Adoption of Email Anti-Spoofing Schemes: A Large Scale Analysis

Sourena Maroofi, Maciej Korczyński, Arnold Hölzel, Andrzej Duda, *Member, IEEE*

***Abstract*—Sending forged emails by taking advantage of domain spoofing is a common technique used by attackers. The lack of appropriate email anti-spoofing schemes or their misconfiguration may lead to successful phishing attacks or spam dissemination. In this paper, we evaluate the extent of the SPF and DMARC deployment in two large-scale campaigns measuring their global adoption rate with a scan of 236 million domains and high-profile domains of 139 countries. We propose a new algorithm for identifying defensively registered domains and enumerating the domains with misconfigured SPF rules by emulating the SPF check_function. We define for the first time new threat models involving subdomain spoofing and present a methodology for preventing domain spoofing, a combination of good practices for managing SPF and DMARC records and analyzing DNS logs. Our measurement results show that a large part of the domains do not correctly configure the SPF and DMARC rules, which enables attackers to successfully deliver forged emails to user inboxes. Finally, we report on remediation and its effects by presenting the results of notifications sent to CSIRTs responsible for affected domains in two separate campaigns.**

***Index Terms*—Email authentication, email spoofing, spam, phishing, SPF, DMARC, DNS, Internet measurements, notifications**

I. Introduction

EMAIL spoofing consists of sending a message with a forged sender address and other parts of the email header so that it appears as sent from a legitimate source. Attackers commonly use this method to mislead the receivers, gain their trust, and eventually, achieve some malicious goals. Phishing and spam campaigns are examples of attacks that rely on email spoofing. Despite tremendous efforts deployed to mitigate this technique, it is still one of the most successful attacks responsible for significant damage. As an example, email spoofing costed US victims more than 1.2 billion dollars in 2018 according to the Internet crime report [1].

Email spoofing comes in two types. The first one consists of *compromising legitimate servers* and using their mail transfer agent to send spoofed emails to victims either by specifying a different `Reply-to:` address or providing a phishing URL in the body of the message. The second type is *domain spoofing* in which attackers send emails on behalf of legitimate domains, e.g., a forged email from *account-security-noreply@accountprotection.microsoft.com* impersonating the Microsoft support team with a fake landing page looking like a real Microsoft login page to steal user credentials [2]. In this paper, we investigate the second type of email spoofing.

The Simple Mail Transfer Protocol (SMTP) for email distribution does not provide support for preventing spoofing [3] so mail systems need to rely on *security extensions* such as the Sender Policy Framework (SPF) [4], the DomainKeys Identified Mail (DKIM) [5], and Domain-based Message Authentication, Reporting, and Conformance (DMARC) [6] to authenticate the sender and decide what to do with suspicious emails. The extensions define a set of rules that specify who is allowed to send emails on behalf of a given domain name and how to deal with suspicious messages. A careful deployment of the extensions can completely **mitigate the problem of domain spoofing**. However, to be effective, both the domain owner and the mail transfer agent of the recipient should implement the extensions: the domain owner needs to correctly set SPF, DKIM, and DMARC rules, and the recipient has to authenticate incoming messages and correctly implement the verification of the SPF and DMARC rules.

In this paper, which is an extension of our previous work [7], we evaluate the extent of the SPF and DMARC deployment and analyze spoofing possibilities enabled by the absence or misconfiguration of their rules. We do not analyze DKIM as it requires access to the email header selector tag, not publicly available (see RFC 6376 for more details [5]).

While previous work already investigated the adoption of SPF and DMARC by the Alexa top-ranked one million domains [8], [9], we consider different datasets as well as threat models. We scan approximately **236 million domain names** including generic top-level domains (gTLD), country-code TLDs (ccTLD) and new gTLDs collected from different sources such as the Centralized Zone Data Service (CZDS)[1] made available by the Internet Corporation for Assigned Names and Numbers (ICANN), OpenData project from Rapid7[2] as well as public and available for download zone files (e.g, .se). The second dataset includes **32,042 high-profile domains** of 139 countries and their **defensive domain registrations**. The high-profile domains are the most popular targets of email spoofing: well-known companies, governmental websites, or financial institutions. To the best of our knowledge, this paper is the first study reporting on the global-scale measurement of the adoption of email authentication extensions.

We investigate the global adoption of SPF and DMARC protocols by scanning each domain in our datasets. Then, we

S. Maroofi, M. Korczyński, and A. Duda are with Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France.

A. Hölzel is with SMT (Simple Management Technologies), the Netherlands.

[1] https://czds.icann.org

[2] https://opendata.rapid7.com

introduce an attack vector in which attackers use subdomains (both existent and non-existent) for email spoofing. We also identify *defensively registered domains* and evaluate their adoption of email anti-spoofing schemes. We show that even if defensive registrations can mitigate some types of attacks like *cybersquatting* and *brand name abuse*, these domains need to be protected against domain spoofing as well.

We extend our previous work [7] and make the following main contributions:

1) we investigate the global adoption of SPF and DMARC for 236 million domain names of different TLDs,
2) in a separate measurement campaign, we evaluate the adoption of SPF and DMARC by top 500 most popular domains of 139 countries including local businesses, national websites, local governments, and financial sectors,
3) we propose a method to find defensively registered domains for top-ranked websites and assess the extent of their adoption of email security extensions,
4) we are the first to measure the extent of SPF and DMARC deployment by the subdomains of the top-ranked websites to gain better insight into how attackers can abuse subdomains to send spoofed emails,
5) we show that it is possible to send forged emails from non-existent subdomains when a DMARC rule is not strict enough regarding subdomains,
6) we demonstrate how syntactically wrong SPF rules may break the trust-based authentication system of selected email service providers by allowing forged emails to land in the user inbox,
7) we present a methodology for preventing domain spoofing based on good practices for managing SPF and DMARC records and analyzing DNS logs,
8) finally, as a proof of concept, we perform an end-to-end email spoofing for subdomains of high profile domain names with misconfigured SPF and/or DMARC.

This paper is an extension of our previous work published in the Network Traffic Measurement and Analysis Conference (TMA) 2020 [7]. Contributions (1), (7), and (8) are entirely novel compared to the previously published version of the paper. Furthermore, we extended the results of previous work by providing more insights on 1) common syntactically wrong SPF rules and 2) the analysis of the large scale notification campaigns, difficulties encountered, and their implications.

To remediate misconfigured SPF rules, we have contacted relevant Computer Security Incident Response Teams (CSIRTs) responsible for misconfigured domains and measured the effectiveness of our notifications. To encourage reproducibility, we make our measurement data available upon request.

The rest of the paper is organized as follows. Section II provides background on SPF and DMARC, followed by possible threat models regarding these protocols. Section III introduces our approach to generate the datasets and find defensively registered domains. Sections IV presents the analysis of the results for scanned domains and subdomains as well as for emulation of SPF rules. In Section V, we study the trust-based authentication issue and Section VI presents a methodology for preventing domain spoofing. Section VII describes our notification campaigns. Section VIII reviews related work and Section IX concludes the paper.

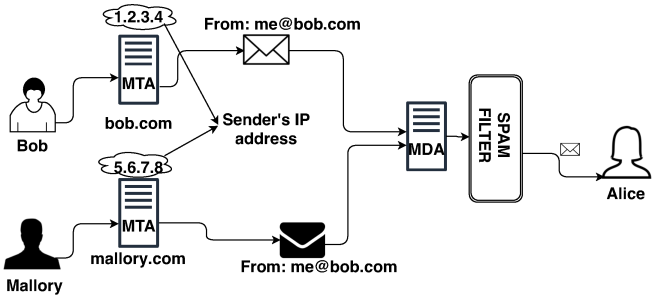

Fig. 1: Email sending and receiving procedure.

II. Background on Anti-Spoofing Schemes

To understand the issue of email authentication, we briefly explain the process of mail delivery. Figure 1 shows Bob (sender) who sends legitimate mails to Alice (receiver). Mallory (attacker) wants to send an email that impersonates Bob to Alice. Mallory and Bob use their respective servers (**`mallory.com`** and **`bob.com`**) to send mails. The Mail Delivery Agent (MDA) on the Alice server delivers two emails with the same sender address (**`me@bob.com`**) but coming from different IP addresses (assuming there is no spam filtering involved). One mail is from Bob (originated from the **`1.2.3.4`** IP address) and the other from Mallory (originated from **`5.6.7.8`**).

An effective anti-spoofing mechanism needs to differentiate the Mallory message from the legitimate Bob's mail. The current first lines of defence to protect end-users from spoofed emails include SPF [4], DKIM [5], and DMARC [6].

