Identifying the Phishing Websites Using the Patterns of TLS Certificates

Yuji Sakurai[1], Takuya Watanabe[2], Tetsuya Okuda[2], Mitsuaki Akiyama[2] and Tatsuya Mori[1,3]

[1]*Waseda University, Shinjuku City, Tokyo, Japan*
[2]*NTT Secure Platform Laboratories, Japan*
[3]*NICT, Japan*
E-mail: s5i@nsl.cs.waseda.ac.jp
[*]*Corresponding Author*

Received 29 December 2020; Accepted 30 December 2020;
Publication 10 April 2021

Abstract

With the recent rise of HTTPS adoption on the Web, attackers have begun "HTTPSifying" phishing websites. HTTPSifying a phishing website has the advantage of making the website appear legitimate and evading conventional detection methods that leverage URLs or web contents in the network. Further, adopting HTTPS could also contribute to generating intrinsic *footprints* and provide defenders with a great opportunity to monitor and detect websites, including phishing sites, as they would need to obtain a public-key certificate issued for the preparation of the websites. The potential benefits of certificate-based detection include (1) the comprehensive monitoring of all HTTPSified websites by using certificates immediately after their issuance, even if the attacker utilizes dynamic DNS (DDNS) or hosting services; this could be overlooked with the conventional domain-registration-based approaches; and (2) to detect phishing websites before they are published on the Internet. Accordingly, we address the following research question: *How can we make use of the footprints of TLS certificates to defend against phishing attacks?* For this, we collected a large set of TLS certificates corresponding to phishing websites from Certificate Transparency (CT) logs and extensively analyzed these TLS certificates. We demonstrated that a *template*

Journal of Cyber Security and Mobility, Vol. 10_2, 451–486.
doi: 10.13052/jcsm2245-1439.1026

of common names, which are equivalent to the fully qualified domain names, obtained through the clustering analysis of the certificates can be used for the following promising applications: (1) The discovery of previously *unknown* phishing websites and (2) understanding the infrastructure used to generate the phishing websites. Furthermore, we developed a real-time monitoring system using the analysis techniques. We demonstrate its usefulness for the practical security operation. We use our findings on the abuse of free certificate authorities (CAs) for operating HTTPSified phishing websites to discuss possible solutions against such abuse and provide a recommendation to the CAs.

Keywords: Phishing, TLS, Certificate.

1 Introduction

The adoption of HTTPS on the Web has increased drastically over the past few years [12, 15]. According to Google's Transparency Report [15], in several countries, such as the United States, Germany, and France, more than 90% of Web traffic has been "HTTPSified." The rate of HTTPSified Web traffic in other countries has also grown over time; for example, in Brazil, Japan, and India, more than 70% of the Web traffic has been encrypted with HTTPS. The primary factors that contribute to the drastic increase in the adoption of HTTPS are continuing HTTPS promotion efforts, such as changes in search engine rankings [17], revisions to security indicators on Web browsers [16, 25], and the publication of useful tools to install or assess HTTPSified websites [21], though the outreach of HTTPS could be widened to impact several other areas [12]. Notably, the cost of the "S" in HTTPS has been significantly reduced in recent times, as reported by Naylor et al. [26]. These changes should have contributed to the widespread adoption of HTTPS.

However, even as the number of HTTPSified websites has drastically increased, phishing websites have also started adopting HTTPS. By adopting HTTPS, an attacker could make his/her phishing website appear legitimate. In addition, the end-to-end encryption mechanism ensures that access to the HTTPSified phishing website can evade network-level detection (e.g., at a web proxy or gateway) that leverages URLs or web content. Furthermore, the recent rise in freely available certificate authorities (CAs), such as Let's Encrypt [21] and cPanel [8], has lowered the barriers to deploying HTTPS on a website. According to the 2019 Q1 Phishing Activity Trends Report

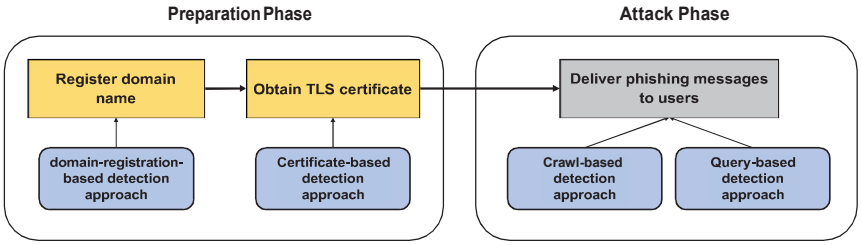

Figure 1 Comparison of phishing-detection approaches based on domain registration, certificate, crawl, and query by their detection phase.

of the Anti-Phishing Working Group (APWG) [2], less than 2% of phishing websites in 2015 adopted HTTPS; however, this number started increasing rapidly since the end of 2016, and reached 74% in 2019.

While HTTPSifying a phishing website may bring several advantages for an attacker, it could also contribute to generating intrinsic *footprints*, which in turn be used to systematically detect the HTTPSified phishing website. The key insight behind this assumption is that by HTTPSifying a website, an attacker must register a valid public key certificate (i.e., TLS certificate) that contains intrinsic features such as issued date, issuer name (CA), and common name (CN). The CN in a certificate is equivalent to the fully qualified domain name (FQDN) of a server. In addition, certificate transparency (CT), which is a standardized framework to publish the public logs of all the issued TLS certificates, plays a vital role for monitoring and auditing TLS certificates. Thus, we expect that we can efficiently detect phishing websites by analyzing TLS certificates.

To understand phishing-detection approaches in terms of a phase in a phishing-attack process, we classified the approaches into domain-name registration, obtaining of issued certificates, and phishing-message delivery to users as shown in Figure 1. The first two approaches comprise the preparation phase, and the last approach comprises the attack phase. The detection approaches in the attack phase, i.e., crawl- and query-based approaches, can find phishing websites only after they are published because these approaches are triggered by the delivered phishing messages (e.g., email/SMS messages and social media contents) or user access to a phishing website. In contrast, the detection approaches in the preparation phase, i.e., domain-registration- and certificate-based methods can find phishing websites in the early phase in which we cannot make use of any phishing messages nor user accesses for detecting phishing attacks. However, the domain-registration-based approach is limited in that not all FQDNs can be found by this approach because

WHOIS records contain only domain names (i.e., websites using DDNS or hosting services cannot be found through this approach). In contrast, the key advantages of leveraging TLS certificates are (1) ability to thoroughly monitor all FQDNs of HTTPSified websites through the issued certificates even if the website owners use DDNS or hosting services, which a domain-registration-based approach may miss, and (2) ability to detect phishing websites before they are published online; these advantages are missed in the conventional crawl- and query-based approaches.

Therefore, in this paper, we address the following research question:

> **RQ**: *How can we make use of the footprints of TLS certificates to defend against phishing attacks?*

Before answering this question, we overview the existing works that attempted to detect phishing websites by using the information contained in the TLS certificates [9, 10, 40]. Among them, the most recent study by Drury and Meyer [10] concluded that distinguishing malicious websites from those that are benign is difficult if they are issued certificates from the same CAs. This is because in the case that both types of websites use certificates issued by common free CAs, such as Let's Encrypt and cPanel, the certificate would have many shared fields, thus complicating their distinction.

To overcome this limitation and address the aforementioned RQ, we focused on the CN, which is the field an attacker can arbitrarily change, and the bulk registration during the survey period.[1] Several previous studies have shown that many attackers generate similar domain names in a short time [22]. In this paper, our extensive analysis of the TLS certificates corresponding to the phishing websites from CT log servers reveals that a *template*, which is a regular expression of CN obtained by analyzing the characteristics of certificates believed to have been generated by the same attacker, can be used for the following promising security applications:

- Discovering previously *unknown* phishing websites.
- Understanding the infrastructure used to generate the phishing websites.

As shown later, our analyses reveal the existence of the phishing-website-generation service with many advanced features, such as a mass mailer to send a huge volume of customizable phishing emails, a notification

[1] Note that many of recent HTTPS client implementations use not only the CN field but also the subject-alternative-name (SAN) field when verifying a TLS certificate; a SAN field may contain multiple hostnames associated with the certificate [5]. We empirically found that in practice, the analysis using the CN field did not differ from that using the SAN field. We will discuss the analysis of SAN in a future study.

mechanism, logging, analytics, and dedicated "marketplace," where customers can buy and even sell the stolen credentials. In this paper, we discuss a possible solution against such undesirable use of free service and provide a recommendation to CAs. Furthermore, we present that 24.8% of the detected phishing attacks utilize DDNS or hosting services, and 88.7% use domain names that are not listed on WHOIS database. Therefore, we can expect that our approach outperforms previous phishing-detection approaches in terms of increasing detection coverage and early detection. In order to demonstrate the effectiveness of our proposed approaches in the practical environment, we develop a real-time monitoring system, which aims to discover HTTPSfied phishing websites, utilizing our techniques and issued certificates. Running the system for around a month, we found that it worked effectively; our system discovered 3,009 of new phishing websites. We will showcase the detected phishing attacks where the attackers targeted famous brands and used sophisticated techniques to successfully deceive users. Note that while our approach does not aim to replace the previous defense mechanisms against phishing attacks, our experimental results indicate that our approach is an appealing complement to the conventional countermeasures against threats of phishing attacks.

