

Intercept and Inject: DNS Response Manipulation in the Wild

Yevheniya Nosyk¹, Qasim Lone⁴, Yury Zhauniarovich², Carlos H. Gañán^{2,5},
Emile Aben⁴, Giovane C. M. Moura^{2,3}, Samaneh Tajalizadehkhoob⁵,
Andrzej Duda¹, and Maciej Korczyński¹

¹ *Université Grenoble Alpes*

² *TU Delft*

³ *SIDN Labs*

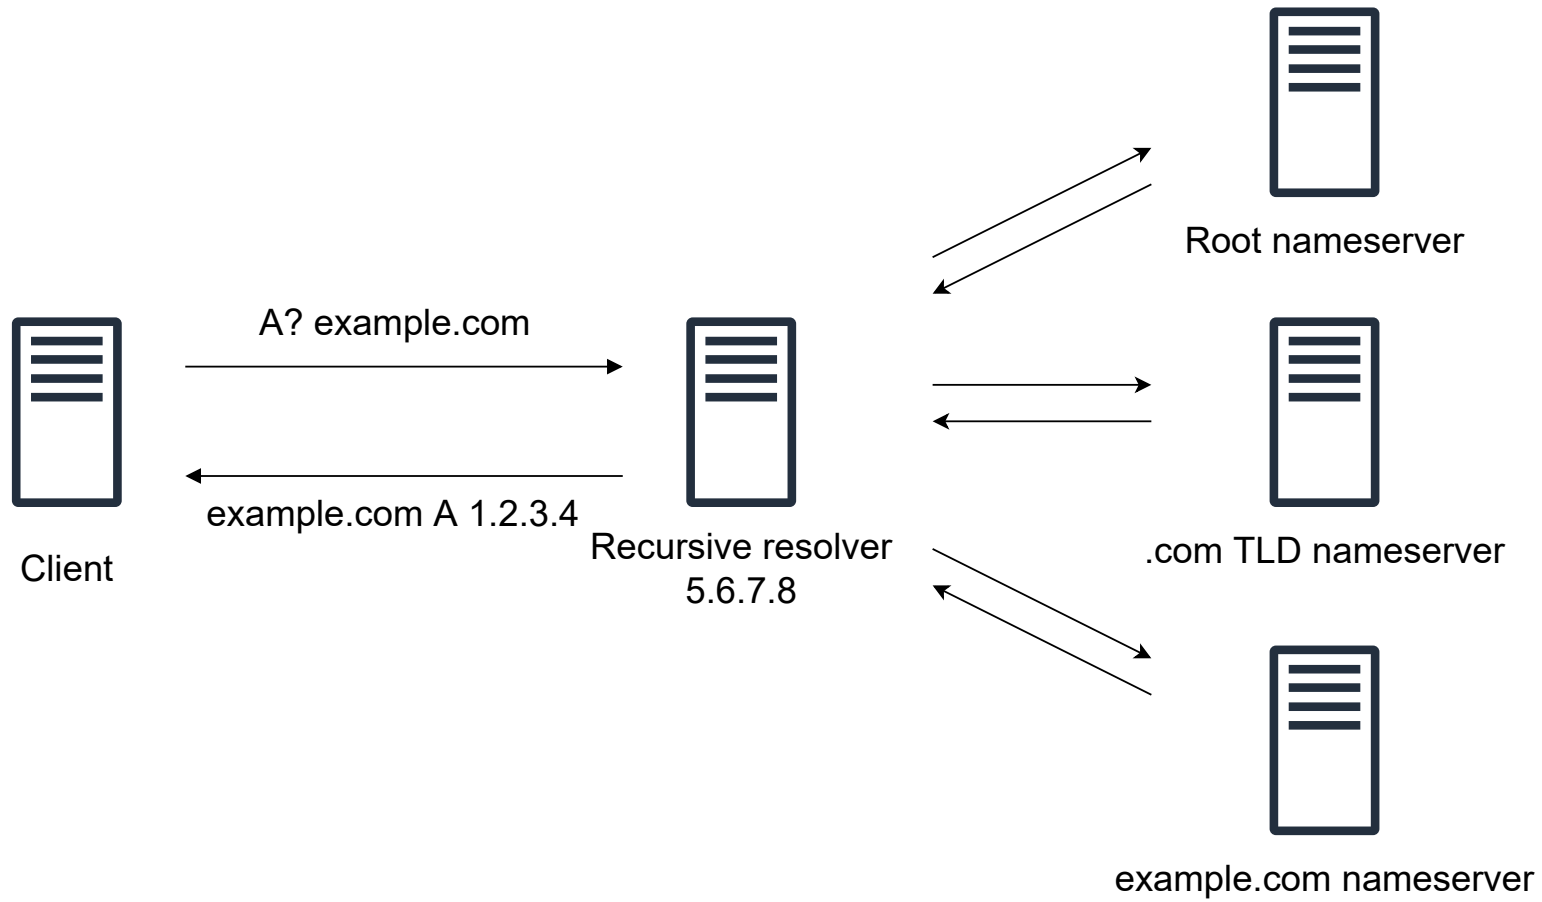
⁴ *RIPE NCC*

⁵ *ICANN*

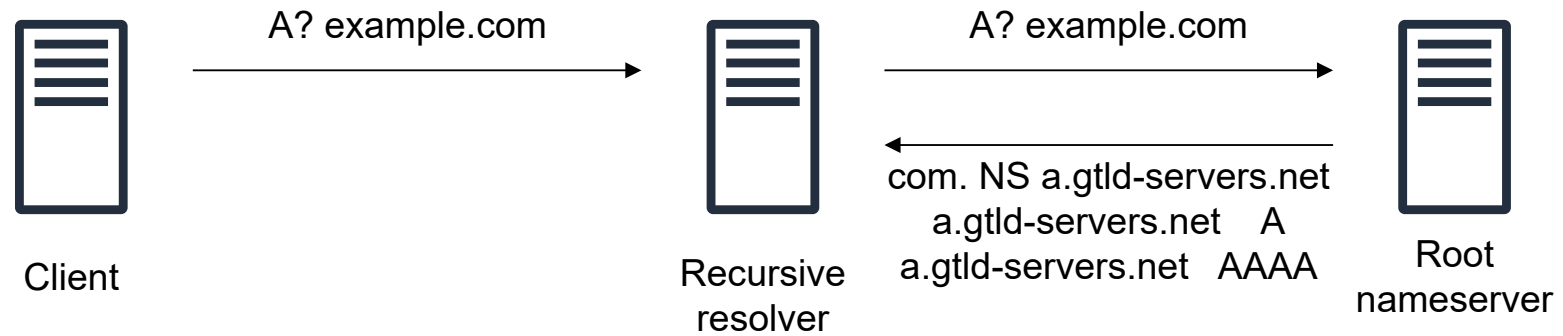
