



CANADIAN CENTRE *for* CHILD PROTECTION®
Helping families. Protecting children.

PROJECT ARACHNID: ONLINE AVAILABILITY OF CHILD SEXUAL ABUSE MATERIAL

*An analysis of CSAM and harmful-abusive content
linked to certain electronic service providers*



Project
Arachnid™

This document is EMBARGOED and is FOR MEDIA USE ONLY. It is current as of June 8, 2021 and contains the results of a review conducted by the Canadian Centre for Child Protection Inc. of its Project Arachnid data. It is being shared to facilitate an advance understanding of the report by media prior to the document being finalized and being publicly released by the Canadian Centre. It is subject to revision, and as such, it is considered a confidential, non-public document by the Canadian Centre.



CANADIAN CENTRE *for* CHILD PROTECTION®

Helping families. Protecting children.

The Canadian Centre for Child Protection would like to thank the international child protection entities working collaboratively within Project Arachnid to scale up its capacity to globally reduce the availability of online child sexual abuse material.

For a full list of those classifying material in Project Arachnid, visit projectarachnid.ca

© June 8, 2021, Canadian Centre for Child Protection Inc. (C3P). All rights reserved. Data relied upon to produce this report is held by C3P, and all analysis was conducted internally by C3P staff. Reasonable efforts have been made to ensure the accuracy of all information herein. The names of the companies reflected in the data are either the names as represented by the host on the site or in terms of service posted on the site, or the name of the URL if no other name could be determined. "CANADIAN CENTRE for CHILD PROTECTION" and "Cybertip.ca" is registered in Canada as a trademark of, and Project Arachnid is used as a trademark of, C3P. All third-party trademarks included within the report are the property of their respective owners, and their inclusion is not meant to imply endorsement or affiliation of any kind. Trademark symbols, if applicable, are not included in any data tables.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	2
ABOUT THE CANADIAN CENTRE FOR CHILD PROTECTION	4
<i>Work with survivor advocacy groups</i>	5
INTRODUCTION	6
PROJECT ARACHNID: A TECHNOLOGY SOLUTION FOR DETECTING CSAM	7
<i>How it works</i>	7
<i>Shield by Project Arachnid</i>	7
ILLEGAL AND HARMFUL-ABUSIVE CONTENT: A C3P FRAMEWORK	9
<i>CSAM and harmful-abusive content</i>	10
<i>Where CSAM and harmful-abusive content is found</i>	11
METHODOLOGY	12
<i>Data collection</i>	12
<i>Image categories</i>	12
<i>Key transparency measures</i>	13
<i>Limitations</i>	14
ANALYSIS AND DISCUSSION	15
<i>Verified media detections</i>	16
<i>Media targeted for removal</i>	20
<i>Removal notices</i>	25
<i>Server locations of media targeted for removal</i>	30
<i>Content removal times</i>	33
<i>Media recidivism</i>	38
<i>Case study: French telecommunications company Free: Project Arachnid's largest source of detected CSAM</i>	42
<i>Case study: Project Arachnid Trichan imageboard campaign</i>	44
RECOMMENDATIONS	47
CONCLUSION	56
APPENDIX	58
<i>Glossary of terms</i>	58
<i>List of acronyms</i>	60

EXECUTIVE SUMMARY

The Canadian Centre for Child Protection (C3P) is issuing this report to highlight how systemic failures of the technology industry and inaction by governments have severely hindered the fight against the proliferation of child sexual abuse material (CSAM) on the internet.

The purpose of this report is to arm governments with key information required to make decisions most likely to be effective in reducing the online availability and distribution of CSAM. The analysis points to a need for consistent enforceable standards that impose accountability requirements on electronic service providers (ESP).

As a survivor-centric organization, C3P has invested resources to build a specialized tool called Project Arachnid that detects CSAM by crawling the open web and issues removal notices to those we believe to have the most immediate control or custody of the media.

Project Arachnid's reach does not generally extend into commonly known social media platforms due to their semi-closed designs. The analysis presented in this report is based on a sub-section of the open web and is therefore a gross underestimation of the true extent of CSAM availability on the internet. As a result, this report highlights the vast networks of lesser-known ESPs that contribute to the problem, and not large mainstream technology companies.

The following are the key findings from the analysis spanning the period of 2018 to 2020:



Project Arachnid has detected and verified more than **5.4 million images** and has issued removal notices to more than **760 ESPs** worldwide.



Nearly half (**48%**) of all media detections are linked to a file hosting service operated by one French telecommunications company.



The vast majority (**97%**) of CSAM detected by Project Arachnid is physically hosted on the clear web. However, the dark web plays a disproportionately large role in directing individuals on where to access CSAM on the clear web.



Project Arachnid is an effective tool that has achieved a median content removal time of **less than 24 hours**. Alarmingly, however, 10 percent of actioned media took more than seven weeks (42 days) before becoming inaccessible.



Overall, images depicting older adolescents (post-pubescent) take **much longer** to be removed compared to images with younger victims (pre-pubescent).



Nearly half (**48%**) of all media Project Arachnid has issued a removal notice on, had previously been flagged to the service provider.



As of the writing of this report, C3P is facing a backlog of more than **32.8 million suspect media**¹ that have yet to be assessed. The rate at which Project Arachnid detects suspect media far outpaces the human resources available to assess the content.

¹ Suspect media is derived only from websites that host known CSAM, and the term refers to any media that is reasonably suspected to be CSAM or harmful-abusive content but which has not been through the assessment process.

These findings, notably the high levels of image recidivism and the often long delays in removal times, suggest many ESPs are not deploying sufficient resources to eliminate, or at least limit, the presence of CSAM and harmful-abusive content.

Even the seemingly more positive results belie one of the core problems: While it is true many ESPs remove media within a day of notification, in the absence of any regulatory requirements, they have no commercial or legal interest in investing in measures to prevent the images from surfacing or re-surfacing in the first place. There are no consequences for inaction on the prevention side. This is laid bare by the correspondingly high image recidivism rates described in this report.

Many ESPs benefit from business models and practices that are currently backstopped by broad immunity protections in the U.S. They also benefit from a general air of uncertainty over jurisdictional issues, as well as deficient regulation in the digital space across the globe.

Given this backdrop, the following set of recommendations may assist governments in developing effective and consistent regulatory frameworks to address the issue:

	1. Enact and impose a duty of care , along with financial penalties for non-compliance or failure to fulfill a required duty of care;
	2. Impose certain legal duties on upstream electronic service providers and their downstream customers;
	3. Require automated, proactive content detection for platforms with user-generated content;
	4. Set standards for content that may not be criminal, but remains harmful-abusive to minors
	5. Mandate human content moderation standards;
	6. Set requirements for proof of subject or participant consent and uploader verification;
	7. Establish platform design standards that reduce risk and promote safety;
	8. Establish standards for user-reporting mechanisms and content removal obligation.

There is a growing public consensus that a largely unfettered digital space void of any meaningful consequences for causing great harm to children is an issue that must be urgently addressed. This report is a road map for governments to develop policy and act in concert on the global and borderless fight against the exploitation of children.

ABOUT THE CANADIAN CENTRE FOR CHILD PROTECTION

The Canadian Centre for Child Protection Inc. (C3P) is a national charity dedicated to the personal safety of all children. C3P operates Cybertip.ca, Canada's tipline to report child sexual abuse and exploitation online, as well as provide other intervention, prevention and education services.

In January 2017, C3P established Project Arachnid — a web platform designed to detect known images of child sexual abuse material (CSAM) and issue removal notices to electronic service providers (ESPs) where possible.

C3P also supports survivors whose child sexual abuse was recorded and distributed online. Through our work with survivors, crucial contextual information about the nature of child sexual abuse is collected and shared with stakeholders committed to the safety and protection of children.



Work with survivor advocacy groups

In addition to our work with individual survivors, we work with several survivor advocacy groups:

The Phoenix 11

For over three years, C3P and the National Center for Missing and Exploited Children (NCMEC) have been working with the Phoenix 11, a group of survivors whose child sexual abuse was recorded, and in the majority of cases, distributed online. This group has banded together as a powerful force to challenge the inadequate response to the prevalence of CSAM.

The Chicago Males

C3P and NCMEC started working with a group of male survivors to learn about their experiences and better understand the unique social stigma males face around sexual abuse. This group is working together to advocate for much-needed change in addressing online child sexual abuse and supporting survivors.

The Aramid Collective

In 2020, C3P was introduced to a group of survivors who have been self-monitoring their own CSAM online and reporting to companies to get it removed. This group is using their knowledge and collective voice to help advocate for survivors and the urgent need to address the images and videos of sexual abuse that exist on many platforms.

Mothers of Child Sexual Abuse Material Survivors

To learn about the hardships families of survivors endure for years after the hands-on abuse has ended, C3P brought together a group of mothers whose children's sexual abuse was recorded and distributed online. We learned from moms there is an emotional continuum long after "the discovery" of the abuse that often includes loss of relationships, financial instability, and a constant worry about their child(ren)'s safety, to name only a few examples. Their insight is crucial to guiding the creation of support resources.

"For the longest time there was nothing to be done. It's on the internet, a lawless black hole of code. Now there is something to combat CSAM, and that solution is Project Arachnid." - *A member of the Phoenix 11*

INTRODUCTION

The purpose of this report is to provide all stakeholders — including governments and ESPs — with key information required to take effective action against the distribution and accessibility of CSAM on the internet.

CSAM perpetuates a cycle of victimization for children by stripping them of their personal safety and right to privacy, while inflicting great and lasting harm. Reducing the availability of this material must be a key pillar of child protection frameworks in our efforts to keep citizens safe.

We know that one of the keys to solving this problem is a deep understanding of the role that internet-based companies — especially those that accept user-generated content — play in facilitating access to and dissemination of abusive and illegal media.

The criminal nature of much of the material itself presents barriers from a research, public awareness and policy development perspective. Over time, these barriers have resulted in a limited understanding of the nature of the material, how it flourishes online and the manner in which it is distributed or accessed.

Primary sources of data are mostly held by ESPs who are privately run, and do not tend to proactively release meaningful information about the distribution, moderation and removal of the content hosted on their platforms. In jurisdictions where mandatory reporting requirements exist, figures reported by private companies are not independently verified and details about the reports are themselves limited.

This lack of transparency has prevented a true understanding of the scale of the threat, and impeded the development of legislative and regulatory responses, as well as remedies for victims and survivors.

Under these circumstances, developing sound evidence-based policies or regulation poses a real challenge. This report fills in some of the gaps using company-specific data on the accessibility of CSAM and harmful-abusive² material linked to certain platforms, all of which is independently collected by Project Arachnid. This report also offers a road map for governments seeking accountability on behalf of children through the responsible regulation of ESPs.

² For a full definition of harmful-abusive content, see page 10.



PROJECT ARACHNID: A TECHNOLOGY SOLUTION FOR DETECTING CSAM

Operated by C3P, Project Arachnid is an innovative tool to combat the growing proliferation of CSAM on the internet.

Launched in 2017, this victim-centred tool crawls the open web³ in search of images of CSAM. When CSAM or harmful-abusive content is detected, a removal request is sent to the ESP most likely to have the most immediate control or custody of the media. This automated process is triggered thousands of times per day.



How it works

Project Arachnid discovers CSAM by crawling specific publicly accessible URLs reported to Cybertip.ca, and also content located at URLs on the clear and dark web known to host this type of material. When media — which includes image, video and/or archive files — displayed at a URL are detected, the system compares its digital fingerprint against a database of fingerprints from previously verified media. If the system detects a match between digital fingerprints, a takedown notice is automatically sent to the content administrator or the hosting provider requesting its removal.

Once a notice is sent, Project Arachnid re-crawls the offending URL every 24 hours, triggering subsequent removal notices until the content is no longer detected. Processing tens of thousands of images per second, Project Arachnid detects content at a pace that far exceeds traditional methods of identifying and addressing this harmful material.

Digital fingerprint values contained in the repository of previously verified media originate primarily from image and video assessments from C3P analysts, teams of analysts working for other child protection tiplines and also from Canadian and international law enforcement.

Shield by Project Arachnid

In addition to actively seeking out harmful material on the clear web, Project Arachnid's platform also provides industry with a no-cost tool to assist with proactive detection of known CSAM.



