

Can a TLS certificate be phishy?

Hageman, Kaspar; Kidmose, Egon; Hansen, René Rydhof; Pedersen, Jens Myrup

Published in:

Proceedings of the 18th International Conference on Security and Cryptography, SECRYPT 2021

DOI (link to publication from Publisher):

[10.5220/0010516600380049](https://doi.org/10.5220/0010516600380049)

Creative Commons License

CC BY-NC-ND 4.0

Publication date:

2021

Document Version

Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Hageman, K., Kidmose, E., Hansen, R. R., & Pedersen, J. M. (2021). Can a TLS certificate be phishy? In S. D. C. di Vimercati, & P. Samarati (Eds.), *Proceedings of the 18th International Conference on Security and Cryptography, SECRYPT 2021* (pp. 38-49). SciTePress. <https://doi.org/10.5220/0010516600380049>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Can a TLS Certificate Be Phishy?

Kaspar Hageman¹, Egon Kidmose¹, René Rydhof Hansen² and Jens Myrup Pedersen¹

¹*Department of Electronic System, Aalborg University, Denmark*

²*Department of Computer Science, Aalborg University, Denmark*

Keywords: Phishing, Digital Certificate, Certificate Transparency, TLS.

Abstract: This paper investigates the potential of using digital certificates for the detection of phishing domains. This is motivated by phishing domains that have started to abuse the (erroneous) trust of the public in browser padlock symbols, and by the large-scale adoption of the Certificate Transparency (CT) framework. This publicly accessible evidence trail of Transport Layer Security (TLS) certificates has made the TLS landscape more transparent than ever. By comparing samples of phishing, popular benign, and non-popular benign domains, we provide insight into the TLS certificates issuance behavior for phishing domains, focusing on the selection of the certificate authority, the validation level of the certificates, and the phenomenon of *certificate sharing* among phishing domains. Our results show that phishing domains gravitate to a relatively small selection of certificate authorities, and disproportionately to *cPanel*, and tend to rely on certificates with a low, and cheap, validation level. Additionally, we demonstrate that the vast majority of certificates issued for phishing domains cover more than only phishing domains. These results suggest that a more pro-active role of CAs and putting more emphasis on certificate revocation can have a crucial impact in the defense against phishing attacks.

1 INTRODUCTION

Decades after its inception in the '90s, phishing remains a significant problem. This scalable form of criminal activity can be characterized by the use of deception in which impersonation is used to obtain information from a target (Lastdrager, 2014). The Anti-Phishing Working Group (APWG) still reports the discovery of tens of thousands of phishing sites monthly (Anti-Phishing Working Group, 2021). This indicates that phishing is far from a solved problem, and that there remains a need for novel and improved detection, prevention and mitigation methods.

A significant effort, from both an academic and commercial perspective, has been made towards the detection and identification of phishing entities, such as URLs, emails, websites and domains. Existing approaches are in many cases based on identifying similarities between suspicious entities and known legitimate ones, as criminals conducting phishing attacks (referred to as phishers) often attempt to deceive victims into believing they are interacting with a legitimate system. However, entities that so far have not received a similar degree of scrutiny are digital certificates. Such certificates are used in establishing a secure communication channel between (among others)

web browsers and web servers. It raises the research question on whether TLS certificates can be labeled as 'phishy' or benign in the same manner URLs and domains have historically been given these labels, ultimately preventing Internet users from interacting with websites serving these certificates. The recent large-scale adoption of two technologies has resulted in a trail of certificates, which potentially can be used for an alternative detection method of phishing attacks:

- HTTPS, *i.e.*, HTTP over Transport Layer Security, by phishing websites, for encrypting traffic between the browser and web server
- The submission of newly-issued TLS certificates to Certificate Transparency (CT) logs by certificate authorities

Certificate Transparency has been developed for third parties to monitor the logs for fraudulently issued certificates, resulting in TLS certificates being inserted in these logs in near real-time and additionally on a global scale. Furthermore, the issuance of certificates is assumed to occur early in the lifecycle of a domain, and thereby also of the phishing attack. Prior research has identified a general short domain lifetime and disposable nature of phishing domains, which we hypothesize to be reflected in the CT log artifacts we can observe:

1. The selected *certificate authority* that phishing domains resort to for issuing their certificates are expected to be cheap and certificate issuance is expected to be automated. In addition, an analysis of CA selection may reveal certain patterns related to the phishing hosting infrastructure of phishers.
2. The *validation level* of certificates is a proxy for the monetary cost that phishers invest into increasing the perceived legitimacy of phishing websites.
3. Phishing attacks may be part of a larger campaign, and the preparation of those attacks may be coordinated. It is hypothesized that this is reflected in certificates covering more than one phishing domain, which in practice would simplify the hosting infrastructure of the phishing attack.

In this work, we test these hypotheses by analyzing a large collection of TLS certificates, hinting towards the ‘phishiness’ of the certificates. The results serve as a preliminary motivation for pursuing a CT-based phishing mitigation system. It is namely important that phishing domains and non-phishing domains handle their infrastructure significantly different from a TLS perspective, in order to rely on them for such as mitigation system.

Prior to testing these hypotheses, we used our collected data to show that for the majority of phishing domains, a certificate is issued before the domain gets blacklisted, which emphasizes the relevancy of a CT log-based protection system. Our main findings are as follows:

- Phishers resort to a relatively small set of CAs for their certificate issuance, and the issuer of certificates reveals information regarding the infrastructure on which services for domains are served, as illustrated by a significant number of cPanel servers for phishing domains.
- Phishers seldom resort to the more expensive EV certificates, and rarely to OV certificates, although these results apply to non-popular domains as well.
- Certificates rarely cover *only* phishing certificates, but a large fraction of certificates issued for phishing domains cover other domains as well.

The remainder of the paper is structured as follows. We present the context of the paper in the background and related work in Sections 2 and 3. This is followed by a description of the methodology and the results in Sections 4 and 5. The overall impact of the results is discussed in Section 6.

