



Phishing in the Free Waters: A Study of Phishing Attacks Created using Free Website Building Services

Sayak Saha Roy
University of Texas at Arlington
Arlington, United States
sayak.saharoy@mavs.uta.edu

Unique Karanjit
University of Texas at Arlington
Arlington, United States
unique.karanjit@mavs.uta.edu

Shirin Nilizadeh
University of Texas at Arlington
Arlington, United States
shirin.nilizadeh@uta.edu

ABSTRACT

Free Website Building services (FWBs) provide individuals with a cost-effective and convenient way to create a website without requiring advanced technical knowledge or coding skills. However, malicious actors often abuse these services to host phishing websites. In this work, we propose FreePhish, a scalable framework to continuously identify phishing websites that are created using FWBs. Using FreePhish, we were able to detect and characterize more than 31.4K phishing URLs that were created using 17 unique free website builder services and shared on Twitter and Facebook over a period of six months. We find that FWBs provide attackers with several features that make it easier to create and maintain phishing websites at scale while simultaneously evading anti-phishing countermeasures. Our study indicates that anti-phishing blocklists and browser protection tools have significantly lower coverage and high detection time against FWB phishing attacks when compared to regular (self-hosted) phishing websites. While our prompt disclosure of these attacks helped some FWBs to remove these attacks, we found several others who were slow at removal or did not remove them outright, with the same also being true for Twitter and Facebook. Finally, we also provide FreePhish as a free Chromium web extension that can be utilized to prevent end-users from accessing potential FWB-based phishing attacks.

CCS CONCEPTS

• Security and privacy → Social engineering attacks; Phishing;

KEYWORDS

Phishing, Cybersecurity, Social media, Abuse, Evasion, Website Builder, Online Safety, Blocklist, Measurements, Framework

ACM Reference Format:

Sayak Saha Roy, Unique Karanjit, and Shirin Nilizadeh. 2023. Phishing in the Free Waters: A Study of Phishing Attacks Created using Free Website Building Services. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC '23)*, October 24–26, 2023, Montreal, QC, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3618257.3624812>



This work is licensed under a Creative Commons Attribution International 4.0 License.

IMC '23, October 24–26, 2023, Montreal, QC, Canada
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0382-9/23/10.
<https://doi.org/10.1145/3618257.3624812>

1 INTRODUCTION

Phishing attacks are responsible for countless data breach incidents and credential fraud, leading to millions of dollars in financial damage [9, 60, 93]. The success of phishing scams is largely dependent on effective social engineering [36], which necessitates the creation of websites that convincingly mimic their legitimate counterparts. This task demands a good understanding of web design, attention to detail, and proficiency in mimicking trusted entities. While phishing kits [86] have emerged to streamline this process, the more effective ones often come with a hefty price tag. Furthermore, the security landscape has evolved, with recent research focused on detecting these phishing kits and the websites they generate [30, 35, 46, 66], rendering them less of a fail-safe approach for generating phishing scams. In addition, the proficiency of anti-phishing entities in detecting regular phishing attacks has notably increased [65, 67]. Therefore, attackers continuously innovate and unearth newer exploits [64, 88, 105] to evade detection. For example, they often invest significantly in procuring domains to host their scams, which is especially burdensome considering the frequent takedowns [67] enacted by anti-phishing entities, leading to a relentless cycle of buying new domains to keep their operations active [8]. Moreover, to enhance the perceived legitimacy of the phishing sites, premium top-level domains (TLDs), such as .com [19] are often used.

Our research closely studies a trend for generating phishing attacks using 17 unique **Free Website Building services (FWBs)**. These services provide a platform that allows hosting websites at no cost and utilize user-friendly drag-and-drop builder interfaces which facilitates creating phishing attacks. These phishing websites are indistinguishable from those generated through manual coding or phishing kits. Also, these FWB services provide several features that allow the creation of phishing scams that can evade anti-phishing detection mechanisms and resist takedowns for extended periods. In summary, using FWBs alleviates many of the challenges in maintaining and scaling phishing operations.

In this work, we propose *FreePhish*, a framework for automatically identifying phishing attacks created using FWBs and shared on social media services, monitoring the reaction of anti-phishing entities against them, and reporting them for removal. Using FreePhish, we characterize the features and effectiveness of over 31.4K such phishing URLs that are shared on Twitter and Facebook over a period of six months.

Despite the widespread use of FWBs to host phishing scams, so far, there has been a limited understanding of these attacks due to only sporadic (and short) reports [1, 2, 15] dedicated to this issue. In response to this knowledge gap, our work pioneers a large-scale, longitudinal study to meticulously characterize and detect these attacks.

The paper is structured as follows: We first establish the need to study these attacks by identifying their pervasiveness on both Twitter and Facebook over the period of two years in Section 2, followed by qualitatively studying these websites in Section 3 to find unique structural characteristics which distinguish them from regular phishing websites (i.e., those which are self-hosted), and make them more evasive to anti-phishing detection. We utilize these findings to develop FreePhish, in Section 4, our automated framework that continuously identifies FWB phishing threats in real-time, reports them to the respective FWBs and social media platforms for removal, and measures the effectiveness of anti-scam entities against them. FreePhish helped us identify more than 31.4K FWB phishing URLs from November 2022 to March 2023, and we dedicate Section 5 to closely investigate how well they are detected by four popular phishing blocklists: PhishTank [5], OpenPhish [69], Google Safe Browsing [4], and APWG eCrimeX [27], 76 security tools, as well as the FWB service providers and the social media networks (Twitter and Facebook). Our comprehensive measurement highlights several gaps in the prevalent anti-phishing ecosystem against FWB-based phishing attacks. Finally, in Section 5.5, we study how attackers further improve their efforts to host more evasive attacks on certain FWB domains, such as using i-frames, multi-step phishing attacks, and malicious drive-by downloads. Moreover, we introduce FreePhish as a Chromium web extension that can proactively prevent users from accessing these attacks. The primary contributions of our work are:

- (1) We have performed the first extensive analysis of phishing websites created using 17 free website creation services (FWBs), revealing features that enable evasion of common anti-phishing methods and website takedown resistance.
- (2) We created FreePhish, a machine learning (ML) framework capable of detecting FWB phishing websites. It auto-identifies, reports, and evaluates these websites, outperforming other ML anti-phishing measures with a 97% accuracy.
- (3) Our framework detected 31.4K new FWB phishing attacks within six months. We assessed the effectiveness of popular phishing blocklists, browser protection tools, social media platforms and the FWBs themselves against these attacks. Our findings show low coverage and slow response time across all these attacks, highlighting significant gaps against these attacks in the present anti-phishing ecosystem.
- (4) We provide FreePhish as a free web extension: <https://github.com/UTA-SPRLab/freephish/> that can be used in any Chromium based web-browser to proactively prevent users from accessing FWB phishing attacks.

2 HISTORICAL Pervasiveness of FWB PHISHING ATTACKS

To underscore the necessity of studying phishing websites created using free website builder services, we documented a rising trend in these attacks over a two-year period. Using the official Twitter API [95] and Meta’s CrowdTangle [13], we compiled 3.1 million URLs from Twitter and 1.4 million URLs from Facebook that contained distinct second-level domains in their URL strings over the period of January 2020 to August 2022. This decision to only collect




Figure 1: Distribution of FWB phishing attacks shared on Twitter and Facebook from Jan. 2020 to Aug. 2022.

URLs with a second-level domain was driven by the intention to analyze websites that are created under another domain. For instance, in the URL `mywebsite.000webhost.com`, `000webhost` serves as the second-level domain. The URLs were then scanned using VirusTotal [26], a service that collates detection scores from 80 different anti-phishing tools. URLs with two or more detections were labeled as phishing, conforming to standards set in previous literature for such identification [71, 74, 87]. In addition to recognizing potential FWB phishing attacks over time, this analysis also helped us identify the FWBs that are being abused by attackers. We labeled 34.7K unique URLs as phishing, out of which 25.2K URLs (16.3k URLs originating from Twitter and 8.9K URLs from Facebook) utilized 17 unique free website-building services. The remaining URLs belonged to Dynamic DNS services (such as DuckDNS, Netlify etc.) which is outside the purview of our study. We refer to these 25.2k URLs as the *initial dataset* or D1 throughout the rest of the paper. Figure 1 displays the distribution of these attacks over time across both social media platforms. Not only does this data highlight a marked escalation in the quarterly frequency of such attacks, but it also reveals a strategic shift in attackers’ preference toward adopting newer hosting services. For each analyzed month, we identify the unique domains that accounted for 80% of the FWB phishing attacks during that period. It is important to acknowledge that the actual scale of FWB phishing attacks is likely larger than our dataset suggests, given the proactive measures both platforms employ to flag or remove malicious content [16, 17]. Our conclusions are further substantiated by Interisle’s in-depth studies conducted between 2020 and 2023 [49–52]. We draw particular attention to their analyses concerning the abuse of Subdomain Providers, which include not just FWB services but also Dynamic DNS vendors like DuckDNS and Netlify although the latter are outside the purview of our research. Interisle’s longitudinal reports corroborate two primary trends: firstly, there is an ascending trajectory in the utilization of FWB services for orchestrating phishing attacks; secondly, attackers are progressively exploiting newer FWB platforms, thereby mirroring the trends observed in our own study.

