EXPERT ROUNDTABLE OUTCOMES BRIEFING

Link-sharing and child sexual abuse: understanding the threat

Hosted by WeProtect Global Alliance in partnership with GCHQ
February 2023

WeProtect Global Alliance and GCHQ hosted a multi-sector roundtable focusing on how to tackle the growing threat of link sharing to child sexual exploitation and abuse online.

This report provides an overview of the issue and what the key themes and challenges are for actors in industry, law enforcement, civil society and government.

Participating organisations

- Argentina Federal Police
- Bitly
- Child Rescue Coalition
- Crisp Thinking
- GCHQ
- Google
- Internet Watch Foundation
- Interpol
- Mega
- Meta
- National Centre for Missing and Exploited Children
- PLDT-SMART
- Snap
- Thailand Department of Special Investigations
- Twitter
- United States Department of Justice
- University of Bristol
- WeProtect Global Alliance

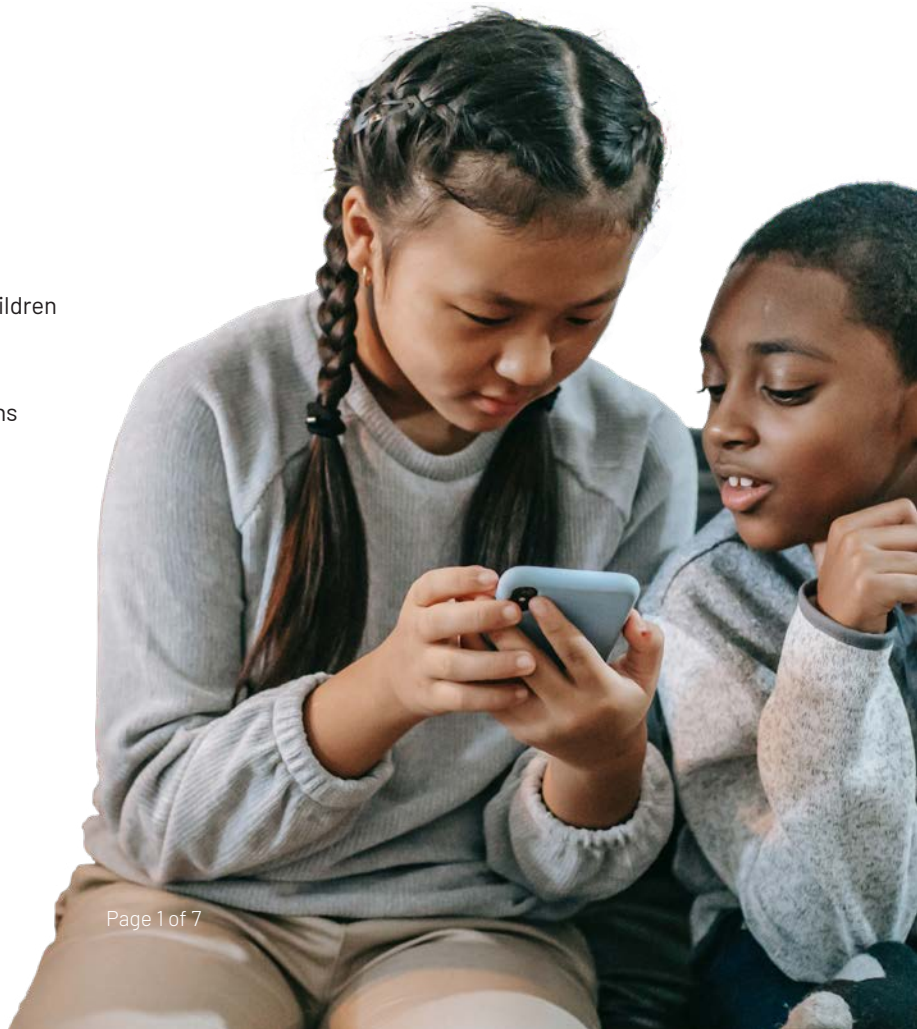

Overview of the threat

The sharing of links to child sexual abuse material or to hosting locations where child sexual abuse material is stored. The term 'links' encompasses original URLs, and those which have been shortened or modified by offenders for obfuscation. Sharing links, modifying links and obfuscating links is just one tactic used by offenders to disseminate child sexual abuse on the internet. It is important to note that this particular activity sits alongside wider harmful and illegal behaviours such as grooming, live streaming, coercing 'self-generated' material and the producing, searching for and sharing of child sexual abuse material online.