The remainder of the paper is organized as follows: In Section 2, we provide a background on phishing-detection and TLS certificates. In Section 3, we present our framework that attempts to discover previously *unknown* phishing websites. Section 4 describes the methods and data used in this work. Section 5 demonstrates the statistical result of discovered phishing websites. We also highlight a case study that reveals the infrastructure used for generating groups of phishing websites. Section 6 describes the overview of the real-time monitoring system we developed and demonstrate how it works with one-month long experiments. In Section 7, we provide a recommendation to CAs as well as the limitations of this study. In Section 8, we review related works and compare our results against theirs. We conclude the paper with Section 9.

2 Background: Phishing and Monitoring

Phishing is one of the most widespread cyber threats. Despite its relatively simple attack vector, the damage caused by phishing attacks is significant. The Internet Crime Complaint Center (IC3) reported that the number of victims of phishing attacks including web phishing, vishing (voice), and smishing (SMS) amounted to 26,379 in 2018, with the damage reaching 48.2 M USD [18]. Such attacks attempt to obtain sensitive information, such as

credentials used for online banking, using a spoofed email address and/or a fake website that looks like an authentic one.

To mitigate the threats caused by phishing attacks, several studies have attempted to make use of features that can characterize such attacks (e.g., domain name [38], URL [4], content [46], and email address [41]). However, these approaches have several intrinsic limitations. By monitoring the registration of new domain names, a defender can proactively detect the domain names that are likely used for phishing in the future. However, as some phishing attacks leverage DDNS or hosting service with specific suffix domain names [31, 33], the approach of monitoring newly registered domain names extracted from WHOIS records will miss those cases because they contain only domain names. Similarly, a method analyzing WHOIS records, which attempts to extract domain names with the same contact information listed on the blacklist, cannot detect attacks that use DDNS or hosting services because the granularity of the analysis comprises domain names, not FQDNs. We also note the GDPR has made it infeasible to use WHOIS information because majority of WHOIS gateways have started masking information such as contact information for privacy reasons. Finally, while the phishing detection methods that leverage URLs, web content, or email messages are expected to achieve high detection accuracy [27], most of them are reactive in nature, that is, these approaches cannot detect all attacks in advance.

In contrast to the aforementioned approaches, our approach aims to proactively detect phishing websites by identifying certificates that are likely used for phishing even when the attackers utilize DDNS or hosting services in which a hostname is generated on the existing domain name. The key idea of our approach is to leverage CT logs [14]. CT is a standardized framework that aims to publish the public logs of all the issued TLS certificates. According to [39], the Chromium project started requiring all public TLS certificates issued to support CT since April 2018. Using the Censys dataset [11], we examine the certificates published after April 2018. Of the 635.7 M certificates, 99.3% of them are issued by CAs that have adopted the CT log mechanism. These CAs include freely available popular CAs such as Let's Encrypt [21] and cPanel [8], implying that all the certificates of the customers using the free certificate service are automatically registered to the public CT log servers. The CT provides the way to monitor and audit the TLS certificates issued by the publicly trusted CAs for everyone and enables defenders to efficiently identify mistakenly or maliciously issued certificates.

In this paper, we compiled certificate data obtained from multiple CT log servers. Newly issued certificates should be registered on one or more

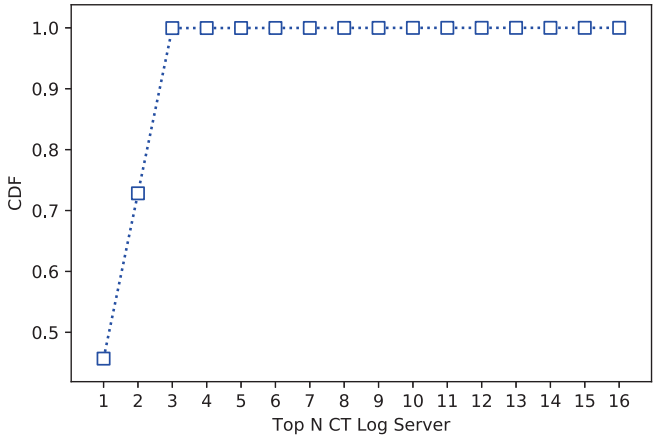

Figure 2 CDF of the number of certificates registered on Top-N CT log Servers. CT log servers are sorted according to certificate record in descending orders.

of the various available CT log servers, such as *argon* operated by Google and *Nimbus* operated by Cloudflare [19]. As CAs arbitrarily choose CT log servers on which to register the newly issued certificate, we need to collect certificate data from multiple CT log servers. We examined CT log servers on which certificates issued by Let's Encrypt or cPanel, both of which tend to be widely used for phishing websites, during the survey period were registered. We found that the certificates were stored on 16 CT log servers. Figure 2 presents CDF of the number of unique certificates registered in Top-N CT log servers. As shown in the figure, when we make use of the data collected from the top CT log server, the coverage is moderate, i.e., 45.6%. However, if we use data collected from the top-3 CT log servers, the coverage becomes 99.9%. As Censys collects certificate data from a number of CT log servers, including the Top-3 servers, we used this database in our study.

3 Framework

In this section, we present our framework for discovering phishing websites. We first provide a high-level overview of the individual methodologies used in our framework. Second, we present the clustering analysis for extracting common characteristics of certificates issued for the phishing websites. Third, we describe a way to extract the intrinsic templates from the clusters. The templates can be used to discover phishing websites that have been

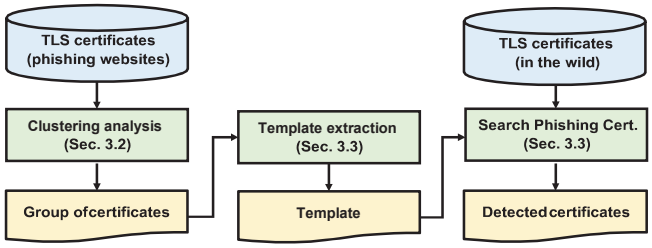

Figure 3 Overview of the framework.

unknown to the security analysts. Finally, we present a method to evaluate the effectiveness of our framework.

3.1 High-level Overview

Figure 3 illustrates a high-level overview of our framework that aims at discovering phishing websites. By analyzing the list of URLs used for phishing attacks, we first collect the TLS certificates from the corresponding websites. Next, we apply the clustering analysis to the certificates and find the group of certificates with similar characteristics (Section 3.2). We subsequently extract the intrinsic templates from the grouped certificates (Section 3.3). By applying the extracted templates to the TLS certificates collected from free CAs, we can discover the certificates that are likely associated with phishing attacks. Finally, we present a method to evaluate the effectiveness of our framework by using a third-party tool (Section 3.4).

3.2 Grouping the Phishing Websites Using Their Certificates

We identify groups of phishing websites that are likely associated with each other. By extracting patterns that are intrinsic among each group, we expect to identify useful characteristics toward discovering phishing websites. To this end, we apply the clustering analysis to the CNs recorded in the certificates.

Before performing clustering analysis, we apply the following data preprocessing. First, we eliminate the substring "`www.`" and top-level domain names (TLDs) such as "`.com`" or "`.io`" from the CN strings, because these substrings are commonly used for all the certificates. Second, after performing filtration, we eliminate the certificates whose CNs are short. The reason for eliminating short CNs is to avoid ambiguities in determining the similarity; for instance, for a CN of short length, such as `apps(.com)`, we will detect many similar CNs, such as `apple(.com)` or `apes(.com)`. However, these

CNs clearly exhibit different semantics, indicating that they are independent domain names. In this work, we empirically derive the threshold as 10.

