Exposing the Roots of DNS Abuse: A Data-Driven Analysis of Key Factors Behind Phishing Domain Registrations

Yevheniya Nosyk
yevheniya.nosyk@korlabs.io
KOR Labs
Grenoble, France

Maciej Korczyński
KOR Labs / Univ. Grenoble Alpes,
CNRS, Grenoble INP, LIG
Grenoble, France

Carlos Gañán
ICANN
Los Angeles, United States

Sourena Maroofi
KOR Labs
Grenoble, France

Jan Bayer
KOR Labs
Grenoble, France

Zul Odgerel
KOR Labs / Univ. Grenoble Alpes,
CNRS, Grenoble INP, LIG
Grenoble, France

Samaneh Tajalizadehkhoob
ICANN
Los Angeles, United States

Andrzej Duda
KOR Labs / Univ. Grenoble Alpes,
CNRS, Grenoble INP, LIG
Grenoble, France

ABSTRACT

Cybercriminals have long depended on domain names for phishing, spam, malware distribution, and botnet operation. To facilitate the malicious activities, they continually register new domain names for exploitation. Previous work revealed an abnormally high concentration of malicious registrations in a handful of registrars and TLDs. However, no existing study systematically analyzed the factors driving abuse, leaving a critical gap in understanding how different variables influence malicious registrations. In this paper, we carefully distill the inclinations and aversions of malicious actors during the registration of new phishing domain names. Having compiled a list of 14.5 k malicious and 15.4 k benign domains, we collect a comprehensive set of 73 features for all the domains encompassing three main latent factors: registration attributes, proactive verification, and reactive security practices. With a GLM regression analysis, we found that each dollar reduction in registration fees corresponds to a 49% increase in malicious domain registrations. The availability of free bundled services, such as web hosting, drives an 88% surge in phishing activities. Conversely, stringent registration restrictions cut down abuse by 63%, while registrars providing API access for domain registration or account creation experience a staggering 401% rise in malicious domains. The results enable intermediaries involved in domain registration to develop tailored anti-abuse practices, yet aligning them with their economic interests.

CCS CONCEPTS

• Security and privacy → Usability in security and privacy; **• Networks → Naming and addressing**;

KEYWORDS

DNS abuse, phishing, domain registration

ACM Reference Format:
Yevheniya Nosyk, Maciej Korczyński, Carlos Gañán, Sourena Maroofi, Jan Bayer, Zul Odgerel, Samaneh Tajalizadehkhoob, and Andrzej Duda. 2025. Exposing the Roots of DNS Abuse: A Data-Driven Analysis of Key Factors Behind Phishing Domain Registrations. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25), October 13–17, 2025, Taipei, Taiwan.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3719027.3744869

1 INTRODUCTION

Cybercriminals extensively exploit the Domain Name System (DNS) for a broad range of illegal or malicious activities, including phishing, spam, botnet C&C or malware dissemination. Domain names serve as pathways that direct victims to servers hosting harmful content. Cybercriminals may register domain names for malicious purposes or use legitimate domains registered by benign users that later fall prey to vulnerabilities in software [49, 53, 73, 77], resulting in their exploitation for hosting malware or phishing websites.

The 2024 Phishing Landscape study [26] revealed that over 1.1 million unique domains were involved in phishing and added to blocklists between May 2023 and April 2024. Despite the costs, most of those (878 k) were deliberately registered by cybercriminals. Once abusive domains appear on blocklists, they are mitigated either by blocking communication at the network level (by ISPs, mail service operators, or DNS resolver operators) or, more efficiently, at the DNS level by registrars or top-level domain (TLD) registries. Thus, attackers have to fulfill a constant need for numerous single-use domain names to maintain their criminal activities effectively [28].

Supply-side dynamics, market competition, and economic strategies all shape abuse patterns. The registration ecosystem is complex,

CCS '25, October 13–17, 2025, Taipei, Taiwan

ACM ISBN 979-8-4007-1525-9/2025/10
https://doi.org/10.1145/3719027.3744869

involving various domain providers: registrars, TLD registries, resellers, as well as regulators. Providers operate under profitability—domain registrations often generate revenue, creating an incentive to maximize volume. For example, price competition is common among new gTLDs, while country-code TLDs (often non-profits) prioritize reputation. Governance is split between ICANN for gTLDs and national bodies for ccTLDs. Registrars combat abuse to protect their reputation, avoid legal disputes (e.g., Namecheap[1]), prevent financial losses, and comply with evolving regulations [11, 37]. Industry initiatives like ICANN Domain Metrica [42] and Netbeacon MAP [57] further reduce information asymmetry by publishing reputation metrics [44] for registrars and TLDs.

These dynamics do not affect all players equally: malicious registrations are not uniformly distributed over different entities of the DNS ecosystem—they tend to be skewed toward certain registrars [12, 27, 28, 45, 51, 82] and TLDs [12, 13]. A notable example is Freenom, which previously managed five ccTLDs (`.tk`, `.ml`, `.ga`, `.cf`, and `.gq`) and offered free registrations. As of 2013, Freenom accounted for 28% of all malicious domain registrations [25].

Previous studies only speculated on factors that make certain registrars appealing to cybercriminals, noting patterns such as malicious domains registered in large campaigns or through registrars offering bulk registration [2, 8, 12, 20, 28, 44, 45, 50, 82]. However, these studies remain descriptive, failing to uncover the underlying mechanisms driving these behaviors. While some research [12, 44, 45, 50] highlighted the prevalence of phishing in low-cost or free TLDs, they fell short of providing a comprehensive statistical model that examines not only a wider range of factors that matter in reality but also their interactions with one another. Moreover, existing data-driven studies have not identified specific measures that effectively increase barriers to DNS abuse while remaining appealing to legitimate clients.

In this paper, we perform a thorough analysis of different factors that, we hypothesize, may influence the choice of certain registrars and TLDs by malicious actors when registering new phishing domains. Our dataset includes 14.5 k maliciously registered and 15.4 k benign domains for which we gather and analyze various TLD attributes and registrar practices. We define 73 features across three main groups of latent factors of DNS abuse: registration attributes (e.g., prices, payment methods, additional services), proactive verification (e.g., checking registrant data), and reactive security practices (e.g., uptimes, the measure of abuse persistence). These features are collected at the domain name level and at the time of registration, offering a novel perspective on both the immediate registration decisions of attackers and the factors shaping their behavior.

We develop two complementary statistical models: the first model aims to estimate the impact of the features on the number of maliciously registered domains while the second one indicates whether the registrar or TLD level features are favored by attackers alone or also by legitimate users. Overall, we make the following contributions:

- We build a Generalized Linear Model (GLM) regression model that empirically analyzes the relationship between various features and the concentration of phishing domain names using *registrar-TLD* pairs as the unit of analysis. This is motivated by the fact that certain variables are inherited from TLD registry practices, while others directly originate from registrars.
- We observe that domain abuse is closely linked to discounts, with each dollar off leading to a 49% increase in malicious domain registrations. Bundled, free services like web hosting result in an 88% rise in the number of phishing domains, while stringent restrictions reduce abuse by 63%. Registrars offering API access for domain registration or account creation see a 401% increase in malicious domains. Mitigation times have little impact, likely because even brief uptimes may provide phishers with valuable credentials and financial gain.
- We propose a second, complementary logistic regression model to analyze factors associated with malicious and benign domain registrations. Using fine-grained regression analysis at the *domain level*, we focus on uncovering and interpreting the relationships between key features and malicious activity, with an emphasis on explanatory insights rather than prediction.
- The analysis shows that discounts attract more malicious users than legitimate ones, while various restrictions reduce the likelihood of abuse by around 19%. This suggests that malicious actors are more sensitive to these factors when choosing a registrar or TLD. These insights can help registrars and registries design policies that deter abuse without discouraging legitimate use.

The remainder of this paper is organized as follows: §2 provides the background on domain registrations and anti-abuse measures. §3 reviews related work. §4 presents the datasets we use to collect the features discussed in §5. We analyze malicious and benign registrations in §6 and evaluate the driving factors of abuse in §7. §8 discusses key insights, §9 reflects on ethics, and §10 concludes the paper.

2 BACKGROUND

This section provides background on the Registrant - Registrar - Registry (RRR) model, DNS abuse, anti-abuse measures, and maliciously registered phishing domains.

2.1 DNS Ecosystem

The Domain Name System (DNS) is one of the fundamental components of the modern Internet providing the mapping between human-readable domain names and IP addresses. Each top-level domain, e.g., `.com` (legacy generic TLD), `.top` (new generic TLD), or `.de` (country-code TLD), is managed by a *registry*—an organization that sets the registration terms and prices, maintains the DNS zone file, and configures DNSSEC. As of January 2025, the DNS Root Zone Database contains 1,591 top-level domains [31].

Registries typically delegate the responsibility of selling domain names to *registrars* that set up contractual agreements with registries and sell domain names under the relevant TLDs. Obtaining ICANN accreditation is essential for selling gTLD domain names, whereas for ccTLDs, accreditation from local registry operators may suffice (e.g., SIDN for `.nl` domains [70]). Finally, a *registrant* is any entity (benign or malicious) that registers a domain name and agrees to the registrar terms providing accurate personal information as required by the registry, ICANN, or both.

[1] https://about.fb.com/news/2020/03/domain-name-lawsuit/

2.2 DNS Abuse

Cybercriminals extensively leverage DNS for a wide panoply of illegal and malicious activities. In 2019, a group of domain registries and registrars, including GoDaddy [22], Tucows [80], Namecheap [55], Public Interest Registry (`.org`) [66], Neustar (`.biz`, `.us`) [58] and Afilias (now part of Identity Digital) [3] voluntarily created the DNS Abuse framework [14]. It aimed to provide a clear definition of DNS abuse and establish guidelines for registries and registrars to combat DNS abuse more effectively and consistently across the industry. They categorized DNS abuse into five different types: malware, botnets, phishing, pharming, and spam (when used to distribute the other threats). These activities exploit the DNS infrastructure as a delivery mechanism for their illicit operations [1, 35].

Abuse handling policies and procedures vary among registrars and the operators of generic TLDs and ccTLDs. ICANN-accredited registrars must adhere to specific abuse handling guidelines. Previously, they were expected to address abuse complaints but specific requirements and timelines for response were not clearly defined. The focus was primarily on reactive abuse measures. As of April 2024, gTLD registries are also required to address abuse proactively due to new contractual amendments [37]. In contrast, ccTLDs are considered national resources with unique characteristics and do not have contractual agreements with ICANN for abuse handling policies. Thus, their procedures depend on local regulations and the voluntary practices of individual registries and registrars.

2.3 Anti-Abuse Measures

TLD registries and registrars undertake various measures to prevent and mitigate abusive domain names, categorized into proactive verification, reactive security practices, and registration attributes. While not strictly preventive, registration attributes may indirectly deter abuse.

