Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs

Maciej Korczyński*, Samaneh Tajalizadehkhoob*, Arman Noroozian*,
Maarten Wullink‡, Cristian Hesselman‡ and Michel van Eeten*
*Delft University of Technology, Netherlands
‡SIDN Labs, Netherlands

***Abstract*—Over the years cybercriminals have misused the Domain Name System (DNS) – a critical component of the Internet – to gain profit. Despite this persisting trend, little empirical information about the security of Top-Level Domains (TLDs) and of the overall 'health' of the DNS ecosystem exists. In this paper, we present security metrics for this ecosystem and measure the operational values of such metrics using three representative phishing and malware datasets. We benchmark entire TLDs against the rest of the market. We explicitly distinguish these metrics from the idea of measuring security performance, because the measured values are driven by multiple factors, not just by the performance of the particular market player. We consider two types of security metrics: occurrence of abuse and persistence of abuse. In conjunction, they provide a good understanding of the overall health of a TLD. We demonstrate that attackers abuse a variety of free services with good reputation, affecting not only the reputation of those services, but of entire TLDs. We find that, when normalized by size, old TLDs like `.com` host more bad content than new generic TLDs. We propose a statistical regression model to analyze how the different properties of TLD intermediaries relate to abuse counts. We find that next to TLD size, abuse is positively associated with domain pricing (i.e. registries who provide free domain registrations witness more abuse). Last but not least, we observe a negative relation between the DNSSEC deployment rate and the count of phishing domains.**

1. Introduction

The Internet is a rather random collection of interwoven actors, pieced together through years of common practices. There was no true regulation of domain names until 1998 when ICANN was established. Before then, one man, Jon Postel, was in charge of administering the Internet Assigned Numbers Authority, later to be incorporated into ICANN. Even today, alternative domain name allocations such as Namecoin[1] and `.onion` domains exist outside of the limits of regulatory authorities such as ICANN.

Each top-level domain or TLD is responsible for its own common practices. While each TLD is ultimately held responsible by ICANN, there is a large variation underneath. The reputation of each TLD is influenced by other operators in the space, such as registrars, webhosting companies and name server operators. There are a number of characteristics that influence the TLD's health in terms of the concentration of cyberciminal abuse. The explanatory factors can be divided into: *i)* generic structural properties of TLDs such as the number of domains in registry or the business model of underlying players and *ii)* properties directly related to security efforts. For example, the concentrations of domain abuse are, to a large extent, a function of the size of the TLD. In other words, the number of domains in a TLD registry can be seen as an approximation of the attack 'surface size' for cybercriminals. The business type of registries or hosting providers is also relevant. For example, the pricing strategy (e.g., free registration and free basic hosting program) combined with a lack of a verification process of the registrant's identity may decrease the unit cost of domain abuse and, as our results suggest, increase immensely the number of domains registered explicitly for malicious purposes.

We propose security metrics to measure and benchmark entire TLDs against the rest of the market and try to distill the relevant factors that may influence domain abuse rates. We develop two types of security metrics for TLDs, using representative phishing and malware domain abuse feeds: *i) occurrence of abuse*, which reflects concentrations of abused domains and depends on both structural properties of TLDs and preventative security practices of their intermediaries, and *ii) persistence of abuse*, which reflects how fast the TLD intermediaries respond once they are notified about abused domains and depends on their reactive security practices. In conjunction, the proposed metrics provide a better understanding of the overall health of a TLD. We view the proposed metrics as health indicators of the 'domain name ecosystem' of individual TLDs. We explicitly distinguish these metrics from the idea of measuring security *performance*, because the measured values are driven by a range of factors, not just by the performance of a particular intermediary.

We also propose a statistical regression model to analyze how the different properties of TLD intermediaries relate to the amount of abuse in their networks. In other words: What features explain the abuse concentrations observed in domain names registered under different TLD?

Overall, our paper makes the following contributions:

1. http://namecoin.info/

- We generate three TLD occurrence metrics measuring the relative 'amount of badness' in a particular TLD in Section 4.1. All reputation metrics are normalized by size which we carefully measure in Section 3.2.1. We find disproportionately low attacker interest in new gTLDs and disproportionately high attacker interest in large gTLDs.
- In addition to the occurrence rate of abused *domains*, we find that the two complementary metrics are useful; the rate of abused *fully-qualified domain names* and *URLs*. These reveal that the profit-maximizing behavior of some attackers lead them to abuse services with good reputation and free domain registration services.
- Next to the occurrence rate of domain abuse, in Section 5 we study the *mean* and *median* uptimes of abused domains via survival analysis. We find that uptime security metrics of certain TLDs are skewed by single phishing incidents.
- In Section 5.3, we compare the median uptime metric of phishing websites and the number of blacklisted domains by the Anti-Phishing Working Group but we do not find a strong correlation as each captures a different aspect of security practices.
- We systematically analyze how different structural and security-related properties of TLD operators influence abuse counts in Section 6; we find that next to TLD size, abuse primarily correlates with domain pricing (free versus paid registrations), efforts of intermediaries (measured through the proxy of their DNSSEC deployment rate), and strict registration policies.

2. Background and Related Work

Domain names are one of the fundamental components of the Internet ecosystem and central to the broad range of Internet scams that seek to attract user traffic to particular websites [1], [2]. Several entities play a role for a domain name to be registered, secured and maintained on the Web. In this section we provide a brief overview of the domain name ecosystem and the various actors that are involved at its different layers.

2.1. Domain Name Ecosystem

The Internet Corporation for Assigned Names and Numbers (ICANN) [3] approves all top-level domains (TLDs) and delegates the responsibility to a particular organization ("registry operators","sponsors" or "delegees") to maintain an authoritative source for registered domain names within a TLD [4]. Domain *registries* manage the registration of domain names within their TLDs and generate zone files that list domain names with their authoritative name servers and registration date.

TLDs can be categorized into three main groups [5]: *i*) country code TLDs (ccTLDS) or TLDs with two letters established for more than 250 countries and country codes. The represented region sets the policies for such TLDs, *ii*) generic TLDs (gTLDs) or TLDs with three or more characters such as `.com` or `.amsterdam`, and *iii*) `.arpa` – a special TLD that is used for technical infrastructure.

Several other entities play a role for a domain name to be registered, secured and maintained on the Web. Domain registrars manage the registration of Internet domain names. Registrars must be accredited by a gTLD registry and/or a ccTLD registry according to ICANN guidelines. Web hosting providers host servers which domain names point to. Finally, Domain Name Aystem (DNS) providers operate DNS servers that map domain and host names to their corresponding Internet Protocol (IP) addresses. Autonomous Aystems (AS) route sets of IP addresses.

Many of these actors play multiple roles. For instance, the `.tk` registry also registers `.tk` domains. Similarly, many large market players such as `GoDaddy` offer registration, web hosting, and DNS services simultaneously.

2.2. Security Metrics

To mitigate domain name abuse more effectively different classes of intermediaries such as registries, registrars or hosting providers should be able to benchmark themselves against their market. Currently, there exists very little empirical information about the security of TLDs and the overall DNS ecosystem. However, a number of studies include security metrics as a part of their analysis.

Levchenko *et al.* find some registrars, ASes and banks which are disproportionately popular among criminals, possibly due to their security practices [6]. Moore and Edelman find a concentration of typosquatted domains on a small number of name servers [7]. Korczyński *et al.* illuminate the problem of non-secure DNS dynamic updates, which allow a cybercriminal to manipulate DNS entries in the zone files of authoritative name servers. They find that 66.2% of vulnerable domains are hosted on the infrastructure of a single broadband Internet Service Provider (ISP). Reconfiguring zone files of just 10 providers would reduce the prevalence of the problem with 88.6% [8]. Ma *et al.* use name server and registrar information to distinguish malicious URls from benign ones [9]. Hao *et al.* observe that 46% of the spam domains come from just two registrars [10]. However, they consider only the `.com` TLD and do not consider the size estimate for smaller registrars which might register a disproportionate amount of malicious domains. Antonakakis *et al.* develop a dynamic reputation system using passive DNS data to classify legitimate and malicious domains and assign a reputation score to the new domains [1]. Our work does not rank individual domains but rather designs reputation metrics for the TLDs.

Numerous studies attribute security incidents to hosting providers by equating them with ASes. The number of incidents is often normalized by the AS size [11], [12]. Mahjoub investigates the concentration of maliciousness in ASes by analyzing hosted content, AS topology and IP space reservation [13]. Noorozian *et al.* present a systematic approach for metrics development and identify the main challenges that plague metric design [12]. In the process,

TABLE 1: Blacklists statistics: unique domains, FQDNs, URLs, or IP addresses for the StopBadware SDP and APWG datasets from 2014 and 2015.

Year	Dataset	# domains	# FQDNs	# URLs	#IPs
2014	StopBadware	694198	1135318	2329437	15127
	APWG	102868	1001626	10473941	3496
2015	StopBadware	701400	997350	8931660	13878
	APWG	233401	1954315	34356755	7180

they answer an urgent question posed to them by the Dutch police: "Which are the worst hosting providers under our jurisdiction?". Other studies identify malicious ASes using AS topology, BGP-related features and by exploring ASes providing transit for malicious ASes [14]–[16].

In the most similar study to our paper, van Eeten *et al.* design security metrics for ISPs by mapping data on the location of infected machines onto broadband markets [17]. They find that the bulk of infected machines are located in well-known, legitimate ISPs in industrialized countries; around 50% of the world's infection comes from only 50 providers. Even providers with similar sizes have differences in infection rates of up to three orders of magnitude, even within a single country such as Germany or the U.S., where the providers operate under similar market and regulatory incentives.

We propose security metrics to measure and benchmark entire TLDs against their market. We consider a range of factors such as TLD size or popularity of sites and support more informative inferences from domain abuse data for security. We explicitly distinguish the metrics from the idea of measuring security performance because the measured values of our metrics are driven by multiple factors, not just by the performance of the particular market player. This is because a TLD is not a single organization but constitutes an entire "domain name ecosystem" of different types of players that all influence the TLD's security posture.

