Intercept and Inject: DNS Response Manipulation in the Wild

Yevheniya Nosyk [1], Qasim Lone [4], Yury Zhauniarovich [2], Carlos H. Gañán [2,5],
Emile Aben [4], Giovane C. M. Moura [2,3], Samaneh Tajalizadehkhoob [5],
Andrzej Duda [1], and Maciej Korczyński [1]

[1] *Université Grenoble Alpes*
[2] *TU Delft*
[3] *SIDN Labs*
[4] *RIPE NCC*
[5] *ICANN*

Passive and Active Measurement Conference 2023

23rd March 2023

DNS resolution

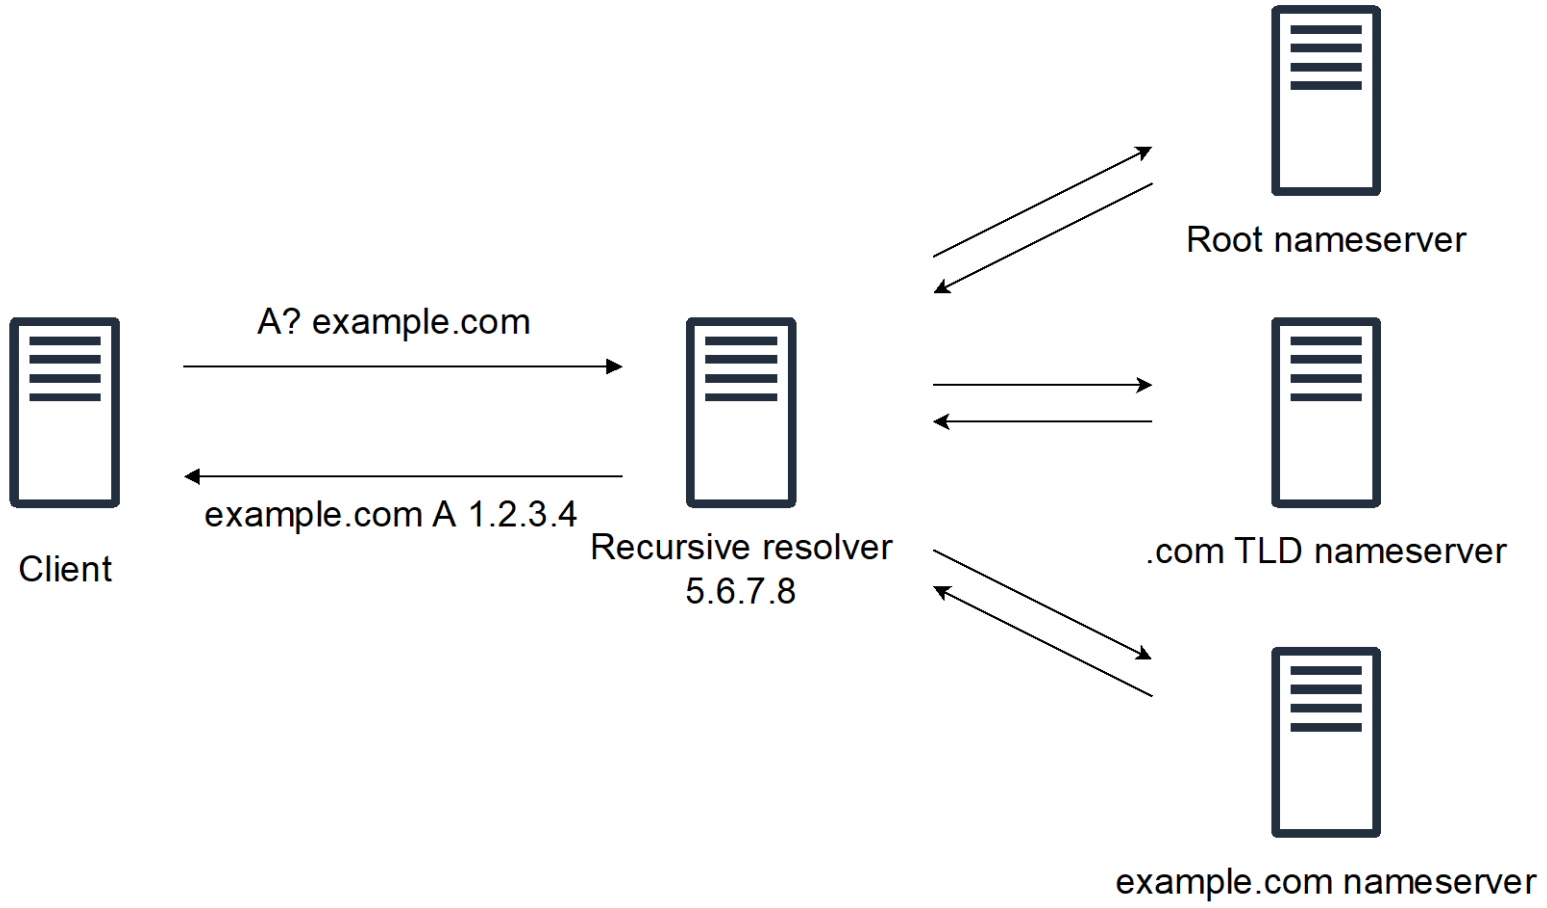

DNS Root Server System

- 12 operators
- 13 root servers (a-m)
- 1,644 anycast instances
- local vs. global

