0345 601 3203
enquiries@swgfl.org.uk
swgfl.org.uk

Ofcom additional measures – Consultation Response SWGfL (October 2025)

Introduction

This consolidated response represents SWGfL's evidence-based position on Ofcom's Additional Safety Measures consultation. Drawing on our unique operational experience through the Revenge Porn Helpline, Report Harmful Content and StopNCII.org, we strongly support Ofcom's proposals to mandate proactive, privacy-preserving technologies for tackling Non-Consensual Intimate Image (NCII) abuse and urge Ofcom to ensure operational parity between NCII and CSAM. We highlight the need for robust enforcement, transparency, and survivor-centred safeguards, and recommend further strengthening of the Codes to close loopholes and future-proof protections.

SWGfL and the Revenge Porn Helpline

The Revenge Porn Helpline (RPH), operated by SWGfL since 2015, is the UK's dedicated support service for adults affected by intimate image abuse. Over the past decade, practitioners have assisted more than 60,000 individuals and facilitated the reporting and removal of over 400,000 intimate images. This frontline caseload provides unparalleled insight into patterns of harm, offender behavior, and the operational realities of removal and prevention. [SWGfL, Revenge Porn Helpline 2024 report]

In December 2021, SWGfL launched StopNCII.org, the world's first privacy-preserving, device-side hashing tool for adult (18+) non-consensually shared intimate content. Survivors create perceptual hashes on their own device; only the hashes are shared with participating companies to detect and block matching content. No intimate imagery is uploaded to StopNCII.org. This data-minimising model is now referenced by regulators and adopted by major services. Recent data shows a 13-fold increase in NCII reports since 2019, with projections indicating over 46,000 cases by 2027 if current trends persist. [StopNCII.org 'How it works' and FAQ].

Ofcom's Additional Safety Measures (ASM) consultation proposes, for the first time, formal expectations to deploy perceptual hash-matching to prevent re-sharing of NCII across:

- providers whose principal purpose is hosting/dissemination of regulated pornographic content
- file-sharing and file-storage services

- user-to-user services with more than 700,000 monthly UK users (where high risk is identified)
- large general search services (for moderation).

We strongly support these proposals. [Ofcom ASM main paper]

Summary of cross cutting issues

We welcome Ofcom's more robust proposals but want to raise some cross cutting themes that need further consideration. We have worked with our partners and colleagues at the UK Safer Internet Centre and the Online Safety Act Network to ensure our response is as comprehensive as possible.

Below we have highlighted some of the key cross cutting areas of concern.

Key Themes and Concerns:

1. Technical Feasibility as a Loophole

- Ofcom's use of the "technically feasible" clause in safety measures allows service providers to self-assess whether they must comply, effectively creating a "get-out clause."
- Services invoking this clause can still claim compliance and benefit from safe harbour without needing to implement compensatory protections.
- There is no transparency or scrutiny mechanism for these self-declared assessments.
- Recommendations include:
    - Removal or reform of the safe harbour provision.
    - Introduction of a "no rollback" clause to prevent services disabling functionality to avoid obligations.
    - Clarity on the boundary between infeasibility and cost.
    - Regular review of technical advances that could render the proviso obsolete.

2. Inadequate Approach to Evidence and Harm

- Ofcom often requires strong evidence of a measure's effectiveness before recommending it but doesn't weigh the consequences of inaction equally.
- There's reluctance to act where evidence is limited, even when harms are well documented.
- Livestreaming is a notable example: Ofcom acknowledges serious risk to children but won't recommend a ban or age-restriction without further consultation evidence, even though some platforms already voluntarily prohibit child livestreaming.

- Ofcom should adopt a precautionary, "safety-first" mindset aligned with the OSA's intention to shift responsibility onto platforms.

3. Delays and Regulatory Lag

- The iterative approach to code refinement is too slow for emergent harms. Ofcom should adopt a "safety-first" mindset and require rapid implementation of proven measures.
- There are now multiyear gaps between identifying harms and implementing protections.
- Meanwhile, children and other users remain exposed to well-evidenced risks, particularly around livestreaming, CSAM, and algorithmic amplification of illegal content.

4. Missed Opportunities for Safety by Design

- Ofcom's understanding of "safety by design" is limited to reactive or technological fixes (e.g., automated moderation, reporting tools).
- True safety-by-design involves systemic, anticipatory design changes, baking safety into platform architecture, not layering it on after harm occurs.
- There is no clear guidance or benchmark for what good design should look like in this context.

