F**k it, let's do it wide!

Security Applications of Fast Internet-Wide Scanning

David Adrian
@davidcadrian

UNIVERSITY OF MICHIGAN


NMAP

+


amazon
web services™

Public IPv4 Crytographic Keys (2012)

	HTTPS	SSH
Live Hosts	12.8 million	10.2 million
Distinct RSA keys	5.6 million	3.8 million
Distinct DSA keys	6,241	2.8 million

Factoring RSA Public Keys

Factoring RSA Public Keys

Factoring is computationally hard

Factoring RSA Public Keys

Factoring is computationally hard

$$n = pq$$

Factoring RSA Public Keys

Factoring is computationally hard

GCD is computationally easy.

$$n = pq$$

Factoring RSA Public Keys

Factoring is computationally hard

$$n = pq$$

GCD is computationally easy.

$$\begin{cases} N_1 = pq_1 \\ N_2 = pq_2 \end{cases}$$

$$(N_1, N_2) = p$$

Broken Keys

1. Gather all RSA moduli
2. Calculate all pairwise GCD
3. See what factors come out.

Factors 0.5% of keys on the Internet!

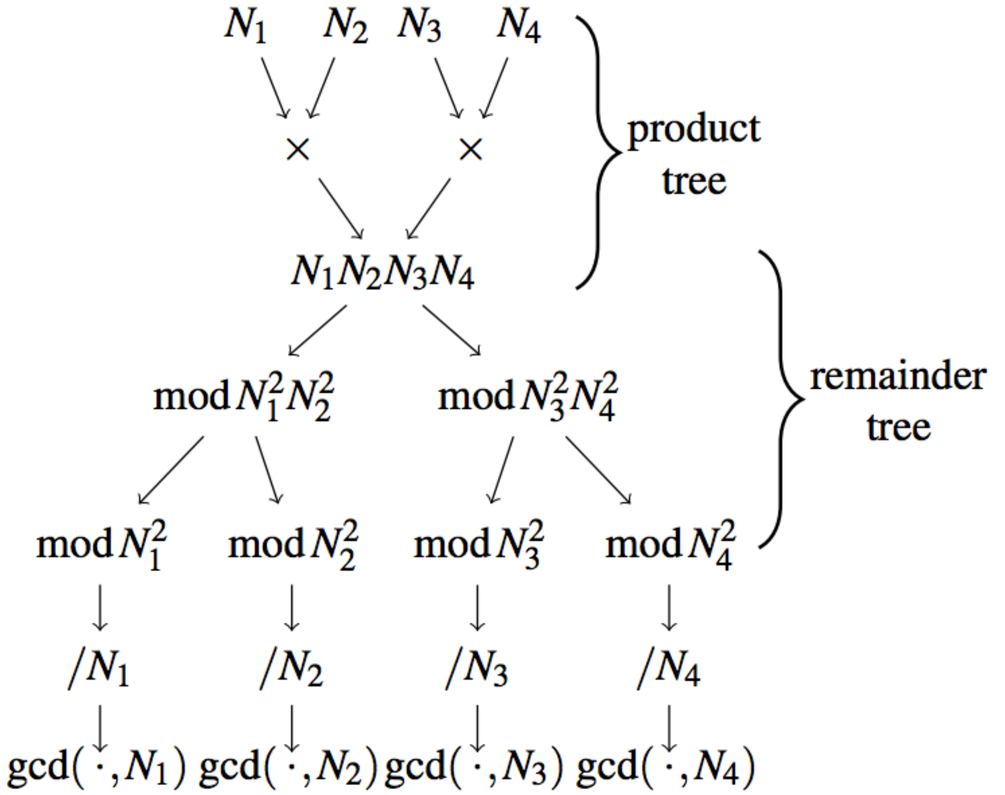


NMAP

+


amazon
web services™

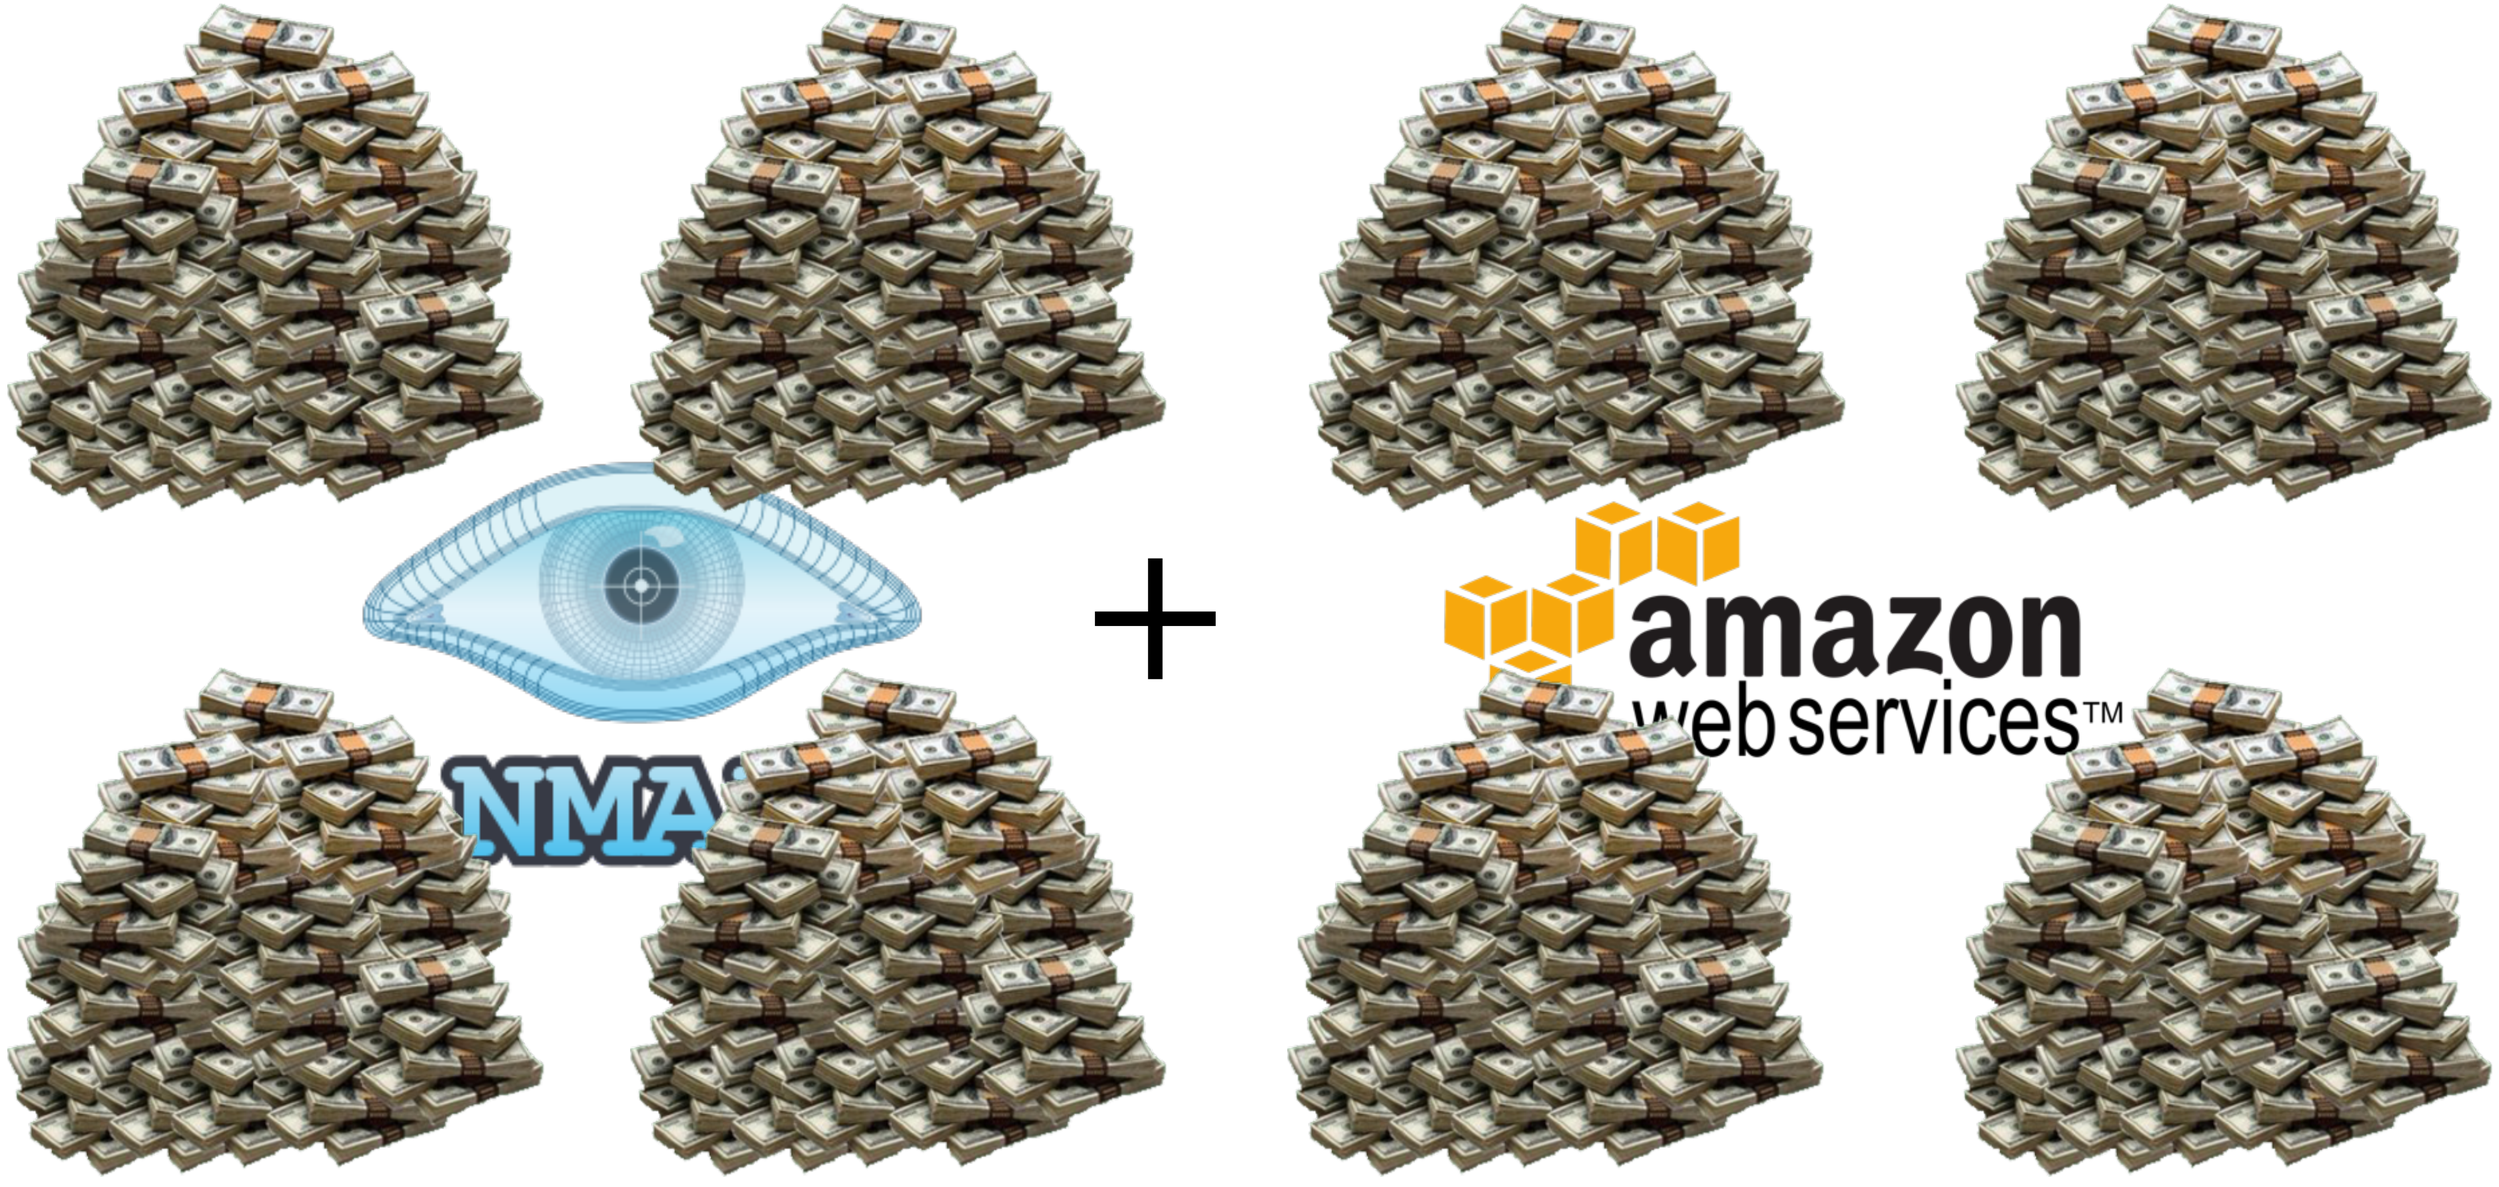
NMA
amazon
web services™

What if Internet survey's didn't require heroic effort?

