SHADOWSERVER

Lighting the way to a more secure Internet

SHADOWSERVER

Internet Spelunking

IPv6 Scanning and Device Fingerprinting

Dave De Coster // Piotr Kijewski
decoster@shadowserver.org // piotr@shadowserver.org

30th June, 2022
2022 FIRST Annual Conference, Dublin

SHADOWSERVER.ORG

149,281,685 Reported IPs

103,553,198,124 UDP Probes
214,618,534,185 TCP SYN
351,695,957 Full Handshakes

Ground Rules

Do no harm

Never exploit

Test, test, test, 1/250th test

Test some more

First, do no harm

- Scans will not compromise, harm, or degrade system performance
  - Use the smallest and most minimal packet possible to get the results
  - Test repeatedly before a full Internet scan occurs
  - 1/250th test
- Only scan what is necessary for remediation
  - Vulnerable or misconfigured systems
  - Specific ports used by criminal infrastructures
- Scans will not break any US laws

SHADOWSERVER

How Did We Get Here?

No (good?) deed goes unpunished.

The Origin

You can all thank Christian Rossow for publishing:

"Amplification Hell: Revisiting Network Protocols for DDoS Abuse"

https://christian-rossow.de/publications/amplification-ndss2014.pdf

The Origin

- Laid out 14 UDP protocols that could be used for a DDoS, including populations and actual amplification of each protocol
  - 11 were the most worrisome
- We focused on seven

SHADOWSERVER

The Origin

Protocol	Port
SNMPv2	UDP/161
NTP	UDP/123
DNS	UDP/53
NetBIOS	UDP/137
SSDP	UDP/1900
CharGen	UDP/19
QOTD	UDP/17

The Origin

- Started with DNS
  - It was easy
  - Miscreants were already abusing it
  - There were already two open DNS scanners available for us to confirm results against
    - Other data sets were deemed too polluted to be used easily for reporting purposes
    - Cleaning other data sets was difficult and the actual methodology of scanning was flawed by both other scanning entities
    - Better to build something new to meet our more narrow scope and mission

The Origin Story

- First scan took 91 hours to complete
- 16.9 million responses (53/udp only)
- 12.25 million openly recursive

DNS Open Resolvers (port 53/udp only)
Start: 2013-04-15 14:45:43
End: 2013-04-19 10:04:39

