Tickets to Hide: An Inside Look into the Anti-Abuse Ecosystem through Internal Abuse Data

Hugo L.J. Bijmans
TU Delft
h.l.j.bijmans-1@tudelft.nl

Michel J.G. van Eeten
TU Delft
m.j.g.vaneeten@tudelft.nl

Rolf S. van Wegberg
TU Delft
r.s.vanwegberg@tudelft.nl

Abstract—**Various governance instruments aim to fight Internet abuse – from legislation to take down copyrighted material to blocklists to stop spam. In turn, these instruments rely on industry standards to handle abuse: reporting abuse to the network owners requesting mitigation. Although many hosting providers swiftly take action to keep the Internet clean, some do not. This raises the question as to what type of abuse receives follow-up and what rationale is behind a decision to either mitigate or ignore reported abuse. Through a unique collaboration with law enforcement in the Netherlands, we were granted access to the operational back-end of a hosting provider with a reputation for abuse. A rare glimpse into its internal abuse handling allowed for the investigation of the mechanisms in the anti-abuse ecosystem that influence anti-abuse actions. We find that client notification rates highly depend on the reporter and the abuse category. CSAM and spam-related abuse reports lead to mitigating actions, whereas reports regarding copyright infringement and port scanning are often neglected. Governance instruments, such as blocklisting, de-peering, and law enforcement inquiries, that could directly impact business continuity, affect client notifications, whereas individual abuse reporting is often easily ignored. We hope our work can inform policymakers on aligning governance repertoire with effective abuse handling in practice.**

I. Introduction

To combat Internet abuse like spam emailing, hosting phishing pages, or sharing copyrighted material, the Internet relies on the practice of abuse reporting, also known as 'notice and takedown'. Here, an abuse report is sent to a service provider, who subsequently notifies its client [1]. Ultimately, either the client or the provider takes action in the case of a legitimate report. For a hosting provider with many clients, handling abuse reports requires effort. Industry best practices recommend swift follow-up for abuse, yet different priorities are assigned to different types of abuse [2]. Adhering to these best practices is voluntary, though some jurisdictions have basic legal obligations for the hoster to evaluate abuse reports. However, evaluating can also mean deciding not to act. New legislation in the European Union – e.g., the Digital Services Act [3] – has codified a 'notice and action' procedure and introduced 'trusted flaggers', which are designated independent entities whose abuse reports should be acted on with higher priority. This indicates that we heavily rely on reporting as a governance mechanism in the fight against online abuse.

Although abuse reporting has been around for years [2], [4], there is limited scientific insight into the abuse-handling processes of hosting providers. Prior studies have taken an external viewpoint and analyzed how abuse events such as phishing sites, spam, and malware C&C servers are distributed across hosters [5]–[7]. Other studies leveraged external characteristics of hosters to identify potentially malicious networks [6], [8]. They found BGP routing dynamics, fragmentation, churn of advertised IP address space, and network size to be strong indicators of malicious activity [6], [8].

However, large amounts of abuse can also occur at legitimate providers, simply because of the size of their infrastructure [6]. Hence, abuse concentrations do not reveal much about the network operators' effort. Are they trying to mitigate the abuse, or are they condoning it? Operators who are perceived not to swiftly respond to abuse reports are referred to as 'bad' or 'bulletproof' hosters – that is, they are seen as impervious to abuse reports [9]. Some characteristics of such hosting have been described in both earlier work [10], [11] and industry reports [12]. Labeling a company as 'bulletproof' assumes that the hoster is at least knowingly ignoring abuse reports – and perhaps even actively enabling abuse [13]. However, examining intent requires an inside view instead of external network characteristics. To the best of our knowledge, only a single study has provided such an inside view on a bad hoster [10], but it did not analyze abuse report handling.

To address this gap, we present a study of the internal processes for handling abuse at a hoster with a reputation for abuse. While it is clear that a single case study has limitations, our understanding of bulletproof hosters (BPHs) critically depends on them. First of all, BPHs are very rare, which stands in tension with the need to study larger samples. Second, larger samples can only be achieved by using external measurements. This means relinquishing access to ground-truth data on the internal workings of BPHs. Such data is only available through seized data obtained in the course of a criminal investigation, which, by necessity, is highly targeted against a specific entity. The external and internal views of BPHs are highly complementary, and both are necessary for effective strategies to combat these criminal organizations. The

Network and Distributed System Security (NDSS) Symposium 2026
23-27 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.240468
www.ndss-symposium.org

leading study that uses an external view is Alrwais et al. [9], which develops a detection approach to identify BPHs. That approach is based on assumptions about how BPHs operate and how they can be detected with external measurements. Without ground-truth data on internal operations, we cannot test the validity of these assumptions and, therefore, determine whether a detection method delivers reliable results or identify effective interventions against such companies. Hence, single case studies and external measurements depend on each other.

In our (single case) study, we aim to answer the following questions: **RQ1**: 'Which type of abuse is followed up on?', **RQ2**: 'What factors influence the decision to follow up on abuse reports?', and **RQ3**: 'How do external network characteristics relate to internal abuse handling?'. We do not disclose the name of the studied hosting provider (and discuss the ethics involved in our research in the Ethics Considerations section at the end of this paper), but it has been listed as one of the top bad hosters since the early 2010s [14], albeit under different brand names. It was referred to as a bulletproof hoster by both law enforcement [15], industry [12], [14], and media [16]. Over the years, the company has regularly changed its name, created new brands, and relocated its registration to the Seychelles, and has remained in operation as of 2025.

In 2020, the Dutch Fiscal Information and Investigation Service (FIOD) raided the company and seized company records [15]. Through a unique collaboration with this agency, we acquired access to its operational back-end, which allowed us to investigate its abuse-handling practices. We collected 1.3M abuse reports from nine years of operation, categorized them, and identified corresponding client notifications. Furthermore, we investigate the time-to-notify and notification rates over time for each abuse category. Lastly, we connect our work to previous studies by relating our findings to the malicious network characteristics identified in prior research.

We find that notification rates highly depend on the reporter and the abuse category. Child sexual abuse material (CSAM) and spam-related reports do result in notifications, whereas copyright and port scanning reports have low client notification rates. Governance mechanisms, such as blocklisting, de-peering, and law enforcement inquiries, that could directly impact business continuity, affect client notifications, whereas individual abuse reporting is often easily ignored. We identify some indicators of malicious networks found in prior research, yet argue that labeling hosters as bulletproof based solely on external data is a tough label to sell.

In short, we make the following contributions:

1) We are the first to report on internal data from a hosting provider, leveraging over 1.3M abuse reports and over 9k client notifications spanning nine years.
2) We find that CSAM and spam-related abuse reports result in mitigating actions, whereas reports regarding copyright and port scanning are neglected.
3) We empirically demonstrate that abuse reports from reporters with a trusted status or those in a position to impact business processes have higher client notification rates than those from individual reporters.
4) Through our longitudinal analysis, we find that government pressure and direct threats resulting from ignoring reports affect client notification rates.

The remainder of this paper is structured as follows. First, we discuss the studied company, the anti-abuse ecosystem, and the public debate on anti-abuse legislation in Section II, followed by an overview of related work in Section III. Thereafter, we describe our dataset in Section IV and discuss our methodology in Section V. We present our results in two sections: take an inside look by quantifying abuse in Section VI and take an outside look through an analysis of malicious network characteristics in Section VII. We contextualize our findings and discuss their limitations in Section VIII. Finally, we conclude our work in Section IX. We reflect on the ethics involved in our work at the end of this paper.

II. Background

This section describes the company on which this study is based, the anti-abuse ecosystem, and elaborates on the public debate on anti-abuse legislation.

The Company: The analysis in this paper is based on data gathered from a CRM system and two mailboxes belonging to a hosting company whose name we do not disclose. It has been operating since the early 2000s, with its own data center located in the Netherlands and offering a variety of (dedicated) servers and Internet connectivity to its clients. Through the years, it served thousands of clients with a relatively small number of employees. The owners hired a handful of technical and support personnel to assist in daily operations. The studied database and mailboxes, which we detail in Section IV, reveal that fewer than five different persons communicated with clients. The company has been listed as one of the top bad hosters since the early 2010s [14], albeit under different brand names. It was referred to as a bulletproof hoster by both law enforcement [15], industry [12], [14], and media [16]. In 2020, the hoster was raided by Dutch law enforcement [15], but has remained in business to the present day. So far, there has been no public statement regarding the outcome of this investigation, and its ASN remains listed on popular blocklists, such as Spamhaus [17].

Anti-Abuse Ecosystem: Similar to the inner workings of the Internet, the ecosystem fighting Internet abuse is also decentralized. In previous work on the anti-abuse ecosystem, Jhaveri et al. [1] defined three roles within this ecosystem: the abuse reporter, the intermediary (i.e., hosting provider), and the resource owner responsible for the abusive resource. Every party has different incentives and possibilities for participation. Abuse reporters have several incentives to voluntarily collect and report abuse data, such as altruism, quid pro quo, or being victims themselves. Intermediaries have a business relationship with the resource owner affected by monetary incentives and can decide to forward abuse reports to their clients. In short, an intermediary can either *ignore* a received abuse report, *notify* the client and wait for it to be fixed, *assist* to fix the problem, *suspend* the server, or ultimately *terminate* the client.

TABLE I: Complaint priorities for abuse according to the M^{3}AAWG anti-abuse common practices [2].

Abuse category	Priority
CSAM / Harmful content	Critical
Botnet C&C / DDoS attacks	High
Malware / Phishing / Brute-force attacks	Medium
Spam(vertising)	Low
Port scanning / Comment spamming	Very low
Copyright / Trademark issues	*Depends*

The Message, Mobile, and Malware Anti-Abuse Working Group (M^{3}AAWG) outlined the anti-abuse common best practices in 2015 to assist intermediaries in their anti-abuse efforts [2]. Their *'Anti-Abuse Best Common Practices for Hosting and Cloud Service Providers'* presents guidelines to keep systems safe (such as vetting new customers and keeping software up to date) and outlines the handling of abuse reports. This includes setting up an abuse email account according to RFC2142 [4], making community abuse reporting straightforward, responding promptly to those reports, and considering trusted reporters to handle certain reports with higher priority. Additionally, the M^{3}AAWG common best practices proposes a complaint prioritization that lists the prioritization for different types of abuse reports – shown in Table I. Here, system abuse is categorized into six different categories. Critical priority is assigned to CSAM or harmful content, followed by high priority for botnet C&C servers and DDoS attacks originating from the network. Malware and phishing hosting are grouped with dictionary/brute-force attacks into the medium priority category. The low-priority category includes spam emailing, as well as spamvertising on either the network or a support network. Web defacement, comment spamming, and port scanning are considered very low priority in the M^{3}AAWG categorization. Lastly, copyright and trademark issues vary in priority due to the location of both the issue and the hosting provider. The M^{3}AAWG would categorize these abuse complaints as high or medium in North America because of DMCA Safe Harbour requirements, whereas their complaints would be categorized as low to very low because of different jurisdictions [2].

