Ofcom
making communications work
for everyone

Exploring Safety in the Online Content Journey

Publication Date: 8 October 2021

Contents

Section

Annex

1. Overview

The online ecosystem is becoming ever more relevant to Ofcom's work, and our remit has recently expanded to explicitly include online content for the first time. The Audiovisual Media Services Regulations, passed in November 2020, introduced requirements for UK-established video-sharing platforms (VSPs) to take appropriate measures to protect children from harmful content and to protect the general public from content containing illegal material and incitement to violence or hatred. Ofcom was appointed the regulator for this new VSP regime and the UK Government has also confirmed its intention to appoint Ofcom as the regulator for online harms and Ofcom is preparing for this planned new role (Ofcom, 2021). Given our new online duties, we are committed to deepening understanding of the wider online ecosystem and sharing our knowledge with people, industry and governments. As part of this commitment, in July 2020 Ofcom together with the Information Commissioners Office (ICO) and the Competition and Markets Authority (CMA) formed the Digital Regulation Cooperation Forum (DRCF) to support regulatory coordination in online services, and cooperation on areas of mutual importance (Ofcom, 2021a). The DRCF also includes the Financial Conduct Authority (FCA), which joined as a full member in April 2021 (UKGOV, 2021).

This report focusses on technological innovations in the internet infrastructure and how these impact on the journey of content – whether harmful or otherwise – to UK users. Its aim is to increase understanding of the way the internet is changing, so we can help ensure these advancements benefit people and that they are protected from harm. Technical changes to the internet may also make monitoring and moderating harmful content more challenging, so it is important for us to anticipate and understand these changes. The report forms a part of our technical programme of research looking at the tools and technologies that can impact the safety of internet users as set out in Ofcom's Plan of Work 2021/2022 (Ofcom, 2021b) and sets the foundation for future online technology projects which will be taking place over this period. This report focuses on current and near future developments: see our *Internet Futures* publication for potential longer-term developments (Ofcom, 2021c).

While we are interested in the general implications that evolutions in internet technology can have for people and society, we do not attempt to assess or describe in this report the policy implications of the technologies we have spotlighted, including service consistency, reliability, privacy and security; consumer harm; and future legislation and regulation.

2. Introduction

The societal benefits of the internet in stripping away geographical barriers, transforming access to information and enabling people to connect and to share information instantaneously are enormous. For millions of people around the world, the internet has become an intrinsic part of their daily life. However, coupled with the many benefits of these innovations, this relatively effortless and immediate access to online services creates challenges related not only to security and privacy of users but also in ensuring their safety.

Whilst from a user perspective, accessing and sharing content and information with people or businesses located anywhere in the world is as easy as tapping an app or opening a browser, behind the scenes a vast network of organisations and technical communication infrastructure - both physical and virtual - is in place to move content from creator to recipient. Different challenges and opportunities for mitigating potential harm exist at different stages along this 'content journey'.

Extant approaches to preventing access to harmful or illegal online content, including commonly used blocking or filtering methods, have relied on the fact that the internet was not designed to be inherently secure. Originating in the 1960s, the internet evolved from a framework of the first computer networks, which consisted of the military ARPANET and university computers connected to each other through it. At that time, commercial use was not envisaged and the original users knew and trusted each other, providing little incentive to prioritise built-in security (InterconnectCommunications, 2011).

Over time, numerous attempts have been made to improve the security of underlying internet architecture, for example by taking steps to secure the Domain Name System (DNS) (ICANN, n.d.) and the Internet Protocol (e.g. IPSec (IETF, 2005)). The revelations by Edward Snowden regarding the extent of global surveillance led to increasing calls for the internet to better protect not just the security but also the privacy of online communications, leading the Internet Engineering Taskforce (IETF) to explore the development of new protocols and standards for encryption at all stages of the online content journey (MAPPINGProject, 2014).

The move towards greater levels of encryption within the internet has implications for the effectiveness of some existing approaches to online safety, suggesting new approaches may be required to ensure continued protection for users. In order to navigate this changing landscape, an in-depth understanding of the underlying architectures, protocols and process flows which make up the foundations of the internet is required. This knowledge must be coupled with an awareness of the implications of new and emergent technical standards and protocols when considering measures which can be taken to enhance safety of internet users, to ensure an effective and proportionate response both now and in the future.

This report provides an overview of the existing and upcoming challenges and opportunities for enhancing online safety through the lens of a typical content journey. It explores how the technical infrastructure and architecture of the internet enables content sharing and identifies the organisations responsible for providing, supporting and administering its technical functions. The report does not aim to provide an exhaustive analysis of the online safety ecosystem – rather, the objective is to identify, at a high level, potential opportunities for further research which will inform our future work in this area.

3. The role of the internet in the online content journey

The content journey starts with a piece of multimedia content which is generated by a user, be it a private individual or a professional creator such as a news organisation or broadcaster. For example,

- A written message is typed into a messaging application.
- A still image is captured using the device camera.
- A voice message is recorded via the user device embedded microphone.
- A video clip is captured using the device camera and microphone.

Once created, the content may be shared via the internet through an on-device application and ultimately be made available to view by others. For the purposes of this report, we will focus on the journey of a piece of content from a user to a social media platform.

A visualisation of a user's typical interaction with social media appears at figure 1 below.

Figure 1. Typical Consumer Interactive Journey with Social Media

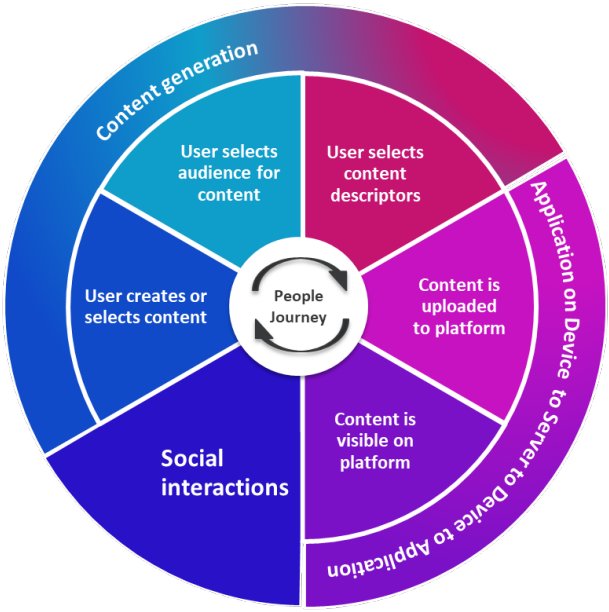

Source: Ofcom

Whilst from the user perspective the act of sharing or accessing content on the internet appears to involve interacting only with a specific platform, service or application, users are relying on a swathe of additional interrelated protocols, organisations and services, including Internet Service Providers (ISPs), Domain Names System (DNS) providers and network operators, which function behind the scenes. In order to understand the online safety eco-system, it is first necessary to have a basic understanding of how the internet operates.

