NSPCC

End-To-End Encryption

Understanding the impacts for child safety online

NSPCC report based on research undertaken by PA Consulting
April 2021

EVERY CHILDHOOD IS WORTH FIGHTING FOR

Contents

Introduction

The world is arguably at a tipping point concerning citizens' attitudes to privacy and safety online. The global proliferation of internet enabled devices and high-speed internet has made social media and messaging services the mainstream means of communication for adults and young people, who increasingly rely on these services for social interaction, news, commerce, banking and accessing many government services. This has increased attention on the need to secure our personal data and safeguard privacy – a trend that has been amplified by inappropriate intrusion, such as the misappropriation of personal data, as illustrated by the allegations around Cambridge Analytica's electoral influence.

The result is widespread and rapidly expanding use of end-to-end encryption by platforms who provide web browsing and interpersonal communication services, such as private messaging and live streaming. Digital privacy rights proponents, including some human rights organisations, argue that people have the right to have private conversations online and service providers have an obligation to protect this privacy, citing examples such as the protection of pro-democracy or human rights groups reporting against excesses by authoritarian regimes. Just as end-to-end encryption creates challenges for identifying threats to child safety, it can greatly enhance the physical safety of journalists, minorities and others in contexts where they could be at serious risk of harm if their communications were vulnerable to hacking or state monitoring.

However, this is a complex area, with many issues to consider at the same time. In October 2019, Facebook CEO Mark Zuckerberg attracted international criticism for his company's plans to expand the encryption measures already applied to WhatsApp to both Facebook Messenger and Instagram. WhatsApp already uses 'End-to-End Encryption' (E2EE), whereby the content of each message is visible only to the sender and recipient. Unscrambling the message requires a private decryption key exchanged between correspondents, so that while the message may be intercepted, it cannot be viewed or monitored by the service provider, law enforcement or any other third party. Apple CEO Tim Cook attracted similar ire in 2015 when his company resisted FBI demands to disclose the passcode of an iPhone belonging to one perpetrator of a terrorist attack in San Bernardino, California, that killed 14 people. While the FBI cast its request as a limited emergency measure, Apple argued that the technique could easily be used again, making iPhone users more vulnerable to spies and thieves. But opponents argue that the privacy protection that encryption affords also extends to the privacy of bad actors (a fact acknowledged by Zuckerberg when he announced Facebook's plans to implement E2EE). In recent years there has been a growing awareness of, and backlash against, the proliferation of online harms and their growing impact on children and young people. Most severe in their scale and impact are the proliferation of online child sexual exploitation and abuse, and of content intended to terrorise or radicalise. But the harms extend to other conduct such as cyberbullying, stalking and harassment, and children's exposure to inappropriate conduct relating to suicide and self-harm.

The purpose of this report is to raise understanding of the impact that ubiquitous end-to-end encryption would have on children's online safety. The NSPCC commissioned PA Consulting to collate the viewpoints of a broad range of stakeholders, representing Civil Society organisations, industry, law enforcement and governments, to identify potential mitigations and trade-offs that should be considered. There are diverse and often conflicting opinions on the extent to which platforms should adopt online privacy and safety features, but the overwhelming majority of participants agree that child safety must remain one of the key considerations at the forefront of that debate. This report aims to provide a balanced narrative based on viewpoints obtained from engagement with experts across the community, with a child-centred focus.

This report examines the ways in which children are exposed to online harms, the current online safety system and legal intrusions permitted to allow citizens to enjoy a safe and secure online experience. It discusses how the industry's plans for expanding E2EE risks tipping the balance in favour of adult users, at the cost of children's safety. It then considers whether privacy and safety can ever practically co-exist and explores how the risks to children of E2EE could be mitigated, whether through technical countermeasures or through 'upstream' interventions that reduce children's exposure to harm.

Methodology

This report draws on interviews with experts in the field, alongside a targeted literature review, commissioned by the NSPCC. The views and statements do not represent the single views of an organisation or the individuals interviewed.

Key documents and reports were located through expert networks, desk-based internet research and recommendations from interview participants. This was an informative, rather than all-encompassing, non-systematic review of the literature of the area. Technical issues, perspectives and evidence from the literature were collated and themed for the report. These focussed on the historical and technical detail of encryption, how it is being implemented and by who, alongside current law enforcement and mitigations in place to protect internet users.

Key themes, issues and technical detail, including those drawn from the literature, were explored in semi-structured interviews with experts in the field from 16 organisations in the UK, USA and Australia. Interviews focussed on the implications of encryption on privacy and safety, with specific attention to the safety of children. Participants were drawn from tech industry, government, law enforcement, civil society and academia.* Individuals and organisations are not attributed to particular views or text in the report, unless stated.

Interviews took place from 4 September to 8 October 2020 and participants were given the opportunity to provide feedback on a draft prior to publication.

* Vivace, Crisp Thinking, Department for Digital, Culture, Media & Sport, Global Partners Digital, Google, Home Office, Internet Watch Foundation, National Crime Agency, National Centre for Missing and Exploited Children, TechUK, Thorn, TikTok, University of South Wales, US Department of Justice, WeProtect Global Alliance

Understanding the context

During this research, three analogies emerged which help to explain the nature of the challenge faced by children's online safety campaigners and the ways other industries have tackled similar challenges.

Online and Real-World Playgrounds

For a generation of parents, carers, teachers and officials who've not grown up as 'digital natives', there is limited understanding of online harms.

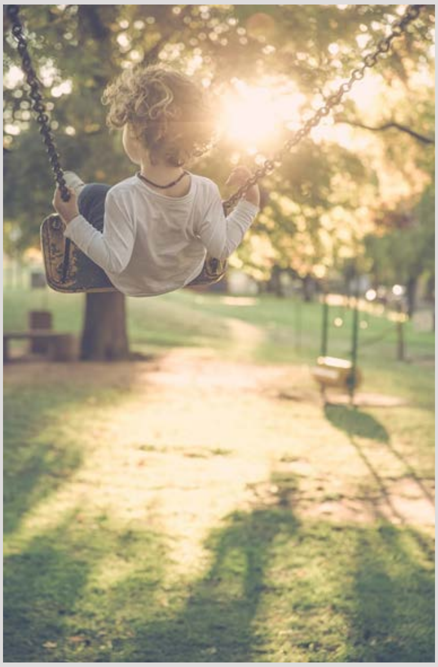

In their real-world (as opposed to virtual) lives, parents and carers have come to expect a level of safety when children visit the local playground. The play equipment, climbing frames and surfaces must conform to current safety standards and industry best practice; carers are alert to the risks of inappropriate contact with adults; and, for organised children's activities, supervisors must have Enhanced Disclosure and Barring Service (DBS) checks. The goal is not to eliminate risk, but to ensure the environment is safe-by-design for children.

Similarly, parents will look to expose children to a graduated set of risks as they get older: they will weigh up the risks and benefits as "no child will learn about risk if they are wrapped in cotton wool."[1] No responsible parent would allow very young children to play unsupervised, late at night, in a hidden playground where the equipment is dangerous and unidentifiable adults are able to interact freely with children. Parents and carers customarily relax restrictions progressively according to children's age, maturity and local perceptions of risk. Yet for many reasons, including the inherent complexities of the internet playground, many parents are unaware of the risks to which their children are exposed when playing online and it can be a highly challenging 'ask' for many, but not all, when considering the highly complex nature of many online harms, differing parental capabilities and the speed at which new platforms and technologies emerge.

The Evolution of Vehicle Safety

It has taken over a century to mitigate risk on roads and safety measures are constantly evolving. The "information highway" will be no different.

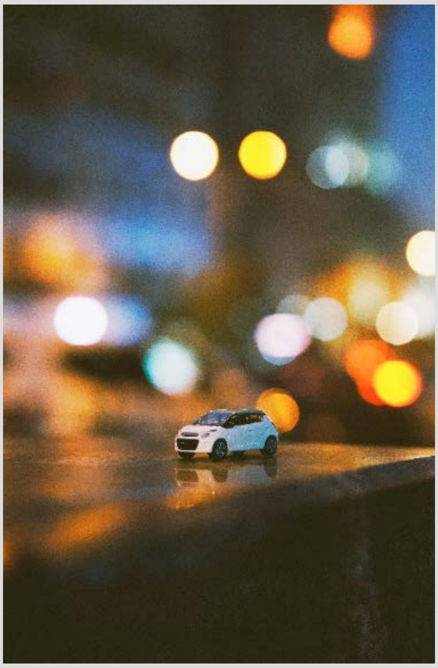

Comprehensive safety regimes typically take decades to build and are continuously evolving. As the automotive industry has evolved, cars have become faster and roads more congested. Fatalities and non-fatal injuries associated with automotive accidents at first grew exponentially, but were eventually reduced and then constrained by a wide range of progressive and complementary safety measures from: driver education (e.g. the theory and practical driving tests, speed awareness courses), vehicle safety measures (e.g. Euro NCAP safety ratings and annual MOT tests, safety belts, passenger airbags and side-impact bars), and highway safety measures (e.g. speed restrictions, lighting, barriers and traffic calming measures). By comparison with the 130+ year evolution of the automotive industry, the internet is in its infancy and citizens have not yet achieved the same level of awareness of the potential for harm, nor passive acceptance of measures designed to improve user safety. Society's awareness of the internet needs to also evolve so that any increased risk and harm is counterbalanced by improved safety measures that protect vulnerable groups (including children), and so that there is broad discussion and consensus around these trade-offs in civil society.

Privacy and Security in Online Banking

Citizens tolerate a balance of privacy and security when their personal financial wellbeing is at stake.