While Ofcom has made progress, especially in stakeholder engagement and willingness to iterate, its interpretation of statutory principles like proportionality and technical feasibility risks undermining user protections. There is too much emphasis on platform cost and operational flexibility, and not enough on user rights and harm prevention. The regulator is not yet delivering on the Act's intention to proactively protect users, particularly children, through timely and systemic safety interventions.

NCII as Illegal content

This consultation confirms that NCII is explicitly categorised as illegal content within the meaning of Schedule 7 of the Online Safety Act. This designation triggers mandatory safety duties for user-to-user and search services. This reinforces that NCII content must not only be removed when identified but must be proactively managed through content detection and preventative technologies.

We also recognise that Ofcom has further articulated the status of NCII content in section 10.38 of the illegal content judgements guidance.

We have long campaigned for NCII content (itself, not just the act of sharing) to be classified as illegal. We have had various organisations state that NCII content is not illegal citing that the illegality is the sharing (or threats to share) intimate content without consent rather than

the content itself and from our perspective this has resulted in extending significant harm. The classification of NCII content being illegal immediately enables further instruments to not only take down illegal content but also to block and disrupt access. We recognise and welcome this change in classification.

We continue to act in the interest of victim-survivors and campaign for the clearest and strongest designation of protection to minimise the harm presented by NCII content.

NCII Statistics and Projections

The Revenge Porn Helpline

Ten years down the line, we are seen as one of the most impactful Helplines in the UK. Over 430,000 intimate images have been reported by us, over 80,000 individual reports and an average of 57% increase in reports every year.

Since its launch, the Revenge Porn Helpline has experienced a sustained and accelerating rise in demand, reflecting both the growing scale of non-consensual intimate image abuse in the UK and increased public awareness of available support. Caseloads have risen from **1,685 reports in 2019** to **22,264 in 2024**, marking a more than **thirteen-fold increase in just five years**. Over the longer term, this represents a **forty-three-fold rise** since 2015, when only 521 cases were recorded. Year-on-year growth has frequently exceeded 100%, with the Helpline recording **106% growth in 2022**, **103% in 2023**, and a further **21% rise in 2024**.

[redacted] This continued escalation not only demonstrates the urgent and evolving nature of image-based abuse but also underscores the Helpline's critical national role as a trusted service for victims. These projections reinforce the need for scalable, proactive, and privacy-preserving solutions such as StopNCII.org and the forthcoming Global Clearing Centre, ensuring that the UK remains equipped to respond effectively to rising demand and emerging forms of harm.

The Scale of Non-Consensual Intimate Image Abuse

Non-Consensual Intimate Image (NCII) abuse is a pervasive and rapidly escalating form of technology-facilitated gender-based violence. Despite national and international efforts to address the issue, the true scale remains largely underestimated due to significant reporting gaps and fragmented enforcement mechanisms. This paper provides a comprehensive, data-driven analysis of NCII prevalence, examining its scope at national, Five Eyes, and global levels.

By applying a modelling approach that extrapolates from UK prevalence rates, we estimate the potential worldwide impact of NCII, revealing its scale to be at least as significant as Child Sexual Abuse Material (CSAM) reporting.

Our findings highlight that NCII affects an estimated 1.42% of adult women annually in the UK, with projected figures indicating that millions worldwide may be impacted each year. Further comparison with CSAM data from the NCMEC CyberTipline suggests that NCII could be at least 76% larger in scale, reinforcing the urgent need for a globally coordinated response. The lack of a dedicated international NCII reporting and enforcement framework leaves victims vulnerable, platforms inconsistent in their responses, and law enforcement agencies under-equipped to address the crisis effectively.

To understand the prevalence of NCII in the UK, we must begin by examining the total adult female population, as women disproportionately experience this form of abuse (Revenge Porn Helpline, 2023). By applying established prevalence rates from governmental and non-governmental sources, we can estimate the true scale of NCII incidents, the volume of imagery being produced and shared, and the gaps in reporting mechanisms. This section will present a data-driven approach to quantifying the impact of NCII, providing insight into the extent of underreporting and the broader implications for policy and victim support services.

Based on UK Government and non-governmental sources:

- Total adult female population (UK): 26,064,118 (ONS, 2023).
- Annual prevalence of gender-based violence (GBV) among women: 8.3% (NPCC, 2023).

This statistic originates from the Call to Action as Violence Against Women and Girls Epidemic Deepens report by the National Police Chiefs' Council (NPCC, 2023). The report underscores the increasing scale of violence against women and girls (VAWG) in the UK, highlighting a systemic failure in tackling domestic abuse, stalking, harassment, sexual offences, and technology-facilitated abuse. It stresses the need for a coordinated response between law enforcement, policymakers, and digital platforms to address the growing crisis. The findings are particularly relevant to this paper as they establish a direct link between NCII and broader patterns of gender-based violence, reinforcing the argument that NCII should be addressed as a critical aspect of VAWG intervention strategies.

