

# Domain Name Lifetimes: Baseline and Threats

Antonia Affinito<sup>\*</sup>, Raffaele Sommese<sup>†</sup>, Gautam Akiwate<sup>‡</sup>,  
Stefan Savage<sup>‡</sup>, KC Claffy<sup>‡§</sup>, Geoffrey M. Voelker<sup>§</sup>  
Alessio Botta<sup>\*</sup>, Mattijs Jonker<sup>†</sup>

<sup>\*</sup> *University of Naples "Federico II"*, Naples, Italy

<sup>†</sup> *University of Twente*, Enschede, the Netherlands

<sup>‡</sup> *CAIDA* / <sup>§</sup> *UC San Diego*, La Jolla, CA, USA

**Abstract**—The domain name system (DNS) is a key component of the Internet. The DNS is essentially a hierarchical and distributed database that involves – and is operated by – many independent parties that fulfill various roles. Top-level domains such as `.com` and `.co.uk` are run by registries. Registrants can register domain names, usually through so-called registrars, but sometimes directly with the TLD registry.

Domain names go through a well-defined life-cycle and names that are only short-lived in ways break expectation. In this paper, we study domain name lifetimes at scale and over a ten-year period. We focus on ten prominent TLDs and observe that under most, the vast majority of lifetimes (95%) last exactly the minimum registration term of one year. The exception to this is `.com`, which sees 40% of lifetimes renewed for at least one more year. We also identify lifetimes that are suspiciously short-lived (e.g., 80% under `.xyz`). Using blocklist data we confirm that about 25% are reportedly malicious and study indicators if names are taken down and how quickly. Finally, we empirically study malicious name registration campaigns and show that this involves registrars that offer bulk registration options.

## I. INTRODUCTION

The Internet Corporation for Assigned Names and Numbers (ICANN) determined a well-defined life-cycle for domain names that nominally leads to domain name lifetimes of yearly granularity. In most cases, the lifetime of a domain name is under the direction of its registrant, with whom rests the decision whether or not to renew the registration. However, there are other possible factors, notably if domain names are used for abusive purposes and taken down.

While the DNS and domain abuse are extensively studied in the literature, the area of domain name lifetimes is arguably still dim. In this paper, we take steps towards closing this gap. We analyze domain name lifetimes under the ten largest top-level domains in CAIDA’s DNS Zone Database [1] across a time span of ten years. To empirically validate the idea that shorter lifetimes can be the result of abuse take-down efforts, we use a large blocklist feed of malicious names and demonstrate that many short-lived names are indeed malicious. We make the following contributions in this paper:

- We perform an analysis of domain lifetimes among 10 of the largest TLDs over a ten-year period, showing that one-year lifetimes predominate (~95% of lifetimes last exactly one year) in most TLDs except `.com`, where 40% of the domains have longer lifetimes;

- Using blocklist data, we evaluate the prevalence of malicious domain names across the TLDs and reveal that a large fraction of malicious names have shorter-lived lifetimes. We also show that malicious names are substantially shorter-lived in some TLDs compared to others (e.g., 80% of malicious `.xyz` names live shorter than the minimum registration term of one year);
- We show signs that malicious names are acted upon and provide insights into take-down times, while we also provide indications that some malicious names are not acted upon and are left to linger;
- We identify a number of malicious registration campaigns and empirically show that such campaigns can include registrars that offer bulk registration options.

All in all, our findings help shed light on domain registration practices and the use of domain names and malicious behaviors. We also shed light on operational practices by studying indicators of the presence (or absence) of take-down efforts.

This paper is structured as follows. We provide background information in Section II and discuss related work in Section III. We detail our methodology in Section IV and our data sources in Section V. In Section VI we present our results and findings. Finally, we conclude in Section VII.

## II. BACKGROUND

In this section, we discuss the implementation of the DNS namespace, domain names, and how users obtain these domain names. We also discuss the typical lifetime of a domain name and the various reasons a domain may be taken down before the end of its contracted lifetime.

### A. DNS Namespace: Top Level Domains

The DNS namespace, first defined in RFC 1034 [2], is a hierarchical inverted tree structure. The root of this inverted tree structure is referred to as the DNS Root. The DNS Root explicitly delegates each individual zone under it, typically referred to as a top-level domain (TLD) (e.g., `.com` or `.nl`) to organizations, called registries, who are responsible for that branch of the namespace, i.e., the TLD zone. Registries are typically responsible for administering authoritative name-servers which provide nameservice for all zones under the TLD. For instance, the registry for `.com`, Verisign, operates

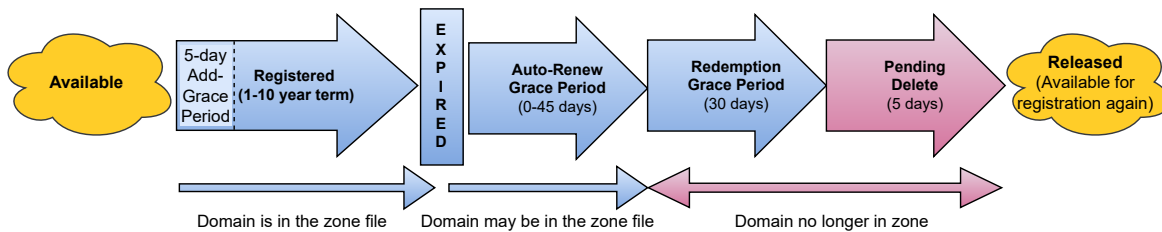


Fig. 1. A visualization of the life-cycle of a domain, along with all possible states

authoritative nameservers which provide nameserver delegations for `example.com` (typically referred to as a second-level domain (SLD) or registered domain). Those nameservers have authority over all zones under `example.com` (e.g., `www.example.com`). Each of these zones can further sub-delegate specific branches of the namespace under it.

The TLDs are typically categorized into two types: generic TLDs (gTLDs) and country-code TLDs (ccTLDs). The gTLDs are further divided into two categories: legacy gTLDs (e.g., `.com`, `.org`, `.net`), and new gTLDs (ngTLDs) (e.g., `.xyz`, `.loan`) introduced by ICANN in 2012 under the new gTLD program [3]–[5]. On the other hand, ccTLDs are assigned to specific countries (e.g., `.nl`, `.uk`, `.de`).

