Not What It Used To Be: Generational Analysis of Top-level Domain Reputation

Janos Szurdi
Palo Alto Networks

Reethika Ramesh
Palo Alto Networks

Ram Sundara Raman
University of California, Santa Cruz

Daiping Liu
Palo Alto Networks

***Abstract*—Over the past decade, ICANN's New gTLD Program has dramatically expanded the DNS namespace, raising persistent concerns about its security implications as another round of applications approaches in 2026. In this paper, we present a large-scale, longitudinal study of both malicious and benign domain usage across four generations of gTLDs—legacy, first-wave, second-wave, and third-wave—alongside country-code TLDs. Using four years of longitudinal data from 2021 to 2025, collected from multiple sources including zone files, active DNS measurements, passive DNS feeds, and domain categorizations from a leading global cybersecurity vendor, we develop three reputation metrics to capture utilization trends: the malicious ratio, the malicious-to-benign ratio, and the non-benign ratio.**

Our analysis shows that newer gTLD generations are substantially more malicious and significantly less utilized for benign purposes than legacy TLDs. Compared to legacy gTLDs, newer generations exhibit malicious-to-benign ratios that are 3.1–9.2× worse, with these ratios worsening rapidly over time: up to 50× growth in malicious-to-benign ratios within four years for the newest gTLDs. We examine contributing factors to show that lower pricing, higher popularity, and certain TLD categories are strongly associated with worse reputation, while defensive registrations account for only a negligible fraction of domain registrations. Finally, we identify a small number of sponsoring organizations that disproportionately operate gTLDs with severe abuse. Our results underscore the need for continued scrutiny and rigorous evaluation of new gTLDs.

I. Introduction

The Internet Corporation for Assigned Names and Numbers (ICANN) [30] is responsible for coordinating the Domain Name System (DNS), including the oversight of generic Top-Level Domains (gTLDs) such as `.com` and `.org` [27]. In the early 2000s, ICANN tested the feasibility of expanding gTLDs through two limited application rounds. The first, in 2000, approved seven "proof of concept" TLDs (*e.g.,* `.biz`, `.info`), while a second round in 2003–2004 added several sponsored gTLDs (*e.g.,* `.jobs`, `.mobi`, `.xxx`) [35], [62]. ICANN launched the formal *New gTLD Program* in 2012, marking the most significant expansion of the DNS namespace to date. During this round, ICANN received 1,930 applications, more than 1,200 of which were approved and gradually introduced into the root zone starting 2013, with the newest being included in October 2025 [21]. We investigate the impact of the expanding gTLD namespace on the DNS ecosystem and the broader Internet community, which is particularly timely given that the next round of gTLD applications is scheduled to open in 2026 [31], and our observation that the reputation of the new gTLDs has significantly deteriorated since the last systematic study of TLD reputation [38], [39].

Researchers have shown that newly added gTLDs are often used for registration abuse [39] and malicious activities [17], while adoption for benign use is slow. Korczynski *et al.* showed that cybercriminals have started to move away from legacy to new gTLDs, potentially due to cheap pricing or the availability of automation features [39], [45]. Worryingly, we see reports of new gTLDs used for malicious campaigns where cybercriminals register domains across many new gTLDs to make their attacks more efficient [40]. However, we are not aware of any study that stratifies new gTLDs into multiple generations and compares them longitudinally, using large-scale data spanning more than four years.

To address this gap, we present the first comparative study of maliciousness in four generations of gTLDs based on when they were released: *legacy* (before 2012), *first-wave* (2012–2015), *second-wave* (2016–2020), and *third-wave* (2021–2025), with addition of ccTLDs (all time). We measure malicious and benign use of TLDs using data from Palo Alto Networks (*PANW*), one of the largest global cybersecurity vendors, which includes categorization of domain names, zone files, a passive DNS dataset, and active DNS data. We obtain these datasets over a four-year time period (June 2021–September 2025) and use them to design and report three reputation metrics across various gTLD generations—the malicious ratio, malicious-to-benign ratio, and non-benign ratio. Prior work has either focused mainly on benign and speculative uses of new gTLDs [17] or on malicious uses [38], [39], [45]. However, none of the studies analyzed reputation metrics to understand malicious and benign domains together. Our study is the first longitudinal, multi-generational study to examine both the malicious and benign use of domains in depth. Moreover, we characterize the effect of pricing, age, and defensive registrations in maliciousness across TLD generations and categories. Overall, our work answers five primary research questions:

- **RQ1.** Are newer gTLD generations more malicious and less benign than legacy gTLDs?
- **RQ2.** How does utilization of gTLDs evolve over time?

Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) 2026
27 February 2026, San Diego, CA, USA
ISBN 978-1-970672-06-0
https://dx.doi.org/10.14722/madweb.2026.23030
www.ndss-symposium.org

- **RQ3.** What role do factors such as age, pricing, popularity, and category play on TLD reputation?
- **RQ4.** Are defensive registrations common in new gTLD generations?
- **RQ5.** Who operates malicious gTLDs?

Updating and confirming prior work [39], [17], our analysis shows both higher malicious rates and lesser benign use of newer generations of gTLDs compared to legacy gTLDs. For instance, we observe that first, second, and third wave gTLDs are 2.3, 2.7, and 4.2 times more malicious than legacy gTLDs respectively, showing that newer gTLDs are increasingly used for malicious purposes. In addition, we show that these newer gTLDs have 1.5–2.8 times less benign content than legacy TLDs, resulting in malicious-to-benign ratios that are 3.1–9.2 times worse for newer gTLDs compared to legacy ones. We further show that defensive registrations have also significantly reduced in newer gTLDs.

Worryingly, we also observe that the utilization of new TLDs compared to legacy TLDs becomes increasingly *worse with time* within each of the newer generations. The malicious-to-benign ratio of the first, second, and third-wave TLDs increased 13.0, 10.6 and 50.4 times, respectively, during our four year study period. Our analysis reveals that new gTLDs are increasingly used for Registered Domain Generation Algorithm (RDGA) campaigns [10], [60], and do not see a proportional increase in registration of benign domains. Our findings warrant further research and potentially a policy-level intervention to remediate the current trajectory of new gTLDs.

In addition to generational and longitudinal comparisons, we also analyze the effect of factors such as age, popularity, pricing, and TLD category on maliciousness across generations of TLDs. We find that TLD categories and generations with lower average prices generally observe higher malicious ratios, reinforcing findings from prior work [45]. The extremely low pricing by many new gTLD operators has led to a surge in malicious registrations—hundreds of thousands of domains registered for single campaigns, as discussed in Section V—driving the apparent popularity of some of these TLDs. We also find that defensive registrations only form a small portion of domain registrations in newer TLDs, and hence do not contribute towards the high rates of low-content domains nor the resulting high non-benign ratios observed in newer gTLDs. Finally, we highlight the TLDs and sponsoring organizations that have disproportionately high malicious ratios: *Elegant Leader Limited* and *Shortdot SA* are the worst offenders, operating the `.xin` and `.bond` TLDs.

Our findings offer a large-scale, multi-generational perspective on malicious activity across TLDs. Rather than examining individual malicious domains, we focus on aggregate patterns of abuse to reveal broader trends in TLD utilization. As a contribution of our work, we maintain a webpage dedicated to TLD reputation metrics and will continue to update it periodically with data from *PANW* [1]. We hope that our work inspires action by ICANN and stakeholders, and further research on the utilization and evolution of TLD generations.

[1]https://github.com/PaloAltoNetworks/tld-metrics

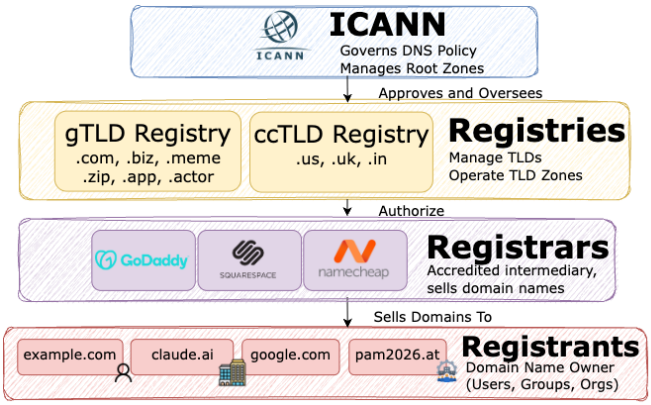

Fig. 1: **Relationships between different stakeholders in the DNS Ecosystem.**

II. Background

The Domain Name System is a foundational component of the Internet which maps human-readable domain names to IP addresses and is managed by a hierarchical ecosystem of stakeholders depicted in Figure 1. At the apex is ICANN [23] governing global DNS policy, managing the root zone, and overseeing the organizations that comprise the domain name market. Operating under contract with ICANN, *registries* are organizations that manage the primary database and technical infrastructure for specific TLDs. Next, registrars are accredited intermediary retailers selling domains to the public. Finally, end users purchasing domain names are called registrants.