As a clustering algorithm, we adopt DBSCAN, which enables us to eliminate certificates that are likely attributed to an attacker who does not belong to any of the existing phishing groups. As a function to measure the distance of two given strings (CNs), we leverage *Ratcliff-Obershelp* similarity [34], which is derived by recursively computing the longest common substring (LCS); for two strings "ABC" and "ADBC," the LCS is "BC." First, we find the LCS of the two given strings. We subsequently split each string using the detected LCS as a separator and attempt to find an LCS again for the pairs of strings at both sides. This operation is performed recursively until there are no characters in common between the split strings. *Ratcliff-Obershelp* similarity for two strings (x, y) is defined as $d(x, y) = 2M/T$, where M is the sum of the lengths of LCSs obtained in the above operation between x and y, and $T = x + y$, where s denotes the length of a string s. This similarity is expressed as a ratio between 0 and 1, and then used as a normalized distance for DBSCAN.

Finally, we require useful heuristics to analyze CNs, which have variable lengths and domain name structures. Thus, we introduce a variable m, which denotes the number of dots in a given CN. For instance, $m = 1$ for `ieee.org` and $m = 3$ for `www.cs.example.edu`. The insight behind these heuristics is that CNs generated by the same attacker are expected to use a fixed domain name structure, implying that the number of dots used for these CNs should be the same.

Figure 4 demonstrates an example of the clustering result. If several certificates have similar CNs, we group them as a cluster. If some certificates have CNs dissimilar to any of those in the found clusters, we eliminate such certificates as outliers. DBSCAN has two parameters. We adjust the first parameter ϵ, which controls the similarity between the CNs in a cluster. We set another parameter *minPts*, which is the minimum number of certificates in a cluster, as *minPts* $= 2$.

Table 1 presents an example of clustering results with different values of ϵ. Here, we select a case in which the number of dots is set to $m = 2$. As shown in the figure, when ϵ is 0.25, all the three CNs in the cluster look similar. For the other cases, the CNs in a cluster contain dissimilar CNs. Thus, as illustrated through this example, we empirically adopt the parameter as $\epsilon = 0.25$ for $m = 2$. For other m, following the same procedure, we empirically derive the thresholds as 0.24 $(m = 1)$, 0.3 $(m = 3)$, 0.33 $(m = 4)$, and 0.35 $(m \geq 5)$, respectively.

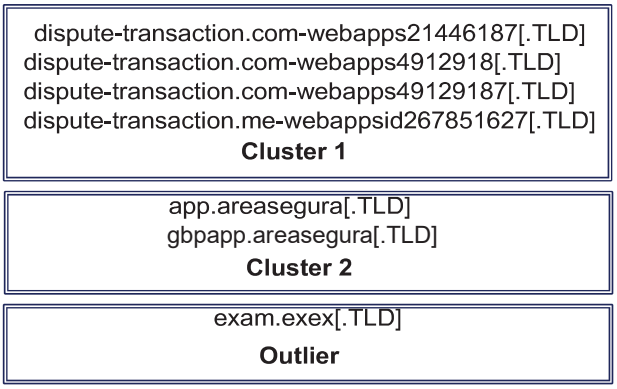

Figure 4 An example of the clustering result.

Table 1 An example of clustering results for CNs with $m = 2$. $\epsilon = 0.25$, 0.30, and 0.35

	Cluster

$\epsilon = 0.25$	`login.portaleprivatimps[.TLD]`
	`secure.portaleprivatimps[.TLD]`
	`accesso.portaleprivatimps[.TLD]`

$\epsilon = 0.30$	`login.portaleprivatimps[.TLD]`
	`secure.portaleprivatimps[.TLD]`
	`accesso.portaleprivatimps[.TLD]`
	`secure.mpsprivati[.]com`

$\epsilon = 0.35$	`login.portaleprivatimps[.TLD]`
	`secure.portaleprivatimps[.TLD]`
	`accesso.portaleprivatimps[.TLD]`
	`secure.mpsprivati[.TLD]`
	`certificazione.areaprivatimps[.TLD]`
	`certificazione.portalemps[.TLD]`
	`certificazione.mpsprivati[.TLD]`

3.3 Extracting Template

Figure 5 illustrates the process of extracting templates from the clusters obtained in Section 3.2. In the case of preprocessing, we eliminate the substring "`www.`" and TLDs in a way similar to what was described in Section 3.2. First, we extract all the substrings common to CNs in a cluster if the substring is three or more characters long. This process is applied to all the strings, which are divided by a dot. Next, we convert the strings other than the common substrings of each CN in the cluster to regular expressions (regexps) and subsequently combine them. When combining regexps, the

common substrings are not modified, and we combine the minimum and maximum lengths of the regular expressions. For example, combining the regexps, [a-z]{3}, [a-z]{5}, and [a-z]{4}, yields the regexp [a-z]{3,5}.

Finally, we verify the genericity of the generated regexps. If a regexp for detecting phishing websites is too generic, it will also detect other legitimate websites, thereby causing large false positives. To test the genericity of a regexp, we adopt *entropy reduction* proposed by Xie [45]. Let e be a regexp. Let $B_e(u)$ be the average number of bits (information entropy) required to encode the representation in binary when using the regexp e to represent a string u. Similarly, let $B(u)$ be the information entropy to represent a string u without using the regexp. Information entropy for a random string can be calculated as $L \log_2 N$, where L is the number of characters that constitute u. N is the number of available characters. It is well known that the information entropy defined in this way is used for measuring the strength of passwords. The information entropy of an original string u and its regexp are calculated as follows.

$$B(u) = L \log 2(A + D) \tag{1}$$

$$B_e(u) = \overline{L_a} \log_2 A + \overline{L_d} \log_2 D + \overline{L_{ad}} \log_2 (A + D) \tag{2}$$

where $\overline{L_a}$, $\overline{L_d}$ and $\overline{L_{ad}}$ are the average number of characters represented by the regular expressions `[a-z]` (alphabet), `[0-9]` (digits), and `[a-z0-9]` (alphabet + digits), respectively. A and D denote the number of characters that can be represented by `[a-z]` and `[0-9]`, with $A = 26$ and $D = 10$.

Next, we introduce a metrics termed as *entropy reduction*, which measures the amount by which a regexp reduces the information entropy to represent a string; i.e., entropy reduction is calculated as $d(e) = B(u)\ B_e(u)$. If a regexp has a small $d(e)$, the information entropy of the regexp e is relatively large, implying the expression is generic. Using a regexp with a large entropy for detecting phishing certificates may result in several false positives owing to its high genericity. Therefore, we extract regexps with d greater than or equal to preset threshold. We empirically derived the threshold as 55. After careful manual inspection, we decided to set a heuristic to eliminate the substrings that are used for the domain names of DDNS or the hosting services. The domain names used by these services are not necessarily limited to use only for phishing websites.

Example: For the purpose of illustration, we present an example of template extraction process. Suppose that we obtain two CNs, `apple-`

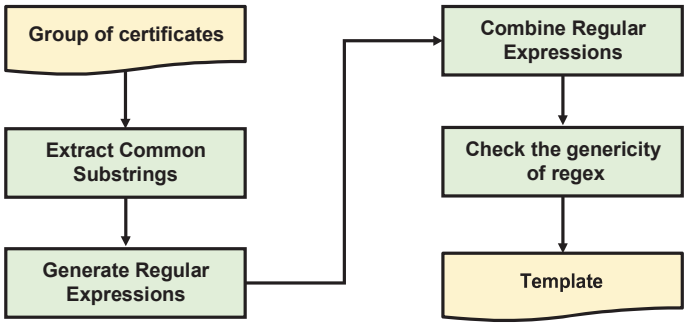

Figure 5 Process of template extraction.

`accountverify123[.TLD]`.[2] The substring common to the two strings is `-accountverify`. The regexps of these CNs are
`[a-z]{5}-accountverify[0-9]{3}[.TLD]` and
`[a-z]{7}-accountverify[0-9]{2}[.TLD]`.
Combining the two regexps yields the following regexp,
`[a-z]{5,7}-accountverify[0-9]{2,3}[.TLD]`.

We now calculate the *entropy reduction*. The number of characters that constitute `-accountverify` is 14, and the average number of characters of the regexp part L_a and L_d are 6 and 2.5, so the total string length L is the sum of these, 22.5. Thus, using Equations (1) and (2), we obtain the following results: $B(u) = 116.3$, $B_e(u) = 36.5$, and $d(e) = 79.8$. Since the entropy reduction exceeds the threshold 55, we adopt the regexp as a template. Using this template, a certificate whose CN is `google-ccountverify37[.TLD]` is detected as the one used by phishing websites. However, although `security-accountverify9[.TLD]` contains the same substring, our approach does not detect it because the number of characters of the regexp is different from those of the template.

