Domain Name Abuse in the US and PR Zones

An investigation and assessment of the extent to which domain names registered in the US and PR zones are associated with abusive activity

Prepared by Interisle Consulting Group, LLC

Interisle
Consulting Group

21 December 2018

Contents

1 Introduction

In order to understand the way in which the US and PR top-level domains are involved in the registration of domain names that are used to support malicious (whether or not criminal) Internet behavior, _____ asked Interisle to conduct a study of abusive activity in those two zones, and to do so quickly (within 3 weeks). This document is the report of that study.

In this context the term "abusive activity" is customarily understood to include phishing, malware distribution, botnet command and control (C2), and spam. The Interisle study focused on spam as the most reliable proxy for abuse in general, as it is frequently the delivery method for attacks in the other three categories. The principal sources of data were reputation data sets compiled by organizations such as Spamhaus and SURBL, augmented by information gathered from Interisle's network of contacts in the global abuse detection and mitigation community.

The context of the study and the data sources on which we relied are described in more detail in Section 2 ("Purpose and Scope"). Section 3 ("Methodology") explains how we gathered and analyzed information. Abuse in the PR zone is described in Section 4; abuse in the US zone is described in Section 5. We summarize our findings in Section 6.

2 Purpose and Scope

The purpose of this study was to investigate and assess the reputations of the US and PR top-level domain registries with regard to abuse in the Domain Name System.

By reputation, we mean *the way in which users and organizations perceive domain names, and by extension the registries from which they are delegated, through the lens of the reputation data that network- and system administrators employ to protect their users, devices, and networks from security threats.*

Users and organizations view reputation data as a means to quantify risk. They want to understand the likelihood that they will be exposed to abuse, financial loss, or other harm should they access web site content or open an email attachment received from a mail server with a hostname delegated from a given registered domain name.

Users and organizations rely upon a large and diverse industry of service providers for reputation data. These providers operate detection and analysis systems that assign a positive or negative characteristic to domain names, hyperlinks or Internet addresses. Reputation data providers often publish or report on abuse activity at a domain registry level. Domain registry reputation is calculated in different ways and for different kinds or types of abuse activities or security threats. Some reputation services compare the incidence of domain names that they associate with an abuse activity against the registry's total delegated domain names. Others calculate a score based on observed abuse domains delegated from a given registry against total observed domains. In other cases, reputation is calculated by scoring abuse domains against delegated domain names that resolve in the DNS (and are thus "weaponized").

For this study, we considered domain names in the US and PR registries that were associated with three types of abuse:

Spam. Spam is more than unsolicited commercial email. It is the preferred delivery method for a variety of financial fraud attacks (*e.g.,* phishing, ransomware, and business email compromise), scams (*e.g.*, advanced fee fraud or credential theft), and sales of counterfeit goods (*e.g.*, licensed software or prescription pharmaceuticals). Spam is also a used to deliver malware attachments or hyperlinks that connect users to hosted ("drive-by") malware. Attackers use malware to infect and recruit devices into botnets. They also use malware to intercept or exfiltrate sensitive data. Spam is not limited to email but is now delivered through mobile, messsaging, and social media services. For this study, we considered the domain names that are used to name the infrastructure components (mail relays and name servers) that transmit spam (*spam domains*).

Botnet or attack infrastructure recruitment and operation. Attackers often utilize algorithmically generated domain names for the command-control components of their botnet or spam infrastructures. When security experts reverse-engineer the domain generation algorithms (DGAs) of malware families, they are able to produce the same complete daily lists of domain names that these malware generate on a daily basis. For this study, we considered the subset of generated names that are registered and used to identify command-control hosts (C2 domains).

Illicit or harmful content hosting. Attackers assign hostnames from domain names they register for several forms of attacks, including phishing, inauthentic news sites, drive-by malware, scams, and sales of counterfeit goods. Many of these domain names are identified through the spam emails that are used to lure victims to the illicit or harmful content. For this study, we considered domain names that are collected via spam and malicious content detection systems and also "typed" as phishing, malware, *etc.* Using DNS and IP threat intelligence data, we also investigated the reputations of the addresses that the delegated host names of our abuse domain name samples resolve to, the networks that use (announce) these addresses, and the organizations that administer the address allocations.

While we do have some phishing data from the Spamhaus reputation data, we did not explore phishing more deeply. Phishing reputation is typically derived from URLs, not domain names. Some URLs associated with phishing attacks are constructed using malicious (abusive) registrations; for example, a phisher may try to lure a victim to a fraudulent bank login page with a deceptive domain name (*e.g.*, http://www.chaase.us/login.php). In other cases, phishing URLs are constructed using the legitimately registered domain names of web sites that have been compromised (*e.g.*, http://myembroidery.us/chasebank/login.php). Discriminating between maliciously registered domain names and legitimately registered domain names that have been exploited to facilitate phishing attacks requires more time than was available for this study.

