Measuring the Deployment of Network Censorship Filters at Global Scale

Ram Sundara Raman[1], Adrian Stoll[1], Jakub Dalek[2], Armin Sarabi[1], Reethika Ramesh[1], Will Scott[3], Roya Ensafi[1]

University of Michigan[1], The Citizen Lab[2], Independent[3]

24 February 2020

NDSS SYMPOSIUM/2020

Content Filtering Technologies

- Filters, DPIs, middleboxes
- **Dual Use Technology**
  - Intended use - Security
  - Side effect - Censorship, surveillance
- **Commoditization of filters** - High availability, low cost, and advanced features
- Very little, but important, information on use of filters

Netsweeper and Citizen Lab

- **Netsweeper -** Canadian filter vendor - Provides carrier grade filtering, dynamic categorization of websites
- **Citizen Lab** conducted investigations of use of Netsweeper products over several years
- **“Alternative Lifestyles”** category used by UAE, others to block LGBTQ content
- Netsweeper **removed the option** to block category

VICE

Canadian Internet Filtering Company Says It's Stopped 'Alternative Lifestyles' Censorship

The UAE was found to be blocking LGBTQ content using a pre-set category in Netsweeper's software. Amid pressure from rights groups, the company says it's disabled that category.

By **Jordan Pearson**

Jan 21 2019, 12:25pm Share Tweet Snap

Auditing filters can drive change!

Proliferation of Filters

FORTINET

FortiGuard Labs

cisco Cisco Umbrella

netsweeper

paloalto NETWORKS

WatchGuard

Freedom House

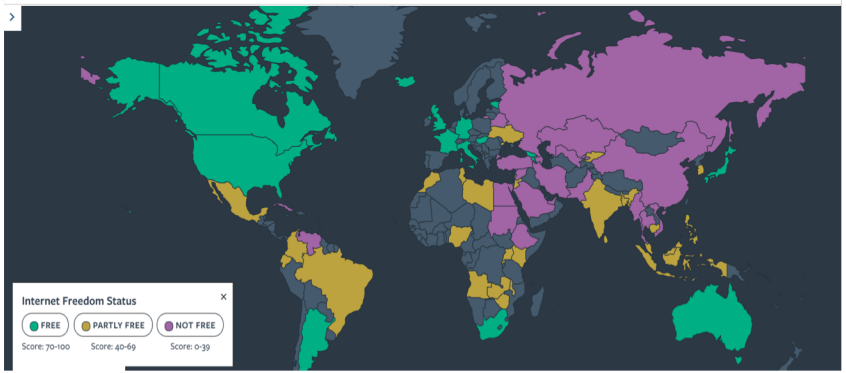

Previous Work

- Biased towards few, well-known filters
- Significant manual effort
  - Physical access
  - In-country collaborators

Blockpages

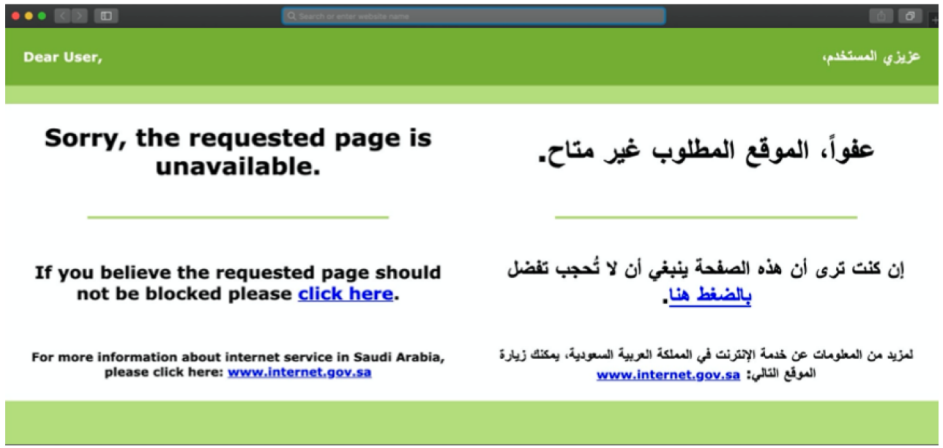

- Filters respond with blockpages
- Rich with information
  - Trademark of the manufacturing vendor
  - Identity of the deploying actor
- **Use blockpages to identify censorship filter deployments**
- Identification using blockpages is consistent and scalable

Objectives

Data Collection

Collect many blockpages from filter deployments

Data Analysis

Identify filters from blockpages

Data Collection

Collect the most comprehensive database of filter blockpages

Data Collection

Censorship measurement techniques frequently observe blockpages

Data Collection

Censorship measurement techniques frequently observe blockpages

OONI

Volunteer measurement
https://ooni.org/

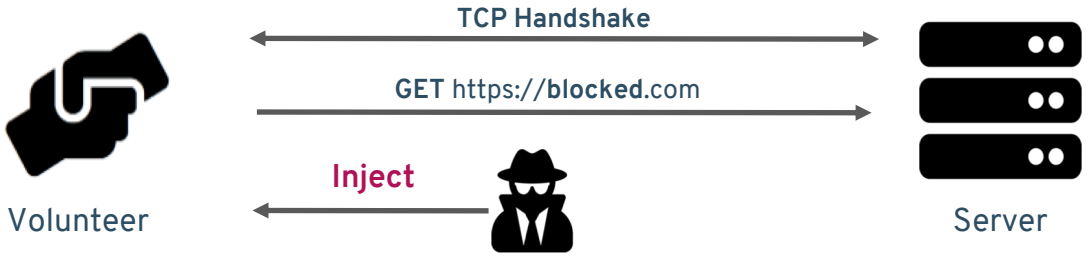

Challenges

- Limited scale and ethical constraints

Data Collection

Censorship measurement techniques frequently observe blockpages

OONI

https://ooni.org

Quack

Remote measurement
VanderSloot et al. [USENIX 2018]

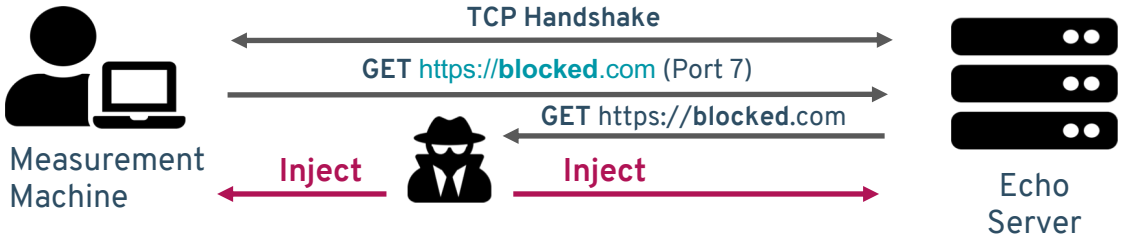

Challenges

- Cannot detect filters on common Port 80/443

Data Collection

Censorship measurement techniques frequently observe blockpages

OONI

https://ooni.org

Quack

Remote measurement

Hyperquack

New remote measurement

- Novel remote measurement technique
- **Web servers** running on ports 80 and 443
- **Idea: Responses from web server when requesting a domain not hosted on the server is predictable**

