LZR: Identifying Unexpected Internet Services

Liz Izhikevich
Stanford University

Renata Teixeira
Inria Paris*

Zakir Durumeric
Stanford University

* Work done while visiting Stanford University

More than 300 security studies have used Internet-wide scanning

A Messy State of the Union:
Taming the Composite State Machines of TLS

Benjamin Beurdouche*, Karthikeyan Bhargavan*, Antoine Delignat-Lavaud*,
Cédric Fournet†, Markulf Kohlweiss†, Alfredo Pironti*,
Pierre-Yves Strub‡, Jean Karim Zinzindohoue§*

Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents

Yang Liu[1], *Armin Sarabi*[1], *Jing Zhang*[1], *Parinaz Naghizadeh*[1]
Manish Karir[2], *Michael Bailey*[3], *Mingyan Liu*[1,2]
[1] *EECS Department, University of Michigan, Ann Arbor*
[2] *QuadMetrics, Inc.*
[3] *ECE Department, University of Illinois, Urbana-Champaign*

Exit from Hell? Reducing the Impact of Amplification DDoS Attacks

Marc Kührer, Thomas Hupperich, Christian Rossow, Thorsten Holz
Horst Görtz Institute for IT-Security, Ruhr-University Bochum

A Large-Scale Analysis of the Security of Embedded Firmwares

Andrei Costin, Jonas Zaddach, Aurélien Francillon and Davide Balzarotti

EURECOM
Sophia Antipolis
France
{*name.surname*}*@eurecom.fr*

Global Measurement of DNS Manipulation

Paul Pearce[◇] *Ben Jones*[†] *Frank Li*[◇] *Roya Ensafi*[†]
Nick Feamster[†] *Nick Weaver*[‡] *Vern Paxson*[◇]

[◇]*University of California, Berkeley* [†]*Princeton University*
[‡] *International Computer Science Institute*

Past studies generally only scan IANA-assigned ports

Internet Assigned Numbers Authority

Port 22

Port 443

- Where are Internet services deployed in practice?

- Where are Internet services deployed in practice?

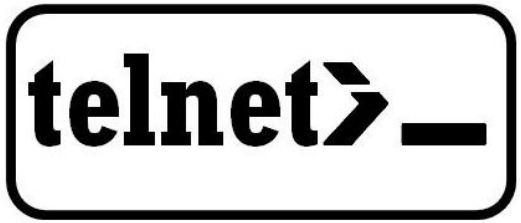

Port ??

- Where are Internet services deployed in practice?

- What is the security posture of services on unexpected ports?

- Where are Internet services deployed in practice?

- What is the security posture of services on unexpected ports?

- Where are Internet services deployed in practice?
- What is the security posture of services on unexpected ports?
- How do we efficiently identify a service's protocol?