What if we could scan every protocol every day?

What if we wrote a whole-Internet scanner from scratch?

ZMap: The Fast Internet-Wide Scanner

ZMap is an Internet-wide port scanner capable of scanning at **95% the maximum theoretical speed** of 10 gigabit Ethernet

ZMap completes a single-port TCP SYN scan of **all of IPv4 in under five minutes**

ZMap

2013 A **1200x performance improvement** over Nmap for an Internet-wide single port TCP scan

2014 Scan the Internet in **under 5 minutes**.

2015 Popular in industry and academia, used by over **104** academic studies

ZMap Architecture

Existing Network Scanners

Reduce state by scanning in batches
- Time lost due to blocking
- Results lost due to timeouts

Track individual hosts and retransmit
- Most hosts will not respond

Avoid flooding through timing
- Time lost waiting

Utilize existing OS network stack
- Not optimized for many connections


ELIMINATE
THE STATE

ZMap Architecture

Existing Network Scanners

Reduce state by scanning in batches
- Time lost due to blocking
- Results lost due to timeouts

Track individual hosts and retransmit
- Most hosts will not respond

Avoid flooding through timing
- Time lost waiting

Utilize existing OS network stack
- Not optimized for many connections

ZMap

Eliminate local per-connection state
- Fully asynchronous components
- No blocking except for network

Shotgun Scanning Approach
- Always send *n* probes per host

Scan widely dispersed targets
- Send as fast as network allows

Probe-optimized Network Stack
- Generate Ethernet frames

Two Cool Tricks To Speed Up Your Address Generation by 1100%!!!!!!

Addressing Probes Trick One

1. Scan hosts according to a random permutation

2. Iterate over multiplicity group of integers modulo *p*

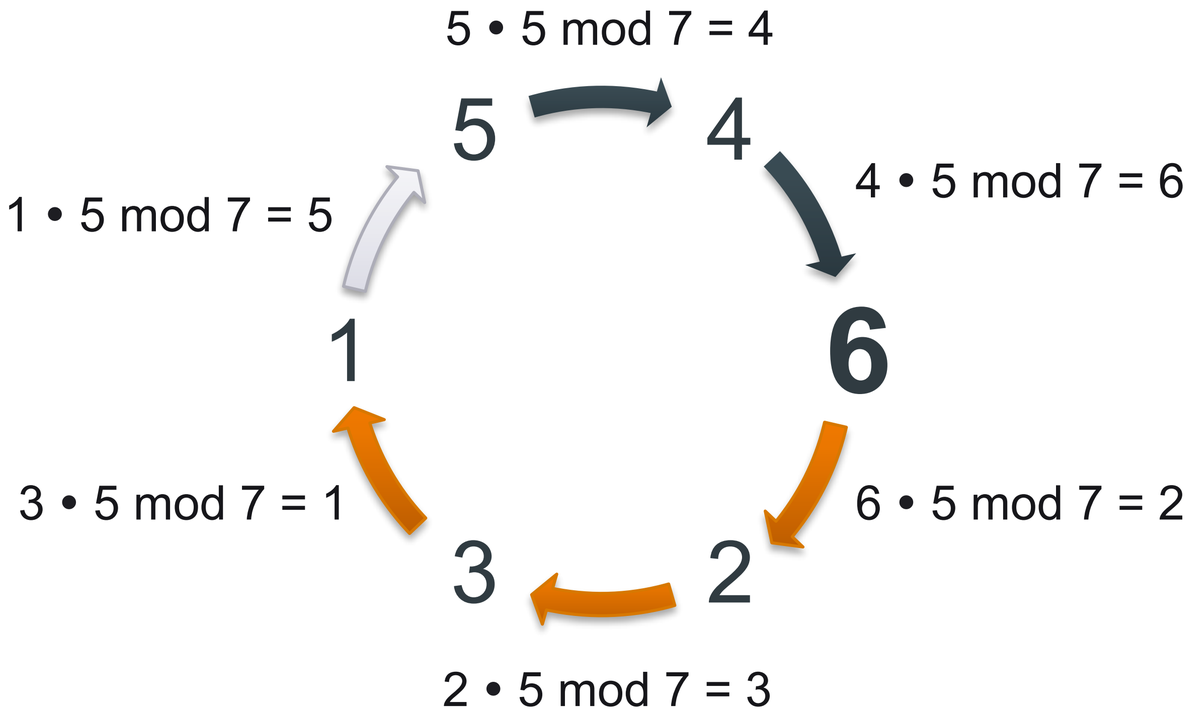

Negligible State

1. Primitive Root
2. Current Location
3. First Address

Addressing Probes Trick Two

Multithreaded iteration over a cyclic group of integers requires a lock

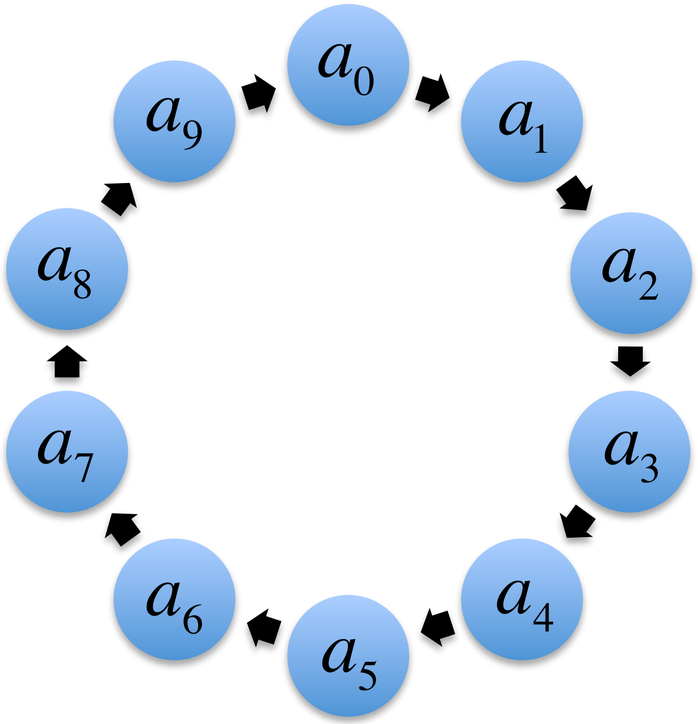

$$a_{i+1} = g \cdot a_i \bmod p$$

Addressing Probes Trick Two

Multithreaded iteration over a cyclic group of integers requires a lock

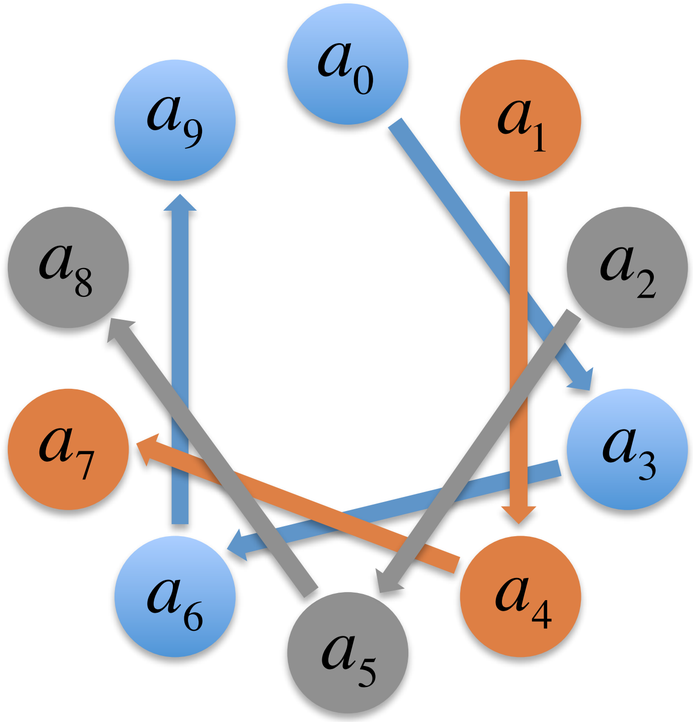

$$a_{i+1} = g \cdot a_i \bmod p$$

$$a_{i+n} = g^n \cdot a_i \bmod p$$

Shard the cycle into disjoint sets

Validating Responses

Encode secrets into mutable fields of probe packets with deterministic effect on responses

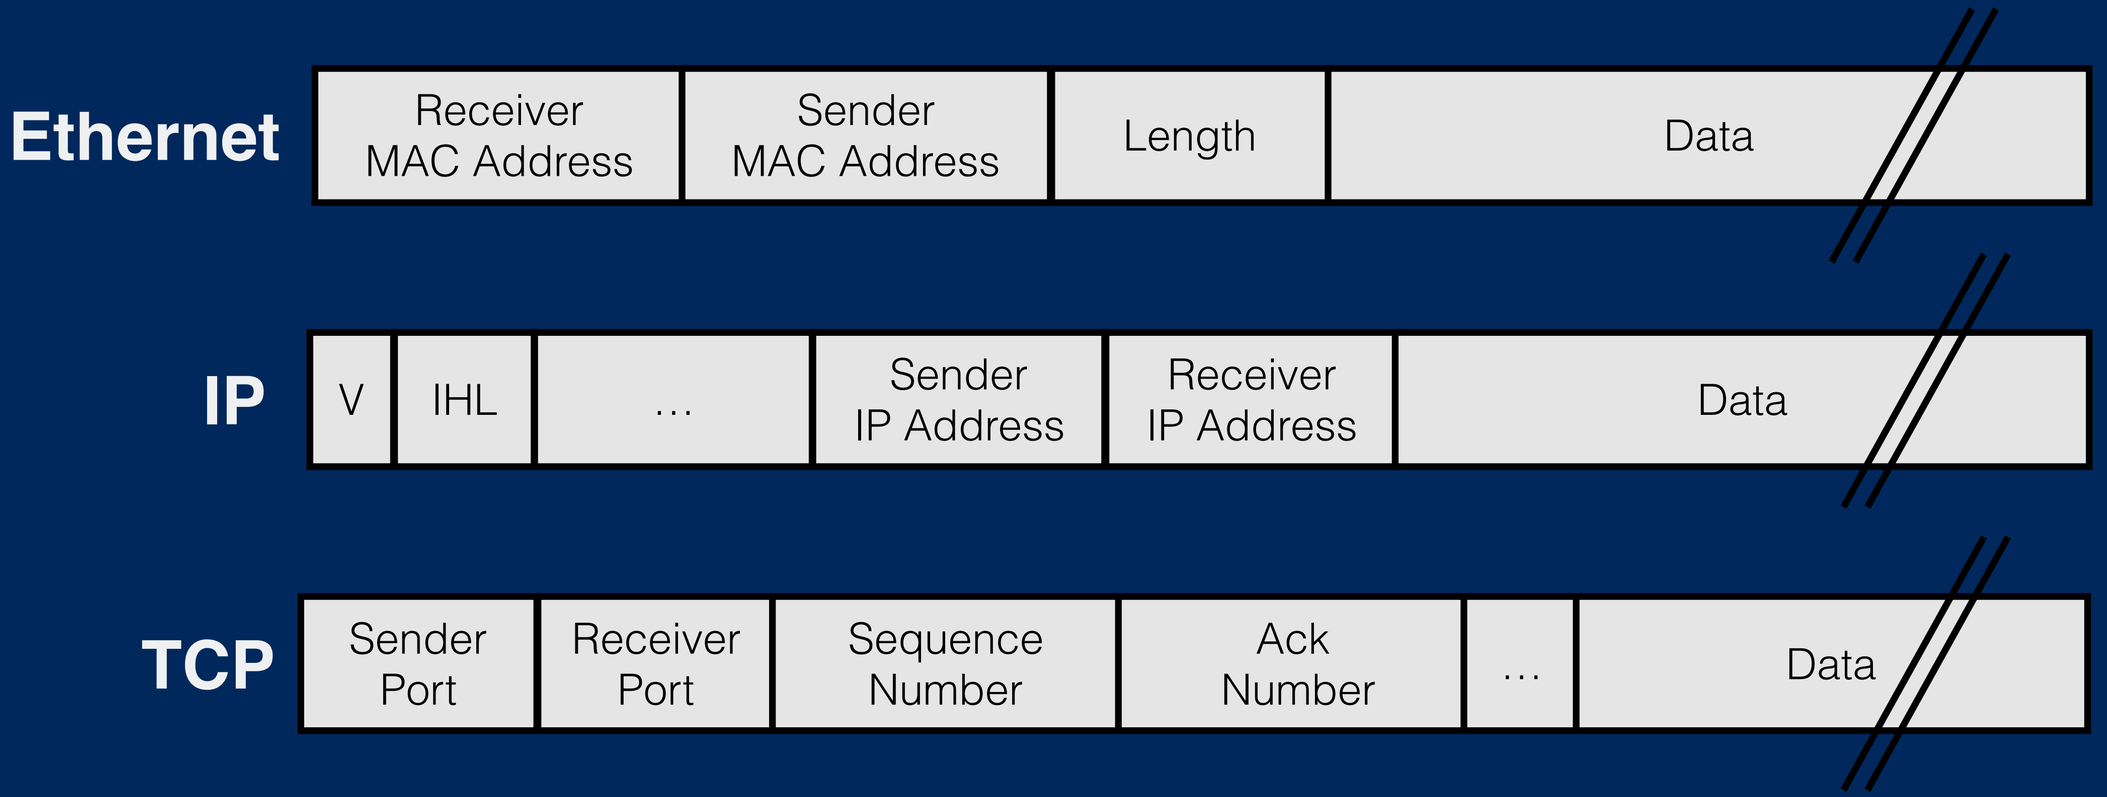

Validating Responses

Encode secrets into mutable fields of probe packets with deterministic effect on responses