Hyperquack

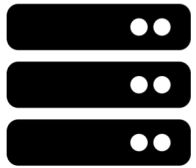

46.43.36.222

Hyperquack

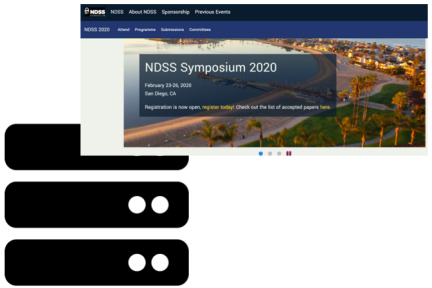

46.43.36.222

Hyperquack

Measurement Machine

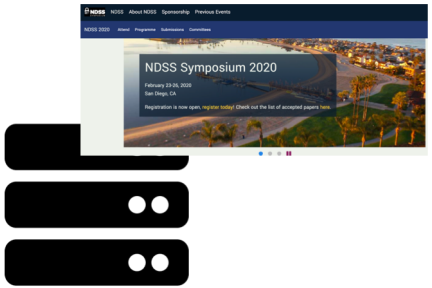

46.43.36.222

Hyperquack

Measurement Machine

GET https://www.ndss-symposium.org

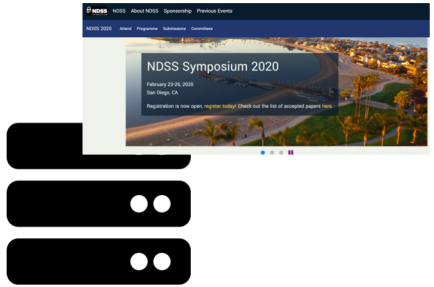

46.43.36.222

Hyperquack

Measurement Machine

GET https://www.ndss-symposium.org

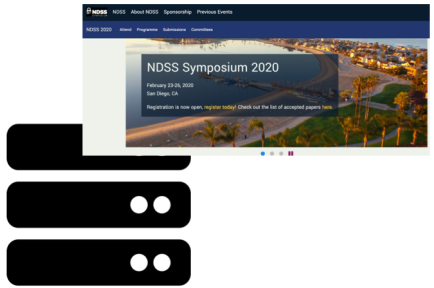

46.43.36.222

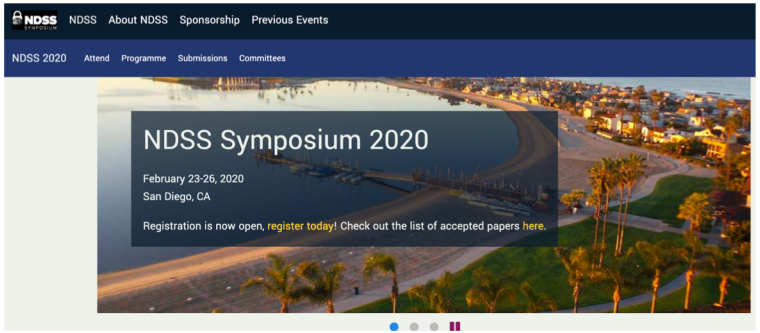

Hyperquack

Measurement Machine

GET https://www.usenix.org

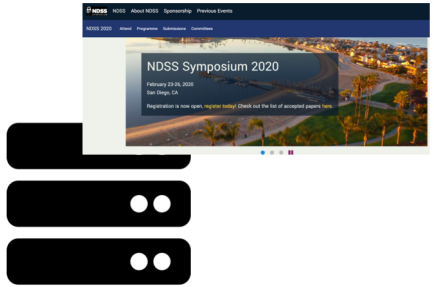

46.43.36.222

Hyperquack

Measurement Machine

GET https://www.usenix.org

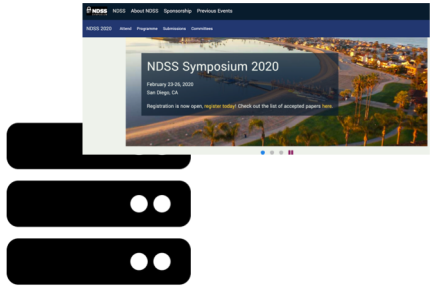

46.43.36.222

Moved Permanently

The document has moved here.

Apache/2.4.25 (Debian) Server at www.usenix.org Port 443

Hyperquack

Measurement Machine

GET https://www.sigsac.org

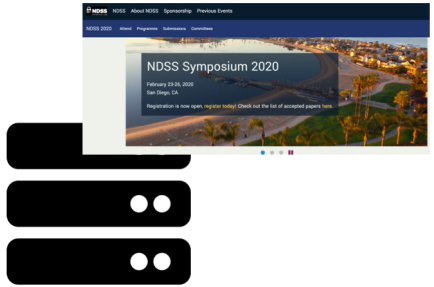

46.43.36.222

Hyperquack

Measurement Machine

GET https://www.sigsac.org

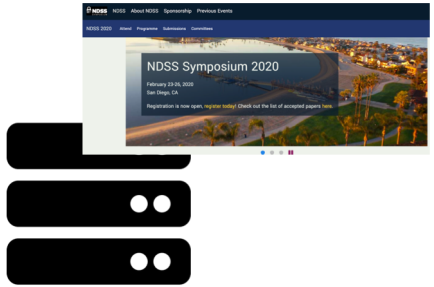

46.43.36.222

Moved Permanently

The document has moved here.

Apache/2.4.25 (Debian) Server at www.sigsac.org Port 443

Hyperquack

Measurement Machine

GET https://www.sigsac.org

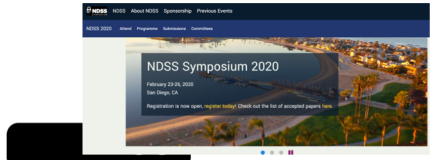

46.43.36.222

Moved Permanently

The document has moved here.

Apache/2.4.25 (Debian) Server at www.sigsac.org Port 443

Canonical Templates

```
<h1>Moved Permanently</h1>
<p>The document has moved
<a href="https://www.consumersinternational.org/
what-we-do/digital/internet-of-things/connect-smart/">
here</a>.</p><hr>
<address>Apache/2.4.25 (Debian)
Server at www.sigsac.org Port 443</address>
```

- Request several bogus but benign domain patterns

  (<www>.example1298.<com>)
- From the response, remove commonly changing elements e.g. date, domain
- If response for all tests match, save as **canonical template**

