SPF Beyond the Standard: Management and Operational Challenges in Practice and Practical Recommendations

Md. Ishtiaq Ashiq[§], Weitong Li[§],
Tobias Fiebig[†], and Tijay Chung[§]

[§]Virginia Tech, [†]Max Planck Institute for Informatics

VIRGINIA TECH

Motivation

From: admin@paypal.com

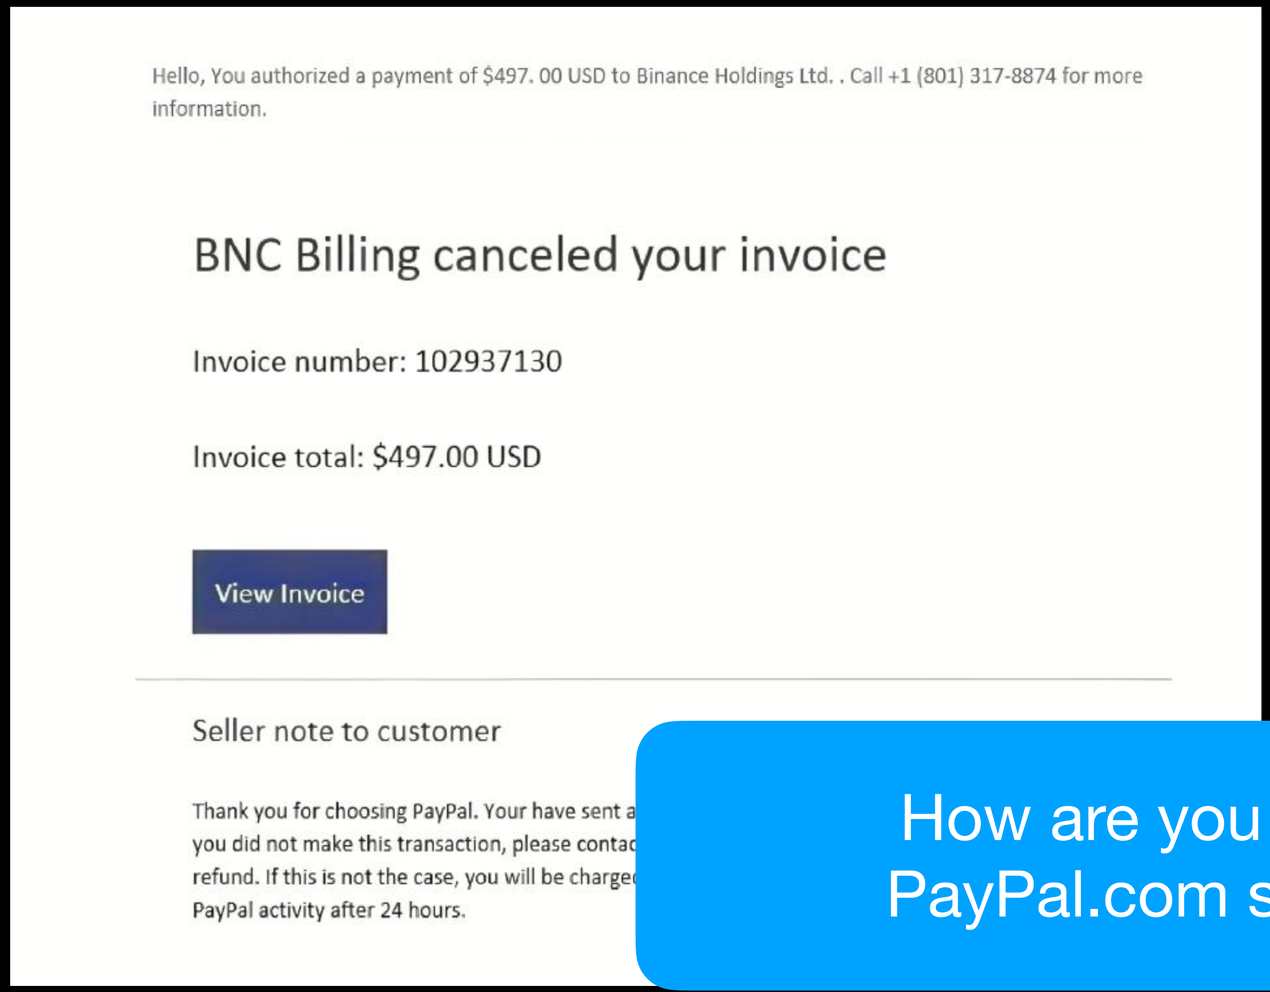

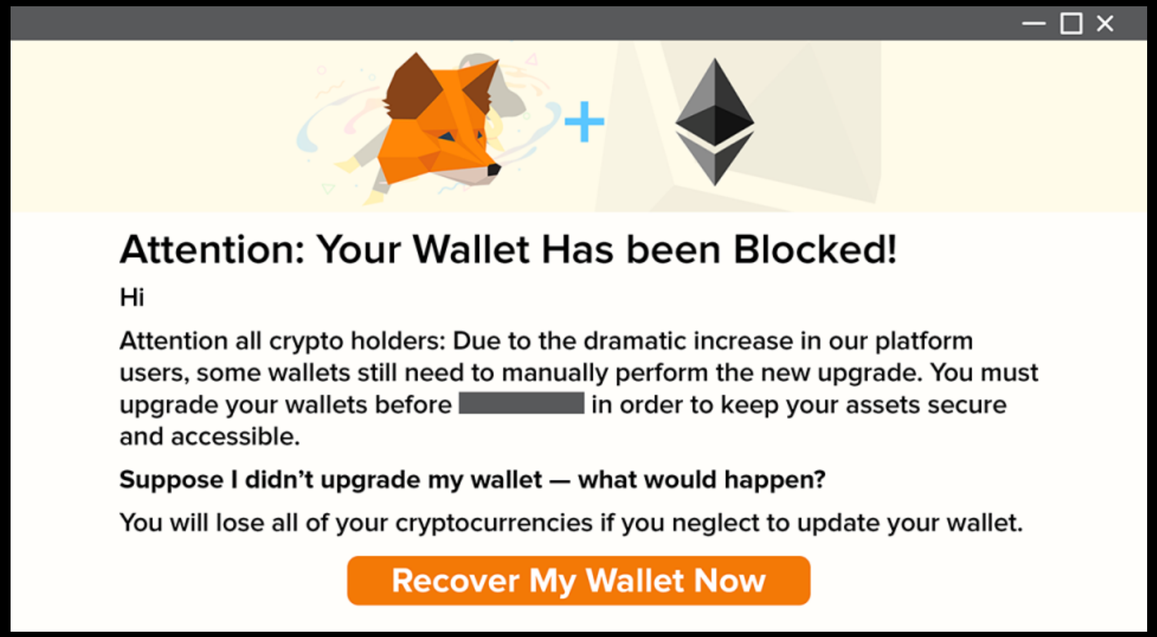

How are you sure that actually PayPal.com sent these emails?

