883,000
	High-cost profile	£997,000	£1,433,000	£1,971,000
Benefits (present value)		£527,000	£2,826,000	£6,224,000

[58] For this analysis, we have not added financing costs (i.e. the cost of capital) for one-off build costs. This reflects that the one-off costs are comparable in magnitude to annual ongoing costs and so only represent a small fraction of total costs. As such, this approach does not materially change any conclusions. When discounting, we have assumed that one-off costs are incurred in the first year and that ongoing costs and benefits occur in subsequent years.

		File-storage and file-sharing service that reaches 70,000 monthly UK users	High-risk service that reaches 700,000 monthly UK users	Medium-risk service that reaches 7 million monthly UK users
Benefit cost ratio	Low-cost profile	7.2	10.8	7.1*
	High-cost profile	0.5*	2.0	3.2

Source: Ofcom analysis, various sources.
**We have presented the estimated costs for all hypothetical services based on both the low-cost and high-cost profiles, but we consider it very unlikely that a service that reaches 70,000 users would have a high-cost profile and that a service that reaches 7 million users would have a low-cost profile.*

A4.43 The table shows that, even if we only consider the benefit of hash matching from reducing contact CSA, the benefits exceed the direct costs for each of the hypothetical services we have considered, under reasonable assumptions about a service's cost profile.

A4.44 While the costs are not insignificant, the benefits arising from a reduction in harm are very substantial. For example, for a hypothetical high-risk service with 700,000 UK users, we estimate hash matching could lead to 32 actionable reports per year leading to the safeguarding of 3 children per year. In monetary terms, this translates into benefits of £2,826,000 when projected over 10 years. This materially exceeds the estimated costs, even if the service has a high-cost profile. The only scenario presented above in which the monetised benefits do not exceed costs is for file-storage and file-sharing services with 70,000 monthly UK users and a high-cost profile, but we consider it very unlikely that a service this small would have a high-cost profile.

Changes to the scope of the measure

A4.45 As outlined in Volume 2, Chapter 4, we received feedback in response to our November 2023 Consultation about the scope of services that the measure applies to. In light of this feedback, we reviewed the evidence on the risk of image-based CSAM on file-storage and file-sharing services. The evidence (as described in Volume 2, Chapter 4) confirmed that many file-storage and file-sharing services are at exceptionally high risk of image-based CSAM, including services with fewer than 70,000 monthly UK users.

A4.46 In Volume 2, Chapter 4, we discuss our consideration of whether it would be appropriate to broaden the scope of the measure to include more file-storage and file-sharing services. For example, we explain that removing the user threshold would reduce the risk of perpetrators switching to file-storage and file-sharing services that are not captured by the user threshold, and therefore the potential risk of harm relating to CSAM on a small file-storage and file-sharing service increasing sharply in a short space of time. We also explain why, based on feedback from stakeholders to the November 2023 Consultation, we are now less concerned that the hash-matching database ecosystem would not be able to cope with the increased demand. This constraint was a key reason why we provisionally proposed having a

user threshold applied to file-storage and file-sharing services in our November 2023 Consultation.

A4.47 We have considered the costs and benefits for file-storage and file-sharing services with fewer than 70,000 monthly UK users. Overall, the analysis suggests that the benefits of the measure would justify the costs for file-storage and file-sharing services that are at high-risk of image-based CSAM.

i) File-storage and file-sharing services with fewer than 70,000 monthly UK users can have a significantly higher prevalence of CSAM than the file-storage and file-sharing service that was modelled. In the analysis set out above, we have used the total number of actionable CSAM reports made by [✂], a file-storage and file-sharing service that already undertakes hash matching for CSAM, as an indicator for the average prevalence of CSAM on a file-storage and file-sharing service. Relative to this service provider however, we are aware that some smaller file-storage and file-sharing service providers (e.g. [✂]), that do not undertake hash matching for CSAM, have been identified as hosting much higher amounts of CSAM by law enforcement agencies and NGOs than the service we used in our model.[59] We expect the benefits of hash matching to be even greater for such file-storage and file-sharing services that do not already undertake hash matching and are at high risk of CSAM.

ii) We have only quantified one benefit of the measure: a reduction in contact child sexual abuse. There are a range of other important and wide-reaching benefits associated with identifying and removing known CSAM that we have not been able to account for quantitatively, but we expect are likely to be substantial.[60] Nevertheless, our analysis suggests that the one benefit we were able to quantify is likely to far exceed the costs for file-storage and file-sharing services with 70,000 monthly UK users and an average prevalence of CSAM.[61] This implies that the benefits would exceed the costs for many file-storage and file-sharing services with fewer than 70,000 monthly UK users. This is particularly true for file-storage and file-sharing services with a high risk of image-based CSAM, as we expect that the benefits of implementing the measure are likely to be far greater for such services.