A. SPF – Sender Policy Framework

SPF is a set of text-form rules in **`TXT`** resource records of the Domain Name System (DNS). SPF specifies a list of servers allowed to send emails on behalf of a specific domain. During mail delivery over the SMTP protocol, the recipient server authenticates the sender Mail Transfer Agent (MTA) using a given **`HELO`** or **`MAIL FROM`** identity based on the published SPF record and the IP address of the sender—SPF needs to contain the domain portion of the **`MAIL FROM`** identity. In our example, the Alice server gets the **`TXT`** records of the **`bob.com`** domain from DNS. Then, it checks whether the sender IP address is on the list of IP addresses allowed to send emails from the **`bob.com`** domain and decides whether the message should be rejected or delivered to Alice.

The decision is made by the **`check_host`** function described in RFC 7208 [4] that takes three arguments on input (IP address of the sender, the domain, the **`MAIL FROM`** or **`HELO`** identity) and returns one of the seven possible results shown in Table I. The third column of the table presents the actions recommended by RFC 7208.

Below, we review the most common SPF rules useful for understanding the threat models presented in the next section (see RFC 7208 for more details). A valid SPF version 1

TABLE I: Possible results of the SPF **check_host** function and their definitions.

Result	Definition	Recommended action
None	1) No valid domain name was extracted from the SMTP session. 2) No SPF record was retrieved from the domain name.	1) The action must be the same as the *Neutral* output.
Neutral	1) There is no definite assertion (authorized or not) about the sender.	1) Depends on the receiver system.
Pass	1) Client is authorized to send emails with the given identity.	1) Whitelist the domain in terms of SPF.
Fail	1) Client is not authorized to send emails with the given identity.	1) Depends on the receiver system. 2) Make decision based on the DMARC policy.
Softfail	1) Client is not authorized to send emails with the given identity. 2) No strong policy specified by the domain owner.	1) Receiver should not reject the message. 2) May mark the message as suspicious.
Temperror	1) A temporary error occurred during retrieving the SPF policy.	1) May defer the message. 2) May deliver the message and mark it.
Permerror	1) Parsing problem in published SPF.	1) May deliver the message and mark it.

record must begin with string **v=spf1** followed by other SPF *mechanisms*, *qualifiers*, and *modifiers*. Mechanisms describe the set of mail servers for a domain and can be prefixed with one of four qualifiers: **+** (*Pass*), **-** (*Fail*), ~ (*SoftFail*), **?** (*Neutral*). If a mechanism results in a match, its qualifier value is used. *Pass* (i.e., **+**) is the default qualifier.

The most common SPF mechanisms are the following:

- **ip4** and **ip6** – they specify an address or a set of IPv4 (or IPv6) addresses to be matched by the **check_host** function with respect to the sender IP address.
- **a** and **mx** – they tell the **check_host** function to perform first a DNS lookup for **A** (or **MX**) records of a given domain and then compare the returned IP addresses with the IP address of the sender.
- **exists** – it indicates a DNS domain name used for a DNS **A** query. If the query returns any **A** record, this mechanism matches.
- **include** – it tells the **check_host** function to include the SPF rule of another domain in the evaluation, which may result in calling the **check_host** function recursively to fetch and analyze the SPF records of the included domains.
- **all** – it always matches, so its corresponding qualifier results in the final decision. For example, **v=spf1 mx -all** means: allow **MX** servers of the domain to send mail and prohibit all others.

The final result of the mechanisms could be *Match*, *No match*, or *Exception*. Qualifiers combined with mechanisms generate the final input for the **check_host** function that evaluates the SPF rule.

Modifiers provide additional information about the SPF records, for instance:

- **redirect=another-domain** – the SPF record for **another-domain** replaces the current record. The redirected domain becomes the target of all DNS queries and evaluations instead of the original domain.

Let us consider the following example:

```
    v=spf1 a ip4:1.2.3.0/24 -all
```

when the **A** record **example.com A 6.7.8.9** is stored in DNS. The SPF rule states that only machines with the IP address of **6.7.8.9** (the **a** mechanism) or with the IP address in the range of **1.2.3.0**...**255** (the **ip4** mechanism) are permitted senders (all others are forbidden). However, by only changing **-all** to **+all**, any machine is permitted to send emails on behalf of the domain **example.com** with the successful SPF *Pass* result.

B. DMARC

DMARC [6] builds on top of SPF and DKIM by explicitly stating the policies to apply to the results of SPF and DKIM. In particular, DMARC binds names checked by SPF with what is listed in the **FROM:** field of the mail header by means of *alignment*, which expresses the fact that these domain names should match (or partially match when using a relaxed setup). For instance, DMARC checks whether the name in the **MAIL FROM** SMTP command and the **FROM:** field of the mail header match or not. In the case of the alignment test failure, a DMARC policy can specify what to do with the message (accept, reject, or quarantine) and where to send reports in case of a mismatch. For a given domain name **domain.tld**, the DMARC policy is stored in the **TXT** record of **_dmarc.domain.tld**. Below, we present selected DMARC tags that an adversary can exploit when they are misconfigured:

- **aspf** (alignment mode for SPF) – it specifies whether the strict (**s** value) or relaxed (**r** value) alignment mode is required by the domain owner. The default value is the relaxed mode. In the strict mode, the domain name used in SPF must be the same as the domain used in the **FROM:** field of the header. In the relaxed mode, any subdomain of the domain can be used in the **FROM:** field of the header and will result in *Pass*.
- **p** (policy) – it specifies the action to be taken by the receiver if the alignment test results in *Fail*. Possible values for this tag are: 1) **none** – no specific action, 2) **quarantine** – the message is suspicious and depending on the mail system of the recipient, it could be delivered as spam, 3) **reject** – the domain owner wishes to reject emails during the SMTP transaction that fail the alignment test.
- **ruf** (reporting URI for failure) – it specifies the email addresses to which message-specific failure information is to be reported. This tag is important since it is the only bridge between the receivers and the true domain owners to fight spam emails [10].

- **sp** (subdomain policy) – it has the same syntax as **p** but applies to subdomains of the domain name. In the absence of this tag, the policy of the **p** tag must be applied to all subdomains [6]. If subdomains are not used to send emails, the owner can set this tag to the **reject** value to prevent subdomain email spoofing.

Let us assume the following DMARC rule of the domain **example.com**: **v=DMARC1; p=none; aspf=r;**. If we have the previously mentioned SPF rule for this domain, an illegal sender with the IP address of **9.10.11.12** can forge emails on behalf of **example.com** or any (existent or non-existent) subdomain of **example.com**, and the delivery decision is up to the receiver since no strict rule has been specified in DMARC. However, changing the DMARC rule to **v=DMARC1; p=quarantine; sp=reject; aspf=s;** tells the receiver to label all the emails that did not pass the SPF evaluation as spam and reject all the emails from the subdomains of **example.com** at the SMTP level.

C. Threat Models

We now consider threats regarding SPF and DMARC in detail. To mitigate mail spoofing, domain owners set up SPF and DMARC rules that are used by inbound mail servers. Therefore, if the recipient MTA does not support the SPF or DMARC check, no matter how strict the rules are, they will not be effective. A misconfigured SPF or DMARC (either syntactically or semantically) rule is as dangerous as the absence of the rules since the output of the evaluation does not lead to a correct decision.

We consider three possible types of threats:

- **Related to domain names**. If a domain uses a misconfigured SPF rule, then it is possible to send forged emails from any IP address with the SPF *Pass* result. For example, we have discovered that **microsoft.com.tr** used the **+all** mechanism in its SPF rule, which made it easy for attackers to send forged emails on behalf of Microsoft from any IP address. Note that after notifying Microsoft, the issue was fixed.
- **Related to subdomains**. Each subdomain should have its own SPF and DMARC rules. Another possibility is to use the **sp** tag in DMARC of the domain name (lower-level domain) to explicitly specify the action to take when receiving messages from subdomains. A possible abuse of subdomains is the following:
  - If a subdomain has no SPF rule (and there is no specified wildcard rule) and no explicit DMARC action, then it is possible to misuse the subdomain for sending forged emails. For example, while **icann.org** has a strict SPF rule, there is no rule specified in **account.icann.org** and no DMARC policy regarding subdomains (also the default action for domains is **none**, which in this case applies to subdomains). Hence, it is possible to send emails with forged sender addresses (e.g., **support@account.icann.org**) with the SPF *Neutral* result.
  - If a subdomain does not exist, the result of the DNS query for the **TXT** record returns a name error (NXDOMAIN). Thus, the **check_host** function returns the *None* result (see Table I). If there is no wildcard **TXT** record that covers non-existing subdomains and there is no DMARC policy specified for subdomains and the domain itself, then again, it is possible to send spoofed emails.
- **Wrong SPF rules**. If the **check_host** function cannot evaluate the existing SPF record of a domain name because of a syntax error, then the result is either *Temperror* or *Permerror*, and a legitimate email will likely arrive in the spam box. However, when the user marks this email as safe, the mail service may also accept spoofed emails from other IP addresses. We show in Section V how syntactically wrong SPF rules may break the trust-based authentication system of email service providers by allowing forged emails to land in the user inbox.