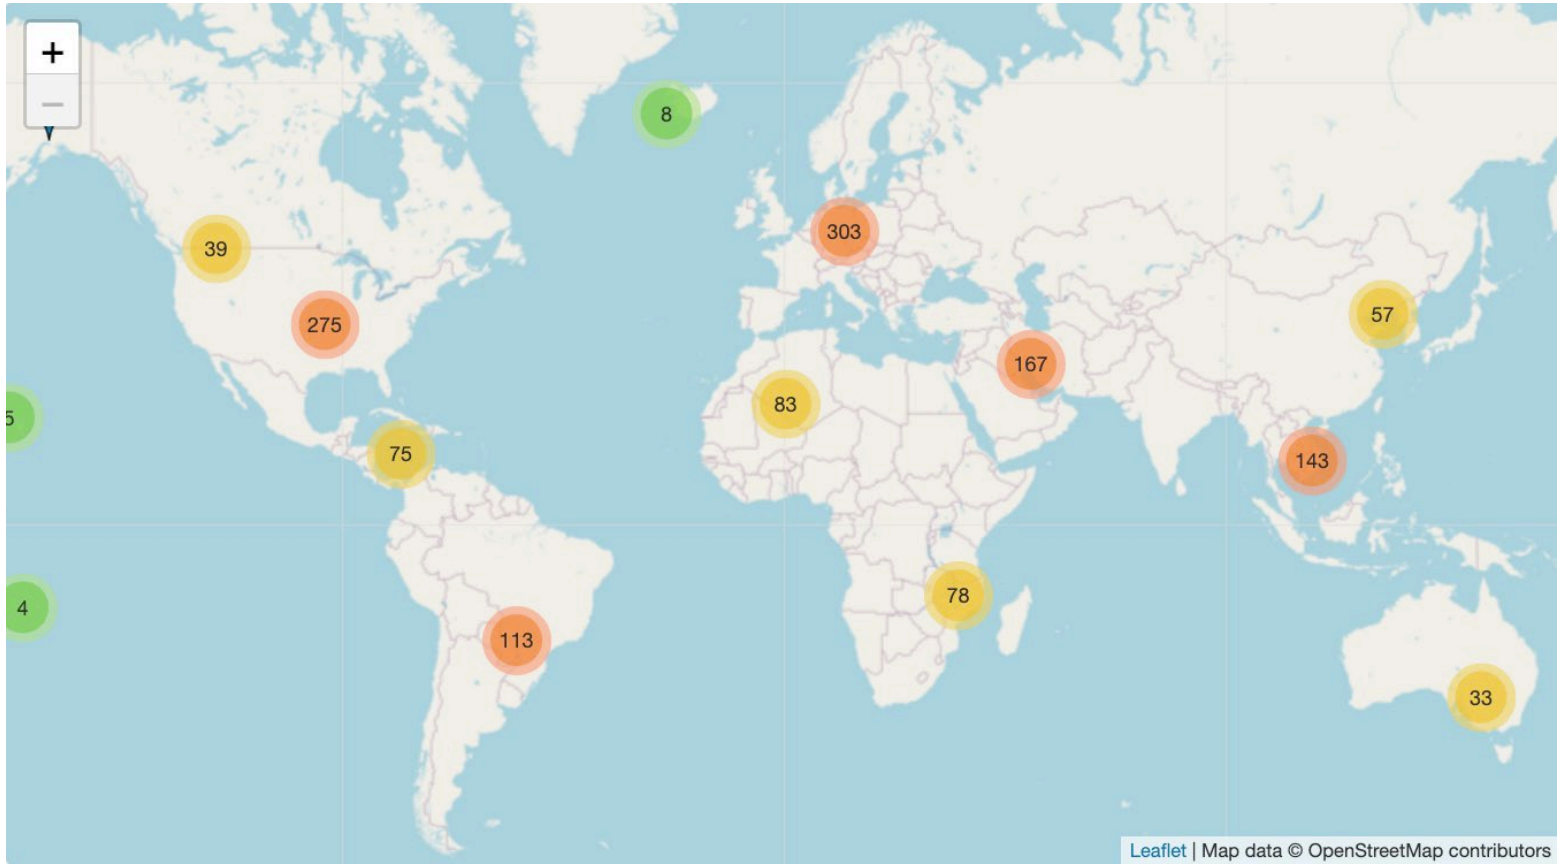

Source : https://root-servers.org

Root Server responses

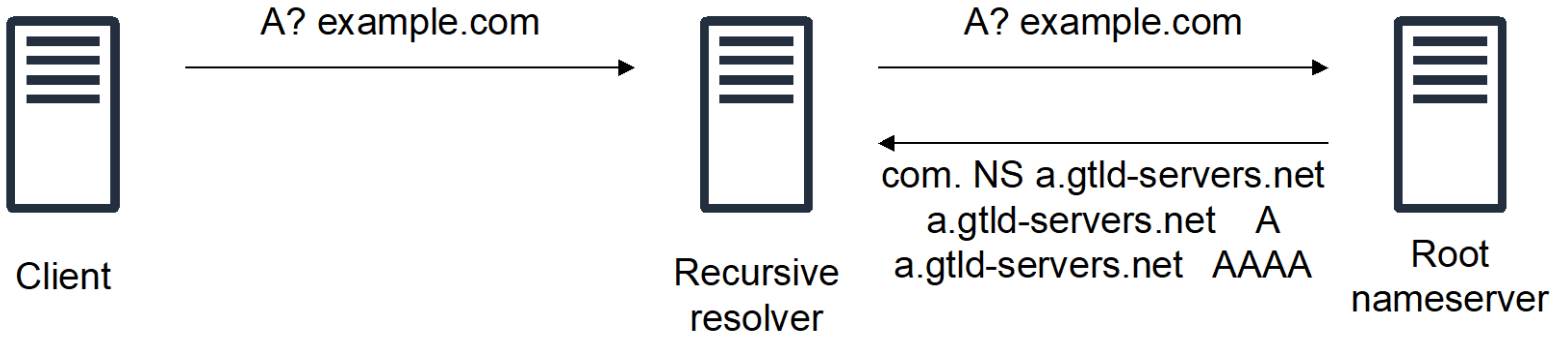

Root Server responses

```
$ dig @k.root-servers.net google.com


; <<>> DiG 9.11.5-P4-5.1+deb10u8-Debian <<>> @k.root-servers.net google.com
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 22260
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 13, ADDITIONAL: 27
;; WARNING: recursion requested but not available


;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;google.com. IN A


;; AUTHORITY SECTION:
com. 172800 IN NS a.gtld-servers.net.
com. 172800 IN NS b.gtld-servers.net.
com. 172800 IN NS c.gtld-servers.net.
com. 172800 IN NS d.gtld-servers.net.
com. 172800 IN NS e.gtld-servers.net.
com. 172800 IN NS f.gtld-servers.net.
com. 172800 IN NS g.gtld-servers.net.
com. 172800 IN NS h.gtld-servers.net.
com. 172800 IN NS i.gtld-servers.net.
com. 172800 IN NS j.gtld-servers.net.
com. 172800 IN NS k.gtld-servers.net.
com. 172800 IN NS l.gtld-servers.net.
com. 172800 IN NS m.gtld-servers.net.
```

```
;; ADDITIONAL SECTION:
a.gtld-servers.net. 172800 IN AAAA 2001:503:a83e::2:30
b.gtld-servers.net. 172800 IN AAAA 2001:503:231d::2:30
c.gtld-servers.net. 172800 IN AAAA 2001:503:83eb::30
d.gtld-servers.net. 172800 IN AAAA 2001:500:856e::30
e.gtld-servers.net. 172800 IN AAAA 2001:502:1ca1::30
f.gtld-servers.net. 172800 IN AAAA 2001:503:d414::30
g.gtld-servers.net. 172800 IN AAAA 2001:503:eea3::30
h.gtld-servers.net. 172800 IN AAAA 2001:502:8cc::30
i.gtld-servers.net. 172800 IN AAAA 2001:503:39c1::30
j.gtld-servers.net. 172800 IN AAAA 2001:502:7094::30
k.gtld-servers.net. 172800 IN AAAA 2001:503:d2d::30
l.gtld-servers.net. 172800 IN AAAA 2001:500:d937::30
m.gtld-servers.net. 172800 IN AAAA 2001:501:b1f9::30
a.gtld-servers.net. 172800 IN A 192.5.6.30
b.gtld-servers.net. 172800 IN A 192.33.14.30
c.gtld-servers.net. 172800 IN A 192.26.92.30
d.gtld-servers.net. 172800 IN A 192.31.80.30
e.gtld-servers.net. 172800 IN A 192.12.94.30