A4.48 For the reasons set out above, we consider that our analysis is likely to significantly understate the benefits associated with hash matching for file-storage and file-sharing services that are high risk for image-based CSAM, and that removing the user threshold for these services would be warranted.

---

[59] C3P, 2021. Project Arachnid: Online availability of child sexual abuse material. [accessed 08 November 2024]. The majority of service providers that C3P's Project Arachnid in 2021 issued removal notices to were file-storage and file-sharing services. Out of these file-storage and file-sharing services, we identified service providers (such as [✂]) that were detected for hosting large amounts of CSAM but that would not meet our 70,000 user threshold based on our analysis of data on their UK audience.

[60] As mentioned above, these benefits include the potential for reduced victim re-traumatisation; fewer people inadvertently viewing CSAM; and the discovery of unknown CSAM, leading to the identification of victims and the arrest of perpetrators.

[61] For such services, our analysis suggests that the benefits of doing hash matching could be around seven times greater than the costs (see Table A5.4 above).

Further analysis for applying safety defaults measure

A4.49 We have undertaken analysis to quantify the costs and benefits associated with our safety defaults measure on reducing grooming harm described in Volume 2, chapter 8: U2U settings, functionalities, and user support. We have focused this analysis on three possible options that were considered in the November 2023 Consultation. These were:

1. Apply the measure to all large services where providers have identified a high or medium risk of grooming in their latest illegal content risk assessment.
2. Apply the measure to:
   i) all services where providers have identified a high risk of grooming in their latest illegal content risk assessment AND at least 25,000 child users; and
   ii) all large services where providers have identified a medium risk of grooming in their latest illegal content risk assessment.
3. Apply the measure to:
   i) all services where providers have identified a high risk of grooming in their latest illegal content risk assessment (with no threshold relating to the number of child users); and
   ii) all large services where providers have identified a medium risk of grooming in their latest illegal content risk assessment.

A4.50 In this section, we set out our assessment of the costs and benefits of Options 1 and 2.[62] Our analysis uses conservative estimates of benefits related only to contact CSA resulting from grooming online. Due to data limitation, we are not able to quantify other negative impacts of grooming when it does not result in contact CSA, such as long term impact on the mental health of victims or self-harm. Our analysis may also understate the proportion of children at risk of being groomed online, as grooming or contact CSA occurrences are widely underreported. This means our analysis will underestimate the benefits of the measure.

A4.51 Despite these limitations, our analysis finds that the benefits outweigh the costs significantly for Options 1 and 2. This supports applying the measure to services that are at high risk of grooming even if they have fewer children on them and is consistent with our decision in Volume 2, chapter 8 that Option 3 is proportionate. We did not receive feedback from stakeholder to our November 2023 consultation suggesting that the scope of the measure is reduced to Options 1 or 2.

A4.52 Below we consider available information on the different costs and benefits associated with each of these options. All the options apply the measure consistently for large services (which are categorised as medium risk services). We have focused our more detailed assessment on the costs and benefits of applying the measure to high-risk services.

Estimating the benefits from reducing contact CSA from grooming

A4.53 Online grooming for child sexual abuse is the method of contacting children and developing a relationship, whether through flattery, emotional connection, sexualisation, bribery, blackmail or coercion, for the purposes of conducting child sexual abuse.[63] Typically, the

[62] The assessment was updated to reflect the salary estimates from ASHE 2023 data and similarly, to reflect monetary values of any benefits to 2023 prices.
[63] See Register of Risks chapter titled 'CSEA'

objective of online grooming is the generation of child sexual abuse material (CSAM) and contact child sexual abuse (contact CSA). Contact CSA can occur in person or can involve the perpetrator remotely forcing the victim to sexually abuse other children or to engage in sexual acts, including penetrative acts. Individuals involved are often forced to share imagery of the abuse with the perpetrator as 'first-generation' CSAM.