Anti-Abuse Legislation: Over the years, decentralized structures designed to combat Internet abuse have been pressured by governments worldwide to safeguard the Internet. The resulting legislation differs significantly. The United States adopted the approach of criminalizing computer fraud and Internet abuse in the Computer Fraud and Abuse Act of 1986 (CFAA) [18]. The bill, last updated in 2008, prohibits intentionally accessing a computer without authorization and committing fraud using a computer, but fails to provide definitions or outline instruments to combat online abuse. As a result, many modern-day Internet activities can be prosecuted under the CFAA with severe punishments [19]. To stop the proliferation of CSAM, the U.S. Senate introduced the 'STOP CSAM' act, which makes reporting such material easier and adds administrative penalties when providers fail to remove CSAM within a certain period [20]. It would also institutionalize NCMEC's CyberTipline by requiring companies to report discovered CSAM material. The introduction was met with both enthusiasm and skepticism, as opponents feared it would jeopardize constitutional rights to privacy and freedom of expression, given the risk-averse nature of companies that might block any sensitive material. Legislation by the European Union has already come into force since February 2024 through the Digital Services Act (DSA), which aims to improve digital safeguards and to prevent illegal and harmful activities online [3]. Besides codifying a 'notice-and-action' procedure, nationally appointed Digital Services Coordinators award a number of organizations the 'trusted flagger' status, whose abuse reports must receive priority. Leveraging expertise from organizations specialized in detecting abusive content could enhance the quality of abuse reporting and improve follow-up actions. Some large online platforms, such as Google and Meta, have created similar programs, albeit under their own terms [21], [22]. The DSA shifts this power to the E.U. member state governments and turns voluntary cooperation into mandatory compliance. Reactions to this new legislation have been mixed. Although it was warmly welcomed by civil rights groups, it was criticized by tech companies for creating a heavy burden and by some politicians and scientists [23] for undermining freedom of speech.

III. Related work

Research on Internet abuse spans several relevant areas: (*i*) external measurements of malicious networks, (*ii*) analyses of bulletproof hosting (BPH) infrastructures, and (*iii*) evaluations of anti-abuse interventions. Although each line of work has contributed substantial insights, none provide visibility into internal abuse-handling processes at hosting providers.

Malicious Networks: There is an enormous amount of research that looks at a specific form of abuse. An example is Levchenko et al. [24], which analyzed the spam value chain and demonstrated that spam affiliate programs employed a highly distributed hosting strategy to remain resilient. Such studies might touch on hosting resources, but they are not focused on provider behavior. Closer to our work, there is research that focuses specifically on identifying networks that host disproportionate amounts of abuse. In 2009, Stone-Gross et al. [7] presented *FIRE* to actively monitor botnets, drive-by-downloads, and phishing website feeds to identify organizations and networks that show persistent malicious behavior. Shue et al. [6] extended this by analyzing additional data sources and examining the BGP behavior of malicious Autonomous Systems (AS). They argue that ASes can be malicious due to either malicious intent by the operator or lax administration and poor security practices. Zhang et al. confirmed this [25]. Shue et al. [6] discovered that ASes with the most malicious activity have a greater number of BGP connectivity changes than benign ASes and that larger ASes are more likely to contain malicious IP addresses. Leveraging these BGP observations, Konte et al. [8] created *ASwatch* to identify malicious ASes by their routing behavior.

They identified additional indicators of malicious networks, including aggressive AS rewiring (characterized by numerous changes in providers and peers), BGP routing dynamics (such as short prefix announcements), and fragmentation and churn of the advertised IP address space.

Statistical studies [5], [26] show that observed abuse volumes often reflect exposure (e.g., IP space size and domains hosted), but also factors like the prevalence of popular content management systems (i.e., WordPress), rather than operator negligence. These studies have not empirically analyzed the impact of anti-abuse mechanisms on provider behavior.

Bulletproof Hosting: There is a specific strand of work that focuses on bulletproof hosting (BPH) as a subset of all malicious networks. BPHs are networks that intentionally participate in abuse, and the current challenge is identifying such BPHs. Prior work has mostly relied on concentrations of abuse. Initially, these concentrations were observed at the AS level. In 2018, Alrwais et al. [9] observed that such networks transitioned from large malicious ASes, such as CyberBunker [27], to fragmented infrastructure located at multiple lower-end service providers through sub-allocations. They report that in 2016, only 19.7% of IP addresses blocklisted by Spamhaus were directly allocated – i.e., managed by its service provider – whereas 80.3% were sub-allocations, half of them owned by a client of a legitimate service provider. As a result, their BPH detection method was based on identifying such concentrations in sub-allocations. Similar work was done by Mahjoub [12]. Investigating these sub-allocations revealed that many legitimate service providers are not responsive to reports of abuse within their networks and frequently rotate IP blocks to evade blocklisting. While these studies label providers as BPHs, they only had an external view of these providers, meaning they could not observe whether these providers knowingly facilitated abuse or not. So far, only a single study has analyzed an inside view of a BPH, namely Noroozian et al. [10]. Leveraging ground-truth data extracted from seized back-end databases, they characterized the business model, supply chains, and clients. This study did not examine whether abuse reports were acted upon or forwarded to the customers. Hence, it does not help us understand if and how anti-abuse mechanisms influence the behavior of the provider.

Anti-Abuse Interventions: There is remarkably limited work on the efficacy of specific anti-abuse interventions, such as blocklisting, peering sanctions, takedown services, and abuse notifications. Most work has focused on blocklisting and determining whether it can keep up with actual abuse, most notably for malware [28] and spam [29]. While this is a rich literature, its focus is on the use of such lists by defenders to prevent traffic to or from the listed resources. It has not examined whether blocklists affect the origin of the abuse, i.e., the hosting provider. Studies of phishing and spam takedowns demonstrate a large variation in effectiveness depending on the credibility of the reporter and the costs imposed on intermediaries [1]. One incentive would be to prevent being blocklisted. A final cluster of work is on notifying hosting providers about compromised systems that are being abused and then measuring whether this abuse is taken down [30]. Various studies found that a substantial fraction of hosting providers do remove abusive resources after being notified [31], [32]. Yet, another fraction does not act. Bulletproof hosters would certainly fall in the latter category. Thus, we need additional insights into what interventions they would be sensitive to.

Our work offers the first internal perspective on abuse handling at a hosting provider. Unlike all prior studies, which infer intent and responsiveness solely from external measurements, our dataset reveals how the provider actually triages, ignores, or acts on reports. This direct evidence closes a long-standing gap in the literature and challenges key assumptions behind BPH detection and abuse-mitigation research.

IV. Data

Our analysis is based on data seized by law enforcement. On September 22nd, 2020, the Dutch Fiscal Information and Investigation Service (FIOD) raided the hoster and copied company records [15]. Through a collaboration with this agency, we were given the opportunity to analyze parts of its back-end systems. This unique internal data enables us to conduct empirical research that would otherwise be impossible. Note that this data assisted daily operations and, therefore, is not structured to support scientific research. Hence, we provide a detailed description of the data, discuss how we assessed its validity, and describe the pre-processing steps we performed. Legal and ethical considerations that come with the use of this data are discussed at the end of this paper.

A. Abuse Mailboxes

As mentioned in Section I, the studied company has operated under various brand names over the years. Law enforcement was able to seize the mail servers of the last two brands and shared a copy of the mailboxes used for handling abuse reports. Both mailboxes were stored in `Maildir` format, which also preserves the folder structure. Table IIa lists the folders, the emails it contains, and whether or not we included this folder in our dataset. Some folder names already suggest default follow-up actions, such as those starting with *ignore*. In both mailboxes, we found folders for handled reports in the *handled* folder and folders related to specific reporters, such as Cloudflare. We omitted the *sent items* folders in both mailboxes, as manual analysis revealed no abuse reports nor client notifications, merely (automated) responses to reporters – e.g., we found 38,439 auto-replies to CyberTip demanding them to stop reporting through email and to use a provided takedown tool instead. The mailbox used by the last brand name contained 55,979 emails from 2019-02-01 until 2020-09-22, and the mailbox used by the second to last brand name contained 2,624 emails in the period 2019-03-07 until 2020-09-22. For every email, we extracted the timestamp, subject header, sender, and the message. Although we included the *deleted messages* folder in our dataset, emails could have been permanently deleted from these mailboxes.

TABLE II: An overview of the two mailboxes and their folders in IIa and the selected database tables in IIb.

Mailbox I (55,979 e-mails)	Items	Incl.	Mailbox II (2,624 e-mails)	Items	Incl.
Handled	1,227	✓	Handled	117	✓
Handling	0	✓	Handling	0	✓
Conversations	100		Conversations	7	
Deleted messages	4,236	✓	CP reports handled	1,682	✓
FMTS	4,965	✓	Deleted messages	12	✓
FMTS.FAPL	324	✓	Ignore	537	✓
Ignore	29,686	✓	Ignore - Cloudflare	276	✓
Ignore - Cloudflare	12,830	✓	Sent items	21,256	
Ignore - Netcraft fakeshop	2,653	✓			
Ignore - PhishLabs unauth host	58	✓			
Sent items	39,581				

(a)

Database table	Count	Missing
`ticket`	2,350,168	6,685
`ticket_post`	2,815,503	6,721
`client*`	31,389	0
`devices`	2,023	0
`devices_events`	488,516	0
`ip_assignments`	8,407	23,468
`packages`	63,193	0

(b) *Including 14,644 accounts registered with the same email address – see *Clients*.

B. *Operational Database*

The other dataset we use is the back-end SQL database of a customer relation management (CRM) system that the company employed to manage its operations. It was copied bit-wise during the raid by law enforcement, resulting in a 101GB database dump. Based on our research questions and in close collaboration with the involved law enforcement officers, it became clear that the most valuable information for our study would be in the *ticket* table. Following the foreign keys within this table led to six other tables containing additional information. As a result, we were granted access to a limited set of seven tables. These are listed in Table IIb and detailed in the following paragraphs.