Structure of the internet

The internet consists of a "network of networks" linked in a hierarchical structure by a multitude of network hardware and fibre-optic cables and other telecom technologies. At the heart of the internet is the internet backbone, where the largest providers (known as "Tier 1" networks) are linked together. The individual networks are referred to as Autonomous Systems (AS), consisting of unique allocations of IP addresses[1] administered independently from each other. The protocol that defines how each AS communicates route information and steers traffic is known as the Border Gateway Protocol (BGP) (Cloudflare, n.d.). BGP makes it possible to establish a route for data to travel from its source to its destination.

Figure 2 – Schematic overview of the internet

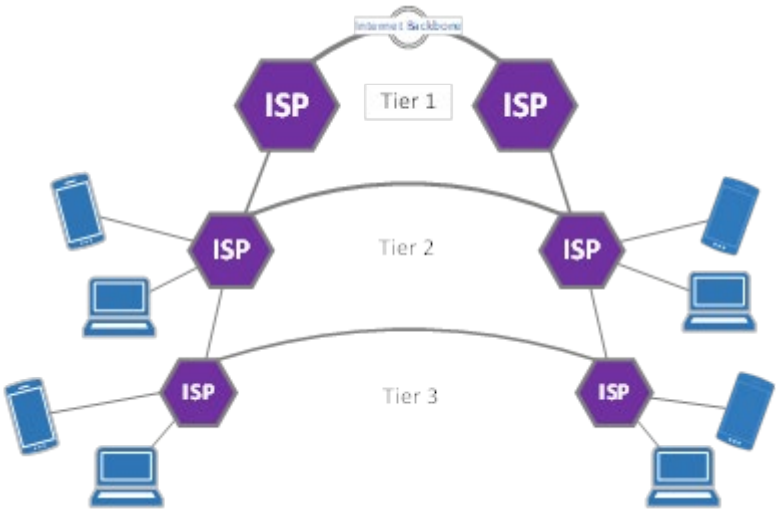

Source: Ofcom

IP addresses and the domain name system (DNS)

During the emergence of the public internet in the 1990's and 2000's, a single, standard format for IP addresses was in use. This format - Internet Protocol version 4 (IPv4) - was sufficient to support approximately 4.3 billion devices. However, due to the growth of the internet, the pool of IPv4 addresses has been exhausted. To meet increased demand, IPv6 was created, sufficient to support approximately 340 trillion trillion trillion devices (TheInternetSociety, n.d.).

As IP addresses are hard to remember, a protocol called the Domain Names System (DNS) is used to translate IP addresses into more memorable domain names. The DNS is often referred to as the "telephone book of the internet" and is a public, decentralised database handling requests for, and responses to, domain name lookups. When a request for a domain name is made, domain name "resolvers" identify the servers responsible for that domain sequentially, starting with the right-most ("top level") domain label. The resolution process typically starts with a query to a root server. These root servers will then respond with a referral to more authoritative servers. The resolver then

[1] IP addresses are numerical designations which identify individual nodes within a network.

iteratively queries these servers until an authoritative answer is received. Figure 3 illustrates the domain resolution process.

Figure 3: Example domain resolution process

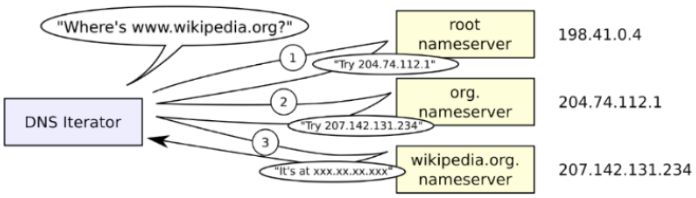

Source: Lion Kimbro (original author); Mess (derivative work), CC0, via Wikimedia Commons[2]

How content is sent across the internet

The individual networks comprising the internet communicate via a standardised set of rules - a protocol suite - referred to as the Transmission Control Protocol and Internet Protocol (TCP/IP) networking model (Cisco, 2011). Initially implemented by the US Department of Defense in 1983 and subsequently delegated to the Internet Engineering Task Force (IETF), the TCP/IP model specifies how data sent via the internet (or indeed across any network) should be packetised, addressed, transmitted, routed and received. The functionality is divided into four abstracted layers[3], with each layer performing a discrete yet interrelated set of tasks to facilitate delivery of data between devices even when at a global distance (see figure 4).

Figure 4 – TCP/IP Networking Model

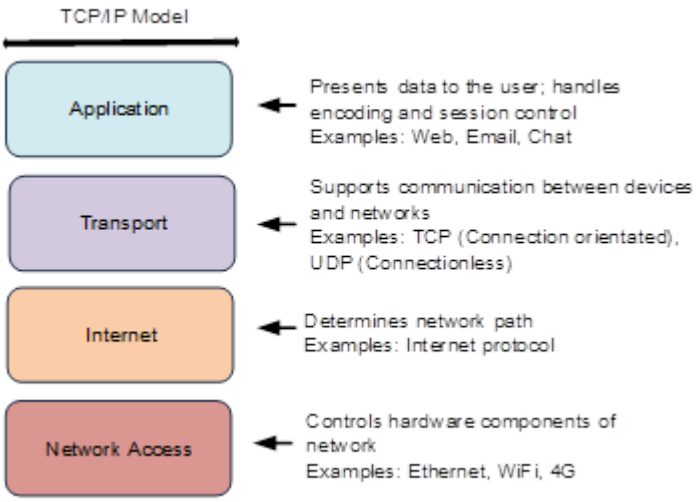

The Network Access layer is concerned with low-level transmission characteristics i.e. electronic signaling, hardware interface with the physical network medium such as, fibre optics or Ethernet.

---

[2] Available at https://commons.wikimedia.org/wiki/File:Example_of_an_iterative_DNS_resolver-it.svg

[3] The TCP/IP model was designed to be hardware independent, and as such some later representations of the model include a notional fifth abstraction layer – a "physical" layer, which sits at the bottom of the stack.

The Internet layer is primarily concerned with the creation of data packets, routing and forwarding of packets to their destination. The Transport layer defines connection requirements or reliability of data sent across the network. Finally, the Application layer defines interaction with the lower Transport layer functionality and the external end-user program making the request, such as a web browser or email client.

The internet is not the same as the World Wide Web (WWW). The WWW is an information-sharing model which is built on top of the internet and consists of the collection of webpages which comprise online platforms and services. As such, online platforms can be visualised as operating "above" the application layer.

Internet governance

The administration and governance of the internet involves a global network of organisations operating at both national and international levels.