A4.54 In our assessment below, we have put monetary values on some of the likely benefits from reducing one aspect of grooming harm, that is, reducing contact CSA. We have also made various assumptions that are often based on limited and uncertain information. We account for uncertainties in part by considering ranges for specific inputs and by developing different scenarios. Given these limitations, our approach estimates a lower bound of the benefits to UK society that may result from the safety defaults measure.

A4.55 We estimate this lower bound of the benefits to society of an individual service provider implementing the safety defaults measure by combining (see equation 1):

a) the expected number of contact CSA occurrences resulting from grooming on a service;
b) the expected percentage reduction in such occurrences resulting from the measure; and
c) the social cost of one contact CSA occurrence.

Equation 1: Expected annual benefits from reducing contact CSA for an individual service[64]

*Annual benefits per service (from reducing contact CSA from grooming) = Expected number of contact CSA occurrences resulting from grooming on a service * expected percentage reduction in such occurrences resulting from the measure * social cost of one contact CSA occurrence*

A4.56 We further estimated the expected number of contact CSA occurrences resulting from grooming on a service by considering (see equation 2):

a) the number of UK child users on that service;
b) the percentage of 'at risk' children subject to grooming leading to contact CSA each year[65];
c) the average number of services visited by UK children where there is a risk of grooming over the course of a year

Equation 2: Expected number of grooming occurrences that lead to contact CSA on an individual service

*Expected number of contact CSA occurrences resulting from grooming on a service = Number of UK child users on the individual service * (% of 'at risk' UK children that experience grooming leading to contact CSA/Average number of services visited by UK children annually where there is a risk of grooming)*

[64] Although we are only applying this equation to grooming that leads to contact CSA, a similar equation could be used for different types of grooming (e.g., non-contact CSA, grooming resulting in fatalities). The combined benefits that result from all equations would then provide an estimate of the total benefits from the measures. We have concentrated our quantitative analysis on grooming resulting in contact CSA because we are not able to quantify the benefits resulting from reducing other grooming harms. Despite this, we have qualitatively considered the potential benefits from other types of harm from grooming when considering the proportionality of recommending this measure in Volume 2, chapter 8: U2U settings, functionalities, and user support.

[65] We define 'at risk' children for the purposes of this analysis as those that use online services where grooming can take place.

A4.57 Combining equations 1 and 2, we quantify the annual benefits for an individual service implementing our recommended safety defaults measure. We combine the estimates of the following to do so:

i) the social cost of one contact CSA occurrence
ii) number of UK child users on the individual service
iii) percentage of 'at risk' UK children that experience grooming leading to contact CSA
iv) average number of services visited by UK children in a year where there is a risk of grooming.
v) expected percentage reduction in grooming occurrences resulting from the measure.

Social cost of contact CSA occurrence

A4.58 As is the case for our measure for CSAM hash-matching, estimating and quantifying the economic and social cost of CSEA offences, and the likely benefits of reducing grooming, is challenging. We recognise that the impact of grooming on each victim and survivor is different and trying to put a monetary value on this is a vast simplification. The only estimate we are aware of that can be related to an aspect of grooming is an estimate by the Home Office of the economic and social cost of contact CSA. This is an estimate of £107,731 per victim at 2023 prices. While this estimate accounts for social costs such as those related to healthcare, police, judiciary, and safeguarding, unsurprisingly, the largest component of this estimate is the physical and emotional harm caused to the victim. However, this is a conservative estimate as it does not reflect the harm where CSA has led to death and does not fully account for the long-term mental health impacts on victims.[66]

A4.59 Because we are only aware of a quantified estimate for contact CSA, we have quantified only this aspect of the many harms that can arise from grooming, and hence the benefit of any reduction in this harm. This means our estimate of the benefits from reducing grooming will be substantially understated.

A4.60 The impact of grooming on victims and survivors is complex and the experience of harm and cost to an individual is a personal one that will vary depending upon many different factors including, but not limited to, personal characteristics, coping mechanisms and the type of harm they have been exposed to. It is recognised that victims and survivors who have experienced similar CSEA harms may describe the effect to them differently, from having minimal impact to having a significant impact on them. For some, the potential impact of grooming can be severe and lifelong, regardless of whether it leads to contact sexual abuse offences. In some cases, online grooming can also lead children to self-harm and taking their own life.[67] We have not included the full spectrum of harms that can arise from grooming. This is due to the complexities of estimating both the social cost of different types of harm