WeProtect Global Alliance's 2021 Global Threat Assessment highlighted that there are signs of offenders moving away from the curation of personal collections and preferring 'on-demand' access to content via the sharing of links that lead to child sexual abuse content. Links to files containing child sexual abuse content are posted across multiple sites and often used as part of offender-to-offender sharing. This creates a raft of challenges for law enforcement. Material is often published and hosted in different jurisdictions, which complicates evidence-gathering. The volume of content in an offender's possession was historically one of several factors used to assess the level of risk they posed, but this is no longer always indicative.

Examples of links:

- Original URLs: *http://www.bad.site/content/123, bad.site/conwtent/123*
- Shortened URLs: *abc.yz/abcd1234*
- Obfuscated URLs: *substitution "bad site/content" or "bad(dot)site(slash)content", removing parts of URL "go to content/123 on the website"*

Some examples of how link sharing may be used:

Anecdotal evidence suggests that offenders may use link sharing in the following ways:

- In an open public space: *forum post, blog post, status update, comment*
- In a space which requires users to 'join' using an account: *social media group, online streaming, online game chat*
- In an invite only space or one which requires approval for requests to join: *group chat, private forum, private stream*
- In a private chat or direct messaging chat with (only 2 people, not a group)

Why do offenders use link sharing?

Anecdotal evidence suggests that offenders may use link sharing for the following reasons:

- Attempt to avoid detection via hash-matching technology[1]/image classifiers[2];
- Perceive greater deniability/distance than when sharing content direct with others; therefore, less discoverable and avoid detection;
- Avoid persistent storage of content on devices;
- Access 'new' material (or at least new to the offender) &/or private spaces;
- Benefit interpersonally – e.g., sense of community/social/status;
- Earn through monetisation of clicks or referral schemes.

This is one of the next key challenges in tackling the sharing of child sexual abuse online. The effectiveness of Internet Service Providers (ISPs) blocking access to illegal sites is deteriorating due to deteriorating due to increased encryption of web traffic. Solutions are needed that target the source of a link being shared and prevent it at that point before it reaches an end user.

Link sharing is a significant problem in tackling child sexual exploitation and abuse online. We convened this roundtable of experts and key stakeholders to build a better understanding of what is currently being done by various organisations to curb this threat, provide a platform to share best practice and help participants identifying potential gaps in their responses.

1 – A process of using databases of hashed child sexual abuse material to detect when the material is re-shared, by matching its hash value against those of already known files.

2 – Classifiers use algorithms informed by machine learning to identify and categorise child sexual abuse material.

Current responses

Different industry players are currently responding to this threat in a number of ways and to varying degrees. There is little available data on how companies are responding, which makes it difficult to assess the efficacy of responses. The discussion tended to focus on activity around link sharing, and the identification and takedown of sites, rather than identifying and blocking links as part of content moderation approaches.

One main challenge for many service providers is how to moderate links on their platform where the content is hosted on a different site. The action taken by industry can depend on where the links take users. For example, a link may take a user to content hosted externally, or link to an image-hosting site or website, or to group chats on group messaging apps and forums. All these may be harmful yet require differing responses.

Collaboration with leading safety technology organisations forms an essential part of the response for leading industry players. Many participants at the roundtable cited the Internet Watch Foundation's (IWF) URL List as a helpful tool in identifying potential harms and blocking access to illicit webpages and material. The IWF is constantly updating and reviewing this list – twice a day – for its members to use in tackling the dissemination of links to child sexual abuse online. In addition to ensuring that access to child sexual abuse material is blocked, the IWF works with relevant actors to ensure that the images and videos at the linked location are removed from the internet.

Participants also highlighted Project Arachnid in Canada as an effective technology to combat link-sharing. Project Arachnid identifies child sexual abuse material by crawling specific publicly accessible URLs reported to CyberTipline, as well as URLs on the surface web[3] and dark web[4] that are proven or known to host child sexual abuse material. It detects URLs that host media and matches content against a database of digital fingerprints. As soon as Project Arachnid detects a match in fingerprints, a removal notice is automatically issued requesting the hosting provider to take it down. It follows up on this request by recrawling URLs linking illegal content every day until the content is taken down.

