

Just in Plain Sight: Unveiling CSAM Distribution Campaigns on the Clear Web

Nikolaos Lykousas
Data Centric
Romania

Constantinos Patsakis, *Senior Member, IEEE*
Department of Informatics
University of Piraeus, Piraeus, Greece
Athena Research Center
Artemidos 6, Marousi, Greece
kpatsak@unipi.gr

Abstract—Child sexual abuse is among the most hideous crimes, yet, after the COVID-19 pandemic, there is a huge surge in the distribution of child sexual abuse material (CSAM). Traditionally, the exchange of such material is performed on the dark web, as it provides many privacy guarantees that facilitate illicit trades. However, the introduction of end-to-end encryption platforms has brought it to the deep web. In this work, we report our findings for a campaign of spreading child sexual abuse material on the clear web. The campaign utilized at least 1,026 web pages for at least 738,286 registered users. Our analysis details the operation of such a campaign, showcasing how social networks are abused and the role of bots, but also the bypasses that are used. Going a step further and exploiting operational faults in the campaign, we gain insight into the demand for such content, as well as the dynamics of the user network that supports it.

Index Terms—Child sexual abuse material (CSAM), illicit campaigns, criminal network analysis, electronic crime

I. INTRODUCTION

Traditionally, communities of perpetrators producing, consuming, and trading child sexual abuse material (CSAM) have been studied primarily within the context of the dark web. The veil of anonymity, along with the lack of censorship and inherent confidentiality of dark web services, facilitates, among others, such illicit exchanges [1], [2], [3]. During the COVID-19 pandemic, CSAM distribution has increased significantly [4]; however, this period coincided with the widespread adoption of messengers supporting end-to-end encryption. The latter is a great enabler for distributing such content, as the material can be shared with a far wider audience without the need for special software. Moreover, it facilitates the exchange and consumption of such content from mobile devices whose use has become ubiquitous.

It should be emphasized that many such communities follow pyramid schemes that follow a hierarchy based on content ownership. Thus, many members are requested to provide new content to be allowed. On the other hand, many of the members simultaneously seek to engage and recruit more users by spamming networks with CSAM. This behavior can be attributed to their need to normalize the consumption of such

content, but also to attract more people to the group who will generate more CSAM that they will consume. As a result, there are often reports about CSAM being distributed via traditional social networks and, of course, users grooming minors over various online platforms.

Unfortunately, the problem of CSAM distribution has reached unprecedented levels. According to the latest report from INHOPE [5], not only has the reported suspicious content skyrocketed, but the illegal content has actually tripled that of the previous year, as shown in Figure 1. Moreover, as can be observed in the figure, there is a steady influx of new content, meaning that more children are subjected to this torture. Most of the content reported to INHOPE is shared in forums (61.07%), followed by image hosting services (20.86%), websites 14.86%, and file hosting services (2.03%).

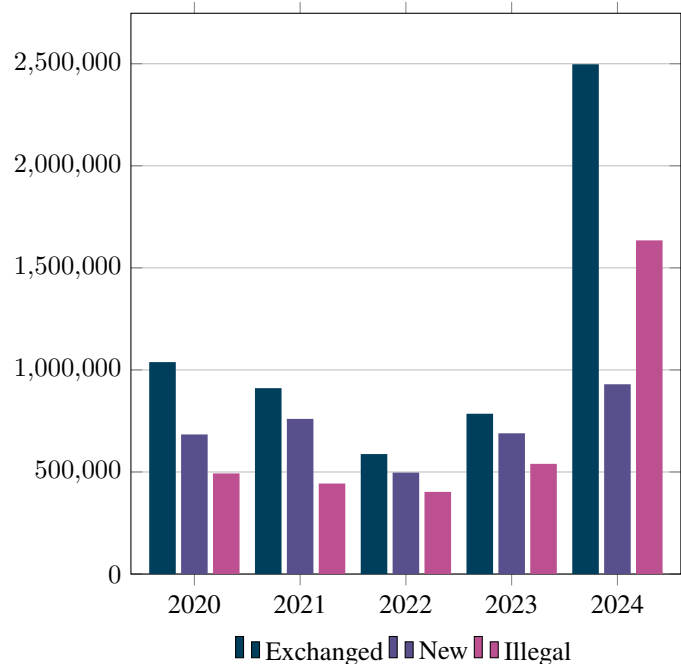


Fig. 1: Reported CSAM record growth according to INHOPE [5].

shocking [6]. Around 1.8 million users worldwide had logged onto the platform between April 2022 and March 2025. When seized, the server contained approximately 72,000 videos, while the platform had around 91,000 videos with an estimated total duration of 6,288 hours.

Contrary to most studies in the literature related to CSAM distribution, which study the availability of such content, the distribution methods, as well as the geographic distribution of the users [7], we study a wide online campaign on the clear web that used 1,026 unique domains. Our study goes beyond statistics by providing insight into the modus operandi of this campaign, detailing how and where the content is distributed. Moreover, we shed light on the dynamics of the referral network, reporting the inviter/invitee dynamics and some characteristics of the CSAM users, who worldwide are on the scale of 738,286. Despite its content and the fact that it is operated on the clear web, it has been ongoing since 2022. It abuses various platforms and is orchestrated through Telegram and other clear web channels to communicate with potential users and buyers, but also to share part of the material. In compliance with ethical standards and legal frameworks, we highlight that we did not access, view, or download any CSAM. Our analysis was limited to metadata, referral mechanisms, domain activity, and platform behavior. While this ensures safety and legality, it limits our ability to verify the exact nature of the shared content.