[66] The original study estimated the cost to be £89,240, in 2018/19 prices and we have adjusted this to 2023 prices. Home Office (Radakin, F., Scholes, A., Soloman, K., Thomas-Lacroix, C., Davies, A.), 2021. *The economic and social cost of contact child sexual abuse.* [accessed 04 September 2024]. The figure is broadly consistent with another study which estimated that the average lifetime cost for victims in the US of non-fatal CSA is $282,734 and $74,691 for female and male victims, respectively. Letourneau et al., 2018. *The economic burden of child sexual abuse in the United States.* [accessed 04 September 2024]

[67] Example of such deaths have been reported in the press. For example, Carrell, S. 2013. Scotland police investigate 'online blackmail' death of Fife teenager, *The Guardian,* 16 August 2013.; Campbell, J. & Kravarik, J., 2022. A 17-year-old boy died by suicide hours after being scammed. The FBI says it's part of a troubling increase in 'sextortion' cases. *CNN*, 23 May 2022.; Yousif, N., 2022. Amanda Todd: Dutchman sentenced for fatal cyber-stalking, *BBC News,* 15 October 2022. [All accessed 04 September 2024]

that are a result of grooming, and also the number of instances of each type of harm which occur each year.

Number of UK child users on an individual service

A4.61 We consider two options for the size of a service where Option 1 applies the measure to relevant large services, and Option 2 applies it to relevant services with at least 25,000 child users (see paragraph A5.49).

A4.62 Our Codes define a 'large' service as one with more than 7 million monthly UK users. To estimate the benefits from the measure we needed to estimate the number of UK children on the service. We assume that a large service would have at least 1 million UK child users aged 3-17 based on an assumption that around 15% of UK internet service users are aged 3-17.[68]

A4.63 Option 2 would apply the measure to relevant services with more than 25,000 UK child users.

A4.64 To test these options, we estimate the benefits and costs for services who have 1 million UK child users and 25,000 UK child users. While we haven't explicitly modelled the benefits and costs for services with fewer than 25,000 UK child users, our analysis suggests that the benefits significantly outweigh the costs even for smaller services with 25,000 UK child users. Thus, it lends confidence to our conclusion in Volume 2, chapter 8 to adopt Option 3 where the measure is applied to even smaller high-risk services.

Percentage of 'at risk' UK children that suffer contact CSA as a result of grooming annually

A4.65 We define 'at risk' children for the purposes of this analysis as those in the UK who use the internet and visit online services where they may be at risk of grooming. Table A5.5 sets out the data sources, and calculations undertaken to arrive at this estimate.

Table A5.5: Estimated percentage of 'at risk' UK children that suffer contact CSA as a result of online grooming annually

	Estimates
UK population projection for 3-17 year olds in 2023[69]	12.2 million
3-17 year olds using social media and other online services[70]	7.8 million
Number of contact CSA cases in England and Wales in 2018[71]	113,000

[68] We have estimated this number to be 15% using Population projections by age category published in January 2022 by the Office of National Statistics and Ofcom, 2023. *Children and Parents: Media Use and Attitudes, pages 21-31,* on internet usage by age. [All accessed 28 September 2024]

[69] Population projections published in January 2022 by the Office of National Statistics. [accessed 28 September 2024].

[70] We only consider the children using online services of the total UK population of 3-17 year olds. This estimate is derived from Ofcom, 2023. *Children and Parents: Media Use and Attitudes, pages 21-31.* [accessed 28 September 2024]. The report states 23% of 3-4 year olds, 30% of 5-7 year olds, 63% of 8-11 year olds, 93% of 12-15 year olds and 97% of 16-17 year olds use social media apps or sites.

[71] Home Office, Dec 2021. *The economic and social cost of contact child sexual abuse,* Section 3.2 [accessed 28 September 2024].

	Estimates
Number of contact CSA cases adjusted to derive UK wide estimate[72]	126,000
Contact CSA cases with an online element[73]	19%
Number of contact CSA cases with an online element in the UK (19% of 126,000)[74]	25,000
Estimated 'at risk' UK children who may experience contact CSA from being groomed online (25,000 of 7.8 million)	**0.32%**

A4.66 Based on our assumptions and available data sources, we estimate that approximately **0.32%** of the 7.8m children in the UK are likely to experience contact CSA as a result of grooming each year. However, we recognise that this assumption is conservative and is likely to understate the true prevalence of contact CSA resulting from grooming in the UK.