The primary objective of proactive measures is to prevent malicious registrations from occurring. Certain TLDs are restricted to specific regions (e.g., `.eu`) or professions (e.g., `.abogado`). Some TLD registries and registrars implement identity verification processes known as Know Your Business Customer (KYBC) [19] (e.g., `.dk` [68], `.cn` [9]). Additionally, other registries use machine learning techniques to identify suspicious domain names during registration (e.g., `.eu` [75], `.nl` [72], `.be` [15]), ensuring they are flagged before being added to the namespace. Another proactive strategy is to block the registration of domains containing keywords associated with well-known brands (e.g., Porkbun and Namecheap prevent the registration of such domains as discussed in Section 6.5).

When a domain name is involved in phishing or malware distribution, the TLD registry or registrar can take action to remove it at the DNS level if there is sufficient evidence of abuse and no legitimate content is being served. However, if the domain itself is legitimate but vulnerable software on the site has been exploited [53], the issue cannot be resolved at the DNS level. Instead, abuse must be handled by the hosting provider or webmaster [77]. Taking action at the DNS level in such cases could cause collateral damage to website visitors and owners of benign domain names.

At the DNS level, several actions can be taken against abusive domains: i) the domain can be deleted from both the DNS zone and the namespace, ii) the domain can be delisted (suspended) from the DNS zone but remain in the namespace, iii) the domain can stay in both the zone and namespace, but its authoritative nameserver can be changed to a dedicated one, managed by the TLD registry specifically for this purpose.

Note that the diverse domain registration attributes proposed by registrars such as varying prices, bundled services (e.g., web hosting), and multiple payment methods (credit cards, PayPal, cryptocurrencies) could act as preventive measures. For example, GoDaddy recently updated its Terms of Use, requiring customers to have 50 or more domains in their accounts to use the Availability API [23, 24]. This change could potentially impact the use of the GoDaddy API for malicious purposes, although its effect on abuse rates is yet to be determined.

2.4 Maliciously Registered Phishing Domains

This study examines phishing abuse and domains registered with malicious intent. Phishing is widely recognized as a significant cyber threat and a prevalent form of DNS abuse. According to the 2023 annual report from the FBI, phishing is the leading type of digital crime, with over 300,000 complaints and losses exceeding $160 million [81]. We specifically focus on phishing because malware delivery URLs are less common and readily detected, spam domains do not always qualify as DNS abuse, and phishing typically involves clear evidence, such as screenshots of fraudulent sites.

While some phishing domains are registered with purely malicious intent (or "attacker-owned" [73]), others are benign but may become compromised through, e.g., vulnerabilities in their content management systems (CMS) [53], etc. Attackers may also exploit free services such as subdomain providers to disseminate malicious content. Current phishing detection methods identify the indicators of ongoing attacks, often conflating maliciously registered and compromised domains into common URL blocklists. Therefore, previous research has proposed methods to distinguish between these two groups [49, 53, 73]. Given the significance and economic impact of phishing attacks, we focus on domains maliciously registered for phishing purposes, rather than benign ones that are later exploited.

3 RELATED WORK

Previous studies speculated on isolated factors that make certain registrars appealing to cybercriminals, noting patterns such as malicious domains registered in large campaigns or through registrars offering bulk registration [2, 8, 12, 20, 28, 44, 45, 50, 82]. However, these studies focused on anecdotal observations rather than systematically uncovering the underlying factors that drive abuse.

Earlier research collected evidence that malicious actors register domains in bulk. Felegyhazi et al. [20] discovered registration clusters from a small seed of known malicious domains. Hao et al. [28] showed that 80% of spam domains were registered in groups, 10% belonging to batches of more than 200 registrations. Similar findings were observed at the `.eu` ccTLD for which 80% of malicious registrations were associated with 20 campaigns [82]. Moreover, Affinito et al. [2] examined two malicious domain registration spikes and noted that the two registrars behind offered bulk registration to their customers. However, these studies tend to stop short by only observing the phenomenon, without uncovering the underlying

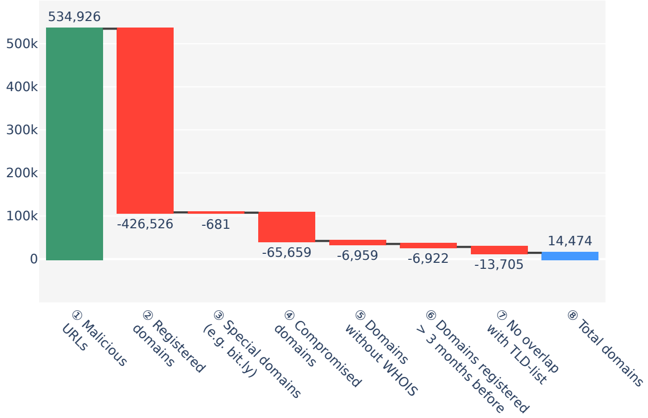

Figure 1: Steps to compile 14.5 k maliciously registered domains from blocklisted URLs using conservative filtering approach which minimizes the inclusion of benign domains.

causes driving such a behavior. In contrast, our research takes a significant step forward by analyzing the factors that foster malicious registrations in the first place. We identify crucial drivers, such as unrestricted API access to domain registration and management services, which makes large-scale malicious operations possible.

Studies like those by Korczyński et al. [44] and Bayer et al. [12] highlighted the prevalence of phishing in low-cost or free TLDs, with Liu et al. [50] additionally exploring the impact of registration fee increases. However, these prior studies continue to face challenges in conclusively demonstrating the impact of pricing on abuse counts, relying primarily on anecdotal evidence while acknowledging several limitations in their analyses. Rather than simply distinguishing free domains (historically served by Freenom) from paid ones [44], we provide a comprehensive analysis of the cost structure at the exact time of domain registration, including retail fees, discounts, renewal, and transfer prices. This analysis further incorporates a rich dataset of additional features, revealing relationships between registration costs and factors like payment methods, bundled services, and registration restrictions. Moreover, while these previous works conducted analyses solely at the TLD level, we adopt the TLD-registrar pair as our unit of analysis and explicitly focus on maliciously registered domains.

Existing research has also focused on the responses of registrars and registries to abuse complaints. For example, Liu et al. [50] documented a successful collaboration between eNom and LegitScript to take down rogue pharmacy domains, while Cheng et al. [8] observed swift actions by Chinese registrars to suspend domains hosting gambling and adult content. However, these studies mainly assess the effectiveness of registrar responses after abuse occurs. Our work, in contrast, explores how malicious actors select registrars, focusing on whether slower mitigation times (i.e., uptimes) influence their choices. Furthermore, while [44] examines the correlation between abuse concentrations and uptimes at the TLD level, they do not incorporate uptimes into a comprehensive analysis of multiple factors. In contrast, our analysis is performed at the registrar-TLD level, recognizing that mitigation efforts may be carried out by either the registry or registrar. We also measure and analyze uptimes with and without notifications to registrars.

Our work advances previous research by specifically focusing on *maliciously registered* domains, distinguishing them from compromised websites. Unlike existing studies that often focus on isolated factors and make anecdotal observations, our approach integrates 73 detailed *domain-level* features collected at the registration time to propose a more comprehensive understanding of the drivers behind malicious registrations. None of the features from existing work directly overlap with the variables collected and analyzed in the paper. Our work advances the field by developing a comprehensive statistical model that systematically examines a wide range of factors influencing registrar-TLD abuse, moving beyond the isolated and anecdotal speculations of previous research [2, 8, 12, 20, 28, 44, 45, 50, 82]. By developing two statistical models, we offer actionable insights into the key factors influencing attackers' choices, helping registrars and registries better address abuse.

4 DATASETS

This section outlines the key datasets: the TLD-List of registrar and TLD features, malicious domains from blocklists, and a sampled set of benign domains for comparisons.

4.1 TLD-List

The TLD-List service [79] has been collecting data on top-level domains and registrars since 2015. With a unique focus on pricing, it includes domain registration costs, discounts, and free features. We subscribed to this service, collecting daily snapshots with data on the registrar payment methods, free features (e.g., SSL/TLS certificates), and prices. Overall, the assembled datasets span 75 domain name registrars and more than 1,500 top-level domains. They enable us to verify the prices and services offered by registrars/registries on a given day.

We validated the data by randomly sampling 20 registrar-TLD pairs across different dates during the analysis period and manually verifying the dataset. While all data was generally accurate, we found seven discrepancies related to payment methods. However, since we aggregate payment methods in our model (see Section 7.1), these minor inconsistencies do not affect our findings.

4.2 Maliciously Registered Phishing Domains

To understand the registration preferences of phishers, we analyze the domain names that satisfy two conditions: i) they were involved in phishing activities and ii) they were deliberately registered by cybercriminals.

Figure 1 details the process of curating the dataset. We first collect 534 k blocklisted URLs (Step ①) from three phishing feeds maintained by the Anti-Phishing Working Group (APWG) [5], PhishTank [61], and OpenPhish [60], spanning the period between August 2023 and January 2024. We selected these feeds because they have been commonly used in prior research [36, 44, 45, 53] and are maintained by reputed organizations. We process all the URLs and extract 108 k registered domains (Step ②), noting that some are benign but have been abused by malicious actors. We

Table 1: Top 20 domains sorted by malicious URLs in our 534 k URL dataset—all are subdomain providers or URL shorteners.

Rank	Domain	Count	Rank	Domain	Count
1.	workers.dev	28,630	11.	duckdns.org	7,994
2.	weeblysite.com	20,540	12.	repl.co	7,771
3.	cf-ipfs.com	13,888	13.	square.site	5,436
4.	web.app	13,491	14.	vercel.app	4,598
5.	firebaseapp.com	12,054	15.	blogspot.com	4,569
6.	dweb.link	11,973	16.	cprapid.com	4,126
7.	r2.dev	11,784	17.	wixsite.com	3,923
8.	github.io	11,145	18.	glitch.me	3,654
9.	pages.dev	9,314	19.	000webhostapp.com	3,495
10.	weebly.com	8,739	20.	ipfs.io	3,484

Table 2: 20 most frequently observed registrar-TLD pairs in our dataset of maliciously registered domain names.

Rank	Registrar/TLD	Count	Rank	Registrar/TLD	Count
1.	NameSilo/`top`	1,807	11.	NameSilo/`buzz`	222
2.	NameSilo/`com`	852	12.	Sav/`com`	211
3.	GoDaddy/`com`	832	13.	Alibaba Cloud/`shop`	197
4.	Hostinger/`online`	764	14.	NameSilo/`us`	191
5.	NameSilo/`info`	513	15.	Hostinger/`site`	179
6.	Hostinger/`com`	479	16.	NameSilo/`life`	178
7.	Namecheap/`com`	479	17.	NameSilo/`sbs`	171
8.	Alibaba Cloud/`com`	327	18.	Hostinger/`shop`	156
9.	NameSilo/`xyz`	233	19.	NameSilo/`cc`	149
10.	Hostinger/`cloud`	225	20.	Alibaba Cloud/`top`	148

begin by excluding the domains associated with 681 URL shorteners[2] (e.g., `bit.ly`), subdomain providers, or file-sharing sites[3] (e.g., `000webhostapp.com`, `ipfs.io`), known to be used for delivering malicious content [47, 48, 59] (Step ③). Overall, these domains accounted for more than 258 k URLs in our dataset, showing how cybercriminals increasingly abuse existing services to disseminate phishing. Table 1 lists the top 20 registered domain names in our dataset of 534 k blocklisted URLs, all of which happen to be subdomain providers, URL shorteners, and file sharing services. These 20 domains alone represent 36% of malicious links analyzed.