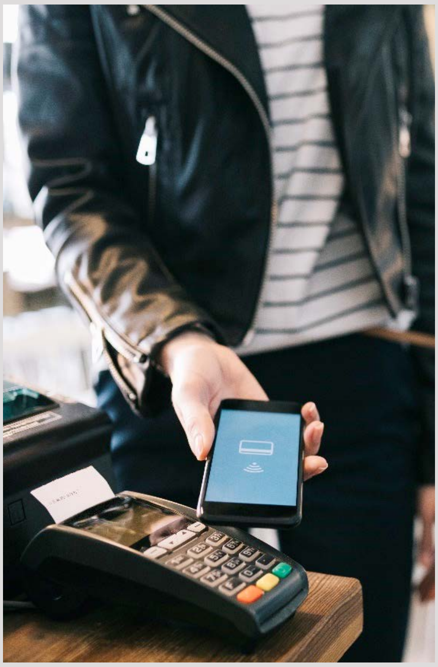

Many tech companies and human rights organisations support the availability of end-to-end encryption, arguing that it is necessary for the protection of human rights, including privacy and freedom of expression. They argue that any restrictions or limitations on private communications (whether for lawful surveillance or safety monitoring) increases individuals' exposure to unlawful monitoring or intrusion and undermines data security, stating that "you cannot make a backdoor that only good guys can go through"[2]. And yet an appropriate and workable balance for our online banking services has been widely accepted by society and legislated for. Through Financial Crime Regulation, banks are legally obliged to monitor transactions and inform law enforcement via Suspicious Activity Reports (SARs) if they suspect anything untoward. Consumers place a high level of trust in the banks to safeguard the security and privacy of their online financial transactions, yet they also tolerate a level of lawful intrusion that allows the banks to monitor for suspicious or fraudulent activity, the international SARs regime, and also to send text alerts or temporarily block transactions when unusual activity is detected. The transaction is secure between the customer and the bank, and between the bank and the payment recipient, but the bank can monitor and block harmful material whilst securing this 'back door' from malicious activity.

How children are exposed to online harms

Understanding the nature of online harms

1 in 3
internet users worldwide is a child[3]

The breadth and scale of harm to which children are exposed online is vast and growing. The DCMS Online Harms White Paper classified three broad types of harm: (1) illegal activity with a clear legal definition, (2) harms with a less clear definition which may or may not be illegal in all circumstances, and (3) underage exposure to legal content. Harmful content may be distributed in text, image, audio, video formats (or a combination, such as memes) or by live-streaming. For the purpose of this report, we have focused on the four ways that children typically experience harm:

1. **Interaction.** When a victim is exposed to a malicious actor online, e.g. grooming. Grooming takes place when an adult interacts with a child online, for the purpose of exploiting or sexually abusing them. Children are often coerced into online abuse or 'in person' contact abuse through blackmail.

2. **Harm Production.** Children may be the victims of exploitation when their abuse is photographed or filmed or through the production of self-generated imagery. This may be as a result of coercion, consensual age-appropriate sharing (often referred to as 'sexting' in an adult context) or through a desire for social affirmation. In some cases, the images or videos may be shared within the boundaries of a consensual relationship, but then distributed maliciously. In all cases the production, distribution and receipt of sexualised images of children is unlawful and this production generates new 'first generation' imagery which is harder to detect or block.

3. **Harm Distribution.** Child Sexual Abuse Material (CSAM) is shared illegally as an online commodity, usually in photo or video form. Whether this is 'first generation' content or images that have already been detected, classified and recorded as child abuse images or material being reshared, the ongoing proliferation results in the persistent revictimisation of children.

4. **Harm Consumption.** This occurs when a child or vulnerable adult is exposed to harmful content online, whether intentionally or by accident. This includes exposure to legal but age-inappropriate content that is harmful to their welfare or development, including legal pornography, content that glorifies or promotes self-harm, or extremely violent and graphic content.

Specific instances of harm are inherently complex and can span multiple platforms and forms of abuse. Those working in online child protection are well versed in the danger, scale and nature of the threat, and they understand how offenders are becoming increasingly sophisticated and organised in their pursuit of children online. However, the stakeholders interviewed in the compilation of this report broadly agreed that public awareness and understanding of the threat is lacking. While there is controversy about the grey area between free speech and hate speech in the context of measures to prevent radicalisation, this is not widely replicated with regard to online child abuse. Many adults may still not understand

Case Study: Understanding how identifying and prosecuting crimes could be made more challenging

Between January 2015 and his arrest in November 2017, Patrick McDonald used multiple Facebook Messenger accounts in which he posed as a teenage girl and contacted teenage boys, inciting them to send him images of themselves performing sexual acts. When interviewed by police, McDonald admitted he had targeted at least 500 boys. This investigation would not have been instigated without the messaging content provided to UK police by Facebook, which was accessed by their own safety systems. Content from Facebook Messenger, including sexualised conversations and the exchange of naked images, enabled police to identify accounts linked to the suspect and associated victims. McDonald was sentenced to four and a half years in prison in January 2018, after pleading guilty to making indecent images of children and inciting children to engage in sexual acts.

Had the content of these messages been end-to-end encrypted, it is possible that Facebook would never have identified this criminal behaviour and been able to make a referral to law enforcement. The offender may also have been less likely to make a confession, in the knowledge that the content of his communications would be much harder to access.[4]

that children's exposure to online coercion can be as traumatising as exposure to in-person contact abuse, that the severity of abuse can be greater (because the abuser has less fear of detection) or that the revictimisation through repeated sharing and viewing of abuse images inflicts a life-long trauma even after the child is rescued from their abuser.[5] This was notably recognised in a 2017 case in Sweden, where an offender was convicted of rape for coercing and blackmailing children to perform sexual acts online.[6] That this was treated as an equivalent offence was a promising recognition by the courts of the extent of trauma caused by online offences against children.

It is vitally important to increase public education around online harms and the importance of safety in product and platform design, in order to improve public engagement and dialogue about where the balance between issues like privacy and safety should be struck. In turn, this may prompt both government and industry to address the need for better safeguards for young people.

Understanding the scale of online Child Sexual Abuse

Child sexual abuse has increased exponentially by almost all measures over the last decade, with the number of offences prosecuted a small fraction of the estimated whole. In April 2017, following campaigning from the NSPCC and others, a new law came into force in England and Wales that criminalised sexual communication with a child.[7] Since then, over 10,000 offences have been recorded.[8] The National Crime Agency estimates that there is a minimum of 300,000 individuals in the UK who are a sexual threat to children.[9]

Globally, around 2 million new images are added to the Child Abuse Imagery database year on year and the Internet Watch Foundation's data tells us that 94 per cent of the CSAM they find online contains images of children aged 13 or younger.[10] Following the first national COVID-19 lockdown in the UK, the NSPCC sought data on online offences through a FOI request,[11] finding that during that time reports of child abuse images online increased by almost 50 per cent.[12]

The vast scale of this challenge is unlikely to ever be fully realised, but studies indicate that the prevalence of adult men who have intentionally accessed sexual images of pre-pubescent children online could be as high as 2-4 per cent.[13,14] Online sexual offences against children often take place across borders,[15] making it even more challenging to quantify the number of victims in the UK, however while this data is stark, it is likely it represents a small number of the true scale of these crimes.

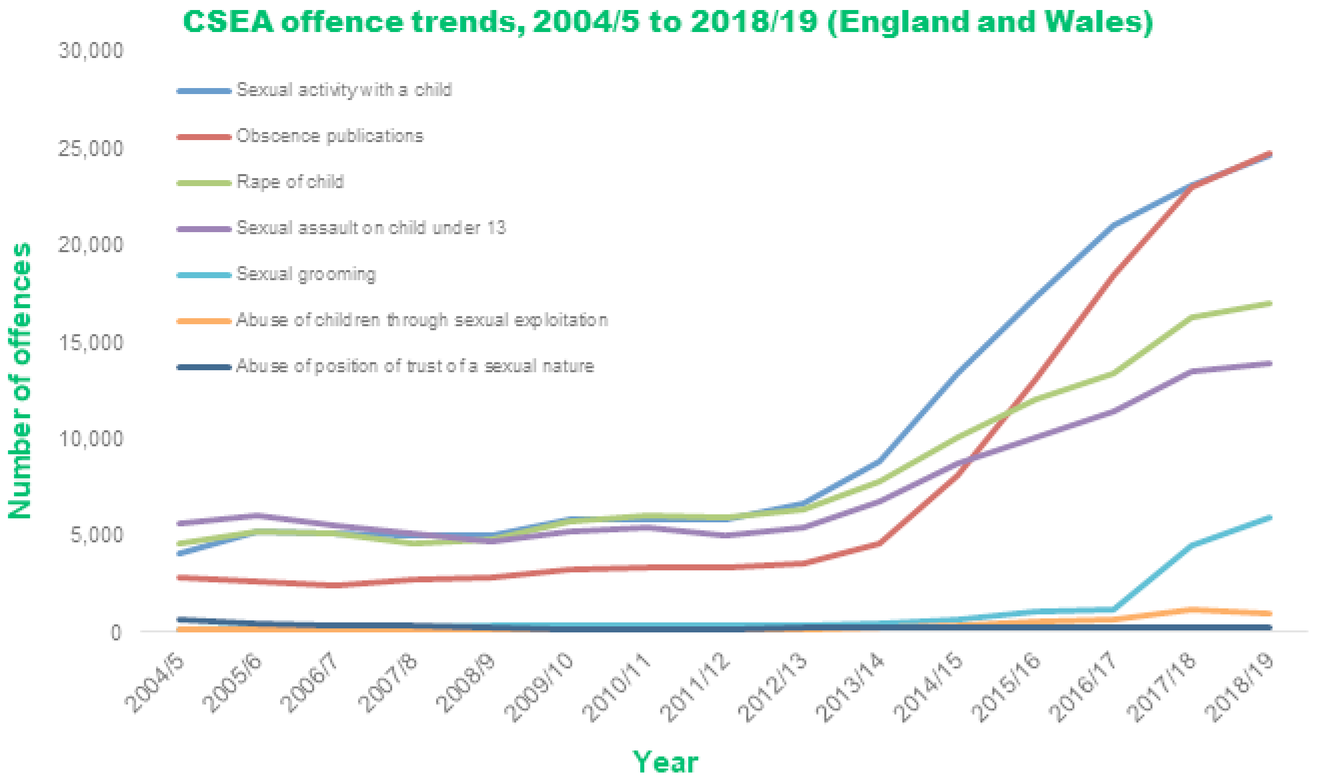

"Obscene publications" is used as a proxy for offences relating to indecent images of children (IIOC). The majority of obscene publications offences have been recorded against codes that cover IIOC offences. These offences will, however, will also include lesser offences, specifically the sending of explicit images between consenting minors. Operation Yewtree in 2012 led to an increase in victims' willingness to report and an increase in recorded crime. Source: Crime Survey for England and Wales 2018/19.