A4.67 While the UK government estimated there were approximately 113,000 cases of contact CSA in 2018 in England and Wales, this estimate was based on a survey of young people and their caregivers, which asked about their experiences of sexual abuse in the year prior to the survey.[75] It is likely that this is underreported, especially for younger children.[76]

A4.68 The proportion of reported contact CSA cases with an online element (19%) may also be underestimated. The reported cases are manually flagged by police officers, and therefore this estimate relies on the officers being aware of the flag and accurately recognising that an online element is present.[77]

---

[72] We increased the number of contact CSA cases in England and Wales (113,000) by the proportion of children across the UK compared to England and Wales. This resulted in an estimate of approximately 126,000 cases of contact CSA per year across the UK.

[73] The percentage of contact CSA cases that have an online element was estimated using data recorded by the police over the period April 2022 to March 2023. DSIT, 2024. Online Safety Act: Enactment impact assessment, paragraph 281 [accessed 28 October 2024]. We assume that all cases with an online element would have involved grooming.

[74] We have rounded this calculation. Our analysis is not sensitive to this rounding.

[75] Home Office, Dec 2021. *The economic and social cost of contact child sexual abuse*, Section 3.1. The estimate in the Home Office study was derived from Radford, L., Corral, S., Bradley, C. and Fisher, H., 2013. The prevalence and impact of child maltreatment and other types of victimization in the UK [all accessed 28 September 2024]. This involved a survey of 6,196 children, young people and caregivers that was conducted in 2009. The victimisation rate from that survey was then applied to the child population in England and Wales in 2018 to estimate the number of occurrences.

[76] For children aged under 11, the survey was answered by caregivers. Therefore, CSA occurrences are unlikely to be included if the caregiver was unaware of the abuse or was the abuser themselves. Source: Home Office, Dec 2021. *The economic and social cost of contact child sexual abuse*, Section 3.1 [accessed 28 September 2024].

[77] The true level of online offences may be higher than 19% due to issues with the flag. First, the flag is typically manually applied by officers, and therefore accurate use relies on officers being aware of the flag, remembering to apply it to specific cases, and recognising that an online element is present. This can be difficult in some cases, such as where online messaging services like WhatsApp or Kik are used, which officers may not recognise as 'internet enabled' or online. Second, the flag differs in usage between forces and is not evenly applied throughout England and Wales. DSIT, 2024. Online Safety Act: Enactment impact assessment, paragraph 282 [accessed 28 October 2024].

Average number of services visited by UK children annually where there is a risk of grooming

A4.69 We identified above, the risk that any individual child is likely to experience contact CSA as a result of grooming each year. An individual child is likely to visit several online services during a year. Each service will only be partially responsible for the cumulative risk of grooming we have estimated per child. Here, we estimate how many services are typically visited by each child, which helps us in calculating the risk of experiencing contact CSA *per service*. For a given risk per child (calculated above as 0.32%), the more services a child visits, the lower the risk per service would be.[78]

A4.70 We used data from Ipsos Iris and our own internal research to estimate the average number of services UK children use each year where they may experience grooming. Ipsos Iris data shows 15-17 year olds visited several thousand [✂] web- and app-based services in December 2022. The data aggregated the visits at the 'organisation' level, that is, where a visit is to services under the same parent company, these were reported at the parent company level. For example, the visits to Google Search, Gmail, YouTube would be reported as a combined figure under Alphabet.

A4.71 We categorised these organisations to have a risk of grooming based on their business category (or service type). This was informed by the evidence in our Register of Risks (chapter titled 'CSEA') where we find an association between specific service types and the risk of grooming harm (and not actual evidence of grooming occurring on a service).[79]

A4.72 From the organisations categorised to potentially have a risk of grooming, we removed organisations where 15-17 year olds spent less than 10 minutes a month. We assumed that these organisations were unlikely to contribute to the child experiencing grooming due to the short time available for any interaction. The analysis suggested 15-17 year olds visit an average of 10 online services a month which pose a risk of grooming.[80]

A4.73 We consider this provides a useful proxy of the number of online services at risk of grooming which a UK child user is likely to visit over a year. It is important to note the following points:

a) The analysis is based on data from surveying fewer than 100 ([✂] 15-17 year olds. There is significant uncertainty in extrapolating results from a survey with a small sample size to be representative of the whole 15-17 population.
b) The analysis only covers 15-17 year olds, which means that it does not provide any information about the number of online services that younger children would typically use. Ofcom research suggests younger children visit fewer online services.[81]
c) The analysis only covers one month. Although we expect that generally child users will visit the same websites from month to month, there may also be some additional services they visit in specific months. This means our assumptions may underestimate the number of sites visited when extending this to an annual estimate.