Next, we perform a set of measurements for all the blocklisted domains during one month after being reported. Specifically, we retrieve registration data (using WHOIS or RDAP protocols) and DNS `A` records. While compromised domains should only have the malicious content removed, maliciously registered ones should result in a takedown action at the DNS level as evidenced by the `NXDOMAIN` DNS response code and the Extensible Provisioning Protocol (EPP) status code set to `clientHold` or `serverHold` [29]. Our analysis includes only domains that were mitigated at the DNS level within the one-month monitoring period, which results in 66 k domains being categorized as compromised (Step ④) and 42 k as maliciously registered, the latter being kept for further analysis.

Furthermore, given that maliciously registered domain names are often used for malicious activities shortly after registration [28], we only include the domain names registered within 90 days prior to being blocklisted [45, 54]. Therefore, we remove 7 k domain names without WHOIS data (Step ⑤) and another 7 k domains registered more than three months before being blacklisted (Step ⑥), resulting in a set of 28.2 k domains. While this rigorous approach—measuring DNS-level mitigations and verifying registrations within 90 days before blocklisting—may miss some malicious registrations, it helps ensure compromised domains are excluded.

Finally, to associate maliciously registered domain names with the daily-collected registration features, we extract the registrar IANA ID and the registration date from the above-mentioned results of the WHOIS scan. We exclusively consider the domain names for which we have the registrar features, referencing the list of registrars supported by the TLD-List dataset (Step ⑦). We excluded 13.7 k domains without such an overlap, including 1,958 under Gname.com Pte. Ltd., 1,932 under NICENIC INTERNATIONAL GROUP CO., LIMITED, and 1,061 registered with PDR Ltd. d/b/a PublicDomainRegistry.com. Overall, we obtained the list of 14,474 maliciously registered domains (Step ⑧) spread across 165 TLDs and 31 registrars. Table 2 shows the 20 most frequently observed registrar-TLD pairs.

4.3 Benign Domains

Some factors that attract attackers, like competitive pricing or free features, may also appeal to legitimate users. To understand the differences, we curated a list of benign domain names as a baseline.

We first gather all registered domains that appear in the Centralized Zone Data Service [34], and Google Certificate Transparency (CT) logs [7]. We then perform a WHOIS scan of them to get registration dates and IANA IDs, keeping only the domains created during the same time window (August 2023 - January 2024) as the maliciously registered ones to ensure temporal comparability. We remove 1 m domains appearing in Spamhaus [78] and SURBL blocklists [76] and keep the list of 19 m domains created at registrars supported by the TLD-List dataset.

Malicious domain names are concentrated, with some registrars exhibiting high levels of abuse despite low overall market share. For instance, GoDaddy represents 55.34% of benign registrations but only 9.18% of malicious ones, whereas NameSilo accounts for 39.73% of malicious registrations despite a smaller market share (3.54%). Thus, using the previously obtained 14.5 k malicious domains as the basis for sampling would inevitably skew the benign collection.

Instead, we need a representative sample of benign domains that takes into account the registrar market share. Therefore, we refer to the ICANN Monthly Registry Reports [39] in which each gTLD registry gives the number of domains managed by each registrar under a particular gTLD. Although these numbers exclude ccTLD domains, they can still serve as an estimate of the registrar market share. Having obtained the market share ratios, we perform the stratified sampling of 19 m benign domains, ensuring that randomness within each stratum minimizes the selection bias and provides a more representative sample of the entire population. We finally collect all the registration and proactive features, which result in a dataset of 15.4 k domains under 259 TLDs originating from 38 registrars.

5 FEATURES

This section provides an overview and rationale for pre-selecting registration attributes and anti-abuse practices for further analysis.

[2] https://github.com/korlabsio/urlshortener
[3] https://github.com/korlabsio/subdomain_providers

Table 3: All the registration features of type Boolean (B) or Numerical (N), most of them retrieved from the TLD-List dataset (T) and some collected manually (M).

#	Feature Set Name	Type	Source
1.	*free_[api, dnssec, web_hosting]*	B	M
2.	*free_[dns, email_account, email_forward, ssl_cert]*	B	T
3.	*api_[create_account, register_domain]*	B	M
4.	*free_bulk_search_number*	N	M
5.	*bulk_discount*	B	M
6.	*payment_[alipay, applepay, banktransfer, bitcoin, cashinperson, cc, check, dinersclub, dwolla, giropay, googlewallet, moneyorder neteller, payeer, paypal, payza, qiwi, skril, topcoin, webmoney westernunion, worldpay, yandexmoney, yoomoney*	B	T
7.	*price_[register, renewal, transfer, whois_privacy]*	N	T
8.	*discount_[register, renewal, transfer]*	N	T
9.	*term_new_customer_only_[register, transfer]*	B	T
10.	*term_limit_per_customer_[register, transfer]*	N	T

Table 4: All the proactive/reactive features of type Boolean (B) or Numerical (N). Apart from the uptime measurements running automatically (A), the remaining features were collected manually (M).

#	Feature Set Name	Type	Source
1.	*[email, phone, address]_syntactically_validated*	B	M
2.	*[email, phone]_operational_validated*	B	M
3.	*random_[warning, prevention]*	B	M
4.	*office365_[warning, prevention]*	B	M
5.	*facebook_[warning, prevention]*	B	M
6.	*restriction_[not_available, local_presence, community_ties age_restriction, infrastructure, group_ties, id_required commitment_required, region_ties, certain_nationals_prohibited professionals_only, org_or_affiliates_only exclusive_registrar, content_restrictions]*	B	M
7.	*uptime_[notified, not_notified]*	N	A

5.1 Registration Attributes

We first describe the pre-selected registration attributes, most of which are derived from the TLD-List dataset. For any missing information not available in these datasets, we manually collect the necessary data. Table 3 provides the summary:

Free API: the registrar APIs enable users to search, purchase, and manage domains, allowing cybercriminals to fully automate the setup of malicious infrastructures. The boolean *free_api* feature indicates whether registrants can access the API without any prerequisites such as a reseller account or a paid subscription. We also define the boolean features *api_create_account* for the account creation and *api_register_domain* for domain registration.

Free DNS service: registrars commonly offer customers a free DNS service, effectively eliminating the need to establish and maintain a custom authoritative nameserver infrastructure. Spam domain owners [28] benefit from such a service as it reduces the overhead required to set up operational domain names. We define the *free_dns* boolean feature.

Free DNSSEC signing: registrars offering DNS services may provide free cryptographic signing of domain names, which may boost the domain reputation, even if not directly relevant to phishing. The *free_dnssec* feature is set to True if the registrar signs the domain without requiring clients to upload custom `DS` records.

Free email: registrars may offer free email boxes and/or email forwarding to their registrants. This service may also be exploited by attackers to deliver malicious content to their victims such as phishing links. Thus, we introduce two boolean features: *free_email_account* and *free_email_forward*.

Free web hosting: previous research indicated that attackers typically do not invest significant effort in creating fully functional websites [53]. They may leverage free hosting plans to host basic content on newly registered (malicious) domain names. If such a service is included for free in each domain registration, we set *free_web_hosting* to True.

Free SSL/TLS certificates: as of August 2020, 77.6% of phishing websites used SSL/TLS certificates [6]. Attackers may value free certificates for malicious domains despite the risk that they appear in CT logs, thereby increasing the chances of phishing detection [69]. We examine the impact of free certificates on the registrar selection using the boolean *free_ssl_cert* feature.

Free bulk search: registering multiple domains names at once may help attackers maintain resilience against quick blocklisting and enable running concurrent campaigns. Research indicates that malicious domains are often registered in batches [2, 43, 62]. Therefore, we examine the capability to search domains in bulk (numerical *free_bulk_search_number*) and any associated discounts (boolean *bulk_discount*).

Available payment methods: malicious actors tend to prioritize anonymity often opting for payment methods harder to trace such as cryptocurrency. For instance, ransomware operators predominantly use Bitcoin to receive payments from victims [30]. We define 24 boolean features for each payment method in the TLD-List dataset including PayPal, Bitcoin, and others (Feature set #6).

Retail pricing: existing research suggests that pricing significantly influences the registration preferences of attackers. They tend to favor domain name registrars and TLDs that offer the most competitive rates. We establish three numerical features: *price_register*, *price_renewal*, and *price_transfer*, given in $.

Discounts: discounts on domain registrations may attract attackers. For instance, Namecheap offers lower prices for bulk registrations of 50 or more domains. The discounts, which vary by TLD, can reduce the cost of building malicious infrastructure. We define three numerical features in $ to capture the discounts: *discount_register*, *discount_renewal*, and *discount_transfer*.

Pricing terms: certain registrars impose specific conditions on domain purchase. For example, discounted pricing might apply only to a limited number of domains or require purchase through an affiliate link. We define these conditions and purchase types using boolean features (*term_new_customer_only_register* and *term_new_customer_only_transfer*) and numerical features (*term_limit_per_customer_register* and *term_limit_per_customer_transfer*).

WHOIS privacy price: WHOIS/RDAP services reveal domain registration data, which may expose malicious actors [82]. The General Data Protection Regulation (GDPR)[10] mandates masking personal details of European Economic Area (EEA) registrants, and some registrars apply this to non-EEA registrations, sometimes for free [52]. Despite GDPR, we include the *price_whois_privacy* feature in our analysis.

5.2 Proactive Verification

To assess proactive measures, we create registrant accounts and add various borderline domain names to a cart, empirically testing the presence of proactive security practices prior to the domain purchase. We review below the examined features (see Table 4):

Syntactic validation of the registrant personal information: the ICANN SSAC Report on Domain Name Registration Data Validation [33] outlines three validation types: syntactic, operational, and identity. We test 38 registrars by attempting to create accounts with syntactically incorrect data, such as missing email symbols, overly long phone numbers, and invalid postal codes. We define three boolean features to indicate whether registrars accept incorrect data without warnings: *email_syntactically_validated*, *phone_syntactically_validated*, and *address_syntactically_validated*.

Operational validation of registrant information: ICANN mandates accredited registrars to collect accurate registrant contact information [32], while a recent EU directive requires verification of at least one contact method [11]. Although ICANN allows verification within 15 days post-registration [38], our focus is on proactive verification. We test whether registrars verify contact email addresses and phone numbers during account creation or before domain purchase. By providing our genuine contact details, we expect verification through email or SMS. The features *email_operational_validated* and *phone_operational_validated* indicate if such verification is performed.

Domain registration warnings and restrictions: certain domain names may trigger suspicion during registration if they include well-known brand names or random character sequences. Registrars may issue warnings or block these domains. We define three labels for such scenarios: i) **a9e86e6d5d4c676441da** (the first 20 characters of the MD5 hash of "DNS abuse"), ii) **office365-my-account**, and iii) **facebook-login-page**. The latter two are among the most targeted brands in our dataset of 534 k phishing URLs.