TCP Scanning Methodology

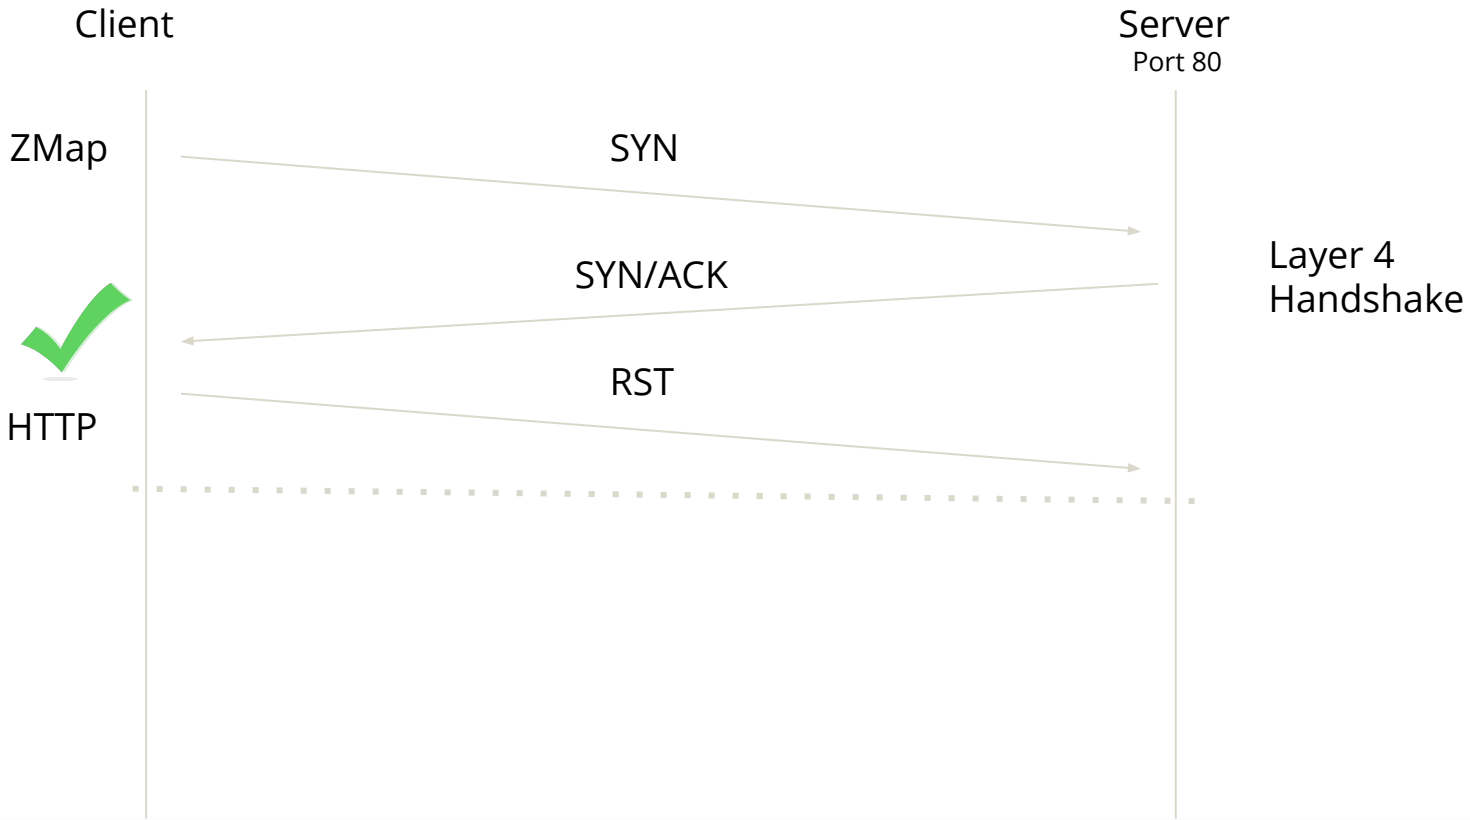

TCP Scanning Methodology

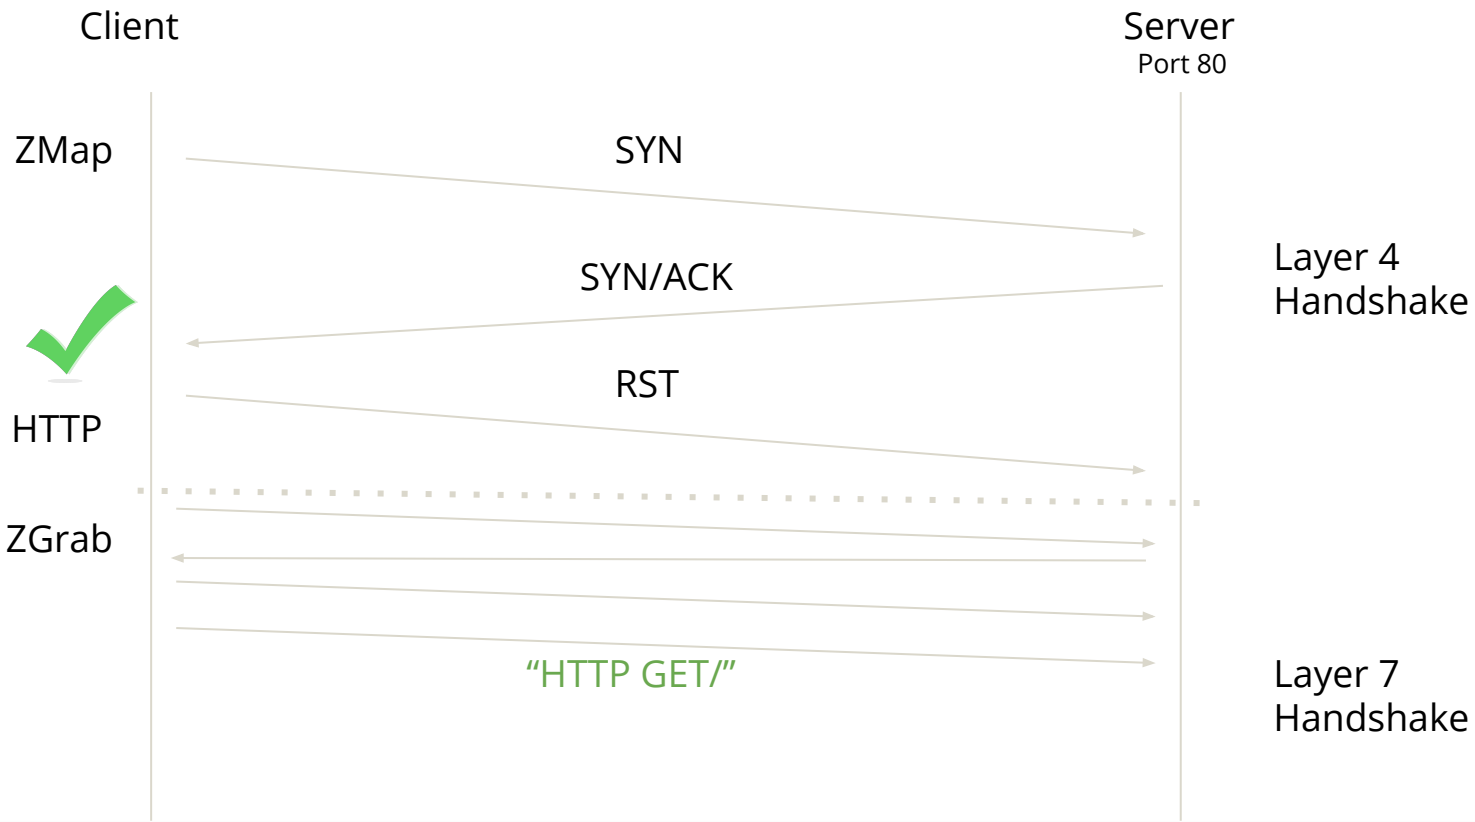

14% of hosts do NOT complete the expected L7 handshake on port 80

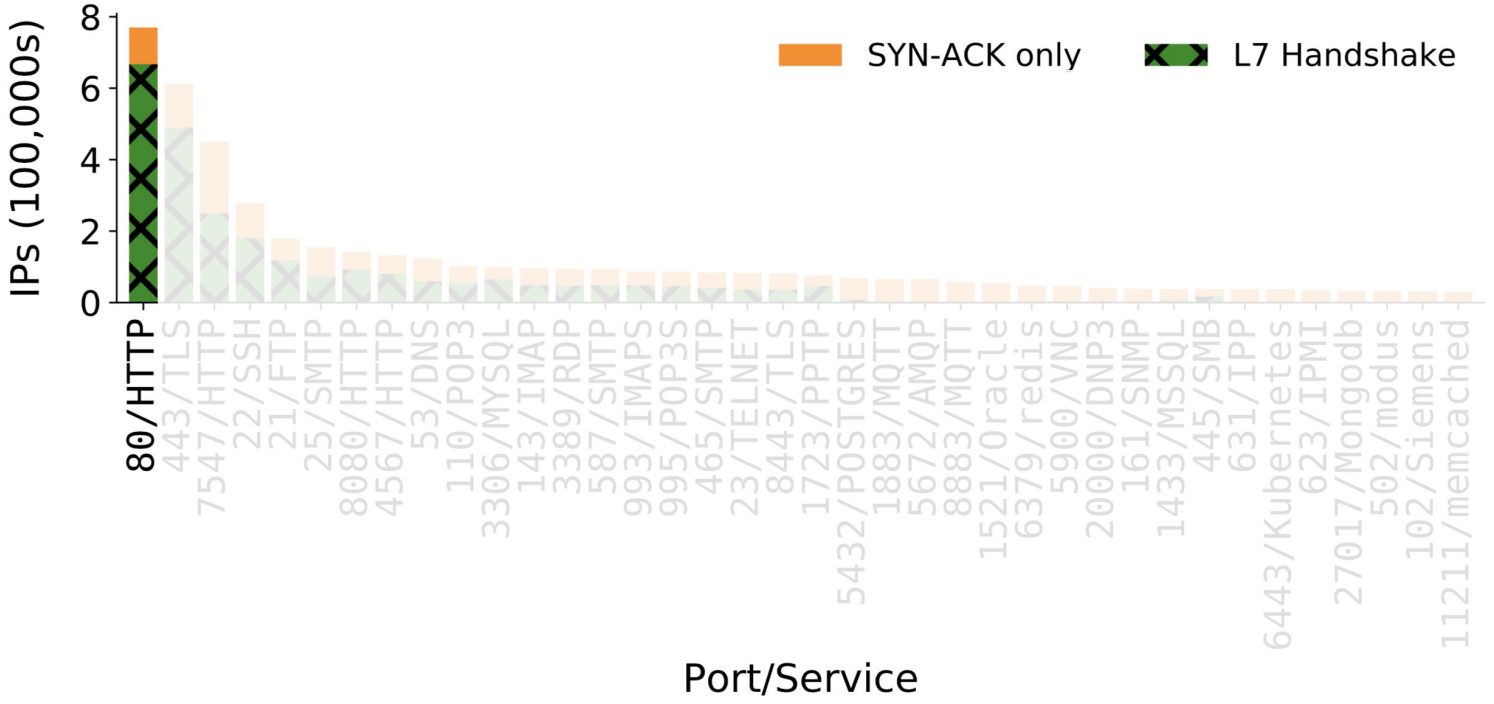