No built-in security!

- In original SMTP protocol,
  - Anyone can send an email impersonating any address!

How SMTP Works

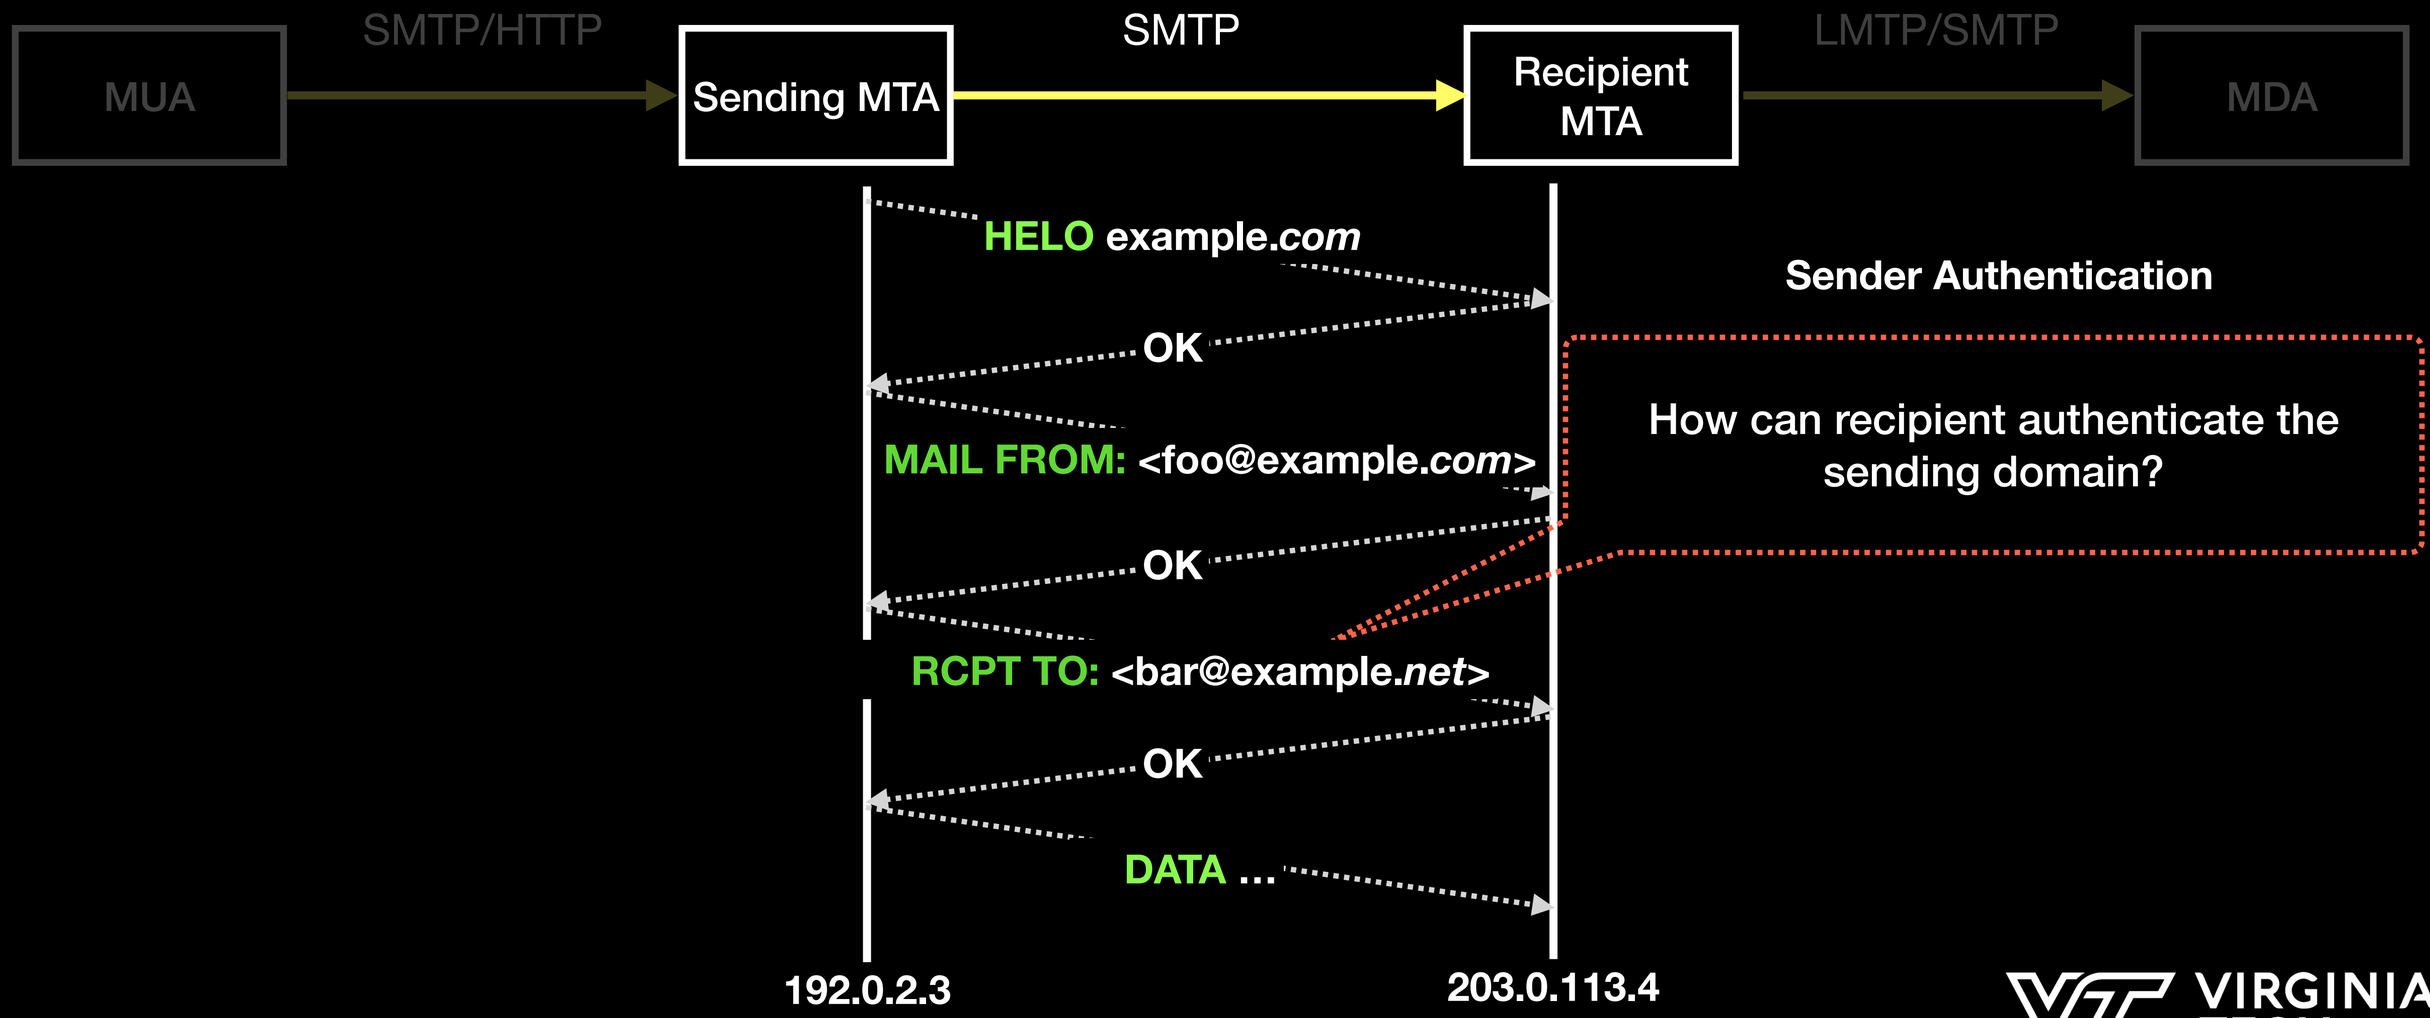

VIRGINIA TECH

SPF (Sender Policy Framework)

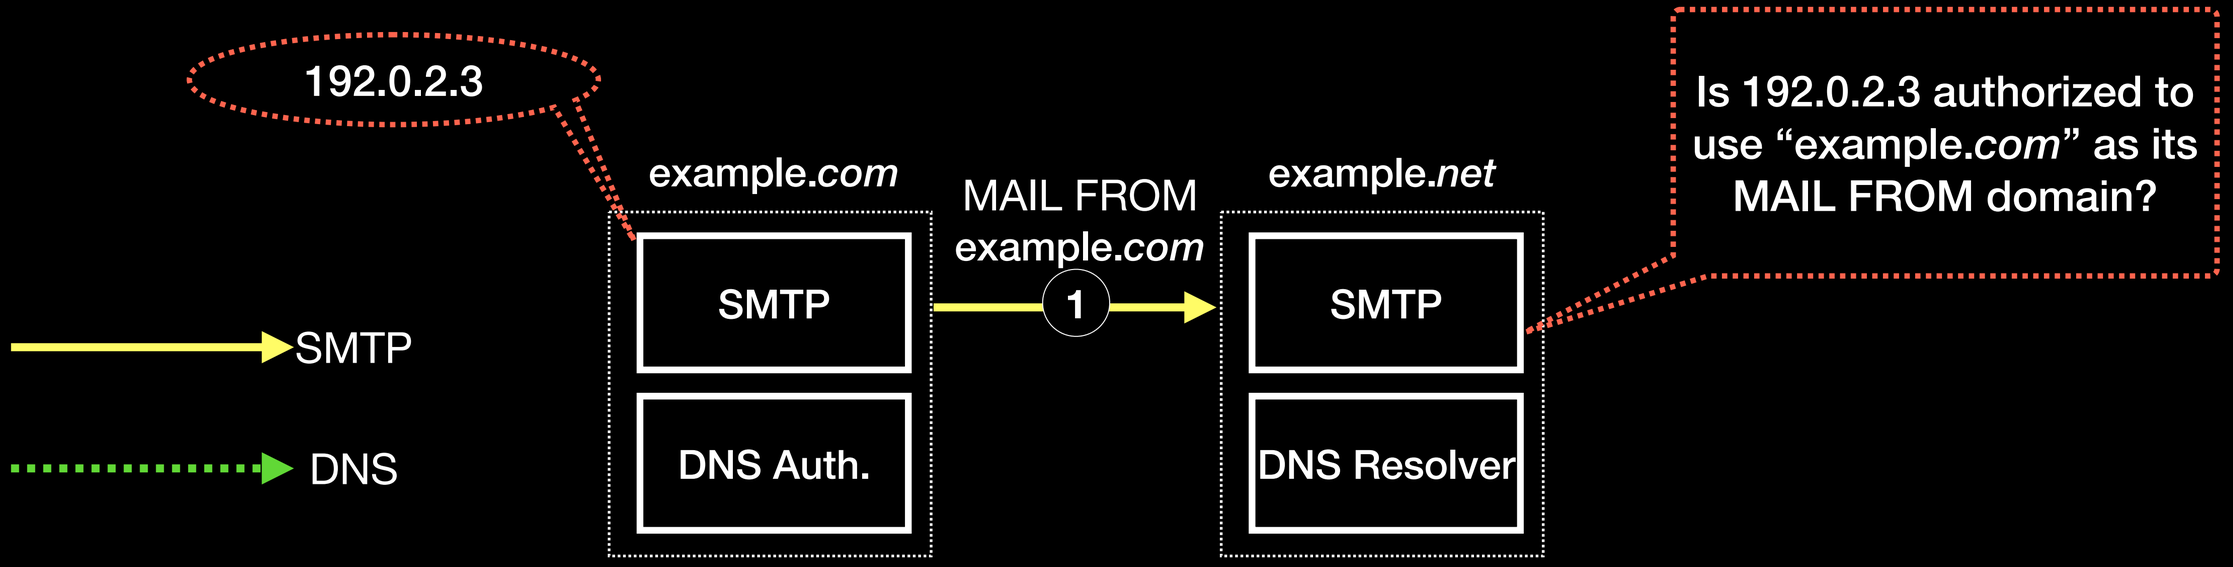

SPF (Sender Policy Framework)