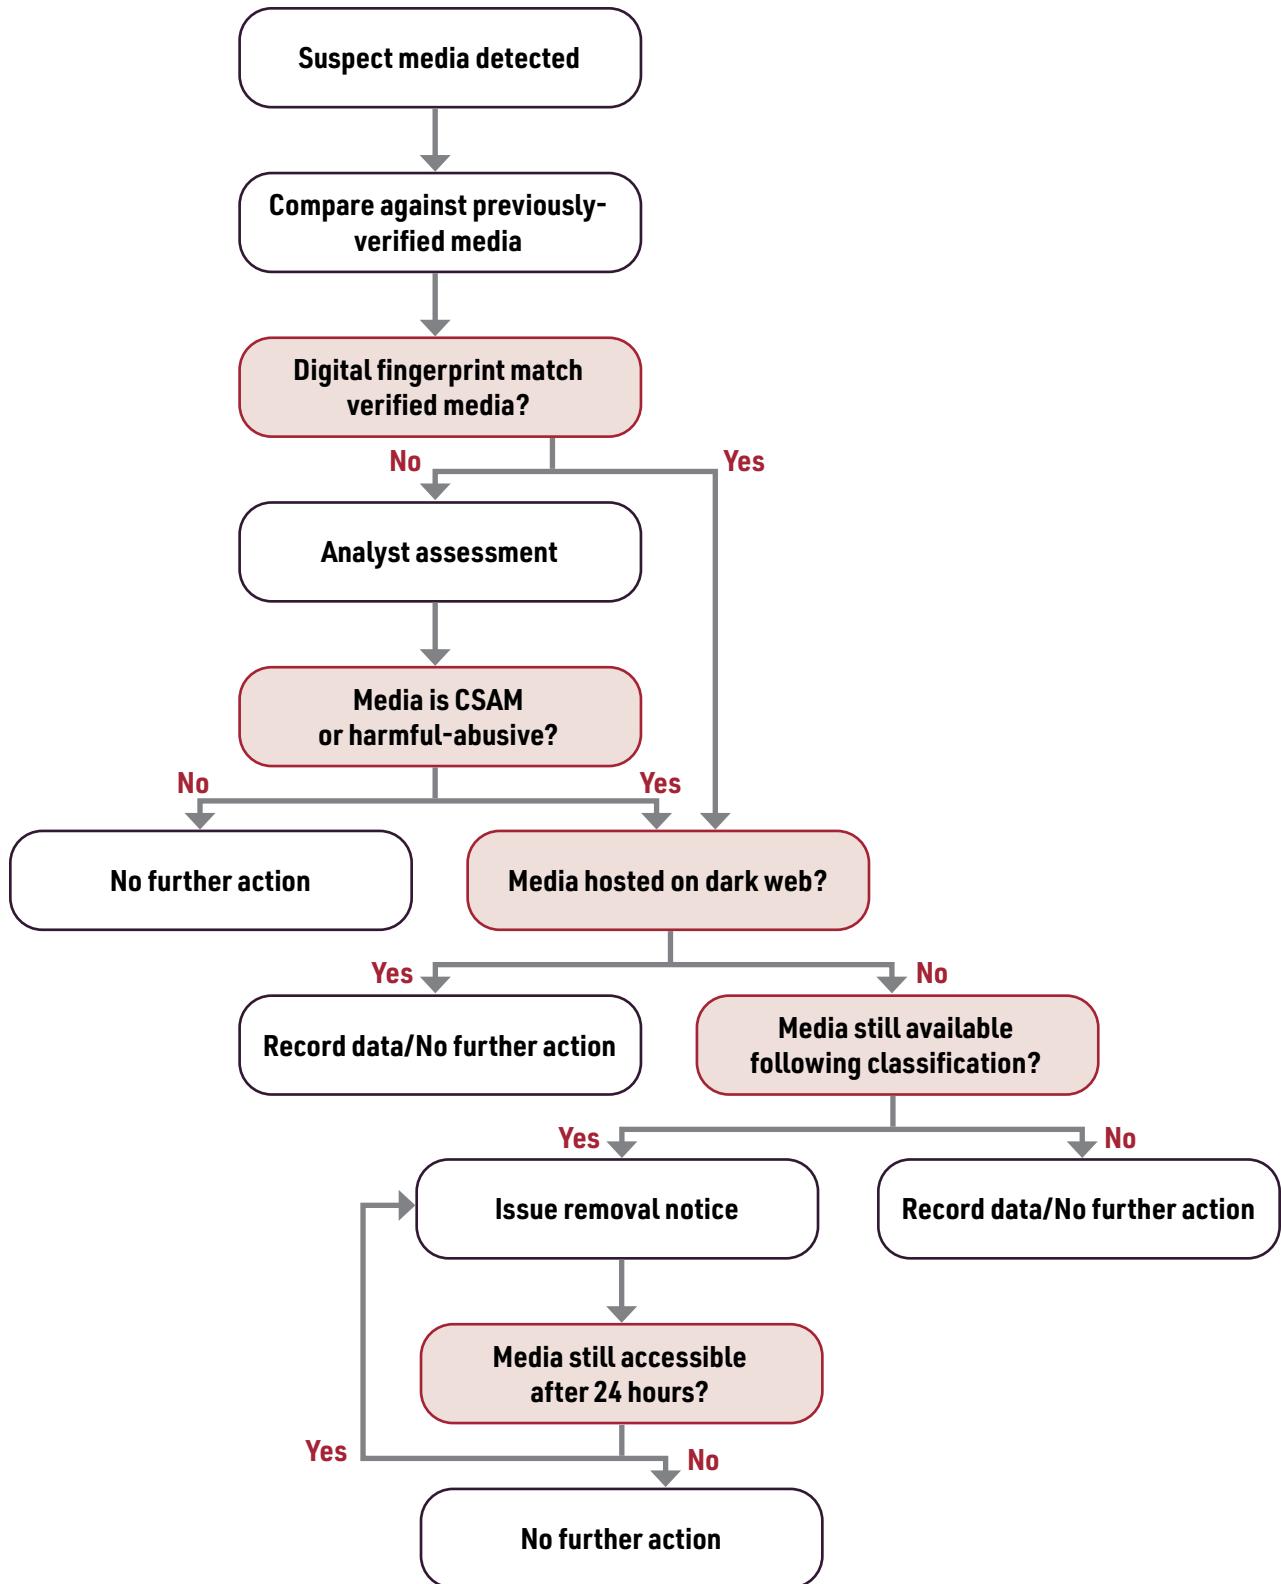
SHIELD
By Project Arachnid

Shield by Project Arachnid is a tool that allows either content administrators or hosting providers to proactively compare incoming or existing media on their service against Project Arachnid's list of digital fingerprints. This tool can be used as part of an ESP's overall content moderation strategy to improve upon and accelerate the detection and removal of CSAM or harmful-abusive content.

For more information on Project Arachnid, visit projectarachnid.ca

³ The open web refers to the publicly accessible areas of the clear and dark web.

Project Arachnid in action

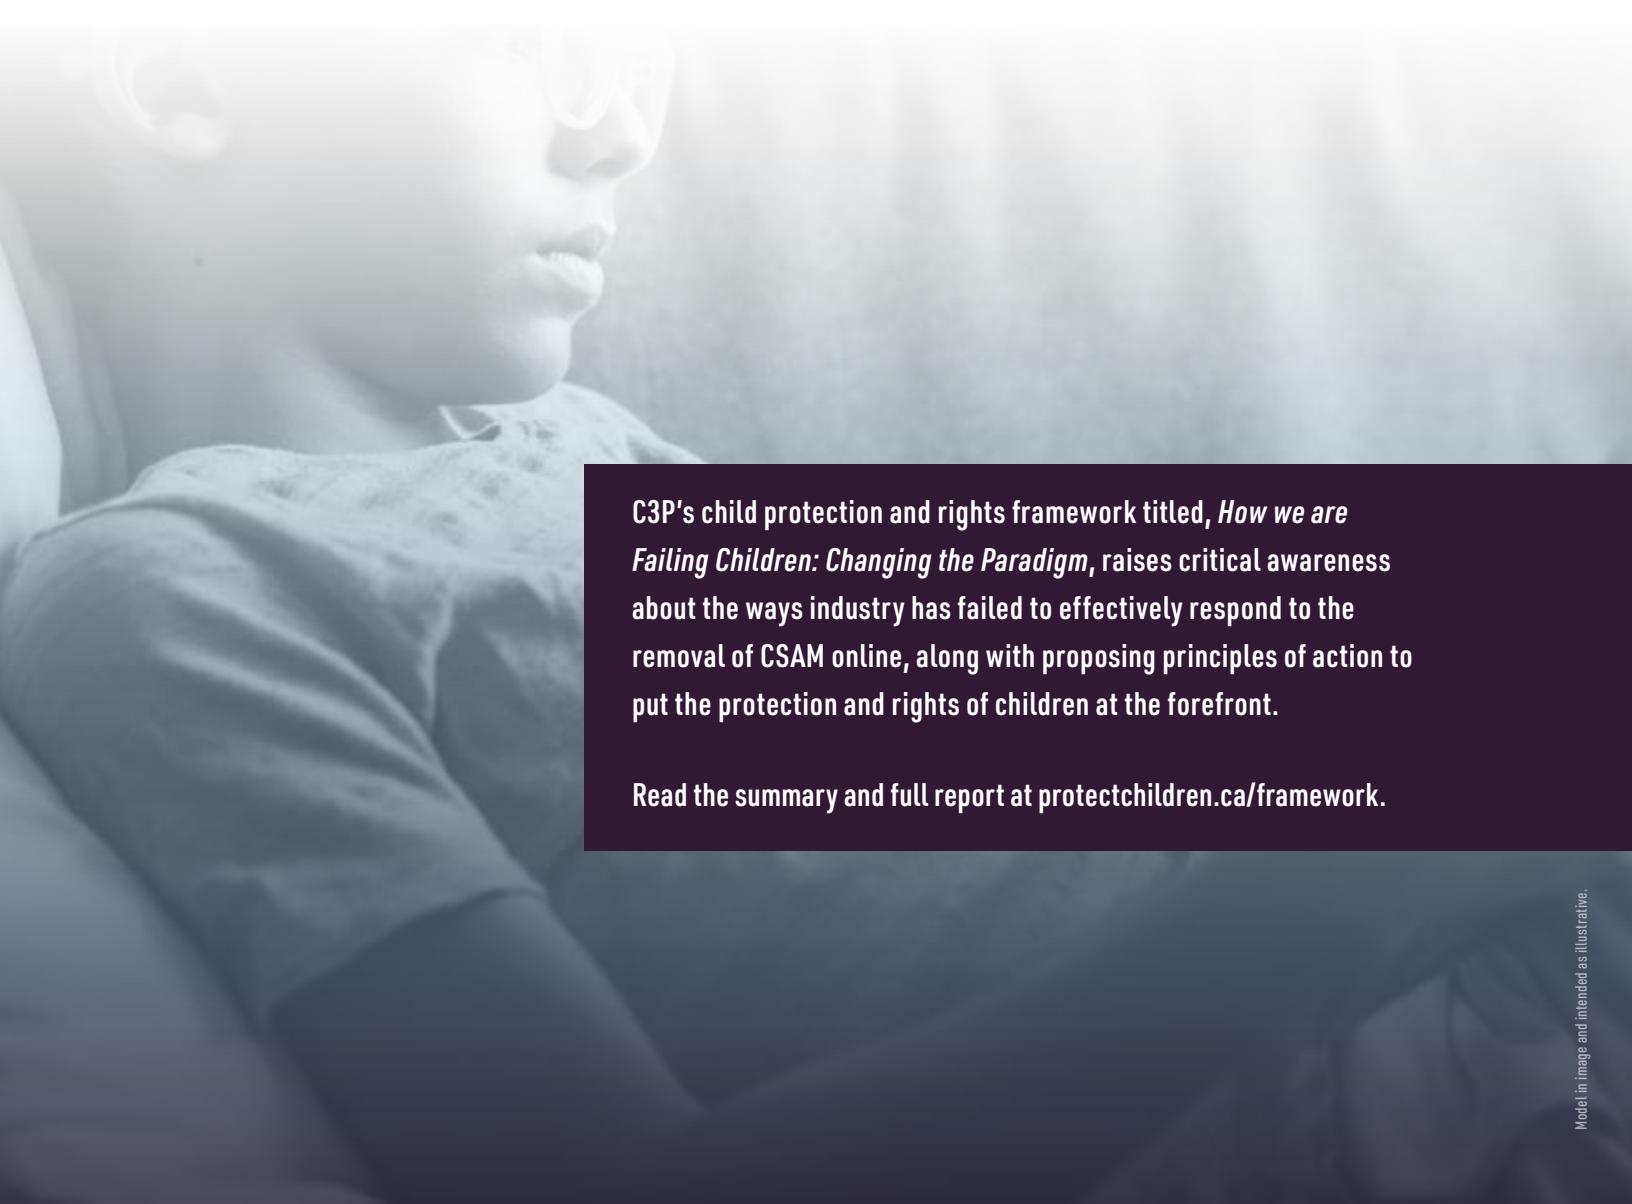


ILLEGAL AND HARMFUL-ABUSIVE CONTENT: A C3P FRAMEWORK

Traditionally, content removal strategies have focused primarily on material that is demonstrably illegal. Unfortunately, this narrow and restrictive approach does not reflect how children are being harmed online. As a result, a wide range of harmful-abusive material circulates with impunity since it does not visually appear to cross a criminal threshold, especially when viewed outside of the broader context of how the media was produced and distributed.

It is clear that confining removal activities solely to what is unarguably criminal has proven to be a short-sighted approach that has failed children worldwide. Informed through close work with survivors and by insight from operating Project Arachnid, C3P developed a child protection and rights framework in 2019 to address this significant gap. The framework establishes a new set of principles for action that places the best interest and protection of children at the forefront of content removal.

As a result of this framework, the scope of Cybertip.ca and Project Arachnid's activities have expanded over time to account for this historically overlooked content.



C3P's child protection and rights framework titled, *How we are Failing Children: Changing the Paradigm*, raises critical awareness about the ways industry has failed to effectively respond to the removal of CSAM online, along with proposing principles of action to put the protection and rights of children at the forefront.

Read the summary and full report at protectchildren.ca/framework.



CSAM and harmful-abusive content

C3P's content-removal and data collection initiatives focus on two core categories of material:

- 1 Images and videos that have been assessed to fall within a criminal definition;
- 2 Harmful-abusive images and videos of children that may not necessarily meet a criminal law threshold.

Harmful-abusive images of children that do not meet a criminal law threshold may nonetheless violate an ESP's terms of service. The images could also be so obviously harmful that no reputable company would host them, especially if they were answerable to any kind of standard or regulatory body.

For example, many platforms have policies surrounding the distribution of unauthorized images of minors, personal information (e.g., doxing) or copyrighted material, and also ban child grooming activities and harassment.

Throughout this report the term "harmful-abusive" refers to images or videos that may be associated with an abusive incident, contain partial nudity and have become publicly available or is being used in a sexualized context. It also includes publicly available images or videos of children being physically abused, tortured or restrained.

Common examples of media that fall within the harmful-abusive category include:

- An image of a child's face covered in what appears to be semen;
- Still images of the initial frames of a known sexual assault video where the child is still clothed or semi clothed that is taken during the progression of the sexual abuse. In isolation these may not necessarily meet a criminal law threshold, but they are part of a larger sequence of illegal material;
- Images of children or adolescents in bathing suits copied from social media accounts and distributed on websites dedicated to the sexualization of children;
- Sexualized content of children that includes images where there is a deliberate attempt to portray adult sexual positions or acts that suggest the sexual availability of the child. The child may be fully or partially clothed.
- Images or videos of a child being physically assaulted or tortured.

Where CSAM and harmful-abusive content is found

It is a common misconception that CSAM and harmful-abusive content are relegated solely to the dark web. In fact, the majority of illegal media detected by Project Arachnid hides in plain sight on the clear web on image/file hosting services, forums, content delivery networks, and also mainstream and fringe adult pornography sites.

The following table provides a general overview of the various areas within the digital ecosystem where Project Arachnid and Cybertip.ca analysts have encountered CSAM or harmful-abusive content.⁴

Category	Examples of services/platforms
Image hosting sites	Imgur®, ImageShack™, Flickr®, PostImage
File hosting sites	Megaupload, Dropbox®, WeTransfer™, dl.free.fr
Cloud service providers, virtual private servers, traditional web hosting	Amazon® AWS™, Microsoft® Azure™, Rackspace®, GoDaddy®, DreamHost®
Content delivery networks (CDN)	Cloudflare®, Fastly®, Akamai®
Dark web forums/chats	Sites primarily hosted as onion services on the Tor network.
Search engine results/cache	Google®, Bing®, Yahoo!®, Yandex®
Forums/chats/messaging	Reddit®, Twitch®, 4chan™, Discord™, WhatsApp®, Kik®
Adult pornography sites (fringe)	Specific fetish/interest, revenge pornography
Adult pornography sites (mainstream)	Pornhub®, XVIDEOS™, YouPorn®
Social media	Twitter®, Facebook®, Instagram®, Snapchat®

⁴ The examples are not exhaustive, and are included only to assist the reader in understanding the category listed.

All trademarks with an ® are registered by the owner in Canada and in the U.S.; all trademarks with a ™ symbol are registered only in the U.S.

METHODOLOGY

Data collection

Project Arachnid detects suspected or known CSAM and harmful-abusive content in three ways:

- 1 Crawling publicly accessible URLs previously reported to Cybertip.ca;
- 2 Crawling publicly accessible URLs/media reported directly to Project Arachnid's API by participating industry members;
- 3 By crawling certain areas of the dark and clear web known to host CSAM.

