



Phishing Landscape 2022

An Annual Study of the Scope and Distribution of Phishing

by

Greg Aaron

Lyman Chapin

David Piscitello

Dr. Colin Strutt

Interisle Consulting Group, LLC

19 July 2022



Table of Contents

Executive Summary.....	3
Introduction	5
Key Statistics and Trends	7
Phishing Activity.....	10
Time Elapsed between Domain Registration and Phishing	12
Phishing Distribution Across Top-Level Domains (TLDs).....	15
Phishing Distribution Across gTLD Registrars	20
Malicious Domain Name Registrations.....	24
Phishing Distribution Across Hosting Networks (Autonomous Systems)	33
List Coverage: The Phish That Get Away.....	38
Targeted Brands.....	40
Cryptocurrency Phishing.....	42
Abuse of Subdomain Service Providers	48
Use of Internationalized Domain Names (IDNs) for Phishing.....	50
Appendix A: Identification of Phishing Attacks.....	51
Appendix B: Distinguishing Maliciously Registered Domain Names from Compromised Domains.....	52
Appendix C: Data Sources and Methodologies.....	54

Executive Summary

Phishing remains a significant threat to millions of Internet users. Phishing attacks lure victims to a web site that appears to be run by a trusted entity, such as a bank or a merchant. The web site, however, is a deception, and the site's content is designed to persuade a victim to provide sensitive information.

Our goal in this study was to capture and analyze a large set of information about phishing attacks, to better understand how much phishing is taking place and where it is taking place, and to see if the data suggests better ways to fight phishing. To do so we determined when attacks occur and how quickly phishers act. We studied where phishers get the resources that they need to perpetrate their crimes — such as where they obtain domain names, and what web hosting is used. This analysis can identify where additional phishing detection and mitigation efforts are needed and can identify vulnerable providers. We also report on the wide range of brands targeted by phishers, and how often they take advantage of the unique properties of internationalized domain names (IDNs).

To assemble a deep and reliable set of data, we collected over three million phishing reports from 1 May 2021 to 30 April 2022 from four widely used and respected threat intelligence providers: the Anti-Phishing Working Group (APWG), OpenPhish, PhishTank, and Spamhaus. From that data we identified 1,122,579 unique phishing attacks.

For this 2022 landscape study, we also analyzed phishing reports collected over a two-year period, a total of nearly five million phishing reports collected from 1 May 2020 to 30 April 2022. By adding biennial measurements and analyses, we began to consider questions like, “How do the yearly trends for phishing attacks, domain names used for phishing, etc., compare to biennial trends?” and “Are phishers “doing business” at the same registry, registration, or hosting services year after year?”

*Our studies identify
opportunities
to effect change*

Our yearly and now, biennial, studies illustrate that phishing is a highly profitable, constantly evolving and expanding industry. Phishing leverages Internet resources, exploits vulnerable technologies, and takes advantage of policy and legislative regimes that are siloed and, by our measurements and analyses, ineffective. From our studies, we observe that

- **Silos frustrate phishing response:** The naming, addressing, and hosting ecosystem exploited by phishers (and cyberattackers generally) is encumbered by vertically isolated (“siloed”) policy and mitigation regimes.
- **Registrars, registries, and hosting providers can take action:** Registries and registrars should identify, “lock”, and suspend domains reported for phishing, and hosting and cloud service providers should remove phishing content or shut down accounts where phishing occurs; all parties should be more responsive to abuse complaints, especially for cybercrimes such as phishing, but they must begin to do so in a more coordinated and determined manner.
- **Regulation and legislation could help:** Changes to or introduction of policy or regulation may be necessary to effectively mitigate phishing. Obliging operators to validate the identity of users and customers, coupled with agreement on a common definition of lawful access that acknowledges the role that the private sector plays in combatting cybercrime, could reduce both the incidence of phishing and the difficulty of responding to it.

Findings

1,122,579 phishing attacks



A 61% increase over the previous yearly study
Monthly reported phishing attacks more than doubled since May 2021

853,987 domain names reported for phishing



A 72% increase over the 2021 study

588,321 malicious domain name registrations



Domains registered by phishers for phishing attacks
increased 83% over 2021 study

New gTLDs domains used disproportionately for phishing

buzztopxyzsho
pliveonlinesitew
orkclubsupport
cyoutokyobar

34% of phishing domains in .COM and .NET
ccTLDs run by Freenom represent 40% of all ccTLD phishing domains

US hosting networks attract phishers



86% of phishing attacks reported against the top 10
hosting networks are using IP addresses in the US

Phishers are targeting more brands



Meta Platforms (Facebook, WhatsApp), Amazon,
Microsoft (Outlook), and Apple

257% increase in cryptocurrency phishing



Wallet brands most targeted
79% of gTLD cryptocurrency phishing domains
were maliciously registered

13% of all phishing attacks used subdomain services



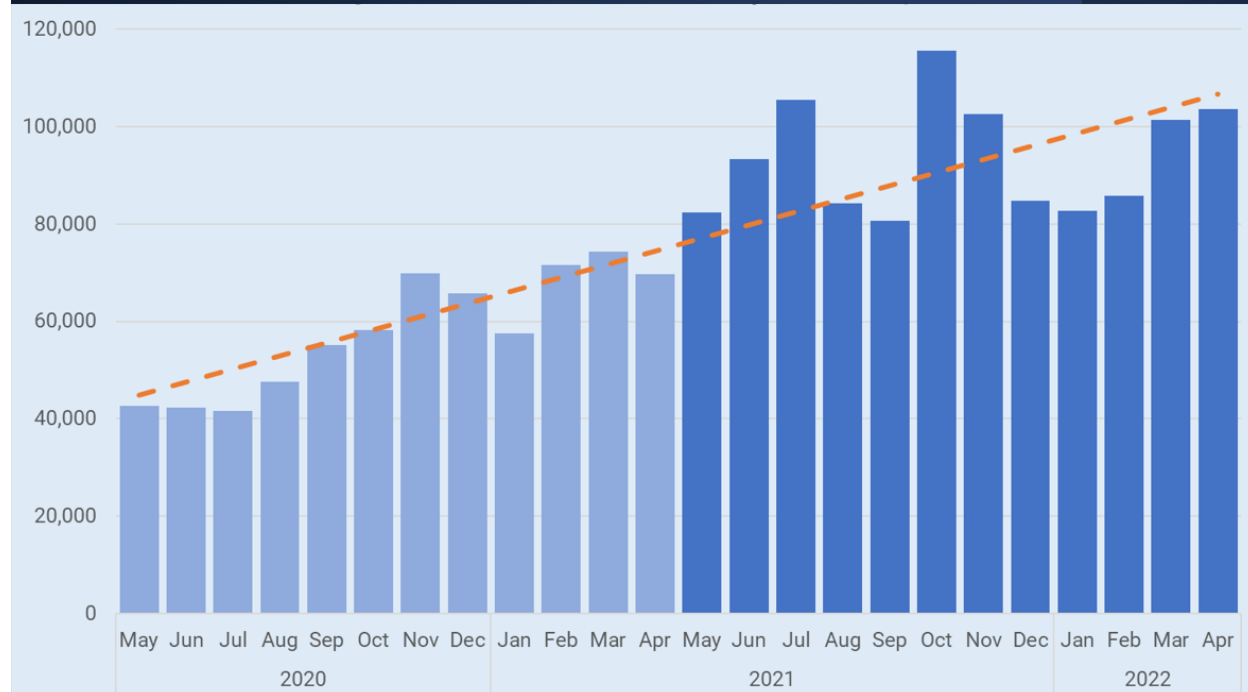
82% hosted on domains operated by
top ten subdomain service providers

Introduction

For this Phishing Landscape 2022 study we analyzed over 3 million phishing reports from phishing data feeds representing more than 1.1 million phishing attacks. We examined phishing activity during the May 2021 – April 2022 period. We looked at attack activity, daily and weekly to determine whether phishers found certain days of the week more opportunistic than others. We compared the dates when our study set of domain names were registered to the dates when the domains were reported for phishing, to understand how phishers prepare for attacks.

We examined where phishing attacks occurred among Top-Level Domain registries, TLD registrars, and hosting providers. We ranked these operators according to raw counts and metrics. We distinguish phishing attacks where domain names were registered by phishers from phishing attacks that were hosted on compromised domains or web sites. This distinction is important because it indicates where additional phishing detection or mitigation efforts could be applied most effectively, and importantly, which operator (registry, registrar, hosting provider) is best positioned to implement these. We completed this study by reporting on brands most targeted by phishers, the role subdomain resellers play in phishing, and how phishers have added cryptocurrencies to their financial fraud target lists.

Monthly number of phishing attacks reported more than doubled since 1 May 2020



The study shows that **the number of phishing attacks reported each month has more than doubled since 1 May 2020 and continues to trend upwards.**

- The trendline shows a 52% growth in phishing attacks over the 24-month period.
- We observed noteworthy spikes in July, October, and November 2021.
- In the most recent months, the number of phishing attacks per month is approximately 2½ times greater than in May 2020.

Throughout the study, we provide year-over-year comparisons of phishing activity. By examining phishing behavior over a 24-month period, from May 2020 to April 2022, we identified domain name registration or hosting patterns that persist over time. From the longitudinal analyses afforded by a multi-year data set, we were able to illustrate (through trendlines) prevailing directions of various phishing metrics.

*Monthly phishing attacks
doubled since May 2021*

The statistics that we present in this report include both absolute metrics (*e.g.*, the number of domain names registered in a particular TLD that appear on a blocklist) and relative metrics (*e.g.*, a phishing score, representing the number of those domain names as a percentage of the total number of domains registered in that TLD). Attention to this distinction is critical to understanding and properly interpreting our analyses and findings.

Key Statistics and Trends

To assemble a deep and reliable set of data, we collected over 3 million phishing reports for a one-year period, from 1 May 2021 through 30 April 2022. We obtained data from four widely used and respected threat data providers: the Anti-Phishing Working Group (APWG), OpenPhish, PhishTank, and Spamhaus (see the section *Phishing Data Sources* in *Appendix C: Data Sources and Methodologies*). We augmented this data set with the data we used for our 2021 study to present year-over-year (comparative) measurements.

Key statistics for this study period (May 2020 to April 2022) compared with the corresponding statistics from the previous study period (May 2021 to April 2022) are:

Measurement	From May 2020 to April 2021	From May 2021 to April 2022	Change
Total number of phishing attacks	695,823	1,122,579	+426,756
Unique domain names reported for phishing	497,949	853,987	+356,038
Maliciously registered phishing domains	322,145	588,321	+266,176
Top-level domains where phishing domains were reported	623	660	+37
Registrars with domains under management reported for phishing	1,008	1,523	+515
Hosting networks where phishing web sites were reported	4,110	4,159	+49

Table 1 Phishing Activity: Key Statistics, Year Over Year

Phishing attacks increased by 61% (year-over-year). The total number of phishing attacks is a sum of the attacks that we identified using the methodology we describe in *Appendix A: Identification of Phishing Attacks*.

Unique domain names reported for phishing increased by 72%. This finding is based on our determination of “the first occurrence of a domain name in a phishing report”. We use this number to account for domains which appeared in multiple phishing attacks in multiple quarters during the yearly period.

Maliciously registered domain names increased by 82%. This finding is based on our determination that a domain name was purposely registered by a phisher to perpetrate a phishing attack using the methodology we describe in *Appendix B: Distinguishing Maliciously Registered Domain Names from Compromised Domains*.

We obtained the numbers of TLDs, TLD registrars, and hosting networks where we observed phishing by counting each operator that appeared in the yearly study data.

Table 2 shows the quarterly key statistics reported at the Cybercrime Information Center.¹

Measurement (by Quarter)	May 2021 – July 2021	August 2021 – October 2021	November 2021 – January 2022	February 2022 – April 2022
Total number of phishing attacks	285,727	296,432	289,996	303,348
Phishing attacks associated with malicious domain registrations	161,294	216,038	155,729	156,353
Unique domain names reported for phishing	222,982	235,697	222,059	223,324
Malicious registered phishing domains	147,198	128,035	140,696	147,036
Top-level domains where phishing domains were reported	506	526	534	519
Registrars with domains under management reported for phishing	473	609	527	453
Hosting networks where phishing web sites were reported	2,329	2,537	2,554	2,617

Table 2 2022 Quarterly Key Statistics, May 2021 – April 2022

To obtain yearly measurements for TLDs or TLD registrars, we performed a de-duplication of domain names or addresses that appeared in more than one quarter. The sum of the four quarterly numbers in Table 2 will thus not be the same the cumulative numbers in Table 1.

Trends of Key Statistics

The trendline in Figure 1 illustrates that phishing attacks, and unique domains reported for phishing all trended up over the two-year period. Phishing attacks and unique domains reported for phishing track along a common trendline until November 2021. We observed a meaningful increase in phishing attacks hosted on subdomain service providers, and examine this later, in the section *Abuse of Subdomain Service Providers*.

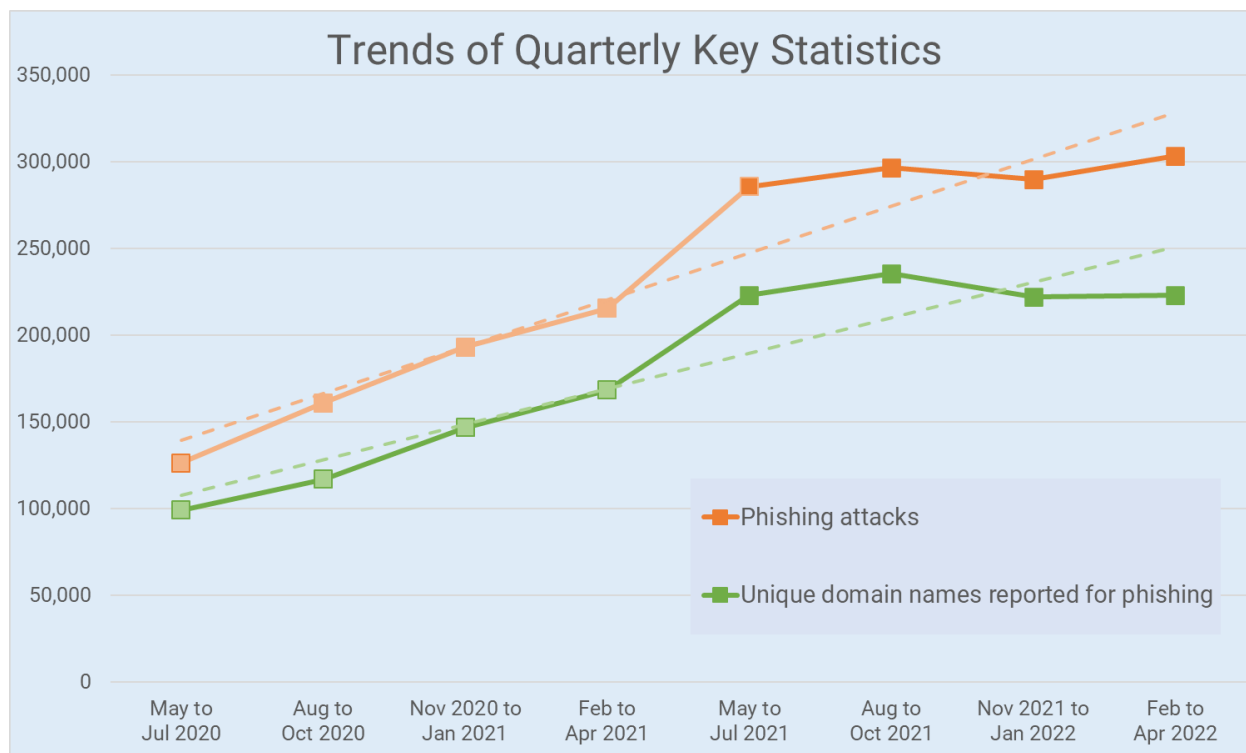


Figure 1 Phishing Domains, Attacks, and Unique Domains Reported for Phishing Trended Up

Readers should exercise care in interpreting the graph of unique domain names reported for phishing and not rush to conclude that phishing decreased in late 2021 and early 2022. This graph only illustrates a drop in the raw count of domain names reported for phishing.

By comparing this graph against the number of phishing attacks, we see that **phishing attacks decreased but not commensurately with domain names reported**. This may indicate a shift in phisher behavior, to using subdomain reseller services more often, or placing multiple phishing websites on one domain.

While the number of phishing domains is relatively flat the number of phishing attacks has continued to increase

Phishing attacks is the most significant measurement to target organizations and victims. Other measurements that identify the resources that phishers employ in their attacks are helpful because they identify criminal tools of trade and from where criminals acquire them. Efforts to reduce criminal use of domain names are important and beneficial but applied without complementary efforts to mitigate phishing on cloud services or hosting abuse (web site, file server, user account and host system vulnerabilities) provides criminals with numerous alternative means to phish. The lamentable state of phishing indicates that more collaboration across the operators whose resources are “tools of trade” is necessary.

Phishing Activity

We define a phishing attack as a phishing site that targets a specific brand or entity. This is a basic measure of how much phishing activity is being observed. In *Appendix A: Identification of Phishing Attacks*, we describe how we determined if multiple phishing reports (or more than one URL) refer to the same phishing site. When we do find this kind of activity, we eliminate duplicates to yield a count of distinct (unique) phishing attacks.

In Figure 2 we show the daily counts of phishing attacks for this study period. The figure shows that **phishing activity continues to be a generally chronic (persistent) problem.**

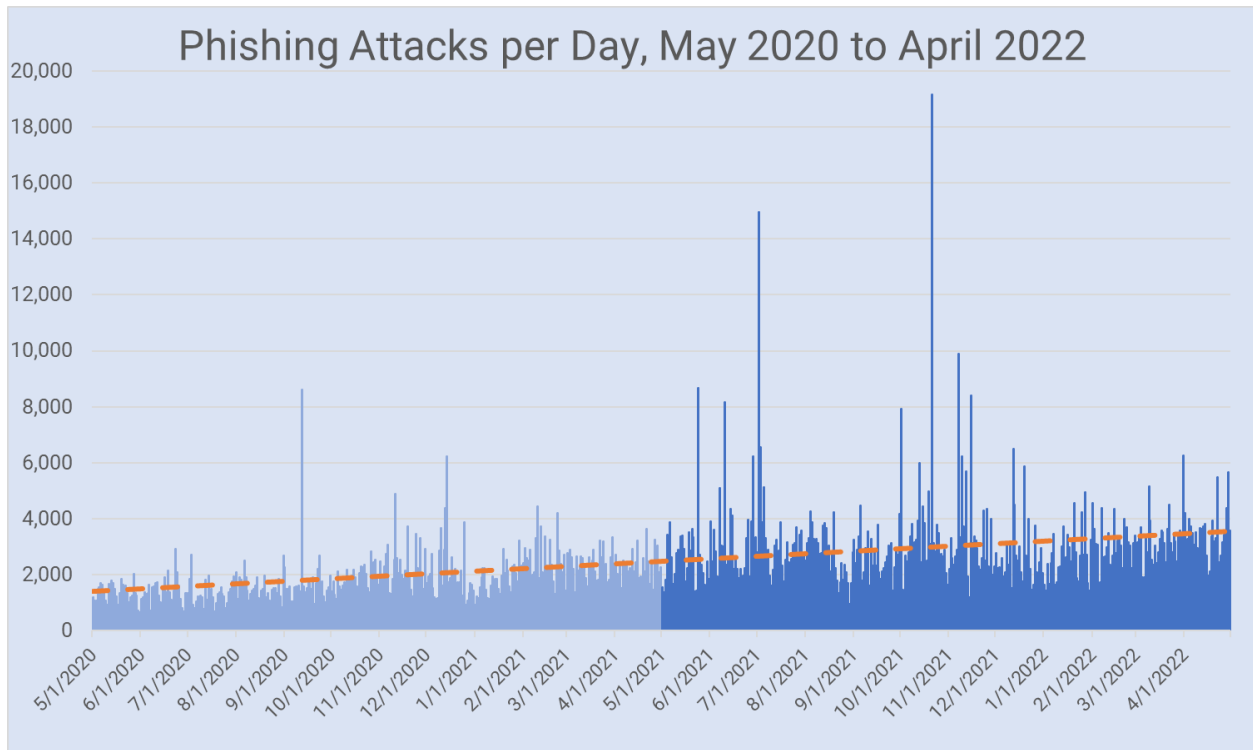


Figure 2 Phishing Attacks per day, May 2020 to April 2022

Instances of acute, daily phishing activity increased significantly

The spikes in the graph of daily phishing activity merit discussion. Daily phishing activity appears to have hovered between 2,000 and 5,000 phishing attacks during the 24 months illustrated in the graph. During the May 2020 – May 2021 period, we observed only two days where phishing activity exceeded 5,000 attacks. By comparison, we observed 20 instances from June 2021 to April 2022 where phishing activity exceeded the 5,000 phishing attacks. We observed the most acute phishing activity during the periods June – July 2021 and October – December 2021.