2 BACKGROUND

Transport Layer Security (TLS) — like its predecessor SSL — is the underlying protocol suite for encrypting communication on the Internet. The secure communication it provides is facilitated by a public key infrastructure, in which the identity of an entity (such as a domain name) is bound to a cryptographic public key. The proof of such a binding is stored and distributed in the form of an X.509 certificate, and consists of the identities, the public key and a cryptographic signature that allows a web browser to verify the identity of the web server it is initiating a TLS connection with. The process of issuing certificates is handled by one of the hundreds of certificate authorities (CAs), which are inherently-trusted, third-party organizations. The X.509 certificate standard supports the binding with multiple identities (Cooper et al., 2008), and these identities are referred to as Subject Alternative Names (SANs). This allows an organization to request a single certificate for multiple domains (or other entity types such as IP addresses), thereby reducing the number of certificates required to secure web traffic towards their infrastructure.

An applicant applies for a certificate at a CA with information about the to-be-created certificate, such as the SANs to cover, the validation level of the certificate, and the period for which the certificate is valid. The CAs are tasked to verify that the requester of a certificate does in fact own all entities that the certificate is about to cover. This verification can be done via a variety of methods depending on the CA, including email verification, an HTTP endpoint, or more thorough background checks. After an applicant proves ownership of all SANs, the CA issues a certificate, signed by their own root (or intermediate) certificate.

Each certificate has a validation level associated with it, indicating the depth - and therefore also the cost - of the verification process the CA and requester went through to get the certificate issued. Domain validated (DV) certificates require the least validation, followed by organization validated (OV) certificates and lastly extended validated (EV) certificates. The latter validation type is reserved for corporations and recognized entities, and requires an in-depth background check. The benefit of these more expensive OV and EV certificates used to be a different indicator in the browser, although most browsers are moving towards removing these differences¹, as different indicators have proven to be ineffective (Thompson et al., 2019).

¹<https://blog.chromium.org/2018/05/evolving-chromes-security-indicators.html>

In addition to the validation level and SANs to be covered in the domain, an applicant must declare the duration of the *validity period* of the certificate. As a security mechanism, certificates expire after a timespan after which a certificate cannot successfully be verified and TLS connections must be rejected. Typically, this period lies between a few months and couple of years, and since 2020, the *CA/Browser Base-line Requirements* restricts the validity period to a maximum of 13 months (CA/Browser Forum, 2020). This period prevents certificates from being abused for long periods of time, in case the private key of the cryptographic key pair of the certificate owner is compromised.

2.1 Certificate Transparency

Fraudulent issuance of a certificate — regardless of its malicious intent — can have a disastrous impact, as demonstrated by two incidents in 2011. Attackers were able to obtain certificates for domains including `google.com` and `microsoft.com`, allowing them to perform large-scale man-in-the-middle attacks (Prins, 2011). The reputation damage eventually resulted in the bankruptcy of one of the CAs. As a direct response to these incidents, the Certificate Transparency project was initiated with the goal of monitoring and auditing the certificate issuance of CAs. The project encourages CAs to submit every issued certificate to publicly accessible, append-only CT ‘logs’. These CT logs can be run by third-party organizations and, due to their public availability, allow anyone to monitor for fraudulent issuance of certificates for their domains. As of February 2021, Chrome’s CT policy recognizes more than a hundred logs being operated by 21 different organizations, and they contain over 12 billion issued certificates².

After submitting a certificate to a CT log, a CA obtains a Signed Certificate Timestamp (SCT), which acts as a proof from the log operator that the certificate is, or will soon be, appended to the log. This SCT is embedded in the certificate, which can be used during the TLS handshake between a browser, and a web server to verify the inclusion of the certificate in the logs. As of April 2018, Chrome has started to require any certificate to comply with its CT policy, effectively requiring all certificates to be logged in at least two CT logs³. A similar policy was introduced by Apple in October 2018⁴. Due to the large combined browser market share of Google and

Apple (an estimated 82.3% as of February 2021⁵), this has resulted in the large-scale adoption of the CT framework. Combined with a generally increased adoption of HTTPS (an estimated 84% for phishing URLs (Anti-Phishing Working Group, 2021)), this makes the CT logs a promising data source for phishing detection.

3 RELATED WORK

Scheitle et al. (Scheitle et al., 2018) recognized the potential of CT logs for phishing detection based on a preliminary experiment, focusing primarily on typosquatting domains, the practice of registering domains looking similar to other domains in order to create confusion. More recently, Fasllija et al. (Fasllija et al., 2019) explored this idea further by implementing and evaluating a classifier based on the certificates contained in the CT logs. Similarly, Sakurai et al. (Sakurai et al., 2020) built domain name templates and match newly issued certificates against these templates to identify new phishing domains. Commercial initiatives such as CertSpotter⁶, PhishFinder⁷, and Facebook⁸ started to provide protection services that alert domain owners whenever a certificate for their domains is submitted to a CT log. These initiatives primarily rely on the lexical properties of the domains and do not explore the TLS-specific information contained in the logs.

Alternative early detection systems for domain abuse have been proposed in prior research. The works of Hao et al. (Hao et al., 2013; Hao et al., 2016) resulted in PREDATOR, a proactive detection system of spamming domains. The Dutch registry, SIDN, uses nDEWS (Moura et al., 2016) as an early detection system for various types of domain abuse, operating at the level of a top-level domain. Lever et al. (Lever et al., 2016) detect domain ownership changes for identifying malicious registrations. These systems primarily rely on the DNS for their data, and could potentially be combined with CT log data analysis for better performances.

The CT logs have been used as a source data set for other application areas besides phishing domain detection. Manousis et al. (Manousis et al., 2016) analyse the impact of *Let’s Encrypt* on the TLS ecosystem, finding early evidence for the adoption of typosquatters and malware hosters. Similarly, Aer-

²<https://www.certificate-transparency.org/known-logs>

³https://github.com/chromium/ct-policy/blob/master/ct_policy.md

⁴<https://support.apple.com/en-us/HT205280>

⁵<https://www.w3counter.com/globalstats.php>

⁶<https://ssllmate.com/certspotter/>

⁷<https://phishfinder.io/>

⁸<https://developers.facebook.com/tools/ct>

sten et al. (Aertsen et al., 2017) identify that in the first year of *Let's Encrypt's* introduction, primarily low-cost domains started to resort to *Let's Encrypt* as their CA. VanderSloot et al. (VanderSloot et al., 2016) investigated the coverage of the entire TLS ecosystem from various viewpoints (in terms of observed certificates). They identified that CT logs at the time already captured over 90% of all certificates, which since then presumably has only increased further.