Percentage of GBV cases involving NCII: 17% (Refuge, 2022). This statistic is drawn from the Unsocial Spaces: Online Harassment and NCII Abuse report by Refuge (2022).

The report highlights the widespread nature of online harassment, with NCII forming a significant component of technology-facilitated gender-based violence. It details how women

are disproportionately targeted, often experiencing a combination of digital abuse, stalking, and coercive control from both current and former partners. The findings underscore the urgent need for improved online safety measures, stronger legislative protections, and increased platform accountability to combat the pervasive threat of NCII abuse.

Annual NCII victims in the UK: 369,272. This figure is derived by applying the NPCC estimate that 8.3% of adult women experience gender-based violence (GBV) annually to the total UK adult female population of 26,064,118 (ONS, 2023). From this group, Refuge (2022) reports that 17% of GBV victims experience NCII abuse, resulting in an estimated 369,272 women affected by NCII annually. This calculation underscores the scale of NCII within the broader context of GBV and highlights the need for targeted intervention strategies. To put this in context, this would suggest that this figure is equivalent to the entire adult female population of the city of Birmingham (ONS, 2023) who are annually victims of NCII.

Estimated NCII images generated per victim: 8.6 images (Revenge Porn Helpline, 2023). This figure is based on data collected from cases handled by the Helpline, which found that, on average, each case involving an adult female victim contained 8.6 non-consensually shared images. This insight highlights the scale of image proliferation per victim, underscoring the compounded impact of NCII beyond individual incidents of abuse. It is important to note that the Revenge Porn Helpline reports a wide range of image numbers per individual which in some cases can be tens of thousands of images being shared without consent. The 2023 Revenge Porn Helpline Annual Report provides further context on how these images are distributed across platforms, with many cases involving multiple platforms and repeated uploads, significantly exacerbating harm.

Total estimated NCII images per year (UK): 3.18 million.

This figure is derived by multiplying the estimated annual number of NCII victims (369,272) by the average number of images per case (8.6), as reported by the 2023 Revenge Porn Helpline Annual Report. This calculation underscores the sheer volume of NCII content being created and distributed each year, highlighting the extensive and ongoing nature of this form of abuse.

This data suggests that NCII affects 1.42% of all adult women in the UK annually. This figure is derived by dividing the estimated number of NCII victims (369,272) by the total UK adult female population (26,064,118), as reported by the ONS (2023). This calculation provides a clearer understanding of the pervasiveness of NCII within the UK population and places its prevalence on par with the scale of CSAM reports received by NCMEC for child victims. To validate this comparison, we examined the NCMEC CyberTipline report (NCMEC, 2023) and

compared CSAM report data from the Five Eyes nations against their respective child populations. The results indicate that the prevalence of CSAM reports relative to child populations ranges from 1.24% (UK) to 1.54% (Canada), with an overall Five Eyes average of 1.47%. This aligns closely with the NCII prevalence estimate of 1.42% among adult women in the UK, reinforcing the argument that NCII is an issue of comparable scale and urgency to CSAM reporting. The full dataset is detailed in the NCMEC report

NCII reporting remains highly underestimated, with the Revenge Porn Helpline assisting around 9,000 adult women who were victims of NCII in 2023, combining direct casework and chatbot support. Given that 369,272 women are estimated to experience NCII annually in the UK, this suggests that for every woman supported, approximately 53 others require assistance but do not seek it. This data underscores the vast discrepancy between known cases and the likely true extent of NCII abuse, highlighting the scale of the 'iceberg effect,' where only a small fraction of incidents come to light while the majority remain hidden beneath the surface.

In a Freedom of Information request submitted to UK police forces for the year 2024, based upon responses from 37 forces there were 6459 complaints made by victims related to NCII. Of those complaints, 1037 arrests were made, and in total 264 resulted in either cautions or charges. Based upon these figures we can state that for those victims making a complaint, only 4% will saw their complaints result in criminal action. When reasons for withdrawal of complaint are explored, in 2180 cases, this was because the victim withdrew support.

