Whacking Moles: Blocklists and Their Role in the Endless Cycle of Malicious Domain Registration

Author: Shawn Reinhart, shawn.reinhart@gmail.com

Advisor: Jonathan Risto

Accepted: December 20, 2023

Abstract

Filtering out the evildoers on the Internet is an endless and often unavailing task. With millions of new domains registered daily, blocklists struggle to distinguish between the good and the bad. How much of the known Internet is evil? How long does it typically take for malicious phishing activity to be identified and verified as dangerous? Free and commercial blocklists exist to help organizations tackle the problem of phishing domains, but can such blocklists respond quickly enough to be effective? This research examines possible ways to answer these questions by combining data freely available from online sources.

1. Introduction

Millions of new Internet domains are registered every day, and the pace at which new registrations occur continues to grow year after year (Domain Name Industry Brief [DNIB], 2023). Among all the new domains registered on any given day, many will be domains registered by ill-intentioned cyber actors solely for use in future phishing or malware campaigns. Distinguishing between legitimate and illegitimate domains has long been a challenge for network defenders, and some critical questions for the cybersecurity community related to malicious domain registration include: How long does it typically take before the actors who register such domains begin using them in their malicious campaigns? How many new domains will eventually be identified and reported as malicious? How long does it typically take for a blocklist organization to verify and block them?

1.1. Domain Registration

The Internet Corporation for Assigned Names and Numbers (ICANN) is responsible for the global assignment of IP addresses as well as the management of the top-level domain (TLD) name system (Internet Corporation for Assigned Names and Numbering [ICANN], 2012). ICANN allocates blocks of IP addresses to Regional Internet Registrars (RIRs), who then further subdivide the blocks of IP addresses, assign Autonomous System Numbers (ASNs), and provide services to Internet service providers (ISPs) in their geographical service region (Number Resource Organization [NRO], 2023). Finally, domain names may be registered with one of the more than two thousand ICANN-accredited registrars or their authorized resellers (ICANN, 2017).

According to the Domain Name Industry Brief's most recent quarterly report, more than 356 million domains were registered in the second quarter of 2023, an increase of 4.3 million compared to the same quarter of 2022 (DNIB, 2023). The pool of available domains also grows yearly, with the total number of Internet TLDs currently standing at 1,458 (Internet Assigned Numbers Authority [IANA], 2023). The continued creation of new TLDs is a boon for malicious actors, as it effectively allows for an infinite number of Internet domains, with many new generic TLDs having either already been added or proposed in 2023 ("New Domain (gTLD) Launch

Schedule," 2023). For example, even Google launched eight new TLDs as recently as May 2023, including .foo, .zip, and .mov (Yeh, 2023).

1.2. Phishing

Phishing can be described as "a fraudulent attempt, usually made through email, to steal your personal information" (PhishTank, n.d.c). In recent years, phishing via SMS messages ("smishing") (Trend Micro, n.d.) and direct voice phone calls ("vishing") (Lenaerts-Bergmans, 2023) has also become prevalent. Phishing emails usually attempt to appear to be from legitimate organizations and can often be identified via the following tell-tale signs:

- Generic greeting: Since phishing emails are often untargeted and sent as part of mass emailing campaigns, messages will often begin with openings such as "dear customer" or a similar generic greeting.
- Requests for personal information: Since the primary purpose of phishing is to trick recipients into providing personal details, phishing messages will often explicitly request a reply with details such as name, address, phone number, age, or Social Security Number.
- Sense of urgency: Phishing emails often persuade recipients to respond quickly, using phrasing such as "final warning" or "your account will be locked" in the subject line and email body. The goal is to get victims to not think about the message's legitimacy and instead provoke an emotional response based on worry or fear.
- Bogus links: Phishing emails will almost always include one or more forged links that appear to be from a legitimate organization but, when moused over, reveal an unrelated domain. Sometimes, they are intentionally crafted to appear to be a legitimate domain, and in other cases, they are entirely unrelated but hidden behind the text in the HTML link.

Despite the growing organizational awareness of phishing and improvements in employee security awareness programs, phishing remains a significant problem for businesses today. According to Proofpoint's most recent State of the Phish report, 84% of survey

respondents had experienced at least one successful email-based phishing attack in 2022, and 58% were forced to deal with three or more attacks (Proofpoint, 2023).

1.3. Tactics and Techniques

Registering domains for malicious use and sending links containing those domains in phishing emails is a well-established cybercrime tradecraft. The MITRE ATT&CK framework (see Figure 1) is a classification system and knowledge base of adversary tactics, techniques, and sub-techniques based on real-world observations, which defenders can use to develop threat models and adversary emulation plans (MITRE, 2023a). The MITRE ATT&CK Enterprise Matrix classifies phishing (T1566) as a technique under the tactic of gaining initial access (TA0001) and describes it as "electronically delivered social engineering" (MITRE, 2023d). Sending malicious links is further classified as a sub-technique (T1566.002), and MITRE describes it as a specific variant of spearphishing in which adversaries send emails containing a malicious link in the hopes of gaining access to victim systems (MITRE, 2023c). MITRE notes that phishing may involve social engineering techniques, such as posing as a trusted source, and evasive techniques, such as manipulating email headers. MITRE also indicates that, generally, these links are accompanied by social engineering text and require a user to actively click or copy and paste a URL into a browser.