TLDs represent the second highest level in the DNS hierarchy, right under the root domain (represented by a single dot). TLDs can be classified into six major types: generic, country-code, sponsored, reserved, test, and infrastructure TLDs. GTLDs are the most common followed by ccTLDs, two-letter domains designated for specific countries and territories. Sponsored TLDs represent specific professional, technical or cultural communities; their registration may be restricted (*e.g.,* `.gov`, `.mil`) or open to the public (*e.g.,* `.app`). There are several reserved TLDs that are permanently unavailable for registration (RFC 2606 [12]). Lastly, there is a single infrastructure TLD (`.arpa`) used for network management and 12 test TLDs that are not installed in the root zone. In this paper, we consider all TLDs as gTLDs if they are not ccTLDs or restricted.

Initially in 1985, the DNS ecosystem had only a limited number of gTLDs such as `.com`, `.org`, and `.net`. Although some TLDs were added incrementally over subsequent decades, the namespace expanded dramatically with ICANN's launch of the New gTLD Program in 2012, increasing the number of gTLDs from 22 to over 1,200 [21]. While this growth was intended to foster competition, it also introduced significant security challenges by creating new vectors for malicious activity and imposing defensive registration burdens

on brand owners [11], [1], [16], [9]. In this paper, our aim is to systematically analyze the security implications of these expansions, especially since the program is set to expand further in 2026.

III. Related Work

A. New TLD measurements

Since the inception of the new gTLD program in 2012, several studies have analyzed these TLDs. Halvorson *et al.* examined `.xxx` in 2014 using ICANN records, zone files, WHOIS data, and web crawling, finding that registrations are primarily speculative or defensive [18]. In a subsequent 2015 study, Halvorson *et al.* examined 502 newly released gTLDs and reported that their registration activity was predominantly fueled by speculative investments and precautionary purchases, rather than organic demand [17]. In 2011, just before the launch of the 2012 new gTLD program and ten years after the `.biz` TLD was introduced, Halvorson *et al.* also conducted a study [19] to compare the usage of `.biz` to the more established `.com` TLD [19]. They found that approximately 20% of the domains in both TLDs were parked pages and 10–25% of all `.biz` registrations were defensive registrations. Zirngibl *et al.* studied domain parking and found that between 20–30% of all registered domains in most gTLDs are parked [63]. We additionally find that most newer TLDs have a significantly higher rate of parked pages. Pfisterer *et al.* performed a comparative analysis of `.org` (a legacy TLD) and `.dev`, a first-wave TLD, finding that the older `.org` is more stable in terms of domain registrations and activity [49].

B. Malicious Activity in New TLDs

ICANN's new gTLD program has not only broadened the namespace, but also opened avenues for abuse. Open registration rules and inexpensive domains provide an attractive environment for adversaries who leverage them for phishing campaigns, malware distribution, and spam. Pouryousef *et al.* demonstrated how the proliferation of gTLDs creates opportunities for both typosquatting and exploitative pricing schemes that target trademark holders [51] which was further explored in other work [56], [2]. Moura *et al.* recently conducted a multi-year investigation of phishing within three country-code TLDs (`.nl`, `.ie`, `.be`), finding that the stricter registration requirements of ccTLDs shift adversaries from new registrations to compromising existing domains [43]. Nosyk *et al.* explored economics factors and showed that domain pricing directly influences abuse, observing that a single-dollar decrease in registration cost correlates with almost 50% rise in malicious registrations [45].

Korczynski *et al.* examined differences in abuse between legacy and new gTLDs over the 2014–2016 period by analyzing security vendor feeds [39]. Their results indicated a rising concentration of spam activity within the new gTLDs, alongside evidence that adversaries were increasingly shifting away from legacy namespaces. Seven years after the last systematic study of new gTLDs, we provide a much needed and timely update, given the impending 2026 gTLD application round and the severe deterioration of new gTLD reputation. Building on previous work, our paper broadens the scope by evaluating three distinct generations of gTLDs and contrasting them to legacy gTLDs and ccTLDs. Korczynski *et al.* also developed reputation-oriented metrics for TLDs, linking abuse to characteristics such as registration cost, namespace popularity and size, and the adoption of security mechanisms like DNSSEC [38]. We improve upon their study by introducing two new TLD reputation metrics malicious-to-benign and low-content ratios. These metrics are important to understand the relationship between benign and malicious utilization of new gTLDs. Importantly, our study provides a longitudinal view of various generations of TLDs to understand how the evolution of their malicious and benign activity compares to legacy and ccTLDs over time.

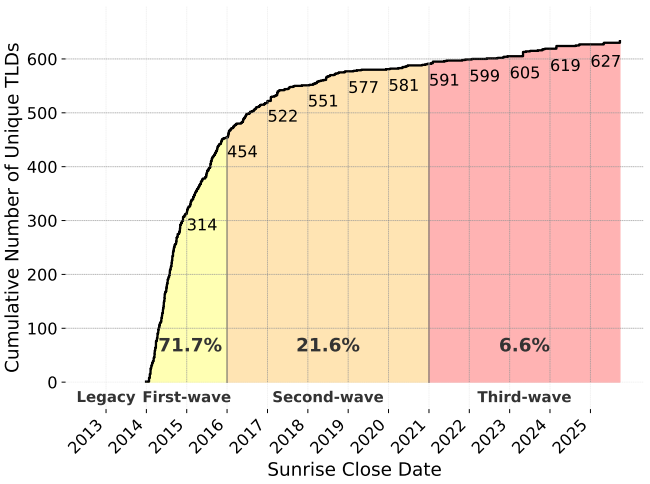

Fig. 2: **The various TLD generations we consider in this paper, and the cumulative distribution of gTLDs across generations.**

IV. Methodology

A. TLD Generations

We define four generations of gTLDs based on the cumulative distribution over their sunrise close date as shown in Figure 2. We use the sunrise close date as it marks the start of gTLD availability to the public [24]. *Legacy TLDs* are the oldest, introduced before ICANN's new 2012 gTLD program [21]. Legacy TLDs are not included in the cumulative distribution in Figure 2, and we consider all 22 legacy gTLDs, including restricted ones as they are insignificant in size and do not affect our results. Legacy TLDs are listed in Table X in the Appendix A.

We next categorize TLDs from ICANN's new gTLD program into three generations. The "First-wave" generation from 2013-2016 captures new gTLDs introduced at a fast pace following ICANN's new gTLD program. In Figure 2, we can observe that the introduction of gTLD starts to slow down by 2016, which we consider to be the end of the first-wave generation. By then, 454 (>71.7%) of all new gTLDs became

available for registration. We split the remaining decade into two equal time periods. The "Second-wave" generation consists of 137 gTLDs, and users could access more than 93% of all new gTLDs by the end of 2021. Finally, the "Third-wave" generation consists of the newest 42 gTLDs that were slowly introduced to the public, until the end of 2025. Furthermore, we also use 305 ccTLDs (detailed in Appendix A) as an additional reference point.

B. TLD Categorization

We split new gTLDs into three popularity categories based on the number of domains registered: *high* (over 1 million domains), *mid* (between 100,000–1,000,000 domains), and *low* (less than 100,000 domains). We analyze different popularity categories to see if new gTLD popularity is correlated with malicious activity *i.e.,* are new gTLDs popular partially because cybercriminals prefer them for their malicious campaigns?

We also split gTLDs into different TLD category types (*e.g.,* commerce) to understand whether certain TLD categories are utilized more for malicious purposes. For categorization, we use tld-list's 24 categories that apply to new gTLDs, including commerce, adult, and professional [57].

C. Data Sources

Studying malicious activity across gTLD generations is inherently challenging due to the lack of a single comprehensive data sources on domain registrations, activity, intent, and maliciousness. We adopt data from multiple large-scale data sources that provide a comprehensive, longitudinal view on malicious and benign domain registrations.

Zone files. To collect domains registered in gTLDs, we extract zone files published by ICANN's Centralized Zone Data Service (CZDS) [20] four times daily. Although previous work has relied primarily on official zone file data [17], [47], they miss short-lived domains (of which malicious domains may especially be a part [6]) due to the low frequency of updates. Moreover, domains that are defensively registered may not be entered into zone files [18].

WHOIS. We use WHOIS data to examine the prevalence of *defensive registrations*, a security defense strategy where organizations preemptively secure domains matching their trademarks and variants to prevent misuse. Prior work by Benjamin et al. in 2024 curated a list of trusted defensive registrars [4], which we leverage in our work to identify defensive registrations. We obtain the latest registrar information for each domain using the WhoisXML API [61].