In Figure 3, we depict when phishing attacks occurred, by day of the week.

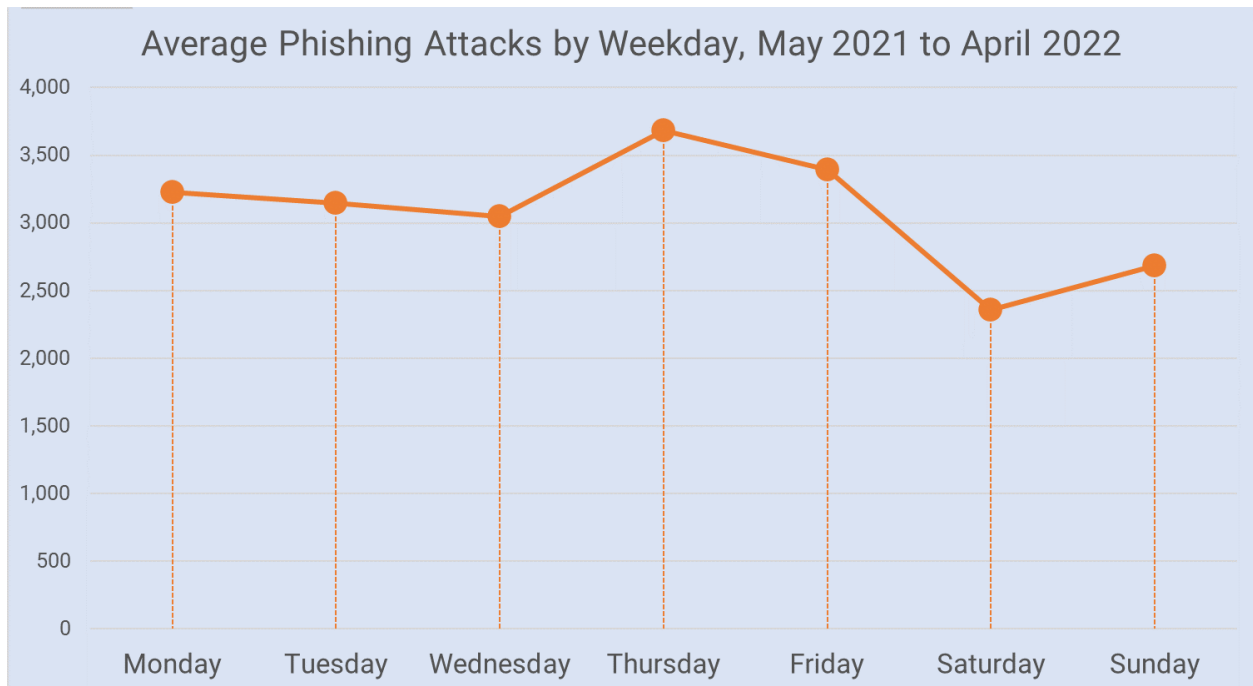


Figure 3 Phishing Attacks by Day of Week, 1 May 2021 to 30 April 2022

In past studies, we observed that phishing activity had been highest in the Monday through Wednesday period, and that blocklistings peaked around Wednesdays. Phishers advertise their attacks via spam mail at what they believe to be an optimal time, *i.e.*, when people check their work and personal email on returning to work or after the weekend is over. We also note that there is a delay between when an attack begins and when it is blocklisted – essentially, attacks do peak a bit earlier than reported.

For this study period, we observed a slight shift in this pattern. **We saw more reports of phishing on Thursdays and Fridays than we have seen in earlier years.** We cannot ascertain from our data whether this shift is a change in phisher behavior, whether phishing investigators are reporting phishing later than they did before, or because they are encountering impediments in their work.

Time Elapsed between Domain Registration and Phishing

We analyzed how many days elapsed between when a domain name was registered and when that domain was associated with a phishing attack by one of the phishing data feeds. For this analysis, we only included domain names for which we were able to obtain a creation date from domain registration data.

Figure 4 shows that during the current study period (chart on right), **41% of domains reported for phishing were used within 14 days following registration and that the majority of these were reported within 48 hours**. This suggests that opportunities exist for TLD registries or registrars to identify and preemptively block suspicious registration attempts. 76% of domain names associated with a phishing attack were reported within the first year of registration. During the prior period (chart on left), 57% were used within 14 days, and 84% were used within the first year of registration.

41% of domains reported for phishing were used within 14 days following registration

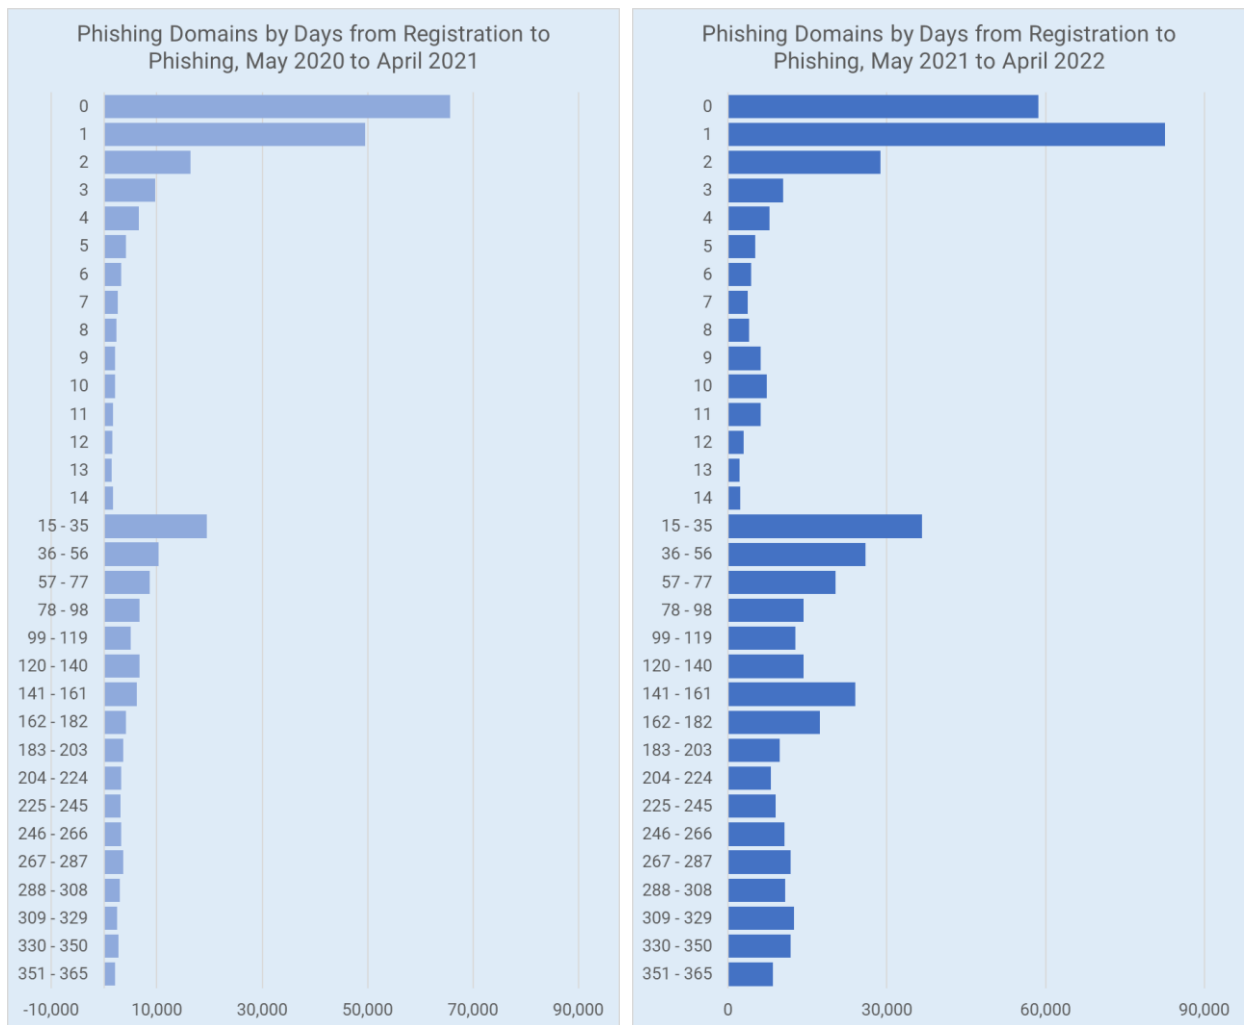


Figure 4 Year-over-year Comparison of Phishing Domains: Days from Domain Registration to Reported for Phishing

Our findings continue to reinforce the conventional wisdom that when **phishers register domains, they tend to use them quickly to avoid detection**. This is consistent with research concerning the risk associated with newly registered domain names.^{2,3}

In the section *Malicious Domain Name Registrations* we explain how we classified phishing domains as maliciously registered – a domain name purposely registered by a criminal to carry out a phishing attack – or compromised – a legitimate web site that was compromised and used for phishing.

Figure 5 shows that, for the yearly study period (chart on the right), 64% of domain names that we classified as malicious were reported for phishing within 14 days following registration and 98% of domain names that we classified as malicious were reported for phishing within the first year of registration. During the prior period (chart on left), 89% were used within 14 days, and 98% within the first year of registration.

The drop from 89% to 64% is significant. Many security systems use domain age as a scoring factor, *e.g.*, they assign higher risk to newly registered domains. Phishers may be registering domains and then warehousing or “aging” them for a few weeks before launching phishing attacks on those domains.

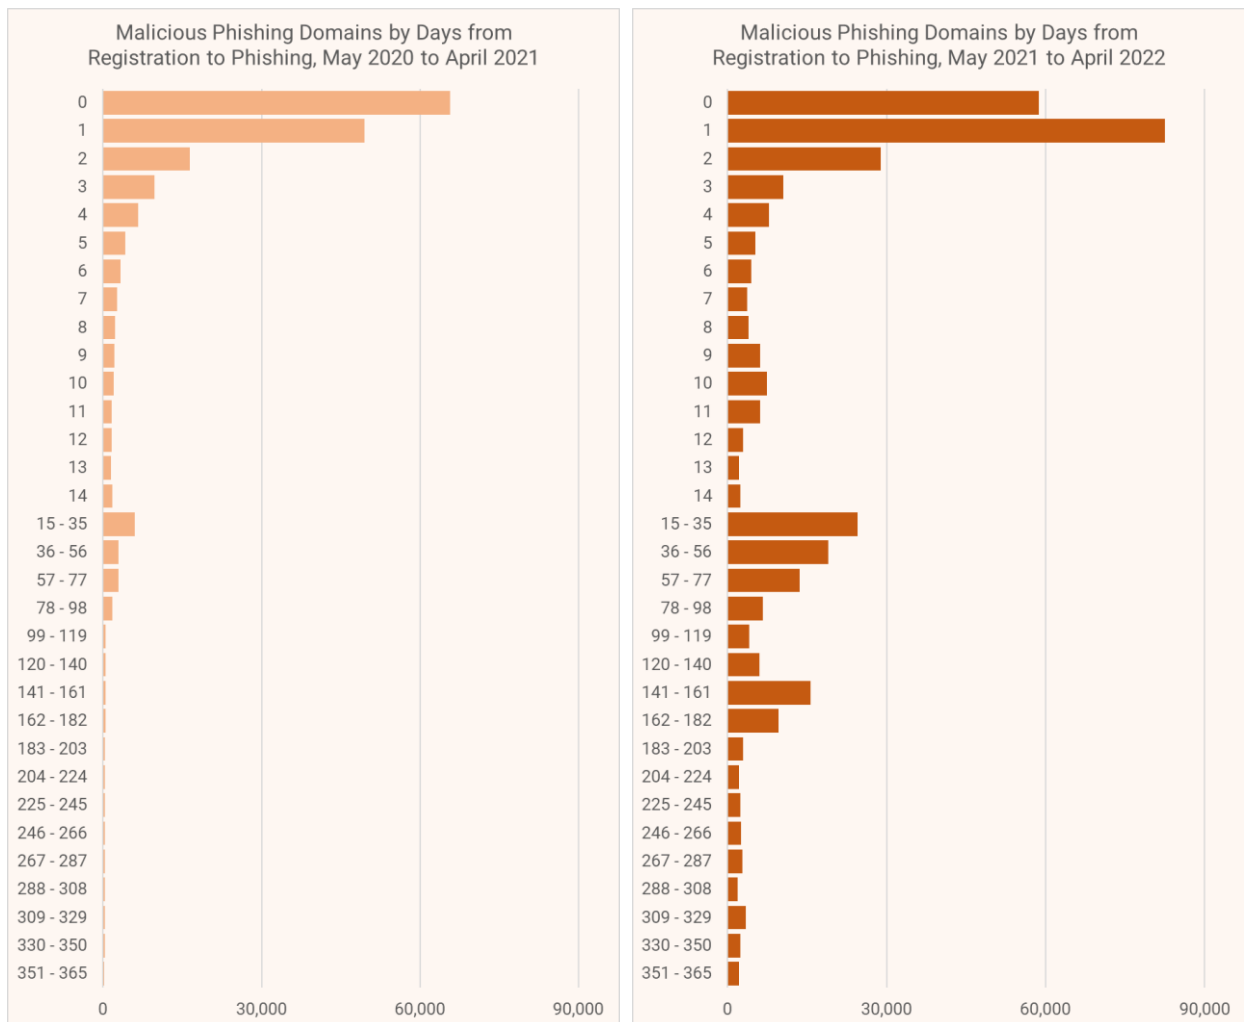


Figure 5 Malicious Phishing Domains, Year Over Year

The yearly as well as year-over-year measurements of malicious phishing domain registrations indicate that phishers are either paying for their domain names with legitimate means, or that the payment processors and the registrars are not recognizing many suspicious or fraudulent transactions at the time of transaction or in the days thereafter.

It also appears that domain registrars are not taking advantage of tools that would allow them to recognize maliciously registered domains in a short time immediately after registration. (These include checks for inaccurate contact data and checks that can identify – or label as suspicious – a domain that was purposely registered for phishing from a domain from a domain that was registered for a legitimate purpose. See *Appendix B: Distinguishing Maliciously Registered Domain Names from Compromised Domains.*)

Our data also show that a small number of domains appear to be maliciously registered but that they were flagged for phishing well past the first year of registration, in some cases several years after registration.

Phishing Distribution Across Top-Level Domains (TLDs)

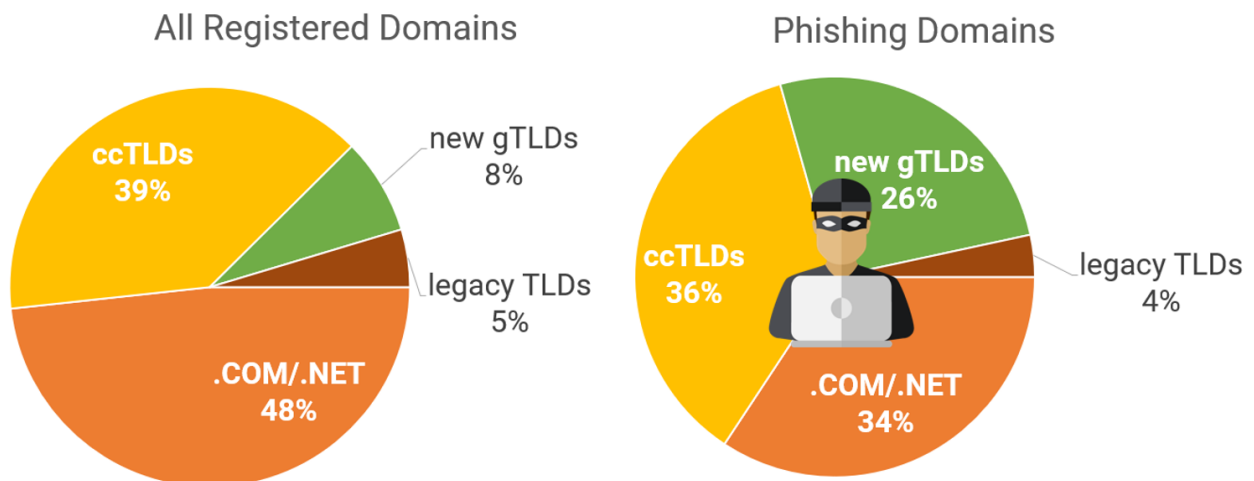
We used Domain Tools⁴ as our source for determining TLD domains under management (DUM).

According to Domain Tools, at the end of March 2022, there were over 358 million registered domains.

For our studies, we divide the overall domain name space into four categories:

- the .COM and .NET registries are operated by Verisign which represented 48% of the domains in the world,
- country-code domains (ccTLDs) which represented 39% of the domains,
- the legacy generic TLDs (those other than .COM and .NET and introduced before 2013, *e.g.*, .ORG, .BIZ, .INFO) which represented 5% of the domains, and
- the new gTLDs (nTLDs) introduced from 2014 to the present (*e.g.*, .ONLINE, .XYZ, .ICU) which represented the remaining 8% of the domains.

We analyzed the phishing domains and attacks to see how they were distributed across the top-level domains.



While we observed phishing in 660 TLDs during the yearly study period, we note that phishing activity continues to be concentrated in just a few namespaces.

34% of all domains reported for phishing were in .COM and .NET. This percentage is significantly smaller than the combined market share (48%) of those TLDs.

New gTLDs are attracting phishers

Domains in the new gTLDs were used disproportionately for phishing and their percentage of phishing domains continues to grow.

In our *Phishing Landscape 2021* study⁵, we reported that in June 2020, new TLDs represented 9% of domain names in the world but 18% of domains used for phishing. We next observed that the new TLDs market share decreased to 6% in March 2021 but reported phishing domains in the new TLDs again increased to 21% during our yearly period. For this study,

the new TLDs' market share grew to 8% but reported phishing domains grew to 26%.

36% of domains used for phishing were in ccTLDs. This is roughly in line with the 39% of the domain name market share represented by ccTLDs. However, setting aside phishing on the Freenom Top-level Domains, the other ccTLDs suffered far less phishing than might be expected based on market share.

Phishing in the ccTLD category continues to be artificially swollen by phishing domains reported in five commercialized ccTLDs run by Freenom (.TK, .ML, .GA, .CF, .GQ), which offers free domain name registrations. Freenom's TLDs represented 19% of the ccTLD-registered domain names but represented 40% of all ccTLD phishing domains reported. Freenom's TLDs represented 8% of all registered domains yet accounted for 14% of phishing domains reported in all TLDs.

Freenom ccTLDs are attractive to phishers

The remaining 4% of phishing was in the legacy TLDs other than .COM and .NET, which is roughly in line with their market share.

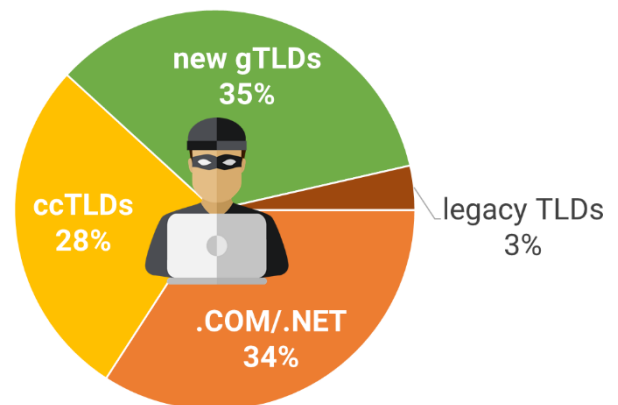
Malicious Domain Registrations Across the Domain Name Space

When we studied where phishers registered domains purposely for phishing, we saw some meaningful differences in percentages from where phishing domains were reported.