III. Methodology for Analyzing SPF and DMARC Deployment

In this section, we describe the methodology for analyzing the deployment of SPF and DMARC. We start with three datasets to perform two different measurements: in one campaign, we use a dataset of approximately 236 million domains from various resources to measure the global adoption of SPF and DMARC. In the second campaign, we use top 500 domains of 139 countries from the Alexa list [11] and online banking systems for all countries provided by FONDY.[3] In the second campaign, our focus is on high-profile domains (well-known companies, governmental websites, and financial institutions) and their defensive domain registrations.

A. Global Measurements

Regarding the global scan of domains for SPF and DMARC, we collected approximately 333 million domains from open zone files, OpenData project of Rapid7, and all the available zone files in Centralized Zone Data Service (CZDS) offered by ICANN. Our data consist of all domains with **.com**, **.net**, **.org**, **.biz** legacy generic TLDs (gTLDs), approximately 1,100 new gTLDs, **.se** and **.nu** country-code TLD (ccTLDs), operated by the Internet Foundation in Sweden, and samples of other domains obtained from Rapid7. Then, we scanned all the domains for **A** resource record using the ZDNS[4] scanner from the ZMap project [12] to keep only alive ones. Finally, our dataset consists of 235,960,991 active domain names in total. We performed the measurement in September 2020.

B. Top 500 Websites of All Countries

The Alexa website ranging system provides top 500 lists of most visited websites for 139 countries, which we collect for the purpose of this study as high-profile domains. Previous work [9], [13] used the Alexa top 1 million domains. However, we are interested in domains that may not be in the top 1M global popularity list but in the top list of each country, e.g., government websites or national businesses. In total, we

[3] https://fondy.eu

[4] https://github.com/zmap/zdns

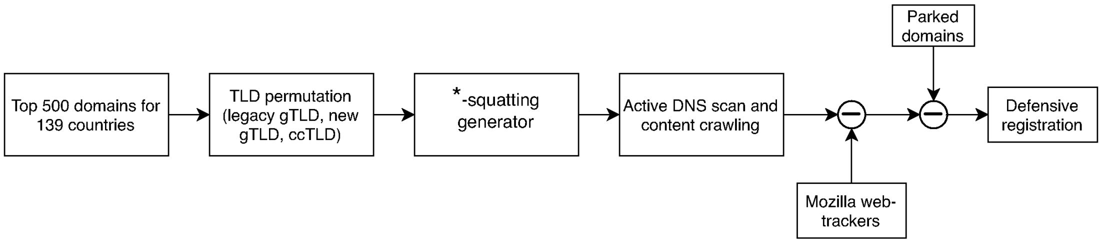
Top 500 domains for 139 countries
TLD permutation (legacy gTLD, new gTLD, ccTLD)
*-squatting generator
Active DNS scan and content crawling
Mozilla web-trackers
Parked domains
Defensive registration

Fig. 2: Generating the list of defensively registered domains.

collect 69,500 fully qualified domain names (FQDNs), which lead to 32,042 unique domains. Domain names are defined as 2nd–level, or lower-level if a given TLD operator provides such registrations, e.g., **example.br** or **example.com.br** [14]. We use a modified version of the public suffix list maintained by Mozilla[5] to get domains from FQDNs. For the purpose of this study, we exclude all private TLDs such as **s3.amazonaws.com** or **blogspot.com**. The dataset consists of 14,084 domains with legacy gTLDs, 1,070 domains with new gTLDs, and 14,084 domains with country-code TLDs. We refer to this list as the TOP500 list.

C. Defensive Registrations

Defensive registration refers to the process of registering domain names (often across multiple TLDs) with different grammatical formats to protect brands from attacks like *typo-squatting* [15]. For example, the **brand.com** company may register **brand.net** and **brand.org**, then redirect them to the original website. Figure 2 shows the algorithm to generate defensively registered domain names from the TOP500 list. We use the following procedure:

- For each domain name in the TOP500 list, we generate the domain names over all the possible TLDs including new gTLDs, legacy gTLDs, and ccTLDs. For example, for **paypal.com**, we generate **paypal.tld** where **tld** refers to all the ccTLDs (e.g, **paypal.in**), legacy gTLDs (e.g., **paypal.net**), and new gTLDs (e.g., **paypal.support**).
- For each domain in the TOP500 list that uses country code TLD or legacy gTLD, we generate *-squatting domains (for *-squatting, we use insertion, deletion, substitution, and internationalized domain names using DNSTwist package[6]). The list consists of 145,250,849 unique domain names.
- We scan all generated domains for **TXT** records with ZDNS. By excluding all DNS error results (e.g., NXDOMAIN, TIMEOUT, and SERVFAIL), we end up with 1,185,167 unique domains. Then, we extract the defensively registered domains based on the following three conditions:
  1) IP address in the requested **A** record of the domain is the same as for the **A** record of at least one corresponding domain in the TOP500 list,
  2) authoritative name server in the **NS** record of the domain is the same as in the **NS** record of at least one corresponding domain in the TOP500 list,
  3) domain part of the automatically visited domain homepage URL is the same as the one of the corresponding domain in the TOP500 list. As a result of this step, the list reduces to 235,508 domain names.
- Some of the domains in the list are related to web trackers [16] and parked domains. For parked domains, we exclude them using the method proposed by Vissers et al. [17], whereas for web trackers and advertising domains, we exclude them by using the Mozilla blacklist for trackers [18].

Our final list contains 55,059 defensively registered domains. For example, we find 226 domain names either registered by Google Inc. for **google.com** or by MarkMonitor[7] on behalf of Google, and 201 domain names related to PayPal Inc.

D. Subdomain Enumeration

We have generated the list of known subdomains for each entry of the TOP500 list using the Spyse[8] API. We only consider 'first-level' subdomains and exclude **www** and name servers since it is more likely that attackers use a first-level subdomain for sending spoofed email since it looks more legitimate. In total, we generate 212,361 subdomains for domains in the TOP500 list.

E. Banks and Financial Websites

For banking and financial websites, we leverage a list of 7,022 domains from the FONDY Github repository[9] and generate 39,310 subdomains using the same method as described in the previous section.

IV. Results on SPF and DMARC Adoption

After collecting all the datasets, we perform three types of scans for all domains and subdomains: 1) find **TXT** records to extract SPF rules, 2) find **TXT** records by prepending **_dmarc** to the domains and subdomains (i.e., **_dmarc.domain.tld**) to retrieve DMARC rules, and 3) analyze SPF and DMARC rules by emulating the **check_host** function [19] using our server IP address as the IP address of the sender (without actually sending emails).

[5] https://publicsuffix.org

[6] https://pypi.org/project/dnstwist/

[7] https://markmonitor.com

[8] https://spyse.com

[9] https://github.com/cloudipsp/all_banks_ips

TABLE II: Scan results for SPF rules.

dataset	total	norecord (%)	noerror (%)	servfail (%)	nxdomain (%)	timeout (%)
TOP500 domains	32,017	29.88	65.92	0.23	0.18	3.78
TOP500 subdomains	212,361	76.15	5.77	0.1	16.31	1.68
Bank domains	7,022	22.39	64.95	1.28	2.75	8.63
Bank subdomains	39,310	70.34	3.53	0.09	22.96	3.09
Defensive domains	55,095	1.2	95.37	0.43	1.03	1.97

TABLE III: Scan results for DMARC rules.

dataset	total	noerror (%)	servfail (%)	nxdomain (%)	timeout (%)
TOP500 domains	32,017	34.32	0.24	63.44	2.0
TOP500 subdomains	212,361	12.61	0.36	82.95	4.09
Bank domains	7,022	35.86	1.21	52.32	10.61
Bank subdomains	39,310	7.95	0.55	87.92	3.58
Defensive domains	55,095	40.08	0.36	57.86	1.7

In this section, we present the results of the scans introduced in Section III. We also study the spoofing possibilities of registered domains and subdomains by i) analyzing the scan results of SPF and DMARC records and ii) performing an end-to-end email spoofing for various email service providers.