OARC 41 (Da Nang, Vietnam)

6th September 2023

DNS resolution



DNS Root Servers



November 2021 k-root event

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

Manu Bretelle [chantr4 at gmail.com](mailto:chantr4@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.

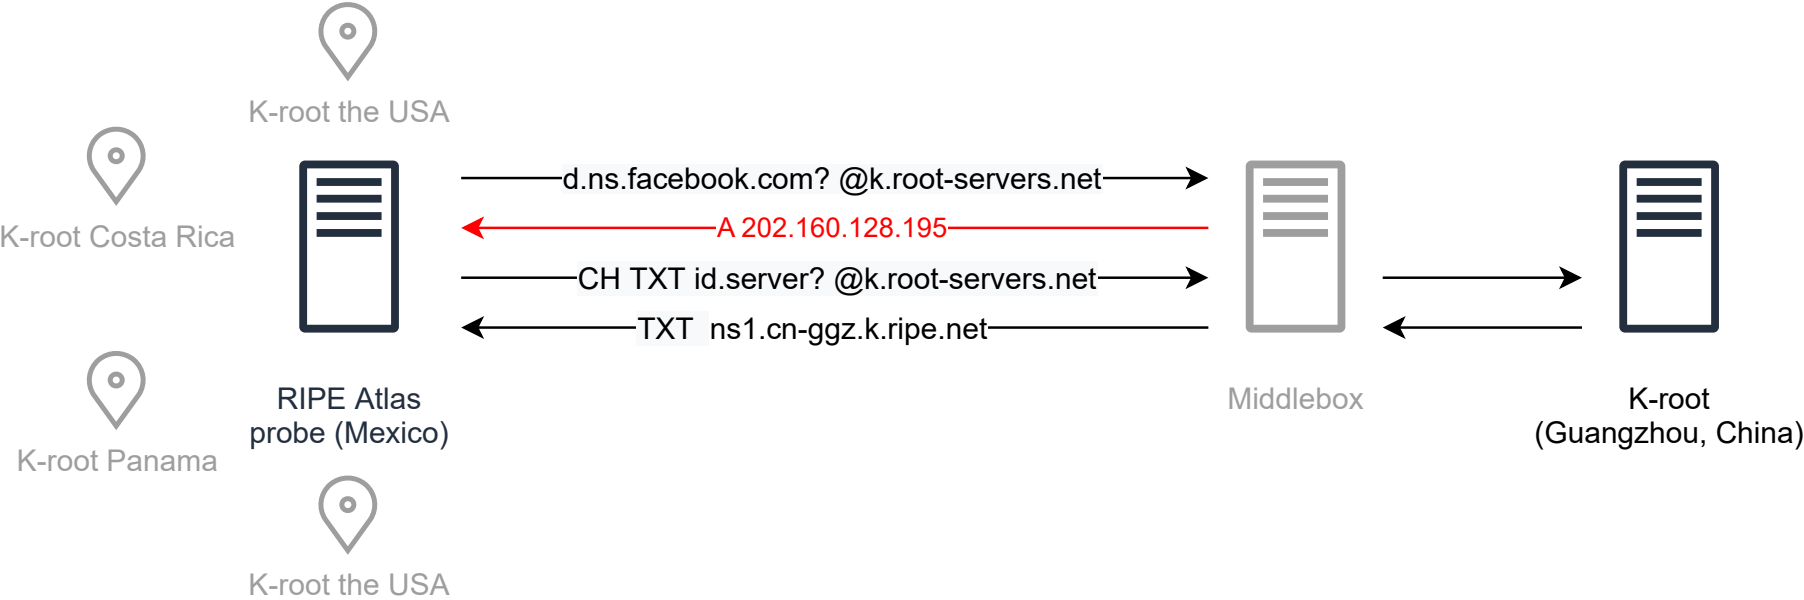
How do we ensure that those are not advertised outside of China so DNS answers are not poisoned by the GFW?