3 Methodology

For this time-limited study we focused on three categories of abuse activity: spam domains, botnets (attack infrastructure), and illicit or harmful content hosting. These are the most widely monitored abuse activities, and accordingly, reputation data is most abundantly available. We were able to obtain license or permission to use commercial or reputation data or privately collected (research) data that we could process and that would provide meaningful insights in the available timeframe. We also used publicly available scoring and ranking data and data from threat intelligence services to which Interisle subscribes. The sets of data we studied included:

- A sample of the Spamhaus Domain Block List[1] (DBL). We obtained API Keys for two methods of API access. We first queried the DBL for a set of domains in the US and PR registries that made "listed" status, *i.e.*, the domains that were captured in Spamhaus's trap network, investigated, and associated with a type of abuse. We obtained 19,555 unique listed domains listed over a fourteen-day window. We queried the DBL to obtain metadata associated with these domains, including registrar, abuse type, creation date, first/last seen timestamps, IP addresses associated with the listed domain, and a confidence score.
- Domain Whois. Where the domain Whois was available, we collected name server data provided by the registrant for the DBL study set. We were unable to acquire point of contact data.
- IP Whois. We collected IP netblock and autonomous system number (ASN) and organization name for those domain names in the DBL study set that were associated with *snowshoe* spam. "Snowshoe" refers to an obfuscation technique: spammers inject spam from many IP addresses to frustrate IP blocking by mail relays.
- The Spamhaus World's Most Abused TLDs list.[2] A ranking system that calculates a *badness index* for all top-level domains based on abuse domains observed and total domains observed for a given TLD.
- The Spamhaus World's Most Abused Registrar list.[3] A ranking system that calculates a *badness index* for registrars based on abuse domains observed and total domains observed for a given sponsoring registrar.

---

[1] Spamhaus Domain Block List, https://www.spamhaus.org/dbl

[2] Spamhaus World's Most Abused TLDs, https://www.spamhaus.org/statistics/tlds

[3] Spamhaus World's Most Abused Registrars, https://www.spamhaus.org/statistics/registrars

- SURBL Most Abused TLDs list.[4] We obtained permission to use a custom data set from SURBL that contains daily abuse domain counts for the Top 50 most abused TLDs from 2015-present.
- Spamhaus Registry of Known Spam Operators (ROKSO).[5] This is a "repository of information and evidence of known persistent spam operations."
- Spamhaus Block List (SBL) Advisories. SBL advisories identify IP addresses (and netblocks) of spam sources. The listings identify operators who are "known" to support spam services, bulletproof hosting, spam operations, or spam gang activities (spampaigns). The listings also include IP addresses (and netblocks) of operators that host or support abuse activities or security threats including botnet controllers (C2), ransomware and malware, or phishing.
- Bambenek C2 data. We received permission to use a list of algorithmically generated domain names from fifty (50) malware families from a fellow researcher. The list includes the domain names that were generated by a known, reverse-engineered DGA that resolved via the DNS from January 1, 2018 – present.
- Seclytics Predictive Threat Intelligence service.[6] Interisle used ASN, IP, and other threat data reported through this subscription service.
- Domain Tools IRIS.[7] Interisle used DNS intelligence provided through this subscription service.

Using reputation data and scoring (ranking) systems, we compared US and PR against other top-level domains, both generic and country code. Specifically:

1. We analyzed a 14-day sample of spam domains in US and PR to identify particular contributing factors to the abuse prevalence such as sponsoring registrar, creation dates, and the types of abuse with which US and PR domains are associated.
2. We compiled a historical comparison of daily spam domain counts in US against gTLDs and ccTLDs from 2015 to present.

---

[4] SURBL Most Abused TLDs list, http://www.surbl.org/tld

[5] Spamhaus ROKSO, https://www.spamhaus.org/rokso

[6] Seclytics, https://www.seclytics.com

[7] Domain Tools IRIS, https://research.domaintools.com/iris

3. We investigated the name server hosting characteristics of the spam domain data from (2).
4. For the snowshoe form of spam, and using IP Whois data, we investigated the reputations of network and hosting operators from which US domains emit spam or host illicit or harmful content.
5. We extracted US domains from a composite list of C2 domains that resolved in 2018.

In the following sections, we report our findings. In consideration of permissions granted or licensing, we provide aggregated or derivative data only. Where appropriate, we represent the data in tables or charts.

Due to the short time available to gather, process, and correlate the data from various sources, we were not able to reconcile all of the oddities and inconsistencies in the data we obtained; for example, some records in our sample data did not contain registrar information, and some IP threat intelligence records did not yield autonomous system numbers. However, we have a high confidence that these data, even though imperfect, are sufficient to support "order of magnitude" conclusions.

4 Abuse in the PR Zone

The data sets we consulted for this study showed almost no sign of abuse in the PR zone. This may be due to the fact that the price to register a name in PR is over US$1,000, and the registration process takes 2 or 3 days. These are barriers that would effectively discourage most abusive registrants and registrars.

5 Abuse in the US Zone

We studied three types of abuse: spam domains, botnet command and control, and URI abuse.

5.1 Spam Domains

We obtained Spamhaus DBL data for a 14-day period ending 13 Dec for the US top-level domain. This sample provided us with 19,555 domains associated with some abuse activity.

The length of our sample period is relevant to Spamhaus DBL: it is our understanding that Spamhaus de-lists abuse domains 14 days after the domain was last seen. Note that any appearance of the abuse domain triggers a reset of the last seen timestamp to fourteen days. Spamhaus treats legitimate domains that are identified as compromised differently: these time out 25 hours after last seen.