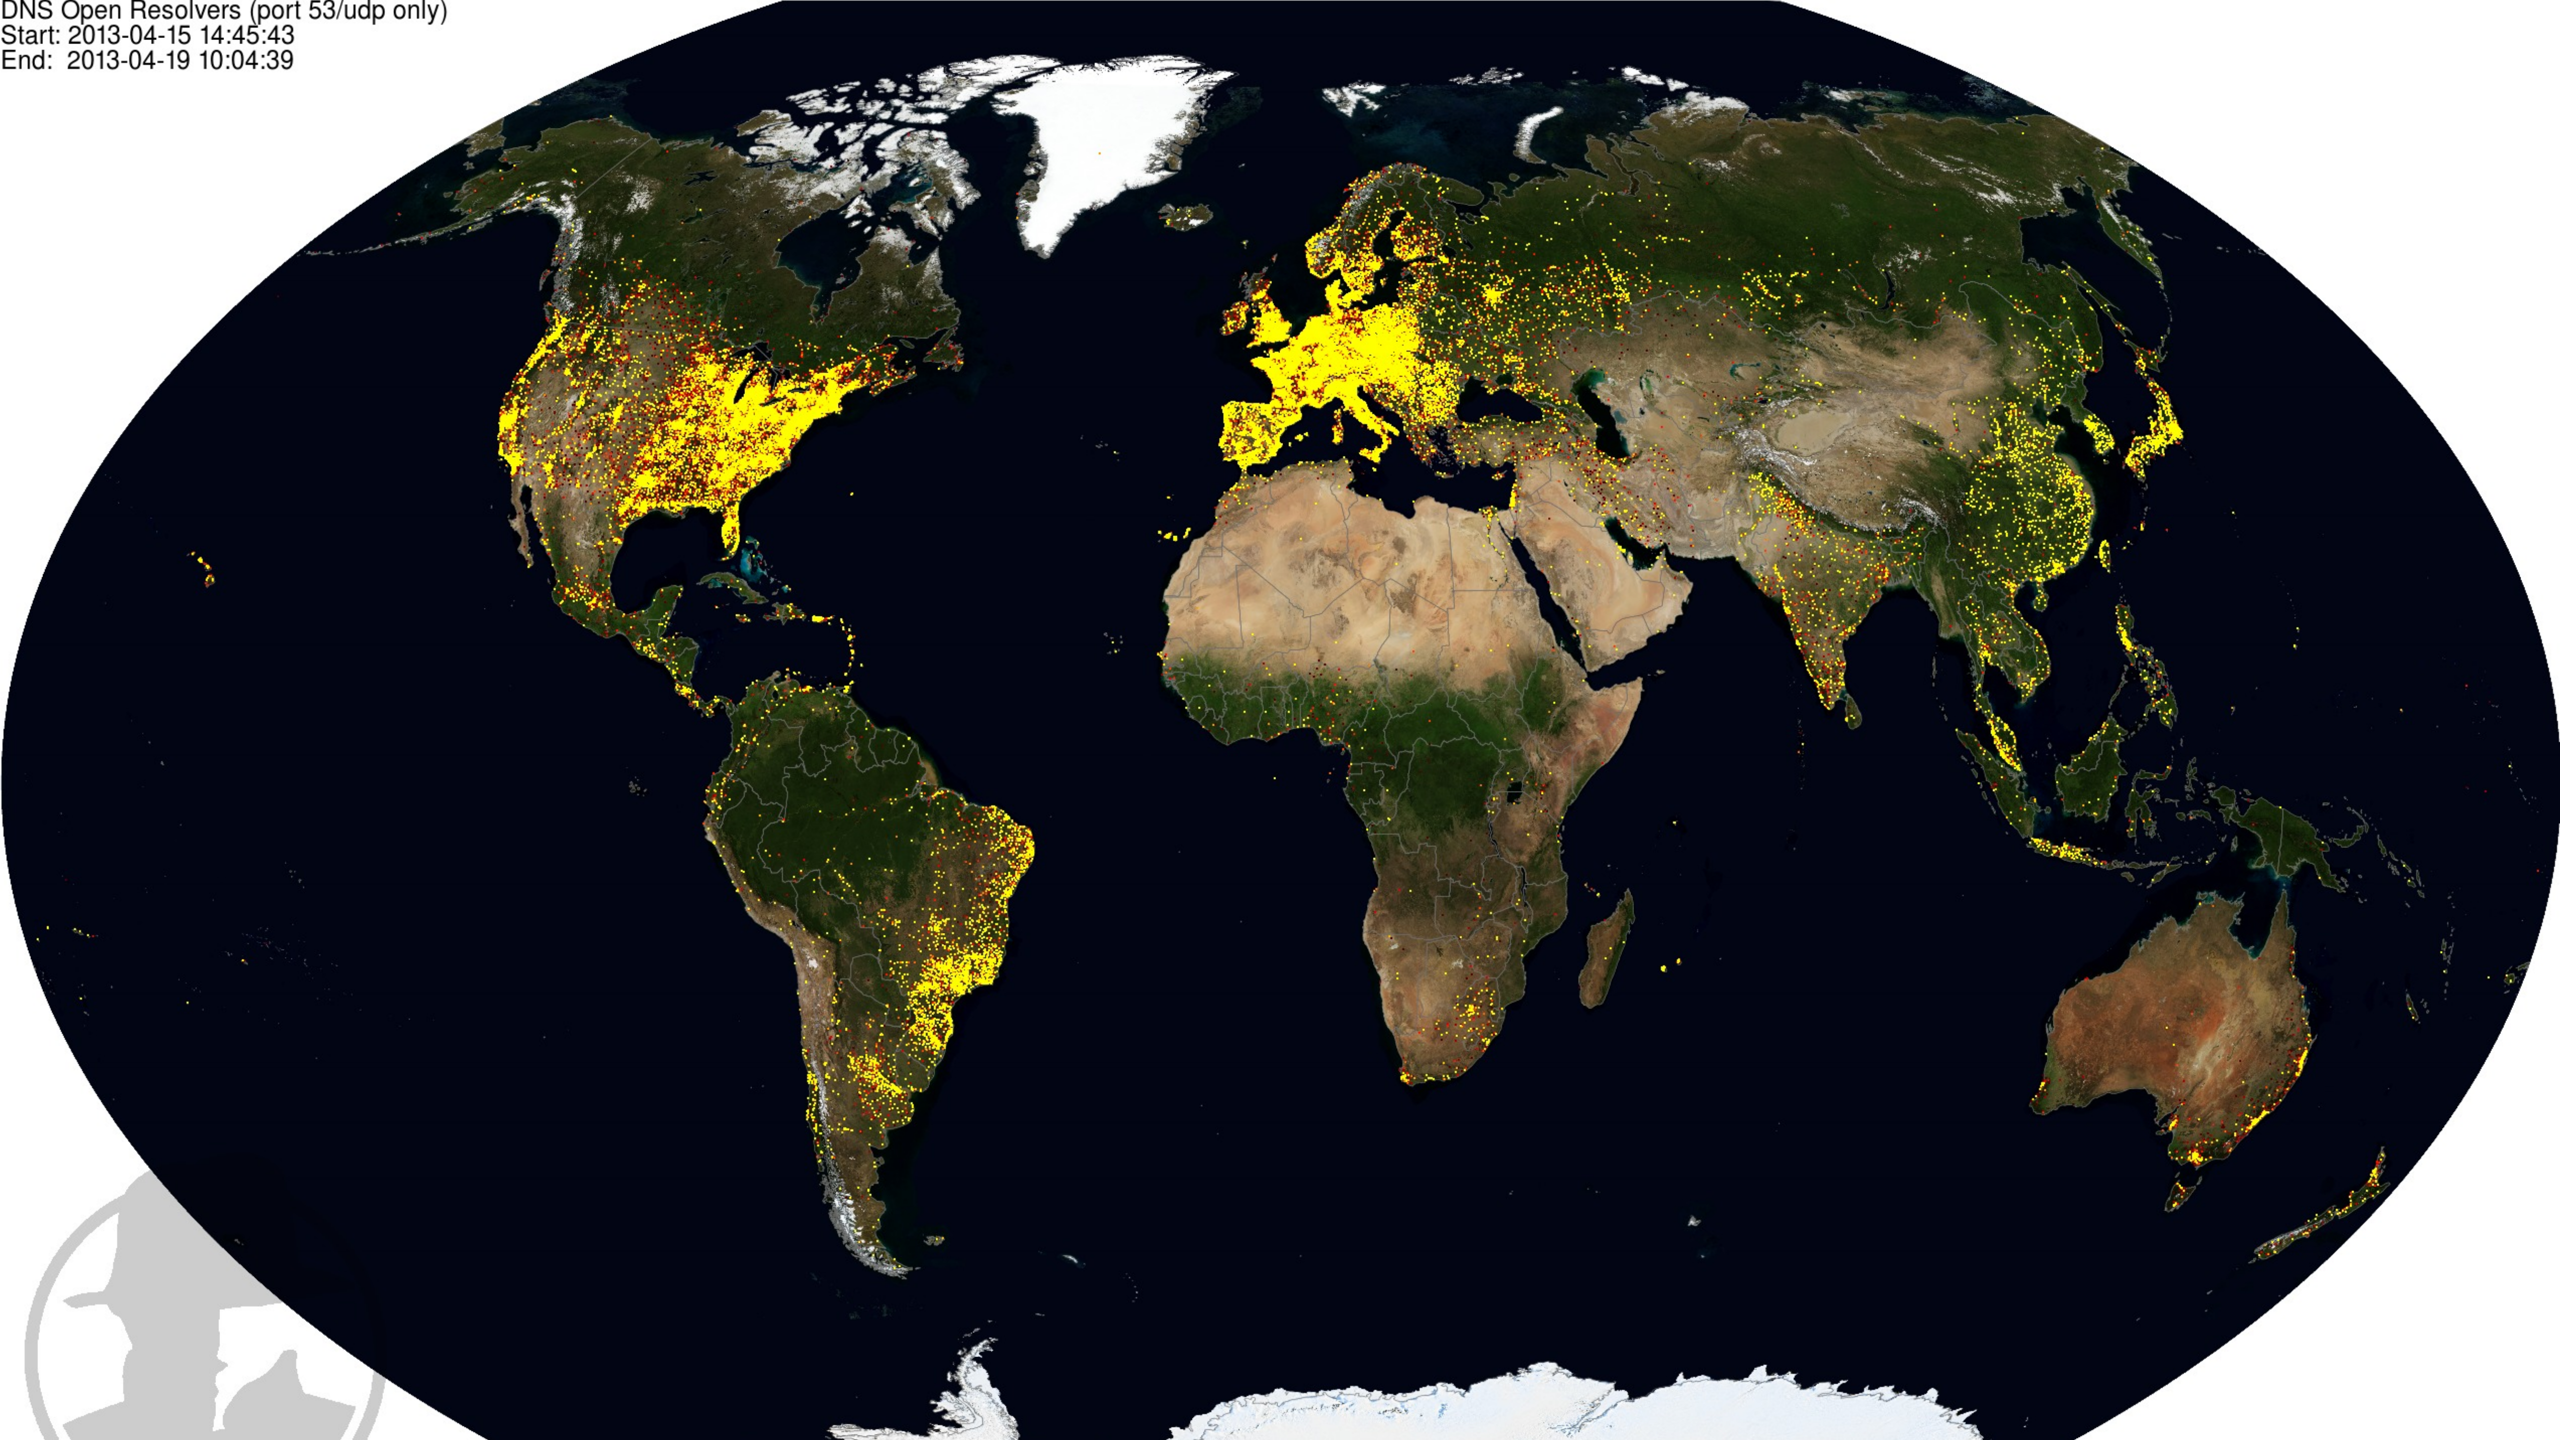

sigh

Fast Forward to curdate()

- The DNS scan now runs in 4 hours
  - 6 million total responses (53/udp only)
  - 1.8 million recursive resolvers

~10.4 million IPs that are no longer abusable

Hey, It worked!

After discovering that the scanning worked, we:

- Acquired more hardware
- Acquired more bandwidth
- Wrote new scanning tools
- Proceeded to implement scans on the rest of the named UDP targets

Something Different!

Smooth sailing until October 2014

- POODLE (SSLv3 Downgrade)
  - Padding Oracle On Downgraded Legacy Encryption

Needed to learn some new tricks..

Discovered that scanning /0 for UDP is *much* easier than TCP

- UDP is just Spray’n’Pray (with some limits)
  - Self DDoS’s can hurt if not controlled and rate limited
- TCP you have to track state and scan twice
  - And you have to talk x509!

SHADOWSERVER

Success! (it took a bit)

- First reported POODLE data:
  - November 2014
  - 15,573,251 IPs vulnerable to a downgrade attack

Fast Forward to curdate()

POODLE (SSLv3) now:

- 2,157,293

- Still a big number, but better

Expansion of the beast

We couldn't let all the lessons we learned sit idle, so we added in a *few* more scans..

SHADOWSERVER

Over 100 Full Scans a Day

Protocol	Port	Protocol	Port	Protocol	Port	Protocol	Port	Protocol	Port	Protocol	Port
AMQP	5672/tcp	DVR DHCPDiscover	37810/udp	HTTPS	5001/tcp	MELSEC-Q	5007/tcp	Omron FINS	9600/udp	SOCKS4/5	1080/tcp
Android Debug Bridge	5555/tcp	ElasticSearch	9200/tcp	HTTPS	4433/tcp	MemCacheD	11211/udp	OPC-UA	4840/tcp	SSDP	1900/udp
Apple File Protocol	548/tcp	EPMD	4369/tcp	HTTPS	6443/tcp	MemCacheD	11211/tcp	PCWORX	1962/tcp	SSH	22/tcp
Apple Remote Management	3283/udp	EtherCAT	34980/udp	HTTPS	447/tcp	Microsoft Exchange	443/tcp	PLEX SSDP	32414/udp	SSH (IPv6)	22/tcp
BACnet	47808/tcp	EtherNet/IP	44818/tcp	HTTPS	4117/tcp	Middlebox	80/tcp	Portmapper	111/udp	SYNful Knock	80/tcp
CharGEN	19/udp	FTP	21/tcp	HTTPS	8080/tcp	Mikrotik (Speed Test)	2000/tcp	ProConOS	20547/tcp	Telnet	23/tcp
cLDAP	389/udp	GE-SRTP	18245/tcp	HTTPS	5443/tcp	Mitel	10074/udp	QOTD	17/udp	Telnet	2323/tcp
CoAP (v1)	5683/udp	Hadoop (DataNode)	50075/tcp	HTTPS	7443/tcp	MODBUS	502/tcp	QUIC	443/tcp	Telnet (IPv6)	23/tcp
CoAP (v2)	5683/udp	Hadoop (NameNode)	50070/tcp	HTTPS (IPv6)	443/tcp	MongoDB	27017/tcp	Radmin	4899/tcp	TFTP	69/udp
CODESYS IEC 61131-3	2455/tcp	HART	5094/tcp	HTTPS	443/tcp	MQTT	1883/tcp	RDP	3389/tcp	Tridium Niagra	1911/tcp
CODESYS IEC 61131-3	1200/tcp	HTTP	80/tcp	ICCP	102/tcp	MQTT SSL	8883/tcp	RDPEUDP	3389/udp	Ubiquiti Discovery Service	10001/udp
CouchDB	5984/tcp	HTTP	8080/tcp	IEC 60870-5-104	2404/tcp	MS-SQL	1434/udp	Redis	6379/tcp	VNC	5900/tcp
Crimson (Red Lion)	789/tcp	HTTP	8000/tcp	IPMI	623/udp	MySQL	3306/tcp	rsync	873/tcp	VNC	5901/tcp
CWMP	7547/tcp	HTTP	80/tcp	IPP	631/tcp	MySQL (IPv6)	3306/tcp	S7	102/tcp	XDMCP	177/udp
CWMP	30005/tcp	HTTPS	8443/tcp	ISAMKP	500/udp	NAT-PMP	5351/udp	SmartInstall	4786/tcp		
DB2	523/udp	HTTPS	9000/tcp	Kubernetes	6443/tcp	NetBIOS	137/udp	SMB	445/tcp		
DNP3	20000/tcp	HTTPS	449/tcp	Kubernetes	443/tcp	Netis	53413/udp	SMTP	25/tcp		
DNS	53/udp	HTTPS	10443/tcp	LDAP	389/tcp	NTP (Monitor)	123/udp	SMTP (IPv6)	25/tcp		
Docker	2375/tcp	HTTPS	8010/tcp	mDNS	5353/udp	NTP (Version)	123/udp	SNMPv2	161/udp		

SHADOWSERVER

How and Why are the next targets chosen

- Topical – new blog comes out with a vulnerability that can be remotely tested
  - Netis, Synfulknock, ISAKMP, etc
- Looking at legacy protocols that really should not be exposed
  - Telnet, rsh, etc
- Current protocols that really should not be exposed
  - MongoDB, Kubernetes, etc
- Someone asked us to look for it

SHADOWSERVER

Some scans are easier than others

- "Banner" services
  - Things that respond to a single packet are easy
  - Telnet, TFTP, et cetera

- Negotiated services
  - Services where you need a HELO or client/server agreement
  - SSL, SSH

- Multi-Step services
  - Services that require a stepwise response to get an answer
  - IEC 60870-5-104

SHADOWSERVER

Fun Facts

We have sent (with daily repeats):

- 209,724,213,326,259 UDP Probes
  - 209.7 Trillion UDP Probes
- 221,639,352,853,200 TCP SYNs
  - 221.6 Trillion TCP Syns
- 508,013,815,018 Full Protocol Connections
  - 508 Billion Connections
- 287,916,573,658 Services for remediation
  - 287.9 Billion Reported

SHADOWSERVER

Sorry for the noise…

The Gear

How the work gets done – Grab the hearing protection

Stack o’ Boxes in a Colo

Just a pile of leftover gear

- 37 x Cisco C220 M3’s
  - 256 GB Memory
  - 5 TB Disk (8 x 1tb RAID 6)
- 2 x 10 Gb/s lines
- 5 x /26 IP blocks (and 1 /24)

Dirtiest CIDRs on the net?

- We scan from 558 IPs:

  184.105.139.64/26
  184.105.247.192/26
  216.218.206.64/26
  74.82.47.0/26
  65.49.20.64/26
  64.62.197.0/24

- Nodes are each assigned 15 IPs
- Evenly split across 2x 10 gb lines

SHADOWSERVER

Scanning Methodology

- TCP and UDP scans are handled differently
  - TCP Scans are:
    - Broken into shards
      - Shard is 1/250th of the IP space to be scanned
      - IPs in a shard are algorithmically determined by a random seed that is supplied to every shard.
    - Will use the entire cluster to scan
    - Performed using commodity software
  - UDP Scans are:
    - Monolithic
    - Run from a single node
    - Performed using custom software

UDP Scans

- Meet “**railgun**”
  - Designed to send a single UDP packet as randomly as possible and as fast as possible to all 3.4B IPs
  - Tuned for sending small packets
  - Will send packets using all available IPs
  - Has very few safety measures

SHADOWSERVER

UDP Scans

- Railgun can usually scan the internet for one service in around four hours.
  - Highly dependent on the number of responding devices.

SHADOWSERVER

TCP Scans

- Commodity tools
  - Assignment of jobs:
    - HTCondor
  - Actual scanning:
    - Zmap performs the initial sweep
    - Zgrab (mostly) performs the connection
    - Other tools for doing custom things

TCP Scans

Each service takes between ten minutes and three hours

- Dependent on the complexity of the scan
  - Things with no crypto (Telnet) are fast
    - 8 minutes in human time
    - 3 hours and 57 minutes in machine time
  - Things with crypto (HTTPS) are much slower
    - 2 hours and 29 minutes in human time
    - 82 hours in machine time

Same From Here

- The raw data is:
  - Parsed (protocol specific)
  - Sanity checked (bad data?)
  - Standardized
  - Shipped off to the Datacenter to get turned into reports

IPv6

You want to scan what?