Prior research has relied on alternative certificate-related datasets. These datasets were often collected by either actively probing TLS servers (*i.e.* active measurements) or monitoring network traffic for TLS handshakes (*i.e.* passive measurements). Active measurements allow researchers to not only capture the certificate, but also parameters exchanged during the TLS handshake (Zhang et al., 2014). Drawbacks of active measurements include its reliance on a list of known domains⁹, partial coverage caused by unavailability of servers and the difficulty of performing continuous measurements. In passive measurements, such as the works by Razaghpanah et al. (Razaghpanah et al., 2017), have no control of which certificate are observed. The degree of coverage of the TLS ecosystem is highly dependent on the quality and quantity of the observation point(s), as a small and homogenous client population is unlikely to query a significant portion of TLS-enabled servers. In contrast to CT logs, neither active nor passive measurements are guaranteed to provide certificates in a timely fashion (*i.e.*, immediately after the certificate has been issued). Note that different measurement types can be used in conjunction (Holz et al., 2019), compensating each other's drawbacks and thereby resulting in a higher quality dataset.

4 METHODOLOGY

To obtain a manageable dataset of certificates (in terms of the cost of analyzing these certificates), we collect certificates issued for a sample of three domain types. An exploratory analysis of all certificates available from the CT logs is considered infeasible, as they contain over twelve billion certificates. Rather than considering Fully Qualified Domain Names (FQDNs), we are primarily concerned with analyzing *root domains*, the part of the FQDN under which

⁹A TLS handshake involves the *Server Name Indication* extension of TLS (Blake-Wilson et al., 2003), *i.e.*, a field that contains the domain name that the client intends to connect to.




Figure 1: The process of retrieving the three samples of domains and their associated certificates. The grey boxes represent the source data sets, the green boxes represent three domain types and the red boxes represent the actions performed by the authors. D = set of domains, U = set of URLs, C = set of certificates.

the domain is registered by its owners¹⁰. Certificates issued for domains under this root domain are generally requested by the domain owner since the domain owner is tasked to demonstrate ownership of the domain to the CA, although newer verification methods (such as an HTTP-based method) challenge this assumption. The aforementioned three domain types used for analysis are defined as follows:

- *Phishing domains* (D_{phish}) are those domains that have historically been used in phishing attacks.
- *Popular benign domains* (D_{pop}) represent the set of domains used for highly popular services that receive the majority of traffic on the Internet.
- *Non-popular benign domains* (D_{nonpop}) are domains that receive a low amount of traffic, yet have not been seen as part of phishing attacks.

For each domain type, we sampled 10,000 domains, resulting in a total number of 30,000 domains for which certificates were collected. Figure 1 shows the process to retrieve the sample of domains and their associated certificates. The *domain sampling* step consists of converting the source domains into a set of sampled, labeled domains. In the *certificate retrieval* step, all certificates for these sampled domains are retrieved, which in turn are analyzed further.

Domain Sampling. The Tranco list (Pochat et al., 2019) combines three other domain lists to provide a robust list of the top one-million popular domains, in terms of the popularity of those domains¹¹. We assume that the domains on the top of this list are inherently benign (*i.e.*, they were not registered for malicious purposes), under the assumption that a maliciously registered domain will never become popular enough to receive high amounts of traffic. The top of

¹⁰For instance, the root domain for the FQDN `www.example.co.uk` is `example.co.uk`.

¹¹The definition of ‘popularity’ differs slightly between the three source lists, but generally represents the amount of traffic the domain receives.

the list therefore represents popular, benign domains.

The APWG’s eCrime eXchange (eCX) platform¹² contains millions of known phishing URLs submitted by contributing organizations, whenever the submitted URL is found to host phishing content. The root domains extracted from these URLs form the base of the phishing domain sample, denoted as D_{ecx} . Some popular domains are often abused for hosting phishing content, such as Facebook or Google Docs. This however does not imply those domains are registered for malicious purposes; the content on these sites is merely user-generated and therefore contains a mix of malicious and benign content. The set of *phishing domains* is defined as $D_{phish} = D_{ecx} \setminus D_{pop}$, or the root domains extracted from eCX URLs, excluding any domain that is in the top 1M Tranco domains.

Lastly, the non-popular, benign domain type is defined as $D_{nonpop} = D_{com} \setminus (D_{ecx} \cup D_{pop})$, where D_{com} is the set of all domains in the .com zone (comprising almost half of all domains globally). None of these domains receive a large amount of traffic and have never been seen in a phishing attack according to the eCX platform.

Certificate Retrieval. We collect any certificate submitted to the CT log ecosystem that covers (at least) one of the domains in the domain sets. This collection is done through a Google BigQuery dataset provided by the Censys search engine (Durumeric et al., 2015). The criterion for a certificate to cover a domain name is that the SANs list of the certificate contains (1) the root domain itself, (2) a subdomain of the root domain, or (3) a wildcard domain under the root domain. In addition to the raw certificates, we also collect several extra fields, including the fingerprint of the certificate, and the validity of the certificate according to different root stores.

5 RESULTS

All four datasets (*i.e.*, the Tranco list, the .com zone file, the eCX URLs, and certificates) were collected in the beginning of 2020. We used the Tranco list from February 20th, 2020¹³. The list of eCX URLs comprises over 8.6 million URLs (belonging to 1.02 million unique root domains), and contains URLs discovered up to February 27th, 2020. The certificates from Censys were collected on March 12th, 2020. An overview of the resulting dataset is shown in Table 1. The vast majority of the 79.1 million collected

Table 1: Overview of the collected dataset.