In addition to these partnerships, industry uses a variety of internal tools and systems to identify links to illegal material and take action. Some technology companies with search functions have introduced disruptive blocking to their systems yielding positive results. Research from 2014 highlights that Microsoft and Google observed a 67% decrease in searches for child sexual abuse material after deploying such techniques. User reports are an essential tool for companies to identify links. Many of those represented said that they had clear reporting tools on their platforms or services with some taking a more aggressive approach to links being reported as directing users to child sexual abuse, whether or not these concerns had been verified.

Some also noted that keywords, which are either searched for or posted, can be useful in tracking down perpetrators who wish to lead other perpetrators and potential offenders to linked content. Many stated that they have dedicated threat-hunting teams who proactively search their services for potential illegal linking activity in addition to technological solutions. When it comes to blocking access, many industry players use internal policies and terms and conditions regarding content when they believe there is a high risk of users sharing links to child sexual abuse online that are new or yet to be proven or added to blocklists – examples may include account violations for inauthentic behaviour or hashtag abuse.

3 – The portion of the web readily available to the general public and searchable with standard web search engines.

4 – The layer of information and pages that you can only get access to through so-called 'overlay networks' (such as Virtual Private Networks (VPN) and peer-to-peer (P2P) file sharing networks), that obscure public access. Users need special software to access the dark web because a lot of it is encrypted, and most dark web pages are hosted anonymously.

5 – It is worth noting that this research was conducted in 2014 and that there have been significant technological developments since publication.

Key themes and challenges in tackling link-sharing

Links to sites known to host child sexual abuse material vs. unknown or unverified sites

Different responses are required for the different types of links that are being shared by offenders. Participants stressed the importance of differentiating between links to sites that are known or have been confirmed to host child sexual abuse material, and sites that are not yet known or remain unverified by the various systems in place.

Identifying and blocking links to sites that are known to host this illegal material is generally easier and more advanced than identifying and blocking links that potentially take users to sites that are not known or remain unverified to host online child sexual exploitation and abuse. URL blocklists that are maintained by organisations such as the Canadian Centre for Child Protection and the Internet Watch Foundation, as well as technology companies themselves, are helpful and deployable tools for organisations and companies who want to proactively identify, block and remove links that have been verified as hosting illegal abuse. Law enforcement authorities and international law enforcement organisations also maintain similar lists, for example, INTERPOL's 'worst of' list (IWOL) supports national police authorities by providing a list of links that disseminate the most severe and abhorrent cases of child abuse material on the internet. These are just a few examples and there are more organisations and institutions that provide blocklisting services.

Some participants from platforms asserted that they had a 'block now, review later' strategy to deal with 'grey zone' links where they have serious reason to believe that the link would take a user through to child sexual abuse material hosted on a site external to their platform, even though the suspect link is not listed on the blocklisting technology that they deploy.

Being able to measure the number of links that have not yet been identified or verified is crucial in assessing the scale of problem. Despite blocklists being constantly updated with newly verified links, knowing the exact volume of links circulating on the clear web alone remains a challenge for service providers, especially when it comes to attempting to quantify the proportion of links to child sexual abuse material yet to be detected and confirmed. Efforts are being made to work out the ratio of verified links to unverified links, but getting the full picture will likely remain unachievable for the immediate future.

Pressures on law enforcement authorities and reporting centres

The sheer volume of child sexual abuse material means that there are implicit pressures on law enforcement authorities. The National Centre for Missing and Exploited Children (NCMEC) reports that their CyberTipline received 29.3 million reports of suspected child sexual exploitation in 2021, which represents an increase of 35% from the numbers reported in 2020. The CyberTipline is accessible to law enforcement authorities across the world and aims to help them to prioritise the urgent and important situations within the truly extensive volume of reports.

It was acknowledged during the roundtable that the capacity, technological capability and resources of law enforcement authorities vary from jurisdiction to jurisdiction. From one side, it was stressed that it is important for technology companies to be as accurate as possible in reporting links to systems such as the CyberTipline because they do not want to overwhelm law enforcement teams with such an abundance of information that it makes it difficult for them to operate in an efficient and effective manner that focuses on getting children out of harm. On the other side, it was emphasised that companies should not use this to underreport links that are suspected but not 100% verified as taking users to child sexual exploitation and abuse sites.

The discussion also focused on the urgent necessity for greater training and resourcing of child protection agencies, specifically law enforcement and victim support services.