Passive DNS. Passive DNS refers to logs of DNS requests and responses collected from multiple vantage points. Our passive DNS dataset contains anonymized data from Farsight's DNSDB [15] and from customers of *PANW*. As of May 2025, this continuously growing dataset contained 692 billion unique resource records, totaling 212 Terabytes of large-scale and real-world DNS resolution behavior. We consider a ccTLD root domain as registered if we observe a DNS response for it in our pDNS data in the past year. We use the pDNS data to estimate whether a domain is registered for ccTLDs, only when zone files are not available.

Active DNS. For every categorized ccTLD root domain not covered by zone files or pDNS, we actively query the A and NS records twice to determine if they are registered.

D. Domain Categorization

To ensure meaningful categorization of all domains, we use the comprehensive domain content categorization data from *PANW*. They crawl and re-categorize newly-registered domains one, three, seven, and 32 days after their registration. We cross-check our ≈363 million registered domains with the categorization of domains available from *PANW*. We note that we remove domains pointing to sinkhole servers as they can bias the reputation metrics for smaller TLDs. In total, we consider 343,148,606 registered domains (for the 2025.09.06 snapshot) that we are able to classify using *PANW*'s data.

Each domain is assigned one of 66 possible categories based on the content hosted on them, using *PANW*'s proprietary categorization tool used in production. For this paper, we classify the 66 categories into three groups: malicious, benign, and low-content, whose distribution is shown in Table I. The malicious category comprises malware hosts (e.g., viruses, trojans), command and control infrastructure, phishing sites, and grayware. We define grayware as non-overtly malicious content, including adware, scams, clickjacking, and potentially unwanted programs (PUPs). Benign domains include categories where we identify benign content such as social networking sites, personal blogs, gambling, and adult websites. A full list of *PANW*'s categories can be found at [46]. Low-content refers to domains where *PANW*'s categorization models are unable to identify any benign content. This includes three sub-categories: insufficient content, parked, and unknown. The insufficient content category refers to cases where no or very little content can be loaded during the scraping of a domain. Domains are categorized as unknown when there is no web service running on the server at the time of scraping, and it is by far the smallest component of the low-content group, only constituting 4.2% of the entire group. Although the low-content group does not mean that a domain is benign or malicious, looking at the entire population of a TLD paints a comprehensive picture of how that TLD is being utilized.

Snapshots. We use eight available snapshots of categorized and registered domains leveraging the same methodology and spanning over four years. The dates of the snapshots are: 2021.06.15, 2022.01.20, 2022.06.16, 2023.02.16, 2023.04.06, 2024.10.17, 2025.02.27, and 2025.09.06. For comparing reputation metrics, we rely on the 2025.09.06 snapshot, but we always contrast it with the earliest 2021.06.15 snapshot. For longitudinal analysis, we use all eight snapshots.

E. TLD Reputation Metrics

We calculate three reputation metrics to better understand the maliciousness and usage of gTLDs. First, we calculate the malicious ratio similar to the metrics in previous work [38]:

$$\text{Malicious Ratio} = \frac{Malicious\ count}{Total\ count} \tag{1}$$

TABLE I: **Fraction of Benign, Malicious, and Low-Content domains** for different TLD generations and ccTLDs on 2025.09.06 (and 2021.06.15 in parentheses).

	Legacy (TLD $n = 22$)	First-wave (TLD $n = 454$)	Second-wave (TLD $n = 137$)	Third-wave (TLD $n = 42$)	ccTLD (TLD $n = 305$)
Benign	0.738 (0.866)	0.493 (0.870)	0.498 (0.868)	0.266 (0.630)	0.809 (0.932)
Malicious	0.038 (0.014)	0.116 (0.016)	0.145 (0.024)	0.238 (0.011)	0.028 (0.012)
Low-Content	0.224 (0.119)	0.391 (0.114)	0.357 (0.108)	0.496 (0.359)	0.162 (0.056)
– Parked	0.171 (0.084)	0.311 (0.062)	0.253 (0.069)	0.403 (0.307)	0.109 (0.015)
– Insufficient content	0.049 (0.032)	0.061 (0.039)	0.078 (0.028)	0.038 (0.019)	0.041 (0.027)
– Unknown	0.004 (0.004)	0.019 (0.014)	0.026 (0.010)	0.054 (0.033)	0.011 (0.014)

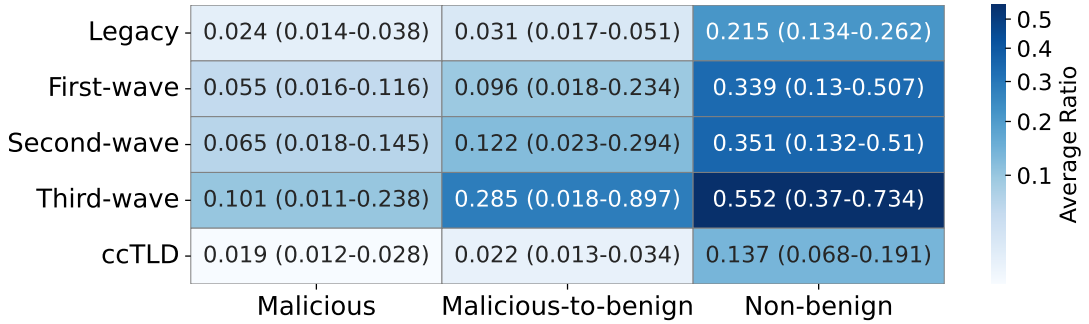

Fig. 3: **Average reputation metrics** for TLD generations and ccTLDs using all eight snapshots between 2021.06.15–2025.09.06. Minimum and maximum in brackets.

In addition, we also calculate two new metrics to identify whether new gTLDs are being used for benign content:

$$\text{Malicious-to-Benign Ratio} = \frac{Malicious\ count}{Benign\ count} \quad (2)$$

$$\text{Non-Benign Ratio} = \frac{Malicious + Low\text{-}content}{Total\ count} \quad (3)$$

We introduce the malicious-to-benign and non-benign ratios because they are especially well suited for measuring adoption of new gTLDs for benign purposes. As the overall registration volume is low in new gTLDs, it is crucial to understand whether registration is driven by a legitimate need for the domain namespace.

V. Analysis

In this paper, we study 960 TLDs across 4 gTLD generations and ccTLDs that comprise 343,148,606 registered and classified domains (per snapshot details are in Table XI in the Appendix). Using our multi-source comprehensive approach, we analyze each of our four research questions below:

A. RQ1: Generational TLD Utilization

Newer gTLD generations exhibit substantially higher malicious and non-benign ratios than legacy TLDs, indicating that abuse is increasingly concentrated within the new gTLD namespace. At the same time, the prevalence of low-content domains suggests limited benign adoption.

The fraction of malicious, benign, and low-content domains across gTLD generations and ccTLDs is shown in Table I. *Overall, we find that newer TLD generations tend to have fewer benign domains*, with the only exception being first- and second-wave gTLDs, which exhibit comparable benign ratios. At the same time, *newer generations consistently exhibit significantly higher malicious ratios than older ones*, indicating a higher density of malicious registrations within the new gTLD namespace. Importantly, this elevated maliciousness is accompanied by a pronounced increase in non-benign registrations driven largely by low-content domains. Compared to legacy TLDs, newer generations contain between 1.6 and 2.2 times more low-content domains, the majority of which are parked. This pattern suggests that many new gTLDs are not only disproportionately abused, but are also sparsely utilized for legitimate content. As a point of contrast, ccTLDs maintain both a higher proportion of benign domains and a lower proportion of malicious domains than legacy gTLDs.

Averaging across all snapshots, Figure 3 further shows that reputation metrics systematically deteriorate across successive gTLD generations. Relative to legacy gTLDs, first-wave gTLDs exhibit 2.3×, 3.1×, and 2.6× increases in malicious, malicious-to-benign, and non-benign ratios, respectively. These trends intensify in later generations: third-wave gTLDs display the poorest overall reputation, with the same metrics increasing by 4.2×, 9.2×, and 2.6× compared to legacy TLDs. Together, these results point to a pronounced concentration of abuse in the newest gTLDs, coupled with limited benign adoption, a theme we explore in greater depth in the following sections by examining contributing factors.

B. RQ2: Evolution of TLD Reputation over Time

> Across more than four years of observation, reputation metrics for all gTLD generations consistently worsen over time, with newer gTLDs deteriorating substantially faster than legacy TLDs.

