BreakSPF: How Shared Infrastructures Magnify SPF Vulnerabilities Across the Internet

by Wang *et al.*, NDSS '24

2024. 4. 4. | MMLAB seminar

Summarized by Subin Song (sbsong@mmlab.snu.ac.kr)

Table of Contents

Introduction

- **SPF (Sender Policy Framework)** is a protocol to prevent email spoofing attacks, by specifying IP addresses allowed to send emails from the domain
- Idea of *BreakSPF* attack
  - Step 1. Obtain IP addresses from cloud service providers, proxy services, etc.
  - Step 2. Find domains with SPF records that allow the IP addresses obtained in Step 1
  - Step 3. Send spoofing emails from IP addresses & domains found in previous steps
- Such vulnerabilities are magnified due to shared infrastructures of SPF records (i.e., `include`, `redirect` mechanisms)

Background

Email Spoofing Attack

- **SMTP** (Simple Mail Transfer Protocol) *doesn't* have a built-in method for "from" address authentication
- Therefore, attackers can forge "from" addresses when sending emails
  - "`MAIL FROM`" in SMTP envelope
  - "`From`" in SMTP header
- Defense: Authentication chain
  - **SPF**, DKIM, DMARC, ARC

    i.e., Sender Policy Framework / DomainKeys Identified Mail / Domain-based Message Authentication, Reporting and Conformance / Authenticated Received Chain

Image source: https://www.proofpoint.com/us/corporate-blog/post/how-does-email-spoofing-work-and-why-it-so-easy

SPF (Sender Policy Framework)

- An IP-based email authentication standard to prevent spam, spoofing, and phishing
- An *SPF record* lists IP addresses that are approved to send emails from the current domain
  - Provided as a DNS TXT record
- The receiving mail server checks the sender's IP address against the SPF record of the received email's "`MAIL FROM`" domain

SPF Record Examples

```
;; QUESTION SECTION:
;mmlab.snu.ac.kr.               IN      TXT

;; ANSWER SECTION:
mmlab.snu.ac.kr.        0       IN      TXT     "google-site-verification=OTGWOX0gv7glWEsfdh8MaH7zdMt-csC-SOeNbCowfMg"
mmlab.snu.ac.kr.        0       IN      TXT     "google-site-verification=p1U8pRdjb1rAa0aOPFEvv76Sa8paxrvQc6WXcVojxF8"
mmlab.snu.ac.kr.        0       IN      TXT     "v=spf1 ip4:147.46.114.27 include:_spf.google.com ~all"
mmlab.snu.ac.kr.        0       IN      TXT     "google-site-verification=ah5kAHfQo1sLKxYp9uNygxfx5iCTDiAn9l9dsxWB9kA"
mmlab.snu.ac.kr.        0       IN      TXT     "v=DKIM1; k=rsa; t=y; p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCaf5DO9dWks
6aa0V+0UFmJ6DhEGzuB8fB7Ioehtx5/SMMr4H2oUHO7dQ56QcA01v0dwJ3GWhP/KkgOVu5amtn432BT4cnytzhPrFzB5NI76crExEnZNG4lH7TIZcTTKbY/vW
9YYBjY4u8VDCuzQV6jzjlNIo9R6A+ZExdQoyKlJwIDAQAB"
```

```
;; QUESTION SECTION:
;snu.ac.kr.                     IN      TXT

;; ANSWER SECTION:
snu.ac.kr.              0       IN      TXT     "google-site-verification=3C9g514zUg8bKZUQdrwHpN2bzdMcTbAo-TokW6z2aRk"
snu.ac.kr.              0       IN      TXT     "google-site-verification=Nza5O8ADLimi4lyjisAF5uIpWYbeo26sve28MeH_IqU"
snu.ac.kr.              0       IN      TXT     "ZOOM_verify_sInFFu9lRcSsWQAe7XJYnw"
snu.ac.kr.              0       IN      TXT     "google-site-verification=fY6W7mOoeSYYmxaVljDkpRBVFeSqKs3R-mk9b1A2c20"
snu.ac.kr.              0       IN      TXT     "v=spf1 include:_spf.snu.ac.kr include:_spf2.snu.ac.kr include:_spf.go
v-dooray.com include:_spf.google.com include:spf.protection.outlook.com ~all"
```

SPF Record Format

- Starts with “v=spf1”, consists of multiple elements in `[qualifier]mechanism:value` format
  - Qualifier : `+` (pass) | `-` (hard fail) | `?` (neutral) | `~` (soft fail)
  - Mechanism : `all` | `include` | `redirect` | `ip4` | `ip6` | `mx` | …