When Project Arachnid detects suspected abuse material, the content is assessed and categorized by trained analysts with C3P or other contributing tiplines.

For each of these files, Project Arachnid stores the actual media file and key data points such as the date of detection and which entities were notified. These records form the primary source of data behind this report's analysis.

Image categories

For the purposes of this report, media is categorized in three simplified reporting categories derived from Project Arachnid's internal image assessment process.

Pre-pubescent CSAM

This category of media refers to content that likely meets a criminal definition of CSAM. It includes images where the depicted victim is pre-pubescent or is in the early stages of puberty.

Post-pubescent CSAM

This category of media refers to content that likely meets a criminal definition of CSAM. It includes images where the depicted victim's status as a child at the time the image was taken has been confirmed and the child is post-pubescent. This category includes media containing victims that are in the later stages of puberty.

Harmful-abusive

Harmful-abusive media are those that do not appear to meet a criminal law threshold across multiple jurisdictions, but may nonetheless violate an ESP's terms of service. These images may also violate the privacy or safety of a child, or be associated with CSAM. For more details refer to the description of C3P's framework (p. 10)

Key Transparency Measures

Volume of media detections, media targeted for removal and removal notices

Project Arachnid generates data that can be analyzed using several different measures. Quantifying the volume of the availability of material and removal notice activities are done in three ways:

Media detections

Media detections is the measure of the total accessibility of CSAM and harmful-abusive content detected by Project Arachnid. It represents all detections of images, videos and multimedia archive files within the areas of the internet Project Arachnid crawls. Throughout this report, the term “media detection” refers to media that has been reviewed by an analyst and constitutes either CSAM or harmful-abusive content.

A single media source may be embedded and displayed across several websites on the internet. Since each online location in which a child’s image is displayed is an independent violation of that child’s right to privacy and dignity, Project Arachnid considers each of these sightings to be a unique media detection. Media detections can also be thought of as individual sightings of the content across the web.

Media targeted for removal

Media targeted for removal represents all media detections that triggered the issuance of a removal notice to an ESP. For reasons explained later in the report, not all media detections can be targeted for removal.

Removal notices

Media targeted for removal may lead to the issuance of one or more removal notices to an ESP. Project Arachnid re-issues removal notices to an ESP every 24 hours, until the media is no longer detected.

Removal times

Removal time calculations are based on the interval of time in days between the issuance of a removal notice for a specific media at a specific URL to an ESP and the last date the specific media was detected by Project Arachnid at that same URL. Since the system re-crawls the actioned media every 24 hours following notification, the elapsed link uptime is accurate to within 24 hours. Media URLs that become inactive in less than 24 hours from the point of notification are rounded up to the day for the purposes of this report.

It is important to note that content may become inaccessible for several reasons, some of which may not necessarily be related to action taken by the targeted ESP or Project Arachnid’s removal notice.

Given the severity of harm to victims caused by the public display of even a single image or video, this report’s focus is primarily on the 90th percentile removal time. This measure represents the maximum time the majority (90%) of URLs remained live on the internet from the moment a removal notice is issued to the ESP.

Media recidivism

This measure represents the number of times an image or video on a specific ESP’s service that was previously the subject of a Project Arachnid removal is re-detected at a later date on the same service, but at a different URL. For the purposes of this measure, establishing the recidivist status of an image is based on the re-emergence of an identical SHA-1 hash value. SHA-1 values are cryptographic hash values (or digital fingerprints) derived from a media file and represent a unique digital fingerprint, distinguishing it from other images. Since SHA-1 matching requires an exact cryptographic match, images that are nearly identical or that are close derivatives are not considered matches under this measure.

When interpreting media recidivism rates in this report, it is important to note that the basis for establishing that a detected file is a recidivist image is based on C3P's original determination that the content in question is CSAM or harmful-abusive content. That original determination is based on a variety of factors that include whether the image or digital fingerprint of the image is within the databases of known CSAM used by Project Arachnid, as well as an independent visual assessment of the physical development and sexual maturation characteristics of the individual(s) in the image. It also includes other environmental cues depicted in the image and the context in which the image is detected.

When an ESP chooses not to remove an image upon C3P's request, the re-emergence of that image would not be viewed by the ESP as a recidivist event.

Limitations

While the data collected by Project Arachnid provides an unprecedented view into the nature and scale of the distribution of CSAM and harmful-abusive content, it is not without limitations.

The following are key notes about the data to ensure the conclusions of this report are kept within their appropriate context.

- 1 It is likely that in nearly all cases, ESP-specific figures throughout this report understate the true volume of CSAM and harmful-abusive content that could be associated with an ESP. For a single image or video to exist online, multiple service providers each play a separate and distinct role.

Since Project Arachnid is currently tooled to target the ESP most likely to take removal action, ESP-specific data captures only one actor in the chain for a specific actioned media. In future reports, C3P expects to adopt a more expansive approach to better represent the full extent to which ESPs are associated with CSAM and harmful-abusive content.

- 2 The data collected for any given ESP represents only what Project Arachnid encountered at specific point in time, based on reports from the public or crawl prioritization. It does not necessarily represent the totality of publicly accessible CSAM on an ESP's service. The volumes of image detections, or the trend of detections over time, on a specific service may not be representative of the total prevalence of CSAM or harmful-abusive content for an ESP.

The semi- or fully closed nature of many websites — especially social media or direct messaging platforms — make them largely inaccessible to Project Arachnid. Therefore, records held by Project Arachnid do not reflect the true extent of CSAM or abusive-harmful content on these ESP's services.

- 3 The volume of media detected for a given ESP is driven by a multitude of factors, such as tips from the public, the nature of the website and the nature of the content. For these reasons, exercise caution when comparing ESP-specific figures.

- 4 ESPs are not necessarily fixed legal entities, nor is it always clear which company is behind the operation of a particular service. ESPs may evolve, merge, split and re-brand over time. Data on platforms operated by related companies are not necessarily combined.

ANALYSIS AND DISCUSSION

Between 2018 and 2020, Project Arachnid's crawling activities have detected more than 5.4 million images or videos of verified CSAM or harmful-abusive (Table 1.1). This content was detected on the services of more than 760 electronic service providers operating across the globe.

Table 1.2 shows C3P analysts and international tipline contributors have collectively assessed more than 4.9 million individual images in the past three years. These assessments contribute to a constantly growing repository of hash values used to enhance future media detection.

As of the writing of this report, C3P is facing a backlog of more than 32.8 million suspect media that have yet to be assessed (Table 1.2). This is because the rate at which Project Arachnid detects suspect media far outpaces the human resources available to assess the content.

Over the three year period, 626,110 media detections were targeted for removal by Project Arachnid (Table 1.1). The significant discrepancy between media detections that were ultimately targeted for removal and the volume of suspect media is due to three factors:

- 1 Archive files that may contain collections of hundreds of images are often treated by Project Arachnid as a bulk removal initiative. This means a removal request may relate to several images but represent only a single record.
- 2 On many occasions issuing a removal notice was no longer required since the offending media was removed or was not accessible by the time it was reviewed. This is a consequence of the assessment backlog.
- 3 Some media were found on the dark web and therefore the identity of the ESP is unknown. No action beyond data collection can occur in these situations.

Table 1.1

At a glance: Project Arachnid activity				
	2018	2019	2020	Total
Verified media detected	1,411,203	2,494,316	1,511,194	5,416,713
Verified media targeted for removal	57,685	301,990	266,435	626,110
Removal notices sent	502,162	1,699,017	1,633,698	3,834,877

The term "verified media" refers to media that an analyst has assessed and evaluated as has been assessed by an analyst and evaluated as being either CSAM or harmful-abusive material.

Table 1.2

Suspect media and assessment backlog	
	Total
Total suspect media detected	37,854,878
Media awaiting assessment	32,899,122
Assessed media	4,955,756

The term "suspect media" refers to any media that is reasonably believed to be CSAM but which has not been through C3P's assessment process.

Verified media detections

Of the 5.4 million verified media detected by Project Arachnid between 2018 and 2020, 2.9 percent (n=158,950) were hosted directly on Tor onion services (a subset of the dark web) with the remaining 5.2 million hosted directly on the clear web (Table 2.1). This represents average daily detections of nearly 5,300 images or videos per day over the three-year period.

As illustrated in Figure 1.1 there is no obvious trend in detection volumes over time. Project Arachnid's crawler is not fixed on a pre-determined set of websites. Rather it crawls areas of the web based on link referrals, public tips and many other factors. In addition, the nature of the type of websites or services that distribute CSAM is such that they may be short lived or frequently change hosting providers as they seek companies willing to tolerate the nature of their content.

In many cases, a single network of websites may generate incredibly high volumes of detections and suddenly go offline as a result of Project Arachnid's efforts or otherwise, causing a sudden drop in detections. For these reasons detection volumes — and by extension Project Arachnid's activities — may vary significantly over different time periods.

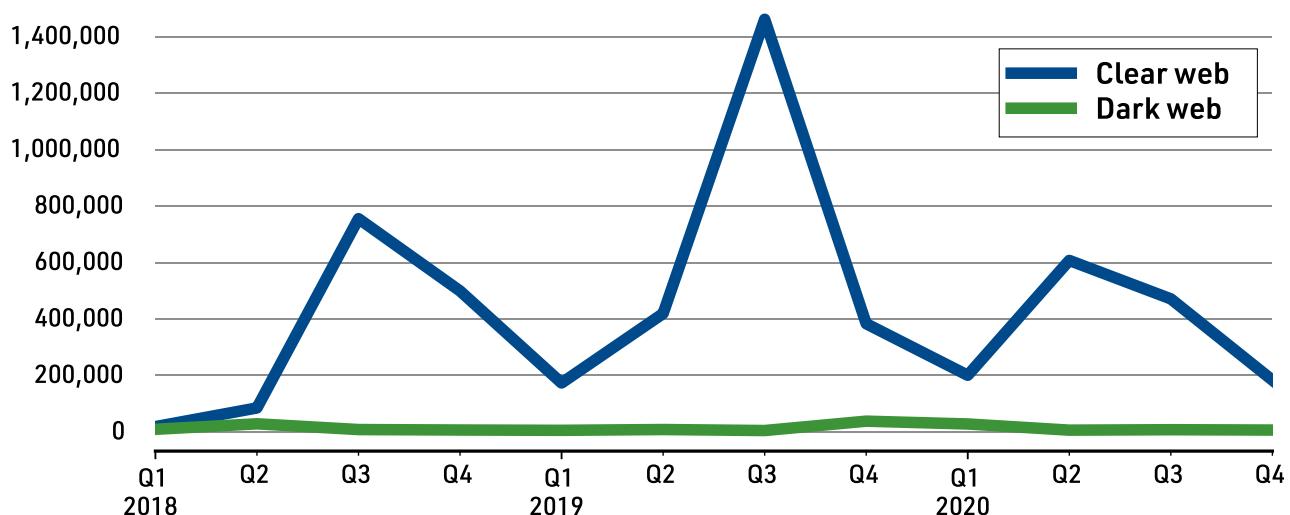
Table 2.1

Verified media detections, by web type				
Web type	2018	2019	2020	Total
Clear web	1,358,109	2,437,230	1,462,424	5,257,763
Tor	53,094	57,086	48,770	158,950

The term "Tor" refers to the largest network within the dark web.

Figure 1.1

Verified media detections, by web type





Model in image and intended as illustrative.

Table 2.2 shows that over the past three years, pre-pubescent CSAM (n=3,403,748) and harmful-abusive (n=1,892,792) have been the two most common image categories encountered by Project Arachnid.

Images or videos containing post-pubescent CSAM represents a small fraction of the overall verified media detections by Project Arachnid (n=120,173).

Table 2.2

Verified media detected, by image category				
Description	2018	2019	2020	Total
Pre-pubescent CSAM	738,378	1,792,639	872,731	3,403,748
Post-pubescent CSAM	27,996	51,432	40,745	120,173
Harmful-abusive	644,829	650,245	597,718	1,892,792

Sum of totals may not reflect totals shown in chart 2.1 due to exclusion of records that have been recategorized.



Discussion

The fundamental measure for capturing the scale of CSAM and harmful-abusive content availability seen by Project Arachnid is the volume of detected media on the internet.

As mentioned earlier in the report, Project Arachnid's reach does not extend into all areas of the internet. As a result, figures provided in this report are certain to be a gross underestimation of the true extent of availability of this material on the internet. The high volumes of content mandatorily reported into NCMEC by large technology platforms, and that are generally not within the scope of Project Arachnid's activities, supports this assertion.

Table 2.2 shows the vast majority of verified media that is CSAM does not depict images of post-pubescent children. This finding however is not representative of the reality of adolescent material on the internet.

Imagery depicting younger victims has historically been and continues to be prioritized in most interventions. From the perspective of law enforcement or categorization efforts, establishing whether an image meets a legal definition of CSAM is more likely to be plainly obvious when the victims involved are younger or pre-pubescent.

However, with post-pubescent victims there often exists a high level of uncertainty in the categorization decision. The need to acquire additional contextual information about the image invariably increases the complexity and resources required to make a final assessment. Consider the challenges in determining, based on visual cues alone, whether an adolescent with full sexual maturation characteristics is a minor or an adult.

This inherent challenge in image categorization with unidentified post-pubescent victims and the patchwork of legal standards related to CSAM across nations has over time skewed image categorizations toward younger victims.

For example, the International Child Sexual Exploitation (ICSE) image and video database, an intelligence tool managed by Interpol, established the “baseline”. The baseline is intended to be a catch-all image category that can be assumed to meet a criminal threshold across nearly all jurisdictions. It is described as, “an international standard to isolate the worst of child abuse materials.”

According to a 2018 report by Interpol and End Child Prostitution and Trafficking (ECPAT) baseline images must depict:⁵

- A real child (not an artificially created image);
- A prepubescent child (no sign or very first signs of puberty, appearing to be younger than 12 or 13 years old);
- A child involved in or witnessing sexual/abuse activities; and
- The media has a clear focus on the child's sexual/anal area.

The net effect of this reality — databases containing the digital fingerprints of known CSAM are skewed heavily toward younger victims and towards media that is on the extreme end of the spectrum. Since most image detection technologies, including Project Arachnid, rely on these digital fingerprint repositories to uncover media on the internet, the nature of what is detected through automation reflects this bias toward extreme content of younger victims.

⁵ INTERPOL. (2018). *Towards a global indicator on unidentified victims in child sexual exploitation material: Technical report*. <https://www.ecpat.org/wp-content/uploads/2018/02/Technical-Report-TOWARDS-A-GLOBAL-INDICATOR-ON-UNIDENTIFIED-VICTIMS-IN-CHILD-SEXUAL-EXPLOITATION-MATERIAL.pdf>

In addition to this, many post-pubescent victims may feel shame or fear related to the creation, distribution and public display of images depicting their abuse or vulnerable moments. C3P Cybertip.ca analysts report when these victims do seek assistance, there is often a desire to avoid triggering a law enforcement or legal response out of continuing fear of the offender who created the images, as well as a strong desire to avoid drawing further attention to themselves.