The existing safeguards for children's online experiences

A multi-layered and multifaceted internet safety and security landscape

The internet safety and security landscape is multi-layered and multifaceted. It consists of both proactive approaches such as safety-by-design, user applied safety tools and content moderation that disrupts abuse before it escalates, as well as reactive approaches including service provider monitoring, independent monitoring, and law enforcement use of investigatory powers. Whilst proactive and reactive approaches are not necessarily equivalent, they work in combination in complex ways to provide a safe and secure online experience. It is worth reiterating that encryption is not a binary choice between E2EE and nothing; at present, messages are encrypted to prevent unlawful access across mainstream social media platforms, the difference being that the platforms retain the capability to access that content to detect illegal activity. The application of E2EE will remove that capability and therefore the proactive monitoring of communication content. To understand the likely impact of E2EE, the existing safety framework is outlined through a five lines of defence model:

1. **Victims or other users reporting online harms:** Users may report harms directly through platform reporting mechanisms, directly to law enforcement, or may have user applied device level safety tools, which include applications using artificial intelligence to review messages as they are typed and flag to the sender that they are at risk of participating in bullying, abuse or grooming. At present, less than 1% of reporting comes from users.[16]

2. **Platforms conducting their own activities to detect, moderate, remove and block harmful content:** Service provider monitoring is done by reporting and reviewing tools on social media platforms, which are able to identify illegal content at scale that is then often manually reviewed by a moderator. These can include on-platform classifiers or photo matching tools. At present, this is voluntary and with variable outcomes.

3. **Third party independent monitoring organisations:** Organisations such as the Internet Watch Foundation (IWF), Canadian Center for Child Protection (CCCP) or National Center for Missing and Exploited Children (NCMEC) conduct independent monitoring, e.g. through web crawling, image detection, classification and hashing, URL blocking, or they supply datasets to companies to allow them to detect known illegal content. On-platform moderation may also be done by third-party independent detection tools such as SAFER, powered by Thorn, which can be applied to a platform to identify and classify both known and new CSAM.[17]

4. **Lawful intrusion by law enforcement:** The Investigatory Powers Act (2017) gives law enforcement the power to intercept some communications to prevent or detect a serious crime with a minimum three-year sentence threshold, which encompasses some (but not all) child sexual abuse offences. The earlier Regulation of Investigatory Powers Act (RIPA, 2000) creates the legal imperative for a suspected offender to provide passwords where necessary to prevent or detect crime. However, in all cases the law enforcement organisation or agency must have sufficient evidence to justify intrusion. For the purposes of protecting privacy, this is rightly challenging, however the opportunity to intercept illegal content is already subject to highly prescriptive safeguards without E2EE. While equipment interference is a new power under the IPA that could mitigate the impact of E2EE in some circumstances, it is a rightly limited tool that requires a warrant for use. Moreover, it deals with abuse after it happens, rather than focusing on proactive early threat detection to safeguard. It will not help to identify the scale of offending that will be lost in E2EE communications.

5. **Regulation:** The upcoming Online Harms Bill is expected to create an Online Harms Regulator that will play a key role in enforcing a statutory duty of care to protect users from harmful and illegal terrorist and child abuse content. At the time of writing, the detail of this is awaiting publication, but expected to be robust. Similar legislation is being developed overseas, such as the Digital Services Act in the EU and EARN IT in the US.

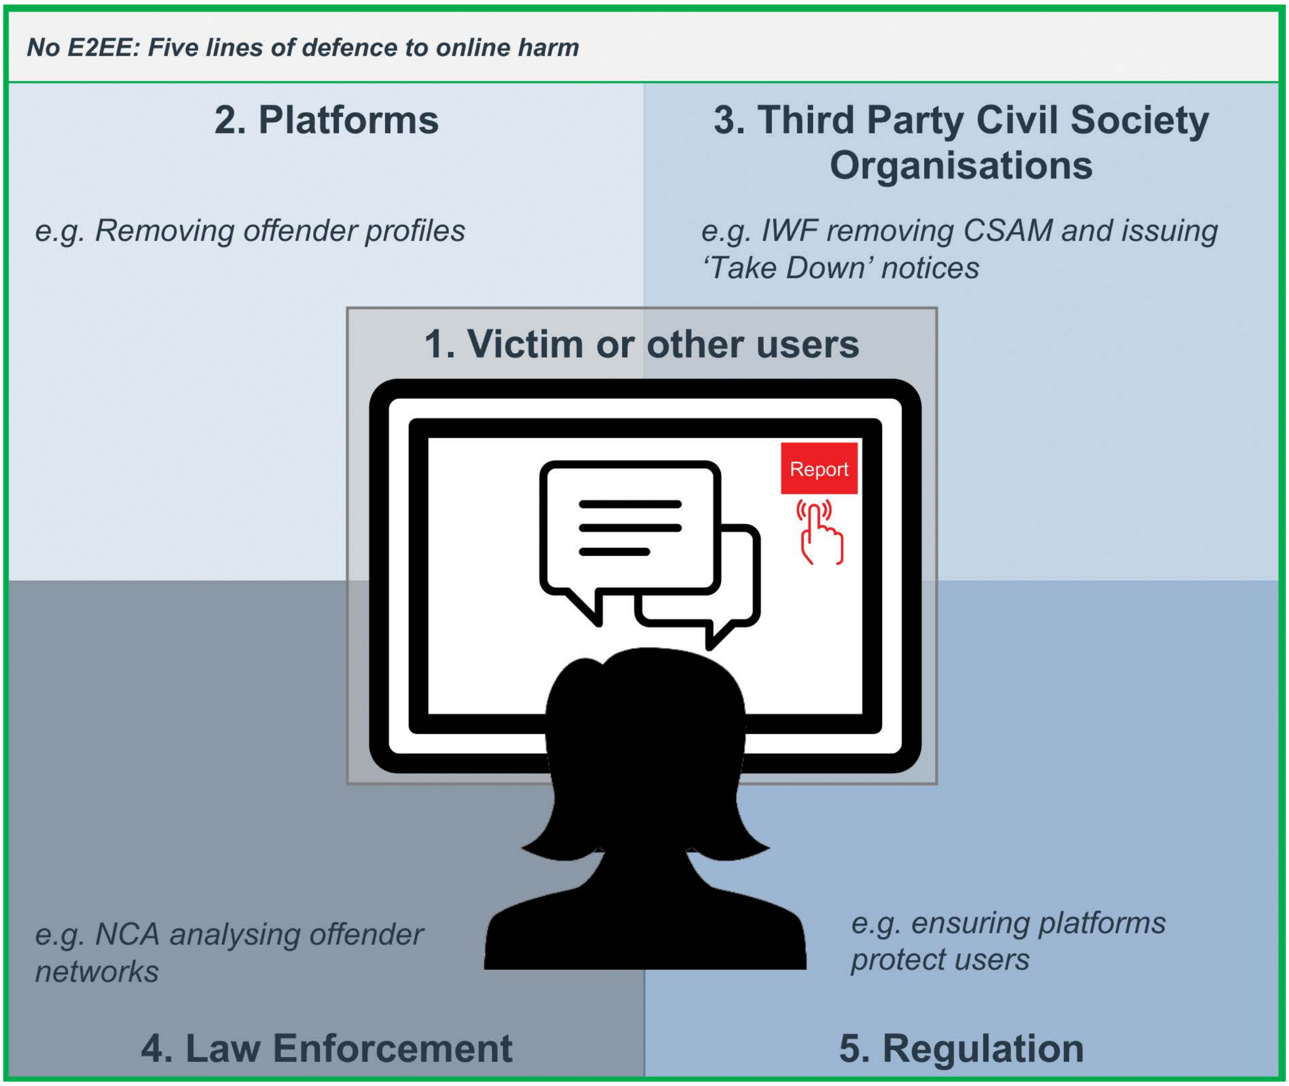
No E2EE: Five lines of defence to online harm
2. Platforms
e.g. Removing offender profiles
3. Third Party Civil Society Organisations
e.g. IWF removing CSAM and issuing 'Take Down' notices
1. Victim or other users
Report
e.g. NCA analysing offender networks
e.g. ensuring platforms protect users
4. Law Enforcement
5. Regulation

Each of these lines of defence plays an integral role in helping ensure a safe and secure online experience. However, each of these also provides different opportunity points for changes to be made to the system, either providing a more enhanced or reduced capability to prevent online harm. It is important to note that the actions of any party could impact the overall ecosystem and ability of another party to deploy safeguarding tools and techniques. Should E2EE become an even more widespread mainstream design choice, these system points provide opportunity for a careful and coordinated whole system approach to mitigations in order to build the best possible collective safeguards.