Tickets: Two tables enable communication with clients through tickets. Such tickets are related to password reset requests, overdue invoices, and abuse reports. The *ticket* table contains 42 columns, including a ticket identifier, client ID (when a ticket is created by a client), timestamp, creator email address, subject, and a message. It also includes the origin of a ticket, as they can be automatically created by e-mails directed towards a set of e-mail addresses (i.e., {abuse, billing, info}@company.com) or can be made by the company itself (i.e., because of late payments). Both the author of the ticket, the involved client, and employees can reply to a ticket. Those reactions are stored in a separate table, *ticket_post*. For our analysis, we joined the *ticket* and *ticket_post* tables to obtain an overview of all tickets and their responses – i.e., each ticket has one or more *posts* attached to it. As shown in Table IIb, we found over 2.3M tickets and identified 6,685 missing tickets (0.3%) thanks to missing auto-incremented ticket IDs. These records were divided randomly in the database, and we found no evidence of record deletion within specific time periods.

Clients: The *client* table contains 57 columns with personal and billing details. We obtained access to only the client ID, email, and registration date to avoid analyzing personally identifiable information (PII), All other columns were removed before access was granted. A total of 31,389 clients were found in this table, with no missing values. Manual inspection of the email addresses used for registration revealed that someone registered 14,644 new accounts using the same email address in 2014. As we found no abuse related to these accounts, we removed them from our analysis in subsequent sections.

IP Assignments & Devices: The remaining four tables can be used to determine the ownership of devices within the company's data center and IP address assignments to clients over time. First, two tables capture the most up-to-date state of all devices and IP assignments. The *devices* table lists all devices and contains 34 columns related to, among others, the location of every device in the data center, its status, and the client associated with it. Upon removal of a device, records are not deleted but nulled. As a result, we found 2,023 devices in this table, with no missing rows. IP addresses assigned to devices at the moment of the raid are stored in the table *ip_assignments*, which does not store any past assignments. When an IP assignment is removed, it is deleted from this table – which explains the 23,468 missing rows listed in Table IIb. Since IP addresses are assigned to specific devices within the data center, we can leverage the *device* table to find which IP address was assigned to which device owned by whom. Again, just the final state is stored in this table. A combination of two other tables is necessary to gather information on historic IP assignments: *device_events* and *packages*. The *device_events* logs every update made to devices – ranging from client changes to power disruptions – in an append-only log containing the before and after states. A total of 488,516 records, without any missing rows, were found. Lastly, some devices are shared by multiple users, each using a different IP address on one Virtual Private Server (VPS). VPS access is offered as a package, and records related to those are stored in the *packages* table, which contains information on both the client using it and the technical management of these packages.

V. Methodology

The following section details our approach to collecting and categorizing abuse reports and client notifications from the aforementioned data sources, depicted in Figure 1.

Collecting Abuse Reports: We extract abuse reports from both the CRM database and the two abuse mailboxes. As mentioned in Section IV-B, abuse reports directed towards a set of mailboxes are automatically stored in the database

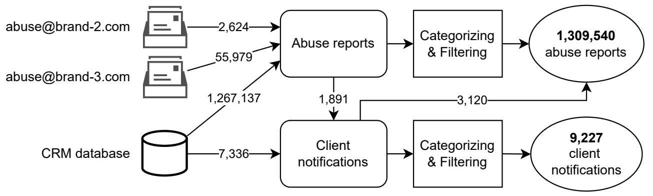

Fig. 1: Overview of the use and processing of data sources.

TABLE III: Dataset descriptives.

Dataset start	August 1, 2011
Dataset end (raid by LEA)	September 22, 2020
Abuse reports	1,309,540
Abuse reports linked to a notification	34,240
Client notifications	9,227
Client notification rate	2.6%
Abuse report senders	9,594
Registered clients	16,030
Abusive clients	3,114

as tickets. To identify and extract those tickets, we selected all tickets originating from outside the company and filtered for tickets directed towards an '@abuse' e-mail address or containing the word 'abuse' in the subject header. This resulted in a set of 1,546,217 tickets, from which we removed 279,080 tickets containing spam emails, newsletters, and failed email delivery attempts, resulting in a set of 1,267,137 tickets. By analyzing the two abuse mailboxes, we obtained another 58,421 abuse reports in the period 02-2019 until 09-2020, which we analyzed similarly. The two data sources – database and mailboxes – were merged for further analysis, thereby omitting the 37 abuse reports that were present in both data sources. Duplicates were identified by comparing abuse reports with a similar subject header, originating from the same sender, and dated within a two-hour time frame of each other.

Categorizing Abuse Reports: As we had to perform our analyses on-premise at law enforcement agencies in secure environments, we adopted a simple keyword-based approach to categorize the collected abuse reports based on their subject, author, and message body. Following the categories of abuse defined by the M^{3}AAWG (listed in Table I), we brainstormed to gather an initial list of five keywords per category. For example, we selected the words *phish*, *malware*, *trojan*, *bruteforce*, and *brute force* as our initial set of keywords for the category encompassing malware, phishing, and brute-force attack abuse reports. Then, we iteratively gathered more related keywords. We applied the initial list of keywords to the abuse reports, manually inspected the top remaining uncategorized reports, sorted by reports per sender (to affect the most significant portion of unlabeled reports), defined the label it should have, and identified new keywords that would only include these yet unlabeled reports. Through this process, we also added the names of organizations focusing on specific types of abuse, such as 'INHOPE', 'IWF', and 'CyberTip' (related to CSAM) or 'SpamCop' (related to spam). Additionally, for certain types of abuse, we added abbreviations, such as 'DMCA' (copyright and trademark issues), or specific subject headers ('clean-mx-trackback' for comment spamming). After each iteration, we selected a random sample of reports per category (from a list aggregated by the sender), checked for errors, and fine-tuned the keywords accordingly. This process was repeated ten times, resulting in the list of words that can be found in Appendix A. Abuse reports can be assigned multiple categories (0.9%) or not assigned any category and labeled as 'Unknown' (3.7%). Comparing 100 randomly sampled and manually labeled abuse reports with the assigned category revealed that our keyword-based approach correctly categorized 95% of them.

Determining Client IP Assignments: To match abuse reports to their responsible clients, a mapping of which clients are using which IP addresses is required at any moment during our measurement period. To gather this information, we leverage the last four tables mentioned in Section IV-B in the following order: *ip_assignments*, *devices* & *devices_events*, and *packages*. We apply a three-step approach to find the associated client based on the trustworthiness of the data sources. First, we search for an active IP assignment for any mentioned IP address at the moment of the incoming abuse report by using the *ip_assignments* table. If this yields no results, we search for historic IP assignments listed in the *device_events* table, combined with the *devices* table, to find the client associated with a historical IP assignment. If neither search yields any results, we use the *packages* table to retrieve a client. However, as this table does not contain an end date for IP address assignments – which are tied to a contract with a set duration, e.g., a month – we consider this data source the least trustworthy. From 87% of these tickets, we could extract a company IPv4 address, which we were able to link to a client in 99.7% of the cases. For unknown reasons, we could not determine the corresponding client for 493 abuse reports containing a valid IP address

Identifying Provider Action: Anti-abuse actions, such as client notifications, are stored in the CRM database in two ways. Either as a post linked to the abuse report ticket or as a separate ticket created by the company itself. Client notifications are not stored in either of the two abuse mailboxes, as a manual investigation of the *sent items* folder in both mailboxes revealed that no client notifications were sent directly from these mailboxes. That folder contained solely replies to abuse reporters. The 1,891 client notifications in the database, linked through posts, are easily matched based on their corresponding ticket identifiers. Matching the separately created client notification tickets is less trivial. To match these to their originating abuse reports, we selected all company-initiated tickets that did not contain a set of 25 keywords related to billing, orders, and maintenance in the subject header. This set of 7,336 tickets was enriched similarly to the abuse reports, extracting the timestamp, category, and IP addresses. Since these tickets are directed toward a client, the corresponding client identifier is always present. We use the timestamp and the mentioned IP addresses in the client notification ticket to search for its underlying abuse report(s).

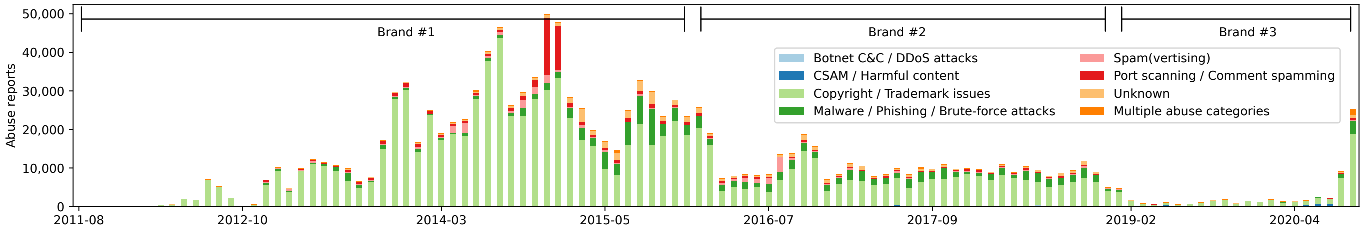

Fig. 2: The number of monthly abuse reports divided by abuse category.

For every client notification, we searched for abuse reports that mention the same IP address within a 336-hour time frame – i.e., 2 weeks – prior to the client notification and did not yet contain a (linked) client notification. This time frame was chosen after a manual inspection revealed that such long time frames were not uncommon. As a result, one client notification can be matched to multiple received abuse reports. We consider this a valid method, as widespread abuse can trigger multiple abuse reports, and it would make sense for a company to group those reports into a single client notification. This is illustrated by a case in which we were able to match one client notification with 1,575 automated spam reports, all received on the same day. Through this process, we were able to match 4,216 (57%) of the separately created client notifications to their corresponding initiating abuse report(s), hereby linking a client notification to 29,229 abuse reports. A notification is linked to 3.98 abuse reports on average, yet the median is 1. Abuse reports could also have originated from phone calls, as one response to an abuse report highlights: *'for urgent cases, please call us, we noticed your email 9h later'*. Hence, we added the remaining 3,120 client notifications, for which we could not find the originating abuse report, to the dataset as well, including every client notification ever sent by the company in our final dataset. Throughout our approach to measuring abuse handling, we apply a conservative take. That is, when in doubt, we assume the company notified the client.