To understand how TLD reputation has evolved over time, we compare different generations between 2021.06.15 and 2025.09.06, spanning more than four years. Figures 4a, 4b, and 4c compare the longitudinal evolution of our reputation metrics for all TLD generations and ccTLDs. For all gTLDs, *we see an overall worsening trend for all reputation metrics*. Only ccTLDs appeared to improve slightly between 2023.04.06 and 2024.10.17, but even their non-benign ratio steadily increases. Worryingly, we find that *all new gTLD generations across all reputation metrics worsen faster than legacy gTLDs.*

In Figures 4a and 4b we see a very sharp increase in the malicious and malicious-to-benign ratios for third-wave TLDs between 2025.02.27 and 2025.09.06. The sharp increase in malicious reputation is a symptom of how the newest TLDs cater to malicious registrations and simultaneously are not used for benign content. In particular, this jump in malicious reputation was caused by hundreds of thousands of Domain Generation Algorithm [13] `.sbs` domain registrations, called Registered DGAs (RDGAs). Uniformly randomly sampled examples of malicious `.sbs` domains can be found in Table XII of Appendix A, where we can see that all are related to RDGAs. During the same period, the second-wave generation reputation metrics stay constant or slightly decrease. At first glance, the lack of increase in malicious activity seems like a positive sign. However, taking a deeper look, we find that this improvement is driven by a single TLD, `.bond`. The number of malicious domains in `.bond` dropped from 1.27 million malicious registrations on 2025.02.27 to to 1.09 million by 2025.09.06. While the reputation metrics steadily increased for all other second-wave TLDs in this period, showing that the small improvement is dominated by a single TLD [10].

Figure 5 shows the overall growth rate of reputation metrics from 2021.06.15 to 2025.09.06. We find that the malicious ratio for the first, second, and third wave new gTLDs increased 2.8, 2.3, and 8.1 times faster compared to legacy TLDs. This comparison becomes significantly worse for the malicious-to-benign ratio, increasing 4.2, 3.4 and 16.3 times faster, respectively. Our findings suggest that the reputation of newer generations of gTLDs from ICANN's program is deteriorating, projecting a grim outlook.

In summary, Figure 4 show that reputation metrics generally only increase (worsen) across successive snapshots for all TLDs, barring a few exceptions. Moreover, we observe that new generations of TLDs are more malicious and have less benign content compared to legacy TLDs. The larger increase in maliciousness and non-benign ratios for newer generations of gTLDs highlights an important question: *Is continued investment in the new gTLD program producing more opportunities for attackers rather than legitimate registration opportunities?* Our findings warrant that we, as a community, need to take action to protect new generations of TLDs from becoming a cesspool of malicious activity. We discuss this further in Section VI.

TABLE II: **Reputation Metrics for TLD Categories.** Darker red indicates worse reputation (column-wise normalized).

Category	Avg Price	Ratio M	Ratio M-B	Ratio N-B
Organizations	2.09	0.38	1.40	0.73
Travel	3.45	0.19	0.81	0.76
Products & Industry	4.41	0.15	0.46	0.67
Money & Finance	4.81	0.17	0.42	0.61
Miscellaneous	1.56	0.16	0.37	0.57
Social & Lifestyle	3.90	0.10	0.19	0.48
Adult	7.88	0.09	0.15	0.44
Internet	1.75	0.08	0.15	0.46
Commerce	1.16	0.07	0.13	0.43
Community	4.36	0.06	0.11	0.44
Media, Art, & Music	5.67	0.07	0.11	0.39
Medical & Health	10.30	0.05	0.11	0.53
Services	3.59	0.06	0.11	0.46
Government	7.75	0.06	0.11	0.44
Food & Drink	6.97	0.06	0.11	0.44
Technology	4.41	0.05	0.10	0.46
Education	8.94	0.06	0.09	0.40
Regional & Cultural	12.01	0.06	0.09	0.40
Sports	12.02	0.05	0.08	0.41
Business	7.67	0.04	0.07	0.40
Religion	5.96	0.05	0.07	0.34
Professional	9.66	0.04	0.07	0.36
Real Estate	8.63	0.04	0.07	0.37
Cities	16.60	0.03	0.05	0.37
(no category)	26.19	0.01	0.02	0.44

TABLE III: **Weighted average of cheapest domain registration price per TLD generation.**

Generation	Pricing*
Legacy	7.76
First-wave	2.91
Second-wave	2.55
Third-wave	1.79
ccTLD	10.42

C. RQ3: Role of TLD Category, Pricing, Popularity, and Age

> We find that malicious activity is systematically higher in cheaper, more popular, and thematically attractive gTLDs. These effects are amplified in newer TLD generations.

TLD Categories. We first evaluate the effect of TLD category types on malicious activity and TLD utilization. As shown in Table II, there is a significant difference between the categories of TLDs in our three reputation metrics. The worst TLD in terms of malicious and malicious-to-benign ratio, *Organizations*, has a 28X higher average malicious-to-benign ratio compared to the best category, *Cities*. We observe that malicious actors prefer to utilize malicious domains in popular categories such as organizations, travel, and finance, which can be part of targeted phishing attacks.

TLD Pricing. Using pricing data from *tld-list.com* [57], we study whether domain registration prices are related to

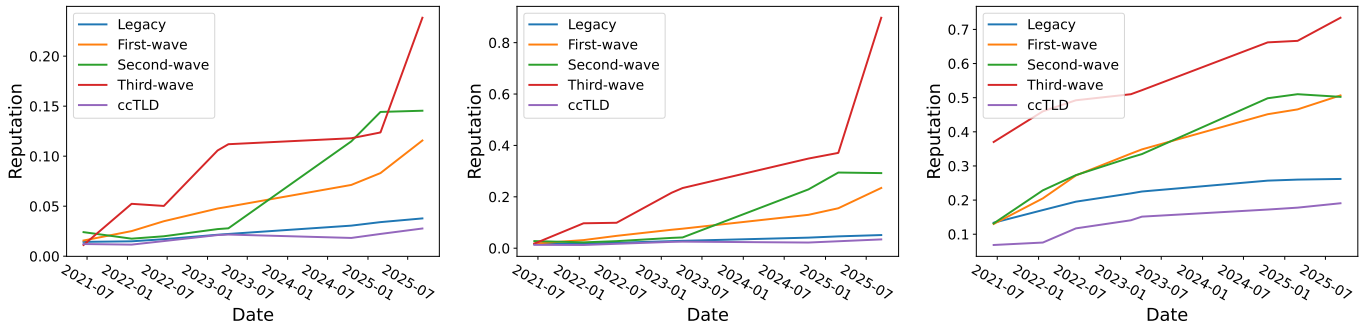

(a) **Malicious ratio** (b) **Malicious-to-benign ratio** (c) **Non-benign ratio**

Fig. 4: **Evolution of TLD reputation metrics over time (2021.06.15–2025.09.06) for different TLD generations.**

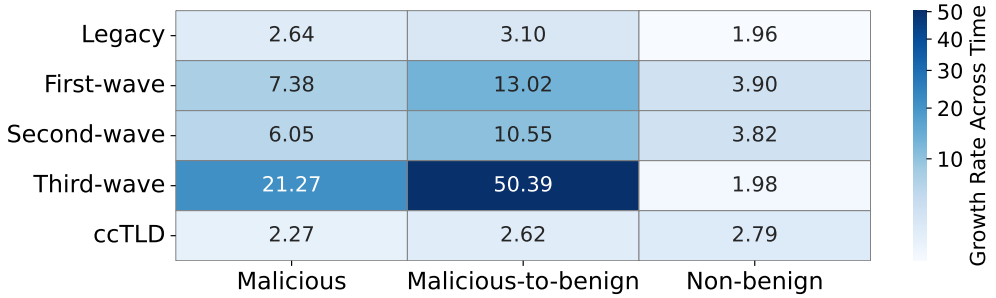

Fig. 5: **Growth rate of TLD reputation metrics across time**. The growth is calculated by dividing the value from 2025-09-06 by the one from 2021-06-15 *e.g.,* a growth rate of 6.05x means the malicious ratio increased from 0.024 to 0.145.