It is crucial to highlight that, due to the age of these URLs (some of these URLs are over two years old), we have refrained from

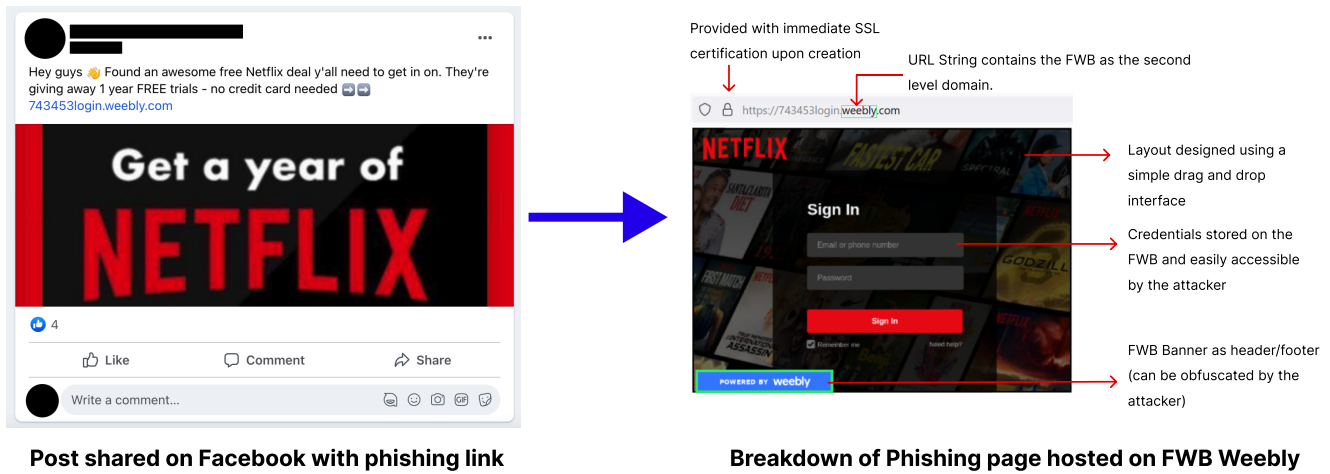


Figure 2: Example of a phishing website created on Weebly, a FWB service, and shared on Facebook. Note - This website has already been taken down, and thus the full URL string is not obfuscated from the image.

utilizing them for any further longitudinal studies, for example, understanding the effectiveness of anti-phishing measures against FWB attacks. Instead, we utilize the zero-day FWB phishing attacks identified by our FreePhish framework in Section 5. We also reported these 25.2k URLs to Twitter, Facebook, and the FWBs directly, thereby facilitating their removal from online platforms.

3 CHARACTERIZING FWB PHISHING ATTACKS

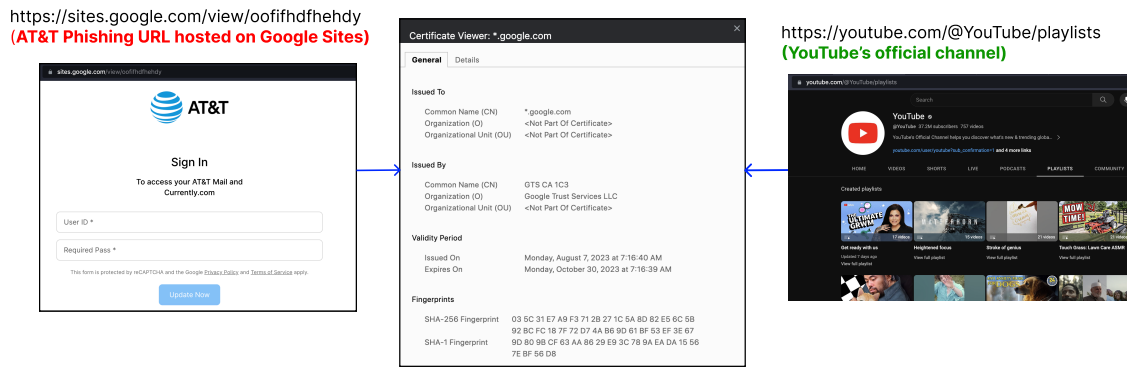
Having established the prevalence of phishing attacks originating from Free Website Builders (FWBs), our next objective was to understand the unique characteristics of these attacks. To do so, we carefully examined a random sample of 5K URLs from Dataset D1 to highlight the specific features that make FWB-created phishing attacks particularly advantageous to the attacker. Figure 2 illustrates the anatomy of a phishing attack created on Weebly, a popular FWB and imitating the HTML login page. Two university students, both specializing in Computer Security with a solid understanding of social engineering scams, were chosen as coders for this qualitative evaluation. The coders primarily assessed whether the URL and website appearance attempted to spoof one of 409 brands targeted by phishing attacks, as reported by OpenPhish in August 2022 [70]. In addition, the coders evaluated whether the websites contained text fields designed for gathering sensitive information (such as email addresses, passwords, SSN etc). For websites that did not meet these criteria but initiated a file download, the coders conducted additional scans using VirusTotal and marked the files as malicious if they generated four or more detections. Prior literature has established this detection threshold as being sufficient for identifying malware samples [71, 87].

Of the 5K randomly sampled URLs, 4,656 were confirmed to be phishing. The initial Cohen’s Kappa inter-rater agreement was 0.78 (i.e., high agreement). Disagreements primarily occurred due to misinterpretation of brand spoofing (i.e., differing views on how

effectively a website mimics a legitimate brand), evasive attacks (Coder #1 failed to recognize two-step phishing attacks as harmful), assessment of text fields (Coder #1 overlooked the relevance of address fields and phone numbers in indicating phishing intent) and language representation (Coder #2 was unable to identify phishing intentions in eight websites that primarily used Spanish and Chinese languages). All disagreements were resolved through discussion and consensus. We further analyzed a sample of the confirmed true positive set and evaluated the design interfaces of the 17 unique FWB services to identify several factors contributing to the convenience and efficiency of launching phishing attacks using FWB services. We will discuss these findings in detail in the following paragraphs.

Initial Investment: Conventionally, phishing attacks have a low life span, getting detected by anti-phishing providers and subsequently removed by the hosting domain within a few hours [67]. In response to the countermeasures, attackers frequently redeploy phishing attacks across multiple domains, which need ongoing investments [65, 67]. However, free website builders have turned this situation in the attacker’s favor, as they allow users to create and host websites on their servers for free, thus allowing a straightforward way to scale attacks. Moreover, if they get blocklisted on one FWB, they can switch to another without any added expenditure.

Easy of creation: Phishing attacks are designed to imitate legitimate websites closely, deceiving users into sharing their sensitive information. Designing the webpage and its related components requires a level of technical expertise. While dedicated phishing kits exist to aid attackers in designing webpages, a recent focus on detecting such kits has lowered the potency of such attacks. On the other hand, FWBs provide easy-to-use drag-and-drop website creation interfaces. These builders do not require any coding experience and can be used to create and publish phishing websites in a matter of minutes. They also handle data entered in the credential fields, provide additional applications like SEO boosting and social



Both Websites share the same SSL certificate with identical Common Name, Organization, Issue/Expiry date and SHA fingerprints.

Figure 3: Example shows FWB websites created using Google Sites share the same certification as YouTube. Note - The phishing website has already been taken down, and thus the full URL string is not obfuscated from the image.

media integration, and allow the embedding of custom HTML code, which can be exploited by attackers.

Immediate SSL Certification: Users are often advised to check if a website has SSL certification to avoid phishing scams [62]. However, SSL based phishing attacks are widespread [38], with more than 49% of phishing URLs now having SSL certification [54]. The registration of a new phishing website typically involves a certification payment and a verification process that can up to several days [11]. More affordable or even free options like Let’s Encrypt [39] and ZeroSSL [103] offer Domain Validation (DV) certificates with a 90-day validity period. While these DV certificates are easier to obtain, they demand additional implementation effort and are generally considered less trustworthy than Extended Validation (EV) or Organization Validation (OV) certificates [33, 91]. The latter types require website owners to undergo a more rigorous verification process, which in turn enhances their credibility. A study by PhishLabs [76] reveals that most SSL-based phishing sites use DV certificates. They also point out an increasing trend of attackers adopting EV certificates to make their fraudulent sites appear more legitimate. In contrast, FWB Services automatically provides all new websites with either an EV or OV certificate, eliminating the need for additional verification steps and easing the implementation process. Figure 3 shows one example where a phishing attack created on FWB service Google Sites has the same SSL certificate as YouTube.com, an official Google platform. Moreover, considering that ML-based phishing detection models often positively co-relate SSL certification to the website being benign [21, 82], having this feature in FWB phishing attacks makes it easier to evade detection.

Premium TLDs: Domain providers price their domains based on the popularity of the top-level domains (TLDs). For example, the .com, .org TLDs are more expensive than the TLDs like .xyz, .live. To save on investment, attackers usually launch phishing websites on cheaper TLDs [89], which have, in turn, tuned blocklists, tools, and automated machine learning tools to use TLDs as a heuristic for detection [77, 92, 102]. Recent studies have found that users have higher trust in .com TLDs [43, 72]. All 14 out of 17 FWBs provide users with a .com TLD for their websites, thus allowing the

attackers to get a premium TLD at no cost. Out of the 4,656 domains that our coders identified as phishing sites, a staggering 4,139 – or approximately 89% were hosted on the 14 FWB services that offer a .com Top-Level Domain (TLD). Thus, along with not appearing in certification logs, this feature can further aid in evading user suspicion and anti-phishing detection.

Longer Domain Age: Phishing websites usually remain active for a very short period of time since they are either blocklisted/taken down or abandoned by the attacker. Thus several anti-phishing tools consider the *domain age* to be an important heuristic feature [66, 73, 85]. However, phishing attacks created using FWBs are essentially sub-domains of the FWB itself, and thus have the same domain age as the FWB, as recorded by registrar information databases such as WHOIS [75]. Since the FWBs (e.g., Weebly, Google Sites, Squarespace) are several years old, detection models which utilize domain age as an important heuristic, might associate FWB-hosted phishing attacks to be legitimate. The median domain age of the 4,656 URLs in our dataset labeled as phishing was 13.7 years, according to their WHOIS domain record. On the other hand, the same number of self-hosted phishing URLs collected from Phish-Tank [5], an open-source anti-phishing blocklist, had a median age of only 71 days.

High Code Similarity Between Legitimate and Phishing Websites: Automated phishing detection measures often assume that the code structure of phishing websites differs significantly from legitimate websites [24, 68]. Moreover, some detection methods tend to compare the source code of new websites with that of known phishing websites to identify potential threats [81]. However, this approach could become less effective in cases where free website builders (FWBs) are used. These platforms allow users to create websites using predefined templates, leading to considerable similarities between the HTML code of different websites hosted on the same service, including both legitimate and phishing sites. In our study, we used the Levenshtein edit distance algorithm, a popular metric for computing similarity between text sequences [25, 78], to measure the extent of similarity in HTML code blocks between benign and phishing websites hosted on these FWBs. The detailed

FWB	# URLs	Median similarity
Weebly	6,309	79.4%
000webhostapp	3,714	68.1%
Blogspot	2,492	63.8%
Google Sites	1,547	72.4%
Wix	1,801	63.7%
Github.io	582	37.4%

Table 1: Website code similarity between FWB phishing and benign websites.

methodology we used for this computation is explained in Appendix A. Table 1 illustrates the median code similarity between phishing and benign websites created using some popular FWBs, showing a high median code similarity for several FWBs. For example, websites created on Weebly have a median similarity of 79.4%. This indicates that attackers can potentially use predefined templates provided by FWBs to make their phishing websites appear more like legitimate websites, thus increasing the chances of evading detection.