A. Global Scan of the SPF and DMARC Rules

As the result of scanning 236 million domain names, we find that only 73,833,342 domains have SPF records set, which is approximately **31%** of all domains. The comparison of the obtained results with the scanning results of the top 1M domains in the Alexa list performed by Hu et al. [9] with 44.9% SPF adoption rate, shows that the global adoption of SPF is approximately **13.9%** lower than in the Alexa top 1M domains. We expected this result because Alexa top 1M domain names are more valuable and well-established in terms of DNS resource records, and therefore, they do not give a representative overall picture of the global SPF deployment.

Regarding DMARC, only 310,185 out of 236 million domains have DMARC corresponding to approximately **0.13%** of the population. For the domains with a DMARC rule, 41% of them have **`p=reject`**, 9.3% have **`p=quarantine`**, and 39.6% have **`p=none`** rule. These figures are also far different from the 5.1% of the domain names in the Alexa top 1M domains with DMARC rules [9], which again confirms that more popular domain names deploy email anti-spoofing schemes on a wider scale.

B. High-Profile Domains and Defensive Registrations

Tables II and III present the results of the scans using ZDNS to retrieve SPF and DMARC rules. Columns contain the following information: 'norecord' – domains exist but there is no SPF rule in the **`TXT`** record of the domains, 'noerror' – the record exists and can be retrieved successfully, 'servfail' – DNS lookup failure, 'nxdomain' – the domain name does not exist in the zone file, 'timeout' – the DNS timeout error. For DMARC, the 'nxdomain' column is the same as 'norecord' column for SPF (if we get 'NXDOMAIN' answer to the DNS query for **`_dmarc.domain.tld`**, it means that **`_dmarc`** subdomain does not exist so there is no DMARC rule).

We can notice in Table II that **29.9%** of the domains in the TOP500 list and **22.4%** of the online banking domains do not have SPF rules at all. As the **`check_host`** function for the domains without SPF rules returns *None* (see Table I), it is up to the receiver of the email to decide on whether to deliver a message and/or mark it as suspicious or not. While this behavior can be acceptable for regular domains, it is insecure for transactional domains (e.g., banking domains) as well as for high-profile domains (e.g., domains in the TOP500 list).

For defensively registered domains, Table II shows that only **1.2%** of them have no SPF rules, which is significantly lower than the results for TOP500 and banking domains. However, evaluating SPF alone is not sufficient since it is up to DMARC policies to make the final decisions about the delivery of messages.

As shown in Table III, as many as **63.4%** and **52.3%** of TOP500 and banking domains have no DMARC rule, which means that even with correctly configured SPF rules, it is still possible to spoof emails. Furthermore, for the domains with a DMARC rule in place (34.3% and 35.9% for TOP500 and banking domains, respectively), we have observed that a large part of them have the tag **`p`** equal to **`none`** (**60%** and **53.8%**, respectively, not shown in the table), which make them prone to email spoofing as well.

For defensively registered domains (see Table III), **57.9%** of them do not have a DMARC rule, which means that it is possible to send spoofed emails. Among 40.1% of the domains with a DMARC rule, **26.7%** have the **`p`** tag equal to **`none`** and 65% have the **`p`** tag set to **`reject`**, which makes them resilient to domain spoofing at the SMTP level.

Overall, we expect much wider deployment of SPF and stricter DMARC rules for defensively registered domains in comparison to high-profile domains—if organizations decide to register domains defensively to avoid domain name abuse, they are also more likely to configure the appropriate SPF and DMARC rules.

C. Analysis of Spoofing Possibilities for Subdomains

Regarding subdomains, the results are worse since **76.1%** of the subdomains related to the domains in the TOP500 list and **70%** of the subdomains related to banking websites do not have SPF records at all (see Table II). While it is not dangerous

TABLE IV: Specified DMARC action for subdomains with no SPF rule in the `TXT` resource record.

data	total	no-DMARC	`none`	`reject`	`quarantine`	invalid rule
TOP500-sub-no-SPF	161,720	108,535 (67.1%)	32,008 (19.7%)	13,286 (8.21%)	7,803 (4.82%)	88 (0.05%)
Bank-sub-no-SPF	27,650	19,070 (68.9%)	4,849 (17.5%)	2,682 (9.6%)	1,023 (3.69%)	26 (0.09%)

TABLE V: Result of the SPF `check_host` emulation.

Result	Global scan	TOP500	bank	defensive	bank subdomains	TOP500 subdomains
None	665,713	10,106	1,956	1,441	37,149	198,615
Neutral	6,340,566	1,497	236	6,220	56	683
Pass	213,112	50	10	114	2	37
Fail	25,040,843	7,083	2,268	22,255	860	4,511
Softfail	33,899,461	10,617	1,591	21,804	354	6,019
Temperror	1,474,437	135	155	523	778	1,485
Permerror	6,199,210	2,529	806	2,738	111	1,011
Total	73,833,342	32,017	7,022	55,095	39,310	212,361

in itself, the absence of strict DMARC rules for subdomains makes them prone to subdomain spoofing. To mitigate this vulnerability, domains need to provide appropriate DMARC rules. The **`sp`** tag (or **`p`** tag in the absence of **`sp`**) in a DMARC rule specifies the default action to be taken upon receiving messages from subdomains with no SPF rule [6].

Table IV shows the DMARC results for subdomains without SPF rules in both TOP500 and banking website lists. To obtain this result, we first scan **`_dmarc.sub.domain.tld`** to extract a **`p`** tag from each subdomain and in case of no DMARC rule in the subdomain, we scan **`_dmarc.domain.tld`** for **`sp`** or (in the case of its absence) **`p`** tags and apply the rule to subdomains (cf. RFC 7489 for more details [6]). In Table IV, **`none`**, **`reject`**, and **`quarantine`** columns correspond to the extracted rules as explained in Section II-B. The 'invalid rule' column refers to the rules that do not follow the syntax specified in RFC 7489 and 'no-DMARC' column corresponds to the domains without DMARC rules in subdomains nor in the domain name. Note that sending emails from a subdomain of any domain with 'no-DMARC' (**67.1%** for TOP500 and **68.9%** for banking websites), with **`none`** rule (**19.7%** for TOP500 and **17.5%** for banking websites), and 'invalid-rule' (less than 0.1% in both cases), regardless of the fact if the subdomain exists or not (non-existing subdomains), does not result in a strict reject decision. This behavior is potentially dangerous for transactional domains as it is possible to send emails with forged sender address using subdomains with no SPF record for as many as approximately **87%** of TOP500 and banking domains.

D. SPF Emulation Results

To analyze the validity of SPF rules using the **`check_host`** function further, we take advantage of **`pyspf`** [19] with our server IP address as the IP address of the mail sender. **`pyspf`** evaluates the SPF rule for a given domain and returns the SPF result. Table V shows the results of the SPF emulation for all domains (see also Table I for the definition of each result and the corresponding recommended action). The reason for the SPF *Pass* result is either the **`+all`** mechanism in the SPF rule or the possible **`redirect`** modifier. For the global scan of SPF and DMARC, 213,112 out of 73,833,342 domains result in SPF *Pass*, approximately 0.28% of the domains with SPF records. We also find that 6,199,210 domains (8.3% of the domains with SPF records) result in SPF *Permerror*.

Among the defensively registered domain names with the *Pass* result (114 domains), we have observed some well-known names like **`microsoft.com.tr`**[10] registered by MarkMonitor Inc.[7] on behalf of the Microsoft Corporation, as well as some major IT companies, local government, and TV channels websites for which we cannot provide the names for security considerations. However, the emulation results are available upon request. We have also noticed 12 different banking websites (1 in Spain and 11 in the United States) with the SPF *Pass* result. Although the number is fairly low, it is still enough for attackers to conduct a successful attack if they obtain the list of customer emails. In the TOP500 list for domains and subdomains, we have found 87 records with the SPF *Pass* result (50 for domains and 37 for subdomains) including several local government websites (mostly in the US), national financial websites, and national mobile operators with thousands of customers.

For as many as 7,195 high-profile domains and subdomains the SPF emulation results in *Permerror* (2.1% of the domains and subdomains with SPF records). As expected, this is less compared to the global domain name population (8.3%), as high-profile domains are more valuable and therefore are better configured. The majority of high-profile domains and subdomains have at least one of the following three problems: i) syntax problem in the published SPF rule (approximately 5,400 records), ii) excessive number of DNS lookups because of too many recursive **`include`** mechanisms [4] (1,131 records), and iii) published more than one valid SPF records (640 samples).

Table VI shows selected syntactically and semantically wrong published SPF records. We can observe that not only the syntax is important to parse an SPF record correctly, but also the number of DNS lookups must be limited to 10 queries based on RFC 7208 (cf. Section 4.6.4). Approximately 91% of the SPF *Permerror* results are related to only three types of misconfigurations with a 70% violation in the number of

[10] The issue was fixed after sending notifications.