Together, .COM and .NET represent 48% of the domain name space overall. We determined that 41% of reported phishing domains in .COM and .NET were purposely registered for phishing. The remaining domain names reported for phishing were likely victims of some form of compromise by phishers, who placed phishing pages on the sites without the owner's knowledge.

New TLDs continued to present attractive registration opportunities for phishers. We determined that 35% of reported phishing domains in the new TLDs were purposely registered for phishing, which is 4.4 times the segment's market share. In our 2021 study, the percentage of malicious registrations in the new TLDs was 3.8 times the segment's market share.

Maliciously Registered Phishing Domains



Ranking of TLDs by Phishing Domains Reported

For the 2022 period, two legacy TLDs, five ccTLDs, and three new TLDs were among the TLDs having the most phishing domains reported. .COM and .NET held the same rankings in 2022 as they did in 2021.

2021 Rank	TLD	Registry Operator	Phishing Domains Reported	2022 Rank	TLD	Registry Operator	Domains in TLD	Phishing Domains Reported
1	com	Verisign	260,636	1	com	Verisign	159,902,632	277,728
2	tk	Freenom	40,002	2	cn	CNNIC	8,980,611	103,869
3	xyz	XYZ.COM	27,532	3	shop	GMO Registry	1,040,404	47,747
4	ml	Freenom	27,284	4	xyz	XYZ.COM	4,130,573	38,604
5	ga	Freenom	21,657	5	tk	Freenom	5,041,535	37,300
6	cf	Freenom	19,187	6	ml	Freenom	5,344,979	28,318
7	gq	Freenom	16,168	7	ga	Freenom	7,049,929	25,717
8	cn	CNNIC	16,052	8	cf	Freenom	5,383,367	17,747
9	top	Jiangsu Bangning	15,129	9	bar	Punto 2012 SAPI	281,575	15,826
10	net	Verisign	14,398	10	net	Verisign	13,170,783	15,083

Of the five ccTLDs that were ranked among the top ten, four (.TK, .ML, .GA, and .CF) are operated by Freenom, a company in the Netherlands that offers free domain registrations in these ccTLDs. These ccTLDs were also in the Top 10 in our 2021 Landscape study.

Two new TLDs, .SHOP and .BAR, joined the Top 10, replacing .TOP and Freenom's .GQ.

Ranking of TLDs by Scoring Metrics

The more phishing domains in a space or portfolio controlled by one company, the greater the opportunity (and need) for that company to take effective anti-abuse measures — including measures to find and suspend malicious phishing registrations early.

Scoring metrics allow for comparisons between TLDs of different sizes. In the quarterly phishing activity published at the Cybercrime Information Center, the metric “Phishing Domains per 10,000” is used to compare whether a TLD has a higher or lower incidence of phishing relative to others. This is a ratio of the number of domain names used for phishing in a TLD to the number of registered domain names in that TLD. We call this metric **TLD Phishing Score**:

$$\text{TLD Phishing Score} = \left(\frac{\text{number of phishing domains}}{\text{total number of domains under management in the TLD}} \right) * 10,000$$

In this report, we use a similar metric to measure the prevalence of phishing in each TLD for the period beginning 1 May 2020 and ending 30 April 2021 (365 days). We take the sum of the four quarters of unique phishing domains reported and divide by the average of the domains under management per TLD for each of the four quarters. We call this metric **Yearly TLD Phishing Score**:

$$\text{Yearly TLD Phishing Score} = \left(\frac{\text{number of unique phishing domains reported across the year}}{\text{number of domains under management per TLD}} \right) * 10,000$$

The numbers for the TLD Phishing Score and for the Yearly TLD Phishing Score are not directly comparable as the latter cover a longer time period.

The following table shows the highest ranking TLDs by Yearly Phishing Domain Score and, for each, shows the percentage of phishing domains that are maliciously registered in that TLD.

Rank	TLD	Registry	Domains in TLD	Yearly Phishing Domain Score	TLDs with high yearly phishing scores also have very high percentages of malicious phishing domain registrations	Rank	TLD	% Malicious Phishing Domains
1	support	Binky Moon	31,288	595		1	support	91%
2	bar	Punto 2012 S.A.P.I.	281,575	562		2	bar	99%
3	shop	GMO Registry	1,040,404	459		3	shop	97%
4	work	Minds + Machines	303,321	340		4	work	98%
5	live	Identity Digital	596,948	208		5	live	88%
6	buzz	DOTSTRATEGY	513,201	196		6	buzz	98%
7	casa	Minds + Machines	30,290	178		7	casa	88%
8	finance	Cotton Cypress	52,709	173		8	finance	74%
9	sbs	Shortdot SA	34,030	164		9	sbs	94%
10	fyi	Identity Digital	40,290	150	10	fyi	95%	

These TLDs also have higher percentages of domain names purposely registered for phishing than the 69% of overall reported phishing domains that we determined were registered maliciously, by phishers.

High yearly phishing scores erode confidence in a TLD

As was the case in our 2021 Study, **TLDs with the highest yearly phishing scores in our 2022 study are all new TLDs**. .BAR, .BUZZ, .CASA, and .LIVE appear in both yearly studies, and all but .CASA's scores worsened year over year.

TLD	Yearly Phishing Domain Score	
	2021 Study	2022 Study
bar	233.3	562.1
casa	199.7	178.0
buzz	190.5	195.5
live	160.6	208.1

A person is more likely to encounter a dangerous domain when they click on a hyperlink in an email message or visit a web site address that contains a domain name registered in a TLD with a high yearly phishing score. Thus, when faced with a 74% or higher likelihood of exposing a user to a maliciously registered phishing domain, risk-averse organizations are likely to blacklist TLDs with persistently high yearly phishing scores *in their entirety*.⁶

The gTLDs that were most used by phishers to purposely register phishing domains (malicious phishing domain registrations) were:

TLD	Phishing Domains	Malicious Phishing Domain Registrations	Percent of Phishing Domains Determined to be Maliciously Registered ▼
bar	15,826	15,670	99%
work	10,315	10,150	98%
buzz	10,031	9,800	98%
shop	47,747	46,379	97%
xyz	38,604	35,665	92%
live	12,420	10,981	88%
top	14,758	12,498	85%
info	13,447	11,382	85%
com	277,727	191,660	69%
net	15,083	9,091	60%

Phishing Distribution Across gTLD Registrars

Of the 3,326,607 phishing reports that we collected during the 1 May 2021 to 30 April 2022 period, 14,489 contained IP addresses, and of these, 2,451 were unique addresses. The remaining phishing reports contained domain names.

Phishers acquire domain names by registering names purposely for phishing. They also break into the domain name management accounts or the hosting accounts of domain name owners. The table below shows that phishers purchase and manage domain names through many gTLD registrars.

Some gTLD registrar services, pricing, or practices appear to be more attractive to phishers than others. We consider this phenomenon in the section *Malicious Domain Name Registrations and gTLD Registrars*.

Ranking of gTLD Registrars by Phishing Domains Reported

Nineteen (19) registrars had portfolios with 10,000 or more gTLD domains reported for phishing from 1 May 2021 to 30 April 2022. In the prior yearly study, there were only four registrars with more than 10,000 gTLD domains reported for phishing (NameCheap, NameSilo, GoDaddy, and PublicDomainRegistry).

Rank	Registrar	Registrar IANA ID	gTLD Domains under Management	Phishing Domains Reported ▼
1	NameCheap	1068	13,645,340	88,643
2	GoDaddy.com	146	66,087,039	44,160
3	NameSilo	1479	4,403,551	42,489
4	DNSPod	1697	1,387,872	30,778
5	ALIBABA.COM SINGAPORE	3775	1,677,681	27,538
6	PublicDomainRegistry	303	4,916,665	21,948
7	REG.RU LLC	1606	726,674	14,472
8	Wild West Domains	440	2,962,240	12,707
9	Wix.com	3817	2,323,890	11,287
10	eNom	48	4,657,282	10,101

Six of the top-ranked gTLD registrars – Namecheap, GoDaddy, Namesilo, Wild West Domains, Wix, and eNom – are headquartered in the United States. DNSPod and Alibaba are headquartered in China. Public Domain Registry is headquartered in India and Reg.RU in Russia. The “lock and suspend” legislation under consideration in the United States⁷, if adopted, could affect future gTLD rankings.

A comparison of gTLD registrar rankings shows a general year-over-year increase in the numbers of phishing domains reported at the top registrars.

2021 Rank	Registrar	Phishing Domains Reported ▼	2022 Rank	Registrar	gTLD Domain Registrations	Phishing Domains Reported ▼
1	NameCheap	79,118	1	NameCheap	13,645,340	88,643
2	NameSilo	37,067	2	GoDaddy.com	66,087,039	44,160
3	GoDaddy.com	35,150	3	NameSilo	4,403,551	42,489
4	PublicDomainRegistry	19,065	4	DNSPod	1,387,872	30,778
5	Tucows Domains	9,972	5	ALIBABA Singapore	1,677,681	27,538
6	Wild West Domains	8,582	6	PublicDomainRegistry	4,916,665	21,948
7	Google LLC	8,413	7	REG.RU LLC	726,674	14,472
8	ALIBABA Singapore	7,883	8	Wild West Domains	2,962,240	12,707
9	Onamae.com	7,276	9	Wix.com	2,323,890	11,287
10	eNom	6,754	10	eNom	4,657,282	10,101

Ranking of gTLD Registrars by Scoring Metrics

When gross numbers alone are used to compare registrars, findings can be biased against registrars with large address delegations. In the quarterly phishing activity published at the Cybercrime Information Center, the metric “Phishing Domains per 10,000” is used to compare whether a gTLD registrar has a higher or lower incidence of phishing relative to others. This is a ratio of the number of domain names used for phishing to the number of registered domain names under management at that gTLD registrar. We call this metric **gTLD Registrar Phishing Score**:

$$\text{gTLD Registrar Phishing Score} = \frac{\text{(number of phishing domains / domains under management at gTLD Registrar)} * 10,000$$

For this report, as we did for TLDs, we use a similar metric to measure the prevalence of phishing in each gTLD registrar for the period 1 May 2020 to 30 April 2021. Here, we take the sum of the four quarters of unique phishing domains reported and divide by the average of the domains under management per gTLD registrar for each of the four quarters. We call this metric **Yearly gTLD registrar Phishing Score**:

$$\text{Yearly gTLD Registrar Phishing Score} = \frac{\text{(number of unique phishing domains reported across the year / number of domains under management at gTLD Registrar)} * 10,000$$

Note that the calculation of these two metrics yields different results (simply put, we use different inputs for the numerators in the division); in particular, one cannot draw any conclusion by comparing the scores from a quarterly phishing score against an annual phishing score. Instead, we encourage comparisons of quarterly phishing scores over time, as well as annual phishing scores over time.

The ranking of gTLD registrars by yearly phishing domain score is:

Rank	Registrar	Registrar IANA ID	gTLD Domains under Management	Phishing Domains	Yearly Phishing Domain Score ▼
1	DNSPod	1697	1,387,872	30,778	221.76
2	OwnRegistrar	1250	313,472	6,632	211.57
3	REG.RU	1606	726,674	14,472	199.15
4	ALIBABA.COM SINGAPORE	3775	1,677,681	27,538	164.14
5	Key-Systems	1345	613,337	8,326	135.75
6	TLD Registrar Solutions	1564	84,936	1,076	126.68
7	Eranet International Limited	1868	260,982	2,752	105.45
8	BigRock Solutions	1495	218,886	2,262	103.34
9	NameSilo	1479	4,403,551	42,489	96.49
10	WebNic.cc	460	848,034	7,872	92.83

Six gTLD registrars appeared in both our 2021 and 2022 yearly studies: TLD Registrar Solutions, NameSilo, Alibaba.com Singapore, BigRock Solutions, and REG.RU. Of these six, only NameSilo showed a year-over-year decrease in its yearly phishing score.

A general description of the 2022 top ranked gTLD registrars and the TLDs where we observed the largest counts of phishing domains registered at these registrars follows:

gTLD Registrar	Profile	Largest counts of Phishing Domains in...
DNSPod ⁸	A Tencent Cloud brand that offers domain name registration, web hosting, and free DNS service for all domain names. Tencent Cloud is reportedly the largest Internet company in China. ⁹	.SHOP, .PRESS, .TOP, .WORK, .COM
OwnRegistrar ¹⁰	Provides domain name services through a large channel partner network, headquartered in New York City, NY, USA. ¹¹	.COM
REG.RU ¹²	A Moscow, Russia-based Internet Service Provider, Website Hosting & Internet-related Services operator. ¹³	.COM, .XYZ, .SITE, .INFO, .ONLINE
Alibaba.com Singapore ¹⁴	Provides domain registrations, DNS service, e-commerce, secure cloud computing and data processing capabilities.	.COM, SHOP, .CYOU, .ICU, .TOP
Key-Systems ¹⁵	A Germany-based registrar that offers domain registration, colocation, virtual cloud and other services.	.BAR

gTLD Registrar	Profile	Largest counts of Phishing Domains in...
TLD Registrar Solutions ¹⁶	A London, UK based registrar that self-identifies as “conceived as a service provider to operators of new TLDs”.	.COM, .APP, .FINANCE
Eranet International Limited ¹⁷	The English-edition website of Todaynic.com, a Hong Kong, CN based registrar and web hosting provider.	.COM, .TOP, .NET.
BigRock Solutions ¹⁸	A provider of web-presence solutions to small businesses, professionals, and individuals.	.COM
NameSilo ¹⁹	Provides hosting, web sites, SSL, premium DNS and email, and claims to provide “the lowest everyday domain prices on the Internet”.	.COM, .XYZ, .TOP, .SHOP, .BUZZ
WebNic.cc ²⁰	Self-identifies as a domain wholesale registrar supplying domain names in bulk with a lower price.	.COM, .NET, .ORG, .XYZ, .TOP.

The following table shows the top gTLD registrars ranked by Yearly Phishing Domain Score (on the right) and the same ranking from the 2021 report.

2021 Rank	Registrar	Yearly Phishing Domain Score ▼	2022 Rank	Registrar	gTLD Domain Registrations	Yearly Phishing Domain Score ▼
1	TLD Registrar Solutions	117.4	1	DNSPod	1,387,872	221.76
2	NameSilo	105.9	2	OwnRegistrar	313,472	211.57
3	ALIBABA Singapore	81.3	3	REG.RU	726,674	199.15
4	Jiangsu Bangning	74.2	4	ALIBABA Singapore	1,677,681	164.14
5	BigRock Solutions	73.7	5	Key-Systems	613,337	135.75
6	NameCheap	71.6	6	TLD Registrar Solutions	84,936	126.68
7	Key-Systems	70.4	7	Eranet International	260,982	105.45
8	NETIM SARL	59.6	8	BigRock Solutions	218,886	103.34
9	REG.RU	59.0	9	NameSilo	4,403,551	96.49
10	Internet Domain Svc BS	54.9	10	WebNic.cc	848,034	92.83

High gTLD registrar phishing scores may indicate that phishers find that the registrar’s processes, pricing, or services are attractive or favorable for registering domain names for phishing. To explore this proposition further, we next consider domain names that have been purposely registered for phishing attacks.

Malicious Domain Name Registrations

We define a maliciously registered domain as a *domain registered to carry out a malicious or criminal act*. For our studies, we distinguish maliciously registered domains from compromised domains, which we define as *domain names that were registered for legitimate purposes but co-opted by criminals* through some form of compromise.

For example, an attacker may hijack a legitimate user's domain registrar account, alter the corresponding DNS entry to resolve a name or URL to a host that the attacker controls; here, the domain and DNS are compromised. An attacker may also exploit a vulnerability at a legitimate domain's web hosting, upload fake or malicious content to the web site; in this case, the web server is compromised.

This distinction is important because it often identifies where investigators should go for assistance with mitigation of the criminal activity:

- If the domain is maliciously registered, an investigator will seek assistance from a domain name registrar, a TLD operator, or the operator that provides DNS for the malicious domain to suspend the domain name registration or name resolution. The investigator may also contact the web hosting provider.
- Suspending a compromised domain would harm the domain's legitimate registrant by bringing down the legitimate site's web site and email. Investigators will contact the hosting provider to have the malicious content removed.

Note that parties that discover phishing pages will do their best to blacklist URLs that identify malicious content to avoid further victimization, whereas they may block maliciously registered domain names (and thus all hostnames and URLs created using this name) to contain the pervasive malicious activity.

To determine maliciously registered domains, we consider:

- **The age of the domain name — the number of days from domain registration to the use of the domain for a malicious purpose.** In general, the older the domain name, the higher the likelihood it will be legitimate. Miscreants tend to use their domains within the first year of registration, before paying for renewal of the domain name. The shorter the time between registration and use for phishing, the more likely the domain was maliciously registered.
- **The content of the domain name.** We apply rules to determine whether the composition of the name contains indicators of misuse or harmful intent, for example, the presence of a famous brand, a misspelled brand or a string intended to resemble a brand.
- **Bulk registered domains.** If many domains are registered within minutes via the same registrar and each reported for phishing, these domains are considered maliciously registered.
- **Extremely long domain names.** If a label is very long, and reported for phishing, it is likely that the domain is maliciously registered.
- **Sequences of similar domain names.** If a sequence of domain names (each of a minimum length) each have a Levenshtein distance²¹ of less than two, and each was reported for phishing, these domains are considered maliciously registered.
- **Sequences of same length domain names.** If a sequence of (a minimum number of) domain names each have the exact same length, and each was reported for phishing, these domains are considered maliciously registered.

When the above criteria identify domains, we then look for clear evidence of common control and usage as an indicator to flag additional domains in a batch.

Prevalence of Maliciously Registered Phishing Domains in TLDs

For the study period, we determined that 69% of the domains reported for phishing across all TLDs were registered maliciously by phishers (588,321 of the 853,987 domains reported for phishing). This percentage is slightly higher than our findings from our 2021 Phishing Landscape study where we found that 65% were maliciously registered. The remaining 31% were domains that we classified as compromised domains, or domains associated with subdomain services (see the section **Error! Reference source not found.**).

69% of phishing domains in all TLDs were maliciously registered

58% of all phishing attacks were hosted on maliciously registered domains

We observed a strong correlation between phishing attacks and maliciously registered domain names. 58% of all phishing attacks were hosted on maliciously registered domains. TLD registries and g registrars could be more proactive to mitigate maliciously registered domains promptly and take preemptive measures against suspiciously composed domains names (*e.g.*, domains that impersonate brands, or bulk registrations of randomly composed strings).

Prevalence of Malicious Registrations by TLD Category

The left-hand chart in Figure 6 shows that maliciously registered phishing domains accounted for a higher percentage of reported phishing domains than compromised domains in the legacy TLDs (a 69:31 ratio). However, from the right-hand chart, we observe that phishers hosted more attacks on compromised domains (a 53:47 ratio). This supports a theory that phishers find compromised hostnames attractive because they are harder to take down.