To obtain our final dataset, we removed all uncategorized abuse reports that did not contain an IP address – 16,438 tickets, mostly spam – and obtained our final dataset of 1,309,540 abuse reports. 87% of these reports contained a company IPv4 address, which enabled us to assign 1,120,201 (86%) abuse reports to 3,114 different clients. Matching client notifications to abuse reports and the linked notifications within tickets enabled us to identify 9,227 distinct client notifications associated with 34,240 abuse reports. Table III lists the final dataset to be used in the remainder of this paper.

VI. Quantifying Abuse & Provider Actions

In this section, we quantify and characterize the abuse reports, categorizing them by frequency and origin. Next, we present insights into the actions taken by the company to combat abuse, addressing our first two research questions.

A. Abuse Follow-up

To answer our first research question, we scrutinize the company's 1,309,540 abuse reports received within the nine years spanning our dataset. An overview of the monthly received abuse reports is depicted in Figure 2. It shows roughly three periods of abuse volumes that line up with changes in brand names. For the first period (until 2016), we observed a slow increase in monthly abuse reports – growing from around 10k in 2013 to almost 50k by the end of 2014. Most of them are related to hosting copyrighted material, but over time, these reports also start to include reports in other categories. During the second period (2015 – 2019), when operating under the second brand name, the number of abuse reports stabilized at around 10k monthly reports. This is primarily due to a decrease in copyright-related abuse reports; the other categories have similar abuse report counts within this time frame. Introducing the third brand name in late 2018 has likely changed the abuse-handling process. We believe that abuse reports are no longer automatically converted into tickets, but are instead handled within the abuse mailbox. This is because we did not find many abuse reports in both the CRM database and in one of the mailboxes (only 37). It is likely that reports have been deleted from the mailboxes during this time frame, and that the volume of abuse reports was higher. As in the last month for which we obtained abuse reports from the abuse mailboxes (September 2020), the number of reports returned to the monthly average of over 20k per month.

As listed in the first row of Table IV, copyright-related abuse makes up the largest portion of reports. Almost 77% of received abuse reports are categorized as such, which involves hosting torrent websites and illegal live sports streams. Other prevalent categories are (comment) spam(ing) and malware, phishing, and brute-force attacks. Reports related to ongoing DDoS attacks, botnet C&C servers, and CSAM are less prevalent. Reporters of copyright-related abuse do so very often, with an average of 609 reports per reporter. This is much lower for all other categories of abuse, which range from eight for botnet C&C servers and DDoS attacks to 40 for port scanning and comment spamming. In total, we found 3,114 clients to be involved in one or more abuse cases. The majority of abusive clients can be found in the malware, phishing, and brute-force category (1,980), followed by the port scanning and comment spamming category (1,575).

The next rows in Table IV report on the client notifications

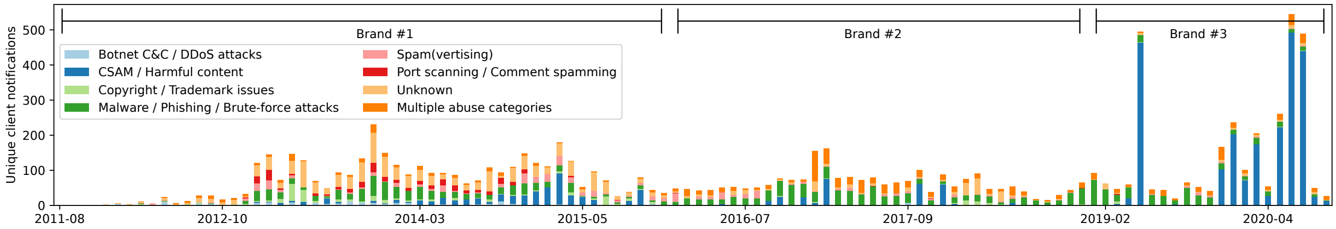

Fig. 3: The number of unique monthly client notifications divided by abuse category.

by the company, presenting both the number of reports with a linked notification and the number of unique client notifications. The latter is also depicted in Figure 3, which shows a very low number of copyright-related client notifications and a high number of CSAM-related client notifications, especially in the last period. A total of 9,227 client notifications have been sent, which we could link to 34,240 abuse reports. Since abuse reports can be assigned to multiple categories, the client notification counts within this row add up to more than the total amount – we discuss the effects of this in Section VIII. The overall client notification rate – i.e., the fraction of abuse reports linked to a client notification – is 2.6%. If we remove the largest category – Copyright & Trademark issues – from this statistic, this number is 9.73%. There are significant differences in client notification rates between the categories, and these rates also change over time. For example, we identified 867 client notifications for copyright-related abuse, linked to 4,200 abuse reports, resulting in a 0.4% notification rate. In contrast, 45.6% of all CSAM-related reports are linked to a client notification. Remarkably, the client notification rate for spam(vertising) (19%) is much higher than for malware, phishing, and brute-force attacks (8.6%), whilst their place in the prioritization according to the M^{3}AAWG in Table I would suggest the inverse. To investigate abuse handling over time, we plotted the client notification rates per year for each abuse category in Figure 4. It shows that the few botnet servers and DDoS abuse reports in 2012 were met with a 68%-client notification rate, which decreased in the years afterward. Client notifications originating from abuse reports related to (comment) spam increased slowly from 8% in 2011 to almost 40% in 2016 and decreased in the years afterward. Other categories remained at stable client notification rates below 20%. The years 2019 and 2020 marked a significant change for CSAM-related abuse reports, as the client notification rate increased to 66% and 85%, respectively.

Takeaway: Abuse reports related to copyright and trademark issues are most common (77%), followed by malware, phishing and brute-force attacks. Abuse reporters in the first category send out large numbers of reports (over 600 per reporter), whereas reporters in the other categories share fewer (9 - 40 reports per reporter). CSAM-related abuse reports are followed up on the most (45%), followed by spam(vertising) (19%). Notification rates fluctuate over the years.

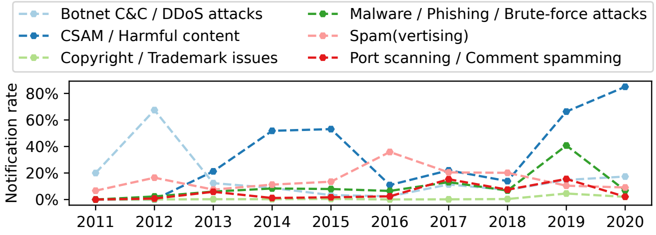

Fig. 4: Yearly notification rates for the different categories.

Ticket conversation on 17/09/2016

Client: On my website, I will stream live sports events. About DMCA, can you confirm that is no problem?

Company: DMCA is USA law. We are not in the USA.

Fig. 5: Conversation regarding the DMCA.

B. Abuse Follow-up Factors

To understand what factors influence client notification rates, we take an in-depth look at each abuse category in increasing order of priority according to the M^{3}AAWG . Like abuse reports, most notifications are identically structured due to the use of templates. To gather more insights into the rationale of the company's operators, we searched for notifications containing non-standard messages or conversations with its customers. Some of them, which characterize the company's stance on abuse, are depicted in various figures as anonymized, translated quotes, some of which edited to improve readability.

Copyright & Trademark Issues: A total of 1,000,730 copyright and trademark-related abuse reports have been received during our measurement period. These reports originated from 1,646 unique reporters; some reporters sent over 100k reports each. Among them are predominantly companies providing content security, anti-piracy services, and Cloudflare. The latter forwards received abuse reports, whereas the others filed the reports themselves. Many of the anti-piracy companies received no response to their reports ever. An example is the now-defunct NetResult, a DMCA takedown service. It has filed 126,085 reports, of which only 411 could be linked to 4 client notifications. When client notification does occur, it happens slowly – as we found a median time of 123.7

TABLE IV: Overview of abuse reports per category, the involved clients and reporters, and corresponding notification rates.

	Botnet C&C DDoS	**CSAM Harmful content**	**Copyright Trademark issues**	**Malware / Phishing Brute-force attacks**	**Spam (vertising)**	**Port scanning Comment spamming**	**Unknown**	**Total**
Abuse reports	9,480	9,247	1,000,730	146,572	40,118	68,546	48,167	1,309,540
Reports with notification	1,050	4,213	4,200	12,614	7,640	2,647	3,621	34,240
Unique reporters	1,150	310	1,646	3,720	1,037	1,696	2,851	9,594
Avg. reports per reporter	8.2	29.8	609.5	39.5	39.0	40.5	16.9	136.8
Client notifications	757	3,125	867	2,862	1,436	1,046	1,587	9,227
Client notification rate (%)	11.1	45.6	0.4	8.6	19.0	3.9	7.5	2.6
Median time to notify (h)	27.5	0.0	123.7	67.7	34.0	46.6	44.5	48.0
Clients involved	1,221	334	832	1,980	907	1,575	1,664	3,114
Avg. reports per client	8.3	36.3	999.5	73.2	42.8	44.0	29.6	367.0

hours (5 days) to notify. Examining the list of abuse reporters whose reports led to notifications, we find several lawyers and BREIN, the private Dutch copyright watchdog. Their reports are not ignored, as multiple client notifications demand immediate action because *'we cannot afford problems with these people'*, according to the company. Clients have only 12 to 24 hours to mitigate these reported issues, and the company threatens to suspend servers if they fail to do so. However, in some cases, the company thinks along with clients to allow their operations to continue. For example, after repeated requests from copyright holders acting for the English Premier League (FAPL), it emails a client to ask permission *'to change the IP address of this server, more offshore due to some issues'* instead of demanding them to take down services. A practice they later had to reconsider once legal action was taken by the FAPL in 2018. Another option it offered was to host 'streaming relay' servers, which do not host any copyrighted material but merely relay it. The company seems willing to facilitate the streaming of any kind of material, as Figure 5 illustrates. Cloudflare, operating its distributed reverse-proxy service, forwarded over 30k abuse reports, 1,247 of them linked to client notifications. Although this is more than any other reporter in this category, we found that client notifications depended on the source of the original report, and not because it originated from Cloudflare. For example, abuse reported by BREIN to Cloudflare is followed up on, whereas other reporters are ignored. The discovery of a folder called *Ignore - Cloudflare* in both mailboxes (see Table IIa) underlines this finding. The company's stance on copyright-related abuse has changed over time. Although client notification rates remain very low, as depicted in Figure 4, it has requested that clients streaming copyrighted material include a takedown tool on their websites since 2015, thereby transferring future abuse reports directly to the clients. Removing itself from the abuse reporting chain – i.e., reporters communicate directly with the client – explains the decrease in copyright-related abuse reports since late 2015.