Surprisingly Familiar

- Like IPv4, just a LOT more of it
- Not feasible to scan it all, so curated lists
  - IPv6 addresses sourced from SSL certificates, IPv6 Hitlist, other.

- Currently scanning 814,675,045 IPv6 addresses

Blindly Scanning is Infeasible

IPv6 space is 3.48x10^38 unique addresses

Time to scan ~6.33x10^32 seconds

Roughly 2x10^25 years

Blindly Scanning is Infeasible

- Use curated lists from:
  - DNS AAAA records (passive DNS)
  - IPv6 Hitlist: https://ipv6hitlist.github.io/
  - Certificate transparency streams
  - Sinkholes
  - Partners

Yet Different…

Fewer options for scanning tools

- zmap6 from https://github.com/tumi8/zmap
- zgrab/zgrab2 have native IPv6 support
- Other tools.. Not so much

And Slower…

IPv6 requires more gentle timings than IPv4

- IPv4: Potential packet loss at > 500,000 pps
- IPv6: Potential packet loss at > 100,000 pps

And Slower…

IPv6 requires more gentle timings than IPv4

- IPv4: Packet loss at > 3500 concurrent senders
- IPv6: Packet loss at > 1500 concurrent senders

And Slower…

Average number of IPs/second that can be processed

- IPv4: 243,116 IPs/second
- IPv6: 58,542 IPs/second

And Doesn’t Like to Share…

IPv4 and IPv6 scans don’t like running at the same time on the same interface

IPv6 Scans

- SSL (443/tcp, 8443/tcp)
- SMTP (25/tcp)
- TELNET (23/tcp)
- SSH (22/tcp)
- HTTP (80/tcp, 8080/tcp)
- MySQL (3301/tcp)
- FTP (21/tcp)

IPv6 Scan Stats

Scan	Port	Responses
SSL	443/tcp	8,192,360
SSL	8443/tcp	75,432
SMTP	25/tcp	407,521
Telnet	23/tcp	25,267
SSH	22/tcp	839,575
HTTP	80/tcp	109,845,303
HTTP	8080/tcp	415,989
MySQL	3306/tcp	1,424,136
FTP	21/tcp	2,622,208

IPv6 Scans (Observations)

SSL

- Fewer hosts with really old ciphers (SSLv3, TLSv1.0, TLSv1.1)
- 3.86% IPv4 vs 0.04% IPv6

FTP

- Far higher ratio of FTP+SSL
- 55% IPv4 vs 91% IPv6

MySQL

- Far fewer hosts with deny rules
- 42% IPv4 vs 4% IPv6

IPv6 Scans

- Always Looking for More Sources of IPv6 Targets

Device Identification

Fingerprinting all things!

Device Identification

- Take all data we collect in all our daily scans
    - match fields, banners and responses to identify device make-and-model
- Classify all IPs by:
    - device_type
    - device_vendor
    - device_model
    - device_version
    - device_sector

Device Identification

- Scan rule engine implemented
- Classifies scan data as it is submitted to the API
- Currently ~1200 scan rules implemented
- Support for detection of devices from 173 vendors
- Daily successfully classifies over 28M devices (excluding desktops/servers, web servers etc)
- More to come!

Device Identification

Scan rules

? | Import scan rules | Export scan rules | Advanced filters | Create new | Create in bulk

Action: Go 0 of 20 selected

☐	Contact	Name	Device model	Device type	Device vendor	Group	Order	Test count	Usage	Enabled	State	Created	Actions
☐	Piotr Kije…	Allegro_Software_RomPag…	RomPager	embedded-sys…	Allegro Software	Allegro Software	100			✔		2021-11-14	View Edit Delete Clone
☐	Piotr Kije…	Allegro_Software_RomPag…	RomPager	embedded-sys…	Allegro Software	Allegro Software	200			✔		2021-11-14	View Edit Delete Clone
☐	Piotr Kije…	Realtron_Embedded_Syst…		embedded-sys…	Realtron	Realtron	100			✔		2022-04-24	View Edit Delete Clone
☐	Piotr Kije…	ASUS_httpd_server_http_…		router	ASUS	ASUS	90			✔		2021-01-29	View Edit Delete Clone
☐	Piotr Kije…	ASUS_by_AiCloud_html_title		router	ASUS	ASUS	90			✔		2022-04-13	View Edit Delete Clone
☐	Piotr Kije…	ASUS_catchall_FTP_Banner		router	ASUS	ASUS	95			✔		2021-02-05	View Edit Delete Clone
☐	Piotr Kije…	ASUS_router.asus.com		router	ASUS	ASUS	100			✔		2020-11-13	View Edit Delete Clone
☐	Piotr Kije…	ASUS_by_ASUSTek_cert		router	ASUS	ASUS	100			✔		2022-04-14	View Edit Delete Clone
☐	Piotr Kije…	ASUS_asuscomm_issuer_…		router	ASUS	ASUS	101			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_asuscomm_lets_en…		router	ASUS	ASUS	102			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_asuscomm_lets_en…		router	ASUS	ASUS	105			✔		2021-02-01	View Edit Delete Clone
☐	Piotr Kije…	ASUS_ASUSWRT_issuer_…		router	ASUS	ASUS	120			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_ASUSWRT_HGG_is…		router	ASUS	ASUS	200			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_Merlin_Koolshare_i…		router	ASUS	ASUS	202			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_ASUSWRT_Merlin_i…		router	ASUS	ASUS	203			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_Merlin_Koolshare_r…		router	ASUS	ASUS	204			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_ASUSWRT_Merlin_…		router	ASUS	ASUS	205			✔		2020-11-23	View Edit Delete Clone
☐	Piotr Kije…	ASUS_Merlin_Koolshare_r…	RT-AX88U	router	ASUS	ASUS	206			✔		2020-11-23	View Edit Delete Clone

Device Identification - Scan rules

- Rule syntax

```
( boolean expression ) -> statement(s)
```

- Rule operators

Name	Operation
and	boolean and
or	boolean or
=	case sensitive string equality
!=	case sensitive string inequality
=~	regex match
!~	regex difference
:=	assignment

Device Identification - Popular matched responses

- SSL Common Names & Organization Names
- HTML body content
- HTTP server names
- HTTP cookies
- SNMP sysdesc, sysname
- FTP, TELNET, SSH banners
- … many more!

Example fingerprinting rule - iRobot Roomba

```
(issuer_common_name =~ /^Roomba/ and
issuer_organization_name = "iRobot")
-> tag := "iot", device_type :=
"home-appliance", device_vendor :=
"iRobot", device_model := "Roomba",
device_sector := "consumer"
```

SHADOWSERVER VARIOT

Device Identification - iRobot Roomba (2022-06-21)