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.

November 2021 k-root event: RIPE Atlas View



November 2021 k-root event: BGP route leak

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

Anand Buddhdev [anandb at ripe.net](mailto:anandb@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.

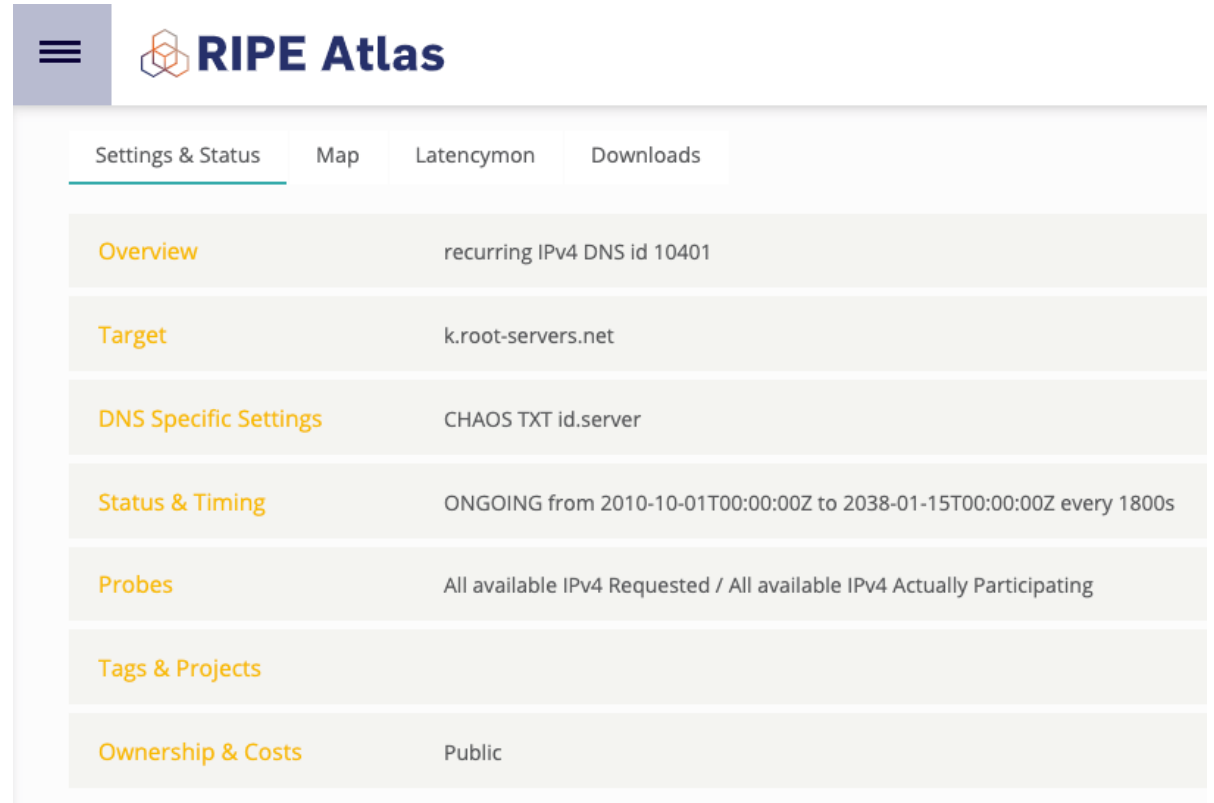
Many people have already said this, but I'd like to make it clear that the K-root server was NOT emitting false responses for Facebook and WhatsApp. The responses were being modified by something between the server and its clients.