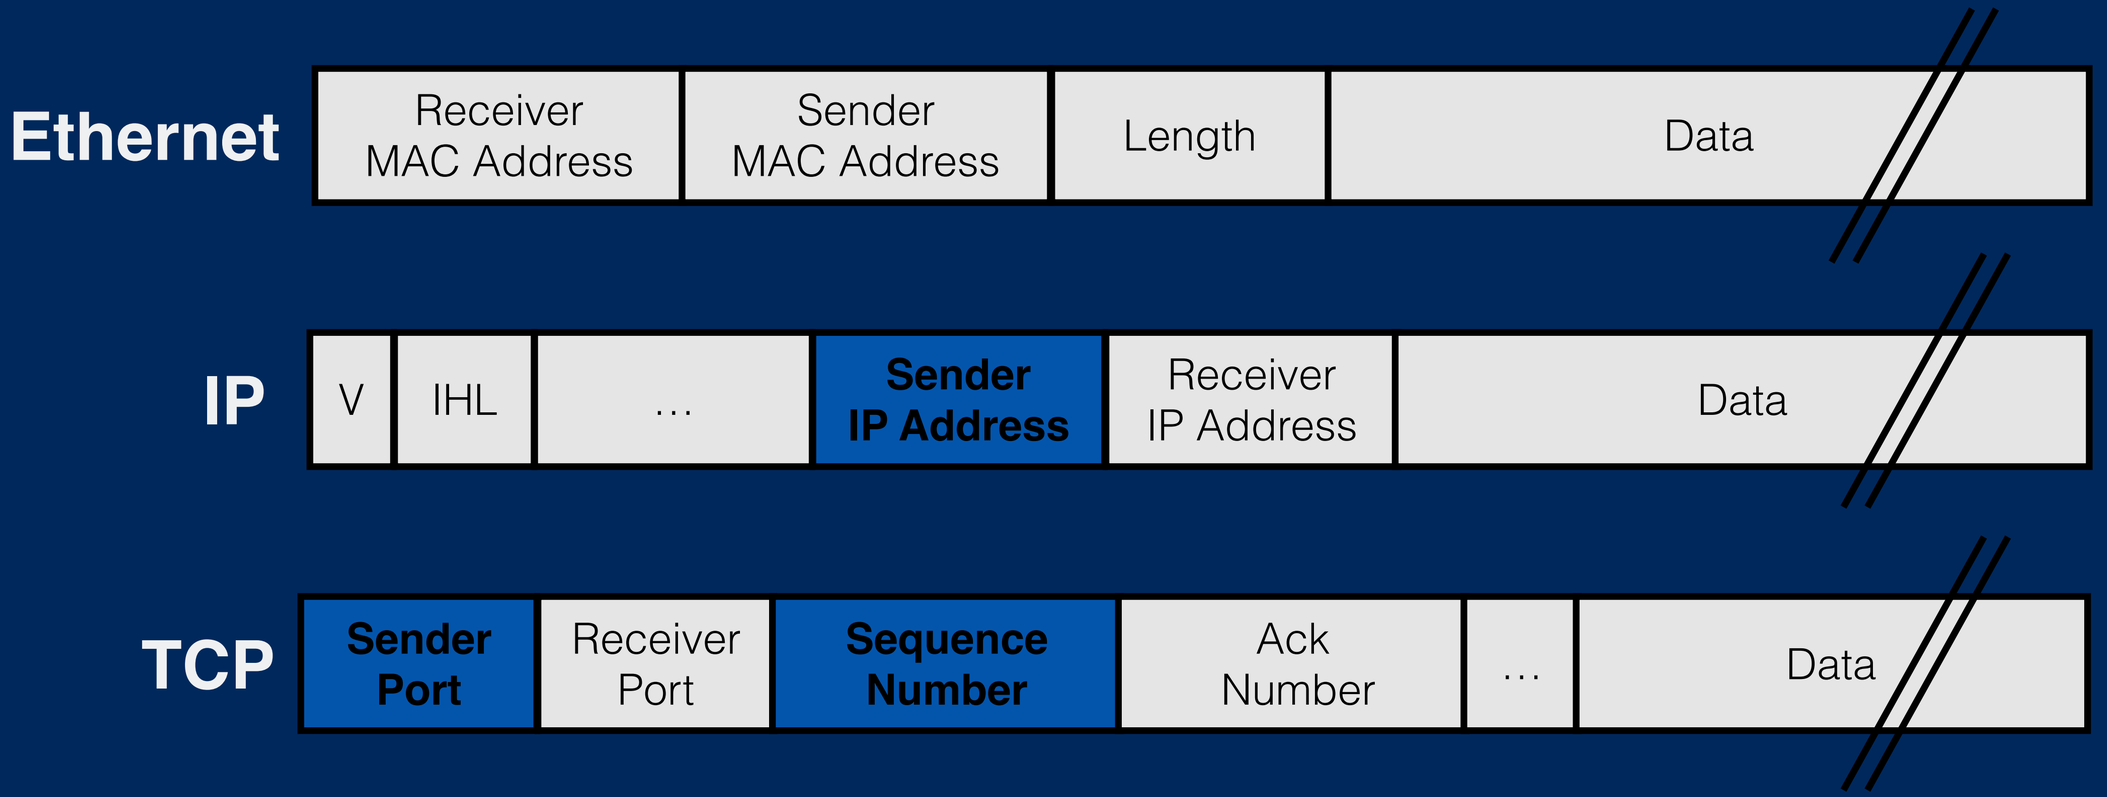

Validating Responses

Encode secrets into mutable fields of probe packets with deterministic effect on responses

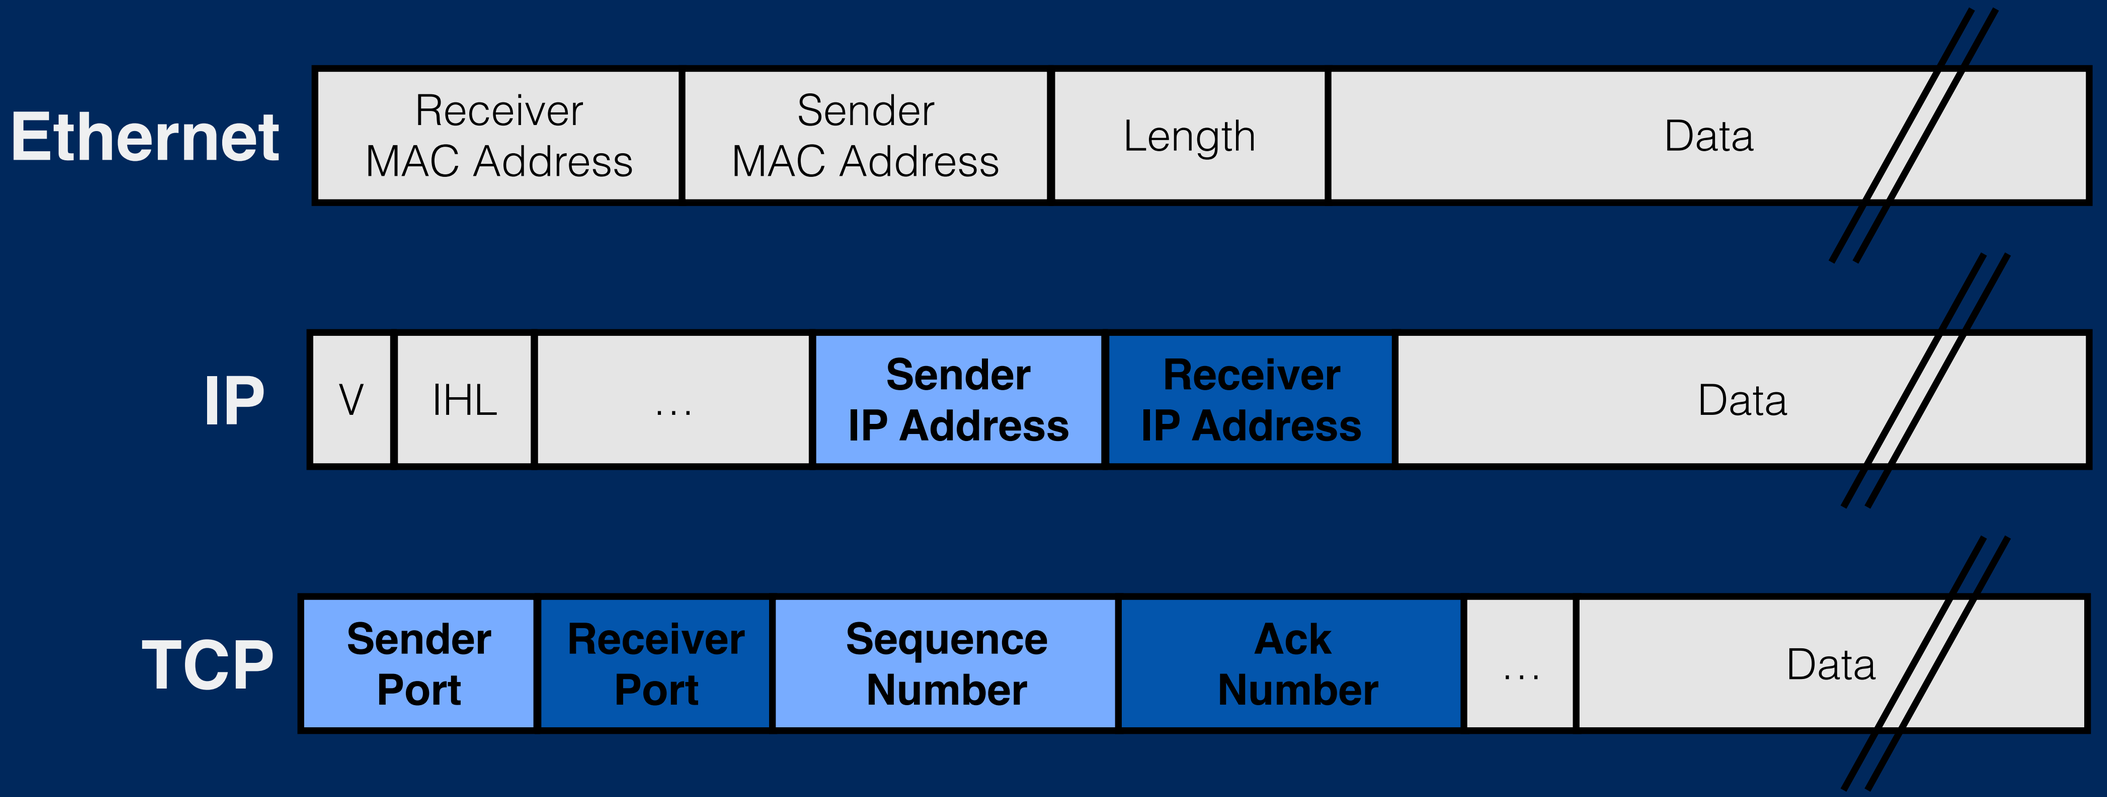

Zero-Copy NIC Access

The Linux kernel is not capable of sending 64 byte packets at 10 GigE linespeed – 14.88 million packets per second

Use the PF_RING ZC library for direct NIC “zero-copy” access to reach linespeed

From 45 minutes to 5 minutes!