Second-level string composition suggests that some abusive registrants employ automatic name generation (*e.g.*, a name suggestion tool) or a DGA; for example, we observe large numbers of strings composed of numbers and letters; for example, 007z.us, 01cn.us, o1qyi.uc. We also see strings composed of numbers and English words (*e.g.*, 01sendfaster.us, 01yourlifeeasy.us) or just English words (kansascitygatewaypage.us, valleyprojecthost.us).

5.1.1 Registrar Intelligence

We sorted the sample data by registrar. The following table shows the top ten of the 30 registrars and the total number of abuse-listed domains under their sponsorship. (Note that we have no registrar data for 152 domains.) An extraordinary 95% of the abuse-listed domains are registered through Namecheap, Inc. or its resellers.

Registrar	*# of Domains*
NameCheap, Inc.	18,587
DNC Holdings, Inc.	410
Dynadot LLC	130
GoDaddy.com, Inc.	89
PDR Ltd. d/b/a PublicDomainRegistry.com	68
TLD Registrar Solutions Ltd.	37
ENOM, INC.	14
<no registrar identified>	152
Grand Total	**19,555**

We next sorted the sample data by abuse type. The majority of abuse domains in the sample are typed as having a bad reputation (14,156 or 75% of the 19,555-domain sample). Snowshoe spam domains account for 24% of the sample.

Type	*# of Domains*
abused-phish	2
abused-spam	9
bad_reputation	14,156
botnet	88
botnetcc	13
malware	1
phish	356
sinkhole	188
snowshoe	4,742
Grand Total	**19,555**

Lastly, we sorted the sample data by registrar and type. The following table shows the top 5 registrars and the totals for each abuse type.

Registrar	*Type*	*# of Domains*
NameCheap, Inc.	abused-phish	2
	abused-spam	3
	bad_reputation	13,952
	botnet	88
	botnetcc	5
	phish	257
	sinkhole	71
	snowshoe	4209
NameCheap, Inc. Total		18,587
DNC Holdings, Inc.	snowshoe	410
DNC Holdings, Inc. Total		410
Dynadot LLC	bad_reputation	1
	phish	6
	sinkhole	111
	snowshoe	12
Dynadot LLC Total		130
GoDaddy.com, Inc.	abused-spam	6
	bad_reputation	32
	malware	1
	phish	12
	sinkhole	1
	snowshoe	37
GoDaddy.com, Inc. Total		89
<no registrar identified>	bad_reputation	104
	phish	3
	snowshoe	45
<no registrar identified> Total		152

The majority of Namecheap's abuse domains are typed as having a bad reputation (13,952, which is 71% of the 19,555 total domains). This tag is used in the API for domain names that are believed to be registered for abusive purposes, based on the reputation of related resources. The second largest type of abuse domain under Namecheap's sponsorship is snowshoe spam domain (4,209, which is 22% of the 19,555 total domains). All of the abuse domains in our sample that were registered through DNC Holdings were used in snowshoe spam campaigns.

A small number of phish domains (257) and botnet domains (88) are present among Namecheap's abuse domain set. Note that the DBL counts sinkhole domains as a separate type. These are not actively engaged in abuse activities, but they remain while they continue to resolve and are observed in Spamhaus's detection systems while infected hosts attempt to connect to the C2 host.

Namecheap is notorious among first responders and law enforcement as a registrar favored by spam gangs. Our finding from the sample we studied suggests that the notoriety is well-deserved.

5.1.2 Addressing and ASN Threat Intelligence

Since we were unable to obtain registrant or other contact information associated with the US domains in our sample, we were unable to study persons of interest or possible conspiracies. We could, however, study hosting behavior using Internet address and autonomous system Whois data and associated addressing and ASN threat intelligence. Pivoting from name space to address space, we investigated the addresses to which the abuse domains in US resolve, and then the network allocations and autonomous systems of these IP addresses to determine whether these addresses lie in known/notorious neighborhoods. We also researched reputation data for these IP addresses to gain an additional understanding of the various kinds of abuse activities that contribute to bad reputations.

Of the 19,555 US domains in the list, 2,571 have more than 4 addresses associated with the spam domain; 283 have more than 10 addresses; 28 have more than 15 addresses; 8 have more than 25; and two have more than 48 addresses. We obtained the AS numbers and hence the organizations and country associated with each IP address associated with each abuse domain. The following table shows the top ten AS organizations and their associated abuse domain counts.

AS organization	AS Country	# of Domains
PIN-AS, Petersburg Internet Network, Limited	Russian Federation	7,002
NTHL, Network Transit Holdings, LLC	United States	4,826
Orange	France	1,589
Timeweb Ltd.	Russian Federation	1,457
Namecheap, Inc.	United States	1,119
Digital Ocean	United States	1,028
VSERVER	Ukraine	950
OVH Net	France	616
CloudFlare Net	United States	581
Online SAS	France	564

We consulted the Spamhaus SBL Advisory to see whether these Autonomous Systems had IP addresses listed as spam sources or security threats. Of the ten AS organizations in the list, those with more than 10 SBL listings are:

- Timeweb has 59 current listings.[8] Most listings identify botnet controllers. Several are spam sources.
- Cloudflare has 44 current listings,[9] mostly botnet controllers, spam, malware downloaders.
- OVH-Net has 35 current listings,[10] for phishing, spam, mostly botnet controllers, snowshoe, bit mining pool (crypto currency abusers).
- Orange has 30 current listings,[11] mostly spam sources, some phishing.
- Namecheap.com has 28 current listings,[12] for phishing, spam, malware downloaders, mostly botnet controllers and scams (Canadian Pharmacy).
- DigitalOcean has 14 current listings,[13] for phishing, mostly spam, and scam sites.