Figure 5: The internet ecosystem

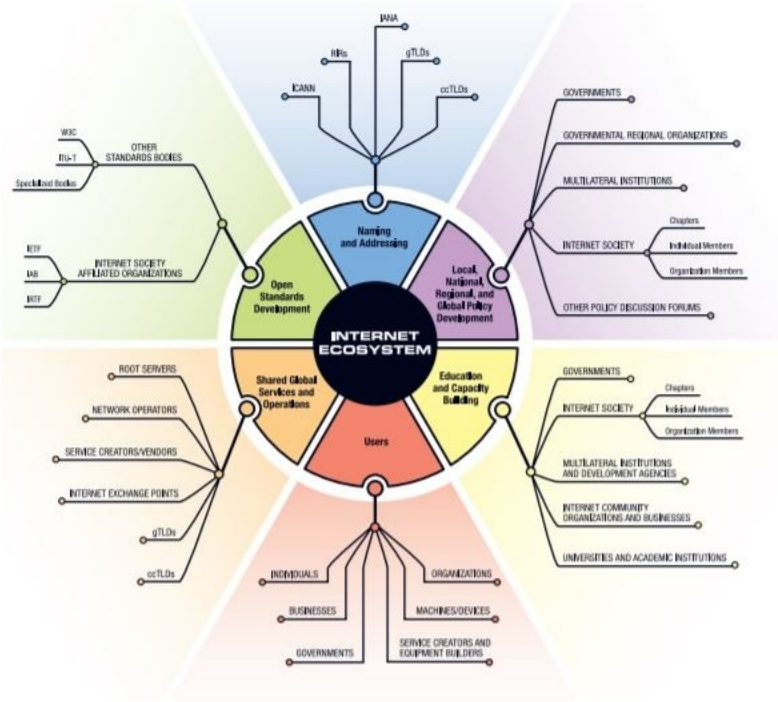

Source: The Internet Society[4]

The distributed nature of the internet means that no single authority has overall responsibility for internet management. Instead the internet is governed in a decentralised, collaborative fashion (TheInternetSociety, 2016).

The technologies, resources and services required to ensure the smooth operation of the internet are all highly interdependent, with each organisation having a specific role. Whilst not an exhaustive list, the Internet Society (TheInternetSociety, 2014), which supports and promotes the development

[4] © The Internet Society 2021. Available at https://www.internetsociety.org/resources/doc/2016/internet-governance-why-the-multistakeholder-approach-works/ (from longer PDF). Under Public License https://creativecommons.org/licenses/by-nc-sa/4.0/legalcode#s6a

of the internet, cites the following organisations as having key roles to play in the technical functioning of internet infrastructure:

- Technical standards bodies including the Internet Engineering Task Force (IETF)[5], the World Wide Web Consortium (W3C[6]) and the Institute of Electrical and Electronic Engineers (IEEE[7]);
- Organisations responsible for global naming and addressing resources including The Internet Corporation for Assigned Names and Numbers (ICANN[8]), which operates the Internet Assigned Numbers Authority (IANA[9]), responsible for the domain names system (DNS); Regional Internet Registries responsible for administering IP address delegation; and domain name registries and registrars; and
- Network infrastructure providers such as DNS providers, network operators, cloud and content delivery network providers and Internet Exchange Points.

The Online Content Journey

When considering how content is transmitted across the internet using TCP/IP, a user's typical interactions take place at the application layer, which can be equated to the perimeter of Figure 1 above. The lower layers in the protocol stack remain largely hidden from the user, however these hidden layers are pivotal in content delivery across the global internet and as such an appreciation of how content is formatted for transmission is fundamental to understanding existing measures implemented to enhance online safety, as will be described in later sections. Figure 6 below shows a high-level overview of the content journey from a user to an online platform.

Figure 6 – Typical content journey from user to online platform

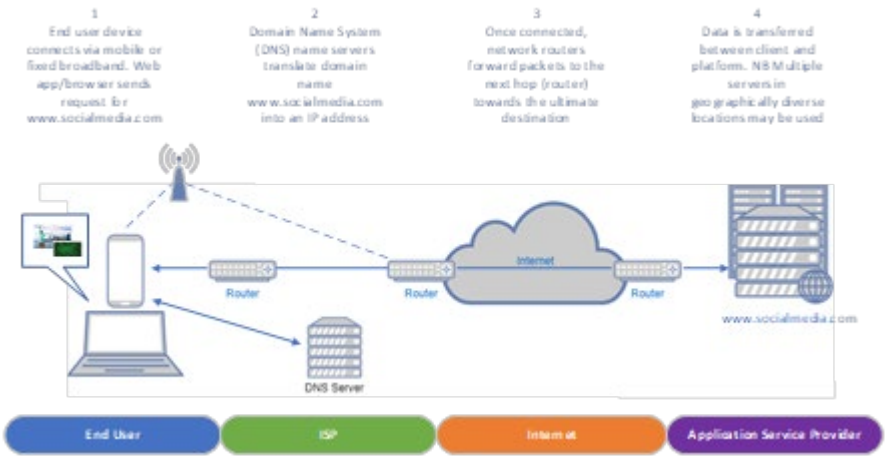

Source: Ofcom

The end user device connects to an access network such as a mobile or fixed broadband network, where in the latter case the end device may connect using wired (Ethernet) or wireless (Wi-Fi) technology.

[5] IETF | Internet Engineering Task Force
[6] World Wide Web Consortium (W3C)
[7] IEEE - The world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
[8] ICANN (icann.org)
[9] Internet Assigned Numbers Authority (iana.org)

Multi-media content is stored on-device in a machine-readable format - i.e. a dataset consisting of 1s and 0s. This digital version of the content is used to display, store and process the original content as well as transfer it across the internet. When the user selects and uploads content, it passes to the application layer, which applies the appropriate protocol to convey the message to the intended recipient. In this example, the data is being passed to a web-based social media platform. An IETF standardised protocol called Hyper Text Transfer Protocol (HTTP) is used to transfer data over the web (W3C, n.d.). When encryption is required, a secure version of HTTP, known as HTTP Secure (HTTPS), is used.

At this point, the data is passed to the transport layer, where the process of packetisation begins. Prior to being sent over the internet, the content is subdivided into smaller blocks and (in the case of our example) a header which contains information such as origin and destination IP addresses is attached to each block. The combination of a data block and its header is called a packet. Each packet is sent independently, typically using one of two protocols TCP or UDP (User Datagram Protocol). TCP is typically used where a reliable connection is required - for example for web browsing or email, whereas UDP is used where a faster connection speed is more important than assuring reliable delivery – for example, voice chat or game streaming.

Figure 7 – Comparison of TCP and UDP

TCP	UDP
Connection oriented	Connectionless
Error recovery, retransmission and acknowledgement	No error recovery, retransmission or acknowledgement
Prioritises reliability over speed	Prioritises speed over reliability
Example uses: Email, HTTP	Example uses: Voice chat, gaming

Source: Ofcom

However, use of QUIC[10], a protocol which combines characteristics of both TCP and UDP to provide secure, reliable, low latency transmission is becoming increasingly common (IETF, 2021b). The QUIC protocol is described in further detail in later sections.

At the internet layer, routing information is added to the packet header to identify the individual hops that each individual packet will take to reach its destination. If the packets are too large for single transmission, they will be broken into smaller fragments which will subsequently be reconstructed by the receiving platform. Finally, the packets pass to the network access layer, which will then initiate the packets' journey across the internet to reach their intended destination.