f.gtld-servers.net. 172800 IN A 192.35.51.30
g.gtld-servers.net. 172800 IN A 192.42.93.30
h.gtld-servers.net. 172800 IN A 192.54.112.30
i.gtld-servers.net. 172800 IN A 192.43.172.30
j.gtld-servers.net. 172800 IN A 192.48.79.30
k.gtld-servers.net. 172800 IN A 192.52.178.30
l.gtld-servers.net. 172800 IN A 192.41.162.30
m.gtld-servers.net. 172800 IN A 192.55.83.30


;; Query time: 5 msec
;; SERVER: 2001:7fd::1#53(2001:7fd::1)
;; WHEN: Thu Mar 09 10:21:31 CET 2023
;; MSG SIZE  rcvd: 835
```

November 2021 k-root event

[dns-operations] K-root in CN leaking outside of CN

Manu Bretelle chantr4 at gmail.com
Sat Nov 6 04:13:53 UTC 2021

- Previous message (by thread): [dns-operations] Request for proposals for implementation for study of RSSAC028
- Next message (by thread): [dns-operations] K-root in CN leaking outside of CN
- **Messages sorted by:** [date] [thread] [subject] [author]

```
Hi all,

Based on https://root-servers.org/, there are a few root servers operated
from Mainland China.

How do we ensure that those are not advertised outside of China so DNS
answers are not poisoned by the GFW?

Are there any contracts that root in CN are supposed to follow to prevent
this? Is the onus put on both the CN ASNs and their respective non-CN ASNs
peers to not advertise/not accept the root range on those specific peering
links? If so, how is it ensured that every operator knows about those rules?
Is there any monitoring performed by root operators to ensure that leaks
are being detected and possibly addressed?

I don't believe this specific leak I am seeing is malicious, but rather is
just a misconfiguration and I really wonder how this could be
prevented/addressed early on.
I have ran some probes in other regions and do not have proof that this is
happening more widely than a specific AS, but this was not exhaustive and I
could have very likely missed something.
```


Source: https://lists.dns-oarc.net/pipermail/dns-operations/2021-November/021437.html