## B. Registries and Registrars

The *registry* is the organization responsible for administration of a TLD. Typically, the administration of TLDs is delegated to a single organization under contract with ICANN [6]. As part of recent transparency initiatives, ICANN now also mandates that the registries operating a TLD make available via the ICANN Centralized Zone Data Service (CZDS) the *TLD zone file* — which includes a list of domains under the TLD and their corresponding nameserver delegations. The TLD zone files obtained from ICANN CZDS and other sources (Section V) are the basis of this work.

The registries contract with *registrars* to provision new domains. Registrars interface between users looking to obtain domains and the registry administering the domain. A *registrant* looking to obtain a domain name under `.com` would contract with a registrar (e.g., Enom) who in turn interfaces with the registry operating `.com`, Verisign, to query the availability of the domain name and then claim it on behalf of the *registrant*. On successful purchase of a domain, the registrar is then responsible for the domain until it expires or is transferred by the *registrant*. In addition to contracts with the *registry*, *registrars* also have to be accredited by ICANN [7].

## C. Domain Name Life-Cycle

The ICANN registry agreement contract that delegates administration of the TLD also lays out in detail the expected life-cycle of a domain, which includes a number of different possible states. Figure 1 provides a visualization of the life-

cycle of a domain name in a generic TLD zone <sup>1</sup>, and illustrates the following states [8], [9]:

- **Available:** A registrant can use a registrar to find the domain names *available* for registration;
- **Registration:** The registrant can purchase the *available* domain name for a period of at least a year. The registration term may be as long as 10 years. The registrant has a 5-day *Add Grace* period during which to undo the registration and receive a refund for the registration fee;<sup>2</sup>
- **Expiration and Renewal:** At the end of the registration term, when a registration is set to expire, the registrant can choose to renew the domain name. On renewal, the registration period (and consequently the expiration date) is extended. The registrant is allowed two grace periods that start after expiration. The first of these grace periods is the *Auto-Renew* period, which ranges from 0 to 45 days. The *Auto-Renew* period allows the registrant to renew the domain name without incurring a penalty;
- **Redemption Period:** After the *Auto-Renew* grace period ends, the *Redemption* grace period starts. In this state, the domain is generally deleted at the registrar, but it still exists in the registry’s database. This period, usually 30 days, allows the registrant to renew the domain name with an additional *redemption fee*;
- **Pending Delete:** If the registrant chooses not to renew, the domain will enter the *Pending Delete* state, which is usually 5 days long and during which it is not possible to renew the domain name;
- **Released** After deletion and release, the name can be re-registered. This state is equivalent to the *available* state.

At registration, a registrant procures a domain for a period of at least one year. However, the registrant may choose a longer registration period — anywhere from one to ten years (but always at the granularity of a year). Note, a registrant may transfer a domain following an initial ICANN policy mandated lock of 60 days, but such a transfer requires purchase of at least an additional year of registration beyond the original registration period [8], [11]. Thus, a domain with a lifetime that is not at a granularity of a year (modulo the ICANN mandated grace periods) indicates an action taken by either a registrar or a registry in response to some complaint.

<sup>1</sup>Since ccTLD registries have wider latitude in how they administer their zone, the life-cycle for domain names in ccTLD zones may differ significantly.

<sup>2</sup>After 2009 a mechanism was introduced to limit abuse of this no-cost grace period, effectively eliminating domain tasting abuse [10].

#### D. Early Take Down of Domains

There are a variety of reasons why a domain may not last the one-year duration in the TLD zone file. While there are legitimate reasons for a domain to disappear from the zone files before one year (*e.g.*, a registrant choosing to withhold their domains from being listed), in most cases the disappearance is indicative of take down in response to illegitimate activity. This illegitimate activity can run the gamut from payment fraud to coordinating botnet activity. These take downs can be roughly bucketed into three categories. The first is the “early take down”: a registrar discovers an irregularity with the domain registration and takes down the domain. For example, a registrar may discover the registrant used a stolen credit card to purchase the domain. These domains are predominantly taken down before they are involved in malicious activity. The second category is “malicious domain take down”: a registrar, or registry takes a domain down in response to abuse reports [12].<sup>3</sup> In this case, the domains are taken down after they are involved in malicious activity. The final category is “co-coordinated legal action”: law enforcement and other organizations seize large numbers of domains. For instance, in 2011, the US Federal Bureau of Investigation (FBI) seized domains related to Coreflood [14], [15]. Typically, these take downs are targeted at Domain Generating Algorithms (DGAs) associated with malware and botnets. Recently, ICANN made efforts to empower *registrars* and *registries* to unilaterally take down domains involved in ongoing security incidents [16], [17]. In this case, some domains may be taken down by *registrars* or *registries* preemptively before they are involved in malicious activity.

While short-lived domains (domains lasting for less than a year) are indicative of malicious activity, it is important to not use these *solely* as a metric for malicious activity. The “early” and the “malicious” take downs are highly dependent on registrars. While the registrars are required to look into abuse as per their ICANN contract, registrars are routinely overwhelmed, at times by false reporting, leading to long resolution times [13] which may result in domains not being taken down. Consequently, our analyses rely on blocklists as an indicator for malicious activity.

### III. RELATED WORK

Domain name abuse is extensively discussed in the literature. For malicious registration detection, several works go beyond blocklists to find additional ways to detect malicious domain names. Sun *et al.* [18] propose a methodology named *HinDom* to detect malicious domain names using a classification based on relationship between clients, domains and IP addresses. Their methodology was able to detect a long-buried botnet and several malicious domains in a real-world scenario. Using an Extreme Learning Machine, Shi *et al.* built a malicious domain detector that uses several features

<sup>3</sup>Note, as per the ICANN Registrar Accreditation Agreement, a registrar must maintain an abuse contact to receive abuse reports involving domain names sponsored by the registrar [12], [13]

(*e.g.*, length of domain, entropy, number of IP addresses) and achieves an accuracy greater than 95% [19]. Hason *et al.* used similar features to build a classifier of malicious domain names, also achieving an accuracy of 95.2% [20].

Bilge *et al.* built a system to detect malicious names, adopting machine learning techniques based on passive DNS data [21]. Combining 15 behavioral features, their system identifies a large number of malicious hosts. Vinayakumar *et al.* assessed the efficacy of using deep learning to detect malicious domain names [22]. They applied CNN (Convolution Neural Network) and RNN (Recurrent Neural Network) approaches to a large volume of DNS logs.