[10] Originally "Quick UDP Internet Connections", however QUIC is now a name and not an acronym (IETF, 2019b)

Figure 8 – Data Packetisation

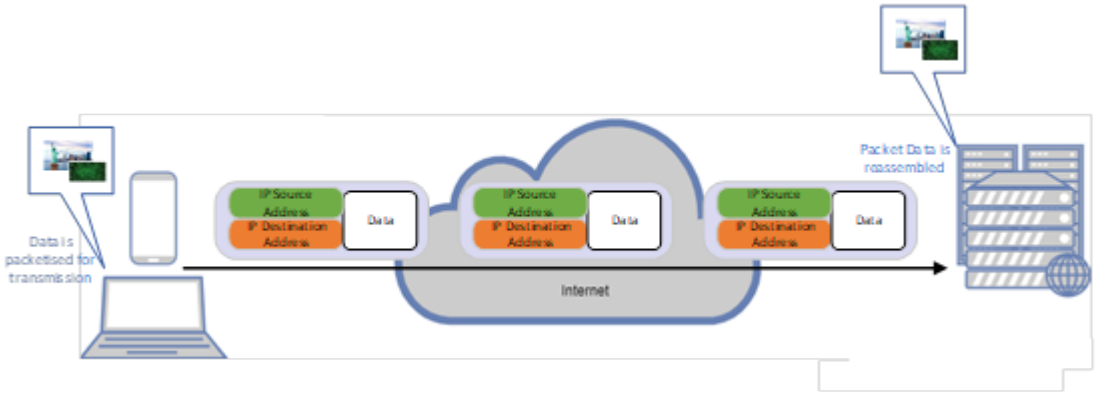

Source: Ofcom

The transmission method depends on how the device is connected to the internet. For example, the packets can be transmitted over the air and to a mobile network if the consumer device happens to be connected to a mobile network. Alternatively, the packets could be sent via Wi-Fi or an Ethernet cable.

Using the control information in the packet header, the receiving platform reconstructs an exact replica of the original contiguous data set which corresponds to the content.

To ensure a low-latency, high availability service many online platforms employ edge computing architectures or use content delivery networks (CDNs) to deliver content to their global users.

In simple terms, both can be visualised as geographically distributed networks consisting of multiple servers or nodes. Content uploaded to a platform is stored, processed and/or replicated across multiple nodes within the network – or, indeed, across multiple networks - and served to users from the network edge which is closest to their location. However, although similar in nature there are key differences between CDNs and edge computing in terms of their functionality. While CDNs simply provide access to stored content, edge computing provides both storage and compute functionality to users.

Platforms may maintain their own CDNs or use one or more third party CDN providers to store content from their platform to facilitate access to their global user base. For example, Netflix, which is based in the United States, deploys content globally using its own CDN, which also includes nodes located physically in the networks of local ISPs (Netflix, 2016).

Figure 9 – Schematic overview of a Content Delivery Network

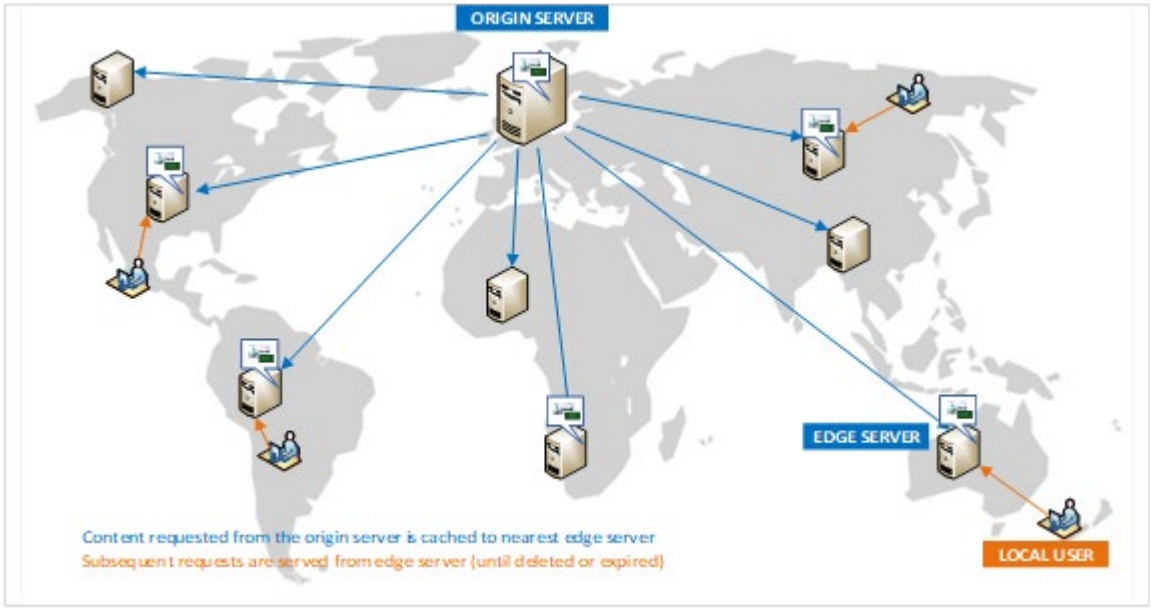

Source: Ofcom

The use of CDNs presents challenges for preventing access to harmful content, as will be described in later sections.

4. Building user safety into the online content journey

The restriction of access to harmful online content is commonly based on the use of blocklists or real-time heuristics[11], for example the use of machine-learning or Artificial Intelligence (AI) based systems which are increasingly deployed by online platforms to assist with content moderation within their services (CambridgeConsultants, 2019). Blocklists typically consist of keywords, domain names, IP addresses or increasingly in recent years cryptographic or perceptual "hashes". A hash can be visualised as a "digital fingerprint" which identifies a specific piece of content - for example the National Center for Missing and Exploited Children (NCMEC), a US non-governmental organisation, provides a list of hashes corresponding to child sexual abuse material which can be used to detect instances of this content within online services (Thorn, 2012).

Network design provides a number of different intervention points for enhancement of user safety. The requirement to have interception mechanisms located somewhere in the network, whether logically or physically, implicates various "points-of-control" within the content journey (IETF, 2018). These control points include: user devices; local networks (e.g. a home or corporate network); the internet backbone; ISPs; content delivery networks (CDNs); and online platforms and applications service providers (collectively referred to here as Application Service Providers (ASPs)). The level of granularity available is dependent on the point within the journey at which mitigations are applied. Figure 10 provides an overview of intervention points within the content journey.