Increased Difficulty of Discovery: Phishing websites created through FWBs are more challenging to discover for multiple reasons. Firstly, once domains obtain their SSL certification, they typically appear on the Certificate Transparency (CT) Log network [94]. As many phishing domains now have SSL certification [54], various anti-phishing methods utilize this log to identify new attacks [37, 40, 84]. However, we found that all websites created using all 17 FWB services inherited the same SSL certification used by the service itself. Notably, three of these services – Weebly [100], Wix [101], and Mailchimp [61] highlight the fact that websites created for free through their platforms are automatically provided with complimentary SSL certification. Having the same SSL certification as their hosting domain renders FWB phishing attacks effectively invisible in Certificate Transparency (CT) logs as they do not appear as individually registered domains. Notably, the standard practice of certificate re-issuance for compromised or vulnerable domains [104] does not apply to FWB-hosted websites, since they naturally inherit the SSL certificates from their primary FWB domain. Secondly, numerous anti-phishing crawlers examine websites indexed by search engines to discover new phishing attacks [44, 48, 80]. However, only a small fraction (4.1%) of the 25.2K FWB phishing URLs in our historical dataset were indexed by Google. Google does not index websites without any incoming links (a common occurrence for phishing attacks hosted under subdomains) or if the website source code contains a `< noindex >` meta tag, instructing search engines not to index the website [12]. Among the FWB-hosted phishing URLs in our dataset, 44.7% contained the `< noindex >` meta-tag, further contributing to the difficulty of their discovery.

4 FREEPHISH FRAMEWORK

In this section, we detail the development of *FreePhish*, a real-time framework designed to detect and combat FWB-based phishing attacks shared on Twitter and Facebook. FreePhish is not only capable of detecting and reporting these attacks to hosting providers

in real-time but also longitudinally monitors the responses of anti-phishing entities against them.

Our framework, illustrated in Figure 4, comprises five core components: (1) *Streaming module*: Operates at 10-minute intervals, monitoring for FWB URLs (whether phishing or non-phishing) shared on Twitter and Facebook. The module also captures the full snapshot of each identified website. (2) *Preprocessing module*: Saves a full snapshot of the website (including screenshots and source code) and extracts several URL and HTML-based features derived from our qualitative analysis in Section 3. (3) *Classification module*: Employs a Random Forest classifier to analyze the extracted features, predicting whether the website is engaging in a phishing attack. (4) *Reporting module*: Reports each URL flagged as a phishing site to the respective FWB service or hosting provider, and subsequently monitors their response. (5) *Analysis module*: Evaluates the effectiveness of anti-phishing entities, monitoring each URL at regular intervals. We discuss each of these components in more detail.

4.1 Streaming and Pre-processing Module

The streaming module utilizes the Twitter and CrowdTangle APIs to collect new posts from Twitter and Facebook every 10 mins. It utilizes a regular expression to extract the URL from the post, which is followed by the pre-processing module. It stores a full snapshot of the website and extracts several URL-based, HTML-based, and FWB features, such as obfuscated banner code, credential fields, presence of `<noindex>` tags, etc. This feature set is passed to the Classification module.

4.2 Classification Module

Like most evasive phishing websites, FWB attacks share significant similarities with self-hosted phishing attacks but introduce subtle modifications to the website structure by using several features of FWBs, allowing them to bypass existing detection mechanisms. The motivation behind crafting this classification module is twofold. Firstly, we want to achieve a high success rate in detecting these attacks. Secondly, we want to ensure time efficiency. Since a large volume of websites is shared on social media throughout the day, a slower classification model can exponentially hamper the framework’s overall performance. To accomplish these objectives, we first build the ground-truth dataset. We test four state-of-the-art ML-based phishing detection models to identify the candidate which provides the optimal trade-off between performance and runtime efficiency. Finally, we augment the model by adding and removing FWB phishing-specific features (Section 4.2), then train and test the augmented model on our dataset and gauge its performance.

Ground truth collection: In order to create our ground-truth dataset, we use the manually verified 4,656 true positive URLs in the initial dataset (D1). We also selected and manually verified an equal number of true negative benign FWB-created URLs from D1 that were shared on Twitter ($n=3,299$) and Facebook ($n=1,357$).

Choosing the optimum model: We compared the performance of four state-of-the-art ML-based phishing detection models. These models included two that rely on the visual features of the website: VisualPhishNet [18] and PhishIntention [59], one that relies on both the URL string and HTML representation of the website:

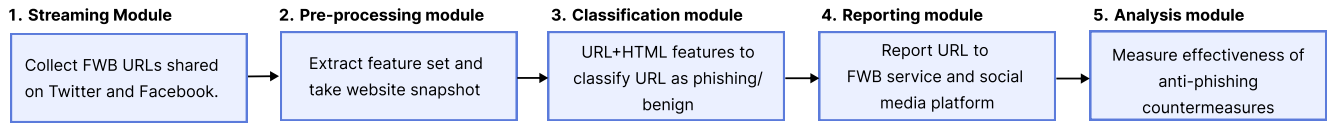


Figure 4: High-level overview of the FreePhish framework

Model	Accuracy	Precision	Recall	F1-score	Total Time Taken (sec)	Median Runtime (sec)
VisualPhishNet	0.76	0.78	0.72	0.75	14,802	5.1
PhishIntention	0.96	0.98	0.94	0.96	32,958	11.3
URLNet	0.68	0.70	0.67	0.68	4,471	1.9
Base StackModel	0.88	0.89	0.87	0.88	8099.4	2.4
Our Model	0.97	0.96	0.97	0.96	8,658	2.8

Table 2: Comparison of various phishing detection models.

StackModel [57], and one that relies on the semantic representation of the URL string only: URLNet [55]. All models were tested across our entire testing set. The models were tested on a system running on an Intel Xeon W Processor with 64GB of RAM and 2x NVIDIA RTX 2080Ti GPU. Table 2 illustrates that URLNet has the lowest median runtime of 1.9 secs but also the lowest recall of 0.68. VisualPhishNet has a median runtime and a recall of 5.1 secs and 0.75, while StackModel has that of only 2.4 secs and 0.87, respectively. PhishIntention performed the best with a recall of 0.97, probably because, unlike the other models, it does not rely solely on visual features or URLs, but incorporates both static and dynamic analyses of the website’s entire workflow [59]. However, its complexity makes it far slower at classifying our samples, with a median runtime of 11.3 secs per URL, which can lead to significant overhead in our framework. We thus choose the StackModel since it strikes the best balance between efficiency and runtime, and we improve upon it for detecting FWB phishing attacks.

Feature extraction: Our model builds on the feature set utilized by the StackModel, which includes 8 URL-based features (such as the usage of suspicious symbols, sensitive vocabulary in the URL string, similar brand names, the length of the URL, etc.) and 12 HTML-based features (such as the number of internal and external links, the presence of empty links, the inclusion of a login form, the length of HTML content). The full list of features can be found in the original publication by Li et al. [57]. Two of the features StackModel uses to detect phishing websites, including the presence of ‘https’ and multiple top-level domains, do not apply to those created by FWBs, because all FWB websites utilize ‘https’ and FWB phishing attacks typically contain only one TLD. Therefore, we did not use them in our model.

In their place, we introduced two new features specifically tailored for FWB phishing detection, drawing on our analysis as discussed in Section 3 - *Obfuscating FWB Footer*: This feature recognizes the trend among FWB services to include a header or footer banner on their free websites. To make their fraudulent sites appear more legitimate, phishers often alter the source code of these websites to hide these banners. They might, for example, add a `< visibility : hidden >` parameter to the `< div >` tag containing

the banner code. *Preventing Indexing*: As detailed in Section 3, attackers frequently employ *noindex* meta-tags to prevent search engines from indexing their URLs, which indirectly reduces traffic from anti-phishing crawlers. We included a feature that checks for the presence of these tags in the website’s source code.

Model training and performance: We employ the methodology utilized by Li et al. [57] for training the StackModel architecture. To summarize, their approach utilizes a two-layer stacking model that combines multiple ML models. The models employed in each layer include Gradient Boosting Decision Tree (GBDT), XGBoost, and LightGBM. A strategy similar to K-fold cross-validation is employed for training, where for each iteration, the dataset (n=4,656 phishing samples, 4,656 benign samples) is partitioned into 70% for the training set and 30% for the testing set. This process continues until each basic model has predicted all samples. The initial layer utilizes both original features and prediction results derived from majority voting as inputs for the second layer. The second layer combines the input features with its output results, and these composite features are then used to train a final GBDT model. This model then predicts which webpages are phishing sites. The proposed model utilizes both URL and HTML-based features for prediction highlighted in Section 4.2. As illustrated in the last row of Table 2, our augmented StackModel has an F1-score of 0.96 and a median run-time of only 2.8 secs.

4.3 Reporting Module

URLs identified by the classification model (as phishing) are reported immediately to the respective FWB service and social media platforms (i.e. Twitter and Facebook). This ensures the timely and ethical disclosure of identified threats from our platform to the sources sharing them online. We used Python Selenium to automate the submission of reports through the Chromium Web Browser, as neither of the platforms have the feature to report malicious URLs through API. According to prior literature, including evidence-based information like website screenshots expedites the malicious URLs evaluation process [65]. As a result, our FreePhish reported URLs to include the full URL, a screenshot of the site, and the targeted organization’s name. For safety reasons, our experiment underwent a thorough review and was approved by our

organization’s Office of Information Security before we submitted the phishing reports. We explore our findings regarding the effectiveness of FWB and social media platforms in removing these threats in Section 5. Contrarily, we refrain from reporting to blocklists. Community-based blocklists, such as APWG eCrimeX and PhishTank, list reported URLs in their feed immediately without any verification. This could undermine our longitudinal measurement’s validity. This also applies to several tools we studied for browser protection, which rely on community threat intelligence.