3. Data Collection

3.1. Abuse Data

3.1.1. Blacklists. To asses the prevalence of maliciously registered and compromised domains per TLD, we use two heterogeneous blacklists provided by StopBadware [18] and the Anti-Phishing Working Group (APWG) [19]. The StopBadware Data Sharing feed consists of blacklists shared by ESET, Fortinet, and Sophos security companies [20]–[22], Google's Safe Browsing appeals results, Internet Identity, Malware Must Die and the StopBadware community. The APWG feed consists of online phishing URL block/white lists with accompanying confidence level indicators submitted by accredited users through the eCrime Exchange (eCX) platform.

Table 1 shows the number of unique domain names, fully-qualified domain names (FQDN), URLs, or unique IP addresses (if domains were not reported) in these data

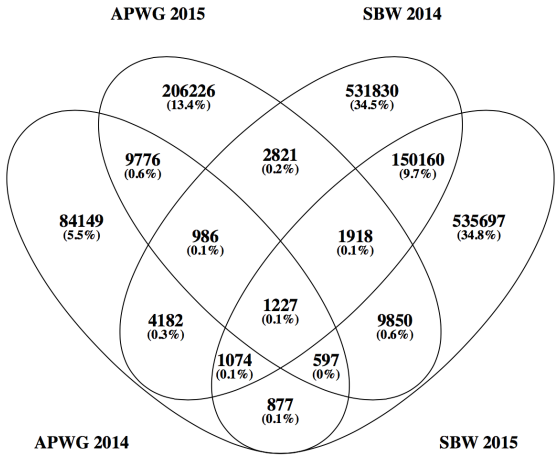

Figure 1: Venn diagram of blacklisted domains in two different datasets.

feeds for 2014 and 2015. Notice that we define domain names as 2nd–level or 3rd–level if a given TLD registry provides such registrations, e.g. `*.com.pl`, `*.net.pl`, `*.gov.pl`, etc. To extract domain names from our feeds, we use a modified version of the public suffix list maintained by Mozilla [23], which for example considers `com.pl` and `net.pl` to be TLDs. For the purpose of this study we have excluded all private TLDs such as `s3.amazonaws.com` or `blogspot.com`. Figure 1 shows the overlap between the two analyzed data feeds in 2014 and 2015 in the form of a venn diagram. In 2015, for example, the StopBadware dataset contains 535,697 domains that were not identified as malicious by either APWG in 2014 and 2015 or StopBadware in 2014. It corresponds to 34.8% of all domains blacklisted by both organization in both periods.

3.1.2. Uptimes. We analyze the uptime of phishing websites with data that was generously provided to us by the Cyscon GmbH security company [24]. The dataset contains the following information about phishing websites: *i)* the 'first time seen' defined as the moment the responsible intermediary has been notified about an abused website, *ii)* the 'last time' the website is seen online and *iii)* a binary variable that indicates whether it has ever been taken offline. The dataset contains 137,577 phishing URLs associated with 48,224 FQDNs. Note that for the websites that are only seen once, the first time seen value is equal to the last time seen, indicating that they were taken down before the second measurement. These are logged as having an uptime of 0 hours. The data contains phishing websites that were collected between June 2015 and January 2016.

3.2. TLD Properties

To analyze the differences in abuse incident counts between TLDs, we collect a number of structural and security effort related properties for the entire population of TLDs. In the upcoming sections of the paper we use these properties

to model phishing abuse counts and further explain the variation of these counts among TLDs.

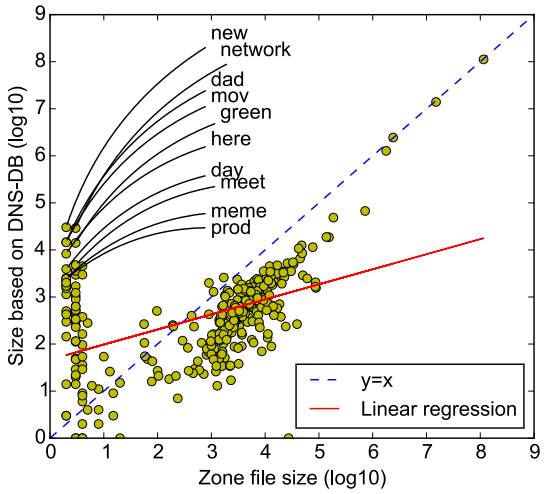

Figure 2: Comparison of TLD size: DNSDB size estimates *vs.* zone file size.

3.2.1. Number of Unique Domains. To obtain a meaningful, quantitative security metric, representing the distribution of blacklisted domains per TLD, we first need to estimate their sizes. The obtained sizes can be used as a normalization factor for the amount of '*badness*' in each TLD. Once normalized, TLDs can be compared in terms of the prevalence of blacklisted domains.

We calculate the size of each TLD by counting the number of 2nd–level domains registered in that TLD. In some cases, however, it is also possible to register 3rd–level domains within a TLD registry. For example, the `.cn` TLD (China) allows domain registration under `com.cn`, `net.cn`, etc. For such cases, the size of the TLD includes 3rd–level domains.

Zone files are the most accurate source for TLD size. We obtain the `.com`, `.net` and `.name` zone files from Verisign [25], [26]. We also perform zone transfers to replicate the DNS databases of the `.nl` zone file under the contract of Stichting Internet Domeinregistratie Nederland (SIDN)—the `.nl` ccTLD registry [27]. Finally, we collect zone files from the `.us` ccTLD, the `.biz` gTLD and all new gTLDs made available by ICANN [28]. Note, however, that not all TLD zone files are openly available. In fact, most of the ccTLD registries are against making their zone files available to third parties. For this reason, we have collected further information on sizes of ccTLDs affiliated with the Council of European National Top Level Domain Registries (CENTR) association [29] from their monthly growth reports and sizes reported in the APWG Global Phishing Survey in early 2015 [30] in addition to size information from Domain Tools published in January 2016 [31]. We have crosschecked the collected size information with the number of DNS `A` records passively observed over the year 2014 in DNSDB – a passive-DNS dataset generously provided to us by Farsight Security [32], [33] – which we use to estimate remaining TLD sizes.

Ordered by accuracy, our datasets for calculating TLD size are: *i*) zone files, *ii*) APWG or DomainTools size information and *iii*) DNSDB.

Figures 2 and 3 provide a comparison between size estimates we calculate based on the DNSDB and other data sources. In Figure 2 we observe an interesting cluster of TLDs for which sizes were largely overestimated based on the DNSDB data. These are due to name collisions in the DNS system related to the introduction of new gTLDs by ICANN. For example, the zone file for the `.new` TLD consist of only 2 domains of which one has been observed in DNSDB. In addition to the one legitimate `.new` domain, we have observed $30,311$ non-unique 2nd–level domains out of which $30,283$ resolved to the special ICANN's IP addresses '127.0.53.53' indicating a name collision occurrence and raising an alert of a potential issue [34].

Figure 3 compares the size estimates for DNSDB with APWG TLD sizes. Apart from some outlying examples such as `.ph` (for which APWG has also estimated its size), the size information of both datasets are relatively consistent.

The presented figures demonstrate that, except for the newly introduced gTLDs, the three TLD size estimates can reliably and interchangeably be used for TLD size calculation. Notice, however, that overestimations of new gTLDs based on the DNSDB data can be easily corrected by filtering out all the passively observed DNS records resolving to a special ICANN's 127.0.53.53 IP address.

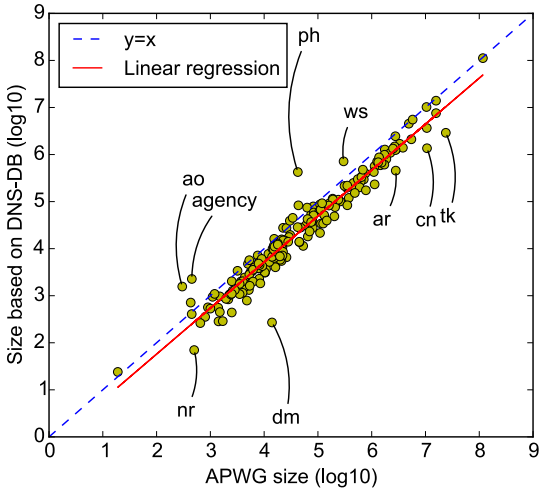

Figure 3: Comparison of TLD size: DNSDB size estimates *vs.* APWG size.

3.2.2. Number of Unique Domains on Shared IPs. We consider an IP address shared, if it hosts more than 10 domain names [35]. We expect that the number of domains on shared IPs to correlate positively with the domain abuse counts of TLDs due to commonly known vulnerabilities of shared hosting services [36] assuming that attackers would

compromise rather than register phishing or malware domains. This variable combined with the number of unique domains in each registry conveys information about the degree to which the business model of hosting providers relies on low-cost shared hosting services. It is calculated based on A records passively observed in DNSDB during 2015.

3.2.3. Domain Popularity Index. To asses the online popularity of domains we use the Alexa top-ranked one million domains [37]. Accredited domain registrars that often provide hosting services are assumed to ensure better security measures if they host popular websites. We calculate the domain popularity index by summing up reversed Alexa ranks for the 2nd–level domains and aggregate them per TLD. The most popular domain gets the rank 1. Its score is calculated as $6-log_{10}(rank)$. We calculate a score per TLD by summing up individual scores of all ranked domains.

3.2.4. URL Shortener Popularity Index. A variety of studies suggest that miscreants tend to misuse legitimate services such as cloud-based file sharing, free hosting and URL shorteners. They are abused for phishing campaigns or drive-by download campaigns [38], [39]. We maintain a list of 332 domains of the most popular URL shortener services such as goo.gl, bit.ly, etc. and calculate scores similar to the domain popularity index. The obtained variable indicates to which extend the level of domain abuse can be explained by the exploitation of legitimate services by the attackers.

3.2.5. Free. Finding reliable pricing information for TLDs is complex mainly because of the many actors involved in domain registration. For example, when someone registers a \$10 .com domain, ICANN charges only \$0.18, the registry fee charged by Verisign, Inc. is fixed and equal to \$7.85. Finally, registrar may charge \$1.97. Price may vary depending on TLDs, registrars or resellers involved in domain registration. Some web hosting companies offer registration discounts if it is combined with the hosting service, etc.