November 2021 k-root event: RIPE Atlas View

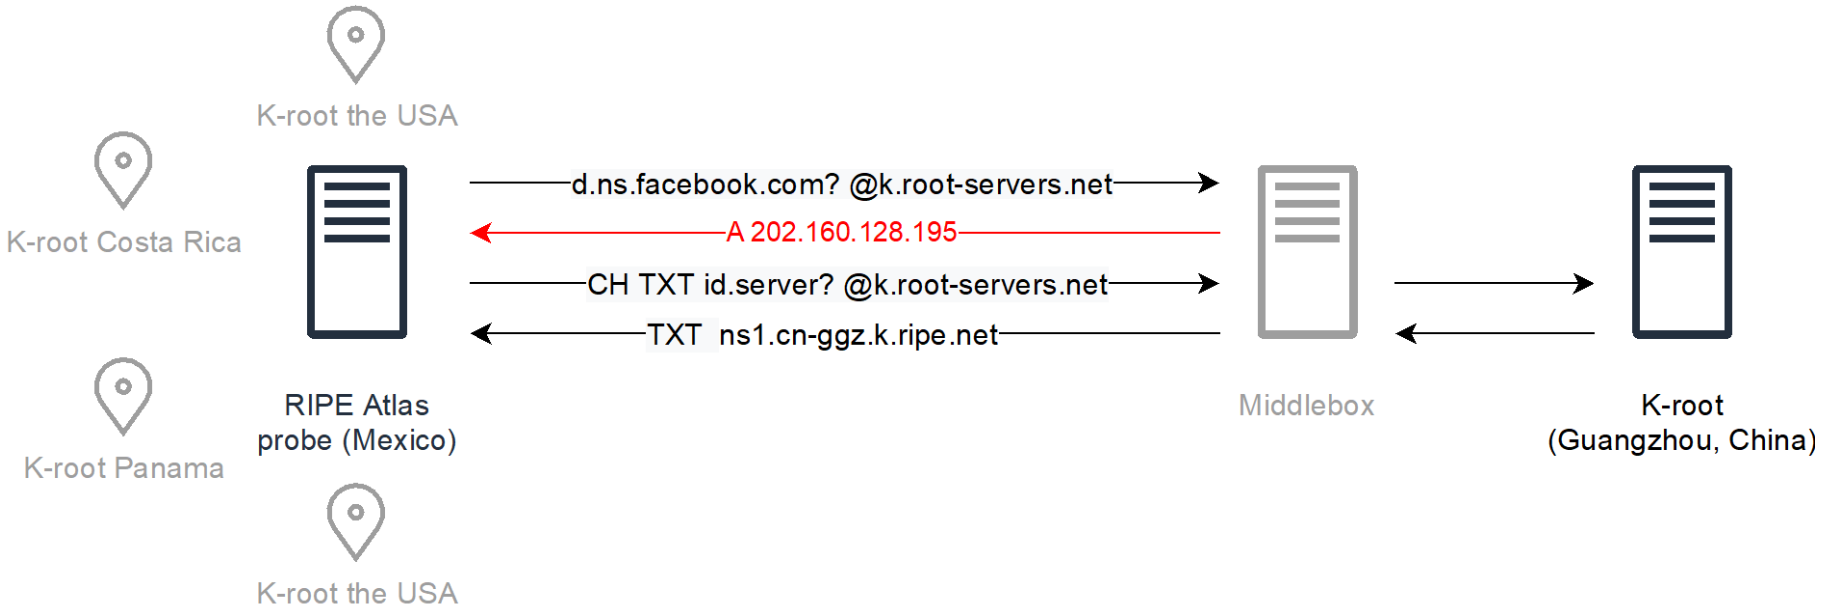

November 2021 k-root event:

[dns-operations] K-root in CN leaking outside of CN

Anand Buddhdev anandb at ripe.net
Mon Nov 8 08:12:40 UTC 2021

- Previous message (by thread): [dns-operations] K-root in CN leaking outside of CN
- Next message (by thread): [dns-operations] K-root in CN leaking outside of CN
- **Messages sorted by:** [date] [thread] [subject] [author]

```
Hi Davey, Manu,

The server we operate in Guangzhou was indeed reachable from outside
China. This is not the intention, of course. On Saturday, when we got
notification about this, we withdrew the prefix from the server, and we
are communicating with the host to solve this.

Many people have already said this, but I'd like to make it clear that
the K-root server was NOT emitting false responses for Facebook and
WhatsApp. The responses were being modified by something between the
server and its clients.

Regards,
Anand Buddhdev
RIPE NCC
```

To what extent the local Guangzhou instance of the k-root is reachable from the outside?

Nameserver identifiers (RFC-4892, RFC-5001)

```
$ dig @k.root-servers.net CH TXT version.bind +short
"NSD"
$ dig @k.root-servers.net CH TXT version.server +short
"NSD"
$ dig @k.root-servers.net CH TXT hostname.bind +short
"ns1.ch-gva.k.ripe.net"
$ dig @k.root-servers.net CH TXT id.server +short
"ns1.ch-gva.k.ripe.net"
$ dig @k.root-servers.net ripe.net +nsid

...

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
; NSID: 6e 73 31 2e 63 68 2d 67 76 61 2e 6b 2e 72 69 70 65 2e 6e 65 74 ("ns1.ch-gva.k.ripe.net")
;; QUESTION SECTION:
;ripe.net. IN A

...
```

RIPE Atlas built-in measurements

- More than 11k active probes at a time
- Instance reachable at least 2 month before being reported from 57 probes in 15 countries (AU, UA, CO, HK, LK, CH, FR, US, KR, DK, MX, ZM, BE, GB, NP, KE)
- Instance occasionally reachable the following 9 months after the fix from 12 probes in 5 countries, but over IPv6 (RU, IL, MX, DK, HK)
- 11 probes receiving bogus responses for facebook.com (IPs of Dropbox and Twitter)