For each registrar-TLD pair, we attempt to add these domains to the cart and proceed through all steps until prompted for payment. If succeeded, we set the corresponding boolean features to True: *random_warning*, *random_prevention*, *office365_warning*, *office365_prevention*, *facebook_warning*, *facebook_prevention*. We do not complete the purchase to avoid brand infringement issues.

Registration restrictions: certain registries rigorously verify registrants to reduce malicious registrations (e.g., KYBC **.dk** [68]). Intuitively, attackers would avoid such TLDs and registrars. However, if these practices were implemented globally, malicious actors might adapt by resorting to identity theft for fraudulent registrations or compromising legitimate websites. We define 14 boolean features related to registration restrictions (see Feature set #6).

5.3 Reactive Security Practices

To evaluate reactive security practices at the registrar-TLD level, we measure abuse mitigation times and, for a subset of our data, notify the corresponding registrars to drive the mitigation of abuse at the DNS level.

Malicious domain name uptimes: successful domain name registration is not the ultimate goal for attackers—they must remain operational to profit. For each unique abusive domain name, we measure uptime (or persistence of abuse [44]), defined as the duration between blocklisting of a malicious URL and mitigation of abuse at the DNS level. Mitigation is confirmed when `A` record queries return `NXDOMAIN` or WHOIS shows the domain placed on hold by the registry (`serverHold`) or registrar (`clientHold`).

While blocklisting (e.g., Google Safe Browsing) significantly reduces user interactions by triggering browser warnings, it does not fully deactivate the domain. Mitigating abuse at the DNS level ensures the domain cannot be resolved, effectively preventing further exploitation. Although malicious domains might remain active for months before blocklisting, our focus is on the post-blocklisting period to evaluate registrar and registry actions—a crucial stage of abuse resolution that blocklisting alone does not address.

Initially, we measure uptime at the instant of acquiring the malicious URL from a blocklist followed by repeated measurements over the next month (approximate times): at 5 min, 15 min, 30 min, 1 h, 2 h, 3 h, 4 h, 5 h, 6 h, 12 h, 24 h, 36 h, and 48 h after blocklisting, and then every 12 h thereafter. Since phishing attacks are typically mitigated within the first day after blocklisting [12], we perform more frequent scans initially and less frequent scans later on. Some URLs from blocklists are already mitigated at the time of the first scan. In these cases, we calculate the time between blocklisting and the first measurement. This period is usually very short and provides a good approximation of the mitigation time. We calculate a median uptime at the registrar-TLD level and create a numerical *uptime_not_notified* feature.

Malicious domain name uptimes with notifications: for a subset of maliciously registered domain names, notifications are sent to registrars at the time of the first measurement using abuse contact information extracted from WHOIS/RDAP records. We then calculate the median uptime (represented by the *uptime_notified* feature) at the registrar-TLD level.

6 DESCRIPTIVE ANALYSIS OF FEATURES

Having collected the features, we analyze the registration, proactive, and reactive security measures used by registrars and TLD registries, focusing on 14.5 k malicious domains deliberately registered by attackers and 15.4 k benign registrations.

6.1 Prices, Discounts, and Fees

Our core assumption is that malicious actors are drawn to lower prices, particularly with discounts or special offers. Figures 2 and 3 show the distribution of registration, renewal, and transfer prices for maliciously registered and benign domains, respectively. Registering a domain is typically cheaper than transferring or renewing it. Since malicious domains usually have short lifespans, attackers are less concerned with transfer or renewal costs. Registration prices of maliciously registered domains range from $0.78 to $69, with nearly 50% costing $2 or less. Examples of expensive domains include **usps.bar** at $69, **support-fb.sh** at $59.99, and **dhlcenter.net** at $56. We hypothesize that while attackers generally prefer cheaper options, the cost may become less of a concern when they have access to a large supply of stolen credit cards or cryptocurrencies. Conversely, the registration prices of benign domains tend to be higher, with a mean cost of $8.62 compared to $4.71 for malicious ones. The price of the four most expensive benign domains peaked at $2998.18, all registered under **.sexy** at Namecheap. The other

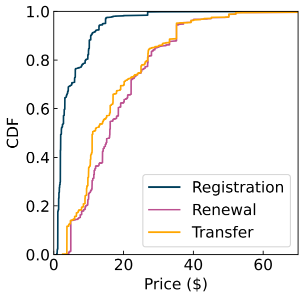

Figure 2: Malicious domain registration, renewal, and transfer prices (in $).

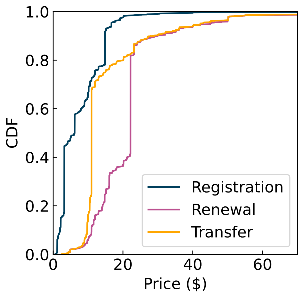

Figure 3: Benign domain registration, renewal, and transfer prices (in $).

costs were even higher, with `cancer.inc` estimated to be worth $4000.17 for a renewal or a transfer.

Registrars may also offer various forms of discounts, either by deducting a fixed amount or a percentage from the regular price. Figures 4 and 5 show the distribution of discount amounts for registration, renewal, and transfer. Most advertised prices lack promotions, especially for renewals. Typically, new registrants are attracted to lower registration prices but pay full rates upon renewal. Discounts for registration and transfer range from $0.01 to $12.95 for both categories of domains analyzed, the transfer promotions being proposed for substantially more malicious rather than benign registrations.

The presented prices may also be subject to various terms, as was the case for 7,168 malicious and 6,855 benign domains. The latter were specifically affected by registration terms, as 6,403 prices were limited to one domain name purchase only (contrary to 4,503 malicious registration prices). Cosmotown and 123 Reg further restricted discounts to new customers only. While attackers might be sensitive to price restrictions, they may not necessarily prevent them from purchasing numerous cheap domains, especially if they can automate the account creation through an API, for example.

6.2 Payment Methods

Out of 24 payment methods known to the TLD-List dataset, 13 are supported by the 38 analyzed registrars. Credit cards and PayPal stand out the most, as they were available for over 97% of both malicious and benign registrations. However, one important consideration for attackers is maintaining anonymity. While they might use stolen credit cards, we hypothesize that they may choose registrars that accept cryptocurrencies or digital wallets, as these add a layer of anonymity to the payment process. Intuitively, Bitcoin and Google Wallet were available for substantially more malicious registrations than benign ones, e.g. 69.85% vs. 22.85% for Bitcoin. It is of no surprise that bank transfer, on the contrary, was proposed for 62.36% benign but only 14.88% malicious registrations, as this payment method is not in line with the anonymization requirements of attackers. Nevertheless, it is recognized that attackers often seek to conceal their identities when purchasing domain names. For example, while Porkbun accepts various forms of cryptocurrencies, they warn that the identity of the registrants may be verified so that they do not "setup a phishing site, a fake store, or some other illegal or otherwise fraudulent/abusive site" [64].[4] Interestingly, only a small fraction of malicious domains were purchased at Porkbun.

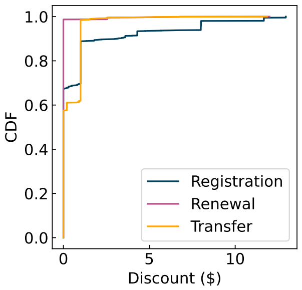

Figure 4: Malicious domain registration, renewal, and transfer discounts (in $).

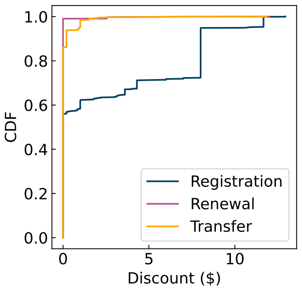

Figure 5: Benign domain registration, renewal, and transfer discounts (in $).

6.3 Free Bundled Services

Whenever purchasing a domain name, registrants may receive various free features and add-ons. DNS service was offered the most commonly, although not always for free. For example, Amazon Route 53 charges for managing hosted zones [4], which likely explains why only 53 malicious and 184 benign domains were registered there. Webnames.ca, the registrar only present in the benign dataset, accounted for eight more domains that require paying for DNS hosting. Free WHOIS privacy was also available for the great majority of benign (94.25%) and malicious (96.57%) registrations. This feature was provided by 22 registrars out of 38. GoDaddy provides the "Free Domain Privacy" service for all the *eligible* registrations but warns that some TLDs prohibit the use of WHOIS privacy services for its domains [21]. Whenever WHOIS privacy is not provided for free, it is billed between $0.31 and $19.66 per year.

A free API is offered by 18 of the 38 examined registrars. We consider it free when provided without prerequisites, excluding cases in which it is only available to domain resellers (e.g., EuroDNS, Instra Corporation, Internet.bs, and Netim). Four registrars permit the automated account creation, either as a new member of an existing organization (Amazon Route 53) or as a sub-account under an existing API user (OVHCloud, Namecheap, INWX). Thirteen registrars allow the automatic registration of domain names. Interestingly, the GoDaddy API was unrestricted during our analysis (August 2023 – January 2024), but as of June 2024, some features are limited to customers with a minimum number of domains or a Discount Domain Club subscription [23].

Attackers often register multiple domains in a single campaign [2, 43, 62]. We found that 23 registrars offer a "bulk search" feature, allowing clients to check availability and prices for multiple domains—ranging from 20 (OVH, Sav), 5 k (Namecheap), to 10 k+ (alldomains.hosting). Additionally, five registrars offer bulk registration discounts, with 101domain.com providing reduced prices for 10+ domains. Dynadot, Internet.bs, and Namecheap require a minimum number of managed domains for discounts, while Above.com and Netim offer them via sales inquiries.

[4] Since the initial test, Porkbun has outsourced its crypto payments to Coinbase and no longer displays this disclaimer on its website.

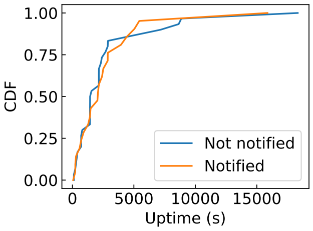

Figure 6: CDF of uptimes for notified and not notified registrars.

DNSSEC signing was available to 63.19% of benign registrations, substantially outnumbering malicious ones (14.56%). In total, 11 registrars offer free DNSSEC signing of domains with a single press of a button. Attackers could also benefit from free email forwarding, accessible to 57.24% malicious registrations but twice as little benign ones. Email account creation, SSL certificates, and web hosting are offered less frequently. Porkbun allows registrants to create up to 20 email aliases and forward emails to existing inboxes [65]. Seven out of 31 registrars provide free email accounts. Free web hosting and SSL certificates facilitate setting up fully-fledged phishing websites but are available for only 254 and 235 malicious registrations, respectively.

6.4 Domain Name Registrant Data

We next analyzed whether registrars *proactively* verify that contact information is operational—ensuring email deliverability and phone number reachability. Registrants may be asked to provide personal information either when creating an account or during domain registration.

Our analysis shows that all registrars reject syntactically incorrect email addresses with a warning. To check verification, we used a valid email, and 29 out of 38 registrars sent confirmation emails. Regtons sends a password to ensure access to the mailbox, while Namecheap and Instra Corporation only send welcome emails without requiring action. Three registrars (DreamHost, Hostinger, Sav) create accounts only at purchase, limiting email validation checks.

