ACM IMC 2024
MADRID

Ten Years of ZMap

Zakir Durumeric, David Adrian, **Phillip Stephens**, Eric Wustrow, and J. Alex Halderman

ZMap: circa 2013

- ZMap is a L4 network scanning tool
- Released in 2013 at USENIX Security
- Improved
  - Convenience
  - Temporal Resolution
  - Host Requirements

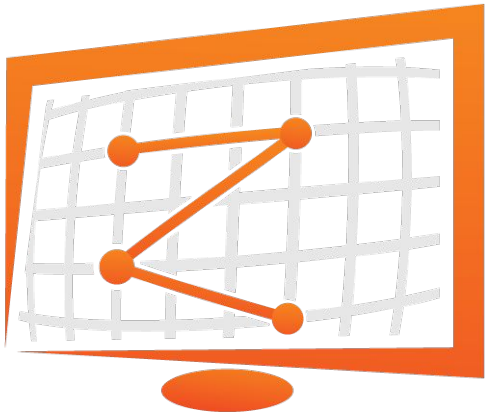

ZMap: 35% of all v4 Internet-wide TCP Scanning

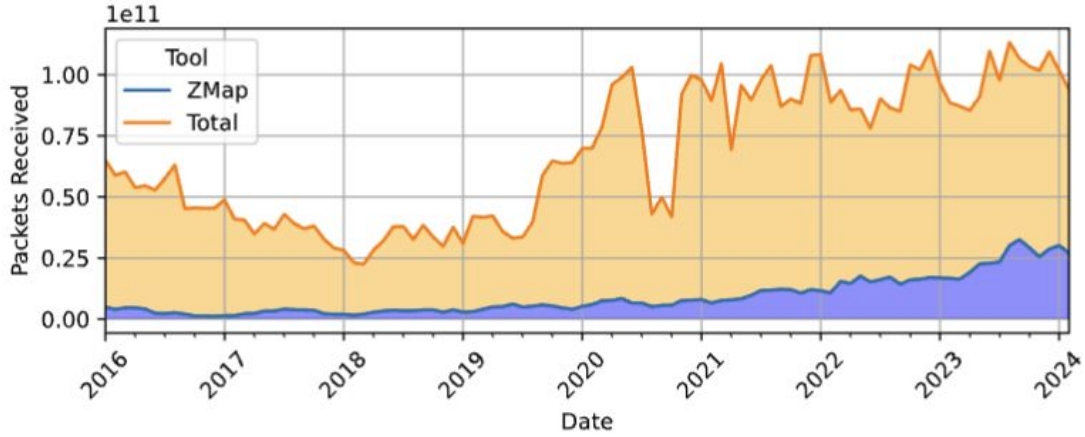

Differences in Targeted Ports

Port Legend

23	- Telnet
80/8080	- HTTP
6379	- Redis Database
8728	- MikroTik Router

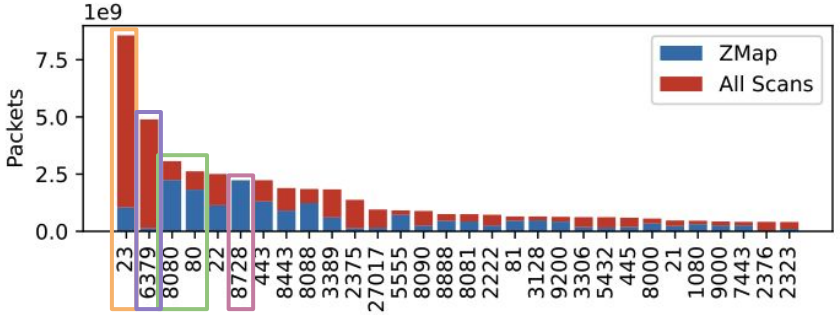

Figure 2: All TCP Scans (Top 30 Ports by Packet)