To this end, the following section presents an overview of the related work. Then, in Section III, we provide an overview of the CSAM campaign, and in Section IV we document the data collection methodology employed for discovering and extracting information from URLs hosting instances of CSAM distribution platforms on the clear web, while also shedding light on the common mechanics used by their operators. We then examine the methods by which these campaigns are disseminated on X/Twitter, focusing on the use of sockpuppet accounts and compromised user profiles. Furthermore, in Section VI we analyze referral network data extracted from these platforms to assess the extent to which the MLM structure is organic, consisting of actual CSAM users who invite other users, perpetuating the cycle. Subsequently, in Section VII we leverage anonymized data captured during a brief window of opportunity, when one of the platforms inadvertently leaked information about visitors' IP addresses and web browsers. Using this data, we provide insights into the visitors' browser footprints, geographic distribution, and other relevant patterns. Since our research has its limitations, we discuss them in Section VIII. Finally, the article concludes, giving an overview of our contributions and analysis, and discussing ideas for future work and mitigation measures.

II. RELATED WORK

Due to the fact that exchanging CSAM is illegal, individuals engaging in such exchanges attempt to implement as many security measures as possible. To this end, users tend to use networks that provide them with as much privacy as possible.

The most obvious method to hide this traffic is through the use of the dark web. Essentially, the traffic in Tor and I2P, the primary dark web networks, not only encrypt all the exchanged traffic but also hide the traffic by, e.g., forwarding the traffic from one node to another, and then using exit nodes to obfuscate where the traffic originates. As a result, dark web services are repeatedly reported for hosting and serving CSAM [8], [9], [10], [6], while some search engines are also reported to promote it [11].

Moreover, peer-to-peer networks offer a veil of anonymity, and the fact that files are broken into smaller chunks that are exchanged among peers facilitates the obfuscation of these exchanges, but also some sort of deniability. Therefore, BitTorrent has been widely used to exchange CSAM [12], [13], [14] with several methods and tools introduced to detect such content, e.g., file names [15]. Nevertheless, other P2P networks, such as IPFS, have been reported to host CSAM [16].

Finally, end-to-end encryption communication platforms are increasingly being abused for CSAM exchange and communication [17]. Beyond strong cryptographic primitives, these platforms provide secrecy and several privacy guarantees in a very compact form, as they can operate in mobile devices, without the need for any expertise.

Given that the detection of most CSAM is based on hashes, the detection of new content is very challenging, e.g., the hash can easily be manipulated to significantly deviate from the original by simply cropping, rotating, or blurring an image or a video frame. Therefore, further research in this direction is necessary, considering the role of AI and machine learning in this area and the potential to automate the corresponding pipelines [18]. For instance, similar to [15], Guerra and Westlake [19] use the names on the websites to detect CSAM, or age estimation [20] to quickly assess the content. Nevertheless, the elephant in the room remains: the identification of victims and offenders.

According to [21], almost half of the sample of CSAM users responded that after watching CSAM, they then sought direct contact with children through online platforms, signifying the impact that watching such material has on its users. In particular, most users reported that they had accidentally been exposed to this material for the first time, usually at a young age. The latter is also aligned with previous studies [7], [11].

III. OVERVIEW OF THE CSAM CAMPAIGN

The CSAM campaign that we examine follows a Multi-level Marketing (MLM) scheme. In essence, the operators employ a 'freemium' model, offering some material for free while reserving more explicit content for higher tiers. The tiers are unlocked either through payments or by inviting other users through personalized referral links. The tiers unlocked with a price that ranges from \$20 to \$120, and according to the platform, provide access to even several terabytes of content.

In this regard, the following paragraphs outline the two main pillars of this campaign, namely the hosting and the distribution of the invitations.

As discussed, the campaign utilizes the clear web, employing multiple domains to host web platforms for user registration and content access. The domains typically have a particularly short lifespan, likely intended to evade attempts by law enforcement agencies to systematically eliminate, investigate, and prosecute perpetrators. Moreover, they use a specific template that gradually changes. A unique characteristic of the campaign is the choice of the domain names. The operators use domain names that exhibit characteristics of Domain Generation Algorithms (DGA) [22], [23]. Practically, the domain name in most cases is a combination of random letters and a few digits, with a preference for specific TLDs, as discussed later.

The referral system of the campaign resembles traditional social media referral mechanisms, enabling the rapid and systematic dissemination across mainstream platforms, such as X, Facebook, and Instagram. Moreover, black-hat marketing methods, such as sockpuppets or posts from compromised social media accounts, further amplify the spread to a broader audience. Interestingly, our analysis demonstrates that while the spread of the campaign is orchestrated by its operators to reach potential buyers, a significant portion of the traffic to these platforms is organic. This traffic originates from offenders actively seeking to unlock higher content tiers by sharing their personalized invite links. To this end, sites that allow for posting without any form of authorization or CAPTCHA mechanism are frequently abused. These include paste-like platforms¹ or comment sections on websites where the posted content is indexed by Google and becomes discoverable by offenders searching for specific keywords. An example of such a site is updownradar.com (e.g. <https://archive.ph/xvZD0>). Such keywords often include the initials “C” and “P”, a coded reference to child pornography, disguised as terms like ‘Club Penguin’ (a discontinued massively multiplayer online game) or phrases such as ‘Caldo de Pollo’ (Spanish for chicken soup). These keywords serve as covert language to avoid detection while used in plain sight. This technique has also been observed in the context of sexually grooming minors in Social Live Streaming Services that employ text-based moderation mechanisms [24]. Notably, the operators of this campaign actively prompt users to amplify and exacerbate its spread by promoting it in popular social platforms, through posting their referral links and using hashtags including such coded terms (i.e., #caldodepollo). For instance, some examples of the relevant interface used in the campaign platforms are shown in Figure 2.

