Detecting intimate image abuse

Online Safety – Additional Safety Measures

Statement

Published 18 May 2026

Contents

Chapter

1. Overview....3
2. Introduction....5
3. Measure on hash matching for intimate image abuse....7

Annex

A1. Impact assessments and statutory tests....58
A2. Legal Framework....68
A3. Glossary....72
A4. Further Detail on Economic Assumptions and Analysis....76

1. Overview

1.1 Ofcom is the United Kingdom's communications regulator, overseeing sectors including telecommunications, post, broadcast TV and radio. Under the Online Safety Act 2023 (the Act) Ofcom is also the UK regulator for online safety.

1.2 Ofcom's job is to make online services safer for the people who use them, especially children. Providers of regulated online services (providers) are required to have effective systems in place to protect all users from illegal content.

1.3 This statement sets out our decision to recommend that certain service providers should use 'hash matching' technology to detect intimate image abuse and reduce its spread online.

1.4 The Act created offences of sharing or threatening to share intimate images without consent (intimate image abuse). Intimate image abuse has profound negative psychological, emotional, professional and relational impacts on survivors and victims, compounded by the challenges they face in getting content removed, and the fear it might resurface online. Intimate image abuse is a widespread harm, evidence suggests it is increasing in prevalence, and the proliferation of generative artificial intelligence (GenAI) tools is causing a further increase. Reflecting the gravity of this development, Parliament has decided to make it an offence to make or supply AI tools that can be used to create intimate images.

1.5 Some service providers already use hash matching to tackle intimate image abuse, and the evidence suggests that it is an effective tool for detecting intimate images that have been shared without consent. More widespread use of hash matching can help to reduce the volume of intimate image abuse being shared and viewed online.

1.6 In June 2025, we consulted on a number of additional safety measures, including a recommendation that certain service providers use hash matching technology to detect intimate image abuse. After careful consideration of the stakeholder feedback, we have decided to confirm this recommendation ahead of the other additional safety measures proposed in the June 2025 Additional Safety Measures Consultation (June 2025 Consultation). In recommending this measure, our aim is that service providers use a hash database that is equivalent to or better than the leading third-party database, StopNCII.org, which will materially reduce the prevalence of intimate image abuse on their services

What we have decided – In brief

In this statement, we set out our decision to confirm our recommendation that the following service providers use 'hash matching' technology to detect intimate image abuse:

- Providers of user-to-user services that are high risk for intimate image abuse and either:
  - whose principal purpose is the hosting or dissemination of regulated pornographic content; or
  - are file-sharing and file-storage services; or
  - have more than 700,000 monthly active UK users.
- Providers of large[1] user-to-user services that are medium risk for intimate image abuse;
- Providers of large general search services.

[1] A large service is defined as a service with more than 7 million monthly active UK users.

We have made some changes to the measure following consultation. We have:

- Expanded the recommendation to use a third-party database to all service providers, unless they have an equivalent or better internal database;
- Revised the criteria for defining an appropriate hash database;
- Clarified the expectations for reviewing matches against a hash database, including the factors that services must consider when using human review to determine whether the content is intimate image abuse and for the purposes of quality assurance;
- Recommended perceptual hash matching (where the database used enables this) or otherwise cryptographic hash matching to detect video-based intimate image abuse; and
- Added a recommendation to notify a database provider of non-intimate image abuse hashes.

We set out our reasoning for these amendments in Chapter 3 of this statement.

Next steps

1.7 The draft amendments set out in this statement to the Illegal Content Codes of Practice (Illegal Content Codes) will be published separately in Summer 2026. They are therefore not yet in force. However, we have published draft consolidated versions of the Illegal Content Codes showing the amendments in context.

1.8 Subject to Parliamentary process, we expect the draft amendments set out in this statement to come into force in Autumn 2026.

1.9 We expect to publish our decision on the other additional safety measures proposed in our June 2025 Consultation in autumn 2026.

1.10 Additionally, in December 2025, the government created two new 'priority offences' under the Act: encouraging or assisting serious self-harm, and cyberflashing. To reflect this, we consulted on proposed changes to the Illegal Harms regulatory documents and guidance in April 2025 and expect to publish our final statement in summer 2026.

2. Introduction

The Online Safety Act 2023 and the Codes of Practice

2.1 The Online Safety Act 2023 (the Act) gives service providers a range of duties in relation to illegal content. These duties require providers to assess and manage the risks arising from the offences specified in the Act.

2.2 Under the Act, Ofcom has duties to produce regulatory documents and guidance to assist providers in complying with their duties. These include the Illegal Content Codes of Practice (Illegal Content Codes), which set out recommended measures that providers can take to comply with their illegal content safety duties and their reporting and complaints duties under the Act. While providers are not required to adopt the measures recommended in the Illegal Content Codes, providers that do so will be treated as compliant with the duties to which they relate.[2]

2.3 We published our first Illegal Content Codes alongside our December 2024 statement on Protecting People from Illegal Harms Online (December 2024 Statement).[3] In June 2025, we consulted on several additional safety measures, including a recommendation that certain service providers use hash matching technology to detect intimate image abuse (June 2025 Consultation).

2.4 We are now amending both the Illegal Content user-to-user Codes and Illegal Content search Codes to introduce an additional measure for tackling intimate image abuse.

2.5 We provide more detail on the statutory basis of Ofcom's role and the issues we must consider when preparing Codes in Annex 2, 'Legal Framework'.

Our approach to impact assessments for the measure

2.6 In designing our measure, we conducted an impact assessment to determine the proportionality of recommending that service providers use hash matching technology to detect intimate image abuse.

2.7 We carefully considered the following factors:

a) The prevalence and impact of harm;
b) The efficacy of the measure to work against the harm;
c) The associated costs and impacts of the measure;
d) Any risks associated with the measure; and
e) The scope of applicable service providers.

2.8 This remains consistent with our approach in the December 2024 Statement.

2.9 The impact assessment for this recommended measure is cumulative. This means we consider the anticipated impacts, costs, and benefits of the measure over and above those that result from existing Code measures.

[2] Section 49(1) of the Act.
[3] Ofcom, 2024. Protecting People from Illegal Harms Online. [accessed 10 March 2026]

Structure of this document

2.10 This statement includes the following chapters that set out our reasoning for the final decisions:

- Chapter 1, 'Overview' gives a high-level summary of our decision.
- Chapter 2, 'Introduction' outlines the structure of statement and corresponding documents.
- Chapter 3, 'Hash matching for intimate image abuse' sets out our recommendation and reasoning.

2.11 We also have four annexes:

- Annex 1: Statutory tests and impact assessments
- Annex 2: Legal framework
- Annex 3: Glossary
- Annex 4: Further detail on economic assumptions and analysis

2.12 We have published draft guidance on appropriate proportion of human review for intimate image abuse hash matching separately.

2.13 We have published our draft consolidated Codes separately:

- Draft consolidated Illegal Content Codes of Practice for user-to-user services
- Draft consolidated Illegal Content Codes of Practice for search services

3. Measure on hash matching for intimate image abuse

Introduction

3.1 In this chapter, we set out our decision to include a new measure in our Illegal Content Codes for user-to-user and search service providers. The measure sets out that service providers should use an automated moderation process known as 'hash matching' to detect intimate images and videos that are shared online without the consent of the person or people depicted. We refer to these images and videos as 'intimate image abuse' for the purposes of this statement.[4]

3.2 Intimate image abuse has profound negative psychological, emotional, professional and relational impacts on survivors and victims. Evidence suggests intimate image abuse is increasing in prevalence, and the proliferation of generative artificial intelligence (GenAI) tools is causing a further increase. There is an urgent need to reduce the spread of intimate image abuse online.

3.3 Recognising the need to address this harm, the UK Parliament recently legislated so that service providers will need to use proportionate systems and processes to take down intimate images shared without a victim's consent within 48 hours, alongside powers to create a statutory register containing hashes of non-consensual intimate images, which designated providers will be required to use.[5] The legislation also bans 'nudification' tools that use GenAI to turn images of real people into fake nude pictures and videos without their permission, and Parliament has already legislated to criminalise the creation of non-consensual sexual deepfakes.[6] We consider these initiatives strengthen the protections against intimate image abuse and complement the decisions we are setting out in this document.

3.4 This chapter sets out a mechanism for service providers to tackle the prevalence of intimate image abuse and reduce the harm to survivors and victims. The chapter is structured as follows:

a) We describe what intimate image abuse is.
b) We describe proactive technology and specifically, hash matching, and the circumstances in which we recommend it.
c) We summarise our decisions in response to the June 2025 Consultation and set out our final recommendation.
d) We assess the impact of the measure, including its effectiveness and benefits, the costs and impacts of the measure, and its impact on the fundamental rights of individuals.

[4] We have provided detailed guidance on how providers should make judgements about whether content is intimate image abuse in Chapter 10 of the Illegal Content Judgements Guidance (ICJG). This includes information about what constitutes 'intimate' and what evidence might suggest lack of consent and either lack of reasonable belief in consent or intent to cause alarm, distress or humiliation. This includes a user report which suggests that the content has been posted non-consensually (absent evidence to the contrary).
[5] Crime and Policing Act 2026.
[6] GOV.UK, 2025. Protecting young people online at the heart of new VAWG strategy. [accessed 27 April 2026].

e) We explain why, on balance, we consider the measure to be proportionate and discuss the service providers that the measure applies to.

3.5 While we have only responded in detail in this statement to the feedback relating to the intimate image abuse proposals, in reaching our decision, we have carefully considered all the feedback from respondents to the June 2025 Consultation.

Intimate image abuse – the law

3.6 As a matter of the criminal law in the UK,[7] intimate image abuse relates to the non-consensual disclosure of, or threats to disclose, intimate images. Intimate image abuse is a priority offence identified in the Act.

3.7 The images need to depict an 'intimate state' or 'intimate situation'. For example, this could include photographs or video of:

- Nudity or a sexual act;
- A person's genitals, buttocks or breasts covered only with underwear; or
- A person in the act of urination or other bodily functions.

3.8 The intimate images can be in the form of a photograph or video.

- A 'photograph' should be understood as any still image which is or appears to be a photograph, whether created by a camera, made by computer graphics or generated in any other way (e.g., a still of a film).
- A 'film' should be understood as any moving image, including automated GIFs.

3.9 The inclusion of the 'appears to show' criteria of the offence means that manipulated images and videos such as AI-generated imagery or deepfakes are within the scope of this offence.

3.10 To be intimate abuse, the image or photograph needs to have been disclosed:

a) Without the depicted person's (Person A's) consent; and
b) When the person disclosing the image (Person B) either
   i) did not reasonably believe that Person A consented; or
   ii) is acting with the intention of causing Person A alarm, distress or humiliation.

3.11 Our Illegal Content Judgements Guidance (ICJG) provides more information on how providers should determine whether content amounts to intimate image abuse. It includes guidance on the nature of an 'intimate state' or 'intimate situation' and about determining consent or a reasonable belief in consent and intent.[8] In relation to consent, the ICJG states that a user report which suggests that the content has been posted non-consensually gives the provider reasonable grounds to believe there was no consent or reasonable belief in consent (absent evidence to the contrary). [9] An example of 'evidence to the contrary' may be where the image itself is clearly an advertisement (e.g. for underwear).

[7] Section 66B of the Sexual Offences Act 2003; section 2 of the Abusive Behaviour and Sexual Harm (Scotland) Act 2016.
[8] Illegal Content Judgements Guidance, paragraphs 10.28-10.46.
[9] Illegal Content Judgements Guidance, paragraph 10.36(a).

Proactive technology – hash matching

3.12 The measure set out in this statement concerns the use of hash matching technology for the detection of intimate image abuse. Hash matching is a form of proactive technology.

Proactive technology

3.13 The term 'proactive technology' is used to describe a set of technologies that can be used as part of a content moderation system. These technologies include content identification technology, user profiling technology, and behaviour identification technology.[10] In practice, this includes rules-based software such as hash matching or keyword detection, automated content classifiers (algorithmic systems which analyse content and associated signals), and other proactive technologies which analyse inputs – including network analysis, user profiling and behavioural analysis.

3.14 Our power to include proactive technology measures in our Codes is subject to certain legal constraints. Under the Act, Ofcom may only recommend the use of proactive technology to analyse user-generated content that is communicated publicly (or metadata related to such content).[11] Consistent with this constraint, the measure described in this chapter applies only to content communicated publicly by means of the service. Service providers may also decide to use proactive technology in relation to content communicated privately by means of the service, including to detect illegal content.

3.15 We must also have regard to the degree of accuracy, effectiveness and lack of bias achieved by the technology when deciding whether to recommend proactive technology.[12] We have considered these matters in the 'Our reasoning' section below.

Hash matching

3.16 Hash matching is a process for detecting content that has previously been identified as being illegal, or otherwise violative, to prevent the upload and sharing of that content.

3.17 It involves creating a hash (an identifying series of characters) of a unique piece of content, storing that hash in a database, and using an algorithm to detect attempts to upload the same or similar versions of that content. Once a match of some form is established, the content can either undergo review – which may include human review – or be removed automatically.

3.18 There are two basic forms of hash matching:

- **Cryptographic hash matching** that only detects files that are an exact match. This means that perpetrators may slightly alter images to circumvent cryptographic hash matching.
- **Perceptual hash matching** that determines whether a given file is similar to a hashed file by setting a threshold that is used to determine when there is sufficient similarity between the hashes to be considered a (near) match. This allows modifications from an original file to be detected.

3.19 Cryptographic and perceptual hashes are generated using different types of hashing algorithms and are not interchangeable. As such, cryptographic hashes cannot be used for similarity-based matching such as perceptual hash matching.

---

[10] Section 231 of the Act.
[11] We have published Guidance on content communicated 'publicly' and 'privately' under the Online Safety Act to assist providers in determining which content on their service is communicated publicly or privately.
[12] Schedule 4 to the Act, paragraph 13(6).

3.20 The types of content that can be hashed include images, videos or files that are generated, uploaded to, or shared on, the providers' services. This means hash matching technology can detect known instances of content generated by artificial intelligence, such as deepfakes.

Measure proposed at consultation

3.21 In our June 2025 Consultation, we proposed including a measure in our Codes that certain service providers (both user-to-user and search) use perceptual hash matching technology to detect intimate image abuse, where technically feasible.

3.22 We proposed that the measure should apply to the following service providers:

a) Providers of user-to-user services that are high risk for intimate image abuse and either:
   i) whose principal purpose is the hosting or dissemination of regulated pornographic content; or
   ii) are file-sharing and file-storage services; or
   iii) have more than 700,000 monthly active UK users.
b) Providers of large user-to-user services that are medium risk for intimate image abuse.
c) Providers of large general search services.

3.23 We also proposed that the following providers should source intimate image abuse hashes from an appropriate third-party database (though they could choose to supplement this with an internal hash database):

a) Providers of user-to-user services that are high risk for intimate image abuse and either:
   i) whose principal purpose is the hosting or dissemination of regulated pornographic content; or
   ii) are file-sharing and file-storage services.

3.24 We said that all other service providers in scope of the measure may use a third-party hash database, an internal hash database or both to source intimate image abuse hashes.

Summary of stakeholder feedback

3.25 We received support from a wide range of stakeholders in response to our proposal.[13] Stakeholders noted there was compelling evidence to demonstrate that hash matching was effective, scalable, and privacy-preserving.[14] One stakeholder emphasised that service providers should use hash matching technology to improve the scalability and effectiveness of hash matching technology and indicated that the proposed measure would facilitate this

[13] Response(s) to our June 2025 consultation: Barnardo's, p.10; British and Irish Law Education and Technology Association, p. 13-14; Canadian Centre for Child Protection, p.11; Commissioner Designate for Victims of Crime for Northern Ireland, p.5; Department of Justice for Northern Ireland, p.9; Domestic Abuse Commissioner, p.1; [✂]; [✂]; Image Angel, p.9; Internet Watch Foundation, p.31; [✂]; National Center for Missing & Exploited Children, p.4; National Police Chiefs' Council, p.6; Online Safety Act Network, p.15; [✂]; Philippine Survivor Network, p.15; Refuge, p.8; Scottish Government, p.3; [✂]; Suzy Lamplugh Trust, p. 4; South West Grid for Learning, p.10-11; The Information Commissioner (ICO), p.15; [✂]. Response(s) to our February 2025 consultation: Centre for Protecting Women Online, p.18-19; Children's Commissioner for England's Office, p. 5; End Violence Against Women Coalition (EVAW), p.16; Equality Now, p.5; Women's Aid Federation of England, p.10.

[14] Response(s) to our June 2025 Consultation: Evangelical Alliance, p.5; Internet Watch Foundation, p.32.

improvement of technology.[15] Another stakeholder highlighted that hash matching enabled quick removal of flagged content and was an effective option to support smaller service providers' efforts to moderate intimate image abuse on their services.[16]

3.26 We also received responses that disagreed with our proposals.[17] The primary concern raised by stakeholders was the use of a database that contains unverified hashes ('unverified' means that the database operator has not checked that the underlying image is intimate). Specifically, several stakeholders noted the higher likelihood that databases that do not verify hashes may contain hashes of content other than non-intimate image abuse content.[18]

3.27 We also received feedback from several stakeholders that recommended alternative online safety measures to tackle intimate image abuse online.[19] Other stakeholders recommended collaboration between relevant stakeholders, including industry and law enforcement, to enable the effective use of hash matching.[20]

3.28 We have carefully considered all the feedback we received from stakeholders in reaching our decision. However, not all feedback is explicitly addressed in this statement.

3.29 We provide more detail on some stakeholder feedback, including how it impacted our decisions, in the 'Our reasoning' section below.

Our decision

3.30 We have decided to broadly confirm the measure proposed in the June 2025 Consultation. This includes the scope of services that the measure applies to.

3.31 We have made some amendments to the measure we consulted on, including:

a) Expanding the recommendation to use a third-party database to all service providers, unless they have an equivalent or better internal database;
b) Revising the criteria for defining an appropriate hash database;
c) Clarifying the expectations for reviewing matches against a hash database, including the factors that service providers must consider when using human review to determine whether the content is intimate image abuse and for the purposes of quality assurance;
d) Recommending perceptual hash matching (where the database used enables this) or otherwise cryptographic hash matching to detect video-based; and
e) Adding a recommendation to notify a database provider of non-intimate image abuse hashes.

3.32 We are also publishing draft guidance alongside this statement to assist providers in determining the appropriate proportion of detected content that should be subject to

[15] Response(s) to our June 2025 Consultation: [✂].
[16] Response(s) to our June 2025 Consultation: British and Irish Law Education and Technology Association, p.14.
[17] Response(s) to our June 2025 Consultation: Computer and Communications Industry Association, p.7; [✂]; Luethje, Y., p.4; Name Withheld 9, p.5; [✂]; Porter, J., p.4; [✂].
[18] Response(s) to our June 2025 Consultation: [✂]; [✂]; [✂]; The Information Commissioner (ICO), p.17-18. Response(s) to our February 2025 consultation: [✂].
[19] Response(s) to our June 2025 Consultation: Domestic Abuse Commissioner, p.2-7; [✂]; Image Angel, p. 9-12; [✂]; Suzy Lamplugh Trust, p. 4; [✂]; UCL Gender and Tech Research Lab, p. 2.
[20] Response(s) to our June 2025 Consultation: Canadian Centre for Child Protection, p. 12; National Police Chiefs' Council, p. 7; [✂]; University College London (UCL) Gender and Tech Research Lab, p. 3.

human review.[21] We will finalise this guidance alongside the amendments to our Codes of Practice following the Parliamentary procedure.

3.33 To contextualise these amendments, we have published the draft consolidated Codes of Practice for user-to-user services and search services. The amendments to the Codes of Practice will be published in early June 2026.

3.34 We discuss in more detail the amendments made to the measure below.

Expanding the recommendation for service providers to use a third-party database

3.35 We have amended the measure to recommend that all service providers use third-party databases to detect intimate image abuse on their services, unless they have an equivalent or better internal database. We set out specific criteria for selecting an appropriate database, which we detail below.

3.36 As a consequence, we have removed the recommendation that user-to-user services that are high risk for intimate image abuse and either (1) whose principal purpose is the hosting or dissemination of regulated pornographic content; or (2) are file-sharing and file-storage services should source intimate image abuse hashes from an appropriate third-party database.

Revising the criteria for defining an appropriate database

3.37 We have amended the measure to remove the specific criteria for a third-party database versus an internal database. We concluded that this distinction is largely unnecessary considering our decision to remove the specific recommendation for certain high-risk user-to-user services to use third-party databases. The measure now sets out criteria that apply to all databases.

3.38 As part of this decision, we have amended the criteria that apply to all databases:

a) The first criterion is that a database needs to "contain hashes of a significant number of original items of content". This is a new criterion. Our aim is that service providers use a database that is equivalent to or better than the leading existing third-party database, StopNCII.org, to materially reduce the prevalence of intimate image abuse on their services.
b) The second criterion is that the database be proactively updated with reasonable regularity. This was included in our consultation and ensures that hashes are regularly added and that the most up-to-date version of the database is used.
c) The third criterion is that the service be secured from unauthorised access, interference or (in the case of a database comprised of verified hashes) exploitation. We have amended this criterion from consultation by removing the reference in our proposed Code measure to securing a database comprised of unverified hashes against "exploitation". While there is a potential risk of people exploiting the database by

[21] The draft Guidance published alongside this Statement applies only to hash-matching measures addressing intimate image abuse. There is an existing separate measure relating to the use of hash-matching technologies for child sexual abuse material (CSAM), in our Codes. In our June 2025 Consultation, we consulted on proposals to amend the CSAM hash-matching measure. We recognise that CSAM presents distinct considerations, including providers' statutory duties to report detected child sexual exploitation and abuse content to the National Crime Agency or a foreign agency. We will consider the role of human review in the context of CSAM hash-matching when we issue our final Statement on the Additional Safety Measures.

submitting hashes of content that is not intimate image abuse, we have concluded that this risk can be managed appropriately by service providers.

3.39 Our reasoning is set out in paragraphs 3.111 to 3.145.

Adding a recommendation to notify a database provider of non-intimate image abuse hashes

3.40 We have added a recommendation that where the provider of a third-party database has made arrangements for service providers to submit either newly identified intimate image abuse content for hashing, or new hashes, they should use such arrangements to provide any relevant content or hashes accordingly.