Previous studies show that miscreants benefit from cheap or free domain registrations via promotions [30]. However, as it is difficult to collect pricing data at scale, we divide TLDs into two groups: those that offer free registrations and those that offer only paid ones. The first group is composed of five TLDs only: .tk, .ml, .ga, .cf and more recently, starting from January 1, 2015, .gq TLD, all operated by Freenom [40].

3.2.6. Restricted Registration. We label TLDs as "restricted" or "unrestricted" depending on specific registration limitations. Some blacklisted domains are registered maliciously, by cybercriminals, rather than hacked. Although they are mainly due to cheap and free domain name registrations, certain restrictions regarding the content of the website or the eligibility of registrants may increase the cost of an online crime and discourage cybercriminals from registering domains with certain TLDs. "Restricted" TLDs consist of:

- Sponsored gTLDs such as .post operated by Universal Postal Union or .gov restricted to government entities in the United States,
- Brand gTLDs like .bmw, .youtube or .cern,
- Community gTLDs such as .abogado restricted to licensed lawyers or .pharmacy for which all the registrants are verified if they meet all applicable regulatory standards [41],
- Geographic gTLDs such as .berlin, .cat or .asia, and certain ccTLDs such as .sk restricted to Slovak companies, organisations and citizens.

3.2.7. Age. Age might be a proxy for how experienced a TLD registry is. However, for example, the founders of new gTLDs such as Donuts Inc. have years of experience in the domain name industry [42]. We would then expect certain new gTLDs to attract less abuse as it is easier to deploy certain security measures on a smaller scale. The variable is expressed in years and it is calculated since a given TLD was introduced. For example, .com TLD has been operational for 31 years.

3.2.8. DNSSEC Deployment Rate. Another important factor that influences the abuse rate of TLDs is the security effort of the registries. As we have no direct observations of these efforts we instead actively measure the DNS Security Extensions (DNSSEC) protocol [43], [44] deployment rate. Although DNSSEC has potential for abuse in Distributed Denial of Service attacks, it undoubtedly strengthens DNS by preventing a number of attacks such as cache poisoning [45].

It is often very challenging to secure hundreds of thousands or even millions of existing domains in the environment composed of multiple stakeholders. To increase the security of .nl domains, SIDN for example, actively encourages and incentivises registrars of domains to deploy DNSSEC by reducing the registration price. For that reason, we use the DNSSEC deployment rate as a proxy for the security effort of the TLD registries.

To actively measure the current deployment of DNSSEC per TLDs, we use zonemaster collector implemented by Patric Wallström [46]. We further sample approximately 2.4 million domains from DNSDB in May 2016 using the following procedure. To determine a representative sample size for each TLD, we use binomial approximation [47]. We first set a narrow total width of confidence interval to $W = 0.01$ with confidence level $CL = 95\%$. Based on our preliminary measurements, we approximate P defined as an expected proportion of DNSSEC-signed domains to all registered domains per TLD. Sample size s is calculated using the following formula: $s = \frac{(Z^2 * P(1-P))}{(0.5*W)^2}$, where Z is the critical value of the normal distribution (for $CL = 95\%$ critical value is equal to 1.96). For example, assuming that for a given TLD the expected DNSSEC deployment rate is 5%, the sample size is equal to 7,299.

To assess the accuracy of our active measurements, we compare the DNSSEC deployment rates of our measurements with the actual ground-truth data obtained from the

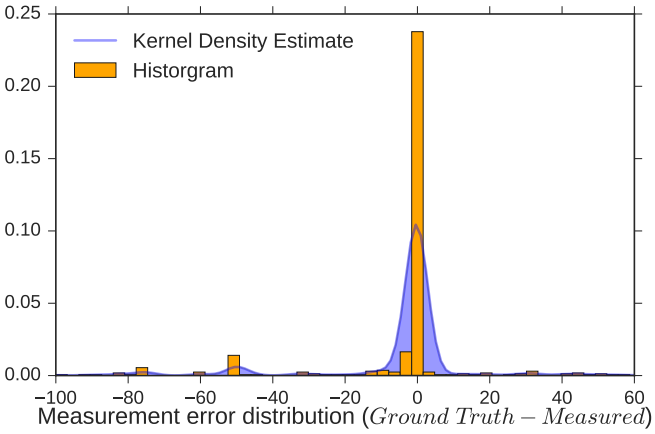

Figure 4: Distribution of differences of the measured DNSSEC deployment rates with the actual deployment rates (in %) based on available zone files.

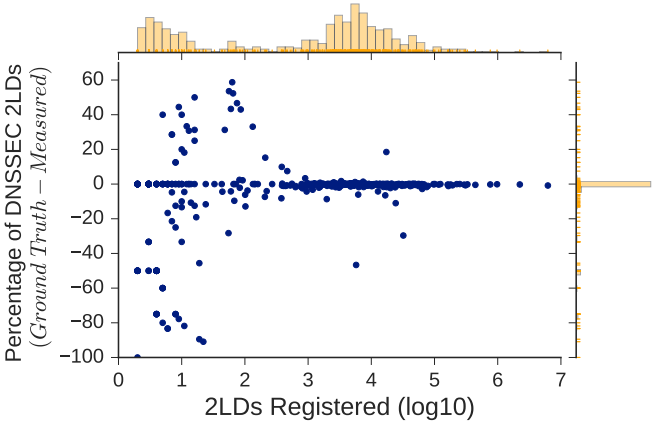

Figure 5: Differences of the measured DNSSEC deployment rates with the actual deployment rates (in %) based on available zone files as a function of TLD sizes.

zone files of 590 TLDs available to us at the time of the measurements. Figure 4 shows the distribution of differences (in %) between the ground-truth data and the measured deployment rates. Although the results indicate a high accuracy of our method, we observe a considerable number of TLDs for which the actual deployment rate does not match with the measured one. We further investigate the differences as a function of the number of 2nd–level domains in a zone file (see Figure 5). Note that our measurements are very reliable for larger TLDs whereas smaller ones are subject to considerable measurement error. To overcome this limitation we use the actual DNSSEC deployment rates when ground-truth data is available and only otherwise use the measured deployment rates in our statistical analyses.

For example, Figure 6 shows the top 20 ccTLDs with the highest DNSSEC deployment rates. Although the `.nl` ccTLD has the largest absolute number of domains with the DNSSEC validation support, in relative terms `.no` and `.se` ccTLDs have higher deployment. Notably, the large gTLDs

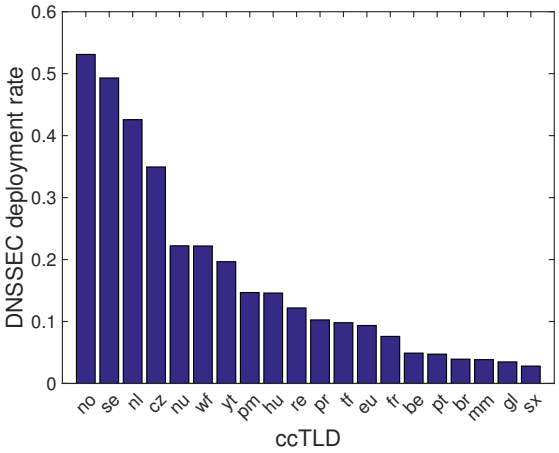

Figure 6: Top 20 ccTLDs with the highest DNSSEC deployment rates.

such as `.com`, `.net` or `.org` have less than 1% DNSSEC-signed domains in registry.

3.2.9. ICT Development Index. The International Telecommunication Union (ITU) provides country-level quantitative indicators that show the development in information and communication technology (ICT) of different countries [48]. Using the delegation details of the top-level domains extracted from IANA root zone database we map each TLD to ICT development index of the country where the TLD operates.

4. Prevalence of Blacklisted Domains by TLD

In this section, we first present the three security metrics that provide insight into the distribution of badness across TLDs and we describe the intuition behind their selection.

4.1. Occurrence Security Metrics

First, we propose to analyze the occurrence of unique *domains*. Although, it is the most intuitive metric, it also has its limitations. It does not give an indication of the '*amount of badness*' coming from a given domain name. For example, modern botnets extensively adopt domain generation algorithms (DGA) to create a large number of domain names and then use their subset for rendezvous points between compromised machines and command-and-control servers (e.g. `123.malicious.com`, `234.malicious.com`, etc.) [49]. Moreover, a single maliciously registered domain name may be used in several phishing campaigns against, for example, different banks [30]. In terms of the number of unique domains, the dynamic reputation system will assign the same reputation score to both TLD registries. To overcome this limitation, we propose a second, complementary metric: the number of unique

FQDNs. We encounter, however, some limitations of the second approach as well. A single FQDN of a compromised website could be used, for example, to distribute malware configuration and binary files or serve as dropzones, etc. using distinctive paths (e.g. `malicious.com/wp-content/file.php`, `malicious.com/wp-content/gate.php`, etc.) [50]. This is why we propose a third, complementary abuse occurrence metric: unique blacklisted *URLs* aggregated by TLDs. It reveals information that is not captured by other two metrics, namely the 'amount of badness' associated with unique FQDN. It stems from our previous work with the Dutch national police. Our analysis of child abuse URLs have revealed that some FQDNs are used more extensively to distribute malicious material. In fact, one FQDN can be used to share one abusive photo whereas another to distribute tens or hundreds of photos. The manual analysis of other types of abuse such as malware or phishing confirms this trend.

Reliable reputation metrics have to account for a commonly observed trend that larger market players such as broadband or hosting providers experience a larger amount of abuse [12]. For that reason, each of the proposed metrics are normalized by the size of the corresponding TLDs which we carefully measure in Section 3.2.1.

In the rest of this section we provide the reader with the qualitative and quantitative analysis of phishing and malware data for TLD reputations and make an attempt to distil a number driving factors for both large and small abuse concentrations.