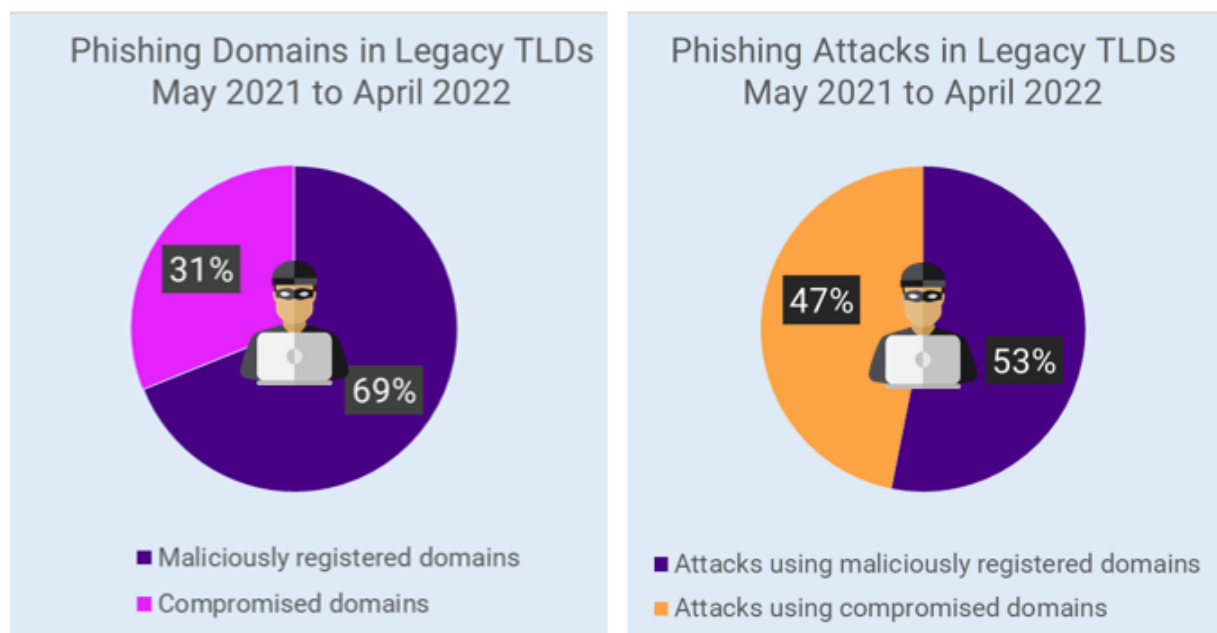


Figure 6 Phishing Domains and Phishing Attacks in Legacy TLDs, 1 May 2021 to 30 April 2022

As most ccTLDs do not provide domain registration data, we are unable to discriminate malicious phishing domains from compromised phishing domains from ccTLDs. We only have registrar information for 461,767 of the 853,987 total phishing domains.

However, in the new TLDs (Figure 7) we see a very different distribution; here, the ratios are extremely biased towards malicious registrations (92:8 compared to 85:15).

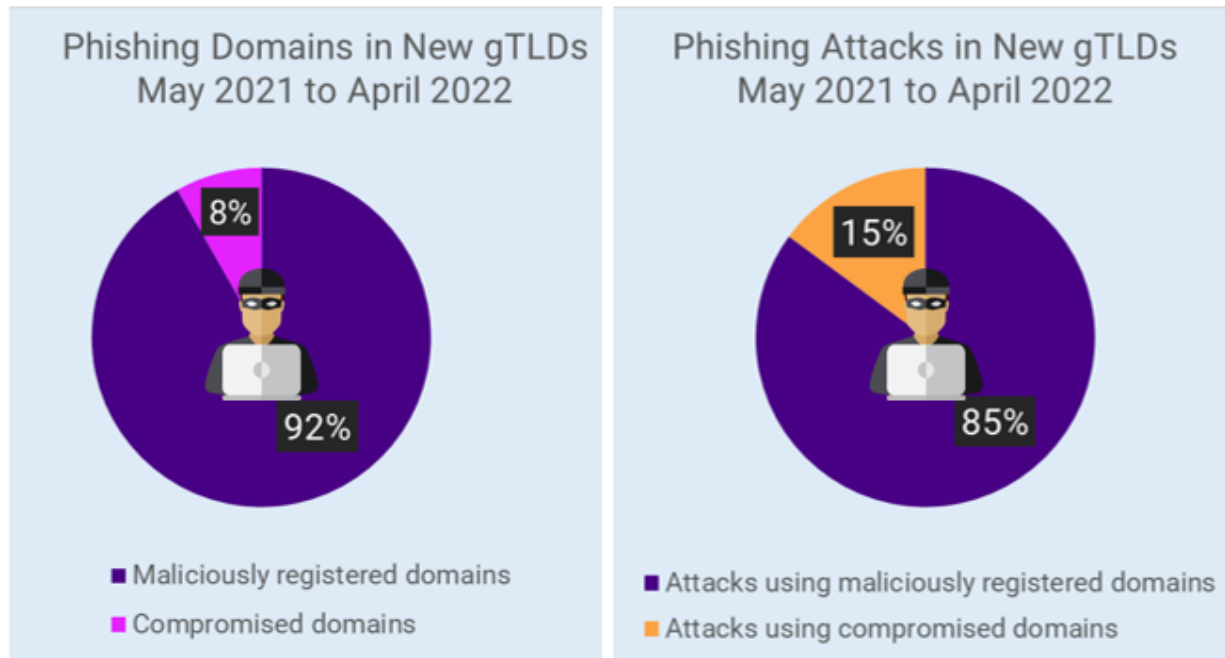


Figure 7 Phishing Domains and Phishing Attacks in New gTLDs, 1 May 2021 to 30 April 2022

Many of the new TLDs where we observe phishing have extraordinarily high percentages of malicious registrations. In some cases, and for extended periods of time, the percentages are so extreme that organizations configure security systems to blocklist the entire TLD.^{22 23}

Rank	TLD	% Malicious Phishing Domains ▼	TLDs with high percentages of malicious phishing domain registrations also have very high yearly phishing scores	Rank	TLD	Registry	Domains in TLD	Yearly Phishing Domain Score ▼
1	support	91%		1	support	Binky Moon	31,288	595
2	bar	99%		2	bar	Punto 2012 S.A.P.I.	281,575	562
3	shop	97%		3	shop	GMO Registry	1,040,404	459
4	work	98%		4	work	Minds + Machines	303,321	340
5	live	88%		5	live	Identity Digital	596,948	208
6	buzz	98%		6	buzz	DOTSTRATEGY	513,201	196
7	casa	88%		7	casa	Minds + Machines	30,290	178
8	finance	74%		8	finance	Cotton Cypress, LLC	52,709	173
9	sbs	94%		9	sbs	Shortdot SA	34,030	164
10	fyi	95%	10	fyi	Identity Digital	40,290	150	

We next look at the effects that malicious domain name registrations have on the levels of phishing activity in TLDs.

Malicious Domain Name Registrations and TLDs

We identified 13 TLDs with 5,000 or more malicious phishing domain registrations from 1 May 2021 to 30 April 2022.

gTLD Rankings: Malicious Phishing Domain Registrations	2022 Rank	2021 Rank	TLD	Registry Operator	Domains in TLD	2022 Malicious Phishing Domain Registrations ▼
	1	1	com	Verisign	159,902,632	191,661
	2	16	shop	GMO Registry	1,040,404	46,379
	3	5	xyz	XYZ.COM	4,130,573	35,665
	4	26	bar	Punto 212 S.A.P.I.	281,575	15,670
	5	9	top	Jiangsu Bangning	1,710,824	12,498
	6	8	info	Digital Identity	3,653,720	11,382
	7	12	live	Digital Identity	596,948	10,981
	8	30	work	Minds + Machines	303,321	10,150
	9	13	buzz	DOTSTRATEGY	513,201	9,800
	10	11	net	Verisign	13,170,783	9,091
	11	13	online	Radix FZC	1,847,551	7,873
	12	15	org	Public Interest Registry	10,614,272	6,104
13	18	site	Radix FZC	1,007,427	5,671	

Counts of phishing domains help us to identify where domain names reported for phishing were registered. We use a complementary analysis – one in which we can consider malicious or criminal

intent on the part of an unknown subject (registrant) – to identify prevention or mitigation opportunities for individual TLDs. Specifically, we discriminate maliciously registered phishing domains from compromised domains (web sites) to identify the parties that are best positioned to combat phishing.

Experience in the field has demonstrated that:

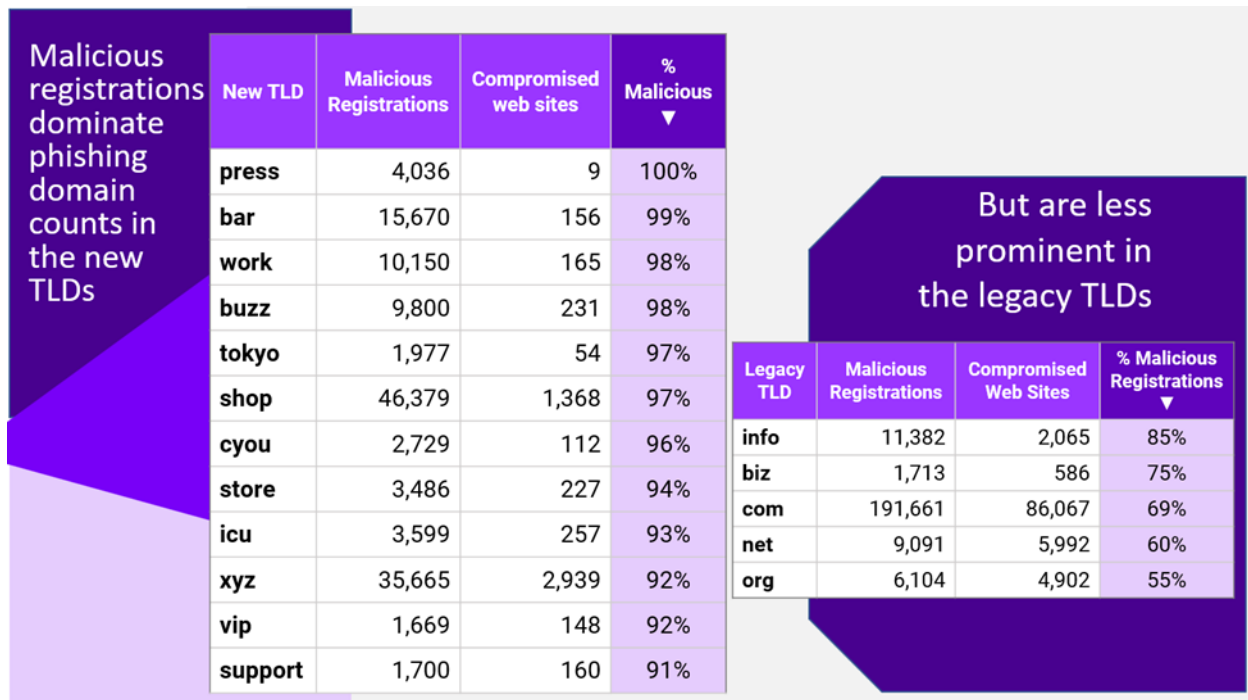
- Maliciously registered phishing domains can be suspended by the registrar or registry operator; this stops the attacks and will not cause any damage or inconvenience to anyone except the phisher.
- Registries with high numbers of maliciously registered domain names can collaborate with their registrars to adopt phishing identification and prevention measures.
- Hosting network operators are best suited to mitigate vulnerabilities for compromised web sites hosting phishing pages. They are also able to deploy measures to detect compromises and to recommend content management practices that can reduce their customers' web vulnerability attack surfaces.
- Phishers are highly unlikely to remove the phishing pages from a hosting server. The responsibility to remove fraudulent phishing content, disable an unauthorized web server, or suspend accounts of subscribers who are perpetrating phishing falls upon hosting operators. Typically, these are violations of the operator's own acceptable use policy.

gTLDs Where Malicious Domain Registrations Dominate in Phishing Reports

In some gTLDs, malicious phishing domain registrations account for the majority of reported phishing domains for the yearly period. This is particularly the case for new gTLDs with a minimum of 1,500 reported phishing domains. **We observed 64 new gTLDs where over 90% of the reported phishing domains were maliciously registered, and 133 new gTLDs where over 80% of the reported phishing domains were maliciously registered.**

Legacy gTLDs with more than 1,500 reported phishing domains had a *lower* percentage of malicious registrations. In these gTLDs, we observed more compromised web sites; in particular, the percent of domain names purposely registered for phishing in .COM was the same as the 69% we estimated for the overall domain name space, and both .NET and .ORG were much lower.

New TLDs have extraordinarily high percentages of malicious registrations intended for phishing



Malicious Domain Name Registrations and gTLD Registrars

Counts of phishing domains help us to identify where domain names reported for phishing were registered. Further analysis is needed to understand what acts of prevention or mitigation are appropriate for gTLD registrars. By identifying characteristics of maliciously registered domain names and distinguishing these from compromised domains, we can identify which parties are best positioned to act to prevent phishing.

The classification *compromised domains* represents the set of domains where the domain name owner who operates a legitimate web site may be a victim. Here, investigators should seek out hosting providers to mitigate phishing attacks (*e.g.*, by having the phishing page and related content removed from the compromised web site).

The classification *maliciously registered phishing domains* represents the set of domains that were purposely registered for phishing, by an actor with criminal intent (to commit fraud). Here, a gTLD registrar is often well positioned to (proactively) identify a domain as “intended for phishing”; for example, *only* a gTLD registrar has the means to:

- examine a domain name such as `amazongjgasb14sjh21saknx.icu`, `appleidsupport-us.com`, or `customersupport-netflix.com` during registration,
- detect a trademark or brand within the domain name (*e.g.*, Amazon, Apple, Netflix), and
- suspend the registration while it reviews the registrant’s contact data to assess the legitimacy of the registration.

The maliciously registered classification also represents the types of domains where investigators should seek the assistance of gTLD registrars to mitigate phishing attacks (*e.g.*, by suspending the domain name or registrant account). For example, when a phishing investigator determines that a phishing campaign is using dozens or more domain names containing random patterns, *only* a gTLD registrar can determine during the early hours of a phishing attack whether the contact data for a set of verified phishing

domains is the same (an historically reliable indicator of a phisher). The gTLD registrar should review the evidence of phishing presented by a phishing investigator quickly and accommodate requests to reveal the contact data of a registrant once they verify the evidence.

The following table shows gTLD registrars with more than 8,000 malicious phishing domain registrations under management from 1 May 2020 to 30 April 2021.

Rank	Registrar	Registrar IANA ID	Malicious Phishing Domain Registrations May 2021 to April 2022 ▼
1	NameCheap	1068	72,905
2	NameSilo	1479	36,513
3	DNSPod	1697	30,446
4	ALIBABA.COM Singapore	3775	26,216
5	GoDaddy.com	146	15,754
6	PublicDomainRegistry	303	12,563
7	REG.RU	1606	11,652
8	Wix.com	3817	11,119
9	Wild West Domains	440	8,572
10	HiChina (www.net.cn)	1599	8,245

We next compared malicious phishing domain registrations to compromised domains, by gTLD registrars. The raw numbers of maliciously registered domains are important — they indicate where phishers registered the most domains.

Malicious registrations directly influence the reputations of the gTLD registrars that are most targeted by phishers when they register domains purposely for phishing. In some cases, a gTLD registrar's malicious domain registration can also have a disastrous effect on the phishing score of a Top-level Domain and consequently on that TLD's reputation. For some TLDs, one gTLD registrar adversely influences a TLD's reported phishing domain counts month after month.

Which gTLD registrars had an adverse effect on which gTLDs, and to what degree? In the following table, we identify gTLD registrars where the registrar's share of all maliciously registered phishing domains in the TLD was at least 40% and thus had a significant influence on a particular gTLD's phishing domain count.

Registrar	TLD	Registrar's share of maliciously registered phishing domains in TLD (%) ▼
NameSilo	buzz	95
DNSPod	work	88
DNSPod	press	87
Key-Systems	bar	81
ALIBABA.COM Singapore	cyou	69

Registrar	TLD	Registrar's share of maliciously registered phishing domains in TLD (%) ▼
NameSilo	top	56
NameCheap	xyz	52
NameSilo	club	48
ALIBABA.COM Singapore	icu	47
DNSPod	shop	42

The new TLDs receive the most negative attention in policy communities and domain industry reports where the topics of DNS abuse or the criminal misuse of domain names and the DNS are discussed.

With measurements of maliciously registered domain names, we can turn attention to registration services. *Something* attracts criminals to certain registrars. Whether pricing, promotions, account security, or other vulnerability, a high percentage of malicious registrations indicates that the reputation of individual TLDs as well as to the new TLD program overall is adversely affected. Organizations and even individual Internet users have and continue to implement measures at firewalls or DNS resolvers to blacklist entire TLDs.²⁴ It is more difficult to blacklist “all the domains under the management of a gTLD registrar” because there is currently no means to timely and reliably identify the registrar of a domain name. But such a measure may be practical and necessary.

Opportunities to Prevent or Mitigate Malicious Registration Activity

Most phishing responses are reactive. Domains that will be used to lure users to phishing sites have already been registered, or they are obtained through some form of compromise attack. The phishing content has been composed and hosted. Spam emails or other means of presenting lures to Internet users have been transmitted. Victims have been harmed.

All gTLD registrars are contractually required by ICANN to have mitigation programs.²⁵ They must:

- maintain an abuse contact to receive reports of abuse and illegal activity, and publish the abuse contact address,
- publish on their website a description of their procedures for the receipt, handling, and tracking of abuse reports,
- document their receipt of and response to all such reports, and
- “take reasonable and prompt steps to investigate and respond appropriately to any reports of abuse.”²⁵

No definition exists for what constitutes a reasonable and prompt response. gTLD registrars have no contractual or globally applicable legal obligation to respond appropriately, for example, to lock or suspend maliciously registered domains pending investigation.

Mitigation programs can reduce harms or losses, but our findings regarding malicious registrations show that phishers can and do register and use large numbers of domains at specific registries and registrars, again and again over time.

These levels of phishing activity might be caused by one or more of the following factors:

- 1) Low pricing, offered as part of a registrar and/or a registry operator's sales strategy. In general, phishers tend to be attracted to low prices.²⁶
- 2) Inattention to abuse problems by the registrar and/or the registry operator. This allows phishers to buy and use domains over time.
- 3) Features at the registrar that facilitate phishing, such as APIs that allow registrations in bulk, or payment methods that offer anonymity or have weak fraud detection. Cybercriminals take advantage of bulk registration services to "weaponize" large numbers of domain names.²⁷

Our purpose for defining a method to distinguish phishing domains as maliciously registered or as hosted on compromised assets is to make clear how phishers acquired resources. If phishers use malicious registrations more frequently than compromised assets, then prevention programs that deal with these registrations more proactively would be most helpful.

There may be opportunities for registry operators and registrars to use the methods that phishing investigators apply when phishing is first seen to suspend domains for malicious or illegal activity before they can victimize users and brands.

gTLD registrars and TLD operators are in an excellent position to identify and suspend malicious domain name registrations with a high degree of accuracy, often at the time of registration, and often by using the same methods that phishing investigators apply when phishing is first seen in the wild. For example, many domains registered by phishers also have telltale characteristics – name composition, common creation dates, similarities in contact data – that an operator can use to identify malicious registrations quickly and with low false-positive rates.

gTLD registrars and TLD operators possess key information – contact data and billing data – that no one else does. This data is highly useful to identify malicious customers at the time of registration.

Access to contact information – the registrant's identity, payment information, IP address, and purchase history – can be essential in a phishing investigation. Traditionally, phishing investigators would use WHOIS contact data to find other domains with similar contact data elements, and thus owned by the same cyber criminals. Only by identifying virtually *all* of a phisher's domain names can investigators hope to fully mitigate a phishing campaign.

gTLD registrars and TLD operators have terms of service that allow them to suspend domains for illegal activity. For example, they can monitor for such activity, and suspend domains for malicious purposes. Many operators have AUPs. Phishing is a recognized manifestation of fraud every jurisdiction in which registrars and TLDs operate. Stringently and uniformly enforcing a prohibition against phishing should result in a reduction in maliciously registered domains.

Phishing Distribution Across Hosting Networks (Autonomous Systems)

We studied where phishing sites were being hosted, to determine if any hosting providers have outsized phishing problems. We collected the IP addresses (A records) that phishing attacks were resolving to. We then looked up the **Autonomous System Number (ASN)**²⁸ containing each IP address. This provides insight into the entities that hosted the phishing attacks.

We are not seeing phishing on IPv6 addresses; the following sections are about IPv4 addresses only.

Ranking of Hosting Networks (ASNs) by Phishing Attacks Reported

We compared rankings of the hosting networks (ASNs) having the most phishing attacks reported in our 2021 study against the rankings for our 2022 study.