3.41 We have also amended the measure to clarify the circumstances in which a service provider should take steps to secure the removal of hashes from a third-party database.

3.42 We provide more information on these additions in paragraphs 3.78 to 3.84.

Clarifying the expectations for how service providers need to review content (including human review)

3.43 We have amended our measure to be clear that where a service provider uses a database with verified[22] hashes, there is not necessarily any need for further review of detected content, except for the purposes of ongoing quality assurance and monitoring of the performance of the technology.[23] Service providers may take moderation action (i.e., take the content down or deprioritise it) without further reviewing the content, depending on the level of assurance the provider has in the detection outcomes achieved by the hash matching technology. We provide more detail on this element in paragraphs 3.69 to 3.72.

3.44 We have amended our measure to set out that where a service provider uses a database that contains unverified hashes, they should review content to determine whether it is intimate image abuse at the first positive match of an image or video to a hash. As per our existing measure on content (ICU C1) and search moderation (ICS C1), this review can be achieved by either human moderators or automated technology (e.g., nudity detection technology).[24] We have made it clearer that subsequent matches to a hash at the same matching parameters may not require re-verification that the image or video is intimate image abuse content, depending on the level of assurance the provider has in the detection outcomes achieved by the hash matching technology. We have added definitions of "first

[22] We use the term 'verified' to refer to a situation where the database operator has determined that the underlying image is intimate image abuse.

[23] Data protection law imposes separate requirements that may necessitate consideration of human involvement in moderation decisions. Services should consult guidance from the Information Commissioner (ICO) for more information.

[24] Service providers should also have regard to the requirements of Articles 22A-22C of the UK GDPR. These provisions relate to automated decision making, including profiling. For the purpose of these provisions, a decision is 'based solely on automated processing' if there is no meaningful human involvement in the taking of the decision (Article 22A(1)(a)); and a decision is a 'significant decision' if, in relation to a data subject, it produces a legal effect for the data subject or it has a similarly significant effect for the data subject (Article 22A(1)(b)). Article 22B puts in place certain restrictions for automated decision-making and Article 22C sets out safeguards for automated decision-making. See also (at the time of this Statement): The Information Commissioner (ICO), 2026. Automated decision-making, including profiling. [accessed 5 May 2026].

match" and "subsequent match" to the measure. We provide more detail on this element in paragraphs 3.73 to 3.77.

3.45 We have amended the measure to provide that, in determining the appropriate proportion of detected content to be subject to human review[25], the provider should have regard to the level of assurance the provider has in the detection outcomes achieved by the technology and any associated systems and processes. We have published draft guidance on this human review for moderation purposes which provides more detail on what factors service providers should consider. Our reasoning is set out in paragraphs 3.167 to 3.177.

Recommending perceptual hash matching (where the database enables this) or otherwise cryptographic hash matching to detect video-based content

3.46 We have amended the measure to recommend that service providers should use perceptual hash matching in relation to video-based content where an appropriate database is available and enables this, but that if there is none, they may use cryptographic hash matching. This will apply to any hash database (i.e., internal or third-party) used by service providers. Our reasoning for this is set out in paragraphs 3.104 to 3.110.

Explanation of the measure

3.47 In this section, we explain in detail the final elements of the measure.

3.48 Our measure sets out two steps that service providers should undergo to support the moderation of intimate image abuse content (as set out in our existing content and search moderation measures) on their services. Those two steps are:

a) Detect suspected intimate image abuse content; and
b) Review of detected content, in some cases, to determine whether it is intimate image abuse content.

3.49 Following these steps, service providers should appropriately moderate content that they consider to be intimate image abuse content. This means that user-to-user services should swiftly take it down; and search services should either swiftly take it down or deprioritise it. We detail this process below.

Step 1 - Detection of suspected intimate image abuse

3.50 Under the measure service providers within scope of the measure should deploy hash matching technology to detect intimate image abuse before, or as soon as practicable after, such content can be encountered by UK users.

3.51 Content already present on the service at the time the hash matching technology is implemented should be analysed within a reasonable time frame.

Sourcing intimate image abuse hashes

3.52 Under the measure, service providers should source intimate image abuse hashes from an appropriate third-party database and/or internal database. We set out below the criteria that databases should meet to be considered appropriate for use.

[25] This is for the purposes of determining whether detected content is intimate image abuse content and for quality assurance.

3.53 We explain our reasoning for recommending the use of third-party databases or an internal database in paragraphs 3.111 to 3.126.

Criteria for an appropriate database

3.54 For the set of hashes to be appropriate, it should:

a) Contain hashes of a significant number of original items of content;
b) Be proactively updated reasonably regularly;
c) Be secured from unauthorised access, interference or (in the case of a database comprised of verified hashes) exploitation (whether by persons who work for that person or are providing a service to that person, or any other person).

3.55 We consider that the current leading hash database, StopNCII.org, meets these criteria and therefore is appropriate for use by service providers. It should be used as the benchmark for determining whether other hash databases, including an internal database, are appropriate to use for this measure. In practice, this means that a service provider should not use an internal database, on its own, unless it is broadly equivalent to StopNCII.org in terms of quantity of original items of content hashed and quality of the hashes. We further explain our reasoning in paragraph 3.141 to 3.145.

Type of detection technology

3.56 For this measure, service providers should use specific types of hash matching technology based on the type of content (i.e., image or video) being detected:

a) Perceptual hash matching to detect image-based intimate image abuse, that is, photographs and other still images.
b) Perceptual hash matching (where the database enables this) or otherwise cryptographic hash matching to detect video-based intimate image abuse.

3.57 We explain our reasoning for recommending these technologies in paragraph 3.104 to 3.110.

Configuration of the technology

3.58 Service providers should ensure that perceptual hash matching technology is configured to strike an appropriate balance between precision and recall. This does not apply to cryptographic hash matching.

3.59 In doing so, service providers will need to consider the following factors:

a) Risk of harm to users from intimate image abuse, reflecting the service's latest illegal content risk assessment, and any information reasonably available to the provider about the prevalence of content that is intimate image abuse;
b) Proportion of content detected as a match by the hash matching technology that is a false positive (e.g. the image is not an intimate image); and
c) Effectiveness of the systems and processes used to identify false positives (e.g., user complaint mechanisms).[26]

3.60 Service providers should review the balance between precision and recall at least every six months, to ensure that it remains appropriate.

---

[26] We consider our existing user appeals measures in the Illegal Content Codes [ICU D1, D2, D8,D9 and D10 (U2U) and ICS D1, D2, D8,D9 and D10] will enable appropriate mechanisms for users to submit complaints to service providers if their content has been wrongly removed (i.e., consensual intimate images or non-intimate images).

3.61 Service providers will also need to keep a written record of their approach to configuring the technology and review it at least every six months.

3.62 Our full reasoning on configuration of the technology is at paragraphs 3.146 to 3.154.

Quality assurance and monitoring of technology

3.63 Service providers should use human review for the purposes of quality assurance and monitoring of the configuration of technology. This review will apply to any technology used for the purposes of this measure, including hash matching technology and other automated technologies (these may be used for the review of content, as set out below).

3.64 This element of the measure reflects amendments to the factors that service providers should take into account when determining what proportion of detected content should be subject to human review. We are also providing additional draft guidance to assist providers.

Step 2 - Review of detected content

3.65 Once content is detected, the next steps depend on whether the service provider is using a database that contains verified or unverified hashes. If a database contains both, the treatment of matches will depend on whether the match is known to be verified or may be an unverified hash.

3.66 A 'verified' hash means the database provider determined that the hash is of intimate image abuse content at the time the hash was added to the database, whereas 'unverified' means that the database operator has not checked this.

Reviewing content matched against a hash database

3.67 The measure sets out different review steps for:

a) review of positive matches derived from verified hashes; and
b) review of positive matches derived from unverified hashes.

Review of verified hashes

3.68 For a match detected against a verified hash, we do not necessarily expect service providers to review a positive match to determine whether it is intimate image abuse content in all cases. We consider that a match against a verified hash may provide reasonable grounds to assume the content is intimate image abuse content, as the included content was reviewed upon submission. However, the extent to which a provider may assume this will depend on its level of assurance in the detection outcomes achieved by the hash matching technology. For example, where the provider has configured the perceptual hash matching technology to prioritise recall, it may be appropriate to review the content.

3.69 In practice, this means that once there is a positive match against a verified hash, the service provider may, depending on the circumstances, treat this as reason to infer that the detected content is illegal content.[27] The lower the provider's level of assurance in the detection outcomes, the more likely it is that further review of detected content will be appropriate.

3.70 Where the service provider does treat the match as reason to infer that the detected content is illegal content, under existing content moderation measures, the provider of a

[27] See ICU C2.3(a) and ICS 1.3(a).

user-to-user service should swiftly take the content down[28] and the provider of a search service should take down or deprioritise the content from search results.[29] Additionally, human review for quality assurance purposes will still be needed.

3.71 Our reasoning is set out in the 'Benefits and effectiveness' at 3.160.

Review of unverified hashes

3.72 For a match against an unverified hash, service providers should treat it as reason to suspect the content may be intimate image abuse. This is because the original item of content has not been reviewed prior to determining this.

3.73 Service providers should review the content in accordance with existing content moderation measures[30] and the ICJG.[31] The purpose of this review is to determine whether the content is intimate image abuse content. Service providers may use either automated technologies (e.g., nudity detection technology) or human moderators (or both) to conduct this review. Whether or not humans review content for content moderation purposes, providers should consider whether humans should review a selection of content for quality assurance purposes, in line with our 'human review' recommendation explained below and accompanying guidance.

3.74 If a service provider determines that the image or video is intimate image abuse content, or if the match is not the first positive match, under existing content moderation measures, the provider of a user-to-user service should swiftly take the content down[32] and the provider of a search service should take down or deprioritise the content from search results.[33]

3.75 In our view, it is only necessary to review the first positive match[34] of the content to determine whether it is intimate image abuse content. Subsequent matches using the same configuration do not necessarily require additional review. However, it may be appropriate to review subsequent matches where a different configuration has been used, depending on the level of assurance that the provider has in the detection outcomes. For example, if the configuration is changed to increase the matching distance threshold,[35] it is more likely to be appropriate to review the subsequently matched content. This is because the revised configuration expands the set of matched content and may increase the risk of false

---

[28] ICU C2.3.
[29] ICS C1.3-4.
[30] ICU C1.3 and ICS1.3.
[31] In relation to consent, the ICJG states that a user report which suggests that the content has been posted non-consensually gives the provider reasonable grounds to believe there was no consent or reasonable belief in consent (absent evidence to the contrary).
[32] ICU C2.3.
[33] ICS C1.3-4.
[34] By 'first positive match', we mean the first time content is detected as a match for an individual reference hash under a given configuration of the hash-matching technology. This can include the first positive match to a hash under the same perceptual matching configuration (or, if used for video content, an exact cryptographic match).
[35] A perceptual hash is calculated in a way that means visually similar inputs produce similar hash values. Similarity between perceptual hashes is determined using a distance metric, which defines how differences between the hash values are measured. The matching distance threshold specifies the maximum permitted distance at which two hashes are classified as a match.

positives. A reduction in the matching distance threshold will not ordinarily require re-review.[36]

3.76 Our reasoning is set out in the 'Benefits and effectiveness' at 3.161.

Appropriate actions based on the content review outcome

3.77 The measure sets out specific recommendations for a service provider depending on the outcome of their content review, which we set out below.

If a service provider becomes aware of content that is intimate image abuse content

3.78 The measure sets out that when service providers identify new intimate image abuse, they should, where possible, add hashes of this content to their internal databases.

3.79 If a service provider is using a third-party database which has made arrangements specifically for providers to submit either newly identified intimate image abuse content for hashing, or new hashes, the provider should use such arrangements to provide such content or hashes to the third-party provider. This is new to the measure because of updates to the technical capabilities of existing third-party databases.[37]

If a service provider becomes aware of content that is not intimate image abuse

3.80 The measure sets out that a service provider should also take reasonable steps to secure the removal of a hash from the hash database where the provider has determined that the hash is not of intimate image abuse content. For this purpose, the service provider may determine that the hash is not of intimate image abuse content in one or more of the following circumstances:

a) The provider has viewed the image used to generate the hash and determined it is not intimate image abuse content.
b) The provider otherwise reasonably believes that the hash is of an image that is not intimate image abuse content.
c) The provider has viewed content matched with the hash using cryptographic hash matching technology and determined that that content is not intimate image abuse content.
d) The provider has viewed content matched with the hash using perceptual hash matching and determined that that content is not intimate image abuse content, provided that the provider reasonably believes that there is an exact match between the hash of the detected content and the hash included on the database.

3.81 In practice, where a service provider uses a third-party database, it will satisfy this component of the measure by notifying the hash database provider of the non-intimate image abuse hash. For example, service providers using StopNCII.org can add a tag to a specific hash to communicate that it is not intimate image abuse. This would meet the criteria set out at 3.67.

3.82 For clarity, our measure does not impose recommendations on third-party database providers to remove hashes from their database that are determined not to be intimate image abuse (we do not have any power to regulate providers of third-party databases). That type of recommendation is not within our remit. We only recommend that a service

[36] This applies only to fine-tuning of the matching configuration (e.g., threshold adjustments). More fundamental changes to the matching pipeline (e.g., changes to preprocessing steps or hash generation) require separate review.
[37] [✂].

provider notify the database provider of non-intimate image abuse. However, we do strongly encourage third-party database providers to remove the non-intimate image abuse hashes that service providers report to ensure greater accuracy of their database. We do not consider it in the interest of third-party databases to retain these hashes.

3.83 If a service provider uses an internal database, taking 'reasonable steps to secure the removal of the hash' would involve removing the hash from that database.

Our reasoning

3.84 In this section, we detail our reasoning for deciding to include the measure in our user-to-user and search Illegal Content Codes. This includes an analysis of the following:

a) The benefits and effectiveness of our measure;
b) The costs and impacts of the measure; and
c) The impact on users' rights from the measure.

Benefits and effectiveness

3.85 The primary benefit of the measure is the reduction in the availability of intimate image abuse content, and, by extension, the harm caused to survivors and victims.

3.86 To reduce the harm, our assessment is that service providers can use proactive technology, specifically hash matching technology, to detect intimate image abuse on their services.

3.87 We discuss the effectiveness of hash matching technology below, including its limitations (i.e., availability and accuracy of databases) and the safeguards available to minimise their impact.

Intimate image abuse is widespread and causes serious harm

3.88 Evidence shows that intimate image abuse is a widespread and increasingly prevalent harm, with AI-generated intimate image abuse growing exponentially. In 2024 over 22,000 reports of intimate image abuse were made to the Revenge Porn Helpline.[38] Additionally, between the year ending June 2022 and the year ending June 2025, recorded intimate image abuse offences in England and Wales rose by 26.9% overall.[39]

3.89 The evidence also indicates that women are disproportionately affected by intimate image abuse, with the Revenge Porn Helpline recording that 98% of the intimate images reported to the service were of women.[40] The prevalence is particularly high among young adults, with 69% of reports to the Revenge Porn Helpline made by adults aged 18 to 34.[41]

3.90 Intimate image abuse has a profound negative impact on survivors and victims, including psychological, emotional, professional and relational harm. For many survivors and victims of this kind of abuse, the impacts are devasting and all-encompassing representing a "social

[38] Revenge Porn Helpline (Papachristou, K.), 2025. Revenge Porn Helpline 2024 Report. [accessed 14 April 2026].
[39] Refuge, 2026. Intimate Image Abuse Research Reveals Rising Reports but Low Charge Rates. [accessed 5th May 2026].
[40] SWGfL (Wright, D. and Mortimer, S.), 2024. Revenge Porn Helpline: Evidence to Women & Equalities Committee. [accessed 10 March 2026].
[41] Revenge Porn Helpline (Papachristou, K.), 2024. Revenge Porn Helpline 2023 Report. [accessed 14 April 2026].

rupture" that changes their lives irrevocably[42] and survivors and victims' struggle with a range of feelings, including shame, isolation and humiliation.[43] The Revenge Porn Helpline reported that 60% of its clients are referred to mental health support services due to the impact of their experiences.[44] Some consultation responses also highlighted the severe consequences and lasting harm caused by this abuse.[45]

3.91 Survivors and victims of intimate image abuse experience a constant fear that images will not be removed or might resurface online.[46] As a result, many survivors and victims can develop hypervigilance[47] and many report feeling "extremely fearful" for their safety.[48]

3.92 Perpetrators can distribute this content on any service with the functionality to post or share images, videos or files. This includes sites dedicated to the non-consensual sharing of intimate images, where perpetrators request, trade and share images. Search services can host or index intimate image abuse. As set out in the Illegal Harms Register of Risks, evidence indicates that services whose principal purpose is the hosting or dissemination of regulated pornographic content and file-sharing and file-storage services are at higher risk of containing intimate image abuse.

Effective complaints handling by service providers is not sufficient to address intimate image abuse and its harm

3.93 Service providers can use their reporting and complaints systems to enable the identification and removal of intimate image abuse.[49] However, reporting and complaints systems alone, even where providers respond appropriately to all complaints, cannot sufficiently address the harm posed by intimate image abuse.

3.94 Survivors and victims of intimate image abuse face ongoing challenges in reporting this content. For example, they may not have knowledge of the content being shared or may only be aware of a threat to share without any details. They also do not necessarily even know that this content exists, as intimate image abuse can include images that were taken or made without consent and/or knowledge, and the growth of AI-generated intimate image abuse increases this risk. Even in cases where survivors and victims are certain that their intimate images have been shared without consent, they may not be able to identify where or by whom (i.e. the account) images were shared.[50]

---

[42] McGlynn, C., Johnson, K., Rackley, E., Henry, N., Gavey, N., Flynn, A. and Powell, A., 2021. 'It's Torture for the Soul': The Harms of Image-Based Sexual Abuse, Social & Legal Studies, 30(4). [accessed 14 April 2026].
[43] Law Commission, 2022. Intimate image abuse: a final report. [accessed 14 May 2026].
[44] Revenge Porn Helpline (Papachristou, K.), 2024. Revenge Porn Helpline 2023 Report. [accessed 14 April 2026].
[45] Response(s) to our June 2025 Consultation: Domestic Abuse Commissioner, p.1; Evangelical Alliance, p.6.
[46] McGlynn, C., Johnson, K., Rackley, E., Henry, N., Gavey, N., Flynn, A. and Powell, A., 2021. 'It's Torture for the Soul': The Harms of Image-Based Sexual Abuse, *Social & Legal Studies,* 30 (4). [accessed 14 April 2026].
[47] McGlynn, C., Johnson, K., Rackley, E., Henry, N., Gavey, N., Flynn, A. and Powell, A., 2021. 'It's Torture for the Soul': The Harms of Image-Based Sexual Abuse, Social & Legal Studies, 30 (4). [accessed 14 April 2026].
[48] Henry, N., Flynn, A. and Powell, A., 2019. Responding to 'revenge pornography': Prevalence, nature and impacts. [accessed 14 April 2026].
[49] The Crime and Policing Act 2026 places a new duty on regulated services to use proportionate systems and processes to remove intimate images shared without the consent of the individual(s) depicted no later than 48 hours after it is reported to them. They will also have a duty to allow users and affected persons to easily report intimate image content to the provider.
[50] Revenge Porn Helpline (Papachristou, K.), 2025. Revenge Porn Helpline 2024 Report. [accessed 14 April 2026].

3.95 There is also a challenge that the intimate images may be shared to an online audience that is unlikely to submit complaints to the service provider. This is because the online audience may not know, or in some cases may not be concerned, that the intimate images are being shared without consent. Even if an individual suspects that the content has been shared without the consent of the person depicted, they may not feel able to raise a concern because they are not depicted in the image.

3.96 Intimate images shared without consent can also be re-shared, re-posted or forwarded to other services or by other individuals. The Domestic Abuse Commissioner for England and Wales noted that perpetrators often use multiple platforms to evade detection.[51] This secondary distribution[52] of images can cause intimate image abuse content to 'go viral', as it becomes more difficult to get images removed when they are repeatedly re-uploaded to the same service or shared to other services. As one respondent highlighted, this means that survivors and victims face a constant challenge of finding and reporting their intimate images, which can cause considerable distress and may lead to re-traumatisation.[53]

3.97 We are not aware of other effective ways to protect victims and survivors of intimate image abuse, besides using proactive technology to detect relevant content quickly.

Hash matching is effective at detecting known non-consensual intimate images

3.98 Many service providers, including large search services, deploy hash matching technology to detect known intimate image abuse and consider it to be effective. For example, one service provider reported that it took action on nearly 270,000 matched intimate images following a pilot implementation of hash matching.[54]

3.99 In response to the June 2025 Consultation, we received support for the measure from many stakeholders,[55] including from service providers who indicated they already use hash matching technology and consider it effective at identifying and removing known intimate image abuse.[56] One stakeholder indicated that the proposed measure will help reduce the need for human review of this distressing material.[57]

---

[51] Response(s) to our June 2025 Consultation: Domestic Abuse Commissioner, p.4.

[52] This secondary content sharing can be done using a service's forwarding functionalities or re-posting functionalities. It can also be done by the user downloading the image to their device and then sharing it on another service. As well as re sharing intimate images to services, it is possible that the intimate images are also being forwarded on private messaging channels.

[53] Response(s) to our June 2025 Consultation: Image Angel, p.8-9.

[54] Microsoft (Gregoire, C.), 2024. An update on our approach to tackling intimate image abuse. [accessed 10 March 2026].

[55] Response(s) to our June 2025 Consultation: Barnardo's, p.10; British and Irish Law Education and Technology Association, p. 13-14; Canadian Centre for Child Protection, p.11; Commissioner Designate for Victims of Crime for Northern Ireland, p.5; Department of Justice for Northern Ireland, p.9; Domestic Abuse Commissioner, p.1; [✂]; [✂]; Image Angel, p.9; Internet Watch Foundation, p.31; [✂]; National Center for Missing & Exploited Children, p.4; National Police Chief' Council, p.6; Online Safety Act Network, p.15; Philippine Survivor Network, p.15; Refuge, p.8; Scottish Government, p.3; [✂]; Suzy Lamplugh Trust, p.4; South West Grid for Learning, p.10-11; The Information Commissioner (ICO), p.15; [✂]. Response(s) to our February 2025 Consultation: Centre for Protecting Women Online, p.18-19; Children's Commissioner for England's Office, p.5; EVAW, p.16; Equality Now, p.5; Women's Aid Federation of England, p.10.