TABLE IV: **The Pearson correlation coefficient between cheapest domain price and per TLD reputation scores.**

Generation	Ratio		
	Malicious	Malicious-to-benign	Non-benign
First-wave	-0.22 (p<0.01)	-0.13 (p=0.07)	-0.39 (p<0.01)
Second-wave	-0.26 (p=0.09)	-0.08 (p=0.62)	-0.25 (p=0.10)
Third-wave	-0.27 (p=0.48)	-0.21 (P=0.58)	0.22 (p=0.56)
All	-0.24 (p<0.01)	-0.06 (p=0.32)	-0.33 (p<0.01)

TABLE V: **The Pearson correlation coefficient between TLD price and reputation per different TLD categories.**

Generation	Ratio		
	Malicious	Malicious-to-benign	Non-benign
First Wave	-0.45 (p=0.03)	-0.40 (p=0.05)	-0.16 (p=0.45)
Second Wave	-0.37 (p=0.08)	-0.32 (p=0.14)	-0.17 (p=0.44)
Third Wave	-0.39 (p=0.14)	-0.34 (p=0.20)	-0.18 (p=0.50)
All	-0.49 (p=0.01)	-0.40 (p=0.05)	-0.37 (p=0.07)

malicious domain registrations. In Table III, we list TLD generations and their corresponding pricing calculated as the weighted average of the cheapest price for domain registration per TLD weighted by the number of domains in the TLD. The pricing of a generation is generally aligned with its reputation: *the lower the price of a TLD generation, the worse its reputation*, while also failing to attract benign users, reinforcing findings from prior work [45].

In Table IV, we show a negative Pearson correlation [48] between *individual gTLD registration prices* and their reputation, indicating that more expensive TLDs have a better (*i.e.,* less malicious) reputation. However, this *per-TLD correlation* is weaker than expected from previous work [45] and not statistically significant for the malicious-to-benign ratio. We hypothesize this is because the abundance of cheap gTLDs leads cybercriminals to prioritize other factors, such as lax identity verification and the ease of bulk registration [39].

Considering the average domain registration price per TLD category, a general trend appears: *categories with cheaper domains have the worst reputation*. Table V shows a strong negative correlation between average price versus reputation ratios for TLD categories. Although the results are not statistically significant for many generations, if we look at the results for all new gTLDs released in the new gTLD program, then there is a substantial correlation between domain price and malicious ratio (where p=0.01). While price is an important factor in domain purchase preferences of cybercriminals, it is not the only factor, for example, some criminal groups might target specific categories for their endeavors (e.g. *Adult*).

TLD Age. A possible explanation for the bad reputation of newer generations of gTLDs is that they did not have time to mature, thus they have less domains with content and more malicious domains. However, Figure 5 already contradicts this theory, as it shows that over time reputation of all gTLDs

TABLE VI: **The Pearson correlation coefficient between TLD age in days and reputation.**

Generation	Ratio		
	Malicious	Malicious-to-benign	Non-benign
First-wave	0.37 (p<0.01)	0.20 (p<0.01)	0.51 (p<0.01)
Second-wave	0.27 (p<0.01)	0.07 (p=0.09)	0.31 (p<0.01)
Third-wave	0.22 (p=0.07)	-0.06 (p=0.64)	-0.16 (p=0.18)
All	0.19 (p<0.01)	0.01 (p=0.64)	0.12 (p<0.01)

TABLE VII: **The Pearson correlation coefficient between TLD popularity (log10 domain count) and reputation.**

Generation	Ratio		
	Malicious	Malicious-to-benign	Non-benign
First-wave	0.16 (p<0.01)	0.12 (p<0.01)	0.37 (p<0.01)
Second-wave	0.39 (p<0.01)	0.19 (p=0.04)	0.51 (p<0.01)
Third-wave	0.26 (p=0.14)	-0.26 (p=0.14)	-0.16 (p=0.36)
All	0.22 (p<0.01)	0.11 (p<0.01)	0.35 (p<0.01)

generations worsens for all metrics. To better understand how age and reputation are connected, we calculate Pearson's correlation between the age and reputation metrics for each generation of gTLD in Table VI. We find that the correlation is nearly always positive for first- and second-wave gTLDs, meaning that TLDs in these generations are getting worse with age. We observe the same finding for the malicious ratio in third-wave gTLDs, but for the other reputation metrics the results are statistically insignificant and negative probably due to the fact that some new TLDs might have very bad reputation very early due to overtly cheap pricing.

TLD Popularity. Figure 6 shows the reputation metrics for different TLD popularity levels. We find that higher popularity of gTLDS corresponds to worse reputation for all metrics. From low to mid popularity, the malicious-to-benign ratio increases by 2.1×, and from mid to high popularity it increases by an additional 1.5×. The connection between popularity and bad reputation signals that at least partially the popularity of new gTLDs stems from malicious registrations.

When we break down the analysis by generation, we observe a similar pattern for first-wave gTLDs: higher popularity is consistently associated with worse reputation metrics. In contrast, for second- and third-wave gTLDs, mid-popularity TLDs exhibit the worst reputation for several metrics, often exceeding even the most popular TLDs. We hypothesize that this reflects the relative immaturity of newer generations, which have not yet had sufficient time to develop many highly popular gTLDs, while a subset of mid-popularity TLDs is disproportionately dominated by malicious registrations, as we discuss further in Sections V-B and V-E.

Table VII reinforces our previous findings as we can see that first- and second-wave generations have a positive statistically significant correlation between their popularity (using log10 unique domain count) and malicious reputations. The only exception is third-wave TLDs, where the results are not statistically significant due to the smaller number of registrations.

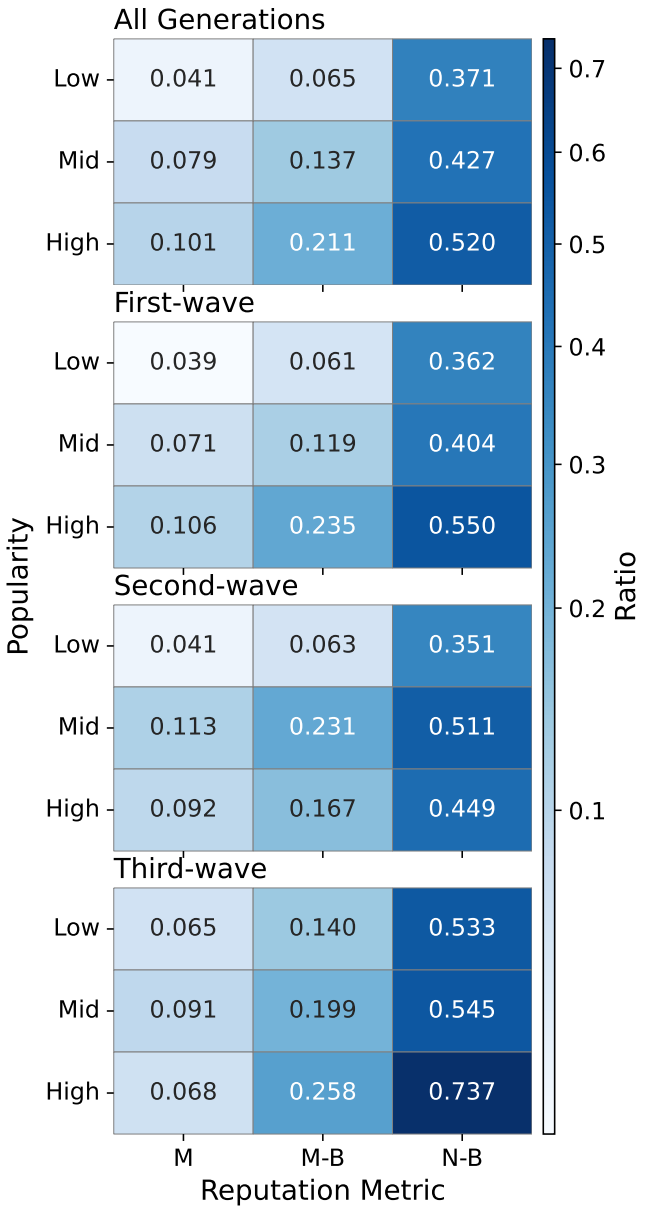

Fig. 6: **Reputation metrics for different popularity levels and gTLD generations on 2025.09.06.**

Moreover, some small or medium gTLDs in the third-wave might be entirely dominated by malicious registrations.

D. RQ4: Investigating Defensive Registrations

> Defensive registrations account for only a negligible fraction of domains in new gTLD generations and decline sharply relative to legacy TLDs.

"Defensive registration" refers to the practice of securing domains to protect intellectual property or trademarks from misuse, such as cybersquatting [4], [55], [22]. In this section, we analyze whether defensive registrations contribute to the high prevalence of low-content domains in new gTLDs.

Typically, defensively registered domains are configured to redirect to the entity's primary service or remain parked or as Registered Non-Resolving (RNX) domains [17], [18], [4]. In the latter case, our categorization would classify them as unknown or insufficient content. To that end, we must determine what proportion of low-content domains are defensive in nature, thereby verifying if defensive practices explain the high ratio of low-content domains in newer gTLDs. While early characterization studies of new gTLDs found that defensive registrations were substantial [17], we investigate whether this trend persists today.

We adopt the methodology introduced by Benjamin et al. [4] in 2024, using their curated list of trusted registrars to identify and quantify defensive registrations in this section. Among the 22 legacy gTLDs, we see that 1.033% of all domains are registered with these 11 defensive registrars. This number includes the primary service website as well as the defensively registered domain names under other gTLDs.