Phone number verification is much less common, with 28 registrars performing syntactic checks and only 5 conducting operational validation. For instance, Above.com allows incorrect numbers during account creation but verifies them during domain registration. Name.com gives users a choice between email or phone verification before registering a domain. Finally, we assess the syntactic validation of physical addresses. Seven registrars flagged this issue, with Cosmotown requiring at least four digits for the postal code. Some registrars, however, do not require or verify postal codes at all.

6.5 Prevention of Suspicious Registrations

Attackers targeting specific brands may create domains that appear trustworthy to deceive victims. They may also employ random sequences of characters to obscure their intentions and avoid detection [26, 63, 74]. To address the former, registrars provide guidelines for dealing with intellectual property violations through the ICANN Uniform Rapid Suspension System (URS) [41] and the Uniform Domain Name Dispute Resolution Policy (UDRP)[40]. However, these are *reactive* measures, implemented after domain registration. We evaluate whether registrars *proactively* block attempts to register domains resembling popular brands (`office365-my-account` and `facebook-login-page`) or random strings (`a9e86e6d5d4c676441da`).

Two registrars blocked both trademarked domains from being added to the cart, while one blocked the `facebook-login-page` string only. Namecheap shows an error: "This domain contains restricted phrase(s) and can't be self-registered. Please contact support." Porkbun displays: "We were unable to add the domain to the cart. Please contact support." Interestingly, an older Porkbun account could still add branded domains, suggesting they might use reputation-based measures to prevent abuse by attackers creating multiple accounts. Interestingly, one registrar provides auto-generated suggestions for domain names. For instance, when attempting to add `facebook-login-page.company` to the cart, it suggests, "This domain is suitable for a website that offers a secure and official login page for Facebook users." Lastly, when attempting to add a domain name containing a random string, none of the tested registrars triggered any error or warning.

6.6 Malicious Domain Uptimes

Shorter uptimes should ideally discourage attackers from using certain TLDs and registrars, as swift suspension might drive them to seek alternatives. However, even brief activity may yield valuable credentials and financial gain, potentially diminishing the impact of reactive security measures on their registrar choices.

Malicious domains are often blocked by registrars after abuse reports or complaints. To investigate, we sampled domains from our dataset, submitted complaints, and compared their uptimes to those of unreported domains. Overall, we notified 22 out of 31 registrars about 768 phishing domains. The uptimes of these domains varied up to 17 hours, with an average of 70 minutes. In comparison, domains not reported had a mean uptime of 61 minutes and a maximum of nearly 18 hours. We found little difference in uptimes between reported and unreported domains.

To gain deeper insight, we analyze the uptimes of reported and non-reported domains aggregated at the registrar level instead of at the domain level as shown in Figure 6. Among the notified registrars, DreamHost, CrazyDomains, and Namecheap all had short median uptimes, taking just a few minutes each. Conversely, the domains registered with Name.com, Alibaba Cloud, and Spaceship exhibited longer median uptimes, at 84 minutes, 1.5 hours, and nearly 4.5 hours, respectively. Maliciously registered domains tend to have similar uptimes regardless of whether abuse reports are filed, likely due to registrars subscribing to reputable abuse feeds or conducting concurrent notification campaigns, enabling swift mitigation of phishing domains.

7 DRIVING FACTORS OF DOMAIN ABUSE

Previously, we have identified various registration attributes and practices that may influence the attacker's preferences when registering domains for malicious activities. In this section, we develop two models to estimate and statistically demonstrate which features have an impact and quantify their magnitude.

We first build a GLM model with negative binomial regression, assessing how features influence the number of malicious domains across registrar-TLD pairs. This model handles overdispersed count data, providing reliable estimates by accounting for outlier registrars or TLDs acting as "super-spreaders". The second model is a multilevel hierarchical logistic regression, evaluating the likelihood of a domain being registered for malicious purposes, with a focus on explanatory insights rather than prediction. It distinguishes whether features at the registrar or TLD level are only preferred by attackers or also used by legitimate users.

Together, these models offer complementary insights: the GLM identifies patterns of abuse concentration, while the logistic regression examines how registrar and TLD features influence malicious domain registrations, separating features used by attackers from those favored by legitimate users.

7.1 Feature Engineering

Given the high dimensionality of the initial feature set, it was essential to undergo a feature engineering phase to refine the model by selecting the most relevant features. During the first phase of the feature engineering process, we merged features that represent similar underlying constructs. For example, the features related to digital wallets (*payment_alipay*, *payment_applepay*, etc.) all represent different forms of digital payment methods. By aggregating these into a single binary indicator (*payment_digital_wallet*), we effectively capture the broader concept of "digital payment method availability" rather than treating each form of digital wallet as an independent predictor. This approach reduces multicollinearity, as similar variables can inflate the variance of coefficient estimates, leading to less reliable models. Similarly, grouping payment methods into categories such as *payment_crypto* and *payment_transfer* consolidates the model to focus on the higher-level types of payment methods rather than individual options. This aggregation maintains the interpretability of the feature and aligns with the idea that different payment methods within a category are likely to have similar effects on the dependent variable.

Additionally, following the same rationale, we aggregated measures preventing registrants from adding suspicious domains to the cart (*prevention*), registration restrictions imposed by registries (*restrictions*), personal data validation (*emailphone_validated*), and API offerings (*API*). These features are likely to have correlated effects on domain abuse, and summing them into composite indicators captures the overall presence or absence of these protective measures rather than assessing them separately. This not only simplifies the model but also aligns with the principle of parsimony in statistical modeling, where the goal is to explain the data with the fewest possible predictors.

To compute the average uptime of maliciously registered domains, both notified and non-notified domains are combined to provide a single representative measure of uptime for each registrar. Uptime, as a feature (*uptime*), could theoretically influence the likelihood of malicious registrations, and having a unified metric simplifies the model without losing relevant information.

Next, to create a parsimonious model, we selected features guided by insights from four experts in domain name abuse, comprising both academic researchers and industry practitioners. The incorporation of expert knowledge in the feature engineering phase was crucial for several reasons. Experts in domain name abuse bring a deep understanding of the actual factors that drive abusive registrations. Their input ensures that the model focuses on the features with the highest impact, which might not be immediately apparent from a purely statistical or automated feature selection process. Moreover, by involving industry practitioners, the model is grounded in real-world practices and challenges, increasing its applicability to current domain abuse mitigation efforts.

The experts identified the following categories of features as the most relevant to malicious domain registrations: *free_dns*, *free_web_host*, *free_ssl_cert*, *free_bulk_search_number*, *price_register*, *discount_register*, *restrictions*, *prevention*, *API*, *payment_digital_wallet*, *payment_crypto*, *payment_transfer*, *uptime*, and *emailPhone_validated*. These features span several important dimensions.

Free bundled services, including variables such as *free_dns*, *free_web_host*, and *free_ssl_cert*, were chosen because offering free services can lower the barriers to entry for malicious actors who seek to register domains at minimal cost. Pricing and discount features, like *price_register* and *discount_register*, were selected based on the hypothesis that lower costs might attract a higher number of abusive registrations, as malicious entities typically operate with limited budgets. Registrar restrictions and prevention measures, including *restrictions* and *prevention*, were chosen to capture the extent to which registrars enforce policies that could mitigate domain abuse. Technical and payment features, such as *API*, *payment_digital_wallet*, *payment_crypto*, and *payment_transfer*, reflect the technical and financial infrastructure that can either facilitate or hinder domain abuse. Finally, the operational feature *emailPhone_validated* was included to assess the operational reliability and the rigor of identity validation procedures, which are crucial in preventing the registration of malicious domains.

7.2 Model$_1$: GLM Negative Binomial Regression

We use a GLM model with negative binomial regression to estimate the impact of features on malicious domain counts per registrar-TLD pairs. This model is well-suited to our analysis due to its ability to handle high variance in domain abuse counts, often caused by "super-spreader" registrars or TLDs (see Table 2). Unlike the Poisson model, which assumes equal mean and variance, the negative binomial model addresses overdispersion and provides more reliable estimates. Additionally, the GLM framework offers clear coefficient interpretation, facilitating communication of findings to stakeholders.

Model$_1$ Results: Table 5 shows the results after estimating the model. It includes 1,066 observations with 14 features and a constant term. It achieves a pseudo R-squared value of 0.7733 indicating that the model explains approximately 77.33% of the variance in the number of malicious domains per registrar-TLD pair. Figure 7 shows the summary of the results. Exponentiating the coefficients of the fitted model allows us to interpret them as multiplicative factors for the dependent variable, the number of maliciously registered domains in this case. In particular, several *registration attributes* have a statistically significant effect:

Table 5: Generalized Linear Model Regression Results

Dep. Variable:	malicious	**No. Observations:**	1066
Model:	GLM	**Df Residuals:**	1051
Model Family:	NegativeBinomial	**Df Model:**	14
Link Function:	Log	**Scale:**	1.0000
Method:	IRLS	**Log-Likelihood:**	-3093.9
Pearson chi2:	1.07e+04	**Deviance:**	2970.6
No. Iterations:	65	**Pseudo R-squ. (CS):**	0.7733

	Coef	**std err**	**z**	**P> $\lvert z\rvert$**	**[0.025**	**0.975]**
Intercept	2.3927	0.355	6.748	0.000	1.698	3.088
Free DNS	1.1134	0.420	2.648	0.008	0.289	1.937
Free Web host	0.6323	0.183	3.458	0.001	0.274	0.991
Free SSL cert	-1.6688	0.198	-8.440	0.000	-2.056	-1.281
Restrictions	-1.0053	0.219	-4.594	0.000	-1.434	-0.576
Prevention	-0.0673	0.200	-0.336	0.737	-0.460	0.325
API	1.6080	0.118	13.585	0.000	1.376	1.840
Payment digital wallet	0.0525	0.264	0.199	0.843	-0.466	0.571
Payment crypto	0.2609	0.109	2.393	0.017	0.047	0.475
Payment transfer	-1.3446	0.133	-10.131	0.000	-1.605	-1.084
EmailPhone validated	-1.2143	0.113	-10.757	0.000	-1.436	-0.993
Free bulk search	-0.0003	5.38e-05	-5.687	0.000	-0.000	-0.000
Price register	-0.0676	0.005	-14.565	0.000	-0.077	-0.058
Discount register	0.3979	0.040	9.957	0.000	0.320	0.476
Uptime	-0.0001	2.54e-05	-4.660	0.000	-0.000	-6.85e-05

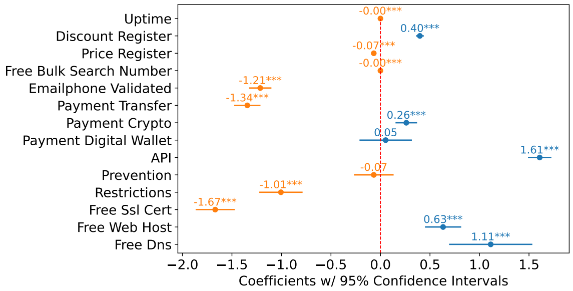

Figure 7: GLM model: estimated coefficients with the 95% confidence intervals.