Synthetic NCII (deepfakes)

Synthetic imagery has rapidly become a major driver of intimate image abuse (The Revenge Porn Helpline has seen a 26% increase from 2023 to 2024) and fundamentally alters the response model. In these cases, victims often do not possess an “original” image to hash, meaning traditional survivor-informed prevention cannot apply. To address this, platforms and trusted NGOs must be able to create and share verified hashes of synthetic NCII through coordinated industry mechanisms. This ensures that once an image is confirmed as abusive, it can then be suppressed across other services, even when no survivor-generated hash exists. The same principle underpins the forthcoming Global Clearing Centre, which will coordinate the reporting and removal of synthetic NCII at scale, moving beyond delisting to full takedown.

Livestreaming

Question 1: Do you have further evidence regarding the harms and risks to users from livestreamed illegal content or content harmful to children, or harms and risks to children from broadcasting livestreams?

We strongly support Ofcom's proposals for human moderators to be available during livestreams and for services to operate crisis protocols. Our casework includes examples where livestreamed content was rapidly captured and redistributed, causing ongoing distress to victims. The IWF's 2023 case study on viral imagery demonstrates the urgency of proactive, real-time moderation and rapid duplicate hunting within the first 24-48 hours of an incident.

Our frontline experience through the Revenge Porn Helpline and StopNCII.org confirms that livestreaming, while not the primary vector for NCII, presents significant risks. Frames from livestreams are easily captured and recirculated as images or clips, perpetuating harm long after the original broadcast.

We also support our partner (IWF's) concerns about limiting the measure to 1to many livestreams and do not feel this takes a safety by design approach

And we call on Ofcom to:

- ensure proactive detection duties for financial services extend to CSEA as currently it's limited to just fraud.
- work with the FCA to produce guidance for financial institutions providers whose electronic payment services are being exploited for the selling and sharing of CSAM.

We also support IWF on welcoing the introduction of human moderation but support them in calling for:

- further clarity on how moderators will respond and intervene will work in practice
- the number of moderators on a platform being proportionate to the risk of harm
- dedicated CSEA reporting channel
- safeguarding moderator welfare

Question 3: Do you agree with our proposals? Please provide your reasoning, and if possible, provide supporting evidence.

We welcome Ofcom's additional proposals aimed at reducing the harms associated with live streaming. The SWGfL-operated Helplines, including the Professionals Online Safety Helpline, Revenge Porn Helpline, and Report Harmful Content service, have extensive experience supporting individuals affected by online harms. Through this work, we are acutely aware of both the severity of harm that can occur via live streaming and the challenges users face when attempting to report such content.

Due to the time-sensitive nature of live streams, harmful content is often difficult to report and remove swiftly enough to prevent further damage. We therefore strongly support the proposal that reporting mechanisms must be easily accessible, clearly signposted, and supported by human moderators who are empowered to intervene in real time.

A recent and tragic example underscores the urgency of this issue: a man lost his life while live streaming, and no action was taken to halt the stream immediately. This incident not only resulted in the loss of life but also exposed viewers to distressing and harmful content. The news article detailing this case can be found here: BBC News.

Our Helplines have also identified further concerns around the moderation of livestream content. In many cases, even when a livestream is removed and an account is moderated, there are insufficient safeguards to prevent the user from creating a new account and continuing to broadcast harmful content. This loophole significantly undermines the effectiveness of current moderation efforts.

It is also critical to emphasise the importance of removing livestreams as quickly as possible. In cases involving non-consensual intimate imagery or child sexual abuse material (CSAM), every second of live broadcast can generate thousands of frames, each a potential image that can be captured, redistributed, and perpetuate harm long after the stream ends.

We urge Ofcom to consider these points carefully and to ensure that any regulatory framework includes robust, real-time intervention capabilities, effective account-level enforcement, and prioritisation of rapid content removal to mitigate long-term harm.

Proactive technology

Question 11: Do you agree with our proposals? Please provide your reasoning, and if possible, provide supporting evidence.

Question 12: Do you have any comments on the Proactive Technology Draft Guidance?

SWGfL acknowledges that hash matching is a form of proactive technology. However, as it involves specific actions that all platforms should be expected to implement, we support the proposal to separate hash matching for IIA from this section as a specific and separate requirement.

Technically feasible

To prevent the 'technically feasible' proviso from becoming an open-ended opt-out under safe harbour, Ofcom should:

- require evidence-based infeasibility claims (including technical constraints, performance benchmarks and alternatives attempted);
- publish its assessment methodology;
- require providers to log configuration choices (precision/recall targets, human review) for audit;
- review accepted infeasibility claims on a fixed cadence (6–9 months) reflecting rapid advances in safety tech.

We support the Online Safety Act's response and analysis in relation to technically feasible.