4.4 Analysis Module

We conduct a longitudinal evaluation to assess the effectiveness of various anti-phishing entities against the discovered FWB phishing attacks. Two key performance indicators were considered: *coverage* - the proportion of FWB-hosted phishing URLs detected or removed within one week of appearing in our dataset, and *response time* - the duration from a URL’s first appearance in our dataset to its detection or removal by the entity. The proceeding paragraphs describe the analyses:

Anti-phishing blocklists: For each FWB phishing URL, we checked if it was present on four anti-phishing blocklists: Google Safe Browsing (GSB), PhishTank, OpenPhish, and APWG eCrimeX, using their respective APIs, at regular intervals of 10 minutes till the URL became inactive. GSB is regularly used on several browsers running on the Chromium engine, such as Google Chrome [42], Firefox [41], Vivaldi [98], etc. PhishTank and OpenPhish regularly contribute their data to several anti-phishing tools and web browsers [3]. APWG’s eCrimeX also shares its data with several organizational anti-phishing entities [65].

Browser protection tool analysis: We used VirusTotal [7] to scan both FWB-based and self-hosted phishing URLs every 10 minutes, from when the URL first appeared in our dataset, up till when it became inactive, or up to a week, whichever was sooner. This is to mitigate the possibility of outdated scores [74]. Unlike our previous analysis, where we evaluated the performance of each blocklist separately, we aggregated the VirusTotal scores, which provided the detection rate for up to 76 different anti-phishing engines and blocklists. VirusTotal data also includes GSB, PhishTank, and OpenPhish, and we excluded the scores of these blocklists to avoid redundancy with our previous analysis.

FWB response analysis: The most efficient way to prevent users from falling victim to phishing attacks is if the website gets removed by the hosting providers. To evaluate domain takedown of FWB-based phishing URLs, we checked if the website was active every 10 mins from when it first appeared in our dataset.

Social media platform response analysis: Figure 1 illustrates how FWB-based phishing attacks are becoming a persistently growing threat on Twitter and Facebook, with attackers exploiting more FWBs to host their attacks. After collecting the initial tweet (containing the phishing link), we used the Twitter Academic API [6] to check if the tweet had been deleted at regular intervals of 10 minutes. To check the same for Facebook posts, we check the URL of each post, identifiable by their unique post id.




Figure 5: Targetted Organizations

5 MEASUREMENT RESULTS

We ran FreePhish for a period of six months, from November 2022 to May 2023, identifying 31,405 zero-day phishing attacks (19,724 URLs from Twitter, 11,681 URLs from Facebook), and created using free website builders. These attacks targeted 109 unique brands and organizations, with the most frequently imitated ones illustrated in Figure 5. In this section, we identify the effectiveness of different phishing countermeasures against these attacks, including Blocklists, Browser protection tools, the respective FWB services themselves, as well as the social media platforms - Twitter and Facebook.

Comparison with self hosted phishing attacks: For each analysis, we also compared the detection coverage and response times for FWB-hosted phishing attacks with that of phishing attacks hosted on self-hosted domains. To find these self-hosted phishing attacks, we ran the base StackModel over a daily stream of Twitter and Facebook posts collected using their respective APIs from November 2022 to May 2023 (The same time period as our measurement for the FWB phishing attacks). We selected a random sample of 31,405 URLs and ensured that the distribution of the sample matched the distribution of the FWB phishing attacks across the two platforms: 19,724 phishing URLs from Twitter and 11,681 phishing URLs from Facebook.

5.1 Blocklist Performance

Low blocklist coverage and response time: Table 3 illustrates the blocklisting capabilities of PhishTank, OpenPhish, Google Safe Browsing (GSB), and eCrimeX against both FWB and regular (self-hosted) phishing URLs. GSB covered 18.4% of all FWB phishing URLs while having a median response time of 6 hours, which is much lower than their performance against self-hosted phishing attacks (74.2% coverage and median response time of 51 minutes). On the other hand, the two open-source blocklists, PhishTank and Openphish performed the worst against FWB-based attacks by only being able to cover 4.1% and 11.7% of all URLs, respectively, while

Method	FWB Phishing websites			Self-hosted Phishing Attacks		
	Coverage	Min/Max (hh:mm)	Median Speed (hh:mm)	Coverage	Min/Max	Median Speed
PhishTank	4.08%	0.10/116:13	07:11	17.4%	0:03/122:03	02:30
OpenPhish	11.70%	0:02/191:30	13:20	30.5%	0:01/146:18	02:21
GSB	18.44%	0:02/148:05	06:01	74.2%	0:01/146:26	00:51
eCrimeX	32.90%	0:07/137:43	08:54	47.9%	0:04/133:05	04:26
Social media Platform	23.06%	0:04/125:53	10:25	50.9%	0:07/114:23	03:41
Hosting domain	29.38%	0:19/158:25	09:43	77.50%	0:08/135:29	03:47

Table 3: Blocklisting performance and response time of anti-phishing entities against FWB phishing attacks.

Domains	URLs	Domain		Platform		PhishTank		OpenPhish		GSB		eCrimeX	
		Removal Rate	Speed (hh:mm)	Removal Rate	Speed	Removal Rate	Speed	Removal Rate	Speed	Removal Rate	Speed	Removal Rate	Speed
Weebly	7031	58.56%	01:39	20.65%	04:41	11.74%	07:16	13.12%	5:38	60.13%	0:30	23.46%	7:13
000webhost	5934	59.04%	00:45	13.82%	07:23	13.88%	05:16	10.70%	4:10	67.98 %	4:02	33.78%	4.75
Blogspot	3156	8.52%	06:51	25.12%	07:03	9.12%	05:00	11.10%	3:57	22.34%	9:12	30.11%	4.07
Wix.com	2338	64.55%	02:16	35.77%	04:35	12.73%	01:29	35.94%	1:26	43.66%	4:18	30.63%	5.08
Google Sites	2247	7.76%	12:22	28.45%	18:08	3.23%	15:43	5.28%	22:14	24.98%	13:55	14.40%	16.8
github.io	942	9.16%	20:34	21.46%	07:05	0.57%	06:01	13.06%	15:52	58.14%	7:40	20.44%	12.5
Firebase	1416	7.22%	14:15	26.86%	09:09	9.40%	14:35	12.09%	10:41	42.72%	3:13	26.08%	11.5
Squareup	1736	18.75%	10:11	34.45%	10:58	8.64%	13:50	6.68%	14:48	46%	11:01	34.22%	19.32
Zoho Forms	498	24.57%	07:11	15.77%	10:30	1.62%	10:24	8.84%	10:12	63.8%	3:59	31.22%	14.56
Wordpress	786	5.09%	20:50	28.96%	17:07	14.14%	13:48	8.18%	47:28	10.98%	14:22	12.47%	19.95
Google Forms	1397	11.96%	06:17	22.56%	31:27	3.87%	07:37	7.59%	29:19	39.45%	4:26	22.52%	11.8
Sharepoint	2181	7.64%	05:07	19.16%	07:41	13.73%	01:37	8.30%	16:28	16.65%	2:08	20.37%	5
Yolasite	601	7.52%	07:05	4.79%	20:37	10.46%	13:28	0	N/A	24.22%	1:31	0	N/A
GoDaddySites	418	5.84%	04:58	16.81%	33:55	0	N/A	2.45%	12:12	32.85%	11:44	0	N/A
MailChimp	183	23.67%	18:11	22.89%	48:07	2.15%	08:16	6.52%	7:02	21.34%	5:19	12.41%	7.26
glitch.me	480	21.31%	34:47	0	N/A	3.10%	10:33	7.08%	9:14	11.67%	16:48	0.00%	0
hpage	61	19.60%	11:45	0	N/A	0	N/A	0	N/A	13.11%	21:27	0	0

Table 4: Coverage and response times of different countermeasures against FWB phishing attacks.

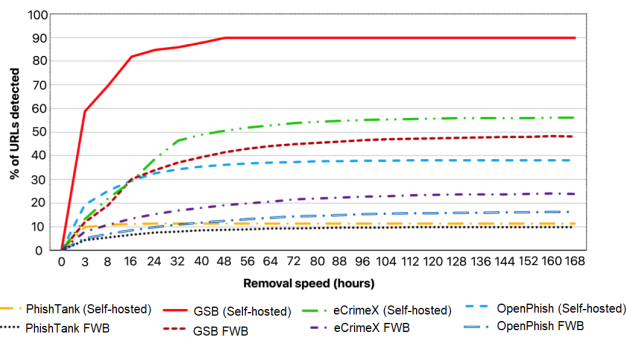


Figure 6: Coverage and speed of blocklists against FWB and self-hosted phishing attacks

having a median response time of 7 hours and 13 hours, respectively. As was the case with GSB, the performance of both these open source blocklists was significantly better against self-hosted phishing URLs, a trend also seen from eCrimeX’s performance, covering only 33% of all FWB URLs, and having a median response time of nearly 9 hours. To look closer at their blocklisting capabilities, we measured their coverage rate over regular intervals up to a week after their first appearance in our dataset, illustrated in Figure 6. We find that GSB covered nearly 60% of all self-hosted phishing URLs within the first 3 hours of their appearance compared to only 11% of FWB-hosted URLs during the same period. Within 24 hours, GSB could block up to 83% of all self hosted phishing URLs while only detecting about 31% of FWB phishing URLs.

On the other hand, eCrimeX had similar coverage for self hosted phishing attacks (11%) and FWB phishing URLs (8%) at the 3 hours mark. The gap is significantly widened near the 24-hour mark, with eCrimeX covering 38% of self-hosted phishing URLs, compared to only 13% of FWB URLs. We also found that all four blocklists had less than 10% coverage for nearly 40% of all FWB URLs. In fact, only GSB has nearly 40% coverage of 80% of the FWB URLs, with the same statistic being only around 8% for PhishTank, 12% for OpenPhish, and 30% for eCrimeX. Thus our measurement indicates that attackers have a lower chance of having their phishing attacks detected when created (and hosted) using an FWB, and even if the website does get detected, it would most likely take a longer time than self-hosted phishing attacks.