However, this percentage falls to 0.275% when we consider domains under the 454 first-wave TLDs. This means that there are 3.75x fewer defensive registrations among the first-wave as compared to the legacy TLDs. For the 137 second-wave TLDs, we see that this number further reduces to 0.174% of all registered domains. However, we note that this category includes two particularly prolific TLDs (`.bond`, `.shop`) which have 8.85 and 3.15 million domains respectively, with only 7,786 domains from defensive registrars (0.06%). These two TLDs are heavily abused by threat actors conducting automated bulk registration of domains for cybercrime. Removing these two TLDs, the defensive registrations amount to 0.216% of all registrations in the remaining 452 TLDs.

Interestingly, among third-wave TLDs, we find that the percentage of defensive registrations increased to 0.27%. Investigating further, we find that a large portion of this is due to two specific gTLDs: `.zip` and `.mov`. In these two gTLDs alone, the percentage of defensive registrations account for 8.06% of the total domains. This effect is noticeable because `.zip` and `.mov` were highly criticized upon release and were heavily misused by attackers due to their names resembling popular file extensions [16]. Removing these two TLDs from consideration in third-wave TLDs, we see that the defensive registration is still very low at 0.177%.

We draw two conclusions from these trends. First, the prevalence of defensive registrations decreases in newer TLD waves. Second, because these registrations constitute a small fraction of the total domain volume, they do not contribute towards the high rates of low-content domains nor the resulting high non-benign ratios observed in newer gTLDs.

E. RQ5: Sponsoring Organizations

A small subset of sponsoring organizations operate gTLDs with extremely high malicious-to-benign ratios, driven by large-scale abuse campaigns and limited benign adoption.

We also study the sponsoring organizations of new gTLDs, which are effectively the owners and operators of these TLDs [57]. In Table VIII, we list the sponsoring organizations with the highest malicious-to-benign ratios. The worst offender is *Elegant Leader Limited* operating the `.xin` TLD. In addition to `.xin` having very few benign registrations, it was also abused for several malicious campaigns. One of the campaigns that heavily affected `.xin` was a SMS phishing campaign [54], [53] also reflected in the uniformly random sample of malicious domains listed in Table XII, Appendix A).

We find that *Shortdot SA* with its three TLDs has 2.5 more malicious domains than benign. We attribute *Shortdot SA*'s extremely high malicious-to-benign ratio to the lack of benign domains and the new trend of RDGA registrations [10], [60], where researchers found that an attacker registered more than 500,000 `.bond` domains [10] for a single campaign (samples in Table XII, Appendix A).

Interestingly, `.sbs` was restricted and originally sponsored by *Special Broadcasting Service Corporation*, which incorrectly appears as the third worst sponsoring organization. It turns out that on May 6, 2021 *Shortdot SA* announced that it had bought `.sbs` and it became generally available in June 2021, around the time of our first reputation snapshot [52]. We conclude that `.sbs` became preferred for malicious registrations after *Shortdot SA's* acquisition of the gTLD.

VI. Discussion

A. ICANN and Policy Recommendations

ICANN operates under a complex multi-stakeholder model [23], in which the economic incentives faced by different stakeholders are often misaligned with the goals of policies intended to reduce abusive domain registrations [55]. ICANN's DNS abuse mitigation program [25] is mainly limited to measurements of abuse and enforcing contractual obligations of registries and registrars. More precisely, the DNS abuse mitigation program consists of three components. First, ICANN runs a program called ICANN Domain Metrica [28] (which replaced DAAR: Domain Abuse Activity Reporting [26]) to measure domain abuse and registration activity in TLDs. Second, ICANN funds a research project called INFERMAL (Inferential Analysis of Maliciously Registered Domains) [45] to understand attackers' preferences regarding DNS abuse and SIFT (Special Interest Forums on Technology) [34] for stakeholders to engage in technical discussions. Third, ICANN enforces contractual obligations of registries and registrars as described in the Base Registry Agreement [33] and the Registrar Accreditation Agreement [32], requiring registries and registrars to take mitigation actions against well-evidenced DNS Abuse, engage in proactive monitoring, and participate in regularly scheduled audits.

Taking a deeper look at ICANN's current DNS abuse mitigation program in light of the empirical evidence from this paper, it appears unlikely that the program, in its current form, will substantially reduce malicious domain registrations. First, a large portion of ICANN's efforts focus on measurement studies which are often limited in their scope. For example, INFERMAL focuses only on 28 thousand phishing domains [45] when malicious actors register hundreds of thousands

TABLE VIII: **Sponsoring organization reputation.** M = malicious ratio; M-B = malicious-to-benign ratio; N-B = non-benign ratio; C-M-B = Change in malicious-to-benign ratio from 2021.06.15–2025.09.06. P-R = original rank on 2021.06.15. * denotes a shortened organization name, listed in the Appendix.

Sponsoring Org.	#TLDs	M	M-B	N-B	C-M-B	P-R	TLDs	Min Price
Elegant Leader*	1	0.765	5.138	0.851	5.106	8	xin	$2.29
Shortdot SA	3	0.524	2.752	0.810	2.710	5	bond, cyou, icu	$0.83–$1.42
Spec Broadcasting*	1	0.285	1.242	0.770	1.242	100+	sbs	$0.94
Dotcfd Registry*	1	0.208	0.747	0.722	0.741	53	cfd	$0.60
dot Date Limited	1	0.272	0.648	0.581	0.544	1	date	$2.99
DOTSTRATEGY*	1	0.258	0.535	0.518	0.518	18	buzz	$1.10
dot Loan Limited	1	0.265	0.445	0.405	0.374	2	loan	$2.99
First Registry*	1	0.195	0.436	0.554	0.403	7	win	$2.99
dotCOOL, Inc.	1	0.249	0.427	0.418	0.421	71	qpon	$1.48
dot Bid Limited	1	0.259	0.422	0.386	0.356	3	bid	$2.99
Jiangsu Bangning*	1	0.225	0.414	0.457	0.401	30	top	$1.17
XYZ.COM LLC	34	0.142	0.384	0.630	0.343	6	car, game, xyz, ...	$1.0–$1995
UNR Corp.	7	0.163	0.378	0.569	0.353	11	click, property, sexy, ...	$0.65–$2500

of domains for a single campaign [10]. Even ICANN's Domain Metrica project considers only 616 thousand malicious domains [28] while *PANW* blocks more than 16 million registered domains for its customers. Second, ICANN requires registries and registrars to take action on well-evidenced registration abuse when reported. While it is important that registries and registrars need to takedown or suspend malicious domains, such a reactive approach cannot hope to keep up with malicious actors who can register cheap domains easily in bulk and keep using different domains from new gTLDs.

In summary, ICANN's current policies appear to be ineffective in curbing abusive domain registrations. At the same time researchers have proposed a variety of domain registration policies that are expected to be significantly more effective, including incentivizing registries and registrars, increased or minimum mandatory registration pricing, bulk registration limitations, and stricter registrant identity verification [39], [55], [45]. Even the ICANN Governmental Advisory Committee's (GAC) report states that pricing and bulk registration plays a major role in enabling abuse [29].

Our research and evidence from previous work [39] show that new gTLDs are less utilized for benign and more for malicious activities compared to legacy TLDs. Taken together, our findings suggest that additional attention and resources are needed to better understand the challenges associated with newly accepted gTLDs, as well as to more carefully evaluate which gTLDs and sponsoring organizations are likely to provide sustained value to the broader Internet community. Most importantly, our findings and recent trends in malicious domain registrations indicate that the pace of abuse is unlikely to slow in the absence of a minimum domain registration price. Under the current low-cost pricing model, the financial impact of domain takedowns remains negligible for cybercriminals, allowing malicious operations to continue largely uninterrupted. While measures such as limiting bulk registrations and strengthening registrant identity verification could help mitigate abuse, enforcing strict identity verification at a global scale remains challenging, which in turn limits the effectiveness of bulk registration restrictions. With ICANN's new gTLD program expecting a new wave of gTLD submission in 2026, our findings indicate that it may be time for the community to reflect whether the opportunities offered by new gLTDs outweigh the risks posed by malicious actors.

B. Limitations

Benign categorization is largely based on web content (except known email and DNS providers), which might miss cases when domains are used for purposes other than web services. Furthermore, while *PANW*'s categorization is of high-quality, it can be imperfect just like any other public sources of domain categorization. These limitations are consistent across gTLD generations, minimizing their effect on our overall results.

Related work [39], [45], [55] investigated how factors other than pricing, such as bulk registration or identity verification, contribute to malicious registrations. We argue that pricing is the primary driver of worsening TLD reputation. While registration fees must *still* be paid, bulk registration is easily automated with AI-powered tools, and identity verification can often be fulfilled at low cost from lower-income countries.