[56] Response(s) to our June 2025 Consultation: [✂]; [✂].

[57] Response(s) to our June 2025 Consultation: Scottish Government, p.3.

3.100 However, two stakeholders noted that hash matching is only effective at detecting previously identified intimate image abuse.[58] Other stakeholders suggested that hash matching is not effective for identifying AI-generated intimate image abuse[59] or called for Ofcom to clarify how the measures apply to this content.[60]

3.101 While we acknowledge that hash matching technology cannot detect 'unknown' intimate image abuse – images that have not been reported as being non-consensual intimate images – hash matching is an effective way to detect known intimate image abuse, and its detection will enable a significant reduction in harm online.

3.102 We also consider the pre-existing use of hash matching technology by some service providers, including user-to-user and search services,[61] reaffirms that the use of these systems and processes is best practice for tackling intimate image abuse. It further means our recommendation is based on proven effectiveness in detecting intimate image abuse on a range of services. We expect the inclusion of this measure in the Codes will reinforce these services' decision to use hash matching technology to detect intimate image abuse and also encourage other service providers to do the same.

Effectiveness of perceptual and cryptographic hash matching

3.103 We conclude that perceptual hash matching is more effective than cryptographic matching at protecting users from harm.

3.104 Perceptual hash matching detects content that is identical or similar to known intimate images. This enables providers to identify images that have been altered in some way to avoid detection – for example, images that have been cropped. Service providers set a threshold to determine when there is a sufficient similarity between the hash in the database and the hash of the piece of content from their service. This allows the fine-tuning of a threshold to produce the greatest accuracy and quantity of detected intimate image abuse.

3.105 Perceptual hash matching is therefore likely to detect a larger volume of matches than cryptographic hash matching and bring about a greater reduction in non-consensual sharing of intimate images. This was reiterated during consultation by stakeholders who emphasised that perpetrators can easily bypass cryptographic hash matching[62] and that perceptual hash matching tackles the issue of minor changes to a file.[63]

3.106 However, we now understand that recommending perceptual hash matching to detect video-based intimate image abuse content may hinder the impact of the measure. This is not because perceptual hash matching technology is ineffective in principle. Rather, it reflects the current capabilities of hash databases, both internal and third-party, for intimate image abuse content only. We explain this below.

---

[58] Response(s) to our June 2025 Consultation: [✂]. Response(s) to our February 2025 Consultation: Do-Ngoc, T. and Carmel, E., p.2.
[59] Response(s) to our June 2025 Consultation: Integrity Institute, p.5; [✂].
[60] Response(s) to our June 2025 Consultation: Online Safety Act Network, p.16; [✂].
[61] We understand that several services use hash matching technology to detect intimate image abuse on their services, as demonstrated by the partnering service providers with StopNCII.org. Source: Industry Partners. [accessed 8 April 2026].
[62] Response(s) to our June 2025 Consultation: [✂].
[63] Response(s) to our June 2025 Consultation: National Police Chiefs' Council, p.6.

3.107 To date, service providers have been using cryptographic hash matching to detect video-based content.[64] This means most existing sets of intimate image abuse hashes include only cryptographic hashes of video-based content. For this reason, if we only recommend the use of perceptual hash matching technology for all intimate image abuse content (i.e., image-based and video-based), as proposed in our June 2025 Consultation, service providers will not be able to effectively detect video-based content.

3.108 Having considered stakeholder responses, we have concluded that restricting our measure to perceptual hash matching technology would significantly reduce the impact of the measure. In many cases, service providers and third-party organisations would need to re-build databases of video-based intimate image abuse content to include perceptual hashes of such content. This would delay the ability for service providers to use such databases to detect intimate image abuse in video form and incur significant additional costs and resources.

3.109 Taking this into consideration, we have amended our measure to recommend that service providers use perceptual hash matching for video-based intimate image abuse where the database enables this, and if not, use cryptographic hash matching. This does not reflect a change in our position that perceptual hash matching is most effective for detecting intimate image abuse. We still expect providers to use perceptual hash matching for the detection of static intimate image abuse content and encourage them, where the database enables this, to use perceptual hash matching to detect video-based content.

Availability of an appropriate hash database, including the accuracy of its hashes, impacts its effectiveness at reducing harm

3.110 Hash matching is only effective if a suitable database containing a set of hashes of intimate images is available to service providers.

3.111 We understand that there are databases that have been developed by service providers themselves (internal databases) and other databases developed by third-party organisations. We also understand that service providers may use one or both of these types of databases. We discuss each of these below.

Service provider's internal hash database

3.112 Service providers may build an internal hash database by including content detected through their daily content moderation processes, including user reports. This means that hashes included in a service provider's internal database are likely to have been reviewed by their content moderators to determine that the content is intimate image abuse content.[65] In other words, these internal databases are more likely to include verified hashes and those hashes are more likely to correspond to content that is illegal or illegal by proxy.[66] This in turn means that the hash matching system is likely to accurately detect intimate image abuse content.

3.113 An internal database is also likely to increase in volume on a regular basis due to newly identified intimate image abuse content via their user reports systems. Service providers

[64] Response(s) to our June 2025 Consultation: [✂]; Google, p.22.

[65] Our Illegal Content Judgments Guidance sets out the basis upon which we expect providers to infer lack of consent to the publication of the image and how to treat forwarding and reposts of such images. See Illegal Content Judgments Guidance paragraphs 10.36 to 10.38.

[66] Content that violates a service provider's terms of service or community guidelines, where these are drafted in a way that is broad enough to capture intimate image abuse content.

should have sufficient user reporting mechanisms in place as part of their existing duties under the Illegal Content Codes of Practice at ICU D1 and ICU D2 for U2U services and ICS D1 and ICS D2 for search services.

3.114 However, the development of an internal database that contains hashes of an appropriate volume of intimate image abuse content may take a significant amount of time and resources and relies on good user reporting systems. This may not be feasible for all service providers. Smaller or less-well known services may face significant challenges to developing and maintaining a suitable internal database due to resource constraints. Additionally, we expect survivors and victims are less likely to report their content to these types of service providers as they may be unaware of their existence.

3.115 There are also risks associated with the use of an internal database. There is a serious risk to fragmentation of survivors and victims' reports of intimate image abuse if each service provider develops and only uses its own internal database. In this scenario, survivors and victims would have to report their intimate images to each service provider individually. The risk of databases becoming fragmented and the benefits of cross-platform hash sharing were raised by several stakeholders in their responses to the June 2025 Consultation.[67] To the extent possible, we would like to reduce this burden on survivors and victims. We discuss in more detail the associated risks below.

3.116 There is also a risk of consensually shared intimate images being included in a hash database. This may result in the content being incorrectly identified as intimate image abuse and removed from the provider's service. We discuss this risk in more detail below.

3.117 Notwithstanding the above, an internal database can be an appropriate option for some service providers. Large or medium sized service providers are more likely to build and maintain a suitable internal database for a variety of reasons, including the number of users, amount of content posted to their services, number of user reports, and available resources to appropriately maintain the database. In particular, a suitable internal database needs to contain hashes of an appropriate number of items of intimate image abuse content, with new ones being added on an ongoing basis, to effectively reduce the spread of such content online. We consider large and medium sized services are more likely to build and maintain a suitable internal database.

3.118 We set out the criteria for an appropriate database in paragraph 3.54.

Third-party databases

3.119 Service providers may also use a third-party hash database. Third parties build their databases by receiving submissions of hashed intimate images from survivors and victims, service providers, and/or other organisations and individuals.

3.120 The current leading third-party hash database for intimate image abuse is StopNCII.org, operated by the Revenge Porn Helpline.[68] This is the only publicly available third-party database that we are aware of, however, other databases could emerge. StopNCII.org allows survivors and victims of intimate image abuse to generate a hash from an image(s) or video(s) on their device. A hash is then uploaded to the database from their device, but not the image or video itself. This means the third-party database provider does not view or

[67] Response(s) to our June 2025 Consultation: British and Irish Law Education and Technology Association, p.14; Domestic Abuse Commissioner, p.6; Online Safety Act Network, p.16-17; Refuge, p.10; UCL Gender and Tech Research Lab, p.1; [✂].

[68] Revenge Porn Helpline is part of SWGfL, an international charity.

retain a copy of the image. StopNCII.org then shares the hash with participating service providers to detect and remove the images or videos, preventing them from being shared on their services. It partners with several large service providers that are at risk for intimate image abuse, including Facebook, Instagram, TikTok, Reddit, OnlyFans, Pornhub, X, and Snap Inc.[69]

3.121 StopNCII.org reported a 130% increase between 2023 and 2024 in the number of hashes submitted to its database,[70] and received over a million hashes of images within a three-year period. We consider this evidence crucial in determining the suitability of the third-party hash database for intimate image abuse content. It demonstrates that StopNCII.org is a well-known tool and regularly receives submissions of presumed intimate image abuse content.

3.122 StopNCII.org does not verify content submitted into its database. It allows survivors and victims of intimate image abuse to generate a hash from an image(s) or video(s) on their device, which is then uploaded to the database. This means the database provider does not view or retain a copy of the image. The only confirmation requested when submitting the hash is that the image or video is (a) depicting themselves; (b) intimate in nature; and (c) they are 18 years old or older in the image or video. However, there is no identity verification, which means this information could be falsely submitted by an individual or organisation. In practice, this means that the database may include hashes of images or videos which are either not intimate images or are consensually shared intimate images.

3.123 We understand that the images associated with the hashes on the database are not verified to safeguard the privacy of survivors and victims. We recognise the importance of privacy to the survivors and victims of intimate image abuse and understand that StopNCII.org does not want to increase the exposure of a survivor or victim's intimate image(s) unnecessarily by verifying the nature of the content itself. We also believe that survivors and victims are likely to be more comfortable and, by extension, more willing to submit hashes of intimate images to a database if the images are not being viewed or retained by the database provider. Ultimately, the objective of the tool is to provide a preventative measure to empower survivors and victims to participate in the online environment and reduce the risk of others viewing their intimate images.

3.124 However, we also recognise that there are risks associated with using a database that contains unverified hashes. Some service providers expressed concerns about this. For example, a service provider suggested that the importance of context in determining whether an image amounts to intimate image abuse means positive matches cannot be reasonably presumed to be violative content and may require additional review to confirm the detected content is violative.[71] Two stakeholders noted that unverified databases increased the risk of inaccuracy and false positives.[72] Another service provider further indicated concerns with the use of perceptual hash matching technology for intimate image abuse due to issues with match quality and accuracy of existing third-party databases.[73] Lastly, a service provider indicated hash matching was not equally effective for all services given different features and reported major reliability issues during initial integration of

[69] StopNCII.org. Industry Partners. [accessed 8 April 2026]
[70] SWGfL, 2024. StopNCII.org Being Used to Protect Over 1,000,000 Images Online. [accessed 5 March 2026].
[71] Response(s) to our June 2025 Consultation: [✂].
[72] Response(s) to our June 2025 Consultation: [✂]; [✂].
[73] Response(s) to our June 2025 Consultation: [✂].

StopNCII.org, resulting in an insignificant amount of true positives and a high amount of false positives.[74]

3.125 We have carefully considered this feedback and agree that the main risk of a database that contains unverified hashes is that it might not be suitably accurate. There is a risk that databases that contain unverified hashes detect content which is not intimate image abuse, that is, intimate images shared consensually or non-intimate images. We analyse these below and discuss how our measure is designed to mitigate these risks.

Risk of hash databases including consensually-shared intimate images

3.126 There is a risk that malicious actors could submit consensually-shared intimate images to a database (that contains either verified or unverified hashes) with the aim of having them removed by service providers.

3.127 Intimate images can be shared online consensually. Determining whether an image has been shared without consent can be challenging without additional context for a service provider. This context is determined in specific ways depending on the type of database:

a) If a service provider uses an internal database, under Ofcom's ICJG, a user report which suggests that the content has been posted non-consensually gives the provider reasonable grounds to believe there was no consent or reasonable belief in consent (absent evidence to the contrary).[75]
b) If a service provider uses a third-party database, we consider that where a hash of an image is included in the database, the service provider should treat this as reasonable grounds to believe that the image was shared non-consensually or without reasonable belief in consent (absent evidence to the contrary).

3.128 An example of 'evidence to the contrary' may be where the image itself is clearly an advertisement (e.g. for underwear).

3.129 Some stakeholders flagged that consensual adult content can be hard to distinguish from intimate image abuse and could be wrongly removed[76] and another expressed concern that hash matching could be used to censor legitimate content.[77]

3.130 We have considered the implications of this for individuals who typically appear in intimate content and other people or organisations who have a commercial interest in that content. This may include, for example, adult content creators, people advertising adult services, people on dating sites who post photos of themselves in their underwear, models, people in movies or other art depicting intimate scenes. It may also include retailers, film studios, art galleries etc. There is a risk that people maliciously target this kind of content or individuals by uploading it to a database to prompt its removal. For example, a commercial organisation could target its competitor's intimate content; or anti–sex work or anti-pornography campaigners could target sex worker or adult content creators' content. This could potentially harm those individuals' freedom of expression, including creative forms of speech, as well as their commercial interests. Some individuals, including some adult content creators, may be mainly or wholly reliant on this content for their livelihood and this kind of targeting could have a significant negative impact on the ability to generate

[74] Response(s) to our February 2025 Consultation: [✂].
[75] Illegal Content Judgements Guidance, paragraph 10.36(a).
[76] Response(s) to our June 2025 Consultation: Google, p.7; We Are Survivors, p.4. Response(s) to our February 2025 Consultation: [✂].
[77] Response(s) to our June 2025 Consultation: Open Rights Group, p.22.

income. This impact could disproportionately affect women given that sex workers[78] and adult content creators are more likely to be female.[79] There is also a risk that ideologically-motivated bad actors use this approach to target the content of groups such as the LGBT+ community, to make it more likely that such content is removed.

3.131 Fully preventing the risk of removing consensual intimate images would require identity verification of both the person depicted and the reporter, which is impractical, invasive for victims, and prohibitively resource intensive for most services.

3.132 Individuals whose consensual content is wrongly removed (e.g., adult content creators) can use existing appeals mechanisms to have their content reinstated.

3.133 As noted above, we are recommending that service providers review the first positive match of each hash on a database that contains unverified hashes to determine whether the content is intimate image abuse content. We acknowledge that this will not fully mitigate the risk that some consensually shared intimate images are taken down. However, we think it is not possible to fully mitigate that risk without significantly undermining the benefits of the measure. This is because mitigating this risk would involve verifying that each intimate image was in fact shared non-consensually. This would take considerable time and would be resource-intensive and very burdensome on victims and survivors. Therefore, we consider this would undermine the measure's ability to address intimate image abuse at scale.

3.134 There is also a risk that individuals or organisations may act maliciously by submitting a large volume of consensual intimate images to a third-party database with the aim of having this content removed from regulated services. We understand that StopNCII.org limits each submission to containing no more than 20 hashes. It also uses automated detection to prevent automated submissions that may result in a large volume of submissions. While this does not prevent malicious bulk submissions, we consider these elements establish frictions for individuals seeking to abuse the database.

3.135 Overall, we consider the risk to those sharing consensual intimate images is proportionate to the significant harms caused by intimate image abuse and the benefits of lowering reporting barriers for victims and having the images removed.

Risk of hash databases including non-intimate images

3.136 There is also a risk that malicious actors could submit non-intimate images to a database that does not verify hashes with the aim of having them removed by service providers.

3.137 For example, a person could add a political opponent or commercial competitor's non-intimate content to the database leading to potential impacts on democratic discourse or competition, respectively. Similarly to the risk mentioned above, ideologically-motivated bad actors could also target non-intimate content associated with groups such as the LGBT+ community. This activity may occur in bulk submissions, meaning individuals submit a large volume of non-intimate images, which would limit the impact of the database. In response to our consultation, South West Grid for Learning (SWGfL) stressed that "there is no

[78] The Home Office commissioned research that identifies sex work as a strongly gendered activity in England and Wales, with women forming the majority of those involved. Source: Home Office, 2019. Prostitution and sex work: The nature and prevalence of prostitution and sex work in England and Wales today. [accessed 05 May 2026].

[79] For example, based on data published in 2026, 78-84% of creators on OnlyFans were female. Source: OnlyFans Statistics: Trends, Demographics, and Growth Insights [accessed 27 April 2026].

evidence of malicious use or systemic false positives" in the StopNCII.org database.[80] It also indicated that the database has some safeguards to prevent the bulk submission of hashes.

3.138 We are not aware of evidence to date that StopNCII.org has experienced abusive use of the database at scale. However, we acknowledge that individuals or organisations may in future abuse the database and, as a result, cause significant impacts on their targets.[81] We therefore consider it necessary that appropriate mitigations are available for service providers to implement in order to reduce the risks and impacts on their users.

3.139 For example, as set out at 3.73 to 3.76, we are recommending that service providers review the first positive match of each hash to determine that the content identified is intimate image abuse content. By "first positive match", we mean the first time content is detected as a match for an individual reference hash under a given configuration of the hash-matching technology. This means the first positive match to a hash under the same perceptual matching configuration (or, if used for video content, an exact cryptographic match). We consider that this review is likely to be effective at identifying when content is not intimate image abuse content and therefore reduce the risk of misuse of the database for the suppression of non-intimate content.

Conclusion on intimate image abuse hash databases

3.140 While we recognise there are risks associated with hash databases, we consider that the deployment of hash-matching technology is likely to be effective to protect victims and survivors against intimate image abuse based on the existing databases, including both internal and third-party databases, available to service providers.

3.141 For an internal database, we consider that service providers should be able to use an internal database if they have the means to build and maintain a database of intimate image abuse hashes, which contains hashes of a significant number of original items of content.

3.142 For a third-party database, we consider the accuracy of third-party databases (including those that contain unverified hashes), and the proportionality of recommending them, depends on whether it is possible for service providers to deploy the database with appropriate safeguards against its misuse. In particular, the database needs to be deployed in a way which minimises the incentive (the incentive being having non-intimate image abuse or non-intimate images taken down or deprioritised) for malicious actors to misuse it. However, we consider that it is difficult to fully mitigate these risks without undermining the benefits of the measure. There need to be appropriate safeguards for those whose content will be wrongfully taken down or deprioritised. We consider there are appropriate safeguards, which we detail below.

3.143 Taking these factors into account, we conclude that service providers can use a database – either internal or third-party – that includes hashes of intimate image abuse content that meets the following criteria:

a) Contains hashes of a significant number of original items of content;
b) Be proactively updated with reasonable regularity; and

[80] Response(s) to our June 2025 Consultation: South West Grid for Learning, p.13.
[81] Evidence from [✂] suggests that [✂], source: Ofcom / [✂], 17 October 2024. Evidence from [✂] also suggests that [✂], source: Ofcom / [✂], 14 April 2025.

c) Be secured from unauthorised access, interference or (in the case of a database comprised of verified hashes) exploitation (whether by persons who work for that person or are providing a service to that person, or any other person).

3.144 We consider that the current leading hash database, StopNCII.org, meets these criteria. It should be used as the benchmark for determining whether other hash databases, including an internal database, are appropriate to use for this measure. In practice, this means that a service provider should not use an internal database, on its own, unless it is broadly equivalent to StopNCII.org in terms of quantity of original items of content hashed and quality of the hashes.

Appropriate configuration of perceptual hash matching technology is necessary to ensure its effectiveness

3.145 Service providers need to configure the perceptual hash matching technology to strike an appropriate balance between precision and recall. This is essential for ensuring the accuracy and effectiveness of the technology and helping to mitigate against the risks of false positives.

3.146 In the June 2025 Consultation, we proposed that service providers should consider certain principles when striking this balance. The balance should take into account the risk of harm from intimate image abuse, the proportion of content that is not intimate image abuse detected as a match by the technology, and the effectiveness of the systems and processes used by the service provider to address this content. We considered that these principles help ensure the accuracy and effectiveness of the technology and should therefore limit the risk of detection of content that is not intimate image abuse.

3.147 We received feedback from the Domestic Abuse Commissioner for England and Wales that recommended mandating ongoing review of technical safeguards, including hash matching precision and recall, with regular reporting to Ofcom.[82] While we acknowledge this suggestion, we ultimately consider it necessary to recommend service providers strike a balance between precision and recall to minimise risks to the accuracy, effectiveness and bias. Providers should review the balance between precision and recall at least every six months. We are also recommending that services maintain a written record of how this balance has been achieved, which Ofcom can request if relevant to the exercise of our functions.

3.148 We also received feedback from several industry stakeholders that the measure was overly prescriptive[83] and suggested service providers should be given greater flexibility over how they implement the measure.[84] However, we are satisfied that the measure provides sufficient flexibility to providers in this regard. We note that the way in which providers approach striking the balance between precision and recall will impact the extent to which the measure interferes with freedom of expression; this means the impact will vary depending on the approach taken by each service provider.

3.149 We considered whether it is possible to be prescriptive about how providers should balance precision and recall. However, we concluded that this is not possible. What constitutes an appropriate balance depends on a range of factors, including the individual service context,

[82] Response(s) to our June 2025 Consultation: Domestic Abuse Commissioner, p.2.
[83] Response(s) to our June 2025 Consultation: [✂]; [✂].
[84] Response(s) to our June 2025 Consultation: Computer and Communications Industry Association, p. 8; [✂]; [✂].

the characteristics and prevalence of intimate image abuse content on the service, and the downstream actions taken in response to detection. In addition, different implementations of perceptual hash matching may use different similarity measures, distance metrics or configuration parameters, meaning there is no single operating point that could be specified that would universally ensure accurate and effective detection of intimate abuse content. Further, providers are expected to review the performance of their perceptual hash matching systems over time, as the appropriate balance between precision and recall may change as content patterns, risks and system performance evolve.

3.150 Providers should therefore determine the appropriate balance between precision and recall by reference to evidence from their testing, evaluation and ongoing monitoring, recognising that the appropriate balance may change over time.

3.151 A decision to optimise for precision may be appropriate, for example, where the provider has assessed its service as being medium risk of intimate image abuse content and there is no evidence that the technology is failing to detect such content. This may particularly be the case where a relatively high proportion of detected content transpires to be false positives and the systems and processes in place to identify and correct false positives are less effective. Conversely, a decision to optimise for recall may be appropriate where the provider has assessed its service as being at high risk of intimate image abuse content, including where there is evidence that known intimate image abuse content is present on the service but remains undetected by the perceptual hash matching technology. This may be particularly relevant where a relatively low proportion of detected content transpires to be false positives and effective systems and processes are in place to identify and address false positives, including through secondary review of detected content.

