An Internet-Wide View of Internet-Wide Scanning

Zakir Durumeric, Michael Bailey, J. Alex Halderman

University of Michigan

Internet-Wide Scanning

We released ZMap at USENIX Security last year

- TCP Scan of full IPv4 in < 45 minutes

Internet-Wide scanning appears to be useful

- 15 studies based on ZMap data

masscan

Internet-Wide Scanning

We released ZMap at USENIX Security last year

- TCP Scan of full IPv4 in < 45 minutes

Internet-Wide scanning appears to be useful

- 15 studies based on ZMap data

Who is using ZMap?

Did ZMap alter the scanning landscape?

Are operators now blocking Internet scans?

masscan

Talk Outline

1. **Broad Overview of Scanning Landscape**
2. Case Studies: Scanning triggered by backdoors in home routers, Heartbleed, and NTP vulnerabilities
3. Defensive reactions against scanning

Detecting Internet-Wide Scan Traffic

Data Collection

- Collected background traffic from a large network telescope at Merit Network during 2013–2014
- Darknet does not host any services — probes are likely part of Internet-wide scans
- Approach will likely miss targeted scanning

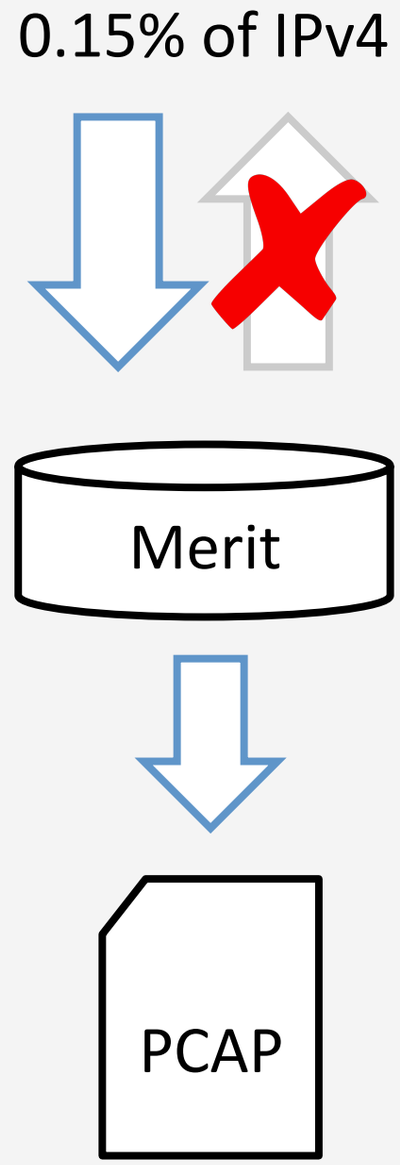

Detecting Internet-Wide Scan Traffic

Data Collection

- Collected background traffic from a large network telescope at Merit Network during 2013–2014
- Darknet does not host any services — probes are likely part of Internet-wide scans
- Approach will likely miss targeted scanning

Estimating Actual Scans

- Assume that scan targets are ordered by a uniform random distribution
- Estimate coverage and scan rate using binomial distribution

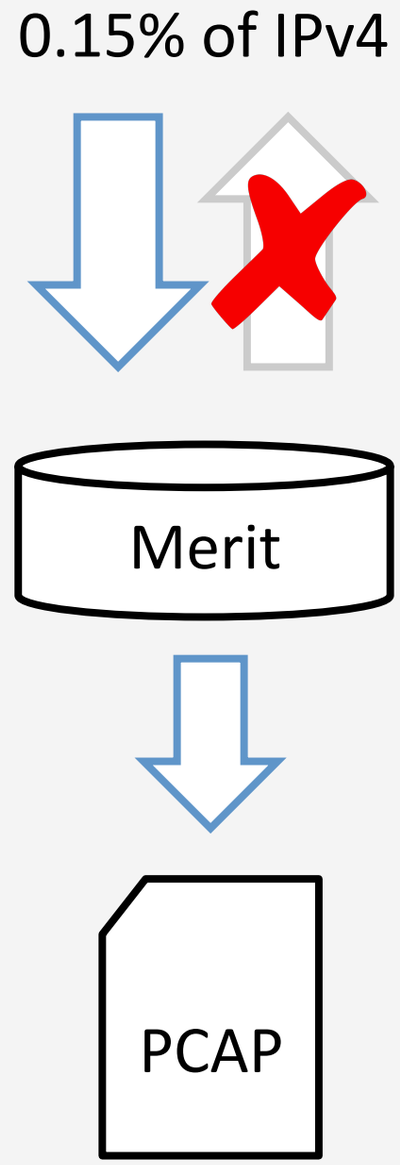

How do we define a “scan”?

Defining a Scan

Destination	targeting a single protocol on a single port
Source	set of contiguous IPs within a single AS
Rate	sending at an estimated rate of 10 pps
Size	reaching ≥100 hosts in our darknet

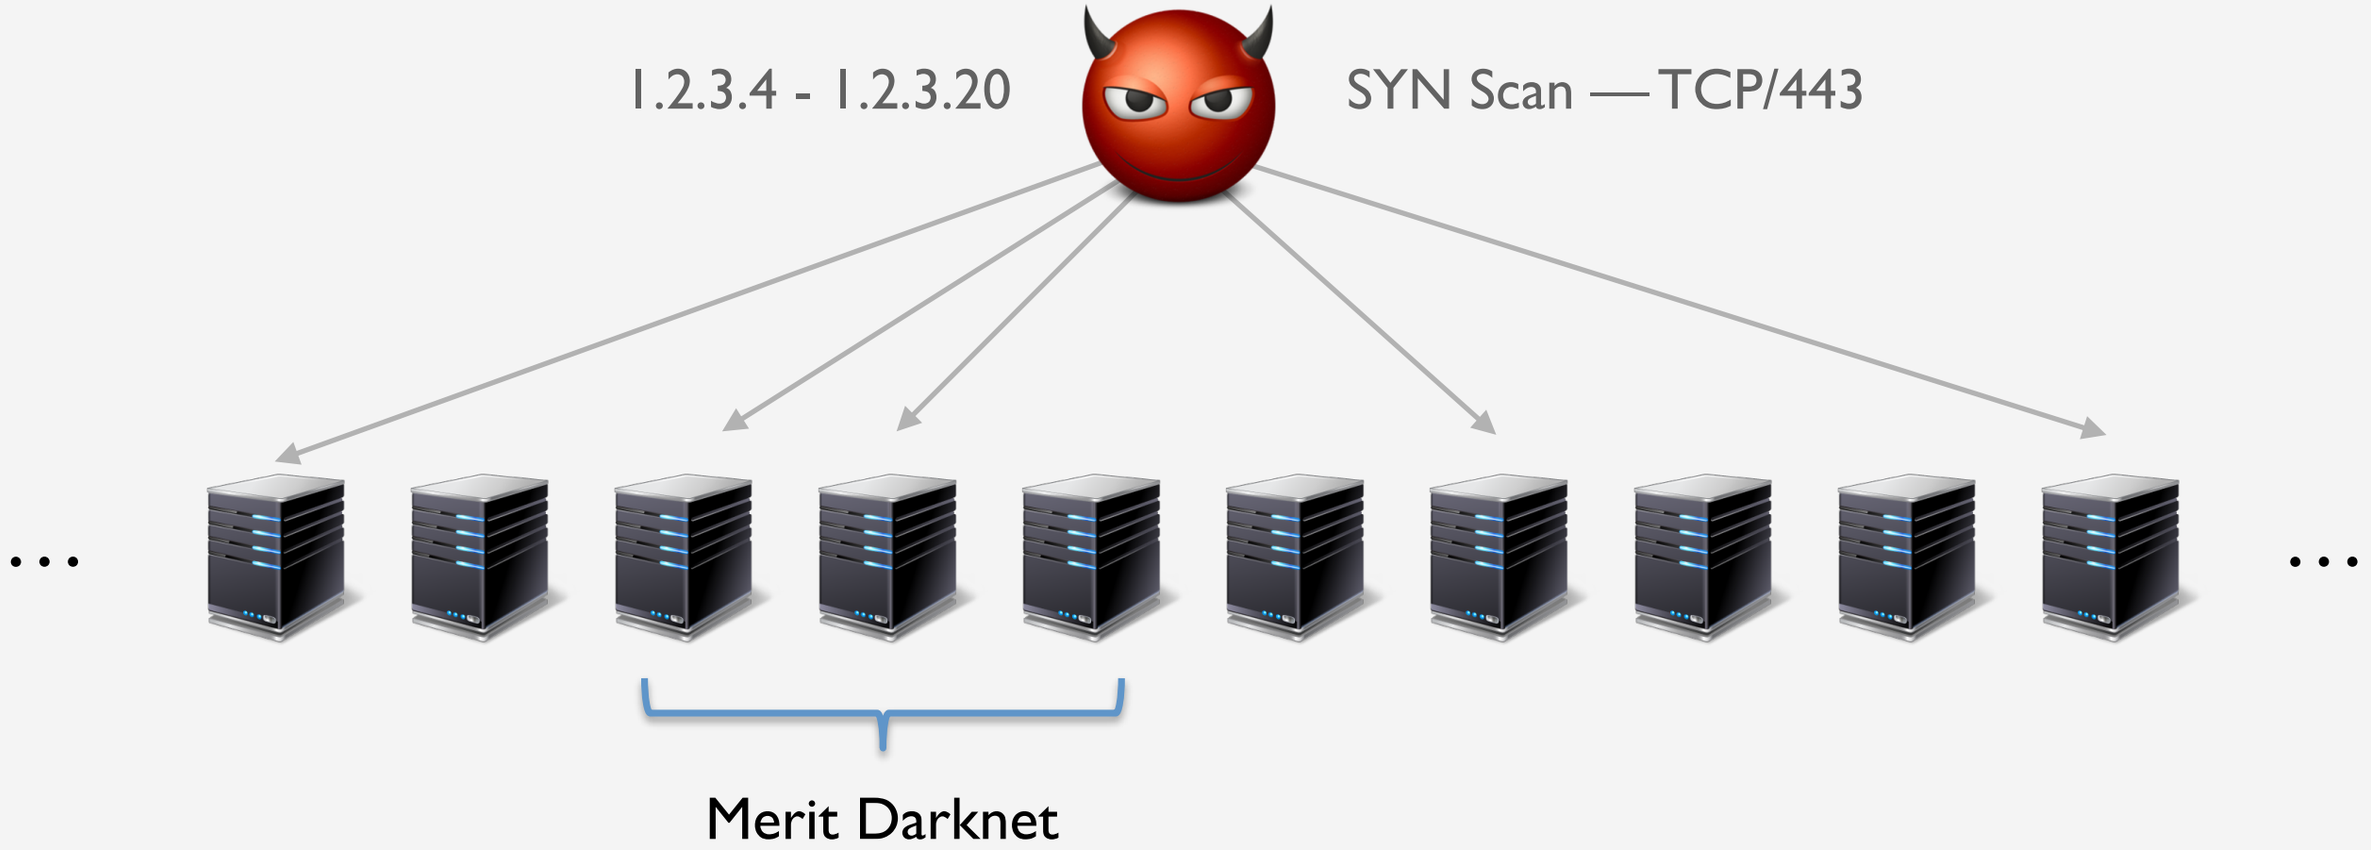

Fingerprinting Scanners

We investigated open source network scanners and created fingerprints for ZMap and masscan

ZMap

- IP ID statically set to 54321

Masscan

- IP ID = dest addr ⊕ dest port ⊕ tcp seqnum

masscan

Network Telescope Traffic Overview

January 2014

- Darknet received an average 1.4 billion packets (55 GB) per day
- Detected 10.8 million scans from 1.8 million unique hosts
- 2,013 ZMap scans and 1,326 masscan scans