[8] https://www.spamhaus.org/sbl/listings/timeweb.ru

[9] https://www.spamhaus.org/sbl/listings/cloudflare.com

[10] https://www.spamhaus.org/sbl/listings/ovh.net

[11] https://www.spamhaus.org/sbl/listings/orange.com

[12] https://www.spamhaus.org/sbl/listings/namecheap.com

[13] https://www.spamhaus.org/sbl/listings/digitalocean.com

From these listings, we can infer that abuse domains registered in US are used in spam distribution and criminal/botnet infrastructures.

We also checked the Seclytics Predictive Threat Intelligence system for some historical data on these operators. Seclytics represents prevalence of abuse in an ASN using a calendar "heat map": abuse reported from multiple reputation feeds are integrated into the map, and the heat (color) intensifies according to the concentration of abuse on a given day increases and highlight the abuse types reported for a sample day. Figures a, b, c, and d below show heatmaps for Timeweb, Cloudflare, OVH-net, and Namecheap.

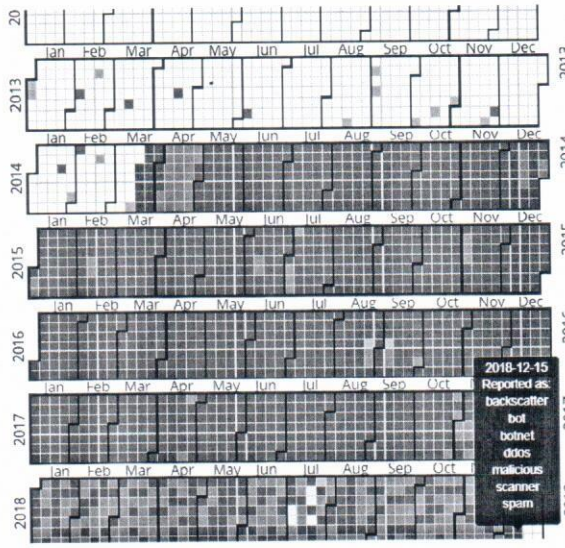

a Namecheap abuse heat map

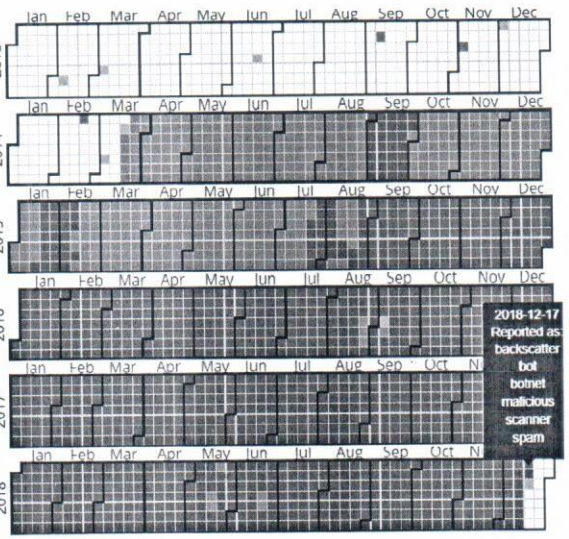

b Timeweb abuse heat map

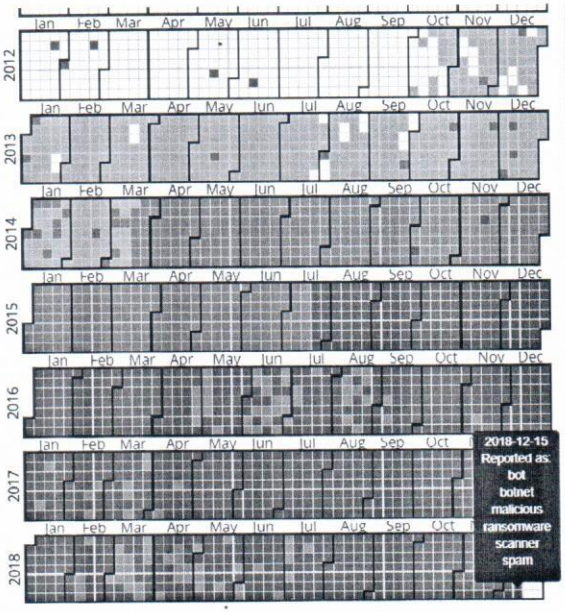

c OVH-net abuse heat map

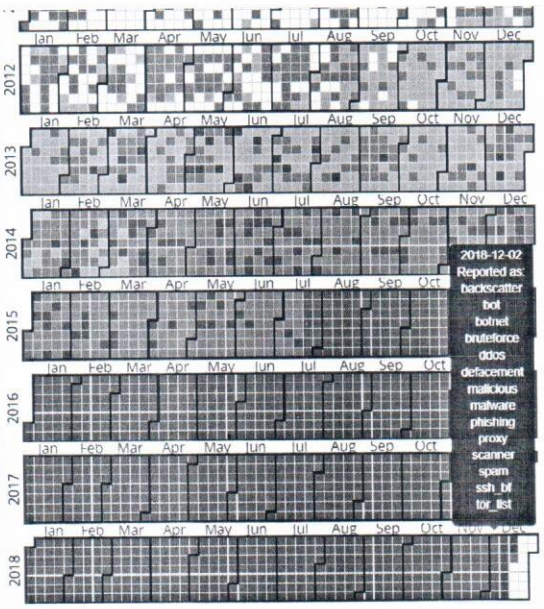

d Cloudflare abuse heat map

From these heat maps, we observe that abuse domains in US frequently resolve to IP addresses in network allocations of ASNs that have significant abuse activities.

5.1.3 Name Servers of Abuse Domains in US

To complete our mapping of the spam or criminal infrastructures, we studied the authoritative name servers of the abuse domains in our US TLD sample.

We queried the DNS using a dig client for the NS records of our 19,555 abuse domain sample. Of these 5,499 had no NS records (as of 17 Dec 2018). We collected the host names and the corresponding (set of) IP addresses (IPv4 and IPv6) for all of the name servers that we discovered, and then used Regional Internet Registry (IP) Whois to obtain the corresponding AS numbers and AS organizations. We identified 34,155 address entries for the 14,056 domains with NS records (19,555 minus 5,499).

The following table shows the top 30 most occurring domain names containing the nameservers:

NS domain	*Occurrences*
registrar-servers.com.	3,800
linode.com.	2,935
dnsdun.com.	1,528
dnsdun.net.	1,526
ns.cloudflare.com.	982
lunariffic.com.	818
csof.net.	436
gozonic.net.	420
dnspod.net.	323
dns.com.	258
nameserver.life.	132
venus.orderbox-dns.com.	96
mercury.orderbox-dns.com.	96
mars.orderbox-dns.com.	96
domaincontrol.com.	96
earth.orderbox-dns.com.	96
clearlinens.com.	86
rack-spaces.com.	82
parkingcrew.net.	80
kryptoslogicsinkhole.me.	71
kryptoslogicsinkhole.com.	71
kryptoslogicsinkhole.net.	71
kryptoslogicsinkhole.org.	71
hostmaze.com.	70
eaav.co.in.	66