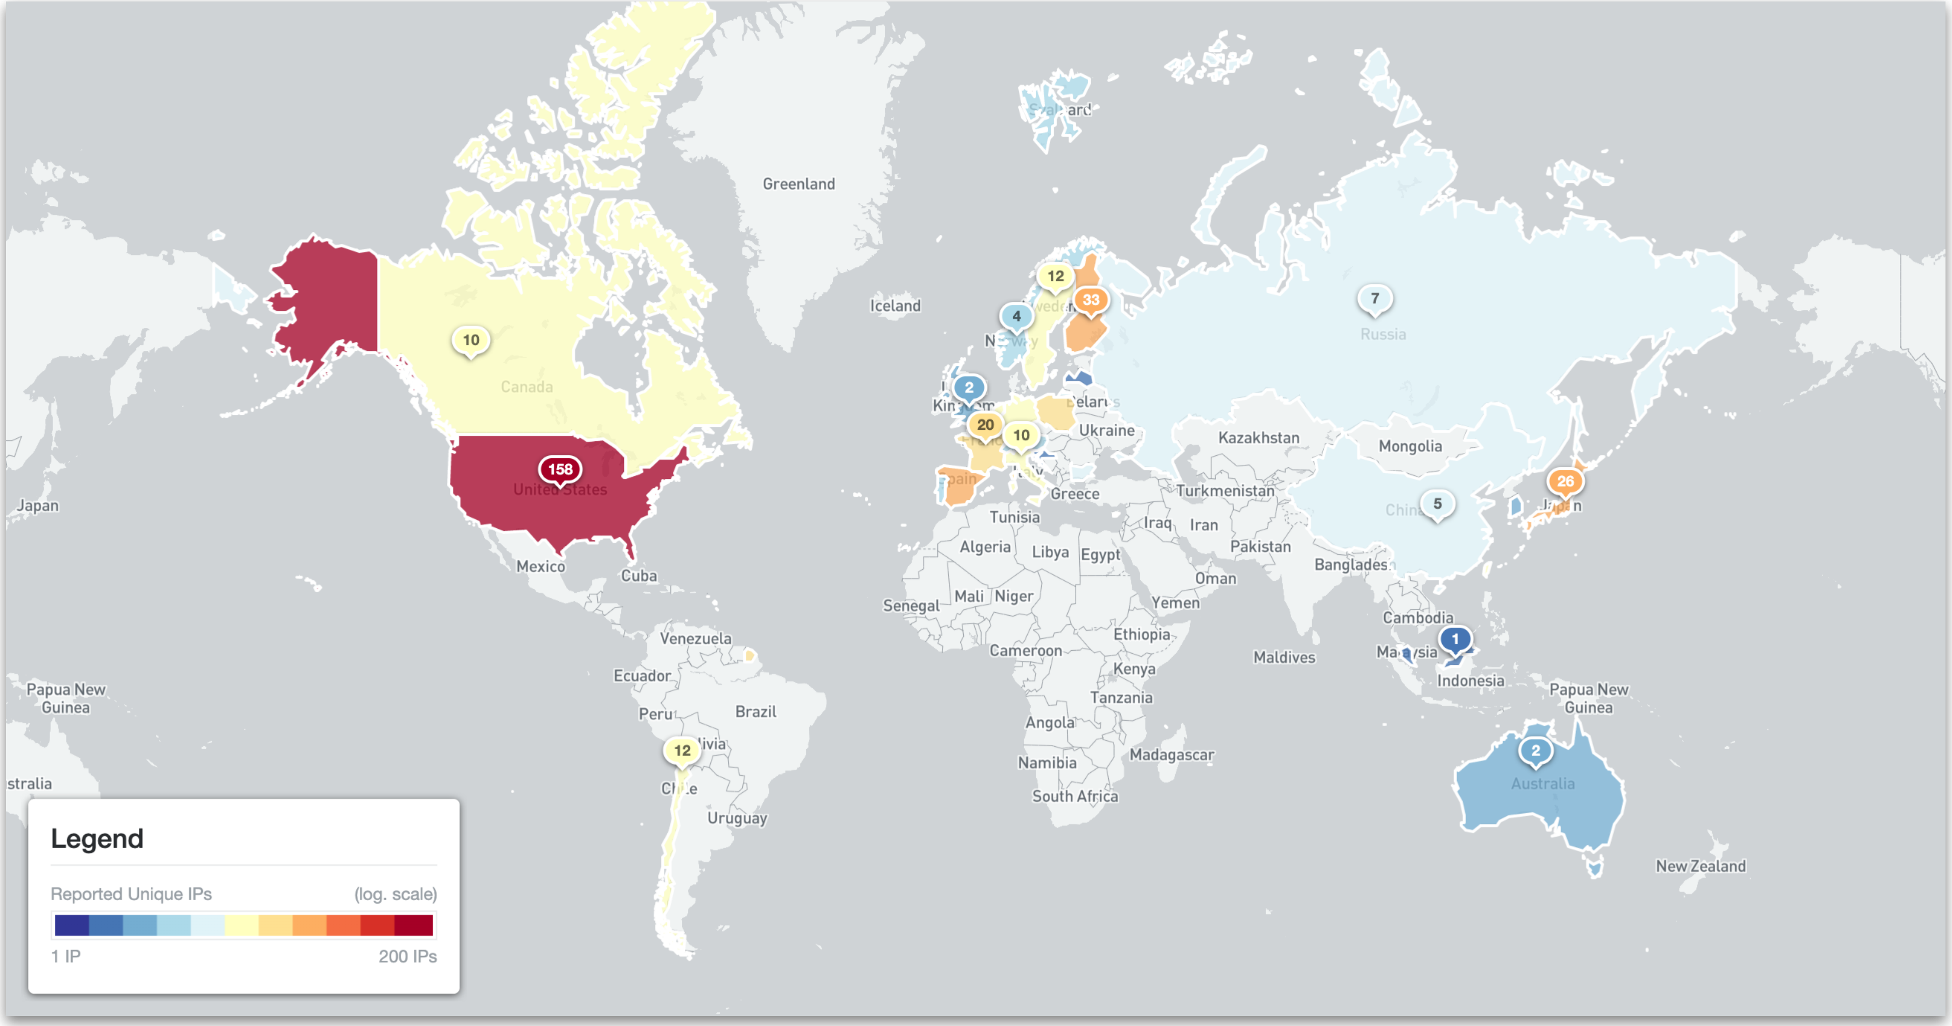

~ 400 devices

Device Identification - Philips HUE (2022-06-21)