Targeted Services

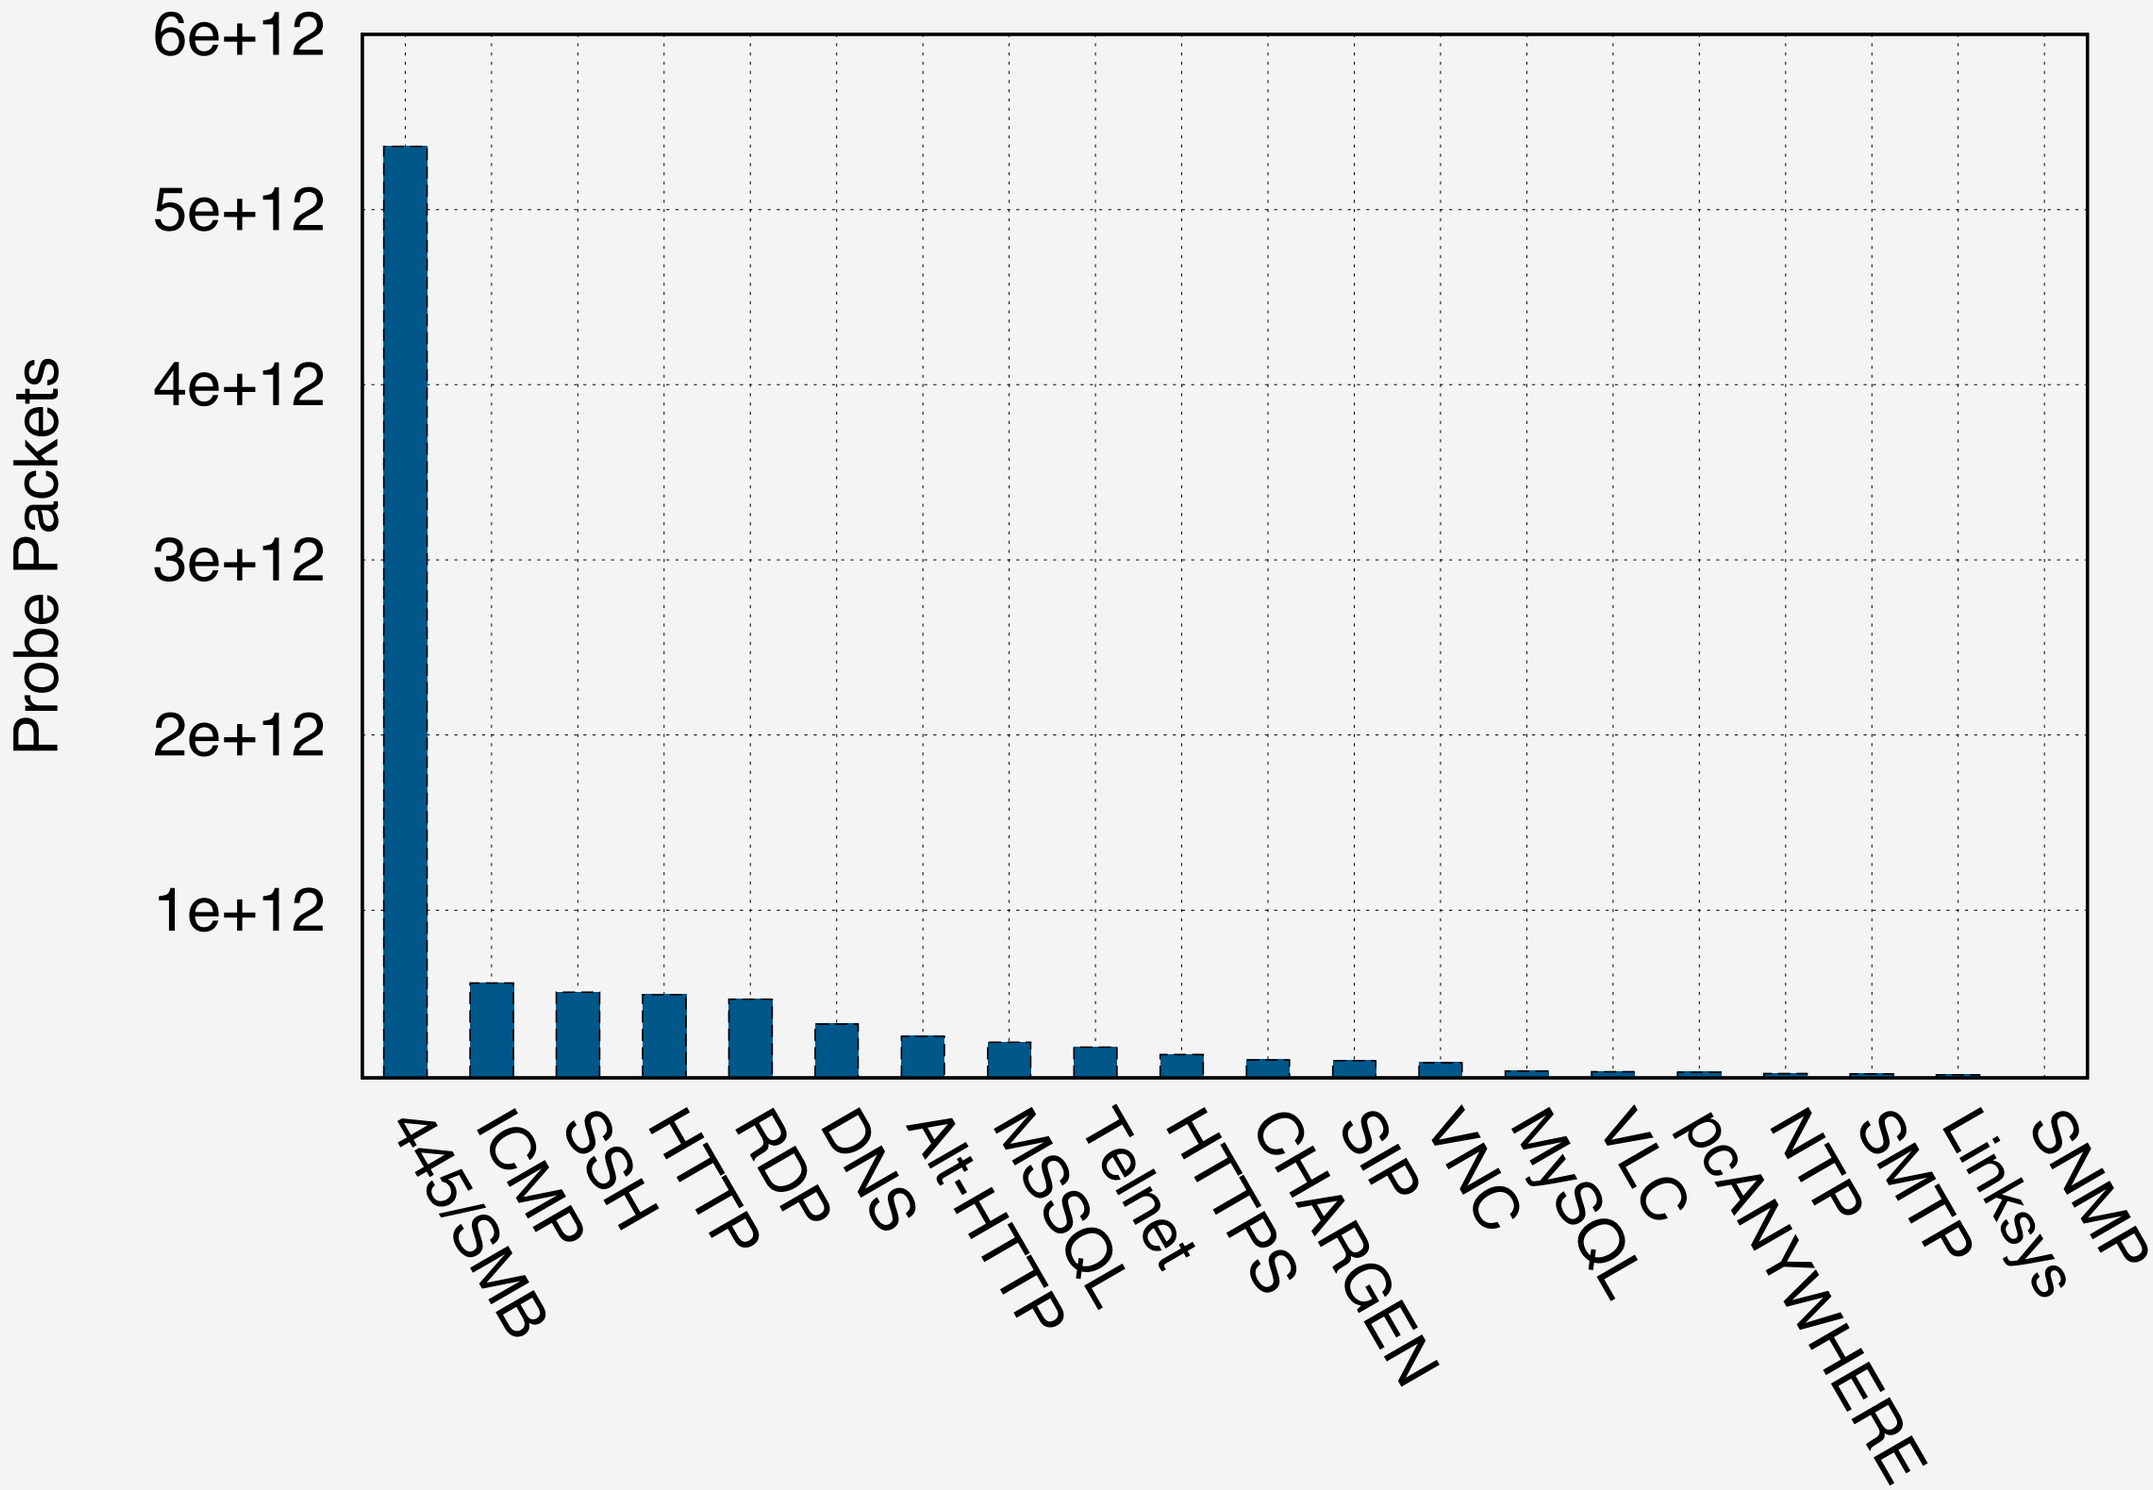

Targeted Services

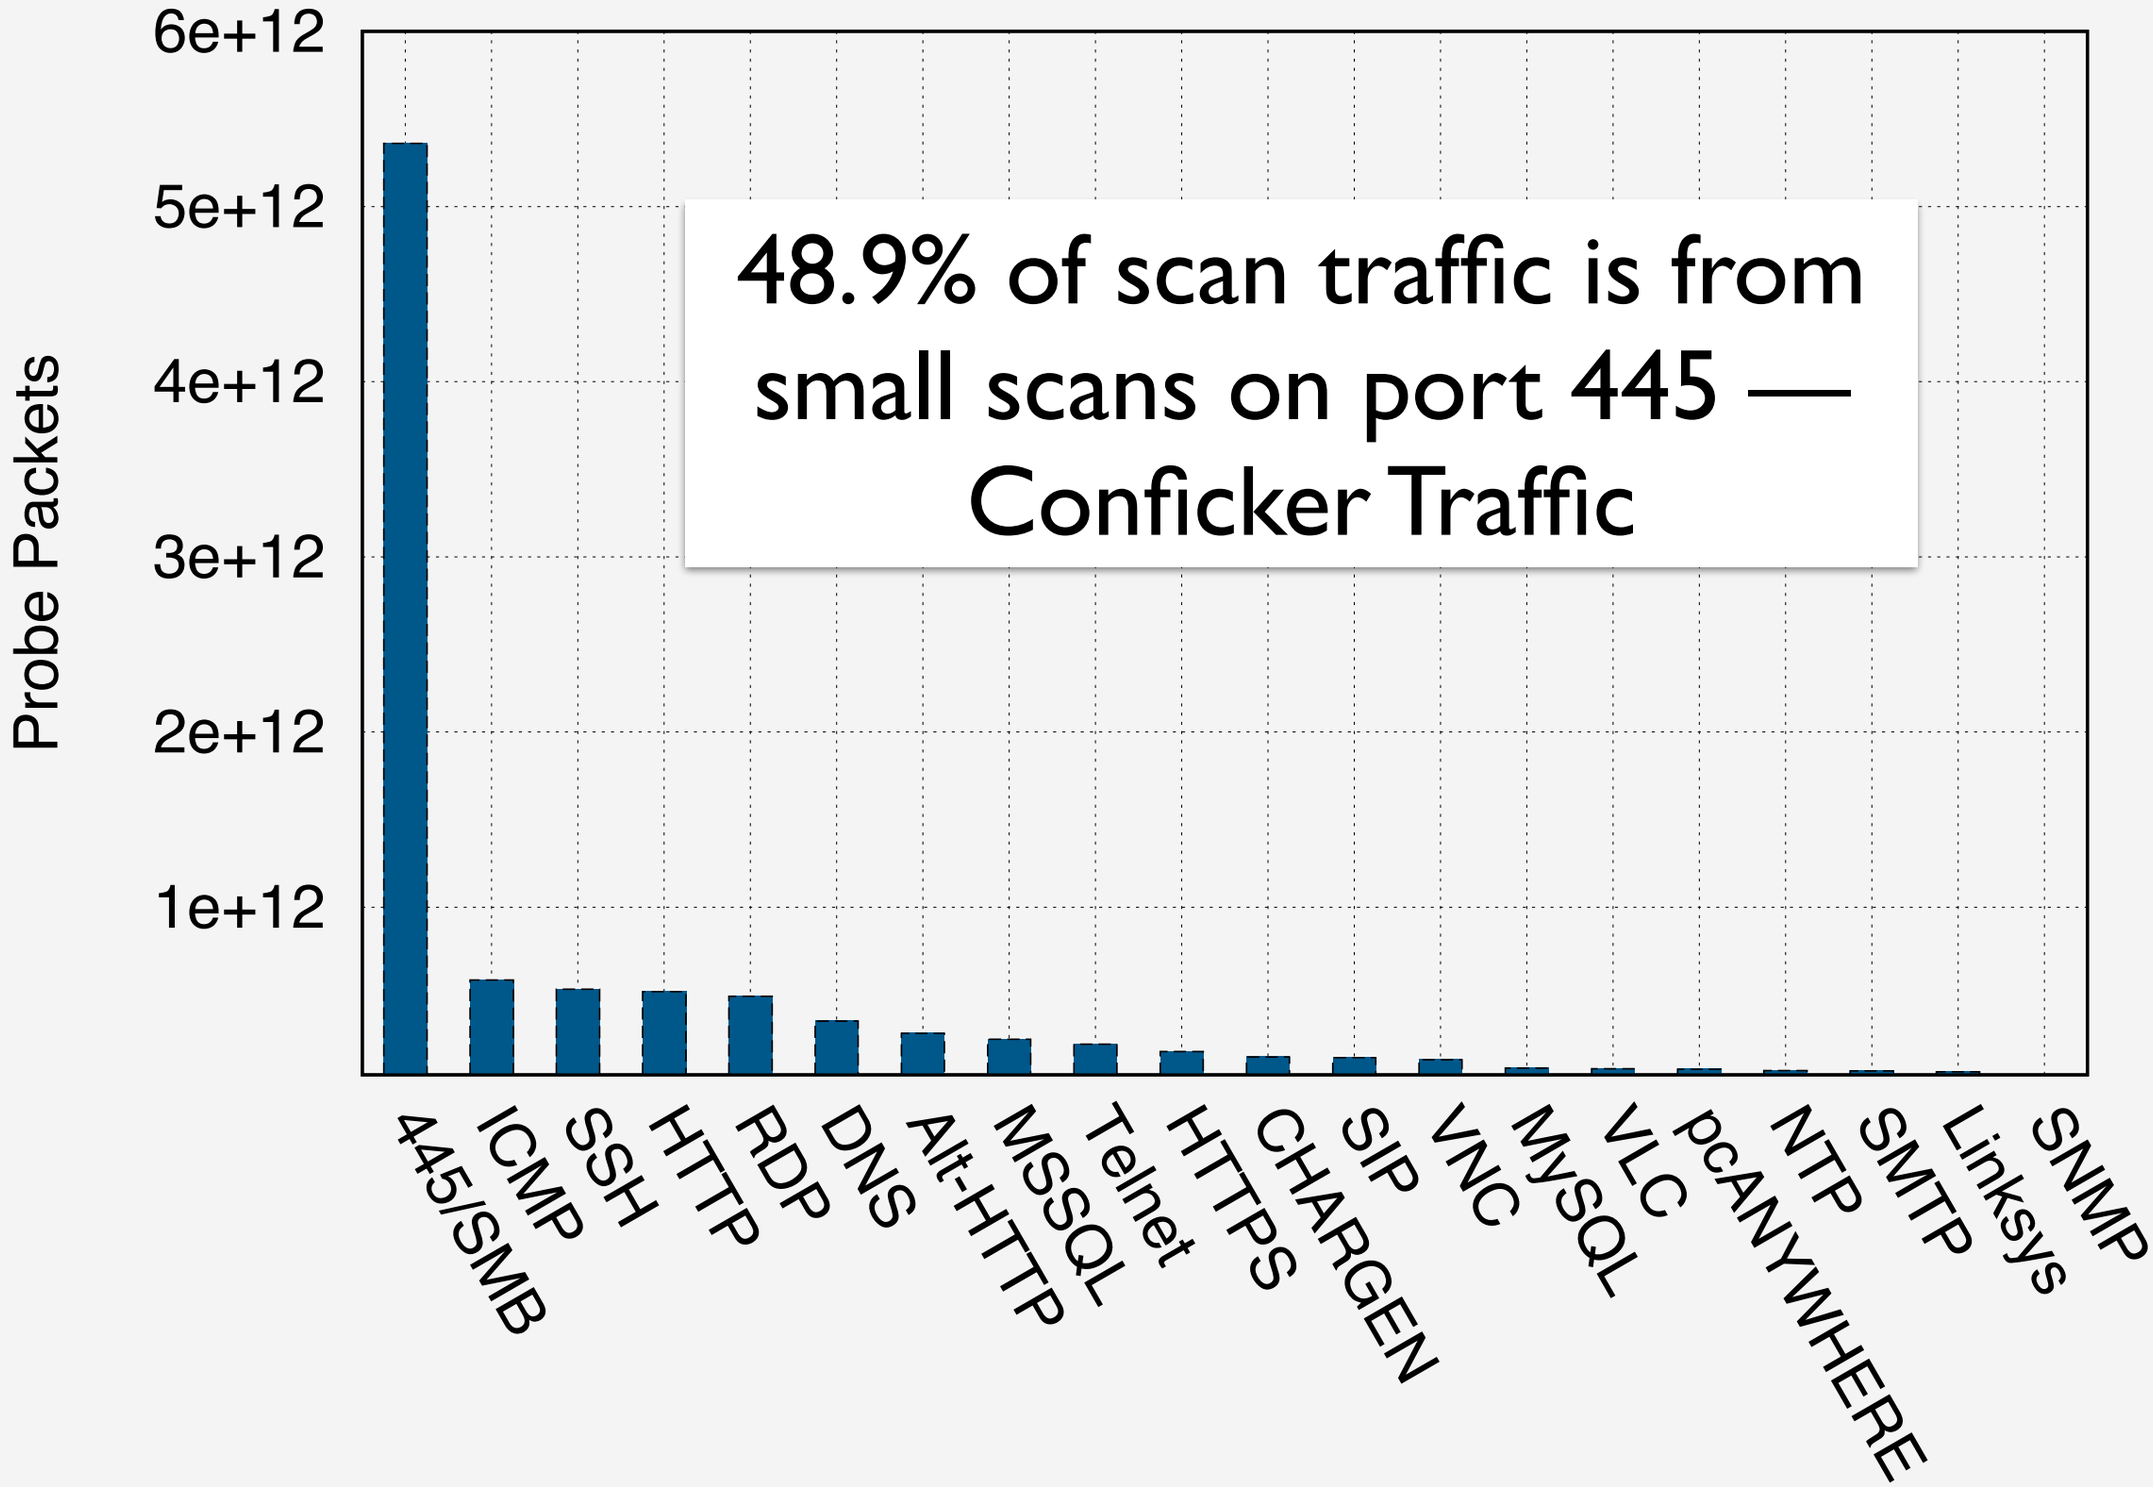

Targeted Services

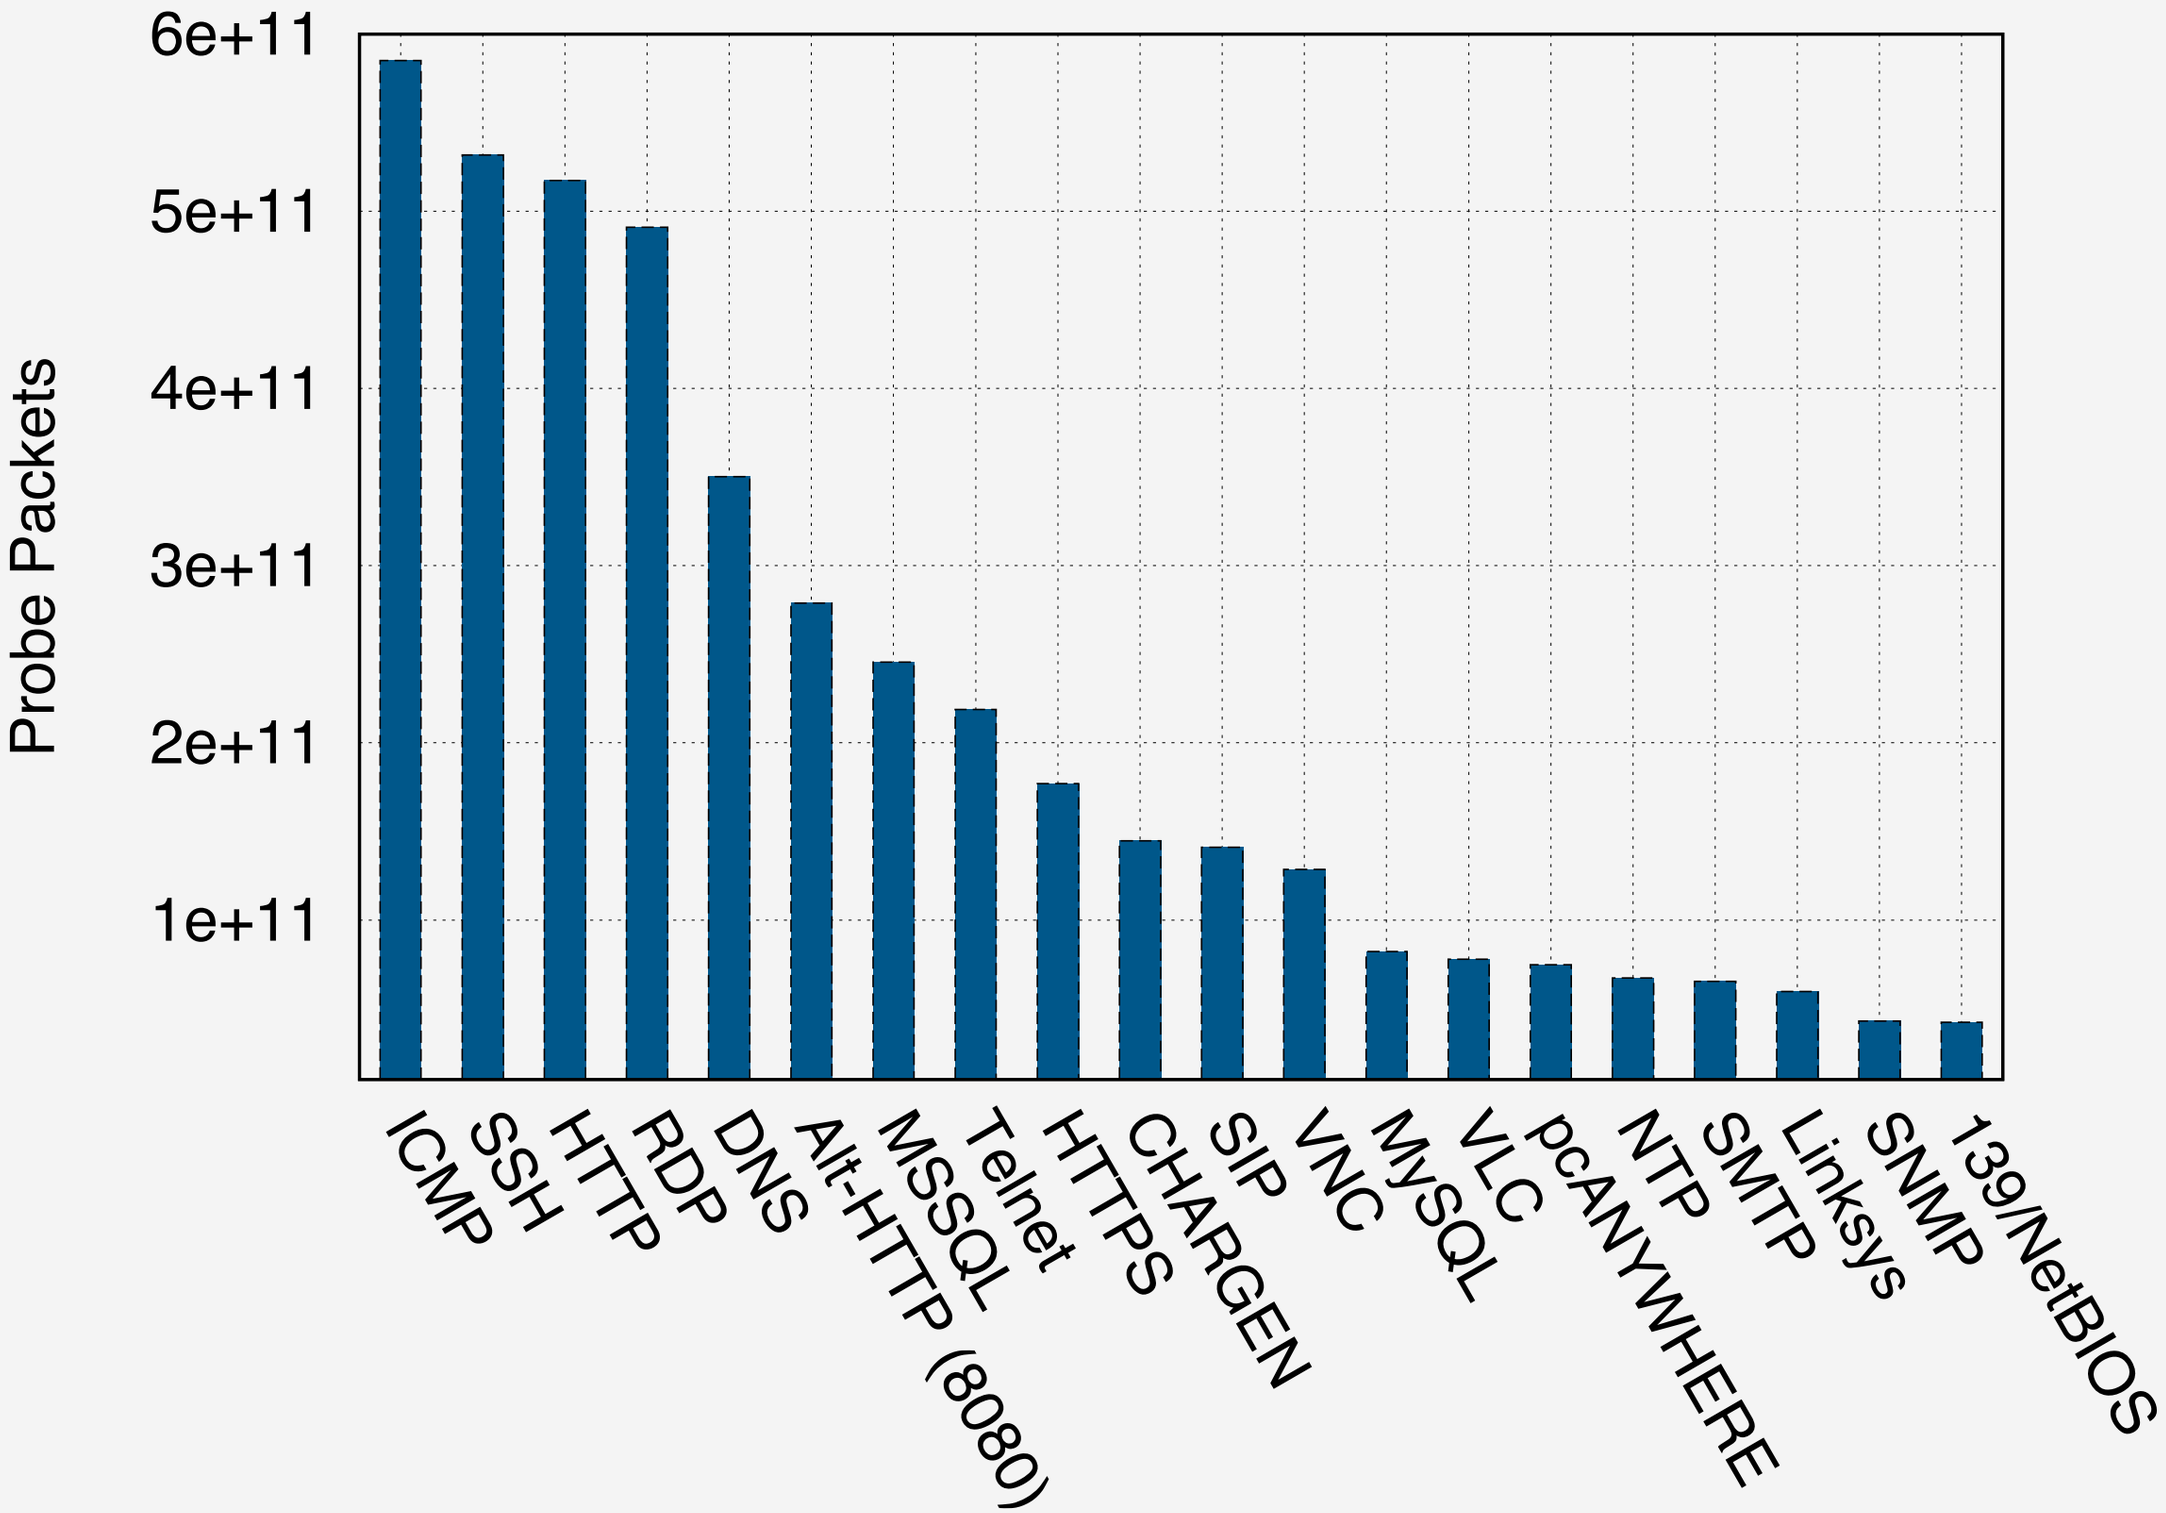