3.152 In practice, service providers have reputational and commercial incentives to encourage them to strike a reasonable balance in their implementation of the measure. We address the measure's impact on human rights in the 'Rights assessment' section below.

3.153 We are therefore not setting out additional guidance on how service providers may configure their technology to strike a balance between precision and recall. The measure is designed to allow flexibility, and this approach enables the measure to be adopted more widely.

Service providers need to review matches of hashes, in certain circumstances, to mitigate accuracy risks

3.154 In our consultation, we said that service providers should treat a positive match against a database as reason to suspect the content is intimate image abuse. Service providers should therefore have regard to the existing user-to-user and search content moderation measures, including reviewing the content and actioning it accordingly.

3.155 We consider that a service provider may need to review content to determine whether it is intimate image abuse. The extent to which this is necessary depends on whether a database containing verified or unverified hashes is being used, and the provider's level of assurance in detection outcomes, as explained further below.

3.156 In this section, we set out two aspects of the review of content:

a) When a service provider needs to review matches; and,
b) How the service provider may review matches.

3.157 We set out the associated factors and our reasoning below.

When does a service provider need to review matches?

3.158 Our review of content recommendations depends on the type of hash database used by service providers and the provider's level of assurance in detection outcomes.

3.159 For a database that contains only verified hashes, we are recommending that service providers may not need to review matches for the purpose of determining whether detected content is an intimate image. This is because the hash will have already been verified to be intimate image abuse. This means, for example, that it may be appropriate for service providers to action matched content immediately – in accordance with appropriate moderation actions set out in our existing content and search moderation measures – if it matches against their internal database that contains only verified hashes. However, whether this will be appropriate depends on the level of assurance the provider has in the detection outcomes achieved by the hash matching technology. It may be appropriate to review the match again if, for example, the service provider has configured its perceptual matching parameters to prioritise recall. In this case, review of the match may be appropriate if the provider notices significant adversarial behaviour in relation to intimate image abuse.

3.160 For a database that contains unverified hashes (partially or entirely), we are now recommending that service providers should review the first positive match for every piece of detected content.[85] The purpose of the review is to determine whether the matched content is intimate image abuse. Any subsequent matches to that specific hash using the same matching parameters are unlikely to require further review. This means a service provider may treat the subsequent matches as illegal content, depending on the level of assurance the provider has in the detection outcomes achieved by the hash matching technology. Where the service provider has reconfigured the matching parameters in a way that increases recall prioritisation, it should have a lower level of assurance in the detection outcomes, and therefore this is more likely to require further review.

3.161 We have added this recommendation to our measure to mitigate the risks associated with service providers using an unverified third-party database.

How does a service provider review a match?

3.162 Service providers may use either automated technology or human moderators to review and determine whether positive matches are intimate image abuse.

Automated technology

3.163 Automated technology can be highly effective at reviewing a large volume of content quickly to determine whether each item is an intimate image. The use of such technology helps safeguard the privacy of individuals depicted in images (including survivors and victims'). However, it requires additional regular quality assurance to test the accuracy, effectiveness, and lack of biases of the technology. This will include the use of human review. We set out our reasoning on the need for human review for quality assurance in paragraphs 3.167 to 3.177.

3.164 Confidence in perceptual hash matches may be further strengthened through the application of additional automated technologies that are also widely available and in widespread use. This could include content classification technologies (such as nudity

[85] We are defining the first positive match as the first time there is a positive match to a particular hash within given perceptual parameters (or an exact match if using cryptographic hash matching for video content).

detection tools[86]), that can provide an additional indication that the matched content is consistent with the characteristics expected of intimate image abuse. While we recognise concerns that such technology may display racial biases,[87] we believe the risk posed by this is mitigated by our further recommendation regarding human review for quality assurance purposes. Service providers may use complementary approaches such as metadata, object or scene analysis, similarity scoring, or broader multimodal analysis of content and context. This will provide additional corroboratory signals, which may be used to support decision making. Human moderation is also accurate in determining the intimate nature of content. However, this type of review requires more resources (including time, people, and financial resources) in comparison to automated technology.

3.165 While automated technology can be highly effective at detecting intimate images, it may be less effective at identifying evidence that an intimate image has been consensually shared. This is a disadvantage of using automated technology for review. However, it is quicker and cheaper compared to human moderation. In this sense it may deliver greater benefits for survivors and victims. Furthermore, while there is a risk of such images being included in a database, it is not clear to what extent this will materialise in practice. Taking these matters into account, and the fact that users can appeal if they believe their content was incorrectly moderated, we conclude that using automated technology to determine whether detected content is intimate image abuse is proportionate.

Human review

3.166 As set out in the 'Explanation of the measure' section, for the purposes of this measure, there are two types of human review:

a) For content moderation purposes (namely, to make decisions about individual items of content); and
b) For quality assurance purposes (namely, to test, monitor and improve the accuracy of moderation tooling and processes).

3.167 In the June 2025 Consultation, we proposed a number of factors to be taken into account when determining the "appropriate proportion" of detected content to be subject to human review, namely:

a) The degree of accuracy achieved by automated systems and/or processes in use for the purposes of content moderation, if any (taking account of data from the service's complaints procedures);
b) The principle that content with a higher likelihood of being a false positive should be prioritised for review; and
c) The importance of understanding the purpose and context of detected content when determining whether it is intimate image abuse.

3.168 In response, two service providers stressed that intimate image abuse is not easily verifiable, which will call into question the effectiveness of human review.[88] Other stakeholders expressed concerns about the costs associated with human moderators reviewing content to determine whether it is intimate image abuse content.[89] In connection

[86] SightEngine, 2026. Detect nudity, porn, suggestive and explicit adult content in Images and Videos; PicPurify, n.d. Nudity detection API; Microsoft Learn, 2026. Detect Adult, racy, or gory content; AWS, 2026, Detecting inappropriate images – Amazon Rekognition. [all accessed 15 April 2026].
[87] Response(s) to our June 2025 Consultation: Glitch, p. 11.
[88] Response(s) to our June 2025 Consultation: Google, p.7; [✂].
[89] Response(s) to our June 2025 Consultation: [✂]; [✂].

with a similar aspect of another proposed proactive technology measure, a commentator noted that the measure does not concretely determine what proportion of detected content should be human reviewed, and argued that this may not be consistent with human rights law.[90]

3.169 In our view, it is not possible to be prescriptive about the proportion of detected content that should be subjected to human review. Rather, the proportion of detected content should depend on the factors we have set out above. For example, all other things being equal, where the provider has configured the perceptual hash matching technology to prioritise recall there is a greater risk of false positives and therefore a greater proportion of human review is needed to mitigate this. If the prevalence of intimate image abuse content on a service is very low, then even very accurate technology could result in a large amount of content being detected which is not intimate image abuse content. We expect that service providers will adjust the configuration of their technology or have human moderators review a higher proportion of matches, if they observe that lawful content is being detected. Regardless of how providers configure their technology, we consider that some human review is always appropriate for quality assurance purposes. We therefore consider that some human review is always appropriate, but the extent of it should depend on the factors identified in our measure.

3.170 We have decided to revise the factors that providers should take into account when deciding what proportion of detected content should be subject to human review, for the reasons set out below.

Combining and revising the first and second factors to, "the level of assurance the provider has in the detection outcomes achieved by the technology and any associated systems and processes"

3.171 The level of assurance a service provider has in the detection outcomes is closely related to the degree of accuracy achieved by automated systems and/or processes and the likelihood of false positive detections. Where the technology and associated systems and/or processes are highly accurate, the risk of false positives is generally lower. Conversely, where there is greater uncertainty as to the performance of the technology, the likelihood of false positives increases, and a higher proportion of human review will likely be appropriate. Indicators such as proximity to a matching threshold, the use of verified or unverified hash databases, and observed error rates should be understood collectively as signals of assurance rather than as standalone factors, informing the extent to which providers can rely on automated outputs without further verification. Service providers are expected to conduct ongoing monitoring and quality assurance (including human review) of the hash matching technology, to assess both the accuracy of the technology and relevant systems and processes, and the prevalence of incorrect detections over time. A service provider should take the outcomes of this monitoring and quality assurance into account when determining its level of assurance in the detection outcomes.

3.172 The appropriate level of human review depends not only on the performance of the hash matching technology but also on the effectiveness of other technical and human review layers. While further review may be warranted where content is detected, it is not necessarily appropriate to prioritise human review solely based on a higher likelihood of

[90] Cyberleagle, 2025. Cyberleagle: Ofcom's proactive technology measures: principles-based or vague? [accessed 13 April 2026]; similar points were raised in response to our June 2025 Consultation by Open Rights Group, p.23-26.

false positives. Providers should instead consider how human judgment can meaningfully add value beyond automated systems and prioritise review where it is most likely to affect content moderation outcomes. Increased quality assurance, including human review where appropriate, is particularly important where assurance is lower or where observed outcomes are inconsistent with expectations.

Removed factor: The importance of understanding the purpose and context of detected content when determining whether it is intimate image abuse.

3.173 Noting consultation respondents' concerns about the efficacy of human review in identifying intimate image abuse content (paragraph 3.169), we have removed the factor asking providers to have regard to the importance of understanding the purpose and context of detected content when determining whether it is intimate image abuse. Once a provider has determined that content is an intimate image, the only additional context a service provider may need is whether the depicted individual consented to the image being shared with others. A provider may have further evidence, but incorporating this into its decisions does not necessarily require human review.

Added new factors for moderation purposes: Severity of harm to those depicted in or encountering content and overall impact of an incorrect decision on the posting user or interested person

3.174 Human review of content usually takes more time than automated processes do. We consider that the appropriate proportion of human review for moderation purposes is affected by the potential severity of the harm. If, for example, a provider is noticing that intimate image abuse content is frequently going viral, or a misogynistic culture is developing on its service, it might be appropriate to reduce the proportion of content which is subject to human review for content moderation purposes temporarily in order to address the problem. We address human review for quality assurance purposes in paragraphs 3.167 to 3.177.

3.175 We consider service providers should consider the impact of an incorrect decision that the content is illegal content on the posting user or interested person when determining the proportion of detected content to human review. This is because the harm caused by an incorrect decision will vary.

Guidance

3.176 We have decided to issue guidance which sets how to apply all the factors in more detail. The draft guidance is set out here.

Conclusion on the benefits and effectiveness of hash matching for intimate image abuse

3.177 The Act requires us to consider the accuracy, effectiveness and lack of bias of proactive technology that we recommend.

3.178 Overall, based on our assessment of the benefits and effectiveness of the measure, as set out above, we consider the use of hash matching technology is effective at detecting intimate image abuse and thereby reducing the spread of this harmful content online. This measure will be effective for both verified and unverified hash databases.

3.179 While we recognise there are risks associated with the use of both verified and unverified hash databases, we consider they can be appropriately mitigated by designing certain safeguards into the measure.

3.180 Overall, we have not identified significant risks of algorithmic bias arising from hash matching technology itself, given its deterministic nature. However, outcomes may still be influenced by the composition and coverage of hash databases and by the use of other automated technologies alongside hash matching (for example, nudity detection tools). The quality assurance and human review recommendations, as set out in this measure, are intended to help providers identify and address any such issues in practice.

Costs and impacts

3.181 In our June 2025 Consultation, we assessed the costs of implementing hash matching technology to detect intimate image abuse across a range of service providers. We indicated that the main costs were likely to consist of:

a) One-off costs related to building a hash-matching system;
b) Annual costs relating to maintaining a hash-matching system;
c) Costs of software, hardware and data; and
d) Costs of reviewing matches, moderating content and dealing with appeals.

3.182 The assessment further noted that the incurred costs would vary according to a service providers number of users, technical complexity and risk of intimate image abuse.

3.183 Some stakeholders supported our cost assessment, including civil society organisations and service providers.[91] Other stakeholders disagreed and, specifically, flagged concerns with regards to certain assumptions impacting the cost-related decisions.[92] We detail this feedback below.

3.184 Having considered this feedback, we have made minor amendments to our assessment of the costs and impacts associated with this measure. These revisions are:

a) Revising our view of the likelihood of cost synergies for service providers already implementing hash matching for other content types; and
b) Updating our assessment to reflect that service providers can now use cryptographic hash matching to detect video-based intimate image (as set out in paragraphs 3.56 to 3.57) – as it may reduce costs for video hash matching.

3.185 We have also revised the quantified cost estimates to reflect updated salary data from the Office for National Statistics (ONS).[93]

Cost synergies with other hash matching measures

3.186 In the June 2025 Consultation, we indicated that our costs estimates may be an overestimate for service providers who had already implemented hash matching technology, including where they had implemented it for different types of illegal content. For example, service providers that are also in scope of the Illegal Content Codes measure (ICU C9) that recommends hash matching to detect child sexual abuse material may incur

[91] Response(s) to our June 2025 Consultation: British and Irish Law Education and Technology Association, p.15; Evangelical Alliance, p.5; [✂]; Philippine Survivor Network, p.16.
[92] Response(s) to our June 2025 Consultation: [✂]; [✂]; [✂]; Pinterest, p.4; [✂].
[93] See Annex 4 for further detail on economic assumptions and analysis.

lower costs as the automated moderation systems and processes have already been implemented.

3.187 However, we received feedback that disagreed with this position. Two stakeholders indicated that implementing hash matching systems and processes for intimate image abuse would be costly regardless of whether the same technology was being used to detect other types of illegal content, such as child sexual abuse material.[94] They said that using hash matching technology to detect each type of illegal content would require a specialised approach and expertise and, as a result, require new systems and substantial investment.

3.188 After careful consideration, we still consider that some service providers will incur lower costs because they have already adopted hash matching processes to address other types of illegal content. However, this may not be universal for all in-scope service providers. This will depend on the technology and tooling that a service provider uses for its existing hash matching capabilities and the nature of its service, such as encrypted services. In some cases, the technologies used by service providers cannot be applied to intimate image abuse content, which will likely result in the provider incurring the full costs of implementing the measure.

3.189 Based on this position, we have updated our impact assessment to reflect that service providers may not incur lower costs if they are using hash matching technology for other types of illegal harms. However, this revision to the assessment does not change our overall proportionality assessment. We set out at consultation the assumption that service providers would not have existing hash matching technology and, as a result, would incur the implementation costs of hash matching technology for the first time. This means the proportionality assessment already weighed the full implementation costs of the proposed measure against the benefits of using hash matching technology to detect intimate image abuse.

Costs of third-party databases

3.190 In the June 2025 Consultation, we proposed that user-to-user services that are high risk for intimate image abuse and either (a) whose principal purpose is the hosting or dissemination of regulated pornographic content; or (b) are file-sharing and file-storage services should source intimate image abuse hashes from an appropriate third-party database. All other in-scope service providers could use either an internal database or a third-party database.

3.191 We set out that the cost for service providers to access the only currently available third-party database, StopNCII.org, starts at £1,000 a year and increases based on provider revenue. We indicated a preference for this approach, as it provided an attainable solution for smaller service providers to use hash matching processes to detect intimate image abuse.

3.192 We received feedback from two service providers highlighting the high costs to access third-party databases of intimate image abuse hashes,[95] with one provider warning that substantial fees may discourage platforms from working with third-party database providers.[96]

---

[94] Response(s) to our June 2025 Consultation: Computer and Communications Industry Association, p.7-8; [✂].
[95] Response(s) to our June 2025 Consultation: [✂].
[96] Response(s) to our June 2025 Consultation: [✂].

3.193 We also received feedback from SWGfL, which provides the StopNCII.org database, that set out specific evidence on implementation time for integrating the database, which is one of the sources of costs to service providers. It referred to a select committee meeting in which OnlyFans had confirmed that it integrated StopNCII.org in under 80 hours and that ongoing maintenance took 4–5 hours per month.[97] These implementation times are lower than the minimum software engineering time we assumed in our cost estimates. Given there may be activities not captured in the OnlyFans figure, we consider this to support our cost assessment.

3.194 A service provider also expressed concern about [✂].[98] For this reason, service providers would need to review positive matches to determine the content is intimate image abuse. The associated costs may become more relevant over time as the databases are more widely used by survivors and victims' but also potentially subject to greater abuse.

3.195 We have decided to amend the measure to recommend that all service providers use third-party databases to detect intimate image abuse on their services, unless they have an equivalent or better internal database. We recognise that the use of a third-party database does incur costs for service providers, which may be higher depending on their revenue. However, we consider this recommendation to be proportionate. We set out our reasoning for recommending third-party-databases, including addressing their availability, in paragraphs 3.120 to 3.126.

3.196 This change does not alter our assessment of the costs for using a third-party database as set out at consultation. We have therefore decided not to revise our assessment.

Costs of hash matching technology

3.197 In the June 2025 Consultation, we proposed that service providers use perceptual hash matching technology to detect intimate image abuse on their services. In response, a service provider indicated that perceptual hashing for videos was more expensive than cryptographic hash matching technology.[99]

3.198 However, as set out in paragraphs 3.56 to 3.57, we have made a minor amendment to allow service providers to use cryptographic hash matching to detect video-based content when there is not an available database that enables perceptual hash matching.

3.199 While we consider this amendment will help mitigate the potential additional costs of perceptual hash matching for video content, we still encourage service providers where possible to use perceptual hash matching for the detection of both image and video-based content where possible. Our assessment of costs relating to the use of perceptual hash matching for intimate image abuse has not changed.

Total costs of the measure

3.200 We have decided to maintain our view of the range of costs estimates set out in the June 2025 Consultation. We have updated our quantified estimates to bring them in line with the latest ONS salary data.[100] We now expect that a relevant service provider would incur the costs summarised in Table 1.

---

[97] Response(s) to our June 2025 Consultation: South West Grid for Learning, p.20.
[98] Response(s) to our June 2025 Consultation: [✂].
[99] Response(s) to our June 2025 Consultation: [✂].
[100] See Annex 4 for further detail on economic assumptions and analysis.

Table 1: Estimated costs to service providers of implementing hash matching for intimate image abuse

Component	Estimated cost[101]
Build cost (one-off)	£21,000 to £369,000
Maintenance costs (annual)	£5,000 to £92,000
Software, hardware and data (annual)	From £2,000 for small services
Human moderation (annual)	£500 to £1,000[102]

Build cost (one-off)

3.201 The one-off costs associated with setting up a hash matching system to detect intimate image abuse will primarily consist of labour costs. Depending on the type of service, the provider can either choose to build a fully in-house hash matching system or implement a hash matching pipeline using pre-existing components or reference architectures.

3.202 For our recommendation on hash matching for child sexual abuse material (ICU C9), we estimated that it could take 2-18 months of full-time work to set up a hash matching system. We expect the upper bound (18 months) to reflect the time it might take for a large service provider to build a fully in-house hash matching system, while the lower bound (2 months) may be more reflective of a smaller service provider implementing hash matching using pre-existing components or reference architecture (e.g., using an open source pipeline such as Hasher-Matcher-Actioner[103] deployed in a cloud environment). We estimate that equivalent time is likely to be needed from other staff such as product managers, analysts and lawyers to support in developing and establishing a relevant policy.

3.203 We expect the labour requirements to be the same for implementing hash matching processes to detect intimate image abuse.

3.204 Based on this, we estimate that the one-off set up costs could range from £21,000 to £369,000, depending on the size and complexity of a service. We expect smaller service providers would incur costs towards the lower end of this range.

Ongoing maintenance cost

3.205 The ongoing maintenance costs will primarily consist of labour costs associated with activities such as applying updates, reviewing the hash matching technology's performance and configuring the technology (at least every six months), ingesting new hashes from the third-party database, and integrating with new service functionalities. This will require input from software engineers as well as other professional occupation staff.

3.206 In line with our standard assumption, we assume that the annual maintenance costs would be 25% of the one-off costs of setting up a hash matching system for intimate image abuse.[104]

---

[101] The cost estimates on labour costs, as set out in this table, are rounded to the nearest £1,000 with the exception of figures below £1,000 which are rounded to the nearest £100.
[102] Based on a service with 80 matches for review per year, see paragraph 3.213 for more details.
[103] Meta (Clegg, N.), 2022. Meta Launches New Content Moderation Tool as It Takes Chair of Counter-Terrorism NGO. [accessed 23 March 2026].
[104] See Annex 4 for further detail on the economic assumptions and analysis.

3.207 Based on this, we estimate that the annual cost of maintaining a hash matching system for intimate image abuse would range from £5,000 to £92,000.

Ongoing software, hardware and data costs

3.208 The ongoing software, hardware and data costs to service providers will mainly consist of the costs associated with accessing hash databases and API solutions from third-party organisations where relevant. As these costs generally tend to scale with a service's user base and/or revenues, we expect large service providers to pay more than smaller service providers.

3.209 As described above, StopNCII.org charges service providers to access its hash database; the amount a service provider must pay is determined based on their annual revenue and can start at around £1,000 per year.

3.210 There are also free open-source third-party solutions such as Hasher-Matcher-Actioner available to service providers, although these solutions require access to computing power to check for content on services. We therefore assume the annual cost of software, hardware and data could start at around £2,000 for a small service provider

Ongoing content and search moderation and appeal costs

3.211 The ongoing content and search moderation and appeal costs will mainly consist of labour costs which will include:

a) Costs associated with the use of either automated technologies or human moderators to review first positive matches of content;[105]
b) Costs associated with human review of an appropriate proportion of detected content to conduct quality assurance of the automated technology used for this measure; and
c) Costs associated with resolving appeals. We expect the costs associated with appeals are likely to scale with the amount of content that is detected by the hash matching technology.

3.212 We expect in many cases that service providers would use automated technology to review first positive matches of content, which will incur significantly lower costs for service providers. However, the measure does allow service providers to use human moderators to review the first positive match, if preferred, which will incur higher costs. For example, we estimate that a service that receives around 80 matches per year for review could spend around £500 to £1,000 on human moderation per year.[106] Some services may have a much lower or higher number of matches. We generally expect that the ongoing cost of content and search moderation will correlate with the risk and amount of intimate image abuse on a service and therefore would be in line with the benefits of our proposed measure.

[105] This is for unverified hashes. The measure sets out different review steps for review of verified hashes, set out from paragraph 3.69 to 3.72.