The acquisition of domains also falls under the MITRE ATT&CK Enterprise Matrix, which classifies it as a sub-technique (T1583.001) under the technique of acquiring infrastructure (T1583) (MITRE, 2023b). In describing this sub-technique, MITRE notes that adversaries may acquire domains for various reasons, including phishing and command and control. MITRE also points out that adversaries may choose domains visually similar to legitimate ones using homoglyphs, different top-level domains, or non-English character sets.

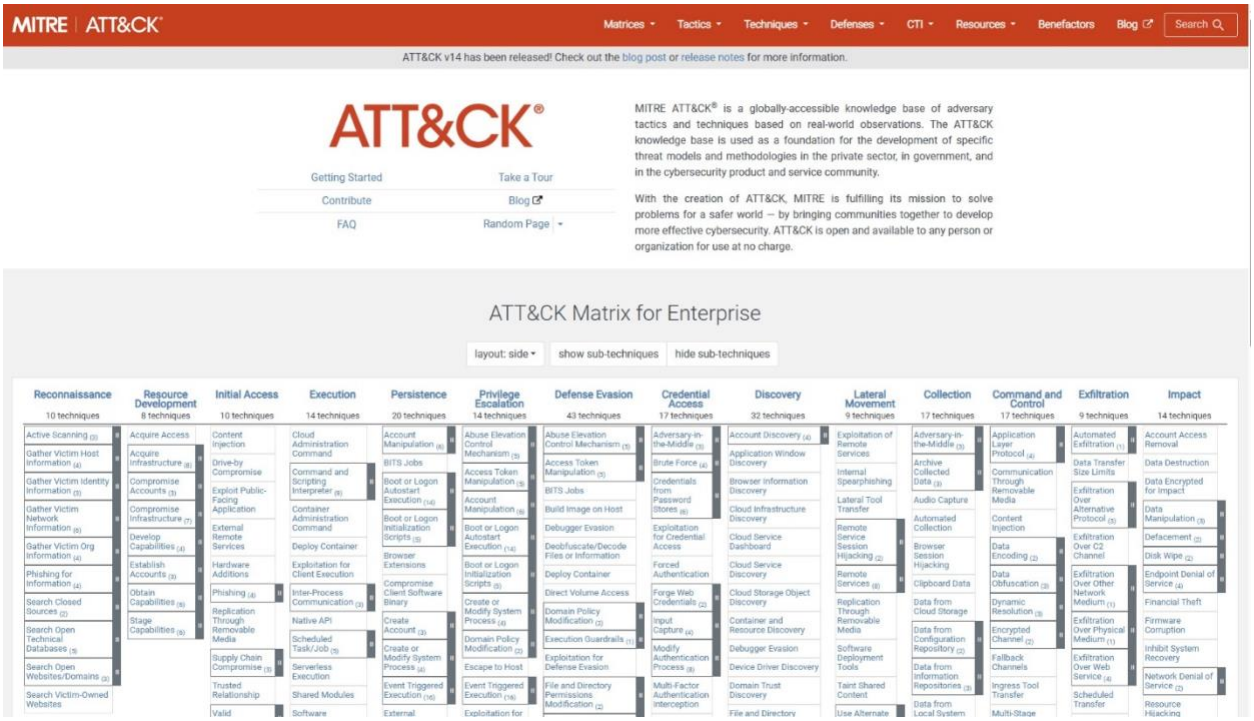

Figure 1. Screenshot of MITRE ATT&CK Matrix for Enterprise

1.4. Blocklists

A blocklist, sometimes known as a Reputation Block List (RBL), can be defined as a collection of source addresses that "have been deemed undesirable, where typically these addresses have been involved in some previous illicit activities" (J. Zhang et al., 2008). Blocklists are created and maintained by commercial service providers, researchers, and non-profit communities and are operated to detect or alert about security threats (Aaron & Piscitello, 2017). Blocklists vary in focus and detection methods, and they can be compiled through different means, such as direct submission by phishing victims or automatically using computer algorithms and data science methods.

Blocklists can be tailored toward specific threat models, and some blocklists have been explicitly designed to block spam, malware, advertising, Internet telemetry and trackers, and other content deemed generally unsafe, particularly for use in parental control mechanisms.

Some examples of well-known blocklists explicitly created to block phishing sites include Google Safe Browsing, OpenPhish, and PhishTank.