C. Ethics Considerations

For our research, we rely on public zone files, DNS queries, anonymized passive DNS, and the classification of domains from *PANW*. The data used is not considered personally identifiable information (PII), and the dataset is processed and analyzed on secure internal servers, following all established best practices of *PANW*'s data use.

VII. Conclusion

We studied reputation metrics of four generations of TLDs to understand how malicious activity in new gTLDs compare to legacy gTLDs. We found that new gTLDs are more malicious and have less benign domains. Worryingly, new gTLD reputation metrics are deteriorating more rapidly over time compared to legacy gTLDs. Our findings warrant further research and action from ICANN to remediate malicious activities in new gTLD generations.

Acknowledgment

The authors are grateful to Zhanhao Chen for his input, and the reviewers for their feedback.

References

[1] New ZIP domains spark debate among cybersecurity experts. Bleeping Computer (2023), https://www.bleepingcomputer.com/news/security/new-zip-domains-spark-debate-among-cybersecurity-experts/
[2] Agten, P., Joosen, W., Piessens, F., Nikiforakis, N.: Seven Months' Worth of Mistakes: A Longitudinal Study of Typosquatting Abuse. In: Network and Distributed System Security (NDSS) (2015)
[3] Wikipedia for .an ccTLD. (2025), https://en.wikipedia.org/wiki/.an
[4] Benjamin, B.C., Bayer, J., Fernandez, S., Duda, A., Korczyński, M.: Shielding Brands: An In-depth Analysis of Defensive Domain Registration Practices against Cyber-squatting. In: IEEE Network Traffic Measurement and Analysis Conference (TMA) (2024)
[5] Wikipedia for .bh ccTLD. (2025), https://en.wikipedia.org/wiki/.bh
[6] Bilge, L., Kirda, E., Kruegel, C., Balduzzi, M.: Exposure: Finding malicious domains using passive dns analysis. In: Ndss. pp. 1–17 (2011)
[7] Wikipedia for .bl ccTLD. (2025), https://en.wikipedia.org/wiki/.bl
[8] Wikipedia for .bq ccTLD. (2025), https://en.wikipedia.org/wiki/.bq
[9] BrandSec: Phishing domains and new gtlds: A growing security concern, https://brandsec.com.au/phishing-domains-and-new-gtlds-a-growing-security-concern/
[10] Centripetal. Revolver Rabbit and the Rise of RDGAs. (2025), https://www.centripetal.ai/threat-research/revolver-rabbit-and-the-rise-of-rdgas
[11] Cloudflare: From .com to .beauty: The evolving threat landscape of unwanted email, https://blog.cloudflare.com/top-level-domains-email-phishing-threats/
[12] D. Eastlake, A.E.: RFC 2606: Reserved Top Level DNS Names (1999), https://www.rfc-editor.org/rfc/rfc2606.html
[13] Domain Generation Algorithm (2025), https://en.wikipedia.org/wiki/Domain_generation_algorithm
[14] Wikipedia for .eh ccTLD. (2025), https://en.wikipedia.org/wiki/.eh
[15] Farsight DNSDB 2.0, https://www.farsightsecurity.com/solutions/dnsdb/
[16] Goodin, D.: Google pushes .zip and .mov domains onto the Internet, and the Internet pushes back. Ars Technica (2023), https://arstechnica.com/information-technology/2023/05/critics-say-googles-new-zip-and-mov-domains-will-be-a-boon-to-scammers/
[17] Halvorson, T., Der, M.F., Foster, I., Savage, S., Saul, L.K., Voelker, G.M.: From .academy to .zone: An Analysis of the New TLD Land Rush. In: ACM Internet Measurement Conference (IMC) (2015)
[18] Halvorson, T., Levchenko, K., Savage, S., Voelker, G.M.: XXXtortion?: inferring registration intent in the .XXX TLD. In: International conference on World Wide Web (WWW) (2014)
[19] Halvorson, T., Szurdi, J., Maier, G., Felegyhazi, M., Kreibich, C., Weaver, N., Levchenko, K., Paxson, V.: The BIZ top-level domain: ten years later. In: Passive and Active Measurement (PAM). Springer (2012)
[20] CZDS: Centralized Zone Data Service, https://czds.icann.org/home
[21] ICANN: Internet Domain Name Expansion Now Underway, Online; accessed 2024-12-16 (2012), https://newgtlds.icann.org/en/about/program
[22] Defensive Registration - ICANNWiki (2021), https://icannwiki.org/Defensive_Registration
[23] Internet Corporation for Assigned Names and Numbers (ICANN), Online; accessed 2024-12-16 (2024), https://www.icann.org/
[24] TLD Startup Information (2024), https://newgtlds.icann.org/en/program-status/sunrise-claims-periods
[25] DNS Abuse Mitigation Program (2025), https://www.icann.org/resources/pages/dns-security-threat-mitigation-2025-11-21-en
[26] ICANN Domain Abuse Activity Reporting (2025), https://www.icann.org/octo-ssr/daar
[27] What Does ICANN Do? (2025), https://www.icann.org/resources/pages/what-2012-02-25-en
[28] ICANN Domain Metrica (2025), https://domainmetrica.icann.org/search
[29] ICANN 81, GAC, DNS Abuse Mitigation (2025), https://gac.icann.org/briefing-materials/public/ICANN81%20-%20GAC%20Briefing%20-%20Session%2015%20-%20DNS%20Abuse%20Mitigation.pdf
[30] ICANN (2025), https://www.icann.org/
[31] ICANN next round of new TLD program for 2026. (2025), https://newgtldprogram.icann.org/en/application-rounds/round2
[32] ICANN Registrar Accreditation Agreement (2025), https://www.icann.org/en/contracted-parties/accredited-registrars/registrar-accreditation-agreement
[33] ICANN Base Registry Agreement (2025), https://www.icann.org/en/registry-agreements/base-agreement
[34] ICANN Special Interest Forums on Technology (SIFT) (2025), https://www.icann.org/octo/sift-en
[35] ICANN new gTLD history. (2025), https://www.icann.org/resources/pages/newgtlds-history-2023-04-05-en
[36] Statistics about .il ccTLD. (2025), https://en.isoc.org.il/il-cctld/number-of-registered-domain-names-2023
[37] Wikipedia for .an ccTLD. (2025), https://en.wikipedia.org/wiki/.il
[38] Korczynski, M., Tajalizadehkhoob, S., Noroozian, A., Wullink, M., Hesselman, C., Van Eeten, M.: Reputation metrics design to improve intermediary incentives for security of TLDs. In: IEEE European Symposium on Security and Privacy (EuroS&P) (2017)
[39] Korczynski, M., Wullink, M., Tajalizadehkhoob, S., Moura, G.C., Noroozian, A., Bagley, D., Hesselman, C.: Cybercrime after the sunrise: A statistical analysis of dns abuse in new gtlds. In: Asia Conference on Computer and Communications Security (ASIACCS) (2018)
[40] Why phishers love new tlds like .shop, .top and .xyz (2024), https://krebsonsecurity.com/2024/12/why-phishers-love-new-tlds-like-shop-top-and-xyz/
[41] Wikipedia for .la ccTLD. (2025), https://en.wikipedia.org/wiki/.la
[42] Wikipedia for .mf ccTLD. (2025), https://en.wikipedia.org/wiki/.mf
[43] Moura, G.C., Daniels, T., Bosteels, M., Castro, S., Müller, M., Wabeke, T., van den Hout, T., Korczyński, M., Smaragdakis, G.: Characterizing and Mitigating Phishing Attacks at ccTLD Scale. In: ACM SIGSAC Conference on Computer and Communications Security (CCS) (2024)
[44] Wikipedia for .mr ccTLD. (2025), https://en.wikipedia.org/wiki/.mr
[45] Nosyk, Y., Korczynski, M., Maroofi, S., Bayer, J., Odgerel, Z., Duda, A., Tajalizadehkhoob, S., Gañán, C.: INFERMAL: Inferential analysis of maliciously registered domains (2024)
[46] PANW: URL Categories, https://docs.paloaltonetworks.com/advanced-url-filtering/administration/url-filtering-basics/url-categories
[47] Park, J., Choi, J., Nyang, D., Mohaisen, A.: Transparency in the new gTLD era: Evaluating the DNS centralized zone data service. Transactions on Network and Service Management (2019)
[48] Pearson correlation coefficient. (2025), https://en.wikipedia.org/wiki/Pearson_correlation_coefficient
[49] Pfisterer, F., Zirngibl, J., Sattler, P.: The evolution of top-level domains: A comparative study of. org and. dev. IITM Seminar (2023)
[50] Wikipedia for .pk ccTLD. (2025), https://en.wikipedia.org/wiki/.pk
[51] Pouryousef, S., Dar, M.D., Ahmad, S., Gill, P., Nithyanand, R.: Extortion or expansion? an investigation into the costs and consequences of ICANN's gTLD experiments. In: Passive and Active Measurement (PAM). Springer (2020)
[52] ICANN Wiki: .sbs. (2025), https://icannwiki.org/.sbs
[53] Resecurity's blog on smishing (2025), https://www.resecurity.com/blog/article/smishing-triad-is-now-targeting-toll-payment-services-in-a-massive-fraud-campaign-expansion
[54] Unit42 timely threat intelligence on smishing. (2025), https://github.com/PaloAltoNetworks/Unit42-timely-threat-intel/blob/main/2025-04-23-IOCs-for-smishing-activity-update.txt
[55] Szurdi, J., Christin, N.: Domain registration policy strategies and the fight against online crime. WEIS (2018)
[56] Szurdi, J., Kocso, B., Cseh, G., Spring, J., Felegyhazi, M., Kanich, C.: The Long "Taile" of Typosquatting Domain Names. In: USENIX Security Symposium (2014)
[57] TLD-List (2024), https://tld-list.com/
[58] Wikipedia for .tp ccTLD. (2025), https://en.wikipedia.org/wiki/.tp
[59] Wikipedia for .um ccTLD. (2025), https://en.wikipedia.org/wiki/.um
[60] Unit42. The Next Level: Typo DGAs Used in Malicious Redirection Chains. (2025), https://unit42.paloaltonetworks.com/typo-domain-generation-algorithms/
[61] The Who Behind Cyber Threat Intelligence (2025), https://www.whoisxmlapi.com/
[62] History of Generic TLDs. (2025), https://en.wikipedia.org/wiki/Generic_top-level_domain
[63] Zirngibl, J., Deusch, S., Sattler, P., Aulbach, J., Carle, G., Jonker, M.: Domain Parking: Largely Present, Rarely Considered! In: TMA (2022)