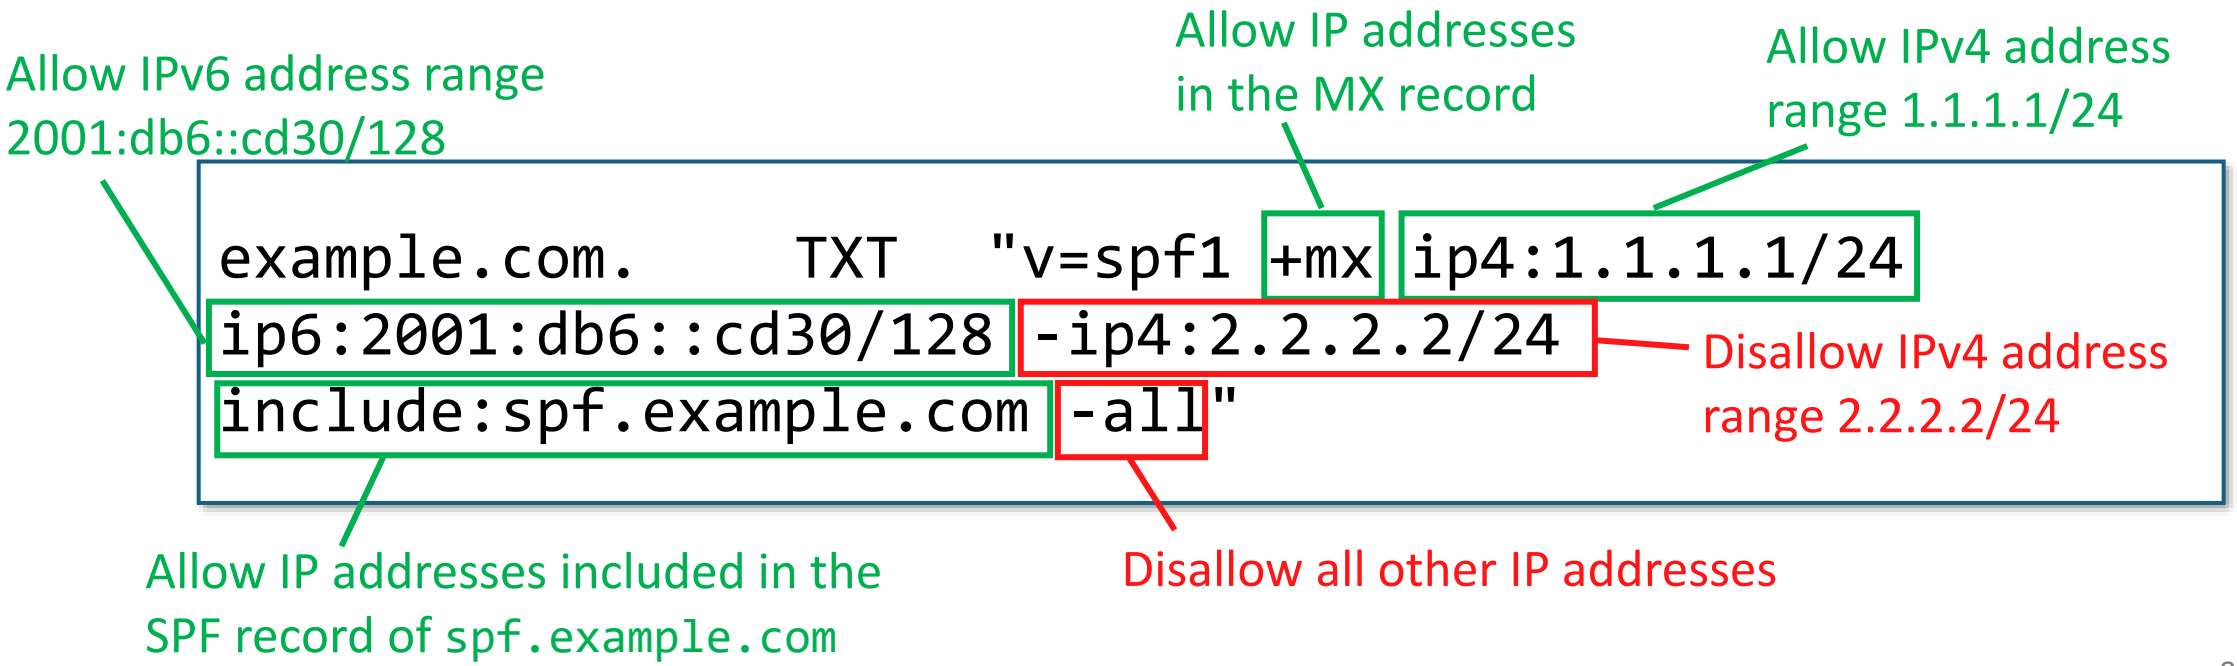

BreakSPF Attack Explained

BreakSPF Attack Model

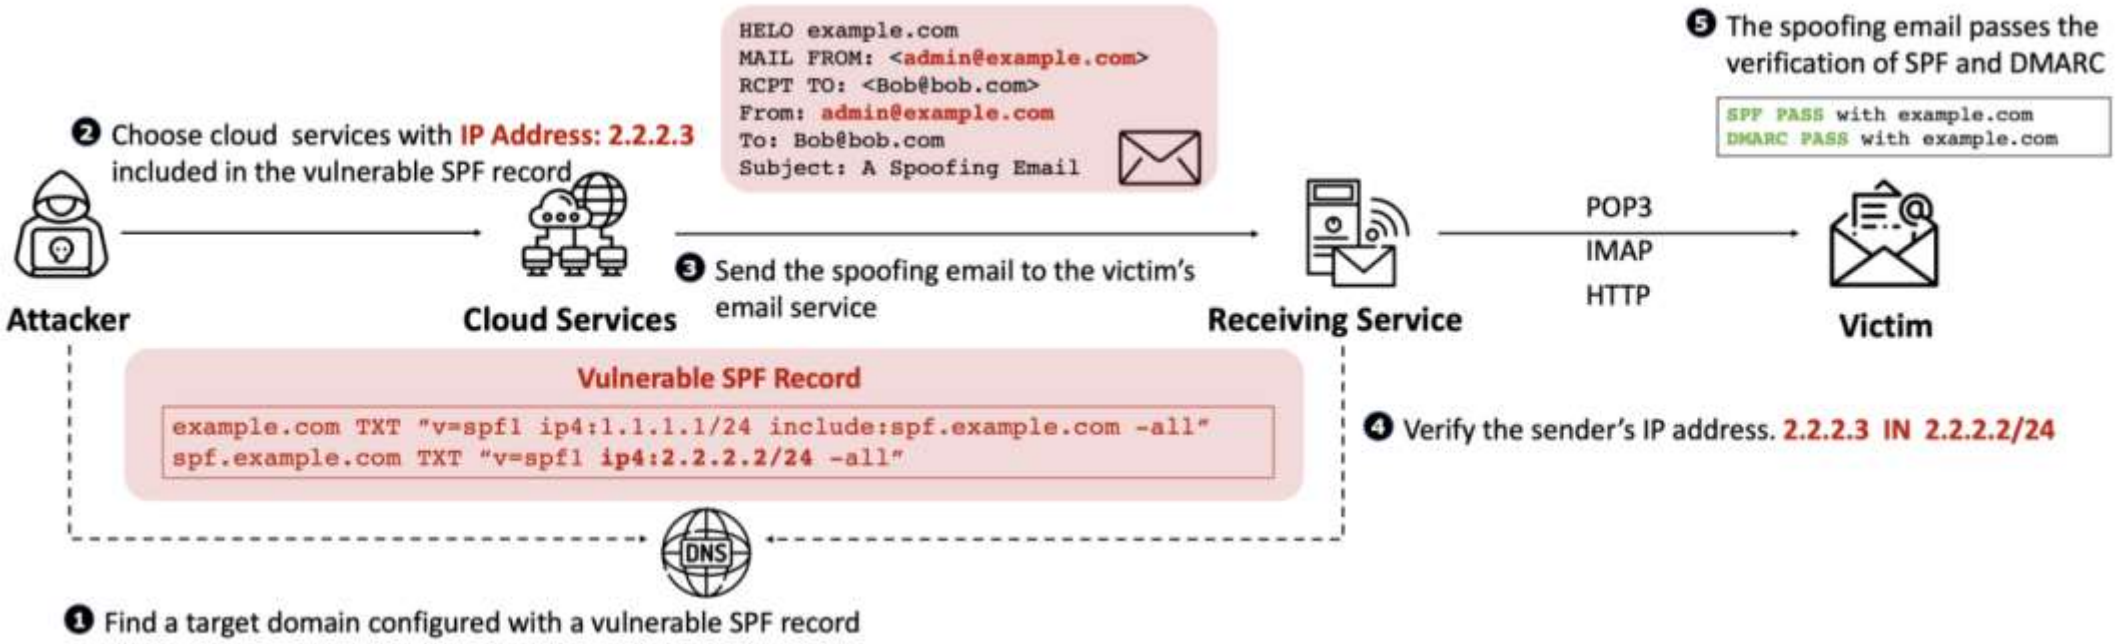

BreakSPF Attack Model

BreakSPF Attack Model

BreakSPF Attack Model

BreakSPF Attack Model

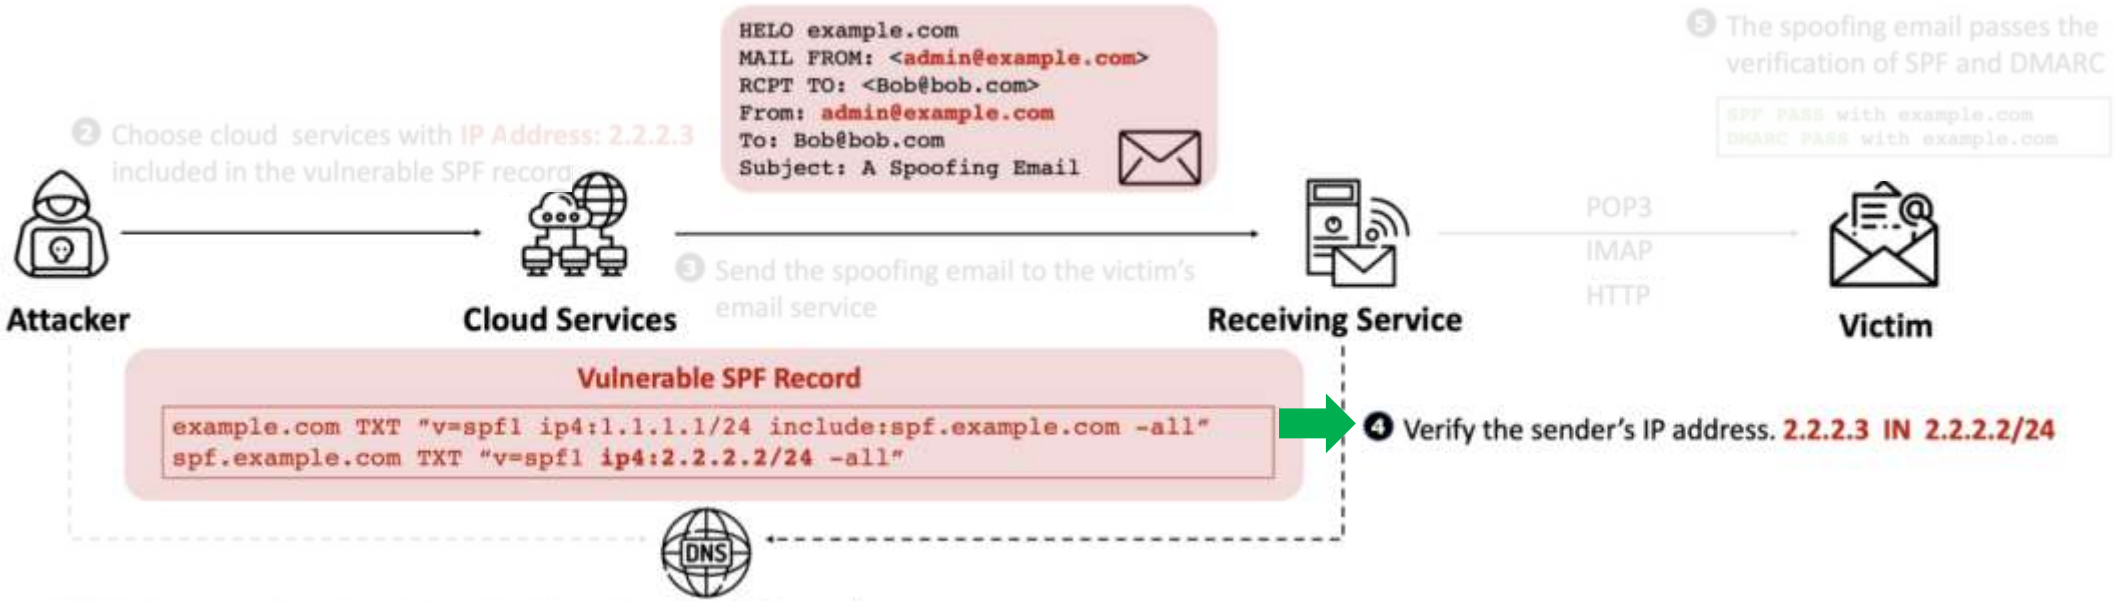

BreakSPF Attack Model

Cross-Protocol Email Spoofing Attack

- Leverages the similarities of HTTP and SMTP
  - Header-body structure
  - Usage of MIME* headers