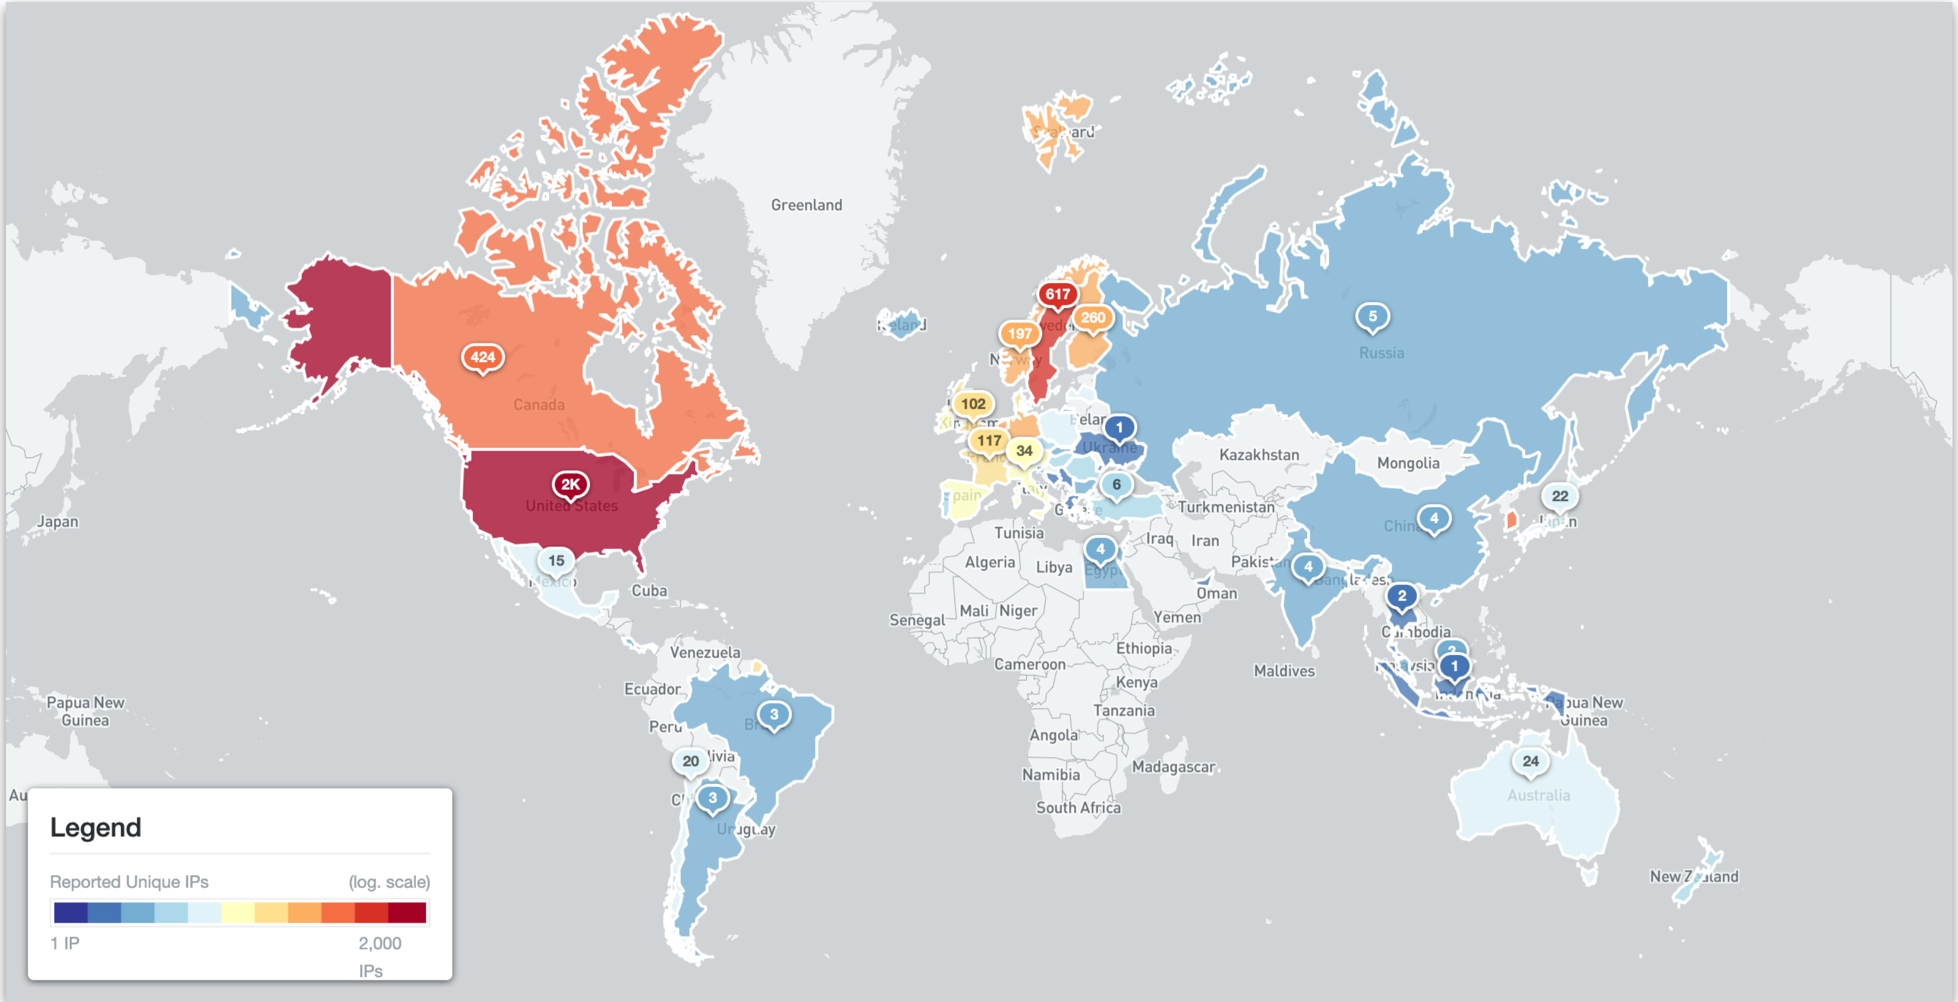

SHADOWSERVER

~ 5300 devices

VARIOT

Device Identification - Siemens SIMATIC S7-300

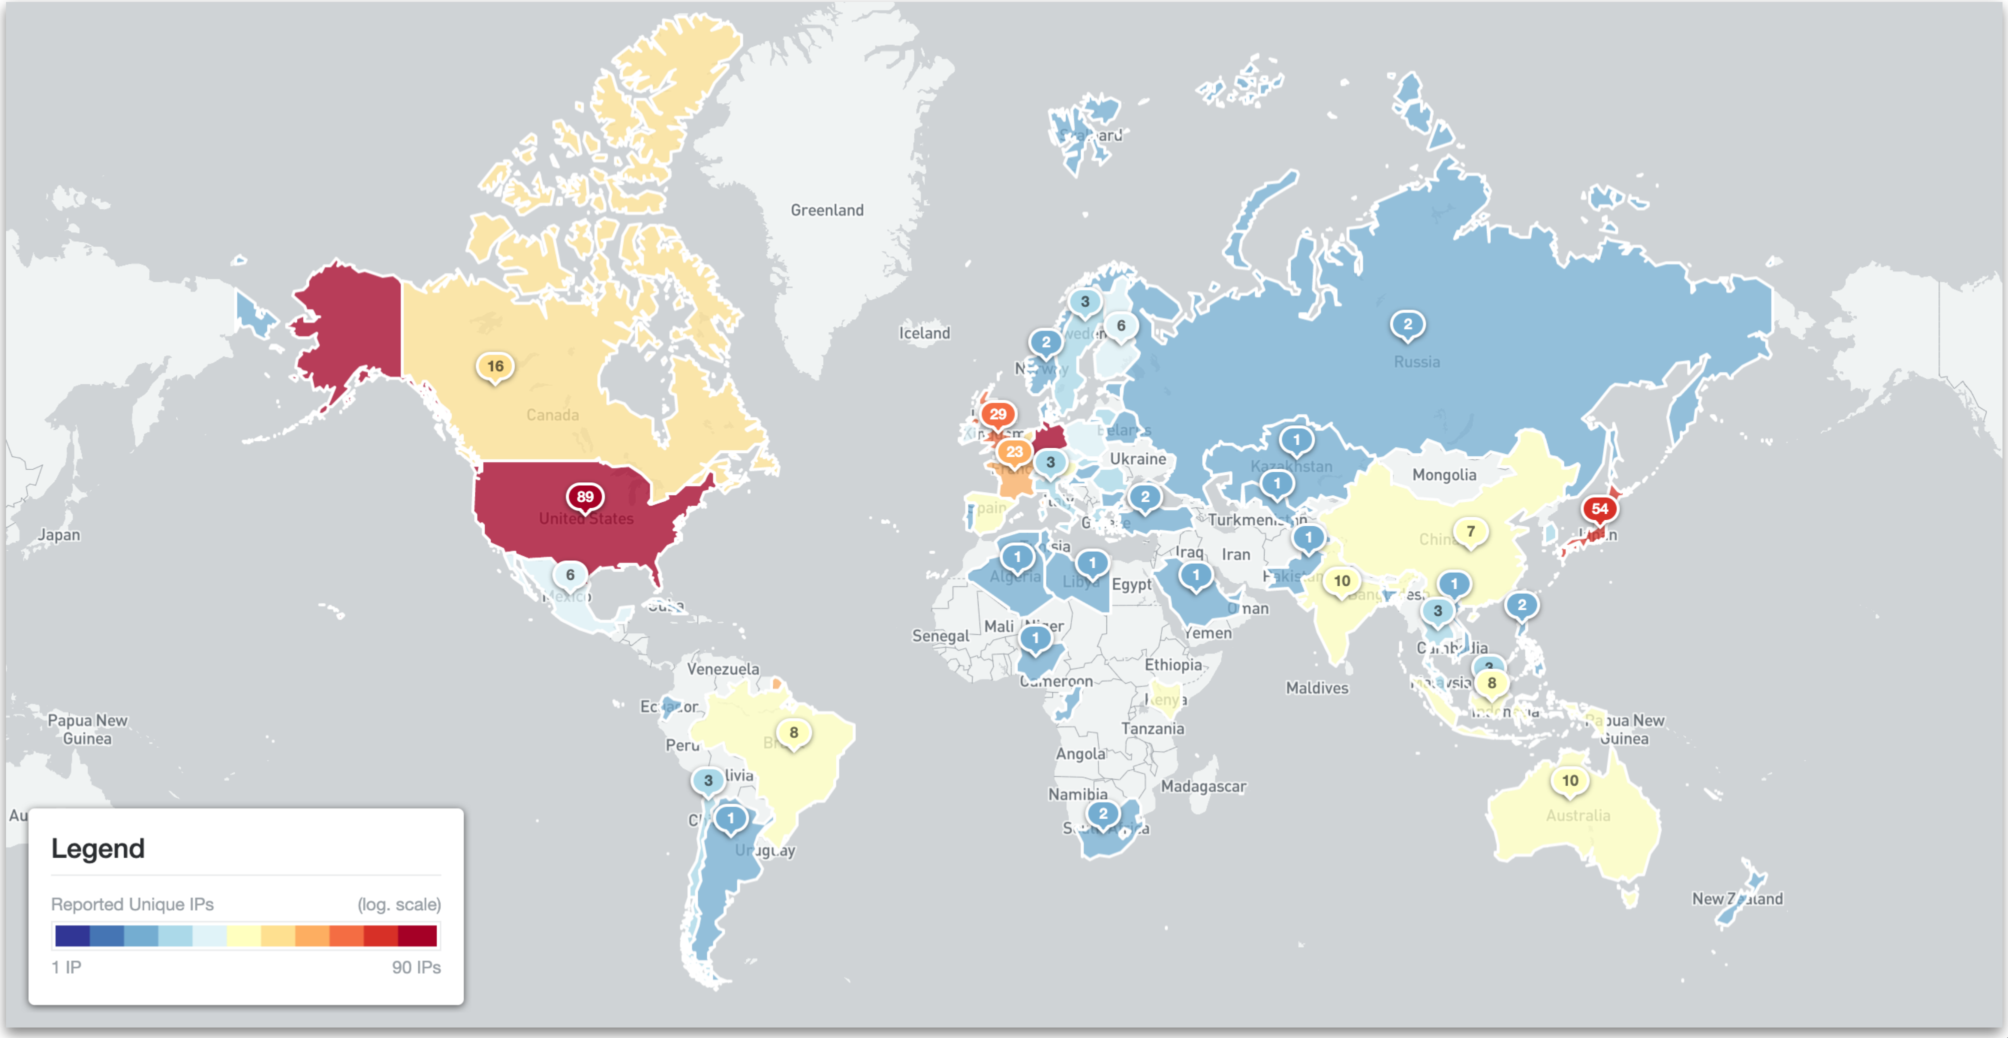

SHADOWSERVER

~ 500 devices (based only on non-native ICS scans)

VARIOT

Device Identification - Siemens SIMATIC S7-300 (2022-06-21)

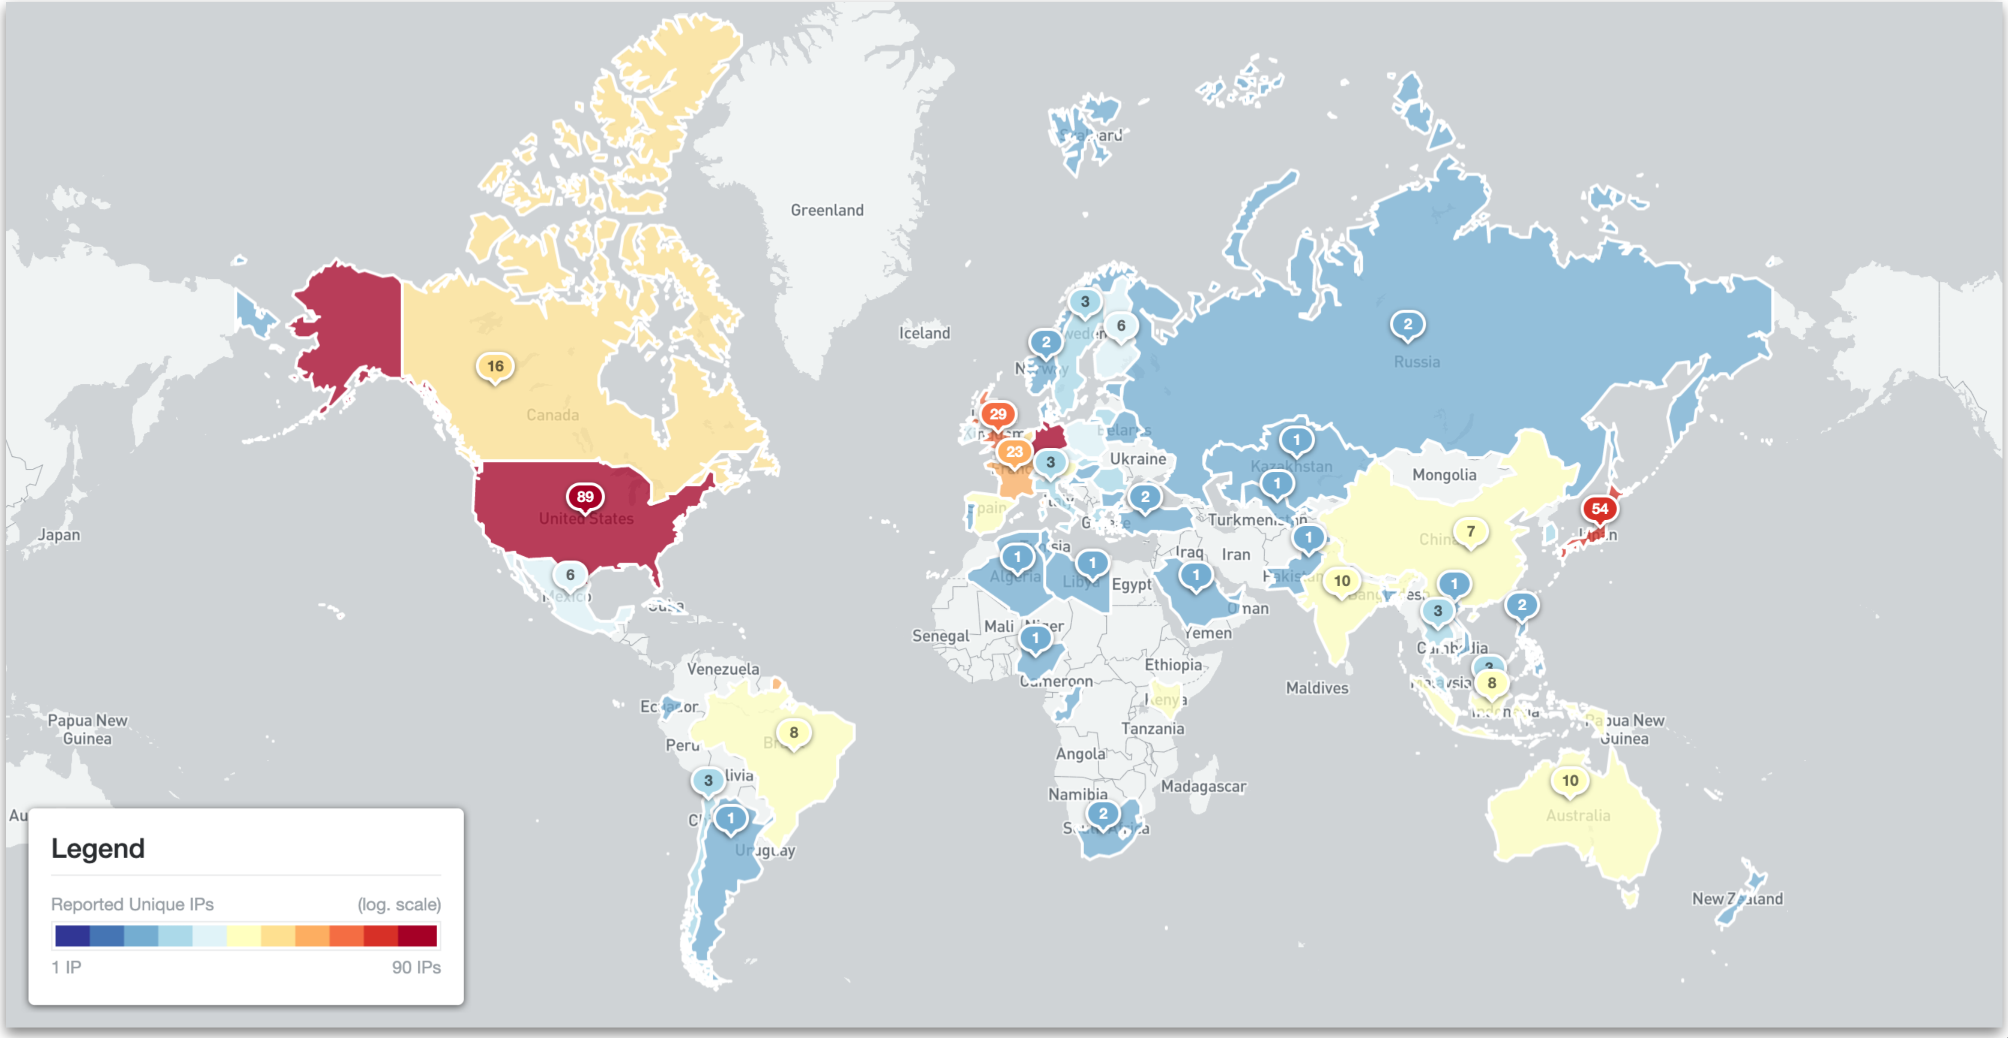

~ 500 devices (based only on non-native ICS scans)