For these reasons, there is no doubt media detected by Project Arachnid dramatically underrepresents the true scale of harm to adolescent victims.

How the dark web facilitates CSAM distribution

Based on Project Arachnid media detections (**Table 2.2**), the volume of content that is directly hosted or accessed on Tor is relatively small compared to the clear web. Without additional context, this finding may lead one to conclude the dark web has a limited role in the distribution of CSAM, when in fact the opposite is true.

Rather than being the place CSAM and harmful-abusive content is physically hosted, dark web networks such as Tor, are often the conduit for directing individuals to the presence of the material on the clear web. Entire communities, emboldened by the anonymity the Tor network affords, congregate within forums where information related to CSAM and other clandestine activities are discussed. Topics often include: where and how to access illegal media, child grooming and abuse tutorials/manuals, encryption, cyber security and evidence destruction strategies.

This relationship between the clear and dark web is important to understand when crafting regulation or for ESPs adopting proactive measures.

Tor, accessible only through specialized browsers, anonymizes the web traffic between a user and the website they are visiting. The process through which traffic is anonymized and encrypted, however, comes at a cost — substantially slower page loading and media download speeds.

For these reasons, those interested in distributing large multimedia collections often choose to upload their content on archive or image hosting services on the clear web, where download speeds are much faster.

Typically, distributors of this material will upload an encrypted, password-protected archive file that may contain hundreds of images or videos onto a free file hosting service that collects limited to no data about its users. Once uploaded, the distributor will then turn to forums on the dark web and provide members with access to the direct download link and password for the archive file.

Content distributed in this fashion accounts for the largest volume of verified media detected by Project Arachnid. As noted earlier in this report, Project Arachnid's reach on the dark web is currently limited to Tor, and so the distribution characteristics of other dark web networks may be different.

Media targeted for removal

Between 2018 and 2020, Project Arachnid's removal notices targeted, on average, 571 images or videos per day. A large increase in media targeted for removal in late 2019 relates to a single ESP — free image hosting provider Imagevenue.com — whose service was being used by hundreds of third-party websites to host CSAM.

Many of the domain names used by these websites contained words that were indicative of child sexual exploitation, including "teen", "cuties", and references to the term "jailbait." This particular grouping of sites were uncovered and crawled by Project Arachnid, leading to a sudden surge in records tied to Imagevenue.com, the service that provided image hosting for the content.

As noted earlier in this report, the number of removal efforts initiated by Project Arachnid shown in Table 3.1 do not correspond with the actual volume of images being targeted in practice due to the existence of file archives that may contain thousands of images.

For the purposes of this report, ESP-specific information presented in Table 3.2 is provided for those with 5,000 or more media or files that have triggered the issuance of one or more removal notices. However, for technical reasons, records related to one ESP in particular — French telecommunications company *Free* — is tracked differently by Project Arachnid and is therefore not reflected in Table 3.2.

Project Arachnid records for the time period of this report show Free, which operates the file hosting website dl.free.fr, hosted at least 18,000 archive files, collectively containing nearly 1.1 million media files of apparent CSAM or harmful-abusive content. Project Arachnid has detected access points to these archives files across many areas of the internet, representing more than 2.7 million media detections. This report provides details related to Free.fr in a standalone case study (See p. 42).

The vast majority of ESPs that have received removal notices from Project Arachnid have been image hosting providers or file hosting services. **Table 3.2** shows the following ESPs have had the greatest volume of media detections targeted for removal:

- **Imagevenue.com:** A domain whose registrant contact information shows the website is based in the Czech Republic, and which operates an image hosting service, but uses hosting provider services for its user-generated content (n=144,000);
- **Serverel:** A U.S.-based company offering hosting service on its own server infrastructure (n=72,412).
- **CloudFlare:** A U.S.-based company offering content delivery network (CDN) services linked to several other ESPs that have many intersections with Project Arachnid (n=49,183).
- **Incrediserve LTD:** A Netherlands-based company that provides hosting services (n=39,400).
- **Trichan forums:** A now defunct network of what appeared to be centrally controlled forums that allowed users to directly host content on its website. Many of the companies providing hosting services to the Trichan forums are reflected in Project Arachnid's records. The figures provided in **Table 3.2** below (n=34,157) under represents the true scale of content on these forums, as repeated non-responses to removal notices led to the adoption of an alternative removal strategy which impacted the record management process.
- **NForce Entertainment B.V.:** A Netherlands-based company offering hosting service on its own server infrastructure (n=23,211).

Table 3.3 reveals that for nearly all the ESPs highlighted in this report, pre-pubescent CSAM is the most commonly actioned type of material by Project Arachnid. The only exception as seen in **Table 3.4** is Serverel which has mostly received removal notices related to post-pubescent CSAM (n=66,824).

A review of the removal notices for Serverel to determine possible reasons for this irregular pattern suggests that many websites using its services are ephemeral adult content websites that host post-pubescent material among legal adult content. Project Arachnid has detected at least 1,200 unique websites displaying these images using Serverel's hosting services.

Table 3.1

Media targeted for removal			
2018	2019	2020	Total
57,685	301,990	266,435	626,110

Archive files containing several images recorded as single entry in table.

Table 3.2

Media targeted for removal, by ESP <i>(Includes only ESPs with 5000+ flagged media detection)</i>					
ESP name	Service type	2018	2019	2020	Total
Imagevenue	Content administrator	6,214	76,579	61,099	143,892
Serverel	Hosting service	826	9,121	62,465	72,412
CloudFlare	Content delivery network	3,117	36,604	9,462	49,183
Incrediserve LTD	Hosting service	15,861	19,353	4,186	39,400
Trichan	Content administrator	7,092	27,065	0	34,157
NForce Entertainment B.V.	Hosting service	789	14,481	7,941	23,211
ImgOutlet.com	Content administrator	0	10,182	8,400	18,582
ImgView.net	Content administrator	96	6,509	4,035	10,640
FranTech Solutions	Hosting service	54	688	8,987	9,729
ImgDew.com	Content administrator	0	5,618	3,574	9,192
Host Sailor	Hosting service	117	6,845	1,778	8,740
ColoCrossing	Hosting service	1,369	5,309	1,131	7,809
ALFA TELECOM s.r.o.	Hosting service	501	6,909	62	7,472
DataWeb Global Group B.V.	Hosting service	598	2,740	3,765	7,103
ImgMaze.com	Content administrator	0	4,541	2,300	6,841
Liteserver Holding B.V.	Hosting service	2	4,125	2,639	6,766
ImageBam	Content administrator	42	2,363	3,934	6,339
OVHcloud	Hosting service	3,104	1,873	1,304	6,281

Records related to the ESP Free not reflected in table.

See case study on page 42 for details on this ESP.

Table 3.3**Media targeted for removal, by image category**

Image category	2018	2019	2020	Total
Pre-pubescent CSAM	51,700	282,500	188,486	522,686
Post-pubescent CSAM	2,581	11,842	68,607	83,030
Harmful-abusive	1,171	2,163	7,208	10,542

Figures do not reflect full contents of archive files containing multiple images.

Table 3.4

ESP name	Pre-pubescent CSAM	Post-pubescent CSAM	Harmful-abusive
Imagevenue	142,449	236	550
CloudFlare	46,033	1273	446
Incrediserve LTD	37,589	118	779
Trichan	32,215	49	624
NForce Entertainment B.V.	23,066	66	58
ImgOutlet.com	18,534	15	28
ImgView.net	10,549	30	32
ImgDew.com	9,139	22	28
Host Sailor	8,689	11	22
FranTech Solutions	7,893	24	1,753
ColoCrossing	7,573	7	57
ALFA TELECOM s.r.o.	7,406	51	13
ImgMaze.com	6,782	33	24
Liteserver Holding B.V.	6,748	0	17
ImageBam	6,189	10	105
OVHcloud	5,543	383	126
Serverel	4,529	66,824	3
DataWeb Global Group B.V.	2,583	4,151	25

Records related to the ESP Free not reflected in table.

See case study on page 42 for details on this ESP.

Removal notices

Between 2018 and 2020, Project Arachnid sent nearly 3,500 removal notices every day. **Table 4.1** shows that over the course of three years, over 3.8 million removal notices were issued to ESPs.

Project Arachnid's system is designed to re-issue removal notices every 24 hours until the media is no longer detectable at the targeted URL. As such, the volume of removal notices issued to an ESP is directly correlated to both the number of images or videos targeted for removal and the length of time the content remains accessible.

A significant rise in removal notices issued in early 2019 relates to a concerted effort to have thousands of images taken down from a now-defunct network of online forums dedicated to child exploitation known as the Trichans. A case study later in the report provides greater insight into this specific initiative (See p. 44).

Broken down by image categorization type, **Table 4.2** shows pre-pubescent CSAM (n=2,986,280) is the most common image category actioned by Project Arachnid. However, the volume of removal notices related to post-pubescent CSAM (n=737,718) is significantly greater than what might be expected given the relatively low volumes of detections shown in **Table 2.2**. This indicates post-pubescent CSAM requires the issuance of a much greater number of removal notices, and consequently, longer timeframes before the media becomes inaccessible.

Table 4.2 shows removal notices related to images classified as harmful-abusive saw a sharp increase in 2020 (n=88,825). This timing coincided with the release of C3P's child protection and rights framework (See p. 9) which led to an expansion of the scope of media triggering removal notices.

Table 4.1

Removal notices issued				
2018	2019	2020	Total	
502,162	1,699,017	1,633,698		3,834,877

Figures include initial and, when required due to non-removal, follow up removal notices to ESPs.

Table 4.2

Removal notices issued, by image category				
Image category	2018	2019	2020	Total
Pre-pubescent CSAM	482,399	1,633,212	870,669	2,986,280
Post-pubescent CSAM	9,934	40,000	687,784	737,718
Harmful-abusive	6,589	12,512	69,724	88,825

Sum of totals may not reflect totals shown in chart 4.1 due to exclusion of records that have been recategorized.

Table 4.3

ESP name	Pre-pubescent CSAM	Post-pubescent CSAM	Harmful-abusive
Trichan	733,927	776	7704
Incrediserve LTD	381,498	641	4,696
NForce Entertainment B.V.	217,068	1,623	525
CloudFlare	170,923	9,646	1,158
Imagevenue	168,448	291	645
ColoCrossing	165,709	170	1,249
FranTech Solutions	94,707	877	15,047
Liteserver Holding B.V.	52,189	0	88
Serverel	44,662	637,631	3
ImgOutlet.com	34,830	34	41
OVHcloud	31,245	2,818	305
Host Sailor	24,768	42	23
ImgView.net	23,693	63	54
ImgDew.com	21,535	41	70
ALFA TELECOM s.r.o.	17,142	133	149
ImgMaze.com	16,257	65	61
ImageBam	6,565	10	105
DataWeb Global Group B.V.	3,778	11,017	26

Records related to the ESP Free not reflected in table.

See case study on page 42 for details on this ESP.



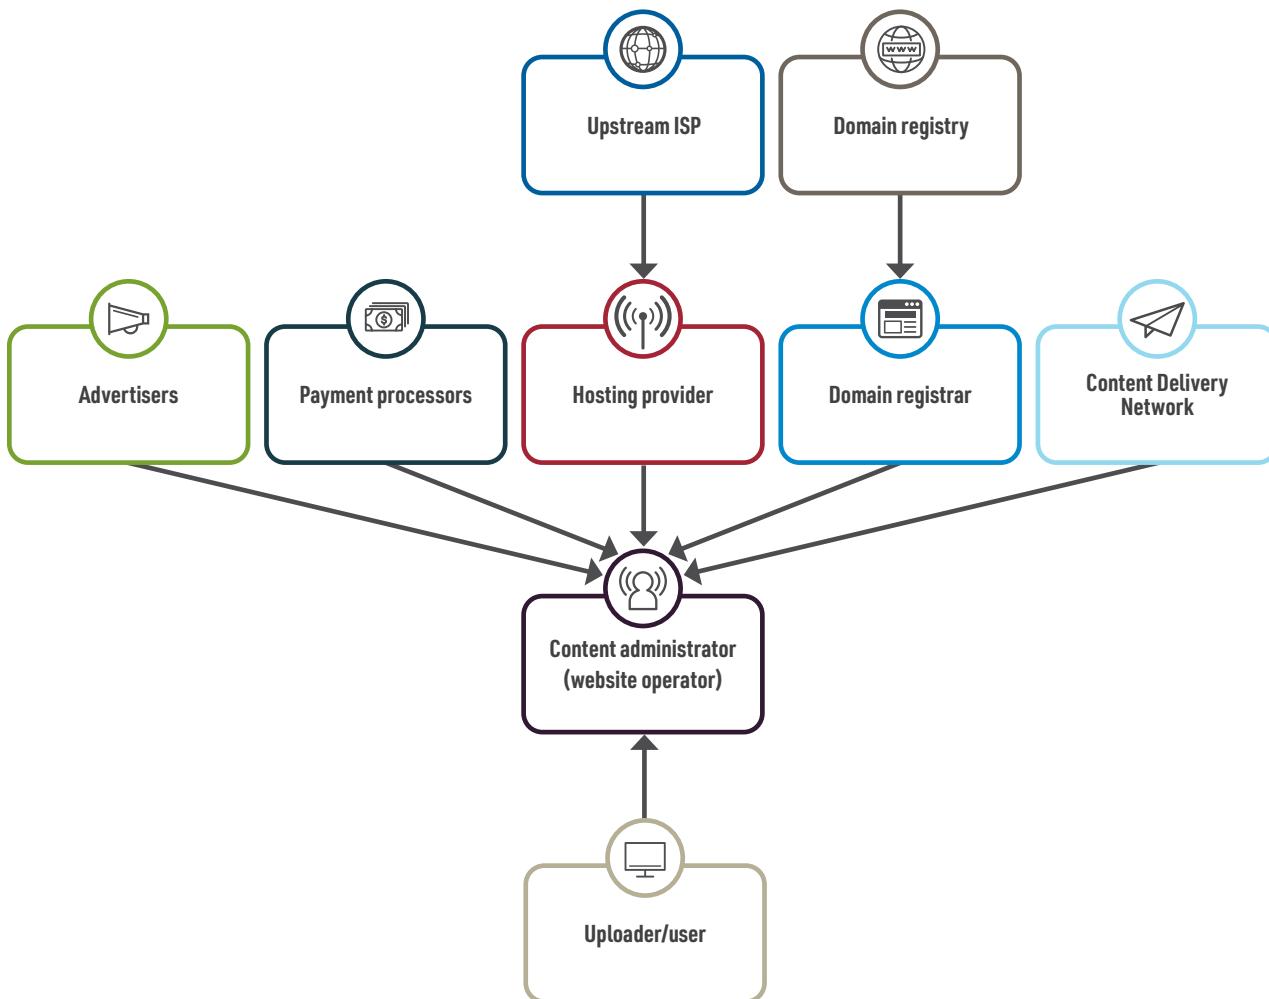
Discussion

While this report highlights the intersections key ESPs have had with CSAM and harmful-abusive content, it must be noted that relying on the presented data alone does not paint a complete picture of the role each ESP's activities play in enabling access to harmful content.