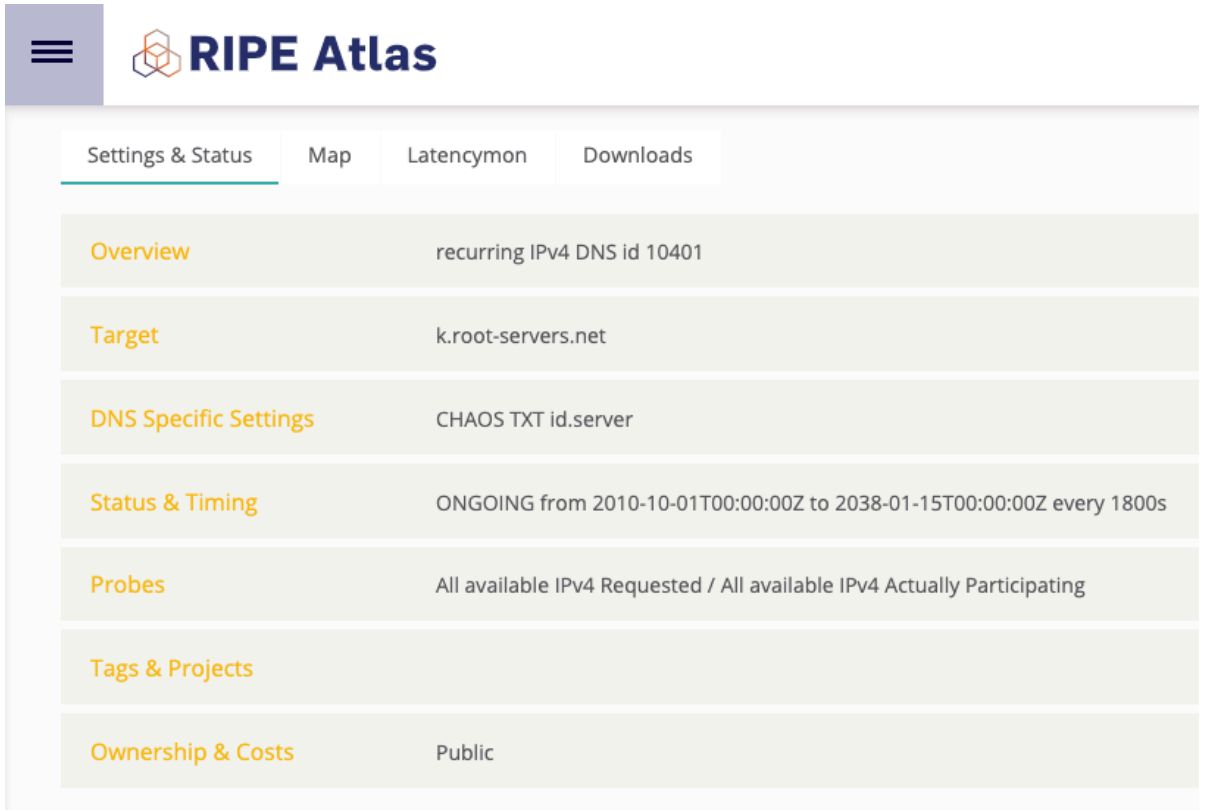

To what extent queries to DNS root servers experience manipulation?

Measurement setup: RIPE Atlas

- February – October 2022
- 1b measurements (312 query types sent every 12h)
- NSID option for identification
- 14.3k RIPE Atlas probes (177 countries and 4,132 ASes)

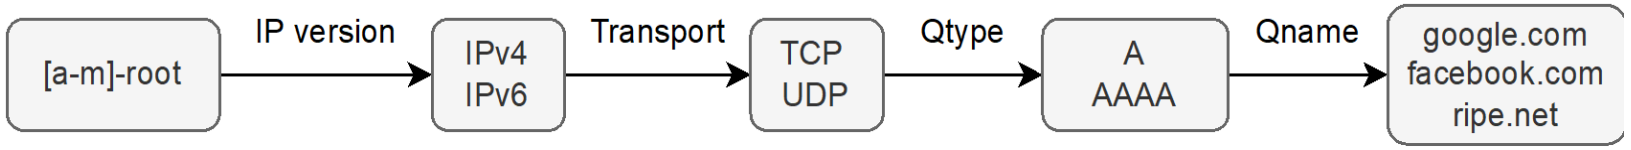

Two types of responses

Non-injected (99.18% or 1b measurements):


```
{"measurement_id": 34848600, "probe_id": 53005, "executed": "2022-01-18 22:36:26+00:00", "response_nsid": ["ns2.nl-ams.k.ripe.net"], "answers": []}

{"measurement_id": 39032627, "probe_id": 27793, "executed": "2022-02-24 15:27:42+00:00", "response_nsid": ["M-ORY-1"], "answers": []}
```


Injected (0.82% or 9m measurements):


```
{"measurement_id": 34848596, "probe_id": 2147, "executed": "2022-01-18 23:28:34+00:00", "response_nsid": ["CleanBrowsing v1.6a - dns-edge-europe-frankfurt-c"], "answers": [{"Name": "google.com.", "Type": "A", "Class": "IN", "TTL": 90, "RDlength": 4, "Address": "142.250.180.238"}]}

{"measurement_id": 34848610, "probe_id": 34903, "executed": "2022-06-29 02:05:38+00:00", "response_nsid": [], "answers": [{"Name": "facebook.com.", "Type": "A", "Class": "IN", "TTL": 600, "RDlength": 4, "Address": "199.59.149.244"}]}
```

Injected responses (11m)

A (2,419 unique IPs, 49% of facebook.com and 89.6% of google.com responses were valid) - 7m responses:

{"measurement_id": 34848595, "probe_id": 31021, "executed": "2022-02-01 11:27:48+00:00", "response_nsid": [""], "answers": [{"Name": "google.com.", "Type": "A", "Class": "IN", "TTL": 235, "RDlength": 4, "Address": "216.58.208.142"}]}

Injected responses (11m)

A (2,419 unique IPs, 49% of facebook.com and 89.6% of google.com responses were valid) - 7m responses:

{"measurement_id": 34848595, "probe_id": 31021, "executed": "2022-02-01 11:27:48+00:00", "response_nsid": [""], "answers": [{"Name": "google.com.", "Type": "A", "Class": "IN", "TTL": 235, "RDlength": 4, "Address": "216.58.208.142"}]}

AAAA (3,221 unique IPs, 64.4% of facebook.com and 98.3% of google.com responses were valid) – 4m responses:

{"measurement_id": 34848597, "probe_id": 33018, "executed": "2022-02-13 12:32:09+00:00", "response_nsid": ["gpdns-sfo"], "answers": [{"Name": "google.com.", "Type": "AAAA", "Class": "IN", "TTL": 174, "RDlength": 16, "Address": "2607:f8b0:4007:80a:0:0:0:200e"}]}

Injected responses (11m)

A (2,419 unique IPs, 49% of facebook.com and 89.6% of google.com responses were valid) - 7m responses:

{"measurement_id": 34848595, "probe_id": 31021, "executed": "2022-02-01 11:27:48+00:00", "response_nsid": [""], "answers": [{"Name": "google.com.", "Type": "A", "Class": "IN", "TTL": 235, "RDlength": 4, "Address": "216.58.208.142"}]}

AAAA (3,221 unique IPs, 64.4% of facebook.com and 98.3% of google.com responses were valid) – 4m responses:

{"measurement_id": 34848597, "probe_id": 33018, "executed": "2022-02-13 12:32:09+00:00", "response_nsid": ["gpdns-sfo"], "answers": [{"Name": "google.com.", "Type": "AAAA", "Class": "IN", "TTL": 174, "RDlength": 16, "Address": "2607:f8b0:4007:80a:0:0:0:200e"}]}

URI (received on 15 probes from Iran) – 42.5k responses:

{"measurement_id": 39032341, "probe_id": 1000185, "executed": "2022-10-04 05:55:01+00:00", "response_nsid": [], "answers": [{"Name": ".", "Type": 256, "Class": 256, "TTL": 107008, "RDlength": 1034, "Rdata": "0a2224"}]}

Injected responses (11m)

A (2,419 unique IPs, 49% of facebook.com and 89.6% of google.com responses were valid) - 7m responses:

{"measurement_id": 34848595, "probe_id": 31021, "executed": "2022-02-01 11:27:48+00:00", "response_nsid": [""], "answers": [{"Name": "google.com.", "Type": "A", "Class": "IN", "TTL": 235, "RDlength": 4, "Address": "216.58.208.142"}]}

AAAA (3,221 unique IPs, 64.4% of facebook.com and 98.3% of google.com responses were valid) – 4m responses:

{"measurement_id": 34848597, "probe_id": 33018, "executed": "2022-02-13 12:32:09+00:00", "response_nsid": ["gpdns-sfo"], "answers": [{"Name": "google.com.", "Type": "AAAA", "Class": "IN", "TTL": 174, "RDlength": 16, "Address": "2607:f8b0:4007:80a:0:0:0:200e"}]}

URI (received on 15 probes from Iran) – 42.5k responses:

{"measurement_id": 39032341, "probe_id": 1000185, "executed": "2022-10-04 05:55:01+00:00", "response_nsid": [], "answers": [{"Name": ".", "Type": 256, "Class": 256, "TTL": 107008, "RDlength": 1034, "Rdata": "0a2224"}]}

SOA (one probe from the USA affected) – 6.7k responses:

{"measurement_id": 39032362, "probe_id": 50347, "executed": "2022-06-08 20:51:58+00:00", "response_nsid": [null], "answers": [{"Name": "facebook.com.", "Type": "SOA", "Class": "IN", "TTL": 30, "RDlength": 62, "MasterServerName": "dns1.dnsfilter.com.", "MaintainerName": "dadmin.dnsfilter.com.", "Serial": 1, "Refresh": 30, "Retry": 30, "Expire": 30, "NegativeTtl": 30}]}

Injected responses (11m)

A (2,419 unique IPs, 49% of facebook.com and 89.6% of google.com responses were valid) - 7m responses:

{"measurement_id": 34848595, "probe_id": 31021, "executed": "2022-02-01 11:27:48+00:00", "response_nsid": [""], "answers": [{"Name": "google.com.", "Type": "A", "Class": "IN", "TTL": 235, "RDlength": 4, "Address": "216.58.208.142"}]}

AAAA (3,221 unique IPs, 64.4% of facebook.com and 98.3% of google.com responses were valid) – 4m responses:

{"measurement_id": 34848597, "probe_id": 33018, "executed": "2022-02-13 12:32:09+00:00", "response_nsid": ["gpdns-sfo"], "answers": [{"Name": "google.com.", "Type": "AAAA", "Class": "IN", "TTL": 174, "RDlength": 16, "Address": "2607:f8b0:4007:80a:0:0:0:200e"}]}

URI (received on 15 probes from Iran) – 42.5k responses:

{"measurement_id": 39032341, "probe_id": 1000185, "executed": "2022-10-04 05:55:01+00:00", "response_nsid": [], "answers": [{"Name": ".", "Type": 256, "Class": 256, "TTL": 107008, "RDlength": 1034, "Rdata": "0a2224"}]}

SOA (one probe from the USA affected) – 6.7k responses:

{"measurement_id": 39032362, "probe_id": 50347, "executed": "2022-06-08 20:51:58+00:00", "response_nsid": [null], "answers": [{"Name": "facebook.com.", "Type": "SOA", "Class": "IN", "TTL": 30, "RDlength": 62, "MasterServerName": "dns1.dnsfilter.com.", "MaintainerName": "dadmin.dnsfilter.com.", "Serial": 1, "Refresh": 30, "Retry": 30, "Expire": 30, "NegativeTtl": 30}]}

CNAME – 4.5k responses:

{"measurement_id": 39032348, "probe_id": 19195, "executed": "2022-10-03 05:53:59+00:00", "response_nsid": [], "answers": [{"Name": "google.com.", "Type": "CNAME", "Class": "IN", "TTL": 49919, "RDlength": 18, "Target": "forcesafesearch.google.com."}, {"Name": "forcesafesearch.google.com.", "Type": "A", "Class": "IN", "TTL": 65939, "RDlength": 4, "Address": "216.239.38.120"}]}