4.1.1. Phishing TLD Reputation. Figures 7a and 7d show the number of unique phishing domains blacklisted in 2014 and 2015 per TLDs as a function of their sizes. Due to the disproportionate nature of TLD market share, we present the results using a logarithmic scale. It is important to note that we also use the log-transformed phishing domain counts in our statistical analysis. The drawback is that a logarithmic scale results in an undefined counts for those TLDs which have zero blacklisted domains. To overcome this limitation, we assign a small positive value close to zero to their phishing domain counts to be able to include such TLDs in our statistical analysis. As expected, the results indicate a moderately strong relationship between the size of a TLD and the number of blacklisted domains ($r^2 = 0.68$ in 2014 and $r^2 = 0.5$ in 2015). This is because in general the majority of blacklisted domains are hacked rather than registered by miscreants [51]. In fact, the number of domains in a TLD registry can be seen as an approximation of the attack 'surface size' for cybercriminals.

There is also a clear difference between ccTLDs (marked in blue) and gTLDs (marked in red). Notice that large gTLDs such as `.com`, `.net` or `.org` (see e.g. Figure 7a and Table 4) in absolute terms are of special interest to miscreants in comparison to 'new' gTLDs (those that have been delegated by ICANN in late 2013 and early 2014 [52]). We find a relatively small number of observed phishing attacks using the new gTLD domains in 2014. We hypothesize that these were not yet included in a criminal business model. This may be due to the higher prices of individual registrations [53] which directly influence malicious registrations rather than compromise frequency. Some new gTLDs might also be less attractive among miscreants because phishing domains with suffixes such as `.bike`, `.tips` or `.photos` may look more suspicious and unlikely to be legitimate bank or luxury brand websites. On the other hand, access to zone files of new gTLDs is open under a special agreement with ICANN [28]; this could be a powerful tool to search for vulnerable domains to compromize.

There are some TLDs, such as `.ml` or `.xyz`, that appear to deviate from other TLDs. By manual inspection of blacklisted `.ml` domains, we find a variety of seemingly unrelated domain names, presumably maliciously registered; `.ml`, `.ga`, `.cf` and `.gq` ccTLDs offer free domain registrations by Freenom [40]. Note that according to Aaron and Rasmussen, Freenom declined to provide the number of domain registrations for `.cf`, `.ml` and `.ga` TLDs for the APWG study in 2014 [51]; instead, the authors used DomainTools to estimate TLD size (which we compare with our own DNSDB estimations in Section 3.2.1).

Figures 7b-c and 7e-f show the two complementary occurrence security metrics: the number of unique fully-qualified phishing domains and URLs per TLDs in 2014 and 2015 (see also overall occurrence security metric scores in Tables 4, 5 and 6). These distributions are more dispersed than those of just domains ($r^2 = 0.6$ and $r^2 = 0.45$ in 2014 and $r^2 = 0.47$ and $r^2 = 0.44$ in 2015, respectively) and may give a better intuition on the 'amount of badness' from a TLD.

By manual inspection of FQDNs we find 17 2nd–level `.ru` domains presumably maliciously registered and 32,596 corresponding FQDNs (3rd–level domains) involved in the same phishing campaign in 2014. Similarly, we find three 2nd–level domains: {`incomparable`, `phenomenal`, `taipei`}`.country` and 8,996 corresponding 3rd–level domains used in phishing campaigns in March and July 2015. Thus the `.ru` and `.country` TLDs stand out in the FQDN reputation metric seen in Figures 7b and 7e, respectively.

We also manually analyze domain abuse hosted on other ccTLDs such as `.pn`, `.nf`, `.vu` or `.vc` as they deviate markedly from the other TLDs. Here, we find that most domain abuse comes from certain legitimate 2nd–level domains acquired by providers that offer free registration (and often free hosting) with selected domain extensions. For example, we find that the great majority of abused `.nf` FQDNs are hosted on `co.nf`, blacklisted `.vc` FQDNs are hosted on `zz.vc`, whereas `.vu` on `de.vu` and `co.vu` 2nd–level domains. As indicated in one of the free hosting provider's website such extension and thus a registered FQDN looks like a real paid domain name, e.g. the domain of United Kingdom (`co.uk`) or Austria (`co.at`). Moreover, while generally TLD registries apply certain measures to protect brands and prevent malicious registrations, it seems that hosting providers offering free domains do not. We observe, for example, maliciously registered domains

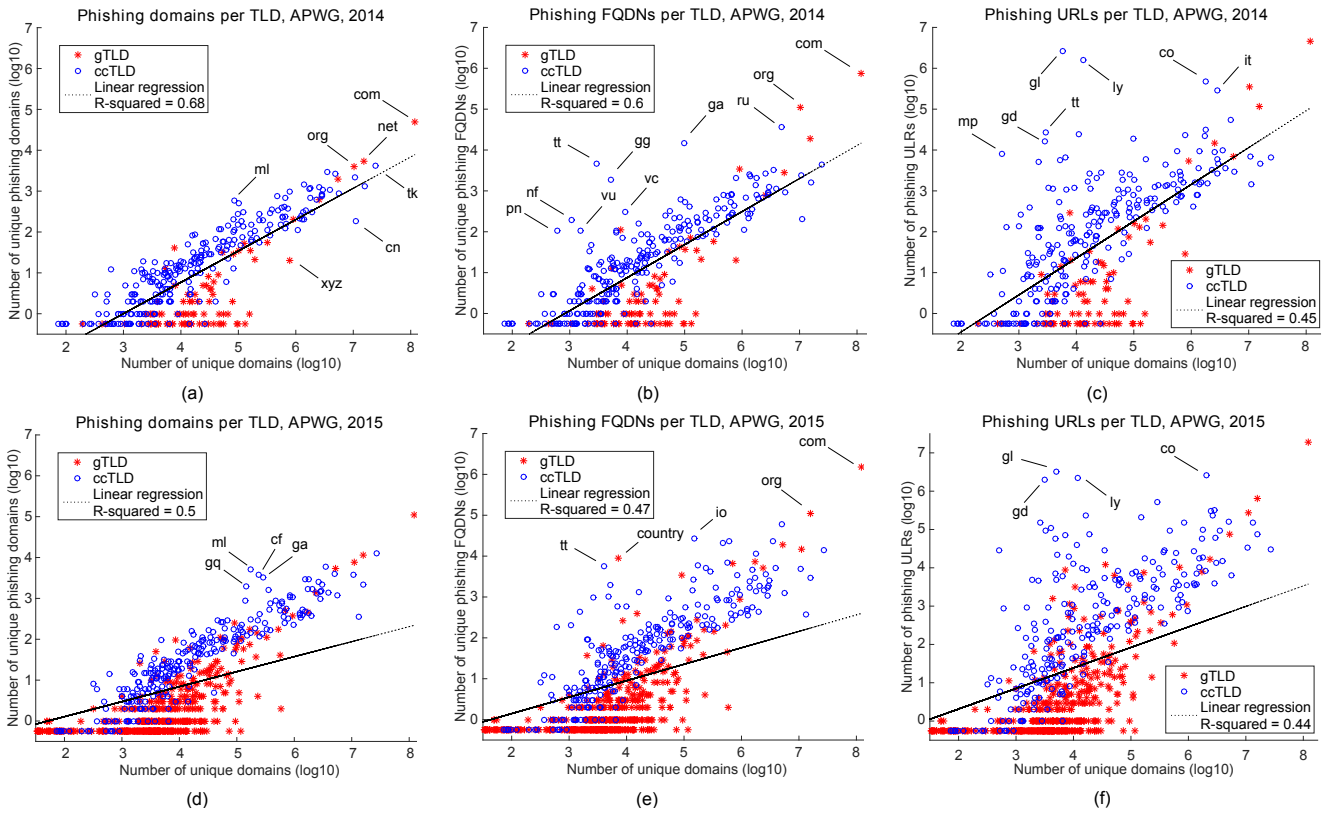

Figure 7: Phishing domains, FQDNs and URLs per gTLD and ccTLDs as a function of TLD size based on the Anti-Phishing Working Group feed from 2014 and 2015.

such as `faacebook.de.vu`, `bankofamerica.co.vu` or `wells-fargo.co.nf`. More importantly, we do not observe any indicators of improved practices over the two-year period of our study.

Through manual inspection of the `.com` domains we find, for example, 44,856 unique `*.s3.amazonaws.com` FQDNs that correspond to an online file storage web service offered by Amazon Web Services (AWS). This indicates that miscreants increasingly use cloud services in their malicious activity. As the majority of cloud services companies use the `.com` prefix, the `.com` TLD will presumably always be an outlier. Nevertheless, an increased number of reported FQDNs of certain services does not always mean that the companies do not properly deal with security incidents. Rather, since the eCrimeX APWG platform is used by known parties and registries to submit and act quickly to shutdown/suspend the phishing domains, the most active entities submitting phishing domains might look worse in terms of a number of reported FQDNs. A very limited number of, for example, `*.s3.amazonaws.com` FQDNs found in StopBadware DSP and in Phishtank [54] datasets seems to support this hypothesis.

We further find 6 unique blacklisted domains on the `.tt` ccTLD for the Republic of Trinidad and Tobago. Notice, that the number of corresponding FQDNs and URLs is significantly higher. First, we find that the content stored on `ge.tt` file sharing and publishing platform is being extensively used for malicious purposes. We count 4,565 unique `*.open.ge.tt` FQDNs and 3,373 unique `ge.tt/*` URLs. Moreover, we find 1,052 reported Dropbox short URLs under the `db.tt` domain. In fact, Dropbox added support for linking to particular files or folders using short URLs that other people are able to access directly from their browsers. We also encounter 88 unique URLs created with the `mf.tt` service to shorten any URLs. This is somehow to be expected that miscreants use file storage and URL shortening services.

The application most extensively used for malicious purposes within the `.tt` TLD is the "If This Then That" (IFTTT) service [55]. It combines more than 150 services like Facebook, Gmail, Dropbox to work together and to create chains of simple conditional actions, called "recipes". A simple example recipe might consist of sending an e-mail message if the user is tagged by someone on Facebook. We encounter 17,787 unique `ift.tt` URLs in the APWG dataset which shows the popularity of this service among miscreants.

For completeness, the toxicity of other TLDs is also heavily affected by popular URL shortening services, namely `t.co` operated by Twitter (377,726), `x.co` by GoDaddy (103,255), `ow.ly` (1,213,302), `bit.ly` (374,283), `adf.ly` (93,994) and `goo.gl` (2,677,239) by Google.