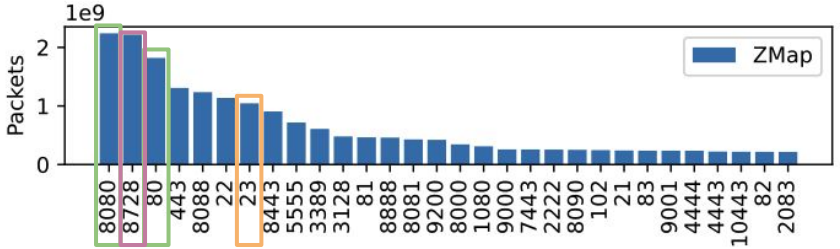

Figure 3: Top Ports Scanned by ZMap (by Packet)

Differences in Regional Usage

US	NL	RU	DE	GB	BG	CN	IN	ZA	HK
66%	33%	0.48%	18%	69%	9%	2%	12%	0.1%	2%

Figure 4: Top 10 Countries by Scan Volume and their Per-Country Scan Traffic from ZMap

How have academics used ZMap?

Academic Usage

- 1,034 References
  - 307 directly using ZMap
- These papers cover a wide variety of topics
  - 38 - Protocol Weaknesses in TLS
  - 25 - Internet-of-Things
  - 14 - Industrial Control Systems
  - 12 - Security-relevant Services
  - 24 - DNS
  - 12 - BGP/RPKI
  - 14 - Censorship
  - 10 - IP usage/NAT