3.4 Evaluation Approach

We present a method of evaluating the correctness of the detected phishing websites. A straightforward approach we present to evaluate the aforementioned correctness is to examine the websites we detected. To this end, several existing tools such as web client type honeypot can be utilized. However,

[2]Throughout this study, we replace the top-level domain part with the string [.TLD] to mask the phishing URLs. and `payment-accountverify55[.TLD]` in a cluster.

among the detected websites, there were extremely few active websites that we could access; this is because malicious websites are usually short-lived. Therefore, we leverage VirusTotal [43], which is the most popular online virus scanner service. VirusTotal inspects a target file or URL with over 70 antivirus software and URL/domain blacklisting services.

Our approach attempts to discover potential phishing websites at the time of TLS certificate issuing phase, implying that we can detect phishing websites before they are actually used. As it may take a considerable amount of time before a domain name is posted to VirusTotal, we performed scanning of the discovered domain names after a certain time of period has passed since the collection of TLS certificates. We note that VirusTotal may have missed several phishing domain names, i.e., it should involve false negatives. Likewise, it should also include false positives. Despite these limitations, we believe that analyzing the outputs of the VirusTotal will provide us with promising means to evaluate the effectiveness of our approach at scale – detecting phishing websites at the time of TLS certificate issue.

4 Data

In this study, we leverage the following two certificate datasets: the blacklisted certificates used for creating templates and certificates issued by free CAs for searching phishing websites in the wild. To collect the certificates of the phishing websites, we use the data collected from OpenPhish [30], a publicly available collection of phishing URL feeds. We note that our analysis is not limited to these data and can be applied to other blacklists such as Phishtank. We collect the phishing URLs from October 2018 to January 2019. For each URL we collect, we obtain the corresponding certificates stored at CT log servers by using the Censys database [1] (See Section 2). In total, we extract 2,638 unique certificates. After we apply the data preprocessing described in Section 3.2, we obtain 1,634 unique certificates, which were reported as having been used for the phishing websites.

Table 2 presents the number of certificates we derive for each *m*, which is the number of dots in a CN. We can see that the majority of certificates had CNs with a small number of dots; $m <= 2$ for more than 87% of the

Table 2 Number of phishing certificates for each m

$m = 1$	$m = 2$	$m = 3$	$m = 4$	$m \geq 5$	Total
956	468	110	70	30	1,634

certificates, while a non-negligible number of certificates had CNs with a large number of dots; $m \geq 4$ for more than 5% of the certificates. The high variability of *m* implies that we need to carefully adjust the thresholds for finding the certificates that have visually similar CNs.

We inspect the CAs that issued the certificates used for phishing websites and found that the majority were issued by two free CAs; 852 (50.9%) of them were issued by Let's Encrypt and 714 (42.7%) are issued by cPanel [8]. Thus, HTTPSified phishing websites can be efficiently identified by searching for the CT logs of these CAs. Given this observation, we collect the certificates issued by these two CAs. We collect 38,669,178 certificates issued by these free CAs; 54.9% of these were issued by Let's Encrypt and the remaining 45.1% by cPanel. We use these data as the basis of our analysis shown in Section 5.3, in which we aim to discover phishing websites.

5 Results

In this section, we present the results using the framework described in Section 3 and the data presented in Section 4. We first present the detected clusters of TLS certificates (Section 5.1), and the *templates* extracted from the clusters (Section 5.2). Next, we present the discovered phishing certificates using the templates (Section 5.3) and then validate them (Section 5.4). Finally, we perform an in-depth analysis of the detected phishing certificate through a case study (Section 5.5). We demonstrate that the analysis enables the learning of the infrastructure of the phishing websites.

5.1 Clusters of Certificates

Applying the DBSCAN algorithm to the CNs of the phishing websites resulted in 106 of distinct clusters. These clusters include 341 (20.8%) certificates out of 1,634, which is the number of certificates covered in this study. We note that the remaining 1,293 of certificates were not grouped into any of clusters due to the configuration of the DBSCAN algorithm, i.e., $minPts = 2$. This observation implies that there are varieties of certificates targeting various websites, using different schemes. We conjecture that by increasing the sample size of phishing websites, the clustering process will generate more clusters.

Table 3 shows the clustering result for each *m*. As we have shown in Table 2, majority of the clusters and certificates were concentrated to small *m*, i.e., $m \leq 2$. Figure 6 presents the distribution of the number of certificates

Table 3 Clustering results

	$m = 1$	$m = 2$	$m = 3$	$m = 4$	$m \geq 5$	Total
#clusters	47	39	8	7	5	106
#certificates	124	122	49	33	13	341

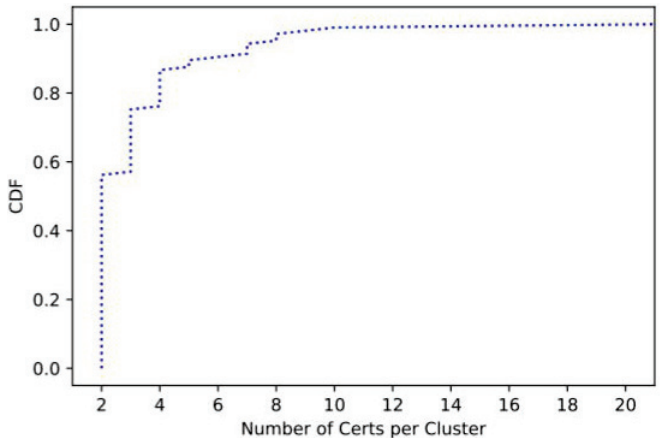

Figure 6 CDF of number of certificates per cluster.

in each cluster. We see that the sizes of each cluster are small in general, while there are non-negligible number of clusters that had a large number of certificates.

5.2 Extracted Templates

Using the 106 clusters, we extracted 69 templates that had the *entropy reduction rate* greater than the pre-determined threshold presented in Section 3. Table 4 presents the examples of domains (CNs) in the two clusters and the extracted templates. We notice that several domains shown in the aforementioned table include those provided by DDNS or hosting services; e.g., `serveirc[.TLD]` and `hoster-test[.TLD]`. The observation shows evidence that attackers leverage DDNS and/or hosting services as the infrastructure of the phishing websites. We found that 10 (14.5%) of the extracted templates contained such domains.

Furthermore, these results suggest some phishing attackers tend to put deceptive strings (for example, “verify-web” and “onedrive” in the table) into

Table 4 Examples of CNs in clusters and the extracted templates

(a) $m = 2$

	Cluster
CN	verify-webapps25476.serveirc[.TLD] verify-webapps72647.serveirc[.TLD] verify-webscrid2678.serveirc[.TLD]
Template	verify-web[a-z0-9]{4,5}.serveirc[.TLD]

(b) $m = 4$

	Cluster
CN	onedrive.liveviewuserauthaspx209hr28jh. srv156794.hoster-test[.TLD] onedrive.liveviewuserauthaspx209hr28jh. srv156816.hoster-test[.TLD] onedrive.liveviewuserauthaspx209hr28jh. srv156797.hoster-test[.TLD] onedrive.liveviewuserauthaspx209hr28jh. srv156796.hoster-test[.TLD]
Template	onedrive.liveviewuserauthaspx209hr28jh. srv156[0-9]{3,3}.hoster-test[.TLD]

all the FQDNs to trick users into believing that the websites are legitimate if they perform similar phishing attacks several times.

5.3 Discovered Phishing Certificates

Using the method described in Section 3, we search for the certificates of the websites that are likely used for phishing attacks. Of the 38.7 M of certificates collected from Let's Encrypt and cPanel, we identified 1,650 certificates that are considered to have been used for phishing. Notably, all the detected certificates had *not* been listed on the OpenPhish blacklist, implying that they were unknown at the phase of certificate issuance. Figure 7 presents the log-log complementary cumulative distribution (CCDF) of the number of detected certificates per cluster. We see that the distribution is *heavy-tailed*; while majority of clusters had a small number of or even zero similar certificates, there are non-negligible number of clusters that had a large number of previously unknown certificates. Specifically, the top-2 clusters had 924 and 395 of the discovered certificates. The existence of clusters with the large number of similar certificates indicates that they likely automate the process of generating the phishing websites.

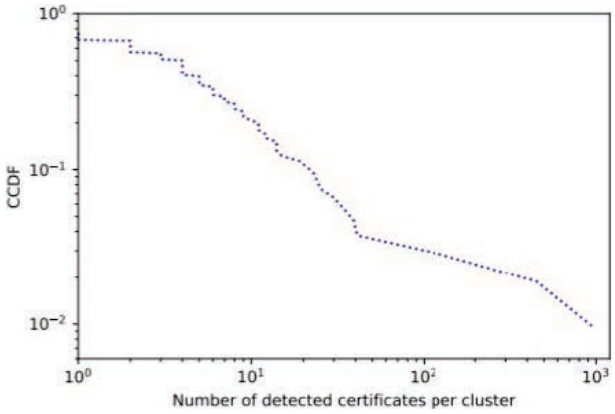

Figure 7 Log-log CCDF of the number of detected certificates per cluster.