96% of hosts do NOT complete the expected L7 handshake on port 502

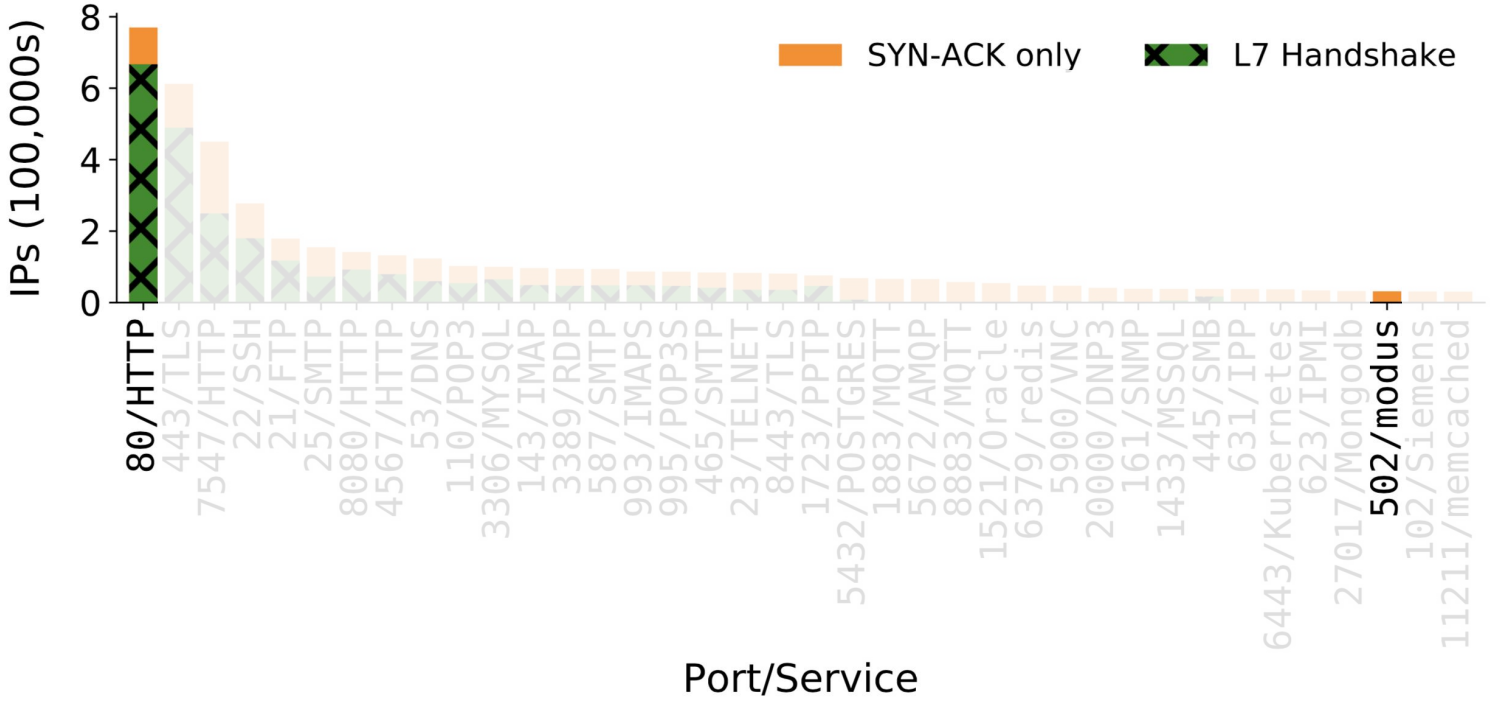

Why are hosts not completing the expected L7 handshake?

Broken TCP Stack?
or
Middleboxes?
or
Unexpected services on the IANA assigned port?