Previous studies have also explored malicious campaigns registered in bulk for large-scale attacks [23]–[25]. Cybercriminals register considerable numbers of domains to quickly replace detected domains and recover from take down efforts [26]. Vissers *et al.* examined malicious campaigns in the registration data related to the .eu TLD [27]. Looking at domain names with the same registrant and registry information, they found that 80.04% of short-lived domain names could be tied to 20 campaigns. Furthermore, they claim that these campaigns differed in terms of duration: from one month to a year and beyond. Their results are in line with ours. Indeed, we detect several campaigns characterised by malicious names with overlap in features. Contrary to their analysis of only the .eu TLD, we investigate a selection of 10 TLDs that represent a sizable part of the global namespace.

Regarding domain name lifetimes, Foremski *et al.* analyzed malicious short-lived domain names, finding that 9.3% of new domains were deleted in the first seven days, with a median lifetime of 4 hours and 16 minutes. Their study leverages the NOD (Newly Observed Domain) service based on passive DNS observation and active DNS measurements [28]. In addition, they inspected several possible causes of deletion, stating that blocklisting is responsible for 6.7% of it. As with Foremski *et al.*, we study domain name lifetimes, focusing on the ten largest TLDs, and we use the DBL blocklist to identify malicious domain names. Unlike this work, we examine all domain name lifetimes (not only the newly observed domains and malicious ones) included in CAIDAs Zone Database over ten years. In addition, to further investigate the causes of their short lifetime, we examine the presence and the lifetimes of malicious domain names in 2018-2021. Barron *et al.* [29] show that early deletions of domain names are significantly correlated to potentially malicious activities, and we have similar findings. The authors also show that short-lived malicious domain names tend to be longer and more pronounceable or prone to typo-squatting. We examine related characteristics of malicious domains and confirm largely similar results.

Finally, Korczynski *et al.* [5] reveal that abuse activity shifted from legacy gTLDs to newer gTLDs, in part due to registration prices. In our work, we show also that legacy gTLDs still include a considerable number of malicious domains. Lauinger *et al.* examined the WHOIS records of domains about to be deleted in DNS zone files during the stages of the expiration and re-registration [30]. They found

that registrars implement different cancellation techniques that are not always compatible with the life cycle of domains. In contrast with our work, they analyze fewer TLDs and do not inspect the expiration and re-registration of malicious domain names. Finally, an interesting study regarding the new domain name registrations related to COVID-19 domain names was conducted by ICANN [31]. They found that these domains were also used for malicious purposes, around 1.8% being flagged.

#### IV. METHODOLOGY

**Lifetime Inference.** Central to this paper is our ability to infer domain lifetimes. We devised a relatively straightforward methodology that uses zone files. Recall from Section II that for a domain name to functionally exist, the parent zone (*i.e.*, registry) typically delegates authority to a name server of choice of the domain name owner (*i.e.*, registrant). We assume that if a domain name is “alive”, its nameserver delegations will be present in the zone file. This assumption does not always hold. There could be cases in which NS records are absent, for example when a domain is parked or in a *grace period*. To account for these blind spots in zone files, our methodology allows for gaps of at most 90 days before considering a lifetime closed. We choose this value arguing that it is sufficient to capture temporary disappearance, *e.g.*, during one or both of the possible grace periods, but not so long as to capture re-registration after release. The 90-days threshold includes a margin of 10 days over the 80 days domain removal scenario defined in subsection II-C, to account for possible errors in zone file collections.

Because of the granularity of our data sources (Section V), we consider lifetimes in terms of multiple days. As we will show in Section VI-E, WHOIS data for malicious domains validate that our assumptions provide a good estimation of domain lifetimes. Note that the lifetimes that we define and consider in this paper are *closed lifetimes*. More specifically, for a given domain, these are the lifetimes for which we are able to observe the start and end, because the domain creation and expiration dates fall within the boundaries of our data.

**Malicious Domain Names.** The other important part of our methodology relates to how we consider and analyze *malicious* domain names. To make a determination of maliciousness, we rely on a blocklist (Section V) as input. To characterize malicious names and study the presence and properties of such names under various top-level domains, we consider the registered domain name part. We extract the registered domains from blocklisted names with Public Suffix List even though they may contain additional labels. This puts the considered entries at the same level as the names (technically, zones) in NS records in TLD zone files, which in most cases do not contain deeper levels of nesting. We note that this choice could lead to classification errors for registered domain names that are in the parent zone to both malicious and non-malicious names (consider, *e.g.*, the shared suffix under Dynamic DNS service providers). Nevertheless, we argue that the number of third-level domain hosting services compared to the number of

TABLE I  
TOP 10 TLDs DATA SET, SHOWING CZDS START AND END DATES, THE NUMBER OF LIFETIMES SEGMENTS INFERRED PER TLD, AND THE NUMBER OF UNIQUE DOMAIN NAMES INVOLVED

TLD	Start Date	End Date	# LT Segments	# Names
.com	2011/04/11	2021/02/14	168.9M	156.5M
.net	2011/04/11	2021/02/14	20.6M	19.4M
.info	2011/06/06	2021/02/14	16.3M	15.7M
.org	2011/05/08	2021/02/14	12.7M	12.1M
.xyz	2014/03/31	2021/02/14	12.8M	12.2M
.top	2014/08/04	2021/02/14	12.1M	11.7M
.icu	2015/06/24	2021/02/14	5.6M	5.6M
.biz	2011/05/06	2021/02/14	4.9M	4.6M
.us	2011/05/06	2021/02/14	4.6M	4.4M
.loan	2015/03/30	2021/02/14	4.6M	4.6M

second-level domains is negligible. In fact, they are managed by established companies that are not likely to have short-lived domain names. We, however, consider registered names that expire, which are less likely to introduce such classification errors.

#### V. DATA SOURCES

We use two data sets in this paper, together with supplementary data. We obtain the primary data sets from two sources: zone file data and malicious domain names.

**Zone files.** We use data from CAIDA’s DNS Zone Database (DZDB) [1], which is built on a sizable collection of TLD zone files and captures the history of domain names, name servers and IP address records. Following the inception of ICANN’s Centralized Zone Data Service (CZDS), most of the newer gTLDs were added to DZDB, which currently contains approximately 210 million names. Our analysis of lifetime behaviors of domain names involves a sizable part of the DZDB data. We consider a time period of roughly ten years, starting at the earliest DZDB data (2011/04/11 – 2021/02/14).