E2EE prevents or restricts the majority of existing lines of defence to online harms

The internet was not built with child safety in mind. The existing safeguards were retrofitted to reduce the risk to which children are exposed in an online environment. Should widespread E2EE be implemented as expected, it will have significant detrimental effects in children's online safety and remove platforms' ability to proactively identify harm within direct communications. The disruption of the Encrochat platform[18] and ensuing debate around the legality of undercover evidence obtained reinforces how E2EE platforms can be abused by criminals, and how difficult it is to uncover this abuse and hold them accountable, particularly when it comes to gathering admissible evidence. Where platforms are entirely encrypted, offenders can "apply Dark Web standard security and anonymity to their Surface Web interactions".[19] The impact of this on the current safety ecosystem could be catastrophic, with NCMEC reporting that 70% of Facebook's total referrals relate to Messenger and are therefore likely to be lost once that service is end-to-end encrypted.[20]

Encryption converts messages or images into a code that can only be translated back to readable content by a decryption key. A truly end-to-end encrypted communication uses "client to client" (device to device) encryption and so is only accessible by the device (and therefore the person) sending and the device (and person) receiving the message; neither the hosting platform nor law enforcement can

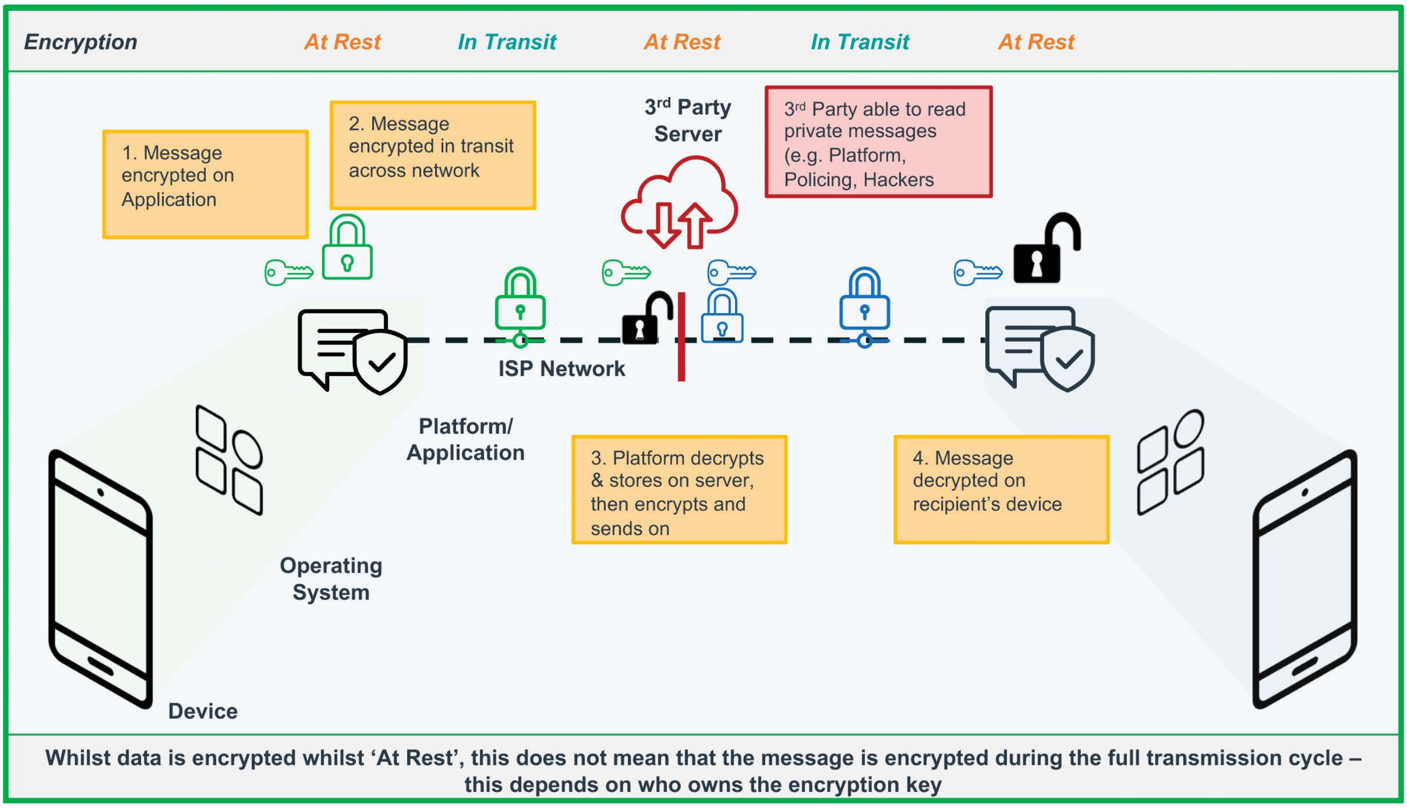
Encryption
At Rest
In Transit
At Rest
In Transit
At Rest
1. Message encrypted on Application
2. Message encrypted in transit across network
3rd Party Server
3rd Party able to read private messages (e.g. Platform, Policing, Hackers
ISP Network
Platform/ Application
3. Platform decrypts & stores on server, then encrypts and sends on
4. Message decrypted on recipient's device
Operating System
Device

Whilst data is encrypted whilst 'At Rest', this does not mean that the message is encrypted during the full transmission cycle – this depends on who owns the encryption key

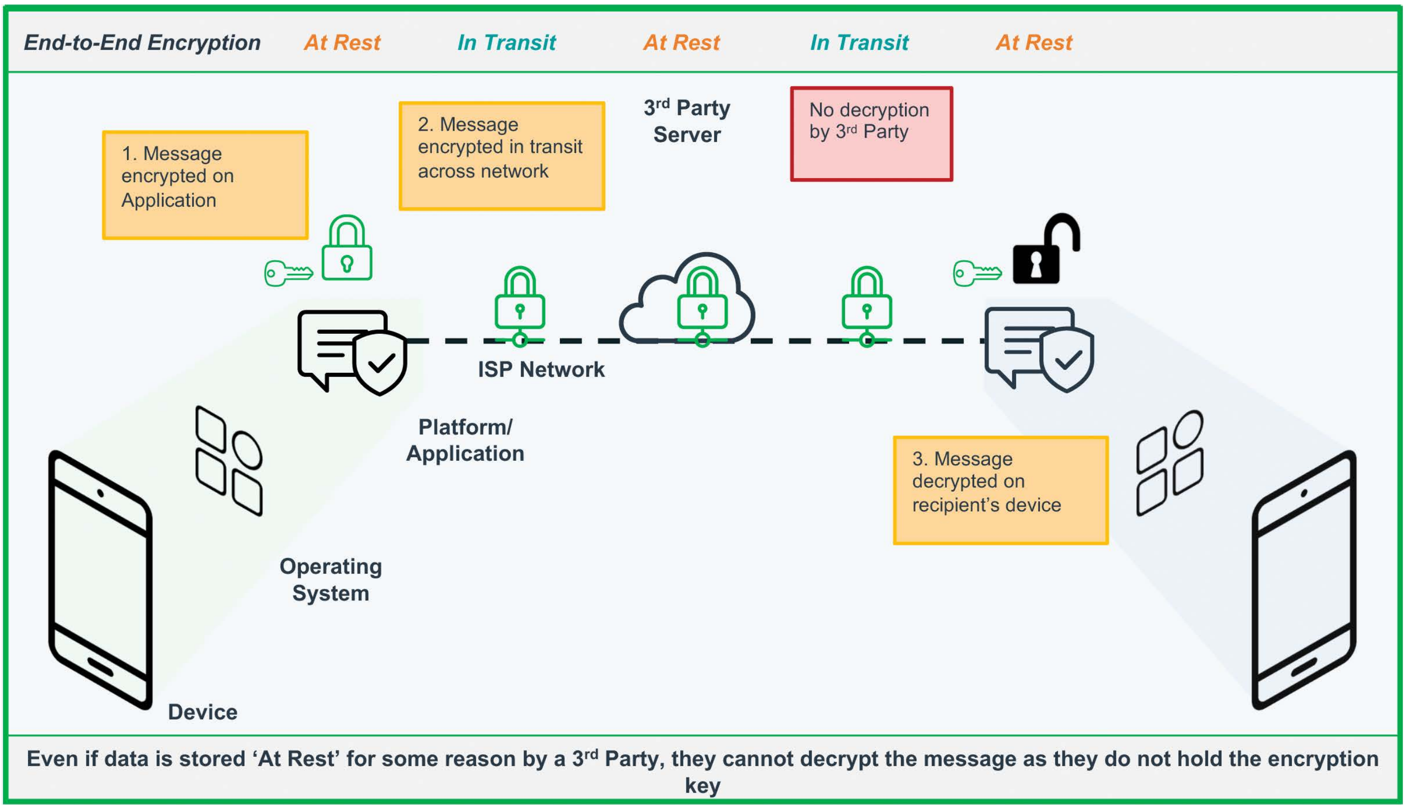
End-to-End Encryption
At Rest
In Transit
At Rest
In Transit
At Rest
1. Message encrypted on Application
2. Message encrypted in transit across network
3rd Party Server
No decryption by 3rd Party
ISP Network
Platform/ Application
3. Message decrypted on recipient's device
Operating System
Device

Even if data is stored 'At Rest' for some reason by a 3rd Party, they cannot decrypt the message as they do not hold the encryption key

see its content. E2EE can take many forms and is continuously developing. Describing an 'end-to-end encrypted future' in general terms ignores its inherent complexities. It is important this is reflected in any proposal for ways to ensure that children are safeguarded on E2EE platforms.

Much of the existing online safety infrastructure cannot be used on E2EE communications. At present, many platforms search messages to detect both CSAM and grooming. They do this using a variety of tools, such as photo matching technologies, or machine learning tools that identify patterns of language and behaviour that appear risky and flag them for human moderation. When the content of messages becomes inaccessible, significant investment will be required to develop workarounds and scan for known CSAM and it is likely to be impossible to use detection technologies to identify grooming behaviour or first-generation imagery. Homomorphic encryption technology* is one possible means of protecting data privacy while analysing its content, however there is debate about its ability to detect CSAM, how robust its privacy measures are and the extent to which it slows down communications. Any mitigations in the near future will almost certainly be less effective than the current ability of many platforms to detect for harmful content and are likely to rely on device-based technologies, use information from metadata, or rely on users to report harmful content.

For law enforcement to access an end-to-end encrypted communication, they are likely to need to identify, locate and physically access the device used. Should a device be locked, law enforcement's ability to access its content is highly restricted without cooperation from a suspected offender. RIPA requires cooperation to unlock a device, however this still requires human cooperation and sufficient evidence to demand it. The scale at which CSAM is shared and children are groomed online is such that a pursuit only approach to identify and prosecute offenders is not a feasible solution. In any case, E2EE drastically limits the evidence of abuse, so that even where a platform may suspect illegal activity and harm, without the content it is incredibly challenging for law enforcement to have sufficient evidence to obtain a warrant and pursue suspected offenders.