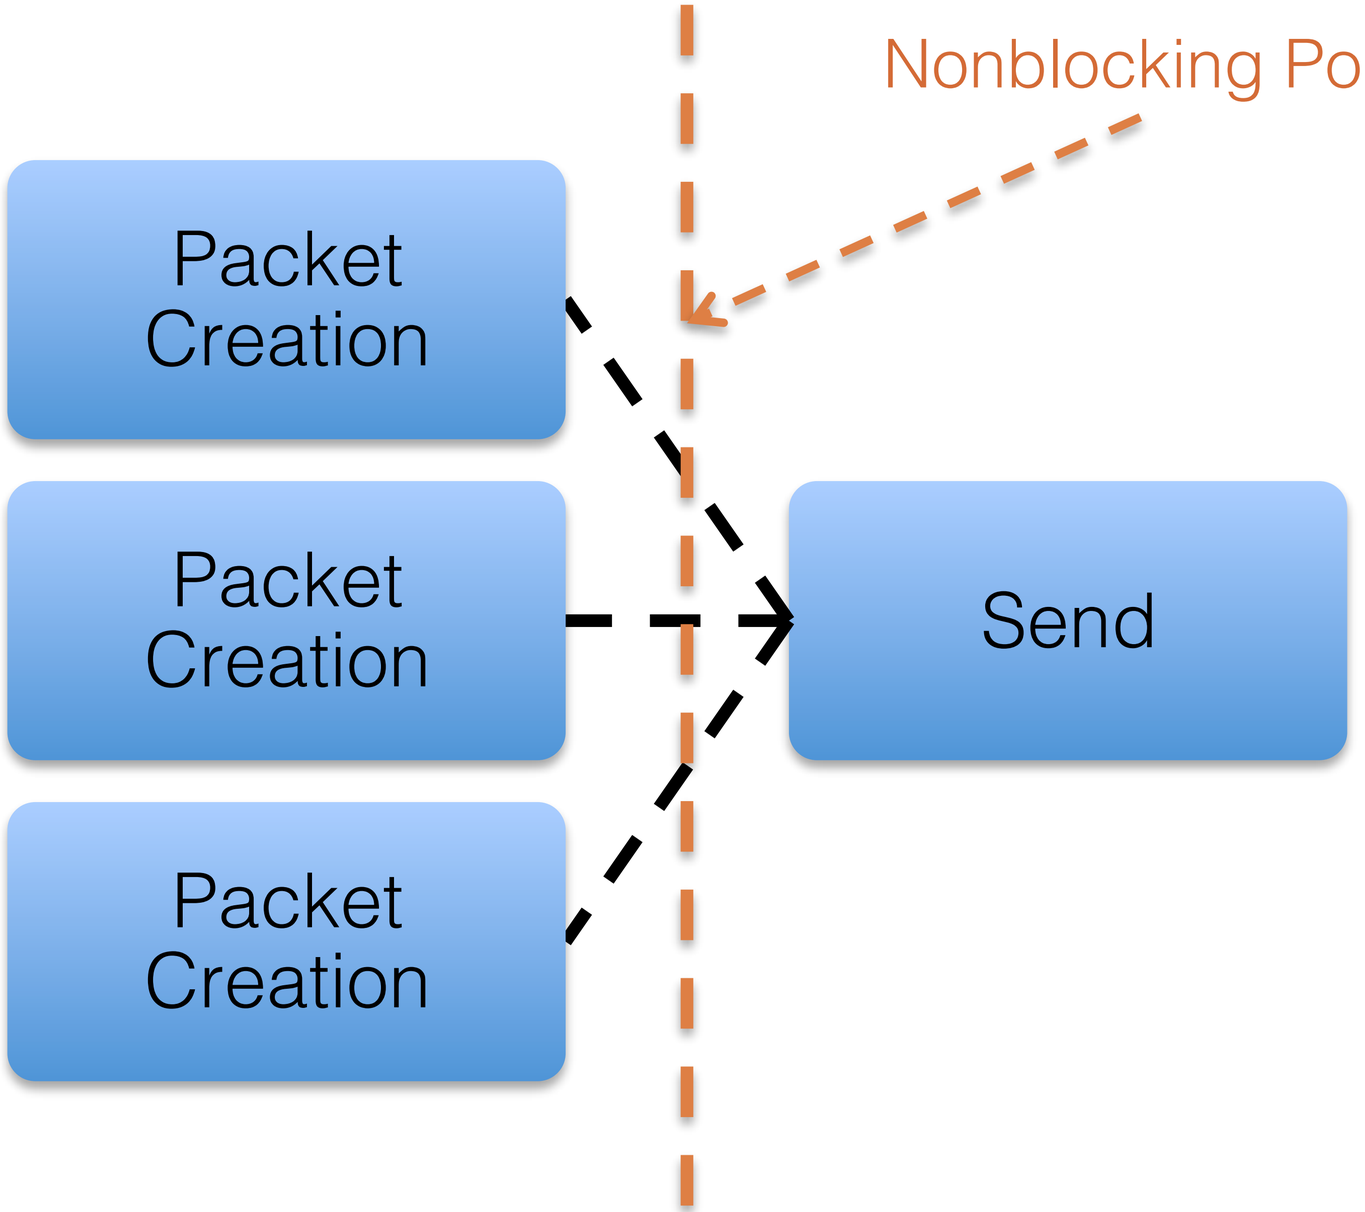
Nonblocking Poll
Packet Creation
Packet Creation
Packet Creation
Send

Internet-Scanning and Email Security

Email Security in Practice

As originally conceived, SMTP had no built-in security.

We've extended with SMTP with new extensions to:

1. Encrypt e-mail in transit
2. Authenticate email on receipt

STARTTLS: TLS Between Hops

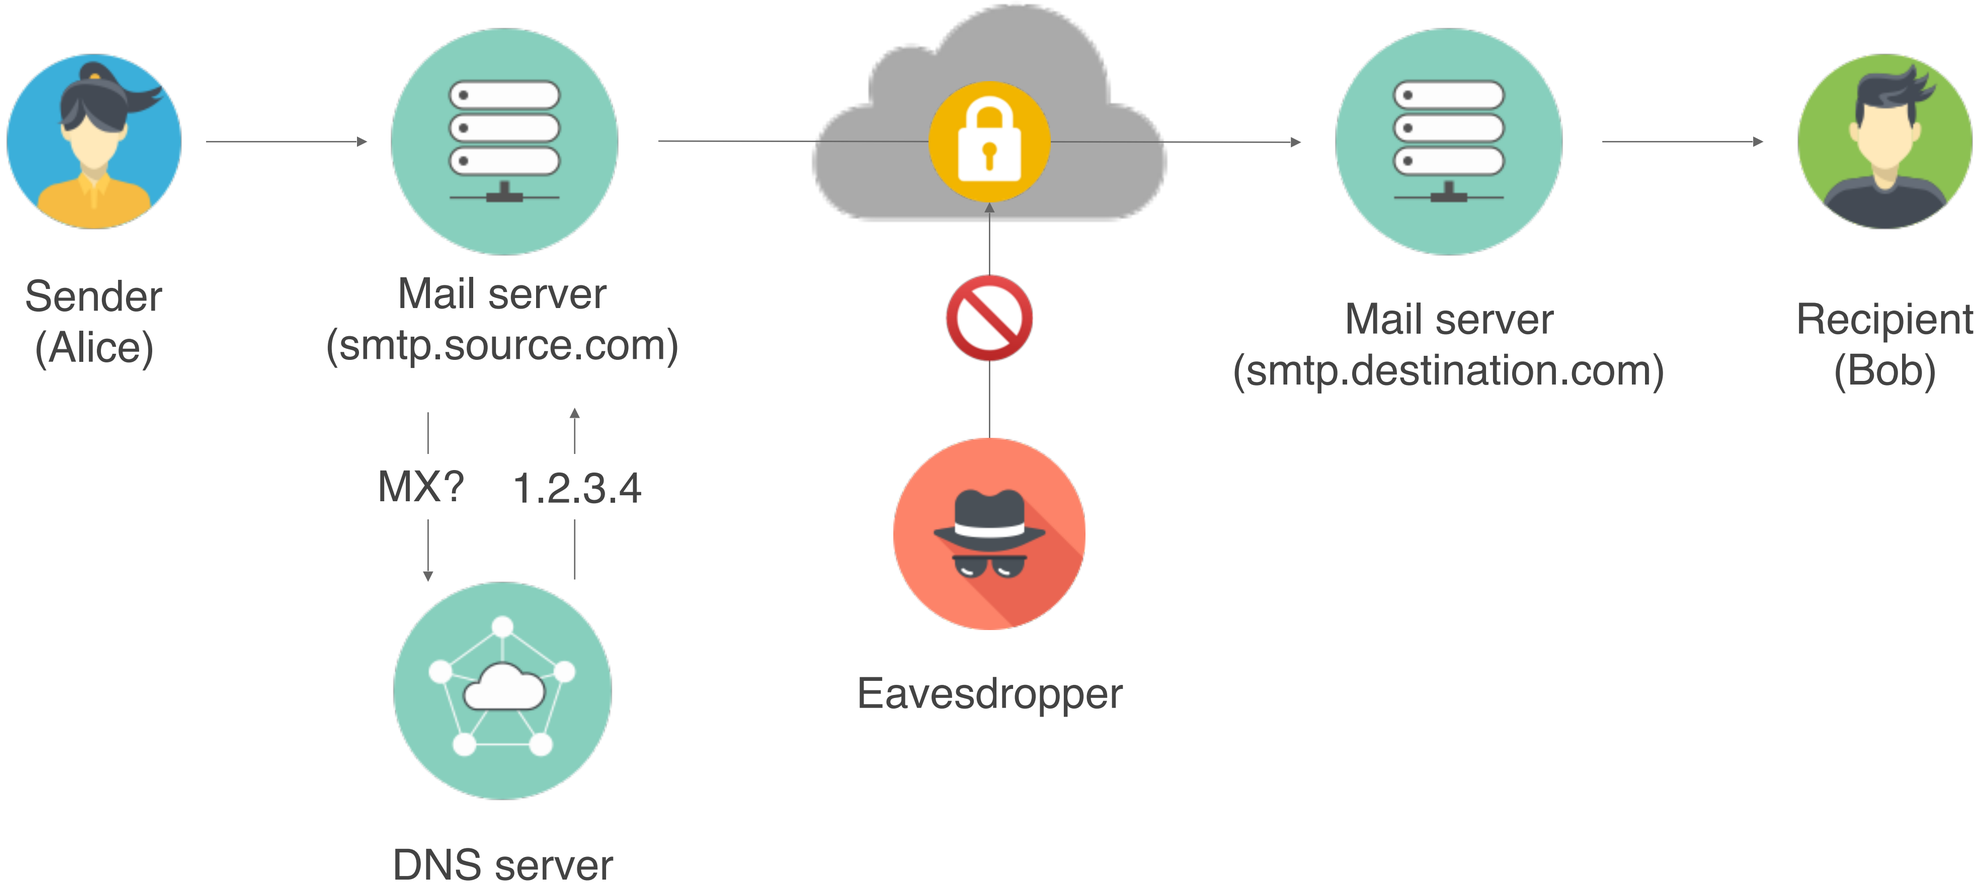

STARTTLS Protocol

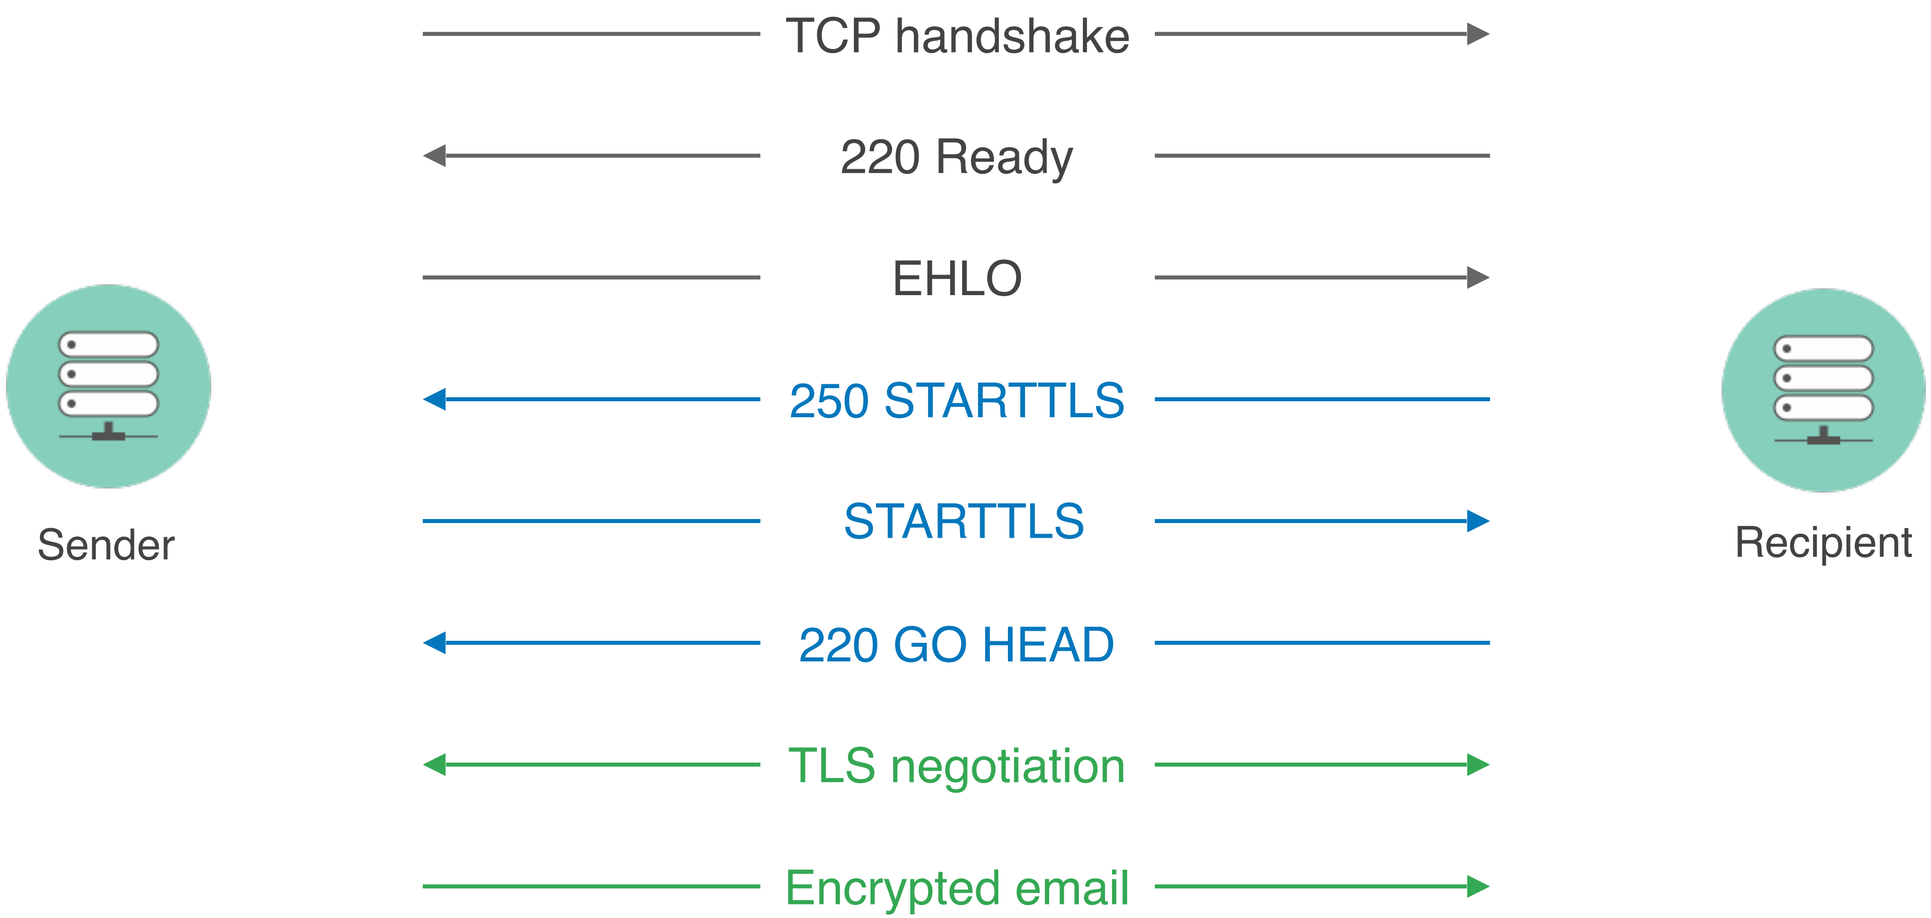

STARTTLS As Seen By Gmail

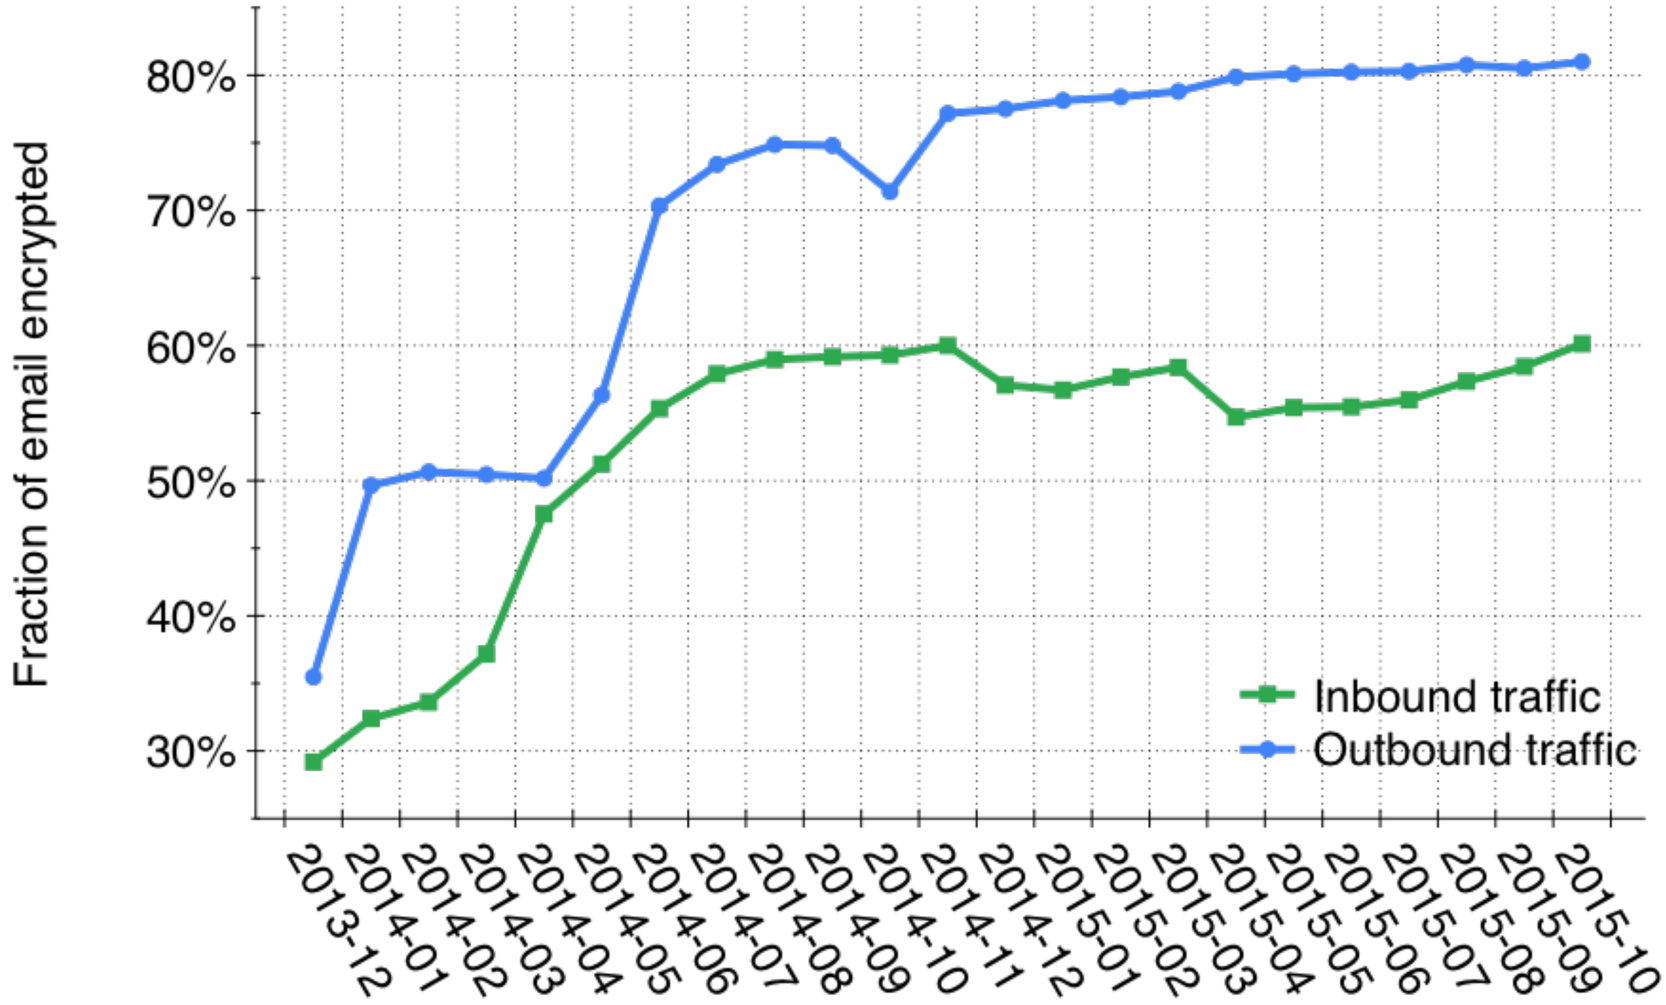

STARTTLS Stripping Attack

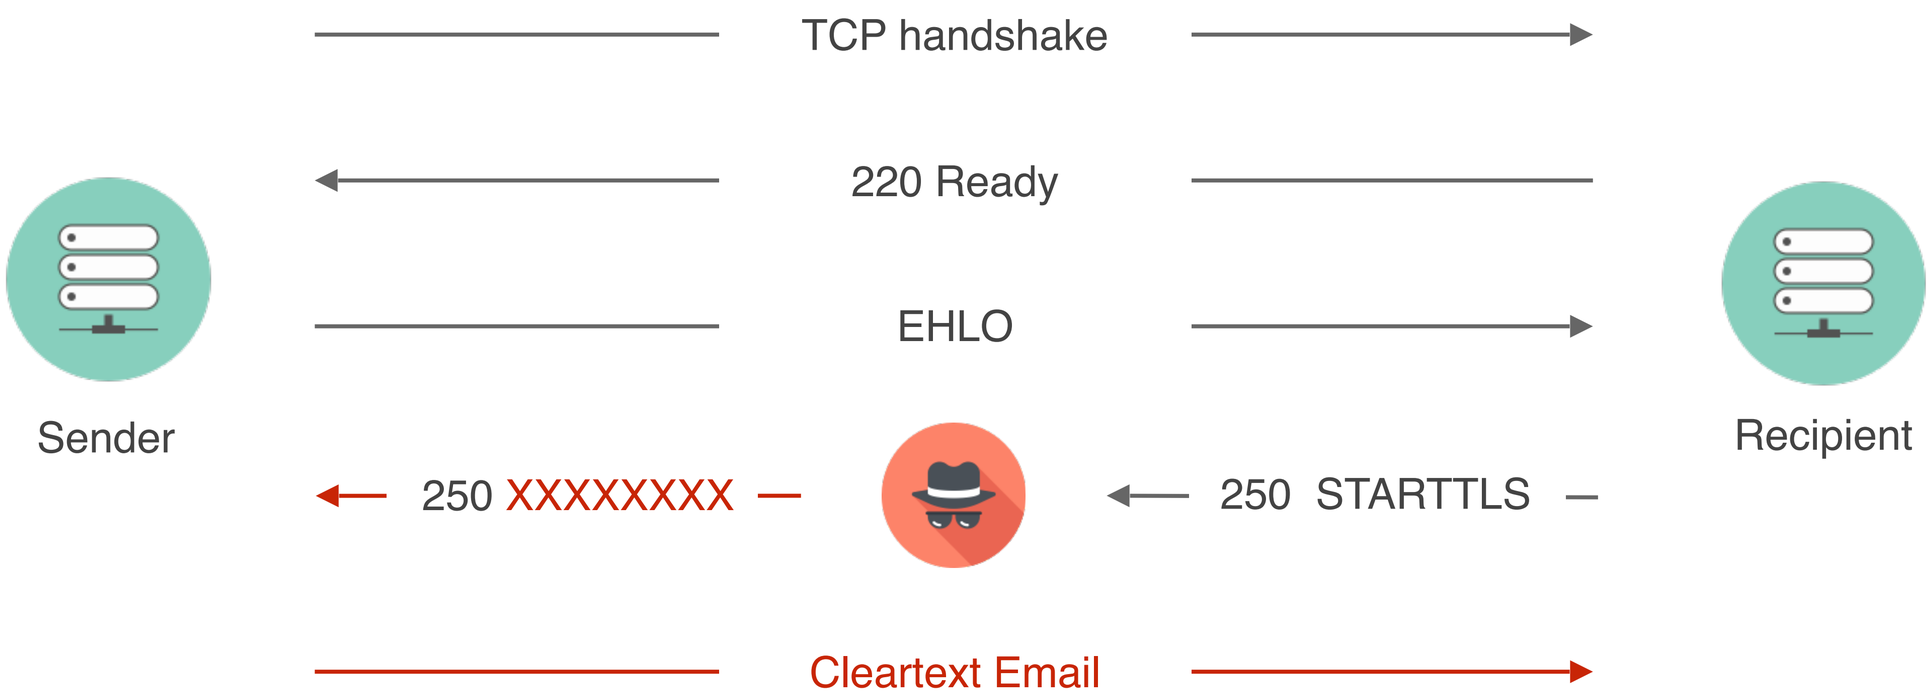

STARTTLS Stripping in the Wild

Country	
Tunisia	96.1%
Iraq	25.6%
Papua New Guinea	25.0%
Nepal	24.3%
Kenya	24.1%
Uganda	23.3%
Lesotho	20.3%
Sierra Leone	13.4%
New Caledonia	10.1%
Zambia	10.0%