Perceptual hash matching and intimate image abuse

Overall, we welcome Ofcom's measures addressing NCII in section 11 of the consultation.

Perceptual hash-matching technology is already in use across a wide range of platforms and jurisdictions, demonstrating clear *technical feasibility, proportionality and cost effectiveness* for platforms. Tools such as StopNCII.org enable individuals to create hashes of intimate images on their own devices without uploading the content, preserving privacy while allowing platforms to detect and block harmful material therefore protecting thousands of victims.

The infrastructure supporting this technology is scalable and has been deployed by 15 partner platforms worldwide (as at point of submission, the full list available here), showing

that implementation is viable for services of varying sizes and types. The system is designed to be survivor-focused and privacy-first, making it suitable for both large-scale and niche platforms.

We challenge the assumption that costs are prohibitive. Integration of StopNCII.org has been achieved in under 80 hours by OnlyFans, with ongoing maintenance of 4-5 hours/month. The societal and individual costs of inaction far exceed the modest technical costs of implementation.

Question 19: Do you agree with our proposals? Please provide your reasoning, and if possible, provide supporting evidence.

Question 20: Do you have any evidence on the relative efficacy of third-party and internal databases for image-based IIA content?

Introduction to StopNCII.org

We strongly support the proposal to mandate perceptual hash matching for NCII. StopNCII.org has proven effective in real-world deployments and is now recognised as the global standard for tackling NCII. Survivor-led, on-device hashing ensures privacy and dignity, while industry hash sharing enables rapid, cross-platform takedown.

StopNCII.org is currently the only third-party hash database available for adult NCII, and its infrastructure is designed to support thousands of platforms. Integration is straightforward, with some providers reporting setup times of under 80 hours. The system is built to minimise technical burden while maximising impact. Given its survivor-centred design and privacy safeguards, StopNCII.org offers a proportionate and effective solution for regulated pornographic, social media, and file-sharing services. The cost of implementation must be weighed against the significant harm caused by NCII, which includes long-term psychological, reputational, and safety impacts on victims.

Partner onboarding process

Each new StopNCII.org partner completes a structured onboarding process designed to ensure legal compliance, technical integrity, and transparent collaboration. The process begins with the agreeing of a Data Sharing Agreement, establishing clear parameters for the purpose and secure exchange and handling of hashed data. Once this agreement is in place, invoices are issued to cover participation costs, and API keys are released to enable technical integration with the StopNCII infrastructure. The integration phase is actively supported by the SWGfL technical team to ensure alignment with platform moderation systems and security protocols. Upon completion, the partnership is publicly announced through

coordinated communications, including press releases and social media, to reinforce transparency and encourage broader industry engagement.

Keily Blair, CEO OnlyFans explained in the Women's Equality Committee oral hearing on the inquiry into non-consensual intimate image abuse "In terms of the timeline from when we signed contracts to when we were able to implement, we started the implementation in January. We finished the first phase of the implementation the same month. It took approximately 80 hours of tech and engineering time to be able to fully integrate phase 1. Phase 2 was implementing the feedback loop, which is enabling us to actually provide feedback to StopNCII about hash matches and things like that. That took a little more time; that was more complex; and that was fully implemented by the end of March. In terms of monthly maintenance hours, I would say, typically, between four and five hours a month of dev and tech time, a little legal advice from time to time comes into things when there is contracting to do, but it is not an overly onerous process. ...For us, it was absolutely worth the time and I would have spent double the amount of time to be able to implement it because of the effectiveness of the technology in protecting victims". https://committees.parliament.uk/oralevidence/14802/pdf/

Privacy persevering principles

StopNCII.org has proven effective in real-world deployments and is now recognised as the global standard for tackling NCII. Unlike CSAM, where expert-sourced databases are common, NCII prevention requires a survivor-focused approach to protect privacy and dignity. Hashes are created 'on-device' by the individual, ensuring no intimate imagery is uploaded. While the hash database ecosystem for NCII is newer than that for CSAM, it is rapidly maturing and already delivering strong outcomes across major platforms.

To build confidence in the integrity of the hashset, StopNCII incorporates eligibility checks and simple questions during hash creation (e.g., *Are you depicted in the image? Does the image include nudity, semi-nudity, or a sexual act?*). These steps strike a careful balance: they provide assurance that the image meets NCII criteria while keeping the process frictionless for victims who need urgent protection. Platforms understand the nature of the hashset and accept the precision/recall trade-off as proportionate to the harm being addressed. Safeguards such as configuration principles, precision/recall optimisation, human-in-the-loop review, and secure database management are embedded to mitigate risks, including potential impacts on freedom of expression.