1.4.1. Google Safe Browsing

Google Safe Browsing was launched in 2005 to protect Internet users from phishing attacks (Google, n.d.a). It is used by Chrome, Google Play, and Gmail, among others, to protect millions of devices by showing warnings when users attempt to access dangerous sites or download malicious files. Safe Browsing also notifies web admins when their sites have been compromised and helps them diagnose and resolve problems. Enhanced Safe Browsing is an option that can further be enabled in Google account settings or the Chrome browser to provide real-time checks against a comprehensive blocklist of known phishing and malware sites. It can also provide additional checks on attachments and web links for Gmail users. Google provides developers with a Safe Browsing API, which is free for non-commercial use (Google, n.d.b).

1.4.2. OpenPhish

The OpenPhish Database is an SQLite database accessed via a free, open-source API (OpenPhish, n.d.b). It is a continuously updated database that contains searchable information on phishing websites, such as hostnames, paths, and SSL certificate metadata. The database also contains metadata that can be used for cyber incident analysis, pattern searching, and training for AI applications. OpenPhish gathers its phishing URLs via autonomous systems using an advanced detection engine that operates without direct human intervention (OpenPhish, n.d.a). OpenPhish's detection engine has a brand identification feature that automatically associates a phishing URL with its targeted brand. OpenPhish provides a monthly list of successfully identified brands, which companies can access to determine whether their company is being impersonated.

1.4.3. PhishTank

PhishTank is a free community blocklist site operated by Cisco Talos Intelligence Group, where anyone may submit phishing data for verification and possible addition to the PhishTank blocklist (PhishTank, n.d.a). They provide a lookup service and a free API for developers and security professionals to download the full blocklist of verified online phishing URLs.

PhishTank relies on a community of human volunteers to manually verify suspected phishing URLs. Members will examine individual URL phishing details, possibly visit the suspicious site, and then submit a vote, with the number of votes required to verify a phishing link depending on the history of those voting. Some well-known organizations using PhishTank data include Yahoo Mail, Mozilla, McAfee, and Kaspersky (PhishTank, n.d.b).

2. Research Method

Domain registration by malicious actors can be observed to follow a "whack-a-mole" game-like cyclical pattern where registration is repeated frequently and consists of distinct stages, in which adversaries first register domains, and then begin using them in phishing attacks. The domains are identified as suspected phishing sites, reported as suspicious, and added to blocklists; eventually, the domains are taken down, and then adversaries launch the cycle once again by registering more domains. This paper will attempt to examine some of the typical timeframes between the following stages:

- when malicious actors first register a new domain;
- when malicious actors begin using the said domain in their operations;
- when phishing victims or monitoring systems first identify and report such a domain to a blocklist as being potentially dangerous; and
- when the blocklist organization finally confirms the domain as malicious and adds it to their blocklist.

2.1. Related Research

Previous research by Bell and Komisarczuk (2020) compared the detection time among different blocklists and specifically looked at the difference in detection time between blocklists using automated detection and those using manual submission and verification methods; the focus of that research, however, was on the numbers of URLs added to the blocklists and their duration in the blocklists. Research by Ramsey (2021) also examined whether automating takedown requests for verified phishing sites could shorten the time between blocklist

verification and sites being taken down or removed. None of this research specifically analyzed the length of time between initial domain usage and blocklist inclusion.

2.2. Methodology

All domains reported as first seen from September 1 to September 30, 2023, were examined for this research. The author then monitored the PhishTank blocklist at multiple points over the following month for the appearance of those domains. The main goal was to determine (1) what percentage of them were eventually confirmed to be malicious, (2) how long, on average, it took for malicious actors to begin using these domains and for victims to report such suspicious activity, and (3) how long it typically took for verification by the blocklist organization.

2.3. Data Sources and Collection

This research primarily relied on two publicly available data sources: the DShield database and the PhishTank verified online blocklist. The information in these databases is made freely available via organization-supported APIs.

2.3.1. DShield Database

To determine new domains first seen on a specific date, the author used the Recent Domains API (see Figure 2), made publicly available as part of the SANS Internet Storm Center's DShield API (SANS Internet Storm Center [ISC], 2023a). The DShield is a distributed intrusion detection system used by the Internet Storm Center for data collection and analysis (SANS ISC, 2023b). It is a volunteer effort, with thousands of sensors around the globe submitting firewall and intrusion detection system (IDS) logs for analysis by humans and machines looking for abnormal trends and activity. The DShield Recent Domains query returns a list of all domains appearing in the DShield database reported as first seen on a given date. For this specific experiment, the author looked at all domains first appearing on each day of September 2023.