    * Multipurpose Internet Mail Extensions

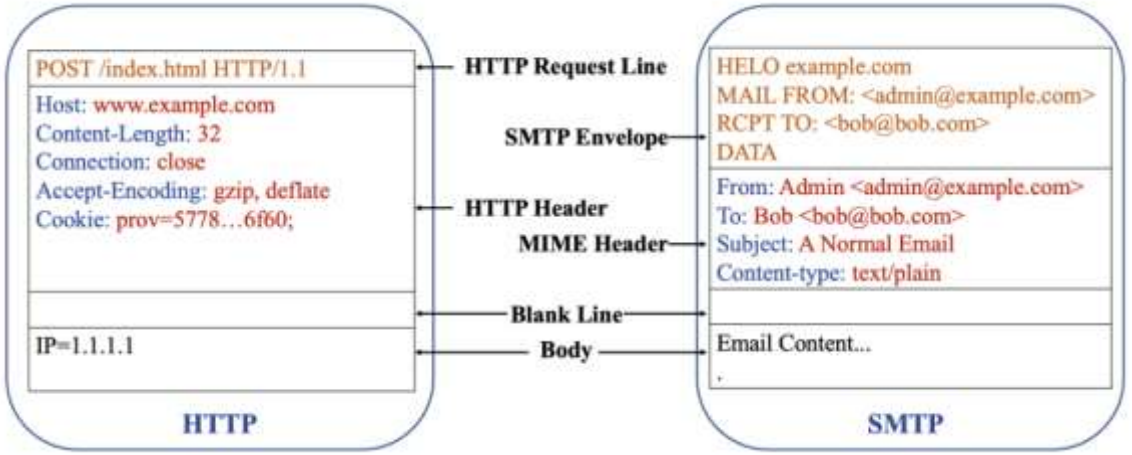

- Use HTTP forwarding services (e.g., HTTP proxy, CDN) to send email packets
  - Email servers are fault-tolerant to some extent, so they can treat HTTP header fields as unidentified SMTP commands and ignore them

Exploitation Workflow

Step 1. Domain collection

- Collect total of 7,183,870 domains (Tranco Top 1M domain names and their subdomains)