TABLE VI: Selected syntactically wrong rules that lead to the *Permerror* result in SPF.

Error type	Example	Correct rule	Frequency
Too many DNS lookups	-	SPF rule must generate less than 10 DNS query	4,349,463 (70%)
Two or more SPF records	-	must set one SPF record for each domain	733,750 (12%)
No valid SPF record for included domain	-	must set one SPF record for included domains	556,811 (9%)
Unknown mechanism found: all.	v=spf1 a mx -all.	v=spf1 a mx -all	153,455 (2.5%)
Invalid IP4 address: ip4:	ip4:xxx.xxx.xxx.xx?all	ip4:xxx.xxx.xxx.xx ?all	72,011 (1.1%)
Empty domain:: a:	v=spf1 mx a: -all	v=spf1 mx a:example.com -all	18,190 (0.2%)

TABLE VII: Measurements of email delivery to inbox (IN), spambox (SP), or no delivery (ND) for five major email providers.

Threat model	Gmail			Yahoo			Outlook			Yandex			Laposte		
	IN	SP	ND	IN	SP	ND	IN	SP	ND	IN	SP	ND	IN	SP	ND
+all in SPF of domain	10	0	0	8	1	1	6	0	4	10	0	0	6	0	4
Defensive registration	9	1	0	9	1	0	3	7	0	9	1	0	9	0	1
Non-existent subdomain	8	2	0	3	0	7	2	8	0	10	0	0	10	0	0
Existent subdomain	7	3	0	7	2	1	4	6	0	10	0	0	10	0	0
Trust-based authentication issue		✗			✓			✓			N/A			✗	

DNS queries, followed by 12% of domains with more than one SPF record.

The domains and subdomains with *Permerror* are important because they may cause serious security problems. Since the domains have SPF records, *Permerror* indicates that they are used by their owners to send legitimate messages to users. However, emails may never get delivered or delivered but labeled as spam (based on the action recommended for *Permerror* as described in Table I). Importantly, we find that any attempt by the end user to detach the spam label from the legitimate email may whitelist all the emails from that domain name with the SPF *Permerror* result including forged emails (see Section V).

Moreover, a wrong implementation of the **`check_host`** function on the receiver without strict limitation of the number of DNS queries, may allow the attacker to put extra burden on the local recursive DNS resolver, which may lead to a Denial of Service (DoS) attack against the DNS server, as explained by Scheffler et al. [20]. Among the domains with syntactically wrong SPF rules, we observe some major IT companies e.g., **`eset.lu`**, the defensively registered domain for **`eset.com`** related to the ESET Internet Security.[10]

The SPF emulation results show that for several major IT companies, government websites, and one of the topmost banking website in the world, it is possible to send spoofed emails from both existent and non-existent subdomains as well as from some of their defensively registered domains due to weak or misconfigured SPF or DMARC rules.

E. End-to-End Spoofing Measurement

To show the possibility of email spoofing based on the different threat models presented in Section II-C, we have tested end-to-end email spoofing from well-known brands to our own registered email addresses at i) Gmail, ii) Yahoo, iii) Outlook, iv) Yandex, and v) Laposte email services. We follow the same steps as Hu et al. [9] to ensure research ethics. Table VII shows the test results. We have considered four different possibilities, namely, a) the SPF record of the domain has **`+all`** in its rule set, b) the defensively registered domain has neither an SPF nor DMARC rule to reject our emails, c) non-existent subdomains (e.g., **`accounts`**.icann.org), and finally, d) an existent subdomain without proper SPF configuration or a restrictive DMARC rule (e.g., **`account`**.icann.org). For ethical reasons, we do not provide the brand names of high-profile domains on behalf of which we sent emails, because for some of them, the problem is still unsolved.

We can observe in Table VII that in the first case (for which there is a **`+all`** in the SPF record), almost all the emails were delivered into the inbox by Gmail, Yahoo, and Yandex. Outlook and Laposte perform slightly better with 60% inbox delivery and 40% rejected emails. For the defensively registered domains, except for Outlook (with 70% delivered into the spam-box), all other email service providers successfully delivered almost all the sent mails into their inbox. Regarding non-existent subdomains, Outlook labeled 80% of the emails as spam while Yahoo rejected 70% of them. Other three services delivered almost all the emails. For the existent subdomain, Outlook performed the best by labeling 60% of the emails as spam. Surprisingly, Yandex delivered 97.5% of all sent emails into inbox, the worst performance in terms of the SPF and DMARC evaluation. The results show that attackers can successfully spoof all the tested email services by sending emails from non-existent subdomains, if domains do not have a strict reject DMARC policy.

V. Trust-based Authentication Issue

In this section, we show how a syntactically wrong SPF rule in a legitimate domain can push users to break the trust-based authentication system by labeling a legitimate email as safe and letting forged emails land in the user inbox. We examine five popular email providers: Outlook, Yahoo, Gmail, Laposte, and Yandex. We explain the issue using the Outlook service as an example, but the process is the same for other email service providers. Table VII presents the summary of results.

First, we register a domain (**`dnsabuse.xyz`**), set up a mail server, and the DNS **`A`** record of the domain. We use **`v=spf1 a aaaa -all`** as the SPF rule in the **`TXT`** record for our registered domain (i.e., syntactically wrong SPF rule because

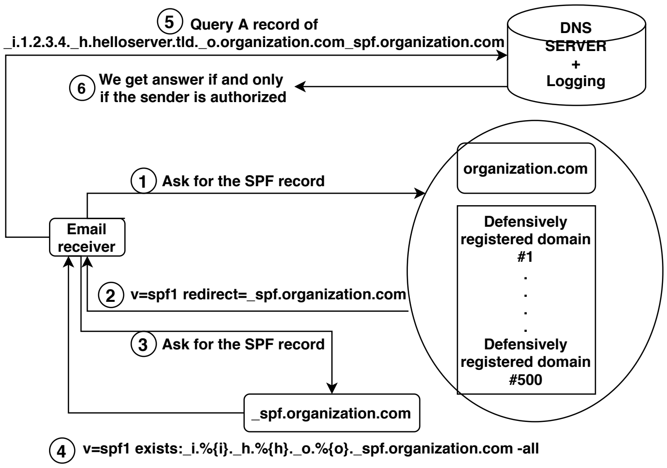
⑤ Query A record of
_i.1.2.3.4._h.helloserver.tld._o.organization.com_spf.organization.com
DNS SERVER + Logging
⑥ We get answer if and only if the sender is authorized
① Ask for the SPF record
organization.com
Email receiver
Defensively registered domain #1
Defensively registered domain #500
② v=spf1 redirect=_spf.organization.com
③ Ask for the SPF record
_spf.organization.com
④ v=spf1 exists:_i.%{i}._h.%{h}._o.%{o}._spf.organization.com -all

Fig. 3: Methodology for preventing domain spoofing.

of a nonexistent `aaaa` mechanism to generate the *Permerror* result). Then, we send a legitimate email with our server to our `outlook.com` email address. Since the SPF record is syntactically wrong and the reputation of our domain is low, the legitimate email lands in the spam box (as we expect) with the SPF *Permerror* result. If the user marks the email as *'safe sender'* (in case of Yahoo, the button label is *'add sender to contacts'*), then the Outlook service considers this email as safe (a correct assumption as it is a legitimate email). However, from now on, Outlook (as well as Yahoo) also accepts spoofed emails from other IP addresses that spoof the domain name.

We suspect that Yahoo and Outlook services whitelist the sender domain name instead of their IP addresses. On the other hand, the Laposte service rejects the sender with SPF *Permerror* at the SMTP level and sends a bounce message informing the sender about the reason for rejecting the mail (i.e., syntax error of SPF). We were not able to evaluate the trust-based authentication for Yandex since both emails (from the legitimate and illegitimate servers) land in the user inbox. Finally, Gmail does not suffer from the issue. We assume that when users detach the spam label from a legitimate email, Gmail only whitelists the IP address instead of the domain name.

VI. Methodology for Preventing Domain Spoofing

In this section, we present a methodology for preventing domain spoofing elaborated based on the experience gained in a study of a real-world scenario related to attacks performed on one of the government financial sectors in a European country. Due to security and ethical considerations, we do not give the name of the organization nor the name of the country.