Another aspect of interest for the specific campaign is the use of AI-generated media and images with embedded text for advertising the specific addresses where the campaign’s platforms are hosted, as well as the Telegram handles of the campaign operators. The fact that these images are AI-generated is demonstrated by the fact that they have obvious faults, e.g., more fingers. These images are used both to promote the campaign on social platforms such as X (possibly to avoid

posting links directly and thus evading detection/moderation mechanisms), and within the CSAM sites themselves, as logos or banners.

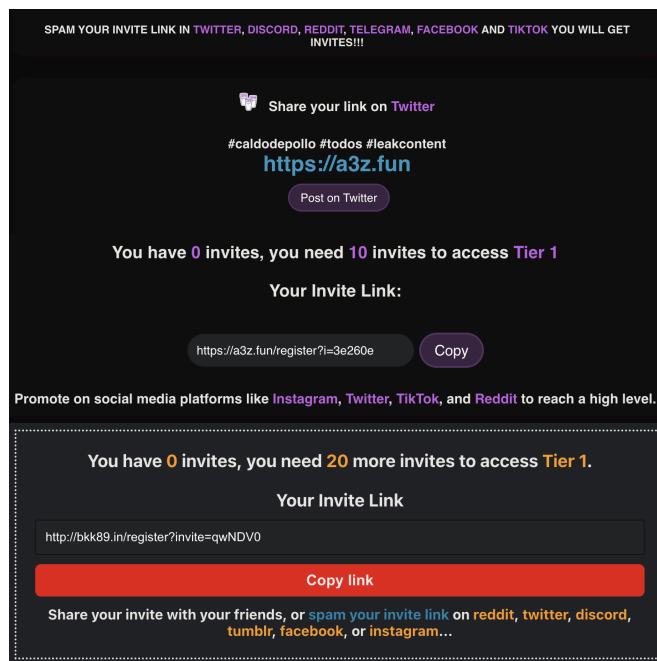


Fig. 2: UI prompting CSAM users to disseminate their referral links in social media.

It must be noted that the first documented observation of this campaign that we have found can be traced back to a Reddit thread² in December 2022, where a user provided some initial insights³ into the modus operandi of the clear-web CSAM distribution platforms:

“They spam themselves on Twitter and dead sub-reddits and it’s structured like a pyramid scheme (kinda). Basically they offer “rewards” in a tiering system if you invite people to their discord, from there it’s filled with links to other discords, their telegram and crypto addresses. Ultimately the invites probably don’t matter since the “free” users will likely never get much but the whales who contact them via telegram and actually start paying are the real prize. They can get taken down on Twitter all they like but the paying customers get to private channels and the pedos get their money. I’m assuming they’ve avoided companies cracking down on them by sheer volume and the fact they tend to use various Russian characters or vague names to disguise their activities.”

IV. DATA COLLECTION AND PROCESSING METHODOLOGY

As the spread of campaign links is highly dynamic, with perpetrators leveraging multiple outlets to post and

¹e.g. <https://archive.ph/8Annc>

²<https://archive.ph/m2b8J>

³<https://archive.ph/G7IXa>

promote their links, we adopt a two-faceted approach to collect URLs hosting such platforms. To capture links posted on paste-like sites or other websites, as previously noted, we constructed a list of queries tailored to match some of the most evidently prominent TLDs by operators to host such platforms. These queries included relevant URL patterns for registering with referral links, such as ``. {cc, de, ru, pw, ws, xyz, city} /invite/i='``. Using SerpAPI⁴, we extracted results indexed by Google that matched these patterns.

Additionally, to retrieve relevant posts from X (formerly Twitter) that include hashtags selected by campaign operators for disseminating links (e.g., #caldodepollo, #clubpenguin, #irlxli, #t33n), we scraped data from the sotwe⁵ platform. This platform utilizes the Twitter API, alleviating the need for costly API access required to conduct such research [25].

Using the methods described above, we constructed a data collection pipeline to identify and extract referral links from these two sources, and subsequently collect the entire source code of the CSAM distribution platforms (if working at crawl time). The automation not only facilitates data processing, but it also prevents access to CSAM content. This pipeline was executed weekly over nine months (June 2024–March 2025), during which we identified 1,026 unique domains hosting various variants of CSAM distribution platforms. Specifically, the SerpAPI approach yielded 745 unique domains, while the sotwe-based approach yielded 281 more domains. It should be noted that for the second approach, almost all domains were not provided in text but embedded in images. To extract them, we employed an OCR pipeline similar to the one described in [26]. We observe that while some domain names seem algorithmically generated, others appear to be hand-picked by the campaign operators, since they include words like *young*, *teen*, etc. Figure 3 shows the most popular sites indexed by Google containing referral links to CSAM distribution sites. We observe that sites allowing users to post comments are frequently abused, including the popular social media platforms, as well as pornographic and pasting websites. The latter finding is a common pattern of disseminating links and information in cybercrime contexts [27], [28]. Notably, the presence of sites such as *urlquery.net* and *urlscan.io* suggests that this campaign has already been flagged and investigated by other researchers or cybersecurity professionals.