The cluster with 395 certificates showed that all made use of hosting services because all the CNs had the domain name suffix of `zap-hosting[.TLD]`. We also detected 15 additional certificates with CNs provided by other DDNS or hosting services; therefore, 410 (24.8%) of the detected attacks used DDNS or hosting services. Furthermore, 88.7% of the CNs of the discovered certificates did not meet the criteria to be listed on WHOIS, i.e., the CNs have one or more labels in addition to effective TLD. Conventional domain-registration-based approaches fail in detecting such attacks as WHOIS records do not contain all FQDNs. In Section 5.5, we perform an in-depth analysis using the cluster containing 924 of certificates.

5.4 Evaluation

We obtained 1,049 unique CNs after eliminating the duplicated CNs; we obtained 1,049 candidates of the phishing websites. Note that some of the discovered certificates had the same CNs because they were issued several times during the survey period. First, the verification with VirusTotal revealed that for 90.8% of the websites we detected, at least one antivirus checker raised alarms while for 72.5% of the websites we detected, at least two antivirus checkers raised alarms. As mentioned in Section 7.1, the possibility remains that the antivirus checkers in VirusTotal overlook malicious ones because malicious websites are usually short-lived. Therefore, some of 9.2% of CNs may include potentially malicious ones, and we cannot determine that they are false positives of our approach. The aforementioned results clearly demonstrate that our

framework was able to find a lot of potentially malicious websites and the majority of them were identified as obviously malicious by third-parties. For example, a template `allegro.pl-login.form-a[a-z0-9]{0,12}.[a-z0-9]{0,6}[.TLD]` detected 10 websites, eight of which were detected as malicious with VirusTotal. Through a careful manual inspection of the corresponding certificates, we conjecture that the other two were also generated by the same attacker. At least, we were unable to identify evidence that the remaining two CNs were used for legitimate services.

However, given actual security operations, our approach should be used to account for potential false positives. A practical usage is a prefilter to extract highly suspicious ones from large certificates and send them for manual inspection.

5.5 In-Depth Analysis

We present an in-depth analysis of the detected phishing certificates using a template that yielded the largest number of phishing websites. The template is `[a-z]{6,8}.runescape.com-[a-z]{1,8}[.TLD]`. Using this template, we detected the following two patterns of domain names:
`secure.runescape.com-[a-z]{1,8}[.TLD]` and
`services.runescape.com-[a-z]{1,8}[.TLD]`.
A simple domain name analysis revealed that these domain names target *Runescape*, which is a massively popular multiplayer online role-playing game (MMORPG). We demonstrate examples of CNs of the certificates in a cluster targeting *Runescape* in the left column of Table 5. We note that 863 (93.4%) out of 924 certificates were issued by the same CA, Let's Encrypt. We also note that in this case, the combinations of TLDs and second-level domains are often different.

Our manual inspection on the discovered certificates revealed that these certificates are generated by a phishing website generation service, which is sold by a rogue company. Although searching the web will reveal such companies in the wild, we refrain from specifying the name of company for the ethical reason. In order to confirm whether the service actually generates certificates with CNs that match the templates we identified, we subscribe to their service to check the certificates in the service. We note that we do not use any of the services provided by the company. As shown in the right column of Table 5, the certificates generated by the kit match to the templates we constructed.

Table 5 Examples of CNs for the Discovered certificates for a cluster and used in a phishing website generation kit targeting *Runescape*

Cluster	Phishing kit
services.runescape.com-an[.TLD]	services.runescape.com-rv[.TLD]
secure.runescape.com-mq[.TLD]	secure.runescape.com-ao[.TLD]
secure.runescape.com-g[.TLD]	secure.runescape.com-vo[.TLD]
secure.runescape.com-l[.TLD]	secure.runescape.com-rs[.TLD]

In addition, we found that the phishing website generation service provides many advanced features to help an attacker perform the phishing attack efficiently; e.g., mass mailer to send a huge volume of customizable phishing email, notification mechanism, logging, analytics, and dedicated "marketplace" where customers can buy and/or even sell the stolen credentials. We note that although previous studies [29, 32] have mentioned the existence of the phishing website generation service, these studies did not provide the deep insight into the ecosystem of the service. Given these results and observations, we may conclude that the analysis of certificates can reveal the infrastructure and ecosystem of the phishing attack.

6 Building A Real-Time Detection System

In the previous section, we showed that a lot of HTTPSified phishing websites were successfully discovered by our approach. However, since we conducted our analyses after a few months had passed since the certificates were issued, many of the websites had become inactive. Therefore, we were not able to conduct a detailed investigation on the detected websites, e.g., the content of the websites. To overcome such shortcomings and evaluate the practicality of our approach, we developed a real-time monitoring system that incorporates the proposed methods and attempts to access the detected websites as soon as we detect them. The system utilizes Puppeteer [6], a Node library which provides a high-level API to control Chrome or Chromium. In the following, we will describe the system overview and demonstrate its effectiveness through a month of experiment.

6.1 System Overview

We describe how our real-time monitoring system works in the practical environment. The system consists of two stages: the preparation stage and detection stage.

In the preparation stage, we retrieve the latest phishing URLs, collect the corresponding certificates, and perform clustering on the CNs with DBSCAN. While majority these processes are implemented as described in Sections 3.2 and 3.3, we changed several processes so that the system runs without involving human intervention. First, we change the process of manually adjusting the DBSCAN parameter ϵ. Specifically, we vary the parameter ϵ from 0.05 to 0.3 in increments of 0.05, which are derived based on our experience in Section 3.2. Second, we use phishing URLs stored in blacklists within the last 90 days of the detection date to create templates. The reason for adopting the time period of 90 days is as follows; after clustering on certificates issued from March to June 2020, we found that 90 percent of certificates in the same cluster were issued within 60 days, with a maximum of 89 days. Third, in order to increase the number of the extracted templates and to detect a large number of malicious websites, the real-time monitoring system retrieves phishing URLs from PhishTank in addition to OpenPhish. Finally, we set the threshold for entropy reduction value at 50; we use templates with entropy reduction values above 50. The reason behind this choice is that we have found that templates with entropy reduction below that value would end up detecting numerous certificates, most of which are less likely to be malicious.

In the detection stage, we first crawl the issued TLS certificates in a real-time manner. We then extract the ones with CNs that match one of the templates built in the preparation stage. To collect TLS certificates, we utilize *CertStream-Server*,[3] which compile the certificates from various CT log servers. Upon detection, we access the website and attempt to obtain their content such as screenshot for the subsequent analysis.

6.2 Performance of the System

We ran the real-time monitoring system roughly for a month, i.e., from 25 August to 30 September 2020. During the period, the system generated an average of 2,079 templates per day. Using those templates, we were able to discover an average of 88.5 malicious websites per day. Those malicious websites were the ones detected by at least one antivirus scanners registered in VirusTotal, amounting to a total of 3,009 websites during the survey period. This result clearly demonstrates that our system can detect the HTTPSified

[3] https://github.com/CaliDog/certstream-server-python.

phishing/malicious websites immediately after their certificates were issued. We also note that our approach worked well in a practical setup.

Once the real-time monitoring system detected the suspicious certificates, it accessed the corresponding websites using a headless browser. In order to understand the content of the malicious website at the time of the certificate issuance, we randomly sampled 200 websites from the 3,009 websites identified as malicious by VirusTotal and manually inspected the screenshots of the websites we found. We classify the websites into the following four categories: (1) legitimate-looking website without a login form, (2) legitimate-looking website with a login form, (3) website indicating it is not ready, such as the default page of website hosting services, and (4) website not retrieved correctly due to various reasons such as errors. The reason for focusing on the presence of a login form placed on websites is that most phishing sites attempt to harvest users' credentials through login forms.

Table 6 shows the result. First, while 98 (49%) of websites did not have a login form, 17 (8.5%) websites had a login form, which implies the presence of the phishing trigger. Of the 17 websites, we found that 13 of the websites were highly likely to be phishing websites because they targeted famous brands such as Instagram and the domain names are different from the legitimate ones. On the other hand, we did not find any known brands that are likely targeted by the remaining four websites, so we could not make a concrete conclusion on those websites. Still, we need to pay attention to these websites although the accuracy of the VirusTotal is not perfect. Second, the most common target brand names among the websites with login forms was Instagram, followed by Rakuten Bank, one of the major Japanese banks. We will present the in-depth analysis on how these phishing websites attempt to steal users' credentials in Section 6.3. Third, 8.5% of the websites were not yet ready. We conjecture that attackers may have issued their certificates before publishing the websites. Therefore, we need to take care of the websites, which could be activated in future. Finally, we were not able to retrieve the contents for the remaining 34% of the websites. Again,

Table 6 Classification results for randomly sampled websites

Category	Number
(1) No Login Form	98 (49.0%)
(2) Login Form	17 (8.5%)
(3) Not Ready	17 (8.5%)
(4) Not Retrieved	68 (34.0%)

we need to pay attention to those websites as they could turn into phishing sites in future.