Takeaway: Enormous amounts of copyright-related abuse reports are ignored, except when there is a threat of legal action from piracy watchdogs like BREIN or lawyers. A compulsory takedown tool for clients operating streaming services reduced the volume of abuse reports by orders of magnitude.

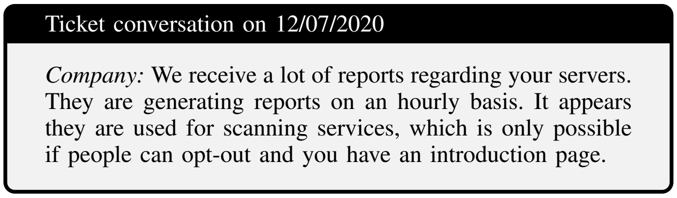

Ticket conversation on 12/07/2020

Company: We receive a lot of reports regarding your servers. They are generating reports on an hourly basis. It appears they are used for scanning services, which is only possible if people can opt-out and you have an introduction page.

Fig. 6: Conversation regarding scanning.

Port Scanning & Comment Spamming: This category comprises 68,546 abuse reports, most of which are related to port scanning. We can associate 1,575 clients with these complaints, with an average of 40.5 reports per client. The majority of reports are very concentrated on a few clients. While one client was responsible for 29,623 abuse reports within two years, this does not seem to have influenced notification rates. Only in the rare case that one client receives an extraordinary amount of reports in a short period of time – as shown in Figure 6 and depicted by the spike in port scanning reports in Figure 2 between 2024 and 2025 – the company does notify. The second-highest number of 1,696 unique abuse reporters stands out because of the many automated abuse reports within this category. Among the top reporters are honeypot operators, intrusion detection systems, and data center network operators, who automatically file abuse reports after a port scan is detected on their servers. Only some of these reports resulted in client notification (3.9%). Spamhaus, the blocklist operator known for its fight against email spam, also lists illicit vulnerability scanners and comment spamming IPs in its blocklist. Such listings do trigger client notifications in 83% of the cases. Another type of abuse within this category is trackback abuse, a form of comment spamming. Blogging systems like WordPress enable the notification of new content on other blogs, which is often abused by spam websites to promote their own content. One honeypot operator monitors abusive trackbacks and reports automatically, which was done 30,467 times without any response or notification.

Takeaway: Numerous abuse reports regarding port scanning and comment spamming from unvetted, automated systems – i.e., Fail2Ban – do not result in many client notifications. Abuse reports from Spamhaus do trigger frequent notifications.

Ticket conversation on 25/12/2011

Company: We hosted your website for a long time. Spamhaus listed your IPs, and we didn't care. After that, they listed a few /23 blocks, now they listed all our IPs because we didn't take any actions. This affects our whole network, thousands of people can't email because of you.

Fig. 7: Conversation regarding Spamhaus listings.

Spam(vertising): We collected a total of 40,118 abuse reports related to spam(vertising), concentrated on several clients operating as resellers. Resellers, as mentioned in earlier work [9], [10], are frequent clients, especially in the category spam(vertising), and resell rented infrastructure to their clients. In doing so, they introduce another intermediary in the abuse notification chain. This is illustrated by one ticket: *'we are resellers, we can't control every client, and we didn't notice all the recent reports'*. Spam-related abuse reports received the second-highest client notification rate of 19%. 1,037 reporters have filed reports regarding spam activity, yet only one reporter received significant follow-up: Spamhaus. The Spamhaus Block List (SBL) contains IP addresses with known spamming activity [33]. As soon as an IP address is listed, the owner of the IP range is notified. If spamming is not handled within a certain period, Spamhaus can escalate the listing to block extended ranges – e.g., a \24 range – or eventually list the entire network of the involved AS. This happened several times, as we learned from the ticket conversations in Figure 7. Such 'escalation listings' bother the company because clients complain that their emails can not be sent or demand new IP addresses outside the listed ranges. As a result, SBL listings are handled swiftly, and temporary solutions are offered, such as email relays through non-blocklisted IP addresses. As a result, 60% of its reports lead to client notifications. Another party that received significant follow-up from its reports is Level 3, a peering partner. Individuals who encounter activities that violate Level 3's acceptable user policy can file a report, which Level 3 forwards to the network operator. The company, possibly afraid to lose connectivity, created 93 client notifications based on Level 3 reports, of which at least 35 were related to spam(vertising).

Takeaway: Although categorized as a low priority by the M^{3}AAWG , spam-related abuse is met with the second-highest client notification rate due to sanctioning by Spamhaus.

Malware, Phishing & Brute-Force Attacks: This category is the second-largest category of abuse within our dataset, totaling 146,572 abuse reports from 3,720 different reporters. This category also involves the most clients, namely 1,980, with 39.5 abuse reports on average. The most abusive client has gathered over 18k reports and has a long business relationship with the company as a reseller offering offshore VPSes. Despite the many abuse reports, this client has never been terminated. The second client on this list, amassing 3,764 abuse reports, is another reseller offering unmanaged VPSes and was threatened with termination. After receiving numerous abuse reports and forwarding only a handful of them, the company decided to terminate all its servers, as shown in Figure 8. However, after some back-and-forth, business continued as usual. Reports regarding dictionary or brute-force attacks come from the majority of reporters within this category and are often the result of intrusion detection systems with automated abuse reporting. Fail2Ban, a popular system to protect (Web)servers from brute-force attacks, can also automatically send an abuse report to the owner of an IP address after a certain number of failed login attempts. At least 57,562 (39%) abuse reports within this category have been the result of this system. Such reports rarely lead to client notifications. Phishing reports originated predominantly from NetCraft and PhishLabs, both take-down services that vet abuse reports thoroughly and provide detailed information, thereby facilitating swift client notifications. Additionally, services like NetCraft monitor reported phishing pages over time to ensure their takedown. Although the folder names in Table IIa would suggest otherwise, 72% of NetCraft phishing reports resulted in client notifications, and 54% of the PhishLabs reports. Malware reports originate from various sources, including community services and country CERTs, and receive varying notification rates.

Ticket conversation on 21/12/2014

Company: Hello, there is no server with legal content in your account. All your servers are running scans and brute-force attacks. We want to stop hosting cybercrime on our network. You're offering bulletproof hosting to your customers, so your company is exit at us.

Client: Your abuse department never sent us reports regarding these servers, and now everything is offline. We want to continue business and settle this.

Fig. 8: Conversation with an abusive client.

Takeaway: The category with the most linked client notifications shows that vetted, trusted abuse reporters are met with higher client notification rates than automated systems like Fail2Ban or individual reporters.

Botnet C&C & DDoS Attacks: We identified a total of 9,480 abuse reports related to Botnet C&C servers and DDoS attacks within our measurement period. A total of 757 client notifications were sent, which we could link to 1,050 abuse reports, resulting in an 11.7% client notification rate. The median time to notify is the second-lowest, namely 27.5 hours, which seems in line with the priority assigned by the M^{3}AAWG [2]. Abuse reports originated from 1,150 different reporters – many filing only a single report – and are evenly distributed between DDoS attacks and botnet C&Cs. Among the top reporters of botnet C&C servers are Spamhaus, a botnet researcher, and a Dutch SOC. Unlike the name suggests, Spamhaus also fights botnets by operating its Botnet Controller List (BCL) [33].

Ticket conversation on 07/01/2012

Company: Hello, your server is sending DDoS attacks for a few hours now. We do not allow our network to be abused for DDoS. To protect other clients, we limited your port speed.

Fig. 9: Conversation about an ongoing DDoS attack.

Similar to the reports related to spam(vertising), the influence of Spamhaus is evident, as 84% of its abuse reports were met with swift client notifications. Clients get just six hours to resolve reported issues and are automatically suspended if there is no immediate reaction. For DDoS reports, there are no reporters who file significant amounts of reports. Most reports originate directly from victims of DDoS attacks when they are attacked by one of the company's servers. The use of automation in abuse reporting causes noise for abuse-handling departments. An example of this is an automated DDoS reporting system that sends out the same abuse report every 15 seconds. Unlike other categories, the company also detects DDoS attacks itself through its data center monitoring systems. When a high volume of outgoing packets is detected – e.g., a client sending spoofed packets – the company steps in and notifies the resource owner since such volumes could damage their network – as depicted in Figure 9. From Figures 3 and 4, we learn that this happened frequently between 2012 - 2015, and diminished in the years after. Four clients were terminated due to DDoS-related abuse, the only category in which we identified client terminations.

Takeaway: DDoS is the only type of abuse predominantly reported by direct victims. Outgoing DDoS attacks harm the company's network by affecting the connectivity of other clients and are, therefore, quickly addressed. Botnet C&C servers listed by Spamhaus are removed rapidly as well.

CSAM & Harmful Content: On top of the M^{3}AAWG priority scheme, we find CSAM and harmful content. We identified a total of 9,247 such abuse reports within our measurement period, associated with 334 clients, having 29.8 reports on average. 3,125 client notifications were created, which we could link to 4,213 abuse reports. Abuse reports originated from 310 unique reporters. Among the very active ones are national hotlines that cooperate within the InHope network, such as the British IWF and the Dutch Meldpunt Kinderporno. Their reports were taken care of to a certain degree (notification rates of 16% and 24%, respectively), whereas reports from individual reporters received almost no follow-up. The clients associated with the reported abuse hosted either forum boards or operated image hosting services. In both cases, clients are given a maximum of 24 hours to handle reports. For example, one client operating multiple image hosting websites is responsible for 585 CSAM-related abuse reports in four years. Most of these reports resulted in client notification and swift action from the affected client. However, after four years of abuse reports, law enforcement stepped in and forced the company to shut down this website. Another client received 382 reports and operated multiple forums from 2017 until 2020. From 2018 onward, all reports related to these forums resulted in client notifications, followed by the deletion of files by the client. In 2020, with the Dutch Justice Department putting more pressure on bad hosting companies [34], the company suggested stopping business with this client, as depicted in Figure 10. Ultimately, no client was ever terminated due to CSAM-related abuse. Government pressure likely resulted in the launch of a website to process takedown requests operated by the company. The effects of this platform are significant, as 94% of the reports filed through this platform resulted in a swift client notification – which explains the median time to notification of 0 hours in Table IV and the increase in client notification rates in Figure 4. Many notifiers successfully utilized it, except for CyberTip, a Canadian initiative aimed at combating CSAM. After repeated messages, the company sets up an autoresponder to instruct CyberTip to use their platform instead of emailing their reports. Despite the 38,425 sent autoresponses (discovered in the *Sent items* folder in one of the mailboxes), CyberTip never did so. In communication with clients, the company's stance is clear: it only takes action when certain parties, such as law enforcement, demand it. They explicitly state this as an excuse to their clients – e.g., *'please understand we only sent you this because authorities demand us, we don't want to play judge ourselves'*.