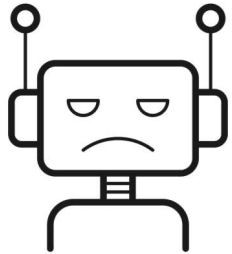

Enabled Reverse-Engineering Nation-State Spyware

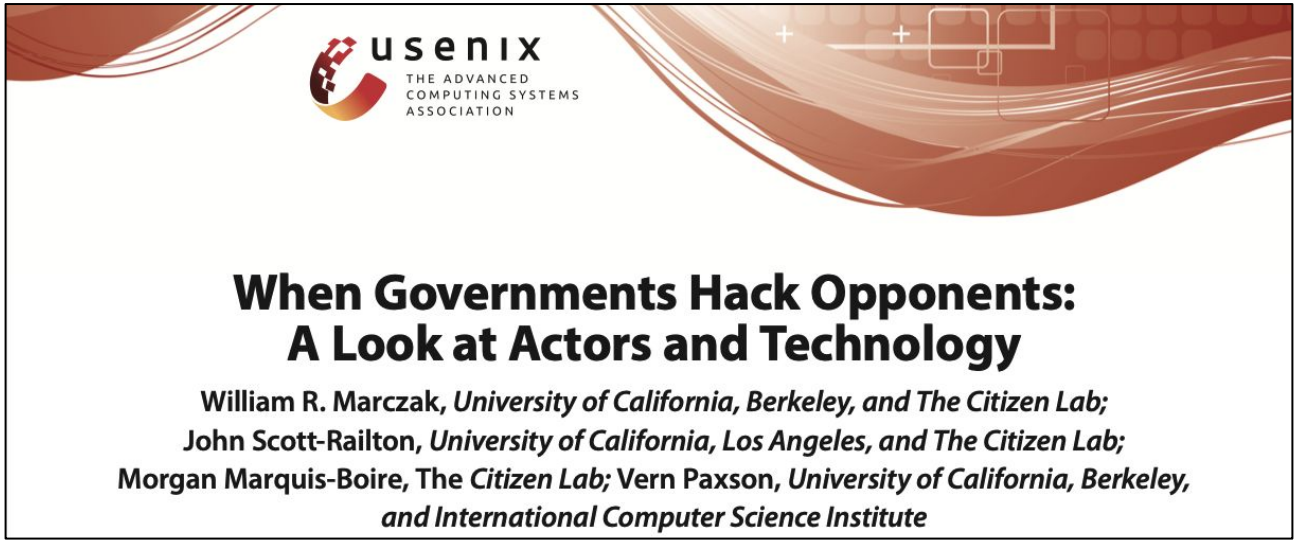

Enabled Democratized LEO Satellite Network Measurement

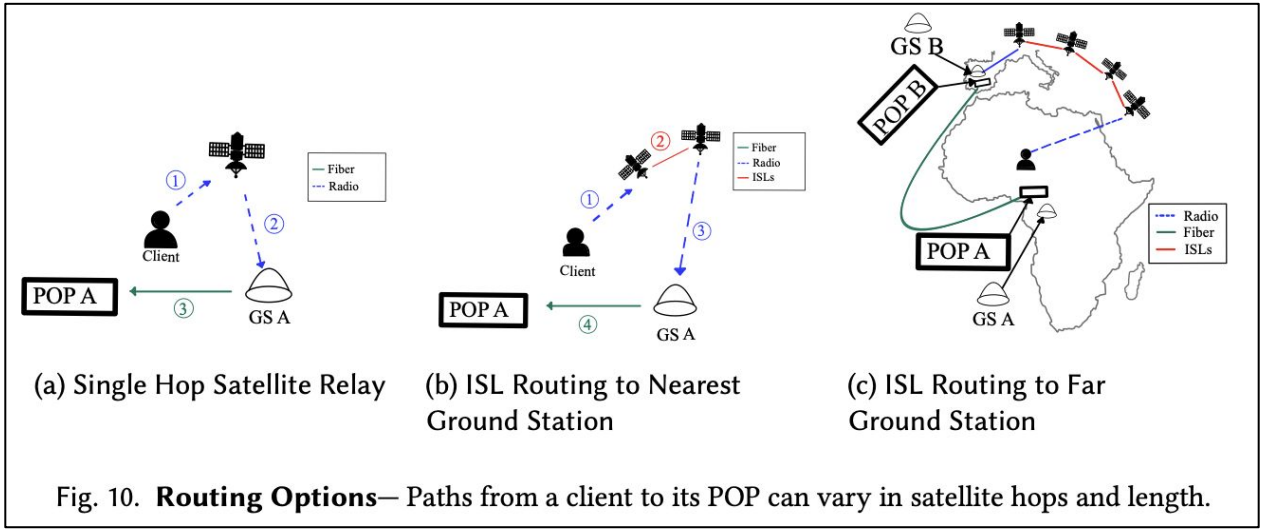

Fig. 10. **Routing Options**— Paths from a client to its POP can vary in satellite hops and length.

Democratizing LEO Satellite Network Measurement, Izhikevich et al., 2024

Enabled DROWN: a novel attack against TLS

Proceedings of the 25th USENIX Security Symposium, August 2016 **https://drownattack.com**

DROWN: Breaking TLS using SSLv2

Nimrod Aviram[1], Sebastian Schinzel[2], Juraj Somorovsky[3], Nadia Heninger[4], Maik Dankel[2], Jens Steube[5], Luke Valenta[4], David Adrian[6], J. Alex Halderman[6], Viktor Dukhovni[7], Emilia Käsper[8], Shaanan Cohney[4], Susanne Engels[3], Christof Paar[3] and Yuval Shavitt[1]

[1]Department of Electrical Engineering, Tel Aviv University
[2]Münster University of Applied Sciences
[3]Horst Görtz Institute for IT Security, Ruhr University Bochum
[4]University of Pennsylvania
[5]Hashcat Project
[6]University of Michigan
[7]Two Sigma/OpenSSL
[8]Google/OpenSSL

Industry Drives Usage

Industry, not academics, driving usage

- None of the top 100 AS's by ZMap traffic are from academic institutions
- Using Greynoise scanning attribution, found usage driven by security companies
  - Attack Surface Management
  - 3rd Party Risk Managers
  - Internet Intelligence

paloalto NETWORKS

BITSIGHT

SHADOWSERVER

ipinfo.io

FICO

Much has changed…

1. Support Multi-Port Scans
   a. Response Deduplication
   b. Randomization Algorithm
2. Changed randomization shard mechanism
3. Changed packet construction
4. General Usability/Bug Fixes