2021 Rank	AS Name	AS number	Phishing Attacks ▼	2022 Rank	AS Name	AS number	# Routed IPv4 Addresses	Phishing Attacks ▼
1	NAMECHEAP-NET	22612	55,903	1	CLOUDFLARENET	13335	2,400,256	120,209
2	CLOUDFLARENET	13335	52,011	2	UNIFIEDLAYER	46606	1,207,808	63,510
3	UNIFIEDLAYER	46606	35,363	3	MICROSOFT	8075	45,502,976	46,995
4	GOOGLE	15169	32,330	4	NAMECHEAP-NET	22612	102,912	40,969
5	DIGITALOCEAN	14061	15,794	5	GOOGLE	15169	23,099,904	30,397
6	AWEX-Hostinger	204915	13,186	6	AMAZON-02	16509	42,667,520	25,591
7	OVH - OVH SAS	16276	12,604	7	ALIBABA (US)	45102	4,955,136	24,242
8	WEEBLY	27647	10,701	8	QUADRANET-GLOBAL	8100	574,208	23,345
9	CONTABO	51167	10,635	9	DIGITALOCEAN	14061	2,701,056	21,791
10	AMAZON-02	16509	10,257	10	FASTLY	54113	457,728	20,541

We found phishing in 4,099 hosting networks. Ten of the top hosting providers accounted for 47% of the 640,487 phishing attacks for which an ASN could be determined, with four of the hosting networks accounted for 30% of those phishing attacks.

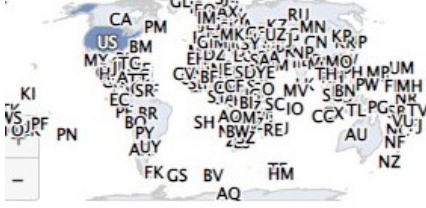






About the top-ranked hosting network operators identified by ASN:

Hosting Network	Description
CLOUDFLARENET	San Francisco CA based Cloudflare provides a DNS redirection service that protects its customers from denial-of-service attacks. Cloudflare's service also prohibits observers from seeing the real hosting locations behind this defense network, and phishers take advantage of this to hide the hosting locations of phishing pages.
UNIFIEDLAYER	A San Francisco CA based provider of managed cloud services, secure enterprise-class cloud, co-location, and disaster recovery services. ²⁹ This ASN rose from 3 rd in 2021 to 2 nd in 2022.
MICROSOFT	One of the largest network providers worldwide. All public and private peering for the Seattle WA corporation's ASNs 8068 and 8069 is handled via this ASN 8075. ³⁰ This ASN was not in the top 10 in 2021 but ranked 3 rd in 2022.

Hosting Network	Description
NAMECHEAP-NET	A Phoenix AZ based domain name registrar that also offers hosting for its customers. Namecheap has the smallest number of routed IPv4 addresses of the Top Ten ASNs, but the 4 th largest number of phishing attacks. In 2021, Namecheap was ranked 1 st .
GOOGLE	Google is also one of the largest network providers. ASN 15169 is delegated to Google LLC, Mountain View, CA and is aliased as Google, YouTube. ³¹
ALIBABA-US-TECHNOLOGY	ASN 45102 is delegated to the Alibaba Group, purportedly the largest cloud option in China. ³² This ASN was not in the top 10 in 2021 but rose to 6 th in 2022. Most of the IP addresses reported for phishing in this ASN are geolocated in China, but some appear in IP prefixes allocated to US hosting (cloud) services.
QUADRANET-GLOBAL	A Los Angeles CA based data center provider, providing colocation, dedicated servers, cluster management, and complex hosting solutions. ³³
DIGITALOCEAN	Known as a developer cloud, the New York City based cloud operator dropped from 5 th in 2021 to 8 th in 2022.
FASTLY	A San Francisco CA based edge cloud platform provider. Fastly achieved notoriety in June 2021 when problems with its content delivery network caused outages to popular sites including websites including Reddit, Twitch, Hulu and The New York Times. ³⁴ Fastly was unranked in 2021 but rose to 9 th in 2022.
AMAZON-02	This Amazon ASN is aliased as Amazon Web Services, one of the largest cloud operations in the world. The Seattle CA company's ASN held rank at 10 th .

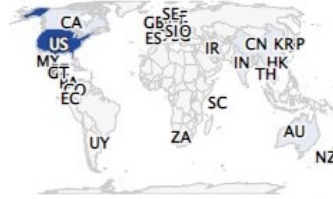
Our data showed that US hosting networks were attractive to phishers. We used RIPEstat geo data³⁵ (per Maxmind GeoLite) to determine the countries where IP addresses reported for hosting phishing attacks for each of ASNs had the most reported phishing attacks.

86% of phishing attacks reported using the top 10 hosting networks (ASNs) are hosted on IP addresses in the US

<p>CLOUDFLARENET AS13335</p> <p>Phishing attacks from IP addresses in US: 118,561</p>	
<p>UNIFIEDLAYER-AS-1 AS46606</p> <p>Phishing attacks from IP addresses in US: 63,300</p>	
<p>MICROSOFT-CORP-MSN AS8075</p> <p>Phishing attacks from IP addresses in US: 29,232</p>	
<p>NAMECHEAP-NET AS22612</p> <p>Phishing attacks from IP addresses in US: 40,138</p>	
<p>GOOGLE AS15169</p> <p>Phishing attacks from IP addresses in US: 25,735</p>	
<p>AMAZON-02 AS16509</p> <p>Phishing attacks from IP addresses in US: 20,589</p>	
<p>Alibaba (US) Technology Co. AS45102</p> <p>Phishing attacks from IP addresses in US: 1,507</p>	

**ASN-QUADRANET-GLOBAL
AS8100**

Phishing attacks from
IP addresses in US: 22,356



**DIGITALOCEAN
AS14061**

Phishing attacks from
IP addresses in US: 17,237



**FASTLY
AS54113**

Phishing attacks from
IP addresses in US: 20,454



Legislation or regulation may be useful or necessary to mitigate phishing attacks effectively

Regulations that mandate accurate contact information from Internet as a Service operators³⁶, or that oblige operators to “lock and suspend”³⁷ a hosting or registration service while an investigation of a malware threat is conducted may provide protections against phishing attacks that currently do not exist across an ecosystem that has no single policy or administrative authority.

Ranking of Hosting Networks (ASNs) by Scoring Metrics

The gross numbers of phishing attacks reported are significant. Here, as with TLDs and gTLD registrars, more phishing attacks means more damage and victimization. A heavily abused ASN can enable many attacks. If it makes improvements to its anti-abuse efforts, it can reduce victimization and make things harder for phishers.

Gross numbers influence how one compares operators who have more or less IP addresses than each other (numbers bias). In the quarterly phishing activity published at the Cybercrime Information Center, the metric “Phishing Attacks per 10,000” is used to compare whether a hosting network (AS) has a higher or lower *incidence* of phishing relative to others. This is a ratio of the number of phishing attacks hosted in an Autonomous System to the IPv4 addresses routed by that hosting network (AS). We call this metric **hosting network (AS) Phishing Attack Score**:

$$\text{Hosting Networks (AS) Phishing Attack Score} = (\text{number of phishing attacks}/\text{IPv4 addresses routed by AS}) * 10,000$$

For this report, and as we did for TLDs, we measured this prevalence of phishing in each hosting networks (ASNs) for the period 1 May 2020 to 30 April 2021. Here, we take the yearly count of unique

phishing attacks reported and divide by the number of IP addresses routed by each hosting network (ASN). We call this metric **Yearly Hosting Network (ASNs) Phishing Attack Score**:

$$\text{Yearly hosting networks (AS) Phishing Attack Score} = \frac{\text{number of unique phishing attacks reported}}{\text{number of IP addresses routed by AS}} * 10,000$$

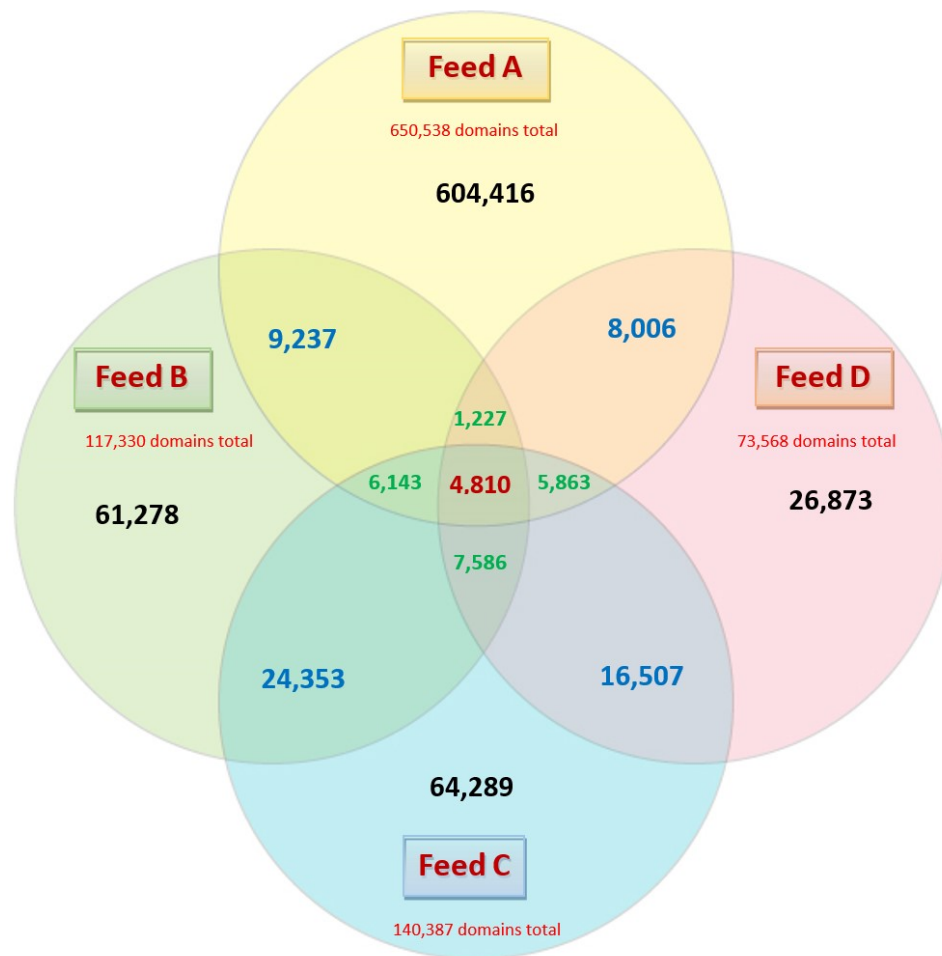
Note that the calculation of these two metrics yields different results (we use different inputs for the numerators in the division); in particular, one cannot draw any conclusion by comparing the scores from a quarterly phishing score against an annual phishing score. Instead, we encourage comparisons of quarterly phishing attack scores over time, as well as annual phishing attack scores over time.

Rank	AS Name	AS number	# Routed IPv4 Addresses	Phishing attacks	Phishing Attack Score ▼
1	NAMECHEAP-NET	22612	102,912	40,969	3980.97
2	CONTABO	40021	52,992	4,162	785.40
3	Domain names registrar REG.RU	197695	91,648	5,331	581.68
4	TimeWeb Ltd.	9123	59,136	3,147	532.16
5	UNIFIEDLAYER	46606	1,207,808	63,510	525.83
6	Hostinger International Limited	47583	124,672	6,278	503.56
7	CLOUDFLARENET	13335	2,400,256	120,209	500.82
8	INMOTI-1	54641	61,440	2,983	485.51
9	PONYNET	53667	63,232	3,061	484.09
10	FASTLY	54113	457,728	20,541	448.76

List Coverage: The Phish That Get Away

By collecting data from multiple sources, we confirmed that there is low overlap between anti-phishing blocklists. For our 1 May 2021 to 30 April 2022 study period, we identified a total of 854,120 unique domain names listed for phishing (again, either URLs on those domains, or the domain itself).

The Venn diagram³⁸ illustrates that most of the domains in the study period were reported via a single feed. Fewer than 1% of the domains – 4,810 out of 854,120 – were reported by all four feeds.



The existence of this coverage problem has been confirmed in a series of studies, which have found similar gaps for cybercrime data generally and for specific types of abuse including phishing.^{39, 40, 41, 42, 43, 44, 45}

What factors contribute to the low overlap? Some factors are common to detecting cybercrime generally, and some are especially relevant to phishing:

1. The Internet is a big place, and each blocklist provider only has a certain window of visibility into it. For example, a provider will have access to only a certain amount of email spam that it can scan for phishing lures. Different observers use different detection or collection methods.
2. The limited duration of phishing attacks provides only a small period in which observers can confirm the presence of a phishing site.

3. Phishers employ a variety of evasive techniques that complicate the confirmation of phishing attacks.^{46, 47, 48} One called “cloaking” notably decreases the likelihood that a phishing site will be blacklisted, and if a URL does get blacklisted, the cloaking substantially delays blocking in browsers.^{49, 50}
 4. The sharing of data is uneven and is not always timely. Some phishing targets do not share data about the phishing that affects them, for fear that it will reflect negatively on their brands. Some anti-phishing vendors do not share their data due to competitive concerns.⁵¹
 5. ICANN policy now allows gTLD domain registrars to redact all domain contact data from publication in WHOIS, even those records not covered by a privacy law such as GDPR. That contact data is a key tool for identifying malicious registrations and differentiating them from compromised domains.
- Over-redaction of WHOIS data continues to contribute to the under-identification of phishing domains and impedes criminal investigations.**^{52, 53, 54}

*Over-redaction of WHOIS data
impairs phishing domain
identification*

How much phishing is not being detected at all? What is the number of “unknown unknown” attacks, and what is the total size (upper boundary) of the phishing problem? No one knows for sure. The factors

*Phishing is a much larger
problem space than is reported*

we list above inhibit even the best detection systems from finding much of the phishing attacks that occur and even the most professional and experienced observers can find only a portion of the phishing that occurs and are challenged to do so in a timely fashion.

A recent case illustrates how phishing is often now caught by monitoring, and how numbers above greatly under-count phishing. In December 2021, Meta Platforms (the parent company of Facebook, WhatsApp, and Instagram, etc.) filed a lawsuit against phishers who were registering subdomains on the NGROK.IO service.^{55, 56} NGROK operates a proxy that relays traffic to websites in a manner that obscured the location of the sites, as well as the identities of the hosting providers and the users who created the subdomains. According to the lawsuit, phishers had registered and used at least 39,000 subdomains on ngrok.io from September 2019 to December 2021. However, only 810 phishing attacks hosted at NGROK.IO were reported by the sources we used to compile this study.

Blocklists are essential tools for cybersecurity: they prevent enormous damage, and all organizations should take advantage of them directly or through their service providers. Organizations should further consider whether they are well served with one blocklist, or whether they would benefit from incorporating multiple sources of threat intelligence in their phishing defenses.

Targeted Brands

Phishers targeted over 2,000 businesses or organizations during the 1 May 2021 to 30 April 2022 period, including banks, social media companies, webmail, and games, national tax services, universities, and cryptocurrency exchanges. The majority of phishing attacks targeted just ten brands.

To identify brands, we used three URL blocklists that identify targets in the metadata included in their phishing reports. Since reports from each phishing feed that we consume vary slightly in granularity and nomenclature, we must compile lists of these variations and normalize spelling as part of our curation; for example, if one feed uses “PayPal” while another uses “PayPal Inc.,” we treat these as one target and normalize our data to a common string (brand name) so that we could analyze brand data.

Using multiple feeds poses classification challenges. For example, WhatsApp is owned by Meta Platforms. One feed may report a phishing URL as an attack against WhatsApp as a separate brand, but a second feed may report the same phishing URL as attacks against Facebook. In these cases, we use the target reported by each feed, with the granularity (discrimination) that feed offers.

In some cases, one source may positively identify a URL as a phish against a specific target, but another source may only report that the URL as a phishing attack against “unknown” or “generic” brand. In these cases, we use the most detailed information available and attributed that attack to the specific brand. In the cases where an attack’s target is not determined by any feed, we set those attacks aside: we do not include them when analyzing brand data.



The following table identifies the most targeted brands, by annual ranking and phishing attacks, and then their quarterly ranking, from May 2021 through April 2022:

Annual Ranking	Brand	Attempted Phishing Attacks ▼	Quarterly Ranking			
			May-July 2021	August-October 2021	November 2021-February 2022	February-April 2022
1	Facebook	52,794	1	3	2	1
2	Amazon	36,656	2	1	4	2
3	Microsoft	28,668	4	2	3	3
4	WhatsApp	15,674	7	7	5	5
5	Apple	14,891	5	4	7	6
6	Crypto/Wallet	13,999	81	25	1	7
7	Instagram	12,606	9	14	12	4
8	Outlook	11,877	6	9	8	8
9	DHL	10,172	8	6	10	11
10	Chase	9,632	15	8	9	10
11	PayPal	9,389	11	11	6	17
12	Adobe	7,323	19	12	16	12
13	PenSam	6,863	3	1,167	1,066	1,234
14	Wells Fargo	5,855	45	13	13	20
15	Netflix	5,777	27	17	15	15
16	AT&T	5,381	31	18	14	21
17	Tencent	5,072	18	19	26	18
18	Citi	4,784	49	32	11	16
19	IRS	4,618	29	16	17	24
20	webmail	4,288	32	23	27	19

A brand can become a phishing target at any time. Phishers constantly look for companies that have potentially lucrative user information, are newly popular, or are not ready to respond to phishing. Phishers also use a variety of ploys to lure Internet users to their phishing pages including,

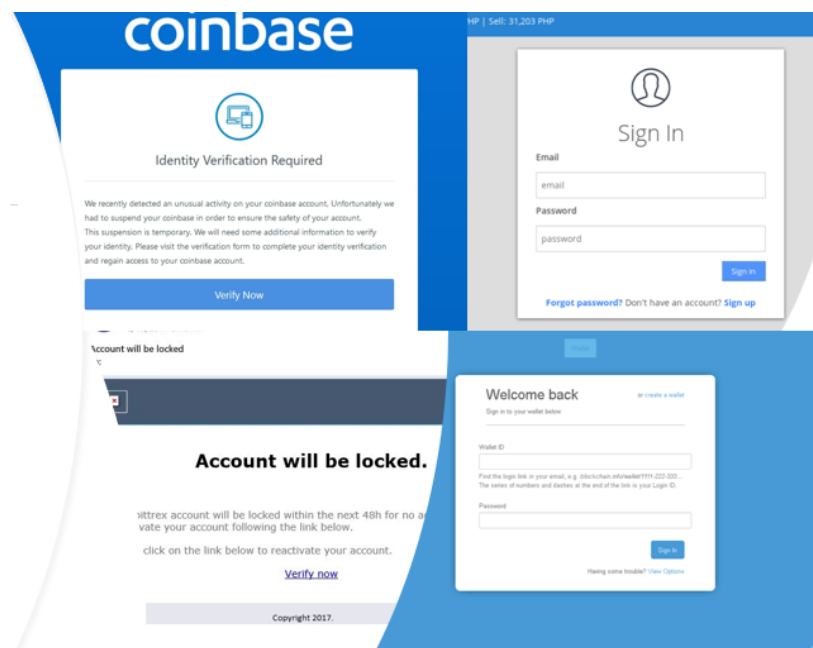
- a new product announcement,
- a critical software update to obtain,
- an issue with a social media account or financial account,
- a problem with a merchant transaction or subscription,
- an inquiry regarding a criminal matter or tax violation,
- a newsworthy or catastrophic event, such as the COVID-19 pandemic, or
- an emerging technology or service such as cryptocurrency.

Cryptocurrency Phishing

The cryptocurrency market value topped two trillion dollars in April 2021.⁵⁷ The bullish interest in the market, which was up over 180% from April 2020 to April 2021, also attracted phishers: crypto-related crimes reportedly exceeded US\$14 billion in 2021, compared to US\$7.8 billion in 2020. Approximately 30% of this value may be attributed to phishing attacks that disclosed private cryptographic keys.⁵⁸

Cryptocurrency phishing objectives are the same as bank phishing: steal money, credentials, and personal identifying information. Many cryptocurrency phishing schemes involve attacks on wallets – a mobile app, browser extension, or hardware device that stores cryptographic keys and allows users to buy, sell, and store cryptocurrency. Phishers use various lures in wallet phishing attacks, for example, one attack used a threat of having a wallet account closed to lure users to a bogus wallet portal, where victims disclosed information used to access MyEtherWallet wallets.⁵⁹ Another attack used a homoglyph: the phisher added an underdot to the first letter “e” in the brand “Ledger” and composed a phishing URL from the domain lędger.com to Ledger customers to a fake site.⁶⁰ Neither of these attack methods is new.