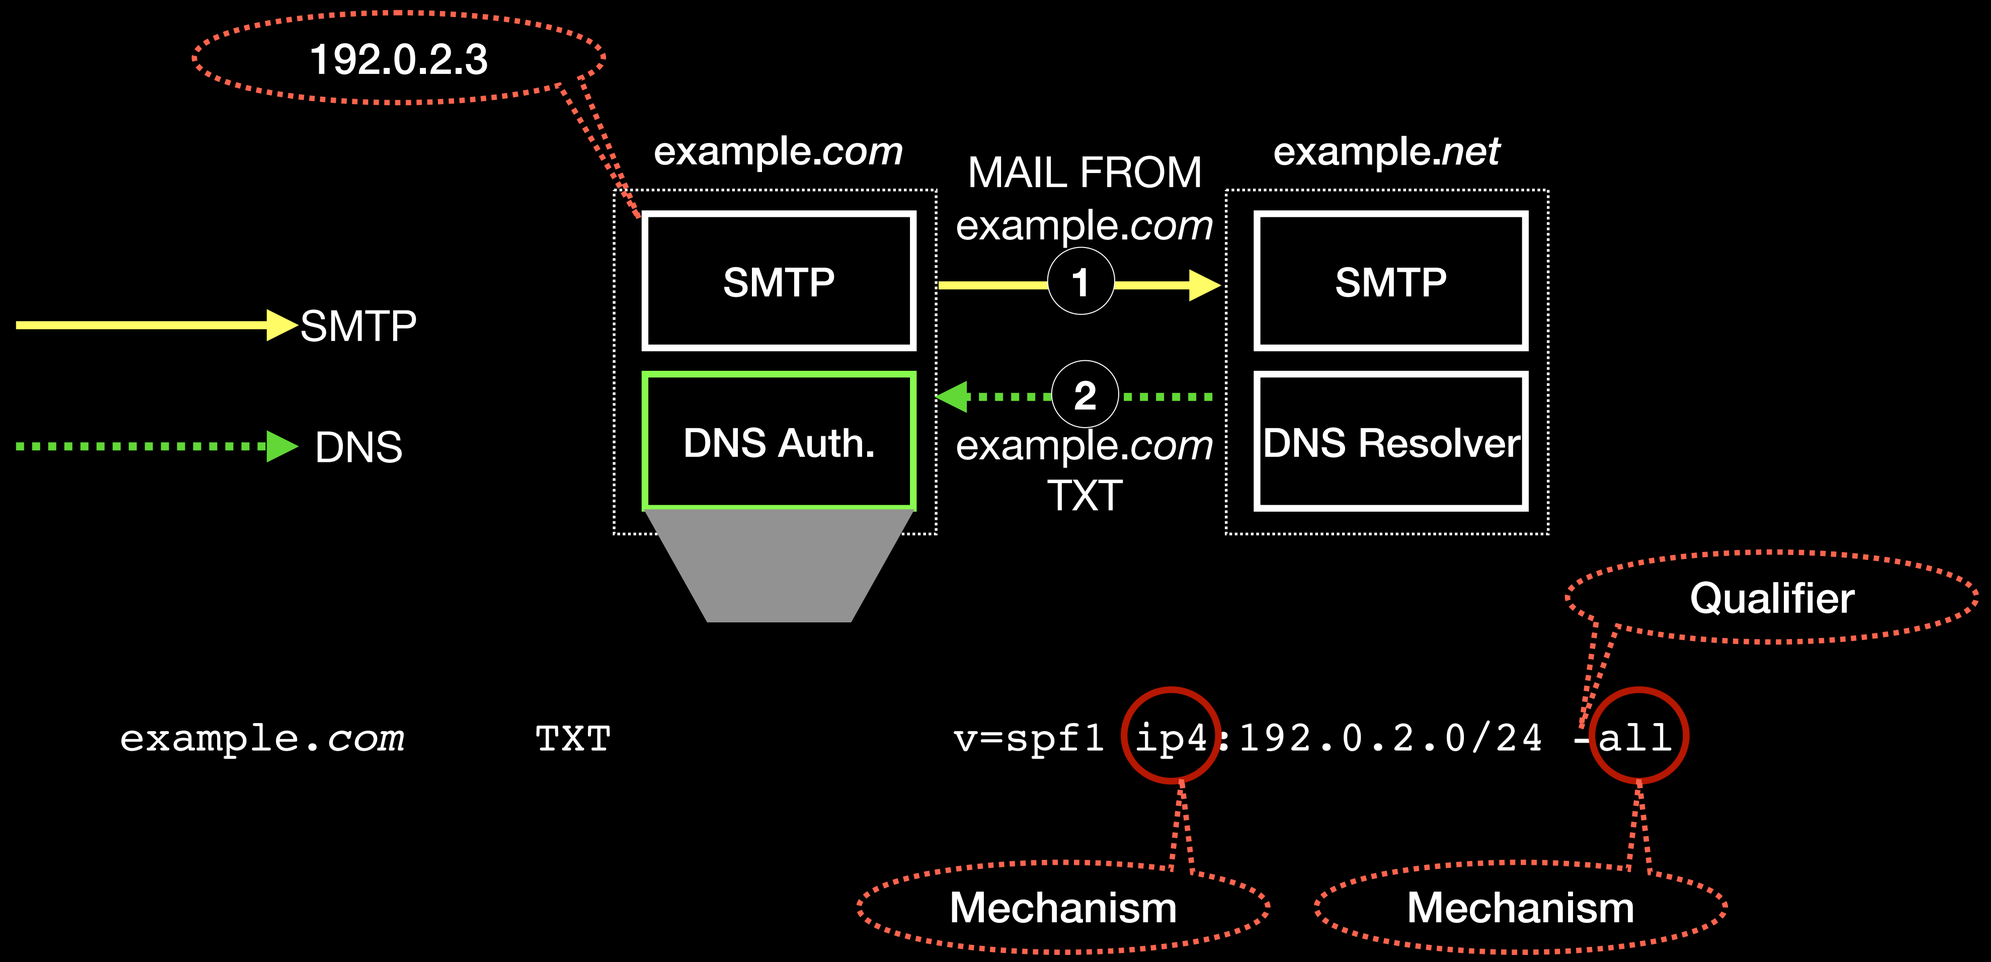

VIRGINIA TECH

Let’s see one example w/ include

Imagine client IP: 35.190.247.227 & sender: user@example.com

example.com	TXT	v=spf1 include:_spf.google.com -all

Let’s see one example w/ include

Imagine client IP: 35.190.247.227 & sender: user@example.com

example.com	TXT	v=spf1 include:_spf.google.com -all
_spf.google.com	TXT	v=spf1 include:_netblocks.google.com include:_netblocks2.google.com include:_netblocks3.google.com ~all

Let’s see one example w/ include

Imagine client IP: 35.190.247.227 & sender: user@example.com

example.com	TXT	v=spf1 include:_spf.google.com -all
_spf.google.com	TXT	v=spf1 include:_netblocks.google.com include:_netblocks2.google.com include:_netblocks3.google.com ~all
_netblocks.google.com	TXT	v=spf1 ip4:35.190.247.0/24 …

Result: Pass

What does this imply?

- Imagine an attacker create an infinite chain of SPF includes in his domain and send email from this domain
- There must be a limit on the number of these resolutions, right?
    - SPF standard dictates that an SPF verifier must not do more than 10 DNS lookups; otherwise, return an error.

Research Question 1

- How many domains require more than 10 DNS lookups to resolve their SPF record?

Dataset Overview

TLD	MX	SPF
.com	75.8M	48M (63.2%)
.net	6.5M	3.5M (53.8%)
.org	5.8M	3.2M (55.2%)
.se	845K	439K (52%)

Data gathered from Nov 2021 to Mar 2023

Answering RQ1

- How many domains require more than 10 DNS lookups to resolve their SPF record?
  - 3,548,014 (6.5%) domains in our latest snapshot require more than 10 DNS lookups
    - Over 99% of them have *include* mechanism

Research Question 2

- Why do so many domains require more than 10 lookups?
  - Is it a misconfiguration or a necessity in today’s world of shared infrastructure?