dynadot.com.	56
mimarketlost.us.	46
iteamserver.com.	44
paynhost.com.	34
whois.com.	32

In total, there were 9,607 distinct nameserver domains. Of these, 9,500 were only used twice. This seems to be the case where such domains specify NS records within the same domain – for example: 4an8uj.us has two nameservers, ns80.4an8uj.us. and ns57.4an8uj.us.

We obtained the AS numbers and hence the organizations associated with each nameserver IP address. The following table shows the top 15 most occurring ASN organizations.

AS Organization	*Occurrences*
TimeWeb Ltd. (*i.e.*, timeweb.ru)	16,860
Cloudflare, Inc.	4,375
Verisign	3,632
DigitalOcean, LLC	1,880
CNSERVERS LLC	1,308
China Telecom	872
Adobe Systems	818
Amazon	496
CHINANET	436
Hangzhou Alibaba Advertising Co., Ltd.	436
Unified Layer	428
Google	322
OVH SAS (France)	191
Namecheap, Inc	174
(no NS record)	5,499

In total, there were 111 distinct ASNs (or sets of ASNs). Looking at Addressing and ASN Threat Intelligence for the name server IP addresses, we see that Timeweb.ru and Cloudflare are again top hosting operations for name servers associated with abuse activities in the US TLD. DigitalOcean, which was also among the top ten AS organizations with SBL listings, is among the top 5. We note that Seclytics heat maps for all 14 of Digital Ocean's ASNs have similarly high concentrations and intensity of abuse activity as the heat maps we showed for Timeweb,

Cloudflare, Namecheap, and OVH net. We also see similar concentrations and intensity of abuse activity for CNSERVERS LLC, China Telecom, CHINANET, Adobe Systems, Amazon, Hangzhou Alibaba Advertising Co., Ltd., Unified Layer, and Google.

With the exception of Verisign, we infer here as we did for hosting addresses that name servers of abuse domains in the US TLD also frequently resolve to IP addresses in network allocations of ASNs that have significant abuse activities.

5.2 Botnet Attack Infrastructure Registration Activity

Domain Generation Algorithms (DGAs) work as follows. Malicious executables installed on compromised hosts that form a botnet contain a DGA procedure. The same DGA procedure is present in the hosts that provide command and control (C2) for the botnet. The DGAs all create tens, hundreds, or more patterned or randomly composed second- or third-level strings daily and prepend these to a set of TLDs that the botnet operator has chosen for ready and typically inexpensive registration. Daily, the botnet operator registers one or several names from the generated list and creates address records for the C2 host name. Daily, the compromised hosts (bots) attempt to resolve the names created until they receive a positive response, and once received, the bots connect to the C2 host for operating instructions or new malicious executables.