[106] Based on our stakeholder engagement, we understand that it is difficult for service providers to specifically identify the time moderators may spend reviewing content detected as intimate image abuse. As a proxy, we instead use global data on the number of content enforcements made and the number of moderators employed by Snapchat, to estimate a ratio for the number of moderators per report (0.0001). Applying this ratio, we expect a service with 80 matches for review to employ the full-time equivalent of around 0.01 moderators. We then apply our assumptions about the salary of a content moderator and the uplift for staff costs, as set out in Annex 4 'Further detail on economic assumptions and analysis'.

3.213 Service providers will also have to use human moderators to review a proportionate amount of detected content to quality assure the automated technology used for this measure. This will incur costs for service providers.

3.214 We outline our recommendations for human review and automated technologies in the 'Explanation of the measure' section.

Indirect impacts

3.215 There is a risk that a hash matching system for intimate image abuse may lead to the removal of consensual intimate images. We discuss in paragraph 3.138 the risks that this could pose to individuals or businesses who may have their content incorrectly removed. For adult content creators, models or pornography companies, this may result in disruption that could negatively impact their ability to generate income.

Rights impact assessment

3.216 In this section we consider how the measure may interfere with human rights, and whether such interference is lawful. As part of this assessment, we have considered stakeholder consultation responses to the June 2025 Consultation which raised issues relating to human rights.

3.217 Ofcom must act lawfully, rationally and fairly and must comply with the provisions of the European Convention on Human Rights (ECHR) when recommending measures in a code of practice, by virtue of section 6 of the Human Rights Act 1998. This includes, among other things, ensuring there is no unjustified interference with the right to privacy in Article 8 (right to respect for private and family life) and 10 (right to freedom of expression) of the ECHR. Pursuant to the ECHR, any interference with the right to privacy or the right to freedom of expression must be in accordance with[107] or prescribed by[108] law and necessary in a democratic society in pursuit of a legitimate aim including prevention of crime or disorder, protection of health or morals and the protection of the rights and freedoms of others. In either case, to be necessary, the restriction must correspond to a pressing social need, and it must be proportionate to the legitimate aim pursued.

3.218 Under the Act, Ofcom must also design its measures for inclusion in a code of practice in light of the importance of protecting users' privacy and freedom of expression and, if appropriate, incorporate safeguards into its measures for the protection of privacy and freedom of expression.[109] [110]

Freedom of expression

Potential interference with freedom of expression

3.219 The measure we are recommending is linked to our wider suite of content/search moderation Code measures which ultimately lead to the take down, removal or deprioritisation of content. We have considered the impact of our measure in this context.

---

[107] Article 8 of the ECHR.
[108] Article 10 of the ECHR.
[109] Paragraph 10 of Schedule 4 to the Act.
[110] Privacy is further defined in paragraph 10(3) as including protecting users from a breach of any statutory provision or rule of law concerning privacy that is relevant to the use or operation of a user-to-user service or a search service including but not limited to any provision or rule concerning the processing of personal data, that is the requirements of the UK General Data Protection Regulation ('UK GDPR')).

3.220 Parliament has defined intimate image abuse content as priority illegal content for the purposes of the Act. It has required that proportionate steps be taken to prevent users from encountering this type of content and to ensure it is taken down or deprioritised when service providers become aware of it.

3.221 However, the way that providers implement the measure may also affect legal content. An interference with users' freedom of expression is likely to arise when lawful content is detected (a 'false positive') and is taken down (for user-to-user services) or deprioritised (for search services) as a result.[111]

3.222 False positives may arise in several ways which we discuss in the 'Benefits and Effectiveness' section above:

a) Hashes of non-intimate images or videos are included in a database. This is a particular risk for a database that contains unverified hashes such as StopNCII.org. This might include, for example:

   i) commercial competitors uploading images associated with their competitors (e.g., logos) with a view to impacting their competitors' livelihood;[112]
   ii) people uploading political or religious content that they do not agree with, which can impact highly protected political or religious speech. Open Rights Group expressed concern about, for example, uploading a picture of a candidate during an election period to secure the removal of that image from searches;[113] or
   iii) people targeting groups that they are ideologically opposed to by uploading content associated with those groups (e.g., LGBTQ+ content).

b) Hashes of intimate images or videos which are consensually shared are included in a database. Again, this is more likely to be a risk for databases that contain unverified hashes. This might include, for example, adult content creators uploading consensually shared intimate images of their competitors, which can impact those competitors' livelihood.

c) Hashes of intimate image abuse content are included on a database, but they match with hashes of lawful content due to the configuration of the hash matching parameters. This is a risk for a database that contains verified or unverified hashes (and databases that contain both).

---

[111] In principle, there could also be a more general 'chilling effect' if some users were to avoid using services which had implemented the measure due to concerns about adverse impacts on users' rights (as noted by Open Rights Group in its response to our June 2025 Consultation in relation to our proposed terror hash matching measure, p.24). However, many UK users already use services which have implemented perceptual hash matching. Any such effect therefore seems unlikely to be significant. Similarly, there could be a chilling effect if providers ceased to provide services in the UK or geo-blocked content so that it cannot be encountered in the UK as a result of the measure. As we know that many providers are already implementing hash matching, and UK users in any event have many thousands of public and private services to choose from, we do not consider such impacts to be likely or significant.

[112] Where the measure has an impact on an individual's income or an entity's revenue, this is likely to engage rights to the protection of property under Article 1 of the First Protocol to the ECHR.

[113] Response(s) to our June 2025 Consultation: Open Rights Group, p.22. While this is a risk, we think the measure is unlikely to have a systematic impact on political content. This is because given the nature of intimate content is not generally speaking likely to include political speech (in contrast to, e.g., terrorist content). We nevertheless recognise the risk that bad actors could in theory exploit a database that contains unverified hashes in the way that Open Rights Group suggested.

3.223 Several stakeholders raised concerns relating to false positives[114] and the risk of removing lawful content.[115] Two stakeholders raised that human review of detected intimate content cannot establish whether or not it was consensually shared.[116] Individual stakeholders expressed concerns about mass censorship,[117] [✂][118] and limiting access to educational and minority viewpoints[119] and information that is important for vulnerable groups.[120] We do not consider that the measure amounts to or will result in mass censorship but agree that the measure could interfere with freedom of expression. We set out below how our Codes mitigate this risk and why we consider any interference to be proportionate.

Prescribed by law

3.224 To be prescribed by law, Ofcom's power to recommend the measure must have a clear basis in domestic law. That legal framework must be 'adequately accessible' to the individuals concerned, and 'foreseeable'. We address these tests below.

3.225 The measure has a clear basis in domestic law stemming from the Act. In this Statement, we explain how we have met the statutory tests for recommending measures relating to proactive technology. The measure is accessible: both the Act and our Codes are published.

3.226 As to the foreseeability of the measure's impact on regulated services and users, a legal commentator[121] noted that another proactive technology measure that we consulted on does not specify thresholds for precision or recall or how much detected content should be subject to human review. They argued this raises questions about the foreseeability of the measure. We consider this comment to be relevant to this measure too.

3.227 The measure is clear about the kind of technology that service providers should use and the type of content to which it relates. Where the measure results in the takedown or deprioritisation of illegal content, this is foreseeable: users should expect that illegal intimate image abuse content will be taken down or deprioritised.

3.228 However, we recognise that the measure may have an impact on legal content too. In this regard, as mentioned in paragraphs 3.146 to 3.154, the measure does give providers flexibility to decide the appropriate balance between precision and recall (provided they take certain factors related to false positives into account). We acknowledge this does not provide absolute clarity to users or services about precisely which items of content should be detected, and that the impact on freedom of expression may vary according to the approach taken by each service provider.

3.229 However, for the reasons set out in paragraphs 3.146 to 3.154 we do not think it is possible to be prescriptive about the appropriate balance between precision and recall in the measure for each one of the services to which it applies. We do not consider that this means the measure is insufficiently foreseeable, because of the context within which

---

[114] Response(s) to our June 2025 Consultation: Computer and Communications Industry Association, p.8; [✂].
[115] Response(s) to our June 2025 Consultation: [✂]; Name Withheld 4, p.4; [✂].
[116] Response(s) to our June 2025 Consultation: We are Survivors, p.4. Response(s) to our February 2025 Consultation: [✂].
[117] Response(s) to our June 2025 Consultation: [✂].
[118] Response(s) to our June 2025 Consultation: [✂].
[119] Response(s) to our June 2025 Consultation: [✂].
[120] Response(s) to our June 2025 Consultation: [✂].
[121] Cyberleagle, 2025. Ofcom's proactive technology measures: principles-based or vague? [accessed 21 April 2026]; similar points were raised in response to our June 2025 Consultation in relation to our proposed terror hash matching measure by Open Rights Group p.23 -26

providers will be making these judgements, and because UK law and our Codes provide for safeguards.

3.230 In practice, service providers have reputational and commercial incentives that should encourage them to strike an appropriate balance in their implementation of the measure. While service providers have incentives to limit the amount of intimate image abuse on or accessed via the service, they also have incentives to limit the costs of moderation (including human review) and the handling of user complaints where content is wrongly taken down or deprioritised. We anticipate that providers will continuously adjust the technical configuration of their hash matching system in light of their experience.

3.231 In relation to the proportion of detected content subjected to human review, as set out at paragraphs 3.167 to 3.177 we have amended the factors that providers should take into account when determining this. We have set out in paragraph 3.170 why we consider it is not possible to be prescriptive about the proportion of detected content that should be subject to human review. We have also given some further guidance on the matters providers should consider in determining the appropriate proportion of human review. We remain of the view that service providers should have a degree of flexibility on this matter, subject to considering these factors, and do not consider that this means that the measure is insufficiently foreseeable.

Necessary for a legitimate aim

3.232 As set out in paragraph 3.218, any interference with the right to freedom of expression must pursue one or more legitimate aims.[122] In considering the measure's impact on individuals' human rights, our starting point is to recognise that there is a substantial public interest in reducing the online prevalence and dissemination of image-based intimate image abuse. That public interest relates to the prevention of the crime of intimate image abuse, the protection of morals of those who would be exposed to such content and the protection of the reputation, rights, freedoms and safety, of those depicted in such content.[123] Intimate image abuse has profound negative psychological, emotional, professional and relational impacts on survivors and victims, compounded by the challenges they face in getting content removed, and the fear that it might resurface online.

3.233 We explained at paragraphs 3.94 to 3.98 why reactive complaints handling is not sufficient to address the harm caused: we are not aware of other effective way to protect victims and survivors of intimate image abuse besides using proactive technology to detect relevant content quickly. We also concluded that perceptual hash matching is more effective than cryptographic matching at protecting users from harm, notwithstanding the greater risk of false positives.[124] It is against this background that we consider the proportionality of the measure's interference with freedom of expression.

Proportionality

3.234 We received several stakeholder comments expressing concern about the proportionality of the measure's interference with freedom of expression. While we recognise these

---

[122] National security, territorial integrity or public safety; the prevention of disorder or crime; the protection of health or morals; the protection of the reputation or rights of others; preventing the disclosure of information received in confidence; and maintaining the authority and impartiality of the judiciary.
[123] Articles 8(2) and 10(2) of ECHR.
[124] For reference, for video-based content we have revised the measure to recommend that service providers use perceptual hash matching in relation to video-based content where an appropriate database enables this, but if there is no such database, they may use cryptographic hashing.

concerns, we remain of the view that the measure is proportionate to address the serious harm to both victims and other users caused by intimate image abuse.

3.235 While there may be an interference with freedom of expression, it is a relatively limited one. It arises where there are mistakes, not as a matter of any particular kind of content being systematically targeted. Incorrect detection as intimate image abuse does not lead to criminal or civil sanctions on those whose content is taken down or deprioritised. The content concerned may be taken down or deprioritised on one service, but this will not necessarily be the case on other services. Finally, the measure is implemented by service providers, and not by the state: it is not a tool through which to limit political speech.

3.236 In addition, the design of the measure, and the wider framework of law, includes a number of safeguards which protect users' freedom of expression. We have highlighted the main safeguards below before turning to our overall assessment of proportionality.

Review of first matches of unverified hashes

3.237 Unverified hashes have a higher risk of being hashes of non-intimate images. To safeguard against the risk this poses, we are recommending that where a service provider detects content using an unverified hash, it should review detected content to confirm it is an intimate image at the first positive match. The purpose of this review is to determine whether the matched content is intimate image abuse content and therefore reduce the risk of lawfully-shared images (which might include political images or commercial logos) being taken down or deprioritised.

3.238 Service providers may use automated technology or humans to conduct this review. In paragraph 3.164 we noted some of the advantages and disadvantages of automated technology, including that it may be less effective than humans at identifying that an intimate image may have been consensually shared (e.g. if it was clearly an underwear advert). However, for the reasons set out in that paragraph, we concluded that using automated technology in this way is proportionate.

Appeals

3.239 An important safeguard in protecting lawful content from being removed is the appeal process.[125] In circumstances where content has been taken down or deprioritised, a user or interested person can appeal to have their content reinstated. This is particularly important for users whose consensual content has been wrongly removed for example, sex workers, adult content creators or others who regularly post intimate content depicting themselves. We discuss this risk in more detail at paragraphs 3.131. Appeals provide a mechanism for those individuals to inform the service provider that the content was shared consensually. This safeguards the speech (and potentially commercial interests) of such individuals.

3.240 We acknowledge that there will be an interference with freedom of expression while content remains unavailable or less visible pending any appeal. However, as set out in paragraphs 3.135 to 3.136 we do not think it is possible to fully mitigate this risk without significantly undermining the benefit of the measure. For example, as noted in paragraph 3.132 to determine whether an intimate image included on a database was shared consensually would require identity verification of both the person depicted and reporter. This approach is impractical, invasive for victims, and prohibitively resource intensive for

[125] See Code Measure D in the Illegal Content user-to-user Codes and Illegal Content search Codes.

most services. It would also undermine the objective of addressing intimate image abuse at speed.

Human review

3.241 The measure recommends human moderators review an appropriate proportion of detected content, taking into account various factors. We discuss this in detail from paragraph 3.172 to 3.176. These factors are:

a) The level of assurance the provider has in the detection outcomes achieved by the technology and any associated systems and processes, as indicated by the ongoing monitoring, evaluation and quality assurance (including human quality assurance) of the performance of the technology; and
b) For content and search moderation purposes:
    i) the potential severity of the harm to those depicted in or encountering intimate image abuse content; and
    ii) the overall impact of an incorrect decision that the content is illegal content, on the user who generated, uploaded or shared the content.

3.242 We have also issued guidance to assist providers in applying this aspect of the measure.

3.243 We consider that the human review factors will help mitigate the risks arising from lack of verification in databases that contain unverified hashes (where they are used), risks arising from the use of other automated technology in the content moderation queue, risks arising from matches with low confidence scores and risks arising from any lack of quality assurance of the configuration of the technology. For example, where there is greater uncertainty as to the performance of the technology (or other technologies used as part of content moderation), the likelihood of false positives increases, and a higher proportion of human review will likely be appropriate. This helps to reduce the risk of providers incorrectly taking down or deprioritising content that is not intimate image abuse.

Review of detected content where verified hashes are used (or after the first match with an unverified hash)

3.244 Where providers detect content using a verified hash, or after a 'subsequent match' with an unverified hash, we recommend that providers may take moderation action without further reviewing the content; however, this depends on the level of assurance the provider has in the detection outcomes achieved by the hash matching technology. In practice this means that the lower the provider's level of assurance in the detection outcomes, the more likely it is that further review of detected content will be appropriate. This helps to safeguard against the risk that content detected using matching parameters configured to place greater weight on recall (which are more likely to produce false positives) is moderated.

Feedback and hash removal

3.245 As set out at paragraph 3.81 to 3.84, the measure sets out that a service provider should take reasonable steps to secure the removal of a hash from the hash database where the provider has determined that the hash is not intimate image abuse content. This mitigates against the risk that the same content will be detected and taken down or deprioritised in the future. This may be particularly important for users whose consensual intimate content was taken down by the service provider (potentially due to a bad actor uploading it to a database), because it enables those users to reduce the risk of this occurring in the future.

Precision and recall configuration

3.246 Our measure also aims to mitigate against the risk of false positives by recommending that providers strike an appropriate balance between precision and recall in configuring their technology. Service providers should take into account certain factors in striking this balance, including the risk of harm from intimate image abuse, the proportion of content which is not intimate image abuse detected as a match by the technology, and the effectiveness of the systems and processes used by the service provider to address this content. These principles aim to help ensure the accuracy and effectiveness of the technology and should therefore limit the risk of detection of content that is not intimate image abuse. We are also recommending that providers review the balance between precision and recall periodically which helps to ensure that providers continue to refine their approach in light of experience.

Other safeguards

3.247 Other existing measures and guidance provide further safeguards against disproportionate interferences with freedom of expression. For example, the ICJG should assist providers when making judgements about whether content is intimate image abuse. We have also recommended that certain providers have internal content/search moderation policies,[126] resource their content/search moderation function to give effect to these policies[127] and provide training to content/search moderators.[128] We acknowledge that these measures only apply to large or multi-risk services[129] or large general search services or multi-risk services[130] and therefore do not apply to all providers in scope of the intimate image abuse hash matching measure. We set out in our December 2024 Statement why we consider this to be proportionate. We also consider that, where a service is not large or multi-risk, the impact of incorrect content moderation decisions is overall likely to be lower, although we acknowledge that the impact may nevertheless be significant for some users.

Data protection law as a freedom of expression safeguard

3.248 Nothing in our Codes or guidance, insofar as it relates to the processing of personal data, either supersedes or derogates from the requirements of data protection legislation. These include the principle of accuracy[131] when processing personal data (e.g., the creation of data indicating that it is reasonable to infer that a user has committed the intimate image abuse offence). Further, the UK General Data Protection Regulation (UK GDPR) restricts the circumstances in which decisions can be made based solely on automated processing of personal data, where the decision has legal or similarly significant effects for the relevant individual.

3.249 Service providers must continue to comply with applicable data protection law in implementing this measure and should have regard to relevant guidance issued by the Information Commissioner (ICO) from time to time.[132] We therefore consider that service providers' data protection obligations will support accurate, effective and free from bias

[126] ICU C3 / ICS C2.
[127] ICU C6 / ICS C5.
[128] ICU C7 / ICS C6.
[129] For user-to-user services.
[130] For search services.
[131] Article 5(1)(d) of the UK GDPR.
[132] For example, at the time of this statement: ICO, 2026. Automated decision-making, including profiling. [accessed 23 April 2026].

decision-making about moderation of user-generated content and therefore safeguard users' rights to freedom of expression in respect of that content.

Overall assessment of proportionality

3.250 When considering the proportionality of the measure, our starting point is that intimate image abuse is a serious harm and reducing its dissemination has clear benefits, including protecting victims' and survivors' privacy rights. We have been particularly mindful of the risk of removal of non-intimate images and consensual intimate images. We have designed the measure to mitigate this impact (as set out above), but it cannot be entirely removed. We need to strike fair balance between the interests of these groups. Our recommendations for the review of content, as well as appeals, provide the most important safeguards in this regard.

3.251 Overall, taking into account the very serious harm caused by intimate image abuse, and the safeguards we have built into the measure, we consider that the risks to freedom of expression that we have identified in this section are proportionate to the very significant benefits that we believe the measure will create in terms of reducing the dissemination and perpetration of intimate image abuse.

Privacy

Data protection considerations in our assessment

3.252 In response to the June 2025 Consultation, the ICO provided feedback on our approach to assessing the impact on privacy and data protection. It recommended that "Ofcom's privacy and data protection rights assessment provides a more detailed overview of the data protection considerations relevant to the use of proactive technologies".[133] 5Rights Foundation recommended our guidance should include greater detail about services providers' responsibilities regarding privacy regulation, including UK GDPR and the UK's Age Appropriate Design Code.[134]

3.253 We have taken these comments into account. The Act requires us to design measures in the light of the importance of protecting the privacy of users[135] and, where appropriate incorporate safeguards for the protection of privacy.[136] Accordingly, while Ofcom's remit is not to regulate data protection or provide guidance on compliance with data protection legislation, we will assess the impact of our measures against the main data protection principles where we consider these are also relevant to privacy. As part of this rights assessment, we have considered some important data protection principles, namely, fairness, transparency, accuracy and data minimisation.[137] We have also considered user rights related to automated decision-making.

General comments about privacy impacts

3.254 We received stakeholder responses relating generally to privacy and data protection concerns. For example, an individual respondent warned that image hashes involved processing large volumes of sensitive data, creating data protection challenges, and that

[133] Response(s) to our June 2025 Consultation: The Information Commissioner (ICO), p.12.
[134] Response(s) to our June 2025 Consultation: 5Rights Foundation, p.9.
[135] This is defined in Schedule 4 paragraph 10(3) as meaning "means protecting users from a breach of any statutory provision or rule of law concerning privacy that is relevant to the use or operation of a user-to-user service or a search service (including, but not limited to, any such provision or rule concerning the processing of personal data)" (underlining added).
[136] Paragraph 10 of Schedule 4 to the Act.
[137] Article 5(a), (c), and (d) of the UK GDPR.

third-party data sharing could violate user privacy.[138] Chayn argued that perceptual hashing could undermine the privacy that keeps survivors safe and could be exploited by undemocratic or repressive regimes to target marginalised groups and individuals including activists, journalists and human rights defenders.[139]

3.255 We consider that our recommended measure is sufficiently clear for providers to implement and appropriately balances the rights of individuals affected by the measure. We set out our reasons for this below.