SPF Deployment

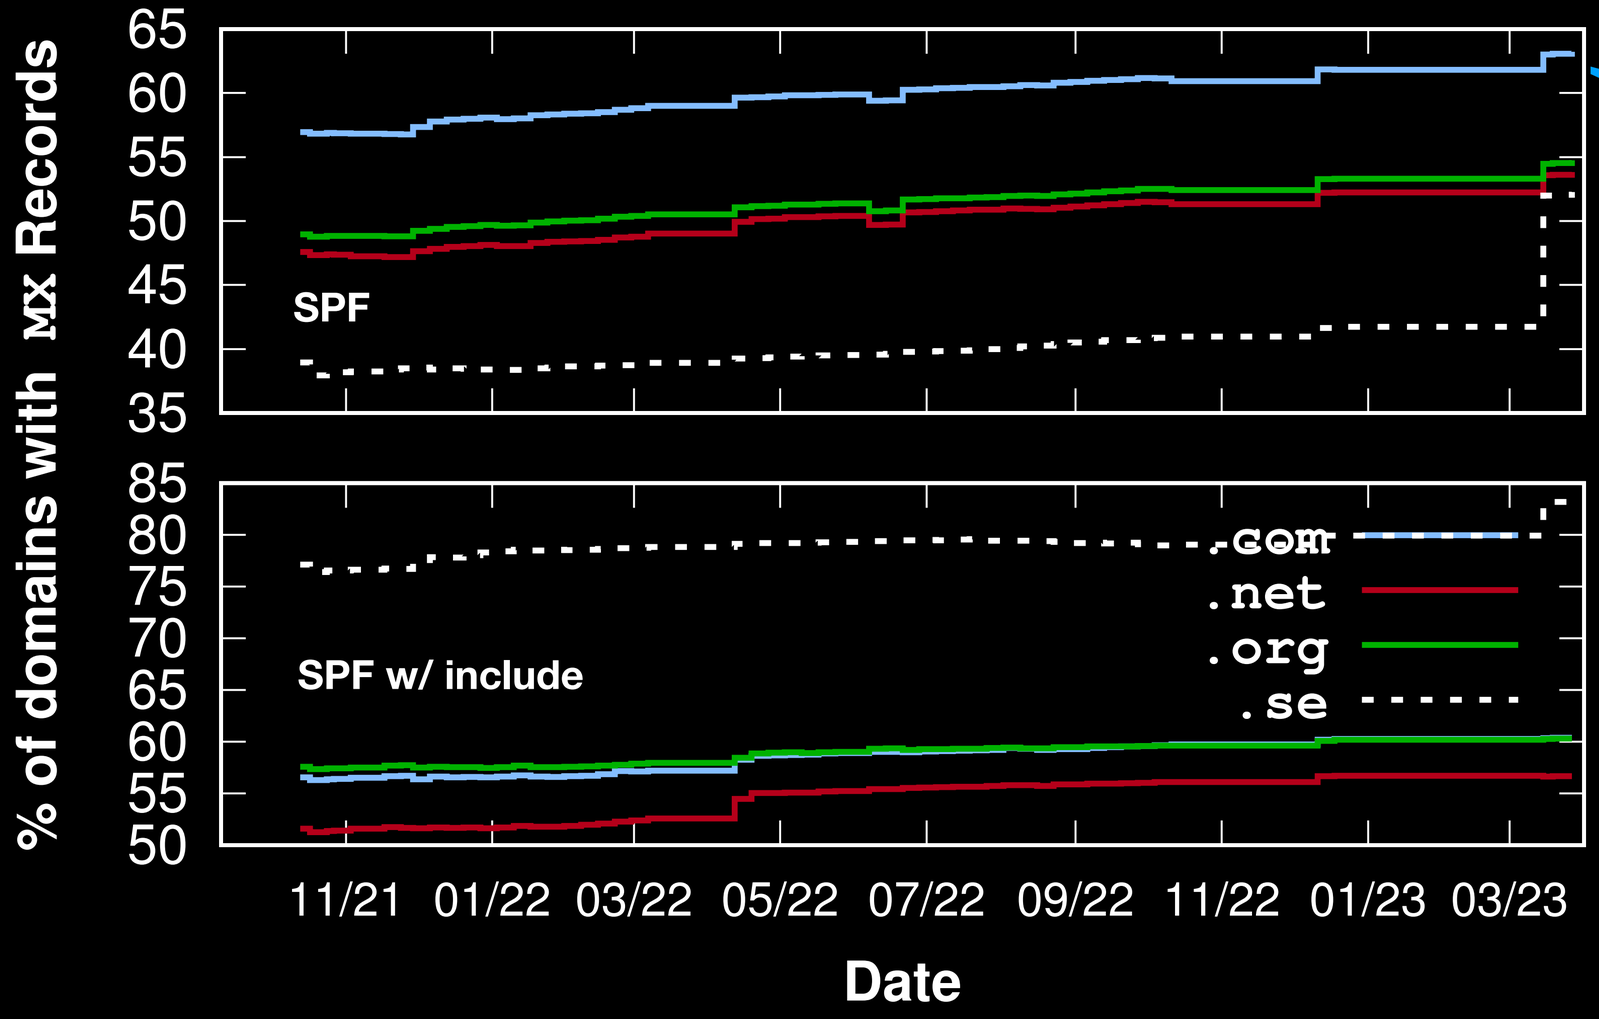

SPF is widely adopted; 63.2% *.com* domains have SPF records as of Mar 2023

~60% of the domains have *include* mechanism in them; 99% of them are external. Email ecosystem is heavily reliant on 3rd-party providers.

VIRGINIA TECH

Let’s find out why?