In 2019, NCMEC received over 15.8 million reports of CSAM from Facebook companies and only 205 from Apple.[21] Facebook Messenger and Instagram Direct are currently not default end-to-end encrypted, while Apple iMessage is. While Facebook receives significant criticism for the volume of reports of CSAM on its services, it also deserves credit for the scale of its reporting that reflects a more comprehensive approach to reporting than many of its competitors. Other social media platforms, with less developed tools for detection, are simply not reporting on the same scale. The true extent of CSAM on social media platforms is unknowable even in the current environment, with a lack of public reporting and transparency from internet service providers having helped create this knowledge gap. This will be exponentially worse when E2EE is implemented at scale. Even one report of CSAM on a user's messages is likely to indicate that when investigated, the user's account will be hosting hundreds of images. CSAE is already an under-reported crime, and a reduction of referrals from social media platforms will further undermine our understanding of the scale of the threat. In this instance, a reduction in reporting does not mean a reduction in prevalence.

The lines of defence for identifying and removing CSAM under E2EE are likely to be impacted as follows:

1. **Victims or other users reporting online harms:** Whilst there are programmes, for example around giving children and young people the confidence to report abuse and harms, victims are often the last to understand that they are being abused because of the dynamics of coercion and control exercised by abusers, and as such a victim-led reporting approach is likely to be fundamentally insufficient as a standalone approach regardless of whether the product is E2EE. While some user-applied safety tools may continue to be workable at the device level when communications are E2EE, this is unlikely to prevent or identify grooming behaviour or to be sufficient to recognise and block first generation CSAM. Evidently, offenders in groups seeking CSAM will not report it.
2. **Platforms conducting their own activities to detect, moderate, remove and block harmful content:** E2EE will likely eradicate a platform's ability to continue deploying current approaches to proactively search for child sexual abuse content in messages, on livestreams, or any other form of communication between users. NCMEC has said that without proactive identification by social media platforms, they will lose over half of their annual reports, as a conservative estimate.[22]
3. **Third party independent monitoring organisations:** Independent monitoring organisations will have similarly limited ability to identify CSAM or grooming. While some identify material in public forums which may remain open, many rely on referrals from platforms or users, which will be reduced as described above.
4. **Lawful intrusion by law enforcement:** Without sufficient evidence, law enforcement agencies cannot obtain a warrant to search a suspected offender's device for content. The barrier for evidence is rightly high, therefore if a platform is able only to share indicators of abuse without any information on the content of an offender's communications, the ability of law enforcement to pursue offenders will be frustrated and drastically reduced.
5. **Regulation:** E2EE does not reduce regulatory powers, however unless regulation expressly targets encryption, the use of E2EE could engineer away the ability to perform moderation, and frustrate or prevent compliance with regulatory requirements.

* Homomorphic encryption is a form of encryption that permits users to perform computations on its encrypted data without first decrypting it. https://en.wikipedia.org/wiki/Homomorphic_encryption

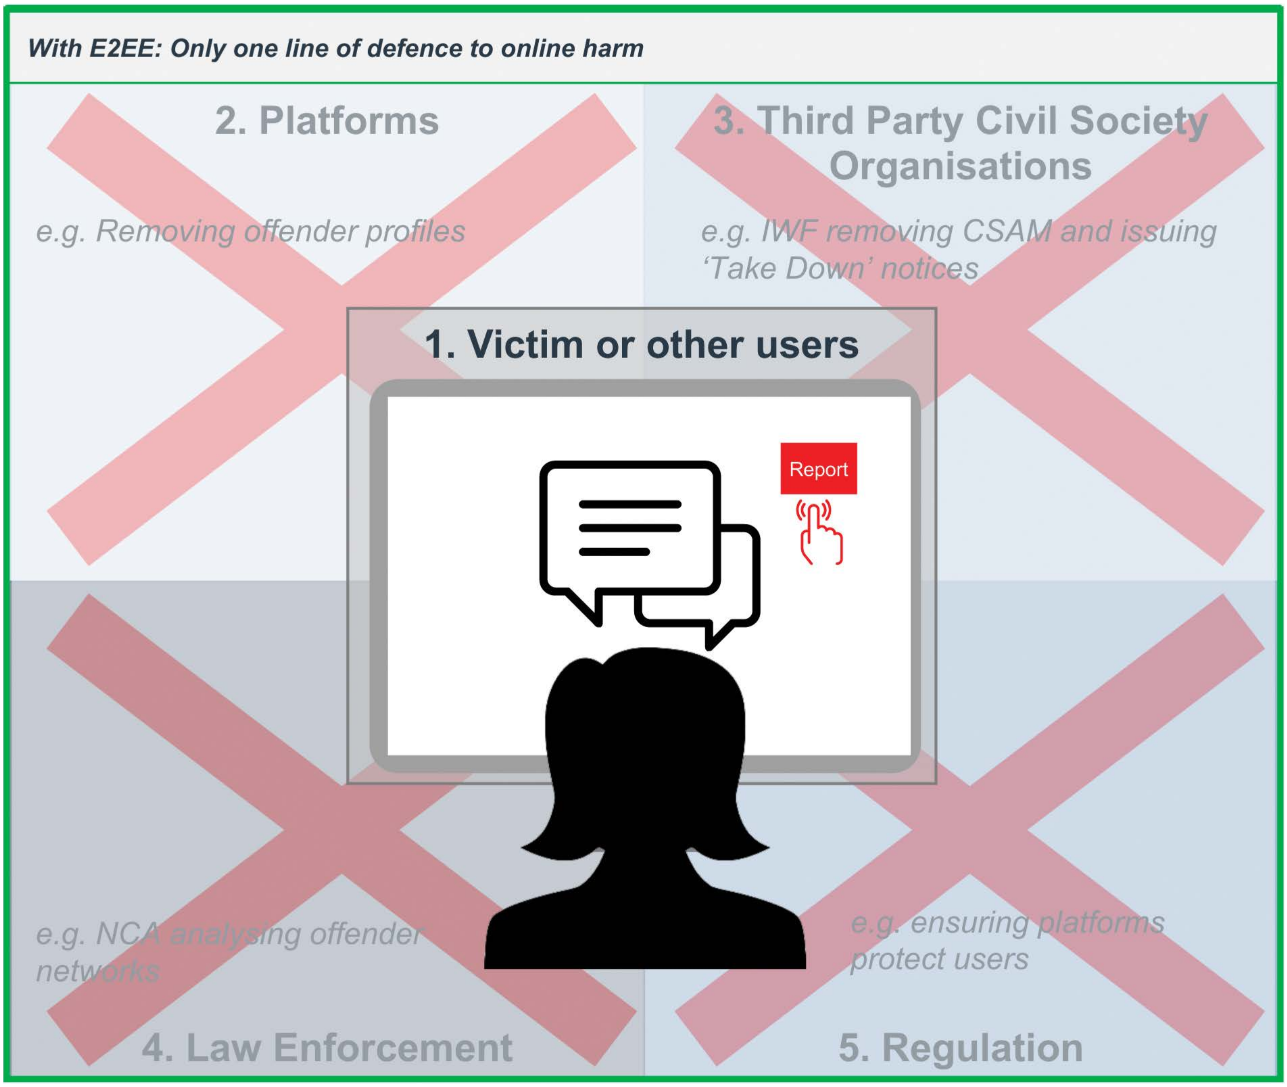

With E2EE: Only one line of defence to online harm

2. Platforms

e.g. Removing offender profiles

3. Third Party Civil Society Organisations

e.g. IWF removing CSAM and issuing 'Take Down' notices

1. Victim or other users

e.g. NCA analysing offender networks

4. Law Enforcement

e.g. ensuring platforms protect users

5. Regulation

Scale of impact of E2EE for online harms

Because the scale of online harms is not fully known, it is challenging to reliably ascertain the true scale of impact that E2EE could have across all online harms to children. However, it is possible to ascertain a sense of scale for the worst type of online harm to children, which is sexual abuse and exploitation.

As set out in the table below, in 2019 over 16.8m reports globally of online child abuse and exploitation on mostly mainstream social media platforms were made to NCMEC's Cyber Tipline, the majority of these from messaging services. A 2019 report by WhatsApp[23] (a Facebook company) reveals that around 250,000 user profiles per month are removed on suspicion of sharing exploitative imagery of children. As WhatsApp messages are end-to-end encrypted, these removals are based on other indicators, such as the use of CSAM in a profile picture or a group name that references CSAM. It is a reasonable assumption that not all groups or users sharing illegal material on encrypted platforms make their actions clear in their user information and users joining groups are unlikely to report.

Figure 1: NCMEC CyberTipline Reports in 2019 (top 10 ESPs out of a total of 148). Note this includes subsidiaries of the companies listed.

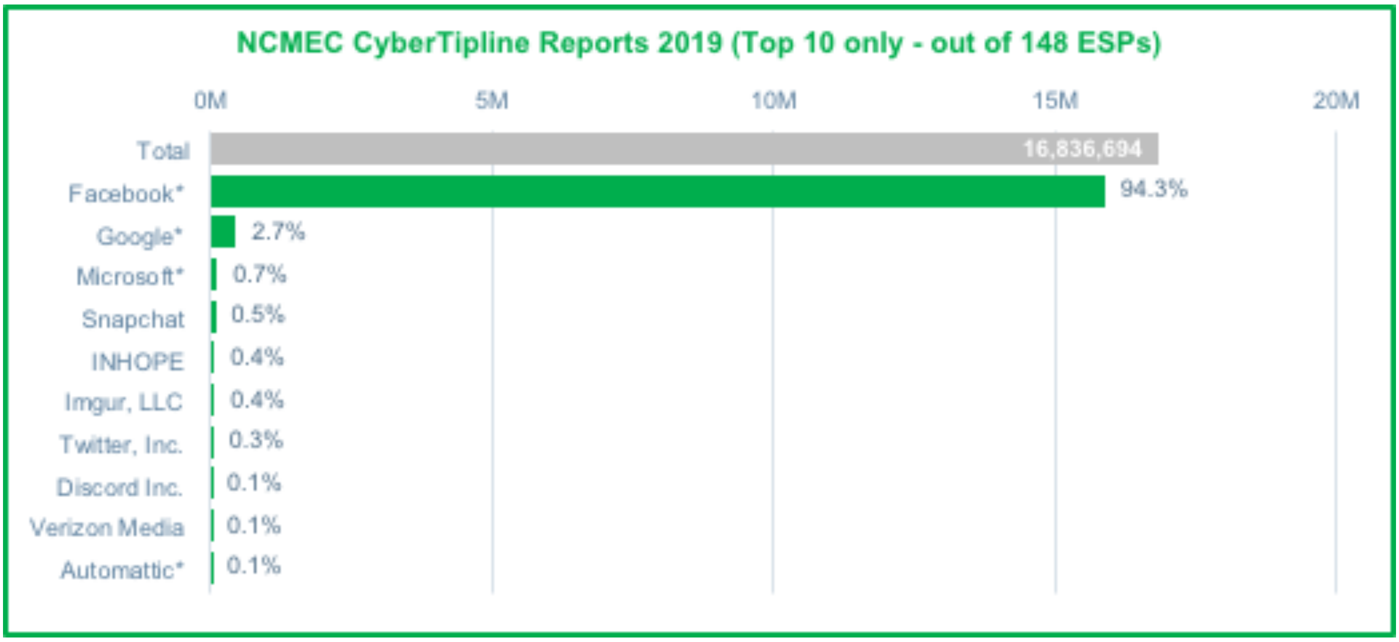

How E2EE risks prioritising adult privacy over child safety

The privacy, security and safety debates are well-established.[24] The risk of E2EE to vulnerable people is undeniable, however society at large is significantly more focused on concerns around data privacy. There are good reasons for concern, with secure communications being highly important for groups such as LGBTQ+ people in authoritarian regimes, religious minorities in theocratic nations and defenders of democracy against authoritarian states.