V. CAMPAIGN SPREAD ON X INSIGHTS

In total, our sotwe approach yielded 225,248 tweets containing one of the keywords that are known to be used by the campaign operators, which were posted by 1,616 unique users. In order to assess the extent to which such tweets were posted in an automated manner, we establish the following threshold: if the same account posted at least two tweets in our datasets within a window of 60 seconds, we label that

account as a bot. Although this is a conservative threshold that likely underestimates the real scale of automation used to spread this campaign, it provides high confidence in our classification effort.

Using this criterion, we identified 731 automated accounts (45.2% of all accounts) that were responsible for 223,980 tweets (99.4% of the entire dataset). This means that for the duration of our data collection, the CSAM campaign was almost entirely spread by bots on the X platform. Notably, such bots posted an average of 306.4 tweets each, while non-automated accounts averaged just 2.0 tweets. Specifically, this observation suggests that the targeted hashtags were primarily used by CSAM campaign operators to communicate CSAM-hosting domains and Telegram channels used for monetization to potential users.

Examining the age distribution of automated accounts, we observe a bimodal pattern. Specifically, 383 bot accounts (52.4%) were created less than one week before posting, contributing 24,954 tweets. On the other hand, 330 accounts (around 45%) were over one year old, responsible for 197,680 tweets (around 88% of all automated content). Almost no mid-age accounts were observed (only 18 accounts between 1 week and 1 year old), suggesting the prevalence of freshly created sockpuppets or compromised/purchased aged accounts. This modus operandi reflects established patterns in illicit ecosystems, where such services are readily available on cybercrime marketplaces [29].

To better understand the differentiating attributes of these two types of accounts, we examine their number of followers. Expectedly, the newly created sockpuppets had practically no followers, while the older accounts averaged 65 followers. Despite this disparity, the two groups achieved nearly the same aggregate reach: sockpuppets' tweets amassed 2,110,270 views (24,954 tweets), and older accounts' tweets 2,213,699 views (197,680 tweets), reaching a total of 4,323,969 views combined. Engagement, however, was effectively zero for both account categories (zero median retweets and favorites). This indicates that visibility is not driven by follower networks, but rather by demand-side discovery of a narrow set of coded CSAM keywords/hashtags. In practice, the campaign's spam reaches far more people than the bots' almost-nonexistent follower bases would predict, with clear implications for recruiting and funneling new potential users into the ecosystem.

Next, we analyze the hashtags found in the collected tweets. In total, there are 298,958 hashtag occurrences, and in almost every tweet, multiple hashtags were used. We plot hashtags with more than five occurrences as a word cloud in Figure 4. We observe that beyond the most prominent tags which were explicitly advertised directly in the CSAM platforms, i.e., #caldodepollo (45,624 occurrences), #irlxli (38,311 occurrences) and #clubpenguin (25,755 occurrences), there are hashtags with explicit terms like #lolipoptwt (32,082 occurrences), #groomertwt (29,672 occurrences) and #anyage (4,555), as well as references to Tele-

⁴<https://serpapi.com>

⁵<https://www.sotwe.com>




Fig. 3: Top domains hosting CSAM referral links from Google results via SerpAPI.



Fig. 4: Hashtags used for disseminating the campaign on X.

GuardI⁶ (#teleguard with 38,112 occurrences), another possible channel of CSAM distribution. Finally, the existence of variations/misspellings like #caldo_de_pollo

(10,890 occurrences), #caldopollo (9,422 occurrences), and #clubpenguin (8,609 occurrences) confirms the adoption of noisy text as a means to circumvent keyword-based detection mechanisms, a technique widely adapted by users

⁶<https://teleguard.com>

involved in online grooming/CSAM ecosystems [24].

Notably, 70.8% of all tweets contained media attachments. This aligns with our earlier observations about AI-generated images containing embedded URLs and Telegram handles, used to bypass text-based moderation. Furthermore, 15% of the tweets contained Telegram invite links, comprising 290 unique channels/accounts. The top channel was referenced in 4,665 tweets posted by eight different bot accounts.

VI. ANALYZING THE REFERRAL NETWORK

Delving into the collected JavaScript code from the CSAM platforms in our dataset, we were able to extract several important findings. For instance, most of the platforms deployed, even across different domains, used the same backend server, and, as a result, the same APIs. Interestingly, consistent across most deployments is the use of a specific API that displays a real-time activity feed of the referral network, showing who invited whom and each user's accumulated invite count. An example is shown in Figure 5. This public ticker mechanism was probably designed to create engagement by motivating users to invite others to unlock access to more explicit CSAM, thus further spreading the campaign.