Missing STARTTLS Marked Insecure

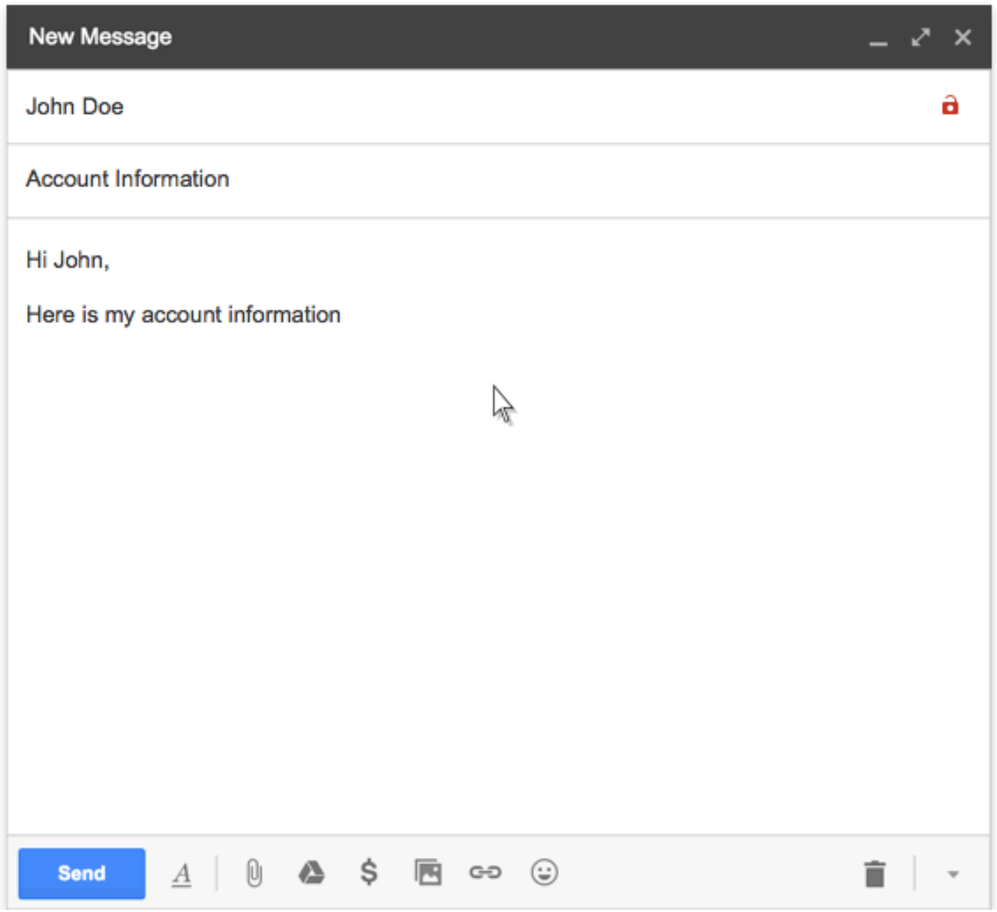

Missing STARTTLS Marked Insecure

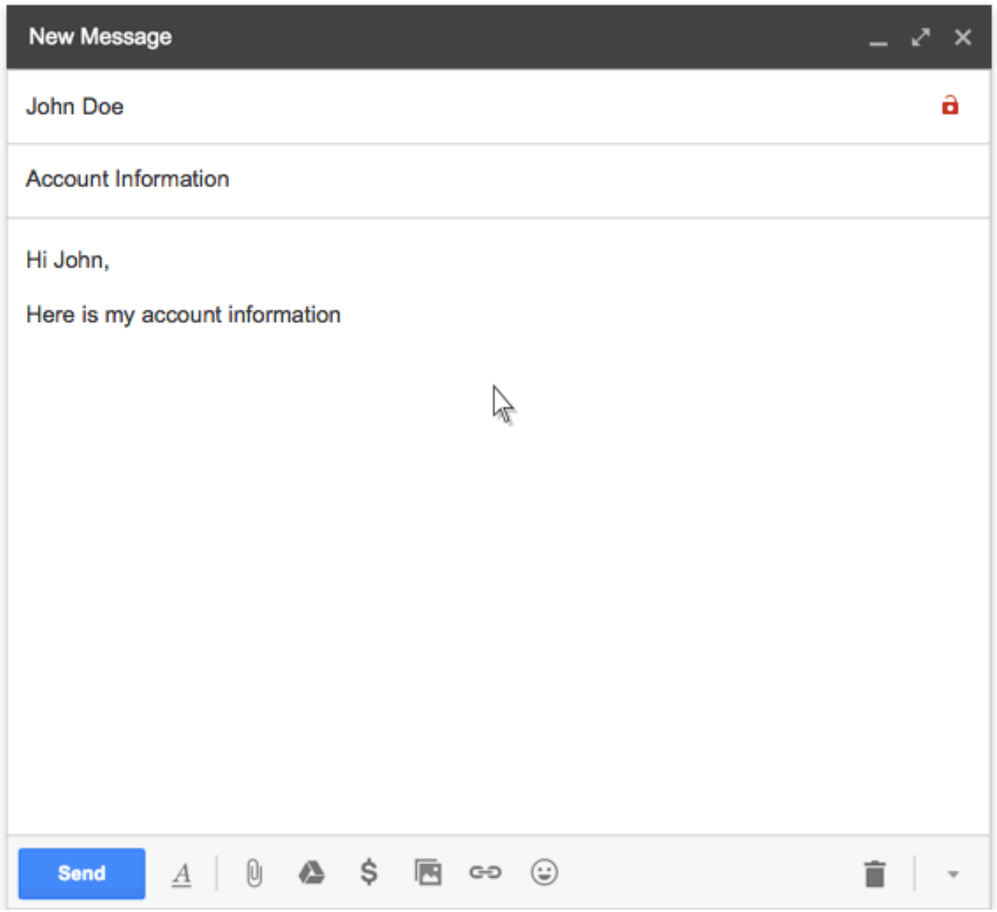

Censys

ZMap Vision

Goals

Enable new and exciting research

Decrease the barriers to entry for Internet-wide surveys

Anyone can scan the entire Internet using a single host

ZMap Vision

Goals	Reality
Enable new and exciting research	Not all researchers can run ZMap
Decrease the barriers to entry for Internet-wide surveys	Negotiate with network administrators for bandwidth and address space
Anyone can scan the entire Internet using a single host	Maintain an opt-out list and respond to complaints

Search engine that allows researchers to **ask questions** about the *devices* and *networks* that compose the Internet

Example

What SMTP servers are having STARTTLS stripped?

25.smtp.starttls.ehlo: *XXXX*

Search

Example

What SMTP servers are having STARTTLS stripped?