	# domains	# domains without cert	# certificates	% certificates trusted
Phishing	10.0k	4.2k	184.9k	95.9
Popular	10.0k	175	78.9M	40.4
Non-pop.	10.0k	6.3k	95.6k	99.0
Total	30.0k	10.7k	79.1M	

certificates were issued for popular domains. This is partially explained by the fact that only a small fraction of popular domains had no certificate issued for it, compared to 4,235 phishing domains, and 6,283 of the non-popular domains. In our further analysis, we discard certificates that are *untrusted* by browsers, as those certificates cause browsers to present users a warning page, instead of the actual content. The purpose of issuing certificates in the first place is to make pages seem legitimate, which is defeated when a warning page is shown. A certificate can be untrusted for various reasons, such as being self-signed or being signed by a root certificate that a particular browser or operating system does not trust. This filtering of untrusted certificates primarily affects the popular domains, with nearly 60% of all certificates issued for those domains being untrusted. The table suggests that phishing domains have adopted HTTPS to a higher degree than non-popular domains, with nearly twice the number of certificates issued for them.

5.1 Temporal Analysis

As previously stated, one advantage of CT log analysis over passive and active measurements is the presumed early submission and publication of certificates to the CT logs. Using CT logs as a data source for early detection is only viable if certificates can consistently be monitored prior to the blacklisting of the domains they cover, and as such we test the following hypothesis:

Hypothesis 1. *Certificate issuance occurs prior to blacklisting of phishing domains*

We analyze the time difference between (1) the blacklisting of a phishing domain and (2) the issuance of its certificate(s). This process is made more difficult because a domain can both have been blacklisted multiple times, and have multiple certificates issued for it. The former challenge is tackled by selecting the first timestamp of blacklisting seen for that specific domain, whereas the latter is tackled by introducing the notion of the *closest certificate*. For each timestamp of first blacklisting, the most recent certificate issuance *before* blacklisting is considered. If no certificate before blacklisting exists, the oldest certificate

¹²<https://apwg.org/ecx/>

¹³Available at <https://tranco-list.eu/list/XWNN>.




Figure 2: The computation of the time difference between domain blacklisting and the closest certificate issuance. For domain d_1 , the most recent certificate *before* the first blacklisting timestamp is taken (resulting in a negative time difference), whereas for domain d_2 the earliest certificate issuance timestamp *after* the first blacklisting timestamp is taken (resulting in a positive time difference).

after blacklisting is considered. The first motivation for this definition is that the most recent certificate before blacklisting is most likely to be issued by the same owner of the domain at the time of blacklisting, as older certificates may have been issued by a previous owner of the domain. Secondly, certificates issued after blacklisting are meaningless for early protection against phishing attacks abusing that domain, hence the low priority of including those certificates. This process is illustrated in Figure 2.

Certificates issued before the first blacklisting date are not necessarily requested by the same domain owner that caused the domain to be blacklisted. Given the decade-spanning time window of the CT logs, it is possible to observe a certificate issued for a domain ten years before it got blacklisted. Since the identification of domain ownership is inherently a difficult task, we instead estimated a lower and upper bound for the number of phishing domains for which we can identify a relevant certificate, issued before the domain got blacklisted.

Upper Bound. For all phishing domains in the dataset, the time difference between the issuance of the closest certificate (t_C) and the earliest URL blacklisting timestamp (t_B) is computed (denoted as $\Delta(t_C, t_B)$). A negative time difference implies that the closest certificate was issued before the URL was blacklisted. For 75.66% of all phishing domains, the earliest certificate timestamp occurs *before* the earliest blacklisting timestamp, which serves as the upper bound of the percentage of domains a CT log-based phishing domain detection system can protect against.

Lower Bound. In order to estimate a lower bound, the registration process of a domain is taken into account. Since domains can change their ownership during their lifecycle, it is vital to only consider certificates issued by the same owner that owned the domain during the time of blacklisting. A domain




Figure 3: Timeline of the re-registration process of a domain name.




Figure 4: The ECDF of the time difference between the issuance of the closest certificate and the earliest URL blacklisting timestamp, for each phishing domain. The figure is zoomed in on the period around the auto-renew and registration period of a domain.

registered under the .com top-level domain (TLD) goes through a process after a domain expires, during which the control of the domain is taken back by the registrar, and after which a domain can potentially change owner¹⁴. An auto-renew period of 45 days is followed by a 30-day redemption period, after which the domain is listed as ‘pending delete’ for 5 days. After this period, the domain is up for re-registration (Lauinger et al., 2017) (see Figure 3). In the auto-renew period, domain owners have the opportunity to renew or sell their domain, and as such a domain cannot change ownership in the 35-day period between the end of the auto-renew period, and the earliest potential to re-register a domain. Given a domain for which a URL was blacklisted at an arbitrary timestamp $t = 0$, all certificates issued in the period $[-35 \text{ days}, 0]$ are requested by the same owner that owned the domain at the time of blacklisting. Figure 4 shows a zoomed-in portion of the Empirical Cumulative Distribution Function (ECDF) of $\Delta(t_C, t_B)$ of all phishing domains, and shows that 56.13% of domains had their closest certificate issued in this period for which there exists the previously described certainty about the ownership of the domain. As such, this percentage is considered the lower bound. The figure

¹⁴Note that other TLDs may have different processes regarding the expiration of a domain, but given the large market share of .com, their expiration process drives our methodology

also illustrates the aforementioned upper bound.

These results indicate that for between 56.13% and 75.66% of phishing domains, a certificate issuance can be observed before a URL for that domain is being blacklisted. Even though this implies that perfect coverage of all domains does not seem feasible, it is important to note that CT logs cover domains across all TLDs and their coverage is arguably only getting higher due to the ever-increasing adoption of HTTPS. Therefore, CT log-based detection of phishing domains could potentially cover a larger set of domains than active or passive measurement-based methods.

5.2 Issuers

Each certificate in circulation has been signed by a CA, which has deliberately been approached by the requester of the certificate (through the submission of a certificate signing request). As such, the choice of CA of a certificate reflects the behavior of the domain owner. CAs have individual differences, such as pricing models, countries in which they operate, the process of verifying the identity of the requester, etc. As these differences may be important considerations for phishers, we test the following hypothesis:

Hypothesis 2. *The CA selection of phishing domains provides insight in the monetary considerations, and the operations of phishers*

In reality, it is not trivial to identify the underlying CA that signed a certificate, as certificates are generally signed by intermediate certificates¹⁵, and a single CA may operate many intermediate certificates. In addition, *cross-signing* of certificates is not uncommon, in which the root certificate of a CA signs the intermediate certificate of another CA, resulting in a chain of trust rooted in more than one CA. We therefore analyze the issuer of certificates rather than identifying the underlying CA, as the issuer selection may provide a similar insight into phishing as the CA selection would.