We continually review and improve reliability through iterative development and transparency measures, such as reporting the number of hashes ingested, matches reviewed, and actions taken. Future enhancements, including the Global Clearing Centre and industry hash sharing, will further strengthen accuracy by enabling verified hashes from

trusted partners and platforms, enriching detection and reducing duplication. This approach is technically feasible, proportionate, and proven in practice, delivering exactly what Ofcom's objectives under the Online Safety Act require.

This privacy-first design is integral to the system's uptake, accuracy, and speed, ensuring that technology remains both effective and survivor-centred. Without these safeguards, far fewer individuals would engage with NCII reporting systems, and the overall efficacy of prevention and removal would decline sharply.

The hash database ecosystem for NCII is still developing, but StopNCII.org has demonstrated success in real-world use. It enables confidential hash sharing across platforms, enriching the dataset and improving detection. Transparency measures, such as reporting the number of hashes ingested and matched, can further strengthen trust. While hashes are not verified by third parties to preserve survivor privacy, there is no evidence of malicious use or systemic false positives. Safeguards such as configuration principles, regular reviews, and secure database management can mitigate risks to freedom of expression and ensure continued effectiveness. To date, none of the StopNCII partners have provided negative feedback on the false positives in the system.

Managing False Positives and Platform Confidence

StopNCII.org has been successfully deployed across major platforms and is widely regarded as a trusted, privacy-preserving solution for tackling the spread of non-consensual intimate image content (NCII). Platforms that use hash matching understand the nature of the hashset and are comfortable managing the matches they receive. They view the current balance between precision and recall as proportionate to the severity of harm being addressed, and they have the moderation capacity to review and act on matches effectively.

To maintain the privacy of survivors, StopNCII does not verify the content of images hashed on a user's device before submission. This privacy-preserving design ensures that intimate material never leaves the survivor's control while still providing platforms with a trusted, auditable dataset from a verified source. False positives are transparently identified through open tagging and platform feedback, ensuring that the overall system maintains an appropriate balance between precision and recall.

As the ecosystem grows to include smaller platforms with more limited resources, we recognise the need for additional flexibility. To support these services, StopNCII.org is introducing mechanisms that allow platforms to prioritise hashes carrying a verification tag, indicating that the content has been confirmed as NCII by a platform or a trusted NGO. This

ensures that even where moderation capacity is constrained, smaller entities can act confidently on verified matches without compromising user safety or rights.

While this approach introduces a short delay, between hash submission, review by a trusted partner, and onward distribution, it provides an important assurance layer for platforms with limited moderation capacity. The full immediacy of survivor-generated hashes is preserved for larger services with robust workflows. Over time, the integration of NGO and industry-submitted hashes will significantly accelerate the verification pipeline, helping smaller services respond more quickly and confidently.

This approach future-proofs the system for a diverse platform landscape. It preserves the immediate power of survivor-generated hashes for larger, well-resourced services while introducing an additional layer of assurance for those that need it. Combined with ongoing improvements, such as the forthcoming Global Clearing Centre and mechanisms for industry hash sharing, StopNCII remains a trusted, adaptable solution that supports Ofcom's objectives under the Online Safety Act and sets a global benchmark for tackling NCII.

The Efficacy of Industry Hash Sharing

StopNCII.org provides a working, globally adopted model of industry hash sharing that demonstrates both efficacy and cost-efficiency in tackling the spread of non-consensual intimate images (NCII) across digital platforms.

Industry partners are now engaged in testing verified-hash exchange workflows through StopNCII.org, with full deployment expected shortly. These pilots demonstrate that the model is technically viable and scalable across diverse service types, providing a clear pathway to operational cross-platform suppression.

Since its launch in December 2021, the tool has seen sustained and growing use by individuals across the world. As of the end of August 2025:

- 743,449 cases have been initiated — each broadly equating to an individual use of the tool by a victim or their representative.
- These cases together generated a total of 1,850,816 unique hashes of intimate images or videos flagged for prevention.

While these headline figures reflect growing user trust and engagement, the more important indicator of efficacy is the number of true positives: instances where NCII content was matched against hashes provided through StopNCII.org, confirmed as violating content, and proactively blocked from being published on partner platforms.

As of August 2025, 37,772 true positives have been registered, each representing a concrete incident of NCII that was identified and prevented from appearing online through the system. These matches have taken place across major global platforms, including those who not only receive but also share verified hashes into the StopNCII ecosystem, strengthening its protective effect.