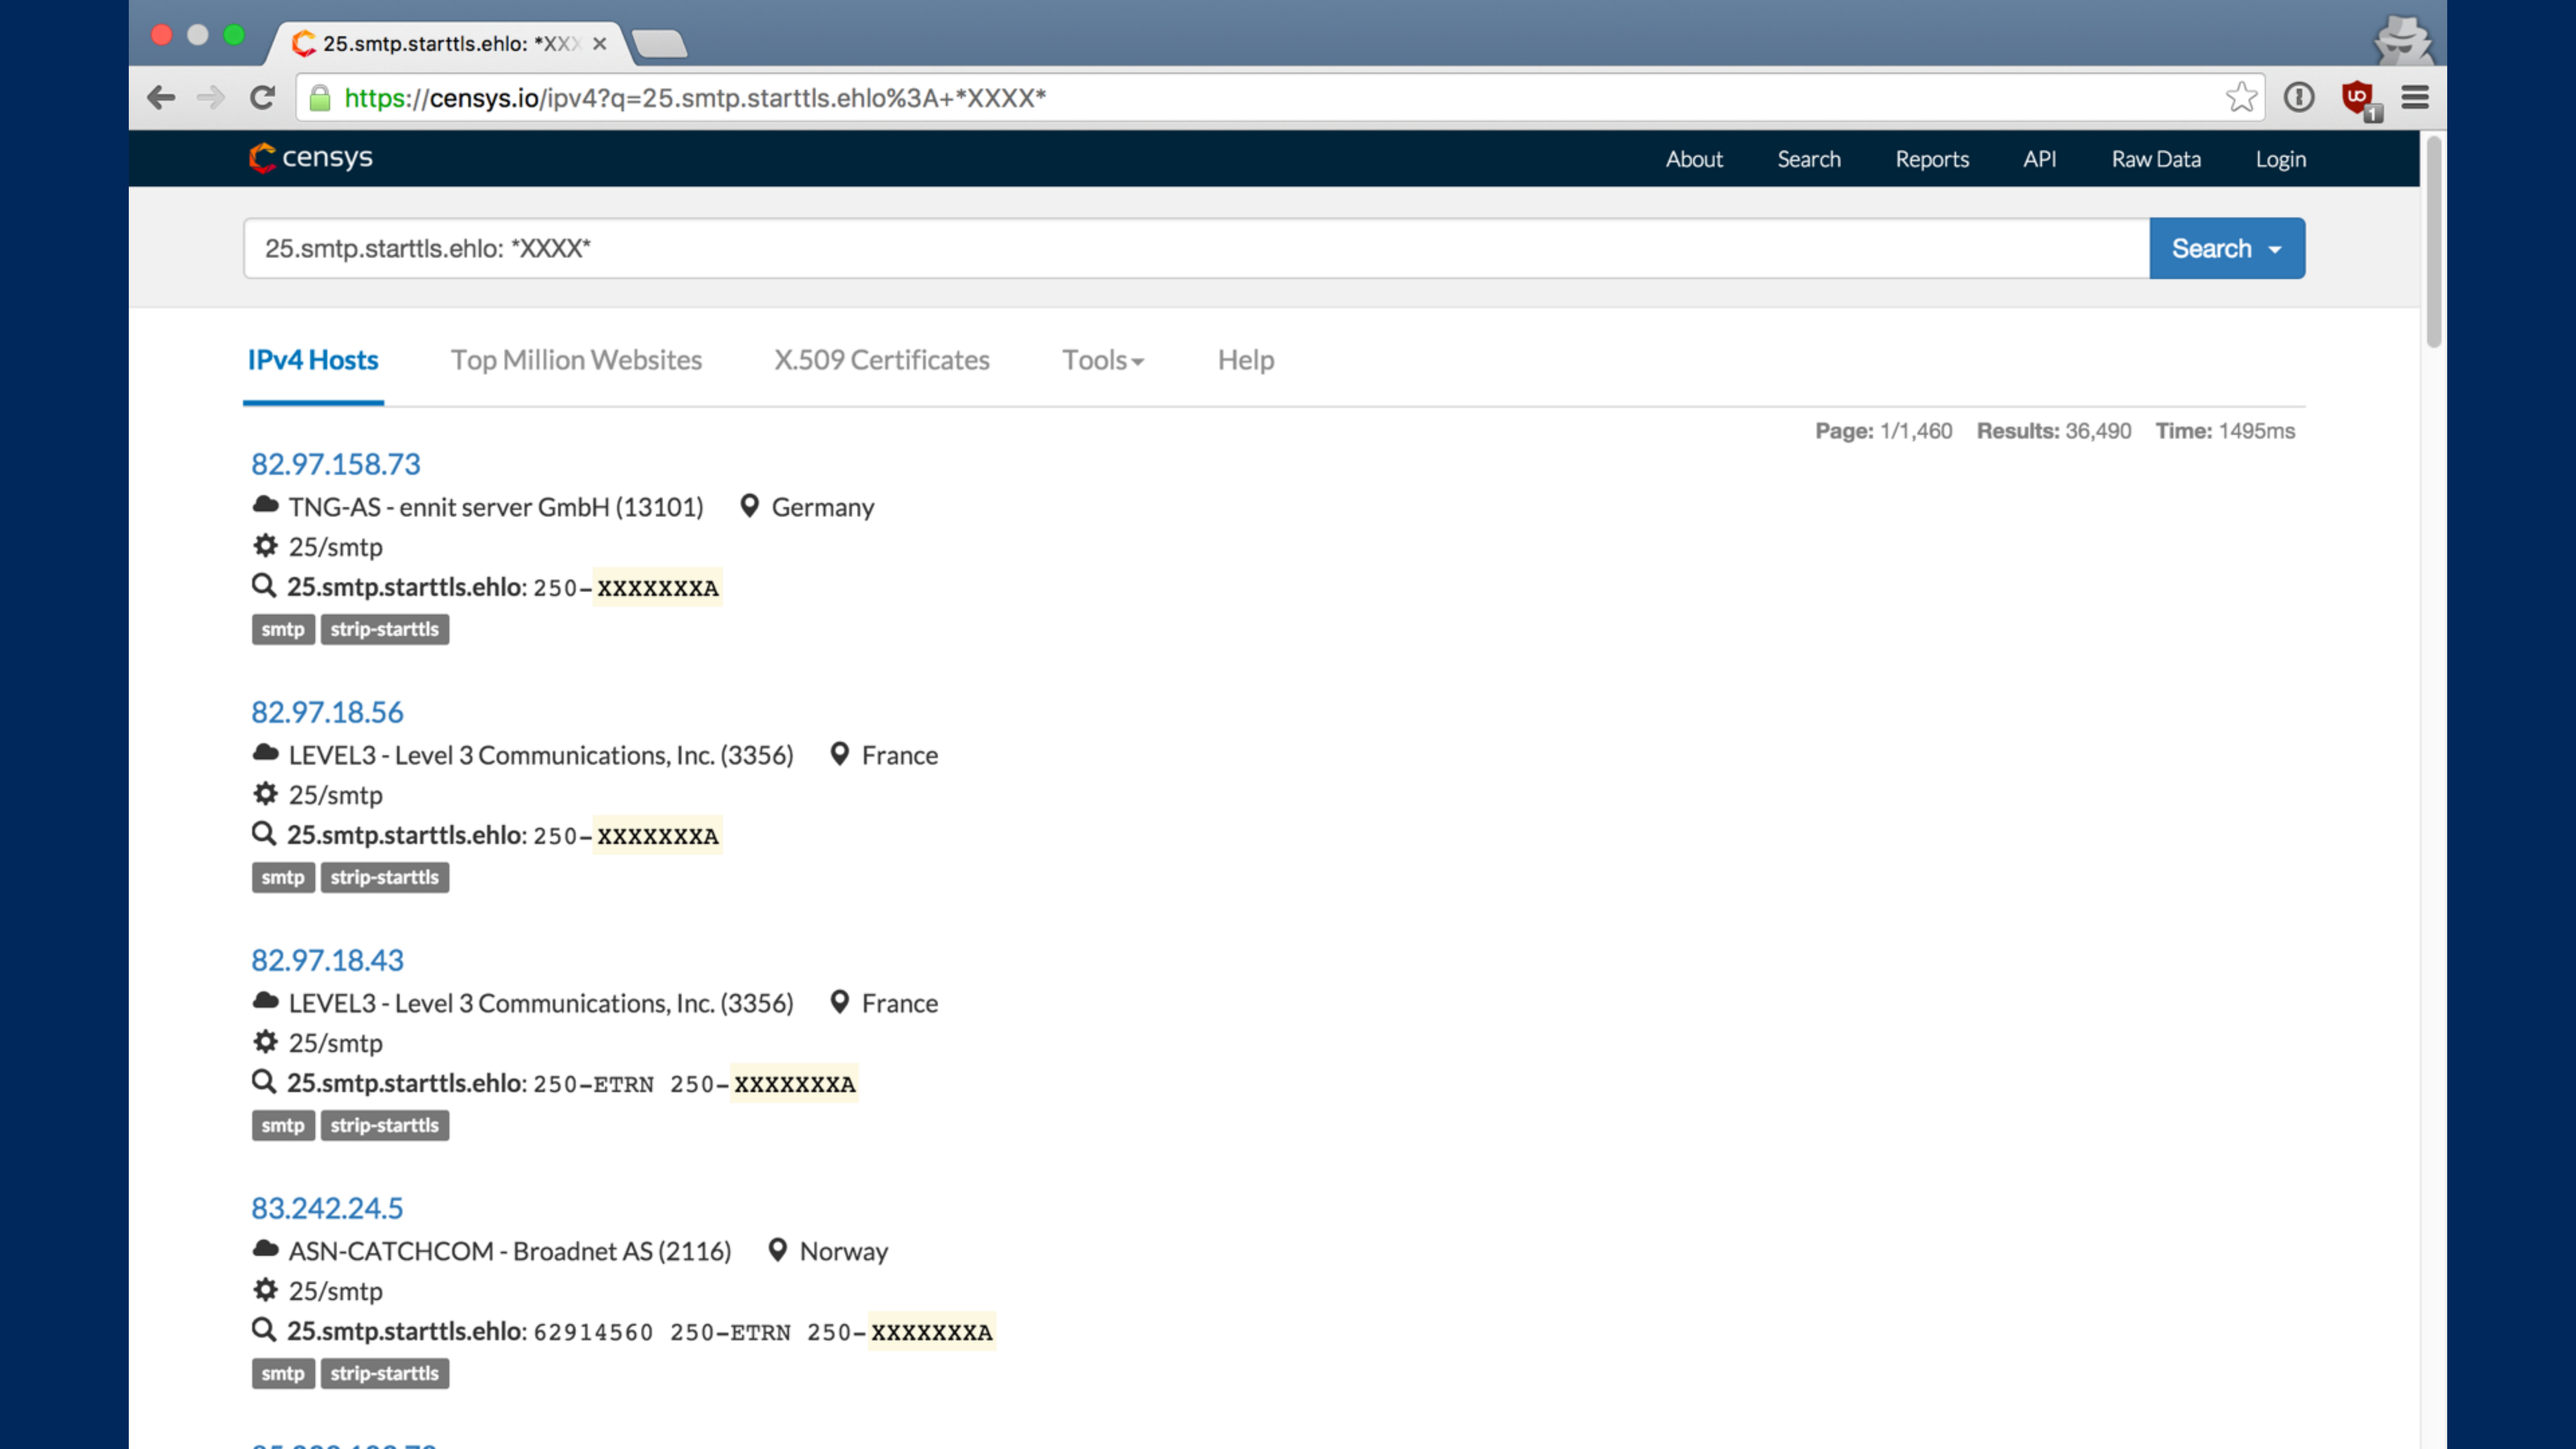
25.smtp.starttls.ehlo: *XXX
https://censys.io/ipv4?q=25.smtp.starttls.ehlo%3A+*XXXX*
censys
About
Search
Reports
API
Raw Data
Login
25.smtp.starttls.ehlo: *XXXX*
Search
IPv4 Hosts
Top Million Websites
X.509 Certificates
Tools
Help
Page: 1/1,460 Results: 36,490 Time: 1495ms
82.97.158.73
TNG-AS - ennit server GmbH (13101)
Germany
25/smtp
25.smtp.starttls.ehlo: 250-XXXXXXXA
smtp
strip-starttls
82.97.18.56
LEVEL3 - Level 3 Communications, Inc. (3356)
France
25/smtp
25.smtp.starttls.ehlo: 250-XXXXXXXA
smtp
strip-starttls
82.97.18.43
LEVEL3 - Level 3 Communications, Inc. (3356)
France
25/smtp
25.smtp.starttls.ehlo: 250-ETRN 250-XXXXXXXA
smtp
strip-starttls
83.242.24.5
ASN-CATCHCOM - Broadnet AS (2116)
Norway
25/smtp
25.smtp.starttls.ehlo: 62914560 250-ETRN 250-XXXXXXXA
smtp
strip-starttls