In order to analyze the selection of CA, it is important to handle any potential biases that can be induced. More specifically, it is imperative that a single domain should not dominate other domains due to the number of certificates issued for it (as demonstrated earlier in Table 1), or a specific issuer to dominate others due to the short validity period of their certificate (thereby requiring customers to frequently re-issue certificates).

The set of certificates issued for domain d_i issued by issuer a_j is formalized as $C_{\text{issuer}}(d_i, a_j)$. The duration

¹⁵These intermediate certificates are in turn signed by a root certificate or another intermediate certificate.




Figure 5: The distribution of the month of issuance of certificates for the three domain types.

of an individual certificate c 's validity period is denoted as $v(c)$. The set of all existing issuers is \mathcal{A} . Using these formulations, we define *domain issuer preference* $\mathcal{P}_{\text{issuer}}(d_i, a_j)$, which represents how much a particular issuer is preferred by a domain:

$$\mathcal{P}_{\text{issuer}}(d_i, a_j) = \frac{\sum_{c \in C_{\text{issuer}}(d_i, a_j)} v(c)}{\sum_{a \in \mathcal{A}} \sum_{c \in C_{\text{issuer}}(d_i, a)} v(c)} \quad (1)$$

This measure is a value between one and zero, and $\sum_{a \in \mathcal{A}} \mathcal{P}_{\text{issuer}}(d_i, a) = 1$. This definition is used to express the *domain type issuer preference* for domain type k towards an issuer a_j :

$$\mathcal{P}_{\text{issuer}}(D_k, a_j) = \frac{\sum_{d \in D_k} \mathcal{P}_{\text{issuer}}(d, a_j)}{\sum_{a \in \mathcal{A}} \sum_{d \in D_k} \mathcal{P}_{\text{issuer}}(d, a)} \quad (2)$$

Where D_k is the set of all domains for domain type k . Similar to Eq. 1, the following holds: $\sum_{a \in \mathcal{A}} \mathcal{P}_{\text{issuer}}(D_k, a) = 1$. Eq. 1 ensures that certificates are weighted proportionally to their validity period, and that each domain within a domain type is equally weighted disregarding the number of certificates issued for the domain, whereas Eq. 2 ensures a proper comparison between domain types.

Using Eq. 2, we compute the domain type preferences for all combinations of issuers and domain types, of which the results are illustrated in Table 2. The top 10 most preferred issuers are shown in the table. In total, 853 issuer certificates were used to sign the set of certificates, of which 846 were used for popular domains, 132 for phishing domains and only 125 for non-popular domains. Generally, the preference for phishing domains and non-popular domains are fairly similar, with the popular domains having a vastly different preference profile.

Although *Let's Encrypt* is the most preferred issuer for popular domains (with 15.7%), it is not nearly as popular as the other two domain types (46.19% and 61.46% for phishing and non-popular domains respectively). One explanation could be that *Let's Encrypt* is a relatively new CA, and popular domains tend to be long-lived, resulting in a stronger preference for CAs that have been operating longer, or for

Table 2: The domain type issuer preference (as percentage) for the ten most preferred issuers and the three domain categories.

a_j	$\mathcal{P}_{\text{issuer}}(D, a_j)$ (in %)		
	$D = D_{\text{phish}}$	$D = D_{\text{pop}}$	$D = D_{\text{nonpop}}$
Let's Encrypt Authority X3	*46.19	*15.70	*61.46
cPanel, Inc. Certification Authority	*33.53	1.23	*10.69
COMODO ECC Domain Validation Secure Server CA 2	*7.94	5.23	*4.51
DigiCert SHA2 Secure Server CA	0.06	*6.88	0.15
Amazon	0.14	*5.91	0.73
Go Daddy Secure Certificate Authority - G2	1.91	4.50	4.21
CloudFlare Inc ECC CA-2	3.65	1.33	4.43
GlobalSign CloudSSL CA - SHA256 - G3	0.32	4.23	0.06
DigiCert SHA2 High Assurance Server CA	0.01	4.09	0.02
Sectigo RSA Domain Validation Secure Server CA	1.38	1.55	4.03

*among the three most preferred authorities within the domain category.




Figure 6: Distribution of the month of URLs blacklisted for phishing domains with cPanel certificates (black). The month of the discovery of several critical vulnerabilities is displayed in red.

now-defunct CAs. Figure 5 shows the distribution of the month of issuance of certificates for the three domain types. The figure shows that the vast majority of certificates are issued after 2018 and that there is no major difference between the domain types, suggesting that the age of *Let's Encrypt* does not play a large role in the issuer preference.

The issuer preference provides information regarding the infrastructure used to serve services using popular domains, of which the high preference for the *Amazon* issuer by popular domains (5.91%, or the third-most preferred issuer) is an example. Amazon issues public certificates for customers of their AWS cloud platform, which specifically focuses on “securing public websites with significant traffic requirements”¹⁶, indicating that part of the infrastructure of popular domains with such certificates operates on AWS. Similarly, we find that phishing domains have a strong preference for the *cPanel* issuer (33.53%) compared to popular (1.23%) and non-popular (10.69%) benign domain types. *cPanel* is a popular web hosting platform with the option for automatic issuance and deployment of TLS certificates. We hypothesize that either (1) phishers rely on *cPanel* for their host-

ing infrastructure on a large scale, or (2) that *cPanel* accounts are compromised at a large scale and repurposed for conducting phishing attacks. Under the assumption that the second hypothesis can only occur on a large scale in case of the discovery of a critical vulnerability, we would expect to see a peak of submissions of phishing URLs on the eCX for domains running *cPanel* software, as *cPanel* hosts are compromised around the time of the vulnerability disclosure. Several critical vulnerabilities were reported for *cPanel* in September 2019¹⁷, but according to Figure 6 this did not coincide with unusual numbers of related URLs being reported. We have found no evidence of peaks of URL submissions for phishing domains running on *cPanel* software, suggesting that the first hypothesis holds true, but this requires further research to confirm.