Censorship Detection

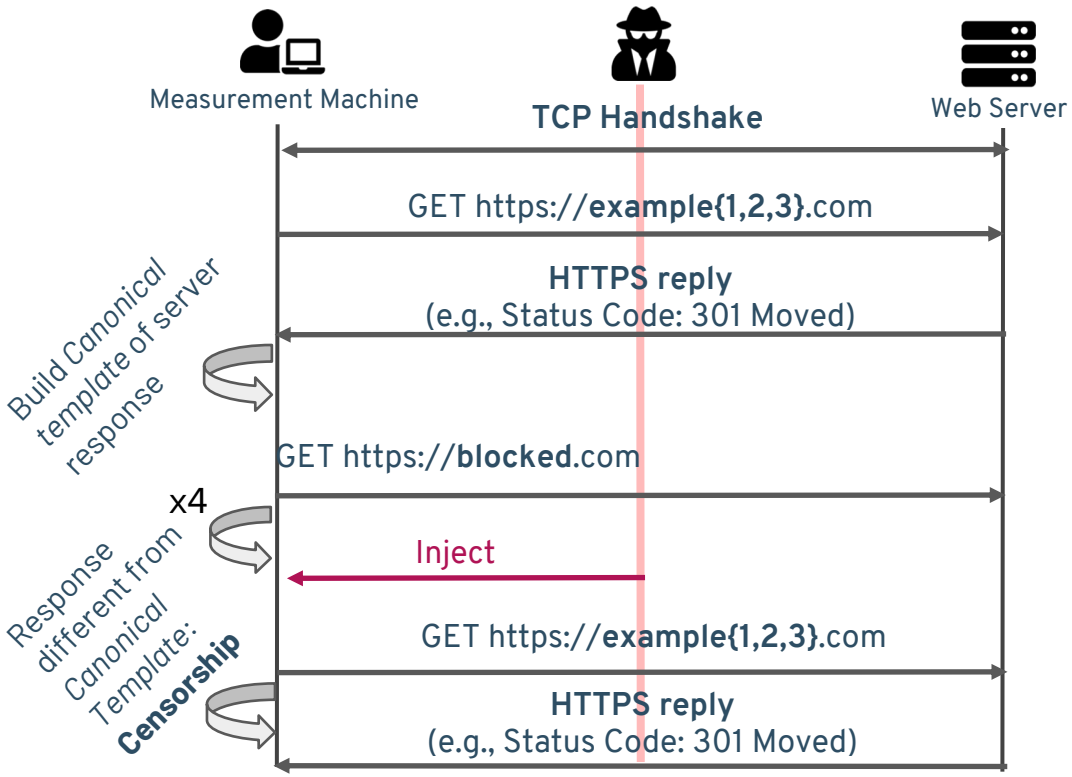

- Send HTTP(S) GET requests for sensitive keywords
- If response different from canonical template, then there is censorship
- Control tests both before and after to ensure consistency

Censorship Detection

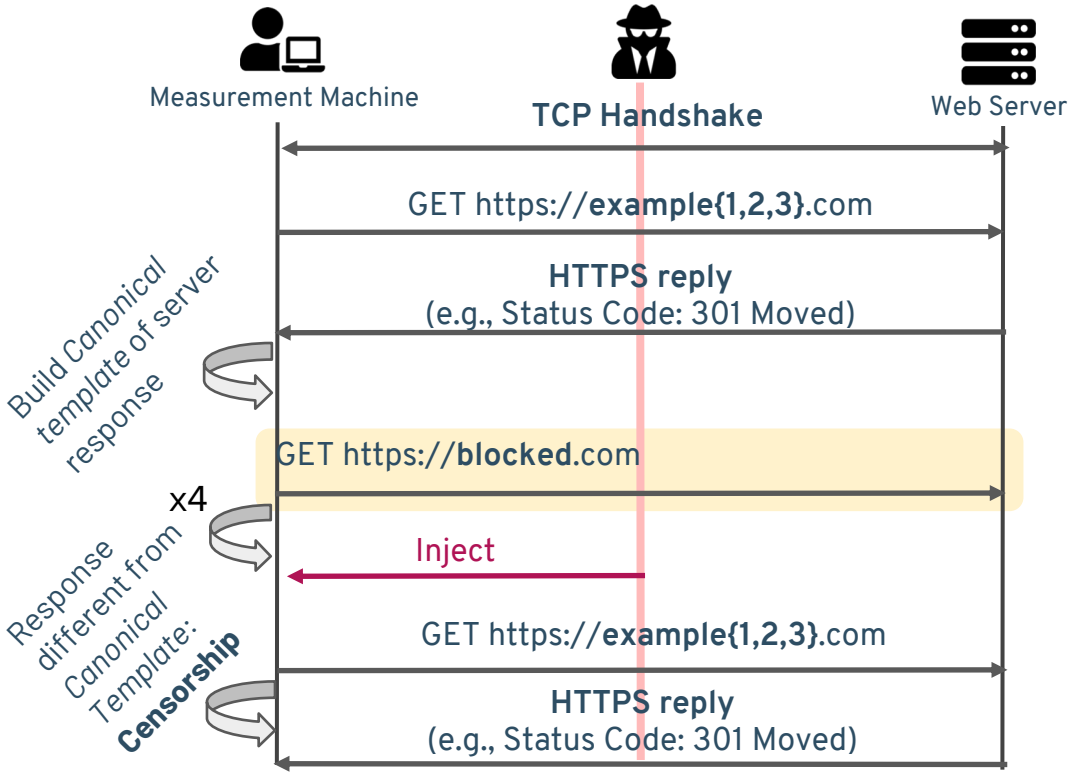

- Send HTTP(S) GET requests for sensitive keywords
- If response different from canonical template, then there is censorship
- Control tests both before and after to ensure consistency

Censorship Detection

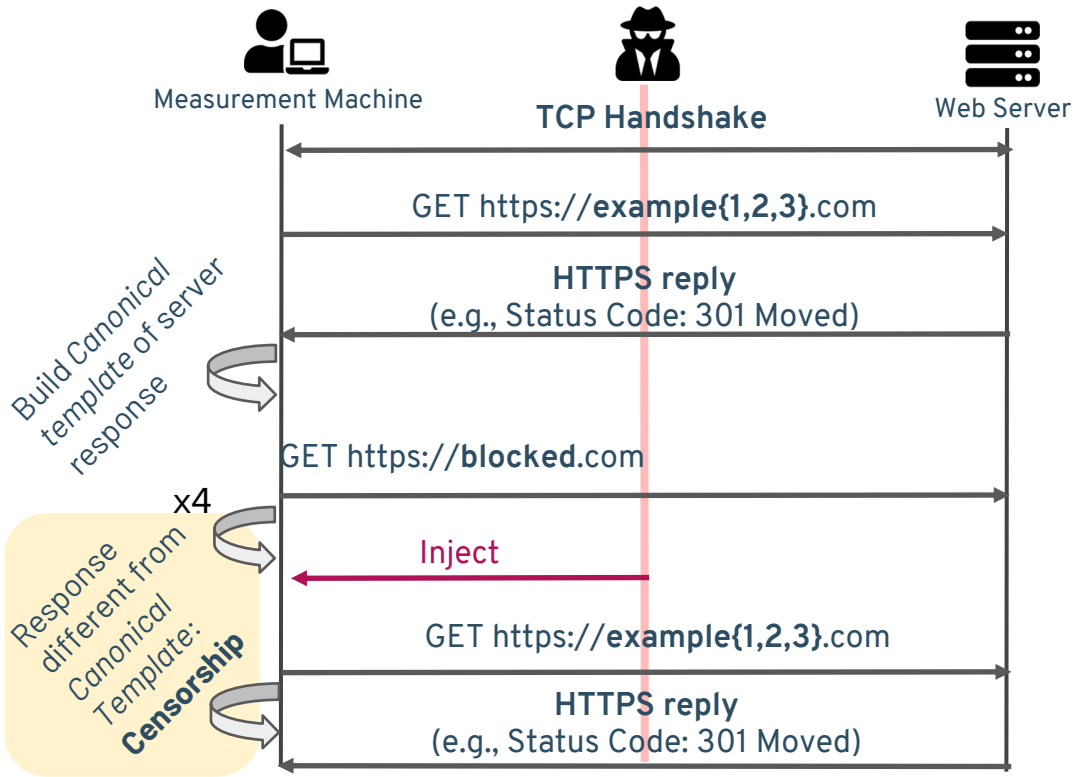

- Send HTTP(S) GET requests for sensitive keywords
- If response different from canonical template, then there is censorship
- Control tests both before and after to ensure consistency