This level of success is made possible because hash sharing is not cost-prohibitive. Participating platforms have integrated the API-based system into their existing moderation workflows, meaning hash matching operates without disrupting user privacy or requiring extensive engineering resources. For NGOs and other trusted partners, hash sharing also enables collective validation of harmful content, improving precision while reducing the moderation burden on any single party.

The continued expansion of platform participation, particularly the uptake of verified hash sharing by both industry and NGOs, further improves the reach and efficiency of this approach. It also creates a trusted, privacy-preserving framework for collaborative moderation that meets Ofcom's expectations for proportionate, technically feasible, and effective additional safety measures under the Online Safety Act.

Quality Assurance

StopNCII.org maintains a strong emphasis on transparency, security, and accountability throughout its operation. Each partner's access to the hashset is automatically recorded whenever data are retrieved through the API, creating a clear and auditable record of system use. The quality of the hash data, including tags that identify confirmed NCII content and false positives, is fully visible to all partner platforms, allowing for shared oversight and collaborative assurance of accuracy.

Security of the platform and associated data handling processes is managed in line with SWGfL's organisational standards and is supported by Cyber Essentials Plus accreditation. Regular data protection reviews are undertaken, at least annually, in consultation with an external Data Protection Officer to ensure continued compliance with data protection legislation and best practice.

StopNCII.org also publishes aggregate transparency data, including the total number of cases supported, hashes created, and the proportion of verified (true positive) matches. These measures collectively demonstrate a proportionate and technically feasible approach to quality assurance, consistent with Ofcom's expectations for proactive technology measures under the Online Safety framework.

[redacted]

[redacted]

[redacted]

[redacted]

Question 21: Do you consider this measure to be effective for file-sharing and file-storage services? Please explain your reasoning and, if possible, provide supporting evidence.

Yes. There is no technical limitation that would prevent hash-matching measures from being effective within file-sharing or file-storage environments. The approach has already been demonstrated at scale across multiple global platforms, including those offering storage and sharing functionalities.

The approach has already been demonstrated at scale across multiple global platforms, including those offering storage and sharing functionalities, confirming that the technology is both technically feasible and operationally effective. This mirrors the success of comparable systems deployed by the Internet Watch Foundation (IWF) for child sexual abuse material, where hash-matching has been proven to work effectively across cloud storage, peer-to-peer, and hybrid file-sharing models.

In the NCII context, StopNCII.org's design provides an additional privacy advantage. Because hashing takes place on the survivor's device, no intimate imagery is ever transmitted or stored. This privacy-preserving process allows file-storage and file-sharing providers to act on verified hashes confidently, without accessing or processing the underlying image content.

Taken together, these factors demonstrate that hash-matching is an effective, proportionate, and privacy-respecting solution for detecting and restricting the distribution of intimate image abuse content on file-sharing and file-storage services.

Question 22: Do you agree with our assessment of the impacts (including costs) associated with this proposal? please provide any relevant evidence which supports your position.

We broadly agree with Ofcom's view that hash-matching for intimate image abuse is both achievable and proportionate, with ongoing maintenance costs that are modest when compared with the harms prevented. Real-world integrations demonstrate that implementation can be delivered rapidly, with limited engineering effort and low operational overhead.

During the Women and Equalities Committee inquiry into Intimate Image Abuse (8 May 2024), OnlyFans' Chief Executive confirmed that the company completed its StopNCII.org integration in two phases, requiring approximately **80 engineering hours** for initial deployment and **4–5 hours per month** for ongoing maintenance. Within weeks of launch, the system had already identified and blocked **dozens of attempted uploads** of known NCII content. At a subsequent evidence session (6 November 2024), Microsoft and Google

outlined similar integration processes and confirmed that perceptual-hashing technologies operate effectively and at scale across their storage and sharing environments. These examples demonstrate that the costs and resources required for adoption are well within reach for mainstream and mid-tier providers.

Ofcom's estimated annual maintenance range of £5,000–£88,000 is therefore reasonable, and likely conservative in many cases. Integration costs for systems such as StopNCII.org are modest compared with the wider societal and individual costs of non-consensual intimate image abuse. Survivors experience severe and enduring harm when platforms fail to act; investment in effective, privacy-preserving technology is both proportionate and in the public interest. For context, the UK Revenge Porn Helpline recorded a rise from 1,685 reports in 2019 to 22,264 in 2024, a more-than-thirteen-fold increase in five years, illustrating the scale and urgency of the problem.