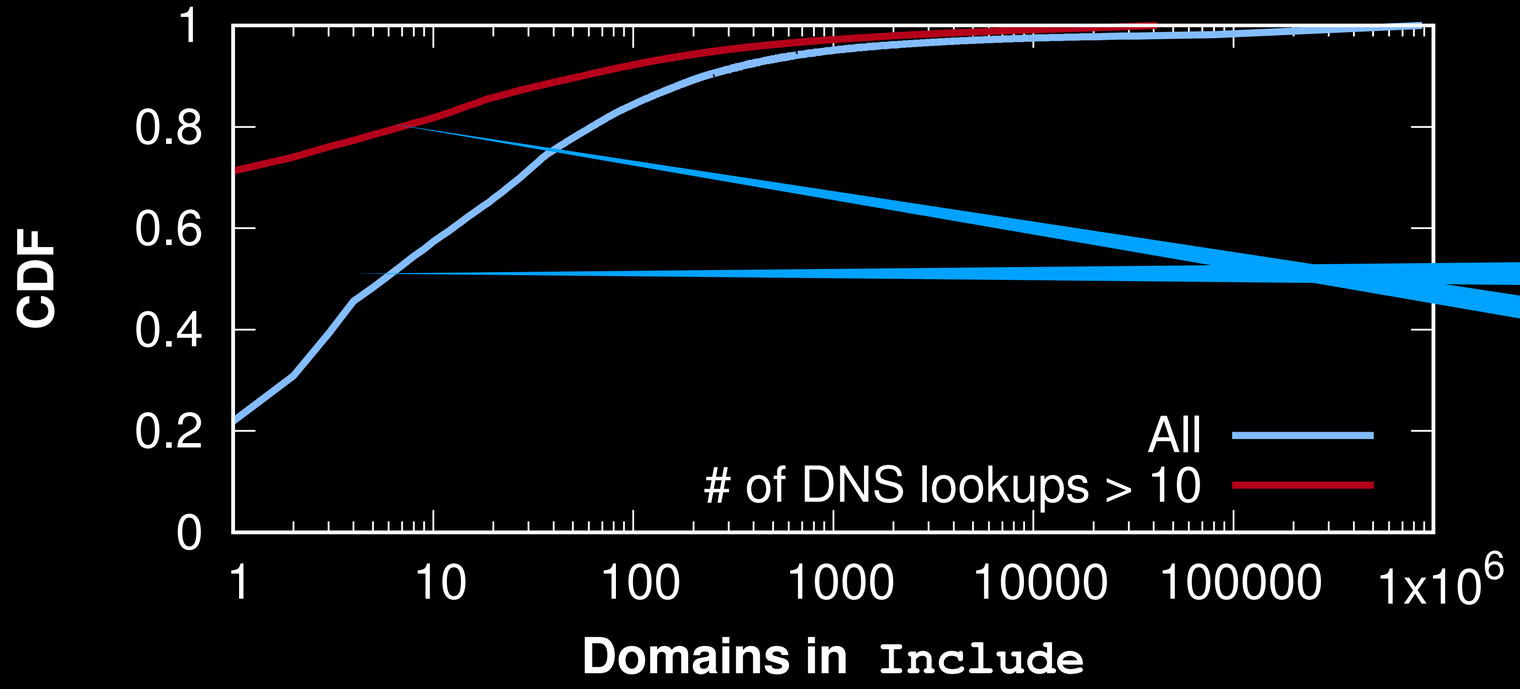

Only 6 domains appear in the include mechanism for 50% of SPF records

Just 2 domains account for 83.7% of the SPF records with more than 10 lookups

VIRGINIA TECH

Let’s find out why?

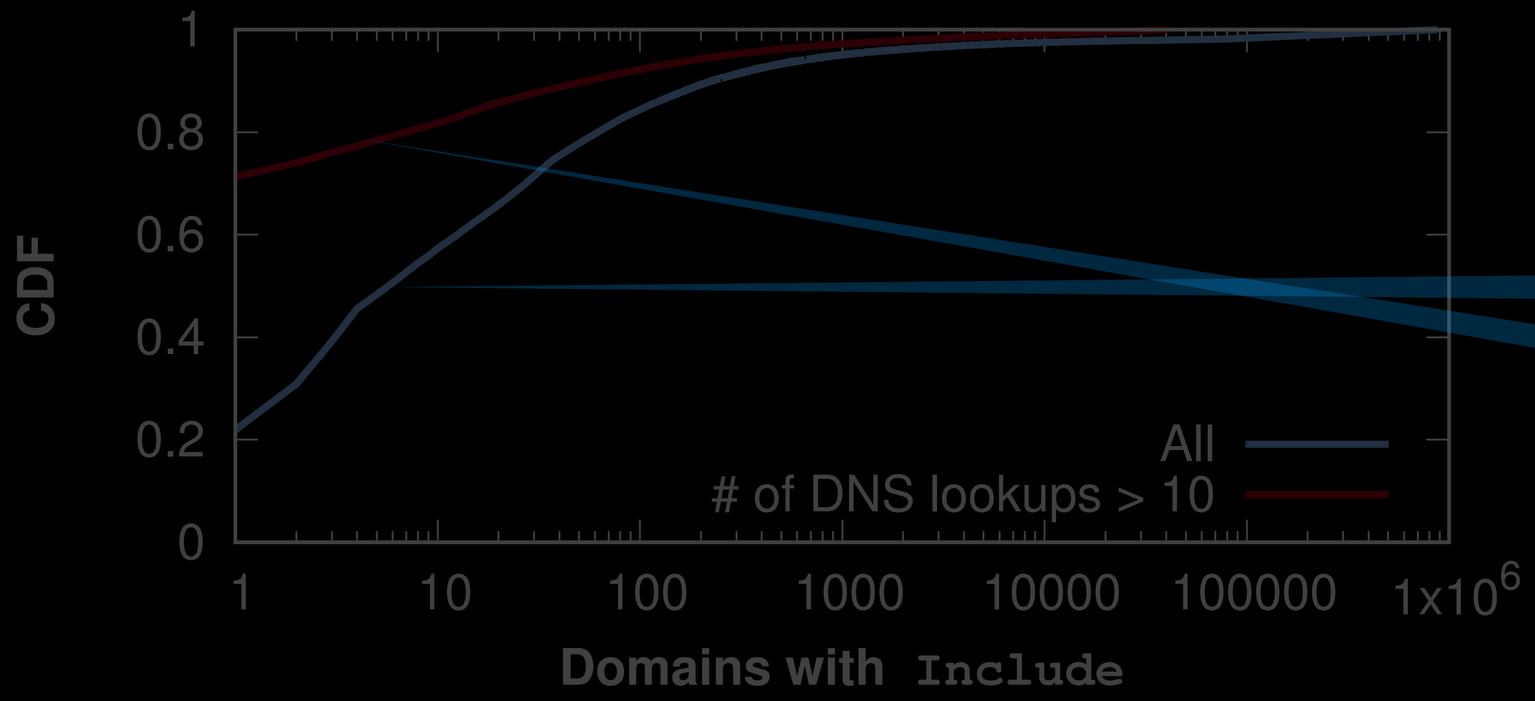

Only 6 domains contribute appear in the include mechanism for 50% of SPF records

Just 2 domains account for 83.7% of the SPF records with more than 10 lookups

What about the rest 16.3% (616,581) domains?

Are all the includes actually being used?

How to detect superfluous includes?

- Leveraged MX records (~350K domains just use *mx* as the only mechanism)

- Computed the likelihood that a domain with a specific MX record also includes an SPF record

$$¶(spf_k \mid mx_m) = \frac{d(mx_m, spf_k)}{\sum_{i=1}^{n} d(mx_m, spf_i)} \; ; d(mx_m, spf_k) = \text{\# of domains containing this tuple combination}$$

Results

- Dataset
  - # of explainable domains: 24,832 domains
    - 20,124 (81%) are burdened with superfluous record

How the lookup # reduces?