Device Identification - Mikrotik (2022-06-21)

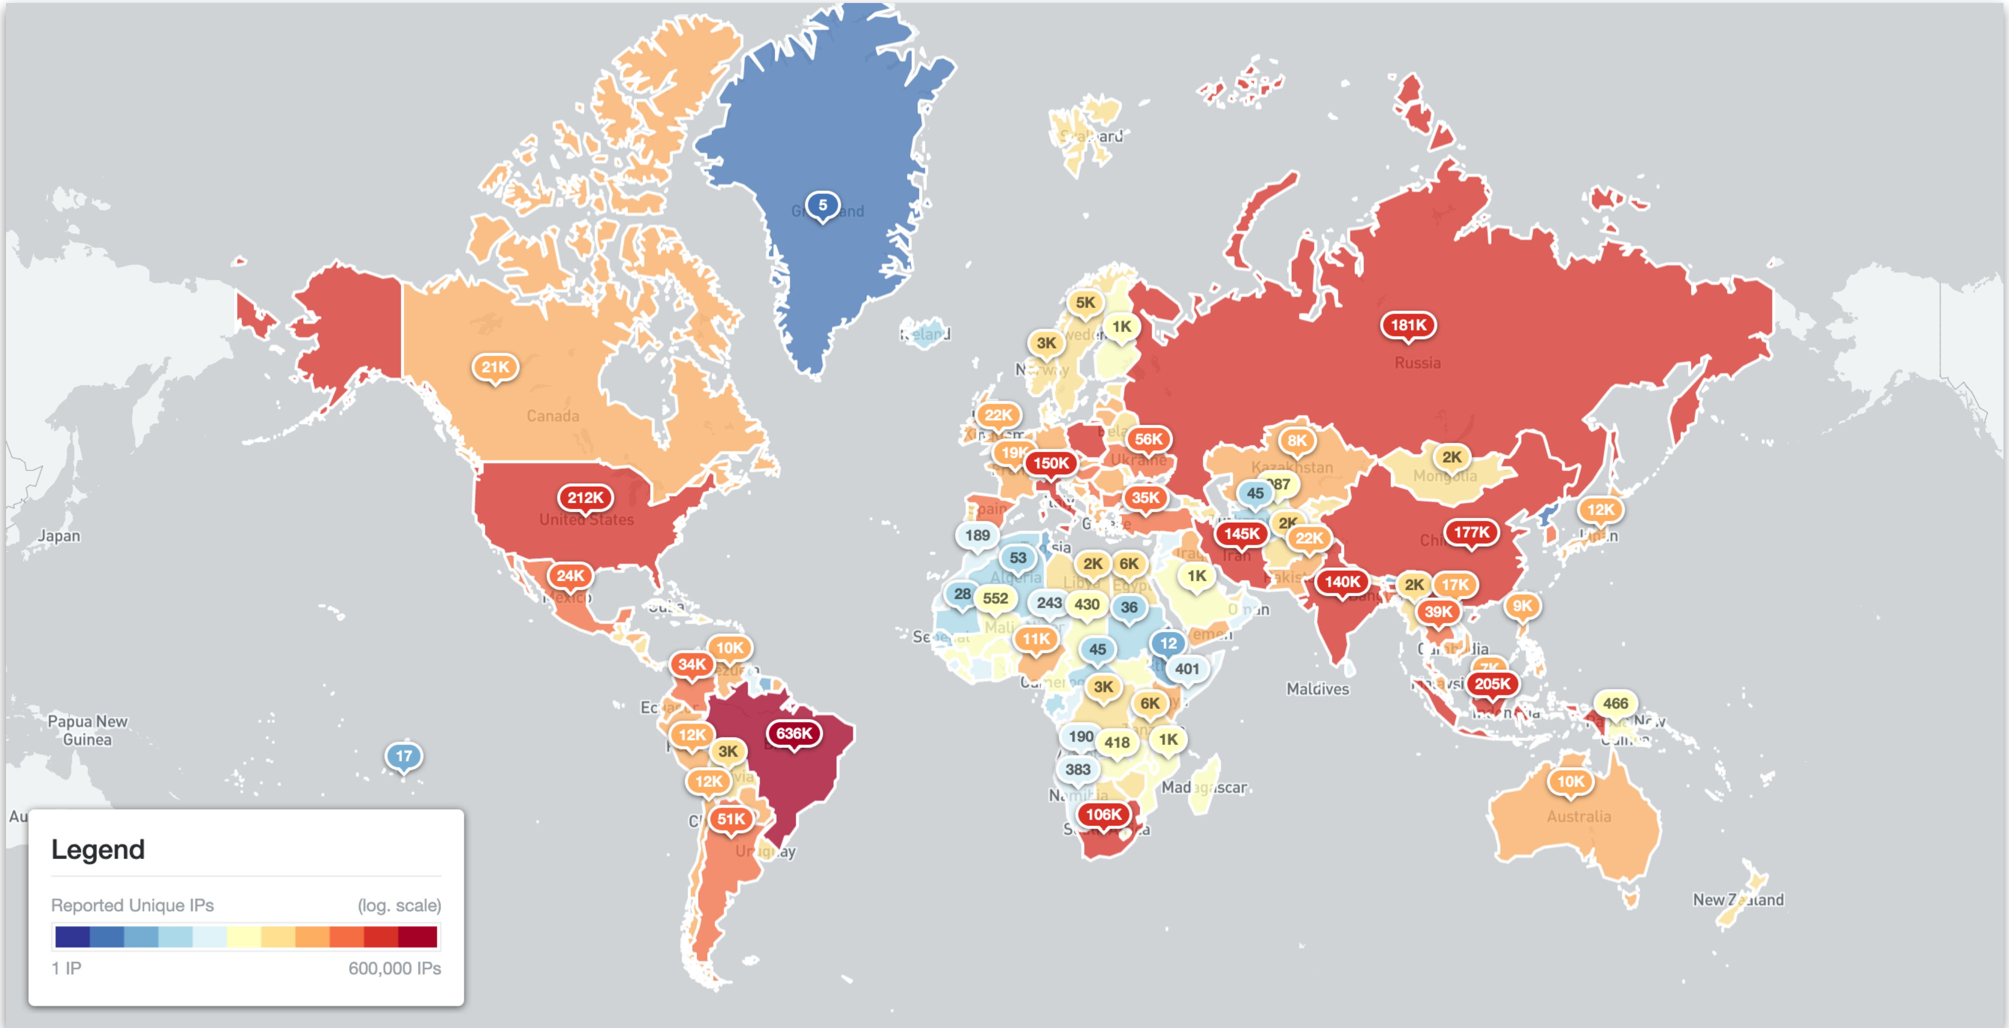

~ 3 200 000 devices

Device Identification - Fortinet (2022-06-21)