Censorship Detection

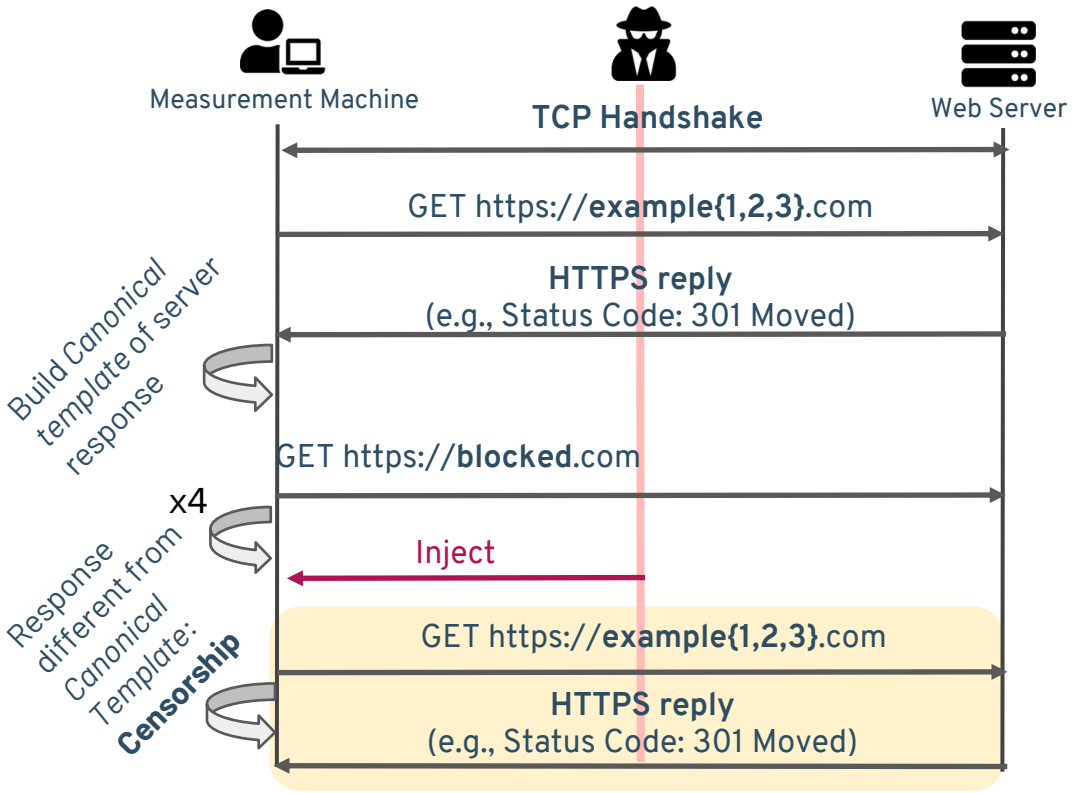

- Send HTTP(S) GET requests for sensitive keywords
- If response different from canonical template, then there is censorship
- Control tests both before and after to ensure consistency

Hyperquack increases scale to millions of vantage points!

53 million public HTTP hosts

Source - censys.io

Vantage Point Selection

- We use **infrastructural servers** to reduce risk
- **PeeringDB** - list of official websites of Internet service providers
- Use servers hosting the website for measurement ~10,000

Vantage Point Selection

- We use **infrastructural servers** to reduce risk
- **PeeringDB** - list of official websites of Internet service providers
- Use servers hosting the website for measurement ~10,000

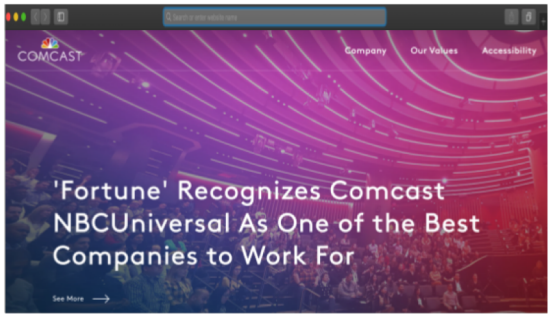

https://corporate.comcast.com/

Vantage Point Selection

- We use **infrastructural servers** to reduce risk
- **PeeringDB** - list of official websites of Internet service providers
- Use servers hosting the website for measurement ~10,000

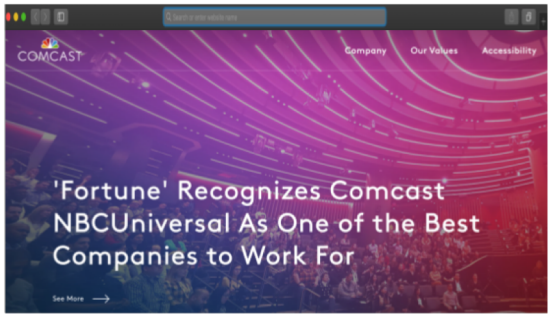

https://corporate.comcast.com/

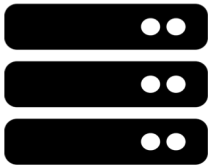

23.219.228.121

Ethics

- Followed all the ethical recommendations made in Quack
- Made it clear that we are running measurements on our website
- Rate limit and close connections
- Make only one measurement at a time to a server
- OONI obtains informed consent

Measurements

- Latitudinal Measurements:
  - 3 weeks in October 2018
  - HyperQuack - 9,223 VPs
  - Quack - 33,602 VPs
  - 18,736 domains - Citizen Lab Test List
  - Added OONI data
- Longitudinal Measurements:
  - HyperQuack and Quack twice a week - November 2018 to January 2019
  - Citizen Lab Global List (~1200 domains) + Alexa Top 1000 domains

Data Analysis

Automate the identification of filters from more than a million disrupted responses

Iterative Classification

- **Insight: Filters often send the same blockpage regardless of the test domain**
- Recursively finds large groups of HTML pages with the same content
- Blockpage clusters are labeled with signatures, a unique subset of the HTML page or header
- Example: <th>Barracuda NextGen Firewall:</th>

Image Clustering

- Cluster pages with **dynamic content** - DBSCAN algorithm
- T**remendously reduce the manual effort** - 1 page in 200 groups