Past methodology for identifying real TCP services is insufficient

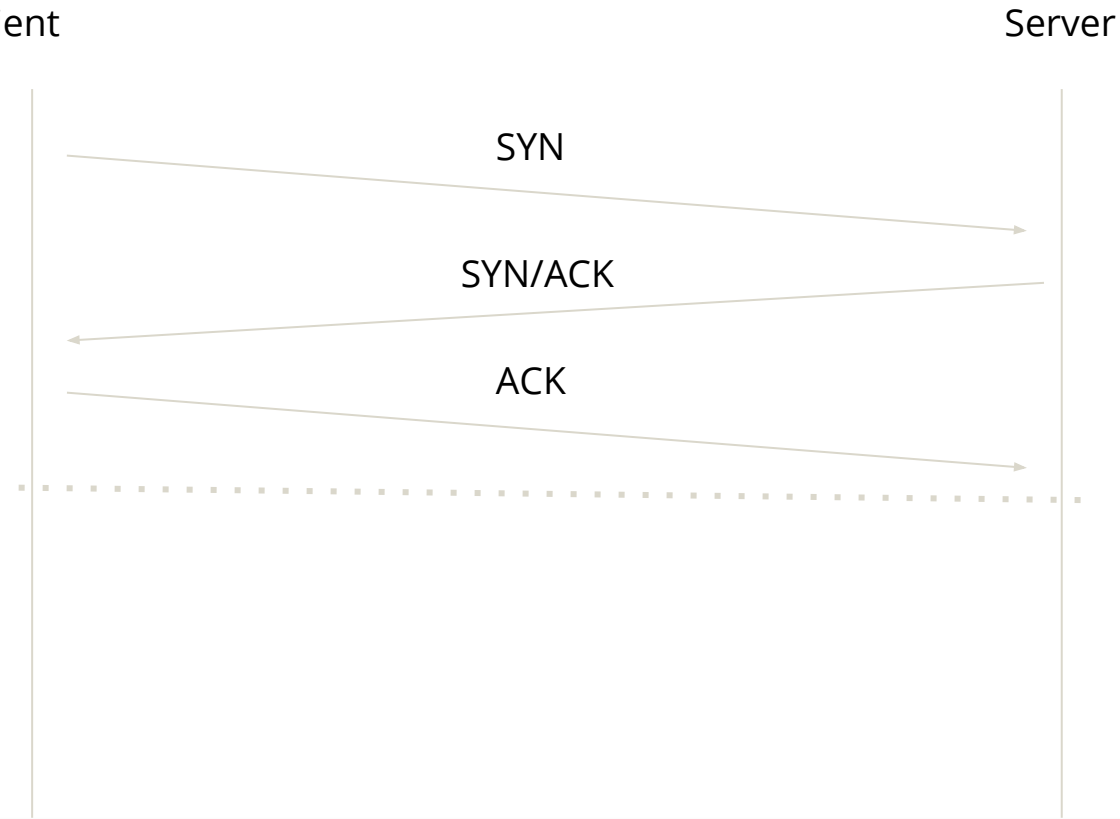

Past methodology for identifying real TCP services is insufficient

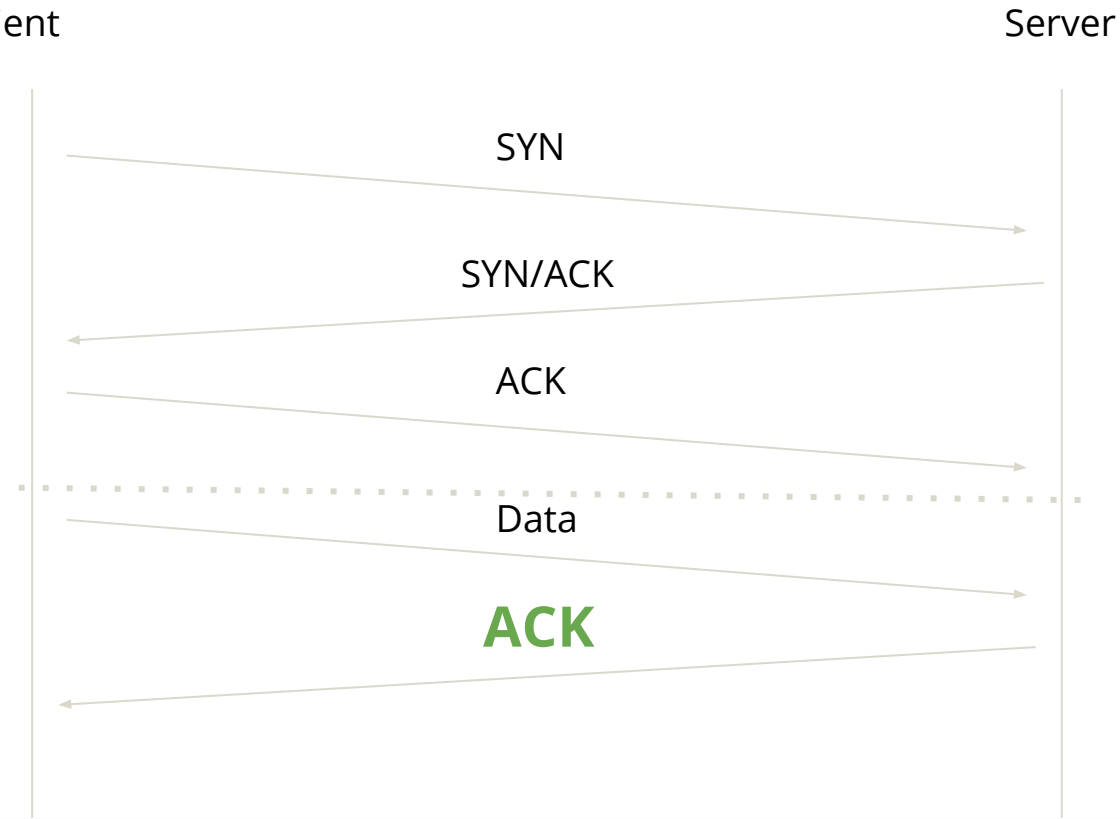

Real Service must:

- accept data
- acknowledge received data

Not all SYN-ACKing IPs acknowledge data

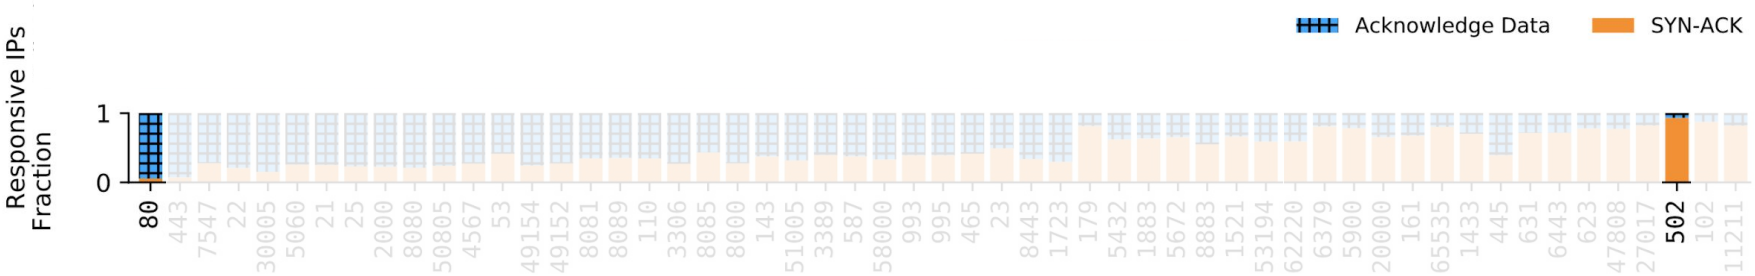

(a) Portion of TCP-responsive hosts that fail to acknowledge data

Not all SYN-ACKing IPs acknowledge data

See paper for 5 reasons why (middleboxes)