- We found 20,124 domains with superfluous record

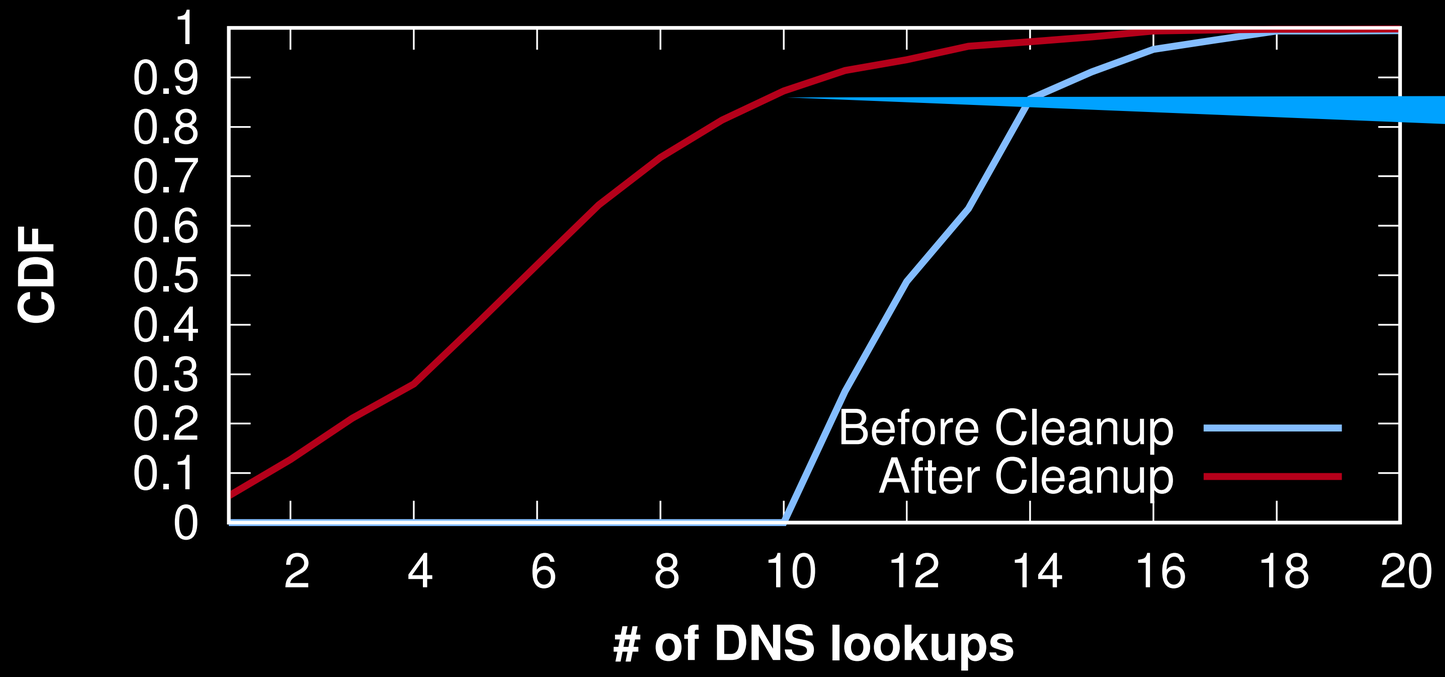

17,554 (87.2%) of these domains can effectively streamline their # of lookups within 10 once extraneous records are removed

VIRGINIA TECH

Why do these superfluous records exist?

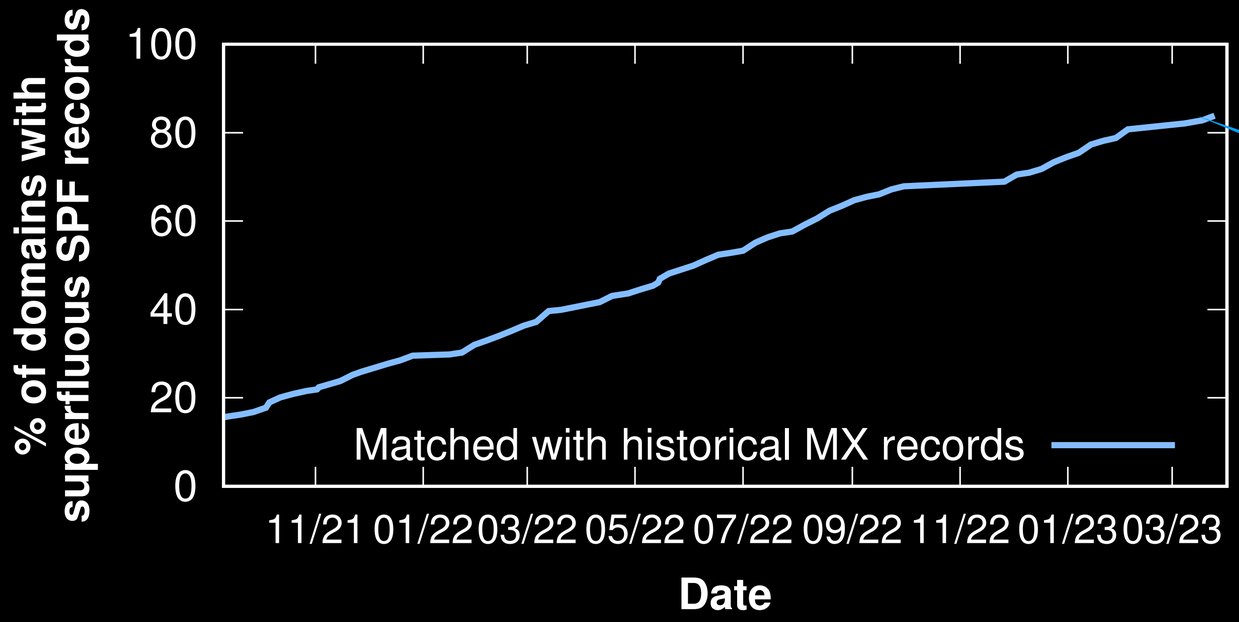

83.9% stem from obsolete SPF entries not removed following mail server migration

VIRGINIA TECH

Research Question 3

- Do SMTP servers in the wild maintain this limit?
  - If not, bad actors can use them as a reflector to launch DoS attacks

What’s the state in the wild?

- Can we conduct a internet-wide scan of SMTP servers to understand whether they are violating this lookup limit?
  - We need to send an email, not ethical
  - Remember that SMTP works based on many commands, but the RFC doesn't define when to check SPF records!
    - So, we can connect, send up to the RCPT command, and quit.

Results

- Initiated a connection attempt to all unique SMTP servers in our dataset (1.89M)
- Connected to 1.2M servers (64%)
- 81K made SPF queries (6.8%)
  - Most opt for validation after the DATA command
- 195 queried all included domains in our SPF structure!

Research Question 4

- Do the existing and popular open-source SPF validators properly comply with the standard?