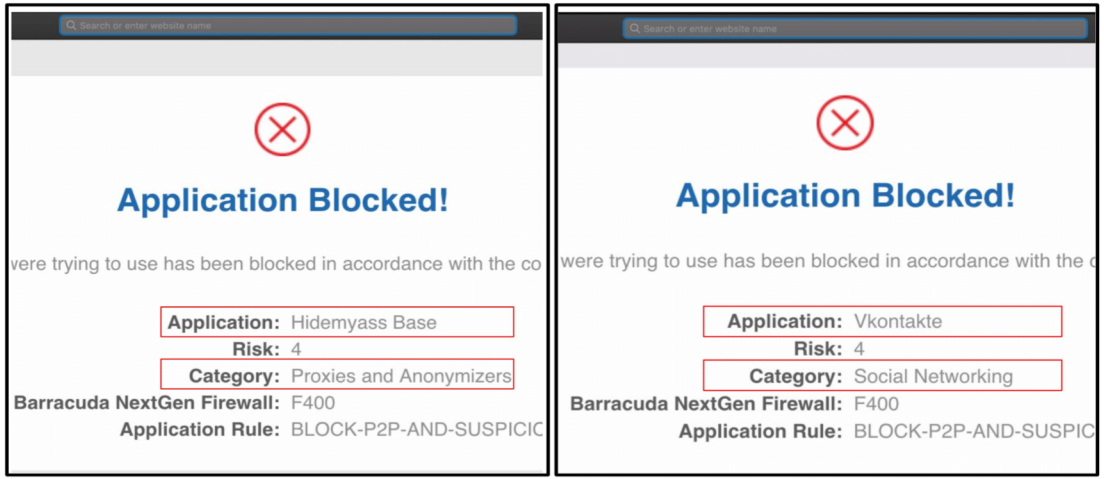

FilterMap

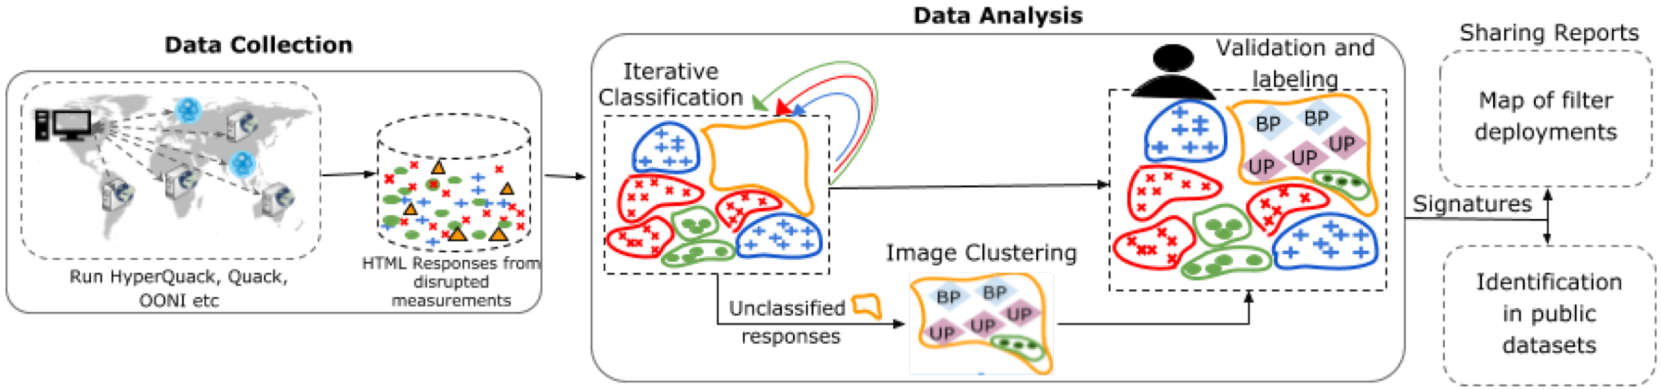

FilterMap enables continuous, sustainable, data-driven view of filter deployment

Results

FilterMap creates a map of filter deployments based on the vantage points measured

FilterMap Results

- FilterMap found **90 blockpage clusters** (Clusters indicate either vendors or actors)
- Filters are deployed in many locations in **103 countries**
- Filter types found - Commercial products, national firewalls, ISP and organizational deployments