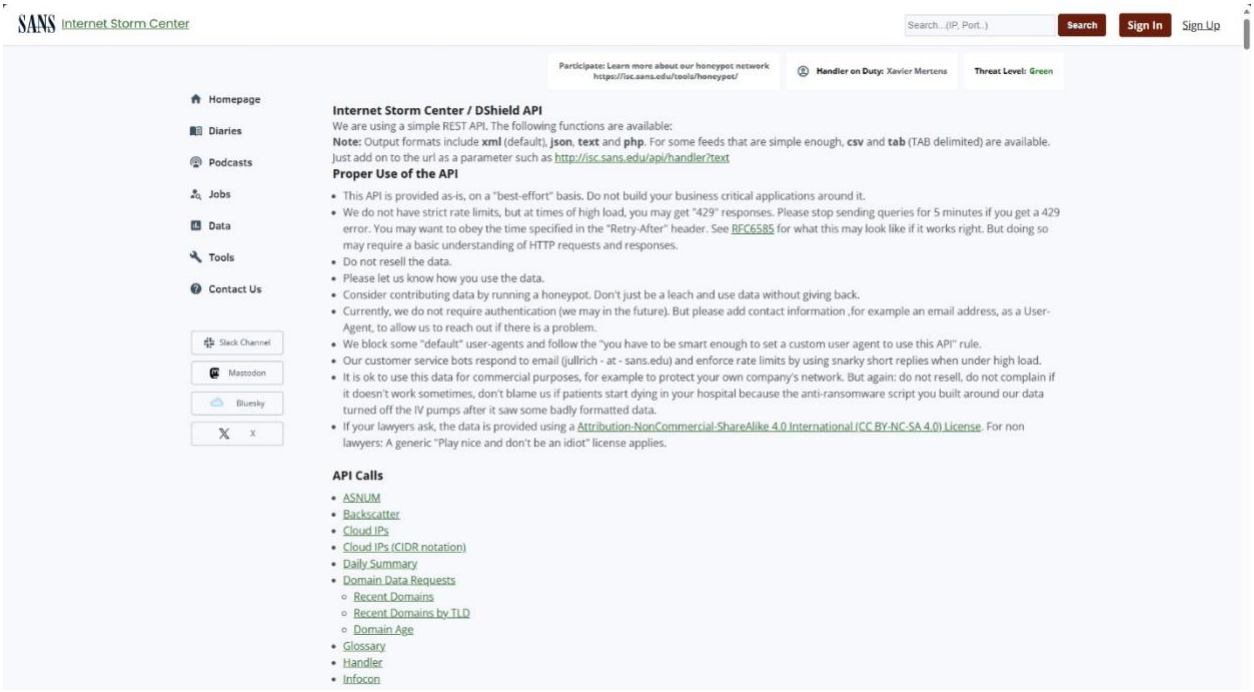

Figure 2. Screenshot of SANS Internet Storm Center

2.3.2. PhishTank Blocklist

The domains found in the results of the DShield database were then compared to the URLs appearing in the PhishTank verified online blocklist data, which is freely downloadable and contains a list of all phishing URLs that have been verified by the PhishTank community and that are currently online. Snapshots of the PhishTank blocklist were downloaded separately on October 15, October 25, and November 1, 2023, and then merged into a single large blocklist after removing duplicate entries. Because the PhishTank blocklist data (see Figure 3) consists of complete URLs, the domain portion of each URL was extracted along with timestamps for the date submitted and the date verified to compare with the DShield data. In the case of multiple blocked URLs containing the same base domain, this research included only the earliest appearance of that domain in the blocklist.