Ticket conversation on 14/07/2012

Company: We can better stop working together. Have a look at your ticket history, you have had over 200 CSAM reports this year, that is way too much. The government is pushing very hard on us to fix this.

Fig. 10: Conversation regarding CSAM reports.

TABLE V: Overview of the hosters included in this analysis.

Hoster	Announc.	IP count	Short announc.	IP churn
The Company	295	27,392	43.0%	28.4%
Bad hoster #1	482	25,344	63.1%	36.0%
Bad hoster #2	49	16,384	45.5%	42.0%
Bad hoster #3	484	179,456	57.3%	45.2%
Good hoster #1	97	144,896	50.6%	1.2%
Good hoster #2	123	19,456	47.1%	43.4%
Good hoster #3	200	16,128	81.3%	20.1%

Takeaway: Most CSAM-related reports are met with swift response. Government actions and trusted notifiers with access to automated takedown portals have an effect, as client notification rates increased massively in 2019 and 2020.

VII. External Network Characteristics & Abuse

We now take an external look to answer our third research question and relate external network indicators found in previous work [8], [9], [12] to internal abuse handling by scrutinizing historical IPv4 prefix announcements. Additionally, we compare these results to those of other bad hosting providers as well as reputable hosting providers.

Methodology: First, we selected three other bad hosting companies with similar IP counts and listed in the same top 50 of bad hosting companies [14] as the company we studied. Next, we selected three reputed Dutch hosting companies with similar IP counts that were active during the same period. For each company, we collect historical IPv4 prefix announcements by the ASes associated with that company within the same time frame (2011 – 2020) by leveraging the RIPE NCC announced prefixes API [35]. While collecting this data, we noticed an abnormally large IPv4 range being announced by the company we studied in this paper. These 262,144 IPv4 addresses were part of the AFRINIC heist [36]. Since we found only 10 abuse reports related to IPs in this large range, and their existence is disputed, we excluded them from further analyses. The other six companies did not have such abnormalities. The resulting dataset consists of 295 prefix announcements for the company we studied, and between 49 and 484 prefix announcements for the other companies. An overview of the characteristics of all the hosters included in this analysis can be found in Table V

Results: In Figure 11, we plotted the number of announced IPv4 prefixes over time per AS for the examined company. We observe a slowly increasing number of 13,000 to over 17,000 announced IPv4 addresses in the period 2011 - 2015, except for two quarters in 2021, which we consider erroneous data. In 2015, the company underwent its first rebranding and relocated its registration to the Seychelles [15], coinciding with significant changes in IPv4 prefix announcements. However, its second brand – which had been active since late 2011 on a different AS – stopped announcing IPv4 prefixes from 2016 to 2019. In early 2019, after another rebranding, all IPv4 prefixes were transferred from Brand #1 to Brand #3. At the same time, the second brand started announcing a few IPv4 prefixes again. Comparing this with the other hosting companies included in this analysis, both the good and the bad hosters, shows similar behavior; all companies are slowly growing in advertised IPv4 size over the years. Konte et al. [8] found (very) short prefix announcements and IP churn to be indicators for malicious networks. We found that 43% of the company's prefix announcements last less than half a year, as shown in the fourth column of Table V. However, this rate seems to be quite average when compared to both good and bad hosters. For this analysis, we only included the prefix announcements from 2013 onward to account for the missing data in 2012. Interestingly, the short prefix announcement rate is above 50% for the first two brands operated by the studied company, whereas the last brand has a significantly lower number of prefix announcements lasting less than half a year (28%). For the company we studied, we found a 28.4% IP churn rate, which also did not deviate significantly from the IP churn rates of both the good and bad hosters in our analysis, as shown in the last column of Table V. However, when zooming in on the yearly additions and deletions of advertised IP space, notable differences between the good and the bad hosters emerge, as shown in Figure 14. Here, we observe that good hosters remove only a very limited number of IPs from their prefix announcements per year, whereas bad hosters – including the company we studied – add and remove significant parts of their IP prefix announcements on a yearly basis. Following the findings of Alrwais et al. [9], this could be the result of rotating IP blocks to evade blocklisting, as we have seen anecdotal evidence for in Section VI. IP space fragmentation, another indicator for malicious networks as identified by Konte et al. [8], is depicted per hoster in Figure 12. Here, we see that over 80% of the announced IPv4 prefixes are indeed small /24 IP ranges. This number is similar to two of the three bad hosters in this analysis, yet also similar to two of the three good hosters.

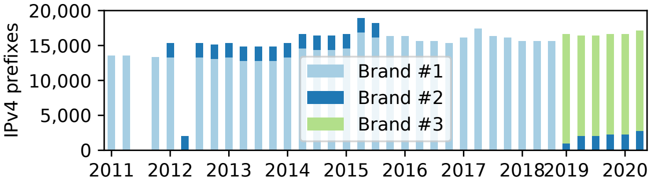

Fig. 11: Quarterly IPv4 prefixes per AS.

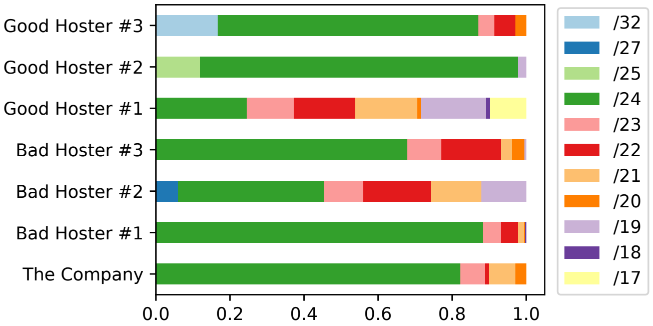

Fig. 12: IP space fragmentation per hoster.

Takeaway: Some indicators for malicious networks are observed at this company, such as frequent rebrands and IP space fragmentation. The yearly additions and removals of advertised IP space resulting in short prefix announcements could indicate evasive actions to prevent blocking. However, IP space fragmentation, IP churn rate, and short prefix announcements are also observed at reputable hosting providers.

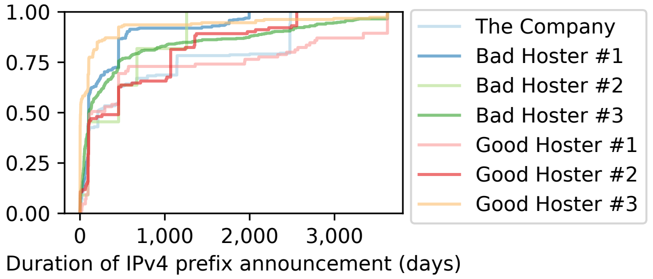

Fig. 13: CDF of IPv4 prefix announcements durations.

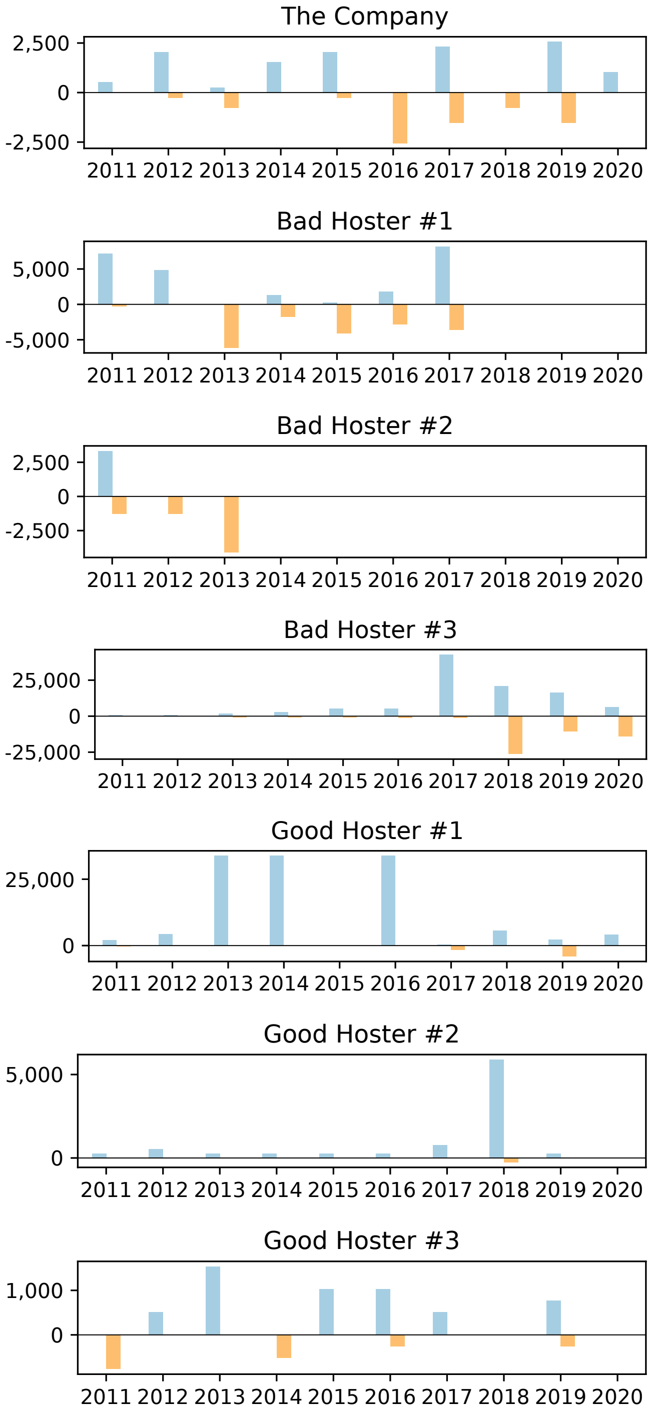

Fig. 14: Yearly IPv4 prefixes additions and removals for each hosting company.

VIII. Discussion

In this section, we revisit the term bulletproof hosting, discuss the public policy takeaways of our findings, and elaborate on the inherent limitations that arise from our work.