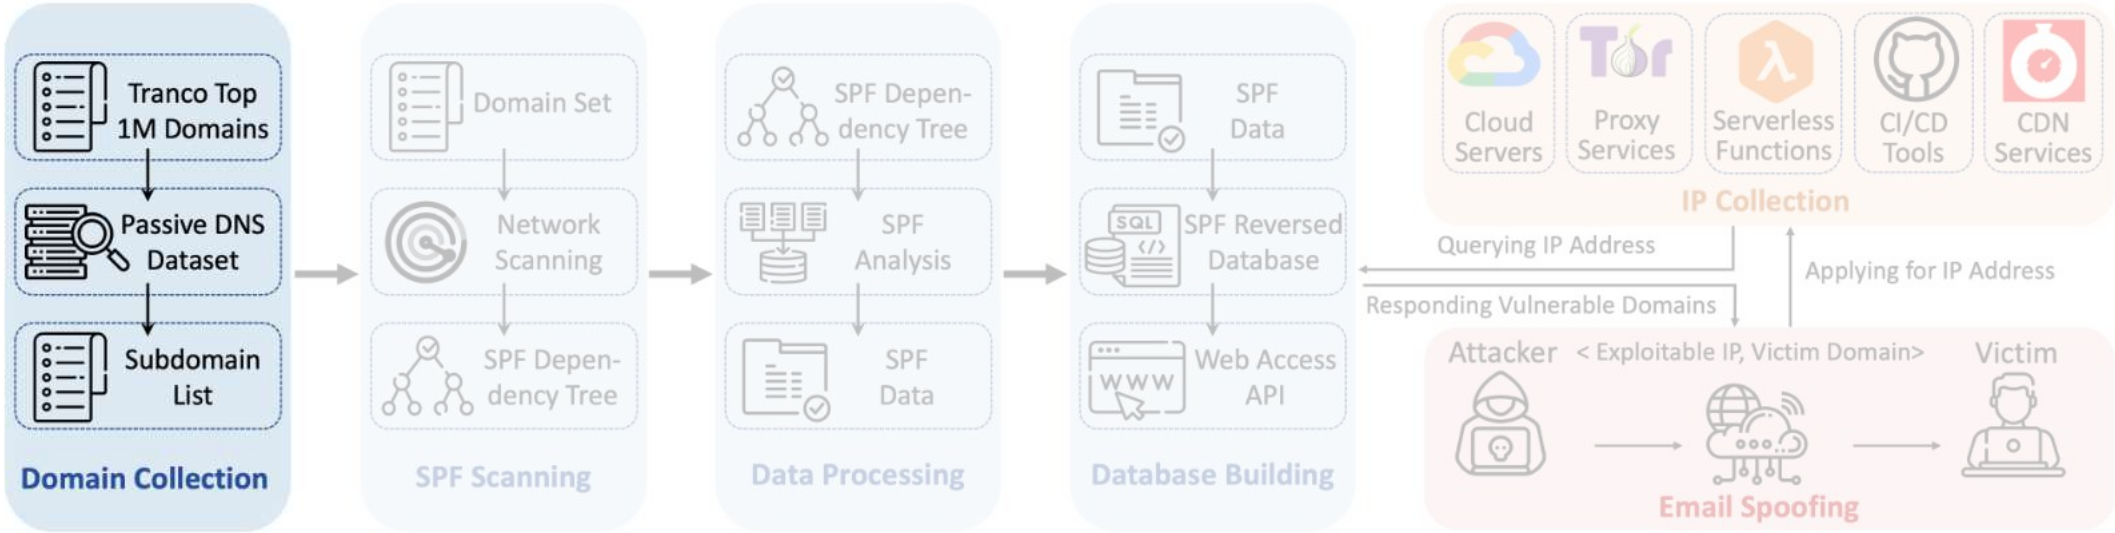

Exploitation Workflow (Cont'd)

Step 2. SPF scanning

- Query the TXT records of the domains & filter out SPF records
- Then, build a SPF dependency tree, based on the `redirect` and `include` relationships in the SPF records

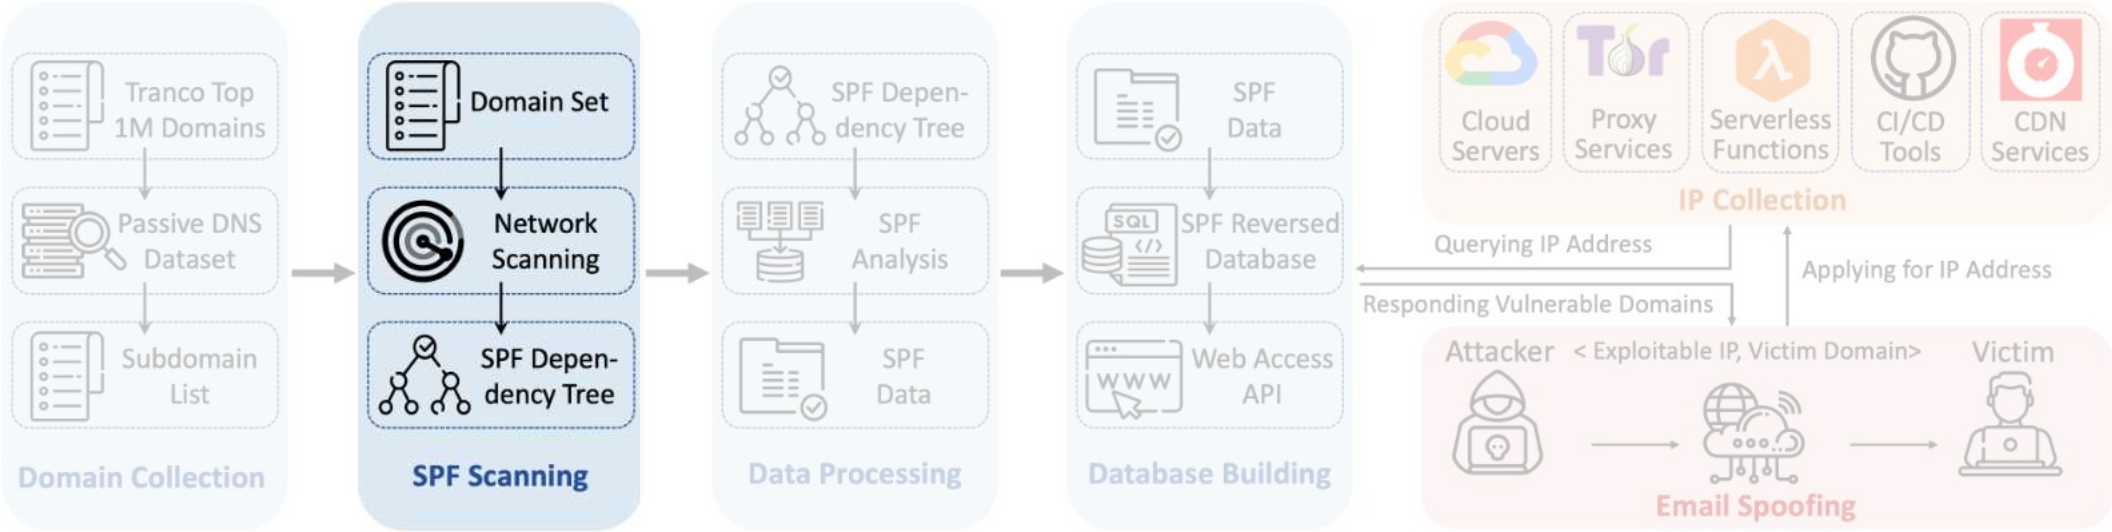

Exploitation Workflow (Cont'd)

Step 3. Data processing

- Analyze the results of SPF scanning
  - e.g., Adoption rate of SPF, grammatical analysis, `include` mechanism analysis, IP coverage of SPF records