In summary, our system discovered active phishing websites displaying live phishing content, which helps in building a prompt alarming for the phishing attack. Our system also detected suspicious websites that are likely in the preparation phase. Keeping track of the activities for those potential malicious websites is an effective way to detect phishing website immediately after they are activated.

6.3 Case Studies

In this section, we present examples of large-scale phishing attacks detected by our system. Attackers targeted famous brands, by generating a large number of domains and corresponding certificates.

Paypal

The system detected phishing websites targeting *Paypal*, which is one of the major online payments services operating globally. The template generated to detect the phishing websites was
`payp[a-z]{0,1}lticket[0-9]{5,8}[.TLD]`.
This template discovered 51 of websites. Of those websites, 36 were detected as malicious by VirusTotal. Our manual inspection revealed that 9 (17.6%) of the detected sites imitated the login screen of Paypal, while the remaining 42 (82.3%) sites displayed "Index of" page. Figure 8 presents a screenshot of a detected phishing site. As all the websites were operating under the domain names with a specific pattern, it is highly likely that they were all generated by a same group of attackers. Thus, we should carefully monitor the inactive websites detected by our system.

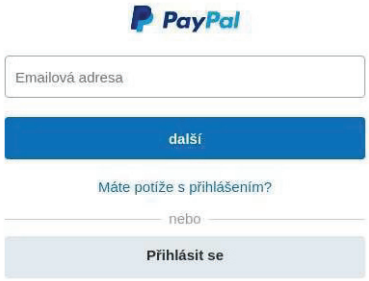

Figure 8 Screenshot of a phishing website targeting Paypal.

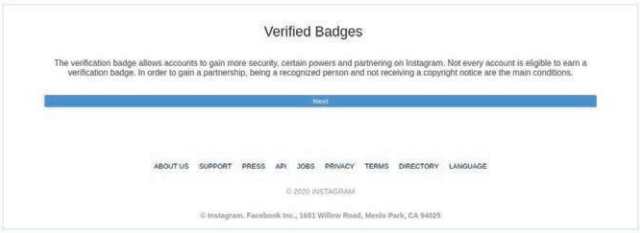

Figure 9 A screenshot of the phishing websites targeting Instagram. The site utilizes verified badge.

Instagram

Our system detected several sophisticated phishing websites, including the ones targeting *Instagram*. We discovered two groups of phishing attacks: those that use "Verified Badge" and those using the typosquatting domains.

Phishing attacks leveraging the verified badge urged users to get a verified badge by entering their credentials. In general, a verified badge appears next to the account name on a SNS website. For SNS users, having an a verified badge means that they can show off their high reputation/influence on the social media. An attacker can increase the success rate of an attack by taking advantage of the user's desire for approval; i.e., they use verified badge as an attractive bait. Figure 9 presents an example of phishing websites leveraging verified badge.

Our system automatically detected phishing domain names using typosquatting; i.e., they include the string "xnstagram," where 'x' is 'l', not 'I'. The templates built by our system include
`[a-z]{1,1}nstagram[a-z]{13,24}[.TLD]` and
`lnstagramhelp[a-z]{3,7}[.TLD]`.
There are 153 websites with that used the templates. 91 of them were detected as malicious by one or more antivirus software registered in VirusTotal. Through the careful manual inspection, we found 49 (32%) websites had login form. Interestingly, of these 49 websites, 27 websites showed copyright infringement to threaten users; i.e., those sites used the same tactics, implying the sites were operated by a same group. There were several designs of those websites, we show an example of the screenshot of the phishing websites in Figure 10. The website states that the users who open the websites violated a copyright, so have to log in and submit a feedback if they do not think they infringed on copyright.

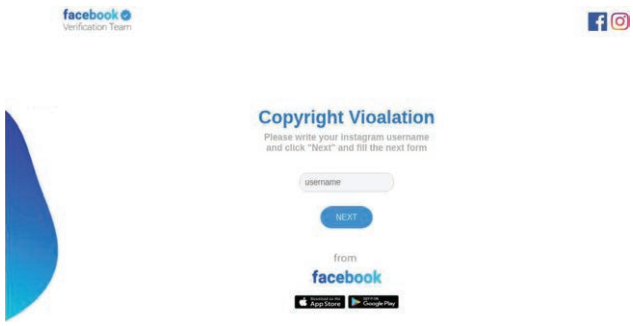

Figure 10 An example of the screenshot of the phishing websites that use typosquatting and target Instagram.

Rakuten Bank

Finally, we present phishing attacks targeting Rakuten *Bank*, one of the major banks in Japan. The real-time monitoring system detected 44 websites likely targeting *Rakuten Bank*. Of those, 32 (72.7%) were detected as malicious by VirusTotal. After the manual inspection, we found 36 (81.8%) of them had login forms masquerading as Rakuten Bank. For the remaining sites, we were not able to retrieve the content due to HTTP errors, etc. Figure 11 presents an example of the screen-shot. The phishing site was indistinguishable from the real site from its appearance. Some examples of the templates built by our system were as follows:

```
[a-z]{5,6}.rakuten.[a-z]{2,2}.[a-z]{2,2}.raku[a-z]
{17,20}[.TLD],
```

and

```
rakuten.[a-z]{2,2}.[a-z]{2,2}.[a-z]{0,1}akuten[a-z]
{6,17}[.TLD].
```

We performed additional analysis using spam emails, which were collected by our spam trap server. We discovered spam emails with URLs matching the template, i.e., "rakuten.co.jp.rakutenycocojp[.TLD]," which matched to the second template shown above. The subject lines of those emails were all "Rakuten Emergency Notification Login Alert", which prompted the users to log in with the following text: "Due to the unusual activity detected in your account, we have suspended your orders and your Rakuten account. You can remove the account suspension by logging into your account and following the instructions on the screen." In addition, since each user's username is listed at the beginning of those emails, they appear to be highly authentic.

Figure 11 An example of the screenshot of the phishing websites targeting Rakuten Bank.

Finally, we study the phishing attacker's behaviour from a chronological perspective. The certificates corresponding to the domain names above were issued on the two consecutive days, i.e., August 31 and September 1, 2020. In the case where an attacker used the certificates issued on September 1, the interval between the certificate issuance and sending of the spam emails was very short; on average, they have sent out the spam emails 3 hours and 1 minute after the issuance, and the shortest interval was only 4 minutes and 21 seconds. Though we do not have data that shows which certificate was used in the phishing attack, our findings clearly demonstrates that some phishing attackers carry out their attacks shortly after the certificates are issued. This observation implies that it is necessary to monitor suspicious websites immediately after the certificate issuance and our certificate-based detection system worked effectively in achieving that goal.

7 Discussion

In this section, we discuss the limitations of this work and the undesirable use of free services such as free CAs, and provide a recommendation to the CAs.

7.1 Limitations

7.1.1 Threats to validity

While the majority of the websites discovered by our method were flagged by VirusTotal, there were other several websites not flagged. Careful manual inspection revealed that among the undetected websites, there were a

few false positives in our approach. If a template contains universal words such as `service` and `communication`, which are often used for benign websites, it becomes difficult to distinguish between a benign website and a malicious website. Here is an example. While template `officespace{1,2}[a-z][.TLD]` detected 31 websites, 12 of them were flagged by VirusTotal. Since the words "office" and "space" are both frequently used, the template detected several benign as well as malicious websites. A promising method of preventing this phenomenon is to collect and list universal words in advance, and reduce the value of the *entropy reduction* accordingly if such universal words are included in the template. This decreases the number of false positives because such highly generic templates with a low value of the *entropy reduction* will be eliminated.

Another false-positive case may occur under the following two conditions: (1) several legitimate websites with similar domain names are mistakenly included in the blacklist and the template is created for them, and (2) certificates of other legitimate websites with domain names to be matched by the template are issued. We note that these events rarely occur and we did not find this case in this work.