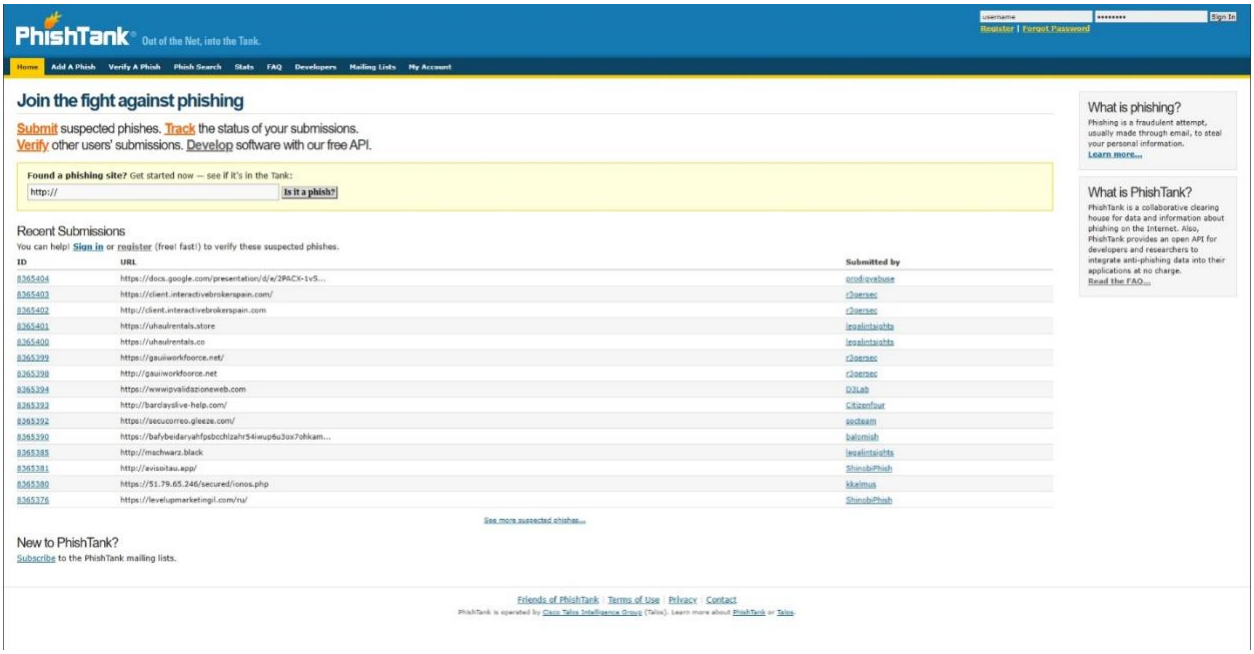

Figure 3. Screenshot of phishtank.org

2.3.3. Limitations of the Data

It is noteworthy that although the DShield employs a distributed network of sensors worldwide, it only observes a sample of the Internet and does not present a complete picture. There are undoubtedly malicious domains active on the Internet that go unseen by the DShield network; therefore, these would not appear in the results of this experiment. It should also be pointed out that the PhishTank blocklist is a real-time snapshot of phishing domains that have not only been verified but are also currently online. For this reason, there may be instances of domains that were submitted, verified, and blocked, taken down by the malicious actors responsible for them, and then removed from the blocklist, which occurred at a time either before, after, or in between the points in time at which the author downloaded snapshots of the blocklist. Such activity would also not be reflected in the results. More frequently collected data could likely provide higher quality results, but gathering such data would require constant monitoring of the blocklist (at least one or more times per day) over a longer duration (perhaps six months or a year, to be comprehensive) and such collection activity fell outside the limited scope of this project.

3. Findings and Discussion

3.1. New Domains

The data from the DShield database revealed that a total of 15.8 million domains were first seen during the month of September 2023, or approximately 527,000 per day.

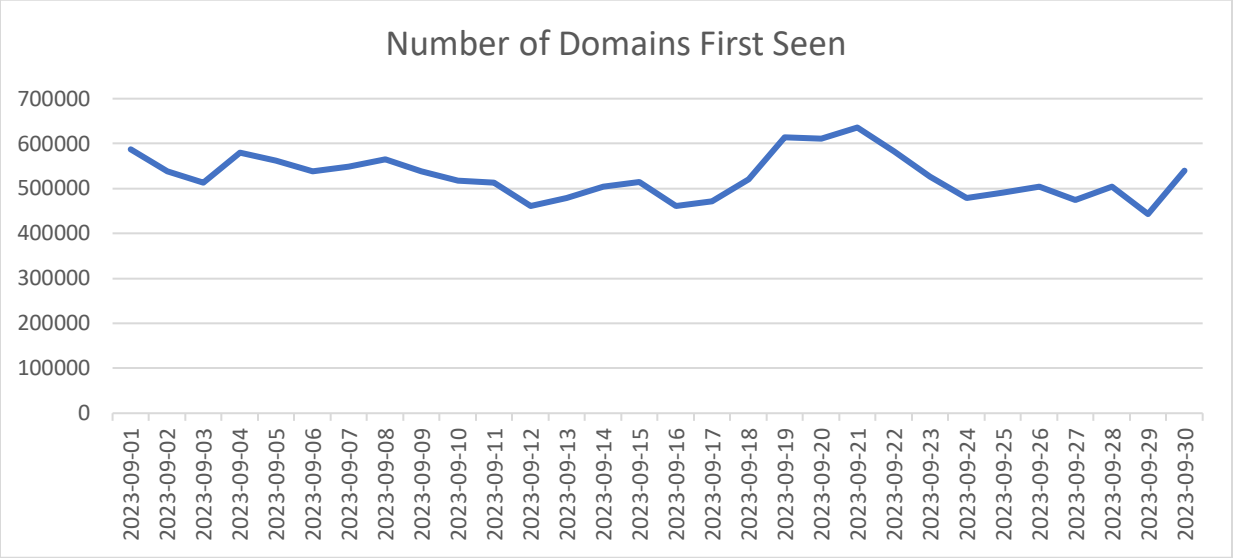

Figure 4. Domains First Seen in September 2023

3.2. Verified Phishing URLs

The PhishTank blocklist is constantly being updated as new phishing URLs are verified and added, and phishing sites that go offline are removed from the list. The samples of the PhishTank blocklist taken on October 15, October 25, and November 1 showed that the blocklist contained, on average, approximately 38,200 verified phishing URLs at any given time.