Cryptocurrency phishers use the same credible fakes to deceive cryptocurrency adopters as they have used to deceive banking customers



Other attacks target trading platforms including Binance, Coinbase, or Kraken. Phishers used text messages in a recent (June 2022) campaign that notified Binance customers of an attempted withdrawal from an unknown IP address. The text directed customers to a fake Binance log in to cancel the withdrawal, where attackers collected credentials from their victims.⁶¹

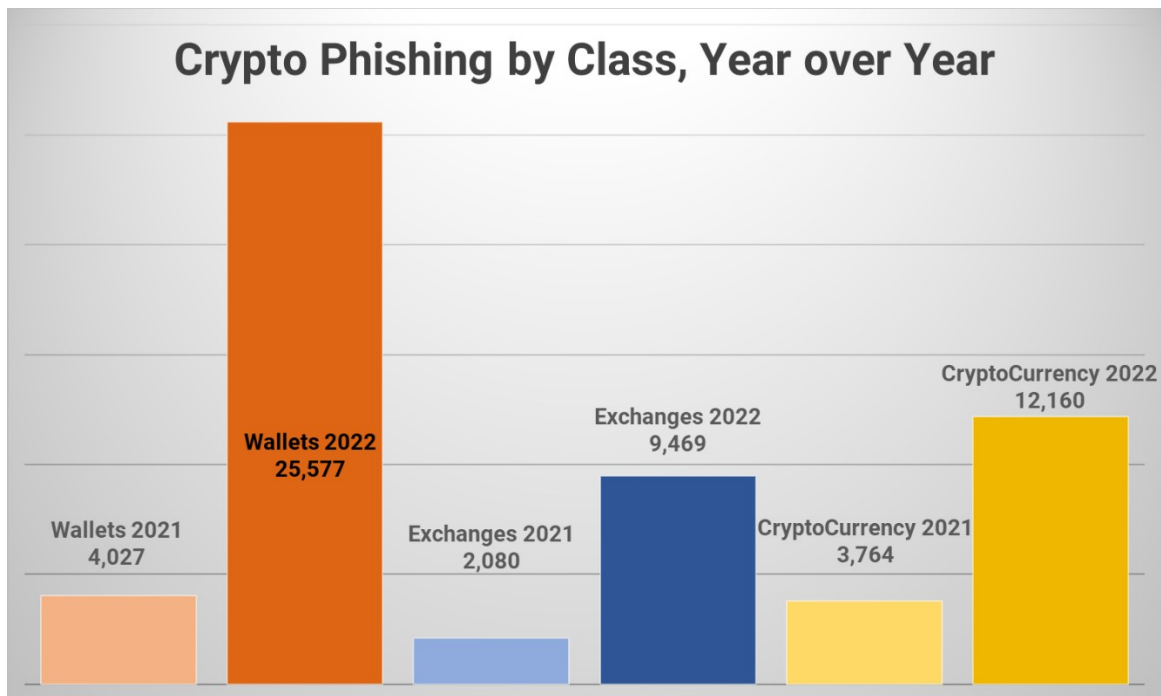
To study phishing activity in cryptocurrencies, we examined phishing URLs for keywords and brands that are common to cryptocurrency and metadata provided by phishing feeds to collect a data set suitable for closer analyses. We associated the URLs with three prominent *classes* of cryptocurrency.

The **Wallet Phishing** class included mobile apps, browser extensions, or hardware devices that store cryptocurrency keys and allows users to buy, sell, and store cryptocurrency. We associated URLs with wallet phishing if they contained keywords (airdrop, wallet) or popular wallet brands (e.g., Ledger, WalletConnect, Trust Wallet).

The **Exchange Phishing** class included platforms such as Binance, Coinbase, or Kraken, that allow cryptocurrency investors to buy, hold or sell cryptocurrencies. We associated URLs with exchange phishing if they contained names of Crypto Exchanges, *e.g.*, Coinbase, Kraken, PancakeSwap, Binance, and Paxful. We also included URLs that contained strings such as exchange, cryptofx, fxcrypto, fxtrade, tradefx, or forex. These strings suggested currency trading and were frequently used in conjunction with an exchange brand name or cryptocurrency.

The **Cryptocurrency Phishing** class included common cryptocurrency vernacular as well as brands and symbols of popularly traded cryptocurrencies. We associated URLs that contained names of popular cryptocurrencies (*e.g.*, Bitcoin, Localbitcoin) and their exchange symbols of the most popular cryptocurrencies (*e.g.*, BTC, ELON). We also included URLs that contained common keywords (*e.g.*, crypto, blockchain, coin) and URLs that were tagged by our phishing feeds as targeting “generic crypto(currency)”. We excluded false matches of keywords and abbreviations (*e.g.*, where btconnect targets BT Connect, not Bitcoin).

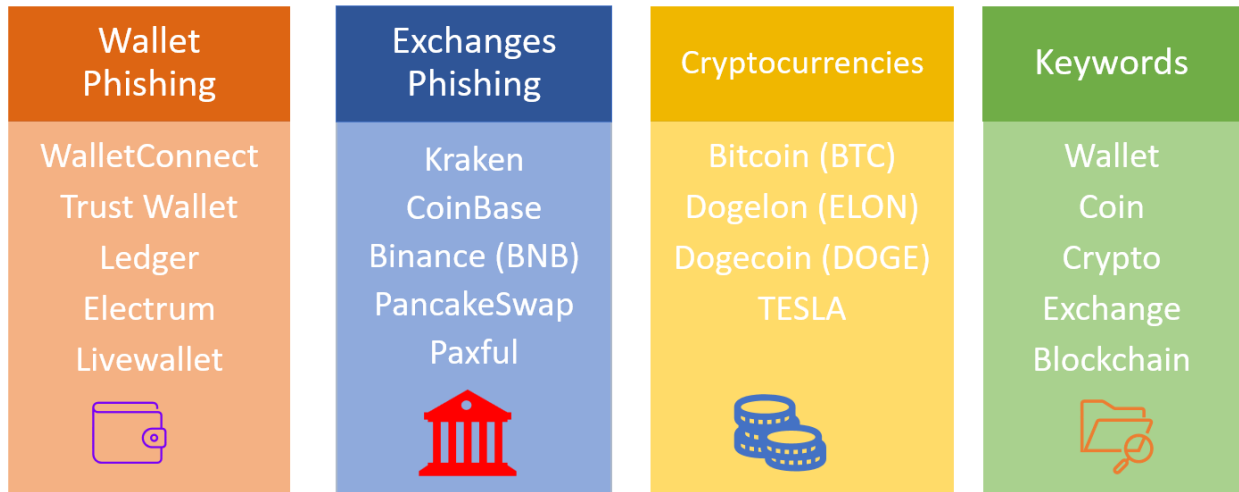
We observed that phishing reporters began tagging cryptocurrency phishing more consistently in our Phishing Landscape 2022 data. We modified our rule sets to include this welcomed metadata and to consider new brands and keywords that appeared in the cryptocurrency world (for example, the emergence in popularity of the PancakeSwap⁶² decentralized exchange, new currencies, etc.).



To fairly compare Phishing Landscape 2021 cryptocurrency phishing numbers against Phishing Landscape 2022, we ran our new rule set against the Phishing Landscape 2021 data to obtain a year-over-year comparison with this common rule set.

Overall, we observed a 257% increase in cryptocurrency phishing. We observed a 630% increase in wallet phishing, 460% increase in phishing that targeted exchanges, and a 323% increase in phishing that targeted cryptocurrencies.

Most Encountered Brands and Keywords



Composition of cryptocurrency phishing domains

We observed that phishers frequently used two or more brands or keywords when they composed domains that they purposely registered for their cryptocurrency attacks. The defanged URL examples that follow illustrate this practice:

```

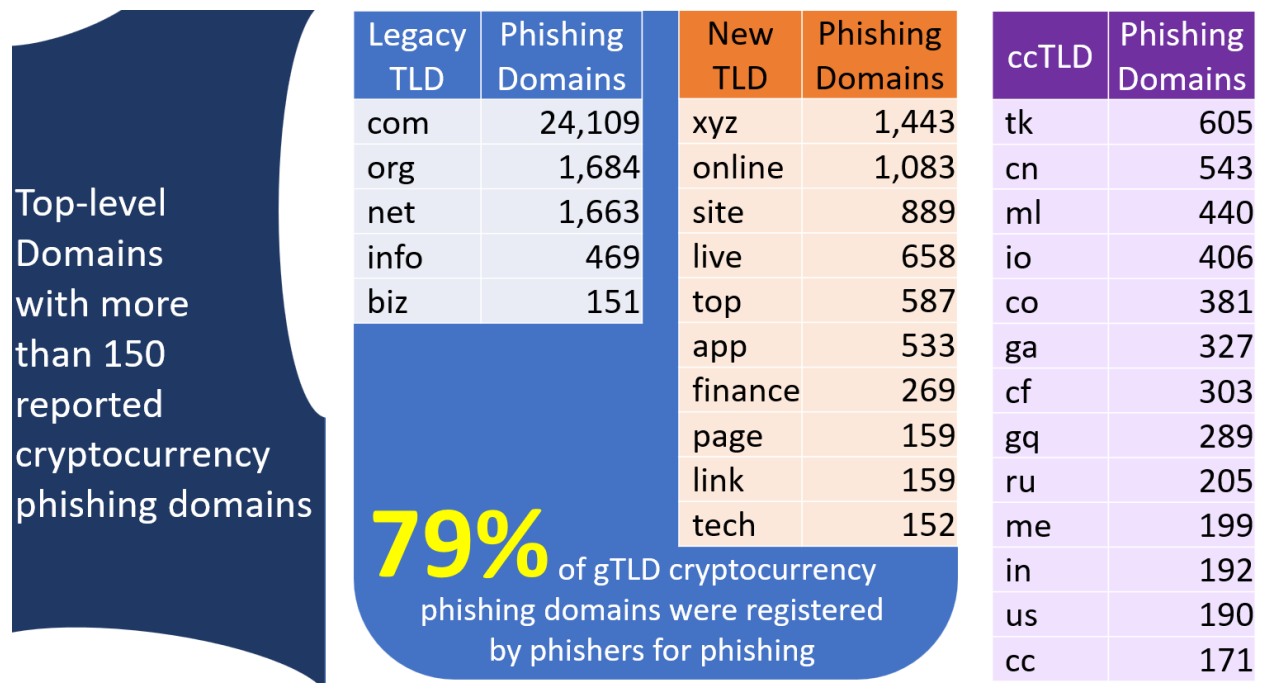
https://www.pancakeswapverify[.]com/swaps/walletconnect/connecting.php
http://binance-trustwallet[.]com/
http://walletpancakeswap[.]com/wallets/trustwallet/
http://pay-localbitcoins[.]reviews-paxful.com/
  
```

The following table shows the cryptocurrency keywords and brands that were most used in our 2022 study data.

Term or Brand	Found in URLs classified as ...		
	Wallet attacks	Exchange attacks	Cryptocurrency attacks
wallet	10,212	208	578
coinbase	24	2,004	2,003
exchange	165	833	75
pancake	844	1,637	15
coin	4	2,326	5,564
crypto	4	147	2,394
blockchain	105	12	611

Cryptocurrency Phishing in TLDs

We identified 67 Top-level Domains with more than 25 unique reported cryptocurrency phishing domains: 24 4,689 reported domains in the 24 ccTLDs, 23,680 reported domains in 43 gTLDs.



We determined that a supermajority (over 75%) of the reported domains in the gTLDs were malicious phishing domain registrations.

The inability to collect domain registration data from all ccTLDs and, to obtain the creation data of domain registrations, prevents us from determining malicious registrations for ccTLDs. We note that ccTLDs under management by Freenom (.TK, .ML, .GA, .CF, and .GQ) occupy 5 of the top 10 rankings.

Cryptocurrency Phishing and Domain Registrars

In the following table we identify the registrars where phishing domains were associated with cryptocurrency phishing attacks were registered. This table includes domain registration businesses identified in domain registration data for all gTLDs as well as ccTLDs for which we were able to collect RDAP or Whois data. Some of the businesses in this list are not ICANN-accredited.

Rank	Registrar	Cryptocurrency Phishing Domains ▼
1	NameCheap, Inc.	5,595
2	Freenom	1,900
3	PDR Ltd. d/b/a PublicDomainRegistry.com	1,743
4	OwnRegistrar, Inc.	1,741
5	NameSilo, LLC	1,436

Rank	Registrar	Cryptocurrency Phishing Domains ▼
6	GoDaddy.com, LLC	1,267
7	Registrar of Domain Names REG.RU LLC	1,188
8	GMO Internet, Inc. d/b/a Onamae.com	514
9	Hosting Concepts B.V. d/b/a Registrar.eu	514
10	Web Commerce Communications Limited dba WebNic.cc	503

Domains delegated from the gTLDs are registered through ICANN-accredited registrars, so we were also able to determine the percent of reported domains that were purposely registered by phishers, for phishing (malicious phishing domain registrations).

Rank	gTLD Registrar	Cryptocurrency Phishing Domains	Malicious Cryptocurrency Phishing Domains	Malicious Registration Percent ▼
1	NameCheap	5,546	4,555	82%
2	OwnRegistrar, Inc.	1,739	1,246	72%
3	Public Domain Registry (PDR)	1,733	1,157	67%
4	NameSilo	1,415	1,130	80%
5	GoDaddy.com	1,224	679	55%
6	Registrar of Domain Names REG.RU	1,182	981	83%
7	GMO Internet d/b/a Onamae.com	514	498	97%
8	Hosting Concepts BV d/b/a Registrar.eu	507	332	65%
9	Web Commerce Comm d/b/a WebNic.cc	502	385	77%
10	Wild West Domains	410	296	72%

Where is Cryptocurrency Phishing Hosted?

We used the DNS to resolve the domain names reported for cryptocurrency phishing to IP addresses on the date reported. We included IP addresses found in reported URLs as well. We then identified the ASN from which the IP addresses were delegated.

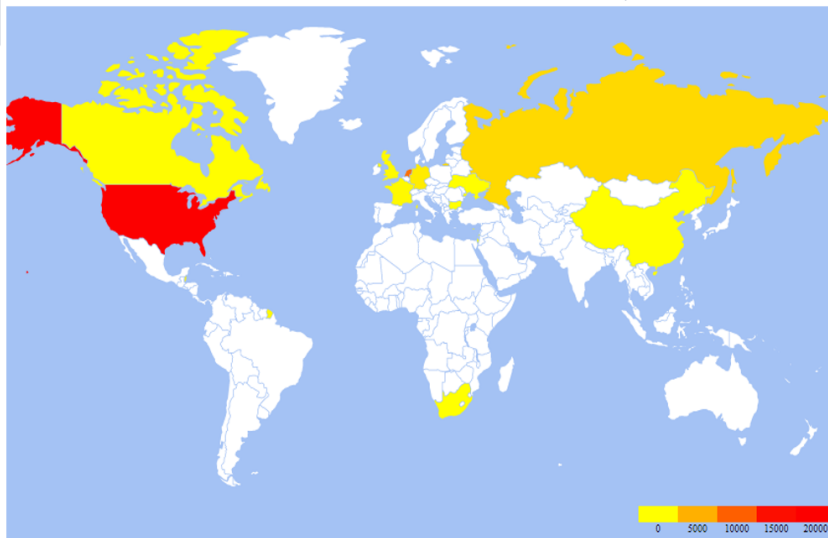
The following table ranks the ASNs by number of unique IP addresses reported hosting (or resolving to domains reported) for cryptocurrency phishing.

Autonomous System	ASN	IP addresses hosting cryptocurrency phishing ▼
MICROSOFT-CORP-MSN-AS-BLOCK	8075	8,561
CLOUDFLARENET	13335	3,751
NAMECHEAP-NET	22612	2,855
DDOS-GUARD - DDOS-GUARD LTD	57724	1,324
UNIFIEDLAYER-AS-1	46606	1,320
AMAZON-02	16509	832
DDOS-GUARD CORP.	262254	809
AS-COLOCROSSING	36352	751
AWEX - Hostinger International Limited	204915	657
FASTLY	54113	621

We next associated the ASNs with a geographic location and created a heatmap to illustrate where cryptocurrency phishing was most frequently hosted during the yearly period. **Red** in the heat scale represents the highest hosting figures, **orange** represents midpoint numbers and **yellow** the minimum. Countries in white had little or no cryptocurrency phishing.

Together, ASNs in the United States and Netherlands host more than 70% of cryptocurrency phishing (24,789 of the 34,672 unique IP addresses reported).

Where is Cryptocurrency Phishing hosted?



Country	Cryptocurrency Phishing Hosts
US	15,847
NL	8,942
RU	2,509
DE	1,296
BZ	795
CY	735
FR	497
GB	482
SC	402
CA	281
HK	266

Abuse of Subdomain Service Providers

Our analysis reveals that 12.8% of all phishing attacks took place using resources at subdomain service providers. This was up from 10.7% in our 2021 report. These phishing attacks are difficult to mitigate and pose persistent problems for phishing targets.

Subdomain services give customers services on a domain name that the provider owns. This gives users their own DNS space, on a third-level domain, of format:

subdomain.domainname.tld

Some of these providers are web hosts; some offer just the third-level domain with free DNS management so the domain owner can point it to other hosting. Yet others offer website-building services. Phishers use the domains and hosting offered by these providers to build and maintain phishing sites.

This use of subdomain services is a challenge for several reasons. Many of these companies offer the services for free. Some offer anonymous registration, with little to no identity validation. (Some don't even validate the user's email address when someone creates an account.) Finally, only the subdomain service providers can effectively mitigate these phishing attacks. These providers often lack effective, proactive measures to keep criminals from abusing their services, and some pay little attention to complaints. Some phishing kits—software used by phishers to launch and manage their phishing sites—integrate the use of subdomains providers, allowing the phishers to sign up for and use subdomains in an automated fashion. This allows the phishers to launch large numbers of attacks, and to abuse these services repeatedly.

We identified 143,506 phishing attacks using subdomains provider services, hosted on just 731 second-level domain names. Of those 143,506 attacks, 82% of them (117,752) occurred on domains operated by just ten providers. **This emphasizes how a service of this type can be used to perpetrate significant amounts of damage, and how important it is for such providers to have proactive and quick anti-abuse monitoring and takedown capabilities.** The top providers were:

Rank	Provider	Domains	Phishing attacks
1	Google	Blogspot domains, appspot.com, firebaseapp.com, web.app, business.site	34,294
2	DuckDNS	duckdns.org	24,236
3	Weebly	weebly.com	17,958
4	Hostinger	000webhosapp.com, preview-domain.com	17,306
5	Glitch	glitch.me	5,545
6	No-IP	multiple	5,466
7	ChangelIP	multiple	3,534
8	Cloudflare	trycloudflare.com, workers.dev	3,471
9	Wix	wixsite.com, usrfiles.com, filesusr.com	3,138

Rank	Provider	Domains	Phishing attacks
10	Plesk	plesk.page	2,804
11	Spaceweb.ru	swtest.ru	2,527
12	Replit	repl.co	2,188
13	IBM	appdomain.cloud	2,173
14	Timeweb.ru	tmweb.ru	2,004
15	Microsoft	azurewebsites.net, cloudapp.net, windows.net	1,798
16	Yola	yolasite.com	1,766
17	CO.VU	co.vu	1,283
18	Sprinthost.ru	xsph.ru	1,195
19	Netlify	netlify.app	1,189
20	Github	github.io	1,183

About the top-ranked subdomain service providers:

- #1 Google had phishers repeatedly take advantage of its services. These included thousands of attacks mounted on Google's Blogspot blog-building service. Thousands of phish also appeared on Appspot.com, a Google product for developing and hosting web applications in Google-managed data centers, and on Google's Business.site business profile service.
- #2 DuckDNS, #6 No-IP, and #7 ChangeI offer dynamic DNS services, which allow one to access devices from the Internet via a simple-to-remember domain name, and can obscure the real location of hosting, and so are attractive to phishers. DuckDNS is a free service.
- #3 Weebly offers a free website builder service, which is used frequently by phishers. Weebly is a subsidiary of Square, the payment processing company.
- #4 Hostinger is a hosting provider that offers free hosting on its 000webhostapp.com domain. This free service has been used prolifically by phishers for years.
- #5 Plesk offers web hosting and automation tools for managing hosting and domain names. Plesk offers free hostnames on its domains plesk.page.