- Registration price has a coefficient of -0.07 (p < 0.001). Exponentiating this coefficient gives $e^{-0.07} \approx 0.94$, indicating that decreasing the registration price by one dollar is associated with a 6.6% increase in the number of malicious domains, suggesting that more affordable registration fees may encourage higher rates of abuse.
- Registration discounts have a positive coefficient of 0.40 (p < 0.001), which suggests that offering a one-dollar discount on domain registration is associated with a 49% increase in malicious domain registrations, highlighting a potential incentive for malicious actors to exploit discounts.
- Cryptocurrency payments show a positive coefficient of 0.26 (p = 0.017), which implies a 30% increase in malicious registrations when cryptocurrency payments are accepted. Conversely, transfer-like payments have a negative coefficient of -1.34 (p < 0.001), suggesting a 74% decrease in malicious domains with the acceptance of bank transfers.
- Free services have a positive coefficient of 1.11 (p = 0.008), which means that they are associated with approximately a 205% increase in the number of maliciously registered domains compared to registrars without them. Similarly, the availability of free web hosting shows a positive coefficient of 0.63 (p = 0.001) indicating that free web hosting is associated with an 88% increase in the number of malicious phishing domains. In contrast, offering free SSL certificates has a negative coefficient of -1.67 (p < 0.001) meaning that it is associated with an 81% decrease in the number of malicious registrations.
- On the technical side, the presence of APIs either to register domains or to create accounts has a positive coefficient of 1.60 (p < 0.001), which indicates that registrars offering API access are associated with a 401% increase in the number of malicious domains.

Focusing on the *proactive* verification, the restrictions implemented by some registrars have a negative coefficient of -1.01 (p < 0.001), which suggests that stringent registrar restrictions are associated with a 63% decrease in the number of maliciously registered domains. Similarly, when the validation of registrant information such as their phone number of email takes place during the account creation or before the domain purchase, it has a significant negative coefficient of -1.21 (p < 0.001) indicating that it is associated with a 70% decrease in malicious registrations.

When it comes to reactive security practices, *uptime* has a small coefficient of -0.0001, which indicates that higher uptimes are weakly associated with a very slight decrease in the number of malicious registrations suggesting that it has a minor impact and its effect on reducing domain abuse is relatively small.

7.3 Model$_2$: Multilevel Logistic Regression

The second model quantifies the impact of expert-selected features on the probability of a domain being registered with a specific registrar with malicious or legitimate intent. This model uses a multilevel hierarchical logistic regression approach, well suited for handling the nested structure of the data, in which domains are clustered within registrars and TLDs. By accounting for this hierarchical structure, the model can estimate more accurately the impact of registrar-specific and TLD-specific practices and attributes on the likelihood of phishing.

In this model, the dependent variable is defined as the binary status of a domain, where `True` indicates that the domain was registered with malicious intent, and `False` means that the domain was registered for legitimate purposes. The independent variables, which include the features identified by experts as potentially influencing domain abuse, are modeled as fixed effects.

The hierarchical structure of the model is captured by including two levels of random effects: one for the registrar and another for the TLD. The registrar-level random effect allows the model to account for variability between registrars that might not be fully explained by the fixed effects such as differences in registrar practices or market strategies. Similarly, the TLD-level random effect accounts for the variability between different TLDs recognizing that the domains within the same TLD might exhibit similar patterns of abuse due to the factors specific to that TLD. By incorporating both registrar-level and TLD-level random effects, the model adjusts for the within-group correlations at each level providing more reliable estimates of the fixed effects.

Model$_2$ Results: Table 6 shows the results after estimating the model. The conditional R^2 for the full model, which incorporates both fixed and random effects, is 0.47, which means that approximately 47.4% of the variance is explained when considering the complete structure of the model, including the effects at both the registrar and TLD levels. Our focus is on explanatory insights rather

Table 6: Multilevel Logistic Regression Results

Features	Coef	CI	P> $\|z\|$
Intercept	0.16	-0.27 – 0.59	0.469
Uptime	-0.01	-0.07 – 0.05	0.815
Discount register	0.01	0.01 – 0.02	<0.001
Payment digital wallet	0.03	-0.13 – 0.18	0.724
Price register	-0.00	-0.01 – 0.00	0.183
Free bulk search	-0.07	-0.19 – 0.06	0.290
Payment crypto	0.23	-0.07 – 0.52	0.130
API	0.12	-0.08 – 0.32	0.227
Free DNS	0.15	-0.27 – 0.57	0.473
Payment transfer	-0.10	-0.27 – 0.07	0.241
Free web host	0.20	-0.20 – 0.60	0.325
Free SSL cert	-0.24	-0.63 – 0.15	0.221
EmailPhone validated	0.03	-0.20 – 0.25	0.828
Restrictions	-0.21	-0.32 – -0.10	<0.001
Prevention	-0.13	-0.40 – 0.14	0.356

Random Effects	Value
σ^2	0.13
τ_{00} TLD	0.03
τ_{00} Registrar	0.07
ICC	0.41
$N_{Registrar}$	38
N_{TLD}	293
Observations	29890
Marginal R^2 / Conditional R^2	0.109 / 0.474

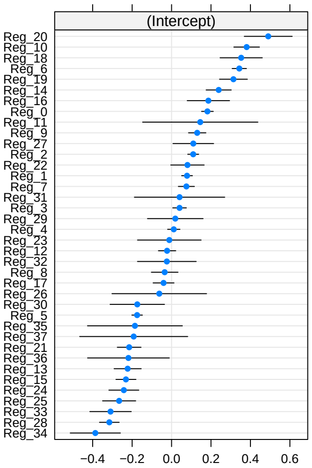

Figure 8: Random effects at the registrar level derived from the second model.

than prediction, so a high pseudo R^2 is not the primary goal. On the other hand, the marginal R^2, which reflects the proportion of variance explained solely by the fixed effects, is considerably lower at 0.11. This difference underscores the significant contribution of the random effects to the model explanatory power, indicating that a substantial portion of the variability in domain abuse is due to the differences at the registrar and TLD levels, beyond what can be captured by the fixed effects alone.

When examining the random effects at the registrar level, the model reveals a variance of 0.06 with a standard deviation of 0.26. Figure 8 shows these estimates. For example, Reg_20 has a random intercept of 0.49, which indicates that this registrar has a higher likelihood of having malicious domain registrations compared to the average registrar. Specifically, this positive value suggests that domains registered through Reg_20 are more likely to be malicious than those registered through registrars with lower or negative intercepts. Conversely, Reg_34 has a random intercept of -0.38, indicating that domains registered through this registrar are less likely to be malicious compared to the average. This negative value implies that the practices or characteristics of Reg_34 are associated with a lower probability of domain abuse, making it a less attractive option for malicious registrants. These intercepts highlight how individual registrar characteristics significantly impact the likelihood of domain abuse.

At the TLD level, the variance of the random intercept is estimated at 0.03, with a standard deviation of 0.16. Although there is some variability in domain abuse likelihood across different TLDs, this variance is modest compared to that observed at the registrar level. The smaller variance at the TLD level indicates that while TLD-specific characteristics do influence domain abuse, their impact is less pronounced than that of registrar-specific factors, which is expected.

On the other hand, the fixed effects in the model reveal key insights into how certain features attract more malicious actors than legitimate registrants. Notably, registration discounts and restrictions emerged as significant features (see Figure 9). Registration discounts have a positive coefficient (0.013, p < 0.001), indicating that the domains registered with discounts are more likely to be malicious, which means that for every unit increase in the discount, the odds of a domain being maliciously registered increase by about 1.3%. When considering legitimate registrations, the significance of discounts in attracting malicious registrations implies that promotions may be less critical for legitimate users. While discounts appear to be a strong motivator for malicious actors—likely because they reduce the financial barrier for bulk domain registrations used in various forms of online abuse—legitimate registrants might prioritize other factors over cost savings.

Contrariwise, restrictions show a negative coefficient (-0.210, p < 0.001), implying that registrars with stringent registration restrictions are associated with a reduced likelihood of malicious domain registrations, which means that the presence of restrictions decreases the odds of a domain being maliciously registered by about 19%. The implication of the results is that stringent registration restrictions effectively deter malicious actors, while legitimate registrants, faced with a more rigorous process, are less likely to choose less restrictive registrars or TLDs.

Other variables such as free DNS, hosting, uptimes, or payment types were not statistically significant, suggesting that they may have a similar or negligible impact on malicious and benign registrants.

8 DISCUSSION

This section synthesizes our findings to propose actionable strategies for intermediaries and aims to align anti-abuse practices with their economic interests.

Economic incentives: The first model demonstrates that economic incentives provided by registrars, such as registration discounts, are linked to an increase in the number of malicious registrations. Even if discounts are limited to new users only, attackers can exploit free unrestricted APIs to automate account creation and register domains at discounted prices. Leveraging low-cost options

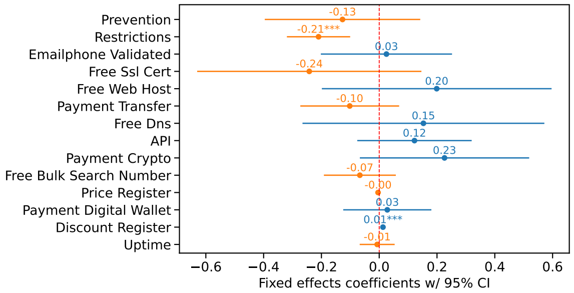

Figure 9: Fixed effect coefficients with 95% confidence intervals derived from the second model.

allows attackers to maximize their return on investment, especially given the short lifespan of these domains before suspension.

Our regression analysis further indicates that discounts attract more malicious actors than legitimate registrants. The statistical analysis supports these findings, revealing that the mean cost of benign domains ($8.62) is higher compared to $4.71 of malicious ones. This highlights attackers' price sensitivity, making them more likely to exploit cheap registration options for large-scale abuse.

Based on our results, registrars and registries could adjust prices incrementally, anticipating a reduction in abuse rates, while closely monitoring their effect on legitimate domain registrations. However, significantly raising registration prices to reduce abuse could conflict with the economic interests of registries and registrars. As an alternative, registries and registrars with lower abuse concentrations could be financially incentivized—such as through reduced domain registration fees provided by ICANN or tiered pricing structures offered by registries—to align their economic interests with anti-abuse efforts better. Similar incentive structures have been deployed by registries in countries such as the Netherlands, Sweden, Norway, and the Czech Republic, where registrars receive lower fees for DNSSEC-signed domains compared to unsigned ones [71].

The Quality Performance Index (QPI) [67], managed by PIR—a `.org` registry operator—evaluates registrars based on Key Performance Indicators (KPIs): Abuse Takedown, Renewal Rates, and Domain Usage, along with secondary KPIs. Weighted scores generate a single QPI score. Registrars meeting the baseline threshold are pre-qualified for promotions. Existing initiatives like Netbeacon MAP [57], ICANN Metrica [42], DAP.LIVE [16], and Domain Risk Score [17] already measure abuse rates and could serve as the foundation for price reduction incentives.

Such practices would encourage intermediaries to develop their own anti-abuse best practices while balancing these policies with economic incentives and allowing for self-regulation.