Removal notices issued by Project Arachnid generally target a specific ESP based on a combination of factors. Evaluations based on which ESP possesses the most immediate control over the targeted media, responsiveness to removal notices and the availability of contact information help guide where notices are ultimately sent. The ESP-specific figures in this report reflect those entities to which notices are sent, which is not reflective of the broader chain of ESPs associated with facilitating the public display of each detected image or video.

Figure 2.0 illustrates a crucial point: The existence of a single image or video on the internet, ultimately requires a coordinated series of services by a number of companies, all of which generally have some ability to mitigate or stop the proliferation of CSAM or harmful-abusive content on specific services.

Figure 2.0





As an illustrative example, the following are common steps required of an individual intent on creating a website that may be dedicated to the sexualization of children.

- 1 An individual must register a domain name with a domain name registrar. The domain name registrar is authorized to sell domains by a domain registry.
- 2 Once the individual has secured a domain name, they must then find a means to upload and make their material available on the internet. This requires the services of a hosting provider. Hosting providers may own or rent physical servers or lease space on another company's servers, potentially spanning many jurisdictions. These ESPs generally have the technical and/or legal ability to shut down any website or a particular server on its services, and are generally able to impose specific and legally binding terms of service on their customers.
- 3 The individual may make use of a content delivery network service (CDN). These services provide website visitors with faster loading times by relaying a mirrored (cached) version of the site to servers distributed around the globe, effectively reducing the distance data must travel. These services appeal to website administrators providing access to CSAM or harmful-abusive content for another key reason: CDNs typically mask the identity of the website administrator and hosting providers.
- 4 The individual may also make use of a virtual private network service (VPN) to mask their origin IP address from their hosting provider during the course of their activities.
- 5 To monetize these efforts the individual may also make use of payment services such as major credit card companies, or other online payment systems. In addition, the content administrator may promote their content in hopes of securing advertising revenues.

Exerting control over the entire digital ecosystem described above are the upstream internet service providers and Tier 1 networks which are effectively the backbone of the internet. Even the largest technology companies are reliant on these companies for their platforms to be accessible globally to users.

Surges in actionable media for certain ESPs

As observed in the data, Project Arachnid removal activities aimed at certain ESPs occasionally surge upward. As noted earlier in the report, detection volumes ebb and flow — sometimes dramatically — depending on the time period. It is important to note that the nature of the activities of a hosting provider's customers (i.e., content administrators) can greatly influence the volume of removal notices they received.

Table 4.3 shows two ESPs in particular, Imagevenue.com and Serverel, received a significant volume of removal notices from Project Arachnid in 2019 and 2020. These volumes were driven in large part by networks of third-party websites that were widely distributing media hosted on their services.

This misuse of file hosting services by third-party actors underscores how important it is for these services to be vigilant and invest in tools to block the uploading of undesirable content and to invest in adequate human moderation resources.

Another key consideration that may impact the volume of Project Arachnid media detection for any given ESP is whether they allow user-generated content to be uploaded and accessed via dark web networks.

Certain ESPs — including some specifically highlighted in this report — allow anonymous users to upload and access content on their platforms from the Tor network. Failing to adopt network security measures to block this type of suspect traffic means users can exploit an ESP's platform for CSAM distribution, especially those whose services require no account or fees.

Server locations of media targeted for removal

The distribution of CSAM and harmful-abusive content is a global issue. The decentralized nature of internet services means that ESPs can have physical or digital footprints across many locations.

Server location information reported by Project Arachnid is based on data available through the services of Maxmind Inc., which reports a 99.8 percent accuracy rate at a country level.⁶

Table 5.1 shows that between 2018 and 2020 nearly 49 percent of removal notices issued by Project Arachnid went to ESPs whose media of concern was hosted on servers located in the Netherlands. The U.S. was second with nearly 33 percent, followed by Canada at 4.7 percent of removal notices.

Table 5.2 shows the top three ESPs (by removal notice volume) for each server location. This table provides insight into which ESPs are making use of server infrastructure, either through direct ownership or through leased space.

Table 5.1

Removal notices issued, by location of server hosting media (Top 15 countries shown)					
GeoIP location	2018	2019	2020	Total	Percent
Netherlands	370,040	1,040,057	468,323	1,878,420	48.8%
United States	54,748	406,420	805,589	1,266,757	32.9%
Canada	15,405	89,059	76,363	180,827	4.7%
Russia	14,204	32,287	51,167	97,658	2.5%
France	15,957	28,021	33,670	77,648	2.0%
Seychelles	5,316	7,797	49,384	62,497	1.6%
Ukraine	10,331	22,834	15,404	48,569	1.3%
Latvia	2	2,329	43,810	46,141	1.2%
Belize	1,273	28,509	4,796	34,578	0.9%
Hong Kong	0	3,984	21,447	25,431	0.7%
Germany	771	4,683	9,890	15,344	0.4%
United Kingdom	2,148	2,279	5,606	10,033	0.3%
South Africa	11	4,126	5,415	9,552	0.2%
New Zealand	332	2,443	5,368	8,143	0.2%
Estonia	91	262	7,569	7,922	0.2%

Server geolocation based on information provided by Maxmind Inc., a cyber-intelligence service.

Percent figures based on all countries, including those not displayed in table.

⁶ For more information on Maxmind Inc.'s accuracy rate, visit www.maxmind.com/en/geoip2-country-database.

Table 5.2

Removal notices based on server location, top 3 ESPs by country					
GeolP country	ESP name	Removal notices	GeolP country	ESP name	Removal notices
United States	Serverel	616,911	Latvia	Telia Latvija SIA	39,204
United States	CloudFlare	183,766	Latvia	FastPic	3,130
United States	ColoCrossing	167,641	Latvia	Telenet Ltd	1,970
United Kingdom	JPG4.NET	3,081	Hong Kong	Amarutu Technology Ltd	21,630
United Kingdom	OVHcloud	1,816	Hong Kong	Tele Asia	3,391
United Kingdom	Trichan	655	Hong Kong	I-Services Network Solution Limited	257
Ukraine	TOV ITT	10,329	Germany	TerraTransit AG	4,616
Ukraine	PE Brezhnev Daniil	9,130	Germany	Koddos/Amarutu Technology Ltd. 2	2,502
Ukraine	ALFA TELECOM s.r.o.	5,996	Germany	imgsrc.ru	1,970
South Africa	Zappie Host LLC	8,742	France	Free.fr	25,551
South Africa	Afrihost	810	France	OVHcloud	25,428
Seychelles	IP Volume	33,957	France	Dedibox SAS	19,428
Seychelles	Incrediserve LTD	26,728	Estonia	Xemu	7,124
Seychelles	Novogara LTD	1,614	Estonia	Estro Web Services Private Limited	312
Russia	imgsrc.ru	17,999	Estonia	GmhostGrupp OU	192
Russia	ALFA TELECOM s.r.o.	11,118	Canada	Imagevenue	169,756
Russia	VDSINA Hosting	5,972	Canada	OVHcloud	6,949
New Zealand	Zappie Host LLC	8,047	Canada	Gayboystube	2,380
New Zealand	Spark New Zealand	96	Belize	Trichan	21,603
Netherlands	Trichan	717,722	Belize	TerraTransit AG	12,077
Netherlands	Incrediserve LTD	362,519	Belize	Koddos/Amarutu Technology Ltd. 2	855
Netherlands	NForce Entertainment B.V.	218,907			

Table does not necessarily reflect where an ESP's operations are based; rather it shows where the servers in use by the ESP are physically located.



Discussion

This report is not intended to provide an examination of jurisdictional issues related to cyber enforcement. However, there is value from a public policy perspective in understanding where content is being physically hosted globally — especially for countries with mandatory reporting requirements for ESPs.

An often cited obstacle by justice and law enforcement officials is the ambiguity over questions of jurisdiction when dealing with internet companies and their activities.

A useful and well-publicized example for illustrating the challenges of establishing whether an organization is subject to a country's laws is the case of adult content website **PornHub.com** and its parent company **MindGeek**⁷.

The parent company has a significant physical footprint in Montréal, Canada with approximately 1,000 employees working out of an office building. And while the company also has offices in Cyprus, England, Romania and the U.S., it claims to consider itself headquartered in Luxembourg where it is legally registered.⁷

However, based on the geolocation information tied to media detected by Project Arachnid, PornHub's content is mostly hosted on U.S.-based servers.

For this one company, jurisdictional arguments could be made for enforcing laws in some or all of Canada, Cyprus, England, Romania, the U.S. and Luxembourg. This underscores the importance for policymakers to establish clear parameters surrounding jurisdictional issues related to ESPs. This is a fundamental prerequisite for the creation of enforceable regulation in the digital space.

⁷ MindGeek. (2021, June2). *MindGeek*. <https://www.mindgeek.com/>
Protection of Privacy and Reputation on Platforms such as Pornhub, House of Commons Canada Standing Committee on Access to Information, Privacy and Ethics, 43^d Parliament, 2^d Session, Meeting 19. (2021). <http://www.ourcommons.ca/DocumentViewer/en/43-2/ETHI/meeting-19/evidence>

Content removal times

Defined as the total elapsed time from the moment of notification to when the targeted media is no longer accessible, content removal time is a crucial measure from a victim-centric perspective.

When considering the results presented in this section, it is important to recognize the calculated removal times presented in this report are based on when a notification was issued to an ESP. In reality, the media being targeted for removal were visible on the internet for an unknown amount of time prior to detection by Project Arachnid. So while the removal time upon notification is known to Project Arachnid, only the ESP knows how long the media were accessible on the internet.

Table 6.1 shows that over the three-year period studied in this report, from the moment a removal notice was issued, 50 percent of media targeted was no longer available the following day. While the median (50th percentile) removal time is 24 hours, the 90th percentile removal time globally was 42 days. Said otherwise, 10 percent of media targeted for removal over the past three years took seven weeks or longer from the point of notification before being removed from the internet. This is a core area of concern.

Significantly slower removal times in 2018 as shown in **Table 6.1** are believed to be the result of a few factors. At the time, Project Arachnid was in its early days of operation and the many challenges and strategies involved in successfully pursuing CSAM removal were not fully appreciated. In addition, Project Arachnid initially issued a portion of its removal notices to ESPs through intermediary regional-specific organizations. A move toward a direct-notice model in 2019 has proven to be a much more efficient and effective process.

However it must be noted that while removal times are down significantly since 2018, they have recently worsened with the 90th percentile removal time increasing from 26 days in 2019 to 38 days in 2020 (**Table 6.1**).

This increase in removal times is in part explained by the fact that 2020 saw an increase in post-pubescent CSAM images being targeted for removal (**Table 3.3**), a category with generally longer removal times compared to pre-pubescent CSAM and harmful-abusive content (**Table 6.2**).

The 90th percentile removal time for pre-pubescent CSAM is 40 days, while post-pubescent CSAM and harmful-abusive media is 56 days and 37 days respectively (**Table 6.2**).

Table 6.3 highlights major differences in removal times across ESPs. Certain ESPs whose platforms host significant amounts of CSAM have been largely unresponsive to removal efforts. The Trichan forums were highly unresponsive to removal requests which drove up the removal times (90th percentile = 138 days). However, when core hosting providers began withdrawing their services for the site, removal times plummeted as the content rapidly became inaccessible.

Table 6.3 also shows certain ESPs have relatively faster removal times. For some, this may be explained by an automated internal mechanism for processing Project Arachnid removal notices to hasten removal times.

Table 6.1

Removal times, all media targeted for removal				
	2018	2019	2020	All years
50 th percentile (median)	5 days	1 day	1 day	1 day
90 th percentile	161 days	26 days	38 days	42 days

Calculation includes elapsed time related to removal efforts initiated during the report period, but not yet concluded by Dec. 31, 2020.

Media is deemed removed when media is no longer accessible at the targeted URL.

Table 6.2

Removal times, by image category		
Image category	50 th percentile (Median)	90 th percentile
Pre-pubescent CSAM	1 day	40 days
Post-pubescent CSAM	2 days	56 days
Harmful-abusive	1 day	37 days

Calculation includes elapsed time related to removal efforts initiated during the report period, but not yet concluded by Dec. 31, 2020.

Media is deemed removed when media is no longer accessible at the targeted URL.

Table 6.3

Removal times, by ESP		
ESP name	50 th percentile (Median)	90 th percentile
Trichan	1 day	138 days
ColoCrossing	27 days	127 days
NForce Entertainment B.V.	8 days	70 days
Serverel	6 days	60 days
Incrediserve LTD	3 days	53 days
Liteserver Holding B.V.	1 day	43 days
FranTech Solutions	13 days	40 days
CloudFlare	1 day	27 days
OVHcloud	3 days	23 days
Host Sailor	1 day	15 days
ImgView.net	2 days	6 days
ImgMaze.com	2 days	6 days
ImgDew.com	2 days	6 days
ImgOutlet.com	2 days	4 days
ALFA TELECOM s.r.o.	1 day	4 days
DataWeb Global Group B.V.	1 day	2 days
Imagevenue	1 day	1 day
ImageBam	1 day	1 day

Calculation includes elapsed time related to removal efforts initiated during the report period, but not yet concluded by Dec. 31, 2020.

Media is deemed removed when media is no longer accessible at the targeted URL.



Discussion

In C3P's discussions with survivors we've learned that the recording of their child sexual abuse and its continued online availability creates an additional layer of trauma which colours every aspect of their lives. Simply knowing such material exist, and that individuals around the world are able to view and take pleasure from them their suffering, evokes a variety of emotions including fear, shame and a pervading sense of powerlessness. This is the fundamental reason why the prompt removal of harmful content is so critical.

The median removal time for content targeted by Project Arachnid is 24 hours. This finding, however, must be considered in the broader context of the problem. In isolation, this statistic is encouraging as it suggests Project Arachnid is an effective tool for achieving relatively prompt image removals for a significant portion of targeted media. However, it belies a core problem in this space: Many ESPs remove media within a day of notification, but in the absence of any regulatory requirements, they have no commercial or legal interest in investing in measures to prevent the images from surfacing or re-surfacing in the first place. There are no consequences for inaction on the prevention side. This is laid bare by the correspondingly high image recidivism rates reported later in the report.

The primary objective for ESPs ought to be to prevent these images from being uploaded onto their service in the first place. However, when this fails, the use of proactive media detection technology can assist with the prompt removal or blocking of previously known material.