---

[78] We assume for the purposes of this analysis that there is an equal risk of grooming on each of these identified services where a child may experience grooming.
[79] Ipsos Iris provides classifies all online services into different categories. Using a manual process, we estimated whether these categories were likely to indicate a site at risk of grooming (e.g., 'Social media', 'gaming', etc.) or limited risk of grooming (e.g., 'Weather', 'Estate agents', etc.).
[80] This is an unweighted average based on the raw data of 15-17 year olds' visits to online services in December 2022.
[81] Our Media Use and Attitudes report suggest that "…16-17-year-olds are branching out in media, using a wider and more diverse diet of apps and sites." Source: Ofcom, 2023. *Children and Parents: Media Use and Attitudes, pages 28.* [accessed 28 September 2024]

d) Our analysis only considers the number of 'organisations' which are visited, not individual online services themselves.[82] This means our analysis may underestimate the number of services visited.

Expected reduction in grooming due to implementation of the measure

A4.74 Our expectation is that the suggested measure will be effective at reducing grooming by introducing friction into the grooming process. However, it is not expected to fully eliminate the incidence of grooming because:

a) the measure recommends default settings which can be changed.
b) the measure will only apply to services, where the service provider can determine a user's age or age group.
c) there are likely to be other pathways for perpetrators to identify children and start a grooming process.

A4.75 For these reasons and because of the uncertainty around the degree of efficacy of the measure, for the purposes of this exercise we assume that all of the elements of the safety defaults measure could reduce the extent of contact CSA as a result of grooming by between 5% and 10%.[83]

A4.76 Based on these assumptions, we have developed a 'high' benefits scenario, where the measure is assumed to reduce contact CSA occurrences from grooming on online services by 10% and a 'low' benefits scenario where the measure reduces it by 5%.

Estimating the benefits and costs for each option

A4.77 We have used the assumptions outlined above to estimate the annual benefits from reduced contact CSA from grooming for each of the options considered. We have focused on the impact on high-risk services,[84] where the difference between the options is as follows:

1. Apply the measure to all large services which have a high risk of grooming.
2. Apply the measure to all services which have a high risk of grooming AND at least 25,000 child users.

A4.78 We present the results of our analysis as follows:

a) a summary of the assumptions in the quantitative analysis used to estimate benefits.
b) a summary of the assumptions in the quantitative analysis used to estimate costs.
c) the results of our quantitative analysis assessing a service with 1 million UK child users.
d) the results of our quantitative analysis assessing a service with 25,000 UK child users.

Summary of input assumptions used to estimate benefits

A4.79 Table A5.6 below summarises the inputs we have used to estimate the benefits for a high and low benefit scenario.

---

[82] 'Organisations' are the parent companies of the groups of websites and apps; for instance, Alphabet organisations include Google Search, Gmail and YouTube; Meta includes Facebook and Messenger, Instagram and WhatsApp. We have undertaken our analysis at an organisation level but recognise that this could lead to an underestimate of the number of individual online services visited. However, we also consider that there could be cost synergies when implementing the measure across an organisation, which provides a further reason why we consider it is reasonable to consider this at an organisation level.
[83] We assume services apply all elements of the safety defaults measure outlined in Volume 4, Chapter 8.
[84] As assessed by service providers in their latest illegal content risk assessment.

Table A5.6: Summary of inputs to estimate the annual benefits from reducing contact CSA from grooming on a service

	Low benefit scenario	High benefit scenario
Social cost of one contact CSA occurrence	£107,731	£107,731
'At risk' UK children that suffer from contact CSA as a result of grooming annually	0.32%	0.32%
Average number of services visited by UK children annually where there is a risk of grooming	10	10
Expected reduction in contact CSA occurrences from grooming due to safety defaults measure	5%	10%

Source: Ofcom analysis

Summary of input assumptions used to estimate costs

A4.80 We outline in Volume 2, chapter 8 in the Costs and Risks section the estimated one-off costs of £10,000 to £325,000 for the safety defaults measure.[85]

A4.81 We applied the following three cost assumptions in the analysis here:

a) **Services with 1 million UK child users**: We assumed that costs are at the top of the range (i.e., £325,000).
b) **Services with 25,000 UK child users (low estimate):** We assumed that costs are at the bottom of the range (i.e., £10,000).
c) **Services with 25,000 UK child users (high estimate):** We assumed that costs are £55,000.