Interference with privacy

Reasonable expectation of privacy

3.256 The measure recommends that service providers deploy hash matching technology to detect intimate image abuse communicated publicly, not content communicated privately.[140] Under the Act, Ofcom is unable to recommend the use of proactive technology to content communicated privately.[141] Individuals posting public content have a lesser expectation of privacy, and therefore the extent of any possible intrusion into privacy is correspondingly reduced.[142] Further, where they post content for an illegal purpose, there can be no reasonable expectation of privacy in relation to it.[143]

3.257 Individuals depicted in content are more likely to have a reasonable expectation of privacy if they have not consented to the sharing of the image. However, where they have submitted a hash of their content to a hash database, they will be doing so on the understanding that the hash database provider and any service providers using the hash database will use it to detect their intimate images and ensure they are taken down or deprioritised. This requires the providers to process the data related to the intimate image and its hash. The persons depicted would also be submitting their intimate images to the hash database on the basis that content matching the hash may need to be reviewed by a service provider to confirm it is an intimate image abuse match and be actioned in the service provider's content moderation process. As such, the data processing would be occurring at the request of the person depicted.[144]

3.258 Service providers should note that intimate images are likely to constitute special category data, as they are likely to relate to the sex life or sexual orientation of a person. As such, they must ensure the processing is carried out in accordance with Article 9 of the UK GDPR and should have regard to the ICO's guidance on special category data.[145]

---

[138] Response(s) to our June 2025 Consultation: [✂].
[139] Response(s) to our June 2025 Consultation: Chayn, p.3.
[140] Paragraph 13(4) of Schedule 4 to the Act.
[141] Paragraph 13(4) of Schedule 4 to the Act.
[142] The ICO also noted that compliance with data protection obligations is required regardless of whether content is communicated publicly or privately: The Information Commissioner (ICO) response to our June 2025 Consultation, p.12.
[143] Sutherland v HM Advocate for Scotland [2020] UKSC 32.
[144] The hash of intimate image abuse content can still be considered personal data where the service provider is able to link the hash back to a specific image or video and, therefore, to an identifiable individual. See (at the time of this Statement): ICO, n.d. Pseudonymisation. [accessed 5 May 2026]. Where this is the case, the individual depicted will continue to have data protection rights in relation to the image or video, and service providers should have regard to their obligations: see Your data protection rights | ICO [accessed 5 May 2026].
[145] ICO, n.d. Special category data. [accessed 21 April 2026].

Interference

3.259 The measure recommends that providers use a suitable database for the purpose of hash matching. The extent to which the inclusion of hashes on a database interferes with privacy varies depending on the database. In respect of StopNCII.org, only hashes are stored in the database as opposed to the intimate image abuse content, therefore the privacy intrusion will be minimal. In relation to internal databases, in cases where the intimate image abuse is stored in the database this is only for the purpose of securing that the content is less visible and therefore protecting victims' and survivors' privacy rights.

3.260 Hash matching involves the automated processing of content communicated publicly on the service. Service providers may also use automated technology to review detected content as part of their content moderation systems. We suggested in our consultation that, where services carry out automated processing in accordance with data protection law, that processing should have a minimal impact on users' privacy. The ICO commented that compliance with data protection law does not, in itself, guarantee that the privacy impact on users will be minimal.[146] We agree; nevertheless, we remain of the view that the hashing of users' content in and of itself involves only minimal intrusion with their privacy rights under the UCHR.[147] The processing is limited to creating a hash of the content and (potentially) analysing the content using technology to determine whether it is an intimate image. Users would have been informed about the provider's use of proactive technology in its terms of service for user-to-user services and in its publicly available statement for search services.[148] Further, only their publicly-communicated content will be hashed, and we expect users to have a reduced expectation of privacy in relation to such content. While we acknowledge there is an interference, we do not consider these processes constitute a significant interference with the privacy rights of people depicted in the content, where the overall purpose is to prevent the harm caused by intimate image abuse and protect the privacy of victims and survivors.

3.261 The measure recommends that an appropriate proportion of detected content be human reviewed for both quality assurance and content / search moderation processes. Therefore, because of the measure, more content (particularly from databases that contain unverified hashes) may be reviewed by human moderators. This is a greater intrusion on users' ECHR privacy rights, and the rights of people depicted in any content than might otherwise have occurred. The extent of the intrusion will also vary depending on the circumstances. For example, the intrusion will be lesser where a person has submitted an image of themselves to a database (or complained about content depicting themselves) and is therefore likely to expect matched content to be reviewed.[149] However, the intrusion will be greater where a person has uploaded a hash of an image of another person without that person's knowledge or consent, or where an image which is not the original image hashed, happens to match.

---

[146] Response(s) to our June 2025 Consultation: The Information Commissioner (ICO), p.15.
[147] We note that there are specific data protection considerations for automated data processing. See (at the time of this Statement): ICO, 2026. Automated decision-making, including profiling. [accessed 5 May 2026].
[148] This is also relevant to the data protection principle of transparency. The ICO has issued guidance in relation to this principle: ICO, 2025. Principle (a): Lawfulness, fairness and transparency. [accessed 23 April 2026].
[149] This is reflected on StopNCII's website, which states "[i]f someone tries to upload a matching image on one of our participating companies' platforms, they will review the content on their platform to check if it violates their policies and take action accordingly". Source: StopNCII.org, 2026. FAQ. [accessed 21 April 2026].

In accordance with law

3.262 We explained above why our measure satisfies the requirement that any interference be 'prescribed by law'. For the same reasons we consider that the measure is 'in accordance with law' for the purpose of Article 8.

Necessary for legitimate aim and proportionality

3.263 In paragraphs 3.233 to 3.234 above, we set out the legitimate aims that the measure pursues, and the serious harm caused by intimate image abuse. We have also explained that we are not aware of other effective ways to protect victims and survivors of intimate image abuse, besides using proactive technology to detect relevant content quickly. We have acknowledged the ways in which this might adversely impact individuals' privacy above, and explain below how we have designed the measure with safeguards to ensure that this impact is proportionate to the benefits of the measure.

3.264 It is important to recognise that the measure's overall objective is privacy-protecting. Intimate image abuse has a negative impact on the privacy rights of victims and survivors depicted in the intimate image abuse. The reduction in intimate image abuse because of this measure should mitigate that negative impact.

Human review

3.265 As noted above, to the extent that our measure recommends human review, this may be viewed as more privacy intrusive than automated review. However, for the reasons discussed above, human review is an important safeguard against taking down or deprioritisation of lawful content; it therefore helps to protect the freedom of expression of users posting lawful content. It also mitigates the risk of inaccurate personal data being created when content is incorrectly determined to be illegal and supports the fairness principle in data protection law.[150] Effectively identifying when images are not intimate images is also necessary to deter misuse of a database containing unverified hashes and protect the rights of those who might otherwise be targeted using it. Overall, we consider that the interference posed by human review is proportionate and that there is no less intrusive way to achieve our policy objective.[151]

3.266 We had regard to the ICO's suggestion that when determining the proportion of content subject to human review, providers should pay particular consideration to the use of a database containing unverified hashes.[152] In our view, this proposed change to the measure is not necessary. This is because we are recommending that the first match on a given set of parameters with an unverified hash is reviewed. We are also recommending that for subsequent matches (and matches with verified hashes) the service should consider its level of assurance in the detection outcomes in deciding whether to further review the content. Finally, we are recommending that providers take certain matters into account when deciding what proportion of detected content to human review, including the level of assurance that the provider has in its content moderation systems and processes as a whole. These safeguards mitigate against various risks, including the risk (identified by the

[150] Response(s) to our June 2025 Consultation: The Information Commissioner (ICO), p.17. The ICO noted that Ofcom's recommendation that services should review an appropriate proportion of detected content is an important privacy and data protection safeguard to guard against content being erroneously detected as intimate image abuse and support the accuracy and fairness principles of data protection law.
[151] Paragraph 11.82 of the June 2025 Consultation.
[152] Response(s) to our June 2025 Consultation: The Information Commissioner (ICO), p.17-18.

ICO) that non-intimate images will be detected due to the use of a databases containing unverified hashes.

Precision and recall configuration

3.267 The measure also recommends that providers configure their technology to strike an appropriate balance between precision and recall, taking into account certain factors. Two stakeholders expressed concern that false positives can lead to potential privacy violations[153] and may breach the accuracy and fairness principles of data protection law.[154] We agree and believe that our recommendation in relation to the configuration of the technology, in combination with the review element referred to above, helps to mitigate the risks of false positives. It also supports the accuracy and fairness principles of data protection law, and provides an important privacy safeguard against content being erroneously moderated.

Appeals

3.268 Users and interested persons can appeal if they believe a provider has taken action against them or their content based on an incorrect illegal content judgement. This enables them to mitigate the risk of their lawful content being taken down or deprioritised in the future, and to rectify inaccurate personal data held by the service provider. We consider this safeguard to be particularly important for individuals who might be the victims of 'malicious' submissions to hash matching databases, for example sex workers, adult content creators or LGBTQ+ individuals (see further paragraph 3.131 and 3.138).

Other safeguards

3.269 Service providers must have particular regard to the importance of protecting users' privacy (among other things) when complying with their safety duties.[155] Users and affected persons can complain to service providers if they consider that the provider is not complying with this duty.[156] This is a further safeguard provided by the Act.

Data protection law as a privacy safeguard

3.270 As set out in paragraphs 3.253 to 3.254 in relation to freedom of expression, the UK GDPR has the effect of protecting privacy to the extent that an interference with privacy may result from a breach of the rules on the processing of personal data.

3.271 Service providers must continue to comply with applicable data protection law in implementing this measure and should have regard to relevant guidance issued by the ICO from time to time. In particular, the data protection principles of data minimisation,[157] accuracy[158] and fairness[159] will serve to protect privacy, and providers should have regard to them (and all other relevant principles) when processing personal data as part of the measure. Automated decision-making is only permitted where service providers have implemented certain safeguards for the data subject's rights, freedoms, and legitimate interests in accordance with the UK GDPR.[160]

---

[153] Response(s) to our June 2025 Consultation: [✂].
[154] Response(s) to our June 2025 Consultation: The Information Commissioner (ICO), p.16.
[155] Section 22 of the Act.
[156] Section 21(4)(b)(iii) and section 32(4)(b)(iii) of the Act.
[157] Article 5(1)(c) of the UK GDPR.
[158] Article 5(1)(d) of the UK GDPR.
[159] Article 5(1)(a) of the UK GDPR.
[160] See Article 22(3) of the UK GDPR and section 14 of the Data Protection Act 2018; see also (at the time of this statement): ICO, 2026. Automated decision-making, including profiling. [accessed 23 April 2026].

Overall assessment of proportionality

3.272 Overall, we consider the measures interference with privacy rights to be proportionate, taking account of the significant harm caused by intimate image abuse, the other rights impacted by the measure and the safeguards identified above.

General monitoring

3.273 Some consultation respondents argued that Ofcom's proposed proactive technology measures risked amounting to prohibited "general monitoring" under EU law.[161] Google warned that expanding the range of content such measures apply to could impose general monitoring obligations beyond Parliament's intent.[162] Open Rights Group raised concerns that the measure could create a conflict of laws and that Ofcom had not assessed the impact of this on cross-border digital trade.[163] European Digital Rights Organisation noted potential conflicts for services operating across both the UK and EU, potentially leading to withdrawal from the UK market.[164]

3.274 The Act empowers Ofcom to recommend proactive technology measures for compliance with safety duties about illegal content without limiting this power to specific categories of illegal content. However, this power is constrained: measures must apply only to publicly communicated content, be proportionate, and take account of privacy and freedom of expression. (A full explanation of how this measure meets the various statutory tests is set out in the 'Statutory Tests' section below.) While we acknowledge that proactive technologies can interfere with privacy, we consider that interference capable of being proportionate in certain circumstances, as explained in this statement.

3.275 The Code measures apply only to service providers' UK operations or as they affect UK users.[165] It is for providers to manage compliance with other jurisdictions' laws. Although some services may choose to restrict UK access, we consider this unlikely to occur, or to be significant if it did, given the widespread existing use of hash matching and the wide availability of alternative services for UK users.

Who the measure applies to

Scope of applicable service providers

3.276 In our June 2025 Consultation, we proposed that hash matching to detect intimate image abuse should be used by the following service providers:

a) Providers of large user-to-user services that are medium risk for intimate image abuse.[166]
b) Providers of user-to-user services that are high risk for intimate image abuse and either:
    i) whose principal purpose is the hosting or dissemination of regulated pornographic content;[167] or

[161] Under EU law, Member States are prohibited from imposing a general obligation to monitor the information which they transmit or store: Article 8 of the Digital Services Act (EU) (Reg 2022/2065).
[162] Response(s) to our June 2025 Consultation: Google, p.4.
[163] Response(s) to our June 2025 Consultation: Open Rights Group, p.20.
[164] Response(s) to our June 2025 Consultation: European Digital Rights Organisation, p.1.
[165] Schedule 4 to the Act, paragraph 11.
[166] 'Large' refers to any user-to-user service with more than 7 million monthly UK users.
[167] Regulated pornographic content is pornographic content which excludes content which consists only of text or text accompanied by one of more of the following: identifying content consisting only of text, identifying content which is not pornographic, a GIF which is not pornographic, or an emoji or other symbol.

ii) are file-sharing and file-storage services; or
iii) have more than 700,000 monthly active UK users.

c) Providers of large general search services.[168]

3.277 We received support from several stakeholders on the scope of applicable service providers,[169] including from respondents that supported the measures application to file-sharing and file-storage service providers[170] and large general search services.[171] One stakeholder indicated that the measure should only apply to file-sharing or storage services when a files are being shared with others rather than being hosted.[172] In this context, it is important to note that our measure only applies to content communicated publicly. Some stakeholders also expressed concerns about whether smaller service providers are in scope of the measure[173] and one recommended extending the measures to all service providers.[174]

3.278 After careful consideration, we have decided to confirm the scope of service providers recommended to implement this measure. We detail our reasoning for applying the measure to each type of service provider below.

Recommendation for certain high-risk services to use third-party databases

3.279 At consultation, we proposed that certain user-to-user services that are high risk for intimate image abuse should source intimate image abuse hashes from an appropriate third-party database. These service providers are those:

i) Whose principal purpose is the hosting or dissemination of regulated pornographic content; or
ii) Who are file-sharing and file-storage services.

3.280 All other in-scope service providers could use either an internal database or a third-party database.

3.281 We received feedback from some stakeholders that highlighted the importance of reducing the need for survivors and victims to report their intimate images to more than one database provider or service provider.[175] One stakeholder emphasised that having to make multiple reports can be exhausting and re-traumatising for survivors and victims.[176] We consider it important to help facilitate reducing the barriers for survivors and victims' getting their non-consensual intimate images being shared online, including the number of times content needs to be reported.

---

[168] 'Large' refers to any search service with more than 7 million monthly UK users.
[169] Response(s) to our June 2025 Consultation: Evangelical Alliance, p.5; Pinterest, p.4.
[170] Response(s) to our June 2025 Consultation: British and Irish Law Education and Technology Association, p.14; Canadian Centre for Child Protection, p.11-12; Internet Watch Foundation, p.32; National Police Chiefs' Council, p.7; Philippine Survivor Network, p.16; Refuge, p.11; South West Grid for Learning, p.20.
[171] Response(s) to our June 2025 Consultation: British and Irish Law Education and Technology Association, p.15; Evangelical Alliance, p.5; Philippine Survivor Network, p.16; Refuge, p.11; South West Grid for Learning, p.21-22.
[172] Response(s) to our June 2025 Consultation: Open Rights Group, p.21-22.
[173] Response(s) to the June 2025 Consultation: Barnardo's, p.10-11; Online Safety Act Network, p.16; Refuge, p.8-9.
[174] Response(s) to our June 2025 Consultation: Department of Justice for Northern Ireland, p.9.
[175] Response(s) to our June 2025 Consultation: Online Safety Act Network, p.16-17; [✂].
[176] Response(s) to our June 2025 Consultation: Refuge, p.10.

3.282 Additionally, we have determined that the accuracy of third-party databases (including unverified), and the proportionality of recommending them, depends on whether it is possible for service providers to deploy the database with appropriate safeguards against its misuse.

3.283 Taking this into consideration, we have amended the measure to recommend that all service providers included in the measure, regardless of their risk level, use third-party databases to detect intimate image abuse on their services, unless they have an equivalent or better internal database.

3.284 Based on this decision, we have concluded that it is no longer necessary to include a recommendation for specific high-risk service providers to use a third-party database.

Assessment of each type of in-scope service provider

3.285 We detail our assessments for including specific types of service providers within the scope of this measure below.

Medium-risk large user-to-user service providers and high-risk user-to-user service providers with over 700,000 monthly active UK users

3.286 We did not receive substantial feedback opposing the inclusion of medium risk large user-to-user services and high-risk user-to-user services with over 700,000 monthly active UK users within the scope of this measure.

3.287 We consider it proportionate to include these services within the scope of this measure. This is based on the wide range of user-to-user service types where intimate image abuse can take place. The Illegal Harms Register describes the types of service where intimate image abuse was likely to be found, including social media services, blogging platforms, image boards and community forums.[177] A risk-based approach is appropriate in this case to identify the user-to-user services which do present a significant risk of intimate image abuse.

3.288 For large user-to-user services in particular, the Revenge Porn Helpline found that Instagram, Facebook (including Messenger) and X accounted for 19%, 15% and 6% of reported cases of intimate image abuse in 2024 respectively. Intimate image abuse was also found on messaging services such as Snapchat, Telegram and TikTok, which combined accounted for 10% of reports.[178]

3.289 For these reasons, we consider it proportionate to recommend this measure for medium-risk large user-to-user services and high-risk user-to-user services with over 700,000 monthly active UK users.

Large general search service providers

3.290 We received feedback from some civil society organisations that expressed their support for applying the measure to large general search services.[179] However, other stakeholders expressed their view that the use of hash matching processes by general search services may result in a higher risk of detecting and removing content other than non-consensual intimate images.[180] Two individuals expressed that this risk is due to the volume of content

[177] Illegal Harms Register, p.172-173.
[178] Revenge Porn Helpline (Papachristou, K.), 2024. Revenge Porn Helpline Report 2024. [accessed 16 March 2026].
[179] Response(s) to our June 2025 Consultation: Evangelical Alliance, p.5; Philippine Survivor Network, p.16; Refuge, p.11.
[180] Response(s) to our June 2025 Consultation: Open Rights Group, p.22.

on search services and the accuracy of hash matching technology.[181] A service provider also suggested Ofcom had not provided sufficient evidence that hash matching technology is effective for detecting intimate image abuse on search services.[182]

3.291 We understand that large general search service providers already use hash matching technology to detect intimate image abuse on their services.. In particular, we understand that Microsoft Bing is a partner of StopNCII.org and uses the database to detect intimate image abuse. Microsoft has reported that it identified and acted on 268,899 images which were being returned in image search results in Bing following a pilot implementation of StopNCII.org's database.[183] Additionally, since we published our consultation, Google has also announced a partnership with SWGfL to use the StopNCII.org hash database.[184] We consider the current use of hash matching technology by large general search services, for the detection of intimate image abuse, reinforces that this technology is effective for search functionalities, particularly when large volumes of content are involved.

3.292 Our position is that large general search services are at risk for widespread circulation of intimate image abuse via their image search results. We take this view regardless of those services' risk level for intimate image abuse.[185] Image search functionalities can display intimate image abuse content from other online services including deepfake imagery.[186] It can allow perpetrators to search for this type of content and reverse image search functionalities can be misused to identify survivors and victims.[187] These capabilities further present significant risks for exposing users to intimate image abuse.

3.293 For these reasons, we consider it proportionate to recommend the use of hash matching technology on large general search services regardless of risk level. In particular, we consider hash matching technology will help these service providers automate their content moderation practices with regards to the detection of intimate image abuse.

File-sharing and file-storage service providers

3.294 We received feedback during the consultation from three individuals that expressed concerns about the inclusion of file-sharing and file-storage services. They indicated that the measure may be ineffective and raised issues around privacy implications and the risks of misusing databases and false positives or negatives.[188]

3.295 We consider do not agree that the measure will be specifically less effective or pose greater challenges for file-sharing or file-storage services. In fact, we consider that applying the measure to file-sharing and file-storage service providers will be beneficial in reducing the spread of intimate image abuse online.

3.296 There is evidence that suggests these services are at particular risk for the sharing of intimate image abuse. For example, in one case, 940 images were found in a collection on a

---

[181] Response(s) to our June 2025 Consultation: [✂]; [✂].
[182] Response(s) to our June 2025 Consultation: Google, p.7.
[183] Microsoft (Gregoire, C.), 2024. An update on our approach to tackling intimate image abuse. [accessed 23 March 2026].
[184] Google (Hunt, G.), 2025. Developing a new partnership to combat non-consensual intimate imagery on Search. [accessed 23 March 2026].
[185] To note, we did not identify search services as high risk for intimate image abuse in our Register. However, we do consider these services to have functionalities that can increase the spread of and, as a result, risks associated with intimate image abuse.
[186] Illegal Harms Registers, p.410.
[187] Illegal Harms Register, p.412.
[188] Response(s) to our June 2025 Consultation: [✂; [✂]; [✂].

file-sharing site and were able to be shared as a single file.[189] As a result, the use of hash matching technology to detect intimate images being shared non-consensually on these services will help reduce their spread online. We set out additional evidence to justify this position in the Illegal Harms Register.[190]

3.297 It is also important to note that the file-sharing and file-storage services in scope of this measure are likely to be in scope of the hash matching for child sexual abuse material measure outlined in the December 2024 Statement due to overlaps in risk profiles. There may be a risk that some of these services withdraw from the UK due to the costs of these measures. However, we consider there may be synergies in the costs between the hash matching measures for child sexual abuse material and intimate image abuse. We also consider the potential impact on the market to be proportionate to the harms being reduced as only file-sharing and file-storage services that are high risk are affected.

3.298 For these reasons, we consider it proportionate to recommend that file-sharing and file-storage service providers that are high-risk for intimate image abuse use hash matching technology to detect intimate image abuse.

Pornography service providers

3.299 We did not receive substantial feedback opposing the inclusion of pornography services that are high risk for intimate image abuse within the scope of this measure.

3.300 We consider it proportionate to apply this measure to pornography services that are high risk for intimate image abuse. Pornography services hosted 25% of the content that was reported to the Revenge Porn Helpline in 2024.[191] The Illegal Harms Register further describes the increased risk of intimate image abuse on pornography services.[192]

3.301 There may be a risk that small pornography services withdraw from the UK due to the costs of this measure. However, because this measure only impacts pornography services that are high risk of intimate image abuse, we consider the potential impact on the market for adult content to be proportionate to the harms being reduced.

3.302 For these reasons, we consider it proportionate to recommend this measure for high-risk service providers whose principal purpose is the hosting or dissemination of regulated pornographic content.