That said, this report is especially focused on drawing attention to much longer removal delays highlighted with the 90th percentile removal times. The victims depicted in images that experience longer delays see greater levels of victimization.

The following table describes the spectrum of responses and behaviours C3P has experienced through its removal activities. It is worth noting that in many cases an ESP may be voluntarily proactive in detecting images but highly resistant to removal requests, meaning the following categories are not mutually exclusive.

Proactive	Companies that actively seek to detect and prevent CSAM from being posted on their service. This typically involves the larger technology companies but can include some smaller ones.
Reactive	Large and small companies that remove CSAM when notified but do not actively seek to prevent it on their service. Those that react to notices also have varying durations in removal time.
Resistant	Companies that debate/push back on removing material, either not being satisfied the image is of a child or not agreeing the image or video is illegal in nature.
Non-compliant	Companies that ignore takedown notifications or simply refuse to remove material that is clearly CSAM .
Complicit	Companies that knowingly allow CSAM on their services and may attempt to protect clients engaged in illegal activities.

Many inferences about changes in ESP behaviour can be made based on the trends observed in the data held by Project Arachnid. Improvements in media removal times are a key indicator for identifying companies that may have adopted a more proactive approach to online harm reduction.

Delays for post-pubescent media removal

As shown in **Table 6.2**, much longer delays in removal occur with post-pubescent CSAM in comparison to pre-pubescent CSAM.

Reasons for these longer delays may be explained by factors such as:

- A perception that post-pubescent CSAM is less serious, and therefore is not prioritized for removal;
- ESPs contesting the assertions the media in question constitutes CSAM;
- ESPs focused on the visual cues of the imagery alone may lack the context surrounding the potentially illegal nature of the media and therefore not be quick to remove.

For illustrative purposes, consider the following event experienced by C3P:

An ESP contested the accuracy of certain image categorizations via a series of email communications. Images of a nude 15-year-old girl were assessed by C3P analysts and linked to a known victim.

The ESP representative indicated they believed C3P's age assessments were incorrect as other online information suggested she was an adult. It had to be pointed out that while she may be an adult at the time the removal notices were issued, the images in question were taken years earlier when she was a minor.

The images were all subsequently removed, though one took nine days. This example underscores the senseless reality of operating in this space. In light of the existing media assessment backlog discussed earlier in this report, these types of one-off interactions with ESPs are a constant drain on limited resources and cause removal delays that are damaging to victims/survivors.

Media recidivism

Unlike evaluating removal times, which is a measure of an ESP's after-the-fact responsiveness to the presence of problematic content, media recidivism provides insight into the preventative measures (or lack thereof) employed by companies.

Forty-one percent (41%) of the 761 ESPs that have received at least one removal notice from Project Arachnid over the period covered by this report have had at least one image or video re-emerge on their service after it had been previously flagged for removal.

Table 7.1 shows that between 2018 and 2020, 48 percent of all media targeted for removal had been previously detected on the respective ESP's service. This same table shows recidivism rates have generally increased over the past three years. **Table 7.1** shows recidivism rates more than doubled from 20.7 percent in 2018 to 54.9 percent in 2020.

It is important to note that the calculated recidivism rates in this report are not necessarily comparable across ESPs since many factors influence the rate. For example, certain websites may be dominated by users who produce CSAM and are therefore more likely to upload previously unknown content, while the user base of other sites may re-upload the same content repeatedly. Under this scenario, the website with users who tend to post previously unknown content would result in a lower recidivism rate, since each media detection by Project Arachnid is more likely to be new material.

For example, an established website that has received many removal notices over the past years would yield a higher recidivism rate than a new website with the exact same media. This is because when Project Arachnid crawls the media on the new website, most removal notices would relate to previously undetected content for that ESP.

Table 7.2 reveals that recidivism rates are significantly higher for post-pubescent CSAM (73.1%) compared to pre-pubescent CSAM (46%) and harmful-abusive content (18%). This suggests that ESPs are more likely to have significantly delayed acting upon (or even ignored) previous Project Arachnid removal requests related to post-pubescent CSAM and/or have been less likely to add the digital fingerprints of this image category to their internal pool of banned images, assuming they actively maintain one to begin with.

Table 7.1

Recidivism rate, all media subject to removal request			
	Recidivist media	All media	Recidivism rate
2018	11,258	54,448	20.7%
2019	103,987	211,470	49.2%
2020	100,464	183,152	54.9%
All years	215,709	449,070	48.0%

Recidivism is established based on matching SHA-1 cryptographic hashes.

Table 7.2

Recidivism rate, by image category				
Image category	2018	2019	2020	All years
Pre-pubescent CSAM	21.7%	49.4%	50.1%	46.0%
Post-pubescent CSAM	23.8%	64.0%	77.4%	73.1%
Harmful-abusive	7.2%	14.1%	21.1%	18.0%

Recidivism is established based on matching SHA-1 cryptographic hashes.

It should also be noted that certain ESPs, while physically hosting the content on their servers, may not necessarily have a view into the material itself. A hosting provider whose customer offers an encrypted service, for example, might not be able to access or view the media itself. Under these circumstances, proactive media detection by the hosting provider is not possible.

Table 7.3 shows ESPs such as Czech Republic-based ALFA Telecom s.r.o., Serverel, Imagevenue and Netherlands-based LiteServer Holding B.V. had calculated recidivism rates exceeding 86 percent. In practice, this indicates these services have repeatedly hosted images that have been flagged multiple times by Project Arachnid.

Table 7.3

Recidivism rate for all media subject to removal notice, by ESP	
ESP name	Percent recidivism, all years
ALFA TELECOM s.r.o.	93.6%
Serverel	93.5%
Imagevenue	87.5%
Liteserver Holding B.V.	86.4%
Host Sailor	68.6%
FranTech Solutions	65.3%
CloudFlare	48.6%
ColoCrossing	35.5%
Incrediserve LTD	34.0%
Trichan	26.2%
OVHcloud	11.4%
DataWeb Global Group B.V.	11.4%
NForce Entertainment B.V.	5.9%
ImgDew.com	5.8%
ImgView.net	5.3%
ImgOutlet.com	4.9%
ImgMaze.com	4.5%
ImageBam	3.2%

Recidivism is established based on matching SHA-1 cryptographic hashes.



Discussion

Digital fingerprint comparisons are the main method for automatically detecting offending media, though not all ESPs use them. When a platform user uploads content onto an ESP's system, these automated technologies compare the unique digital fingerprint of the media to databanks of fingerprints that relate to previously identified CSAM or harmful-abusive media. If a match is found, the content is either blocked or removed. While this process is highly effective at curbing the distribution of known imagery, it cannot prevent newly created content from being uploaded.

If deployed properly, content administrators that employ digital fingerprint comparison technology and do not offer a fully encrypted service should see a limited amount of previously removed media re-emerging on their service.

Unfortunately, a large number of companies do not appear to be using these basic tools, with nearly 41 percent of ESPs having at least one recidivist image detected by Project Arachnid on the basis of a SHA-1 hash value match.

As noted in the previous section, some ESPs may physically be in possession of content on their servers but not necessarily have the ability to view or directly detect media that is managed by their customer. Under these circumstances, certain hosting providers have told C3P that proactive detection is not technically possible, and therefore they cannot prevent repeat media from re-emerging on their servers using this approach.

It is C3P's view that while certain technical limitations do exist, such as those described above, nothing prevents ESPs from adopting contract-based solutions that impose legal requirements on their customers as a condition of service. Such requirements could include a requirement that customers use specified media detection technologies, block file uploads from the dark web, maintain certain levels of human moderation and remove media within a specified time from detection or notification.

These types of practical solutions are described in greater detail later in the report as part of a list of recommendations.

SHA-1 vs. PhotoDNA image matching

A SHA-1 image match is a match based on a unique digital cryptographic hash value. This means the digital file, down to the binary level, exactly matches another image. Any modification to the image or to the file's underlying metadata will result in a different SHA-1 hash value, and therefore no longer match with the previous version of the image. Examples of modifications to an image that yield different SHA-1 hash values include:

- Modifications to the colour;
- Resizing the image;
- Saving in a different file format;
- Modifying the metadata (i.e., Exif data);
- Removing or adding a single pixel;
- Taking a screenshot of an existing image.

In practice, many images detected by Project Arachnid are slight derivatives of previously verified images, but the differences may be imperceptible to the human eye. Despite being visually the same image, their unique digital SHA-1 fingerprints are different for reasons explained above.

In order to match these non-identical images, approximate image matching technology can be used such as Microsoft's PhotoDNA algorithm, a popular tool used in this space. While Project Arachnid does use PhotoDNA as part of its operations, image recidivism in this report is tied to exact SHA-1 matches, and does not take into account variants of what appear visually as the same image. This means that in practice, the recidivism rates presented in this report are likely to be very conservative figures.

Post-pubescent recidivism

As noted in **Table 8.2**, post-pubescent CSAM has significantly higher recidivism rates when compared to pre-pubescent CSAM and harmful-abusive content.

Some of the same factors discussed earlier regarding possible reasons why this class of media experiences longer removal times may also be driving the higher rate of recidivism. It is also important to note that in addition to the preventative actions that can be taken by an ESP, image recidivism rates are also driven in part by the nature of the content uploaded by an ESP's user base. For example, a review of a collection of newly produced content would yield relatively low recidivism rates compared to a collection of popularly traded historical media.

Unlike pre-pubescent CSAM, post-pubescent CSAM is often found intermixed with adult pornography or on otherwise popular platforms with large user bases that permit adult pornography (such as Twitter). This may indicate individuals uploading post-pubescent media are satisfied the apparent age of the individual depicted in the images is uncertain enough to provide plausible deniability should they be questioned. It may also indicate a lack of understanding about what constitutes CSAM and the consequences of distributing this type of imagery.

Combined, these factors may lead to a perception by content administrators and users that the uploading (and future re-uploading) of post-pubescent CSAM is a relatively low-risk endeavour.

Building upon the above points, in the absence of information suggesting that certain sexual images depict minors, website administrators may opt to ignore removal notices and to not add flagged images of post-pubescent victims to their banned media lists.

This resulting attitude toward post-pubescent CSAM could explain in part higher volumes of image recidivism for this category.



CASE STUDY:

French telecommunications company Free: Project Arachnid's largest source of detected CSAM

Over the past three years, Project Arachnid has issued removal notices to more than 760 ESPs. The records gathered as a result of those interactions clearly demonstrate certain ESPs directly or indirectly contribute in more significant ways to the distribution of CSAM and harmful-abusive content on the internet.

Project Arachnid records on actionable media detections, moreover, point to a single ESP whose service has been used for hosting and sharing a very significant volume of CSAM and harmful-abusive content: French telecommunications giant Free, owned by the Paris-based parent company Iliad Group.

From 2018 to 2020, Project Arachnid detected more than 18,000 archive files, collectively containing nearly 1.1 million image or video files of apparent CSAM or harmful-abusive content. These were, or in some cases continue to be, hosted directly on Free's public file hosting service.

In many cases, Project Arachnid's web crawler has detected links to these archived files across many areas of both the clear web and Tor sites. Given these many access points to the media archives, the total known availability of CSAM and harmful-abusive images or videos on Free's hosting service is more than 2.7 million media detections.

Past controversy over Free's hosting service

The file hosting service – found at the address dl.free.fr – came under fire in October 2007 when France's then-Minister of Culture singled out the company, alleging they were failing to take adequate steps to curb the illegal distribution of copyrighted material.

In an official statement, the minister expressed concerns over recent changes made to the company's file sharing service that permitted, "users to anonymously upload large volumes of pirated material on dl.free.fr."⁸

The statement also asked that then-Chief Executive Officer of parent company Iliad Group make use of Free's "undisputed technical mastery" to enforce existing laws by either highly restricting the use of the file hosting service or by outright deleting it.

The minister's statement also reiterated the details of a recent court decision that required Free to block access to 14 binary newsgroups, a type of decentralized online community forum often used for sharing media between users. This came after several groups of rights holders voiced their concerns over the service. Ultimately the French government went as far as to threaten to withhold the company's 3G license application over the file hosting service.

⁸ Ministère de la Culture. (2007, October 12). *Christine Albanel demande à Free de lutter plus activement contre le piratage*. <http://www2.culture.gouv.fr/culture/actualites/communiqu/albanel/free07.html>

Details of the hosting service

Cached versions of the company's free file hosting web page can be found as far back as early 2006, around the time the service was experimentally launched.⁹

The file hosting service offered through the company's website appears to be popular among individuals intent on distributing large volumes of media anonymously, with discussions on dark web forums recommending Free's file hosting service for CSAM distribution.

Possible reasons behind the service's popularity include:

- The service requires no account, registration, contact information or payment to upload content and generate a link to the media that can then be shared with anyone, anywhere;
- Despite the minimalist design of the service, it provides a very generous file size limit, allowing for large media collections to be uploaded and distributed;
- An uploader can password protect an archive file, meaning only a recipient with knowledge of the password can access the media. These links and passwords are often found on the dark web.

Of note is the fact that the specific web page users access to upload content on Free's service is outdated and has not been graphically updated since 2008.¹⁰ In addition, the company's main site uses secure hypertext transfer protocol (HTTPS) – a standard in modern websites, but the image hosting portal uses the outdated unsecured hypertext transfer protocol (HTTP).

In addition, as of May 18, 2021, when on the file hosting service web page, the link provided to users for reporting "illicit content" leads to a default 404 error page, meaning the requested web page does not exist.

This all suggests the company has not prioritized or given this service much consideration in recent years.

How Free's file hosting service is misused for CSAM distribution

Individuals intent on distributing CSAM and harmful-abusive content have taken advantage of Free's hosting service to anonymously store media online, and then disseminate the direct download link on forums across the internet.

Based on Project Arachnid's crawling records, a significant number of dl.free.fr download links are detected on Tor-based discussion forums where an unknown number of anonymous users are able to view the direct download links and passwords required to unlock and access the content.

Rather than viewing content ephemerally embedded on a web page, users must download the media, generating new copies of the files on their local computers. In this context, even if the media hosted at the source is eventually removed, several other privately held copies are likely to exist and may very well re-emerge at a later date on the internet.

Communications with Free representatives

Beginning in 2018, C3P began corresponding directly with company officials, providing them with lists of direct links to the file archives containing CSAM being hosted on their system.

Project Arachnid has continued to detect and issue notices on newly uncovered CSAM and harmful-abusive media to the company. As of May 18, 2021, nearly 3,000 archives for which removal notices were issued between 2018 and 2020 (inclusively) were still publicly accessible, according to Project Arachnid records.

⁹ Free. (2006). *Conditions d'utilisation du nouveau service experimental* <http://dl.free.fr>. <https://web.archive.org/web/20060126211229/http://dl.free.fr/>

¹⁰ Free. (2008). *Service d'envoi de fichiers* <http://dl.free.fr> <https://web.archive.org/web/20081106103820/http://www.dl.free.fr/>



CASE STUDY:

Project Arachnid Trichan imageboard campaign

Another prominent host of CSAM encountered by Project Arachnid was a collection of internet forums known as the Trichan imageboards. The now-defunct forums were primarily dedicated to the sexual exploitation of children and had been in operation for at least seven years, largely hosted out of the Netherlands.