In considering feasibility and proportionality, it is essential that "technical feasibility" is not treated as a reason for inaction. The existence and successful deployment of StopNCII.org across multiple global services demonstrate that proactive, privacy-preserving hash-matching is already technically possible and operationally proven. Regulated services should therefore be expected either to implement comparable measures or to evidence alternative mitigations that achieve the same safety outcomes. Allowing claims of infeasibility to substitute for genuine engagement would risk undermining the intent of the Online Safety Act and slowing innovation in user protection.

Overall, the evidence indicates that the costs Ofcom identifies are proportionate and that the measure is both technically and operationally deliverable. The relatively low cost of adoption, set against the substantial and enduring harms it prevents, makes hash-matching a high-impact and cost-effective safety intervention.

Question 23: Do you consider this measure to be effective for large general search services? Please explain your reasoning and, if possible, provide supporting evidence.

Yes. Hash-matching for intimate image abuse is already demonstrably effective for large general search services, particularly for image-based NCII. Microsoft Bing and Google Search have both adopted StopNCII.org hashes, confirming technical feasibility and impact at scale.

Current implementation

- Microsoft (Bing): Since March 2024, Microsoft has been using StopNCII.org image hashes "to prevent this content from being returned in image search results in Bing," and by August 2024 had "acted on 268,899 images as part of [its] StopNCII integration." Microsoft

confirmed that the approach "helps to prevent further harm to victims while protecting user privacy," and that it will continue to expand the partnership. Additional information.

- **Google (Search):** Google has publicly announced that it will begin using StopNCII.org hashes to identify and remove policy-violating NCII imagery from Search "over the next few months," describing this as "an important step to reduce the burden on survivors of intimate image abuse." Additional information.
- **Parliamentary evidence (Women & Equalities Committee, 6 Nov 2024):** Both Microsoft and Google confirmed that StopNCII-style hash-matching is operationally reliable at search scale and integrates efficiently with existing moderation systems. Additional information.

These live and forthcoming integrations demonstrate that hash ingestion and enforcement at search-engine scale are already technically and operationally proven.

Why hash-matching complements URL detection for search

URL detection remains a valuable safety measure where specific, verified URLs are available and can be quickly removed from search results. However, it is limited to individual addresses and becomes less effective once intimate-image-abuse material is re-uploaded, mirrored, or renamed across multiple domains.

Hash-matching provides an essential complementary layer. By identifying the underlying image content rather than the file location, it enables search services to detect and act on duplicates and modified versions that would otherwise evade URL-based removal. When used together, URL detection can target known web locations while perceptual hash-matching ensures continued detection of the same imagery wherever it reappears.

This combined approach strengthens protection for survivors, reduces re-victimisation, and ensures that search results remain free from non-consensual intimate imagery even as content migrates across the web.

Recommended progression

Building on the effective use of hash-matching in image search, we recommend that Ofcom encourage search providers to:

1. **Maintain image-hash ingestion and enforcement.** Continue to prevent hashed NCII images from appearing in image search results, and apply hash-based signals to down-rank or delist matching web results where policy allows. *Status:* Implemented by Microsoft (Bing); onboarding in progress for Google Search.

2. **Report source URLs to the Global Clearing Centre.** Delisting reduces discoverability but does not remove the underlying content. Search providers should report the source URL(s) of identified NCII to SWGfL for inclusion in the forthcoming Global Clearing Centre, which will coordinate host and ISP takedown requests and prevent re-indexing.
   Status: Proposed enhancement.

Conclusion

Hash-matching for image-based NCII is **effective and already in operation** within major general search services. By extending beyond delisting to include source-URL reporting to the Global Clearing Centre, search providers can move from merely reducing visibility to enabling permanent removal of abusive content. This approach delivers a proportionate, privacy-preserving, and scalable response to intimate-image abuse within the search ecosystem.

Crisis response

Question 49: Do you agree with our proposals? Please provide your reasoning, and if possible, provide supporting evidence.

When NCII starts to trend, providers should activate surge moderation and 'hash-as-you-remove' so hashes can be shared during the incident window to prevent cross-platform spread. The IWF's case study on preventing viral imagery shows why rapid, proactive duplicate hunting within the first 24–48 hours is decisive. [IWF 2023 Annual Report case study; Ofcom ASM crisis response].

We also support our partners response (IWF) on amending the harms in scope of the crisis response protocol to include CSEA.

Broadening appeals

Question 54: Do you agree with our proposals? Please provide your reasoning, and if possible, provide supporting evidence.

We support broadening appeals to include 'illegal content proxy' determinations. For NCII, appeals should be time-bound with transparent rationales, and providers should retain match artefacts (hash IDs, decision notes) for Ofcom scrutiny, consistent with Annex 6's deployment/monitoring/feedback lifecycle.