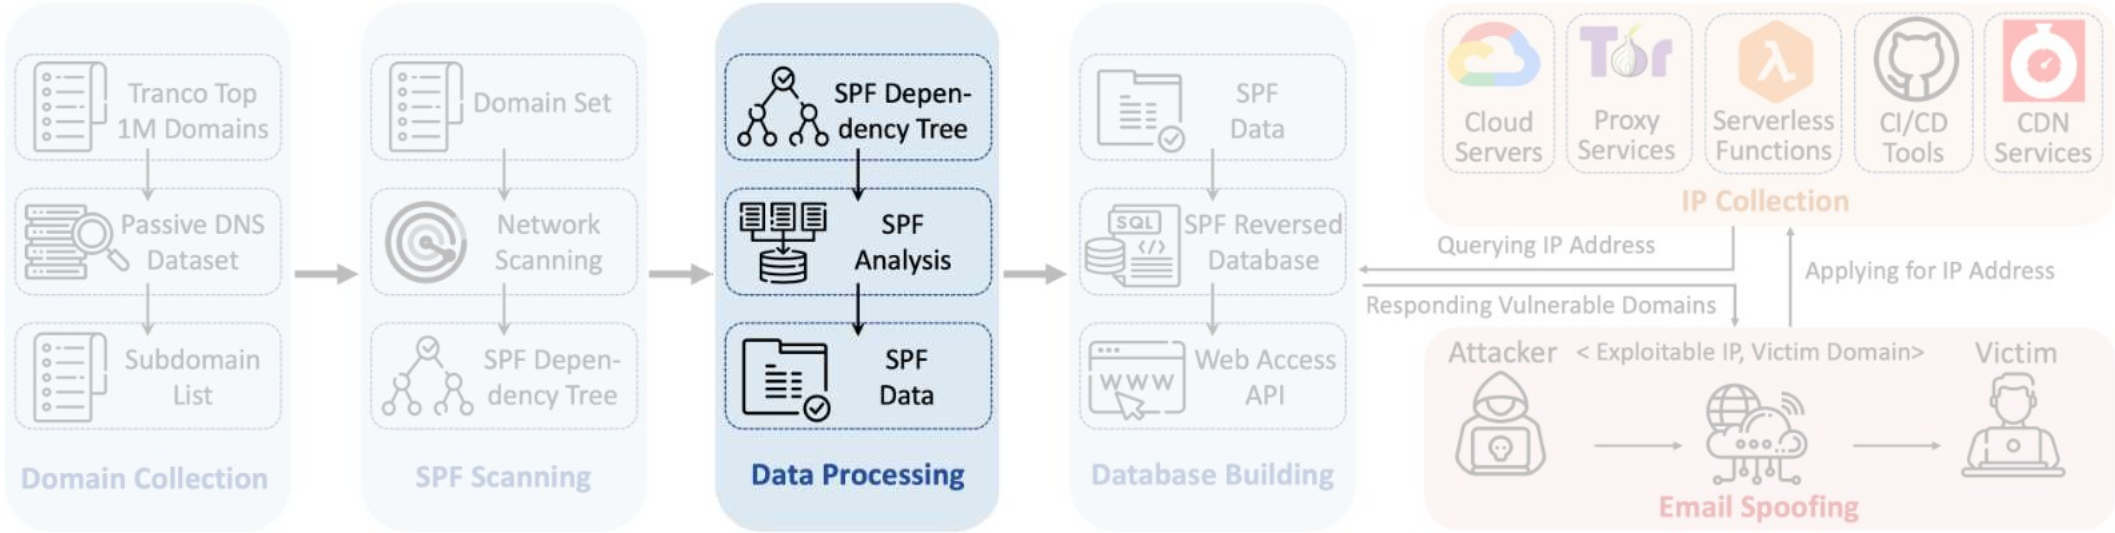

Exploitation Workflow (Cont'd)

Step 4. Database building

- Construct a SPF reverse lookup database (i.e., given an IP address, we can find out which domains include the IP address in their SPF records & which other domains depend on those domains)

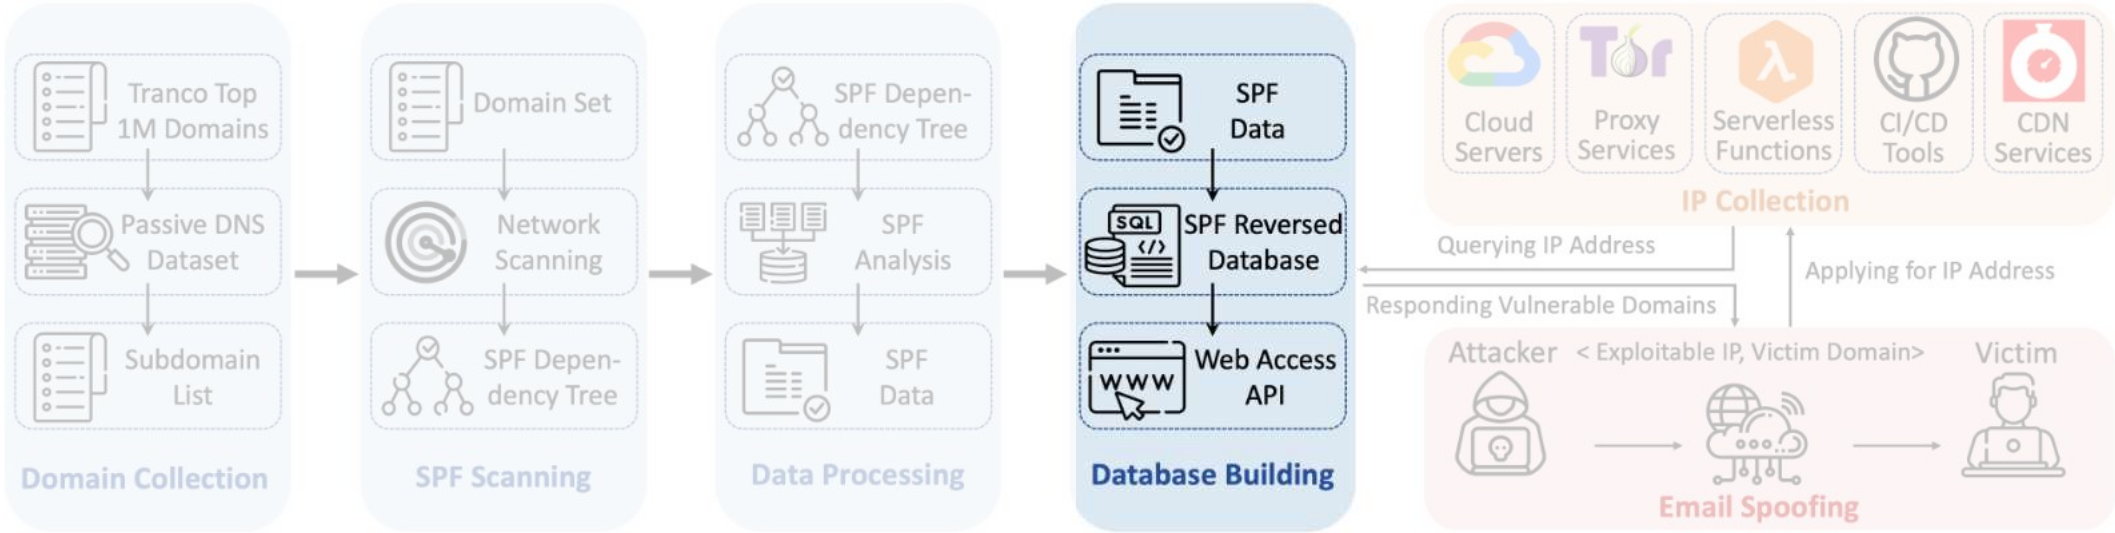

Exploitation Workflow (Cont'd)

Step 5. IP address collection

- Try to obtain as many IP addresses as possible
  - Cloud servers, proxy services, serverless functions, CI/CD tools, CDN services
- Then, use the previously constructed SPF reverse lookup database to identify domain names vulnerable to spoofing using these IP addresses