4.1.2. Malware TLD Reputation. We now study the malware activity reported by the StopBadware DSP. We refer the reader to Tables 7, 9 and 8 for overall occurrence security metric scores. As before, we observe a moderately strong relationship between the TLD size and the number of blacklisted domains ($r^2 = 0.68$ in 2014, $r^2 = 0.53$ in 2015), FQDNs ($r^2 = 0.63$ in 2014, $r^2 = 0.52$ in 2015) and URLs ($r^2 = 0.6$ in 2014, $r^2 = 0.52$ in 2015) . We again observe a clear difference between ccTLDs and new gTLDs which have less attacker interest.

We notice very different results for TLDs offered by Freenom. Namely, the `.tk` TLD experiences very little malicious registrations, whereas the `.ml` and `.ga` TLDs experience many more relative to their size. In fact, we observe a similar absolute number of blacklisted malware domains for all the three TLDs even though the sizes of the former two are few hundred times smaller than the `.tk` TLD size. Notice that Freenom offers some accredited interveners, such as Facebook, the Anti-Phishing Alliance of China or Internet Identity the ability to suspend malicious domains in the registry. According to Aaron and Rasmussen, at the time Freenom did not provide a similar tool to alleviate the problem of blacklisted domains registered on `.tk` on the other TLDs, i.e. `.ml`, `.ga` and `.cf` [51]. This could explain smaller occurrence scores for `.tk` domains. However, a chief executive officer of Freenom Joost Zuurbier shared with us that the accredited interveners have the ability to suspend domains of all the TLDs operated by their company. Thus, another possible explanation for different relative abuse counts could be that the `.tk` TLD contains many more expired and parked domains (by Freenom itself) than the other four TLDs. This leads to an overestimation of its size, leading to lower metrics on the concentration of abuse. To reliably explain the variance in abuse data between certain TLDs, we need to consider other TLD properties than just size such as domain registration price or other TLD-specific registration policies.

By manual inspection of the `.edu` TLD (one of few gTLD outliers, see Table 8), we do not find any significant difference in the amount of malware domains in different educational institutes; this contrasts with the APWG feed where 65 out of 112 bad `.edu` domains in 2014 come from the Harvard University Library (`*.hul.harvard.edu`).

Finally, we observe a relatively large volume of `*.tt` FQDNs and URLs corresponding to only 8 unique blacklisted domains in 2014. Most of them are subdomains of the `ge.tt` domain—a file sharing platform, as in the APWG feed. More specifically, we observe 1,231 `*.open.ge.tt` FQDNs, 1,234 `*.blob#.ge.tt` FQDNs, and 1,098 unique `ge.tt/*` URLs out of all 3,433 unique FQDNs. (Note that * refers to a string, whereas # represents a digit). We explore the common part of these domains blacklisted by APWG and StopBadware, and do not find any common domain.

5. Persistence of Phishing Domains in TLDs

In this section, we present two security metrics that reflect how fast the TLD intermediaries respond once a phishing domain under their TLD is abused.

5.1. Uptime Security Metrics

Another important health indicator of a TLD is the persistence of domain abuse or *uptime*. It indicates how fast the responsible intermediaries respond once the domain is abused or, in other words, how long a malicious domain stays online before the 'take-down' procedure, i.e. domain suspension or quarantine. The uptime of malicious domains is used in other security research as a standard metric for studying the lifetime of individual domain names in relation to various attack types and affected hosting providers [35], [56]–[58].

In this paper, we study various security metrics related to the uptime of phishing domains per TLD. Our first metric, proposes the use of the *mean uptime* of unique domains per TLD. While it is the most intuitive metric, it comes with its limitation, that is, it may be skewed by long-lived malicious domains. To overcome this limitation, we further consider the use of *median uptime* of abused domains. Regardless of the specific type of abuse, we expect the median to be considerably less than the mean uptime.

TABLE 2: Descriptive statistics of uptimes (in hours) for selected TLDs based on phishing websites reported by Cyscon.

TLD	Min	Mean	Median	Max	SD	SE
BR	0	199.81	24.09	4,586.27	506.75	6.36
CL	0	174.76	27.47	4,119.50	448.41	10
COM	0	182.96	24	4,858.38	494.60	1.50
CZ	0	164.60	0	2,859.14	493.27	42.61
ML	0	165.32	29.50	3,283.28	408.84	21.98
NET	0	181.83	24	4,635.54	515.99	5.87
NL	0	196.08	0	3,997.50	603.15	30.35
PL	0	199.94	0.61	4,714.04	595.17	15.15
TK	0	281.42	37.37	3,866.72	693.96	36.42
TR	0	58.43	0	4,679.64	343.39	8.38
FJ	3,418	3,418	3,418	3,418	–	–
RW	0	2,272	2,969	3,848	2,016	1,164
DANCE	2,376	2,376	2,376	2,376	–	–

5.2. Uptime of TLDs

To calculate uptime, we use the phishing dataset provided to us by Cyscon security company as explained in Section 3.1.2. We define uptime of a phishing website as the number of days between the first and last time the website is observed online as reported by Cyscon. Table 2 contains descriptive statistics of the uptime of the phishing websites that had been taken down by the end of our measurement period in a set of selected TLDs (see Table 10 for top 20 TLDs with the highest median uptime). As expected, the median uptime is often much less than the mean uptime. We find that, for example, phishing websites of the `.pl` and `.nl` TLDs remained available on average

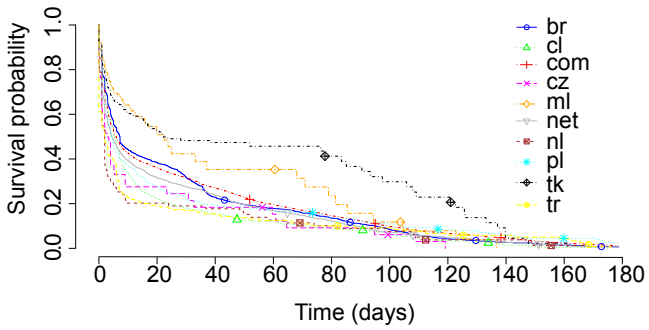

Figure 8: Survival analysis of phishing websites of ten selected TLDs based on data provided by Cyscon.

for more than 199 and 196 hours respectively, with a very small median or median equal to zero. It means that the majority of blacklisted `.nl` websites were cleaned as soon as the responsible providers were notified, even before the subsequent measurement. Notably, the mean and median uptime of certain TLDs are skewed by single abuse incidents, especially those TLDs with the highest uptimes (see Table 2 and 10). For example, through manual inspection of the only blacklisted `.fj` domain, we find that between June and November 2015, it was used in two phishing campaigns: one against Yahoo and another one tagged as "generic auto". One of the blacklisted URLs (`wp-includes/js/crop/mene/Indexz.html`) indicates that the miscreant might have exploited a vulnerable WordPress installation (or one of its plugins) rather that having registered a domain for malicious purposes.

In addition, a number of phishing websites may remain online when our measurement period is finished, which leaves their uptime unknown. In order to correctly account for such instances, we use *Survival Analysis* which allows for taking such artifacts into account through right-censoring of the data points. Let survival function *S(t)* be the probability that a phishing website is online at time *t* during the measurement period. *S(t)* is calculated using a standard Kaplan-Meier estimator [59].

Figure 8 displays the survival curves for phishing websites of 10 selected TLDs. The shapes of the curves differ considerably for the selected TLDs, which further indicate their differences in taking down phishing websites. The two most deviant TLDs are `.tk` and `.ml`, which are far from others in terms of the rate with which the phishing websites are taken down: after approximately 70 days 45% and 35% of their phishing websites remain online, respectively.

5.3. Occurrence Versus Persistance

Figure 9 compares the median uptime of phishing websites reported by Cyscon between June 2015 and January 2016 and the number of phishing domains blacklisted by APWG in 2015. Similar to before, we assign a small positive value close to zero (0.01), to TLDs with no domains blacklisted by APWG. These values are transformed to a negative value on the logaritmic scale. As the figure indicates, there is a weak relationship ($Pearson\ r = 0.27$) between the phishing occurrence and uptime metrics which was expected as each captures a different aspect of security practices (i.e. proactive vs reactive efforts). While providers within certain TLD ecosystems are more *reactive* to abuse notifications, they might not necessarily invest resources in *proactive* security measures such as vulnerability and patch management. Clearly, some TLD ecosystems experience large amounts of abuse but also manage to quickly take their malicious domains down (upper left region of the plot). Interestingly, the results seem to suggest that there are no TLDs that "perform" consistently bad: experience large amount of abuse and be slow in taking them down (upper right region of the plot is empty).

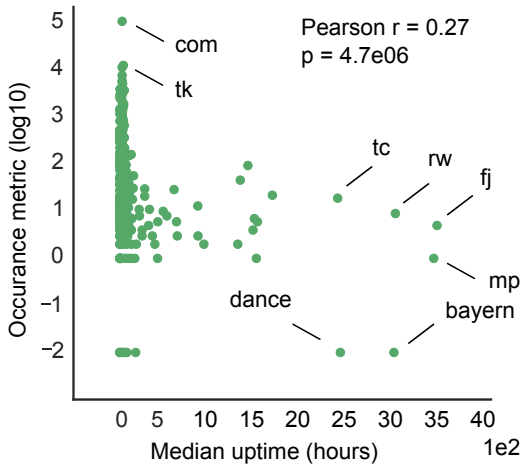

Figure 9: Comparison of phishing occurrence metric based on domains reported by APWG and median uptime metric based on phishing websites reported by Cyscon.

6. Regression Analysis of Abuse in TLDs

In Section 4, we demonstrated that domain abuse counts relate to the size of TLDs with relatively high $Pearson\ r$ value. However, size estimates come with their own measurement errors and biases (i.e. under/over counts of domain names registered within a TLD). In addition to the size of a TLD, there might exist other properties that drive the abuse counts. Using a regression model, one can examine the amount of variance that those properties can collectively explain, from the total observed variance in the abuse counts.

Previous research has proposed an approach to study the variation of abuse across the population of hosting providers, using regression models that carefully decompose different drivers of variation in abuse counts, such as size, pricing, time-in-business and the amount of WordPress sites on phishing abuse [60]. Here, we use a similar statistical regression approach, to analyze how the different properties of TLD intermediaries relate to the amount of abuse in their networks.