For our analyses, we consider the Top 10 TLDs in DZDB in terms of total size ranking since 2011. The Top 10 is representative (they cover 87% of all the SLDs in our entire data set) and allow us to provide insights into administration policies for individual TLDs. Table I shows the Top 10 TLDs, a summary of DZDB data available for them, and the number of lifetimes that we infer. Taking .com as an example, we infer 169M lifetimes throughout the ten-year period. Relative to the total of 157M unique .com names, this shows that for some names we infer altogether new registration (and another lifetime), as per our methodology (see Section IV).

**Blocklists.** As an indicator of malicious activity, we rely on the Domain Block List (DBL) maintained by the Spamhaus project. Our data set consists of daily snapshots of the DBL feed from 2018/01/01 to 2021/02/14. While a single blocklist is a narrow window into malicious domain related activity, we find this window illuminating. Since our DBL data set starts in 2018, we only consider DZDB data from 2018 onwards for our analysis in Section VI-B. However, a limitation of this data set is that it does not include the type of malicious activity associated with the domain name. Consequently, we cannot

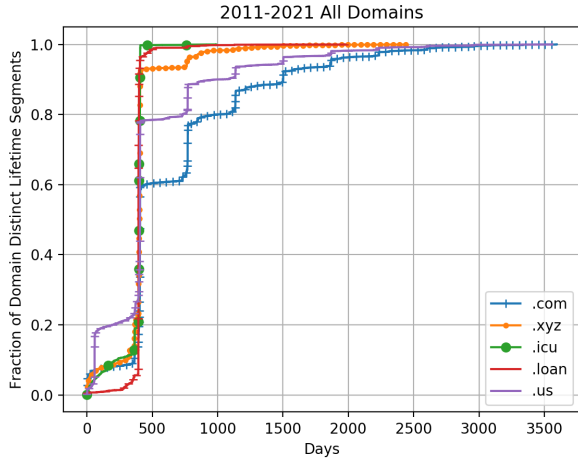


Fig. 2. Domain name lifetime in selected Top 10 TLDs under consideration

show the trends of the lifetimes by varying the malicious activity. To overcome this limitation, we relied on the “Domain Status and Categorization” API provided by Cisco Umbrella Investigate [32], but it did not give us enough data to support this analysis.

**WHOIS data.** We rely on data provided by Cisco Umbrella to investigate malicious domain name registration campaigns (Section VI-E). Their Investigate API gives a complete view of domain name, IP address, ASN, and malware file details to help identify misused infrastructure and to predict future threats [33]. Relevant to our work, the provided WHOIS data includes registration information for domain names, including creation date; registrant organisation, city and country; and registrar name and IANA ID.

## VI. RESULTS

In the section, we present our results. We start with an overview of lifetimes for domains under the Top 10 most populous TLDs in our data (Section VI-A). We then investigate malicious name lifetimes (Section VI-B) and suspicious short-livedness (Section VI-C). Next, we look at post-blocklist life and possible take down actions (Section VI-D). Finally, we investigate malicious registration campaigns (Section VI-E).

### A. Lifetime of Domain Names

As explained in Section II-C, a domain name can go through five different life-cycle states, of varying lengths, which together form the *lifetime* of a domain name. We expect most domain names to be visible in the zone files while *registered*. This expectation allows us to evaluate the lifetime of domain names as the time from when a domain is first and last seen in the zone file. In our methodology (Section IV), we consider a lifetime to have ended when, after appearing, it is absent from the zone file for 90 days or longer. We treat reappearance beyond this point as an altogether new registration.

We infer domain name lifetimes for the Top 10 TLDs. Figure 2 shows CDF plots for domain names under `.com`,

`.icu`, `.xyz`, `.loan`, and `.us`. For this analysis, we consider domain names in zone files that have valid first seen and last seen values in the period 2011/01/01 through 2021/02/14, capping the lifetime at roughly 3700 days. Therefore, our analysis does not include domain names still active at the last collection time. For clarity, we do not plot the other five TLDs, but they display similar trends as we further detail below. The results show that a considerable number of domain names are registered for lifetimes of one year in most TLDs, with all TLDs showing a sharp increase around 410 days: one year plus the *Auto-Renew* grace period of 45 days.<sup>4</sup> Moreover, zones also see lifetimes that are under the minimum registration term of one year, which may be the result of take-down efforts (see Section II-D).

For `.com` domains, 60% of their lifetimes are at most a year (101M of 169M lifetimes in Table I), and about 20% of `.com` names involve lifetimes of three years or longer. The TLDs `.org`, `.net`, `.info` and `.biz` all show similar trends (not plotted). These zones belong to the first set of ICANN gTLDs, originally created between 1985 and 2001 [34]. A comparable trend occurs for the domain names of the ccTLD `.us`. In this case, 78% of their lifetimes last at most one year (3M of 4.6M lifetimes), and around 20% of `.us` lifetimes are longer than three years. Moreover, `.us` includes roughly 20% of domain names with a lifetime less than 70 days, in contrast with the `.com` and the analogous TLDs where this value is significantly lower (0.08% of 169M lifetimes). In contrast, for `.xyz`, about 93% of lifetimes (11.9M) are about one year, 95% of at most about two years, and only a small percentage of domains remain registered three years or longer. The `.top` TLD (not plotted) presents a similar trend to `.xyz`. Indeed, around 2019, these were the new gTLDs with the most number of registrations [35].

An interesting behavior seen relates to `.icu`. This TLD was created in 2015, but the first domain names under it were registered around 2018. Therefore, we have a three-year observation period for this TLD. The trend that becomes apparent is that most lifetimes are one year. The same applies to `.loan`, except that `.loan` includes fewer domains with a duration of less than 400 days than `.icu`. We have also evaluated the number of domains still active at the end of the collection period, and found that `.com` is still the TLD with the highest number of domains.

*Key takeaway: a significant number of lifetimes last exactly one year (the minimum registration term). Furthermore, a non-negligible number is also shorter-lived.*

### B. Malicious Domain Names

To better understand possible causes for patterns in lifetimes, and considering that lifetimes can be cut short as a result of take-down efforts (see Section II-D), we match domain names from the zone files with those included in the Spamhaus DBL blocklist. Note that, for now, we consider *any* malicious

<sup>4</sup>The edge is slightly slanted because registered names may take a few days to appear in the zone file, as we will show in Section VI-E.