The organization had one official registered domain (with ccTLD) and more than 500 defensively registered domains to protect the official one. In 2019, the domain administrators realized that a quite considerable number of attacks targeted their organization using different attack vectors: i) sending a forged email on behalf of the main domain, ii) sending emails with the `MAIL FROM` address of the defensively registered domains, and most importantly iii) sending emails from nonexistent subdomains of either the main domain or defensively registered domains. The main problem was that the targeted organization had no control over any part of the attack scenarios. They did not know anything about the sender, which could be the attacker or a compromised machine sending spoofed emails on behalf of the attackers, nor anything about the receivers of the emails. Thus, to solve the problem, not only they had to identify the sender but also inform possible recipients so that they do not accept incoming messages and potentially send a report related to these emails. Figure 3 illustrates the resulting methodology for preventing domain spoofing, a combination of good practices for managing SPF and DMARC records and analyzing DNS logs.

Assume that the IP address of the attacker is `1.2.3.4`, the host name used in the SMTP `HELO/EHLO` command is `helloserver.tld`, and the `MAIL FROM` field used in the spoofed email is `organization.com`, the same as the targeted brand. In this scenario, the SPF rule of the main domain, as well as all the defensively registered domain names, point to a single subdomain (`_spf.organization.com`) under the control of the organization using `redirect` modifier. When the receiver receives a spoofed email on behalf of `organization.com` (or of any defensively registered domain), it asks for the `TXT` record, retrieves the SPF rule of the domain (step ①), and gets the following answer: `v=spf1 redirect=_spf.organization.com` (step ②). In step ③, the receiver again asks for the SPF record of the specified domain name in the `redirect` modifier and receives a macro specified by the `exists` mechanism (step ④). The `exists` mechanism tells the receiver to create the domain name based on the specified rules and query the generated domain for the `A` resource record. The receiver can make the final decision based on step ⑥. If the domain name in step ⑤ resolves (no matter to which IP address), it means that the email is legitimate in terms of SPF. However, if the query returns no result (e.g., NXDOMAIN) not only the SPF will fail but also the DNS server logs the IP address of the attacker (or the compromised machine used by attacker) as well as the targeted domain (the main domain of the organization or one of the defensively registered domains). In addition, the receiver sends an extra `TXT` query for the DMARC policy to make the final decision about the received email. By specifying the `ruf` field in the DMARC rule, the domain administrators will receive a copy of the rejected email (e.g., phishing email) for further forensic analysis both to identify bugs in their mail software and gain better insight into the possible phishing/spam attacks on their domains. After one year of using the methodology to protect the targeted organization, the results show that this technique can effectively reduce the number of phishing attempts on the organization and help identify the malicious email senders.

VII. Remediation

Notifying the owners of the affected domains with misconfigured or missing SPF and DMARC rules is highly problematic since there is no straight way to retrieve the contact information of the domain owners [21], [22]. Public availability of the domain WHOIS data is affected by the introduction of the General Data Protection Regulation (GDPR) and "Temporary Specification for gTLD Registration Data" adopted by ICANN [23]. It obliges generic TLD registries and registrars to redact the Registrant and Administrative Contact in the public WHOIS.

Therefore, we decided to perform notifications through the Computer Security Incident Response Teams (CSIRTs). We use the following bottom-up approach to send notifications—we send email notifications if there is a CSIRT responsible for: 1) the domain name, 2) the TLD of the domain (mostly in case of private TLDs), 3) the IP range to which the IP address of the domain belongs to, 4) the autonomous system of the IP address for that domain, or 5) the national CERT responsible for the TLD (in case of country-code TLD) or the entire IP address space. We used this approach to perform two notification campaigns: the first one for high-profile domains, which are more critical to be fixed as soon as possible, in December 2019, and the second campaign related to the global scan in September 2020.

A. Results of the First Notification Campaign

Regarding high-profile domains, we have sent 128 emails to notify CSIRTs responsible for 7,653 domains with SPF *Pass* or *Permerror* results. We were not able to find any abuse contact address of responsible CSIRTs for 573 domains. For some high-profile domains prone to phishing attacks, e.g., **microsoft.com.tr**, we manually visited their websites and contacted them directly. In the first 5 days after sending notifications, we repeated our scans and found that 160 domain owners re-configured their SPF rules. The quickest clean-up action was initiated by the US government CERT (50 domains), national CERT of Austria (7 domains), Spain (7 domains) followed by CERT Polska, French CERT (ANSSI) and Danish CERT (CFCS-DK): 5 domains each.

Re-scanning the same set of domain names in October 2020 shows that 1,734 domain names changed their status from *Permerror* to *Softfail* (663 domains), *Fail* (569), *Neutral* (83), *None* (361), *Pass* (2), and *Temperror* (56). Moreover, 43 out of 152 high-profile domains changed their status from *Pass* to another status. Note that it is challenging to assess the effectiveness of our notification campaign because administrators may replace, for example, one misconfiguration by another (e.g., *Permerror* by *Pass*), however overall, after notifying CSIRTs responsible for misconfigured domains, as many as **23.2% (1,777 out of 7,653) were re-configured**.

B. Results of the Second Notification Campaign

Regarding the global scan, we found the total number of 6,412,322 misconfigured domains, 213,112 with SPF *Pass* results, and 6,199,210 with SPF *Permerror* results. For 23,116 domain names, we were not able to find any contacts to responsible CSIRT. Using the same above-mentioned notification approach, we sent emails to 110 CSIRTs. For some CSIRTs, due to the large size of the attachment files, we had to send two separate emails, one related to domains with SPF *Pass* and the other one related to SPF *Permerror*. Then, we re-scanned the domains every week to see how CSIRTs react to our notifications. After one week, we observed changes in the SPF results of 11,552 domains and another 917 domains after the second scan (we did not observe any major change after the third scan). For those domains that changed their SPF results, 567 changed from *Pass* to *Fail*, 56 domains to *Neutral*, 8,792 domains to *None*, 2,344 to *Softfail*, and others to *Temperror* and *Permerror*. We also did not observe any major changes in domains with *Permerror*. Overall, after notifying CSIRTs responsible for affected domains, **0.2% (12,469 out of 6,412,322) were re-configured.**

The differences between the remediation rates of the first and second campaign are to be expected and likely caused by: 1) the importance of vulnerable domains (high-profile domains are more likely to be fixed), 2) the magnitude of vulnerable domains (the number of vulnerable domains in the second campaign was three orders of magnitude larger). The magnitude of vulnerable resources is important since obtaining contact information at scale is highly problematic (for researchers, security companies, or CSIRTs), especially after the introduction of GDPR, and there is no alternative method suitable for large-scale notifications [22].

C. Notes on Notification Campaigns

We present below more insight into our notification campaigns and summarize major problems we encountered.

Figure 4 (see Appendix) shows the email template of the first notification campaign we sent to CSIRTs about vulnerable/misconfigured SPF records. Although we did not explain the problem in detail, we received many replies from the CSIRTs in the first 24 hours after sending notifications either thanking us for notifying them (we only consider manually typed emails rather than automatic replies) or with followup questions about the problem, e.g., whether we can prove it by sending a spoofed email. Figure 5 (see Appendix) shows one of the replies we received from one of the CSIRTs stating that they do not understand the problem and they think that the receiving MTA should be "smart enough" to handle *Permerror* responses. After providing the proof of concept, they notified the domain owners and fixed all the SPF records. On the other hand, in the second campaign, we used a more detailed email template and explained more about the problem (for each domain, we specified the reason for misconfiguration, i.e., *Permerror* or *Pass*).

Note that re-configuring domain names does not necessarily mean that the domain owners permanently solved the problem. As mentioned earlier, for 8,792 domains, the SPF result changed from *Pass* to *None*, which means that either the administrators removed the SPF record (possibly thinking that removing the record is better than setting a wrong one) or the domain just expired and was not registered anymore.

Sending large scale notifications present its own difficulties already discussed in previous work [21], [22], [24]. In addition to them, we encountered three major problems with sending emails to CSIRTs:

- Some countries do not have an official CSIRT to notify.
- Some CSIRTs do not have an officially published email address. Therefore, to notify them, one needs to fill an

online form on their websites making it impractical for large scale notifications.

- Finally, some CSIRTs changed their email addresses so that we received bounced emails.

Overall, our experience from the two notification campaigns shows that reporting vulnerabilities through CSIRTs can be effective but depends on its possible impact and magnitude of affected resources.

VIII. Related Work

In this section, we review previous work on measuring and analyzing email security extensions.

Durumeric et al. [25] measured the adoption of SMTP security extensions and their impact on end users. They studied SMTP server configurations for the Alexa top one million[11] domains and SMTP connections to and from Gmail gathered over a year. They reported the existence of a long tail of over 700,000 SMTP servers, of which only 35% successfully configure encryption, and only 1.1% specify a DMARC authentication policy.