Regards,
Anand Buddhdev
RIPE NCC

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

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)



The screenshot displays the RIPE Atlas web interface. At the top, there is a navigation bar with a hamburger menu icon and the text "RIPE Atlas". Below this, there are four tabs: "Settings & Status" (which is selected and underlined), "Map", "Latencymon", and "Downloads". The main content area shows a list of settings for a specific instance, each with a label in orange and a corresponding value:


Overview	recurring IPv4 DNS id 10401
Target	k.root-servers.net
DNS Specific Settings	CHAOS TXT id.server
Status & Timing	ONGOING from 2010-10-01T00:00:00Z to 2038-01-15T00:00:00Z every 1800s
Probes	All available IPv4 Requested / All available IPv4 Actually Participating
Tags & Projects	
Ownership & Costs	Public

**To what extent queries to
DNS root servers
experience manipulation?**

RIPE Atlas :

custom measurements

- 1b measurements (312 queries sent every 12h)
- February – October 2022
- 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, "RDLlength": 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, "RDLlength": 4, "Address": "199.59.149.244"}]}
```

Injected responses (1/5)

- Type A
- 7m responses out of 11m
- 2,419 unique IPs
- 49% of facebook.com and 89.6% of google.com responses were valid

```
{  
  "Name": "google.com.",  
  "Type": "A",  
  "Class": "IN",  
  "TTL": 235,  
  "RDLlength": 4,  
  "Address": "216.58.208.142"  
}
```

Injected responses (2/5)

- Type AAAA
- 4m responses out of 11m
- 3,221 unique IPs
- 64.4% of facebook.com and 98.3% of google.com responses were valid

```
{  
  "Name": "google.com.",  
  "Type": "AAAA",  
  "Class": "IN",  
  "TTL": 174,  
  "RDLlength": 16,  
  "Address": "2607:f8b0:4007:80a:0:0:0:200e"  
}
```

Injected responses (3/5)

- Type URI
- 42.5k responses out of 11m
- Received by 15 probes from Iran
- “0a2224” Rdata for facebook.com and “ef2678” Rdata for google.com

```
{  
  "Name": ".",  
  "Type": 256,  
  "Class": 256,  
  "TTL": 107008,  
  "Rdlength": 1034,  
  "Rdata": "0a2224"  
}
```

Injected responses (4/5)

- Type SOA
- 6.7k responses out of 11m
- Received by 1 probe from the USA

```
{  
  "Name": "facebook.com.",  
  "Type": "SOA",  
  "Class": "IN",  
  "TTL": 30,  
  "RDLlength": 62,  
  "MasterServerName": "dns1.dnsfilter.com.",  
  "MaintainerName": "dadmin.dnsfilter.com.",  
  "Serial": 1,  
  "Refresh": 30,  
  "Retry": 30,  
  "Expire": 30,  
  "NegativeTtl": 30  
}
```

Injected responses (5/5)

- Type CNAME
- 4.5k responses out of 11m
- Always google.com to forcesafesearch.google.com
- Received by 6 probes from the USA, Spain, the Netherlands, and Russia.



```
{  
  "Name": "google.com.",  
  "Type": "CNAME",  
  "Class": "IN",  
  "TTL": 49919,  
  "RDLlength": 18,  
  "Target": "forcesafesearch.google.com."  
}
```

```
{  
  "Name": "forcesafesearch.google.com.",  
  "Type": "A",  
  "Class": "IN",  
  "TTL": 65939,  
  "RDLlength": 4,  
  "Address": "216.239.38.120"  
}
```


NSIDs of injected responses

- do not contain NSIDs of root servers
- 78% empty
- public resolvers (CloudflareDNS, OpenDNS, Quad9, Google DNS)
- filtering services (CleanBrowsing)
- unclassified

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 R. Dolmans P. Hoffman
AFNIC NLnet Labs 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
- 7% of RIPE Atlas probes from 66 countries are affected
- Injected data is not always bogus
- Transparent to end users
- May introduce collateral damage
- BGP leaks stay unnoticed

Thank you!

yevheniya.nosyk@univ-grenoble-alpes.fr