Figure 10 – Potential intervention points for user safety in the content journey

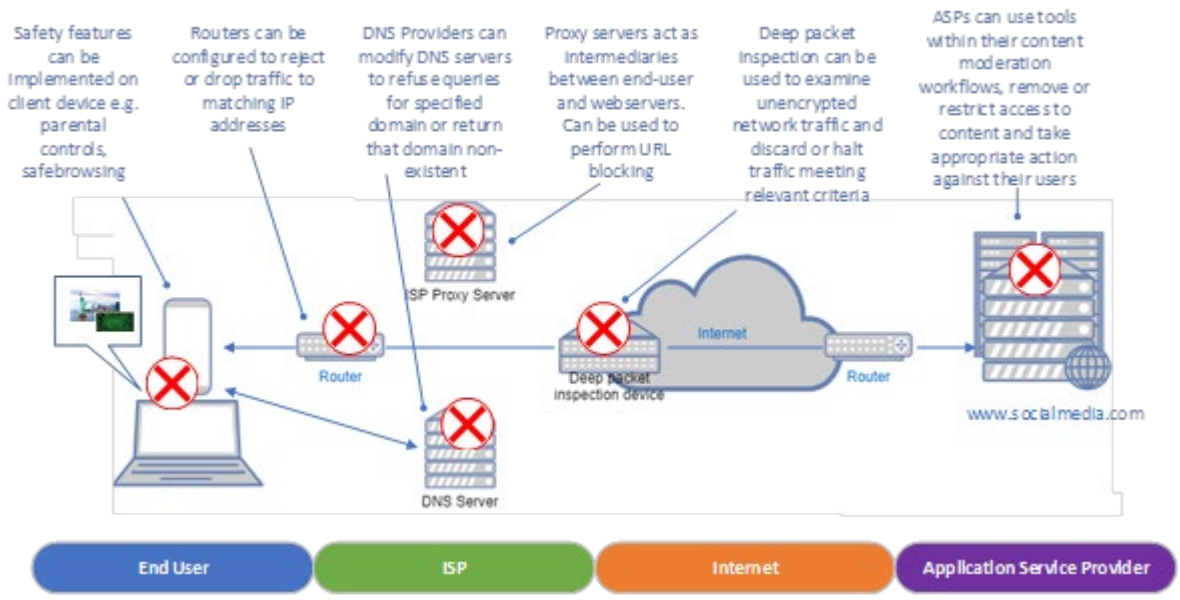

Source: Ofcom

[11] Heuristic programming underpins the field of machine learning/artificial intelligence and involves solving problems using experience-based rules or protocols.

End users

The user's online journey, irrespective of the platform, begins and ends with the device, be it a mobile phone, tablet, laptop or desktop. Numerous approaches to on-device safety are already in existence. For example, for the mobile phone, which has become a dominant platform for access to the web (Ofcom, 2021d), the two largest operating system providers are Apple (through iOS) and Google (through Android). Both Apple and Google provide and promote on-device safety features such as parental controls, which extend through account management to other devices and applications and provide enhanced protection to children as they navigate the online environment. In terms of protecting users of all ages, features such as Google's "Safe Browsing" (Google, n.d.), which warns when an attempt is made to browse to a website identified as containing harmful content such as malware or phishing, are implemented as a local database within the browser on each device. Where properly implemented, such features may have significant benefits in enhancing online safety, for example by ensuring children are unable to access websites containing age-inappropriate content (McAfee, n.d.) (Google, 2021).

As the internet moves towards a more encrypted architecture, the user device may become increasingly important as a point of implementation for safety technology, which can be applied prior to content being encrypted and before it is sent and/or viewed. A recent example of such an approach includes Apple's plan to implement a system for identification of child sexual abuse material being uploaded to iCloud partially through use of a locally stored hash database (Apple, 2021).

Fixed and mobile internet service providers (ISPs)

There are four primary methods which may be used either in isolation or in combination by internet service providers and mobile operators to block user access to harmful content:

- **IP blocking:** network equipment is configured such that internet traffic destined for a specific IP address/website is discarded.
- **DNS blocking/alteration:** the DNS service that translates IP addresses into domain names is modified such that the requesting device is advised either that the website does not exist, or redirects to an alternative website, for example one which explains that access to the site has been blocked.
- **Uniform Resource Locator (URL) blocking:** A URL identifies a specific resource, such as a webpage or image. Using this technique, it is possible to block access to individual instances of harmful content without impeding access to the website as a whole. For example, the majority of UK ISPs already block URLs supplied by the Internet Watch Foundation that display child sexual abuse imagery.
- **Packet inspection:** blocking techniques based on analysis of network traffic either at high level (Shallow Packet Inspection (SPI)) or at a detailed level (Deep Packet Inspection (DPI)), to identify the presence of harmful content.

A report produced by Ofcom in 2011 (Ofcom, 2011) examined blocking in the context of preventing access to copyright infringing content and provided a more detailed analysis of the implementation and overall effectiveness of each of these techniques, however, a key takeaway is that at all levels of

the network hierarchy the intervention points are essentially the same – packets are analysed to identify harmful content and a blocking or shaping mechanism is then implemented to prevent or impede access.

While blocking can be circumvented by site operators and end users and is therefore an imperfect solution for preventing access to harmful online content, it nevertheless introduces barriers and prevents inadvertent user exposure. For example, in 2020, the Internet Watch Foundation reported that 8.8m attempts by UK users to access URLs containing child sexual abuse material had been prevented by use of their blocklist (InternetWatchFoundation, 2021).

However, existing techniques employed by ISPs and mobile operators to block or filter access to harmful content rely on the ability to examine traffic within their networks. As previously noted, this has largely only been possible because the internet was not initially designed with built-in privacy and security (InterconnectCommunications, 2011). In recent years, a number of new privacy-enhancing protocols and initiatives have been proposed and /or implemented which have implications for existing measures used by ISPs and mobile operators to block or impede access to harmful content. Several examples are explained in the following sections.

Transport layer security 1.3 (TLS 1.3)

Transport Layer Security (TLS) is a cryptographic protocol which provides end-to-end security of data sent between internet applications by encrypting the communications link so that only the sender and receiver can see what is being transferred (IETF, 2018b). TLS is used to secure a wide number of applications including email, instant messaging, audio/video conferencing and web browsing – for example, HTTPS is TLS encryption implemented on top of the HTTP protocol.

To set up a TLS connection between server and client, communicating nodes establish a shared secret key through a "handshake", and once connected they can securely exchange data. The first stage of this handshake, called the "Client Hello", typically contains the name of the destination website in a field called the "Server Name Indication" (SNI) (IETF, 2003), which is sent unencrypted, enabling requests for harmful websites or services to be identified and blocked/rerouted. However, the latest version – TLS 1.3 – supports an extension called Encrypted Client Hello (ECH), whereby the entire handshake is encrypted using a key acquired from the DNS (IETF, 2021). When used in conjunction with DNS over HTTPS (DoH) or DNS over TLS (DoT) (described below), it is not possible for an ISP to identify the destination website.

DNS over HTTPS (DoH) and DNS over TLS (DoT)

The DNS is a network protocol which is over 35 years old, and which was initially designed without privacy or integrity in mind. DNS queries were therefore sent unencrypted, meaning that it was possible for operators of DNS servers (commonly the local ISP) to identify which websites a user was requesting.

Both DoH and DoT are methods of encrypting DNS queries, and operate in broadly the same way, taking the unencrypted DNS protocol and placing it into a secure encrypted channel between the end-user and the encrypted DNS service provider (figure 11). The main difference between DoH and

DoT is that the latter typically uses the dedicated network port[12] 853, whereas DoH uses port 443, the same port as the HTTPS protocol, which is typically used to secure users session on websites. As such, whilst DoT traffic is more readily identifiable by network operators, DoH traffic is indistinct from other web traffic and therefore effectively impervious to modification or surveillance.