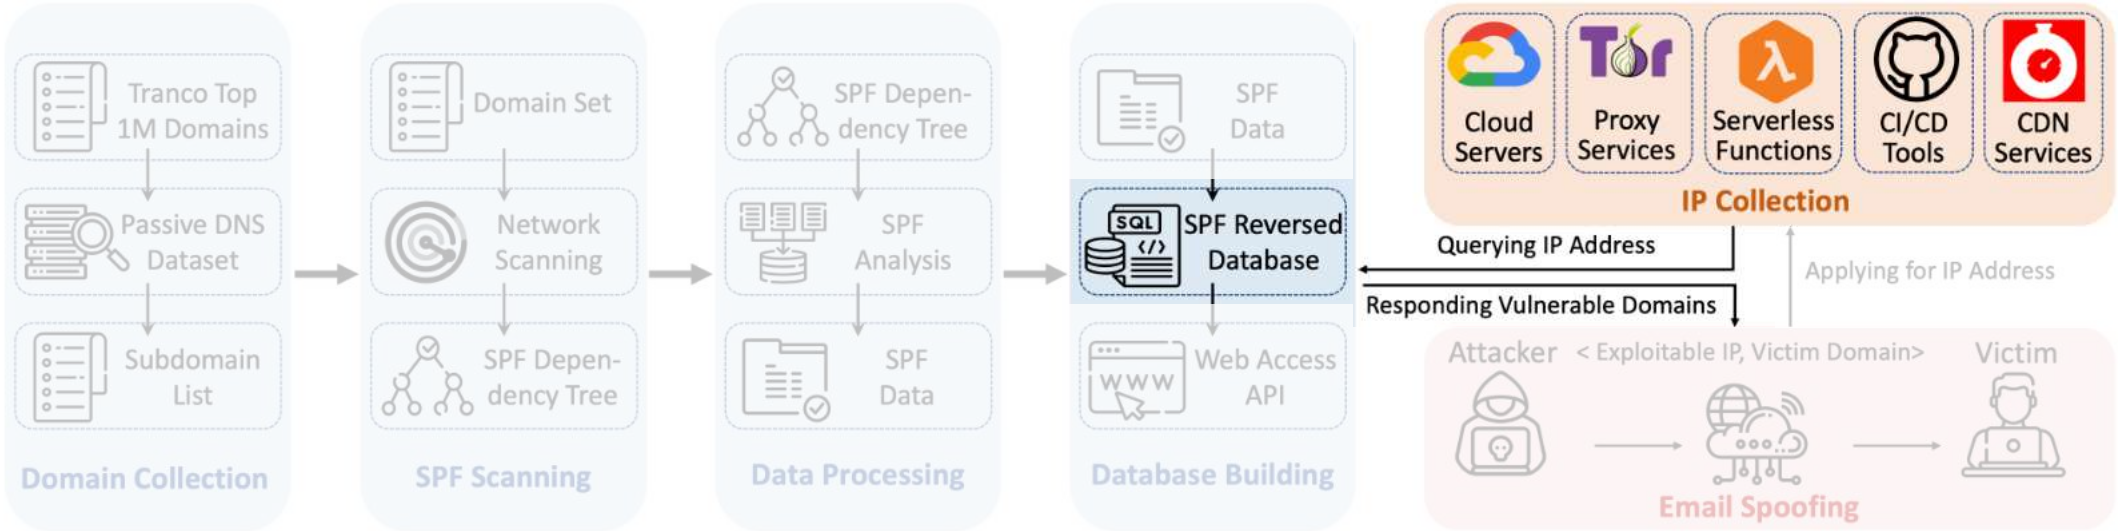

Exploitation Workflow (Cont'd)

Step 6. Email spoofing attack!

- Conduct cross-protocol email spoofing attacks using IP addresses collected in the previous step

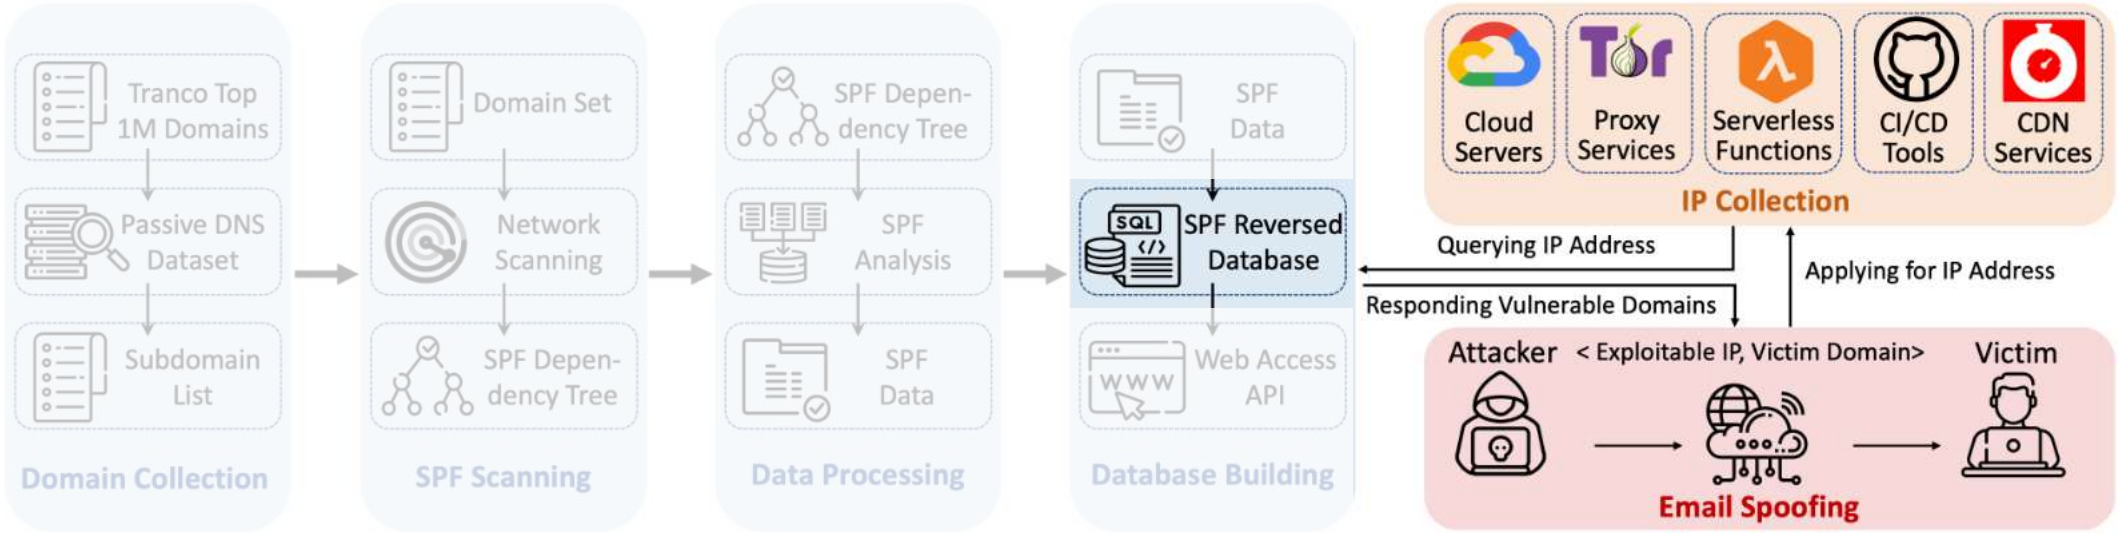

Results

SPF Deployment Analysis Results

- Adoption rate of SPF
  - **79.4%** of email domains* have SPF records
  - **72.7%** of email domains* have *valid* SPF records

* Domains with MX records, or that provide email services on port 25 in their A records

Status	Top1M Domains # (%)	Email Domains[1] # (%)
Total domains	1000000 (100.0 %)	738310 (100.0 %)
w/ SPF	609,236 (60.92 %)	586,316 (79.41 %)
w/ valid SPF	559,296 (55.93 %)	536,976 (72.73 %)
Soft Fail	311,277 (31.13 %)	305,326 (41.35 %)
Hard Fail	205,181 (20.52 %)	189,984 (25.73 %)
Neutral	25,997 (2.60 %)	25,266 (3.42 %)
Pass	742 (0.07 %)	670 (0.09 %)
w/ Include	417,144 (41.71 %)	410,899 (55.65 %)
w/ Redirect	13,737 (1.37 %)	13,520 (1.83 %)