TABLE II

TOP 10 TLDs DATA SET WITH MALICIOUS NAMES, SHOWING THE NUMBER OF LIFETIMES SEGMENTS INFERRED AND UNIQUE NAMES IN CZDS DATA FOR 2018+, AS WELL AS THE MALICIOUS FIGURES

TLD	# Total LT Segments	# Total Names	# Malicious LT Segments	Malicious Names (%)
.com	32.7M	32.2M	2.5M	7.53%
.net	2.6M	2.6M	201K	7.62%
.info	2.0M	2.0M	256K	12.41%
.org	1.7M	1.6M	51K	3.05%
.xyz	3.5M	3.4M	233K	6.62%
.top	6.2M	6.1M	1.3M	21.56%
.icu	5.7M	5.6M	244K	4.27%
.biz	928K	950K	270K	28.46%
.us	794K	790K	156K	19.67%
.loan	2.0M	2.0M	240K	12.20%

name, regardless of the duration of its lifetime. In a later section we will focus on short-lived names in particular.

Our overall lifetime analysis and Figure 2 capture a ten-year period. As we obtained DBL data from Jan 1, 2018 onward, we can only match domain names registered after this date against DBL inclusion. For this reason, going forward we consider zone files data for 2018 and onward.

The lifetime of malicious domains is usually considerably shorter than that of benign names [19], [36]. Malicious names are deactivated once revealed or because hackers want to minimize blocklist interference. For example, many spam domains are only active for one day, in an attempt to avoid detection and from being added to blocklists [23], [27].

We calculate the percentages of malicious domains in the Top 10 TLD data for 2018 and beyond and extract malicious lifetimes. Table II summarizes the results. We show the total number of names and lifetimes inferred as before (Table I). The .biz TLD contains the highest percentage of malicious domain names (28.46% of 950K), followed by .top and .us. While lower, .loan and .info are still above 10%. Under the largest TLD .com, 7.5% of domains are malicious. Spamhaus estimates an abuse score for each TLD based on the prevalence of malicious domains<sup>5</sup>. Our findings are largely in line with these scores: the current Spamhaus scores identify .biz, .top and .us as most-abused, and .org as least.

Figure 3 relates specifically to the lifetimes of malicious domain names. We show only the CDFs related to .com, .xyz, .icu, .loan, .top, .biz. Lifetimes for malicious names under the other TLDs show trends similar to the counterparts of these TLDs we reported in Section VI-A. Malicious domain names in .xyz generally have shorter lifetimes than those under other TLDs. The TLD .icu is next in rank. The TLD .loan sees a considerable number of malicious domain names that have a lifetime of around one year, followed by .biz. The TLD .top includes a high percentage of malicious domain names (e.g., 21.56%) with longer lifetimes than the other TLDs (i.e., 12% of malicious .top domain lifetimes are shorter than 365 days). More specifically, 97–

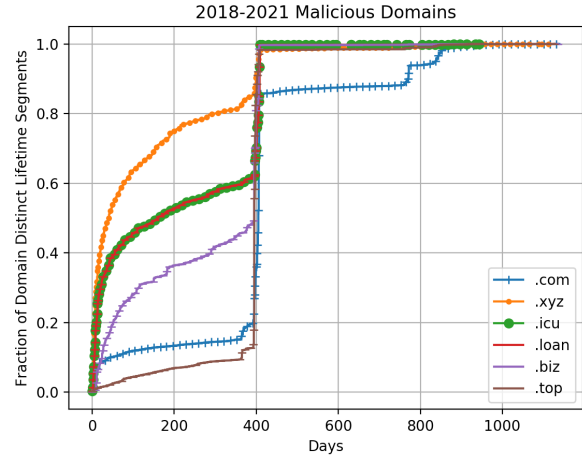


Fig. 3. Domain name lifetime in various Top 10 TLDs under consideration for names that are reportedly malicious

99% of malicious .xyz, .icu, .loan, .top, .biz domain lifetimes are shorter than 410 days. For .com it is 86%. The .xyz TLD could stick out for multiple reasons. First, we see that malicious .xyz domains are less likely to be renewed in general (Figure 2). Second, as we show in Section VI-D, malicious .xyz names are acted upon quicker compared to other TLDs.

*Key takeaway: A large fraction of malicious domains have a lifetime shorter than one year, which is indicative of take down efforts or otherwise technical removal from the zones.*

### C. Short-Lived Domain Names

We now focus on domain name lifetimes of 364 days or shorter. We chose this threshold because it captures domain names that live less than the minimum registration term, considering the minimum of 0 days under grace (Section II-C). For the overall DZDB data (i.e., starting at 2011), 6.19% of lifetimes are 364 days or shorter. For 2018 onward, which aligns with the DBL data available to us, the percentage is 19.57%: 11.3M lifetimes involving 11.0M unique domain names. We cannot make a strong inference from the relative increase in percentages, but do note that anecdotal evidence suggests increases in domain name abuse [37]. In addition, although the ICANN report shows an increase in the number of registrations and a decrease in the number of abuses from 2017 to 2022, we see a drop in the number of new registrations from 2018 to 2020 [38]. Furthermore, the percentage of lifetimes less than 364 days is 15.5% in 2018 and 16.2% in 2019. We cannot estimate this percentage in 2020 because our data set lasts until February 2021. Considering DBL data, we confirm that 24.27% of short-lived lifetimes involve malicious domain names. These  $\sim 1.3$ M lifetimes involve almost the same number of domain names, and hence we rarely encounter malicious names for which we infer multiple (short-lived) lifetimes.