We model the number of abused domains as a dependent variable using Negative Binomial Generalized Linear Model (GLM) regression with a Log link function. The specific choice of model is due to the dependent variable constituting a 'count' and due to the interdependence of events and non-equal mean and variance of the dependent variable. Our regression model is built using the APWG dataset from 2015, as explained in Section 3.1.1. We define our dependent variable Y_i as the number of blacklisted domains in APWG for $i = 1, \ldots, n$, with n being the total number of TLDs where Y_i follows a Negative Binomial distribution. Our independent variables in the model are TLD properties introduced in Section 3.2. The selected variables are the 'Number of unique domains', 'Number of unique domains on shared IPs', 'Free', 'Domain popularity index', 'URL shortner popularity index', 'Age', 'DNSSEC deployment rate', 'Restricted registration' and 'ICT development index'. In order to prevent dependence in the residuals due to presence of 'ICT development index' as a country-level predictor in the model, we add the 'country' variable as a fixed-effect, by fitting a separate dummy variable as a predictor for each class of the variable.

Our Negative Binomial GLM model with the country fixed-effects is defined as follows:

$$\ln(Y_i) = \beta_i x_i + \ldots + \delta_i, \tag{1}$$

where i refers to different measurements within each TLD, β_i are the estimated coefficients for x_i, x_i are the independent variables collected for TLDs and δ_i is the "fixed-effect" parameter [61].

Table 3 shows the summary of the regression results with the estimated coefficients and their significance levels produced by the regression analysis using this model. One should note that the final model in Table 3 is chosen from a stepwise addition of the variables into a baseline model with a single explanatory variable. Table 3 illustrates the model with the maximum Log likelihood, θ values and minimum Akaike information criterion (AIC) value.

The results indicate that in addition to a TLD size, the other independent variables contribute significantly to the variation in the number of phishing domains, as the dependent variable. The size of the coefficients display the extent to which they contribute.

As expected, the 'Number of unique domains' variable (TLD size) explains a significant part of the variance in phishing counts. Holding the other independent variables constant, if we increase the 'Number of unique domains' by one unit (equivalent to multiplying the size of a TLD by 10 since it is in the log_{10} scale), the number of phishing domains is multiplied by $e^{(1.416)} = 4.120$. The 'Number of unique domains on shared IPs' however, plays a very week significant role in explaining the variance in phishing counts and can be interpret in a similar fashion to the 'Number of unique domains'.

As hypothesized, we see a very strong significant positive relation between the price of domain registration and phishing counts in the model. As defined in Section 3.2.5, the variable 'Free' is equal to 1 if the TLD offers free domain registrations and 0 if not. Hence, the results indicate that free registration within TLDs multiplies the number of phishing domains by $e^{(5.804)} = 331.623$, while holding all other variables constant. While the number of domains in the registry is related to the number of hacked websites and indicates the 'attack surface' for cybercriminals, free domain registration is another means by which a miscreant can gain access to a domain.

The variables 'Domain popularity index' and 'URL shortner popularity index' also have significant effects. While the resulting coefficient for the 'Domain popularity index' suggests that more popular TLDs are attacked less, the 'URL shortner popularity index' coefficient indicates that more popular URL shortners are used more in phishing attacks, which is inline with our own observation explained earlier in this paper. Naturally, we would expect a much stronger correlation between the 'URL shortner popularity index' and an alternative abuse occurrence metric such as the number of blacklisted URLs aggregated per TLD.

In addition, the 'Age' of a TLD plays a significant positive role in explaining the variance of phishing domains - more precisely - the older the TLD, the more its domains are used for phishing. As mentioned before, one possible explanation could be that it is more difficult for 'older' TLDs such as the legacy `.com` or `.net` gTLDs to deploy security

TABLE 3: Negative Binomial GLM regression models with 'Log_e' link function for number of phishing domains per TLD.

	Response Variable:
	Number of phishing domains in APWG
Number of unique domains [log scale]	1.416***
	(0.068)
Number of unique domains on shared IPs	0.00000*
	(0.000)
Free	5.804***
	(1.150)
Domain popularity index	−0.0001**
	(0.000)
URL shortner popularity index	0.367**
	(0.141)
Age	0.120***
	(0.009)
DNSSEC deployment rate	−2.900***
	(0.508)
Restricted registration	−0.900***
	(0.173)
ICT development index	0.027
	(0.361)
Constant	−2.507
	(2.096)
Observations	891
Log Likelihood	−2,061.0300
θ	0.768*** (0.058)
Akaike Inf. Crit.	4,412.059
Note:	*p<0.05; **p<0.01; ***p<0.001
	Standard errors in brackets

measures at scale. The results confirm our findings indicating that the new gTLDs contain less malicious domain names (see Figure 7).

Moreover, we see a negative significant relation between the variable 'DNSSEC deployment rate' and the count of phishing domains. Note that we use this variable as a proxy for the security efforts of TLD registries as it does not prevent phishing attacks.

As expected, the variable 'Restricted registration' shows that TLDs with restricted registration policy such as `.museum`, `.pharmacy` or `.nyc` contain less phishing domains. On one hand, restricted registration policies require more effort for a cybercriminal to register a domain name for malicious purposes. On the other hand, domains of new brand gTLDs, which are not available for registration, are also much more likely to be secure and often remain unused or redirect to a different domain name [62].

Finally, the 'ICT development index' does not play a significant role in explaining phishing abuse counts, after fixing the country variable and controlling for country variations.

7. Conclusion

In this paper, we propose security metrics to measure and benchmark entire TLDs against their market using representative malware and phishing datasets. We explicitly distinguish the metrics from the idea of measuring security *performance* because a TLD is not a single organization but constitutes an entire 'domain name ecosystem' of different types of intermediaries. Registrars add domains to the list kept by the registry while hosting companies provide web servers for domains.

We devise three different measures of 'badness' per TLD (domain, FQDN and URL) which reflect attackers' profit-maximizing behavior. Attackers abuse free legal services, affecting the reputations of such associated services. We find bad URLs using third party domains on the `.com` TLD such as Dropbox's `dropbox.com` and the `.tt` TLD such as *If This Then That*'s domain `if.tt`. We also find that the attackers abuse free domain registration provided by Freenom such as on the `.tk` and `.ml` TLDs.

Our findings suggest that new generic gTLDs such as `.new` and `.pharmacy` have disproportionately less attacker interest. This may be due to the higher prices of individual registrations at the beginning of the life cycle of the most of new gTLDs (i.e. during the "landrush" period) when registrants can get any domain name for a price premium; or due to a number of safeguards that were introduced by ICANN to mitigate rates of malicious, abusive and criminal activity in new gTLDs [63]. This contrasts with existing TLDs such as legacy `.com` and `.net` gTLDs which have disproportionately more attacker interest.

Another important indicator of TLD security we have looked into is the persistence of domain abuse or uptime. We find that in general the median is much less than the mean uptime as the latter is skewed by long-lived malicious domains. We have also analyzed how the uptime of TLDs relates to the amount of abuse in their networks, with the assumption that both indicate the security measures taken by TLDs. Notably, we found no correlation between the occurrence of abuse and median uptime metric suggesting that the majority of intermediaries act fast once they are notified about domain abuse but are less proactive to prevent them.

Our approach, however, comes with certain limitations which are planned to be improved in our future work. Our TLD metrics conflate maliciously registered domains and hacked websites. Distinguishing between those two different categories is important because they require different mitigation actions by different intermediaries. For example, hosting providers have a larger role to play in cleaning up content of compromised websites whereas domain registrars are more responsible for suspending maliciously registered domains.

Our paper also tries to answer the following question: to which extent domain abuse levels of TLDs are driven by TLD related properties? As expected, the number of unique domains or, in other words, the TLD size explains a significant part of the variance in phishing counts. In other words, the more domains in the registry, the bigger the 'attack surface' for cybercriminals. Moreover, we observe a very strong significant positive correlation between the price of domain registration and phishing counts. Last but not least, we observe a negative significant relation between the DNSSEC deployment rate and the count of phishing domains. As DNSSEC does not prevent phishing attacks, we interpret it as a proxy for the security efforts of the TLD registries.

We identify a number of recommendations for TLD registries. We recommend to sustain the effort to proactively acquire new feeds of phishing, malware, spam, and other sources of maliciously registered and compromised domains. Registries should send real-time or periodical notifications about abused domains to their clients, i.e. registrars. TLD registries should consider introducing positive and/or negative incentive structures for market players. For example, domain registrars/hosting providers running outdated versions of website software could be encouraged to update it before a domain renewal.

Finally, to facilitate monitoring of the three proposed occurrence security metrics indicating the concentration of 'badness' per TLD in comparison to its peers we maintain the website: http://remedi3s-tld.sidnlabs.nl.

Acknowledgments

This work was supported by SIDN, the .NL Registry and by NWO (grant nr. 12.003/628 .001.003), National Cyber Security Centrum (NCSC). Authors thank the contributors of data to Farsight Security's DNSDB. We would like to thank StopBadware, Anti-Phishing Working Group, and Cyscon GmbH for providing access to their data. Authors thank Wojciech Mazurczyk for his valuable comments.

References

[1] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster, "Building a Dynamic Reputation System for DNS," in *USENIX Security Symposium*. USENIX Association, 2010, pp. 273–290.

[2] H. Liu, K. Levchenko, M. Felegyhazi, C. Kreibich, G. Maier, G. M. Voelker, and S. Savage, "On the Effects of Registrar-level Intervention," *Proceedings of 4th USENIX LEET*, 2011.

[3] "Internet Corporation for Assigned Names and Numbers (ICANN)," https://www.icann.org.

[4] B. Rechterman and T. Ruiz, "Method for Gathering Domain Name Registration Information From a Registrant via a Registrar's Web Site," 2004, US Patent App. 10/408,050. [Online]. Available: http://www.google.com/patents/US20040199608

[5] "Top-Level Domains (gTLDs)," http://archive.icann.org/en/tlds.

[6] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage, "Click Trajectories: End-to-End Analysis of the Spam Value Chain," in *Proceedings of the 2011 IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2011, pp. 431–446.