Date Sampled	Number of Verified URLs
2023-10-15	38,968
2023-10-25	37,552
2023-11-01	38,172
Average	38,231

Figure 5. Number of Verified Phishing URLs

Reviewing the different URLs appearing on the blocklist across the three samples, the total number of unique URLs found was 48,364. After extracting only the domain portion of the URLs and removing duplicates, the total number of unique malicious domains was 32,298.

Type of Unique Total	Total
Total Number of Unique URLs	48,364
Total Number of Unique Domains	32,298

Figure 6. PhishTank Unique URLs and Domains

3.3. Comparison Results

Comparing those domains from the DShield database that were reported as first seen in September 2023 to the domains seen in the combined PhishTank verified online blocklist data, the results showed that, as of November 1, 2023, only 558 unique domains had been verified as phishing sites or 0.0035% of new domains.

Value Calculated	Result
Total Number of New Domains Reported Seen	15,819,931
Total Number of New Domains Found in Blocklist	558
Percentage of New Domains Found in Blocklist	0.0035%

Figure 7. Comparison of Newly Seen Domains to Blocklisted Domains

Somewhat confusingly, the PhishTank data also revealed that many of the domains from the DShield database that appeared in the blocklist, 221 to be exact, were first submitted and verified on dates *before* the date they were reported as first seen in the DShield data; possible reasons for this will be examined in the discussion section of this paper, but to eliminate any chance of abnormalities in the results due to using improper data, data related to these domains was removed from the subsequent calculations regarding the time until submission and time until verification.

3.4. Time Until Submission

For the remaining 337 domains that were submitted to PhishTank for suspected phishing on or after the date they were first seen, the average number of days between the date first seen and the date submitted was 13.7 days, with the minimum being zero days (meaning the domain was submitted to the blocklist the same day it was first seen), and the maximum being 60 days.

Time Measured	Number of Days
Minimum Time Until Submission	0
Maximum Time Until Submission	60
Average Time Until Submission	13.7

Figure 8. Minimum, Maximum, and Average Time Until Submission

It is important to note that since the last date the PhishTank blocklist was monitored was November 1, 2023, and because the research only compared domains first seen on or after September 1, 2023, the maximum possible time until submission value in this dataset would be 61 days. Any new domains that took more than 61 days to be submitted would not appear in the results. Likewise, since the earliest snapshot of the PhishTank data that the author recorded was taken on October 15, 2023, as noted in the data limitations section above, any domains first seen on or after September 1 but that underwent the complete cycle of being registered, actively used, reported as suspicious, and then taken offline and removed from the blocklist before October 15 also would not appear in the collected blocklist data and therefore would be excluded from the results.

3.5. Time Until Verification

As for PhishTank's voting-based, manual verification process, the average time until verification was 398 minutes, or 6 hours and 38 minutes. It is worth noting, however, that there was an extensive range of values seen in the time until verification results, which included some rather obvious outliers; the results showed that the shortest time until verification was only 1 minute, but the longest was 18,160 minutes (12 days, 14 hours, and 40 minutes).

Time Measured	Number of Minutes
Minimum Time Until Verification	1
Maximum Time Until Verification	18,160
Average Time Until Verification	398

Figure 9. Minimum, Maximum, and Average Time Until Verification

Instances of verification taking more than 24 hours were relatively rare, with 95.5% (322 out of 337) of all submitted domains verified within the first 24 hours. By removing the fifteen domains with unusually long verification times and focusing only on those domains that were verified within a day, the average time until verification shrinks to only 64 minutes.

Time Measured	Number of Minutes
Minimum Time Until Verification	1
Maximum Time Until Verification	1,408
Average Time Until Verification	64

Figure 10. Minimum, Maximum, and Average Time Until Verification for Domains Verified within the First 24 Hours after Submission

For further comparison, looking specifically at the fifteen outliers that took more than 24 hours to verify, the average time until verification for those domains was 7,556 minutes (5 days, 5 hours, and 56 minutes), well over one hundred times longer compared to those verified within the first 24 hours after submission.

Time Measured	Number of Minutes
Minimum Time Until Verification	1,599
Maximum Time Until Verification	18,160
Average Time Until Verification	7,556

Figure 11. Minimum, Maximum, and Average Time Until Verification for Domains Verified More Than 24 Hours after Submission

The reason behind the much longer verification times for these few domains is unclear; one possibility is that the malicious actors responsible for the domain for some reason (possibly because they noticed the domain was starting to get blocked, perhaps due to being added to some other blocklist) took the domain offline shortly after the period of initial use in which it was first reported as suspicious, which then delayed verification until some later date when the domain's owner brought it back online and PhishTank users could finally verify it. Regardless of the

reason, since these instances fall well outside the normal range of values for time until verification, the author has excluded these values from the conclusions; they are included here only for full transparency.