Targeted Services

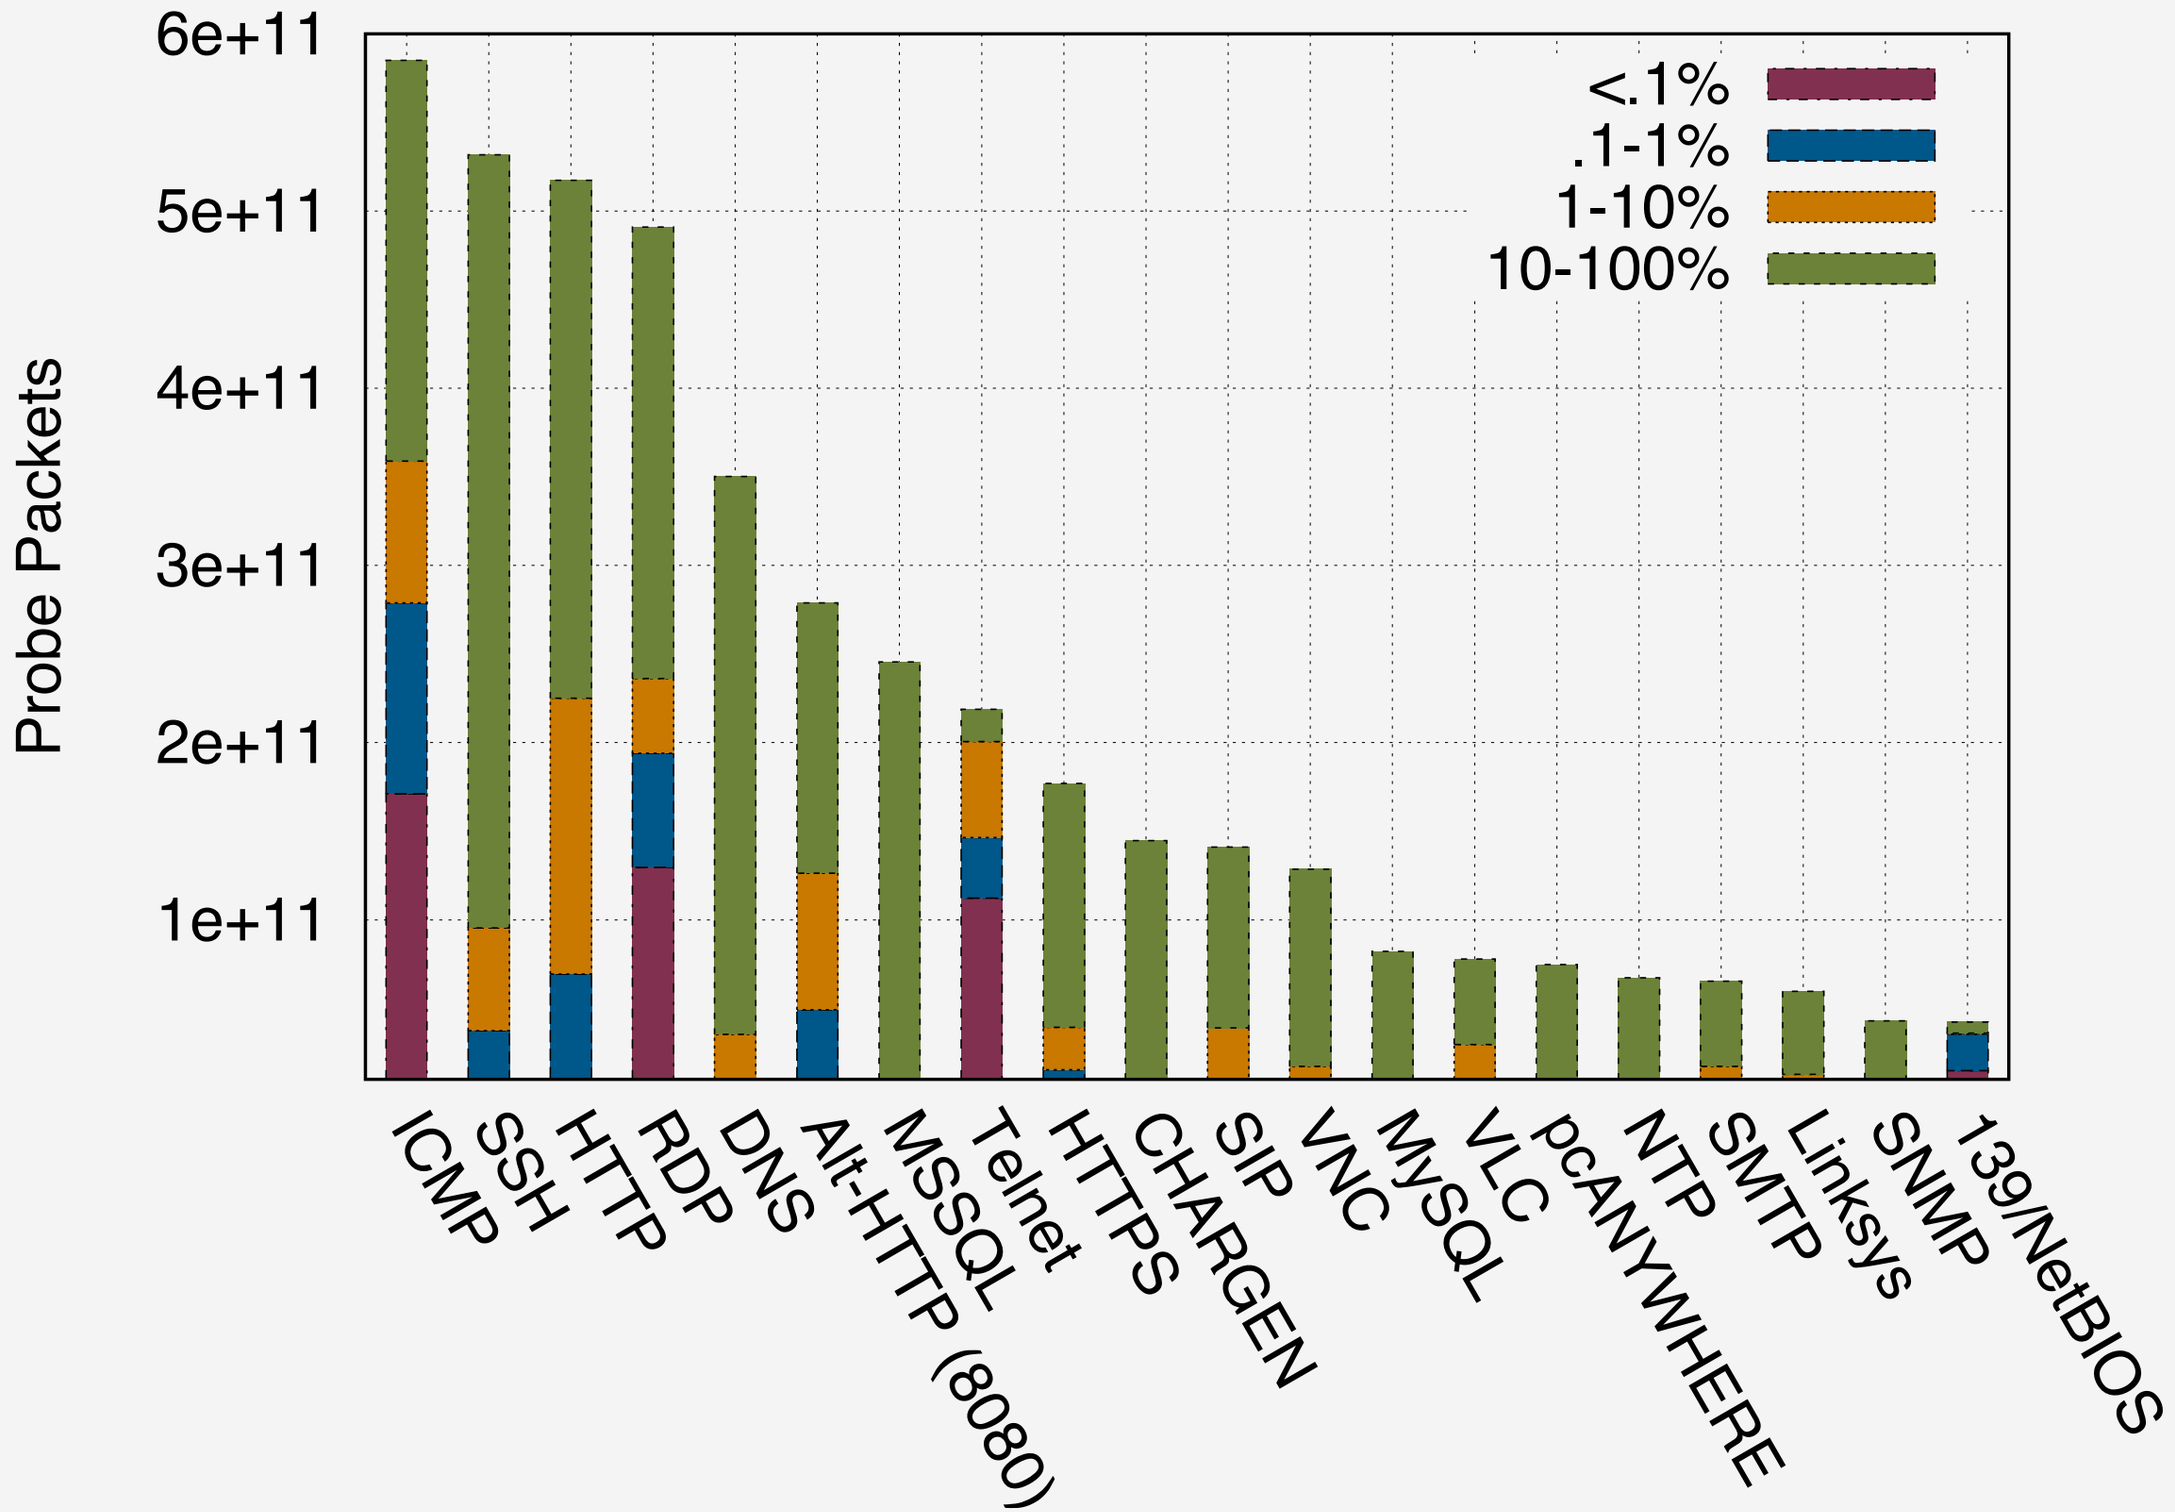

Large Scans

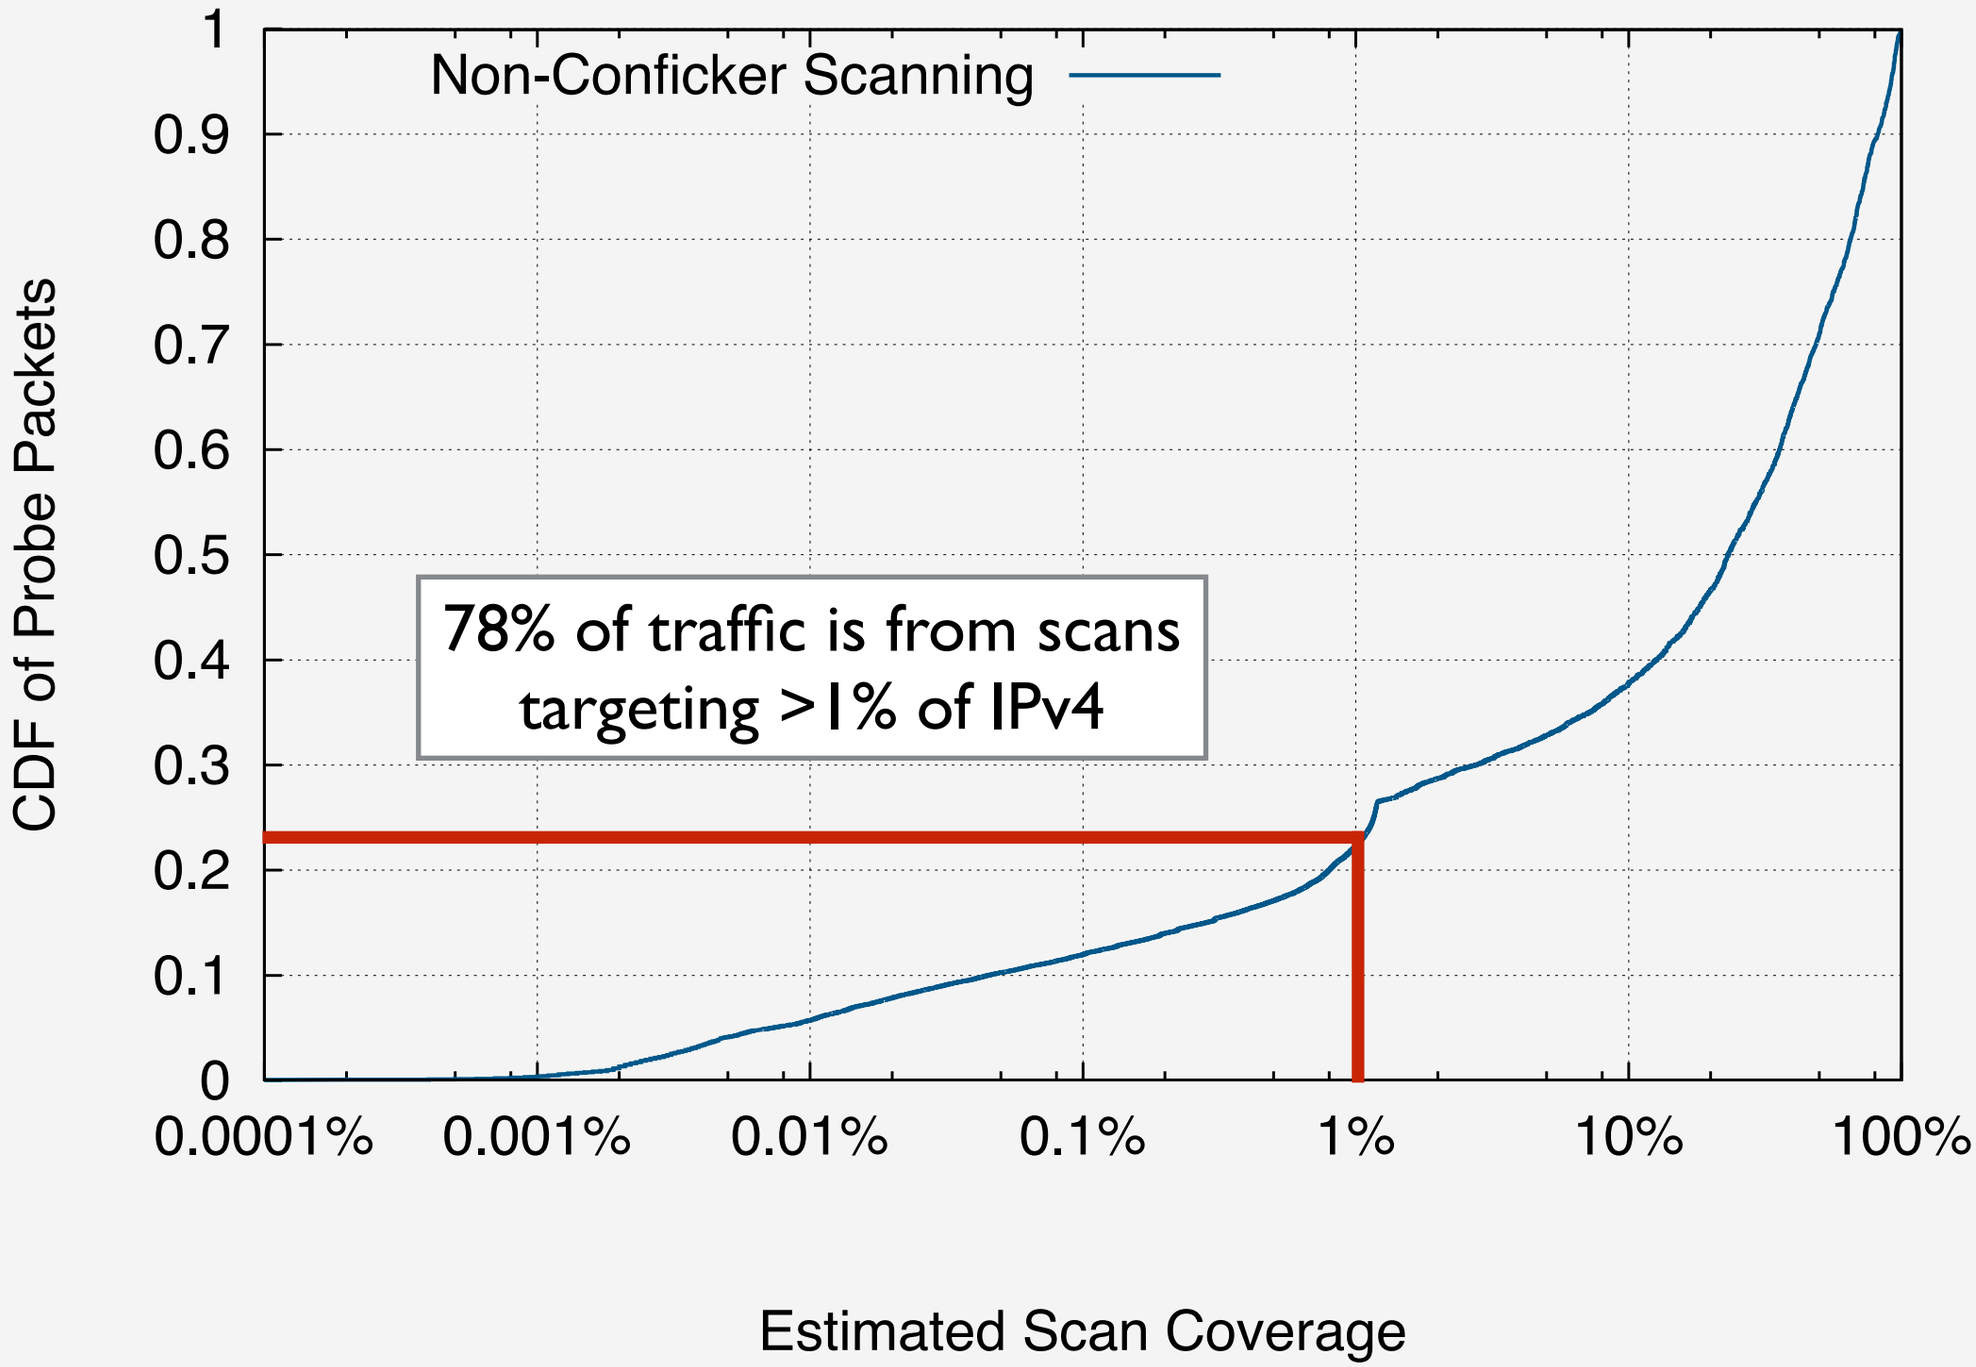
Non-Conficker Scanning
78% of traffic is from scans targeting >1% of IPv4
CDF of Probe Packets
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.0001%
0.001%
0.01%
0.1%
1%
10%
100%
Estimated Scan Coverage