Blocklist coverage and response time varies on a per domain basis: As evident from our findings in the previous section, detection efforts towards FWB phishing attacks were poor across all blocklists that we tested. We then investigated whether blocklisting varied on a per FWB service basis. Table 4 illustrates blocklists’ coverage and response time over each of the FWBs found in our dataset. Our results indicate disproportionate performance by the blocklists based on which FWB the phishing website was created/hosted on. Interestingly, Weebly, 000webhost, and Wix, who collectively contributed to more than 48% of all URLs in our database, had higher coverage from all the blocklists when compared to phishing websites created on other FWBs. For example, GSB was able to detect 60.13% of all URLs hosted on Weebly and had a median response time of 30 mins, compared to a less popular FWB Yolasite, whose URLs have a coverage of only 24.22% with a median response time of 91 mins. However, an exception to this rule is phishing websites hosted on Blogspot, Google Sites, Google Forms, and Sharepoint,




Figure 7: Cumulative distribution of anti-phishing detections for FWB and self-hosted phishing websites

which despite being popular FWBs for hosting phishing attacks (collectively contributing to 28.5% of our dataset), had low coverage and high response time over different blocklists. For example, while PhishTank shows a relatively poor performance across all popular FWBs, it is particularly evident in the case of phishing attacks hosted on Google Sites, where they covered only 3.23% of the URLs and had a median response time of 15 hrs and 43 mins. We explore the reasons these attacks hosted on these FWBs have both low coverage and high response time at the end of our measurement study in Section 5.5. Blocklists also had negligible or no coverage for some less popular FWBs. For example, PhishTank had no coverage for phishing websites created on GoDaddy Sites and hpage, while having negligible coverage for github.io (0.57%), Zoho Forms (1.62%), Mailchimp (2.1%) and glitch.me (3.1%). A similar trend was also observed in OpenPhish and eCrimeX. While GSB covers URLs from all FWBs, it also has low coverage for several of them. The lack of blocklist coverage for a particular FWB might entice attackers to more frequently abuse that service.

5.2 Browser Protection Tool Coverage

We conducted a longitudinal study to evaluate the effectiveness of 76 third-party anti-phishing tools on our dataset. These tools are included in security software or are available for download as browser extensions [14, 58, 99]. Figure 7 presents the distribution of anti-phishing detection for FWB-based URLs in comparison to self-hosted phishing URLs one week after their initial appearance on the respective social networks. Our findings reveal that approximately 50% of FWB attacks disseminated via these two social media platforms had a median detection of four tool detections, while self-hosted phishing attacks boasted a median of nine detections, indicating that FWB attacks are less frequently detected than their self-hosted counterparts. We find that FWB URLs, regardless of the social media platform share a similar tool detection trajectory, a trend that is also observed for self-hosted phishing URLs. We also compared the proportion of phishing attacks detected with the number of detections on a daily basis up to a week following their initial appearance, illustrated in Figure 8. We discovered that nearly 75% of FWB-based URLs posted on Twitter had only 2 detections (the minimum threshold for inclusion in our dataset), and after

one week, 41% of these URLs still had four detections or fewer. By contrast, 82% of FWB URLs shared on Facebook had two detections on the first day, 43% only had four detections or fewer after a week. 32% of self-hosted phishing URLs shared via Twitter had 2 detections after one day, and 11% of these URLs had four or fewer detections after a week. For Facebook URLs, 34% of self-hosted phishing URLs had 2 detections after one day, and 8% of these URLs had four detections or fewer after a week. While the graph provides more detailed insights, one key takeaway is that over time, FWB phishing URLs accrue fewer detections from anti-phishing tools than self-hosted phishing URLs.

5.3 Coverage and Response of FWB Services

Considering the low blocklist coverage of FWB phishing attacks, proactive removal of these websites is vital to prevent potential victims from visiting these websites. This action is usually taken by the domain registrar. As illustrated in Table 3, we found only 29% of the websites were removed by the respective FWB service after two weeks since they first appeared in our dataset, at a median speed of 9 hours and 43 mins. Table 4 presents the coverage and response time of phishing attacks on a per FWB basis. Weebly and 000webhost having the highest number of reported URLs also exhibit relatively high domain removal rates, 58.56% and 59.04%, respectively, suggesting that these FWBs are heavily targeted by phishing attacks, prompting more proactive measures. Their domain response speeds are also relatively quick with a median speed of 1 hr 39 mins and 45 mins respectively, indicating efficient phishing attack mitigation processes. Wix.com, despite having fewer URLs than Weebly and 000webhost, boasts the highest domain removal rate at 64.55%, showcasing its effectiveness in dealing with such threats. Their response speed, however, is longer than Weebly and 000webhost but still faster than most other FWBs.

FWBs such as Blogspot, Google Sites, Github.io, Firebase, Sharepoint, Yolasite, and GoDaddySites have lower domain removal rates. This could point to less effective handling of phishing attacks. Interestingly, their response speeds vary widely, with Github.io and Firebase having particularly long response times with median removal times of 14 hrs 15 mins and 20 hrs 34 mins respectively, perhaps indicating less efficient processes for dealing with such threats. On the other hand, Squareup, Zoho Forms, Google Forms, MailChimp, glitch.me, and hpage exhibit a moderate domain removal rate. Their response speeds are also quite varied, suggesting differences in their effectiveness in dealing with phishing attacks. Thus, our findings indicate that the handling of phishing attacks varies widely among FWBs, with some services demonstrating a strong commitment to mitigating threats that are created on their platform while others lag behind. This variation can potentially have a significant impact on the overall security of these domains, and potentially lead to attackers favoring FWBs which have a lower response. **Response to reporting:** There are several factors that can influence the coverage and removal speed of phishing attacks by FWBs. Some of these are discussed towards the end of this Section in 5.5. However, one factor that might significantly impact their effectiveness is how each FWB registrar responds to our reports. WordPress, GoDaddySites, Firebase, Google Sites, Sharepoint, and Yolasite did not respond to any of our reports, which

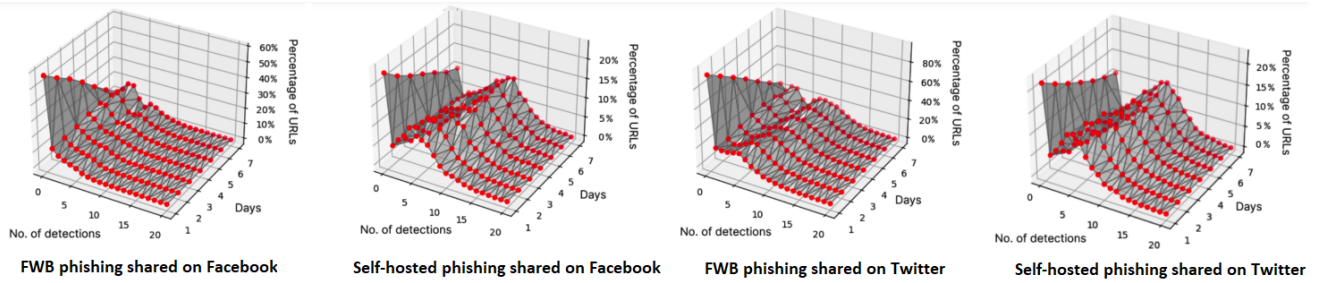


Figure 8: The percentage of FWB-based phishing URLs detected by anti-phishing engines over seven days.




Figure 9: Coverage and speed of platforms against FWB and self-hosted phishing attacks

raises concerns about their commitment to dealing with such issues. On the other hand, Squareup (23.7%), Github.io (37.4%), Google Sites (15.2%), and Blogspot (28.3%) acknowledged some of the initial reports, creating a support ticket but did not follow up with a resolution. Values inside parentheses denote the percentage of URLs where we received an initial response. Incidentally, these FWBs had the lowest coverage rates and longest response times, as evident from our findings in Table 4. This trend raises questions about whether they have implemented suitable measures to validate and respond to phishing URLs reported by our framework. On the contrary, FWBs with high coverage rates, such as Weebly (71.6%), Wix (65.3%), 000webhost (82.7%), and Zoho Forms (70.4%), were very responsive to our reports and followed up with additional information and the subsequent removal of both the website and user account associated with the phishing attack. Values inside parentheses denote the percentage of URLs where we received an initial response and a follow-up. Given that these FWBs are frequently targeted by phishing attacks, they seem to adopt a more alert stance towards these threats, emphasizing rapid action for their removal. Thus, the responsiveness of FWBs to these reports can significantly impact their coverage and response speed, affecting the overall security of these platforms.

5.4 Platform Effectiveness

By tracking the activity of posts that shared the URLs over the course of a week, we found that both of the platforms, i.e., Twitter and Facebook, collectively cover about 23% of all FWB URLs,

in comparison to 71.9% of self-hosted phishing URLs. The possibility exists that the tweets may have been deleted by the users themselves, not the platforms. However, prior research [28, 56] has shown that less than 2% of tweets are typically removed on social media platforms, translating to minimal noise in our measurement data.

The second column of Table 4 displays the platforms' coverage rate and response time to FWB attacks per FWB service basis. Unlike anti-phishing blocklists, which vary in their proficiency in detecting certain domains, or domain providers who demonstrate varying efficiency in removing FWB URLs, platform coverage (though low) does not differ significantly between FWBs. We do observe a higher median response time for popular FWBs, like Google Sites and Forms (18.08 hours and 31.27 hours, respectively), WordPress (17.07 hours), and Sharepoint (7.41 hours). Conversely, phishing websites hosted on less popular FWBs, such as Yolasite (20.37 hours), glitch.me (not applicable), GoDaddySites (33.55 hours), and MailChimp (48.07 hours) take longer to remove.

To illustrate the coverage and removal speed of FWB URLs shared on Twitter at a granular level, we monitored the coverage rate of these URLs over multi-hour periods up to a week (168 hours) after their initial appearance in our dataset. This is displayed in Figure 9. Within the first three hours, Twitter and Facebook removed only about 10% and 6% of FWB-based phishing URLs shared on their respective platforms, compared to the removal of 32% and 47% of self-hosted phishing URLs within the same time frame. After 16 hours, Twitter managed to remove over 70% of self-hosted phishing URLs while removing only 21% of FWB attacks. In the same period, Facebook removed over half of the self-hosted phishing URLs but only 10% of FWB URLs. Thus, phishing attacks hosted using FWBs on Twitter and Facebook tend to persist for much longer than self-hosted attacks.

From the end-users point of view, Figure 10 showcases the warning page Twitter used to display when a user attempted to navigate to a known malicious site. However, since Twitter's transition to the brand "X" in July 2023 [96], this warning mechanism seems to have been discontinued, even though the platform still actively removes posts featuring phishing URLs. Unlike Twitter, to the best of our knowledge, Facebook does not employ user-facing warning pages and instead directly deletes posts containing malicious links.