5.3 Validation Level

In addition to the CA selection, the selection of a validation level provides insight in the monetary investment in conducting phishing attacks, which leads to the following hypothesis:

Hypothesis 3. *Phishers resort to certificates with higher validation levels to increase the perceived legitimacy of phishing websites.*

Similar to the issuer analysis, the analysis of validation level differences requires a normalization step taking into account the number of certificates and the validity duration of certificates. The set of all certificates issued for domain d_i with a validation level l_j is denoted as $C_{\text{val}}(d_i, l_j)$. The set of possible validation levels (i.e., DV, OV, EV, and unknown) is denoted as \mathcal{L} . The domain validation level preference $\mathcal{P}_{\text{val}}(d_i, l_j)$

¹⁶<https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html>

¹⁷We used a CVE record with a score of 8 or higher as the definition of a critical vulnerability, obtained from https://www.cvedetails.com/vulnerability-list/vendor_id-1766/product_id-3023/Cpanel-Cpanel.html

Table 3: Values of $P_{val}(D, l)$ for all domain categories and validation levels (as percentage).

$\downarrow D \rightarrow l$	DV	OV	EV	Unknown
D_{phish}	92.94	6.65	0.11	0.31
D_{pop}	49.41	39.17	4.03	7.38
D_{nonpop}	91.99	7.28	0.14	0.59

represents how much a particular validation level is preferred by a specific domain:

$$P_{val}(d_i, l_j) = \frac{\sum_{c \in C_{val}(d_i, l_j)} v(c)}{\sum_{l \in \mathcal{L}} \sum_{c \in C_{val}(d_i, l)} v(c)} \quad (3)$$

Again, $v(c)$ denotes the duration of the validity period of certificate c , and this measure is a value between one and zero, and $\sum_{l \in \mathcal{L}} P_{val}(d_i, l) = 1$. Then, the domain type validation level preference for each of the domain categories D_k (i.e., D_{phish} , D_{pop} and D_{nonpop} for phishing, popular and non-popular sampled domains respectively) for a specific validation level l_j is defined as follows:

$$P_{val}(D_k, l_j) = \frac{\sum_{d \in D_k} P_{val}(d, l_j)}{\sum_{l \in \mathcal{L}} \sum_{d \in D_k} P_{val}(d, l)} \quad (4)$$

We rely on the reported validation levels from the Censys dataset¹⁸ to compute the preferences. Table 3 shows the results of the computation of these values, for all domain type and validation level combinations. Unsurprisingly, popular domains have a stronger preference for OV and EV certificates, as popular domains have a larger incentive and budget for protecting the legitimacy of their brand. Phishing and non-popular benign domains have a highly similar preference.

Notably, we identified 133 EV certificates issued for 16 unique phishing domains. Further inspection shows that only two of those 16 domains are marked by VirusTotal as malicious, questioning whether the other 14 are really phishing domains in the first place. As such, we conclude that EV certificates are virtually unused by phishers in our dataset, and this suggests that EV certificates are in fact a meaningful method for marking websites as trustworthy, even though browsers stopped presenting such indicators to users.

5.4 Certificate Sharing

Based on the following hypothesis, there is an expectation that phishing domains may be deployed on

¹⁸We verified the correctness of these values by computing a validation level based on the existence of specific Object Identifiers (OIDs) in the X.509 certificate extensions and matching those against the Censys values. Our method in general led to more inconclusive results.




Figure 7: ECDF of the number unique roots of certificates covering at least a single phishing domain ($n = 975,426$). The red line represents the fraction of covering only a single unique root domain.

a shared infrastructure, which would be reflected by *shared certificates*, or certificates that cover more than a single phishing domain.

Hypothesis 4. *Individual phishing attacks may be part of a larger campaign, and the deployment of those attacks is coordinated.*

In order to test this hypothesis, we collect all certificates in our dataset that cover at least a single phishing domain. Note that we consider all phishing domains, not only the 10,000 domains in the initial sample. In total, 975k certificates were identified. For each certificate, we extract the unique root domains, and the ECDF of the count is shown in Figure 7. 37.57% of certificates issued for phishing domains cover only a single unique root, indicating that the remaining 62.43% (or 608,934 certificates) is potential evidence that certificate sharing is rampant in phishing attacks.

Surprisingly, only 0.90% (or 5,452) of these certificates cover *only* phishing roots, meaning that the remaining certificates cover not only phishing roots, but also domains of unknown nature. The implication of this is that there is a small portion of certificates that can be considered unambiguously ‘phishy’, whereas the vast majority of certificates cannot. The existence of these ‘ambiguous’ certificates could be interpreted in several ways, including:

1. Certificates are requested by a third party to which domain control is delegated by the actual owner, and this third party requests certificates that cover benign and phishing domains.
2. There exists a vast number of domains that have not yet been discovered to be phishing domains, which means that in reality these ambiguous certificates are in fact ‘phishy’.
3. Not all domains owned by a particular domain owner are used for phishing attacks. This includes for example benign domain owners whose websites are hacked and repurposed for phishing attacks, or phishers who manage domains for non-phishing activities, and domains with user-generated content of mixed nature.

We investigate the first point through a case study:

Cloudflare. Cloudflare’s Content Delivery Network (CDN) enables customers to encrypt traffic to their websites by proxying traffic through Cloudflare servers. This requires them to change the DNS name server of their domain to Cloudflare, effectively delegating the control of their DNS configuration to Cloudflare, which is used by the ACME protocol to verify the ownership of domains. Rather than requesting a single certificate for each domain, a single certificate is requested for a set of domains, all from different owners. In addition, these certificates contain a Cloudflare-specific domain (usually the first domain name in the list of SANs), matching the structure `sni{6 digits}.cloudflaressl.com`. Naturally, none of these certificates can unambiguously be assumed to be phishy, as they cover domains from many different owners. This set of Cloudflare certificates alone already comprises 91,369 certificates, or 9.37% of the full set of certificates issued for phishing domains. We identified other similar certificate structures to Cloudflare certificates, that include SANs such as `statuspage.io` and `incapsula.com`, leading us to believe that those domains are used for similar purposes as the `sni{6 digits}.cloudflaressl.com` domains. Further research is needed to fully differentiate ‘phishy’ certificates from these CDN-type certificates.