Large Scans

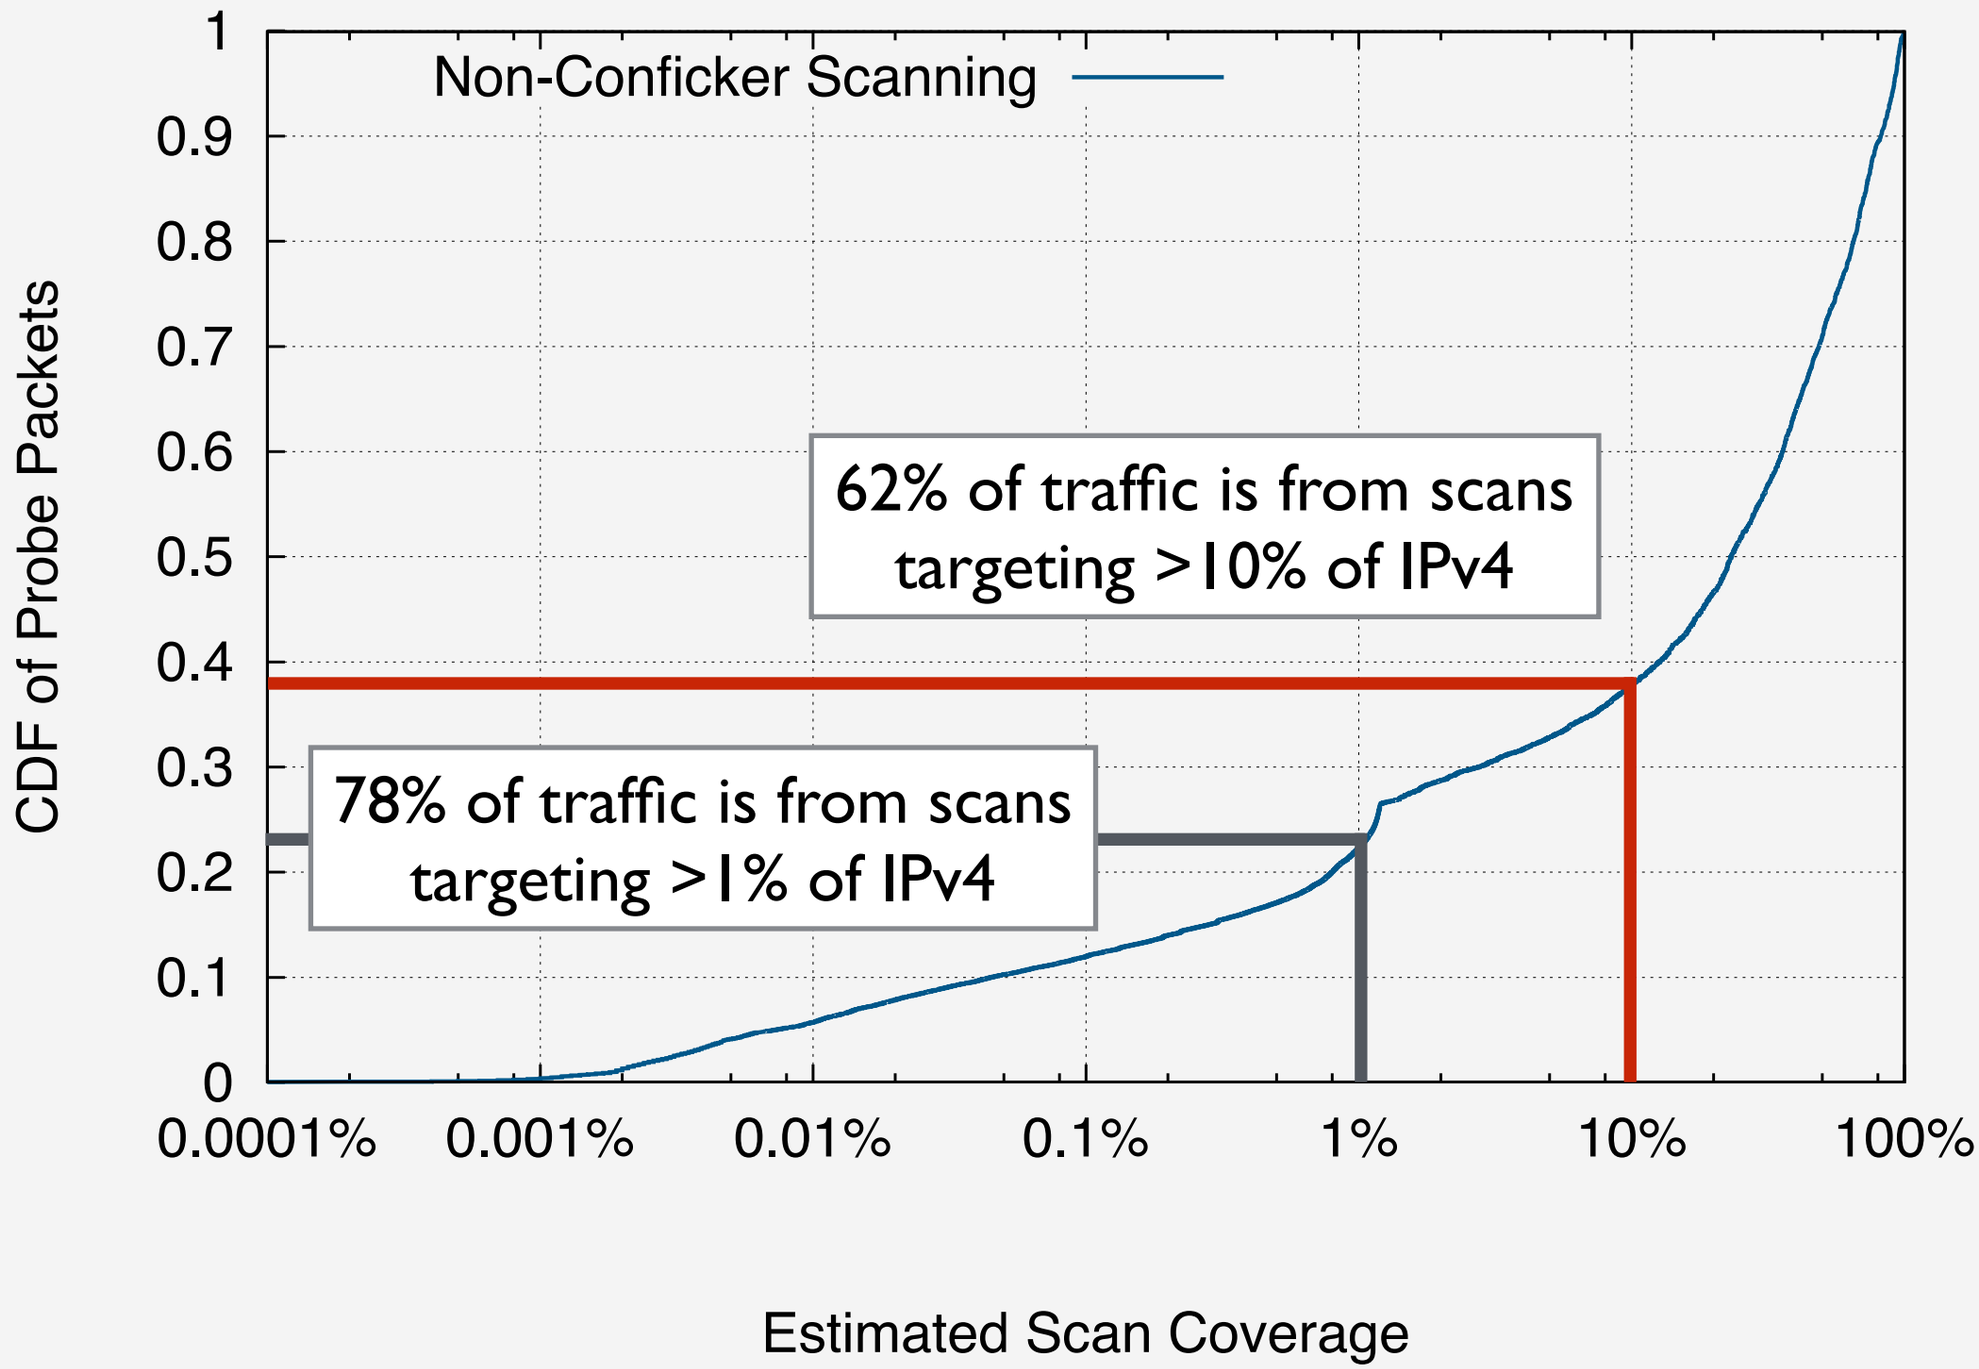
Non-Conficker Scanning
62% of traffic is from scans targeting >10% of IPv4
78% of traffic is from scans targeting >1% of IPv4
CDF of Probe Packets
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.0001%
0.001%
0.01%
0.1%
1%
10%
100%
Estimated Scan Coverage

Scan Dynamics

January 2014

- 18,000 scans (0.28%) targeted ≥1% of the IPv4 address space
- 2,700 scans (0.04%) targeted ≥10% of the IPv4 address space
- 100 ASes responsible for 85% of this scan traffic

Scan Dynamics

January 2014

- 18,000 scans (0.28%) targeted ≥1% of the IPv4 address space
- 2,700 scans (0.04%) targeted ≥10% of the IPv4 address space
- 100 ASes responsible for 85% of this scan traffic

Four types of scanning stand out:

- Academic and industry research groups
- Regularly scheduled scans from Chinese ASes
- Unidentifiable scans from bullet-proof hosting providers
- ShodanHQ Search Engine

Research Groups and Security Consultants

Many of the networks responsible for the most scan traffic are academic institutions and consultants performing regular scans

Primarily focused on amplification attacks (NTP, DNS) and cryptographic ecosystems (SSH, HTTPS)

In almost all cases, studies appear to be conducted responsibly and allowed easy exclusion

Regular Chinese Scans

Regular daily scans of ICMP, SSH, SQL Server, and TCP/0

TCP/0 — non-standard-compliant port frequently used to fingerprint network stacks and bypass firewalls

Responsible for the majority of ICMP, SQL Server, MySQL, and ICMP traffic — far more than other countries