Bulletproof hosting: The term 'bulletproof' was first coined by industry reports [11] and later found its way into academic work [8], [27] to describe hosting providers that systematically and intentionally ignore abuse reports. The analyses of such work are dominated by external viewpoints. We argue that bulletproof hosting – i.e., a behavioral pattern of purposely ignoring abuse reports – cannot be deduced solely from an external perspective, as it is impossible to measure intent without knowing the internal abuse-handling processes. Yet, these external measurements do have value. Measuring takedown rates of abusive content does indicate how willing a hosting provider is to fight abuse. However, without an inside look, it remains unknown to what extent neglecting abuse reports is a result of an intermediary not notifying its clients or clients not removing their abusive content. In our case, the studied company is often referred to as a bulletproof hoster by both law enforcement [15] and researchers [12]. Through similar analyses as performed by earlier work [8] in Section VII, we do see indicators of malicious networks as well. However, our analysis of its abuse-handling processes shows that it was not immune to abuse reports, as it did act upon a portion of the received reports, albeit a small one. Although some may question the morality of the company's decisions, our analysis and many of the tickets do not indicate an upfront intent to enable abuse. We do see, however, that it puts an absolute minimum effort into anti-abuse actions and solely prioritizes minimizing negative business effects. As a result, Spamhaus listings (which could harm client connectivity), Level3 reports (which could lead to de-peering), and CSAM-related abuse (with legally binding consequences) are met with swift client notifications. In contrast, individual phishing, spam, or port scanning reports are not. Hence, our analyses show that the term 'bulletproof,' when relying solely on external measurements, is a tough label to sell. The term implies intent, which can only be accurately captured through an insider's view.

Public Policy Takeaways: Here, we reflect on our case-study findings and report on public policy takeaways that apply to the anti-abuse ecosystem as a whole. From Section VI we learn that certain instruments lead to abuse follow-up, whereas others do not. Some reporters within the industry have gained significant power and have thereby obtained de facto trusted reporter status. Abuse reports from these trusted reporters, such as CSAM hotlines, or reporters who can pressure the company into taking action, like escalated Spamhaus blocklisting, result in more client notifications than individual reporters. This implies that individual abuse reporting, either manual (e.g., after receiving a phishing email) or automated (e.g., Fail2Ban abuse reports), seems less effective. Abuse reports originating from automated reporting systems operated by individual networks result in many similar yet unstructured and less detailed abuse reports that are more likely to pollute abuse mailboxes than assist abuse-handling personnel. Security practitioners who want to report abuse could, therefore, consider reporting to trusted or powerful reporters instead. Our analyses have shown this to be the case in the instances of phishing (NetCraft), spam (Spamhaus), and CSAM (InHope hotlines). Recent E.U. regulations to institutionalize and appoint 'trusted flaggers' through the DSA [3] seem a deliberate action to make the Internet safer. However, we question why these trusted flaggers are appointed on a national level. It makes sense to appoint trusted flaggers for copyright-related abuse on a per-member-state basis, as copyright laws differ per jurisdiction. However, to fight spam or malware-related abuse, there is no need for 27 different nationally appointed trusted flaggers performing similar work as Spamhaus is currently doing. Here, European or even worldwide trusted flaggers would seem more effective.

Limitations: First, our study is focused on a single hosting provider that has encountered friction with law enforcement. This inherently limits the generalizability of our findings across hosting providers. If the company has been taken down, our case study might suffer from a version of survivorship bias, since takedowns are an exception, where most bad providers persist. In our case, however, the company is still functioning and a part of the hosting market. Although our internal data is historical (2011 – 2020), the company seems to persist in its disreputable abuse handling, reflected by its listing on prominent blocklists like the Spamhaus DROP list [17]. Yet, we are very reluctant to generalize. Our goal was to better understand whether and how the governance mechanisms of the anti-abuse ecosystem function in the context of a 'bulletproof' provider that was explicitly willing to condone abuse on its network. We do not argue that the observed numbers are generalizable to the entire hosting industry. Furthermore, our case study provides a rare ground-truth case to inform research on methods for identifying 'bulletproof' or bad hosting providers.

Second, we encountered missing or deleted data. As listed in Table IIb, 6,685 tickets (0.3%) have been deleted from the database, randomly distributed over time. Given the small size of this fraction, we do not believe it could significantly impact the results. We suspect that emails related to abuse reports were deleted from the two mailboxes used for handling such reports. This could be due to a change in abuse handling processes, as we expect the company to handle at least part of the abuse reports directly from the mailbox, rather than from their ticketing system, as they did previously. The 3,120 client notifications that we were unable to match to an abuse report represent an indication of these missing abuse reports. However, by adding them to the dataset and treating them as notifications with a corresponding abuse report, we prevented these deletions from impacting our findings. Finally, some abuse reports, as well as client notifications, may have never been included in our dataset. Some of them could have been handled by phone, as we have seen anecdotal evidence for this in Section V. Given the size of the operations and the limited number of staff, we do not think it is likely that a significant fraction of cases were handled outside of the main systems.

Third, our method of categorizing abuse reports is straightforward. This rigid categorization is not always correct and could have influenced our results. For example, from 2017 onwards, the company notified abusive clients 220 times via emails with the subject header `URGENT: MALWARE / PHISHING / SCANS / SPAM`, resulting in notifications matching all of these categories. However, as only 0.9% of the abuse reports were assigned to multiple categories, we consider this impact to be low.

IX. Conclusions

This study empirically investigated internal abuse data of a hoster with a reputation for abuse to study governance instruments in the anti-abuse ecosystem. Analysis of 1.3M abuse reports and 9,227 client notifications showed large differences in client notification rates among abuse reporters and categories. CSAM and spam-related reports result in client notifications, whereas reports regarding copyright and port scanning have low client notification rates. We find that reporters with either a trusted status – e.g., NetCraft, InHope hotlines – and governance instruments like blocklisting (Spamhaus), de-peering (Level3), or governmental pressure that could directly hurt business continuity affect these client notification rates. In contrast, individual reporters of abuse are often ignored. Next, we observe a mismatch between the severity of certain abuse types and their corresponding anti-abuse governance instruments. We identify some previously found indicators of malicious networks at this company; yet we argue that labeling this company 'bulletproof' based solely on external measurements is a tough label to sell.

Ethics Considerations

We discuss the ethics involved in our work extensively in the following paragraphs. First, we will detail the ethical considerations and privacy-preserving steps we took in handling the seized data. Then, we will use the Menlo report [37] to outline how we addressed the sensitive nature of our data in our analyses.

Dataset

In line with applicable laws and regulations, Dutch authorities were able to seize company records, including the mailboxes and CRM database. While we use data from a legal seizure, one should not assume that users were engaged in illegal behavior or that this was a factor in deciding to use this data for our research. Note that providing any evidence of any kind for any law enforcement effort is not the purpose of this study. Before back-end data was made accessible to us for academic research purposes, public prosecutors weighed, among other things, the impact of the work on the rights and privacy of all parties. A Dutch law enforcement privacy officer vetted that our data subset was limited, contained only data vital to our research, and contained no personally identifiable information (PII). All of our analyses were conducted on-site at Dutch law enforcement agencies, where the data was stored and protected under their safety and security guidelines. We conferred with our IRB beforehand, and they viewed this work as outside of their jurisdiction, yet were satisfied with the assessments and applied procedures outlined above stemming from the public prosecutors and the law enforcement privacy officer.

Analyses

We discuss further ethics considerations using the principles identified in the Menlo Report [37].

Respect for persons: In order to protect the privacy of the company's clients, we took great care not to analyze PII – i.e., the data was stripped of all PII. This process was outlined by the involved privacy officer, following strict regulations that exceed the requirements of GDPR or IRB institutional frameworks, and then implemented by Dutch law enforcement. As a result, we only had access to data essential to our analyses, which had been stripped of any PII by law enforcement before we were granted access. Moreover, in this paper, we only report on aggregated values and use (translated) excerpts of anonymized conversations in abuse tickets. Extracting aggregate data points for our tables and figures was conducted under strict supervision through one specific, monitored channel. To respect the privacy of all individuals involved, we do not refer to any user – neither clients nor employees – in particular. With this approach, the data was cleared by Dutch authorities for this research, in accordance with Dutch privacy law.

Beneficence: We believe that our analysis does not create further harm as we did not partake in or stimulate any criminal business model – by buying criminal services or in any other way contributing to its ecosystem. The authors and involved law enforcement professionals believe that the benefits of a comprehensive understanding of the workings of the anti-abuse ecosystem outweigh the potential costs associated with making our work public.

Justice: The benefits of our work are distributed to the wider public, in terms of helping to reduce abuse through identifying improvements for the anti-abuse ecosystem. It especially helps to protect vulnerable individuals from victimization – e.g., harmful content, like CSAM. We see no direct impact on individuals, as we do not report on any specific individuals and have stipulated that one should not assume any criminal wrongdoing on their part.

Respect for law and public interest: This study has been conducted with the approval of, and in collaboration with, involved law enforcement professionals and public prosecutors. It is essential to note that while seized information may suggest certain illegal conduct, this paper does not provide nor seek to provide any legal evidence of criminal conduct.

References

[1] M. H. Jhaveri, O. Cetin, C. Gañán, T. Moore, and M. V. Eeten, "Abuse Reporting and the Fight Against Cybercrime," *ACM Computing Surveys*, vol. 49, no. 4, pp. 1–27, Dec. 2017. [Online]. Available: https://dl.acm.org/doi/10.1145/3003147

[2] M3AAWG, "Anti-Abuse Best Common Practices for Hosting and Cloud Service Providers," Mar. 2015. [Online]. Available: https://www.m3aawg.org/sites/default/files/document/M3AAWG_Hosting_Abuse_BCPs-2015-03.pdf

[3] European Parlement & the European Council, "Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act)," Oct. 2022. [Online]. Available: http://data.europa.eu/eli/reg/2022/2065/oj

[4] IETF, "Mailbox Names for Common Services, Roles and Functions," May 1997. [Online]. Available: https://www.ietf.org/rfc/rfc2142.txt

[5] S. Tajalizadehkhoob, R. Böhme, C. Gañán, M. Korczyński, and M. V. Eeten, "Rotten Apples or Bad Harvest? What We Are Measuring When We Are Measuring Abuse," *ACM Transactions on Internet Technology*, vol. 18, no. 4, pp. 49:1–49:25, Jul. 2018.