Identifying responding services with NSIDs

More than 12k unique NSID strings:

- 2k root servers (**always** return valid responses)
- 9.9k public resolvers (CloudflareDNS, OpenDNS, Quad9, Google DNS)
- 211 unclassified
- 41 filtering services (CleanBrowsing)
- empty strings (78% of all the injected responses)

Affected probes

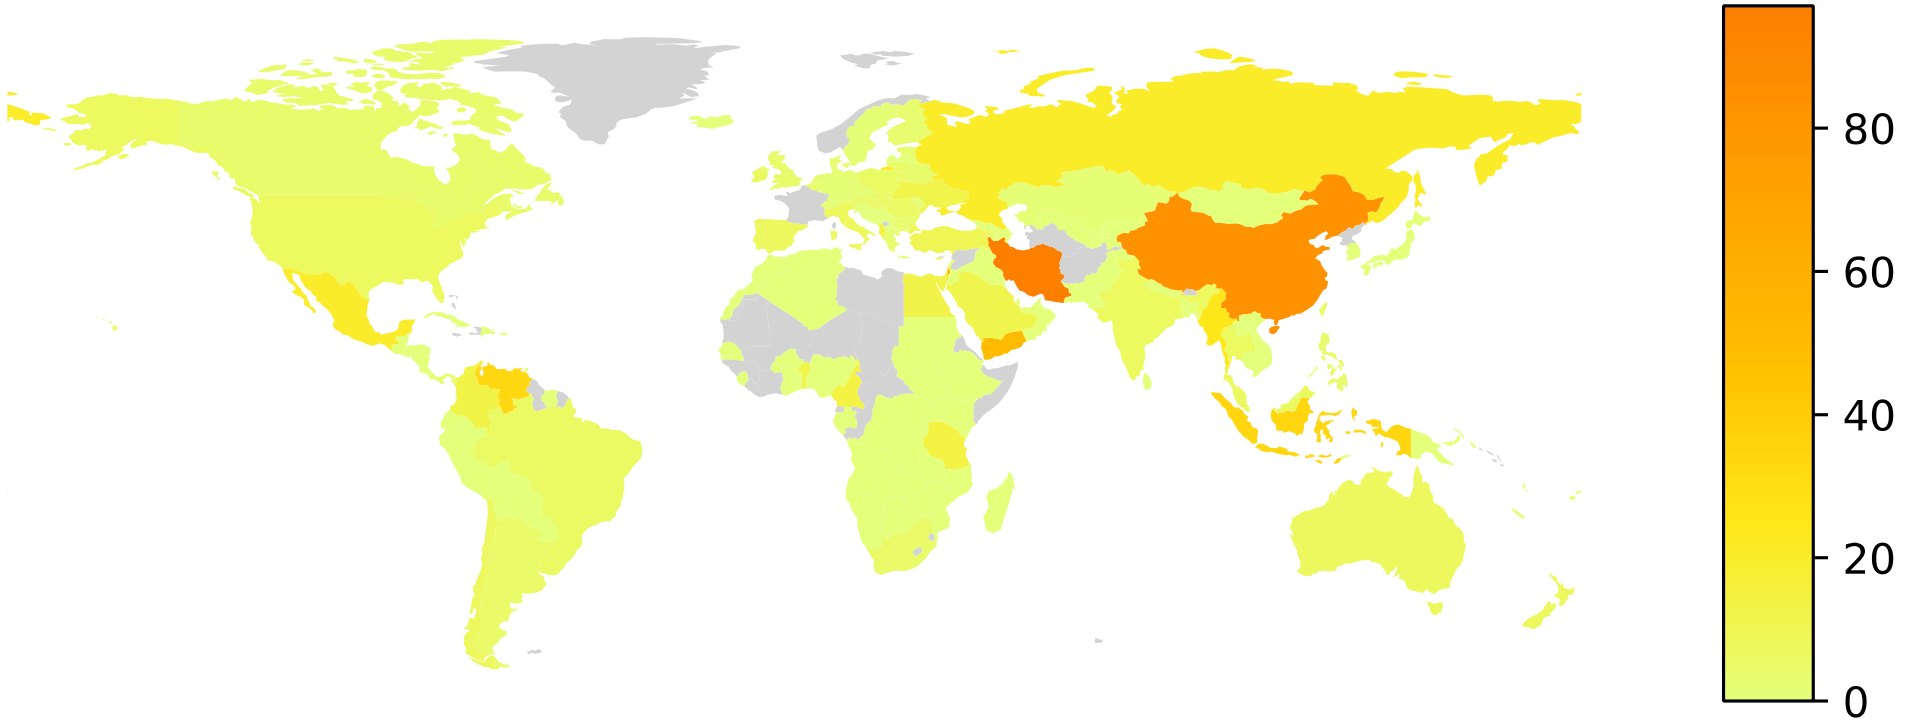

Persistence

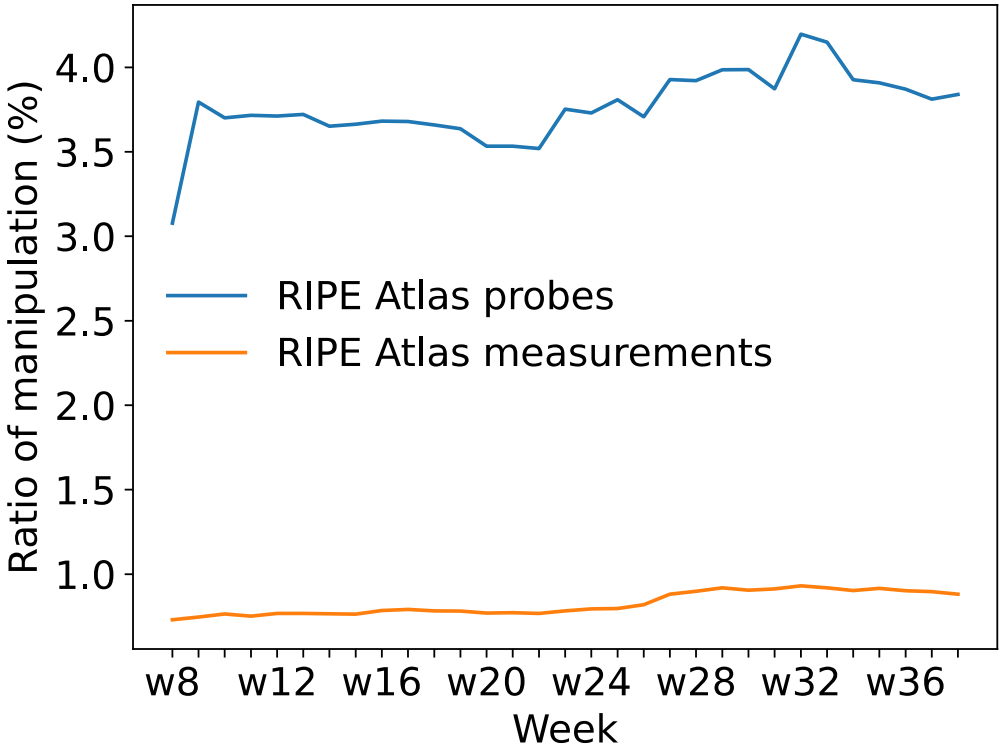

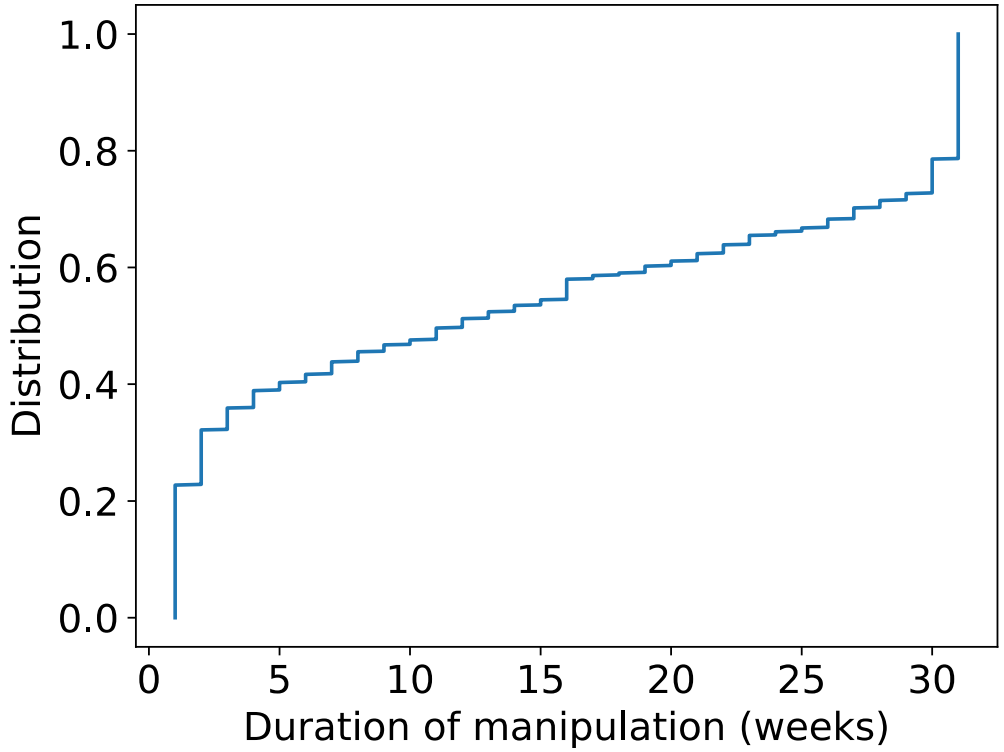

Countermeasures

- BGP communities [1]
- QNAME minimization
- Encrypted DNS
- DNSSEC

Status:	Proposed Standard
Obsoletes:	7816
More info:	Datatracker | IPR | Info page

Stream:	Internet Engineering Task Force (IETF)
RFC:	9156
Obsoletes:	7816
Category:	Standards Track
Published:	November 2021
ISSN:	2070-1721
Authors:	S. Bortzmeyer (*AFNIC*), R. Dolmans (*NLnet Labs*), P. Hoffman (*ICANN*)

RFC 9156
DNS Query Name Minimisation to Improve Privacy

Abstract

This document describes a technique called "QNAME minimisation" to improve DNS privacy, where the DNS resolver no longer always sends the full original QNAME and original QTYPE to the upstream name server. This document obsoletes RFC 7816.

[1] *Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet anycast: performance, problems, & potential. In Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18). Association for Computing Machinery, New York, NY, USA, 59–73. https://doi.org/10.1145/3230543.3230547*

Key takeaways

- DNS root queries are manipulated
- Injected data is not always bogus
- Transparent to end users
- May introduce collateral damage
- BGP leaks stay unnoticed

Acknowledgments

We thank root server operators for validating the nameserver identifiers. This work was partially supported by RIPE NCC, Carnot LSI, the Grenoble Alpes Cybersecurity Institute (under the contract ANR-15-IDEX-02), and the French Ministry of Research (PERSYVAL-Lab project under the contract ANR-11-LABX-0025-01 and DiNS project under the contract ANR-19-CE25-0009-01).

Thank you!

yevheniya.nosyk@univ-grenoble-alpes.fr