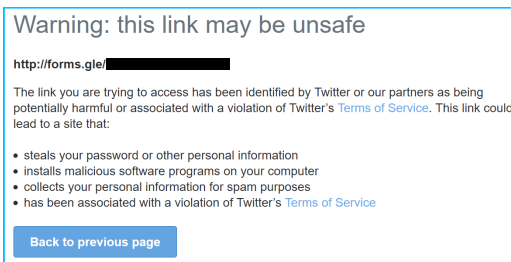


Figure 10: User alert displayed when attempting to access a website flagged as malicious by Twitter

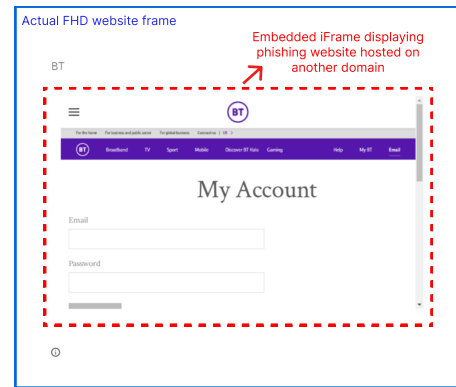


Figure 12: FWB containing an embedded i-Frame for another attack on an external domain.

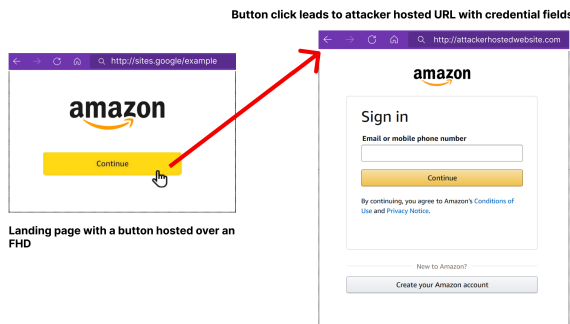


Figure 11: Two-step phishing example: A landing page on hosted FWB features only a button. Clicking on the button redirects users to a separate, attacker-controlled site that asks for login details.

5.5 Qualitative Analysis of Evasive Attacks

Nearly 14.2% of all URLs in our dataset (n=4,459 URLs) did not contain credential fields. Several of these URLs belonged to popular FWBs such as Google Sites, Google Forms, Blogspot, and Sharepoint. To investigate this issue closely, we randomly sampled 1K URLs from this set and qualitatively analyzed their content (both appearance and codebase) to identify the attack vectors, which led us to identify three variants of malicious websites that deviate from regular phishing strategies. We developed heuristics to automatically identify these attack vectors across our dataset’s FWB phishing attacks.

Linking to other phishing pages: We found 539 (about 24%) URLs hosted on Google Sites containing a button, clicking on which led to a phishing website on a different domain. This design serves two evasive purposes - Avoiding credential fields on the website, to more effectively evade detection by anti-phishing bots, and - by not directly sharing the linked phishing page accessed through user clicks, attackers lower the risk of discovery and detection. These URLs’ low coverage and response time are evident from the fact that few approaches [59] have been able to identify this evasion tactic. For 174 URLs, the attackers used another FWB as its linked page, which can be more effective at evasion. An example of this attack is illustrated in Figure 11. We also found 349 (about 16%) URLs on Sharepoint, 293 (about 21%) on Google Forms, and 447 (about 14%) on Blogspot.

i-Frames Embedding External Attacks: In 427 URLs on Google Sites and 473 URLs on Blogspot, attackers embedded a phishing site’s i-frame into a benign website’s code hosted on an FWB. This concealed i-Frame activates either by button click on the parent website or by loading alongside it. Figure 12 illustrates an example on Google Sites. These attacks could evade prevalent anti-phishing tools, as i-frame content is client-side rendered and unavailable to tools scanning for malicious code [65]. We discovered these FWB domains comprise 62% of all i-frame phishing attacks in our dataset.

Malicious drive-by downloads: 725 (23%) of the URLs on Blogspot, 651 URLs (29%) on Google Sites, and 1,177 URLs (54%) on Sharepoint were used to share malicious downloads hosted on third-party websites. For each URL that triggered a download, we scanned the file using VirusTotal. If the file had four or more detections (a threshold proposed in literature [71, 87]), the FWB URL was marked as distributing malicious drive-by downloads. Combining this with the absence of credential-requiring fields in such websites hosting malicious downloads might further indicate the low coverage of these attacks from anti-phishing entities. Interestingly, 741 malicious URLs were hosted on Sharepoint, mimicking OneDrive or Office 365 documents, highlighting targeted spoofing of Microsoft Office products.

6 RELATED WORK

Anti-phishing systems: Phishing attacks pose a significant threat to users [31, 53], prompting the integration of anti-phishing tools [4, 14, 58] in web-browsers. These tools rely on blocklists, databases updated with known phishing sites taken from community reports [5], third-party data [69], or web crawler frameworks [65]. Researchers have improved phishing detection using automated models incorporating machine and deep learning [79, 83], and heuristics [63, 90].

Nevertheless, recent studies have highlighted the vulnerability of blocklists and automated tools to evasive attacks. For instance, Oest et al.[65] unveiled the significant impact of server-side cloaking attacks on blocklist detection, while Acharya et al.[20] and AlEroud et al. [23] outlined ways attackers use browser-based fingerprinting obfuscations and meta-URL features to bypass automated anti-phishing crawlers.

Our work explores how attackers leverage features of several Free Website Builder services to evade detection. We characterize unique aspects of these attacks, enabling the creation of a real-time automated machine learning model for detecting such threats in the wild.

Anti-phishing measurement: Anti-phishing measurements at scale are instrumental in revealing emerging threats and gaps in the current detection ecosystem. Oest et al's frameworks, Phish-Time [65] and Golden Hour [67], focus on real-time phishing threat identification and traffic monitoring, stressing the need for fortified mobile security and evidence-based reporting protocols. On another front, Bijman et al [30] sheds light on the influence of off-the-shelf phishing kits in the Dutch financial sector and suggested policy enhancements. Further, Bell's et al [29] examination of three phishing blacklists underscores the necessity for improved blacklist management, given their varying size, update frequency, and potential effectiveness. Our framework, FreePhish contributes to this collective effort by identifying and measuring the impact and efficacy of phishing websites created using Free Website Building services against anti-phishing tools, blocklists, hosting services, and sharing platforms.

Phishing Attack Costs: Phishing kits [10] have reduced the technical expertise needed to execute phishing attacks as they can automate the creation and hosting of phishing sites and can incorporate evasion strategies[30, 66]. However, they are typically costly for attackers [47], and free versions often contain malware [34]. Furthermore, attackers typically acquire inexpensive web domains with non-traditional top-level domains (TLDs) from hosting services, such as *.store*, *.top*, or *.live*[45], given the short lifespan of most phishing sites[67]. Yet, the high volume of phishing attacks on these TLDs invites more scrutiny from anti-phishing blocklists and hosting services [22, 32, 97], thereby reducing the attacks' uptime. On the other hand, our research highlights how attackers leverage Free Website Builders to alleviate the financial burdens typically associated with phishing attacks. These platforms offer a host of advantages, including premium top-level domains, instant SSL certification, and user-friendly design interfaces, thereby making it possible to host sophisticated phishing sites at no cost.

7 CONCLUSION

Our study investigated and described the tactics used to construct phishing websites using free website builders. We discovered that attackers exploit various features offered by these services to create phishing websites that can evade detection by anti-phishing entities. To combat this, we developed a framework, *FreePhish*, designed to monitor and promptly identify FWB phishing attacks shared on social media websites. This framework reports these attacks for immediate takedown and assesses the efficiency of anti-phishing entities in countering these threats. In six months, we identified and reported over 31.4K newly generated FWB phishing URLs, thus providing both FWB services and social media platforms a tool to mitigate and remove these threats more effectively. We also FreePhish as a free web extension that can be installed on Chromium-based browsers to protect users from accessing FWB phishing URLs. Figure 13 illustrates the extension in action where it has blocked a malicious FWB phishing attack. The web extension

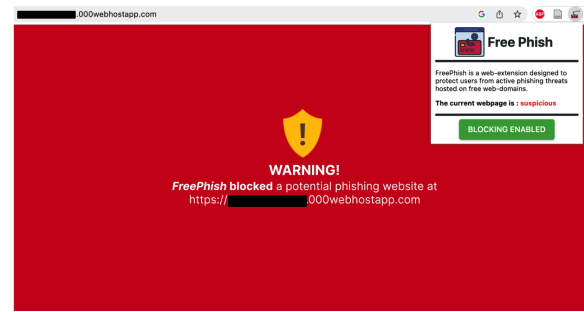


Figure 13: The FreePhish web-extension in action

and source code of the FreePhish framework source code is available at: <https://github.com/UTA-SPRLab/freephish/>.

8 ETHICS

We address the sensitivity and potential risk of phishing attacks examined in our study by abstaining from publicly sharing the specific URLs utilized. But to help the research community to evaluate these attacks, our initial dataset will be available upon request. At the conclusion of our measurement, we have also reported our URLs to the APWG eCrimeX blacklist [27], and plan to continue reporting FWB phishing attacks found by our FreePhish framework to the respective FWBs and social media platforms. The qualitative analysis of this data was conducted securely, using a designated computing system situated within our research laboratory. The two coders involved had access to the URLs strictly within this controlled environment. We have also not retained or made use of any personally identifiable information from the posts gathered from Twitter and Facebook. The data has been used exclusively in an aggregated manner, specifically for relevant analysis within the scope of our work.

ACKNOWLEDGMENT

We express our deep gratitude to the anonymous reviewers and our shepherd, Dr. Cigdem Sengul, for their constructive feedback. We also thank Arjun Dahal for helping us implement the algorithm for calculating code similarity between FWB websites and Ryan Lahlou for contributing to the FreePhish web extension. This work was partially supported by the Comcast Innovation Fund.