To understand whether or not criminals register algorithmically generated domain names from the US and PR TLDs for the C2 components of their botnet or spam infrastructures, we obtained permission to use a list of algorithmically generated domain names from fifty (50) malware families from a fellow researcher. The list includes the domain names that were generated by a known, reverse-engineered DGA *and* that resolved via the DNS from January 1, 2018 to present.

Our C2 sample contained 2,471 DGA names delegated from 25 TLDs. No PR domains were found. COM and NET are typically the most abused gTLDs. 74 of the C2 domains in our sample were delegated from the US registry, representing three percent (74 of 2,471) of C2 domains across all TLDs found in the sample and 24% of ccTLDs found in the sample.

TLD	Count of C2 domains
net	1,157
com	798
eu	141
info	87
us	74
pw	67
xyz	45
top	23
biz	17
org	9
click	8
support	6
bid	6
co	5
work	5
tv	4
cc	3
ga	3
cx	2
online	2
bz	2
ac	2
mx	2
me	2
ru	1
Grand Total	**2,471**

ccTLD	Count of TLD
eu	141
us	74
pw	67
co	5
tv	4
cc	3
ga	3
bz	2
me	2
ac	2
mx	2
cx	2
ru	1
Grand Total	**308**

From the sample, we can conclude that the US registry is "in play" among botnet operators but not much more than that. First, our sample is derived from a modest set of malware family DGAs. We cannot extrapolate from our sample to whether or not the US registry is routinely included in other DGAs, nor can we extrapolate to whether or not the same percentage of domains from US are generated when US *is* included. We cannot reliably collect Whois for the domains in our sample: we received a list with no registration metadata. Whois for domain names that have been suspended or deleted is no longer available and some Whois has been changed to reflect change in registration from botnet operator to a sinkhole operator. We are also limited by the currently redacted Whois service. The inability to collect registrant data impedes our ability to associate the domains to an operator or set of conspirators. A future study including samples from more DGA families, including the sets of IP addresses and domain registration data captured while the domain resolved in the DNS, might provide additional intelligence.

5.3 URI Abuse

We obtained, from SURBL,[14] a series of files comprising daily data from 30 Jan 2014 through 12 Dec 2018 (a total of 1,659 days' data). (Data was not available for all days in 2014 after 30 Jan – data was available for only 217 days in that year.)

Each day's data comprised a list of the top 50 TLDs on the block list and the number of domains in that TLD that day. However, we found a number of anomalies or inconsistencies in the data. In some cases, an IP subnet (3 quads) was provided in place of a TLD; in other cases an SLD was also included (*e.g.*, tumblr.com or co.za). IP address entries were ignored. Note, therefore, that any TLD which did not fall in the top 50 for any day did not get included in any roll-up count.

We analyzed the data as two sets, first looking at all TLDs and then looking at just the ccTLDs. The following table shows for the top 20 TLDs the average number of domains in those TLDs each month on the block list, highlighting US in green and the legacy gTLDs in yellow:

TLD	*Average # domains /month*
com	453,227
men	117,637
top	112,215
net	98,174
biz	96,749
us	74,226
info	59,163
gq	50,908
org	49,758
ml	47,786
gdn	39,709
cf	34,675
win	34,511
science	33,407
ru	32,638
ga	27,032
tk	26,197
link	24,052
loan	22,799
shop	22,366

[14] https://www.surbl.org

The following chart shows for the top 10 TLDs the average number of domains in those TLDs over each month, with US highlighted in red:

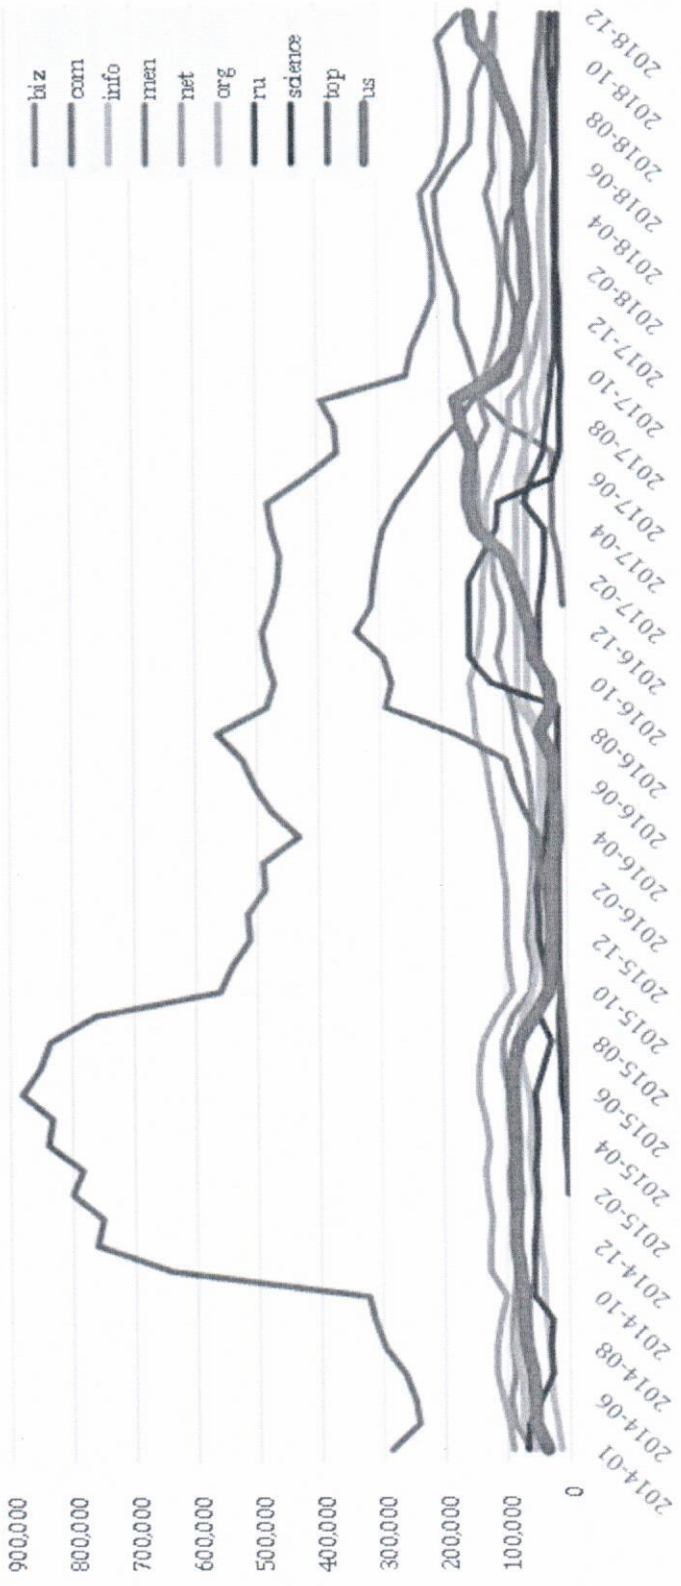
biz
com
info
men
net
org
ru
science
top
us
900,000
800,000
700,000
600,000
500,000
400,000
300,000
200,000
100,000
0
2014-01
2014-06
2014-08
2014-10
2014-12
2015-02
2015-04
2015-06
2015-08
2015-10
2015-12
2016-02
2016-04
2016-06
2016-08
2016-10
2016-12
2017-02
2017-04
2017-06
2017-08
2017-10
2017-12
2018-02
2018-04
2018-06
2018-08
2018-10
2018-12

The following table shows for the top 20 ccTLDs the average number of domains in those ccTLDs each month on the block list, highlighting US in green:

ccTLD	*Average # domains /month*
us	74,226
gq	50,908
ml	47,786
cf	34,675
ru	32,638
ga	27,032
tk	26,197
co	15,486
cc	9,209
pw	7,862
eu	6,840
me	4,784
in	3,273
cn	2,485
pl	2,304
uk	2,131
fr	1,782
ro	1,653
br	1,547
ws	1,436

The following chart shows for the top 10 ccTLDs the average number of domains in those ccTLDs over each month, with US highlighted in red:

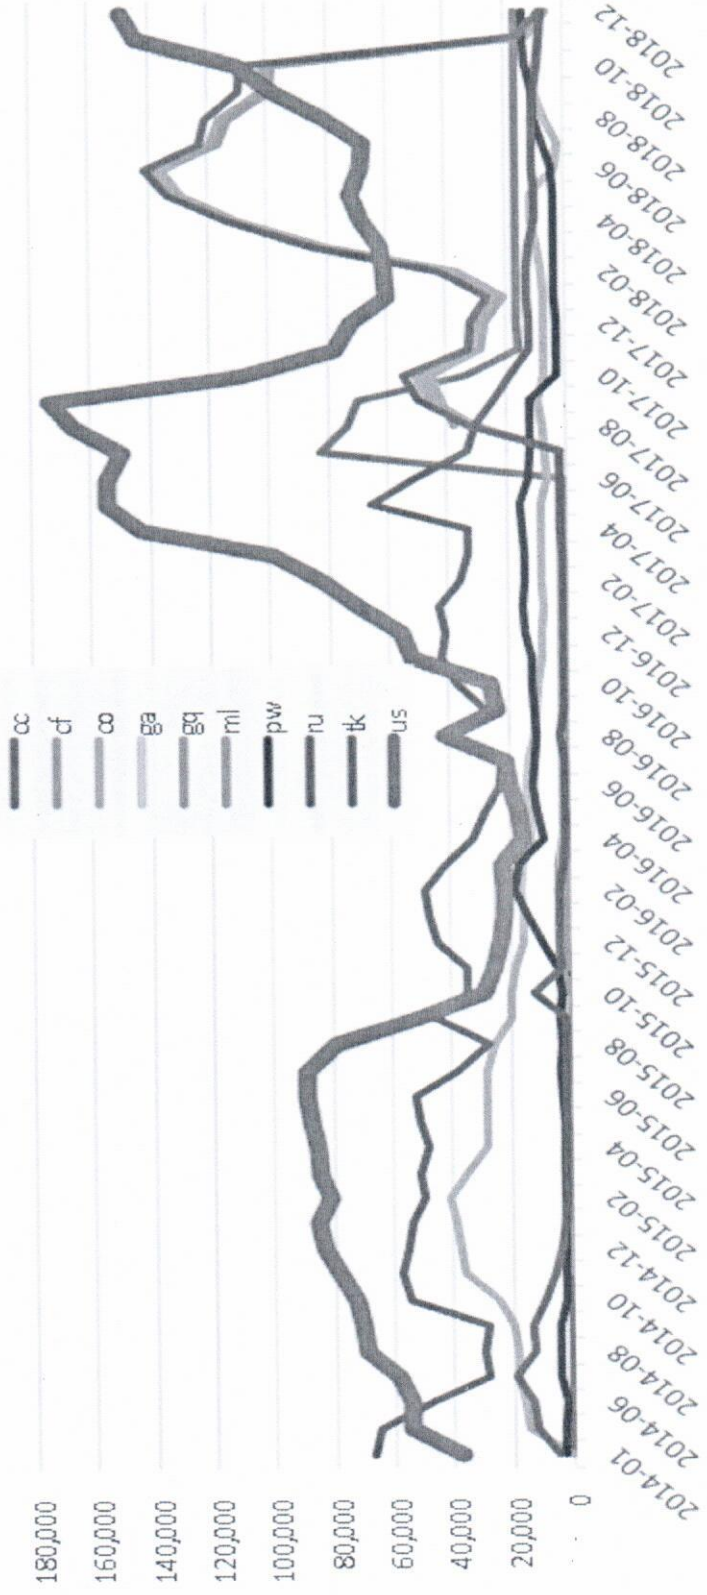
cc
cf
co
ga
gq
ml
pw
ru
tk
us
180,000
160,000
140,000
120,000
100,000
80,000
60,000
40,000
20,000
0
2014-01
2014-06
2014-08
2014-10
2014-12
2015-02
2015-04
2015-06
2015-08
2015-10
2015-12
2016-02
2016-04
2016-06
2016-08
2016-10
2016-12
2017-02
2017-04
2017-06
2017-08
2017-10
2017-12
2018-02
2018-04
2018-06
2018-08
2018-10
2018-12