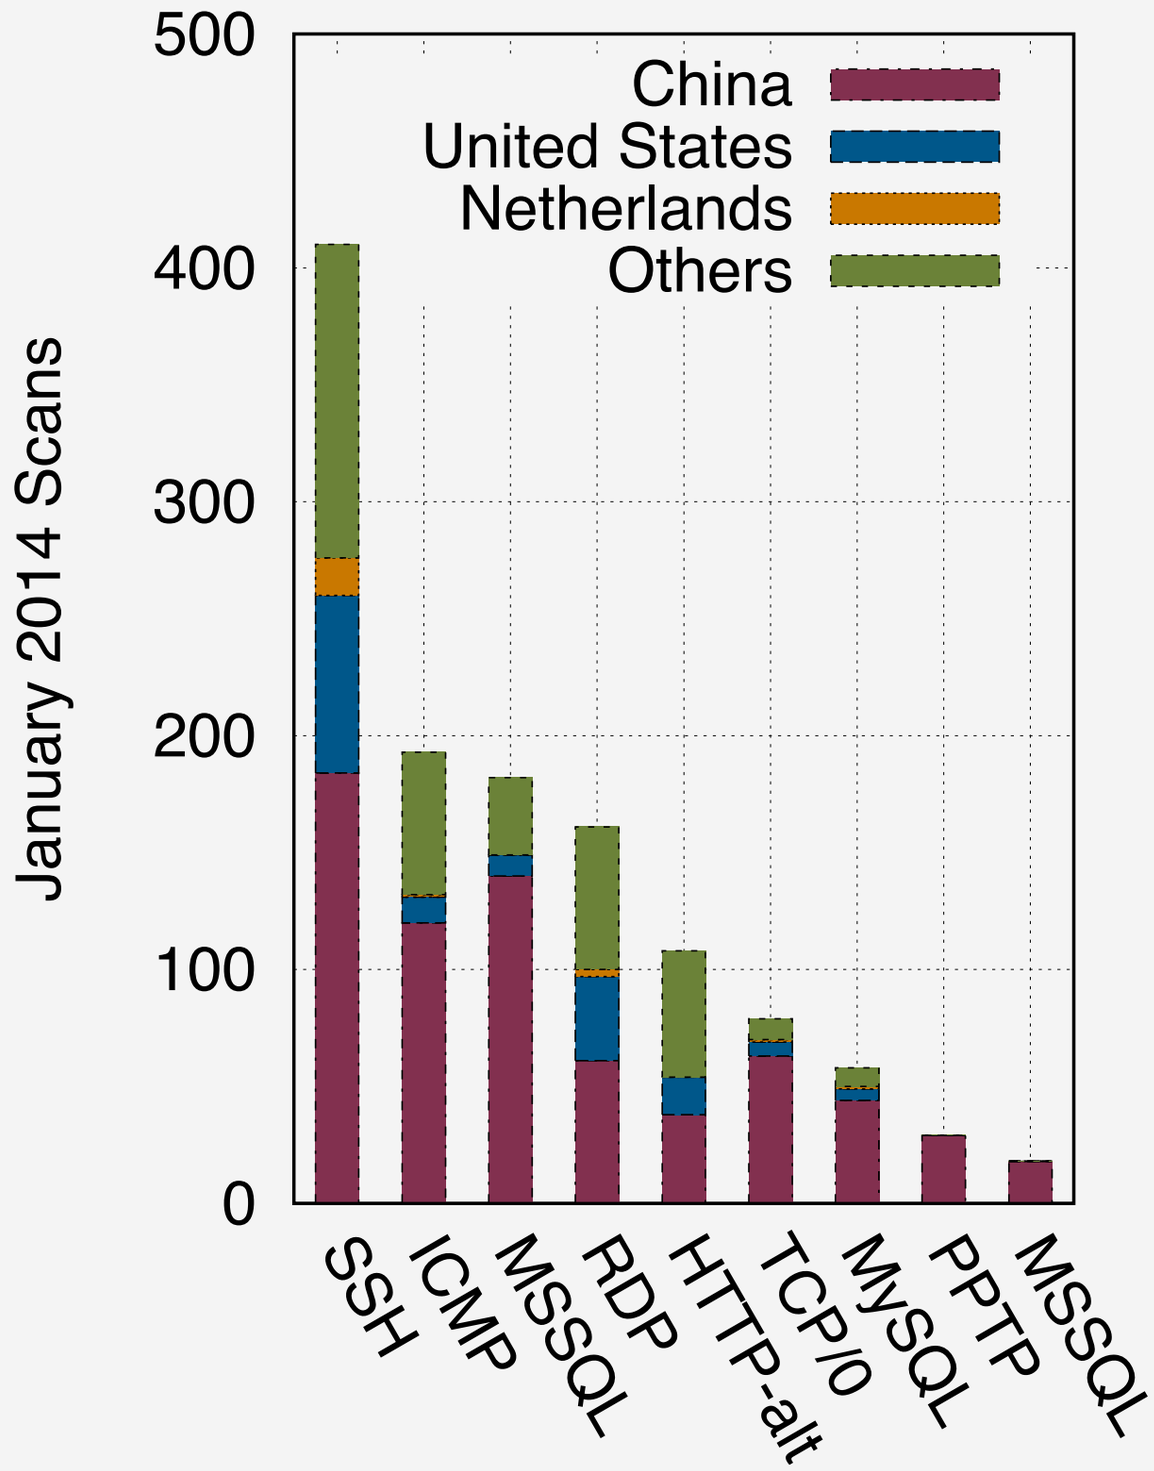

Large Hosting Providers

50% of the top 100 ASes responsible for scan traffic were large hosting providers

Many were bullet-proof hosting providers

Bullet-Proof Hosting Providers

- Advertise turning a blind-eye to malicious behavior
- Scanning for almost every common protocol
- Very rarely any identifiable information about owners

Top Scanning Providers
Ecatel Network (NL)
Plus Server (DE)
Slask Data Center (PL)
Single Hop (US)
CariNet, Inc. (US)
Server4You (DE)
OVH Systems (UK)
Thor Data Center (IS)
Psychz Networks (US)

Talk Outline

1. Broad Overview of Scanning Landscape

2. **Case Studies: Scanning triggered by backdoors in home routers, Heartbleed, and NTP vulnerabilities**