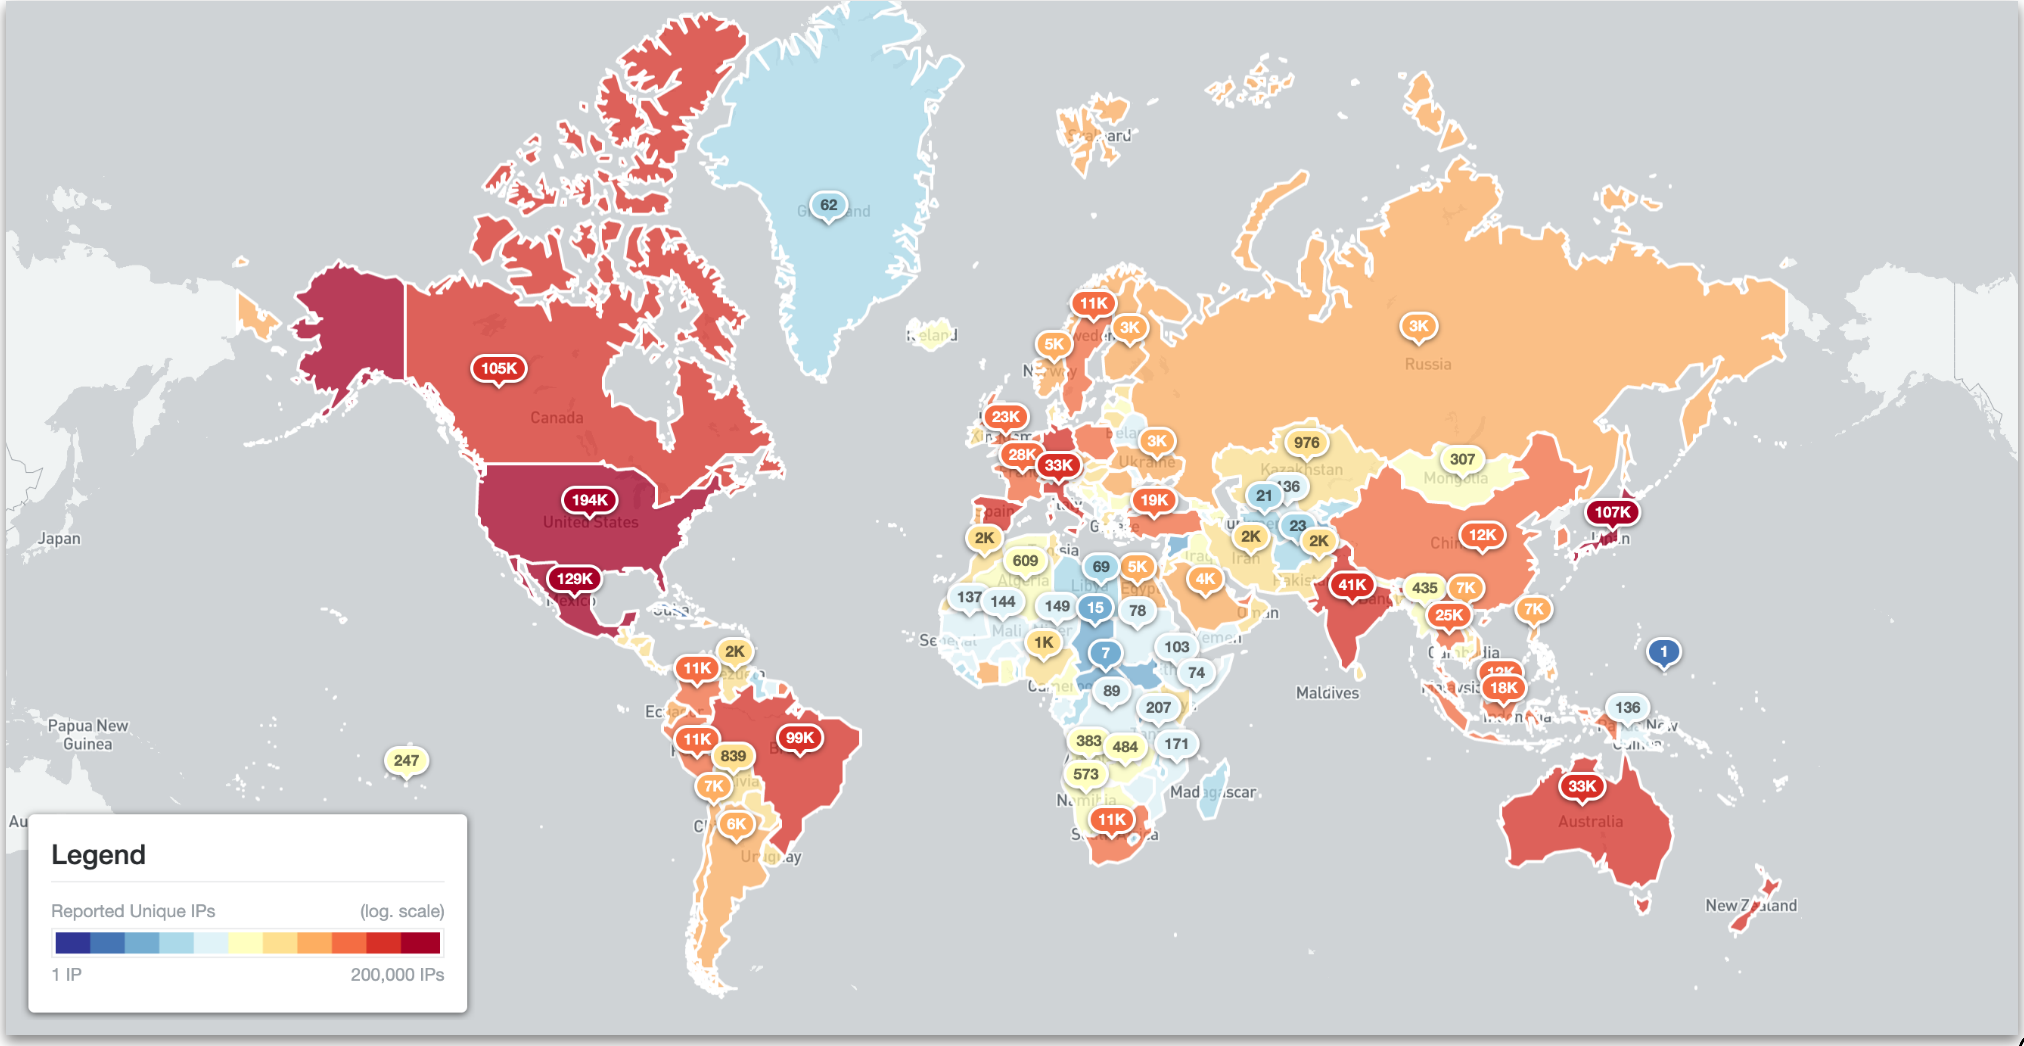

SHADOWSERVER

~ 1 400 000 devices

VARIOT

Devices identified by country (2022-06-21)

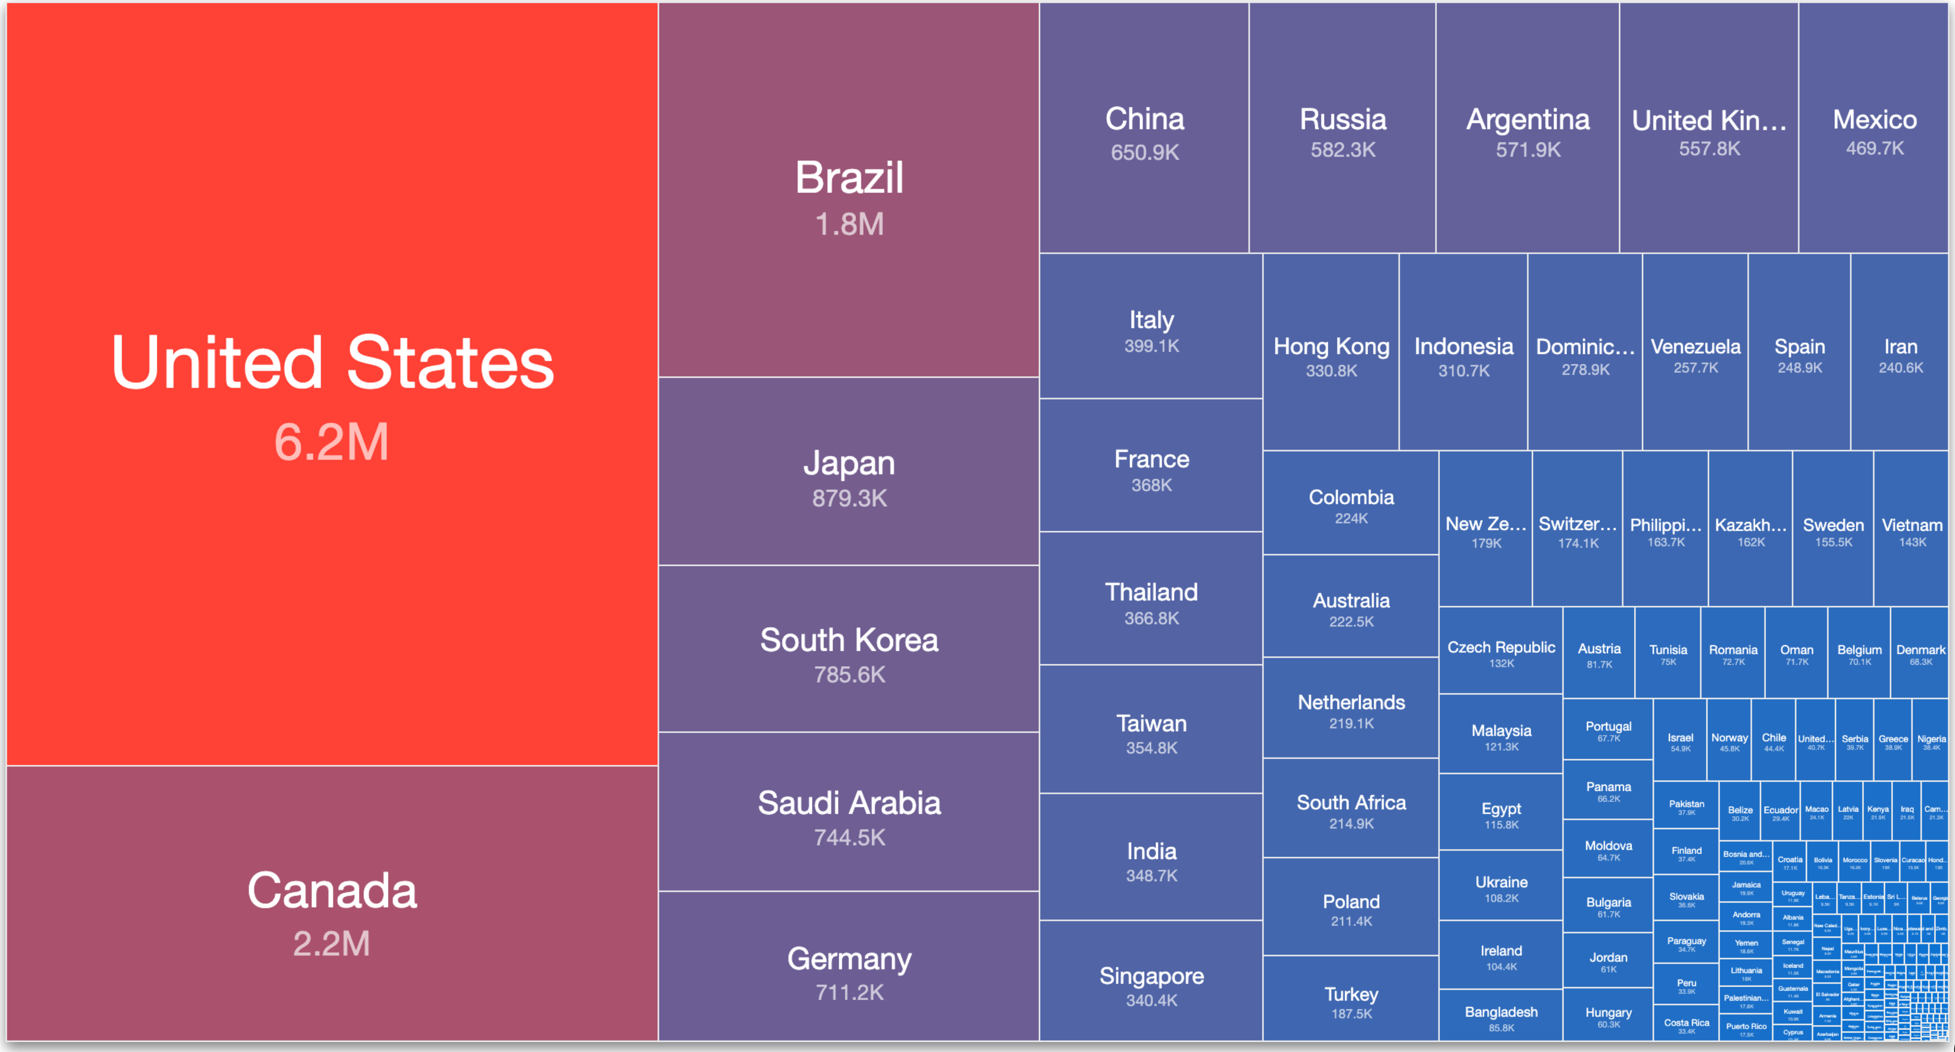

SHADOWSERVER

(Excluding desktop/servers & web servers)

VARIoT

Device Identification - Vendors (2022-06-21)