7.1.2 Wild card certificate

Even if an attacker generates multiple CNs and performs phishing, it is difficult to investigate them using our method if the attacker uses wild card certificates. However, using such certificates can be a disadvantage for attackers because if one of the hosts they use is blacklisted, the other hosts will probably be disabled by antivirus software, Google Safe Browsing, etc. Hence, if attackers intend to generate many similar CNs and perform phishing, they would benefit by changing the domain part (as does the phishing kit discovered in this study) and issue certificates accordingly.

7.2 Detection Evasion

An attacker could efficiently perform phishing attacks under the constraints on time and strings that are effective for creating phishy URLs. For the time constraint, a large amount of the certificates for a phishing website are issued for a short period. For the string constraint, the URL of phishing website must include deceptive strings to make victims believe that the prepared website is genuine. The examples of deceptive strings are specific brand names, generic terms (service, account, etc.), and actions (login, pay, registration, etc.). Our analysis works based on these attacker constraints. The attacker ignoring the

above-mentioned constraints may fail to deliver efficient phishing attacks (e.g., issuing certificates over a long period and using fully randomized domain names). This is why we especially focused on phishing websites among malicious activities. A possible method, especially against a string constraint, to complicate our analysis is to use "leetspeak."

Suppose that there is a benign website with the domain name `login-account-service[.TLD]`, and an attacker tries to impersonate the website and issues three certificates, the CNs of which are `login-account-serv1ce[.TLD]`, `l0g1n-acc0unt-serv1c3[.TLD]`, and `l0gin-4ccount-s3rv1ce[.TLD]`. In this case, the following two problems may occur: (1) the created template is too generic to use and (2) we detect the benign website by the template. As for the first problem, even if an attacker uses leetspeak, those certificates will surely be incorporated into the same cluster by DBSCAN clustering because of their high degrees of similarity; the average *Ratcliff–Obershelp* similarity between them is 0.778. However, the template we obtained from these CNs (`[a-z0-9]{10}unt-s[a-z0-9]{6}[.TLD]`) after applying the proposed method is too generic for detecting phishing sites as the value of the entropy reduction is 36.19, which is much lower than the preset threshold. As a countermeasure, we can decode leetspeak by using a tool, such as Universal Leet Converter [35], during template extraction and phishing detection. In this example, we obtain `login-account-service[.TLD]` as a template by decoding the leetspeaks. On the other hand, this template generation may create false positives because this generated template simply matches the legitimate one. To eliminate such false positives, we can use such template for matching only certificates with CN including leetspeaks. Incorporating these improvements into the detection system is left for future work.

7.3 Recommendations

The number of phishing websites with HTTPS have been increasing. Some countermeasures are essential considering that most use free certificates issued by Let's Encrypt or cPanel. Sectigo Ltd. [36], which operates cPanel, specifies in its certificate practice statement (CPS) that if a certificate is found to have been used for illegal purposes, such as phishing and malware, they will revoke it within 7 days [37]. On the other hand, Let's Encrypt terminated efforts to confirm websites were not malicious using Google Safe Browsing, because they consider that domain validation (DV) certificates are only intended to secure communications between the client and server, not to

ensure the safety of the website. However, as shown in Section 5.5, 93.4% of the certificates of *similar CNs* that are considered to be generated by the phishing kit targeting *Runescape* were issued by Let's Encrypt, and CNs with those specific patterns can be found easily using our method. Considering these findings, the CA should identify such CNs using the approach presented in this study and revoke the certificates.

8 Related Work

In this section, we review related works and discuss the comparison between them and our research.

8.1 Detecting URLs and Contents

Multiple studies have shown that features extracted from URLs and content can be used as clear indicators to detect phishing websites. The features are created through the following expert knowledge: lexical anomalies in URLs (e.g., blacklisted words, hyphens used instead of dots, and brand/service names in the URL path) [20, 42], IP address used as the domain name in the URL [13, 42], many dots in the URL [13, 20], inconsistent brand names/logos (e.g., the *brand-X* name not on a *brand-X* domain name) [24], similarity among contents [7, 23], and so on. CANTINA and CANTINA+ are complementary approaches using the above heuristics. They examine the content to determine whether the website is legitimate or not by using search results of important terms in the content extracted by the term frequency-inverse document frequency (TF-IDF) algorithm [44, 46]. These approaches based on URL and content features successfully detect phishing websites. However, their limitation is that they detect only visited websites or listed websites (e.g., the URLs in delivered emails). In other words, detecting never-accessed/-listed websites is beyond the scope of these approaches. Our method does not face such limitations, because it relies on the certificates that anyone can comprehensively list via the certificate transparency (CT) log server [14] or a repository of Internet scannings such as Censys [11].

8.2 Detecting Certificates

Given the rapid increase in the number of HTTPSified phishing websites, there have been some attempts to detect phishing websites using the certificates [9, 10, 40]. Torroledo et al. [40] and Dong et al. [9] proposed methods

for identifying malicious use of certificates based on the features included in the fields of the certificate. However, Drury and Meyer mentioned some fields are very similar (or the same) for all certificates issued by the same issuer and concluded that it is generally difficult to differentiate certificates of phishing websites from those of benign websites if the certificates of both phishing and benign websites are provided by the same issuer [10]. While previous studies make use of features from the certificates to identify differences between the certificates of benign and phishing websites, our study reveals that we can make valued use of the information obtained from the certificates; that is, we can discover previously unknown phishing websites, systematically find targeted websites, and understand the infrastructure used for generating such phishing websites.

8.3 Phishing Kit and Evasion

Criminals create phishing kits, which are packages used to deploy a phishing website on a web server. They sell phishing kits in underground marketplaces and accept custom requests for kit creation [3, 29]. Phishing kits include server-side and client-side evasion techniques using server directives (`.htaccess` files), server-side scripts, and JavaScript to interfere with detection by the security community [28, 29]. The evasion is carried out based on a client IP address, referrer, and user agent. If the accessing client environment is detected by evasion techniques, the content would not be available. We emphasize that in most cases, our analysis is not affected by such evasion techniques because our approach leverages the characteristics of TLS certificates, which can be collected from the publicly available CT logs.

9 Conclusion

This work focuses on the fact that phishing websites have started adopting HTTPS; this could expose their intrinsic features that could be used to detect them in a systematic manner. Compared to conventional phishing-detection approaches, certificate-based approach has the following advantages: it allows a defender (1) to comprehensively monitor all HTTPSified websites through the issued certificates, even if the attacker utilizes DDNS or hosting services, and (2) to detect phishing websites before they are published on the Internet.

Although some previous studies have reported that distinguishing a benign website from a malicious website by using certificate information

alone is difficult, we established a framework to discover unknown phishing websites through a *template* extracted using attackers' bulk registration and CN in certificates. We demonstrated that the template can be applied not only to discover phishing websites with low false positives but also to understand the infrastructure used to generate the phishing websites, e.g., phishing-website-generation kit. We also demonstrated that our proposed approach can find several types of phishing websites that existing approaches cannot detect in nature because these websites make use of DDNS or hosting services, which are not listed in domain name-based database such as WHOIS. Furthermore, using the methodologies and findings obtained so far, we have developed a real-time monitoring system. Through a month-long experiment, we have demonstrated that the system can detect HTTPSified phishing websites in a real-time manner, and that it can efficiently detect websites that are not currently active as phishing sites, but require special attention as they could turn into phishing sites in future. The templates automatically generated by the system are effective for efficiently detecting phishing sites, and can be applied to firewall rules. We believe that our approach contributes to complement the lack of the various existing phishing-detection techniques and sheds new light on the *footprints* of TLS certificates as a key to understanding the origin of threats.

References

[1] Censys. https://censys.io/.
[2] APWG. Phishing activity trends report 3rd quarter 2019. https://docs.apwg.org/reports/apwg_trends_report_q4_2019.pdf.
[3] Dominik Birk, Sebastian Gajek, Felix Grobert, and Ahmad-Reza Sadeghi. Phishing phishers – observing and tracing organized cybercrime. In *Proc. of ICIMP '07*.
[4] Aaron Blum, Brad Wardman, Thamar Solorio, and Gary Warner. Lexical feature based phishing url detection using online learning. In *Proc. of ACM AISec*, pages 54–60. ACM, 2010.
[5] Sharon Boeyen, Stefan Santesson, Tim Polk, Russ Housley, Stephen Farrell, and Dave Cooper. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280, May 2008.
[6] Chrome DevTools team. Puppeteer. https://github.com/puppeteer/puppeteer.
[7] Igino Corona, Battista Biggio, Matteo Contini, Luca Piras, Roberto Corda, Mauro Mereu, Guido Mureddu, Davide Ariu, and Fabio Roli.

Deltaphish: Detecting phishing webpages in compromised websites. In *Proc. of ESORICS*, 2017.

[8] cPanel. https://cpanel.net/.