Figure 11: DNS request using DoH

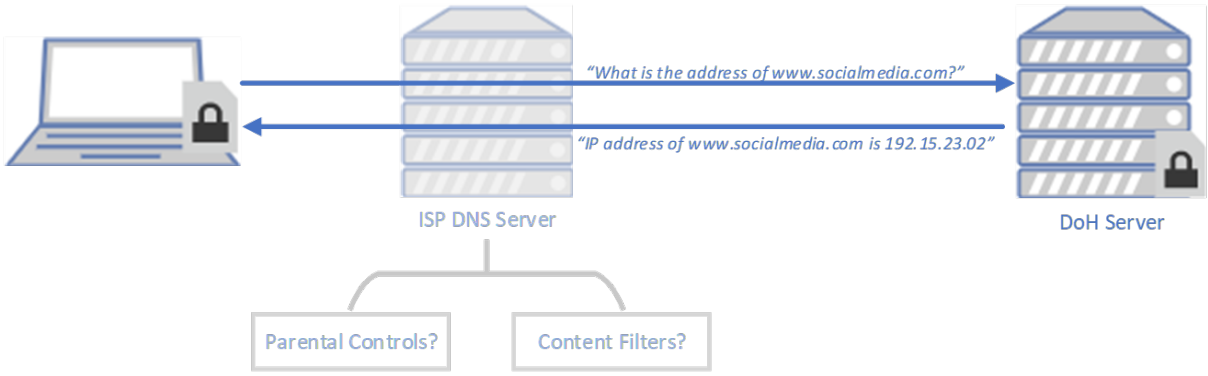

Source: Ofcom

Governments , via ISP and mobile operator DNS services, have come to rely on the ability to modify DNS responses to implement court mandated based blocking (for example, in relation to copyright infringing websites (BBC, 2012)) and blocking on a voluntary basis (for example for child protection purposes or to prevent access to criminal content). In bypassing local DNS servers in favour of those on a prescribed list of "trusted" DNS resolvers, both DoT and DoH may have the effect of preventing ISPs and mobile operators from implementing existing family friendly filters or provision of network performance metrics; it may even impact on their ability to carry out common user support tasks such as setting up devices or routers, or providing mobile top-ups.

While there are clearly legitimate privacy benefits to the implementation of DoH, careful consideration will need to be given to any unintended consequences for network security as well as for user safety. While Google, Mozilla and Microsoft have implemented DoH support within their products, it is still a work in progress. An IETF working group was set up in 2019 to examine and develop a consensus on how the protocol should be implemented (IETF, 2019). Additionally, an industry working group, the Encrypted DNS Deployment Initiative, has been formed with the objective of ensuring encrypted DNS technologies are adopted in a manner which ensures "the continued unimpaired functionality of security protections, parental controls, and other services that depend upon the DNS" (EncryptedDNS, 2019).

Oblivious HTTP/Oblivious DoH

The introduction of protocols and standards such as DoH and DoT has given rise to concerns that, until there is wider deployment amongst ISPs, the centralisation of the DNS could introduce single points of failure into the network. Additionally, there is concern that rather than increasing user privacy, such centralisation could have the opposite effect, channeling all user queries through a

[12] A port is a virtual point where network connections start and end. Ports are standardised across all network-connected devices with each being assigned a number. Most ports are reserved for certain protocols, which relate to specific applications or services – for example, HTTP messages go to port 80.

small number of DNS resolvers who could then easily correlate the queries to IP addresses for profiling purposes (Cloudflare, 2020). To respond to these concerns, a protocol called Oblivious DoH (ODoH) is under development at IETF, which provides for separation of the requesting IP address from the DNS query (Kinnear, et al., 2021). ODoH works by adding a layer of encryption to the request and introduces a network proxy between clients and DoH servers. This combination means that only the user has access to both the DNS messages and their own IP address.

Similarly, when a user makes a request for a web resource using HTTP, information about the client's identity is revealed to the server, enabling the destination platform to correlate requests over time to assemble profiles of user behaviour. Whilst users can to an extent currently avoid this by using third party proxies such as Tor or Virtual Private Networks (VPNs) to decouple their IP address from their requests, this can introduce latency, and in the case of one-hop proxies the provider still has sight of user traffic. The IETF is therefore exploring the development of a new protocol, Oblivious HTTP (OHTTP) (Thomson & Wood, 2021) which delinks the client IP from the HTTP request.

An example of a proposed implementation of "oblivious" technologies is Apple's iCloud Private Relay (Apple, 2021b), announced in June 2021, which it is claimed will enable users to browse the web in such a way "that no one in this chain - not even Apple - can see both the client IP address and what the user is accessing." (Pauly, 2021). Figure 12 below shows how the process works.

Figure 12 – iCloud Private Relay

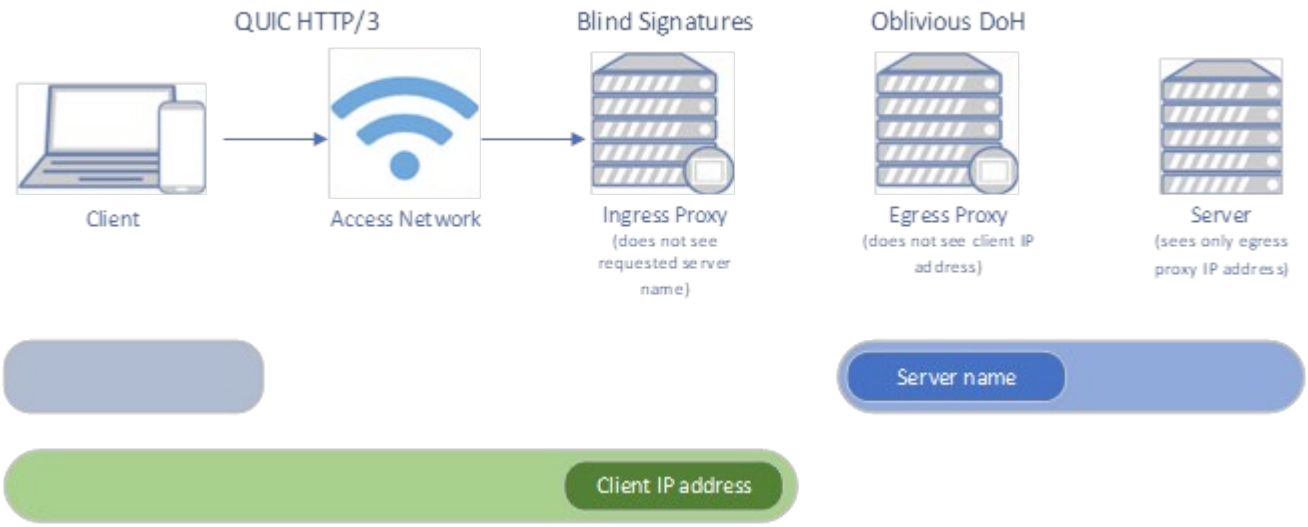

Source: Ofcom

The use of such "oblivious" protocols has similar implications to DoH in terms of the impact on network based filtering, but may also impact more widely on user safety – for example, at this stage the level of interoperability with third party and on-device parental controls is unclear.