3.6. Discussion

Based on the Domain Name Industry Brief's reported 356.6 million new domains registered in Q2 2023 (DNIB, 2023), one can conclude that nearly four million domains are registered daily. While the DShield data revealed many new domains, more than half a million per day on average, this highlights a possible limitation of the system and shows that while it is relatively comprehensive, it still only observes a partial sample of the overall traffic on the Internet. Some of this also can be attributed to the DShield only monitoring logs of actual online activity, and some newly registered domains are likely never brought online and are parked indefinitely.

URL and domain-based blocklists are constantly changing and require continuous pruning and updating. Advanced adversaries will use automated domain name generation schemes for domain registration and other techniques such as DNS fast flux ("What is DNS Fast Flux?", 2023) to bypass the effectiveness of blocklists, which struggle to keep up. Some URLs are also re-added multiple times to a blocklist after having previously been removed; in their analysis of the PhishTank blocklist, Bell and Komisarczuk (2023) found that 3% of all URLs were added a second time. This research found that within ten days (October 15 to October 25), the PhishTank blocklist removed 6,454 URLs and added 5,038 new URLs, a significant turnover.

As for the earlier question about the number of new domains confirmed as malicious, the research showed that only some are officially recognized as phishing sites. Of all the new domains observed, only a tiny fraction, less than one-hundredth of a percent, were eventually added to the blocklist. Therefore, one might be led to conclude that very few domains are dangerous; however, the more likely inference is that very few malicious domains are correctly identified and reported.

Regarding how long it takes on average before suspicious domains are verified, despite requiring participation by human volunteers, PhishTank's vote-based verification system proved

to be surprisingly effective, with the average time until verification taking roughly an hour. Machine-based validation methods would likely prove even quicker, making this subject a suitable candidate for future research.

The research revealed a substantial number of anomalous results, where the PhishTank data showed that domains that supposedly had been first seen in September had, in fact, already been submitted for verification weeks or months earlier. To investigate further, domain registration information was checked for some of these domains to determine their initial registration dates. For example, the domain smartwaypayment.com was first seen by the DShield on September 5, 2023, but was submitted to PhishTank months earlier, on May 14, 2023. A Whois record lookup of the domain revealed that the Creation Date for this domain was September 3, 2019, meaning the domain had existed for several years before being first seen. This brief examination highlights that there is no definitive source for, or straightforward way to pinpoint, the exact start time for when malicious actors begin using a domain operationally; using the first seen date as reported in the DShield is, at most, a best guess estimate.

J. Zhang et al. (2008) define the required components of a high-quality blocklist as having a high hit rate, timely inclusion of addresses, and proactive inclusion of addresses not previously encountered by members of the blocklist's consumer network. On the first point, the PhishTank blocklist would score highly, as it contains tens of thousands of verified true positive phishing URLs, but on the second point of timely inclusion, it is likely lacking. The research results showed that it takes, on average, 13.7 days, or approximately two weeks, between when a malicious domain is first seen on the Internet and when it is first submitted for phishing verification. As for the third point of proactive inclusion, PhishTank likely fails completely, as it relies entirely on user submissions of URLs after direct observation.

Researchers have proposed several ideas for automating blocklist generation to overcome the deficiency in proactive inclusion facing most blocklists that rely on manual submission. J. Zhang et al. (2008) proposed a "highly predictive" blocklisting system that combines a relevance-based ranking of attack sources with a severity ranking to create individualized blocklists for organizations. Soldo et al. (2010) developed an alternative method of predictive blocklisting by formulating the problem as an implicit recommendation system rather than a

link-analysis problem. CANTINA is an interesting heuristics-based phishing detection approach based on analyzing the actual content of websites to determine their legitimacy (Y. Zhang et al., 2007).

Finally, it should also be noted that, despite the frequent use of the term throughout this paper, there is, in fact, no such thing as a "malicious domain"; a domain is simply a tool of the Internet, and it is the intentions of those who are registering the domain that should be viewed as malicious. In other words, how a domain is used after registering it truly determines whether a registrant's intentions were malevolent.

4. Conclusion

Regarding the percentage of new domains that are labeled as "malicious," the results of this project showed that, overall, despite hundreds of thousands of new domains appearing daily, very few of these newly observed domains ended up being recognized as verified phishing sites. Of all the new domains monitored in this experiment, less than one-hundredth of a percent were eventually seen to be added to the examined phishing blocklist. As for how long it typically takes before suspected domains are verified, the existing blocklist verification method proved relatively efficient, with verification typically taking just over an hour.

The answer to the question of how long it typically takes for phishing domains to be flagged and reported as dangerous, on the other hand, could have been more satisfying. The research showed that, on average, about two weeks pass between when a malicious domain is first seen and when it is first submitted for suspicious activity. This gap in time means that those who rely on blocklists as a primary defensive tool remain woefully unprotected during the first few weeks after a new malicious domain becomes active. With millions of new domains registered daily, the danger of hidden malicious phishing sites becomes an ever-present threat. Cyber defenders must identify dangerous URLs and domains as quickly as possible and ensure that channels for promptly reporting them exist. Participation in distributed monitoring networks like the DShield would be an excellent first step for any organization.