3. Defensive reactions against scanning

Linksys Router Backdoor

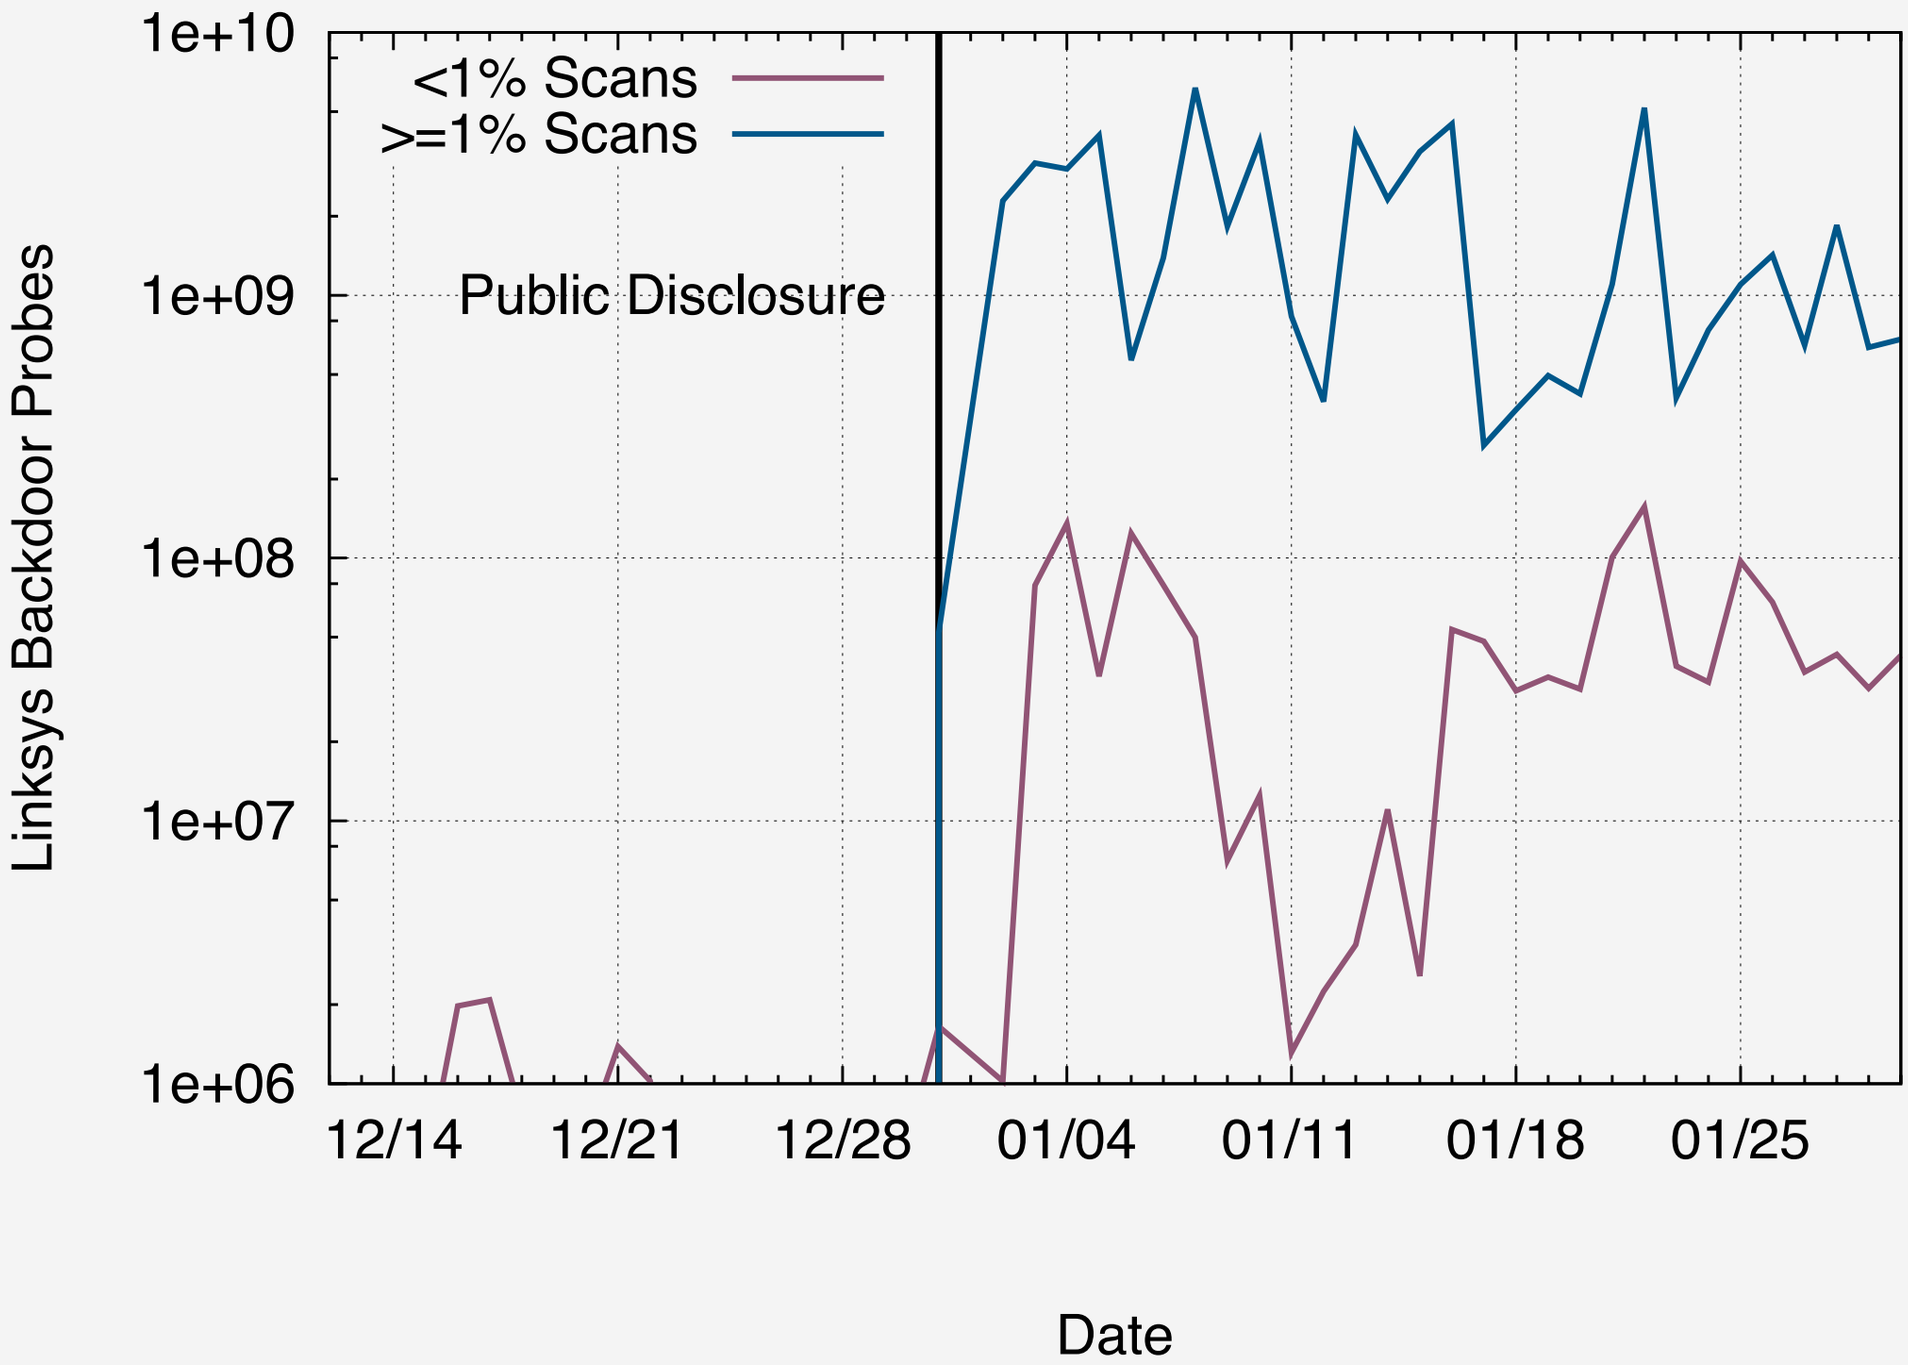

Linksys Router Backdoor

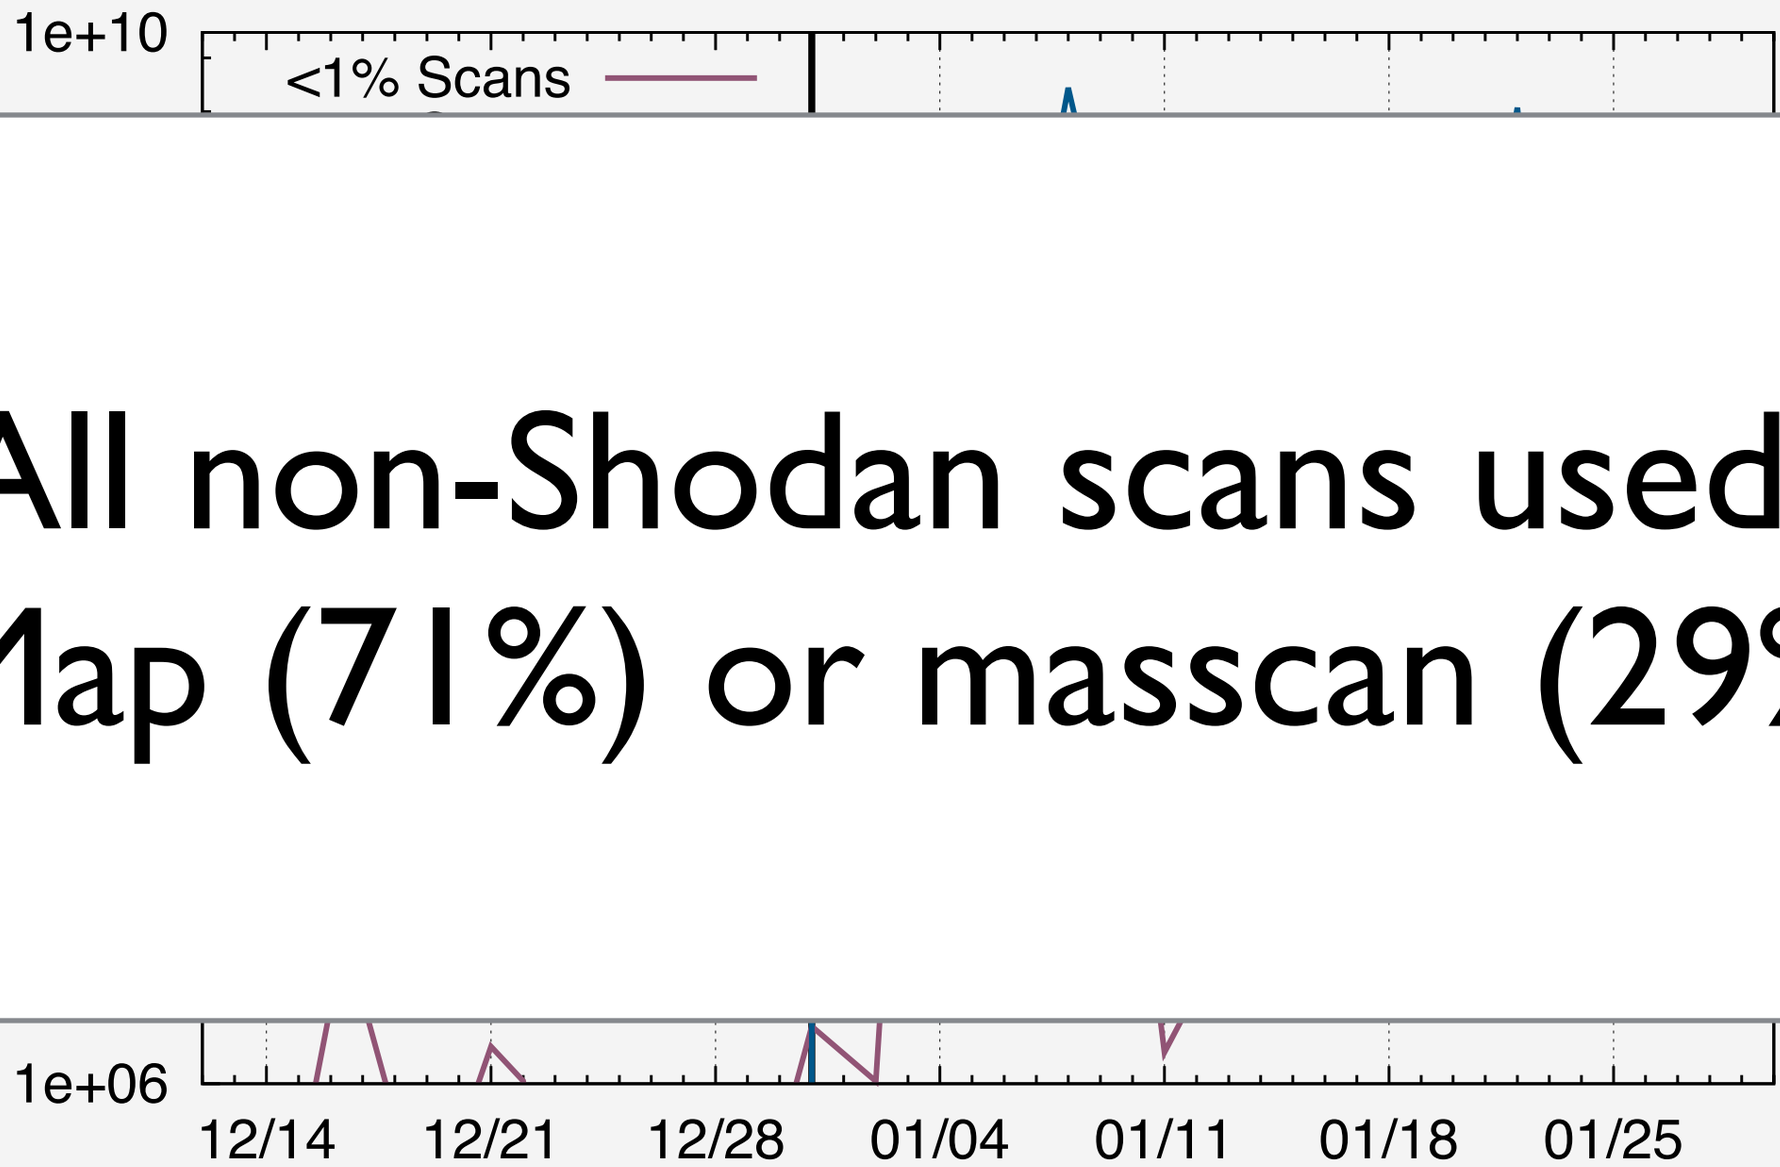

All non-Shodan scans used ZMap (71%) or masscan (29%)

Open NTP Resolvers

97.3% of probe traffic is part of large scans (targeting >1% of IPv4)

Primarily scanned from bullet-proof hosting providers.

50% of scans used ZMap or Masscan

Not certain that scanners are malicious, but absolutely appear so

“#yolo”

“#lulz”

“Openbomb Drone Project”

http://ra.pe

Heartbleed Vulnerability

Scans began <24 hours after disclosure

53 scans from 27 hosts in the week following disclosure

38% of scans originated from China

Scans occurring from bulletproof hosting providers

95% of scans used ZMap or Masscan

Heartbleed Vulnerability

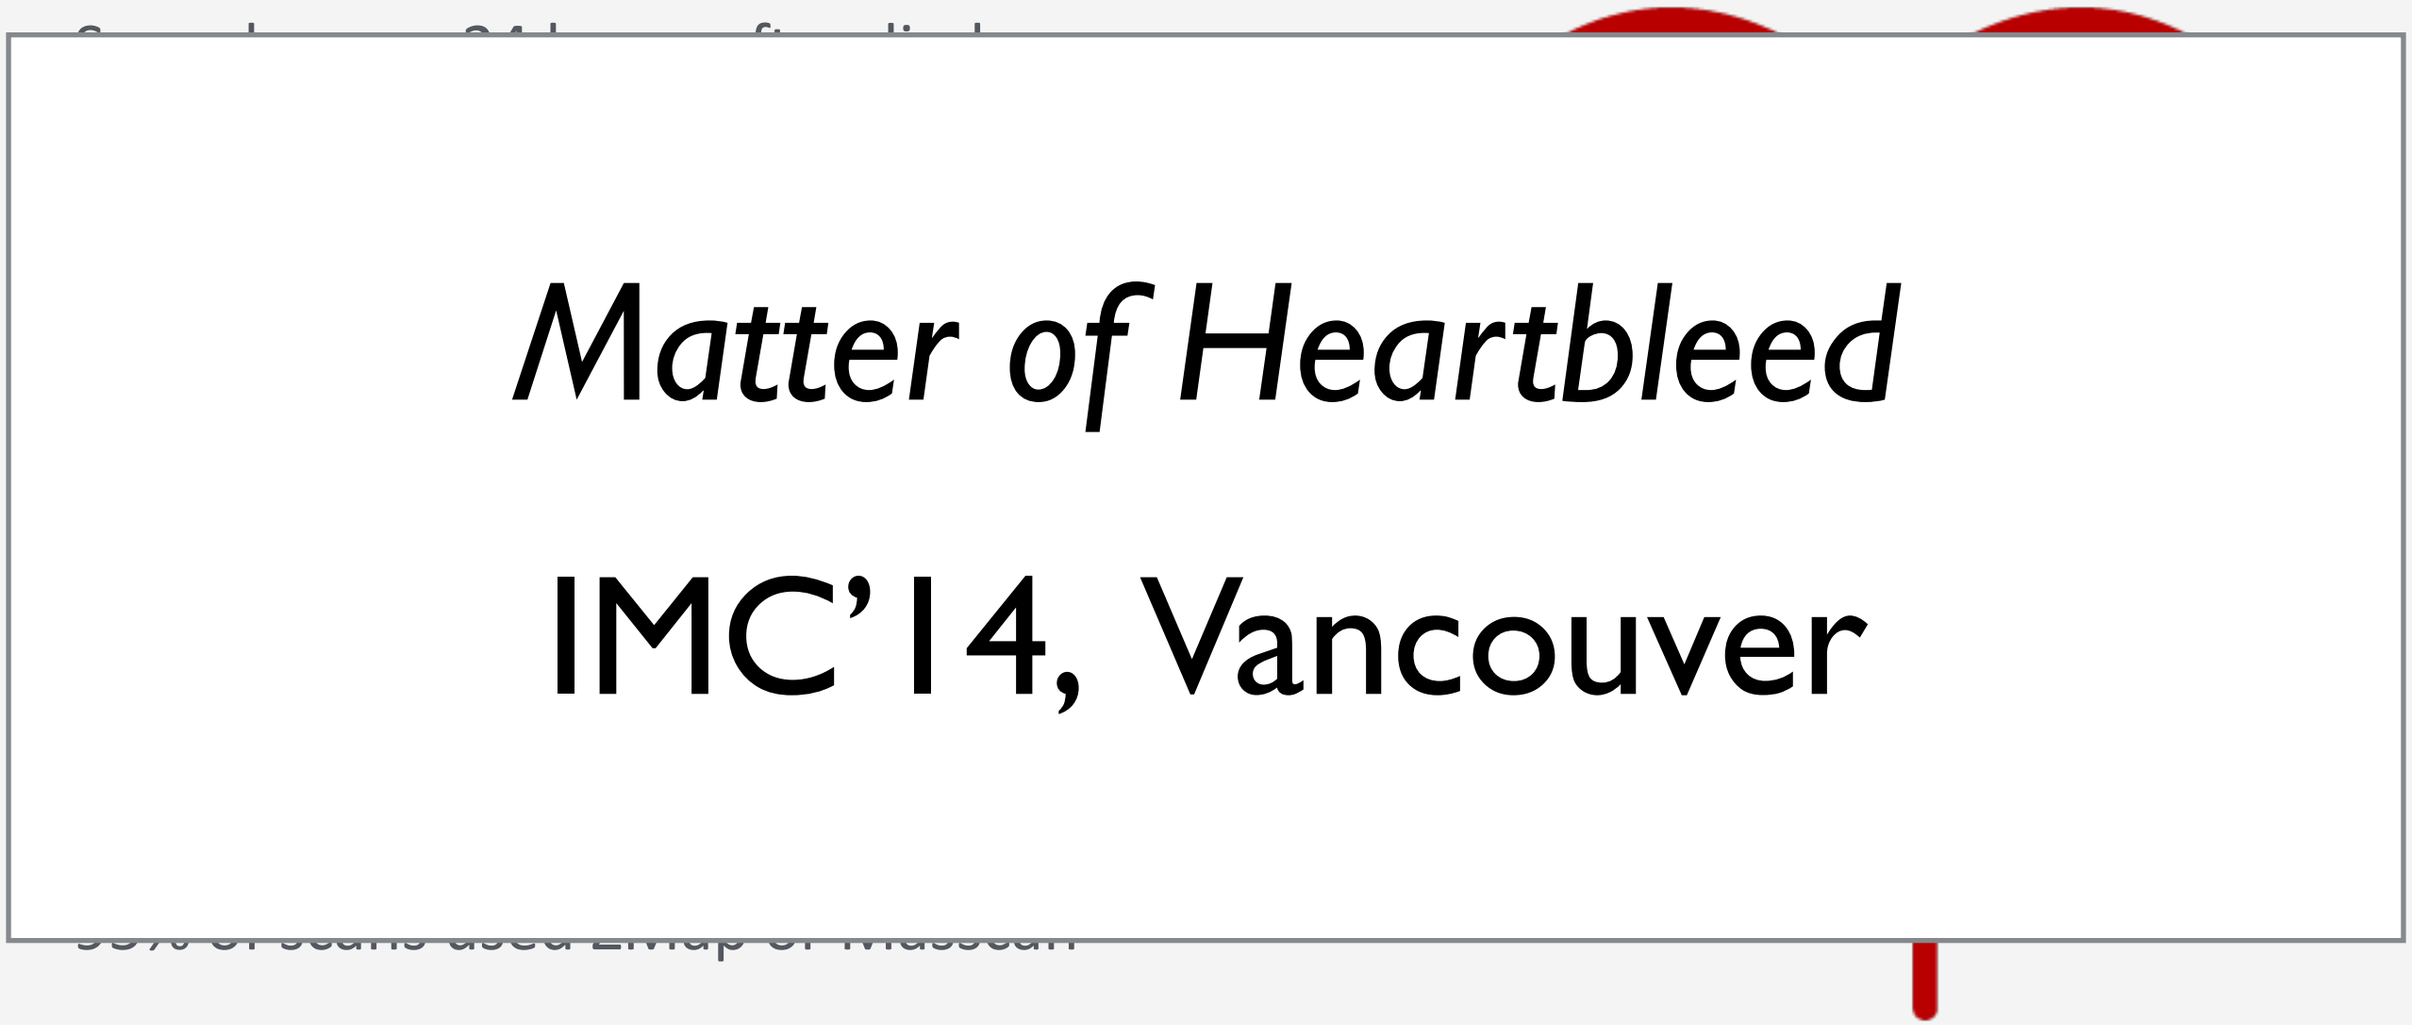

So what about ZMap?