- Grammatical analysis of SPF records
  - **8.4%** of SPF records have grammar errors
    - **63.2%** of them are "too many DNS lookups"
      - More than 10 DNS queries per resolution → error!
    - **30.7%** of them are "multiple SPF records per domain"

Misconfiguration Type	# Domain	%
Too Many DNS Lookups	32,254	63.15%
Double SPF Records	15,700	30.74%
Format Errors	2,838	5.56%
Spelling Errors	986	1.93%
Coexisting `all` and `redirect`	612	1.20%
Total	51,076	100.00%

SPF Deployment Analysis Results (Cont'd)

- `include` mechanism analysis
  - **73.5%** of domains with SPF records contain `include`
  - **20%** of all SPF records recursively include `outlook.com`, and **15.7%** include `google.com`

- IP coverage of SPF records
  - **51.7%** of SPF records have more than 655,536 (2^{16}) IP addresses included

Rank	Email Providers	# Included	%
1	outlook.com	181,544	20.07%
2	google.com	142,317	15.73%
3	amazonses.com	44,466	4.92%
4	sendgrid.net	44,200	4.89%
5	mandrillapp.com	38,437	4.25%
6	mcsv.net	38,260	4.23%
7	mailgun.org	34,790	3.85%
8	zendesk.com	30,869	3.41%
9	mailchannels.net	20,837	2.30%
10	salesforce.com	20,692	2.29%

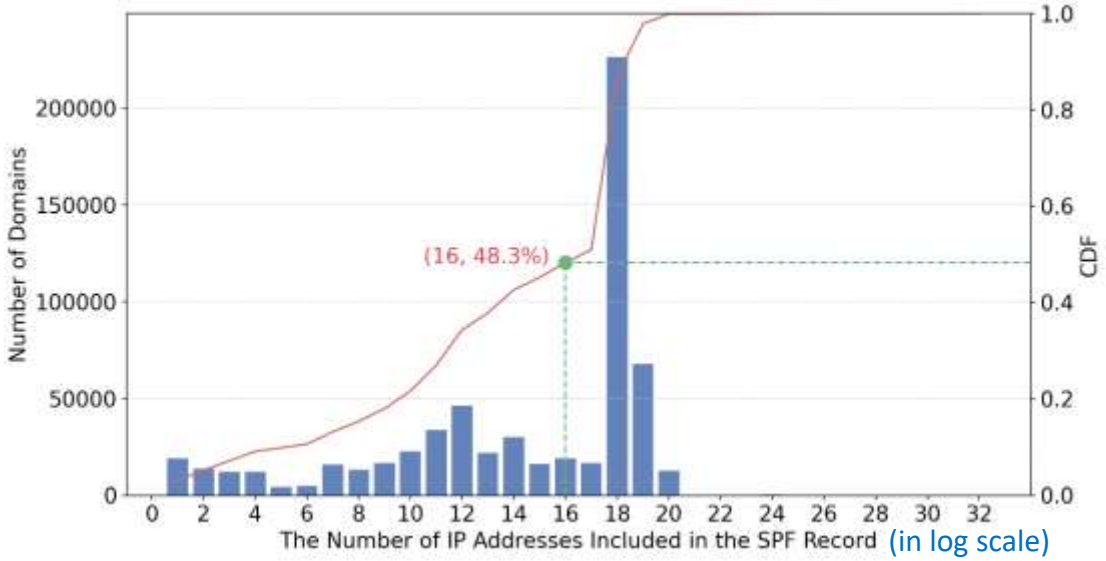

Shared IPs Collection Results

- Obtained total of **87,430** IP addresses
  - Cloud servers, proxy services, serverless functions, CI/CD platforms, CDN services
- Low cost
  - On average, less than $0.01 per IP address
  - This is because most service providers offer free tiers & credits

Shared IPs Collection Results (Cont'd)

of IP addresses that are included in *some* domain's SPF record

Services		IP Obtained	Unique IPs	Successful Hit	IP diversity /8	IP diversity /16	IP diversity /24	IP diversity ASN	Port 25	Port 465
Cloud Servers	Alibaba	1,028	909	887	19	55	721	2	◑	●
	Amazon	9,680	9,679	8,788	21	449	7,304	2	◑	●
	Azure	33,580	30,498	6,255	22	376	10,998	1	◑	●
	Digitalocean	987	976	967	34	55	822	1	●	●
	Google	1,036	216	216	7	88	215	1	◑	●
	Linode	1,017	989	977	28	45	426	1	●	●
	Tencent	1,009	996	944	25	65	730	2	◑	●
	Vultr	307	282	277	31	46	232	1	◑	●
Proxy Services	VPN	389	339	309	102	282	306	101	◑	●
	Open Proxy	68,653	3,061	13,704	189	1,811	2,713	1,985	●	●
	RESIP	30,000	23,876	22,468	193	8,063	16,533	2,851	●	●
	Tor	1,213	1,208	1,068	108	378	592	238	◑	◑
Serverless Function	Alibaba	3,269	39	33	4	13	33	2	●	●
	Amazon	100	3	1	2	3	3	1	●	●
	Azure	1,879	13	0	1	3	4	1	●	●
	Baidu	60	3	3	2	2	3	1	●	●
	Google	46	4	4	2	2	4	1	●	●
	Huawei	234	6	6	5	5	6	3	●	●
	Tencent	7,398	62	32	8	9	38	2	●	●
CI/CD Platforms	Circleci	4,446	377	329	13	147	372	1	●	●
	Github	5,000	3,648	1,388	14	148	2,578	1	●	●
	Vercel	3,209	3,198	2,196	4	50	2,405	1	●	●
CDN Service	Gcore	13,514	200	87	18	35	74	1	●	●
	Verizon	11,157	1,097	989	4	4	13	1	●	●
	Alibaba	14,615	549	546	11	12	23	5	●	●
	Fastly	16,917	5,127	4,838	9	9	113	1	●	●
	Tencent	14,385	70	61	23	33	48	10	●	●