However, the specific safeguarding needs and digital rights of children and young people are often excluded from discourse around internet governance, with the discussion around adults primarily focused on or advocating for personal privacy above all else. In reality, privacy and safety need not be diametrically opposed, and there is urgent need for greater nuance and a broader, more informed civil society dialogue that considers all of the implications of E2EE in a rounded and considered way.

Children's safety and human rights

The UN Convention on the Rights of the Child[25] establishes both the right to be protected from sexual exploitation and the right to privacy, though without providing for scenarios in which these appear to contradict one another. The European Convention on Human Rights (ECHR) implies positive obligations on the state to protect children from abuse, which are grounded in the prohibition of inhuman or degrading treatment and the right to privacy.[26] In particular, the state is expected to protect children from sexual abuse which amounts to inhuman or degrading treatment when it knows, or should have known, that a child is at risk. This includes where there are repeated warnings or other weighty evidence of abuse, with an emphasis on the need to proactively respond and respond where a there is a suggestion that harm is taking place. Ultimately, privacy is not an absolute right and there are (limited) lawful reasons for state intervention in the interests of public safety in the offline world. The European Court of Human Rights has found that in weighing the interest of the child against the protection of potential abusers, the fight against child abuse should be given significant weight.[27] It is right that these protections are replicated for online harms.

Although encryption may offer children privacy benefits, there is a clear risk that E2EE could significantly impair their safety and may also, somewhat counterintuitively, weaken their privacy. Where E2EE is not applied, resharing of known child abuse imagery can be significantly limited through the use of photo hashing tools that crawl the internet to identify and report content that has already been classified and taken down by law enforcement. The victims seen in this material continue to suffer, with some reported to spend huge amounts of time searching for and reporting their own pictures.[29] In the Canadian Center for Child Protection's 2017 Survivors Survey, 20 per cent of those interviewed reported having been identified by someone who had seen their abuse online.[30] Child abuse survivors also have a right to privacy, which because of the ease with which abuse images could be circulated in end-to-end encrypted environments, is threatened by end-to-end encryption.

What is hash matching?

Image identifying technologies, such as Photo DNA, use databases of known illegal media files to detect CSAM and other illegal content. They create unique hashes to represent each image, which can be matched with copies of that image to identify where it has been reuploaded or distributed online.[28]

Understanding the impact: a survivor's story

The Phoenix 11 is a group of survivors of child sexual abuse, whose abuse was recorded and distributed online. The following statement from a survivor discusses the impact of the abuse on her life and her concerns about end-to-end encryption.

"Every day I live with the knowledge that there are images of my abuse, rape, and torture as a child being viewed on the internet and these images continue to be distributed to this day. As a teenager and young adult I've been stalked by paedophiles who have viewed these images both via social media and in real life. I've gone through periods where I've been scared to leave my house after some scary experiences. I have had to change my name and take extreme measures to ensure my own personal safety and the safety of my children from paedophiles who stalk and threaten me. I worry about how to keep my children safe. It's hard to trust them with anyone or let them out of my sight. I worry about what will happen when they are old enough to be on the internet, the things that they might discover about our family history, and how to tell them about the dark things that have happened in my life before they find out another way.

When I think of changes to privacy with tech companies implementing end-to-end encryption, it scares me. I already feel such a lack of control as to who views and shares the images of my abuse. These images of my abuse have already been viewed and traded so much, I do not want that to be made easier in any way. I don't think there should ever be a trade-off when considering the safety of children. I would feel much safer knowing that companies held off implementing end-to-end encryption until they had the appropriate tools for monitoring child sexual abuse images within that framework."[31]

An appropriate balance between children's and adult's rights

Tempering a child's experience in the interests of their safety is not a new concept; this is ideologically uncontroversial in the offline world, in line with social norms and the expectation of increased freedoms over time. This is also a broadly understood concept in respect of many aspects of online services, for example the age-gating of certain sites, and the Information Commissioner's Children's Code requirements for a user appropriate experience dependent on age.

However, in crucial aspects of technology design and internet governance, the needs of and risks to children are often not appropriately balanced, or we see a highly polarised debate. The needs of, and risks posed to, all users must be considered, but not at the expense of those who are least represented in policy making forums and without appropriate protections being in place.

In the absence of regulation, or where there is unhelpful ambiguity, we have often seen industry choose to prioritise one side of the trade-off between privacy and safety, for example focusing predominantly on data privacy benefits rather than seeking a more balanced approach. Any future regulatory approach must enable a more balanced approach that factors in children's needs and gives appropriate prioritisation to the duty of care for child safety online.

When considering potential mitigations for E2EE, there are also nuances based on whether they're focused on the device (e.g. an iPhone) or service (e.g. WhatsApp) level, as well as difference between encryption products that affect data "at rest" and data "in motion" (data sitting on a device, or in transit between devices). Lawful access debates can quickly reach a stalemate, because it's hard to define what an "acceptable" level of intrusion to a child's communications looks like in order to keep them safe. A parent can apply nuance and intuition in the offline world, with graduated choices based on well-established norms of societal conduct. Industry has yet to successfully replicate that approach on social media platforms, such that parents can reasonably feel their children are being kept safe from harm.

Mitigating the risks associated with end-to-end encryption

Technical countermeasures

End-to-end encrypted communications remain encrypted from a device controlled by the sender to one controlled by the recipient, where no third parties, not even the service provider or the host platform, can access the content in between. A third party in this context means any organisation that is not the sender or recipient directly participating in the conversation. So, where 'true' E2EE is applied, this means that all the current mainstream safety strategies used by social media companies, other than user reporting, are no longer possible.

Instead, platforms rely on unencrypted metadata[32] and behavioural analysis to detect anomalies ('signals') which may indicate harmful behaviour, although these are highly limited in their ability to detect, assess and respond to harm. Regardless of other uses, metadata indicators would not solve the problem of adult offenders trading CSAM with each other.

In its '*Strategy for a more effective fight against child sexual abuse*', published in July 2020, the European Union has called for "solutions which could allow companies to detect and report child sexual abuse in end-to-end encrypted electronic communications. Any solution would need to ensure both the privacy of communications and the protection of children from sexual abuse and sexual exploitation, as well as the protection of the privacy of the children depicted in the child sexual abuse material."[33] Under the EU Internet Forum, the Commission has launched an expert process with industry to map and preliminarily assess, by the end of 2020, possible technical solutions to detect and report child sexual abuse in encrypted electronic communications, and to address regulatory and operational challenges and opportunities in the fight against these crimes.

There are three possible types of technical solution that permit automatic content detection to continue: 1) device related, 2) server related, and 3) encryption related solutions. Their application is based on a legal interpretation that detection tools don't infringe upon privacy because the algorithms don't 'understand' the content – they simply block its transmission or flag it for human review if it matches a digital hash signature for known child sexual abuse imagery, or specified keywords. Some privacy groups strongly oppose automatic detection as a major infringement of people's fundamental right to privacy (for example, the ongoing debate about the temporary derogation from the ePrivacy Directive in the EU),[34] even if its intent is limited. The following options each afford distinctly different levels of security, privacy and safety.

Device-related solutions

There are two mainstream 'on-device' approaches, both of which rely on hashing and matching technology to work. This "creates a unique digital signature (known as a "hash") of an image" which is compared against a database of hashes to find copies of the same image.[35] For example, when a known image of child abuse is sent on a system using this technology, this can be flagged so that the image is removed and the offender referred to law enforcement. In its current state, hashing and matching only works on communications that are not E2EE. In order to manage this, industry could move towards conducting all forms of abuse detection (hashing and matching) on the device, or to separate on-device hashing and on-server matching.

On-device hashing and matching affords relatively high safety levels by enabling proactive auto-detection of harmful content before the message is transmitted or received, whilst enabling fully encrypted peer-to-peer communication of appropriate content. It could, in principle, be incorporated into the application or the operating system level, although the feasibility of storing up-to-date hash libraries and matching algorithms on device is questionable, and there are risks associated with users subverting or reverse-engineering detection tools.

Moving the matching function to the server reduces the technical complexity. This method sends both hashes and the encrypted message to the server, and the message is only released if there are no matches with the hash database. In this case, the processing of the user's hashed content at the server increases the risk to the privacy of the communications, and the on-device hashing algorithm remains vulnerable to subversion. In general, on-device solutions offer the lowest interference with private communications but the highest risk of subversion. This approach works best for user/parental-applied safety measures such as smart keyboards and image filters that incorporate safeguarding assistants.

Server-related solutions

The baseline approach is to apply a server-side 'back door' to enable the Platform/Service provider (or authorised public authorities) to decrypt and assess the content of a specific communication. However, this method affords poor safety, privacy and security, principally because it precludes preventative auto-detection for harmful content (it would be neither feasible, nor proportionate to screen every message). It also creates a vulnerable access point for malicious actors to exploit.