QUIC and HTTP/3

QUIC is a new transport layer protocol developed on top of UDP. As briefly outlined above, QUIC combines the reliability of TCP (ordered and error checked delivery of packets) with the low latency of UDP (connection-less transmission – sometimes referred to as "fire and forget" (Postown, 2020)). It does this by adding a "connection ID" to each packet header. The connection ID is used at the receiving end to track and reorder the received packets. Using QUIC, connections are maintained

even where the network changes – for example if moving between mobile networks – as there is no requirement to renegotiate a secure connection. A typical handshake using TCP and TLS 1.3 takes two round trips to complete. QUIC reduces the round-trip time (RTT) required when establishing a connection by integrating the security handshake with the connection set up, reducing latency and fully encrypting the handshake by default.

HTTP/3, which uses QUIC as a transport layer, is the latest version of HTTP (Duke, et al., 2021), the application layer protocol which is used to transport data over the web. It will benefit from QUIC's zero round-trip times (0-RTT) and default encryption. While at time of writing the HTTP/3 protocol is still officially in draft form, it is reportedly already supported by 73% of running web browsers (CanIUse.com, 2021).

Online platforms/application service providers (ASPs)

The ever-increasing volumes of content requiring moderation mean it is now common practice for applications service providers (ASPs) to use automated content classification systems (ACCs) as part of the moderation workflow within their online platforms. Such ACCs aim to identify the presence of harmful or prohibited content at a highly granular level.

ACCs take many forms, ranging from highly advanced machine learning solutions (sometimes referred to as "artificial intelligence" or "AI") through to less sophisticated filtering systems which may remove or flag content for additional review based on the presence of specific keywords or character strings, or hash based searches for known unacceptable imagery. Given the highly nuanced nature of online content, many workflows are not fully automated, but include a "human" layer.

Moderation may be proactive, in that it seeks to prevent harmful material appearing on the platform at all, or reactive, for example, in response to escalations made by platform users (CambridgeConsultants, 2019). The mechanisms used to detect harmful content will to an extent be dictated by the nature of the content itself but also the delivery medium – for example, specific challenges exist in relation to content which is live-streamed, as only the immediate and preceding segments are available for review and harm may quickly escalate.

Changes to the underlying internet infrastructure as outlined in previous sections suggest that ASPs may play an increasingly important role in ensuring user safety within the online content journey in the future, given their ability to respond to identified instances of harmful content in a highly granular and targeted manner.

As such, Ofcom is currently engaging in a wider programme of research examining in more detail specific tools and implementations available to ASPs to assist in content moderation within their platforms. This will include increasing our understanding on the potential of identifying known harmful content within end-to-end encrypted environments including by metadata analysis (Kamara, et al., 2021) and the use of homomorphic encryption (HE). Like other forms of encryption, HE uses a public key to encrypt data. Unlike other forms of encryption, it uses an algebraic system which allows functions to be performed on the encrypted data. The widespread application of homomorphic encryption is currently limited due to the high computational overhead involved. However, it is currently reported that research is being carried out by platforms to explore the potential use of forms of HE to continue to deliver advertising based on the content of encrypted

communications without having visibility of the message content (Holt, 2021). HE may also have scope to be used to identify the presence of illegal content such as child sexual abuse material (CSAM) (Singh & Farid, 2019).

Content delivery networks (CDNs)

The use of CDNs presents challenges for preventing access to harmful content as multiple copies are distributed across the network. As such, it is necessary for ASPs to ensure that the content is removed not only at a single source, but also that any remaining copies are purged from all stored locations.

CDNs also present challenges for ISPs in blocking access to harmful content, as in many cases webpages being served by CDNs will share the same IP addresses. This means that should IP blocking be used to prevent access, there is a risk that access to legitimate web content sharing the same IP address will also be blocked. While theoretically other forms of blocking such as URL filtering could be implemented, CDNs typically use encryption (HTTPS) to secure connections which renders this method infeasible (Zolfaghari & Houmansadr, 2016).

5. Conclusion

The online content journey provides several intervention points and opportunities for enhancing user safety at all stages of transmission across the internet.

However, a number of existing methods for detecting and managing exposure to potentially harmful content rely on the fact that the internet was not designed to be inherently secure. Over time, initiatives and protocols have been implemented to address these shortcomings and better protect the privacy and security of users. The current direction of travel towards a more privacy-respecting online experience, while providing enhanced protections in some areas, also creates challenges for the way potentially harmful content has historically been identified and managed. This includes the increasing use of encryption to protect data and communications, such as DoH/DoT and QUIC; and the development of technologies and protocols designed to protect user anonymity – for example Oblivious HTTP and Oblivious DoH.

As such, an understanding of the underlying internet infrastructure, not just the platforms and applications with which users directly interact, is necessary to fully manage approaches to online safety.

Additionally, while platforms and application providers are innovating and implementing technical measures to ensure the safety of their users, this gives rise to challenges relating to transparency and consistent standards and application which need to be recognised and addressed.

Ofcom is building expertise in these areas to track technical developments of relevance to online safety and provide support for evolving regulation to keep the internet safer for the people and businesses of the UK.

A1. References

Apple, 2021. Expanded Protections for Children. [Online] Available at: https://www.apple.com/child-safety/

Apple, 2021b. Prepare Your Network or Web Server for iCloud Private Relay. [Online] Available at: https://developer.apple.com/support/prepare-your-network-for-icloud-private-relay#:~:text=Prepare%20Your%20Network%20for%20iCloud%20Private%20Relay%201,for%20network%20traffic%20audits.%20...%206%20Additional%20resources

BBC, 2012. The Pirate Bay must be blocked by UK ISPs, court rules. [Online] Available at: https://www.bbc.co.uk/news/technology-17894176

CambridgeConsultants, 2019. USE OF AI IN ONLINE CONTENT MODERATION. [Online] Available at: https://www.ofcom.org.uk/__data/assets/pdf_file/0028/157249/cambridge-consultants-ai-content-moderation.pdf

CanIUse.com, 2021. HTTP/3 protocol. [Online] Available at: https://caniuse.com/http3

Cisco, 2011. The TCP/IP and OSI Networking Models. [Online] Available at: https://www.ciscopress.com/articles/article.asp?p=1757634&seqNum=2

Cloudflare, 2020. Improving DNS Privacy with Oblivious DoH in 1.1.1.1. [Online] Available at: https://blog.cloudflare.com/oblivious-dns/

Cloudflare, n.d. What is BGP? | BGP routing explained. [Online] Available at: https://www.cloudflare.com/de-de/learning/security/glossary/what-is-bgp/

Duke, M., Sarker, Z. & Westerlund, M., 2021. A new era in Internet transport. [Online] Available at: https://www.ietf.org/blog/new-era-transport/

EncryptedDNS, 2019. About the Encrypted DNS Deployment Initiative. [Online] Available at: https://www.encrypted-dns.org/about

Google, 2021. Google Transparency Report. [Online] Available at: https://transparencyreport.google.com/safe-browsing/overview?hl=en_GB

Google, n.d. Keeping over four billion devices safer. [Online] Available at: https://safebrowsing.google.com/

Holt, K., 2021. Facebook is reportedly trying to analyze encrypted data without deciphering it. [Online]Available at: https://www.engadget.com/facebook-analyze-encrypted-messages-ad-targeting-175739715.html

ICANN, n.d. DNSSEC – What Is It and Why Is It Important?. [Online] Available at: https://www.icann.org/resources/pages/dnssec-what-is-it-why-important-2019-03-05-en#:~:text=DNSSEC%20protects%20the%20user%20from%20getting%20bad%20data,resolvers%20can%20be%20assured%20of%20getting%20good%20data.