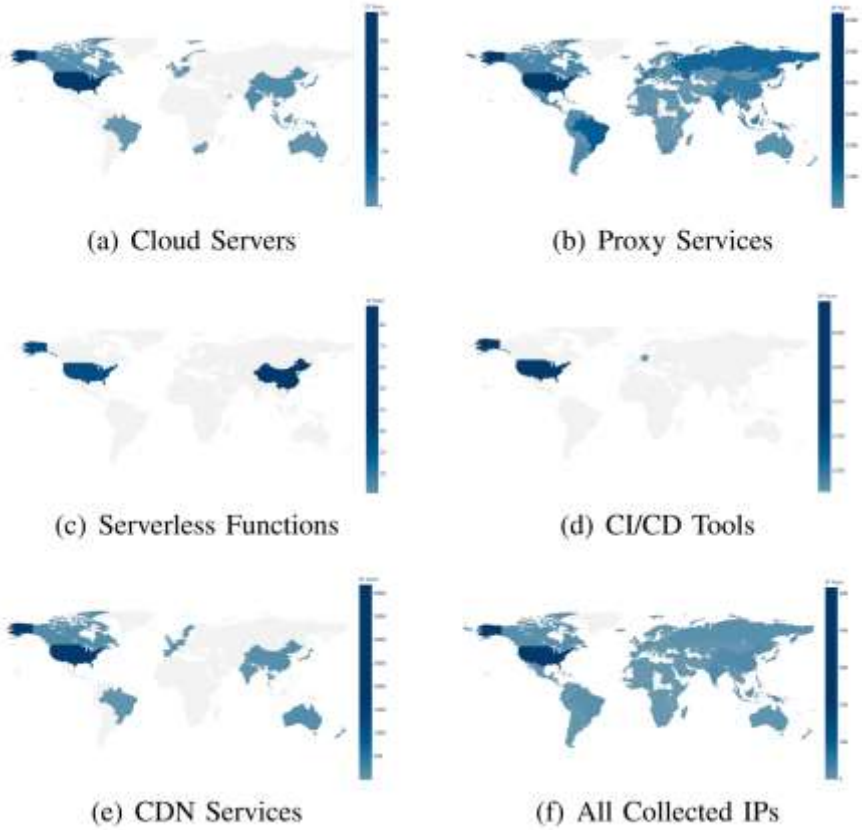

Global distribution of collected IP addresses

BreakSPF Attack Results

- Well-known domains like `microsoft.com`, `qq.com`, `godaddy.com`, and `ieee.org` were vulnerable to BreakSPF attacks

- A single IP address could be used to perform BreakSPF attack for up to 10,000 domains

Domain	Rank	IP	Source
microsoft.com	5	20.*.*.30	CI/CD Platforms
qq.com	11	114.*.*.86	Cloud Servers
csdn.net	76	114.*.*.86	Cloud Servers
huanqiu.com	110	114.*.*.86	Cloud Servers
godaddy.com	142	72.*.*.69	Tor
rednet.cn	306	114.*.*.86	Cloud Servers
mama.cn	311	114.*.*.86	Cloud Servers
zhihu.com	420	114.*.*.86	Cloud Servers
ieee.org	523	201.*.*.173	RESIP
ucla.edu	610	131.*.*.85	VPN

10 well-known domains vulnerable to BreakSPF attack

Rank	IP	# Domain[1]	Source	Provider	Representative Domain
1	162.*.*.128	11,408	Proxy Service	HTTP Proxy	websitewelcome.com
2	114.*.*.153	4,604	Cloud Server	Tencent	qq.com
3	213.*.*.46	4,580	Proxy Service	HTTP Proxy	batmanapollo.ru
4	116.*.*.140	1,189	Proxy Service	RESIP	mailcontrol.com
5	161.*.*.149	411	Cloud Server	Alibaba	shopee.ph
8	80.*.*.207	240	Proxy Service	Tor	mailbox.org
9	154.*.*.131	131	Proxy Service	RESIP	netblocks.aserv.co.za
10	185.*.*.2	110	Proxy Service	Tor	octopuce.fr
11	133.*.*.61	97	Proxy Service	HTTP Proxy	myasp.jp
13	81.*.*.68	74	Proxy Service	HTTP Proxy	jino.ru

Top 10 IP addresses that can attack multiple domains

BreakSPF Attack Results: Example

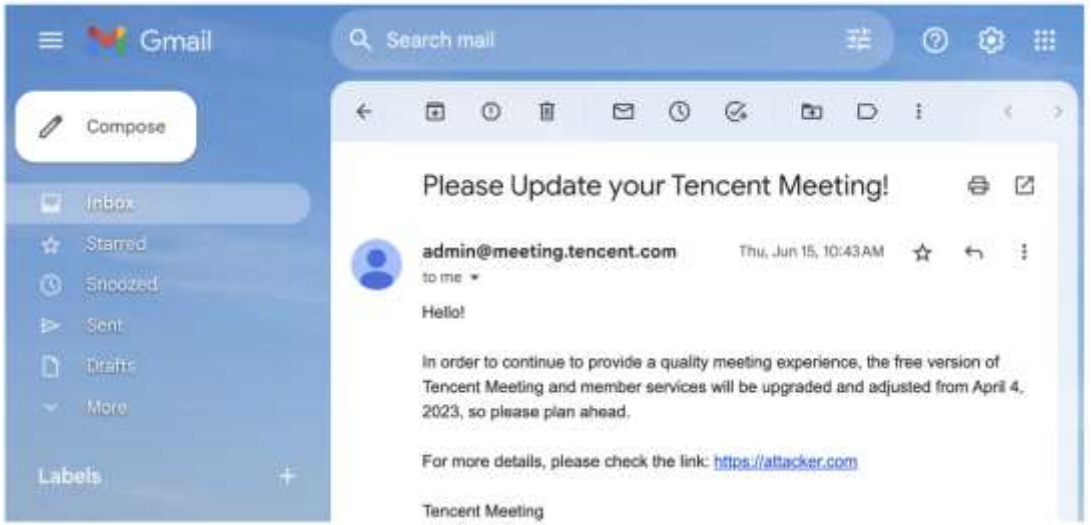

A spoofed email sent with BreakSPF attack
(From: admin@meeting.tencent.com)

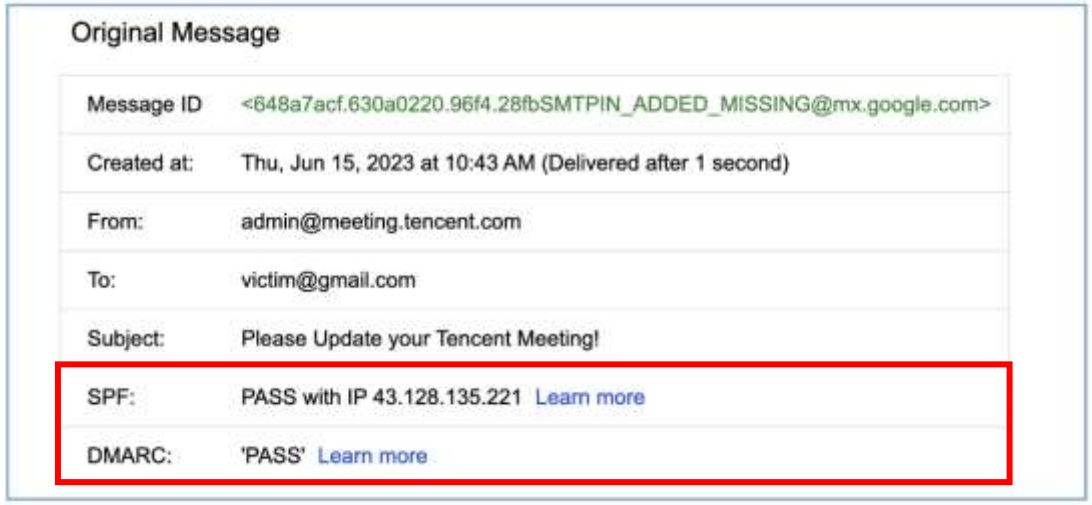

Validation result of the spoofed email

Mitigations

- **Port management.** Cloud services, proxy services, etc. should restrict egress communication to port 25, 465, etc.
- **Online detection services.** https://breakspf.cloud/
- **DMARC* reports.** Recipients who receive emails can aggregate & send validation results to the domain owner.

** **DMARC** (Domain-based Message Authentication, Reporting and Conformance)*
: A protocol that aligns the domain name in “From” header and the authenticated “MAIL FROM” address from *SPF* or *DKIM (DomainKeys Identified Mail)*