A better approach would be to establish a 'secure enclave' on the Cloud, that can decrypt the communication and perform the same operations and checks as done in unencrypted communications, but in a secure, closed off environment where neither the user's data nor detection operations are visible to the platform/service provider. Whilst this is not strictly 'end-to-end' encryption, it affords the equivalent level of privacy unless the third-party server is compromised.

Encryption-related solutions

At present, the only advanced technology which offers the potential to balance security and privacy is homomorphic encryption, a form that enables calculations to be performed on encrypted data without decrypting it first. In this model messages could be transmitted using E2EE, using device-level homomorphic encryption of the images and videos, which can then be hashed and matched at the server during the course of the message transmission. A proof of concept for images exists but additional research and development is needed to reduce processing times as these are currently slow for images and infeasible for video files.

Safety by Design

Until the legal arguments are resolved and the regulatory framework is agreed, international governments including the UK have been urging technology firms not to develop systems and services in ways that empower criminals or put vulnerable people at risk, but to prioritise the protection of their users and the wider public when designing services. The Australian eSafety Commissioner's Safety by Design principles are one example of government working constructively with industry to achieve this.[36] There is increasing need for a holistic approach in engineering schools, where at present students are taught to build secure apps but not to build safe apps. This is critical to building a resilient ecosystem.

Case Study: How TikTok is using a Safety-by-Design approach

In April 2020, TikTok announced changes to its Direct Messaging policy, removing access for under 16s.[37] They also block users from sending unsolicited messages to those they're not friends with, or from sending any photos or videos attached to a DM. In January 2021, further safety measures were announced, including setting accounts for under 16s to private by default and additional measures to encourage young people to communicate only with those they know.[38] TikTok's Head of Child Safety Public Policy in Europe, Alexandra Evans, tells us that these are common sense decisions that support the platform's mission to "spread joy and inspire creativity", which can only happen when users have a safe space in which to share content.

Child protection "is a democratic good, not a commercial prerogative"[39] and for online safety to become enshrined in the fabric of how tech industry operates, and how it takes product decisions, there must be clear commercial and regulatory drivers. Some of the contributors to this report told the NSPCC they believe that significant change will only come when there is a financial incentive to prioritise safety at every stage of the product and device design process, through fines, shareholder behaviour, or similar. It is possible that upcoming legislation in the UK will make steps in this direction.

Policy and regulation

Although tech firms may choose to implement adequate technology focused mitigations for end-to-end encryption, regulation may be considered or determined to be necessary. The internet is one of the most unregulated industries in the world, described as the "product of the 90s zeitgeist of laissez-fair neoliberalism".[40] The lack of accountability for industry is widely seen as the most significant challenge the child safety community faces and it's not a new problem. Section 230 of the US Communications Decency Act 1996[41] provided tech companies with the indemnity that has allowed illegal content to proliferate on the internet ever since[42].

In anticipation of Online Harms Regulation, the UK is at a pivotal moment. It is expected that new emphasis will be put on industry's Duty of Care to protect children online, but the detail will be key.

Case Study: What can be learnt from the approach to Modern Slavery?

Modern Slavery and Human Trafficking legislation required that all organisations which meet certain criteria publish an annual statement, outlining the steps taken to deal with modern slavery risks in the supply chain.[43] This must be signed off at director level.

At present, no equivalent approach exists to ensure that companies are not working with those who trade in or allow CSAM to proliferate on their platform. The modern slavery legislation has not only increased transparency around this challenge, but also allowed the UK to play an ambassadorial role in encouraging similar approaches overseas.

As internet and device access proliferates across the Global South, developed nations must take a collective responsibility to protect children who are likely to become increasingly at risk in the developing world.[44] The UK will have a critical ambassadorial role in encouraging other nations to adopt similar regulatory approaches, and as with modern slavery legislation, may be able to make significant strides.

Other countries will need to decide whether they too take a legislative approach, as is currently happening in the United States. There is likely to be a strong aversion to regulation there and civil society's ability to work with lawmakers and influence public interest in this area is critical. The EARN IT Bill has the potential to be game changing in the demands it makes of industry, and interestingly it is mostly technology and service agnostic. If passed, particularly alongside similar legislation in the UK, the global impact is likely to be significant.

Education and culture change

We have some way to go in terms of public understanding of the complexity of child sexual abuse online. In consultation for this report, one government official described public awareness of the online CSAE problem as "terrifyingly naïve". Wider public understanding of the privacy/safety trade-offs is needed in order to inform the response. Should widespread E2EE go ahead, then the need for an informed public will be even more critical to enable broad civil society and public discourse on where the appropriate balance between safety and privacy lies, in order to inform the appropriateness of responses and to help move the response upstream.

There is a need for large scale public awareness raising and renewed efforts to educate children, their parents and teachers about the reality of CSAE and how to take a sensible a measured approach to safe interactions online, including E2EE. Regulatory and technological changes will only emerge when there is sufficient public support and demand for them, and should reflect a balanced understanding of the risks and opportunities which need to be considered.

Those working with survivors estimate that only 5 per cent of children who appear in CSAM are ever known to authorities,[45] meaning that the vast majority of victims are never safeguarded and do not receive the psychological and physical health support they need. The societal cost of the ensuing challenges is difficult to quantify, though conservative estimates reach into the billions of pounds annually for the UK alone.[46] It is much easier to process the danger that a child faces crossing the road than it is to understand the reality of the risk they might face in their bedroom, playing on their phone.

The vulnerability created by children being inadequately supported in the real world, while potentially exposed to damaging content online, is significant. The inquest process into the death of Molly Russell[47] highlights one example of the type of harm that children can be exposed to, and this vulnerability almost certainly extends to child sexual abuse as well.

Part of any move forward in education around online sexual abuse will involve positive education around sexual behaviour and relationships. Multiple stakeholders engaged for this report identified a broader range of harms, including the normalisation of sexual violence more generally as being contributors to the child sexual abuse problem online.[48] Sex and relationship education must demonstrate what a healthy, consenting relationship looks like, and some stakeholders have suggested that exploring sexual relationships in an age-appropriate setting forms part of the graduated approach to giving children increased autonomy online, but only where privacy can be adequately protected too.

A whole system response is required for incremental gains at every opportunity.

Technology based countermeasures, underpinned by safe design principles, policy and regulation, are the most tangible way to ensure the strongest safeguards are in place for online child safety. In addition, greater education and cultural awareness should provide additional layers that collectively help to build a whole system approach to keeping children safe from sexual abuse online.

Among the suggested mitigations for UK government is the need to approach online CSAM as we would a comparable threat, the obvious one being counter terrorism. The counter terror response in the UK is much more coordinated, despite there still being multiple agencies working on the same problem (and indeed, entire agencies focused almost exclusively on it). It's not clear why this isn't the case for CSAM, particularly when the number of victims is comparatively much greater.

There are existing functional models for tackling a threat of this nature, however government appears to be lacking the will to take this approach. Through all stakeholder conversations for this report, a common theme prevailed. Any approach must be rooted in collaboration and shared responsibilities. Technology is the final line of defence to protect children online, but it cannot resolve the challenge presented by over 300,000 individuals in the UK who want to access child abuse material and if systems aren't safe by design they enable rather than prevent abuse. There are significant opportunities in culture and policy change that shouldn't be overlooked. There is no silver bullet; multiple layers of safeguarding must be built up to block harm at every opportunity, through an incremental gains approach.

Case Study: how the WePROTECT Global Alliance uses a multi-stakeholder approach to protect children online

WePROTECT Global Alliance (WPGA) brings together 98 countries alongside 41 private sector companies, and 44 civil society organisations and international institutions with a common goal: to end the sexual abuse of children online. Since 2014, WPGA has worked across borders to encourage a global response to the threat and to develop collaborative frameworks to support governments and industry to keep children safe.

An independent organisation with international reach, WPGA has an unmatched ability to bring together diverse stakeholders and empower marginalised voices in discourse on combatting CSAE. As an inherently global threat, this is critical to making progress; no single entity, no matter their resource or size, can solve this problem alone. Collective action will form the foundation of finding a way forward on E2EE and WPGA's Model National Response[49] and Global Strategic Response[50] frameworks provide a blueprint for working collaboratively to protect children.

A whole system approach that delivers incremental interventions at every level can best safeguard children online