First encounter

In March 2019, Project Arachnid encountered a sharp rise in detected media on these forums. The sudden surge in volume was such that C3P's internal technology experts had to retool the crawler to manage the flow of data being detected.

As the system began issuing takedown notices, it quickly became apparent the owner of the Trichan sites was generally not prepared to take action. Despite repeated emails to the websites' contact addresses, the vast majority of actioned media persisted on the site. Eventually, notices sent by email simply bounced back as undelivered.

Despite issuing numerous unsuccessful requests for removal, Project Arachnid continued to gather records on detected CSAM on the website. Internally, C3P analysts began assessing the images and concluded that, based on a sample of 51,917 images, it was likely that nearly 34 percent of images on the website were CSAM, with the remainder being possible CSAM and harmful-abusive images.

Armed with this data, and in light of the fact the Trichan administrator essentially ignored Project Arachnid notifications, C3P approached the service's upstream service providers in an attempt to have the content removal requests acted upon.

Initially, this approach was met with a great deal of resistance. Some ESPs further up the chain suggested C3P redirect its removal notices to the Netherlands-based hotline. After much back and forth with the Trichan upstream service providers, some of the providers took action to null-route (block) the IP addresses of the Trichan sites.

Over the subsequent weeks, the Trichan forums repeatedly changed hosting providers. With each change, C3P would engage with the new host and present them with data regarding the nature of the content hosted within the Trichan forums.

Deploying evasive techniques

Eventually, the websites went offline for a period of about three months. However, when the forums later resurfaced, it was quickly determined that the Trichans appeared to be employing new techniques to disrupt the automated detection of CSAM on their service.

The technique involved automatically injecting "noise" into the images by randomly offsetting pixels in ways that are imperceptible to the human eye. The forums were then set to cycle through modified versions of the same image when a new website visitor loaded the content.

This tactic made Project Arachnid comparisons of hash values against libraries of previously actioned hash values more challenging. However, with the use of approximate image matching technology such as Microsoft's PhotoDNA algorithm, Project Arachnid successfully continued to identify images and issue removal notices.



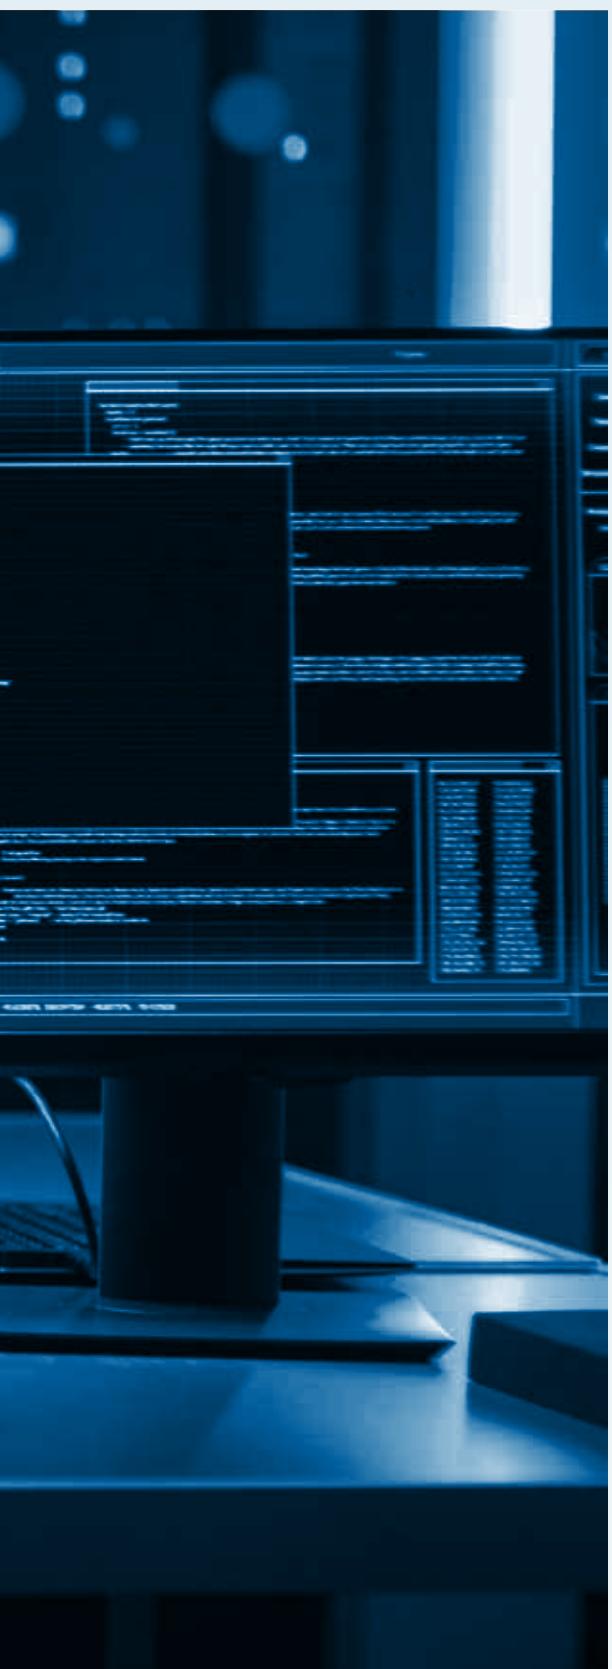
These two images look identical to the human eye, but they represent two completely different digital signatures. The red blotches represent collections of pixels that have been slightly offset from one image to the next.

Once again, C3P communicated with the latest hosting provider, and the website was taken offline thereafter. Seemingly out of options for finding a new hosting provider that would tolerate the nature of their content, the Trichan site administrator eventually capitulated. In an announcement posted on their homepage a site moderator bemoaned the relentless efforts to have their content removed before stating the site would be shutting down permanently:

"It's been a wonderful seven years and we would've loved to go for another seven, but antis are hunting us to death with unprecedented zeal, and after being shut down more than two dozen times and serving more than 100,000 brothers from all over the world daily we don't have the finances go on any longer."

Over approximately two and a half years, Project Arachnid detected more than 1.5 million verified media on the Trichan forums.





Discussion

This case study highlights the real-world barriers faced by organizations that operate in the CSAM removal space. It also foreshadows the many challenges governments seeking to introduce a regulatory framework for internet-based content will inevitably encounter.

As noted in a recent study¹¹ focused on the Trichan campaign, the forums' continued operation was made possible by a hierarchy of internet companies, some of which were unaware of the CSAM, and others of which were seemingly unconcerned by the fact their clients made CSAM publicly available.

The resistance of certain upstream ESPs to take action, even when presented with evidence of rampant CSAM problems on a client's service, is a key issue policymakers must address.

C3P was ultimately successful in its efforts to have the CSAM removed by contacting the upstream service providers of the Trichan forums. As Salter and Richardson (2021) note, this intervention highlights the effectiveness of focusing on power relations between ESPs and coordinating with influential nodes within the overall digital network.

The data collected by Project Arachnid enables C3P to largely map out the relationships between the higher and lower order ESPs and how the actions and inactions of these parties can directly affect the existence of CSAM and harmful-abusive content on the internet.

This case study highlights that, "providers of internet transit and other key services are revealed to be routinely entering into commercial arrangements with service providers and clients involved in abuse material."¹² And while these commercial arrangements are central to the distribution of CSAM on the internet, Salter and Richardson's report notes there exists no legal obligation for ESPs to deny service to a customer engaging in these abusive and possibly illegal activities.

11 Salter, M., & Richardson, L. (2021). The Trichan takedown: Lessons in the governance and regulation of child sexual abuse material. *Policy & Internet*, 13(2). Advance online publication. <https://doi.org/10.1002/poi3.256>

12 Salter, M., & Richardson, L. (2021). The Trichan takedown: Lessons in the governance and regulation of child sexual abuse material. *Policy & Internet*, 13(2). Advance online publication. <https://doi.org/10.1002/poi3.256>

RECOMMENDATIONS

The analysis contained in this report highlights several key issues that warrant close attention and immediate action by both ESPs and policymakers.

The findings also suggest that relying upon ESPs to voluntarily invest adequate resources in content moderation and adopting a vision that prioritizes the safety and privacy of children is simply not working.

Other signs of failures in this space are the lack of comprehensive reporting requirements across jurisdictions, the patchwork of moderation measures employed by companies and the deluge of victims and survivors coming forward about their struggles to have their abusive material removed.

Over the course of its content removal operations, C3P has gathered significant information through survivor surveys, processing tips from the public and the operation of Project Arachnid that highlights weaknesses of wholly inadequate regulatory environments. This real-world insight puts our organization in a unique position to provide recommendations on a regulatory response that will have the best possible outcomes for children.

The following list of recommendations are rooted in C3P's extensive experience in reducing the availability of CSAM and harmful-abusive content on the internet. Policymakers should view these as critical components in the development of effective regulation of ESPs as it relates to the online protection of children.

RECOMMENDATION 1:

Enact and impose a duty of care, along with financial penalties for non-compliance or failure to fulfill a required duty of care

ESPs that do not comply with regulatory requirements or fail to prioritize the safety of children online must face financial penalties, proportionate to the level of harm.

Penalties should factor in, at minimum:

- The volume of content;
- The number of users who viewed the media;
- The number of times the content was re-published (i.e., shared);
- Delays in removal time;
- The severity of the content;
- Number, ages and visibility of victims depicted in the content.

In addition, once notified of problematic content, upstream ESPs must be held financially accountable for media distributed by their downstream clients who may be in violation of regulatory requirements.

RECOMMENDATION 2:

Impose certain legal duties on upstream electronic service providers and their downstream customers

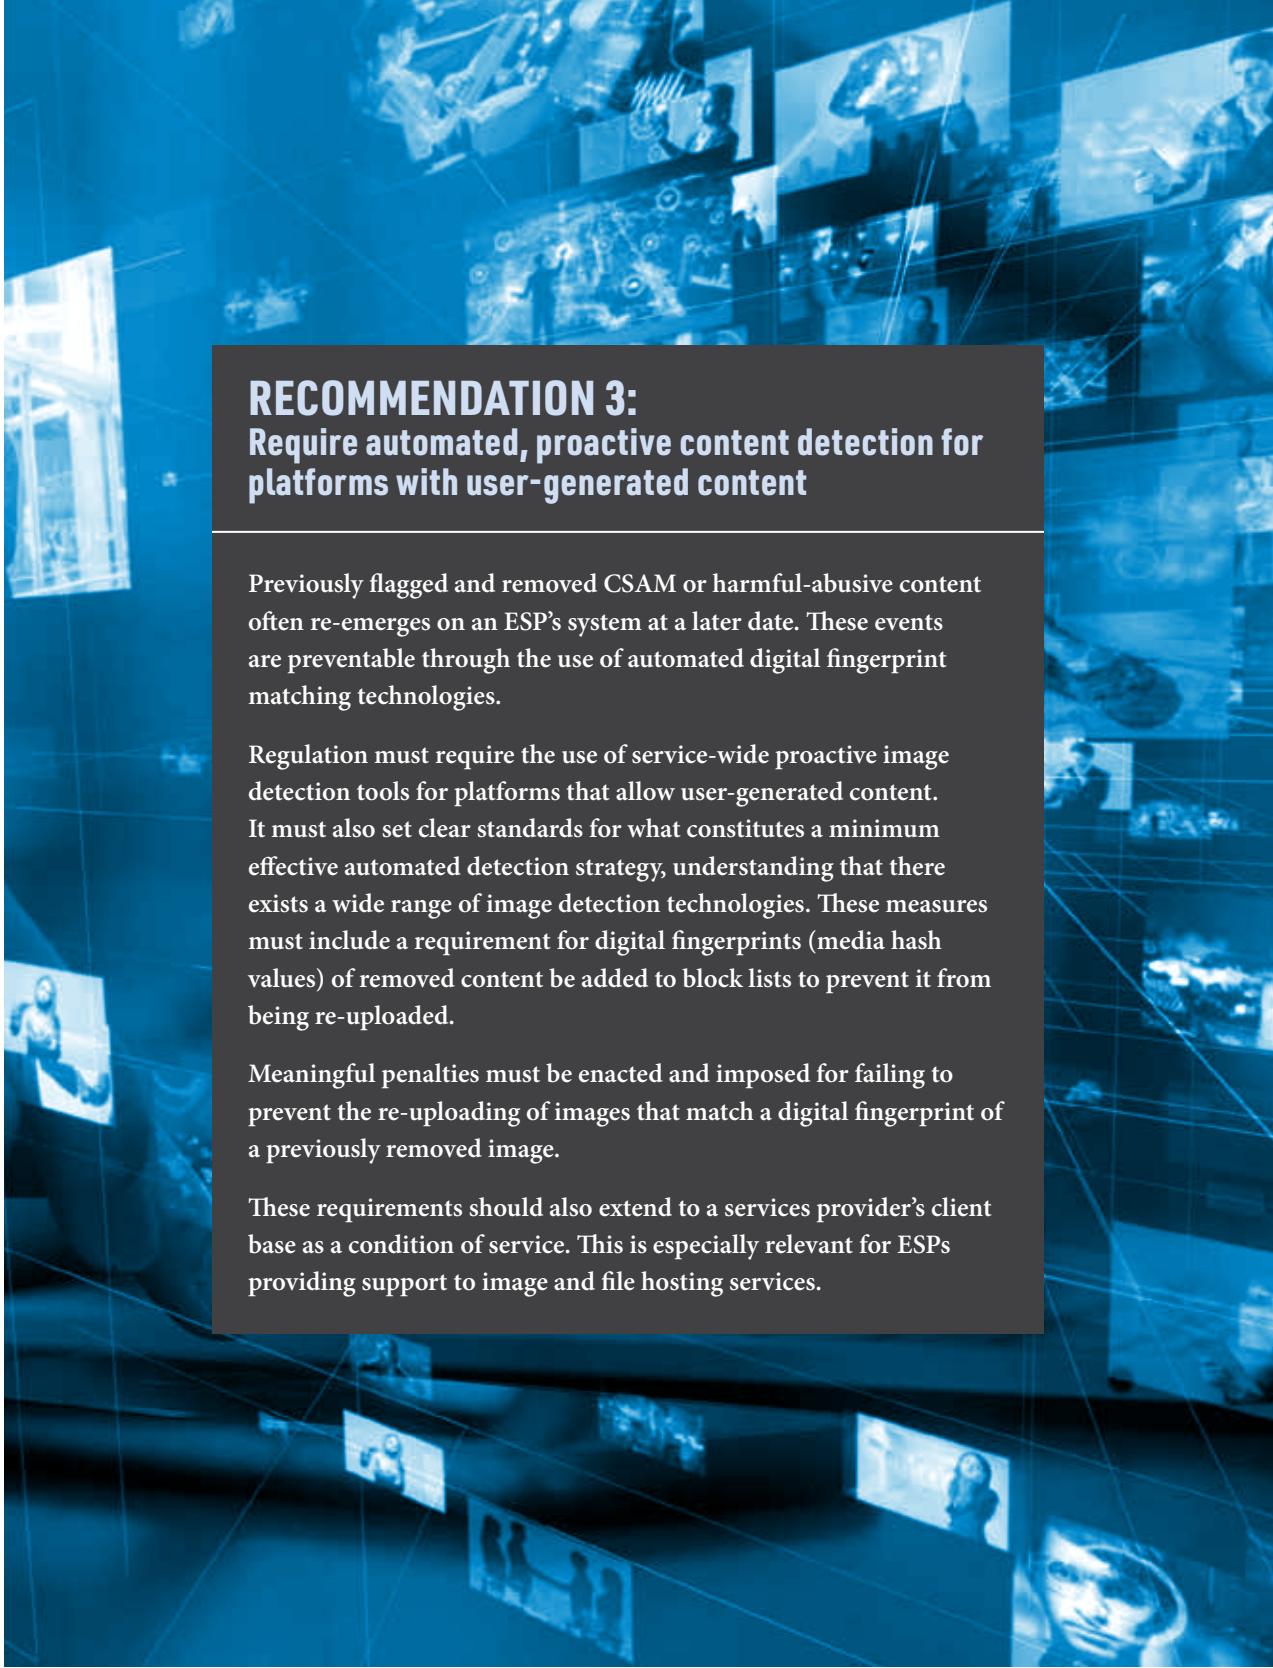
The operations of the internet traverse multiple jurisdictions and there are little to no coordinated regulatory or legislative requirements for internet based companies. Instead, the digital ecosystem is primarily structured through a myriad of complex and interrelated contracts made between various entities, each of which may be in different legal jurisdictions and have different tolerance levels for illegal content.