In 2017, Durumeric [8] measured the extent of SPF and DMARC adoption for one million top domains in the Alexa list. His results showed that 40.1% of the domains have published SPF records while only 1.1% of them have valid DMARC records. Hu and Wang [9] reported similar statistics in 2018 with the results of 44.9% published SPF records and 5.1% published DMARC records showing approximately 5% of increase in one year. In their end-to-end experiment, they spoofed 30 high-profile domains and reported the ratio of emails that reached inboxes of selected email providers. We perform a similar analysis for both SPF and DMARC records but in two different phases. First, we analyze the global adoption of SPF and DMARC rules for different TLDs and then, we focus on more prominent domains (with transactional emails) including banking websites, government portals, national and international businesses as well as defensively registered domains and their subdomains. We also consider end-to-end spoofing but just as a proof of concept for our defined threat models and only for 10 high-profile domains.

Foster et al. [13] evaluated the security extensions using a combination of measurement techniques to determine whether major providers support the Transport Layer Security (TLS) protocol [26] at each point in their email message path, and whether they support SPF and DKIM on incoming and outgoing mail. They reported that while the use of SPF is common, enforcement was limited. Scheffler et al. [20] investigated the consequence of a wrong implementation of the **`check_host`** function at the receiver, which lets attackers perform denial-of-service (DoS) attacks on a local DNS resolver. While our goal is not to evaluate the SPF abuse, we show that 4,349,463 domains in the global scan, 1,131 high-profile, and defensively registered domains have published SPF records that require more than 10 DNS lookups. Therefore, such misconfigured records may lead to abuse of local DNS resolvers.

Finally, Hu et al. [27] investigated the reasons behind the low adoption rates of anti-spoofing protocols. They conducted a user study involving email administrators and showed that they believe the current protocol adoption lacks the crucial mass due to the protocol defects, weak incentives, and practical deployment challenges.

IX. Conclusion

It is paramount for high-profile domains and defensively registered domains to establish appropriate SPF and DMARC policies to reduce the chance of successful spear phishing attacks. In this paper, we evaluate the adoption of the SPF and DMARC security extensions by domain names in two phases and analyze spoofing possibilities enabled by the absence of their rules or their misconfigurations. The results show that a large part of the domains do not correctly configure the SPF and DMARC rules, which enables attackers to successfully deliver forged emails to user inboxes. In particular, we show that for top 500 domains of 139 countries, the adoption rate of SPF and DMARC records are 65.9% and 34.3%, respectively. For banking websites, we obtain almost the same results (64.9% and 35.9%) as for the TOP500 list. However, for defensively registered domains, the results are significantly higher especially in terms of published SPF records with 95.37% adoption and 40.1% for DMARC. We also, for the first time, investigate the problem of subdomains in the anti-spoofing techniques and their possible abuse to send forged emails.

We also emulate the SPF **`check_host`** function not only to evaluate *Pass* and *Fail* results but also obtain all the possible results such as *Permerror*, *None*, and *Neutral* for both domains and subdomains. The investigation shows that syntactically wrong SPF rules may break the trust-based authentication system of email service providers (e.g., Outlook and Yahoo) by allowing forged emails to land in the user inbox. To improve deployment of SPF and DMARC, we have presented a methodology for managing SPF and DMARC records and analyzing DNS logs that may prevent domain spoofing.

For remediation, we have sent the total of 238 emails to notify the CSIRTs responsible for 6,419,975 domains. Within the first two weeks after the notification campaigns, they managed to inform domain owners and re-configure SPF records of 12,629 vulnerable/misconfigured domains. More importantly, as many as 23.2% of high-profile domains were re-configured at the end. Our experience shows that disclosing vulnerabilities through CSIRTs can be effective, especially for valuable domain names. Finally, while we do not publish the scan data because of ethical concerns, we make the data available upon request to encourage reproducibility.

Appendix

In this section, we present the email template of the first notification campaign we sent to CSIRTs about vulnerable/misconfigured SPF records (see Figure 4) and the exchange of mails with one of the CSIRTs that led to fixing the misconfigured SPF records (see Figure 5).

Acknowledgments

This work has been partially carried out in the framework of the COMAR project funded by SIDN, the .NL Registry

[11] https://www.alexa.com/topsites

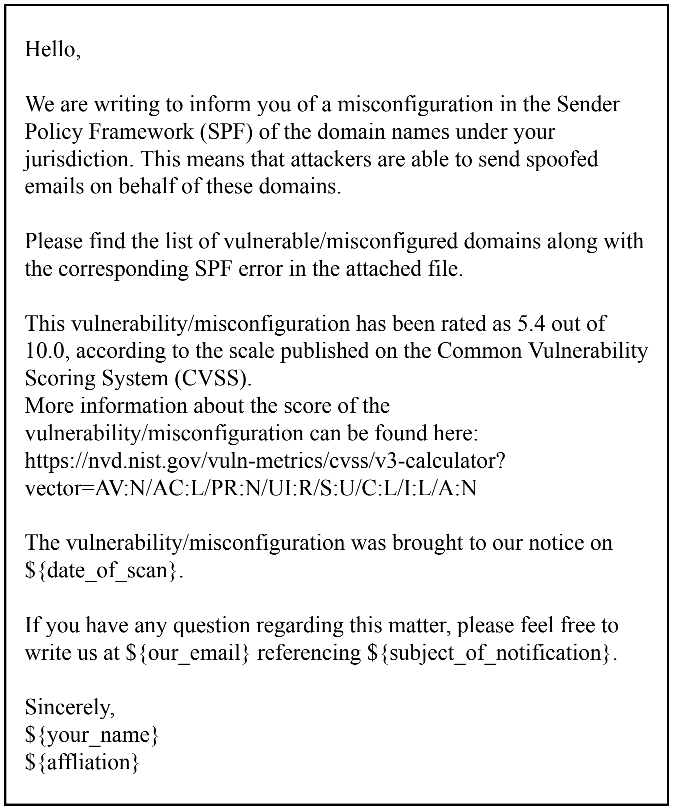

Hello,

We are writing to inform you of a misconfiguration in the Sender Policy Framework (SPF) of the domain names under your jurisdiction. This means that attackers are able to send spoofed emails on behalf of these domains.

Please find the list of vulnerable/misconfigured domains along with the corresponding SPF error in the attached file.

This vulnerability/misconfiguration has been rated as 5.4 out of 10.0, according to the scale published on the Common Vulnerability Scoring System (CVSS).
More information about the score of the vulnerability/misconfiguration can be found here:
https://nvd.nist.gov/vuln-metrics/cvss/v3-calculator?vector=AV:N/AC:L/PR:N/UI:R/S:U/C:L/I:L/A:N

The vulnerability/misconfiguration was brought to our notice on ${date_of_scan}.

If you have any question regarding this matter, please feel free to write us at ${our_email} referencing ${subject_of_notification}.

Sincerely,
${your_name}
${affliation}

Fig. 4: Content of the email for the first campaign.

and AFNIC, the .FR Registry. It was also partially supported by the PrevDDoS project funded by IDEX UGA IRS and the ANR projects: the Grenoble Alpes Cybersecurity Institute CYBER@ALPS under contract ANR-15-IDEX-02, PERSYVAL-Lab under contract ANR-11-LABX-0025-01, and DiNS under contract ANR-19-CE25-0009-01.

References

[1] (2019) Internet Crime Report. [Online]. Available: https://www.fbi.gov/news/stories/2019-internet-crime-report-released-021120

[2] (2019) A Phishing Campaign Reports Unusual Activity on Your Microsoft Account. [Online]. Available: https://www.logitheque.com

[3] J. Klensin, "RFC 5321: Simple Mail Transfer Protocol," Internet Requests for Comments, 2015. [Online]. Available: http://tools.ietf.org/html/rfc5321

[4] S. Kitterman, "RFC 7208: Sender Policy Framework (SPF) for Authorizing Use of Domains in Email," Internet Requests for Comments, 2014. [Online]. Available: http://tools.ietf.org/html/rfc7208

[5] D. Crocker, T. Hansen, and M. Kucherawy, "RFC 6376: DomainKeys Identified Mail (DKIM) Signatures," Internet Requests for Comments, 2011. [Online]. Available: http://tools.ietf.org/html/rfc6376

[6] E. Kucherawy, M. Zwicky, and E. Zwicky, "RFC 7489: Domain-Based Message Authentication, Reporting, and Conformance (DMARC)," Internet Requests for Comments, 2015. [Online]. Available: http://tools.ietf.org/html/rfc7489

[7] S. Maroofi, M. Korczyński, and A. Duda, "From Defensive Registration to Subdomain Protection: Evaluation of Email Anti-Spoofing Schemes for High-Profile Domains," in *Proc. Network Traffic Measurement and Analysis Conference (TMA)*, 2020.