[6] C. A. Shue, A. J. Kalafut, and M. Gupta, "Abnormally Malicious Autonomous Systems and Their Internet Connectivity," *IEEE/ACM Transactions on Networking*, vol. 20, no. 1, pp. 220–230, 2012.

[7] B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda, "FIRE: FInding Rogue nEtworks," in *2009 Annual Computer Security Applications Conference*. Honolulu, Hawaii, USA: IEEE, Dec. 2009, pp. 231–240. [Online]. Available: http://ieeexplore.ieee.org/document/5380682/

[8] M. Konte, R. Perdisci, and N. Feamster, "ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*. London United Kingdom: ACM, Aug. 2015, pp. 625–638. [Online]. Available: https://dl.acm.org/doi/10.1145/2785956.2787494

[9] S. Alrwais, X. Liao, X. Mi, P. Wang, X. Wang, F. Qian, R. Beyah, and D. McCoy, "Under the Shadow of Sunshine: Understanding and Detecting Bulletproof Hosting on Legitimate Service Provider Networks," in *2017 IEEE Symposium on Security and Privacy (SP)*. San Jose, CA, USA: IEEE, May 2017, pp. 805–823. [Online]. Available: http://ieeexplore.ieee.org/document/7958611/

[10] A. Noroozian, J. Koenders, E. van Veldhuizen, C. H. Ganan, S. Alrwais, D. McCoy, and M. van Eeten, "Platforms in Everything: Analyzing Ground-Truth Data on the Anatomy and Economics of Bullet-Proof Hosting," in *28th USENIX Security Symposium*. Santa Clara, CA, USA: USENIX Association, 2019, pp. 1341–1356.

[11] M. Goncharov, "Criminal Hideouts for Lease: Bulletproof Hosting Services," TrendMicro, Tech. Rep. 28, 2015. [Online]. Available: https://www.trendmicro.com/vinfo/pl/security/news/cybercrime-and-digital-threats/bulletproof-hosting-services-cybercriminal-hideouts-for-lease

[12] D. Mahjoub, "Behaviors and Patterns of Bulletproof and Anonymous Hosting Providers," 2017. [Online]. Available: https://www.usenix.org/conference/enigma2017/conference-program/presentation/mahjoub

[13] B. Collier, R. Clayton, A. Hutchings, and D. R. Thomas, "Cybercrime is (often) boring: Maintaining the infrastructure of cybercrime economies," in *Workshop on the Economics of Information Security (WEIS 2020)*, Brussels, Belgium, Dec. 2020.

[14] J. Armin and C. Everett, "Top 50 Bad Hosts & Networks 2010 Q1," Hostexploit, Tech. Rep., Mar. 2010. [Online]. Available: http://hostexploit.com/downloads/top_50_bad_hosts_201003.pdf

[15] Openbaar Ministerie, "FIOD doet onderzoek naar bedrijf dat 'bulletproof' internetdiensten aanbiedt," Sep. 2020. [Online]. Available: https://www.om.nl/actueel/nieuws/2020/09/25/fiod-doet-onderzoek-naar-bedrijf-dat-%E2%80%98bulletproof%E2%80%99-internetdiensten-aanbiedt

[16] B. Krebs, "Naming and Shaming 'Bad' ISPs," Mar. 2010. [Online]. Available: https://krebsonsecurity.com/2010/03/naming-and-shaming-bad-isps/

[17] Spamhaus Project, "Don't Route Or Peer Lists (DROP)," Nov. 2025. [Online]. Available: https://www.spamhaus.org/blocklists/do-not-route-or-peer/

[18] Federal Government of the United States, "Computer Fraud and Abuse Act (CFAA)," 1986.

[19] T. Curtiss, "Computer Fraud and Abuse Act Enforcement: Cruel, Unusual, and Due for Reform," *Washington Law Review*, vol. 91, no. 4, pp. 1813–1850, 2016.

[20] R. J. Durbin, "STOP CSAM Act," Apr. 2023. [Online]. Available: https://www.congress.gov/bill/118th-congress/senate-bill/1199/text

[21] Google, "About the YouTube Priority Flagger program," 2024. [Online]. Available: https://support.google.com/youtube/answer/7554338?hl=en

[22] Meta, “Bringing local context to our global standards,” 2023. [Online]. Available: https://transparency.meta.com/policies/improving/bringing-local-context/

[23] L. Sivetc and M. Wijermars, “The Vulnerabilities of Trusted Notifier-Models in Russia: The Case of Netoscope,” *Media and Communication*, vol. 9, no. 4, pp. 27–38, 2021.

[24] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Felegyhazi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, He Liu, D. McCoy, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage, “Click Trajectories: End-to-End Analysis of the Spam Value Chain,” in *2011 IEEE Symposium on Security and Privacy*. Berkeley, CA: IEEE, May 2011, pp. 431–446. [Online]. Available: http://ieeexplore.ieee.org/document/5958044/

[25] J. Zhang, Z. Durumeric, M. Bailey, M. Liu, and M. Karir, “On the Mismanagement and Maliciousness of Networks,” in *Proceedings 2014 Network and Distributed System Security Symposium*, vol. 14. San Diego, CA: Internet Society, 2014, pp. 23–26. [Online]. Available: https://www.ndss-symposium.org/ndss2014/programme/mismanagement-and-maliciousness-networks/

[26] A. Noroozian, M. Ciere, M. Korczynski, and M. van Eeten, “Inferring the Security Performance of Providers from Noisy and Heterogenous Abuse Datasets,” in *16th Workshop on the Economics of Information Security (WEIS)*, San Diego, CA, USA, 2017.

[27] D. Kopp, E. Strehle, and O. Hohlfeld, “CyberBunker 2.0 - A Domain and Traffic Perspective on a Bulletproof Hoster,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. Virtual Event Republic of Korea: ACM, Nov. 2021, pp. 2432–2434. [Online]. Available: https://dl.acm.org/doi/10.1145/3460120.3485352

[28] M. Kührer, C. Rossow, and T. Holz, “Paint It Black: Evaluating the Effectiveness of Malware Blacklists,” in *Research in Attacks, Intrusions and Defenses*. Cham: Springer International Publishing, 2014, vol. 8688, pp. 1–21. [Online]. Available: http://link.springer.com/10.1007/978-3-319-11379-1_1

[29] B. Stone-Gross, T. Holz, G. Stringhini, and G. Vigna, “The Underground Economy of Spam: A Botmaster’s Perspective of Coordinating Large-Scale Spam Campaigns,” in *Proceedings of the 4th USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET ’11)*, San Francisco, CA, USA, 2011. [Online]. Available: https://www.usenix.org/legacy/events/leet11/tech/full_papers/Stone-Gross.pdf

[30] T. Moore and R. Clayton, “The Impact of Incentives on Notice and Take-down,” in *Proceedings of the Seventh Workshop on the Economics of Information Security (WEIS 2008)*, M. E. Johnson, Ed., Hanover, NH, Germany, 2008. [Online]. Available: https://weis2008.econinfosec.org/papers/Moore.pdf

[31] O. Çetin, M. Hanif Jhaveri, C. Gañán, M. Van Eeten, and T. Moore, “Understanding the role of sender reputation in abuse reporting and cleanup,” *Journal of Cybersecurity*, vol. 2, no. 1, pp. 83–98, Dec. 2016. [Online]. Available: https://academic.oup.com/cybersecurity/article-lookup/doi/10.1093/cybsec/tyw005

[32] F. Li, G. Ho, E. Kuan, Y. Niu, L. Ballard, K. Thomas, E. Bursztein, and V. Paxson, “Remedying Web Hijacking: Notification Effectiveness and Webmaster Comprehension,” in *Proceedings of the 25th International Conference on World Wide Web*. Montréal Québec Canada: International World Wide Web Conferences Steering Committee, Apr. 2016, pp. 1009–1019. [Online]. Available: https://dl.acm.org/doi/10.1145/2872427.2883039

[33] Spamhaus Project, “Spamhaus - Strengthening trust and safety across the internet,” 2024. [Online]. Available: https://www.spamhaus.org/

[34] Rijksoverheid, “Grapperhaus spreekt ICT-bedrijven aan op kinderporno op hun servers,” Apr. 2020. [Online]. Available: https://www.rijksoverheid.nl/actueel/nieuws/2020/04/30/grapperhaus-spreekt-ict-bedrijven-aan-op-kinderporno-op-hun-servers

[35] RIPE NCC, “RIPEstat Announced Prefixes,” Aug. 2024. [Online]. Available: https://stat.ripe.net/docs/02.data-api/announced-prefixes.html

[36] Jan Vermeulen, “The big South African IP address heist – How millions are made on the “grey” market,” Sep. 2019. [Online]. Available: https://mybroadband.co.za/news/internet/318205-the-big-south-african-ip-address-heist-how-millions-are-made-on-the-grey-market.html

[37] E. Kenneally and D. Dittrich, “The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research,” *SSRN Electronic Journal*, 2012. [Online]. Available: http://www.ssrn.com/abstract=2445102

Appendix

TABLE VI: List of words used to categorize abuse reports.

Abuse category	**Words**
CSAM / Harmful content	child, csam, cybertip, klpd, (CP), underage, iwf, kinderporno, inhope
Botnet C&C / DDoS Attacks	command-and-control, botnet server, c&c, ddos, dos, udp flood, portflood, dns-attack, dos attack, botnet controller, spamhaus botnet controller list
Malware / Phishing / Brute-force attacks	phish, fraudulent, malware, trojan, bruteforce, brute force, intrusion, hack, breakin attempt, break-in attempt, network attack, unauthorized login attempts, possible malicious activity from this host, unauthorized access, security incident, unallowed network access, service: ssh, auth_login authenticator failed for, sshd:auth, was found attacking, sasl login authentication failed, clean-mx-phishing, clean-mx-viruses, spyeye, malware distribution
Spam(vertising)	spam, spam email, unsolicited, clean-mx-spam, spam emitter, spam sites, clean-mx-portal, spamcop, spamming
Port scanning / Comment spamming	scanning, scanner, portscan, port scan, netscan, objectionable traffic, badbot regbot, clean-mx-trackback, port_scanning
Copyright / Trademark issues	piracy, dmca, copyright, piracy, infringement, urgent live stream escalation, notice of unauthorized use of, unauthorized sound recordings, notice of infringing activity, notice of claimed infringement, infringing material, fapl, counterfeit, balenciaga, yves saint laurent, tag heuer, dior, gucci, moncler, chanel, longchamp, trademark infringement, fake shop, torrent, warez, brand protection