All of the companies bound by these contractual arrangements are necessary to make content ultimately accessible to an end user. As a result, to address a particular problem, every entity within the system must be bound by enforceable contractual terms that address the problem and also be required to impose and enforce similar contractual terms against its own customers. If any entity in the chain is not bound by such terms, or is not willing or able to enforce its own terms against its customers, that gap can be exploited thereby enabling the problem to flourish.

Similar to the way in which many nations have adopted legislative and regulatory control to ensure consumer protection in the areas of insurance, sale of goods and privacy issues, so too must they establish a framework to manage what the internet has become. Legislative and regulatory requirements that set out minimum base standards that are non-negotiable must be enacted. Each and every provider in the chain must be legally obligated to adhere to those base standards within their own operations, and to impose those same standards on their customers. Each ESP must be capable of being linked to at least one real person and nations must put an end to the endless legal loopholes that have enabled companies to evade legal liability for the harm they are facilitating by coordinating rules across jurisdictions.

The rules must apply, at a minimum, to those that provide image or file hosting services and include at least the following elements:

- Prescribed definitions and removal requirements for CSAM and harmful-abusive content;
- Required accountability measures to be taken by the provider in the event of illegal or harmful-abusive content being hosted by the customer of the provider;
- Significant and meaningful liability/penalties for any provider that fails to take certain actions when its customer violates the removal requirements.



RECOMMENDATION 3:

Require automated, proactive content detection for platforms with user-generated content

Previously flagged and removed CSAM or harmful-abusive content often re-emerges on an ESP's system at a later date. These events are preventable through the use of automated digital fingerprint matching technologies.

Regulation must require the use of service-wide proactive image detection tools for platforms that allow user-generated content. It must also set clear standards for what constitutes a minimum effective automated detection strategy, understanding that there exists a wide range of image detection technologies. These measures must include a requirement for digital fingerprints (media hash values) of removed content be added to block lists to prevent it from being re-uploaded.

Meaningful penalties must be enacted and imposed for failing to prevent the re-uploading of images that match a digital fingerprint of a previously removed image.

These requirements should also extend to a services provider's client base as a condition of service. This is especially relevant for ESPs providing support to image and file hosting services.

RECOMMENDATION 4:

Set standards for content that may not be criminal, but remains harmful-abusive to minors

There are fundamental problems with using, in isolation, criminal law definitions of child sexual abuse images to determine what images/videos should be removed from public view. When those restrictive definitions form the basis of a regulatory framework, a significant proportion of images that are harmful-abusive to children are left to propagate online.

Some examples of harmful-abusive content that may not meet a criminal law definition of CSAM in all jurisdictions:

- A series of images, some of which were taken prior to or after the act of abuse was recorded;
- Images of children in bathing suits distributed on forums dedicated to sexualizing children;
- Images of children urinating;
- Imagery depicting clothed or semi-clothed children in provocative poses, sometimes inaccurately labelled as “child modelling”;
- Images of children being physically assaulted or tortured;
- Information related to grooming and/or abuse tactics;
- Written content describing or advocating/counselling child sexual abuse;
- Sexual commentary related to an image or video of a child;
- Releasing of personal information about a child.

Regulation must clearly define and capture this type of material and include it under the definition of CSAM or child abuse as part of any broader child protection regulatory framework or initiative.

RECOMMENDATION 5: **Mandate human content moderation standards**

Automated proactive detection relies on comparing incoming media to databanks of previously removed content. This technology is therefore ineffective against newly created or previously unknown content, since there are no comparative images against which a match can be made.

Human moderation is therefore a critical component of a platform's defenses against CSAM and harmful-abusive content when user-generated content is accepted.

Regulation must establish a clear set of expectations related to:

- The proper supervision of content moderation teams;
- Frequent moderator training, including education related to sexual maturation assessment;
- Standards for staffing levels given a service's incoming content volume.

Regulation must also establish requirements that all user-generated content on platforms that allow pornography or nudity as part of their terms of service be manually reviewed prior to publication.

Critically, moderation practices must correspond with overall regulatory framework definitions of CSAM and harmful-abusive content.



RECOMMENDATION 6:

Set requirements for proof of subject or participant consent and uploader verification

Platforms that lack moderation and allow content uploaded by anonymous users are often exploited for the distribution of CSAM and harmful-abusive content over time.

ESPs that allow user-generated content — especially those that focus on, or partially cater to, adult pornographic content and nudity — are at greater risk of intersecting with CSAM and harmful-abusive material.

A regulatory framework related to user verification and consent must:

- Set clear standards for verification requirements for content uploaders that are appropriate given the risk level of the site;
- Define what constitutes verification and set storage, access and disclosure requirements for those verification records;
- In the case of pornographic or sensitive content, set clear requirements for establishing the age of the subjects appearing in the image or video;
- In the case of pornographic or sensitive content, set clear requirements for establishing that all subjects consented to the recorded acts and also consent to the distribution of the content.

RECOMMENDATION 7: Establish platform design standards that reduce risk and promote safety

In addition to proactive and reactive moderation measures, platforms must further reduce the prevalence of CSAM or harmful-abusive content by cultivating an environment that discourages users from exploiting their service for this purpose.

Regulation should establish requirements for:

- Prohibiting user-generated content where the uploader originates from an IP address associated with a Tor exit node, VPN service or other IP concealment techniques;
- Blocking search terms and forum/chat names that are associated with CSAM or harmful-abusive content;
- Removing or suspending accounts that distribute or access CSAM or harmful-abusive content;
- Segregating children and adults in the digital space by design. When not feasible, additional rules and protections must be implemented;
- Requiring platforms to provide an easily accessible and responsive mechanism for users to contact content administrators for lodging complaints;
- Measures, such as user age verification, for preventing children from accessing adult or mature content.

RECOMMENDATION 8:

Establish standards for user-reporting mechanisms and content removal obligations

Moderation practices may not always successfully detect CSAM or harmful-abusive content. For this reason, ESPs must have user interfaces designed to facilitate content reporting and complaint submissions, paired with specific removal requirements.

Regulation should establish clear standards that include:

- A requirement that all content types (e.g., images, videos, users, web pages, comments, posts, etc.) be directly reportable;
- Clear and unambiguous issue-specific reporting categories — including for CSAM — to ensure higher-risk content can be prioritized for review;
- Specifically in the case of reported CSAM or harmful-abusive images, a requirement that flagged content be automatically suspended/made unavailable until it can be assessed, rather than allowing the media to remain online pending review;
- Prescribed assessment and removal times for content upon receiving a complaint;
- Record retention requirements related to the image, uploader, communications with the complainant and any actions taken related to complaints;
- Mandatory reporting of actioned content to a specified authority or tipline, including transparency requirements about removal/non-removal decisions.

CONCLUSION

Many internet companies are failing to prioritize the safety and privacy of children online. A digital ecosystem enabled by jurisdictional uncertainty, along with a lack of clear regulation or transparency, has significantly contributed to the proliferation of CSAM and harmful-abusive content on the internet.

The findings contained in this report, which is based on three years of data collected by Project Arachnid, analyzed details on 5.4 million images or videos of CSAM and harmful-abusive content related to more than 760 ESPs.

The report established there exist high levels of image recidivism and often long delays in removal times for many internet companies. This suggests many ESPs are not deploying sufficient resources to ensure their platforms are free of, or dramatically limit, the presence of CSAM and harmful-abusive content on their services.

Other key insights discussed in this report include:

- The role the dark web plays in facilitating access to CSAM on the clear web;
- How a relative few ESPs can have a significant impact on the availability of CSAM on the internet;
- Why statistics related to adolescent victims dramatically underrepresent the true scale of harm they experience;
- The central role lesser-known ESPs play in making CSAM and harmful-abusive content available on the internet;
- The importance of considering the broader chain of ESPs that facilitate the availability of CSAM on the internet.

The report strongly suggests expecting industry to voluntarily invest in resources to prevent the spread of CSAM and harmful-abusive content has been a failure. It points to a pressing need for consistent, enforceable and global standards that impose accountability requirements on ESPs.

Flowing from the findings, a set of eight key evidence-based recommendations are presented for governments seeking to reduce the availability and distribution of CSAM on the internet, and to adopt measures that prioritize the safety of children.

This report is both a road map and an opportunity to properly extend the duty of care we owe to children in the online world.





APPENDIX

Glossary of terms

Assessed media

A term describing media that has been assessed by an analyst. Assessed media is not necessarily CSAM or abusive-harmful content.

Child/children

Any person under the age of 18.

Content

Refers to any media.

Content delivery network (CDN)

Refers to a network of servers that are geographically dispersed to enable faster web performance by locating copies of web content closer to users. These services typically mask details of the underlying hosting provider information for a website making use of CDN services.

Clear web

The clear web (sometimes referred to as the “Clearnet”) refers to the publicly accessible internet whose web pages are largely indexed on search engines.

Content administrator

Refers to websites or web-based services. With the exception of large-scale ESPs, most content administrators do not own or operate their own physical servers. Websites that provide individual file hosting services typically fall under this category.

Dark web

A catch-all term to refer to the series of networks not viewable using a standard web browser. These networks, which include Tor, are generally configured to encrypt internet traffic and provide anonymity and privacy for users.

Detections

Refers to the discovery, or “sightings”, of media on the internet by Project Arachnid. Detections serve to measure the availability of media. A single hosted image that is embedded (and therefore visible) on several websites would result in multiple media detections if encountered by Project Arachnid’s crawl.

Electronic Service Provider (ESP)

A catch-all term to reference any entity that provides a service in the digital space, including content delivery networks, hosting providers, cloud service providers, content/website administrators, internet service providers, etc.

Exif data

Exchangeable image file format (Exif) is a standard that defines specific information (metadata) related to an image or other media captured by a digital camera. This can include, but is not limited to, information on the creation date, the image aspect ratio, the resolution, the location the image was taken.

Harmful-abusive

An image category that refers to images depicting children that does not appear to meet a criminal law threshold across multiple jurisdictions, but may nonetheless violate an ESP's terms of service. These images may also violate the privacy or safety of a child, or be associated with CSAM. For more details refer to the description of C3Ps framework (p. 10).

Hash value

A digital fingerprint (or signature) that uniquely identifies a computer file. Hash values are derived by computer algorithms.

Hosting provider

Refers to a business that provides the technologies and services needed for a website or web page to be accessible and viewable on the internet. Websites are hosted (or stored) on servers operated by hosting providers.

Media

Refers to all content types processed by Project Arachnid. Typically, this refers to images, videos and/or archive files (that contain images or videos).

Minor

Any person under the age of 18.

PhotoDNA

PhotoDNA is an image comparison technology used for detecting matches between modified versions of the same image or images with similar features. It is sometimes referred to as "fuzzy matching" or "perceptual hashing". The technology was developed by Microsoft in partnership with Dartmouth College.

Post-pubescent CSAM

An image category that refers to content that likely meets a criminal definition of CSAM. It includes images where the depicted victim's age has been confirmed and is post-pubescent. This category also includes media containing victims that are in the later stages of puberty.

Pre-pubescent CSAM

An image category that refers to content that likely meets a criminal definition of CSAM. It includes images where the depicted victim is pre-pubescent or is in the early stages of puberty.

Removal notice

The process by which C3P's Project Arachnid notifies an ESP of the presence of CSAM or harmful-abusive content on their servers and requests the removal of the media in question.

SHA-1

Stands for "Secure Hashing Algorithm". It is a specific cryptographic hash value assigned to media processed by Project Arachnid to assign a unique digital signature to content using a computer algorithm.

Suspect media

Refers to any media that is reasonably believed to be CSAM, but has not been through the assessment process.

The Onion Router (Tor)

Tor, short for "The Onion Router", is an open source privacy network that permits users to browse the web anonymously. Tor is generally considered a subset of what is commonly referred to as the dark web.

URL

Short for “universal resource locator”, a URL refers to the direct path address to a web page or media on the internet.

Verified media

A term describing media that an analyst assessed and evaluated as being either CSAM or harmful-abusive content.

Virtual private network (VPN)

A private network in which two end points create a single, private connection, or tunnel, while using a larger network infrastructure such as the internet or wide area network. Several ESPs offer commercial VPN services for users.

List of acronyms

C3P: Canadian Centre for Child Protection

CSAM: Child sexual abuse material

ECPAT: End Child Prostitution and Trafficking

ESP: Electronic service provider

Exif: Exchangeable image file format

HTTP: Hypertext transfer protocol

HTTPS: Hypertext transfer protocol – secure

Interpol: International Criminal Police Organization

NCMEC: National Center for Missing & Exploited Children (U.S.)

SHA1: Secure Hash Algorithm 1

Tor: The Onion Router

URL: Uniform Resource Locator

VPN: Virtual private network





CANADIAN CENTRE for CHILD PROTECTION®
Helping families. Protecting children.

 protectchildren.ca

 [@CdnChildProtect](https://twitter.com/CdnChildProtect)

 Canadian Centre for Child Protection

 [@cdnchildprotect](https://www.instagram.com/cdnchildprotect)