[9] Zheng Dong, Apu Kapadia, Jim Blythe, and L. Jean Camp. Beyond the lock icon: Real-time detection of phishing websites using public key certificates. In *Proc. of APWG Symposium eCrime*, 2015.

[10] Vincent Drury and Ulrike Meyer. Certified phishing: Taking a look at public key certificates of phishing websites. In *Proc. of USENIX Symposium SOUPS*, 2019.

[11] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. A search engine backed by Internet-wide scanning. In *ACM CCS*, 2015.

[12] Adrienne Porter Felt, Richard Barnes, April King, Chris Palmer, Chris Bentzel, and Parisa Tabriz. Measuring HTTPS adoption on the web. In *26th USENIX Security Symposium*, 2017.

[13] Ian Fette, Norman Sadeh, and Anthony Tomasic. Learning to detect phishing emails. In *Procc of International Conference WWW*, 2007.

[14] Google. Certificate transparency. https://www.certificate-transparency.org.

[15] Google. Https encryption on the web. https://transparencyreport.google.com/https/overview?hl=en.

[16] Google. A secure web is here to stay. https://security.googleblog.com/2018/02/a-secure-web-is-here-to-stay.html.

[17] Google. Webmaster central blog, https as a ranking signal. https://webmasters.googleblog.com/2014/08/https-as-ranking-signal.html.

[18] Internet Crime Complaint Center. 2018 internet crime report. https://pdf.ic3.gov/2018_IC3Report.pdf.

[19] B. Laurie, A. Langley, and E. Kasper. Certificate transparency. RFC 6962, RFC Editor, June 2013.

[20] Anh Le, Athina Markopoulou, and Michalis Faloutsos. Phishdef: Url names say it all. In *Proceedings of 2011 IEEE INFOCOM*, 2011.

[21] Let's Encrypt. https://letsencrypt.org/.

[22] X. Li, G. Geng, Z. Yan, Y. Chen, and X. Lee. Phishing detection based on newly registered domains. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 3685–3692, 2016.

[23] J. Mao, W. Tian, P. Li, T. Wei, and Z. Liang. Phishing-alarm: Robust and efficient phishing detection via page component similarity. *IEEE Access*, 5:17020–17030, 2017.

[24] Samuel Marchal, Kalle Saari, Nidhi Singh, and N. Asokan. Know your phish: Novel techniques for detecting phishing sites and their targets. In *Proc. of IEEE ICDCS 2016*.

[25] Mozilla. Communicating the dangers of non-secure http. https://blog.mozilla.org/security/2017/01/20/communicating-the-dangers-of-non-secure-http/.

[26] David Naylor, Alessandro Finamore, Ilias Leontiadis, Yan Grunenberger, Marco Mellia, Maurizio Munafò, Konstantina Papagiannaki, and Peter Steenkiste. the cost of the "s" in https.

[27] L. A. T. Nguyen, B. L. To, H. K. Nguyen, and M. H. Nguyen. A novel approach for phishing detection using url-based heuristic. In *2014 International Conference on Computing, Management and Telecommunications (ComManTel)*, pages 298–303, April 2014.

[28] A. Oest, Y. Safaei, A. Doupé, G. Ahn, B. Wardman, and K. Tyers. Phishfarm: A scalable framework for measuring the effectiveness of evasion techniques against browser phishing blacklists. In *2019 IEEE SP*.

[29] Adam Oest, Yeganeh Safei, Adam Doupé, Gail-Joon Ahn, Brad Wardman, and Gary Warner. Inside a phisher's mind: Understanding the anti-phishing ecosystem through phishing kit analysis. In *APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–12, 2018.

[30] OpenPhish. Openphish faq. https://openphish.com/faq.html.

[31] Peng Peng, Chao Xu, Luke Quinn, Hang Hu, Bimal Viswanath, and Gang Wang. What happens after you leak your password: Understanding credential sharing on phishing sites. In *Proc. of ACM CCS*, 2019.

[32] Peng Peng, Chao Xu, Luke Quinn, Hang Hu, Bimal Viswanath, and Gang Wang. What happens after you leak your password: Understanding credential sharing on phishing sites. pages 181–192, 07 2019.

[33] PHISHLABS. 2019 phishing trends and intelligence report the growing social engineering threat. https://info.phishlabs.com/hubfs/2019PTIReport/2019PhishingTrendsandIntelligenceReport.pdf.

[34] John W Ratcliff and David E Metzener. Pattern-matching-the gestalt approach. *Dr Dobbs Journal*, 13(7):46, 1988.

[35] Robert Ecker. Universal leet (l337, l33t, 1337) converter. http://www.robertecker.com/hp/research/leet-converter.php.

[36] Sectigo. https://sectigo.com/.

[37] Sectigo. Sectigo certification practice statement (version 5.1.1). https://sectigo.com/uploads/files/Sectigo-CPS-v5.1.1.pdf.

[38] Hossein Shirazi, Bruhadeshwar Bezawada, and Indrakshi Ray. Kn0w thy doma1n name: Unbiased phishing detection using domain name based features. In *Proc. of ACM SACMAT*.

[39] The Chromium Project. Chromium Certificate Transparency Policy. ht tps://github.com/chromium/ct-policy.
[40] Ivan Torroledo, Luis David Camacho, and Alejandro Correa Bahnsen. Hunting malicious tls certificates with deep neural networks. In *Proc. of ACM AIsec*, October 2018.
[41] Amber van der Heijden and Luca Allodi. Cognitive triaging of phishing attacks. In *28th USENIX Security Symposium (USENIX Security 19)*, Santa Clara, CA, August 2019. USENIX Association.
[42] Rakesh Verma and Keith Dyer. On the character of phishing urls: Accurate and robust statistical learning classifiers. In *Proc. of ACM CODASPY 2015*.
[43] VirusTotal. Virustotal. https://www.virustotal.com/.
[44] Guang Xiang, Jason Hong, Carolyn P Rose, and Lorrie Cranor. Cantina+: A feature-rich machine learning framework for detecting phishing web sites. *ACM TISSEC*, 2011.
[45] Yinglian Xie, Fang Yu, Kannan Achan, Rina Panigrahy, Geoff Hulten, and Ivan Osipkov. Spamming botnets: Signatures and characteristics. *ACM SIGCOMM CCR*, 38:171–182, 01 2008.
[46] Yue Zhang, Jason I. Hong, and Lorrie F. Cranor. Cantina: A content-based approach to detecting phishing web sites. In *Proceedings of the 16th International Conference on World Wide Web (WWW)*, 2007.

Biographies

Yuji Sakurai received a B.E degree in computer science and communication from Waseda University, affiliated with the Network Security Lab. His interests include Network Security, and Internet Measurement.

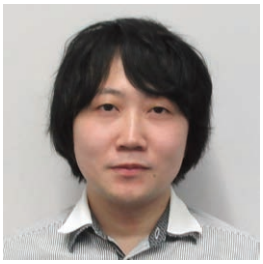

Takuya Watanabe received B.E. and M.E. degrees in computer science and engineering, and Ph.D. degree in engineering from the Waseda University, in 2014, 2016, and 2020, respectively. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2016, he has been engaged in research of consumer security and privacy. He is now with the Cyber Security Project of NTT Secure Platform Laboratories.

Tetsuya Okuda received B.E. and M.E. degrees in aeronautics and astronautics from the Tokyo University in 2009, 2011, respectively. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2011, he has been engaged in research and engineering of data security and web services. He is now with the Data Security Project of NTT Secure Platform Laboratories.

Mitsuaki Akiyama received his M.E. and Ph.D. degrees in information science from Nara Institute of Science and Technology, Japan in 2007 and 2013. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2007, he has been engaged in research and development on cybersecurity. He is currently a Senior Distinguished Researcher with the Cyber Security Project of NTT Secure Platform Laboratories. His research interests include cybersecurity measurement, offensive security, and usable security and privacy. He is a member of the IEEE, IPSJ, and IEICE.

Tatsuya Mori is currently a professor at Waseda University, Tokyo, Japan. He received B.E. and M.E. degrees in applied physics, and Ph.D. degree in information science from the Waseda University, in 1997, 1999 and 2005, respectively. He joined NTT lab in 1999. Since then, he has been engaged in the research of measurement and analysis of networks and cyber security. From Mar 2007 to Mar 2008, he was a visiting researcher at the University of Wisconsin-Madison. He received Telecom System Technology Award from TAF in 2010 and Best Paper Awards from IEICE and IEEE/ACM COMSNETS in 2009 and 2010, respectively. Dr. Mori is a member of ACM, IEEE, IEICE, IPSJ, and USENIX.