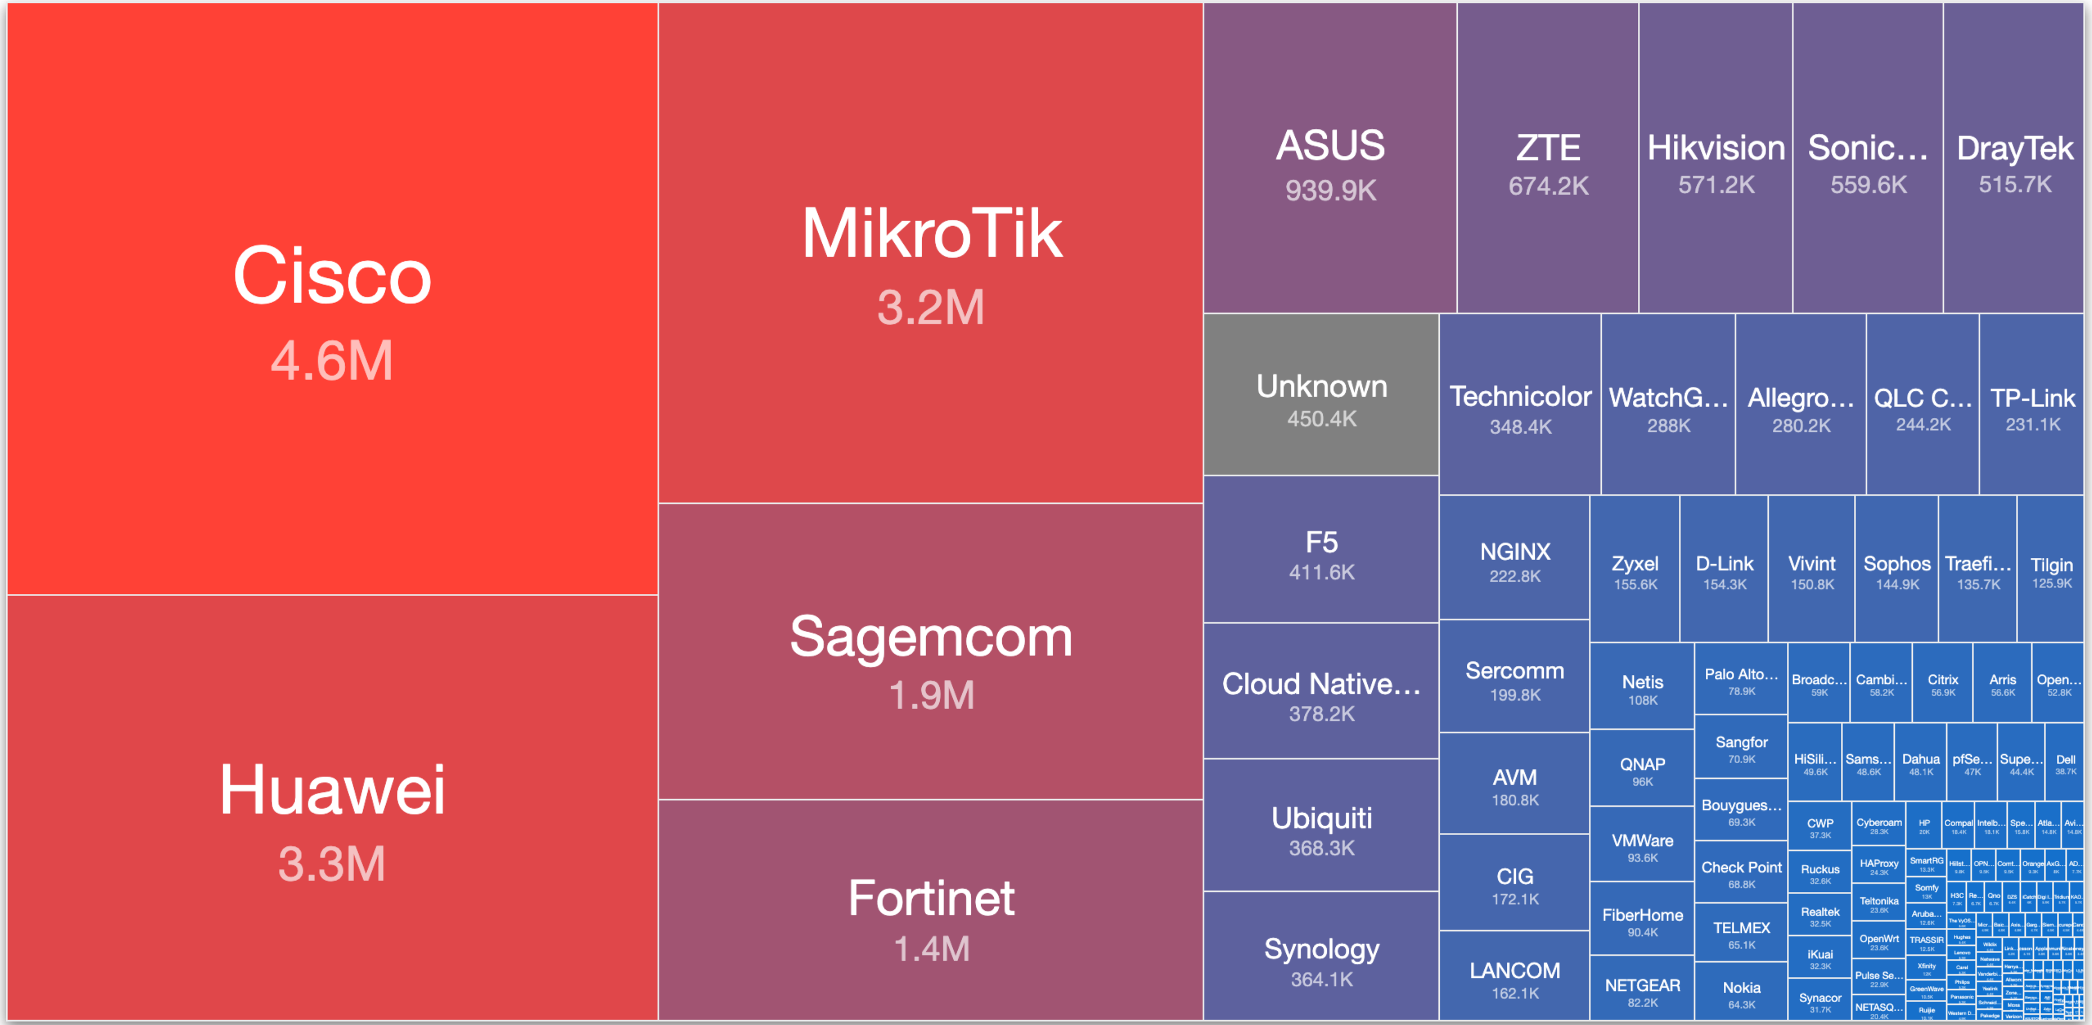

(Excluding desktop/servers & web servers)

HaDEA CEF - VARIoT Project

- July 2019 - Oct 2022
- Shadowserver role: focused on improving
  - scanning of IoT devices
  - observations of IoT attacks
  - collection & analysis of IoT malware
  - sharing of statistics as open data
  - https://variot.eu

Co-financed by the Connecting Europe Facility of the European Union

Subscribe to free daily threat feeds!

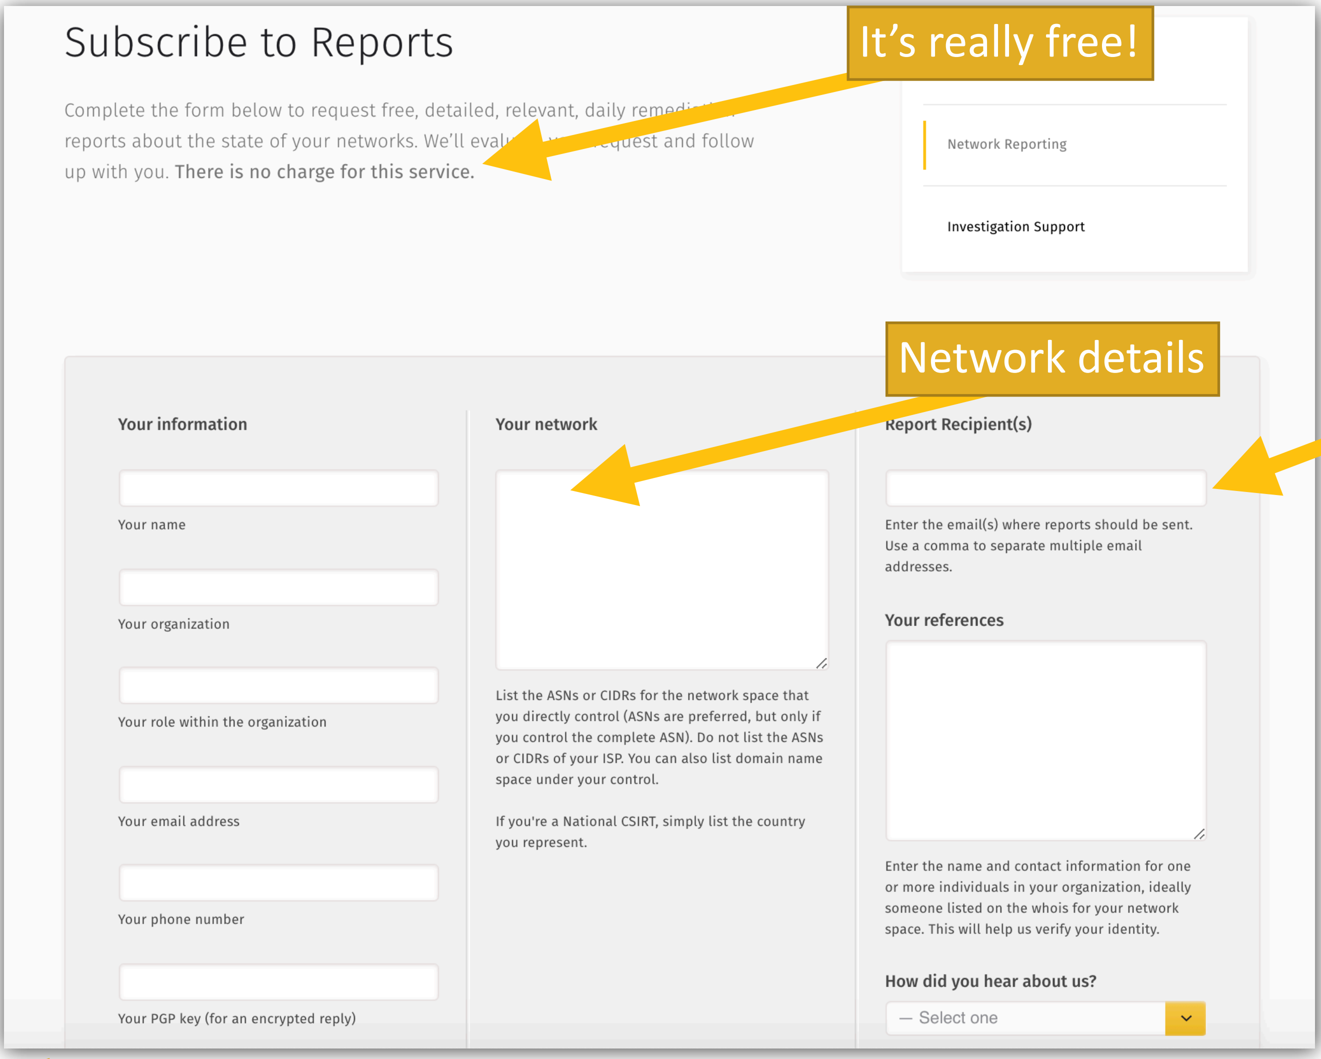
Subscribe to Reports

Complete the form below to request free, detailed, relevant, daily remedi
reports about the state of your networks. We'll evalu request and follow
up with you. **There is no charge for this service.**

Network Reporting

Investigation Support

Your information

Your name

Your organization

Your role within the organization

Your email address

Your phone number

Your PGP key (for an encrypted reply)

Your network

List the ASNs or CIDRs for the network space that you directly control (ASNs are preferred, but only if you control the complete ASN). Do not list the ASNs or CIDRs of your ISP. You can also list domain name space under your control.

If you're a National CSIRT, simply list the country you represent.

Report Recipient(s)

Enter the email(s) where reports should be sent. Use a comma to separate multiple email addresses.

Your references

Enter the name and contact information for one or more individuals in your organization, ideally someone listed on the whois for your network space. This will help us verify your identity.

How did you hear about us?

— Select one

It's really free!

Network details

Email address where reports or download links will be sent

SHADOWSERVER

https://www.shadowserver.org/what-we-do/network-reporting/get-reports/

Questions?

SHADOWSERVER

Lighting the way to a more secure Internet

@shadowserver

decoster@shadowserver.org, piotr@shadowserver.org