IETF, 2003. Transport Layer Security (TLS) Extensions. [Online] Available at: https://datatracker.ietf.org/doc/html/rfc3546

IETF, 2005. IP Security Protocol (ipsec). [Online] Available at: https://datatracker.ietf.org/wg/ipsec/charter/

IETF, 2018. A Survey of Worldwide Censorship Techniques. [Online] Available at: https://tools.ietf.org/id/draft-hall-censorship-tech-05.html

IETF, 2018b. The Transport Layer Security (TLS) Protocol Version 1.3. [Online] Available at: https://datatracker.ietf.org/doc/html/rfc8446

IETF, 2019. Adaptive DNS Discovery (add). [Online] Available at: https://datatracker.ietf.org/wg/add/about/

IETF, 2019b. QUIC: A UDP-Based Multiplexed and Secure Transport. [Online] Available at: https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-19#section-1.2

IETF, 2021. TLS Encrypted Client Hello [Online] Available at: https://datatracker.ietf.org/doc/draft-ietf-tls-esni/

IETF, 2021b. QUIC in the Internet industry. [Online] Available at: https://www.ietf.org/blog/quic-industry/

InterconnectCommunications, 2011. MC / 080:DNSSEC Deployment Study. [Online] Available at: https://www.ofcom.org.uk/__data/assets/pdf_file/0015/19131/domain-name-security.pdf

InternetWatchFoundation, 2021. Millions of attempts to access child sexual abuse online during lockdown. [Online] Available at: https://www.iwf.org.uk/news/millions-of-attempts-to-access-child-sexual-abuse-online-during-lockdown

Kamara, S. et al., 2021. Outside Looking In: Approaches to Content Moderation in End-to-End Encrypted Systems. [Online] Available at: https://cdt.org/wp-content/uploads/2021/08/CDT-Outside-Looking-In-Approaches-to-Content-Moderation-in-End-to-End-Encrypted-Systems.pdf

Kinnear, E. et al., 2021. Oblivious DNS Over HTTPS. [Online] Available at: https://www.ietf.org/id/draft-pauly-dprive-oblivious-doh-07.html

MAPPINGProject, 2014. Snowden revelations make IETF rethink security. [Online] Available at: https://observatory.mappingtheinternet.eu/item/snowden-revelations-make-ietf-rethink-security

McAfee, n.d. Parental Controls. [Online] Available at: https://www.mcafee.com/en-gb/parental-controls.html?id=eula&culture=en-gb&pir=1

Netflix, 2016. How Netflix Works With ISPs Around the Globe to Deliver a Great Viewing Experience. [Online] Available at: https://about.netflix.com/en/news/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience

Ofcom, 2011. Site blocking to reduce online copyright infringement. [Online] Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/78095/Ofcom_Site-Blocking-_report_with_redactions_vs2.pdf

Ofcom, 2021. Ofcom to regulate harmful content online. [Online] Available at: https://www.ofcom.org.uk/about-ofcom/latest/features-and-news/ofcom-to-regulate-harmful-content-online

Ofcom, 2021a. Digital Regulation Cooperation Forum. [Online] Available at: https://www.ofcom.org.uk/__data/assets/pdf_file/0021/192243/drcf-launch-document.pdf#:~:text=The%20DRCF%20is%20a%20non-statutory%20body.%20It%20is,and%20%E2%80%A2%20%20The%20Information%20Commissioner%E2%80%99s%20Office%20%28ICO%29.

Ofcom, 2021b. Ofcom Statement of Work 2021/2022. [Online] Available at: https://www.ofcom.org.uk/__data/assets/pdf_file/0019/216640/statement-plan-of-work-202122.pdf

Ofcom, 2021c. Internet Futures: Spotlight on the technologies which may shape the Internet of the future. [Online] Available at: https://www.ofcom.org.uk/research-and-data/internet-and-on-demand-research/internet-futures

Ofcom, 2021d. Online Nation 2021 report. [Online] Available at: https://www.ofcom.org.uk/__data/assets/pdf_file/0013/220414/online-nation-2021-report.pdf

Pauly, T., 2021. Apple Developer Videos. [Online] Available at: Prepare Your Network or Web Server for iCloud Private Relay - Support - Apple Developer

Postown, H., 2020. Network traffic analysis for IR: UDP with Wireshark. [Online] Available at: https://resources.infosecinstitute.com/topic/network-traffic-analysis-for-ir-udp-with-wireshark/

Singh, P. & Farid, H., 2019. Robust Homomorphic Image Hashing, s.l.: s.n. Available at: https://farid.berkeley.edu/downloads/publications/cvpr19/cvpr19c.pdf

TheInternetSociety, 2014. Internet Ecosystem: Naming and addressing, shared global services and operations, open standards development. [Online] Available at: https://www.internetsociety.org/internet/who-makes-it-work/

TheInternetSociety, 2016. Internet Governance – Why the Multistakeholder Approach Works. [Online] Available at: https://www.internetsociety.org/resources/doc/2016/internet-governance-why-the-multistakeholder-approach-works/

TheInternetSociety, n.d. Frequently Asked Questions (FAQ) on IPv6 adoption and IPv4 exhaustion [Online] Available at: https://www.internetsociety.org/deploy360/ipv6/faq/

Thomson, M. & Wood, C. A., 2021. Oblivious HTTP. [Online] Available at: https://www.ietf.org/archive/id/draft-thomson-http-oblivious-00.html

Thorn, 2012. Sharing hashes across industry. [Online] Available at: https://www.thorn.org/reporting-child-sexual-abuse-content-shared-hash/

UKGOV, 2020. The Audiovisual Media Services Regulations 2020. [Online] Available at: https://www.legislation.gov.uk/uksi/2020/1062/made

UKGOV, 2021. The Digital Regulation Cooperation Forum [Online] Available at: The Digital Regulation Cooperation Forum - GOV.UK (www.gov.uk)

W3C, n.d. HTTP - Hypertext Transfer Protocol. [Online] Available at: https://www.w3.org/Protocols/

Zolfaghari, H. & Houmansadr, A., 2016. Practical Censorship Evasion Leveraging Content Delivery Networks, s.l.: s.n. Available at: https://people.cs.umass.edu/~amir/papers/CDNReaper.pdf