We calculated the percentages within each Top 10 TLD to investigate how they compare. We find that .biz has the

<sup>5</sup><https://www.spamhaus.org/statistics/tlds/>

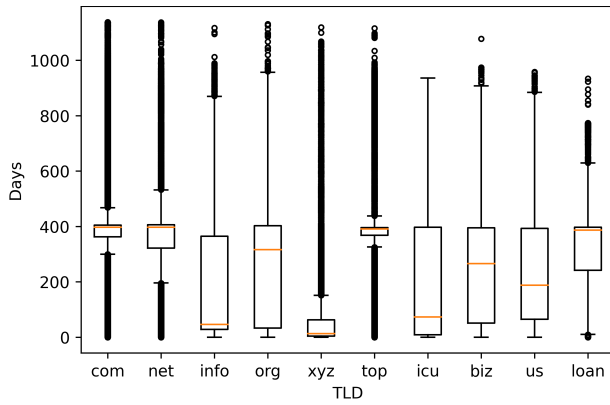


Fig. 4. Number of days elapsed between the insertion of malicious names on the blocklist and their removal from the zone file

highest percentage: 34%. Recall from Section VI-B that this TLD also sees the highest percentage of malicious names. The `.icu` and `.top` TLDs contain the lowest percentages of short-lived domain names. At the same time, however, if we consider strictly malicious domain names in these TLDs, we see that many fit the short-lived criterion (see Figure 3).

*Key takeaway: a non-negligible number of domain names that have short-lived lifetimes are also demonstrably used for malicious purposes.*

#### D. Post-Blocklist Life and Removal

We investigate how much longer domain names live after appearing on the blocklist, noting that removal can be the result of take-down efforts. To this end, we look at the number of days between DBL insertion and removal from the zone file.

**First**, we consider any malicious name (*i.e.*, not necessarily short-lived ones), including names that naturally expire.

Figure 4 shows the resulting boxplots for the Top 10 TLDs. The TLDs `.com`, `.net`, and `.top` see median deletion times of 379, 379, and 387 days, respectively. These values are close to 410 days (one year plus the auto-renew grace period), which is the minimum lifetime of a domain if it is not renewed. Therefore, this plot shows that these three TLDs include most blocklisted names that may have naturally expired rather than being acted upon (*e.g.*, by registries or registrars). The TLD `.xyz` shows the opposite: a median of just 13 days. With the exception of `.xyz`, the upper quartiles are close to the one-year mark, suggesting that a long tail of names under most TLDs naturally expire. Finally, looking at 95-percentiles, we see that there are malicious domains that live for multiple years before expiring.

**Second**, we consider short-lived malicious names, postulating that malicious names that do not live for the minimum registration term of one year are likely to have been taken down. Figure 5 shows the resulting boxplots. The `.xyz` TLD again shows the lowest median value (10 days here), indicating that malicious domain names are removed from this zone shortly after being blocklisted. The short boxplot for `.xyz`

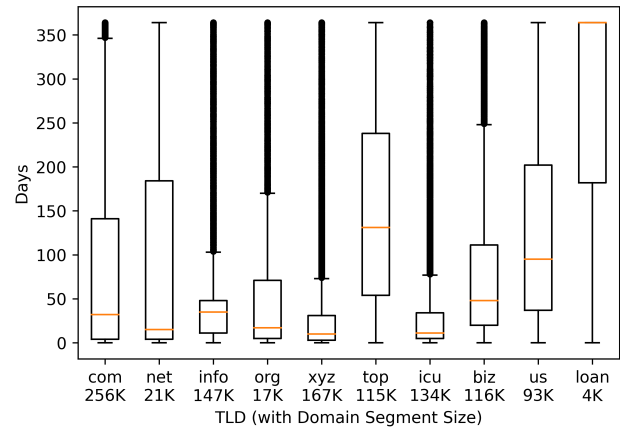


Fig. 5. Number of days elapsed between the insertion of short-lived, malicious names on the blocklist and their removal from the zone file

also suggests that few malicious domains live anywhere near the minimum registration term.

We observe different behavior for `.top`, which sees a high median value of 131 days. Its relatively tall plot and upper quartile shows that some malicious names live for a considerable amount of time after being blocklisted. Similar observations can be made for several other TLDs such as `.net` and `.us`, although not as pronounced. With the exception of `.top`, the results are comparable for the situation in which we considered any malicious name, regardless of whether they are short-lived. Considering Figure 5, we conclude that, for suspected take down efforts, the median removal time is largely between 0 and 2 months. Given that the 4k short-lived malicious names represent only a tiny fraction of the malicious `.loan` names (0.02%), we do not consider its results representative.<sup>6</sup>

As only 24.27% of the short-lived lifetimes involve malicious names in DBL data, we consulted two parties — a ccTLD registry and a large global registrar — about other possible reasons for domains being short-lived. The registry stated that the blocklist perspective only accounts for a subset of short-lived domains, but what is missed is still due to abuse. The registrar indicated that malicious domains can be re-registered with them after being taken down and after expiration of the redemption period. Finally, we note that some registrars, such as Freenom, provide an API to security researchers to immediately take down free domains following signs of abuse<sup>7</sup>. We do not know if such mechanisms are available for the TLDs that we considered. However, it could help explain differences in take-down timings.

*Key takeaway: We see indications that malicious names — especially short-lived — are taken down. Names under some TLDs are seemingly acted upon quickly. However, many names are also left to naturally expire.*

<sup>6</sup>As we show in Section VI-B, the malicious names that we found in `.loan` are typically longer lived. One possible reason is that it does not react to abuse notifications.

<sup>7</sup>[https://www.freenom.com/en/antiabuse\\_api.html](https://www.freenom.com/en/antiabuse_api.html)

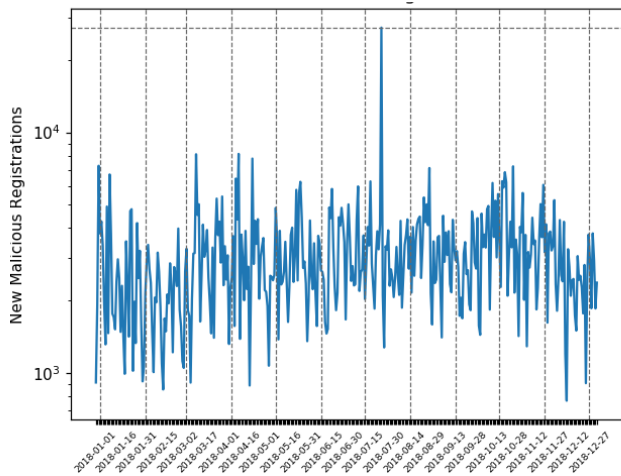


Fig. 6. New Malicious Registrations in 2018 - .com TLD

### E. Investigating Malicious Campaigns

Some cybercriminals register a considerable number of domain names for malicious purposes at once [27], [39]. There are registrars that make this possible by offering bulk registration options. To investigate, we study malicious name registrations over time and look for signs of bulk registration. We cross-reference DZDB and DBL data and calculate how many new malicious registrations occur every day. Figure 6 shows the results for .com, which usually sees 1K – 10K malicious registrations daily and also contains a pronounced spike on July 28, 2018. For other TLDs (not plotted) we observe lower daily averages and also also occasional spikes.