The majority of scan traffic is not generated by ZMap

Research groups are using ZMap responsibly

Evidence that attackers are starting to take advantage of ZMap and Masscan

Ultimately lowers the barrier of entry for both groups

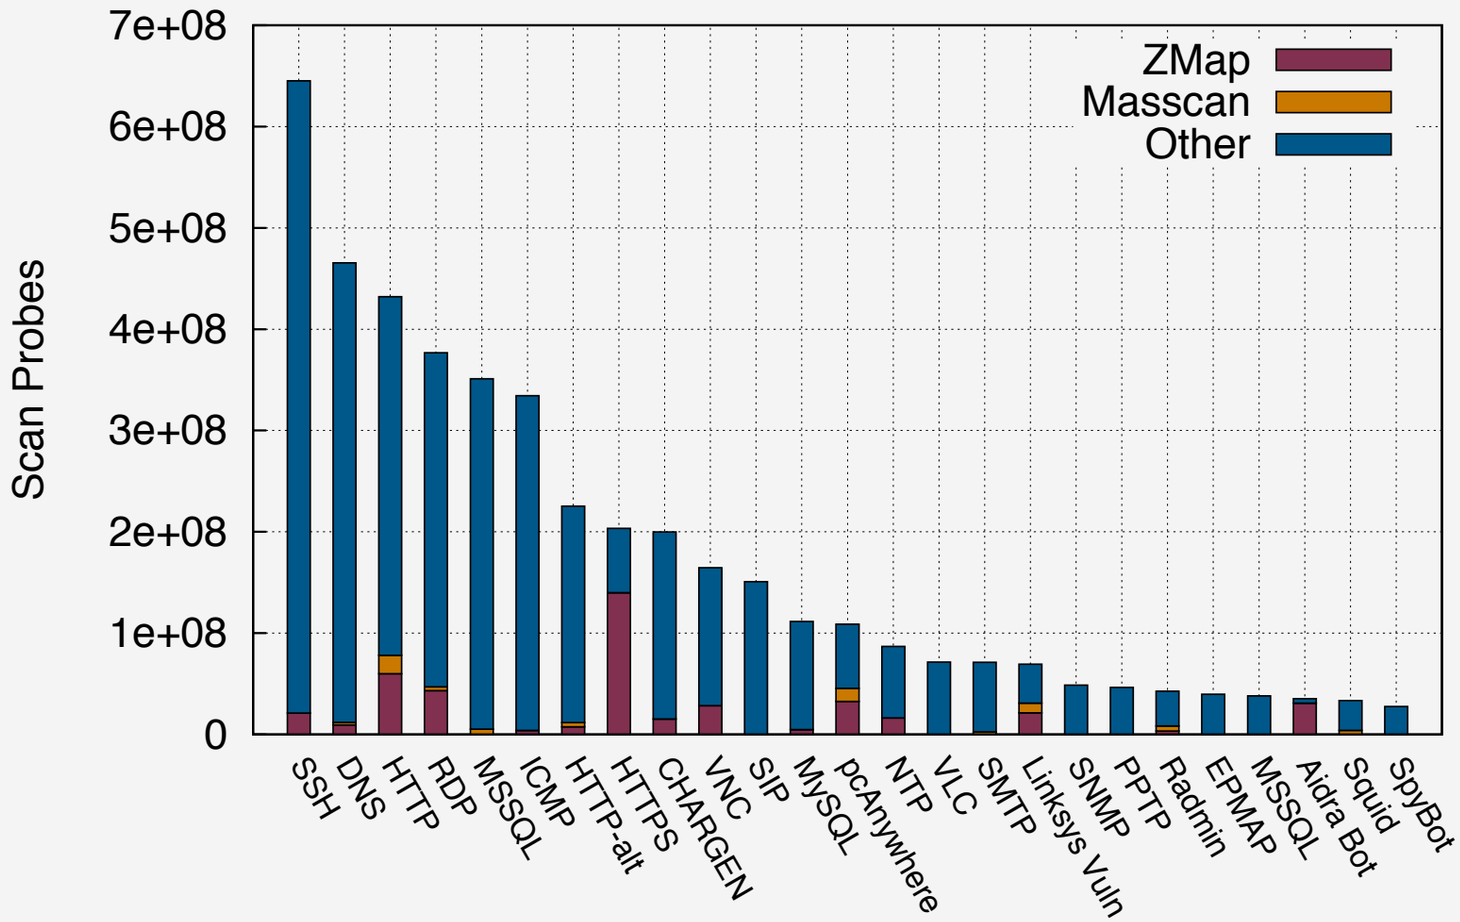

Talk Outline

1. Broad Overview of Scanning Landscape

2. Case Studies: Scanning triggered by backdoors in home routers, Heartbleed, and NTP vulnerabilities

3. **Defensive reactions against scanning**

Do networks drop scan traffic?

Michigan Engineering AS is responsible for 3rd most scan traffic

Performed simultaneous scans from Georgia Tech and Michigan to detect blocked traffic

Scanned using same randomization seed —reduce hosts lost due to churn

Do networks drop scan traffic?

Estimated 0.05% of IPv4 address space is no longer accessible

208 exclusion requests — 0.15% of IPv4 address space

Dropped traffic and excluded networks have a minuscule impact

When do networks drop scan traffic?

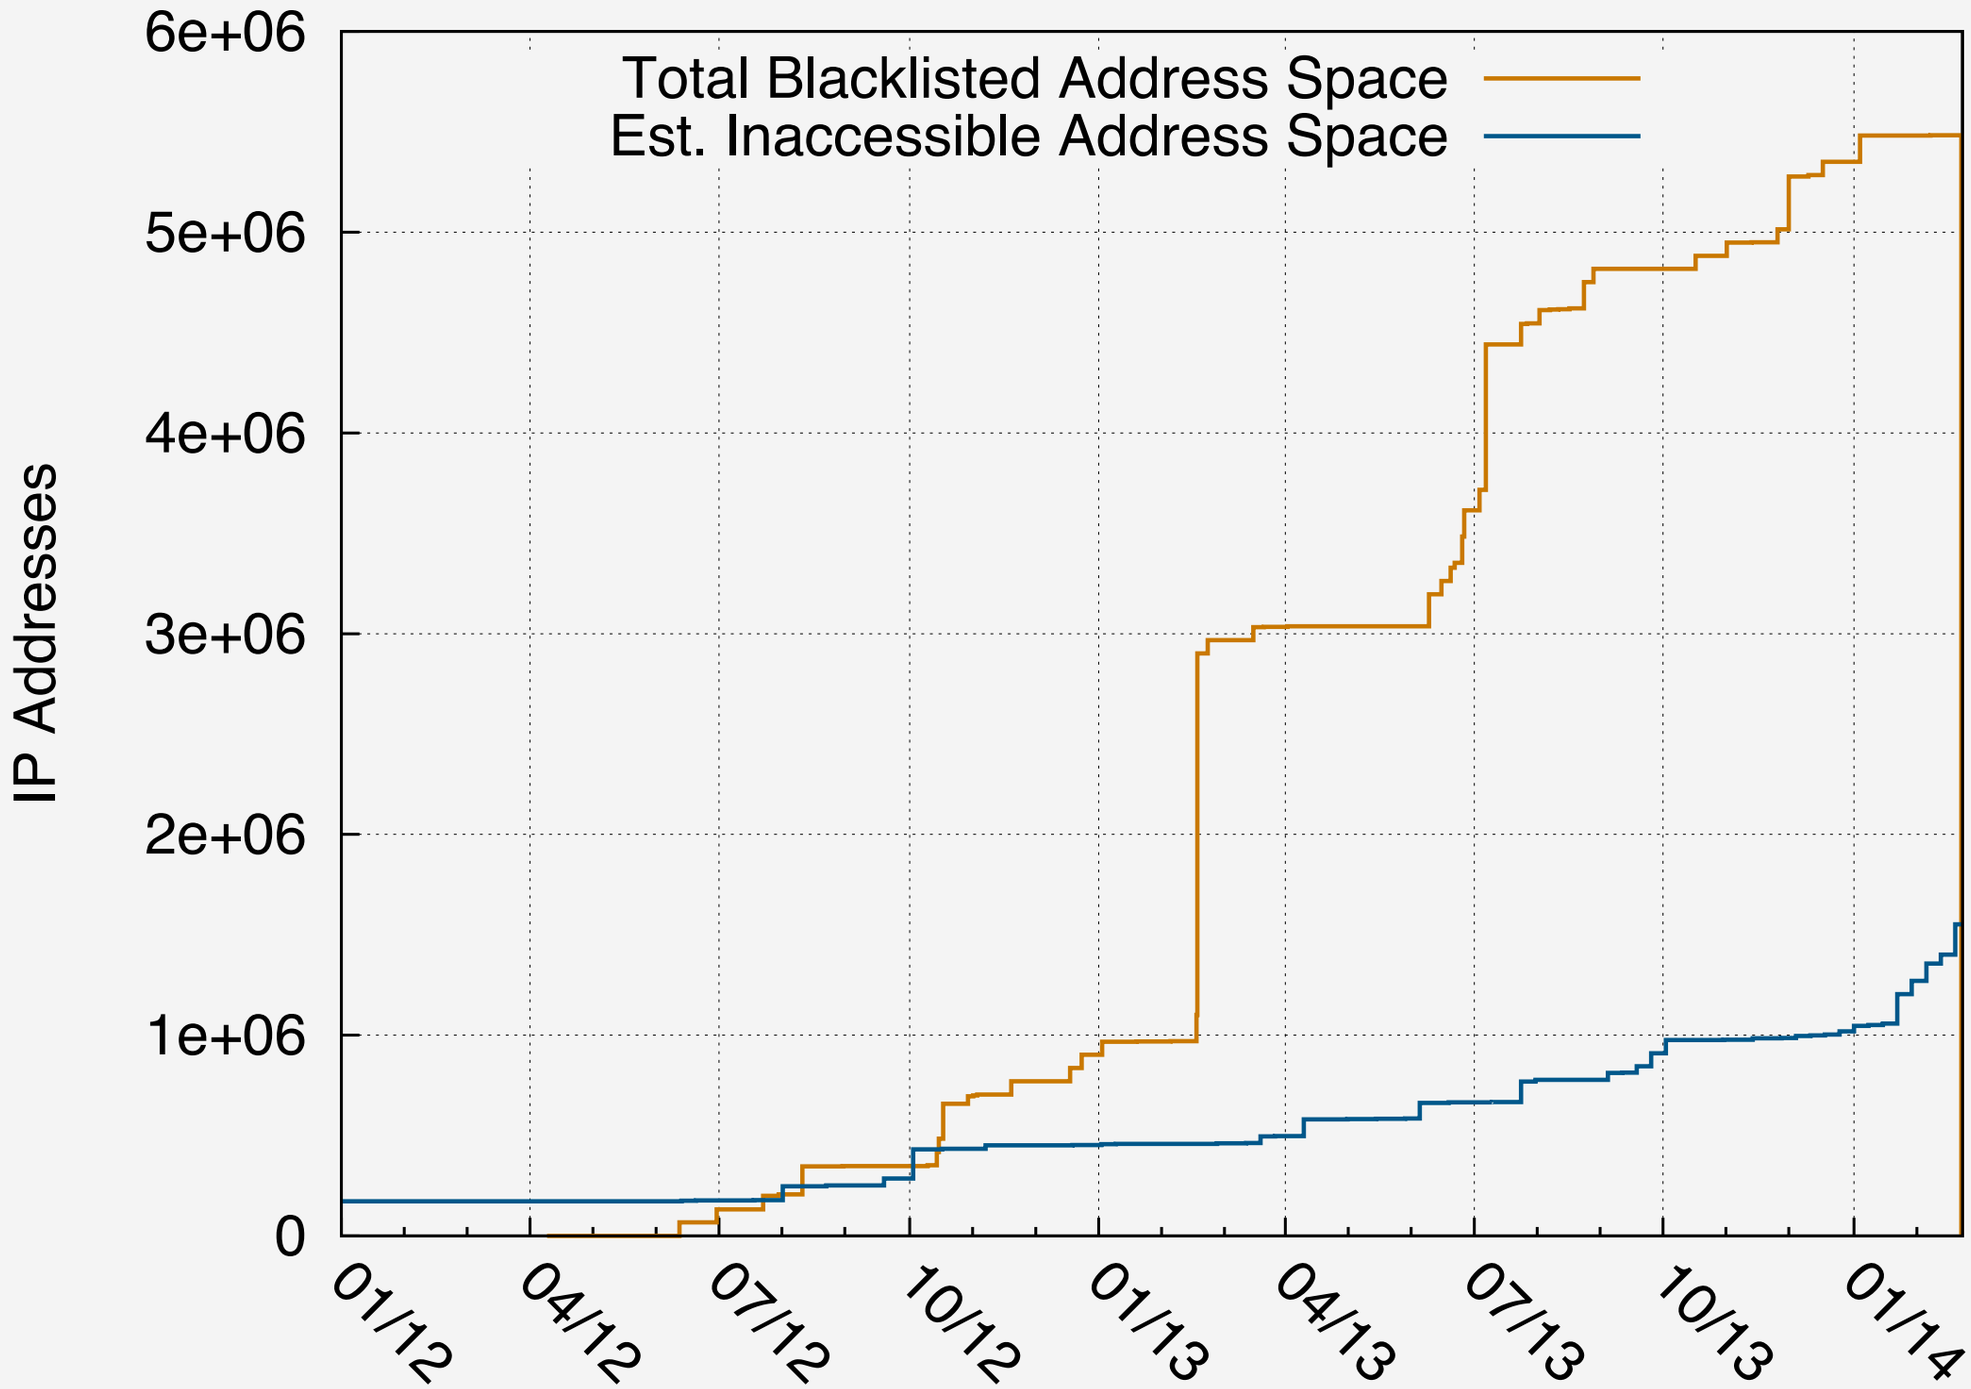

How are organizations noticing?

Detection Mechanism	Organizations
Firewall Logs	22 (34%)
Web Server Logs	14 (22%)
IDS Logs	10 (16%)
Invalid SSH or OpenVPN Handhshake	10 (16%)
Public Blacklists	2 (3%)
Other	6 (9%)

Future Work

Exclusion standard

Understand defensive reactions

Correlating distributed scanners

Determining scan intent

Conclusion

Scanning landscape has shifted — large horizontal scans are now common

Internet-Wide scanning is a combination of both researchers and attackers taking advantage of new tools

Network operators have been slow to respond to scanning despite scanning being easy to detect

Internet-Wide scanning remains a valid methodology

Questions?

An Internet-Wide View of Internet-Wide Scanning

Zakir Durumeric, Michael Bailey, J. Alex Halderman

University of Michigan

scanning-team@umich.edu