[7] T. Moore and B. Edelman, "Measuring the Perpetrators and Funders of Typosquatting," in *Proceedings of the 14th International Conference on Financial Cryptography and Data Security*, ser. FC'10. Springer-Verlag, 2010, pp. 175–191.

[8] M. Korczyński, M. Król, and M. van Eeten, "Zone Poisoning: The How and Where of Non-Secure DNS Dynamic Updates," in *Proceedings of the 2016 ACM on Internet Measurement Conference*, ser. IMC '16. ACM, 2016, pp. 271–278.

[9] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, "Beyond Blacklists: Learning to Detect Malicious Web Sites from Suspicious URLs," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '09. ACM, 2009, pp. 1245–1254.

[10] 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 (IMC'13)*. ACM, 2013, pp. 63–76.

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

[12] A. Noroozian, M. Korczyński, S. Tajalizadehkhoob, and M. van Eeten, "Developing security reputation metrics for hosting providers," in *Proceedings of the 8th USENIX CSET*, 2015, pp. 1–8.

[13] D. Mahjoub, "Sweeping the IP Space: the Hunt for Evil on the Internet." Virus Bulletin Conference, 2014. [Online]. Available: https://www.virusbtn.com/pdf/conference/vb2014/VB2014-Mahjoub.pdf

[14] B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda, "FIRE: FInding Rogue nEtworks," in *Proceedings of the 2009 Annual Computer Security Applications Conference*, ser. ACSAC '09. IEEE Computer Society, 2009, pp. 231–240.

[15] M. Konte, R. Perdisci, and N. Feamster, "ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*. ACM, 2015, pp. 625–638.

[16] C. Wagner, J. François, R. State, A. Dulaunoy, T. Engel, and G. Massen, "ASMATRA: Ranking ASs Providing Transit Service to Malware Hosters," in *IFIP/IEEE International Symposium on Integrated Network Management*. IEEE, 2013, pp. 260–268.

[17] M. van Eeten, J. M. Bauer, H. Asghari, S. Tabatabaie, and D. Rand, "The Role of Internet Service Providers in Botnet Mitigation an Empirical Analysis Based on Spam Data." TPRC, 2010.

[18] "StopBadware: A Nonprofit Anti-malware Organization." https://www.stopbadware.org.

[19] "Anti-Phishing Working Group (APWG): Cross-industry Global Group Supporting Tackling the Phishing Menace," http://www.antiphishing.org.

[20] "ESET: Security Software," http://www.eset.com.

[21] "Fortinet: Network & Computer Security," http://www.fortinet.com.

[22] "Sophos: Computer Security, Antivirus," http://www.sophos.com.

[23] "Public Suffix List," https://publicsuffix.org.

[24] "Cyscon GmbH." http://www.cyscon.de.

[25] "Verisign, Inc. Internet Security and Web Domain Names," https://www.verisigninc.com.

[26] "TLD Zone File Access for .com, .net and/or .name," https://www.verisign.com/en_US/channel-resources/domain-registry-products/zone-file/index.xhtml.

[27] "Stichting Internet Domeinregistratie Netherland (SIDN) Labs," https://www.sidn.nl.

[28] "ICANN: Centralized Zone Data Service," https://czds.icann.org.

[29] "Council of European National Top Level Domain Registries (CENTR)," https://centr.org.

[30] G. Aaron and R. Rasmussen, "Anti-Phishing Working Group (APWG) Global Phishing Survey: Trends and Domain Name Use in 2H2014," http://internetidentity.com/wp-content/uploads/2015/05/APWG_Global_Phishing_Report_2H_2014.pdf, 2015.

[31] DomainTools, "Domain Count Statistics for TLDs," http://research.domaintools.com/statistics/tld-counts.

[32] "Farsight Security," https://www.farsightsecurity.com.

[33] "DNSDB," https://www.dnsdb.info.

[34] "ICANN Approves Name Collision Occurance Management Framework," https://www.icann.org, August 2014.

[35] S. Tajalizadehkhoob, M. Korczyński, A. Noroozian, C. Gañán, and M. van Eeten, "Apples, Oranges and Hosting Providers: Heterogeneity and Security in the Hosting Market," in *Network Operations and Management Symposium (NOMS)*. IEEE/IFIP, 2016, pp. 289–297.

[36] N. Nikiforakis, W. Joosen, and M. Johns, "Abusing Locality in Shared Web Hosting," in *Proceedings of the 4th European Workshop on System Security*. ACM, 2011, pp. 1–7.

[37] "Alexa: Actionable Analytics for the Web," http://www.alexa.com.

[38] F. Maggi, A. Frossi, S. Zanero, G. Stringhini, B. Stone-Gross, C. Kruegel, and G. Vigna, "Two Years of Short URLs Internet Measurement: Security Threats and Countermeasures," in *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 2013, pp. 861–872.

[39] N. Gupta, A. Aggarwal, and P. Kumaraguru, "bit.ly/malicious: Deep dive into short url based e-crime detection," in *2014 APWG Symposium on Electronic Crime Research (eCrime)*, September 2014, pp. 14–24.

[40] "Freenom: Registry Operator of the .TK, .ML, .GA, .CF, and .GQ ccTLDs," http://www.freenom.com.

[41] ".Pharmacy gTLD," http://nic.pharmacy.

[42] "Donuts Launches Domain Namespace Expansion with 307 gTLD Applications," http://www.donuts.domains/donuts-media/press-releases/donuts-launches-domain-namespace-expansion-with-307-gtld-applications.

[43] Donald E. Eastlake 3rd, "Domain Name System Security Extensions," Internet Requests for Comments, RFC 2535, March 1999. [Online]. Available: http://www.rfc-editor.org/rfc/rfc2535.txt

[44] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," Internet Requests for Comments, RFC 4033, March 2005. [Online]. Available: http://www.rfc-editor.org/rfc/rfc4033.txt

[45] R. van Rijswijk-Deij, A. Sperotto, and A. Pras, "DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, ser. IMC '14. ACM, 2014, pp. 449–460.

[46] P. Wallström, "Zonemaster: A new DNS Delegation Testing Tool." Presented as the DNS-OARC's Spring 2015 Workshop, 2015. [Online]. Available: https://indico.dns-oarc.net/event/21/contribution/21/material/slides/0.pdf

[47] S. B. Hulley and S. R. Cummings and W. S. Browner and D. Grady and T. B. Newman, *Designing Clinical Research: an Epidemiologic Approach. 4th ed.* Philadelphia, PA: Lippincott Williams & Wilkins, July 2013.

[48] International Telecommunication Union (ITU), "Measuring the Information Society Report 2014," https://www.itu.int/en/ITU-D/Statistics/Documents/publications/mis2014/MIS2014_without_Annex_4.pdf, 2014.

[49] D. Plohmann, K. Yakdan, M. Klatt, J. Bader, and E. Gerhards-Padilla, "A Comprehensive Measurement Study of Domain Generating Malware," in *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Aug. 2016, pp. 263–278.

[50] "ZeusTracker: A Nonprofit Organization Tracking ZeuS C&C Servers." https://zeustracker.abuse.ch.

[51] G. Aaron and R. Rasmussen, "Anti-Phishing Working Group (APWG) Global Phishing Survey: Trends and Domain Name Use in 1H2014," http://docs.apwg.org/reports.

[52] "ICANN: New Generic Top-level Domains," http://newgtlds.icann.org.

[53] "CENTR Report: Myth & Fact: New gTLDs, Registries & Registrars," https://centr.org, October 2014.

[54] "PhishTank: A Nonprofit Anti-phishing Organization," http://www.phishtank.com.

[55] "IFTTT: If This Then That Service," https://ifttt.com.

[56] T. Moore and R. Clayton, "Examining the impact of website take-down on phishing," in *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*. ACM, 2007, pp. 1–13.

[57] C. Gañán, O. Cetin, and M. van Eeten, "An Empirical Analysis of ZeuS C&C Lifetime," in *Proceedings of the 10th ACM Symposium (ASIACCS)*. ACM, 2015, pp. 97–108.

[58] S. Tajalizadehkhoob, C. Gañán, A. Noroozian, and M. van Eeten, "The Role of Hosting Providers in Fighting Command and Control Infrastructure of Financial Malware," in *Proceedings of the 12th ACM Symposium (ASIACCS)*. ACM, 2017.

[59] E. L. Kaplan and P. Meier, "Nonparametric Estimation from Incomplete Observations," *Journal of the American statistical association*, vol. 53, no. 282, pp. 457–481, 1958.

[60] S. Tajalizadehkhoob, R. Böhme, C. Gañán, M. Korczyński, and M. van Eeten, "Rotten Apples or Bad Harvest? What We Are Measuring When We Are Measuring Abuse," 2017. [Online]. Available: https://arxiv.org/abs/1702.01624

[61] A. C. Cameron and P. K. Trivedi, *Regression analysis of count data*. Cambridge university press, 2013, vol. 53.

[62] 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 ACM Conference on Internet Measurement Conference*, ser. IMC'15. ACM, 2015, pp. 381–394.

[63] "New gTLD Program Safeguards Against DNS Abuse ICANN Operations and Policy Research," https://newgtlds.icann.org/en/reviews/dns-abuse/safeguards-against-dns-abuse-18jul16-en.pdf, July 2016.

Appendix

TABLE 4: Occurrence Security Metrics: Top 10 large gTLDs with the highest relative concentration of unique phishing domains, FQDNs and URLs.