A4.82 We have included two cost scenarios for services with 25,000 UK child users because although generally we would expect smaller services to be towards the lower end of the range and larger services to be towards the top, there are also several factors not related to size which would affect the level of costs that might apply to different services.[86]

A4.83 There might be some providers of relatively small services who would have costs higher than the bottom of the range, but for which the top of the range is also unlikely to be appropriate. This is because the upper end of our cost range includes significant overhead and coordination costs[87] that are likely to be much lower for providers of smaller services. Therefore, we have included an additional 'high estimate' for services with 25,000 UK child users where we assume the one-off cost is £55,000.

[85] Note that for the quantitative analysis we do not estimate any indirect costs. We have updated the estimates of direct costs since the November 2023 Consultation in line with the latest wage data released by ONS. However, since our cost estimates are rounded, the estimates may not necessarily have changed when using the updated wage assumptions.
[86] For example, the complexity of their systems.
[87] For more explanation on overhead and coordination costs, please see Volume 2, chapter 8.

A4.84 As well as the one-off costs we also include ongoing costs[88] associated with implementing the measures. The one-off cost will be incurred once as the measure is implemented, while the ongoing costs that arise from implementing the measures will occur annually.

A4.85 To ensure that costs can be considered alongside annual benefits, we estimated an annualised cost by calculating the total cost of the measures over a 10-year period, including the financial costs that we expect to be incurred for one-off investments,[89] and then converted to a flat annuitized rate to estimate the annualised costs of the investment.[90]

A4.86 We do not include any indirect costs within our analysis as we are unable to appropriately quantify them.[91] However, we consider indirect costs (for both service providers and users) will be lower for smaller services, as they will tend to scale with the number of users. As we describe below, the case for applying the measure for large services (Option 1) is very strong, and we are confident it is proportionate to apply to such services even though we have not quantified the indirect costs. When considering the case for recommending for smaller services (as with Options 2 and 3), these indirect costs will be much smaller and so not including them is less important.

A4.87 Table A5.7 outlines the cost assumptions we use in different scenarios including the estimate of costs on an annual basis.

Table A5.7: Estimated cost of safety defaults measure under different scenarios

	Costs per service		
	Low cost (small service) scenario	**High cost (small service) scenario**	**High cost (large service) scenario**
Initial one-off cost	£10,000	£55,000	£325,000
Ongoing annual cost (after initial cost)	£2,500	£13,750	£81,250
Annualised cost	**£4,000**	**£22,000**	**£128,000**

Source: Ofcom analysis

Quantitative assessment of services with 1 million UK child users (Option 1)

A4.88 We estimated the annual benefits for a service with 1 million child users by inserting the assumptions in Table A5.6 into equations 1 and 2. We used this together with the annualised

[88] We assume that annual ongoing costs are equal to 25% of the one-off costs. For an explanation, please see paragraph A5.12.

[89] We have assumed a real pre-tax financial cost of capital of 7% for these companies, but note our findings presented below are not very sensitive to the assumed cost of capital. A WACC of 7% is broadly consistent with an estimate the CMA used for Google and Facebook as part of their 2020 Market Study on online platforms and digital advertising. *2020 Online Platforms and Digital Advertising market study final report, Appendix D*, [accessed 15 September 2024].

[90] We annualise the costs by assuming that the one-off costs are akin to an investment that is paid back at a consistent rate in real terms over 10 years. This payment includes both the financing costs and repayment of the initial investment cost.

[91] For more explanation on indirect costs, please see Volume 2, chapter 8. These costs could include impact on user experience and engagement, how child users make friends online, how adult users interact with children they know online.

cost from Table A5.7 to estimate illustrative benefit cost ratios for a service of this size when applying the measure.

A4.89 The estimated benefit cost ratios for both the high and low benefit scenarios for a service with 1,000,000 UK child users are given in Table A5.8.