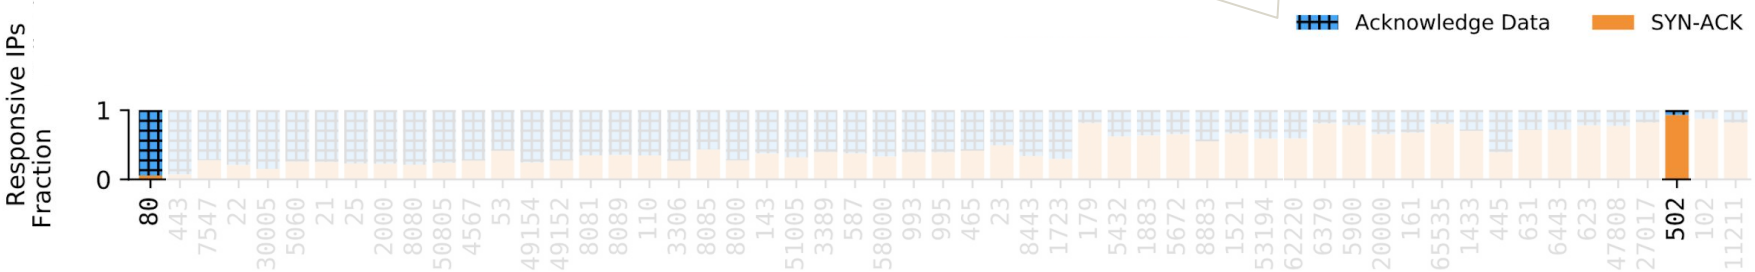

(a) Portion of TCP-responsive hosts that fail to acknowledge data

What fraction of services that acknowledge data are unexpected?

Experiment:

- Scan all 65,535 ports with 30 unique protocols across 0.1% of IPv4
- Filter for services that acknowledge data

What fraction of services that acknowledge data are unexpected?

Experiment:

- Scan all 65,535 ports with 30 unique protocols across 0.1% of IPv4
- Filter for services that acknowledge data

Result:

- 27% of services on popular ports and 63% of services on unpopular ports are unexpected.

HTTP and TLS dominate unexpected services

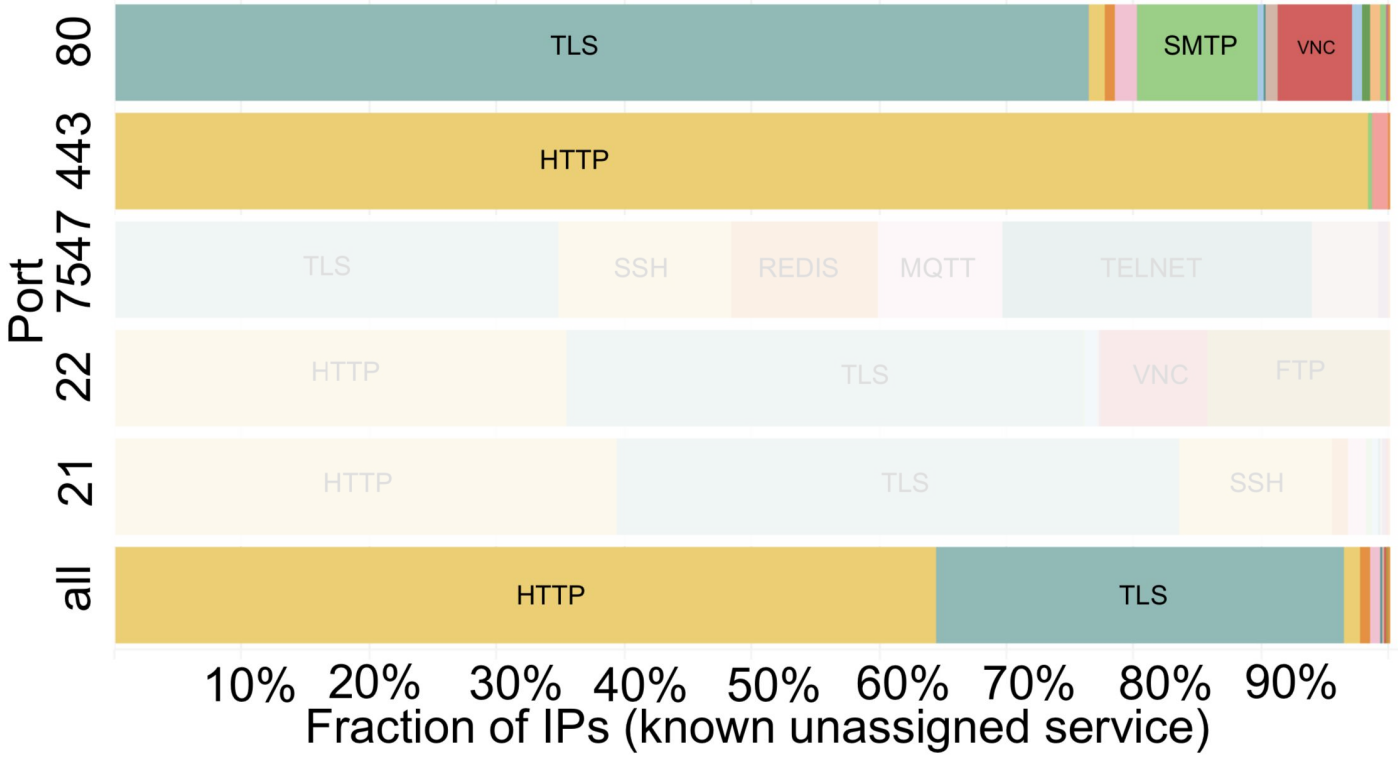

IANA-Assigned protocols are diffuse

- Only 3% of HTTP
  → Port 80
- Only 5.5% of Telnet
  → Port 23
- Only 6.4% of TLS
  → Port 443