The numbers above greatly under-count the number of subdomains used for phishing. In December 2021, Meta (the parent company of Facebook and Instagram) filed a lawsuit against phishers who were registering subdomains on the NGROK.IO service.^{63, 64} NGROK operates a proxy that relays traffic to websites in a manner that obscured the location of the sites, as well as the identities of the hosting providers and the users who created the subdomains. According to the lawsuit, phishers had registered and used at least 39,000 subdomains on ngrok.io from September 2019 to December 2021. However, only 810 of those attacks were reported by the sources we used to compile this study. For more about the reasons that phishing is under-reported, please see the section *List Coverage: The Phish That Get Away* above.

Use of Internationalized Domain Names (IDNs) for Phishing

Data continues to show that the unique characteristics of Internationalized Domain Names (IDNs) are not being used to facilitate phishing in meaningful numbers. Phishers do occasionally take advantage of them =, though, because they can fool the human eye, and they can evade automated detection by security programs that do not recognize the words they are constructed to represent.

IDNs are domain names that contain one or more non-ASCII characters. Such domain names can contain letters with diacritical marks such as ã and ü, or be composed of characters from non-Latin scripts such as Arabic, Chinese, or Cyrillic. Over the past sixteen years, IDNs have been available at the second and third levels in many domain name registries. IDN TLDs allow the entire domain name to be in non-Latin characters, including the TLD extension.

The IDN homographic attack is a means by which a phisher seeks to deceive Internet users by exploiting the fact that characters in different language scripts may be nearly (or wholly) indistinguishable, thereby allowing the phisher to spoof a brand name. These look-alike domains can be displayed in browser address bars if IDN display is enabled.

In our data set we saw 1,907 IDN domain names, used in 1,988 attacks. (That is up from 1,222 IDN domain names used in 1,469 attacks in our 2021 report.) The 2021-2022 total is just 0.22% of the domains used for phishing overall.

- 1,821 domains were on non-IDN TLDs, such as .COM: xn--blockchain-5v3e.com
- 176 domains were in eighteen IDN TLDs, mostly in xn--p1ai (the Russian “рф”)

We classified 682 as true homographic attacks, for example:

xn--instgrm-5fgc.gq → instagram.gq

and

xn--Intrbank-d1a.pe → Intérbank.pe

This is far above the 125 homographic attacks we observed in our 2021 report. Of the 681, 362 of them targeted a small list of cryptocurrency service providers: Blockchain.com, TrustWallet, Trezor.com, AtomicWallet, Metamask, LocalBitcoins, Coinbase, plus NFT market OpenSea. It’s possible that these domains were registered by one or two phishers who sought to steal cryptocurrency and also used IDNs to avoid having their domains detected.

xn--atmicwallet-m9b.com → atōmicwallet.com

The other 220 targeted companies such as Amazon, PayPal, BestChange, Santander, and Wells Fargo.

Additional IDNs had strings that were misleading, but the domain did not feature a brand name, such as:

xn--web-secure-verification-m8b.com → web-secure-vérification.com

Appendix A: Identification of Phishing Attacks

Phishers commonly point many URLs to one phishing site and use wildcarding⁶⁵ and redirection⁶⁶ techniques to hide the location of the phishing site from investigators. They may use a single domain name to host several discrete phishing attacks against different companies or may use multiple URLs for any given phishing site to host multiple pages.

To identify unique attacks from this diverse environment of domains, hostnames, and URLs, we examined URLs and metadata associated with URLs. We applied a set of rules to compare URLs for similarities; for example, if the hostname in two or more URLs is the same, and if the report dates for those URLs fall within 7 days of each other, and if the target across those URL reports was the same, then we treated this set of URLs as involved in one phishing attack.

Phishers use a wide variety of URL construction methods, so we formulated additional rules to group URLs into attacks based on observed cases. When we prepare our reports, we perform a final round of manual examination to find additional batches of related URL. For example, some phishers generate multiple subdomains as part of one attack. In some cases, phishers register large numbers of pseudo-randomly generated domain names (see Automating Detection of “Random-looking” Algorithmic Domain Names⁶⁷). In such cases, if the date of the abuse report and the target (brand) were the same, and the reporting feed was the same, then we grouped all those URLs as part of one attack.

Our methodology may result in underreporting the number of attacks. Others who apply a similar methodology may independently arrive at slightly different (higher) numbers; for example, if one were to use the report date window of 30 days from the research paper, COMAR: Classification of Compromised versus Maliciously Registered Domains⁶⁸, but in all other respects apply our rules, the results might identify more attacks.

Appendix B: Distinguishing Maliciously Registered Domain Names from Compromised Domains

A maliciously registered domain is defined as a domain registered by a criminal to carry out a malicious act — in this case phishing. Compromised domains are domains registered by innocent parties; an attacker leverages a vulnerability, usually in the web hosting setup, to upload a phishing page on the domain. Because they are dedicated to abuse, maliciously registered domains can be blocklisted in their entirety, and can be suspended by the domain name's registrar or registry operator. Compromised domains generally should not be approached the same way — domain suspension would affect the legitimate services on the domain. When compromised domains appear on blocklists, it is usually a specific URL that is listed, so that URL only can be blocked and prevent collateral damage to legitimate uses of the domain.

To differentiate between compromised and maliciously registered domains, operational security professionals and researchers have relied primarily on two factors:

1. The content of the domain string.
2. The age of the domain name — the number of days elapsed between domain registration and the use of the domain for a malicious purpose. In general, the older the domain name, the higher the likelihood it will legitimate. Miscreants tend to use their domains within a short time after registration to avoid detection of their registrations, and almost always within the first year of registration, before they must pay for renewal.

For this study, we refined the algorithm we used on our 2020 study. In the present study, we considered a domain to be maliciously registered if it appeared on a blacklist within fourteen days of being registered, or if a blocklisted domain had a famous brand name or misleading string in it, subject to certain time limits. We also applied additional rules that indicated common control and risk.

Our approach was at its core similar to the COMAR methodology, which was designed by researchers at two security-minded ccTLD operators, SIDN (.NL) and AFNIC (.FR).⁶⁹ COMAR's inputs are "public data," in that it is freely available or can be purchased commercially and does not contain personally private data, such as registrant data. Our data shared those characteristics.

In one way our method is more conservative than the COMAR method, which considers a domain to be maliciously registered if it appeared on a blacklist within *three months* of its registration time, or if it has a famous brand name/misleading string in the domain name. COMAR found that among compromised domains used for phishing, only 12% of the domains get compromised within three months of their registration. The implication is that a new domain name is unlikely to be compromised; it usually takes some time for a phisher to discover new domains on vulnerable hosting.

COMAR uses additional criteria to ferret out compromised domains, such as the number of web pages on a suspicious site, the use of SPF records, and a TLD maliciousness score. These additional checks help to find more maliciously registered domains than our fewer criteria; they also refine out border cases. For its phishing data, COMAR used OpenPhish, APWG, and PhishTank — three of the four sources we used.

Neither we nor the COMAR program had access to one of the most useful pieces of data available: domain name contact data, *i.e.*, information about who registered the domain name. Recent changes in ICANN policy allow registrars to redact contact data at will. Falsified contact information is an excellent indicator of bad faith on the part of the registrant, and there are ways for registrars and registry

operators to validate accuracy to various degrees of rigorousness. Also, registrars possess additional detailed data that can help them detect suspicious registrations: the registrant's payment information, the registrant's IP address, and the registrant's purchase history. These are highly useful factors to determine whether a registration is risky, and whether the registrant customer has been honest about its contact information.

Like the COMAR project, we looked for misspellings of brand names. COMAR used dnstwister and Levenshtein distance (with distance = 1) to find misspellings of brand names. They identified 231 brand names mostly targeted by attackers in phishing attacks (*e.g.*, PayPal, Amazon, Yahoo, or Gmail), and looked to see if those strings were contained in the domain name.

We created a list of more than 500 brand names that were targeted in phishing attacks. We used these as the basis for creating a list of misspellings that were distinctive enough to avoid false positives.⁷⁰ For example, we decided that "Uber" is not distinctive enough, since it is a common word in German. We complemented this list with misspellings that we encountered in our phishing URLs. We then compared that list to the domains used for phishing. We also looked for variations contained within the domain name, and this identified domains such as feddexx.com, facebaak.gq, and faceb00k-seecuurity-dept.com. Similar to COMAR, we also looked for a short list of misleading words within the domain name designed to fool victims, such as "verification" and "login".

We then performed an examination of remaining domain names. Here we relied on some additional evidence:

- We found *evidence of common control and intent*. The tests above sometime led us to *batches* of domains that were registered, used for phishing, and hosted together, indicating *common control and intent*. We also identified long strings of random and meaningless characters, whereas most domains intended for a useful purpose signify some sort of meaning.
- The Spamhaus DBL phishing feed contains a "return code" indicating whether Spamhaus considers a domain compromised (127.0.1.104, "abused legit") or a domain that may be malicious (127.0.1.4).

Our methodology and the more involved COMAR methodology created generally comparable results. One reason is that many malicious registrations are simply "beyond the pale" — they are designed to fool users and were used for phishing within a week of registration.

Appendix C: Data Sources and Methodologies

Phishing Data Sources

The use of DNS blocklists as a way to track and measure Internet abuse has a long history, and collating data reported by multiple sources is a standard procedure in academic and professional cybercrime studies.^{71, 72, 73, 74, 75} To find phishing attacks, blocklist operators use several techniques, including capturing spam email lures, reports from user, and heuristics that examine a variety of data and signals.

The following sources of phishing-specific data were chosen because they are used by a wide variety of organizations to protect users, have low false-positive rates, and have meta-data that is useful for studies such as ours.^{76, 77, 78}

- **Anti-Phishing Working Group eCrime eXchange (eCX) phishing feed.**⁷⁹ The eCX phishing feed is a repository of URLs reported to the APWG by APWG members, who are companies and government and academic investigators. Metadata associated with each uniquely identified URL includes the discovered date, targeted organization (brand) if identified, a confidence level, status (active, inactive), the discovered date, and the date of the last modification of the record.
- **OpenPhish Phishing Intelligence, premium level.**⁸⁰ The OpenPhish feed is a commercial source that contains phishing URLs discovered by OpenPhish or reported to OpenPhish directly and then verified. Metadata associated with each uniquely identified URL includes the IP address where phish was hosted, targeted brand, discovered timestamp, name of the ASN operator from which the IP address is delegated, hostname of the phish, country where the IP address is geo-located, and Top-level domain (TLD) from which the domain name in the URL was delegated.
- **PhishTank (API).**⁸¹ PhishTank is operated by OpenDNS, and publishes phishing URLs discovered by and confirmed by PhishTank community contributors. Metadata associated with each uniquely identified URL includes submission time (discovered), verification data (verified, yes/no, and verification time), status (online, yes/no), and details including IP address(es), IP network/prefix, ASN, RIR that delegated the ASN and IP allocations, and country.
- **Spamhaus Domain Block List (DBL).**⁸² The DBL is an rsync feed of registered domain names that have been associated with a malicious or criminal activity. For this study, we used only DBL-listed domains that were associated with two return codes: phish domain (127.0.1.4) and abused legit phish domain (127.0.1.104). We used as the discovery date the timestamp of each rsync access.

We collected data covering the period 1 May 2020 to 30 April 2021. We collected and analyzed only newly found phishing incidents reported during that time. We downloaded updated data from PhishTank and Spamhaus three times a day, and APWG and OpenPhish once a day. The APWG, OpenPhish, and PhishTank feeds allow the downloading of historical listings, and contain timestamps of when the listing was created. Thus we did not miss any listings that appeared between the daily downloads and did not have to worry about a delay of hours between the time the blocklist provider add an entry to its list and when we downloaded those blocklist updates. The Spamhaus DBL is stateful and does not offer “time-of-listing” time stamps, and it is possible that we missed some short-lived listings there.

These sources provide data about attacks that targeted the general public; they do not quantify “spear-phishing” attacks, which are directed at a few specific individuals and are therefore difficult to detect and count reliably.

Confidence Levels

We used only high-confidence reports in our collected data set.

- OpenPhish reports only URLs that are verified to support phishing attacks.
- The PhishTank API feed contains only phishing URLs that have been verified as supporting phish. It does not contain URLs that were reported to PhishTank but had not been verified.
- The APWG feed contains a confidence level provided by the reporting APWG member company. We used only APWG reports at the 90% level (verified by heuristics) and 100% level (verified by a human).
- The Spamhaus phishing feed does not offer confidence ratings. We consider them to be of high confidence because the Spamhaus Domain Blocklist is maintained as a “near-zero false positive list,” only containing domains that Spamhaus recommends be blocked in their entirety because they are considered dangerous. See the previous section for more about Spamhaus return codes.

Data Normalization and DNS Data

We collected reports from each feed at least once per day to find new entries. This *collected data set* then required curation to allow data from different sources to be stored together and compared. Each time a URL (or plain domain) was reported, we stored that as a separate report. Some URLs were reported by more than one source.

It was necessary to normalize certain metadata such as target (brand). For example, different sources reported slight variations of target names (“Microsoft” vs. “Microsoft Corp” vs. “Microsoft Corporation”). We normalized such examples to a common form of the company name.

UTC time is the time convention used by the four data sources, and in all gTLD registry and registrar systems including WHOIS. We used UTC.

Some sources provided IP (A record) data and AS data. For every domain reported, CIC queries DNS and separately stored the A record and determined the AS. We relied upon RIPE-NCC’s WHOIS⁸³ to find ASN name, organization, and IP prefix. When we list the number of IPv4 addresses in an AS, that is a count of routed addresses.

To identify TLDs we used the IANA root zone list.⁸⁴ We used the Public Suffix List⁸⁵ to identify registered domain names (zones in which registries offer third level registration, such as example.co.uk).

The “legacy generic TLDs” introduced before 2013 (other than .COM and .NET) are: .AERO, .ASIA, .BIZ, .CAT, .COOP, .INFO, .JOBS, .MOBI, .MUSEUM, .NAME, .ORG, .POST, .PRO, .TEL, .TRAVEL, and .XXX.

For gTLD domain names CIC obtains registry WHOIS to identify the sponsoring registrar, along with the registrar’s IANA ID⁸⁶ for normalization. Some gTLD registries severely rate-limited⁸⁷ our queries and made it impossible to obtain basic data about their domain names, including the domain registration date and the identity of the domain’s sponsoring registrar. For this reason, some gTLD domain names were not attributable to registrars and do not appear in the phishing-by-registrar tables and could not be included in the analysis of registration-to-phishing times. We did not use WHOIS for ccTLD domains due to limited access and non-uniformity of WHOIS output. Also, ccTLD registrars are not identified via a uniform identifier across ccTLD registries, making the compilation of by-registrar statistics difficult.

In the tables, the number of gTLD domains sponsored by each gTLD registrar are from the monthly ICANN reports for January 2022, the latest month available when we began writing the report.⁸⁸ The gTLD and ccTLD domain counts are gathered from DomainTools.⁸⁹

Target Identification

The APWG, OpenPhish, and PhishTank feeds identify target brand for each report; the Spamhaus DBL does not provide target information but classifies the domains according to the type of threat the domain is used to perpetrate. The sources determine target by either heuristics (which parses the content of the email phishing lure, and /or identifies the logos and wording on the phishing site), or by manual verification.

Each feed varies slightly in its granularity and nomenclature. We normalized variations in spelling — for example one feed used “PayPal” while another called it “PayPal Inc.” and so we consolidated those. Some feeds present classification differences. For example, WhatsApp is owned by Facebook. Some sources report WhatsApp as a separate brand, but another source reported the same WhatsApp phishing URLs as attacks against Facebook. In that case we accepted both and counted those as two brands.

In some cases, a source would positively identify a URL as a phish against a specific target. Another source would then report the same URL as an attack against an unknown or “generic” brand. In such cases we attributed that attack to the specific brand. In the cases where an attack’s target was still unknown, we set those attacks aside when analyzing brand data.

AS Rankings

We took into consideration previous work done to develop security reputation metrics for hosting providers.^{90, 91, 92, 93} That work notes that rankings are one way of unifying the scales on which normalized abuse is measured and allows cross comparisons, and that normalized abuse is an indicator of security performance by itself. Per the work of Noroozian *et al*⁹⁰, our work has some useful features, namely that our approach considered second-level domain-IP pairs as a unit of abuse, and that normalized abuse is abuse-type specific (because we considered phishing only).

In an AS, there may be multiple organizations which use a part of the IP space, and in the future, we wish to refine approaches to that issue. In the end we believe that our initial effort points to interesting concentrations of abuse in IP spaces under common control and are useful indicators for additional study.

About the Authors

Greg Aaron is an internationally recognized authority on the use of domain names for cybercrime, and is an expert on domain name registry operations, DNS policy, and related intellectual property issues. Mr. Aaron is Senior Research Fellow for the Anti-Phishing Working Group. As a member of ICANN's Security and Stability Advisory Committee (SSAC), he advises the international community regarding the domain name and numbering system that makes the Internet function. He works with industry, researchers, and law enforcement to investigate and mitigate cybercrime, and is also a licensed private detective. He was the Chair of ICANN's Registration Abuse Policy Working Group (RAPWG) and has been a member of ICANN's EPDP Working Group, which examined registration data access policies. He was the senior industry expert on a team that evaluated the policy and technical qualifications of more than one thousand new TLD applications to ICANN in 2012-2013. He has created products and services used by organizations to discover and track Internet-based threats, and has managed large top-level domains around the world, including .INFO, .ME, and .IN. He is President of Illumintel, Inc., a consulting company. Mr. Aaron is a *magna cum laude* graduate of the University of Pennsylvania.

Lyman Chapin has contributed to the development of technologies, standards, and policy for the Internet since 1977, and is widely recognized and respected as a leader in the networking industry and the Internet community. Mr. Chapin is a Life Fellow of the IEEE, and has chaired the Internet Architecture Board (IAB), the ACM Special Interest Group on Data Communication (SIGCOMM), and the ANSI and ISO standards groups responsible for Network and Transport layer standards. Mr. Chapin was a founding trustee of the Internet Society and a Director of the Internet Corporation for Assigned Names and Numbers (ICANN). He currently chairs ICANN's Registry Services Technical Evaluation Panel (RSTEP), which is responsible for assessing the impact of new Domain Name System (DNS) registry services on the security and stability of the Internet, and the DNS Stability Panel, which evaluates proposals for new Internationalized Domain Names (IDNs) as country code top-level domains (ccTLDs). He is also a member of ICANN's Security and Stability Advisory Committee (SSAC). He has written many other papers and articles over the past 40 years, including the original specification of the Internet standards process operated by the IETF. Mr. Chapin holds a B.A. in Mathematics from Cornell University.

David Piscitello has been involved in Internet technology and security for more than 40 years. Until July 2018, Mr. Piscitello was Vice President for Security and ICT Coordination at ICANN, where he participated in global collaborative efforts by security, operations, and law enforcement communities to mitigate Domain Name System abuse. He also coordinated ICANN's security capacity-building programs and was an invited participant in the Organisation for Economic Co-operation and Development (OECD) Security Expert Group. Dave is an Associate Fellow of the Geneva Centre for Security Policy. He served on the Boards of Directors at the Anti-Phishing Working Group (APWG) and Consumers Against Unsolicited Commercial Email (CAUCE). He is the recipient of M3AAWG's 2019 Mary Litynski Award, which recognizes the lifetime achievements of individuals who have significantly contributed to making the Internet safer.