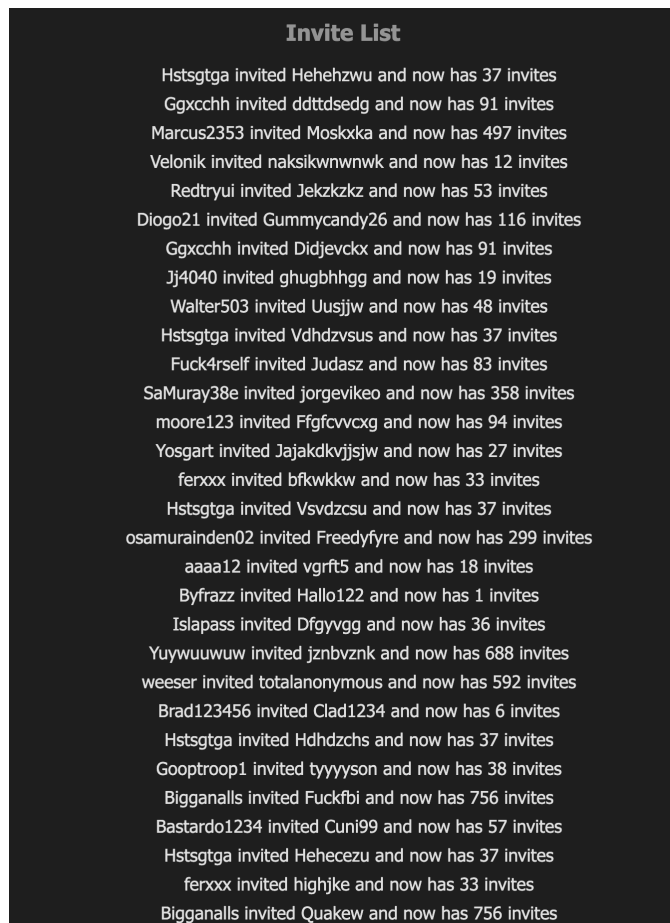


Fig. 5: Referral ticker.

As the live expansion of the referral network was made available through this API, which did not require authenti-

cation, we queried it on a daily basis to collect information about inviters and invitees and to unveil the referral dynamics. This data collection effort also lasted for the same 9-month duration, but it should be noted that there were sparse periods of downtime related to the specific API, which in total account for 57 days where no data was collected, resulting in a complete snapshot of the referral network formation over 247 days.

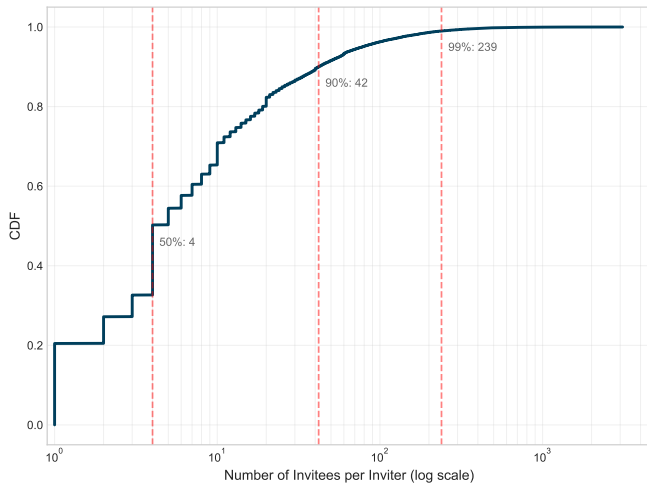
In the collected data, we observed 700,531 new user registrations via the referral system, which were generated by 37,755 unique inviters. Notably, 22,431 invited users (3.2%) became inviters themselves, further cascading the referral chain. We construct the inviter-invitee graph and plot the cumulative distribution function (CDF) and the frequency distribution of invitees per inviter on a logarithmic scale in Figure 6. We observe that 50% of inviters recruited four or fewer users, 90% recruited 42 or fewer, yet the top 1% (377 users) recruited more than 239 users each. Specifically, the top 1% of inviters generated 22.5% of all registrations, while the top 10% (3,775 users) accounted for 64.5% of the campaign's network growth. Furthermore, the most prolific inviter single-handedly recruited 3,103 users. The frequency distribution is consistent with the power law distributions observed in scale-free networks [30], indicating that the observed recruitment/engagement of CSAM users is organic, meaning that the registered users are real, rather than artificial. Concretely, the distribution reveals three distinct user categories: (1) casual participants comprising the majority who recruit fewer than 10 users, likely seeking content access; (2) motivated affiliates in the top 10% driving substantial growth; and (3) super-spreaders who may be campaign operators or highly incentivized distributors.

VII. CHARACTERIZING CSAM USERS


Next, we analyze the data to extract behavioral patterns and characteristics of the users participating in the CSAM referral network. On average, $\approx 2,836$ new users joined per day in our 247-day dataset, indicating sustained demand for this illegal content. In Figure 7, we plot the total number of registered users per day. Evidently, the most popular registration days are Friday, Saturday, and Sunday, with 348,010 total registrations. These three days average 116,003 registrations per day compared to 88,130 for other weekdays.

As we analyzed the data, we identified a brief window during which the API leaked additional user information, providing us with further insight into user characteristics. For approximately 30 days during our data collection period, the returned data regarding registration contained additional fields for each user, including their country, timezone, and User Agent. Although this appears to be a catastrophic operational security failure by the campaign operators, it provided invaluable data about the user base. In total, we captured these attributes for 65,261 users before the leak was patched.

We present the geographic distribution of the users of the CSAM platform in Figure 8. The United States dominates with 38,382 users (58.8% of the total with country attribute),



(a) Cumulative distribution of invitees per inviter.