Large gTLD, APWG

2014			2015		
TLD	**# Domains**	**Score**	**TLD**	**# Domains**	**Score**
COM	48159	41	INFO	5520	106
ORG	3954	38	COM	109381	88
INFO	1999	37	NET	11406	72
NET	5503	35	ORG	7748	70
BIZ	632	25	BIZ	1410	59
TLD	**# FQDN**	**Score**	**TLD**	**# FQDN**	**Score**
ORG	111789	1065	COM	1482304	1200
COM	756068	637	NET	112556	712
NET	18667	120	INFO	18487	356
INFO	2826	52	BIZ	5075	214
BIZ	797	31	ORG	14414	131
TLD	**# URL**	**Score**	**TLD**	**# URL**	**Score**
COM	4499882	3789	COM	19140097	15499
ORG	350741	3342	NET	636241	4025
NET	117029	752	ORG	279650	2548
BIZ	15104	589	INFO	76538	1476
INFO	6902	127	BIZ	16496	697

TABLE 5: Occurrence Security Metrics: Top 10 ccTLDs with the highest relative concentration of unique phishing domains, FQDNs and URLs.

ccTLD, APWG

2014			2015		
TLD	**# Dom**	**Score**	**TLD**	**# Dom**	**Score**
PN	12	1993	KP	1	4347
NE	3	928	ML	5158	2946
TL	19	863	NE	8	2476
VU	11	724	PN	14	2325
ML	585	680	JE	94	1926
SD	12	576	GP	38	1791
TG	3	562	KI	6	1612
TO	82	522	CF	3727	1539
CI	18	514	TL	30	1480
KE	87	500	GQ	1918	1325
TLD	**# FQDN**	**Score**	**TLD**	**# FQDN**	**Score**
TT	4590	153975	TT	5626	140579
GG	1845	35480	GG	1976	38368
NF	198	18574	TF	967	37049
PN	105	17441	PN	185	30730
GA	14999	15305	VU	436	28703
VU	106	6978	JE	1019	20881
VC	308	3311	NF	204	19136
TF	47	1958	IO	26889	17751
SH	104	1637	MU	601	8577
TL	33	1500	TO	1312	8067
TLD	**# URL**	**Score**	**TLD**	**# URL**	**Score**
GL	2672798	46891192	GD	2032729	65068149
LY	1614324	12417876	GL	3290493	64494178
MP	8197	1600976	LY	2160713	18195477
TT	26556	890841	TF	152721	5851379
GD	16292	581857	MP	28079	5484179
HT	5000	227272	GY	90928	2796924
ST	24681	226120	TT	109272	2730434
GG	5789	111326	TO	229928	1413810
SH	6543	103039	JE	58327	1195225
GY	1162	35742	GS	69916	982932

TABLE 6: Occurrence Security Metrics: Top 10 smaller gTLDs with the highest relative concentration of unique phishing domains, FQDNs and URLs.

Smaller gTLDs, APWG

2014			2015		
TLD	**# Dom**	**Score**	**TLD**	**# Dom**	**Score**
POST	1	5263	HERE	10	500000
EDU	41	540	ANDROID	4	400000
GOV	13	243	ZIP	4	400000
COOP	5	63	PLAY	2	200000
LINK	31	58	SONY	5	166666
CODES	2	52	NEW	3	150000
PRO	53	41	NICO	1	100000
HOST	1	40	BOO	1	50000
CAT	29	35	DAD	1	50000
XXX	35	34	EAT	1	50000
TLD	**# FQDN**	**Score**	**TLD**	**# FQDN**	**Score**
POST	1	5263	ANDROID	8	800000
EDU	111	1462	HERE	12	600000
MOBI	3361	369	ZIP	4	400000
GOV	13	243	PLAY	2	200000
HOST	4	161	SONY	5	166666
COOP	5	63	NEW	3	150000
LINK	33	62	COUNTRY	8997	125955
PRO	72	56	NICO	1	100000
CODES	2	52	BOO	1	50000
CAT	42	51	DAD	1	50000
TLD	**# URL**	**Score**	**TLD**	**# URL**	**Score**
POST	1	5263	ANDROID	8	800000
EDU	293	3860	HERE	12	600000
MOBI	5287	580	ZIP	4	400000
HOST	9	363	PLAY	2	200000
GOV	14	261	SONY	5	166666
COOP	20	253	NEW	3	150000
LINK	118	222	COUNTRY	8997	125955
PHOTOS	34	198	NICO	1	100000
XXX	158	154	BOO	1	50000
CLUB	209	130	DAD	1	50000

TABLE 7: Occurrence Security Metrics: Top 10 large ccTLDs with the highest relative concentration of unique malware domains, FQDNs and URLs.

ccTLD, Stopbadware

2014			2015		
TLD	**# Dom**	**Score**	**TLD**	**# Dom**	**Score**
VU	725	47728	VU	3703	243778
PN	23	3820	MV	14	1564
GA	3291	3358	PN	9	1495
TF	61	2541	ML	2535	1448
CU	31	2073	TO	218	1340
ML	1659	1929	NF	13	1219
ZA	1971	1925	TH	711	1138
NF	19	1782	CF	2562	1058
MV	15	1675	TF	27	1034
TH	924	1421	LS	8	998
TLD	**# FQDN**	**Score**	**TLD**	**# FQDN**	**Score**
VU	918	60434	VU	3712	244371
TT	1581	53035	TT	2781	69490
NR	123	24600	NR	113	22600
PN	47	7807	GG	414	8038
GA	4445	4535	PN	22	3654
SU	4761	4060	RU	123533	2458
NU	8870	3805	TO	332	2041
ML	3252	3781	NF	21	1969
NF	39	3658	TL	39	1924
KI	13	3494	TH	1055	1689

ccTLD, Stopbadware

2014			2015		
TLD	**# URL**	**Score**	**TLD**	**# URL**	**Score**
TT	2522	84602	VU	6466	425674
VU	1060	69782	TL	2546	125666
HT	787	35772	TT	4123	103023
NR	170	34000	GG	3081	59825
GL	641	11245	ID	65497	46964
PN	58	9634	EG	3393	36099
TL	198	9000	NF	350	32833
LY	1055	8115	TR	121391	32647
MZ	305	7093	RU	1314685	26167
TH	4199	6460	GE	6783	25922

TABLE 8: Occurrence Security Metrics: Top 10 smaller gTLDs with the highest relative concentration of unique malware domains, FQDNs and URLs.

Smaller gTLDs, Stopbadware

2014			2015		
TLD	**# Dom**	**Score**	**TLD**	**# Dom**	**Score**
INT	7	3589	ADS	1	100000
EDU	270	3557	EDU	223	2634
GOV	63	1178	WIN	9909	1774
MUSEUM	3	696	SITE	1387	1732
COOP	32	406	INT	2	1025
CAT	275	338	BID	804	834
PRO	307	241	LOAN	540	498
AERO	22	232	WORK	445	480
ASIA	656	206	GOV	21	380
TRAVEL	35	179	RACING	80	295
TLD	**# FQDN**	**Score**	**TLD**	**# FQDN**	**Score**
EDU	392	5164	ADS	1	100000
INT	9	4615	EDU	322	3804
MUSEUM	11	2552	ASIA	6728	2978
GOV	73	1365	COUNTRY	181	2533
COOP	35	444	WIN	9937	1779
NAME	778	405	SITE	1391	1737
CAT	286	352	INT	2	1025
ASIA	949	299	BID	805	835
PRO	343	269	WORK	721	778
AERO	22	232	MUSEUM	8	662
TLD	**# URL**	**Score**	**TLD**	**# URL**	**Score**
MUSEUM	54	12529	ADS	4	400000
EDU	907	11949	MUSEUM	713	59072
INT	13	6666	EDU	1679	19839
XXX	4276	4176	NINJA	9958	19175
GOV	117	2188	ASIA	28050	12418
TRAVEL	152	780	WEBSITE	11420	9522
CAT	560	689	ROCKS	4852	8029
COOP	47	596	LOL	2187	6633
NAME	1113	579	RED	3649	6164
PRO	690	542	WORK	5622	6066

TABLE 9: Occurrence Security Metrics: Top 10 large gTLDs with the highest relative concentration of unique malware domains, FQDNs and URLs.

Large gTLD, Stopbadware

2014			2015		
TLD	**# Domains**	**Score**	**TLD**	**# Domains**	**Score**
COM	320718	270	COM	358118	290
ORG	26080	248	ORG	26914	245
NET	36170	232	NET	37628	238
BIZ	5391	210	INFO	11520	222
INFO	10787	198	BIZ	4885	206
TLD	**# FQDN**	**Score**	**TLD**	**# FQDN**	**Score**
COM	598861	504	COM	470505	381
NET	62377	401	ORG	35475	323
ORG	34759	331	NET	49081	310
BIZ	8089	315	BIZ	7030	297
INFO	12209	224	INFO	12765	246
TLD	**# URL**	**Score**	**TLD**	**# URL**	**Score**
COM	1178453	992	BIZ	82373	3483
NET	138956	893	COM	4077811	3302
ORG	92629	882	INFO	163014	3145
INFO	36132	665	ORG	302093	2752
BIZ	13606	531	NET	423201	2677

TABLE 10: Uptime Security Metrics: Top 20 TLDs with the highest median uptime (in hours) based on phishing websites reported by Cyscon between June 2015 and January 2016.

TLD	**Min**	**Mean**	**Median**	**Max**	**SD**	**SE**
FJ	3,418.18	3,418.18	3,418.18	3,418.18	–	–
MP	0	2,555.56	3,383.15	4,505.47	1,817.38	605.79
RW	0	2,272.63	2,969.54	3,848.35	2,016.61	1,164.29
BAYERN	2,953.66	2,953.66	2,953.66	2,953.66	–	–
DANCE	2,376.27	2,376.27	2,376.27	2,376.27	–	–
TC	0	1,919.85	2,347.26	3,412.28	1,745.83	1,007.96
CX	0	1,643.56	1,643.56	3,287.12	2,324.35	1,643.56
WF	0	1,407.89	1,482.41	3,345.55	1,279.74	522.45
LIMO	461.07	1,471.04	1,471.04	2,481.01	1,428.31	1,009.97
PHOTOGRAPHY	0	1,449.40	1,449.40	2,898.80	2,049.76	1,449.40
GM	0.21	1,356.72	1,435.16	4,357.53	865.54	96.17
JE	556.93	1,379.68	1,379.68	2,202.42	1,163.54	822.74
LA	0	1,483.53	1,297.08	3,829.01	1,075.59	260.87
SM	1,270.34	1,270.34	1,270.34	1,270.34	–	–
CODES	664.97	905.21	905.21	1,145.44	339.74	240.23
ESTATE	842.06	842.06	842.06	842.06	–	–
GURU	839.99	839.99	839.99	839.99	–	–
GF	0	618.61	618.61	1,237.22	874.84	618.61
BUSINESS	610.44	610.44	610.44	610.44	–	–
VC	0	939.10	583.80	3,999.93	1,128.59	273.72