Table A5.8: Indicative benefit cost ratios for services with 1,000,000 UK child users

	Annual estimated benefits[92]	Annualised cost (high cost large service)[93]	Indicative benefit cost ratio[94]
10% reduction in contact CSA (high benefits scenario)	£3,447,000	£128,000	27
5% reduction in contact CSA (low benefits scenario)	£1,724,000	£128,000	13

Source: Ofcom analysis

A4.90 The benefit cost ratios in Table A5.12 show benefits to be much greater than costs for large services with above 1 million child users, even though our estimate of benefits only considers the benefits that would arise from a reduction in contact CSA. As there are other benefits beyond what we have been able to account for quantitatively, the total benefits from the measure for a service of this size are likely to be significantly higher.

A4.91 Although the analysis does not specifically cover medium risk services, the benefits for large services with high risk of grooming are so high that this is likely to hold true even if the benefits for medium risk services are significantly lower than for high-risk services.

A4.92 We also assumed that the expected costs were at the top of our estimated range. This conservative approach combined with the high ratio of benefits to costs gives us significant confidence that applying the measures to all large services which have a high risk of grooming is likely to deliver significant benefits and supports the conclusion that Option 1 would be proportionate.

Quantitative assessment of services with 25,000 UK child users (Option 2)

A4.93 We estimated the annual benefits for a service with 25,000 UK child users by inserting the assumptions in Table A5.6 into equations 1 and 2. We used this together with the annualised cost from Table A5.7 to estimate illustrative benefit cost ratios for a service of this size when applying the measure.

A4.94 The estimated benefit cost ratios for the relevant four scenarios are given in Table A5.9.

[92] Figures based on using estimates presented in Table A5.6 and inputting them in equations 1 and 2 with 1,000,000 child users.
[93] Figures based on Table A5.7
[94] Note the values in this column may not precisely match with the presented benefit and costs in the other columns due to rounding.

Table A5.9: Indicative benefit cost ratios for services with 25,000 UK child users

	Annual Estimated benefits[95]	Annualised cost (small service)[96]	Indicative benefit cost ratio[97]
10% reduction in contact CSA (high benefit)/low cost small service scenario	£86,000	£4,000	22
5% reduction in contact CSA (low benefit)/low cost small service scenario	£43,000	£4,000	11
10% reduction in contact CSA (high benefit)/high cost small service scenario	£86,000	£22,000	4
5% reduction in contact CSA (low benefit)/high cost small service scenario	£43,000	£22,000	2

Source: Ofcom analysis

A4.95 Table A5.9 shows that the estimated benefits are greater than costs across all four scenarios we have assessed for services with 25,000 UK child users. We consider these results support the conclusion that, at the very least, it would be proportionate to apply the measure to services with 25,000 UK child users or more.

A4.96 Overall, the quantitative analysis illustrates that the measure is likely to result in benefits that are significantly in excess of the costs for both Options 1 and 2. The quantitative analysis does not explicitly model the benefits and costs of applying the measure under Option 3 where services may have even lower numbers of UK child users. However, the analysis set out in Table A5.9 would suggest that benefits are likely to outweigh costs for significantly fewer child users than 25,000.

A4.97 Not only does the modelling above on its face suggest that it would be proportionate to apply the measure to services with significantly fewer than 25,000 child users, but we consider that the model is conservative in some of its assumptions. Our analysis of benefits is limited to what we have been able to quantify. We have only estimated the social cost of contact CSA offences resulting from grooming, which we acknowledge represents a small proportion of all grooming occurrences. Our estimate of children at risk of grooming may also be underestimated due to underreporting.

[95] Figures based on using estimates presented in Table A5.6 and inputting them in equations 1 and 2 with 25,000 child users.
[96] Figures based on Table A5.7.
[97] Note the values in this column may not precisely match with the presented benefit and costs in the other columns due to rounding.

A4.98 Given these limitations, it is likely that the true benefit would be higher than what we have been able to estimate, and potentially significantly higher than what we have modelled for a service with 25,000 UK child users. This being the case, we conclude that it would be proportionate to target the measure at a far broader range of services than the modelling presented here, would on its face imply.

A4.99 As explained in Volume 2, chapter 8, there are other strong reasons to apply this measure to all high-risk services regardless of size that the quantitative analysis here is unable to account for. For instance, the risk that if we confined ourselves to improving the protections on services of a certain size, perpetrators would move to smaller services. This would in turn compromise our ability to protect children from grooming. Taking into account the assessment of benefits and costs which we have discussed above, and our assessment of the displacement risks associated with exempting small services from this measure, we conclude that the measure is proportionate for all services which have a high risk of grooming (irrespective of size) and all large services which have a medium risk of grooming.