Role of free services: The provision of free DNS services and web hosting significantly increases the number of malicious registrations, as these bundled services lower entry barriers for attackers, enabling them to maintain malicious websites at minimal cost. However, the results of the second model indicate that these bundled features are attractive to both malicious and legitimate users.

Disallowing such bundled services would not be practical, as they are often essential to the business models of registrars, who typically operate with thin profit margins on domain registrations. Managed hosting services, in particular, serve as a key revenue stream for registrars, emphasizing the need for balanced anti-abuse measures that do not disproportionately impact their financial viability.

More importantly, unrestricted APIs for domain registration and account management are strongly correlated with a higher volume of malicious registrations, with registrars offering API access linked to a staggering 401% increase in malicious domains. APIs simplify and automate the account and domain registration process, facilitating large-scale phishing campaigns.

Therefore, registrars should consider implementing restrictions on API access for reputable entities, such as long-standing registrants and resellers. This approach, already adopted by certain market players as highlighted in our descriptive analysis, can effectively deter automated account creation and domain management by malicious users without disrupting legitimate registrants.

Proactive restrictions and verifications: Our study demonstrates that implementing proactive restrictions is strongly associated with a decrease in malicious domain registrations, underscoring the effectiveness of upfront checks in curbing abuse.

A widely discussed proposal among stakeholders is the adoption of Know Your Business Customer (KYBC) procedures [12, 46]. Registrant identification could be standardized through harmonized frameworks. For instance, in the EU, KYBC could use eID authentication under the eIDAS Regulation [12]. However, the introduction of KYBC requires careful consideration. The community must anticipate how attackers might adapt. While our findings suggest that proactive restrictions could significantly increase barriers to abuse, they may also introduce new risks, such as digital identity theft (e.g., stolen ID cards), which could facilitate malicious registrations.

For example, `.dk` domains require KYBC verification and exhibit lower abuse rates, potentially due to these measures or other factors like higher registration costs [12]. In contrast, `.cn` domains also mandate registrant verification, yet the `.cn` TLD and Chinese registrars remain among the most abused [56]. These nuances highlight the need for context-sensitive and balanced policy interventions.

Reactive measures: Promptly suspending malicious domain names is crucial for mitigating potential harm. However, our analysis, which focuses on the post-blocklisting period to assess registrar and registry actions, reveals that longer uptime has only a marginal effect on the concentration of malicious domains and minimal influence on attackers' choice of registrar or TLD. We hypothesize that even short periods of operation can yield valuable credentials and financial rewards for attackers, limiting the effectiveness of reactive security measures focused on reducing uptime.

8.1 Limitations

The presented methodology, while rigorously designed, has several inherent limitations. Some of them stem from our desire to avoid any ethical issues as detailed in Section 9.

While our primary source, the third-party registration data, has been validated and proved to be highly accurate, it may not fully reflect the attributes actually chosen by registrants of both malicious and benign domains. Some features, like certain restrictions, are defaults, while others may be used optionally rather than specifically selected. Without comprehensive ground truth data, we cannot estimate the extent of any potential bias in this assumption.

Our methodology relies on domain registration data but some TLDs lack WHOIS information (e.g., `.gr`), have restrictions (e.g., `.es`), or offer limited details (e.g., `.de`). Moreover, our methodology relies on IANA IDs and the TLD list that provide limited visibility of the non-ICANN-accredited registrars, so the dataset is skewed toward the ICANN-accredited registrars. However, our combined dataset of malicious and benign domains spans 1,067 registrar-TLD pairs and nearly 30 k domains, capturing a wide range of feature combinations and providing a robust basis for analysis.

We adopted a stringent approach to identifying malicious domains, which may have led to missing some aged domains or those mitigated at the hosting level. However, it was a deliberate choice to: i) exclude compromised phishing domains from our analysis and ii) filter out benign domains that were mistakenly added to various external blocklists (i.e., false positives). Given previous research indicating that malicious domains are typically used shortly after registration and registrars have strict mitigation obligations, we believe this approach does not significantly bias our results.

9 ETHICS AND ARTIFACTS

Our study was designed with a strong emphasis on minimizing ethical risks. The potential for harm was carefully evaluated using a consequentialist approach ensuring that our research positively contributes to the domain ecosystem without introducing new risks.

The TLD-list used in our analysis was secondary data collected by a third party comprising publicly available information. To complement the available datasets, we performed a series of active DNS and WHOIS measurements, following the best current practices as outlined by Durumeric et al. [18]. Specifically, we issued a minimum number of requests required to fulfill our research goals, thus limiting the impact on the destination systems.

We deliberately avoided methodologies that could raise ethical or legal concerns, such as using fake registrant data, registering trademarked domains, or simulating phishing to test registrar responses. These approaches risk legal issues, misuse of resources, or interfere with real abuse mitigation efforts.

While the findings of this paper are primarily intended to support registrars and registries, they could also inadvertently inform cybercriminals seeking to refine their attack strategies. However, we believe that transparent knowledge sharing outweighs the risks of withholding findings under a "security through obscurity" approach. We will make all data analysis code and the two machine learning models available to interested researchers upon request. These tools can support studies on the factors that drive domain name abuse.

10 CONCLUSIONS

Our study reveals critical insights into the factors influencing malicious domain registrations by analyzing over seventy registrar- and TLD-based features. We find that registrars offering discounts and bundled services are more likely to attract malicious actors, while those enforcing stricter registration requirements tend to experience lower abuse rates. Features that enable rapid and automated registrations, such as unrestricted APIs, also correlate with increased abuse. Some of the features associated with abuse appear to be less important to legitimate users, suggesting that targeted interventions could raise barriers for malicious actors without significantly affecting the benign use.

More broadly, our findings support a comprehensive, evidence-based approach to DNS abuse mitigation that extends beyond cost-focused strategies. Operational characteristics such as unrestricted API access, aggressive discounting, and minimal identity verification present tangible opportunities for intervention. Registrars and registries may consider implementing the measures such as rate-limiting automated registrations, or introducing identity validation for high-risk registrations. However, the viability of these approaches should be carefully assessed in terms of the implementation effort, the operational cost, and the potential risk of discouraging legitimate customers.

ACKNOWLEDGEMENTS

This work has been funded by ICANN. We would like to thank the Anti-Phishing Working Group, PhishTank, OpenPhish for providing data access for this analysis and to Spamhaus and SURBL for granting access to their blocklists, which helped us filter out malicious domains from the benign list we sampled. We thank Rowena Schoo, Graeme Bunton (NetBeacon Institute), Siôn Lloyd (ICANN), as well as the broader ICANN community and anonymous reviewers for their valuable comments.

REFERENCES

[1] I&J Policy Network. Domains & Jurisdiction Program: Operational approaches norms, criteria, mechanisms, 2019. https://www.internetjurisdiction.net/uploads/pdfs/Papers/Domains-Jurisdiction-Program-Operational-Approaches.pdf.

[2] Antonia Affinito, Raffaele Sommese, Gautam Akiwate, Stefan Savage, Kimberley Claffy, Geoffrey M. Voelker, Alessio Botta, and Mattijs Jonker. Domain Name Lifetimes: Baseline and Threats. In *TMA*, pages 1–9, Austria, 2022. IFIP.

[3] Afilias. We are the Top Level Domain (TLD) Registry Services and DNS Solutions Experts, 2019. https://web.archive.org/web/20191231103541/https://afilias.info/.

[4] Amazon Route 53. Amazon Route 53 Pricing, 2024. https://aws.amazon.com/route53/pricing/.

[5] Anti-Phishing Working Group. Unifying The Global Response to The Cybercrime, 2023. https://apwg.org.

[6] APWG. Phishing Activity Trends Report, 2020. https://docs.apwg.org/reports/apwg_trends_report_q2_2020.pdf.

[7] Cali Dog Security. Certstream, March 2022. https://calidog.io.

[8] Yanan Cheng, Yali Liu, Lianmiao Wang, Zhaoxin Zhang, Tingting Chai, and Yuejin Du. Evaluating the Effectiveness of Handling Abusive Domain Names by Internet Entities. *Electronics*, 11(8):1–31, 2022.

[9] CNNIC. How to register the CN domain names, 2024. https://www.cnnic.com.cn/IS/CNym/cnymyhfaq/#3.

[10] European Commission. 2018 reform of EU data protection rules, May 2018. https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf.

[11] European Commission. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive) (Text with EEA relevance), 2022. https://eur-lex.europa.eu/eli/dir/2022/2555.

[12] European Commission, Content Directorate-General for Communications Networks, Technology, J Bayer, Y Nosyk, O Hureau, S Fernandez, S Paulovics, A Duda, and M Korczynski. *Study on Domain Name System (DNS) abuse – Technical report. Appendix 1.* Publications Office of the European Union, Luxembourg, 2022.

[13] Marco Cova, Corrado Leita, Olivier Thonnard, Angelos D. Keromytis, and Marc Dacier. An Analysis of Rogue AV Campaigns. In *RAID*, pages 442–463, 2010.

[14] DNS Abuse Framework. Framework to Address Abuse, 2020. https://dnsabuseframework.org/media/files/2020-05-29_DNSAbuseFramework.pdf.

[15] DNS Belgium. Registrant Verification, 2024. https://docs.dnsbelgium.be/be/general/registrantverification.html.

[16] DNS Research Federation. Introduction to DAP.LIVE, 2025. https://dnsrf.org/docs/dap-live/introduction-to-dap-live/index.html.

[17] DomainTools. Domain Risk Score, 2025. https://www.domaintools.com/resources/api-documentation/domain-risk-score/.

[18] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. ZMap: Fast internet-wide scanning and its security applications. In *USENIX Security*, 2013.

[19] EURid. Know-Your-Customer, 2024. https://eurid.eu/en/register-a-eu-domain/data-quality/#nav_kyc_project.

[20] Mark Felegyhazi, Christian Kreibich, and Vern Paxson. On the Potential of Proactive Domain Blacklisting. In *LEET*, page 6, USA, 2010. USENIX Association.

[21] GoDaddy. GoDaddy - Domain Name Proxy Agreement, 2023. https://www.godaddy.com/legal/agreements/domain-name-proxy-agreement.

[22] GoDaddy. Domain Names, Websites, Hosting & Online Marketing Tools, September 2024. https://www.godaddy.com/.

[23] GoDaddy. Get Started, July 2024. https://developer.godaddy.com/getstarted.

[24] GoDaddy. Get Started, April 2024. https://web.archive.org/web/20240418021546/https://developer.godaddy.com/getstarted.

[25] Interisle Consulting Group. Phishing landscape 2023: A study of the scope and distribution of phishing. Technical report, Interisle Consulting Group, 2023.

[26] Interisle Consulting Group. Phishing landscape 2024: A study of the scope and distribution of phishing. Technical report, Interisle Consulting Group, 2024.

[27] Shuang Hao, Alex Kantchelian, Brad Miller, Vern Paxson, and Nick Feamster. PREDATOR: Proactive Recognition and Elimination of Domain Abuse at Time-Of-Registration. In *CCS*, page 1568–1579, New York, NY, USA, 2016. ACM.

