

CERTainty: Detecting DNS Manipulation at Scale using TLS Certificates


Elisa Tsai, Deepak Kumar*, Ram Sundara Raman, Gavin Li, Yael Eiger, Roya Ensafi

University of Michigan, Stanford University*



July 13, 2023

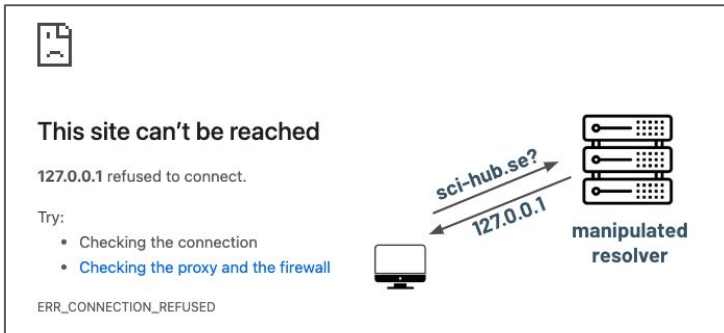
Motivation: Detecting Global DNS Manipulation



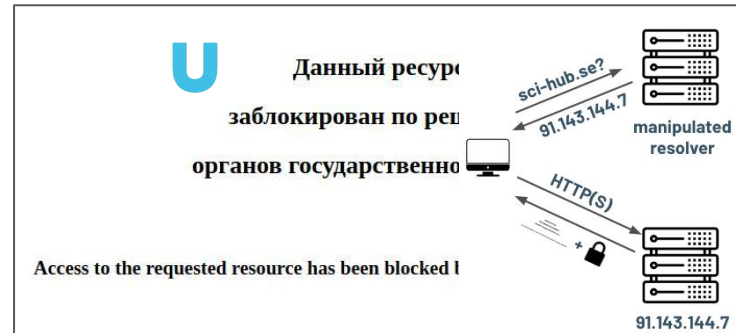
N



K



R



U

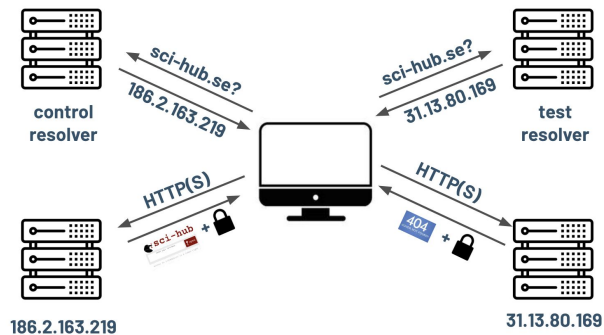


DNS Manipulation is diverse on a global scale

Challenges in global DNS manipulation measurement:

1. Website localization
2. Difference in censor behaviors
3. Lack of clear signals of manipulation

Prior Work: Consistency-Based Detection



In situ:

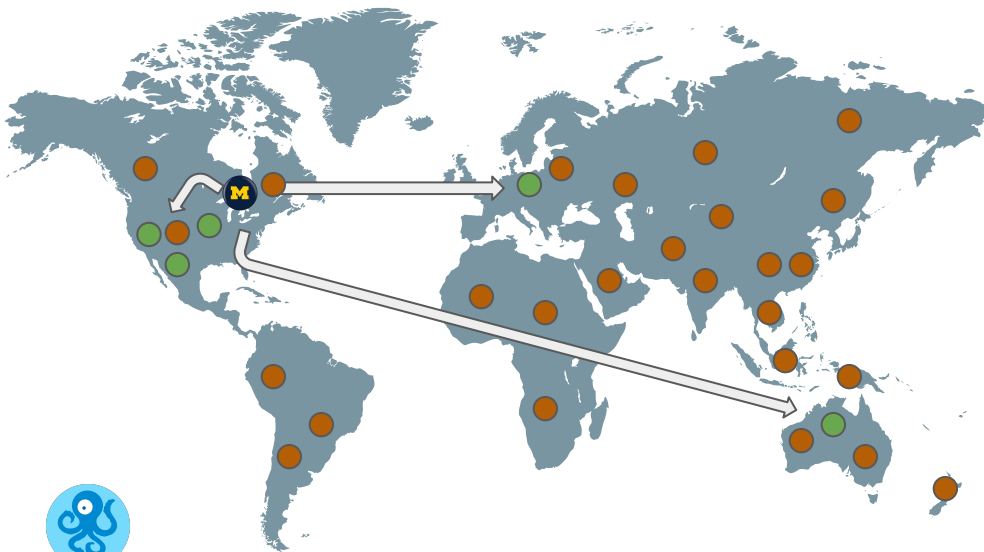
1. Rely on volunteers or 3rd party services (VPN, VPS)
2. Direct access to vantage points in residential networks

Platforms: OONI, IClab, REMeDy, UBICA...