(b) Distribution of invitees per inviter.

Fig. 6: Invitees per inviter distribution.




Fig. 7: Total user registrations per day of the week.

followed by a second level comprising the United Kingdom (6,085), Brazil (6,079), Germany (3,513), Mexico (3,320), and Canada (3,046). In particular, Western Europe, the Americas, and Australia account for the vast majority of users, while regions with lower Internet adoption or restrictive policies show proportionally fewer participants. However, the users of the campaign platforms during the leak period were distributed in 110 different countries, indicating that this campaign has a global reach. A limitation of this analysis is that we cannot determine the extent of proxy or VPN/Tor usage, which may impact the actual geographic distribution of users.

To understand the devices used to access these platforms,

we parsed the User Agent strings using the `user-agents`⁷ Python library. Fig. 9 shows the distribution across operating systems. We observe that almost half of the users were using iOS devices, followed by Android. Overall, mobile devices accounted for 84.1% of usage. In contrast, desktop users account for a small fraction, as the combined share of desktop operating systems of Windows, macOS, and Linux represents less than 15% of users. Finally, we tested the usernames in the complete dataset against a regular expression matching email patterns, identifying 1,606 distinct email addresses among the usernames. This means that, despite the criminal nature of the content, a subset of users still used email addresses for registration, even though the platform did not require or validate them. Thus, everyone could use arbitrary emails, even ones that did not belong to them.

VIII. LIMITATIONS

Despite the scale and depth of our findings, this study is subject to several limitations, which, for transparency, we list below.

First, our data collection was constrained by the visibility of indexed content on public platforms and the limitations of API-based querying. While our methodology identified more than 1,000 domains and captured thousands of tweets and site metadata, it is highly likely that we only observed a fragment of the entire CSAM ecosystem operating on the clear web. Domains hosted on private forums, non-indexed sites, or behind paywalls may have been missed by our detection pipeline.

Moreover, several behavioral insights, such as user geography and browser data, were derived from a temporary API misconfiguration that leaked identifiable metadata. Although these data were collected passively and ethically, the leak period spanned only 30 days, representing less than 15% of our

⁷<https://pypi.org/project/user-agents/>

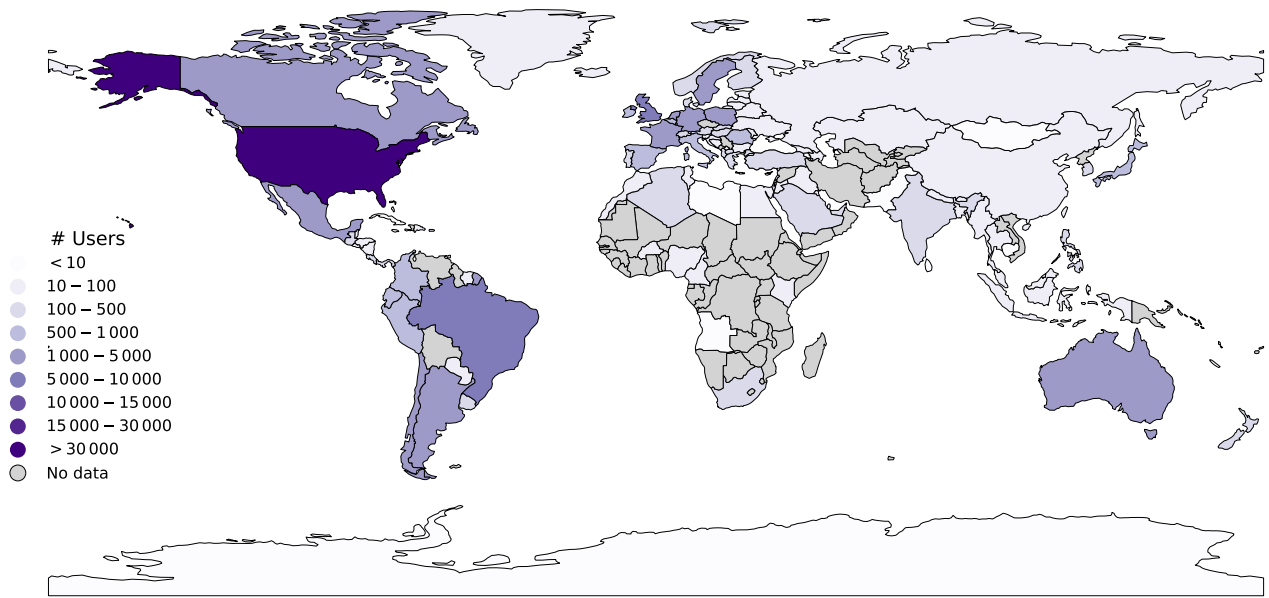


Fig. 8: Geographic distribution of CSAM platform users.