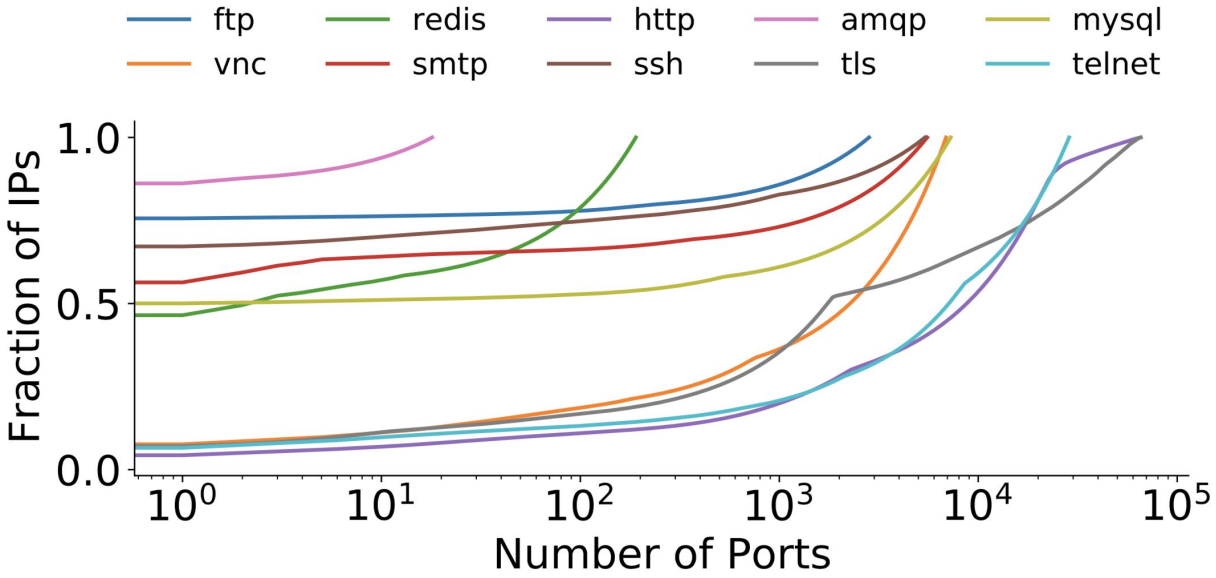

50% of unexpected TLS belongs to IoT

35% of 8000/TLS in Korea Telecom

5% of 8443/TLS across Korean Networks

38% of 80/TLS across 1% of all ASes

Unexpected services are more vulnerable than assigned services

- Ports hosting unexpected TLS host 2x more certificates with a known private key compared to Heninger et al.[1] and Hastings et al.[2]
- 23% of ports hosting unexpected TLS are more likely to host shared public keys than 443/TLS
- Ports hosting unexpected SSH are 2.4 times more likely to allow non-public key authentication

[1]N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman. Mining your Ps and Qs: Detection of widespread weak keys in network devices. In *21st USENIX Security Symposium*, 2012.

[2]M.Hastings,J.Fried,and N.Heninger.Weak keys remain widespread in network devices. In *ACM Internet Measurement Conference*, 2016.

Security studies should scan unexpected services

How do we scan to find unexpected services?