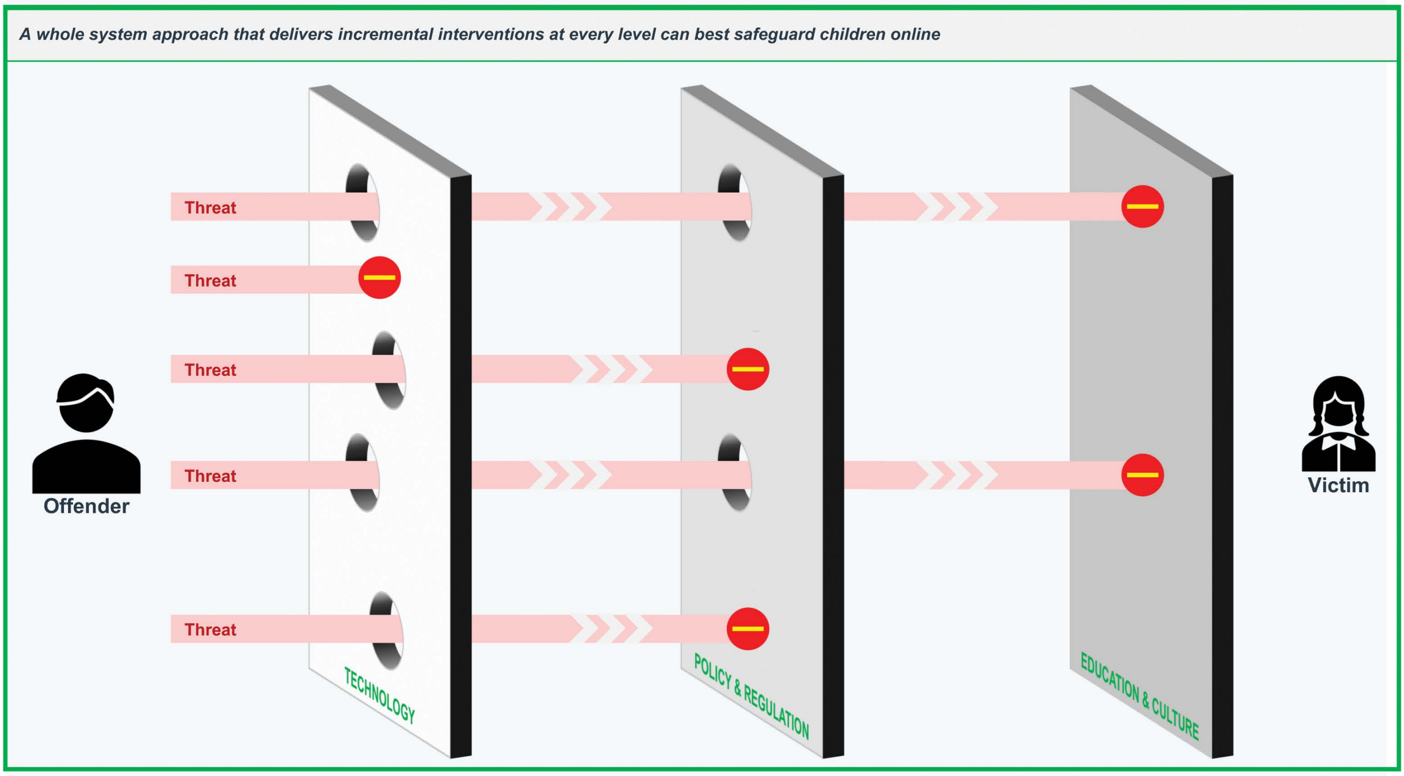

Conclusions

At the time of writing, the imperative to bring stakeholders together in the best interests of child protection is as critical as ever. Just as we are recording stark increases in instances of online abuse connected to COVID-19 lockdowns, Google is the latest major tech player to announce plans to test end-to-end encryption in messaging and debate in the EU is escalating over an agreement on a child safety exemption to new privacy rules.[51]

The good news is that it is entirely possible to design and build safer platforms. The analogies set out at the start of this report – the development of safety standards in cars, the expectation of safety on playgrounds and the coexistence of high standards of safety and security in financial transactions – serve as a reminder that the ingenuity exists to find a way through challenges of this nature. However, as discussed throughout, there are complicated trade-offs to be made in order to reach an appropriate balance, and it is unlikely that industry will get there alone. Regulation that incentivises industry to put the needs of children and young people at the heart of their design approach is not necessarily an impediment to commercial success but is a much-needed driver of culture change. Any regulatory approach should be principles-based, in order to futureproof progress to the greatest extent possible.

The impact that end-to-end encryption has on our online ecosystem is not a fait accompli. Humans participate in every stage of the process, from design to implementation and we have the ability to build safe platforms if we choose to. Should we choose not to, we run the risk of creating an online world in which children themselves are the main line of defence against their own abuse and we rely on victim reporting alone. For us to adequately safeguard children, safety must become a line item in everyone's agenda, every time an update or new service is planned. Through collaboration, shared pragmatism and continuous improvement to the innovation that is applied to keep children safe on the internet, safety must become baked into what tech industry does.

The discourse on E2EE needs to move beyond a polarised debate that sees privacy and safety pitted against each other, towards a broad and balanced understanding of the risks and opportunities faced. This must reflect both the needs of children and adults, in recognition that one third of UK internet users are under the age of 18[52] and that this is an inherently vulnerable population. There is a need to greatly increase engagement from the public and to ensure that civil society provides a balanced view to inform responses, including from tech firms, governments and regulators, to ensure it strikes the right balance and that we have the right protections in place.

References

1 https://www.hse.gov.uk/entertainment/childrens-play-july-2012.pdf

2 https://www.telegraph.co.uk/news/2019/10/04/facebooks-zuckerberg-defends-decision-encryption/

3 https://www.unicef.org.uk/press-releases/unicef-make-digital-world-safer-children

4 Case study provided by DCMS and the National Crime Agency, November 2020

5 https://learning.nspcc.org.uk/research-resources/2017/impact-online-offline-child-sexual-abuse

6 https://www.independent.co.uk/news/world/europe/online-rape-conviction-bjorn-samstrom-grooming-webcams-sex-acts-victims-uk-us-canada-uppsala-court-a8086261.html

7 https://www.gov.uk/government/news/new-crackdown-on-child-groomers-comes-into-force.

8 https://www.nspcc.org.uk/about-us/news-opinion/2019/recorded-online-sexual-grooming/

9 https://www.nationalcrimeagency.gov.uk/who-we-are/publications/437-national-strategic-assessment-of-serious-and-organised-crime-2020/file

10 Data provided by DCMS, November 2020

11 https://www.nspcc.org.uk/about-us/news-opinion/2020/instagram-grooming-crimes-children-lockdown/

12 Data provided by DCMS, November 2020.

13 Dombert, B., Schmidt, A. F., Banse, R., Briken, P., Hoyer, J., Neutze, J., & Osterheider, M. (2016). How common is men's self-reported sexual interest in prepubescent children? Journal of Sex Research, 53(2), 214-223.

14 Seto, M., Hermann, C. A., Kjellgren, C., Priebe, G., Svedin, C. G., & Långström, N. (2015). Viewing child pornography: Prevalence and correlates in a representative community sample of young Swedish men. Archives of Sexual Behavior, 44(1), 67-79.

15 https://link.springer.com/referenceworkentry/10.1007%2F978-94-024-1555-1_43

16 Data provided by DCMS, November 2020

17 https://www.thorn.org/blog/how-safers-detection-technology-stops-the-spread-of-csam/

18 https://news.sky.com/story/encrochat-what-it-is-who-was-running-it-and-how-did-criminals-get-their-encrypted-phones-12019678

19 WePROTECT Global Threat Assessment 2019, p12

20 https://questions-statements.parliament.uk/written-questions/detail/2021-03-22/173224

21 https://www.missingkids.org/content/dam/missingkids/gethelp/2019-reports-by-esp.pdf

22 PA Consulting interview with NCMEC, 30 September 2020

23 https://faq.whatsapp.com/general/how-whatsapp-helps-fight-child-exploitation/?lang=fb

24 https://carnegieendowment.org/2019/09/10/moving-encryption-policy-conversation-forward-pub-79573

25 https://www.ohchr.org/en/professionalinterest/pages/crc.aspx

26 https://www.echr.coe.int/documents/convention_eng.pdf

27 Nordisk Films v Denmark, no 40485/02, 2005 (Where the Court found that a national court's decision to compel a journalist to provide footage that assisted in the prosecution of paedophilia was proportionate) and Juppala v Finland, no 18620/03, 2008 (where the Court held that defamation laws which had a chilling effect on reporting of potential child abuse violated the right to freedom of expression)

28 https://www.microsoft.com/en-us/photodna

29 PA Consulting interview with Dr Michael Salter, UNSW, 10 September 2020

30 https://protectchildren.ca/pdfs/C3P_SurvivorsSurveyFullReport2017.pdf

31 Phoenix 11 Survivors' Story provided by the Home Office, November 2020

32 Metadata is data relating to a communication (but excluding the content of a communication) indicating e.g. the origin, destination, route, format, time, date, size, duration, or type of underlying service.

33 https://ec.europa.eu/home-affairs/sites/homeaffairs/files/what-we-do/policies/european-agenda-security/20200724_com-2020-607-commission-communication_en.pdf

34 https://www.eurochild.org/news/eurochild-signs-joint-letter-asking-to-adopt-a-temporary-derogation-to-the-eprivacy-directive/

35 https://www.microsoft.com/en-us/photodna

36 https://www.esafety.gov.au/about-us/safety-by-design

37 https://www.bbc.co.uk/news/technology-52310529

38 https://www.theguardian.com/technology/2021/jan/13/toktok-to-tackle-grooming-with-curbs-for-young-users

39 PA Consulting interview with Dr Michael Salter, UNSW, 10 September 2020

40 PA Consulting interview with Dr Michael Salter, UNSW, 10 September 2020

41 https://www.law.cornell.edu/uscode/text/47/230

42 https://www.irishtimes.com/news/court-rules-aol-not-responsible-for-child-porn-1.376990

43 https://www.gov.uk/guidance/publish-an-annual-modern-slavery-statement#who-needs-to-publish-a-statement

44 WePROTECT Global Threat Assessment 2019

45 PA Consulting interview with Dr Michael Salter, UNSW, 10 September 2020

46 WePROTECT Global Threat Assessment 2018, pg 24

47 https://www.bbc.co.uk/news/uk-england-london-54307976

48 WePROTECT Global Threat Assessment 2018, pg 15

49 https://protect-eu.mimecast.com/s/v1VrCj2AGs3KM90cW4JLO?domain=weprotect.org

50 https://protect-eu.mimecast.com/s/kbTVCk56JFr86NvtV0KZn?domain=static1.squarespace.com

51 https://www-politico-eu.cdn.ampproject.org/c/s/www.politico.eu/article/europes-thermonuclear-debate-on-privacy-and-child-sexual-abuse-2/amp/

52 Livingstone, S.,Carr, J. and Byrne, J. (2016) One in Three: Internet Governance and Children's Rights. Innocenti Discussion Paper No. 2016-01, UNICEF Office of Research, Florence, Page 15

NSPCC

Everyone who comes into contact with children and young people has a responsibility to keep them safe. At the NSPCC, we help individuals and organisations to do this.

We provide a range of online and face-to-face training courses. We keep you up-to-date with the latest child protection policy, practice and research and help you to understand and respond to your safeguarding challenges. And we share our knowledge of what works to help you deliver services for children and families.

It means together we can help children who've been abused to rebuild their lives. Together we can protect children at risk. And, together, we can find the best ways of preventing child abuse from ever happening.

But it's only with your support, working together, that we can be here to make children safer right across the UK.

nspcc.org.uk

EVERY CHILDHOOD IS WORTH FIGHTING FOR