Dr. Colin Strutt has published and spoken extensively on networking technology, name collisions, enterprise management, eBusiness, and scenario planning, and has represented the interests of Digital Equipment, Compaq, and the Financial Services Technology Consortium in national and international industry standards bodies. He holds six patents on enterprise management technology and brings more than forty years of direct experience with information technology, as a developer, architect, and consultant, with recent work including design and operation of a regional public safety network, providing technical expertise relating to patents, and analysis of world-wide Internet use. Dr. Strutt holds a B.A. (with First Class Honours) and Ph.D. in Computer Science from Essex University (UK).

About Interisle Consulting Group, LLC

Interisle's principal consultants are experienced practitioners with extensive track records in industry and academia and world-class expertise in business and technology strategy, Internet technologies and governance, financial industry applications, and software design. For more about Interisle, please visit: www.interisle.net

Acknowledgments

The authors extend thanks to:

- Spamhaus and OpenPhish, for their kind contribution of data for this study.
- The PhishTank and the APWG eCrime Exchange communities, for their collaborative efforts to identify phishing.
- John Levine, for his role in developing software for the Cybercrime Information Center.
- Domain Tools, for access to historical and bulk parsed WHOIS.
- Saeed Abu-Nimeh for access to the Seclytics Predictive Threat Intelligence platform.
- Artists Against 419, for access to its fake sites database.
- Rod Rasmussen, who co-created the Global Phishing Survey with Greg Aaron, published via the Anti-Phishing Working Group from 2007 to 2017.
- All the security personnel who fight phishing.

End Notes

¹ Cybercrime Information Center
<https://cybercrimeinfocenter.org>

² P. Foremski and P. Vixie. "Modality of Mortality in Domain Names: An In-depth Study of Domain Lifetimes." 2018.
<https://www.farsightsecurity.com/assets/media/download/VB2018-study.pdf>

³ Palo Alto Networks' cybersecurity research team studies large numbers of newly registered domain names found in zone files and concluded that 70% of them are "malicious" or "suspicious" or "not safe for work." Palo Alto Networks, Unit 42. "Newly Registered Domains: Malicious Abuse by Bad Actors." 20 August 2019.
<https://unit42.paloaltonetworks.com/newly-registered-domains-malicious-abuse-by-bad-actors/>

⁴ DomainTools
<https://www.domaintools.com/>

⁵ Phishing Landscape Study 2021, Interisle Consulting Group
<https://interisle.net/PhishingLandscape2021.pdf>

⁶ EOP/ATP – Need to block entire TLD, Microsoft Community
<https://answers.microsoft.com/en-us/msoffice/forum/all/eopatp-need-to-block-entire-tld/9f587b1c-9f1f-416a-8b18-c16a42a231cc>

⁷ Executive Order 13984 of January 19, 2021, Taking Additional Steps to Address the National Emergency with Respect to Significant Malicious Cyber-Enabled Activities
<https://www.federalregister.gov/executive-order/13984>

⁸ DNSPOD
<https://www.dnspod.cn/?lang=en>

⁹ Tencent Cloud - 腾讯云
<https://cloud.tencent.com/>

¹⁰ OwnRegistrar
<https://ownregistrar.com>

¹¹ Who is OwnRegistrar at Zoominfo
<https://www.zoominfo.com/c/ownregistrar-inc/347359162>

¹² REG.RU
<https://www.reg.ru/>

¹³ Who is REG.RU at Zoominfo
<https://www.zoominfo.com/c/regru/426544778>

¹⁴ Alibaba.com Singapore E-commerce
<https://www.alibabacloud.com/>

¹⁵ Key-Systems
<https://key-systems.net>

¹⁶ TLD Registrar Solutions
<https://www.tldregistrarolutions.com/>

¹⁷ EraNet International Ltd.

<https://www.eranet.com/>

¹⁸ BigRock Solutions

<https://www.bigrock.com/>

¹⁹ NameSilo

<https://www.namesilo.com/>

²⁰ Webnic.cc

<https://www.webnic.cc/>

²¹ Levenshtein Distance – a metric for measuring the number of characters that differ between two strings (the number of insertions, deletions, or substitutions).

https://en.wikipedia.org/wiki/Levenshtein_distance

²² How to: Block Top Level Domains, Cisco Umbrella

<https://support.umbrella.com/hc/en-us/articles/115005974587-How-to-Block-Top-Level-Domains>

²³ How to Block TLDs, Domains or Email Addresses with SpamAssassin and DirectAdmin

<https://www.vpsbasics.com/cp/how-to-block-tlds-domains-email-addresses-with-spamassassin-and-directadmin/>

²⁴ How Far Will Email Operators Take Blocklisting to Prevent Spam?

<https://www.securityskeptic.com/2017/11/how-far-will-email-operators-take-blocklisting-to-prevent-spam-.html>

²⁵ ICANN. 2013 Registrar Accreditation Agreement. Section 3.18.

<https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en>

²⁶ For examples of pricing-related issues, see the APWG Global Phishing Survey series of papers, 2007-2017, at <https://apwg.org/globalphishingsurvey/>. Further study of the effect of pricing on domain name abuse would be welcome. Such studies are hard to carry out because the sales specials and rebate programs offered by registry operators, and the bulk sale prices offered by some registrars, make it difficult to establish point-in-time wholesale and retail prices for specific domain names.

²⁷ D. Piscitello and C. Strutt. "Criminal Abuse of Domain Names: Bulk Registration and Contact Information Access." 17 October 2019.

<http://interisle.net/sub/CriminalDomainAbuse.pdf>

²⁸ Definition of hosting network (ASN), Cybercrime Information Center

<https://www.cybercrimeinfocenter.org/terminology#hostingnetwork>

²⁹ Unifiedlayer at Crunchbase

<https://www.crunchbase.com/organization/unitedlayer>

³⁰ PeeringDB: Microsoft Corporation

<https://www.peeringdb.com/net/694>

³¹ PeeringDB: Google LLC

<https://www.peeringdb.com/net/433>

³² Alibaba Cloud

<https://www.alibabacloud.com/>

³³ Quadranet: About Us

<https://www.quadranet.com/about>

³⁴ What is Fastly, the company behind the worldwide internet outage? New York Times
<https://www.nytimes.com/2021/06/08/business/fastly-internet-outage.html>

³⁵ RIPE NCC provides details for IP address, IP prefix, ASN, country code, and hostname. The following URL provides data about ASN "AS<num>"

<https://stat.ripe.net/ui2013/AS<num>#tabId=at-a-glance>

³⁶ H.R. 6352: DRUGS Act

<https://www.govtrack.us/congress/bills/117/hr6352>

³⁷ Executive Order 13984 of January 19, 2021, Taking Additional Steps to Address the National Emergency with Respect to Significant Malicious Cyber-Enabled Activities

<https://www.federalregister.gov/executive-order/13984>

³⁸ Note that this two-dimensional Venn diagram does not show the sets that were shared by just A and C, and by B and D. For this reason, if one adds up the seven numbers contained in a circle, the result will be less than the total number of domains found by that source.

³⁹ A. Pitsillidis, C. Kanich, G.M. Voelker, K. Levchenko, S. Savage. "Taster's Choice: A Comparative Analysis of Spam Feeds." Proceedings of the 2012 Internet Measurement Conference, 427-440.

<https://cseweb.ucsd.edu/~apitsill/papers/imc12.pdf>

This paper compares the contents of ten distinct feeds of spam-advertised domain names (which includes phishing URLs).

⁴⁰ L. Metcalf and J. Spring, "Everything You Wanted to Know About Blacklists But Were Afraid to Ask." Publication CERTCC-2013-39.

https://resources.sei.cmu.edu/asset_files/WhitePaper/2013_019_001_83445.pdf.

This study compares the contents of 25 different common public-internet blacklists to discover any patterns in the shared entries.

⁴¹ L. Metcalf, J. Spring, "Blacklist Ecosystem Analysis: Spanning Jan 2012 to Jun 2014." CERT Division, Software Engineering Institute, Carnegie-Mellon University. Proceedings of the 2nd ACM Workshop on Information Sharing and Collaborative Security, 12 October 2015, 13-22.

<https://discovery.ucl.ac.uk/id/eprint/10037798/>

This study compared the contents of 86 Internet blacklists, include how many lists an indicator is unique to, list sizes, expanded list characterization and intersection, etc.

⁴² L. Metcalf, E. Hatleback, J. Spring. "Blacklist Ecosystem Analysis: 2016 Update." Software Engineering Institute, CERT Coordination Center, Pittsburgh, PA. March 2016.

https://resources.sei.cmu.edu/asset_files/WhitePaper/2016_019_001_466029.pdf

This follow-study confirms that blacklists generally fail to overlap substantially with each other, suggesting that available blacklists present an incomplete and fragmented picture of the malicious infrastructure on the Internet.

⁴³ V. Guo Li, M. Dunn, P. Pearce, D. McCoy, G. Voelker, S. Savage, K. Levchenko. "Reading the Tea Leaves: A Comparative Analysis of Threat Intelligence." Proceedings of the 28th USENIX Security Symposium. August 14–16, 2019, Santa Clara, CA, USA.

https://www.usenix.org/system/files/sec19-li-vector_guo.pdf

This recent study systematically characterizes a broad range of public and commercial sources of threat intelligence. Although it concentrates on IP address and file hashes, the findings seem generally applicable to domain name-based threat intelligence feeds.

⁴⁴ H.Griffioen, T. Booij and C Doerr. "Quality Evaluation of Cyber Threat Intelligence Feeds." International Conference on Applied Cryptography and Network Security (ACNS), May 2020.
https://www.researchgate.net/publication/341385656_Quality_Evaluation_of_Cyber_Threat_Intelligence_Feeds

⁴⁵ FireHOL.

<http://iplists.firehol.org/#comparison>

This is a comparison of IP blocklists, noting overlaps.

⁴⁶ C. Kanich, N. Chachra, D. McCoy, C. Grier, D.Y. Wang et al. "No Plan Survives Contact: Experience with Cybercrime Measurement." CSET, 2011.

<http://damonmccoy.com/papers/cset11kanich.pdf>

⁴⁷ A. Oest, Y. Safaei, A. Doupé, G. Ahn, B. Wardman, and G. Warner. "Inside a Phisher's Mind: Understanding the Anti-Phishing Ecosystem Through Phishing Kit Analysis". In Proceedings of the 2018 APWG Symposium on Electronic Crime Research (eCrime), pages 1–12, May 2018.

<https://docs.apwg.org/ecrimeresearch/2018/5349207.pdf>

⁴⁸ Palo Alto Networks, Unit 42. "Newly Registered Domains: Malicious Abuse by Bad Actors." 20 August 2019.

<https://unit42.paloaltonetworks.com/newly-registered-domains-malicious-abuse-by-bad-actors/>

⁴⁹ A. Oest, Y. Safaei, A. Doupé, G. Ahn, B. Wardman, and K. Tyers. "PhishFarm: A Scalable Framework for Measuring the Effectiveness of Evasion Techniques Against Browser Phishing Blacklists." In: 2019 IEEE Symposium on Security and Privacy (SP), 19-23 May 2019.

<https://ieeexplore.ieee.org/document/8835369>

⁵⁰ A. Oest, Y. Safaei, P. Zhang, B. Wardman, et al: "PhishTime: Continuous Longitudinal Measurement of the Effectiveness of Anti-phishing Blacklists." Proceedings of the 29th USENIX Security Symposium, August 12–14, 2020.

<https://www.usenix.org/system/files/sec20-oest-phishtime.pdf>

⁵¹ T. Moore and R. Clayton. "How Hard Can it Be to Measure Phishing?" Mapping and Measuring Cybercrime, 2010.

<https://www.cl.cam.ac.uk/~rnc1/cyberbias.pdf>

⁵² Europol European Cybercrime Centre (EC3). "The Indispensable Role of WHOIS for Global Cybersecurity: Statement by the EC3 Advisory Group on Internet Security." 25 January 2018.

<https://www.icann.org/en/system/files/files/gdpr-statement-ec3-europol-icann-proposed-compliance-models-25jan18-en.pdf>

⁵³ Messaging, Malware and Mobile Anti-Abuse Working Group (M3AAWG) and Anti-Phishing Working Group. "ICANN GDPR WHOIS Policy Eliminates Pre-emptive Protection of Internet Infrastructure Abuse; Obstructs Routine Forensics to Cybercriminals' Advantage." 24 October 2018.

<https://www.m3aawg.org/rel-WHOISsurvey2018-10>

⁵⁴ D. Piscitello. "Facts & Figures: WHOIS Policy Changes Impair Blocklisting Defenses." 8 March 2019.

<https://www.securityskeptic.com/2019/03/facts-figures-whois-policy-changes-impair-blacklisting-defenses.html>

⁵⁵ "Taking Legal Action Against Phishing Attacks." 20 December 2021. <https://about.fb.com/news/2021/12/taking-legal-action-against-phishing-attacks/>

⁵⁶ "Meta Files Federal Lawsuit Against Phishing Operators." 21 December 2021.

<https://www.darkreading.com/attacks-breaches/meta-files-federal-lawsuit-against-phishing-operators>

-
- ⁵⁷ Cryptocurrency market value tops \$2 trillion for the first time as Ethereum hits record high, <https://www.cnn.com/2021/04/06/cryptocurrency-market-cap-tops-2-trillion-for-the-first-time.html>
- ⁵⁸ The 2022 Crypto Crime Report, Chainalysis <https://go.chainalysis.com/rs/503-FAP-074/images/Crypto-Crime-Report-2022.pdf>
- ⁵⁹ Twitter post by @MyEthernWallet warns of phishing attack <https://twitter.com/myetherwallet/status/995043917878902784>
- ⁶⁰ Ledger owners lose 1.1 million XRP to scam site <https://cointelegraph.com/news/ledger-owners-lose-1-1-million-xrp-to-scam-site>
- ⁶¹ Binance Phishing Attack Is Underway, Warns CEO Changpeng Zhao, Finance Magnates <https://www.financemagnates.com/cryptocurrency/binance-phishing-attack-is-underway-warns-ceo-changpeng-zhao/>
- ⁶² What is PancakeSwap and How does it Work? Bybit.com <https://learn.bybit.com/defi/what-is-pancakeswap-and-how-does-it-work/>
- ⁶³ "Taking Legal Action Against Phishing Attacks." 20 December 2021 <https://about.fb.com/news/2021/12/taking-legal-action-against-phishing-attacks/>
- ⁶⁴ "Meta Files Federal Lawsuit Against Phishing Operators." 21 December 2021 <https://www.darkreading.com/attacks-breaches/meta-files-federal-lawsuit-against-phishing-operators>
- ⁶⁵ Definition of Wildcarding <https://tools.ietf.org/html/rfc4592>
- ⁶⁶ DNS hijacking and redirection <https://www.imperva.com/learn/application-security/dns-hijacking-redirection/>
- ⁶⁷ Automating Detection of "Random-looking" Algorithmic Domain Names <https://www.farsightsecurity.com/blog/txt-record/automatingdetection-20190517/>
- ⁶⁸ COMAR: Classification of Compromised versus Maliciously Registered Domains http://mkorzynski.com/COMAR_2020_IEEEEuroSP.pdf
- ⁶⁹ Maroofi, M. Korczynski, C. Hesselman, B. Ampeau, A. Dud, "COMAR: Classification of Compromised versus Maliciously Registered Domains." 2020 IEEE European Symposium on Security and Privacy (EuroS&P) http://mkorzynski.com/COMAR_2020_IEEEEuroSP.pdf and <https://comar-project.univ-grenoble-alpes.fr/>
- ⁷⁰ W. Wang and K. Shirley, "Breaking Bad: Detecting Malicious Domains Using Word Segmentation," In Proc. 9th Workshop on Web 2.0 Security and Privacy, 2015.
- ⁷¹ A. Oest, Y. Safaei, A. Doupé, G. Ahn, B. Wardman, and K. Tyers. "PhishFarm: A Scalable Framework for Measuring the Effectiveness of Evasion Techniques Against Browser Phishing Blacklists." In: 2019 IEEE Symposium on Security and Privacy (SP), 19-23 May 2019. <https://ieeexplore.ieee.org/document/8835369>
- ⁷² D. Piscitello, G. Aaron. "Domain Abuse Activity Reporting (DAAR) System Methodology". Internet Corporation for Assigned Names and Numbers (ICANN). November 2017. <https://www.icann.org/en/system/files/files/daar-methodology-paper-30nov17-en.pdf>

-
- ⁷³ Dietrich C.J., Rossow C. (2009) Empirical research of IP blacklists. In: Pohlmann N., Reimer H., Schneider W. (eds) ISSE 2008 Securing Electronic Business Processes. Vieweg+Teubner.
https://doi.org/10.1007/978-3-8348-9283-6_17
- ⁷⁴ S. Maroofi, M. Korczynski, C. Hesselman, B. Ampeau, A. Dud, "COMAR: Classification of Compromised versus Maliciously Registered Domains." 2020 IEEE European Symposium on Security and Privacy (EuroS&P). http://mkorczynski.com/COMAR_2020_IEEEEuroSP.pdf and <https://comar-project.univ-grenoble-alpes.fr/>
- ⁷⁵ Pitsillidis, C. Kanich, G.M. Voelker, K. Levchenko, S. Savage. "Taster's Choice: A Comparative Analysis of Spam Feeds." Proceedings of the 2012 Internet Measurement Conference, 427-440.
<https://cseweb.ucsd.edu/~apitsill/papers/imc12.pdf>
- ⁷⁶ D. Piscitello. "Reputation Block Lists: Protecting Users Everywhere." 1 November 2017. Internet Corporation for Names and Numbers (ICANN).
<https://www.icann.org/news/blog/reputation-block-lists-protecting-users-everywhere>
- ⁷⁷ B. Greene. "What Makes a Good 'DNS Blacklist'?"
<https://blogs.akamai.com/2017/08/what-makes-a-good-dns-blacklist.html> and <https://www.akamai.com/us/en/products/security/enterprise-threat-protector.jsp>
- ⁷⁸ G. Aaron, D. Piscitello. "Domain Abuse Activity Reporting (DAAR) System Methodology". Internet Corporation for Assigned Names and Numbers (ICANN). November 2017.
<https://www.icann.org/en/system/files/files/daar-methodology-paper-30nov17-en.pdf>
- ⁷⁹ Anti-Phishing Working Group eCrime Exchange
<https://apwg.org/ecx/>
- ⁸⁰ OpenPhish
<https://openphish.com>
- ⁸¹ PhishTank
<https://www.phishtank.com/>
- ⁸² The Spamhaus Project
<https://www.spamhaus.org/>
- ⁸³ RIPE-NCC
<https://stat.ripe.net/> and <https://www.ripe.net/manage-ips-and-asns/db/tools>
- ⁸⁴ IANA root zone list
<https://www.iana.org/domains/root/db>
- ⁸⁵ Public Suffix List
<https://publicsuffix.org/>
- ⁸⁶ IANA Registrar IDs
<https://www.iana.org/assignments/registrar-ids/registrar-ids.xhtml>
- ⁸⁷ ICANN Security and Stability Advisory Committee (SSAC): *SAC101v2: SSAC Advisory Regarding Access to Domain Name Registration Data*. 12 December 2018
<https://www.icann.org/en/system/files/files/sac-101-v2-en.pdf>
- ⁸⁸ ICANN Monthly Registry Reports
<https://www.icann.org/resources/pages/registry-reports>

⁸⁹ DomainTools, Domain Count Statistics for TLDs

<https://research.domaintools.com/statistics/tld-counts/>

⁹⁰ A. Noroozian, M. Korczynski, S. Tajalizadehkhoob, and M. van Eeten, "Developing Security Reputation Metrics for Hosting Providers," in Proc. Workshop on Cyber Security Experimentation and Test (CSET), 2015

<http://mkorczynski.com/UsenixCSETNoroozian.pdf>

⁹¹ M. Korczynski, S. Tajalizadehkhoob, A. Noroozian, M. Wullink, C. Hesselman, and M. v. Eeten, "Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs," in IEEE EuroS&P, 2017, pp. 579–594

<http://mkorczynski.com/SPEurope2017Korczynski.pdf>

⁹² B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda. "FIRE: Finding Rogue nEtworks". In: ACSAC. 2009, pp. 231–240

https://sites.cs.ucsb.edu/~chris/research/doc/acsac09_fire.pdf

⁹³ J. Armin, editor. "World Hosts Report 2014"

<http://hostexploit.com/?p=whr-201403>