REFERENCES

- [1] 2017. Phishing attack abuses Microsoft Azure, Google Sites to steal crypto. <https://www.infoworld.com/article/3187346/phishing-scammers-exploit-wix-web-hosting.html>.
- [2] 2019. Drive-By Malware Uses Google Sites for Drive by Download Attacks. bit.ly/3WJMC55.
- [3] 2020. Friends of PhishTank [Infographic]. <https://www.phishtank.com/friends.php>.
- [4] 2020. Google Safebrowsing. <https://safebrowsing.google.com/>.
- [5] 2020. PhishTank. <https://www.phishtank.com/faq.php>.
- [6] 2020. Twitter Developer. <https://developer.twitter.com/en/docs>.
- [7] 2020. VirusTotal. <https://www.virustotal.com/gui/home/>.
- [8] 2020. Weaponizing Domain Names: how bulk registration aids global spam campaigns. bit.ly/3Q0whcK.
- [9] 2021. 2021 Cybersecurity threat trends: phishing, crypto top the list. bit.ly/3owhpWv.

- [10] 2021. Have Money for a Latte? Then You Too Can Buy a Phish Kit. <https://www.proofpoint.com/us/blog/threat-insight/have-money-latte-then-you-too-can-buy-phish-kit>.
- [11] 2021. How long does it take to get an SSL certificate? <https://comodossllstore.com/ssl-validation-process/how-long-to-get-an-ssl-certificate>.
- [12] 2022. Block Search indexing with noindex. <https://developers.google.com/search/docs/crawling-indexing/block-indexing>.
- [13] 2022. Crowdtangle. <https://www.crowdtangle.com/>.
- [14] 2022. McAfee WebAdvisor. <https://www.mcafee.com/en-us/safe-browser/mcafee-webadvisor.html>.
- [15] 2022. Phishing attack abuses Microsoft Azure, Google Sites to steal crypto. bit.ly/3C08g1Y.
- [16] 2023. Community Standards - Cybersecurity. <https://transparency.fb.com/en-gb/policies/community-standards/cybersecurity/>.
- [17] 2023. Phishing, spam, and malware links. <https://help.twitter.com/en/safety-and-security/phishing-spam-and-malware-links>.
- [18] Sahar Abdelnabi, Katharina Krombolz, and Mario Fritz. 2020. VisualPhishNet: Zero-day phishing website detection by visual similarity. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*. 1681–1698.
- [19] Lawrence Abrams. 2021. These are the top-level domains threat actors like the most. *BleepingComputer* (2021). <https://www.bleepingcomputer.com/news/security/these-are-the-top-level-domains-threat-actors-like-the-most/>
- [20] Bhupendra Acharya and Phani Vadrevu. 2021. {PhishPrint}: Evading Phishing Detection Crawlers by Prior Profiling. In *30th USENIX Security Symposium (USENIX Security 21)*. 3775–3792.
- [21] Sadia Afroz and Rachel Greenstadt. 2011. Phishzoo: Detecting phishing websites by looking at them. In *2011 IEEE fifth international conference on semantic computing*. IEEE, 368–375.
- [22] Pieter Agten, Wouter Joosen, Frank Piessens, and Nick Nikiforakis. 2015. Seven months' worth of mistakes: A longitudinal study of typosquatting abuse. In *Proceedings of the 22nd Network and Distributed System Security Symposium (NDSS 2015)*. Internet Society.
- [23] Ahmed AlErroud and George Karabatis. 2020. Bypassing detection of URL-based phishing attacks using generative adversarial deep neural networks. In *Proceedings of the Sixth International Workshop on Security and Privacy Analytics*. 53–60.
- [24] Mona Ghotiaish Alkhozai and Omar Abdullah Batarfi. 2011. Phishing websites detection based on phishing characteristics in the webpage source code. *International Journal of Information and Communication Technology Research* 1, 6 (2011).
- [25] Manuel Álvarez, Alberto Pan, Juan Raposo, Fernando Bellas, and Fidel CACHED. 2007. Using clustering and edit distance techniques for automatic web data extraction. In *International Conference on Web Information Systems Engineering*. Springer, 212–224.
- [26] VirusTotal APL. [n. d.]. <https://developers.virustotal.com/v3.0/reference#overview>.
- [27] APWG. [n. d.]. eCrimeX. <https://apwg.org/ecx/>.
- [28] Lujo Bauer, Lorrie Faith Cranor, Saranga Komanduri, Michelle L Mazurek, Michael K Reiter, Manya Sleeper, and Blase Ur. 2013. The post anachronism: The temporal dimension of Facebook privacy. In *Proceedings of the 12th ACM Workshop on Workshop on Privacy in the Electronic Society*. 1–12.
- [29] Simon Bell and Peter Komisarczuk. 2020. An analysis of phishing blacklists: Google safe browsing, openphish, and phishtank. In *Proceedings of the Australasian Computer Science Week Multiconference*. 1–11.
- [30] Hugo Bijmans, Tim Booi, Anneke Schwedersky, Aria Nedgabat, and Rolf van Wegberg. 2021. Catching phishers by their bait: Investigating the dutch phishing landscape through phishing kit detection. In *30th USENIX Security Symposium (USENIX Security 21)*. 3757–3774.
- [31] Jan-Willem Hendrik Bullée, Lorena Montoya, Wolter Pieters, Marianne Junger, and Pieter Hartel. 2018. On the anatomy of social engineering attacks—A literature-based dissection of successful attacks. *Journal of investigative psychology and offender profiling* 15, 1 (2018), 20–45.
- [32] CircleID. [n. d.]. Which Domains Stand the Strongest Against Phishing Attacks? <https://bit.ly/3h0IkWr>.
- [33] Comodo. 2022. A DV vs EV Certificate — Why Should I Choose One Over the Other? <https://comodossllstore.com/resources/a-dv-vs-ev-certificate-why-should-i-choose-one-over-the-other/>.
- [34] Marco Cova, Christopher Kruegel, and Giovanni Vigna. 2008. There Is No Free Phish: An Analysis of “Free” and Live Phishing Kits. *WOOT* 8 (2008), 1–8.
- [35] Qian Cui, Guy-Vincent Jourdan, Gregor V Bochmann, and Iosif-Viorel Onut. 2021. Proactive detection of phishing kit traffic. In *Applied Cryptography and Network Security: 19th International Conference, ACNS 2021, Kamakura, Japan, June 21–24, 2021, Proceedings, Part II*. Springer, 257–286.
- [36] Julie S Downs, Mandy Holbrook, and Lorrie Faith Cranor. 2007. Behavioral response to phishing risk. In *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*. 37–44.
- [37] Arthur Drichel, Vincent Drury, Justus von Brandt, and Ulrike Meyer. 2021. Finding phish in a haystack: A pipeline for phishing classification on certificate transparency logs. In *The 16th International Conference on Availability, Reliability and Security*. 1–12.
- [38] Vincent Drury and Ulrike Meyer. 2019. Certified phishing: taking a look at public key certificates of phishing websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 211–223.
- [39] Let's Encrypt. 2023. A nonprofit Certificate Authority providing TLS certificates to 300 million websites. <https://letsencrypt.org/>.
- [40] Edona Faslija, Hasan Ferit Erişer, and Bernd Prünster. 2019. Phish-Hook: Detecting phishing certificates using certificate transparency logs. In *International Conference on Security and Privacy in Communication Systems*. Springer, 320–334.
- [41] Mozilla Firefox. [n. d.]. How does built-in Phishing and Malware Protection work? <https://support.mozilla.org/en-US/kb/how-does-phishing-and-malware-protection-work>.
- [42] Google. [n. d.]. Use Safe Browsing in Chrome. <https://bit.ly/3UYV1ZF>.
- [43] GrowthBadger. 2021. Domain Extensions: .com vs. .org, .net, .io & 4 Other TLDs (Study). <https://growthbadger.com/top-level-domains>.
- [44] Brij B Gupta and Ankit Kumar Jain. 2020. Phishing attack detection using a search engine and heuristics-based technique. *Journal of Information Technology Research (JITR)* 13, 2 (2020), 94–109.
- [45] Srishiti Gupta and Ponnurangam Kumaraguru. 2014. Emerging phishing trends and effectiveness of the anti-phishing landing page. In *2014 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 36–47.
- [46] Xiao Han, Nizar Kheir, and Davide Balzarotti. 2016. Phisheye: Live monitoring of sandboxed phishing kits. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 1402–1413.
- [47] Cormac Herley and Dinei Florêncio. 2010. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In *Economics of information security and privacy*. Springer, 33–53.
- [48] Jun Ho Huh and Hyoungshick Kim. 2011. Phishing detection with popular search engines: Simple and effective. In *International Symposium on Foundations and Practice of Security*. Springer, 194–207.
- [49] Interisle. 2020. Phishing Landscape 2020. <https://interisle.net/PhishingLandscape2020.pdf>.
- [50] Interisle. 2021. Phishing Landscape 2021. <https://interisle.net/PhishingLandscape2021.pdf>.
- [51] Interisle. 2022. Phishing Landscape 2022. <https://interisle.net/PhishingLandscape2022.pdf>.
- [52] Interisle. 2023. Phishing Landscape 2023. <https://interisle.net/PhishingLandscape2023.pdf>.
- [53] Marianne Junger, Lorena Montoya, and F-J Overink. 2017. Priming and warnings are not effective to prevent social engineering attacks. *Computers in human behavior* 66 (2017), 75–87.
- [54] KrebsonSecurity. [n. d.]. <https://krebsonsecurity.com/2018/11/half-of-all-phishing-sites-now-have-the-padlock/>.
- [55] Hung Le, Quang Pham, Doyen Sahoo, and Steven CH Hoi. 2018. URLNet: Learning a URL representation with deep learning for malicious URL detection. *arXiv preprint arXiv:1802.03162* (2018).
- [56] Linfeng Li, Eleni Berki, Marko Helenius, and Salla Ovaska. 2014. Towards a contingency approach with whitelist-and blacklist-based anti-phishing applications: what do usability tests indicate? *Behaviour & Information Technology* 33, 11 (2014), 1136–1147.
- [57] Yukun Li, Zhenguo Yang, Xu Chen, Huaping Yuan, and Wenyin Liu. 2019. A stacking model using URL and HTML features for phishing webpage detection. *Future Generation Computer Systems* 94 (2019), 27–39.
- [58] BitDefender Traffic Light. [n. d.]. <https://www.bitdefender.com/solutions/traffilight.html>.
- [59] Ruofan Liu, Yun Lin, Xianglin Yang, Siang Hwee Ng, Dinil Mon Divakaran, and Jin Song Dong. 2022. Inferring Phishing Intention via Webpage Appearance and Dynamics: A Deep Vision Based Approach. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*.
- [60] Security Magazine. [n. d.]. Phishing could have cost businesses \$354m in potential direct losses. bit.ly/30LwJ5N.
- [61] Mailchimp. 2023. Compare Marketing plans. <https://mailchimp.com/pricing/marketing/compare-plans/>.
- [62] Mozilla. [n. d.]. <https://blog.mozilla.org/firefox/6-tips-for-safe-online-shopping/>.
- [63] Luong Anh Tuan Nguyen, Ba Lam To, Huu Khuong Nguyen, and Minh Hoang Nguyen. 2013. Detecting phishing web sites: A heuristic URL-based approach. In *2013 International Conference on Advanced Technologies for Communications (ATC 2013)*. IEEE, 597–602.
- [64] Adam Oest, Yeganeh Safaei, Adam Doupé, Gail-Joon Ahn, Brad Wardman, and Kevin Tyers. 2019. Phishfarm: A scalable framework for measuring the