Remote:

1. Rely on identifying ethical open resolvers on the global scale
2. Enhanced consistency, continuity, and coverage.

Platforms: Iris, Censored Planet



Design DNS manipulation detection heuristics (contd.)

Consistency

Intuition: shared infrastructural signals in global deployments:

- IP
- HTTP content hash
- HTTPS certificate hash
- AS number and name
- PTR (CDN)
- Threshold: domains → IP
- TTL

Verifiable Signals

- Blockpage matching

	Consistency								Verifiable Signals		
	IP	h(Cert)	h(HTTP)	ASN	ASNa	PTR	TTL	Thres	Cert	Page	Manual
OONI (2012)	•			•		•				•	
Censored Planet (2020)	•	•	•	•	•	•					
IClab (2020)	•			•				•		•	
Yadav et al. (2018)	•			•							•
Iris (2017)	•	•	•	•	•	•			•		
REMeDy (2017)	•			•			•				
UBICA (2015)	•										
Verkamp et al. (2012)	•					•					

Challenges with consistency heuristics:

- Rise in popularity of CDNs and cloud providers



Insight:

Move from consistency-based heuristics → verifiable signals

Verifiable signals:

- Certificate
- Blockpages fingerprinting

-> **valid TLS certificates can only be issued by the domain owners**

Data

- **Timespan:** 7 months (mid May to Nov, 2022)
- **Frequency:** Twice per week
- **Volume:**
 - **DNS:** 2,000+ domains measured on 25,000+ open resolvers, 50 M per snapshot
 - **Page:** 4 M per snapshot

Blockpage Fingerprint Dataset



Curated Blockpage Fingerprints

Category	Product	National	ISP	Corporation	Unknown	General
Count	26	92	38	14	15	30

Blockpage fingerprints open-sourced: community can easily integrate into their systems

Certificate misissuance

	<p>www.dtic.mil Issued by: DOD SW CA-60 Expires: Tuesday, August 23, 2022 at 8:36:26 AM Eastern Daylight Time ✘ "www.dtic.mil" certificate is not standards compliant</p>
	<p>sni.dreamhost.com Self-signed root certificate Expires: Friday, August 8, 2025 at 2:24:23 PM Eastern Daylight Time ▲ This certificate has not been verified by a third party</p>
	<p>www.kcna.kp Issued by: www.dprk.gov.kp Expired: Thursday, August 19, 2021 at 11:51:02 PM Eastern Daylight Time ✘ "www.kcna.kp" certificate is not trusted</p>

Stats: 1.3% control certificate are invalid, represents 3.24% of the ~2,000 domains in the test list

Certificate Validity


- The certificate chains to a **trusted** root in the Mozilla NSS Root Store (used by Mozilla Firefox)
- The **hostname** in the certificate (either in the common name or the subject alternative name) matches the domain we are attempting to reach, following the rules as specified in RFC 612

Certificate as Proxy of DNS Manipulation Detection

- 0. **Valid certificate:** confirms correct DNS resolution.
 - a. Strong signal that the IP address is *not* manipulated
 - > no pages come with a valid cert is a known blockpage

- 1. **Untrusted Root With Matched Hostname:**


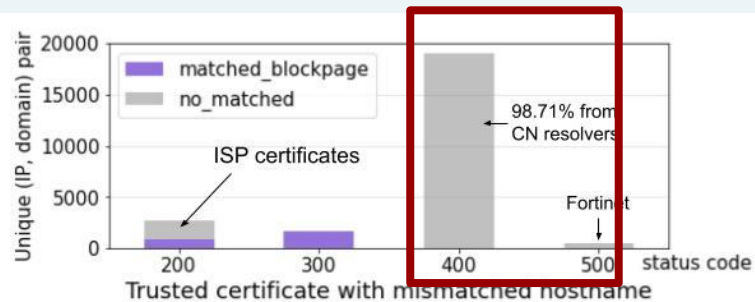
- a. **Blockpage matching: 86.25%**
(2,521 out of 2,923) of the certificates come with blockpages
- b. **The rest:** TLS proxies - Keweon, WebTitan and Mimecast
- c. 451 Unavailable For Legal
Reasons: SkyDNS and SafeDNS



Certificate as Proxy of DNS Manipulation Detection

2. Trusted Root With Mismatched Hostname:

- a. **Blockpage matching:** 10.48% (2,518/24,029) certificates match a blockpage
- b. **No matched blockpage:**
 - i. **200 OK:** largely ISP informative certificates
 - ii. **400+ status code:** 98.71% from China
 - iii. **500+ status code:** Fortinet



3. Untrusted Certificate With Mismatched Hostname:

- a. **Blockpage matching:** 92.31% (4,167/4,514) of the certificates come with blockpages
- b. **Informative certificates**
- c. **Potential misconfiguration:** common name as “textexp”, “test” and “Plesk”.