We investigate the suspicious spike, which involves 27k malicious names, for possible causes. Related work has shown that maliciously registered names in bulk can involve overlap in WHOIS features [27]. For this reason, using Cisco Umbrella data, we look for overlap in: the registrant organization, city, and country; and the registrar name and IANA ID. This identifies a campaign characterized by 4362 domain names, which can be tied to a single registrant organization in Malaysia, and the registrar GoDaddy. Furthermore, a considerable number of domains related to this peak (around 17K) were registered by the registrar Alibaba. Both registrars offer bulk registration. We looked for visually prominent spikes for other Top 10 TLDs as well (not plotted). Two peaks occurred on 2018/03/01 and 2018/05/16, respectively, involving 83k and 52k malicious domain name registrations, 93% and 97% of which are under .top. In both cases, Alibaba was also the registrar used, and the names share a single Chinese registrant. Consequently, all malicious spikes analyzed were triggered by a significant number of registrations performed by the Alibaba registrar.

We also looked beyond spikes and examined 12 “average” days of malicious .com registrations, one per month, equally spaced over a year. Figure 6 marks the dates with dashed vertical lines. We identify several smaller campaigns in an average of about 3k daily registrations. We find registrar overlap for GoDaddy, GMO Internet, PDR Ltd. d/b/a, or Xin Net

Technology Corporation. Finally, we looked at the malicious names to further confirm commonality. Using Levenshtein distances we observe that in some campaigns, names differ by only a few characters.

**Lifetime inferences.** We extracted from WHOIS data the creation dates for domain names involved in the 12 snapshots and three peaks, and compared with the registration date that we infer from zone files. This comparison reveals that 80% of domains were registered up to one day and 97% up to two days earlier. This shows that, in most cases, our zone file approach to inferring the date on which a domain name’s lifetime starts is reasonable. (Recall from Section II-C that domain name owners may withhold names from the zone files.)

*Key takeaway: We empirically identify campaigns that can be linked to specific registrars and registrants. Registrars offering bulk registration can appear in multiple campaigns.*

## VII. CONCLUSION

In this paper, we analyzed domain name lifetimes. We showed that, among a representative selection of TLDs, initial ICANN gTLDs (e.g., .com) exhibit a higher renewal rate than newer gTLDs (e.g., .icu). We also see signs that a non-negligible number of domain names do not live as long as the minimum registration term of one year. To investigate possible causes, we examined the presence and lifetimes of malicious names. Half of the TLDs considered involve substantial numbers of malicious names (i.e., 12.20–28.46%). Moreover, malicious names in some TLDs live longer than in others. We see indications that domains are subjected to take down efforts, finding also that in some TLDs this takes place quickly after domains have appeared on a blacklist. Finally, we looked at malicious registration campaigns. We empirically identified a number of them on the basis of WHOIS feature overlap (e.g., registrant or registrar) and also found indicators that some registrars are used regularly to this end. We believe that the investigation of the malicious campaign may be applied also in the threat intelligence or cybersecurity fields. Specifically, the security level of a domain may be pre-estimated by observing its registration features and also whether it belongs to a bulk registration. Future work can extend coverage of TLDs to less popular ones beyond the Top 10, and increase the coverage of malicious names, e.g., considering other blocklists like VirusTotal and Cisco Umbrella Domain status.

## ACKNOWLEDGMENTS

We thank our shepherd Oliver Gasser and the anonymous reviewers for their valuable comments and suggestions. We also thank Moritz Müller, Jakob Dhondt, Anna Sperotto, and Roland van Rijswijk-Deij for their helpful insights. This work was supported in part by: the NWO-DHS MADDVIPR project (628.001.031/FA8750-19-2-0004); the EU H2020 CONCORDIA project (830927); MIUR through the “ICT for Health” project, Dipartimento di Eccellenza (2018-2022) “Ingegneria Elettrica e delle Tecnologie dell’Informazione”; Cisco Systems through the Sponsored Research Agreement “Research Project



for Industry 4.0”; and support from the UCSD Center for Networked Systems.