Link disruption, obfuscation and networks

Offenders are constantly on the move online. Many are dedicated to their cause, always making new channels, accounts and profiles knowing that there is a high risk of one or more of their points of dissemination being taken down or blocked. Many perpetrators use sophisticated techniques to spread themselves over multiple platforms and multiple accounts to share links to child sexual abuse online. In some cases, a link is taken down only for another (or multiple others) to be modified and uploaded to the same or a different account within minutes.

In the discussion, offender behaviour on social media platforms was likened to that of 'bot farms'[6], but for child sexual abuse material. Links that take users to spam-like bots that then either bombard the user with links to websites that host adult pornography and subsequently child sexual abuse material, or such bots can also take users through layers of links until they reach a site hosting child sexual abuse. Such bots often run ads and profit from these techniques along the way. With just a few additional clicks, offenders arrive at their content.

One of the most challenging technological developments utilised by offenders is the use of link-altering technologies. The shortening and modification of links can make it difficult for platforms, service providers and those curating blocklists to stay on top of blocking bad links. Therefore, solutions to the link sharing problem must seek to improve content moderation methods which may specifically address modified or shortened links.

The more significant problem lies with new content, which is advertised and distributed through peer-to-peer networks. Some participants have seen examples of offenders often using 'linguistic noise'[7] to obfuscate illegal content on peer-to-peer networks. Others have seen cipher[8] and similar techniques – with specific signposting applied – to obfuscate URLs that host child sexual abuse. In addition to cipher, offenders often have specific related keywords and abbreviations – their own "language" – which is used to avoid detection. Such behaviour helps offenders to avoid being caught by technologies that identify keywords and verified links.

However, it was also flagged that sometimes some offenders don't even make the effort to hide what they are doing online, with one participant citing that there are millions of examples of files shared on peer-to-peer networks with no attempt to obfuscate the language or the link at all. The reason for this was suggested that there is a lack of capacity or resource at a law enforcement level in some countries that hinders the ability identify and arrest offenders. It was suggested that more research is done to investigate how links shared in these peer-to-peer networks can be connected and integrated with photo and video analysis tools.

Preventative measures

Making it hard for users searching for child sexual abuse and disrupting offender behaviour is an important element of the response in reducing the sheer scale of URLs linking to child sexual abuse on the clear web. It was generally agreed that service providers need to be as disruptive and inhospitable as possible to bad actors. To do this, the focus must not solely be on tackling the supply of links to offenders and potential offenders, but also the wider demand for links to child sexual abuse online. It was noted that different platforms and regions have varying levels of capacity and infrastructure when dealing with prevention and reducing demand.

Some service providers lead users to help boxes or deterrence messaging when they think that the user is at risk of engaging in activities related to child sexual exploitation and abuse online. In such cases, users are given a prominent warning that child sexual abuse imagery is illegal, with information on how to report this content to trusted organisations or get help. As referenced in the Global Threat Assessment, deterrence messaging and help boxes can have positive effects. It was noted that warnings and help boxes are mandatory in some jurisdictions. Along with help boxes and promoting signposts to preventative support, the use of 'pop-up chatbots' was highlighted as a potential new innovative way to disrupt perpetrator behaviour and make them think twice about their activities.

One of the difficulties in disrupting offending is that it is largely difficult to measure the impact of such interventions. Concerns were also raised and understood that disrupting behaviour might push offenders to other platforms or the dark web. However, the majority of child sexual abuse material is still accessed through the clear web, end-to-end encrypted apps, or via peer-to-peer sharing. It was stressed that it is important to put a stop to the normalisation of child sexual abuse material being hosted on everyday platform that are easy access, user-friendly who often have large swathes of users.

6 – Bot Farms generate – often large volumes of - activity and interactions on the internet by using bots (autonomous software that carry out tasks) rather than people.

7 – Linguistic noise is when users modify their language through regular shifts in spelling, grammar, with slang and more.

8 – Cipher is an algorithm that can be used encrypt or decrypt online services.

It was also highlighted as a concern for some that if companies are too transparent with their policies, offenders may read up on their rules and identify workarounds to continue carrying out their illegal activities. However, it was also emphasised that a lack of transparency on policies and processes may increase or exacerbate user concerns surrounding privacy and security. Sharing useful and helpful information regarding these policies and processes in a thoughtful manner doesn't necessarily mean that offenders will learn new tradecraft.