[8] Z. Durumeric, "Fast Internet-Wide Scanning: A New Security Perspective," Ph.D. dissertation, University of Michigan, 2017.

[9] H. Hu and G. Wang, "End-to-End Measurements of Email Spoofing Attacks," in *Proc. 27th USENIX Security Symposium*, 2018, pp. 1095–1112.

[10] (2018) Dmarc overview. [Online]. Available: https://dmarc.org/overview/

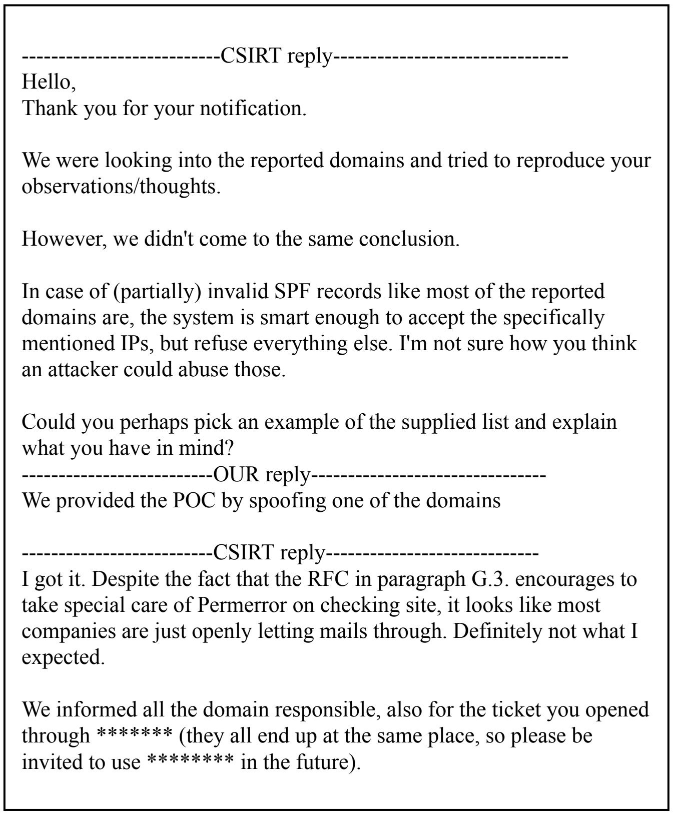

--------------------------CSIRT reply--------------------------------
Hello,
Thank you for your notification.

We were looking into the reported domains and tried to reproduce your observations/thoughts.

However, we didn't come to the same conclusion.

In case of (partially) invalid SPF records like most of the reported domains are, the system is smart enough to accept the specifically mentioned IPs, but refuse everything else. I'm not sure how you think an attacker could abuse those.

Could you perhaps pick an example of the supplied list and explain what you have in mind?
-------------------------OUR reply--------------------------------
We provided the POC by spoofing one of the domains

-------------------------CSIRT reply-----------------------------
I got it. Despite the fact that the RFC in paragraph G.3. encourages to take special care of Permerror on checking site, it looks like most companies are just openly letting mails through. Definitely not what I expected.

We informed all the domain responsible, also for the ticket you opened through ******* (they all end up at the same place, so please be invited to use ******** in the future).

Fig. 5: Sequence of the emails we have exchanged with one of the CSIRTs.

[11] (2019) The Top 500 Sites on the Web. [Online]. Available: https://www.alexa.com/topsites/countries

[12] Z. Durumeric, E. Wustrow, and J. A. Halderman, "ZMap: Fast Internet-Wide Scanning and its Security Applications," in *Proc. 23rd USENIX Security Symposium*, 2013, pp. 605–620.

[13] I. D. Foster *et al.*, "Security by Any Other Name: On the Effectiveness of Provider Based Email Security," in *Proc. 22nd ACM CCS Conference*. ACM, 2015, pp. 450–464.

[14] M. Korczyński, S. Tajalizadehkhoob, A. Noroozian, M. Wullink, C. Hesselman, and M. van Eeten, "Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs," in *Proc. Euro S&P*, 2017, pp. 579–594.

[15] J. Szurdi, B. Kocso, G. Cseh, J. Spring, M. Felegyhazi, and C. Kanich, "The Long "Taile" of Typosquatting Domain Names," in *Proc. 23rd USENIX Security Symposium*, 2014, pp. 191–206.

[16] S. Schelter and J. Kunegis, "Tracking the Trackers: A Large-Scale Analysis of Embedded Web Trackers," in *Proc. 10th International AAAI Conference on Web and Social Media*, 2016.

[17] T. Vissers, W. Joosen, and N. Nikiforakis, "Parking Sensors: Analyzing and Detecting Parked Domains," in *Proc. NDSS Symposium*. Internet Society, 2015, pp. 53–53.

[18] (2019) Shavar Tracking Protection Lists. [Online]. Available: https://github.com/mozilla-services/shavar-prod-lists

[19] (2019) Python SPF Package. [Online]. Available: https://pypi.org/project/pyspf/

[20] S. Scheffler, S. Smith, Y. Gilad, and S. Goldberg, "The Unintended Consequences of Email Spam Prevention," in *Proc. PAM Conference*. Springer, 2018, pp. 158–169.

[21] O. Cetin, C. Ganan, M. Korczyński, and M. van Eeten, "Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning," in *Workshop on the Economy of Information Security*, 2017.

[22] S. M. Wissem Soussi, Maciej Korczyński and A. Duda, "Feasibility of Large-Scale Vulnerability Notifications after GDPR," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2020.

[23] ICANN. (2018, May) Temporary Specification for gTLD Registration Data. ICANN. [Online]. Available: https://www.icann.org/resources/pages/gtld-registration-data-specs-en

[24] F. Li, Z. Durumeric, J. Czyz, M. Karami, M. Bailey, D. McCoy, S. Savage, and V. Paxson, "You've Got Vulnerability: Exploring Effective Vulnerability Notifications," in *USENIX Security*, 2016.

[25] Z. Durumeric *et al.*, "Neither Snow Nor Rain Nor MITM...: An Empirical Analysis of Email Delivery Security," in *Proc. ACM IMC Conference*. ACM, 2015, pp. 27–39.

[26] E. Rescorla, "RFC 8446: The Transport Layer Security (TLS) Protocol Version 1.3," Internet Requests for Comments, 2018. [Online]. Available: https://tools.ietf.org/html/rfc8446

[27] H. Hu, P. Peng, and G. Wang, "Towards Understanding the Adoption of Anti-Spoofing Protocols in Email Systems," in *Proc. SecDev Conference*. IEEE Computer Society, 2018, pp. 94–101.

Sourena Maroofi is a PhD student at Université Grenoble Alpes, France. He obtained his master's degree from Sadjad University of Technology, Mashhad, Iran, and his bachelor's degree in telecommunication engineering from Shahid Bahonar University of Kerman, Iran. He is currently working on different aspects of DNS. His main research interests include large-scale measurements of the Internet and DNS-related topics.

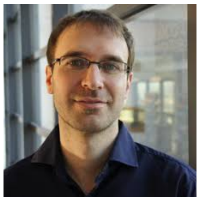

Maciej Korczyński is an Associate Professor at Grenoble Institute of Technology. He received the PhD degree in computer science from the Grenoble Alps University (2012). He was a post-doctoral researcher at the Rutgers University, USA (2013-2014) and Delft University of Technology, the Netherlands (2014-2017). His main research interests include Internet-wide passive and active measurements for cybersecurity, domain name abuse, incident data analysis, vulnerability notifications, economics of cybersecurity, and security of Internet protocols, with a focus on DNS.

Arnold Hölzel studied Information and Communication Technology in Amsterdam, the Netherlands. He is now Senior Security Consultant working with SMT, a data-driven solution provider located in the Netherlands. He is mostly operating within governmental Security Operations Centers. He has been doing all sorts of security-related work for about 15 years, but sees it more as an out-of-control hobby. He loves to analyze and dig through multiple terabytes of data to find that one outlier of hidden treasure. He always tries to broaden his knowledge, either through training or self-study.

Andrzej Duda is a Full Professor at Grenoble Institute of Technology. He received his PhD from the Université de Paris-Sud and his Habilitation diploma from the Grenoble University. Previously, he was an Assistant Professor at the Université de Paris-Sud, a Chargé de Recherche at CNRS, and a Visiting Scientist at the MIT Laboratory for Computer Science. In 2002-2003, he was an Invited Professor at EPFL (Swiss Federal Institute of Technology in Lausanne). He published over 180 papers in the areas of performance evaluation, distributed systems, multimedia, and networks.