The SURBL data show that

- the US registry has been abused by spammers persistently;
- the raw counts of spam domains frequently placed US at the top of the list of most abused ccTLDs; and
- prior to the introduction of new TLDs, US was frequently among the top five most abused among *all* TLDs.

Collectively, these views offer strong indications that US remains a targeted TLD for spam and related abuse activities.

Spamhaus Most Abused TLD data looks at abuse activity from a different perspective to derive a *badness index*. Spamhaus examines domains that are in active use and appear in mail feeds and related DNS traffic observed by their detection systems. For each TLD, Spamhaus counts the number of abuse domains they observe, counts the total number of active domains they observe and computes a *bad domains fraction* that takes into consideration (weighs) the TLD's size. On December 20, 2018, the US registry had a badness index of 19.4%, which translates into a bad score of 2.01. Sampling other TLDs, we see on the same date:

- loan = 90.1% bad (score 9.35)
- gq = 89.7% bad (score 10.09)
- tk = 77.8% bad (score 8.77)
- work = 77.2% bad (score 8.72)
- biz = 56.2% bad (score 6.04)
- us = 19.4% bad (score 2.01)
- info = 19.4% bad (score 2.01)
- com = 7.7% bad (score 1.00)
- net = 10.5% bad (score 1.16)
- eu = 4.5% bad (score 0.36)
- org = 4.3% bad (score 0.40)
- uk = 2.1% bad (score 0.17)
- pr = 1.9% bad (score 0.01)
- se = 0.1% bad (score 0.00)
- edu = 0.0% bad (score 0.00)

US is not among all worst abused TLDs: many new TLDs and ccTLDs have considerably higher badness indexes or scores. COM, NET, and ORG often serve as benchmarks for reputation (for no obvious reason other than that they are well-known and large), so comparing US to these legacy TLDs, the US registry's reputation has room for improvement.

6 Summary of Findings

While recognizing the limitations inherent in a tightly time-constrained study, we are confident that the findings described in this report are qualitatively valid. They are summarized here:

1. The PR zone is almost completely free of domain name abuse. The high cost and long processing time associated with registration in PR appear to effectively discourage abusive activity.
2. The US zone is persistently abused by spammers, and has been for many years. Raw counts of spam domains frequently place US at the top of the list of most abused ccTLDs, and prior to the introduction of new gTLDs, US was frequently among the top five most abused among *all* TLDs.
3. The US zone is clearly "in play" among botnet operators, but because our sample is derived from a small set of malware family DGAs, the extent of botnet-related abuse in the zone cannot reliably be determined. We cannot extrapolate from our sample to whether or not the US registry is routinely included in other DGAs, nor can we extrapolate to whether or not the same percentage of domains from US are generated when US *is* included.
4. Abuse domains in US—and their name servers, which are often hosted outside the geographic U.S.—frequently resolve to IP addresses in network allocations of ASNs that have significant abuse activities.
5. Based on the data we examined from Spamhaus, 95% of the abusive domains in US are registered through Namecheap, Inc. or its resellers. Namecheap is notorious among first responders and law enforcement as a registrar favored by spam gangs; our finding from the sample we studied suggests that the notoriety is well-deserved. DNC Holdings was a very distant second on the abusive domain registrar list.