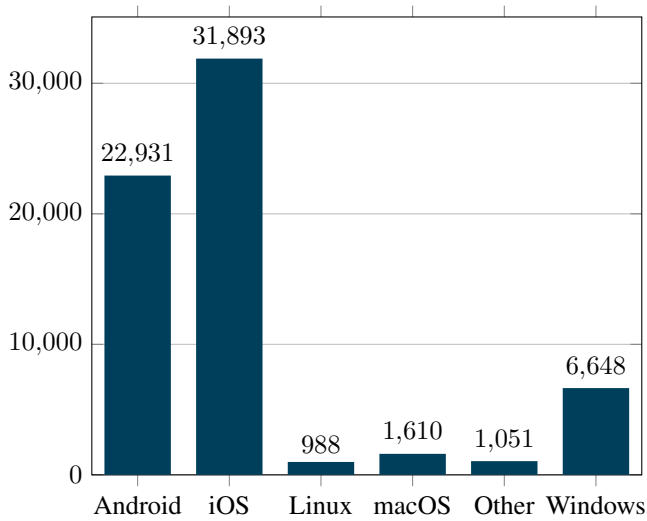


Fig. 9: Operating Systems used by CSAM users.

total collection window. As such, the resulting characterization of users may not be applicable to the entire user base.

As discussed, to estimate the use of bots on social media platforms, we employed a conservative rule: accounts posting two or more tweets within a 60-second window were labeled as bots. While this minimizes false positives, it may underestimate the true scale of bot activity and overlook more sophisticated automation techniques that simulate human behavior.

Finally, in the same vein of thought, the structure of the referral system suggests genuine viral spread; however, it is possible that some users fabricated referrals to unlock higher content tiers, thereby inflating the number of invites. Similarly,

the use of arbitrary or fake email addresses during registration undermines any direct mapping between user identities and real individuals.

IX. CONCLUSION

In this paper, we presented the first longitudinal, data-driven investigation of a large-scale CSAM referral campaign operating openly on the clear web. Our research shows that CSAM distribution is a persistent and evolving threat that plagues traditional social media platforms, whose content moderation and reporting mechanisms are easily circumvented. Moreover, the large-scale bot activity facilitates the global dissemination of content. The fact that the campaign has been operating for so long on the clear web rather than hidden on encrypted or anonymized platforms illustrates both the limitations of current enforcement mechanisms and the existence of a large, engaged user base.

By combining multiple data collection pipelines and exploiting operational flaws in the infrastructure of CSAM campaigns, we were able to reveal in detail the structure and behavior of this ecosystem, from domain lifecycles and user acquisition funnels to device/browser usage patterns and geolocation data, without directly accessing illicit material. Our work provides evidence that a multi-level marketing-style system fuels much of the CSAM circulation on the clear web, and that the majority of engagement appears to stem from genuine users, rather than synthetic or inflated interactions. Nonetheless, our research reveals that illicit campaigns encounter operational issues that can be leveraged to gain deep insight into what is exchanged and to better understand their modus operandi, and aid future research and law enforcement efforts.

We argue that the huge and continuous increase in the CSAM distribution, coupled with the use of AI, will exacerbate the problem. The growing realism of synthetic content raises new questions regarding legality, ethics, and the identification of victims. As a result, law enforcement authorities will struggle to determine whether there is an actual victim in the seized content and identify it.

These developments underscore the urgency for improved technical safeguards, rapid response pipelines, and coordinated action between platforms, researchers, and regulatory bodies. We hope that our work serves as both a warning and a foundation for future work. Continued interdisciplinary action is essential if we are to curb the proliferation of CSAM as new content is continuously generated and distributed, affecting thousands of victims worldwide.

ACKNOWLEDGMENT

This work was partially supported by the European Commission under the Horizon Europe Programme, as part of the project SAFEHORIZON (Grant Agreement No. 101168562).

The content of this article does not reflect the official opinion of the European Union. Responsibility for the information and views expressed therein lies entirely with the authors.