[28] Shuang Hao, Matthew Thomas, Vern Paxson, Nick Feamster, Christian Kreibich, Chris Grier, and Scott Hollenbeck. Understanding the Domain Registration Behavior of Spammers. In *IMC*, page 63–76, New York, NY, USA, 2013. ACM.

[29] Scott Hollenbeck. EPP Domain Name Mapping. RFC 5731, 2009.

[30] Danny Yuxing Huang, Maxwell Matthaios Aliapoulios, Vector Guo Li, Luca Invernizzi, Elie Bursztein, Kylie McRoberts, Jonathan Levin, Kirill Levchenko, Alex C. Snoeren, and Damon McCoy. Tracking Ransomware End-to-end. In *IEEE S&P*, pages 618–631, Los Alamitos, CA, USA, 2018. IEEE Computer Society.

[31] IANA. Root Zone Database, 2024. https://www.iana.org/domains/root/db.

[32] ICANN. 2013 Registrar Accreditation Agreement, 2013. https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en.

[33] ICANN. SAC058: SSAC Report on Domain Name Registration Data Validation, 2013. https://itp.cdn.icann.org/en/files/security-and-stability-advisory-committee-ssac-reports/sac-058-en.pdf.

[34] ICANN. Centralized Zone Data Service, March 2022. https://czds.icann.org.

[35] ICANN. DNS Security Threat Mitigation Program, 2023. https://www.icann.org/dns-security-threat.

[36] ICANN. ICANN's Domain Abuse Activity Reporting (DAAR) Project, 2023. https://www.icann.org/octo-ssr/daar-faqs.

[37] ICANN. 2024 Global Amendments to the 2013 Registrar Accreditation Agreement (RAA) and Base gTLD Registry Agreement (Base RA), 2024. https://www.icann.org/resources/pages/global-amendment-2024-en.

[38] ICANN. ICANN Organization Enforcement of Registration Data Accuracy Obligations Before and After GDPR, 2024. https://www.icann.org/resources/pages/registration-data-accuracy-obligations-gdpr-2021-06-14-en.

[39] ICANN. Monthly Registry Reports, 2024. https://www.icann.org/resources/pages/registry-reports.

[40] ICANN. Uniform Domain-Name Dispute-Resolution Policy, 2024. https://www.icann.org/resources/pages/help/dndr/udrp-en.

[41] ICANN. Uniform Rapid Suspension (URS), 2024. https://www.icann.org/urs-en.

[42] ICANN. ICANN Domain Metrica: A Measurement Platform, 2025. https://www.icann.org/octo-ssr/metrica-en.

[43] Interisle Consulting Group. Criminal Abuse of Domain Names: Bulk Registration and Contact Information Access, 2019. https://interisle.net/sub/CriminalDomainAbuse.pdf.

[44] Maciej Korczyński, Samaneh Tajalizadehkhoob, Arman Noroozian, Maarten Wullink, Cristian Hesselman, and Michel Van Eeten. Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs. In *IEEE EuroS&P*, pages 579–594. IEEE Computer Society, 2017.

[45] Maciej Korczyński, Maarten Wullink, Samaneh Tajalizadehkhoob, Giovane C. M. Moura, Arman Noroozian, Drew Bagley, and Cristian Hesselman. Cybercrime After the Sunrise: A Statistical Analysis of DNS Abuse in New GTLDs. In *ASIACCS*, page 609–623, New York, NY, USA, 2018. Association for Computing Machinery.

[46] KYBC. Digital Services Act: Know Your Business Customer obligations must apply to all intermediary service providers to offer a meaningful tool for tackling illegal activities and products online, 2022. https://www.kybc.eu/wp-content/uploads/2022/02/KYBC-letter-18-Feb-2022.pdf.

[47] Nhien-An Le-Khac and Tahar Kechadi. Security Threats of URL Shortening: A User's Perspective. *Journal of Advances in Computer Networks*, 2015.

[48] Sophie Le Page, Guy-Vincent Jourdan, Gregor V Bochmann, Jason Flood, and Iosif-Viorel Onut. Using URL Shorteners to Compare Phishing and Malware Attacks. In *APWG eCrime*, pages 1–13, New York, NY, USA, 2018. IEEE.

[49] Sophie Le Page, Guy-Vincent Jourdan, Gregor V. Bochmann, Iosif-Viorel Onut, and Jason Flood. Domain Classifier: Compromised Machines Versus Malicious Registrations. In *Web Engineering*, pages 265–279. Springer International, 2019.

[50] He Lonnie Liu, Kirill Levchenko, Márk Félegyházi, Christian Kreibich, Gregor Maier, and Geoffrey M Voelker. On the Effects of Registrar-Level Intervention. In *LEET*, pages 1–8, Boston, MA, 2011. USENIX Association.

[51] Suqi Liu, Ian Foster, Stefan Savage, Geoffrey M. Voelker, and Lawrence K. Saul. Who is .com? Learning to Parse WHOIS Records. In *IMC*. ACM, 2015.

[52] Chaoyi Lu, Baojun Liu, Yiming Zhang, Zhou Li, Fenglu Zhang, Haixin Duan, Ying Liu, Joann Qiongna Chen, Jinjin Liang, Zaifeng Zhang, Shuang Hao, and Min Yang. From WHOIS to WHOWAS: A Large-Scale Measurement Study of Domain Registration Privacy under the GDPR. In *NDSS*, 2021.

[53] Sourena Maroofi, Maciej Korczyński, Cristian Hesselman, Benoît Ampeau, and Andrzej Duda. COMAR: Classification of Compromised versus Maliciously Registered Domains. In *IEEE EuroS&P*. IEEE Computer Society, 2020.

[54] Giovane C. M. Moura, Thomas Daniels, Maarten Bosteels, Sebastian Castro, Moritz Müller, Thymen Wabeke, Thijs van den Hout, Maciej Korczyński, and Georgios Smaragdakis. Characterizing and Mitigating Phishing Attacks at ccTLD Scale. In *CCS*, page 2147–2161, 2024.

[55] Namecheap. Buy a domain name - Register cheap, 2024. https://www.namecheap.com.

[56] NetBeacon Institute. June 2024 Monthly Analysis, 2024. https://netbeacon.org/wp-content/uploads/2024/06/MAP-Report-June-2024-.pdf.

[57] Netbeacon Institute. NetBeacon Measurement and Analytics Platform (MAP), 2025. https://netbeacon.org/map-analytics/.

[58] Neustar. Registry Solutions, 2019. https://web.archive.org/web/20191230213941/https://www.home.neustar/registry-solutions.

[59] Nick Nikiforakis, Federico Maggi, Gianluca Stringhini, M Zubair Rafique, Wouter Joosen, Christopher Kruegel, Frank Piessens, Giovanni Vigna, and Stefano Zanero. Stranger Danger: Exploring the Ecosystem of Ad-based URL Shortening Services. In *WWW*, pages 51–62, New York, NY, USA, 2014. ACM.

[60] OpenPhish. Phishing Intelligence, 2023. https://openphish.com.

[61] PhishTank. Join the Fight Against Phishing, 2023. https://phishtank.org.

[62] Dave Piscitello. Weaponizing Domain Names: how bulk registration aids global spam campaigns, 2020. https://www.spamhaus.org/news/article/795/weaponizing-domain-names-how-bulk-registration-aids-global-spam-campaigns.

[63] Victor Le Pochat, Tim Van Hamme, Sourena Maroofi, Tom van Goethem, Davy Preuveneers, Andrzej Duda, Wouter Joosen, and Maciej Korczynski. A Practical Approach for Taking Down Avalanche Botnets Under Real-World Constraints. In *NDSS*. The Internet Society, 2020.

[64] Porkbun. Buy Account Credit with Crypto, 2023. https://porkbun.com/products/crypto.

[65] Porkbun. Free Mail Forwarding, 2024. https://porkbun.com/products/email_forwarding.

[66] Public Interest Registry. .ORG - Power Your Inspiration. Connect Your World., September 2024. https://pir.org.

[67] Public Interest Registry. Quality Performance Index (QPI), 2025. https://pir.org/our-impact/qpi/.

[68] Punktum dk. Terms and procedures, 2024. https://punktum.dk/en/articles/terms-and-procedures.

[69] Quirin Scheitle, Oliver Gasser, Theodor Nolte, Johanna Amann, Lexi Brent, Georg Carle, Ralph Holz, Thomas C. Schmidt, and Matthias Wählisch. The Rise of Certificate Transparency and Its Implications on the Internet Ecosystem. In *IMC*, 2018.

[70] SIDN - The Foundation for Internet Domain Registration in the Netherlands, 2024. https://sidn.nl/en.

[71] SIDN Labs. DNSSEC adoption heavily dependent on incentives and active promotion, 2020. https://www.sidn.nl/en/news-and-blogs/dnssec-adoption-heavily-dependent-on-incentives-and-active-promotion.

[72] SIDN Labs. Assessing the Risk of New .nl Registrations Using RegCheck, 2023. https://www.sidnlabs.nl/en/news-and-blogs/assessing-the-risk-of-new-nl-registrations-using-regcheck.

[73] Ravindu De Silva, Mohamed Nabeel, Charith Elvitigala, Issa Khalil, Ting Yu, and Chamath Keppitiyagama. Compromised or Attacker-Owned: A Large Scale Classification and Study of Hosting Domains of Malicious URLs. In *USENIX Security*, pages 3721–3738. USENIX Association, August 2021.

[74] Aditya K. Sood and Sherali Zeadally. A Taxonomy of Domain-Generation Algorithms. *IEEE Security & Privacy*, 14(04):46–53, July 2016.

[75] Jan Spooren, Thomas Vissers, Peter Janssen, Wouter Joosen, and Lieven Desmet. Premadoma: An Operational Solution for DNS Registries to Prevent Malicious Domain Registrations. In *ACSAC*, pages 557–567, New York, USA, 2019. ACM.

[76] SURBL BV. SURBL Lists. http://www.surbl.org/lists.

[77] Samaneh Tajalizadehkhoob, Tom van Goethem, Maciej Korczyński, Arman Noroozian, Rainer Böhme, Tyler Moore, Wouter Joosen, and Michel van Eeten. Herding Vulnerable Cats: A Statistical Approach to Disentangle Joint Responsibility for Web Security in Shared Hosting. In *CCS*, pages 553–567. ACM, 2017.

[78] The Spamhaus Project. Domain Blocklist (DBL). https://www.spamhaus.org/dbl.

[79] TLD-List. Compare Prices of All Top-Level Domains, 2022. https://tld-list.com.

[80] Tucows. Making the Internet better, September 2024. https://www.tucows.com.

[81] US FBI, Internet Crime Complaint Center. Internet Crimer Report. https://www.ic3.gov/Media/PDF/AnnualReport/2023_IC3Report.pdf, 2023.

[82] Thomas Vissers, Jan Spooren, Pieter Agten, Dirk Jumpertz, Peter Janssen, Marc Van Wesemael, Frank Piessens, Wouter Joosen, and Lieven Desmet. Exploring the Ecosystem of Malicious Domain Registrations in the .eu TLD. In *RAID*, 2017.