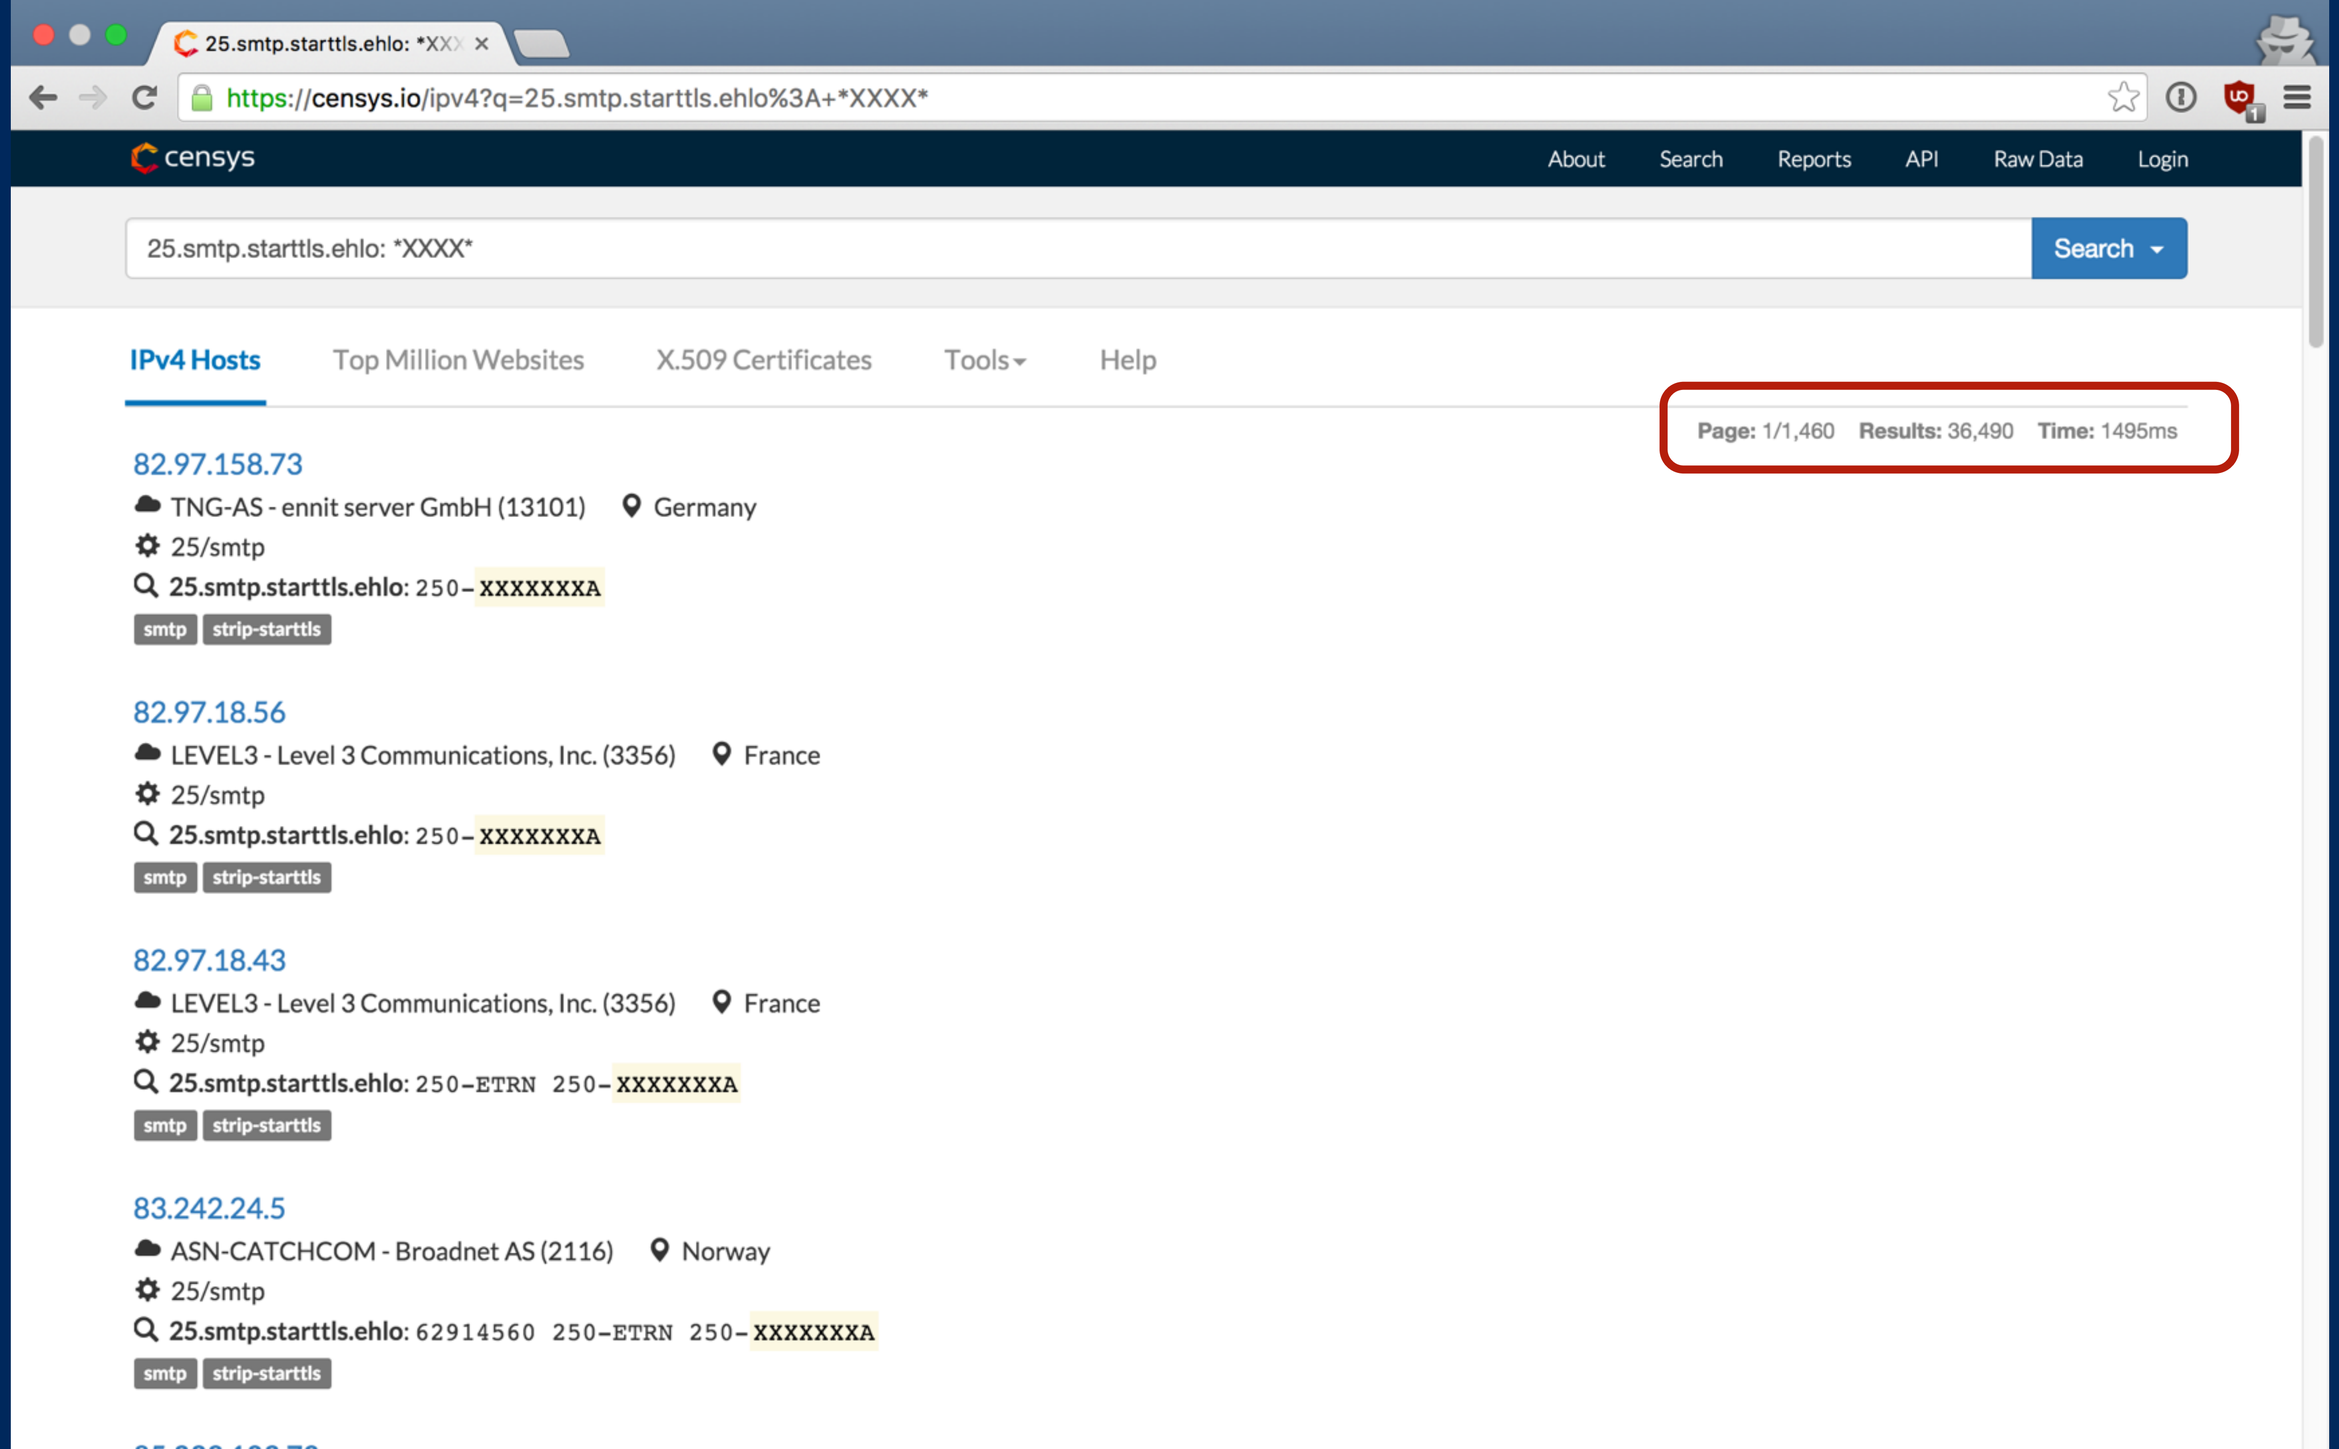
25.smtp.starttls.ehlo: *XXX
https://censys.io/ipv4?q=25.smtp.starttls.ehlo%3A+*XXXX*
censys
About
Search
Reports
API
Raw Data
Login
25.smtp.starttls.ehlo: *XXXX*
Search
IPv4 Hosts
Top Million Websites
X.509 Certificates
Tools
Help
Page: 1/1,460 Results: 36,490 Time: 1495ms
82.97.158.73
TNG-AS - ennit server GmbH (13101)
Germany
25/smtp
25.smtp.starttls.ehlo: 250-XXXXXXXA
smtp
strip-starttls
82.97.18.56
LEVEL3 - Level 3 Communications, Inc. (3356)
France
25/smtp
25.smtp.starttls.ehlo: 250-XXXXXXXA
smtp
strip-starttls
82.97.18.43
LEVEL3 - Level 3 Communications, Inc. (3356)
France
25/smtp
25.smtp.starttls.ehlo: 250-ETRN 250-XXXXXXXA
smtp
strip-starttls
83.242.24.5
ASN-CATCHCOM - Broadnet AS (2116)
Norway
25/smtp
25.smtp.starttls.ehlo: 62914560 250-ETRN 250-XXXXXXXA
smtp
strip-starttls

Censys

https://censys.io/ipv4/report?q=25.smtp.starttls.ehlo%3A+*XXXX*&field=location.country_code&max_buckets=10

censys

About Search Reports API Raw Data Login

location.country_code 10 Build Report

Host Report

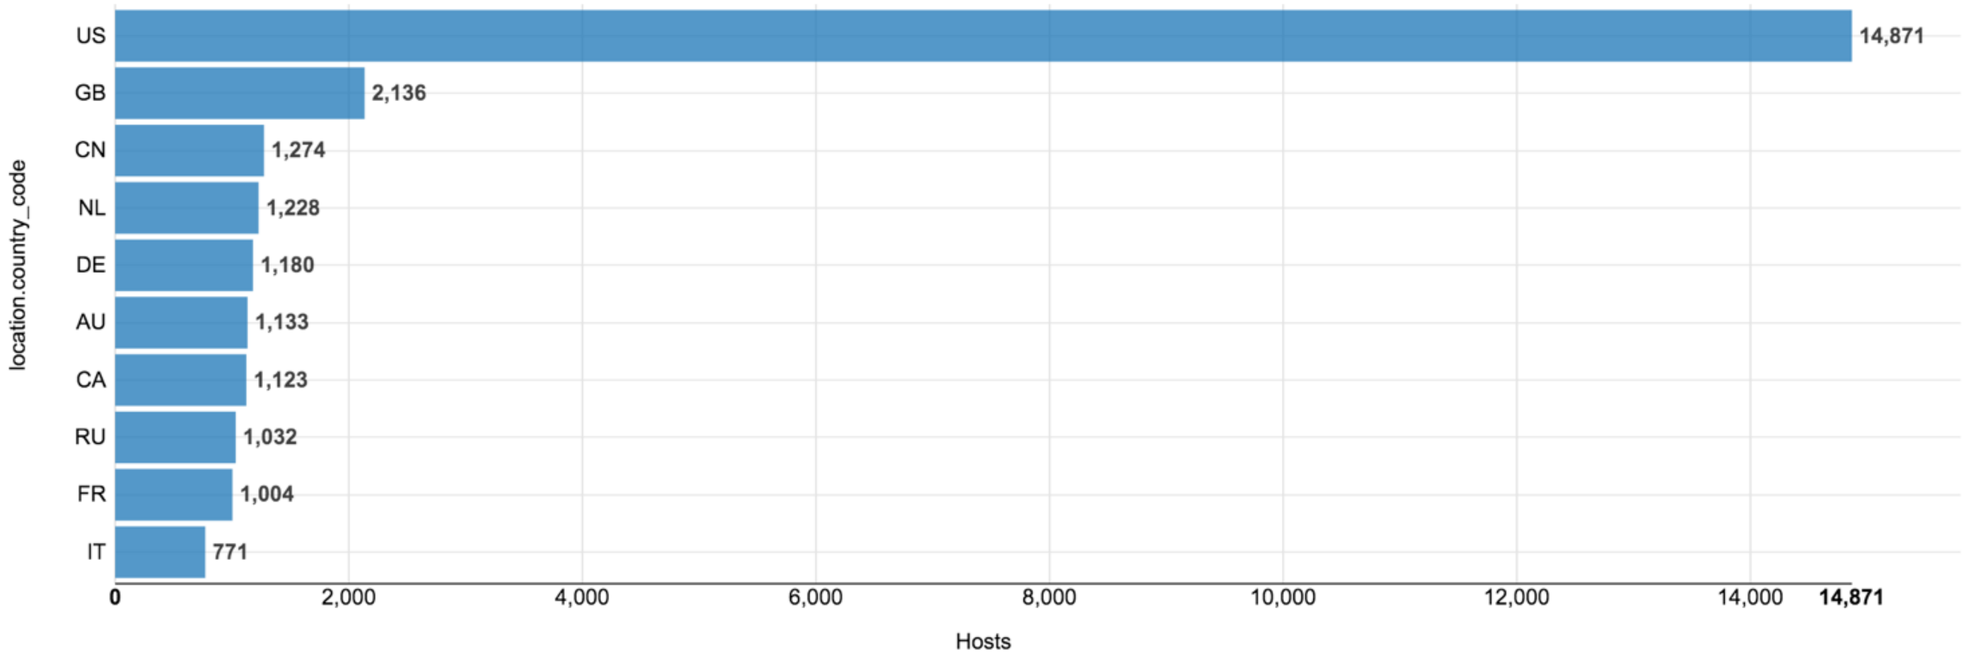

Raw Data

location.country_code	Hosts	
US	14,871	40.75%
GB	2,136	5.85%

Full-text search

Current and historical data

SQL

API

https://censys.io

Contributing

Are you extending ZMap, ZGrab, or another scanner with a new protocol?

Do you have annotations to add to our framework?

We'll work with researchers to add new scan modules to Censys

https://github.com/zmap/zmap

https://github.com/zmap/zgrab

https://github.com/zmap/ztag