Results

Name	Version	# of Allowed Lookups
libspf2	latest	10
Mail::SPF:Query	< 1.8	∞
Mail::SPF::Query	> 1.8	10
pyspf	latest	10
milter-greylist	latest	10
spfmilter	latest	10
Mtpolicyd	latest	10
policyd-spf	latest	20
iRedAPD	< 5.1	∞
iRedAPD	>= 5.1	20
SpamAssassin	latest	20
RSpamD	latest	30

Invoking SPF Resolution Timeout on Valid Emails

- Some SPF extensions do not handle incoming emails concurrently
  - With a domain and customized DNS authoritative server, adversaries can exploit these extensions
- Impact
  - Interruptions in valid incoming emails at the victim MTA

Invoking SPF Resolution Timeout on Valid Emails

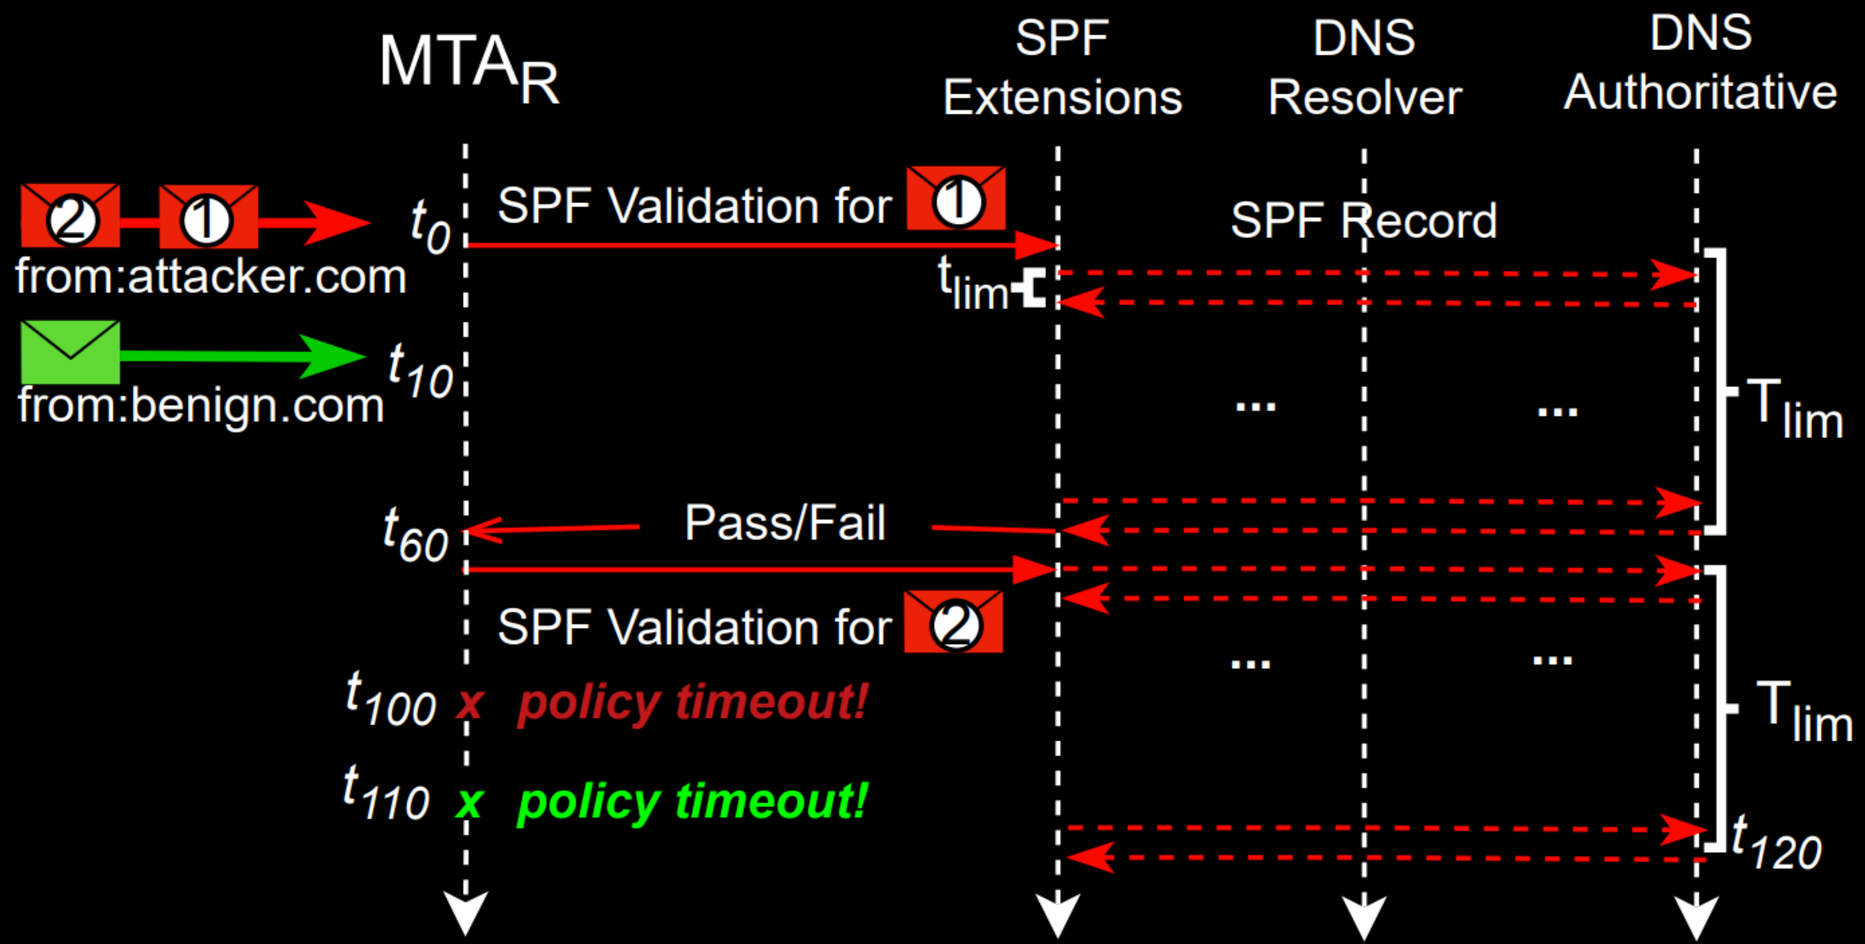

VIRGINIA TECH

Using CNAME expansion to bypass limits

example.com	TXT	“v=spf1 include:_spf.example.com -all”
_spf.example.com	CNAME	_spf2.example.com
_spf2.example.com	TXT	“v=spf1 include:_spf3.example.com"
_spf3.example.com	CNAME	_spf4.example.com

...

Evaluation

- Found one popular SPF extension to be vulnerable to this attack
  - Given the default policy service timeout of *Postfix* (100s), just sending 2 emails can create a ~20s time window, where all valid emails will likely be rejected

Conclusion

- First measurement study to deep-dive into the reasoning behind excessive DNS lookups
- Identify vulnerable open-source SPF verifiers and SMTP servers in the wild
- Show how non-parallel SPF verifiers can be misused and exploited
- Qualitative study
- Recommendations for future iterations on RFC7208


iashiq5@vt.edu

ashiq5.github.io