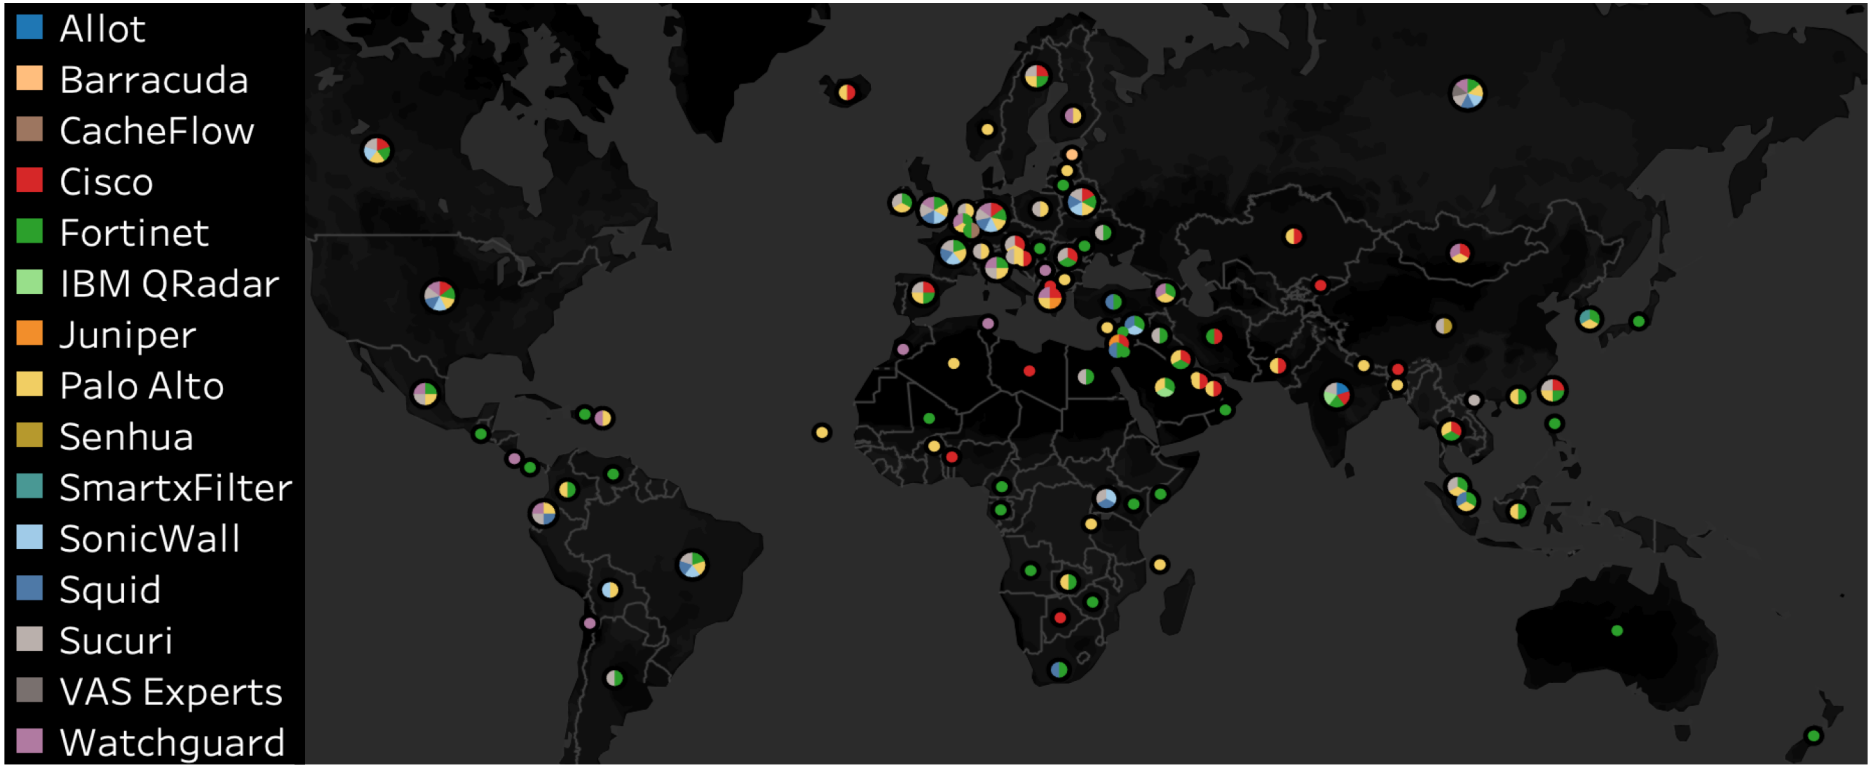

Commercial Filters

Commercial Filters

- **15 commercial filters used in 102 countries**
- Sold by companies in the US
- Filters found in **36 out of 48 countries** labelled as “Not Free” or “Partly Free” by Freedom House
- **Pornography**, **gambling**, **provocative attire and anonymization tools** most commonly blocked

FilterMap Results

- 4 National Firewalls - Iran, Saudi Arabia, Bahrain and South Korea

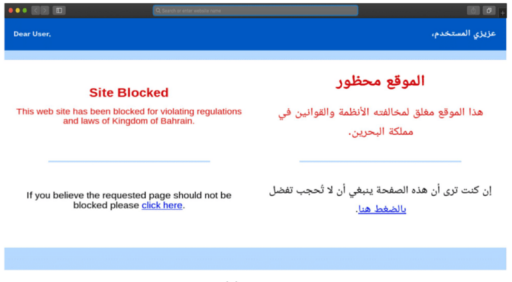

(a)

(b)

(c)

(d)

FilterMap Results

- 4 National Firewalls - Iran, Saudi Arabia, Bahrain and South Korea
- Large number of filters in ISPs, especially in Russia

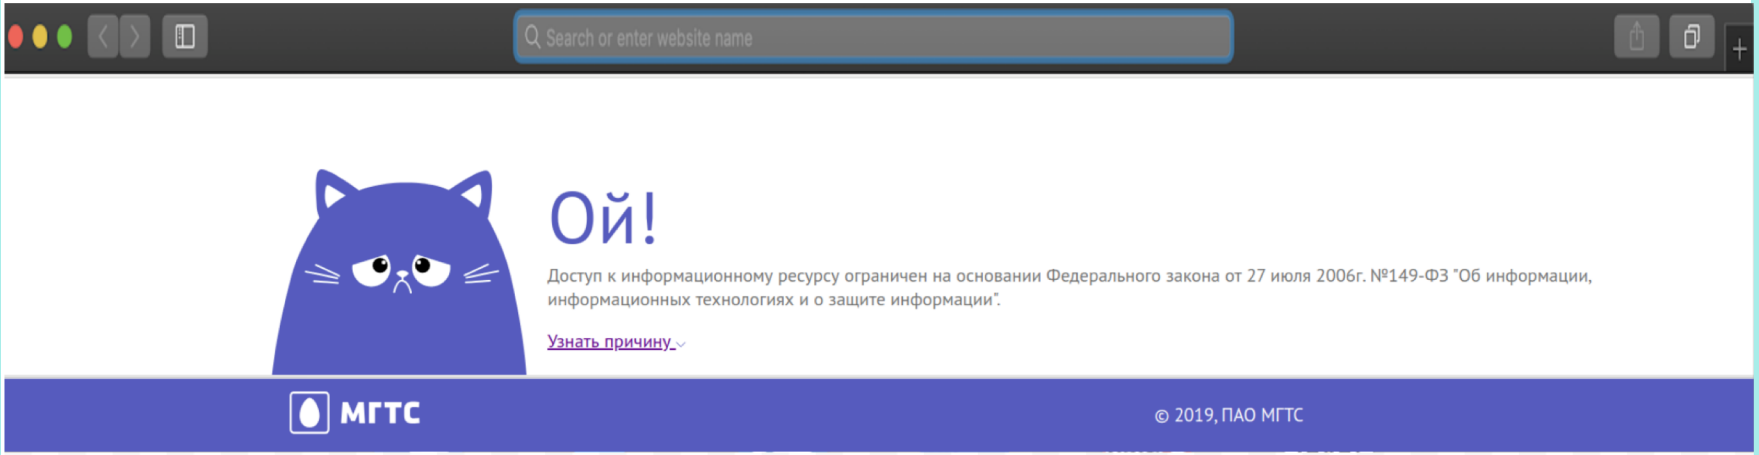

FilterMap Results

- 4 National Firewalls - Iran, Saudi Arabia, Bahrain and South Korea
- Large number of filters in ISPs, especially in Russia
- Of the 90 blockpage clusters -
  - 70 - Latitudinal
  - 20 additional - Longitudinal
- FilterMap can continuously track filter proliferation

Limitations and Future Work

- Blockpages as a source
  - Future work - Certificate, TCP/IP header
- Evasion - Possible but unlikely
- Exact filter location in network is unknown

Implications

- Unrestricted transfer - Easier to deploy and harder to circumvent
- Million-dollar fines and increased regulation
- FilterMap is maintained as source of longitudinal data
- Accountability to filter manufacturers

Summary

- Crucial to collect information about the use of dual-use technologies for censorship
- FilterMap - Framework for semi-automatically measuring filter deployments continuously and sustainably
- Found widespread use of filters for blocking access to content
- Data and Results available at **https://censoredplanet.org/filtermap**

Measuring the Deployment of Network Censorship Filters at Global Scale

Ram Sundara Raman[1], Adrian Stoll[1], Jakub Dalek[2], Armin Sarabi[1], Reethika Ramesh[1], Will Scott[3], Roya Ensafi[1]

University of Michigan[1], The Citizen Lab[2], Independent[3]