References

Aaron, G. & Piscitello, D. (2017, November 1). *Reputation block lists: Protecting users everywhere*. Internet Corporation for Assigned Names and Numbering. https://www.icann.org/fr/blogs/details/reputation-block-lists-protecting-users-everywhere-1-11-2017-en

Bell, S., & Komisarczuk, P. (2020, February). An analysis of phishing blacklists: Google Safe Browsing, OpenPhish, and PhishTank. *ACSW '20: Proceedings of the Australasian Computer Science Week Multiconference.*

Domain Name Industry Brief. (2023). *The domain name industry brief quarterly report: Q2 2023 data and analysis*. https://dnib.com/articles/the-domain-name-industry-brief-q2-2023

Google. (n.d.). *Safe Browsing*. https://safebrowsing.google.com/

Google. (n.d.). *What is Safe Browsing?* https://developers.google.com/safe-browsing/

Internet Assigned Numbers Authority (2023, November 15). *TLDs alpha by domain*. Retrieved November 15, 2023, from https://data.iana.org/TLD/tlds-alpha-by-domain.txt

Internet Corporation for Assigned Names and Numbering. (2012, February 25). *Welcome to ICANN!* https://www.icann.org/resources/pages/welcome-2012-02-25-en

Internet Corporation for Assigned Names and Numbering. (2017, June 20). *Registering domain names*. https://www.icann.org/resources/pages/register-domain-name-2017-06-20-en

Lenaerts-Bergmans, B. (2023, May 4). *Vishing: Definition and prevention*. CrowdStrike. https://www.crowdstrike.com/cybersecurity-101/vishing/

Loshin, P. (2021, November). *What is ICANN?* TechTarget. https://www.techtarget.com/whatis/definition/ICANN-Internet-Corporation-for-Assigned-*Names-and-Numbers*

MITRE ATT&CK. (2023). https://attack.mitre.org/

MITRE ATT&CK. (2023, March 30). *Acquire infrastructure: Domains*. https://attack.mitre.org/techniques/T1583/001/

MITRE ATT&CK. (2023, September 6). *Phishing: Spearphishing link*. https://attack.mitre.org/techniques/T1566/002/

MITRE ATT&CK. (2023, September 8). *Phishing.* https://attack.mitre.org/techniques/T1566/

New domain (gTLD) launch schedule. (2023). Webnames. Retrieved November 15, 2023, from https://www.webnames.ca/domain-registration/new-domain-name-launch-schedule.aspx

Number Resource Organization. (2023). *The Internet registry system*. https://www.nro.net/about/rirs/the-internet-registry-system/

OpenPhish. (n.d.). *Knowledge base*. https://openphish.com/kb.html

OpenPhish. (n.d.). *OpenPhish database*. https://openphish.com/phishing_database.html

PhishTank. (n.d.). *Developer information*. https://phishtank.org/developer_info.php

PhishTank. (n.d.). *Friends of PhishTank*. https://www.phishtank.com/friends.php

PhishTank. (n.d.). *What is phishing?* https://www.phishtank.com/what_is_phishing.php

Proofpoint. (2023). *2023 state of the phish: An in-depth exploration of user awareness, vulnerability and resilience.* https://www.proofpoint.com/sites/default/files/threat-reports/pfpt-us-tr-state-of-the-phish-2023.pdf

Ramsey, S. (2021, October 6). *Can we move past blocklists to automated takedowns?* SANS Information Security Reading Room. https://www.sans.org/white-papers/can-we-move-past-blocklists-to-automated-takedowns/

SANS Internet Storm Center (2023). *DShield API*. https://isc.sans.edu/api/

SANS Internet Storm Center (2023). *ISC history and overview*. https://www.dshield.org/about.html

Soldo, F., Le, A., & Markopoulou, A. (2010, March). Predictive blacklisting as an implicit recommendation system. *INFOCOM '10: Proceedings of the 29th Conference on Information Communications.*

Shawn Reinhart, shawn.reinhart@gmail.com

Trend Micro. (n.d.). *What is smishing?* https://www.trendmicro.com/en_vn/what-is/phishing/smishing.html

What is DNS fast flux? (2023). Efficient IP. https://efficientip.com/glossary/what-is-dns-fast-flux/

Yeh, C. (2023, May 3). *8 new top-level domains for dads, grads and techies*. Google Registry. https://blog.google/products/registry/8-new-top-level-domains-for-dads-grads-tech/

Zhang Y., Hong J., & Cranor L. (2007, May). CANTINA: A content-based approach to detecting phishing web sites. *WWW '07: Proceedings of the International Conference on World Wide Web.*

Zhang, J., Porras, P. A., & Ullrich, J. (2008, July). Highly predictive blacklisting. *SS '08: Proceedings of the 17th Conference on Security Symposium.*