- effectiveness of evasion techniques against browser phishing blacklists. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1344–1361.
- [65] Adam Oest, Yeganeh Safaei, Penghui Zhang, Brad Wardman, Kevin Tyers, Yan Shoshitaishvili, and Adam Doupe. 2020. Phishtime: Continuous longitudinal measurement of the effectiveness of anti-phishing blacklists. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*. 379–396.
- [66] Adam Oest, Yeganeh Safaei, Adam Doupe, Gail-Joon Ahn, Brad Wardman, and Gary Warner. 2018. Inside a phisher’s mind: Understanding the anti-phishing ecosystem through phishing kit analysis. In *2018 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 1–12.
- [67] Adam Oest, Penghui Zhang, Brad Wardman, Eric Nunes, Jakub Burgis, Ali Zand, Kurt Thomas, Adam Doupe, and Gail-Joon Ahn. 2020. Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale. In *29th USENIX Security Symposium (USENIX Security 20)*.
- [68] Chidimma Opara, Bo Wei, and Yingke Chen. 2020. HTMLPhish: enabling phishing web page detection by applying deep learning techniques on HTML analysis. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [69] Openphish. [n. d.]. Phishing feed. "https://openphish.com/faq.html".
- [70] OpenPhish. 2022. OpenPhish Monthly Brands List. https://openphish.com/monthly_brands_list.csv.
- [71] Alina Oprea, Zhou Li, Robin Norris, and Kevin Bowers. 2018. Made: Security analytics for enterprise threat detection. In *Proceedings of the 34th Annual Computer Security Applications Conference*. 124–136.
- [72] VARN original research. [n. d.]. How your domain name could be damaging your business. https://varn.co.uk/10/27/70-of-people-dont-trust-newer-website-domains/.
- [73] Srushti Patil and Sudhir Dhage. 2019. A methodical overview on phishing detection along with an organized way to construct an anti-phishing framework. In *2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS)*. IEEE, 588–593.
- [74] Peng Peng, Limin Yang, Linhai Song, and Gang Wang. 2019. Opening the blackbox of virustotal: Analyzing online phishing scan engines. In *Proceedings of the Internet Measurement Conference*. 478–485.
- [75] Richard Penman. [n. d.]. Python-WhoIS. https://pypi.org/project/python-whois/.
- [76] PhishLabs. 2021. Top 10 TLDs Abused. https://www.phishlabs.com/blog/top-10-tlds-abused/.
- [77] Pawan Prakash, Manish Kumar, Ramana Rao Kompella, and Minaxi Gupta. 2010. Phishnet: predictive blacklisting to detect phishing attacks. In *2010 Proceedings IEEE INFOCOM*. IEEE, 1–5.
- [78] Shama Rani and Jaiteg Singh. 2017. Enhancing Levenshtein’s edit distance algorithm for evaluating document similarity. In *International Conference on Computing, Analytics and Networks*. Springer, 72–80.
- [79] Routhu Srinivasa Rao and Alwyn Roshan Pais. 2019. Detection of phishing websites using an efficient feature-based machine learning framework. *Neural Computing and Applications* 31, 8 (2019), 3851–3873.
- [80] Routhu Srinivasa Rao and Alwyn Roshan Pais. 2019. Jail-Phish: An improved search engine based phishing detection system. *Computers & Security* 83 (2019), 246–267.
- [81] S Roopak and Tony Thomas. 2014. A novel phishing page detection mechanism using html source code comparison and cosine similarity. In *2014 Fourth International Conference on Advances in Computing and Communications*. IEEE, 167–170.
- [82] S Roopak, Athira P Vijayaraghavan, and Tony Thomas. 2019. On effectiveness of source code and SSL based features for phishing website detection. In *2019 1st International Conference on Advanced Technologies in Intelligent Control, Environment, Computing & Communication Engineering (ICATIECE)*. IEEE, 172–175.
- [83] Ozgur Koray Sahingoz, Ebubekir Buber, Onder Demir, and Banu Diri. 2019. Machine learning based phishing detection from URLs. *Expert Systems with Applications* 117 (2019), 345–357.
- [84] Yuji Sakurai, Takuya Watanabe, Tetsuya Okuda, Mitsuki Akiyama, and Tatsuya Mori. 2020. Discovering HTTPSified phishing websites using the TLS certificates footprints. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 522–531.
- [85] Nuttapon Sanglerdsinlapachai and Arnon Rungsawang. 2010. Using domain top-page similarity feature in machine learning-based web phishing detection. In *2010 Third International Conference on Knowledge Discovery and Data Mining*. IEEE, 187–190.
- [86] Securelist. 2022. What are phishing kits? https://securelist.com/phishing-kit-market-whats-inside-off-the-shelf-phishing-packages/106149/
- [87] Mahmood Sharif, Jumpei Urakawa, Nicolas Christin, Ayumu Kubota, and Akira Yamada. 2018. Predicting impending exposure to malicious content from user behavior. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*. 1487–1501.
- [88] Fu Song, Yusi Lei, Sen Chen, Lingling Fan, and Yang Liu. 2021. Advanced evasion attacks and mitigations on practical ML-based phishing website classifiers. *International Journal of Intelligent Systems* 36, 9 (2021), 5210–5240.
- [89] Spamhaus. [n. d.]. The World’s Most Abused TLDs. https://www.spamhaus.org/statistics/tlds/.
- [90] Jayesh Sreedharan and Rahul Mohandas. 2016. Systems and methods for risk rating and pro-actively detecting malicious online ads. US Patent 9,306,968.
- [91] SSL.com. 2022. DV, OV, IV, and EV Certificates. https://www.ssl.com/article/dv-ov-and-ev-certificates/.
- [92] Tanmay Thakur and Rakesh Verma. 2014. Catching classical and hijack-based phishing attacks. In *International Conference on Information Systems Security*. Springer, 318–337.
- [93] CNBC Tom Huddlestone Jr. [n. d.]. How this scammer used phishing emails to steal over \$100 million from Google and Facebook. bit.ly/45qZXTM.
- [94] Certificate Transparency. [n. d.]. Working together to detect maliciously or mistakenly issued certificates. https://certificate.transparency.dev.
- [95] Twitter. [n. d.]. https://developer.twitter.com/en.
- [96] The Verge. 2023. Twitter is being rebranded as X. https://www.theverge.com/2023/7/23/23804629/twitters-rebrand-to-x-may-actually-be-happening-soon.
- [97] Rakesh Verma and Keith Dyer. 2015. On the character of phishing URLs: Accurate and robust statistical learning classifiers. In *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*. 111–122.
- [98] Vivaldi. [n. d.]. Privacy and Security Settings. https://help.vivaldi.com/desktop/privacy/privacy-settings/.
- [99] Norton Safe Web. [n. d.]. https://safeweb.norton.com/.
- [100] Weebly. 2023. TLS Certificate. https://www.weebly.com/online-store/tls-certificate.
- [101] Wix. 2023. Wix homepage. https://www.wix.com/.
- [102] Guang Xiang, Jason Hong, Carolyn P Rose, and Lorrie Cranor. 2011. Cantina+ a feature-rich machine learning framework for detecting phishing web sites. *ACM Transactions on Information and System Security (TISSEC)* 14, 2 (2011), 1–28.
- [103] ZeroSSL. 2023. Free SSL Certificates and SSL Tools - ZeroSSL. https://zerossl.com/.
- [104] Liang Zhang, David Choffnes, Dave Levin, Tudor Dumitras, Alan Mislove, Aaron Schulman, and Christo Wilson. 2014. Analysis of SSL certificate reissues and revocations in the wake of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. 489–502.
- [105] Penghui Zhang, Adam Oest, Haehyun Cho, Zhibo Sun, RC Johnson, Brad Wardman, Shaown Sarker, Alexandros Kapravelos, Tiffany Bao, Ruoyu Wang, et al. 2021. Crawlphish: Large-scale analysis of client-side cloaking techniques in phishing. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1109–1124.

A CODE SIMILARITY BETWEEN FWB WEBSITES

In Section 3, we calculate the code similarity between phishing attacks and legitimate websites that are created on FWBs. To calculate the similarity between two websites (Website A and Website B), we first extracted the tag elements from each website. For each tag element T in Website A, we calculated its Levenshtein distance (LV) to every tag element in Website B. The tag from Website B with the lowest Levenshtein distance was considered the most similar to the T tag of Website A. This is denoted as T_{\max} . Thus,

$$T_{\max} = \min\{LV(T, T_B) \mid \text{for each } T_B \in \text{Website B}\}$$

We then calculated the similarity between Website A and Website B ($\text{sim}_{A \text{ to } B}$) as the median of all T_{\max} values.

$$\text{sim}_{A \text{ to } B} = \text{median}(T_{\max} \mid \text{for all } T \in \text{Website A})$$

We also calculated the similarity from Website B to Website A ($\text{sim}_{B \text{ to } A}$) using the same procedure.

$$\text{sim}_{B \text{ to } A} = \text{median}(T_{\max} \mid \text{for all } T \in \text{Website B})$$

Finally, the similarity between the two websites ($\text{sim}_{A,B}$) was calculated as the mean of $\text{sim}_{A \text{ to } B}$ and $\text{sim}_{B \text{ to } A}$.

$$\text{sim}_{A,B} = \text{mean}(\text{sim}_{A \text{ to } B}, \text{sim}_{B \text{ to } A})$$