30+ Handshakes/Port is too intrusive and costly

Scanning Insight: 8/30 protocols identify themselves first to the client

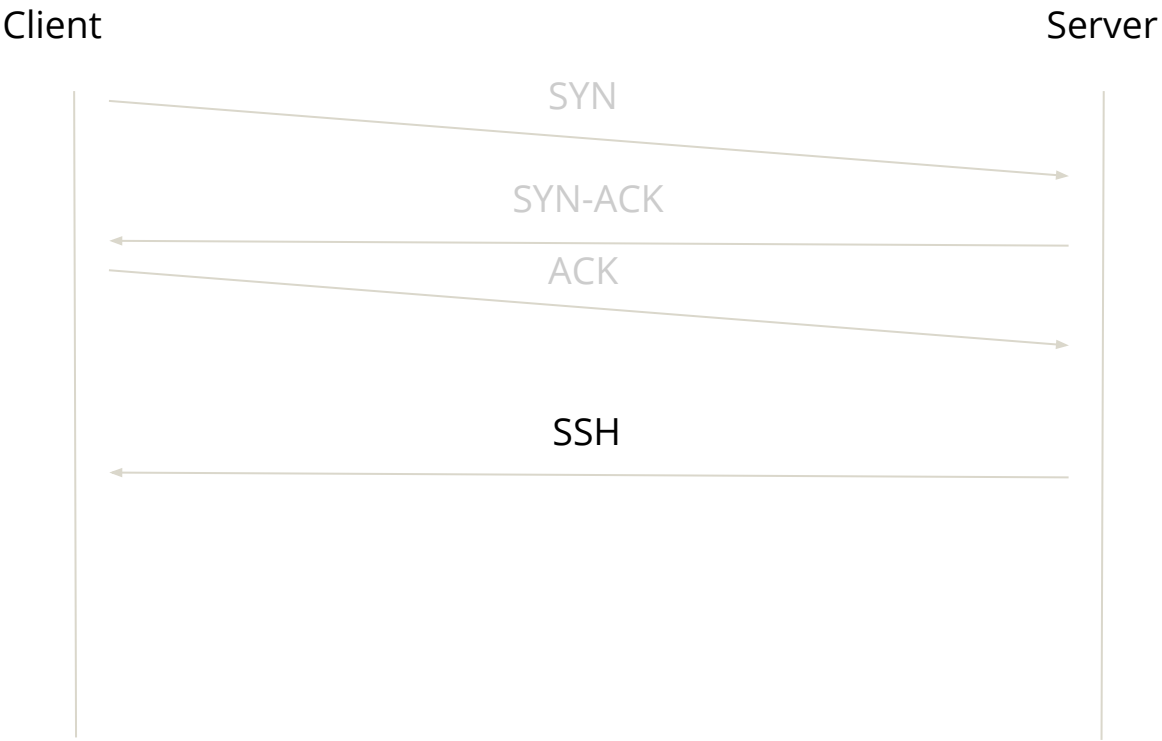

Scanning Insight: 10/30 protocols identify themselves to the wrong handshake

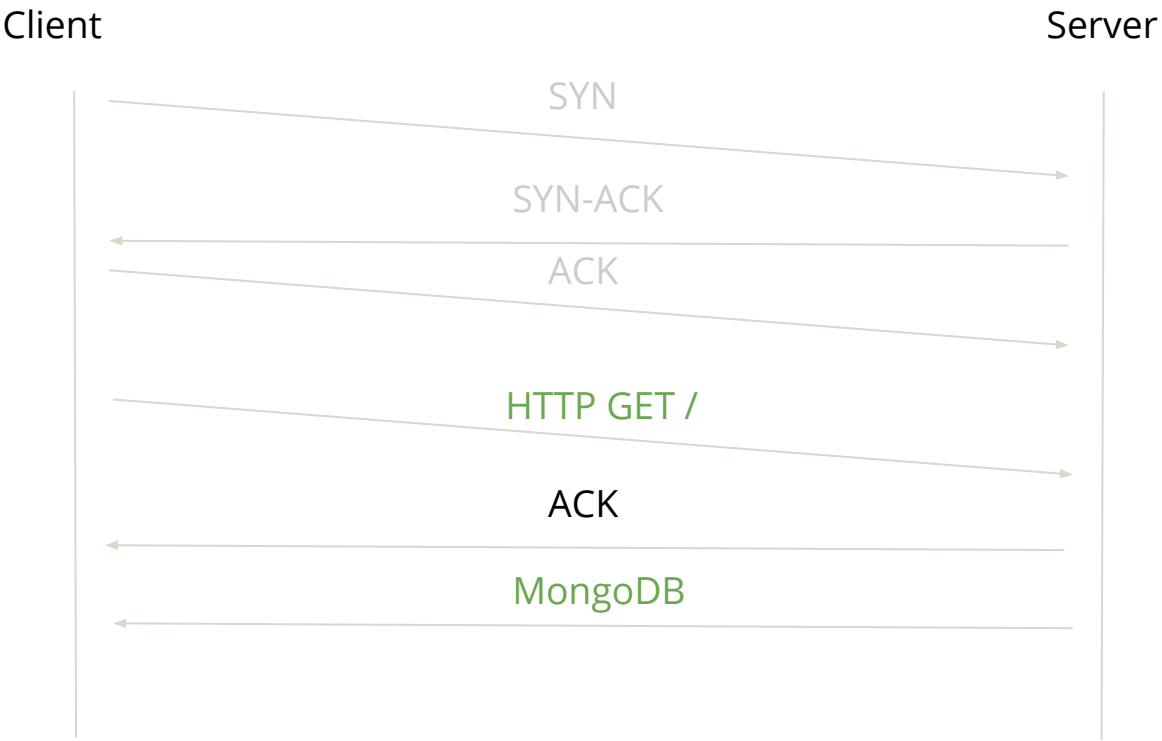

LZR: A system for efficiently identifying services

LZR's framework

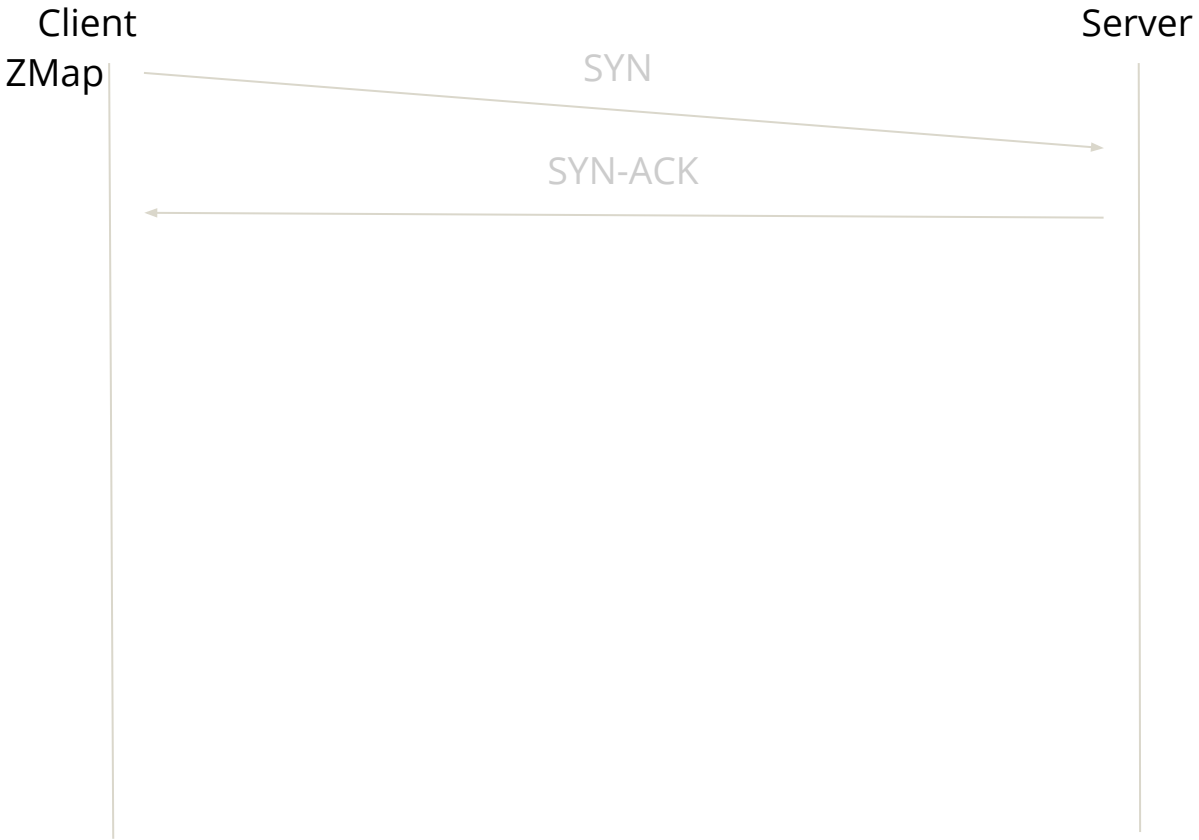

LZR's framework

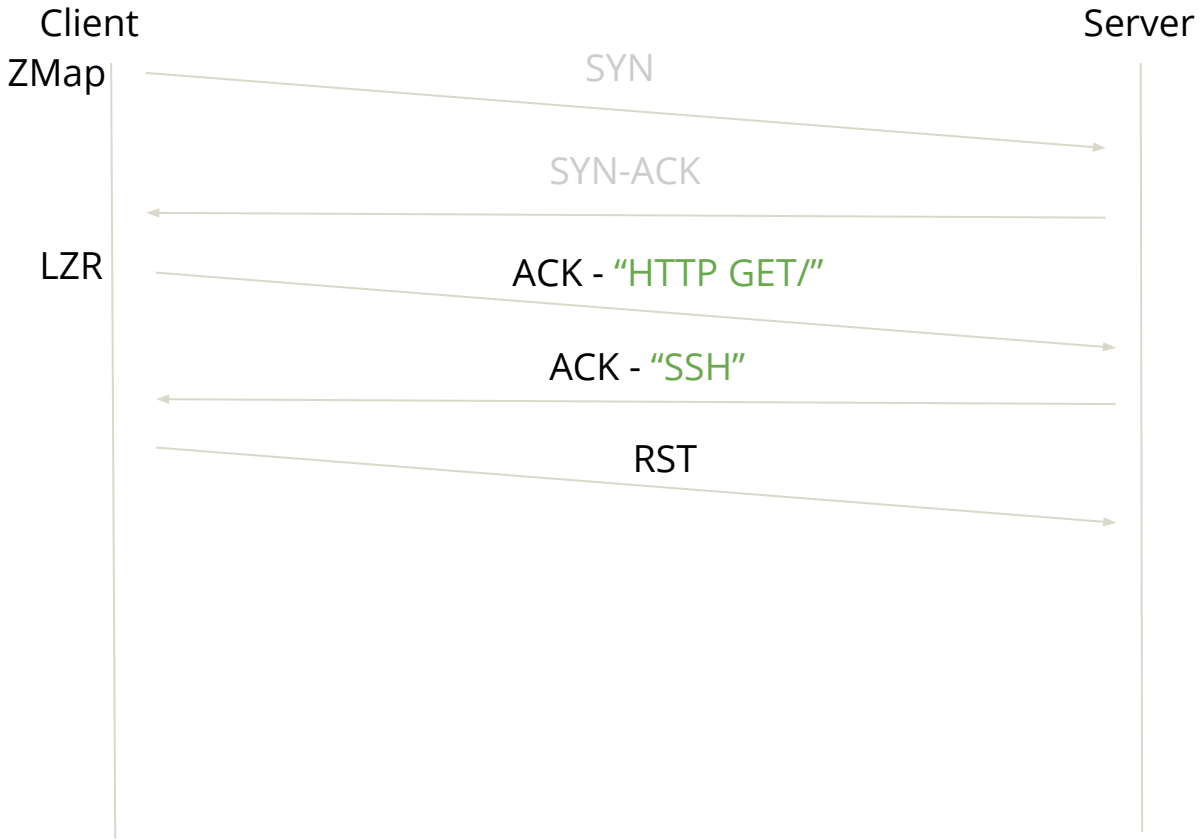

LZR's framework

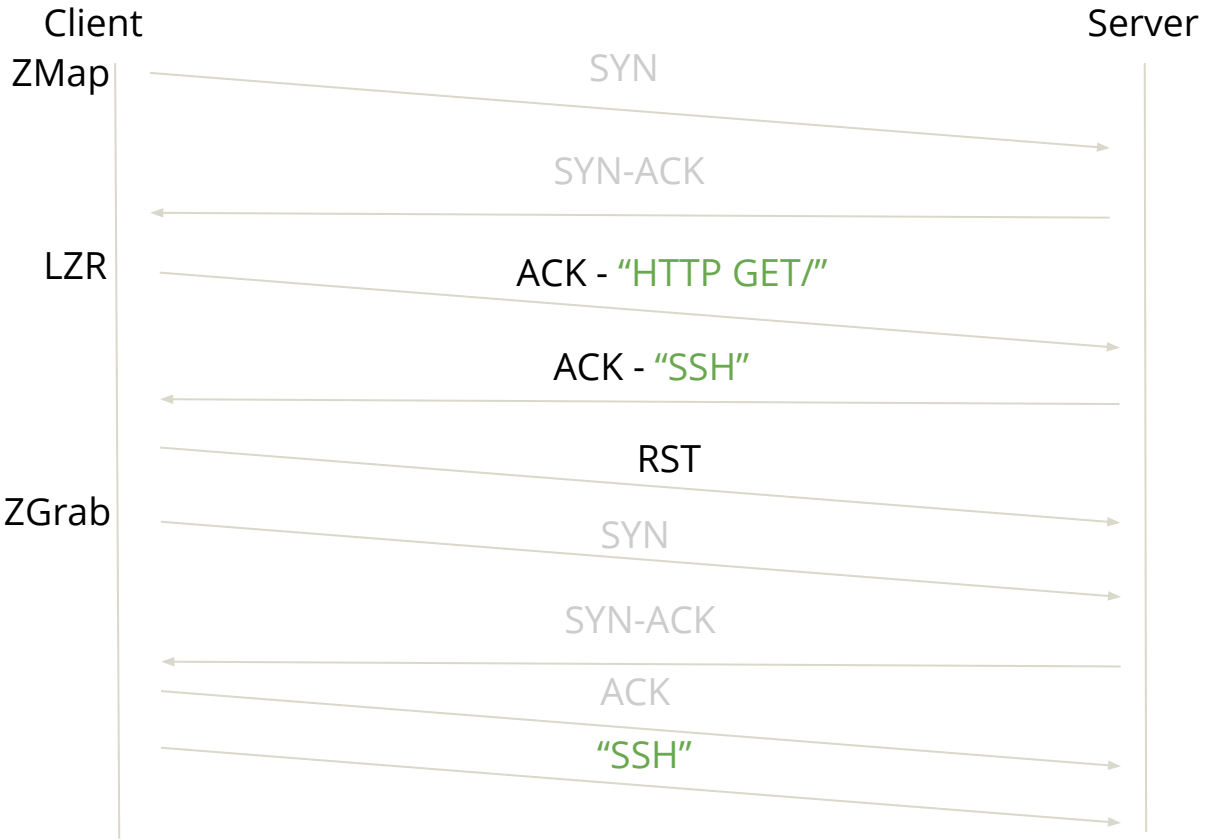

LZR key features

Fingerprints All Server Responses

Ignores Non-Acknowledging Hosts

LZR key features

Fingerprints All Server Responses → Finds Unexpected Services

Ignores Non-Acknowledging Hosts → Reduces Scanning Time!

LZR is up to 55x faster than ZGrab

Port Protocol(s) (Consecutively Scanned)	80 HTTP	443 TLS	21 FTP	23 TEL	5672 AMQP	5900 VNC	27017 Mongo	62220 HTTP	80 HTTP TLS	443 TLS HTTP	47808 HTTP TLS
Offline ZMap + LZR	4.1×	4.1×	5×	10.7×	11.4×	13.3×	55×	25.3×	5.6×	3.4×	29×

LZR finds up to 18 unique protocols in one scan

Port Protocol(s) (Consecutively Scanned)	80 HTTP	443 TLS	21 FTP	23 TEL	5672 AMQP	5900 VNC	27017 Mongo	62220 HTTP	80 HTTP TLS	443 TLS HTTP	47808 HTTP TLS
Unique Unexpected Protocols Found	18	16	10	10	11	8	14	12	18	16	14

Takeaways

- A SYN-ACK != Real Service
  - Scanning studies must scan Layer 7 to find real services

Takeaways

- A SYN-ACK != Real Service
  - Scanning studies must scan Layer 7 to find real services

- IANA-assigned protocols are diffuse across all 65K ports
  - Scanning studies should scan for protocols across all ports

Takeaways

- A SYN-ACK != Real Service
  - Scanning studies must scan Layer 7 to find real services
- IANA-assigned protocols are diffuse across all 65K ports
  - Scanning studies should scan for protocols across all ports
- Unexpected services are more likely to be vulnerable
  - Security studies should scan for protocols across all ports

Takeaways

- A SYN-ACK != Real Service
  - Scanning studies must scan Layer 7 to find real services
- IANA-assigned protocols are diffuse across all 65K ports
  - Scanning studies should scan for protocols across all ports
- Unexpected services are more likely to be vulnerable
  - Security studies should scan for protocols across all ports
- LZR is an open-sourced scanner that efficiently finds unexpected services
  - https://github.com/stanford-esrg/lzr

Questions?
lizhikev@stanford.edu