Certificate validation is an effective proxy to detect DNS manipulation.

1. Quick automated detection of DNS manipulation
2. It reveals critical information when the middleboxes and ISPs choose not to return blockpages
3. Discover covert DNS manipulation (no blockpage)

Evaluation

	Consistency									Verifiable Signals		
	Measurement Range	IP	h(Cert)	h(HTTP)	ASN	ASNa	PTR	TTL	Thres	Cert	Page	Manual
OONI (2012)	Global (200 countries)	●			●		●				●	
Censored Planet (2020)	Global (220 countries)	●	●	●	●	●	●					
IClab (2020)	Global (62 countries)	●			●				●		●	
Yadav et al. (2018)	India	●			●							●
Iris (2017)	Global (151 countries)	●	●	●	●	●	●			●		
REMeDy (2017)	Local ISPs	●			●			●				
UBICA (2015)	Pakistan, South Korea and Italy	●										
Verkamp et al. (2012)	Global (11 countries)	●					●					

CP/Iris False Positives: Consistency-based Heuristics

Iris Manipulated				Iris Unmanipulated			
Comparison	<i>CERTainty</i> Result	Count	Percentage	Comparison	<i>CERTainty</i> Result	Count	Percentage
Same with <i>CERTainty</i>	Invalid Cert	95,624	13.98%	Contradict with <i>CERTainty</i>	Invalid Cert	11,097	0.13%
	HTTP Blockpage	15,492	2.27%		HTTP Blockpage	840	0.01%
Contradict with <i>CERTainty</i>	Valid Cert	495,532	72.45%	Same with <i>CERTainty</i>	Valid Cert	7,529,487	88.85%
Unconfirmed by <i>CERTainty</i>	HTTP Only	33,592	4.91%	Unconfirmed by <i>CERTainty</i>	HTTP Only	186,627	2.20%
	Connection Error	38,407	5.61%		Connection Error	551,179	6.50%
	Malformed Cert	5,275	0.77%		Malformed Cert	194,390	2.29%

- **False Positives:** 72.45%
- **Reason:**
 - Coverage of Control
 - Metadata tagging: best effort
 - cert hash: 30.3%
 - HTTP hash: 93%
 - AS: 99%

ASN	AS Owner	Count	Percentage	Type
AS3303	Swisscom	86,115	13.63%	CDN
AS9498	Airtel	82,099	13.00%	CDN
AS20940	Akamai	63,592	10.07%	CDN
AS1299	Arelion	33,763	5.35%	CDN
AS139341	Aceville Pte	18,183	2.88%	Cloud Provider
AS54113	Fastly	16,153	2.56%	CDN
AS24940	Hetzner	12,524	1.98%	Cloud Provider
AS9121	Türk Telekom	11,815	1.87%	Telecom
AS9002	RETN	10,380	1.64%	Telecom

Top 10 ASes of False Positives

CP/Iris False Negatives: Consistency-based Heuristics

Iris Manipulated				Iris Unmanipulated			
Comparison	<i>CERTainty</i> Result	Count	Percentage	Comparison	<i>CERTainty</i> Result	Count	Percentage
Same with <i>CERTainty</i>	Invalid Cert	95,624	13.98%	Contradict with <i>CERTainty</i>	Invalid Cert	11,097	0.13%
	HTTP Blockpage	15,492	2.27%		HTTP Blockpage	840	0.01%
Contradict with <i>CERTainty</i>	Valid Cert	495,532	72.45%	Same with <i>CERTainty</i>	Valid Cert	7,529,487	88.85%
Unconfirmed by <i>CERTainty</i>	HTTP Only	33,592	4.91%	Unconfirmed by <i>CERTainty</i>	HTTP Only	186,627	2.20%
	Connection Error	38,407	5.61%		Connection Error	551,179	6.50%
	Malformed Cert	5,275	0.77%		Malformed Cert	194,390	2.29%





- **False Negatives:** 9.7%
- **AS and CDN (PTR):** experiential constraint - blockpages pages can be hosted on big CDNs
- **HTTP and cert hash:** general error page and CDN certificates
































Matched Heuristics	HTTP hash	Cert hash	ASN	AS name	CDN
Count	372	460	10,388	10,384	11,937
Percentage	3.12%	3.85%	87.02%	86.99%	100.00%

False negatives introduced by consistency-based heuristics

Findings:

Filtering Product Vendors

- **Stats:** 17 DNS manipulation filtering product vendors, 52 countries
- **Different deployment strategies:**
 - Page info:
 -  (red square) - legal blockpage
 -  (red circle) - general blockpage
 - Root cert:
 -  (black triangle) Trusted root - informative leaf cert
 -  (red triangle) Untrusted root - MitM
- **Centralized IP pool for decentralized deployment:**
 - Fortinet: one IP (208.91.112.55, AS40934)

	Product	Origin	Block Page	Root Cert	Country of Deployment
Observed in one country	Cira	CA			CA
	WebTitan	US			US
	OneDNS	CN			CN
	JusprogDNS	DE			DE
	Infoblox	US			US
	NextDNS	US			US
	Comodo	US			US
	Zyxel	CH			CH
	WatchGuard	US			US
	Securly	US			US
Observed in multiple countries	OpenDNS (Cisco)	US			AR, AU, BR, CA, CL, CN, CR, CZ, DE, ES, FR, GR, ID, IE, IN, IT, JP, KR, KZ, MX, NZ, RO, SE, SK, US, ZA
	AdguardDN:	CA			GB, BY, CY, FR, ID, LV, NZ, RU
	SafeDNS	US			AU, NL, US
	Kewoen	DE			AU, DE, FR, GB, JP, NL, US
	SkyDNS	RU			RU, UA, KZ
	CloudVeil	US			CA, US
	Fortinet	US			AR, AT, AU, BD, BF, BR, CA, CH, CL, CN, CZ, DE, DK, FR, GB, HK, ID, IN, IQ, IT, JP, KR, KW, MR, MY, NL, PH, PL, SV, TH, TR, TT, TW, US

Findings:

ISP DNS Manipulation


- **Stats:** 26 countries via cert validation
- **Different deployment strategies:**
 - Leaf cert:
 - ● issued by ISP
 - ■ issued by ISP for blocking
 - Page info:
 - ■ legal blockpage
 - ● general blockpage
 - Root cert:
 - ▲ Trusted root
 - ▲ Untrusted root

Country	AS number of returned IPs	Leaf Cert	Block Page	Root Cert
Belgium	AS2611	●	●	▲
	AS5432, AS8717	●		▲
Denmark	AS35158	●	●	▲
Italy	AS29050	■		▲
Columbia	AS35158	●	■	▲
Greece	AS6799	■		▲
Switzerland	AS3303	●	■	▲
Germany	AS24940		■	▲
Australia	AS16509	■	■	▲

Country	AS number of returned IPs	Leaf Cert	Block Page	Root Cert
Russia	AS12616, AS44347, AS44587, AS49505, AS34241	■	■	▲
	AS25549, AS31483, AS34757	●	●	▲
	AS12389, AS50466		■	▲
	AS42071, AS42071	●		▲
	AS57571, AS43287, AS49469	■	■	▲
	AS8395	●	■	▲
Ukraine	AS42546	●	■	▲
	AS42546	●	●	▲
Indonesia	AS58396, AS45287, AS45287, AS45287, AS38758	●	●	▲
	AS9341, AS9341, AS5578, AS9341		●	▲
	AS16276, AS141626, AS141626, AS7713	●	●	▲
	AS58495, AS132634	■		▲
	AS140413, AS136873	■	■	▲
	AS56241	●	●	▲
Nepal	AS63991	●		▲
	AS140973	■	●	▲
Thailand	AS23969	■		▲
Singapore	AS3758, AS3758	●		▲
Belarus	AS6697		■	▲
Lithuania	AS212531	●	■	▲
Romania	AS31313		■	▲
	AS12302			▲

Case study: Covert DNS Manipulation

- **Signal:** 400+ status code page with trusted cert
- **Stats:** 98.71% of those IPs are returned by DNS resolvers in China.
- **IP ownership:**
 - Facebook (66.30%)
 - Twitter (29.10%)
 - Cloudflare (3.36%)
 - other blocked CDN services: Fastly and Akamai (less than 0.08%)
- **Potential censorship leakage:** 14 surrounding countries shared some overlaps



Summary

- Consistency-based heuristics are **error-prone**:
 - 72.45% of the manipulated DNS responses identified by the current state-of-the-art are **false positives**.
 - Experiential constraints like AS matching also introduce **false negatives** (9.7%).
- Should actively look for **verifiable signals** of DNS manipulation
- Identified 17 TLS proxy vendors deployed in 52 countries, as well as 26 countries with ISP-level DNS manipulation -> **pinpoint the deployer of DNS manipulation**
- Identified covert cases of DNS manipulation.
- Open-sourced 200+ unique DNS blockpage fingerprints=
- Collaborating with other platforms to **improve data quality**



<https://www.geo3550.org/>