## REFERENCES

- [1] CAIDA, “DZDB,” [https://catalog.caida.org/details/software/dzdb\\_api](https://catalog.caida.org/details/software/dzdb_api), accessed: 2022/03/03.
- [2] P. Mockapetris, “Domain names - concepts and facilities,” 11 1987. [Online]. Available: <https://www.rfc-editor.org/info/rfc1034>
- [3] Y.-D. Song, A. Mahanti, and S. C. Ravichandran, “Understanding evolution and adoption of top level domains and dnssec,” in *2019 IEEE International Symposium on Measurements Networking (M N)*, 2019.
- [4] T. Halvorson, M. F. Der, I. Foster, S. Savage, L. K. Saul, and G. M. Voelker, “From .academy to .zone: An analysis of the new tld land rush,” in *Proceedings of the 2015 Internet Measurement Conference*, 2015.
- [5] M. Korczyński, M. Wullink, S. Tajalizadehkhooob, G. Moura, A. Noroozian, D. Bagley, and C. Hesselman, “Cybercrime after the sunrise: A statistical analysis of dns abuse in new gtlds,” in *ASIACCS 2018 - Proceedings of the 2018 ACM Asia Conference on Computer and Communications Security*, 2018.
- [6] ICANN, “Generic Top-Level Domain (gTLD) Registry Agreements,” <https://www.icann.org/en/registry-agreements>, (Accessed on 2022/03/03).
- [7] “2013 Registrar Accreditation Agreement,” <https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en>, (Accessed on 2022/03/03).
- [8] ICANN, “Life cycle of a typical gTLD domain name,” <https://www.icann.org/resources/pages/gtld-lifecycle-2012-02-25-en>, (Accessed on 2022/03/03).
- [9] “Domain name life cycle: Life of a typical top-level domain. - connectreseller,” <https://www.connectreseller.com/blog/domain-name-life-cycle-life-of-a-typical-top-level-domain/>, (Accessed on 12/14/2021).
- [10] ICANN, “AGP (Add Grace Period) Limits Policy,” <https://www.icann.org/resources/pages/agp-policy-2008-12-17-en>, (Accessed on 2022/03/03).
- [11] DNSimple, “ICANN 60-Day Lock After Change of Registrant,” <https://support.dnssimple.com/articles/icann-60-day-lock-registrant-change/>, (Accessed on 2022/03/03).
- [12] ICANN, “Registrar abuse reports,” <https://www.icann.org/resources/pages/abuse-2014-01-29-en>, (Accessed on 2022/03/03).
- [13] Namecheap, “Our fight against fraud is just getting started,” <https://www.namecheap.com/blog/namecheaps-fight-against-fraud-is-just-getting-started/>, (Accessed on 2022/03/03).
- [14] Wikipedia, “Coreflood,” <https://en.wikipedia.org/wiki/Coreflood>, (Accessed on 2022/03/03).
- [15] US District Court for the District of Connecticut, “Government’s Supplemental Memorandum in Support of Preliminary Injunction,” <https://www.justice.gov/archive/opa/documents/coreflood-govt-supp.pdf>, (Accessed on 2022/03/03).
- [16] ICANN, “Expedited Registry Security Request Process,” <https://www.icann.org/resources/pages/ersr-2012-02-25-en>, (Accessed on 2022/03/03).
- [17] Domain Incite, “Registrars to get more domain takedown powers,” <https://domainincite.com/26917-registrars-to-get-more-domain-takedown-powers/>, (Accessed on 2022/03/03).
- [18] X. Sun, M. Tong, J. Yang, L. Xinran, and L. Heng, “HinDom: A robust malicious domain detection system based on heterogeneous information network with transductive classification,” in *22nd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2019)*, 2019.
- [19] Y. Yu Shi, G. Chen, and J. Li, “Malicious domain name detection based on extreme machine learning,” *Neural Processing Letters*, vol. 48, pp. 1347–1357, 2017.
- [20] N. Hason, A. Dvir, and C. Hajaj, “Robust malicious domain detection,” in *Cyber Security Cryptography and Machine Learning*, 2020, pp. 45–61.
- [21] L. Bilge, S. Sen, D. Balzarotti, E. Kirda, and C. Kruegel, “Exposure: A passive dns analysis service to detect and report malicious domains,” *ACM Trans. Inf. Syst. Secur.*, 2014, p. 28.
- [22] V. Ravi, S. Kp, and P. Poornachandran, “Detecting malicious domain names using deep learning approaches at scale,” *Journal of Intelligent and Fuzzy Systems*, pp. 1355–1367, 2018.
- [23] S. Hao, M. Thomas, V. Paxson, N. Feamster, C. Kreibich, C. Grier, and S. Hollenbeck, “Understanding the domain registration behavior of spammers,” in *Proceedings of the 2013 Conference on Internet Measurement Conference*, 2013.
- [24] Spamhaus, “Weaponizing Domain Names: how bulk registration aids global spam campaigns,” <https://www.spamhaus.org/news/article/795/weaponizing-domain-names-how-bulk-registration-aids-global-spam-campaigns>, (Accessed on 2022/03/03).
- [25] Y. Zhauniarovich, I. Khalil, T. Yu, and M. Dacier, “A survey on malicious domains detection through dns data analysis,” *ACM Comput. Surv.*, p. 36, 2018.
- [26] G. Aaron, L. Chapin, D. Piscitello, and D. C. Strutt, “Phishing landscape 2021 an annual study of the scope and distribution of phishing,” *Interisle Consulting Group, LLC*, 2021.
- [27] T. Vissers, J. Spooren, P. Agten, D. Jumpertz, P. Janssen, M. Van Wese-mael, F. Piessens, W. Joosen, and L. Desmet, “Exploring the ecosystem of malicious domain registrations in the .eu tld,” in *Research in Attacks, Intrusions, and Defenses*, 2017, pp. 472–493.
- [28] P. Foremski and P. Vixie, “The modality of mortality in domain names,” *Virus*, p. 1, 2018.
- [29] T. Barron, N. Miramirkhani, and N. Nikiforakis, “Now you see it, now you Don’t: A large-scale analysis of early domain deletions,” in *22nd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2019)*, 2019.
- [30] T. Lauinger, K. Onarlioglu, A. Chaabane, W. Robertson, and E. Kirda, “Whois lost in translation: (mis)understanding domain name expiration and re-registration,” in *Proceedings of the 2016 Internet Measurement Conference*, ser. IMC ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 247253. [Online]. Available: <https://doi.org/10.1145/2987443.2987463>
- [31] “An 18 month summary of icanns dnsticr project,” <https://www.icann.org/en/blogs/details/an-18-month-summary-of-icanns-dnsticr-project-2-9-2021-en>, (Accessed on 05/12/2022).
- [32] Cisco, “Umbrella Investigate API: Domain Status, Risk Score,” <https://developer.cisco.com/docs/cloud-security/#investigate-getting-started/getting-started>, (Accessed on 2022/03/03).
- [33] “Cisco Umbrella Investigate,” <https://umbrella.cisco.com/products/umbrella-investigate/>, (Accessed on 2022/03/03).
- [34] ZookNIC, “Domain name counts,” <http://www.zooknic.com/Domains/counts.html>, (Accessed on 2022/03/03).
- [35] E. Roser, “On the podium of the new gTLDs: .top, .xyz and .club,” <https://www.internetx.com/en/news-detailview/on-the-podium-of-the-new-gtlds-top-xyz-and-club-1/>, (Accessed on 2022/03/03).
- [36] J. M. Spring, “Modeling malicious domain name take-down dynamics: Why ecrime pays,” in *2013 APWG eCrime Researchers Summit*, 2013, pp. 1–9.
- [37] Spamhaus, “The most abused top-level domains in 2018,” <https://www.spamhaus.com/resource-center/the-most-abused-top-level-domains-in-2018/>, (Accessed on 2022/03/03).
- [38] “The last four years in retrospect: A brief review of dns abuse trends,” <https://www.icann.org/en/system/files/files/last-four-years-retrospect-brief-review-dns-abuse-trends-22mar22-en.pdf>, (Accessed on 05/20/2022).
- [39] M. Felegyhazi, C. Kreibich, and V. Paxson, “On the potential of proactive domain blacklisting,” in *Proceedings of the 3rd USENIX Conference on Large-Scale Exploits and Emergent Threats: Botnets, Spyware, Worms, and More*, 2010, p. 6.