6 DISCUSSION

The fact that between 56.13% and 75.66% of phishing domains observe a relevant certificate being issued before a URL is blacklisted is a promising result and a strong motivation for pursuing the development of CT log-based abuse prevention systems (accepting Hypothesis 1). Given the constant growth of phishing attacks being conducted over HTTPS (over 84% of identified phishing attacks in the last quarter of 2020 according to the APWG (Anti-Phishing Working Group, 2021), compared to only 10% in the first quarter of 2017), CT log’s early coverage of phishing domains is only expected to grow in the future. Simultaneously, the introduction of more traffic encryption methods, such as ECN¹⁹ could render passive measurements less effective, emphasizing the relevance of CT logs even more.

We have demonstrated that phishing domains and non-popular benign domains have similar preference profiles for the issuer selection of their certificates (except *cPanel*), but a vastly different profile compared to popular domains. Our results provide some

insight in the operations of phishing domains (*e.g.*, preference for *cPanel*), thereby we accept Hypothesis 2. Further research is required to draw stronger conclusions regarding the domains relying on *cPanel*, but these results emphasize the importance of the position of CAs in the fight against phishing. CAs have the opportunity to interfere with the TLS ecosystem through the revocation of certificates and in fact some CAs state in their policies that malicious activity is grounds for certificate revocation²⁰. Although the effectiveness of certificate revocation have historically been limited (Liu et al., 2015), our results are an argument for better handling of certificate revocation.

Our results for the validation level of certificates show that phishers rarely resort to EV certificates (thereby rejecting Hypothesis 3). Given the relatively high cost of OV and EV certificates (\$27.44 and \$72.18 per year at COMODO respectively for example²¹), this suggests phishers are not willing to significant amounts of money, or are actually rejected during the vetting process. Even though browser indicators have been demonstrated to be ineffective from the perspective of users (Thompson et al., 2019), higher validation levels could be an effective signal for identifying certificates that are not used in phishing attacks. EV certificates could act as a white-listing method for identifying ‘benign’ certificates.

Lastly, we found few unambiguous cases of shared certificates between phishing domains, with 5,452 certificates *only* covering phishing domains. The vast majority of certificates that cover at least a single phishing domain also covers domains whose nature is unknown, which makes it difficult to extract meaningful information from the certificates about the decisions of the phishers. A substantial part of this challenge can be explained by the practice of domain owners delegating control to CDNs, which aggregate many domains from many owners, and obtain certificates covering many domains of various types, including phishing. We found that 9.37% of the phishing domains had been mixed with other domains by Cloudflare, and indications that other CDNs apply similar practices. Consequently, CDNs have a role and a responsibility in the fight against phishing, when they offer the service of managing certificates, and in general when applying practices that allows phishers to mix with benign domain owners. We cannot reject nor accept Hypothesis 4, which remains inconclusive.

This work provides preliminary characteristics of

²⁰Example of Sectigo: <https://sectigo.com/uploads/files/Sectigo-CPS-v5.2.2.pdf>

²¹<https://comodossllstore.com/resources/dv-vs-ov-vs-ev-ssl-which-certificates-are-good-for-site-security/>

¹⁹<https://tools.ietf.org/html/draft-ietf-tls-esni-10#page-6>

phishing certificates, suggesting that there is a certain degree of 'phishiness' that can be assigned to a certificate. Even a simple-heuristics based warning system can potentially be useful for identifying candidate phishing domains, by for example identifying *cPanel* certificates covering several domains, including a domain that previously already was blacklisted by the eCx.

Limitations. Our methodology differentiates between three domain types, and considers phishing domains a homogenous group of domains registered for phishing purposes. Even though we accounted for domain ownership changes in the collection of certificate for phishing domains, we do not address the potential of domain compromise. Maroofi et al. (Maroofi et al., 2020) manually collected a dataset of phishing URLs and identified 58% to be maliciously registered and 42% to be compromised. This implies that there is a likelihood that our analysis of phishing domains leads to conclusions for compromised domains, rather than maliciously registered domains. Identifying whether a domain is compromised in itself is already a challenging task, which is significantly more difficult with historical data, where the website may not be online anymore.

It is possible that the domain type samples contain false positives of domains that in reality are placed in the wrong class. One cause of this could be the (lack of) vetting of URLs in the eCX platform, which could result in URLs submitted to the platform that are in reality not phishing URLs. In addition, it is unlikely that the eCX URLs are fully complete, covering every phishing URL in existence, as not all phishing attacks are detected and reported. Additionally, the eCX relies on member contributions, and these are unlikely to be fully complete. As a result, there are likely to be domains in the non-popular domain set that are in reality phishing domains. Unfortunately, this is a core limitation of the CT framework for phishing prevention, as TLS certificates are issued on a domain basis instead of on a URL basis.

7 CONCLUSIONS

This paper addresses the potential of using the phishiness of digital certificates as a method to identify phishing domains early in their lifecycle. By comparing the certificates issued for three distinct domain sets, we identify relevant patterns in the differences across these domain sets. Firstly, our temporal analysis shows that for 56.13% to 75.66% of phishing domains, a certificate is issued before the domain is

being blacklisted, indicating the scale at which CT-based mitigation can protect against phishing. Furthermore, our results show that phishing domains resort to a relatively small group of issuers, particularly gravitating to *cPanel*, which emphasizes that stronger adherence to certificate revocation lists produced by these issuers can be highly valuable. We have also shown that phishers are unlikely to resort to (expensive) EV certificates, which could suggest that domains that *do* employ them could serve as a whitelist for non-phishing domains. Lastly, we found that certificates are unlikely to be unambiguously phishy or benign, given the set of phishing domains they encompass. Although we identified only a few certificates that *only* cover phishing domains, the majority of certificates issued for phishing domains cover multiple phishing domains, which is a hopeful takeaway. These results led us to provide several suggestions for changes to the TLS ecosystem.