TABLE IX: TLD reputation per category and generation.

			Malicious Ratio				Malicious to Benign				Non Benign Ratio			
Category	**Count**	**Avg, Price**	**First Wave**	**Second Wave**	**Third Wave**	**All**	**First Wave**	**Second Wave**	**Third Wave**	**All**	**First Wave**	**Second Wave**	**Third Wave**	**All**
Organizations	38	2.09	0.02	0.43	0.28	0.38	0.03	1.57	1.23	1.40	0.19	0.73	0.77	0.73
Travel	22	3.45	0.04	0.22		0.19	0.06	1.34		0.81	0.34	0.84		0.76
Products & Industry	37	4.41	0.04	0.20	0.14	0.15	0.07	1.01	0.26	0.46	0.41	0.80	0.44	0.67
Money & Finance	42	4.81	0.17	0.02	0.13	0.17	0.43	0.02	0.57	0.42	0.61	0.22	0.77	0.61
Miscellaneous	50	1.56	0.16	0.02	0.06	0.16	0.37	0.02	0.16	0.37	0.57	0.13	0.65	0.57
Social & Lifestyle	81	3.90	0.11	0.09	0.11	0.10	0.21	0.16	0.23	0.19	0.49	0.47	0.53	0.48
Adult	7	7.88	0.03	0.10		0.09	0.03	0.19		0.15	0.21	0.49		0.44
Internet	48	1.75	0.08	0.07	0.04	0.08	0.15	0.11	0.08	0.15	0.46	0.39	0.46	0.46
Commerce	39	1.16	0.12	0.07	0.17	0.07	0.20	0.12	0.60	0.13	0.40	0.43	0.72	0.43
Community	34	4.36	0.06	0.10	0.07	0.06	0.10	0.20	0.16	0.11	0.43	0.50	0.59	0.44
Media, Art, & Music	38	5.67	0.07	0.07	0.06	0.07	0.11	0.12	0.13	0.11	0.39	0.41	0.49	0.39
Medical & Health	21	10.30	0.05	0.02	0.07	0.05	0.09	0.05	0.26	0.11	0.41	0.50	0.74	0.53
Services	61	3.59	0.06	0.06	0.17	0.06	0.11	0.11	0.33	0.11	0.44	0.47	0.50	0.46
Government	10	7.75	0.06			0.06	0.11			0.11	0.44			0.44
Food & Drink	24	6.97	0.07	0.04	0.00	0.06	0.12	0.07	0.00	0.11	0.44	0.41	0.40	0.44
Technology	16	4.41	0.05	0.05		0.05	0.08	0.11		0.10	0.41	0.50		0.46
Education	14	8.94	0.06		0.03	0.06	0.09		0.06	0.09	0.40		0.54	0.40
Regional & Cultural	29	12.01	0.05	0.06		0.06	0.09	0.09		0.09	0.42	0.29		0.40
Sports	27	12.02	0.05	0.04		0.05	0.08	0.07		0.08	0.41	0.40		0.41
Business	30	7.67	0.04	0.07	0.22	0.04	0.06	0.12	0.63	0.07	0.41	0.38	0.66	0.40
Religion	9	5.96	0.05	0.06		0.05	0.07	0.08		0.07	0.34	0.31		0.34
Professional	28	9.66	0.04	0.02	0.02	0.04	0.07	0.04	0.07	0.07	0.36	0.44	0.63	0.36
Real Estate	26	8.63	0.02	0.08		0.04	0.04	0.14		0.07	0.35	0.41		0.37
Cities	33	16.60	0.03	0.03		0.03	0.05	0.05		0.05	0.38	0.32		0.37
(no category)	10	26.19		0.01	0.03	0.01		0.02	0.06	0.02		0.44	0.49	0.44

TABLE XII: **Uniformly random sample of malicious domains** from gTLDs.

.xin	**.sbs**	**.bond**
com-dpuj.xin	xy831.sbs	cybersecuritydegreesonline.bond
com-mcob.xin	ya2twm.sbs	cleaning-gel-82085.bond
com-ticketli.xin	ideaakdi.sbs	portable-power-station-88407.bond
com-ytcr.xin	intesasanpaolo97it.sbs	247-nurse-92799.bond
com-tieuy.xin	eiylfrm.sbs	telegelrs.bond
org-tlj.xin	luis625.sbs	warehouse-inventory-48771.bond
paymentxtt.xin	9880005com10xl01.sbs	prefabricated-homes-67793.bond
txtagwxw.xin	trmhsn.sbs	home-care-17857.bond
telegeltla.xin	qweftjtp.sbs	air-condition-98954.bond
com-web.xin	zuekvnua.sbs	hr-outsourcing-66318.bond

TABLE X: **Legacy generation TLDs.**

Timeline	**Legacy generation TLDs**
up to 2000	.com, .net, .org, .gov, .mil, .edu, .int
2000 - 2004	.info, .biz, .museum, .name, .coop, .pro, .aero
2004 - 2012	.asia, .cat, .jobs, .mobi, .tel, .travel, .xxx, .post

TABLE XI: **The number of categorized and registered domains for each snapshot date.**

Snapshot Date	**# Categorized Domains**
2021.06.15	288,281,521
2022.01.20	310,002,907
2022.06.16	319,845,662
2023.02.16	327,581,776
2023.04.06	329,279,168
2024.10.17	329,777,226
2025.02.27	326,985,654
2025.09.06	343,148,606

APPENDIX

In Table VIII, the names of the sponsoring organization that we shortened are the following: Elegant Leader* is Elegant Leader Limited, Spec Broadcasting* is SPECIAL BROADCASTING SERVICE CORPORATION, Dotcfd Registry* is DOTCFD REGISTRY LTD, DOTSTRATEGY* is DOTSTRATEGY CO., and First Registry* is First Registry Limited.

Table IX enumerates all TLDs by TLD category type, average price and presents the three reputation metric per TLD generation.

Table X lists the legacy gTLDs. Our list of 305 ccTLDs does not include three discontinued TLDs (.an, .tp, and .um) [3], [58], [59], four TLDs where users cannot yet register domains (.bl, .bq, .eh, and .mf) [7], [8], [14], [42], and five internationalized TLDs where their original two ASCII letter ccTLD also exists and registrations or traffic is so low that they do not appear in our datasets (xn–4dbrk0ce, xn–mgbcpq6gpa1a, xn–mgbah1a3hjkrd, xn–mgbai9a5eva00b, and xn–q7ce6a) [37], [36], [5], [44], [50], [41].

Table XI provides the number of categorized and registered domains per snapshot used to analyze the reputation of TLDs.

Table XII lists uniformly randomly selected malicious domains from selected new generation TLDs.