Combined impact assessment

3.303 In the previous section, we assessed the impacts of the recommended measure on a standalone basis and concluded that it is proportionate. Here we consider the combined impact of the recommended measure in combination with existing Illegal Content and Protection of Children Code measures.[193] We consider whether this measure has distinct

[189] Huber, A. and Ward, Z., 2024. Non-consensual intimate image distribution: Nature, removal, and implications for the Online Safety Act, *European Journal of Criminology, 22*(1). [accessed 23 March 2026].
[190] Illegal Harms Register, p.174.
[191] Revenge Porn Helpline (Papachristou, K.), 2024. Revenge Porn Helpline Report 2024. [accessed 16 March 2026].
[192] Illegal Harms Register, p.173.
[193] While this measure is only an addition to the Illegal Content Codes, we recognise and have considered as part of our assessment that many services in scope of these codes will also be in scope of our Protection of Children Codes.

benefits to contribute to the overall package and whether the overall impact on services is proportionate, particularly for smaller services.

3.304 In our June 2025 Consultation, our Combined Impact Assessment provisionally concluded that the proposed package of measures was proportionate. We received a range of views from stakeholders on this, including stakeholders that supported the assessment, but also those that disagreed with our package of measures. Some stakeholders specifically disagreed with our package of measures for smaller services, with some arguing it could place a disproportionate burden on smaller or low-risk providers and create a negative impact on competition in the market.[194]

3.305 Our view is that the benefits of this recommended measure for using hash matching to detect intimate image abuse does not substantially overlap with existing measures. This measure substantially builds on existing content moderation measures to recommend a specific technology that is effective in identifying known intimate image abuse and we consider it will bring a distinct benefit in targeting this harm. We have taken a risk-based approach and only recommended this measure for smaller services when they are at high risk of intimate image abuse, and in these cases, we consider the cost of this measure to be proportionate to the benefits.

Growth duty

3.306 We have considered the impact of recommending this measure, in addition to existing code measures, further to our growth duty. While this measure is likely to increase regulatory compliance costs for service providers, including some smaller services, we consider that innovation, investment and competition are unlikely to be significantly impacted. We also do not expect any significant adverse impact on adjacent markets or the wider economy from these measures.

3.307 We note that the measure is targeted at services which are at risk of this illegal content and is expected to have a significant positive impact on safety, while our assessment indicates that any adverse effects on growth are likely to be minimal. After assessing the impact considering the growth duty, and the wider economic impacts of our decision, we remain of the view that recommending this measure is proportionate.

Conclusion

3.308 Taking into consideration this analysis, we have decided to confirm our measure to recommend that specific service providers use hash matching processes to detect intimate image abuse on their services. This confirmation includes the amendments set out in paragraph 3.31.

3.309 We consider the measure to be proportionate and that it will significantly reduce the spread of intimate image abuse online.

[194] As this feedback relates to the whole package of measures which we proposed in our June 2025 consultation, these responses will be addressed in more detail when we publish our Combined Impact Assessment as part of the statement for the remaining package of measures.

A1. Impact assessments and statutory tests

Our approach to the impact assessment

A1.1 In developing our recommendations, we are obliged to carefully consider the impact and proportionality of the measure on both service providers and online users.

A1.2 We assess the impact of our recommended measure based on the following factors:

a) The prevalence and impact of the harm the measure is combatting;
b) The efficacy of the measure in combatting this harm;
c) The direct and indirect costs of the measure;
d) The impact the measure would have on privacy and freedom of expression; and
e) Any risks associated with the measure.

A1.3 We also consider the proportionality of applying the measure to either all, or specific subsets, of service providers. This remains consistent with our approach to impact assessments in the December 2024 Statement.

A1.4 However, for the purpose of this statement, our assessment took into consideration the anticipated impacts, costs, and benefits of the measure over and above those which are expected to arise from existing measures.

A1.5 We consider that our cumulative impact assessment captures the additional costs and benefits because of the recommended measure set out in this statement.

A1.6 The detailed impact assessment, including the cumulative assessment, is set out in Chapter 3 of this statement.

Equality impact assessment

A1.7 We are also required to carefully consider the impact of our recommendations on individuals and communities with protected characteristics. This is called an equality impact assessment (EIA).

A1.8 We conduct the EIAs in accordance with Ofcom's legal obligations under:

a) Section 3 of the Communications Act 2003 (the 2003 Act)
    i) To further the interests of citizens in relation to communications matters; and
    ii) Secure the adequate protection of citizens from harm presented by content on regulated services;
b) Section 149 of the Equality Act 2010, requiring public authorities to have due regard to the need to:
    i) Eliminate unlawful discrimination, harassment and victimisation;
    ii) Advance equality of opportunity; and
    iii) Foster good relations between different groups; and
c) Section 75 of the Northern Ireland Act 1998, which extends these obligations to include political opinion, marital status, and dependants.

A1.9 We did not receive stakeholder comments on our EIA related to this measure.

A1.10 In preparing this statement, we have carefully considered the potential impacts of our measure on people sharing protected characteristics (including age, sex, race, disability, religion or belief, sexual orientation, gender reassignment, pregnancy and maternity, and marriage and civil partnership, as well as dependants and political opinion in Northern Ireland). Due to overlap with the protected characteristics in the Equality Act 2010 and the Northern Ireland Act 1998, we have had regard as part of our Equality Impact 3 Assessment to the vulnerability of those whose circumstances put them in need of special protection and the needs of persons with disabilities and of the elderly.

A1.11 The measure described in this statement is designed to reduce the risk of the priority illegal harm known as 'intimate image abuse'. It includes stronger protections against intimate images shared non-consensually. We consider our measure will primarily have positive impacts for people with protected characteristics – and in particular women and girls, given that intimate image abuse is predominantly perpetrated against them.

A1.12 We have considered whether the measure could have negative impacts on equality of opportunity or the fostering of good relations. We have considered the risks associated with the use of databases that contain unverified hashes, including that they could be exploited by people who want certain content associated with protected characteristics (e.g., LGBT+ content) taken down. We have also considered the risk that people could submit consensually shared intimate content to a database, such as content associated with sex workers or adult content creators. There is a risk this content could be submitted by competitors seeking to cause commercial harm or who ideologically oppose, for example, sex work. As noted above, these impacts are more likely to be felt by women, given that they make up most sex workers and adult content creators.

A1.13 We have built in safeguards to mitigate these risks. For example, providers should review the 'first positive match' of content matched with an unverified hash; this mitigates the risk of non-intimate content being taken down or deprioritised. We recognise that consensually shared intimate content is particularly vulnerable to being incorrectly moderated. To some extent this risk is mitigated by users' ability to appeal if they think their content is wrongly taken down; however, we recognise that this is not necessarily a quick process. We ultimately concluded that fully preventing the risk of removing consensual intimate images would undermine the benefits of the measure, which is itself likely to benefit groups with protected characteristics given that most victims and survivors of intimate image abuse are women and girls.

Welsh language assessment

A1.14 The Welsh language has official status in Wales. To give effect to this, certain public bodies, including Ofcom, are required to comply with the Welsh Language Standards.[195] Accordingly, we have considered:

a) The potential impact of our policy proposals on opportunities for persons to use the Welsh language;

[195] The Welsh language standards with which Ofcom is required to comply are available on our website. Source: Ofcom, 2023. Our work in the Welsh language. https://www.ofcom.org.uk/about-ofcom/structure-and-leadership/welsh-language [accessed 10 April 2026].

b) The potential impact of our policy proposals on treating the Welsh language no less favourably than the English language; and
c) How our proposals could be formulated to have, or increase, a positive impact, or not to have adverse effects or to decrease any adverse effects.

A1.15 We did not receive stakeholder comments on our Welsh language impact assessment related to this measure.

A1.16 Ofcom's powers and duties in relation to online safety regulation are set out in the Online Safety Act 2023 (the Act) and must be exercised in accordance with our general duties under section 3 of the t 2003 Act. In formulating our proposals in this consultation, where relevant and to the extent we have discretion to do so in the exercise of our functions, we have considered the potential impacts on opportunities to the Welsh language and treating the Welsh language no less favourably than English. We have not identified any likely negative impacts on the Welsh language arising from our decision to introduce this measure. We recommended in our December 2024 Statement on Protecting People from Illegal Harms Online (December 2024 Statement) and our April 2025 Statement on Protecting Children from Harms Online, that service providers should have regard to the needs of their UK user base in considering what languages are needed to effectively resource their content moderation, complaints handling, terms of service and publicly available statements. These ancillary measures will continue to apply in relation to the intimate image abuse hash matching measure. For example, service providers have duties to include information in their terms of service about proactive technology (including hash matching) used to comply with an illegal content safety duty; in doing so they should have regard to the language needs of their UK user base.

Statutory tests

A1.17 In designing the Codes, the Act requires us to have regard to a number of principles and objectives, set out in Schedule 4 to the Act. The 2003 Act also places a number of duties on Ofcom in carrying out our functions, including requiring us to have regard to the risk of harm to citizens presented by content on regulated services.

A1.18 In Chapter 3 of this statement, we set out our final recommended measure for service providers. We consider that our recommended measure meets the requirements set out in Schedule 4 to the Act and section 3 of the 2003 Act. In this section, we take each of the relevant requirements in turn and briefly set out how we have met them in reaching our recommended measure.

Duties and principles

The Communications Act 2003

> Section 3(1): It shall be the principal duty of Ofcom, in carrying out their functions: a) to further the interests of citizens in relation to communication matters; and b) to further the interests of consumers in relevant markets, where appropriate by promoting competition.

A1.19 The measure will mitigate the risk of harm resulting from intimate image abuse. This includes harm to various groups of citizens and consumers, in particular harm to victims and survivors (who are predominantly women and girls) and harm to viewers encountering this illegal content. We have also considered the rights of different groups affected by the measure.

Section 3(3): In performing their duties under subsection (1), Ofcom must have regard in all cases to (a) the principles under which regulatory activities should be transparent, accountable, proportionate, consistent and targeted only at cases in which action is needed, and (b) any other principles appearing to Ofcom to represent best regulatory practice.

A1.20 In the interest of transparency, accountability and fairness, we consulted stakeholders and have set out the evidence and reasoning informing our decision. Our impact assessment of the measure considers effectiveness, costs, rights, and so on, and explained why we consider the measure proportionate. Our measure is informed by our assessment of the risk posed by intimate image abuse content. We have targeted the measure at the kinds of services which we think should be deploying them because this would lead to the greatest benefits, given the risks they pose.

Section 3(2)(g): In carrying out our functions, Ofcom are required to secure the adequate protection of citizens from harm presented by content on regulated services, through the appropriate use by providers of such services of systems and processes designed to reduce the risk of such harm.

A1.21 The measure is intended to mitigate risks to users from illegal intimate image abuse content. It is informed by our assessment of the risk of harm, as set out in our Illegal Harms Register of Risks and Risk Profiles.

A1.22 In relation to matters to which section 3(2)(g) in the 2003 Act is relevant, section 3(4A) sets out that in performing their duties under subsection (1), Ofcom must have regard to such of the following as appear to them to be relevant in the circumstances:

(a) The risk of harm to citizens presented by content on regulated services.

A1.23 Our Illegal Harms Register set out the risks of harm posed by content on regulated service providers, including intimate image abuse. These risks, alongside findings from services providers' risk assessments, largely inform whether a service should implement the measure to address the risk of harm. The measure's application is limited to services with a certain risk, type or size.

(b) The need for a higher level of protection for children than for adults.

A1.24 Our Codes as a whole already ensure a higher level of protection for children than for adults. The intimate image hash matching measure which is the subject of this Statement does not specifically target the protection of children.

(c) The need for it to be clear to providers of regulated services how they may comply with their duties set out under the Act.

A1.25 Our recommendations, and the explanation in this document of how the measure works, aim to provide clarity and tangible steps that service providers can take to meet their duties in the Act. We have explained that service providers which choose to implement the measure in the Codes will be considered as complying with relevant duties. We have also explained that service providers may seek to comply with their safety duties by choosing to take alternative measures. Finally, we have published guidance to assist providers in deciding what proportion of detected content to human review.

(d) The need to exercise their functions so as to secure that providers of regulated services may comply with such duties by taking measures, or using measures, systems or processes,

which are (where relevant) proportionate to (i) the size or capacity of the provider in question, and (ii) the level of risk of harm presented by the service in question, and the severity of the potential harm.

A1.26 We have clearly identified in the Code measures the types and sizes of services we recommend adopt each measure, for the reasons given in each chapter. The Act requires us to ensure measures are proportionate, and we recognise that the size, capacity, and risks of services differ widely. We therefore do not take a one-size-fits-all approach. Instead, we have set out what types of service providers we think should use specific safety measures to comply with their duties, with the most extensive expectations placed on the riskiest services.

(e) & (f) The desirability of promoting the use by providers of regulated services of technologies which are designed to reduce the risk of harm to citizens presented by content on regulated services; and the extent to which providers demonstrate, in a way that is transparent and accountable, that they are complying with their duties.

A1.27 In this document we set out our decision to implement a measure involving the use of such technologies. We think the use of hash matching technology will enable intimate image abuse to be addressed at scale.

Section 3(4) : Ofcom must also have regard, in performing those duties, to such of the following as appear to them to be relevant in the circumstances […] (b) the desirability of promoting competition in relevant markets, […](d) the desirability of encouraging investment and innovation in relevant markets; […] (h) the vulnerability of children and of others whose circumstances appear to Ofcom to put them in need of special protection; (i) the needs of persons with disabilities, of the elderly and of those on low incomes; (j) the desirability of preventing crime and disorder; (k) the opinions of consumers in relevant markets and of members of the public generally; (l) and the different interests of persons in the different parts of the United Kingdom, of the different ethnic communities within the United Kingdom and of persons living in rural and urban areas.

A1.28 We have had regard to the desirability of promoting competition and encouraging investment and innovation. Our measure provides flexibility for services to decide how to achieve compliance. As set out above, we have considered the interests of consumers in relevant markets as part of our impact assessments of the measure. The measure is recommended for the purpose of complying with illegal content safety duties, which aims to prevent crime and disorder. In relation to the opinions of consumers in relevant markets and of members of the public generally we have taken account of stakeholder responses to our consultation. In our EIAs, we have considered the needs of persons of protected and listed characteristics. We have also considered our Welsh language obligations.

Schedule 4 – Online Safety Act 2023

Paragraph 2(a): providers of Part 3 services must be able to understand which provisions of the code of practice apply in relation to a particular service they provide.

A1.29 In our Codes and this statement, we have been clear about which services this measure applies to.

Paragraph 2(b): the measures described in the code of practice must be sufficiently clear, and at a sufficiently detailed level, that providers understand what those measures entail in practice.

A1.30 We have been as clear as possible and included an appropriate level of detail in this statement and the amended Codes. We have sought to be sufficiently detailed and precise while ensuring our measure is technically feasible and proportionate for the wide range of services in scope of the measure. Our approach to the amended Codes strikes the balance between providing certainty about what providers need to do and allowing them flexibility to implement the measure in a way that works in the context of their own services and is proportionate. While we recognise that some aspects of the measure are not prescriptive, we have explained why this is not possible and have provided supporting guidance to assist service providers.

Paragraph 2(c): the measures described in the code of practice must be proportionate and technically feasible: measures that are proportionate or technically feasible for providers of a certain size or capacity, or for services of a certain kind or size, may not be proportionate or technically feasible for providers of a different size or capacity or for services of a different kind or size.

A1.31 We have considered proportionality and technical feasibility, where appropriate, as part of our assessment of impacts across this document. The hash-matching measure recommends that the provider should ensure that, only where technically feasible, hash matching technology is used to analyse content to assess whether it is intimate image abuse content.

Paragraph 2(d): the measures described in the code of practice that apply in relation to Part 3 services of various kinds and sizes must be proportionate to Ofcom's assessment under section 98 of the risk of harm presented by services of that kind or size.

A1.32 The measure aims to combat that harm caused by intimate image abuse. We have addressed above in this statement why the measure is proportionate in addressing this harm. As required by section 3(4A) (d)(ii) of the 2003 Act, in considering proportionality we have had regard to the severity of the potential harm as well as the level of risk of harm, as identified in the Illegal Harms Register.

A1.33 This measure is targeted at the kind of services which we think should be deploying hash-matching technology because they pose greater risks in relation to intimate image abuse. Having regard to the desirability of encouraging investment and innovation in the markers for regulated services and these technologies, our measure provides sufficient flexibility for services.

Ofcom's online safety objectives

User-to-user services

A1.34 As required by paragraph 3 of Schedule 4 to the Act, we have ensured that our Codes measures are compatible with the pursuit of the applicable online safety objectives for user-to-user services as set out in this sub-section.

A1.35 In line with paragraph 11 of Schedule 4 to the Act, this measure relates only to the design or operation of a Part 3 service (a) in the United Kingdom, or (b) as it affects United Kingdom users of the service.

Paragraph 4(a)(i): a service should be designed and operated in such a way that the systems and processes for regulatory compliance and risk management are effective and proportionate to the kind and size of service.

A1.36 The Codes already include measures related to governance and accountability, and we are not changing these.

Paragraph 4(a)(ii): a service should be designed and operated in such a way that the systems and processes are appropriate to deal with the number of users of the service and its user base.

A1.37 We have explained in paragraph 3.97 that we are not aware of other effective ways to protect victims and survivors of intimate image abuse, besides using proactive technology to detect relevant content quickly. We have explained above who the measure applies to, including our view that it is proportionate for the measure to apply to some non-large services which are at particular risk for intimate image abuse.

A1.38 Further, our existing content moderation, automated content moderation and reporting and complaints measures were set having regard, among other things, to the number of users of the services and its user base.

Paragraph 4(a)(iii): a service should be designed and operated in such a way that United Kingdom users (including children) are made aware of, and can understand, the terms of service.

A1.39 Our Codes recommend measures related to terms of service and we are not amending or introducing additional such measures. This measure recommends the use of proactive technology. In this regard, the Act requires service providers to include provisions in their terms of service giving information about proactive technology used by a service provider.

Paragraph 4(a)(iv): a service should be designed and operated in such a way that there are adequate systems and processes to support United Kingdom users.

A1.40 The removal of intimate image abuse supports UK victims and survivors of intimate image abuse and more generally UK users from encountering this illegal content. As set out in paragraph 3.132, individuals whose consensual content is wrongly removed can use existing appeals mechanisms to have their content reinstated.

Paragraph 4(a)(vi): a service should be designed and operated in such a way that the service provides a higher standard of protection for children than for adults.

A1.41 As set out above, our Codes as a whole already ensure a higher level of protection for children than adults. The intimate image abuse hash matching measure which is the subject of this statement does not specifically target the protection of children.

Paragraph 4(a)(vii): a service should be designed and operated in such a way that the different needs of children at ages are taken into account.

A1.42 In our December 2024 Statement at Chapter 10: 'Terms of service', and Chapter 8: 'U2U settings, functionalities, and user support', we set out how we had regard to the different needs of children at different ages. Further, in our April 2025 Statement, we noted that service providers have a duty, as part of their children's risk assessment, to assess their user base, including separate consideration to children in different age groups on the

service and assessing how the design and use of the service affects the level of risk of harm to children and that we therefore consider that this objective will be secured in particular via the children's risk assessment duties and the Children's Risk Assessment Guidance.

Paragraph 4(a)(viii): a service should be designed and operated in such a way that there are adequate controls over access to the service by adults.

A1.43 Our Illegal Content Codes include measures to limit the access of a terrorist organisation proscribed by the UK government. In our December 2024 Statement we set out why we do not consider it appropriate to restrict access to services generally by adults.

Paragraph 4(a)(ix): a service should be designed and operated in such a way that there are adequate controls over access to, and use of, the service by children, taking into account use of the service by, and impact on, children in different age groups.

A1.44 In our April 2025 Statement on our Protection of Children Codes we explained how our Age Assurance take into account the use of the service by, and impact on, children in different age groups. In relation to our content moderation measures and recommender systems measures we also explained that providers should consider children's ages as a factor when designing the part of their system relating to the appropriate action.

Paragraph 4(b): a service should be designed and operated so as to protect individuals in the United Kingdom who are users of the service from harm, including with regard to:

- algorithms used by the service,
- functionalities of the service, and
- other features relating to the operation of the service.

A1.45 The measure is intended to mitigate risks to users from illegal intimate image abuse content. Our measure recommends that content identified on user-to-user services as intimate image abuse in line with content moderation should be swiftly taken down. This protects both victims and survivors and users of the service more generally from seeing the illegal content.

Search services

A1.46 As required by paragraph 3 of Schedule 4 to the Act, we have ensured that our recommendations, in the context of our existing Codes, are compatible with the pursuit of the applicable online safety objectives as set out in this sub-section.

A1.47 In line with paragraph 11 of Schedule 4 to the Act, relate only to the design or operation of a Part 3 service (a) in the United Kingdom, or (b) as it affects United Kingdom users of the service.

Paragraph 5(a)(i): a service should be designed and operated in such a way that the systems and processes for regulatory compliance and risk management are effective and proportionate to the kind and size of service.

A1.48 The Codes already include measures related to governance and accountability, and we are not changing these.

Paragraph 5(a)(ii): a service should be designed and operated in such a way that the systems and processes are appropriate to deal with the number of users of the service and its user base.

A1.49 We have explained at 3.102 why we consider proactive technology is needed to protect victims and survivors of intimate image abuse by detecting such content at scale. We have explained above who the measure applies to, including our view that it is proportionate for this measure to apply to only large general search services.

A1.50 As shown above, intimate image abuse is widespread and causes serious harm and hash matching in line with this recommendation is a proportionate way of dealing with the harm.

Paragraph 5(a)(iii): a service should be designed and operated in such a way that United Kingdom users (including children) are made aware of, and can understand, the publicly available statement referred to in sections 27 and 29.

A1.51 In our December 2024 Statement and our April 2025 Statement, we set out recommended measures that, among other things, are designed to ensure the publicly available statement is accessible to children.

Paragraph 5(a)(iv): a service should be designed and operated in such a way that there are adequate systems and processes to support United Kingdom users.

A1.52 The deprioritising of intimate image abuse supports UK victims and survivors from this harm. It also more generally supports UK users as its stops them from viewing this illegal content. Further, in our December 2024 Statement and our April 2025 Statement, we have set out recommended measures that support this goal including in relation to reporting and complaints.

Paragraph 5(a)(v): a service should be designed and operated in such a way that the service provides a higher standard of protection for children than for adults.