Thank you

https://censoredplanet.org/filtermap

Backup Slides

Netsweeper

Canadian Filter Vendor

Summary of Data Collection Techniques

	Pros	Cons
OONI	In-depth measurements close to the user (Volunteer -> Site)	Scale, Continuity, Ethics
Quack	Scale - 33,000 vantage points	Only Port 7 measurements
Hyperquack	Port 80 and Port 443 measurements	Can only detect filter if it acts in both directions (MM -> VP)

Blockpages as Identifiers

- Goes against the purpose of the censor to remove blockpages
- Vendors rarely have any incentive to remove trademarks
- Modified blockpages can still be detected
- Identification using blockpages is scalable
- Work can be extended to include other identifiers such as TCP/IP headers, DNS records, certificates

Unexpected Responses

- Observation - Disrupted measurements could either be filter **blockpages** or **unexpected responses** - Server not found errors, DDoS checks
- Similar to blockpages, Analysis also identified groups of unexpected responses

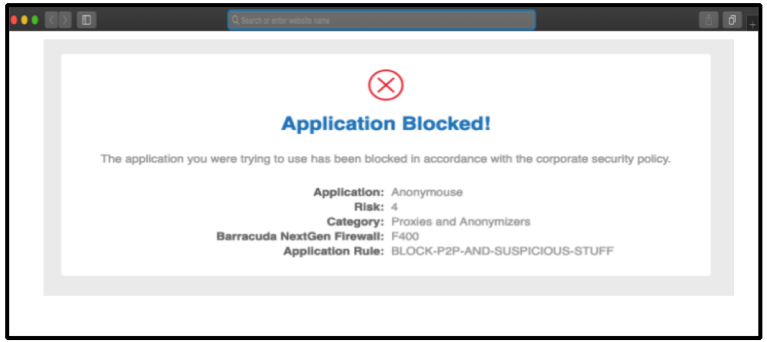

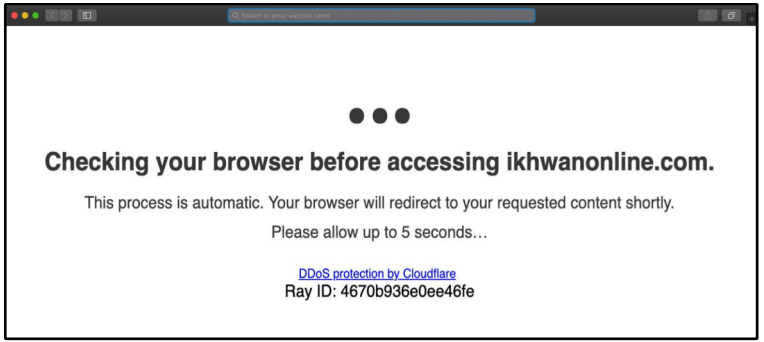

The page length metric

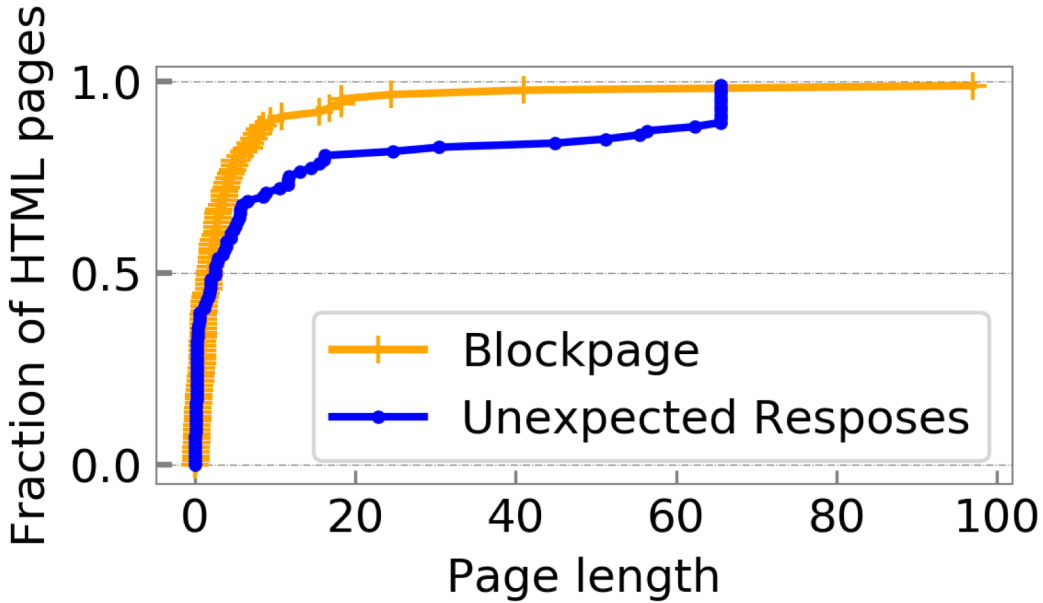

Data Collection

Censorship measurement techniques frequently observe blockpages

OONI

Volunteer measurement
https://ooni.org/

Quack

Remote measurement
VanderSloot et al. [USENIX 2018]

Hyperquack

New remote measurement

OONI

Direct measurement technique

Pros

- In-depth, user view

Challenges

- Limited scale
- Ethical constraints

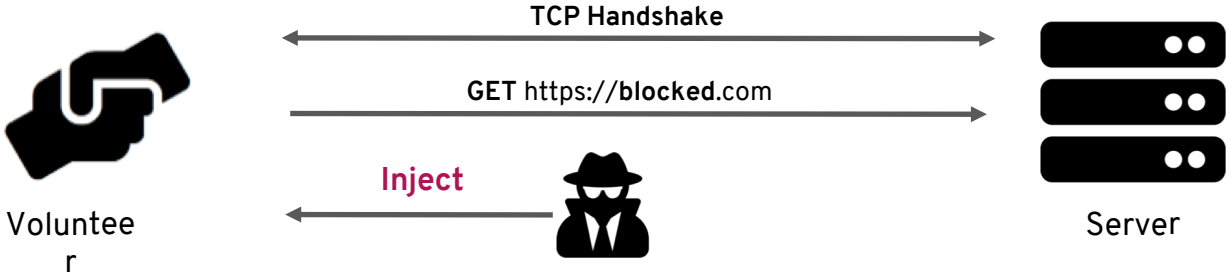

Quack

Remote measurement - TCP port 7 (Echo)

Pros

- 33,000 usable Echo servers

Challenges

- Cannot detect filters on common Port 80/443

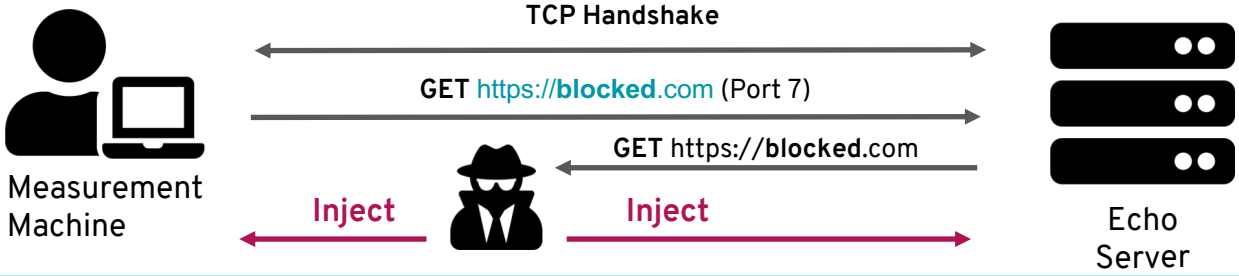

Hyperquack

- Novel remote measurement technique introduced in this study
- Uses **web servers** running on port 80 and port 443
- **Idea: Responses from web server when requesting a domain not hosted on the server is predictable**