Our work opens up several pathways for future work. Firstly, given the preliminary nature of our results, we encourage the research community to integrate our results in novel or existing domain classification methods. An alternative promising direction is the disambiguation of certificates, which - if successful - could lead to a very effective way to propagate the 'phishiness' of a certificate to all the domains it covers. Additionally, the potential impact of CAs (*cPanel* in particular) and certificate revocation could be explored in more detail.

ACKNOWLEDGEMENTS

This research is carried out under the SecDNS project, funded by Innovation Fund Denmark. We thank Censys.io for sharing their CT log data with us, and we are thankful for the APWG for granting us access to their eCX platform.

REFERENCES

- Aertsen, M., Korczyński, M., Moura, G. C. M., Tajalizadehkhoob, S., and van den Berg, J. (2017). No domain left behind: is let's encrypt democratizing encryption? In *Proceedings of the Applied Networking Research Workshop*. ACM.
- Anti-Phishing Working Group (2021). Phishing Activity Trends Report - 4th Quarter 2020. https://docs.apwg.org/reports/apwg_trends_report_q4_2020.pdf. Accessed: 2021-02-12.
- Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and Wright, T. (2003). Transport layer security (tls)

- extensions. RFC 3546, RFC Editor. Accessed: 2021-02-12.
- CA/Browser Forum (2020). Baseline requirements for the issuance and management of publicly-trusted certificates (version 1.7.3). Technical report, CA/B Forum. Accessed: 2021-02-08.
- Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and Polk, W. (2008). Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280, RFC Editor. <http://www.rfc-editor.org/rfc/rfc5280.txt>.
- Durumeric, Z., Adrian, D., Mirian, A., Bailey, M., and Halderman, J. A. (2015). A search engine backed by internet-wide scanning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 542–553. ACM.
- Fasllija, E., Enişer, H. F., and Prünster, B. (2019). Phishhook: Detecting phishing certificates using certificate transparency logs. In *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 320–334. Springer International Publishing.
- Hao, S., Kantchelian, A., Miller, B., Paxson, V., and Feamster, N. (2016). PREDATOR: Proactive Recognition and Elimination of Domain Abuse at Time-Of-Registration. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1568–1579. ACM.
- Hao, S., Thomas, M., Paxson, V., Feamster, N., Kreibich, C., Grier, C., and Hollenbeck, S. (2013). Understanding the domain registration behavior of spammers. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, page 63–76, New York, NY, USA. Association for Computing Machinery.
- Holz, R., Amann, J., Razaghpanah, A., and Vallina-Rodriguez, N. (2019). The era of TLS 1.3: Measuring deployment and use with active and passive methods. *CoRR*, abs/1907.12762.
- Lastdrager, E. E. (2014). Achieving a consensual definition of phishing based on a systematic review of the literature. *Crime Science*, 3(1).
- Lauinger, T., Chaabane, A., Buyukkayhan, A. S., Onarlioglu, K., and Robertson, W. (2017). Game of registrars: An empirical analysis of post-expiration domain name takeovers. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 865–880, Vancouver, BC. USENIX Association.
- Lever, C., Walls, R., Nadji, Y., Dagon, D., McDaniel, P., and Antonakakis, M. (2016). Domain-z: 28 registrations later measuring the exploitation of residual trust in domains. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 691–706. IEEE.
- Liu, Y., Tome, W., Zhang, L., Choffnes, D., Levin, D., Maggs, B., Mislove, A., Schulman, A., and Wilson, C. (2015). An end-to-end measurement of certificate revocation in the web's pki. In *Proceedings of the 2015 Internet Measurement Conference*, IMC '15, page 183–196, New York, NY, USA. Association for Computing Machinery.
- Manousis, A., Ragsdale, R., Draffin, B., Agrawal, A., and Sekar, V. (2016). Shedding light on the adoption of let's encrypt.
- Maroofi, S., Korczynski, M., Hesselman, C., Ampeau, B., and Duda, A. (2020). COMAR: Classification of compromised versus maliciously registered domains. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE.
- Moura, G. C. M., Muller, M., Wullink, M., and Hesselman, C. (2016). nDEWS: A new domains early warning system for TLDs. In *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*, pages 1061–1066. IEEE.
- Pochat, V. L., Goethem, T. V., Tajalizadehkhoob, S., Korczynski, M., and Joosen, W. (2019). Tranco: A research-oriented top sites ranking hardened against manipulation. In *Proceedings 2019 Network and Distributed System Security Symposium*. Internet Society.
- Prins, J. (2011). DigiNotar Certificate Authority breach “Operation Black Tulip”. <https://media.threatpost.com/wp-content/uploads/sites/103/2011/09/07061400/rapport-fox-it-operation-black-tulip-v1-0.pdf>. Accessed: 2021-02-12.
- Razaghpanah, A., Niaki, A. A., Vallina-Rodriguez, N., Sundaresan, S., Amann, J., and Gill, P. (2017). Studying TLS usage in android apps. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, pages 350–362. ACM.
- Sakurai, Y., Watanabe, T., Okuda, T., Akiyama, M., and Mori, T. (2020). Discovering HTTPSified phishing websites using the TLS certificates footprints. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 522–531. IEEE.
- Scheitle, Q., Gasser, O., Nolte, T., Amann, J., Brent, L., Carle, G., Holz, R., Schmidt, T. C., and Wählisch, M. (2018). The rise of certificate transparency and its implications on the internet ecosystem. In *Proceedings of the Internet Measurement Conference 2018*, IMC '18, page 343–349, New York, NY, USA. Association for Computing Machinery.
- Thompson, C., Shelton, M., Stark, E., Walker, M., Schechter, E., and Felt, A. P. (2019). The web's identity crisis: understanding the effectiveness of website identity indicators. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pages 1715–1732.
- VanderSloot, B., Amann, J., Bernhard, M., Durumeric, Z., Bailey, M., and Halderman, J. A. (2016). Towards a complete view of the certificate ecosystem. In *Proceedings of the 2016 Internet Measurement Conference*, IMC '16, page 543–549, New York, NY, USA. Association for Computing Machinery.
- Zhang, L., Choffnes, D., Levin, D., Dumitraş, T., Mislove, A., Schulman, A., and Wilson, C. (2014). Analysis of ssl certificate reissues and revocations in the wake of heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, page 489–502, New York, NY, USA. Association for Computing Machinery.