Overall, it was acknowledged that the prevention of link-sharing to child sexual abuse material requires a whole-system approach that targets and addresses each layer of the threat – from improving the accuracy of the identification of links to illegal content to accelerating the speed at which they are blocked and reported. Beyond the bigger picture, there are a wide range of specific activities and behaviours that need tailored and nuanced responses in order to reverse this damaging trend.

Economic and social factors

The threat of link-sharing has similarities and differences across societies and economies. All around the world links to livestreamed child sexual exploitation and abuse are a particularly difficult problem to solve. Participants stressed that poverty was a huge factor in terms of understanding the crime of link-sharing in some regions, where it was not uncommon to see instances of direct family members or relatives coercing children into exploitative activities and sharing links to the crime for financial gain.

In families without the financial means to access devices and connectivity, there may often be a middle person or an "investor" who provides them equipment in order to facilitate content-making – often live streaming – and link-sharing. When it comes to links to livestreaming, it was suggested that the response should prioritise the safeguarding of the child or children at risk, with detection and analysis being subsequent focal areas. This was stressed as being of paramount importance when implementing a victim and survivor-centric approach.

Ensuring a response that protects all children is of paramount importance. It was stressed by one participant that they had noticed an increase in link-sharing to children who identify as lesbian, gay, bisexual and transgender, with a particularly noticeable increase in links to child sexual abuse materials depicting gay boys and young people.

Opportunities

This roundtable discussion was just one of several preparatory activities for the second round of the cross-UK government Safety Tech Challenge Fund (STCF). Following on from the learnings and outcomes from the first STCF, in which participants developed innovative prototypes that detect child sexual abuse material in and around end-to-end encrypted environments (whilst upholding user privacy), the second iteration of the fund seeks innovative tools to tackle the sharing of links to child sexual abuse content. Until Monday 20 March, UK-registered organisations were able to apply for a share of up to £700,000 for projects that help protect children by identifying and disrupting the sharing of links to child sexual abuse material online. The key opportunity here lies in content moderation methods which tackle the problem of link-sharing.

Recommendations

This discussion identified that there are multiple layers to solving the problem such as identifying links shared in public spaces, reporting links, identifying links in private text-based chats/messaging apps and detecting and disrupting link shortening/modification. There is currently no single type of response followed in this space with each company dealing with the issue of link-sharing in different ways. However, there were several recommendations that became apparent through the discussion:

- Experts agreed upon the importance of creating as hostile an environment as possible for offenders and potential offenders. Despite service providers being at different stages in their online safety journeys, there is a clear opportunity when it comes to disrupting the ease of use of tools on the clear web for offenders at every stage in the process.
- Constant technological innovation and the increased deployment of artificial intelligence will be required to respond to the scale and complexity of the threat. Many organisations have effective tools to tackle the spread of links to child sexual exploitation and abuse online. Some organisations are already deploying technology to confront the issue of links to general spam, fraud or malware and best practice can be taken from this space to better protect children online.
- Increased collaboration is a key foundation to an improved response to this threat. Developing a "global chain of trust" and cooperation – between internet service providers, telecommunication companies, technology companies, safety tech, law enforcement authorities, security agencies, reporting centres, hotlines and victim support services.
- It is important to acknowledge that the sheer volume of nefarious links reported can risk overwhelming law enforcement authorities and possibly reporting centres, leading to discussions about the need for greater training and resourcing of child protection agencies, specifically law enforcement and victim support services. There was a consensus, however, that the volume of reporting and material must not be a reason for not expanding industry efforts in tackling links shared and hosted on digital services.

In summary, the volume of link-sharing related to child sexual abuse online is increasing and therefore must be a key focus for preventative efforts to stop the spread of new and known illegal content.

Participants agreed that solutions to disrupt link-sharing to child sexual abuse material in public and private spaces need to be wide-reaching, cross-border and multi-sectoral.

This report has been written by WeProtect Global Alliance and reflects the conversation held under Chatham House rules.

The examples and opinions expressed in this document are reflective of the contributions from representatives who took part in the discussion and do not necessarily reflect the views of GCHQ, WeProtect Global Alliance or participating organisations.

For further information, please contact Eleanor Linsell, Advocacy Manager, at WeProtect Global Alliance: eleanor.linsell@weprotectga.org