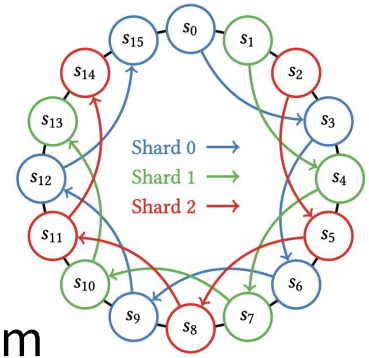

(a) Interleaved (Old)

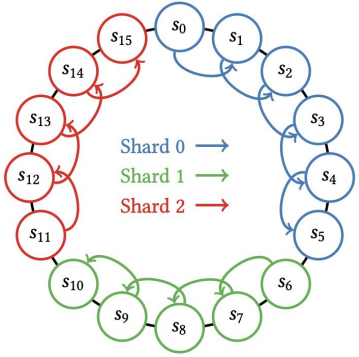

(b) Pizza (New)

Lessons over a Decade

Simple Tools Working Together

- ZMap was initially built as a measurement framework
  - Poor usage fit - No one wanted to use a complicated framework

- Simple, easy-to-understand tools that work well together

ZMap: a tool in a pipeline

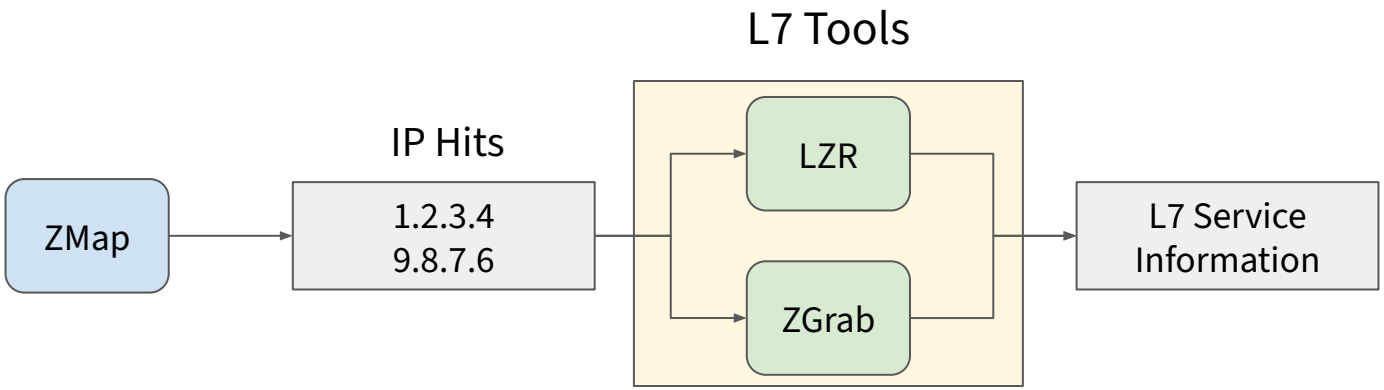

Library + CLI Wrapper

- CLI-only applications are natural, but difficult to integrate into larger systems

- A Library + CLI wrapper lets companies or larger organizations incorporate your tool

Provide Metadata for Post-Scan Forensics

- Most UNIX tools have two types of output: *stdout* and *stderr*
  - We recommend data, logs, and metadata
- Metadata is important to know if a scan was problematic without analyzing the data.
  - NIC drops packets
  - Gateway was inaccessible

Scanning Recommendations

- Investigate existing datasets

- Publish new datasets if existing is insufficient

Looking ahead

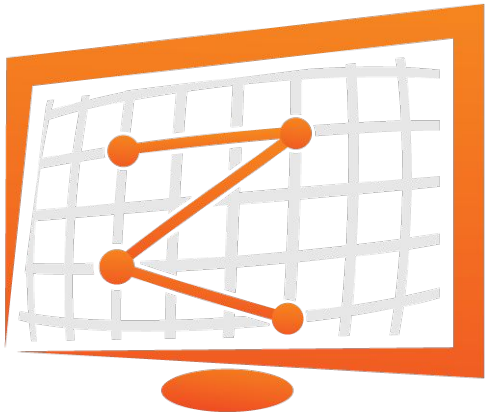

Thank you!

For additional questions, email:
phillip@cs.stanford.edu

Stanford University