A1.53 As set out above, our Codes as a whole already ensure a higher level of protection for children than adults. The intimate image abuse hash matching measure which is the subject of this statement does not specifically target the protection of children. In our December 2024 Statement and our April 2025 Statement, we set out recommended measures that support this goal. We note that the measures in the Protection of Children search Code are designed to secure a higher level of protection for children than for adults.

Paragraph 5(a)(vi): a service should be designed and operated in such a way that the different needs of children at different ages are taken into account.

A1.54 In our December 2024 Statement and our April 2025 Statement we set out recommended measures to support this goal including in the 'Reporting and complaints' and 'PAS' chapters of our December 2024 Statement and the draft Protection of Children Codes for search services.

Paragraph 5(b): a service should be assessed to understand its use by, and impact on, children in different age groups.

A1.55 As set out in our April 2025 Statement, service providers have a duty, as part of their children's risk assessment, to assess their user base, including the number of children in different age groups on the service. Additionally, service providers must assess the impact of the risk of harm to children in different age groups on their services.

Paragraph 5(c): a search engine should be designed and operated so as to protect individuals in the United Kingdom who are users of the service from harm, including with regard to:

- algorithms used by the search engine,
- functionalities relating to searches (such as a predictive search functionality), and
- the indexing, organisation and presentation of search results.

A1.56 As explained in our June 2025 Consultation, we are recommending that certain providers of search services implement hash matching to detect intimate image abuse content. This will protect users of the service from harm (as well as other individuals who are not users of the service).

Other duties

A1.57 Under section 92(2) of the Act, Ofcom must have regard to the Secretary of State's statement of strategic priorities for online safety when carrying out online safety functions. We have had regard to this statement, which was designated on 2 July 2025, shortly after we published our consultation. We responded to the statement on 25 July 2025[196] and have taken it into account in finalising our decisions set out in this statement.

[196] Ofcom to the UK Government Director for Security and Online Harm on the Statement of Strategic Priorities for Online Safety, 25 July 2025.

A2. Legal Framework

A2.1 We set out the statutory basis of Ofcom's role and the issues we must consider when preparing Codes.

A2.2 This document further discusses our approach to proposing a new measure in the Codes, and the impact assessment of the measure in the Codes.

The Online Safety Act 2023 and the Codes of Practice

A2.3 The Online Safety Act 2023 (the Act) is a set of laws designed to protect all UK users online, including children, by placing duties on a range of user-to-user and search service providers. The duties on service providers include identifying, mitigating and managing the risk of harm caused by illegal content and activity, including intimate image abuse.

A2.4 The Act establishes Ofcom as the regulator responsible for online safety. It places a requirement on us to prepare and issue Codes, which are a package of measures recommended for service providers to comply with their safety duties.

A2.5 The safety duties in the Act only apply to the design, operation and use of services in the United Kingdom (UK).[197] They also apply to the design, operation and use of a service as it affects UK users (i.e., duties that relate to users).[198] The Code measures must therefore relate to the design or operation of a service that operates in the UK and/or as it affects UK users of the service.[199]

A2.6 In December 2024[200] and April 2025,[201] we published our Illegal Content Codes of Practice and Protection of Children Codes of Practice, which set out a series of measures that aim to protect users online. Service providers should implement these measures to be compliant with their obligations in the Act. These Codes are a "safe harbour", as set out in the Act, which means that service providers who implement all applicable measures will be treated as compliant with their relevant duties under the Act.

A2.7 However, service providers are permitted, under the Act, to comply with their safety duties by implementing "alternative measures." In these cases, a service provider needs to retain a record of relevant decisions (i.e., the alternative measures) and explain how the relevant safety duties have been met. They are also expected to carefully consider the rights of users, including the right to freedom of expression and privacy.

A2.8 Under the Act, we are required to prepare and issue the following sets of Codes for user-to-user and search service providers (i.e., Part 3 services):

a) A Code covering terrorism content (relating to the offences set out in Schedule 5);
b) A Code covering child sexual exploitation and abuse content (relating to the offences set out in Schedule 6); and
c) One or more Codes for the purpose of compliance with other relevant duties (including but not limited to those relating to the offences set out in Schedule 7).

[197] This is defined in section 4 of the Act.
[198] Section 8(3) of the Act.
[199] Schedule 4 to the Act, paragraph 11.
[200] Ofcom, 2024. Statement: Protecting People from Illegal Harms Online. [accessed 10 March 2026].
[201] Ofcom, 2025. Statement: Protecting Children from Harms Online. [accessed 10 March 2026].

A2.9 We have issued four Codes that collectively meet this obligation: (1) the Illegal Content user-to-user Codes; (2) the Illegal Content search Codes; (3) the Protection of Children user-to-user Code; and (4) the Protection of Children search Code. Taken as a whole, the measures set out in these Codes address a range of illegal harms, for both adults and children, that are identified as priorities in the Act.

A2.10 We are committed to further strengthening these Codes to protect all users. We have therefore decided in this statement to amend both the Illegal Content user-to-user Codes and Illegal Content search Codes. These amendments will not have an impact on the Protection of Children Codes.

Ofcom's duties and online safety functions

A2.11 This section sets out the statutory basis of Ofcom's role and the issues we must consider when preparing Codes.

Ofcom's general duties under the Communications Act 2003

A2.12 Ofcom is the independent regulator for communications services. We have regulatory responsibilities for the telecommunications, post and broadcasting sectors, as well as for online services.

A2.13 As a public authority, Ofcom must act lawfully, rationally and fairly.

A2.14 The Communications Act 2003 (the 2003 Act) places duties on Ofcom that need to be fulfilled when exercising our regulatory functions, including for online safety. The 2003 Act sets out our principal duty is:

a) To further the interests of citizens in relation to communication matters; and
b) To further the interests of consumers in relevant markets, where appropriate by promoting competition.[202]

A2.15 In performing that principal duty, the 2003 Act sets out that Ofcom's regulatory activities need to be transparent, accountable, proportionate, consistent and targeted only at cases where action is needed.[203] We further need to ensure that UK citizens are properly protected from harm caused by content on regulated services. We achieve this by requiring service providers to use suitable systems and processes that help minimise the risk of harm.

A2.16 The 2003 Act further requires[204] that Ofcom must have regard to several factors, as they appear to us to be relevant in the circumstances,[205] and as a result we have carefully considered the following in making our decisions:

[202] Section 3(1) of the 2003 Act.
[203] We must also have regard to any other principles appearing to us to represent best regulatory practice.
[204] Section 3(4A) of the 2003 Act.
[205] In relation to matters to which section 3(2)(g) is relevant. The 2003 Act sets out other matters to which Ofcom must, to the extent they appear to us relevant in the circumstances, have regard, in performing our duties. They include: the desirability of promoting competition and encouraging investment and innovation in relevant markets; the vulnerability of children and of others whose circumstances put them in need of special protection; the needs of persons with disabilities, the elderly and of those on low incomes; the desirability of preventing crime and disorder; the opinions of consumers and of members of the public generally; and the different interests of persons in the different parts of the UK and of the different ethnic communities within the UK. See Annex 1– Schedule 4 tests.

a) The risk of harm to UK citizens presented by regulated services;
b) The need for a higher level of protection for children than for adults;
c) The need for it to be clear to providers of regulated services how they may comply with their duties under the Act;
d) The need to exercise our functions to secure that providers may comply with such duties by taking or using measures, systems or processes which are proportionate to the size or capacity of the provider and the level of risk (and potential severity) of harm presented by the service;
e) The desirability of promoting the use of technologies which are designed to reduce the risk of harm to citizens; and
f) The extent to which providers demonstrate, in a way that is transparent and accountable, that they are complying with their duties.

A2.17 In line with our additional duties under the 2003 Act,[206] we have also considered the vulnerability of children and of others whose circumstances put them in need of special protection. We have considered:

a) The needs of disabled people, older people, and of those on low incomes;
b) The opinions of consumers and of members of the public generally;
c) The interests of persons in the different parts of the UK; and
d) The interests of the different ethnic communities within the UK.

Schedule 4 and additional Illegal Content Codes considerations

A2.18 The Act sets out that Ofcom must consider the appropriateness of the measures we recommend to different kinds and sizes of services and to providers of differing sizes and capacities.[207]

A2.19 We must also have regard to the principles that:

a) Providers must be able to understand which measures apply to their service;
b) The measures must be sufficiently clear, and at a sufficiently detailed level, that providers understand what they entail in practice;
c) The measures must be proportionate and technically feasible; and
d) The measures must be proportionate to our assessment of the risk of harm presented by services of that kind or size.

A2.20 We must also ensure that the measures described in the Codes are compatible with the pursuit of the list of online safety objectives set out in Schedule 4, and that we include measures relating to each of the areas specified in sections 10(4) and 27(4) (concerning illegal content).

A2.21 Under the 2003 Act, we are also required to conduct impact assessments when preparing a Code or amendment to a Code, including an assessment of the impact on small and micro-businesses.[208]

A2.22 We consider that assessing measures based on our impact assessment criteria is the right approach to ensuring the Codes protect UK users from illegal content while also protecting their rights and enabling service providers to operate and innovate in the market.

[206] Section 3(4) of the 2003 Act.
[207] Schedule 4 of the Act.
[208] Section 7 of the 2003 Act, as amended by section 93 of the Act.

The Deregulation Act 2015

A2.23 Ofcom is required, when exercising our regulatory functions, to have regard to the desirability of promoting economic growth, including through considering the importance of ensuring the regulatory action we take is necessary and proportionate.[209] This duty is referred to as the "growth duty".

A2.24 We are also required to have regard to the Government's statutory guidance on the growth duty.[210] That guidance explains that the duty needs to be considered alongside our other statutory duties, and that its purpose is not to achieve or pursue economic growth at the expense of necessary protections.[211] Among other things, it also identifies particular drivers of economic growth, including innovation, investment and competition.

A2.25 The growth duty has applied to our online safety functions since 6 April 2026, following the end of a time-limited exclusion for these functions. As such, we did not expressly consider the growth duty in our June 2025 Consultation. For this statement, we have therefore reassessed our impact assessment of this measure and the combined impact, by considering the wider economic impacts of our decision (see paragraph 3.306 to 3.307).

[209] Section 108 of the Deregulation Act 2015.
[210] Section 110(3) of the Deregulation Act 2015.
[211] Department for Business and Trade, 2024. Growth Duty: Statutory Guidance – Refresh.

A3. Glossary

A3.1 This glossary sets out definitions of terms used throughout the Statement.

Terms	Definitions
Act	The Online Safety Act 2023 (c.50.).
Content	Anything communicated by means of an internet service, whether publicly or privately, including written material or messages, oral communications, photographs, videos, visual images, music and data of any description.[212]
Cryptographic hash matching technology	Image matching technology which compares the similarity between hashes created from images by means of an algorithm known as a cryptographic hash function, to assess whether those images are cryptographically similar to each other. This does not include technology which compares similarity through the use of machine learning.
Deepfake	A form of audio-visual content that has been generated or manipulated using AI, and that misrepresents someone or something. Deepfakes are usually intended to cause harm by deceiving an audience into believing that something happened when it did not. In some cases, deepfakes consist of wholly new content, whereas in others they take the form of existing content that has been manipulated in some way.
ECHR	The European Convention on Human Rights (incorporated into domestic law by the Human Rights Act 1998).
False positive	Detected content that is not illegal content.
File-storage and file-sharing service	A service whose primary functionalities involve enabling users to (i) store digital content, including images and videos, on the cloud or dedicated server(s); and (ii) share access to that content through the provision of links (such as unique URLs or hyperlinks) that lead directly to the content for the purpose of enabling other users to encounter or interact with the content.
First positive match	Any first positive match to a hash within given perceptual parameters (or cryptographic parameters if used for video content).
General search service	A service that enables users to search the web by inputting search requests. It derives search results from an underlying search index (developed by either the provider of the service or a third party). Search results are presented using algorithms that rank based on relevance to a search request (among other factors). (For the avoidance of doubt, this term includes combined services, the public search engine part of which enables users to search the internet in the manner described above.)

[212] Section 236 of the Act

Terms	Definitions
Generative artificial intelligence (GenAI)	AI models that can create text, images, audio and videos, typically in response to a user prompt.
Harm	This means physical or psychological harm. References to harm presented by content, and any other reference to harm in relation to content, have the same meaning given to it by section 234 of the Act.
Hash matching	Hash matching is a process for detecting content which has previously been identified as illegal or otherwise violative. It allows service providers to prevent the upload, re-upload, download, viewing or sharing of illegal or harmful content. It involves matching a hash of a unique piece of known illegal content stored in a database with user-generated content. Hashing is an umbrella term for techniques to create fingerprints of files on a computer system. An algorithm known as a hash function is used to compute a fingerprint, known as a hash, from a file.
Illegal content	Content that amounts to a relevant offence.
Illegal Content Judgement Guidance (ICJG)	The guidance about making illegal content judgements that Ofcom is required to produce under section 193 of the Act. It is available at Illegal Content Judgement Guidance.
Illegal content proxy	For user-to-user services, this means content that has been assessed and identified as being in breach of the service's terms of service, where the provider is satisfied that the terms in question prohibit the types of content that include illegal content (including but not limited to priority illegal content). For search services, this means search content that has been identified in the provider's publicly available statement for the service as being subject to appropriate moderation action, where the provider is satisfied that illegal content is included within that kind of content (including but not limited to priority illegal content).
Illegal content safety duties	The duties in section 10 of the Act (user-to-user services) and section 27 of the Act (search services).
Illegal harm	Harms arising from illegal content and the commission and facilitation of priority offences.
Intimate image	A photograph or film which shows, or appears to show, a person in an intimate state (as defined in section 66B of the Sexual Offences Act 2003) or an intimate situation (as defined in section 3 of the Abusive Behaviour and Sexual Harm (Scotland) Act 2016).
Intimate image abuse	An offence of sharing or threatening to share intimate images.
Intimate image abuse content	Content amounting to the offence of intimate image abuse.

Terms	Definitions
Large general search services	A general search service which has more than 7 million monthly active UK users.
Large service	A service with more than 7 million monthly active UK users.
Moderation	When a service provider reviews and assesses content to determine whether it is harmful to children or not, or whether it is in breach of the terms of service or publicly available statement of the service, and takes appropriate action based on that determination. We use 'content moderation' when referring to user-to-user services, and 'search moderation' when referring to search services.
Perceptual hash matching	Perceptual hash matching determines whether a given file is likely to be perceived as similar to a hashed file. It does this by comparing the similarity between the hashes of the files, assessed using a 'distance metric'. A threshold is set to determine when there is sufficient similarity between the hashes for the files to be considered a (near) match – that is, perceptually similar to each other. This allows modifications from an original file to be detected.
Pornography services	Services whose principal purpose is the hosting or dissemination of pornographic content and who host user-generated pornographic content. These services are subject to the risk assessment duties and the children's safety duties. Pornography that is published or displayed by the provider of the service is subject to different duties set out in Part 5 of the Act and Ofcom has published separate guidance for providers subject to these duties.
Priority illegal content	Content which amounts to a priority offence.
Publicly available statement	A statement that search services are required to make available to members of the public in the UK, often detailing various pieces of information on how the service operates.
Record Keeping and Review Guidance	The guidance that Ofcom is required to produce under section 52(3) of the Act to help services to comply with their record keeping and review duties under sections 23 (user-to-user) and 34 (search) of the Act, available at Record-Keeping and Review Guidance.
Illegal Harms Register of Risks	The assessment of the risks of harm from illegal content on user-to-user and search services that Ofcom is required to prepare under section 98 of the Act, available at Illegal Harms Register of Risks.
Search content	Content that may be encountered in or via search results of a search service. It does not include paid-for advertisements, news publisher content, or content that reproduces, links to, or is a recording of, news publisher content. Content encountered "via search results" includes content encountered as a result of interacting with search results (e.g., by clicking on them) and does not include content encountered as a result of subsequent interactions.

Terms	Definitions
Service	A regulated user-to-user or search service, that is, only the user-to-user or search part of the service.
Smaller service	A service which is not a large service.
Terms of Service	All documents comprising the contract for use of the service (or of part of it) by United Kingdom users.
User report	A specific type of complaint about content, submitted through a reporting tool.
User-to-user service	An internet service by means of which content that is generated directly on the service by a user of the service or uploaded to or shared on the service by a user of the service, may be encountered by another user, or other users, of the service.

A4. Further Detail on Economic Assumptions and Analysis

A4.1 This annex provides further information related to the economic analysis used to support our decision on the hash matching for intimate image abuse measure assessed in this statement.

A4.2 We have made some general cost assumptions which are set out here. These general assumptions are combined with assumptions specific to the measure that are set out in the 'Costs and impacts' section in Chapter 3.

Labour costs

A4.3 We used data from the Annual Survey of Hours and Earnings (ASHE), to develop our estimates for the labour cost required to implement the code measure. All quantified estimates of costs are provided in 2025 prices, unless otherwise stated, which is the most recent data available.[213] Our June 2025 Consultation used 2024 prices which was the most recently available data at that time.[214]

A4.4 We used the ASHE 2025 gross median earnings[215] for the following Standard Occupational Classification ('SOC') 2020 references:

- Programmers and software development professionals' salary (2134) to estimate the cost of 'software engineer' time used when developing our cost estimates;
- Database administrators and web content technicians' salary (3133)[216] to estimate the cost of 'content moderator' time when developing our cost estimates; and
- Professional occupations salary (2) to cover a range of professions that are employed at various online services and might be required to implement code measures.[217] This could be legal employees, operations, product managers, and so forth.

A4.5 We chose these occupations because they are likely to develop and/or manage the systems and processes that in-scope service providers will need to have to comply with the duties.[218]

---

[213] Office for National Statistics ('ONS'), 2025. Earnings and hours worked, occupation by four-digit SOC: ASHE Table 14. Data from 2025 is provisional at time of writing. [accessed 24 March 2026]

[214] Our December 2024 Statement on Protecting People from Illegal Harms (December 2024 Statement) and our April 2025 Statement on Protecting Children from Harms Online (April 2025 Statement) used 2023 prices.

[215] We note that in our June 2025 Consultation Annex we stated that we had used the 'gross median *full-time* earnings', however this was not correct. We used gross medium earnings 'for all employee jobs' in our estimates and the salary figures presented in the Annex.

[216] We note that in our June 2025 Consultation annex the salary code was wrongly reported as 3113, not 3133. 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.

[217] We note that in our June 2025 Consultation Annex we stated that we had used the 'Business, median and public service professionals (24)' salary; however, this was not correct. We used the 'Professional Occupations (2)' salary in our estimates and the salary figures presented in the Annex. This is also what was used in our cost estimates for our December 2024 Statement and April 2025 Statement.

[218] ASHE documentation does not explicitly state that gross salaries include bonuses, but our understanding is that the gross pay includes bonuses, tips and other payments.

A4.6 For some service providers median UK wage rates may differ from actual salary rates.

a) This may be especially the case for larger service providers based in the US, who may have higher salary levels. The salary costs of some types of staff, such as software engineers with certain specialisms, may vary and may be considerably higher in some cases. To take account of this, we have calculated a higher salary estimate, which is double the value of our lower estimate.
b) Conversely, some service providers may outsource some relevant work to locations where average pay is lower than the UK, which may reduce costs. To the extent this is the case, our salary range may tend to overstate costs.

A4.7 We applied a 22% uplift to the gross wage costs to account for non-wage labour costs, such as employers' National Insurance contributions.[219] This is calculated based on the most recent 2024 Blue Book data. Estimates in our June 2025 Consultation used a 21% uplift based on 2023 data, which was the most recent data at the time of that publication.

A4.8 Table A4.1 shows the 'low' and 'high' labour cost estimates for different time periods, including the 22% uplift, for each of the three occupations. The figures are based on annual labour costs, and we have calculated the monthly, weekly and daily estimates.[220]

Table A4.1: Low and high range – Estimates of labour costs

Occupation	Low	High
	Annual labour cost estimates	
Software engineer	£67,816	£135,632
Content moderator	£43,938	£87,877
Professional occupations	£55,210	£110,420
	Monthly labour cost estimates	
Software engineer	£5,651	£11,303
Content moderator	£3,662	£7,323
Professional occupations	£4,601	£9,202

[219] ONS, 2021. *Uplifts from wages and salaries to total employment costs: a note on data*. [accessed 24 March 2026]. ONS recommends dividing the 'employer's social contributions (D.12)' by 'wages and salaries (D.11)' to arrive at the uplift. Both series are published as part of the annual UK National Accounts: Blue Book time series. They provide economy wide estimates of D.11 and D.12 annually. At time of writing, the most recent data available is for 2024.

[220] When producing cost estimates for the measure, we have used resourcing estimates based on different time periods (e.g. days/weeks/months) suitable for the particular measure. The annual wages are derived from the ONS, 2025 Earnings and hours worked, occupation by four-digit SOCs, Table 14.7a. [accessed 24 March 2026]. Gross median annual pay for all employee jobs, 2025 provisional estimates. Monthly, weekly and daily wages are all derived from this annual figure. The monthly wages are derived from dividing the annual wages by the number of months in a year (12). The weekly wages are derived by dividing the annual figure by 45.54 (the number of working weeks per year, calculated by dividing the number of working days by the number of working days in a week). 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 five days a week and that there are eight bank holidays and on average people take an additional 25 days leave a year.

	Weekly labour cost estimates	
Software engineer	£1,489	£2,978
Content moderator	£965	£1,930
Professional occupations	£1,212	£2,425
	Daily labour cost estimates	
Software engineer	£298	£596
Content moderator	£193	£386
Professional occupations	£242	£485

Source: ONS, 2025, Annual Survey of Hours and Earnings. Includes 22% uplift. Calculations are performed based on the underlying median gross salary (the 'low' estimate, before uplift is applied) and then uplifted by 22%.

Non-engineering costs for system changes

A4.9 Where system or other software changes associated with a measure involve a software cost, we typically match the amount of engineering time with an equivalent amount of non-engineering time for work carried out by people in professional occupations. This is to account for labour time that a business might need to spend on a system change, for instance, legal or project management.

Maintenance costs for system changes

A4.10 Where system or other software changes associated with a measure involve an initial cost, we have assumed an ongoing annual maintenance cost of 25% of the initial cost. These ongoing costs reflect work likely required to ensure the system continues to operate as intended. We have applied this assumption in the absence of actual information about the ongoing maintenance costs. This is consistent with the approach we took in Our December 2024 Statement on Protecting People from Illegal Harms Online and our April 2025 Statement on Protecting Children from Harms Online.