Ethics

- OONI provides good summary of risk and obtains informed consent
- Only use organizational servers in Quack and Hyperquack
  - Servers of ISPs
  - Echo servers having NMap labels such as routers, switches etc.
- Discussed the study with colleagues inside and outside the community

Ethics

- Set up WHOIS records and web page
- Spread our requests over many servers, make a single request at a time, add delays, and use a round-robin schedule
- Fresh TCP connections and close all states
- Average - triggered filters 99 times a day

Vantage Point Characterization

	HTTP	HTTPS	Quack
Initial Set	9223	6200	36000
Experiment Set	9063	6070	33602
Number of Countries	215	204	75
Median / Country	11	13	151
Number of AS	4558	3442	3463

Iterative Classification Evaluation

	BP (%, #)	UR (%, #)	UC (%)	# of Iterations
HTTP	(56.51%, 27)	(39.39%, 105)	4.10%	3
HTTPS	(3.48%, 5)	(83.83%, 67)	12.70%	1
Quack	(93.08%, 34)	(4.8%, 116)	2.12%	2
OONI	(13.02%, 16)	(43.27%, 44)	43.71%	2

FilterMap Results - Data Collection

- Hyperquack - 38 signatures - Mostly commercial products
- Quack - 49 signatures - Mostly ISP deployments
- OONI - 21 signatures - Mostly ISP and organizational deployments
- Hyperquack detected deployments in three times as many countries as Quack and OONI

FilterMap Results - Blockpages

- Blockpages in 14 languages - Majority of blockpages were in English
- Most blockpages cited a legal concern for blocking access to content
- Many blockpages were served from redirects

FilterMap Results - Manufacturing Country

Country of Origin	Commercial filter
Israel	Allot
China	Senhua
Republic of Korea	SmartxFilter
Russia	VAS Experts
United States	Barracuda, CacheFlow, Cisco, Fortinet, IBM QRadar, Juniper, Palo Alto, SonicWall, Squid, Sucuri, WatchGuard

FilterMap Results - Categories

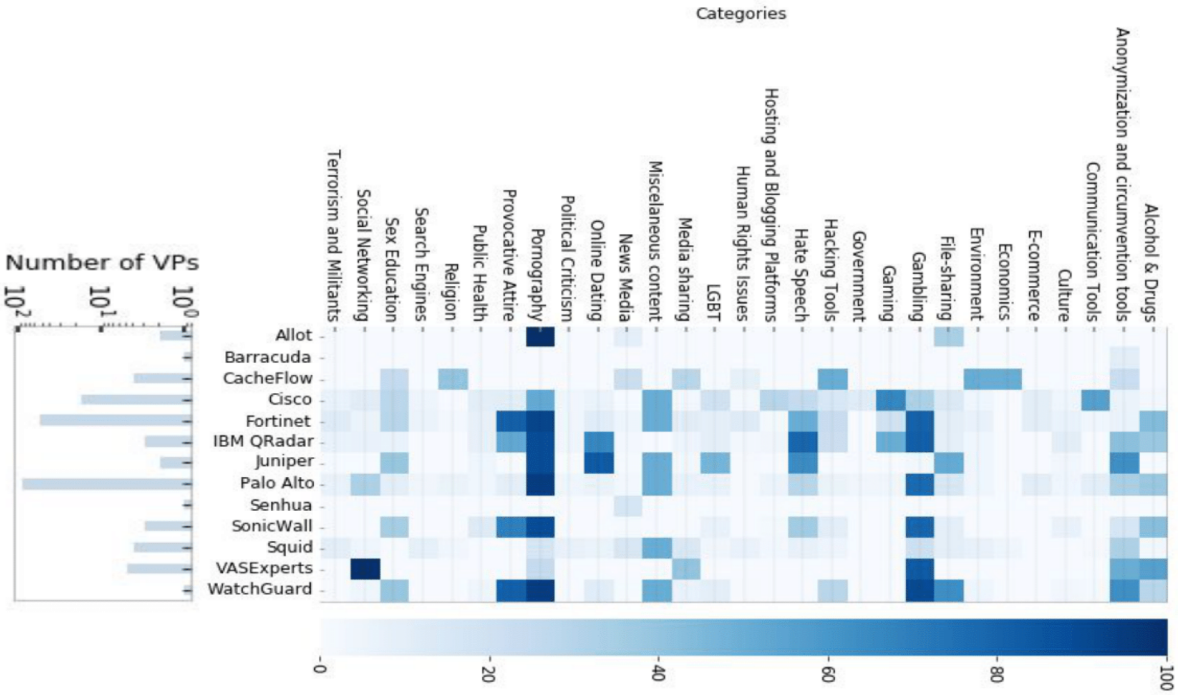

FilterMap Results - Longitudinal

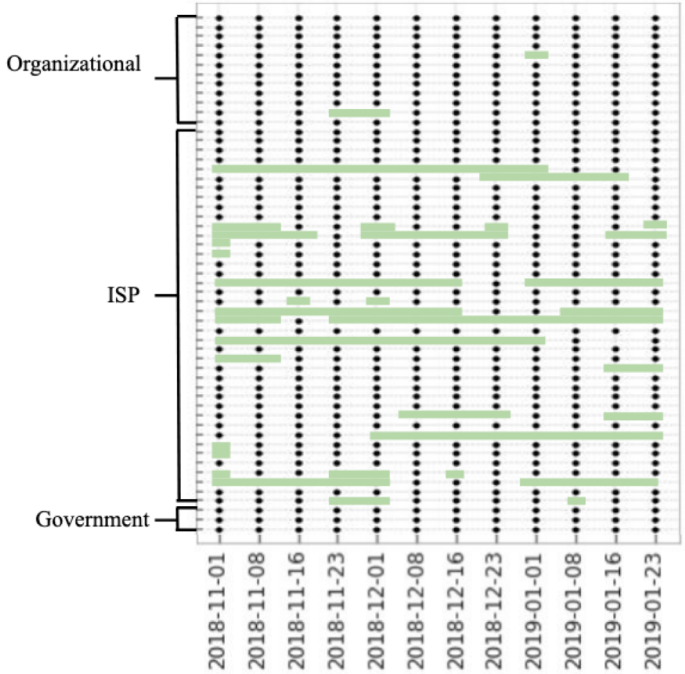

FilterMap Results - Censys

Filter	# of IPs	# of countries
Barracuda	29	4
Fortinet	10,748	151
Juniper	41	2
Palo Alto	3,087	72
Watchguard	211	28
Cisco	1,434	63
IBM QRadar	22	5
SmartxFilter	33,639	2
Sucuri	24	8
Squid	1	1