REFERENCES

- [1] C. Guitton, "A review of the available content on Tor hidden services: The case against further development," *Computers in Human Behavior*, vol. 29, no. 6, pp. 2805–2815, 2013.
- [2] C. Gannon, A. A. Blokland, S. Huikuri, K. M. Babchishin, and R. J. Lehmann, "Child sexual abuse material on the darknet," *Forensische Psychiatrie, Psychologie, Kriminologie*, vol. 17, no. 4, pp. 353–365, 2023.
- [3] V. M. Ngo, R. Gajula, C. Thorpe, and S. McKeever, "Discovering child sexual abuse material creators' behaviors and preferences on the dark web," *Child Abuse & Neglect*, vol. 147, p. 106558, 2024.
- [4] Europol, "Internet Organised Crime Threat Assessment (IOCTA) 2020," 2024. [Online]. Available: https://www.europol.europa.eu/cms/sites/default/files/documents/internet_organised_crime_threat_assessment_iocta_2020.pdf
- [5] International Association of Internet Hotlines - INHOPE, "Annual report 2024," 2025. [Online]. Available: <https://inhope.org/EN/articles/inhope-annual-report-2024>
- [6] Europol, "Global crackdown on Kidflix, a major child sexual exploitation platform with almost two million users," 2025. [Online]. Available: <https://www.europol.europa.eu/mediapress/newsroom/news/global-crackdown-kidflix-major-child-sexual-exploitation-platform-almost-two-million-users>
- [7] R. Wortley, D. Findlater, A. Bailey, and D. Zuhair, "Accessing child sexual abuse material: Pathways to offending and online behaviour," *Child Abuse & Neglect*, vol. 154, p. 106936, 2024.
- [8] R. Liggett, J. R. Lee, A. L. Roddy, and M. A. Wallin, "The dark web as a platform for crime: An exploration of illicit drug, firearm, csam, and cybercrime markets," in *The Palgrave handbook of international cybercrime and cyberdeviance*. Springer, 2020, pp. 91–116.
- [9] J. Woodhams, J. A. Kloess, B. Jose, and C. E. Hamilton-Giachritsis, "Characteristics and behaviors of anonymous users of dark web platforms suspected of child sexual offenses," *Frontiers in Psychology*, vol. 12, p. 623668, 2021.
- [10] M. van der Bruggen, I. van Balen, A. van Bunningen, P. Talens, J. N. Owens, and K. Clapp, "Even 'lurkers' download: The behavior and illegal activities of members on a child sexual exploitation Tor hidden service," *Aggression and violent behavior*, vol. 67, p. 101793, 2022.
- [11] J. Nurmi, A. Paju, B. B. Brumley, T. Insoll, A. K. Ovaska, V. Soloveva, N. Vaaranen-Valkonen, M. Aaltonen, and D. Arroyo, "Investigating child sexual abuse material availability, searches, and users on the anonymous Tor network for a public health intervention strategy," *Scientific Reports*, vol. 14, no. 1, p. 7849, 2024.
- [12] M. Rutgaizer, Y. Shavitt, O. Vertman, and N. Zilberman, "Detecting pedophile activity in bittorrent networks," in *International Conference on Passive and Active Network Measurement*. Springer, 2012, pp. 106–115.
- [13] Y. Shavitt and N. Zilberman, "On child sex abuse presence in bittorrent networks," *IEEE Internet Computing*, p. 1, 2013.
- [14] J. Wolak, M. Liberatore, and B. N. Levine, "Measuring a year of child pornography trafficking by us computers on a peer-to-peer network," *Child Abuse & Neglect*, vol. 38, no. 2, pp. 347–356, 2014.
- [15] A. Panchenko, R. Beaufort, and C. Fairon, "Detection of child sexual abuse media on p2p networks: Normalization and classification of associated filenames," in *Proceedings of the LREC Workshop on Language Resources for Public Security Applications*, 2012, pp. 27–31.
- [16] S. Sokoto, L. Balduf, D. Trautwein, Y. Wei, G. Tyson, I. Castro, O. Ascigil, G. Pavlou, M. Korczyński, B. Scheuermann *et al.*, "Guardians of the galaxy: Content moderation in the [InterPlanetary] file system," in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 1507–1524.
- [17] Europol, "Internet Organised Crime Threat Assessment (IOCTA) 2024," 2020. [Online]. Available: <https://www.europol.europa.eu/cms/sites/default/files/documents/Internet%20Organised%20Crime%20Threat%20Assessment%20IOCTA%202024.pdf>
- [18] S. Singh and V. Nambiar, "Role of artificial intelligence in the prevention of online child sexual abuse: A systematic review of literature," *Journal of Applied Security Research*, vol. 19, no. 4, pp. 586–627, 2024.
- [19] E. Guerra and B. G. Westlake, "Detecting child sexual abuse images: Traits of child sexual exploitation hosting and displaying websites," *Child Abuse & Neglect*, vol. 122, p. 105336, 2021.
- [20] M. Castrillón-Santana, J. Lorenzo-Navarro, C. M. Travieso-González, D. Freire-Obregón, and J. B. Alonso-Hernandez, "Evaluation of local descriptors and cnns for non-adult detection in visual content," *Pattern Recognition Letters*, vol. 113, pp. 10–18, 2018.
- [21] T. Insoll, A. K. Ovaska, J. Nurmi, M. Aaltonen, and N. Vaaranen-Valkonen, "Risk factors for child sexual abuse material users contacting children online: Results of an anonymous multilingual survey on the dark web," *Journal of Online Trust and Safety*, vol. 1, no. 2, 2022.
- [22] D. Plohmann, K. Yakdan, M. Klatt, J. Bader, and E. Gerhards-Padilla, "A comprehensive measurement study of domain generating malware," in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 263–278.
- [23] F. Casino, N. Lykousas, I. Homoliak, C. Patsakis, and J. C. Hernandez-Castro, "Intercepting hail hydra: Real-time detection of algorithmically generated domains," *J. Netw. Comput. Appl.*, vol. 190, p. 103135, 2021.
- [24] N. Lykousas and C. Patsakis, "Large-scale analysis of grooming in modern social networks," *Expert Syst. Appl.*, vol. 176, p. 114808, 2021.
- [25] H. Ledford, "Researchers scramble as Twitter plans to end free data access," *Nature*, vol. 614, no. 7949, pp. 602–603, 2023.
- [26] F. Casino, N. Totosis, T. Apostolopoulos, N. Lykousas, and C. Patsakis, "Analysis and correlation of visual evidence in campaigns of malicious office documents," *Digital Threats: Research and Practice*, vol. 4, no. 2, pp. 1–19, 2023.
- [27] M. Chertoff and T. Simon, "The impact of the dark web on internet governance and cyber security," 2015.
- [28] R. Madarie, S. Ruiter, W. Steenbeek, and E. Kleemans, "Stolen account credentials: an empirical comparison of online dissemination on different platforms," *Journal of Crime and Justice*, vol. 42, no. 5, pp. 551–568, 2019.
- [29] N. Lykousas, V. Koutsokostas, F. Casino, and C. Patsakis, "The cynicism of modern cybercrime: Automating the analysis of surface web marketplaces," in *2023 IEEE International Conference on Service-Oriented System Engineering (SOSE)*. IEEE, 2023, pp. 161–171.
- [30] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.