

SIDN Labs

<https://sidnlabs.nl>

February 6th, 2023

Peer-reviewed Publication

Title: Intercept and Inject: DNS Response Manipulation in the Wild

Authors: Yevheniya Nosyk, Qasim Lone, Yury Zhauniarovich, Carlos H. Gañán, Emile Aben, Giovane C. M. Moura, Samaneh Tajalizadehkhoob, Andrzej Duda, and Maciej Korczyński

Venue: In Proceedings of 2023 Passive and Active Measurement Conference (PAM2023)

DOI: TBA

Conference dates: March 21–23, 2023

Citation:

- Yevheniya Nosyk, Qasim Lone, Yury Zhauniarovich, Carlos H. Gañán, Emile Aben, Giovane C. M. Moura, Samaneh Tajalizadehkhoob, Andrzej Duda, and Maciej Korczyński. Intercept and Inject: DNS Response Manipulation in the Wild. Proceedings of the 2023 Passive and Active Measurement Conference (PAM2023), Virtual Conference, March 2023

- Bibtext:

```
@inproceedings{Nosyk23a,  
  author = {Nosyk, Yevheniya and Lone, Qasim and Zhauniarovich,  
    Yury and Gañán, Carlos H. and Aben, Emile and Moura,  
    Giovane C. M. Tajalizadehkhoob, Samaneh and Duda,  
    Andrzej and Korczyński, Maciej},  
  title = {{Intercept and Inject: DNS Response Manipulation in the Wild}},  
  booktitle = {Proceedings of 2023 Passive and Active Measurement Conference},  
  year = {2023},  
  address = {Virtual Conference},  
  publisher = {Springer},  
}
```

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 Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, Grenoble, France

`first.last@univ-grenoble-alpes.fr`

² TU Delft, Delft, The Netherlands

³ SIDN Labs, Arnhem, The Netherlands

⁴ RIPE NCC, Amsterdam, The Netherlands

⁵ ICANN, Los Angeles, California, USA

Abstract. DNS is a protocol responsible for translating human-readable domain names into IP addresses. Despite being essential for many Internet services to work properly, it is inherently vulnerable to manipulation. In November 2021, users from Mexico received bogus DNS responses when resolving `whatsapp.net`. It appeared that a BGP route leak diverged DNS queries to the local instance of the `k-root` located in China. Those queries, in turn, encountered middleboxes that injected fake DNS responses. In this paper, we analyze that event from the RIPE Atlas point of view and observe that its impact was more significant than initially thought—the Chinese root server instance was reachable from at least 15 countries several months before being reported. We then launch a nine-month longitudinal measurement campaign using RIPE Atlas probes and locate 11 probes outside China reaching the same instance, although this time over IPv6. More broadly, motivated by the November 2021 event, we study the extent of DNS response injection when contacting root servers. While only less than 1% of queries are impacted, they originate from 7% of RIPE Atlas probes in 66 countries. We conclude by discussing several countermeasures that limit the probability of DNS manipulation.

Keywords: DNS · root servers · DNS manipulation · DNS censorship · BGP route leaks

1 Introduction

The Domain Name System (DNS) [41,42] is one of the core Internet protocols. It was introduced to translate human-readable domain names (e.g., `example.com`) into IP addresses (e.g., `2001:db8::1234:5678`), but has gone far beyond this basic service. It is now a large-scale distributed system comprising millions of recursive resolvers and authoritative nameservers—the two main components of the DNS infrastructure. It was designed in a hierarchical manner so that no single entity stores the data about the entire domain name space. Each authoritative

nameserver is only responsible for a subset of the domain tree and it is the role of recursive resolvers to follow the chain of delegations and find authoritative query responses.

Nonetheless, DNS is prone to manipulation. The original specification does not ensure data integrity or authentication, allowing on-path entities (Internet Service Providers, national censors, attackers, etc.) to intercept plain-text DNS traffic and inject responses (whether bogus or not). A latter standard, Domain Name System Security Extensions (DNSSEC) [55], provides data integrity, but its usage and deployment remain optional and far from being universal [17,10].

In November 2021, Meta engineers reported that users in Mexico were receiving bogus A records when querying `whatsapp.net` and `facebook.com` [15]. A closer look revealed that those queries were routed to the local China-located anycast instance of the `k-root`, despite having several other points of presence nearby [1]. As the Great Firewall of China (GFW) is known to inject bogus responses when detecting sensitive domains [28], Mexican users might have experienced collateral damage from DNS censorship. The `k-root` operator (RIPE NCC) later confirmed that a Border Gateway Protocol (BGP) route leak made the local root server instance globally available. This outage stemmed from a series of unfortunate events but nevertheless rendered both domains unavailable. The Internet has already seen similar events in the past [59,56] and researchers reported on detecting DNS root manipulation [32,44,21,35]. However, the prevalence of this phenomenon has not been systematically analyzed across all the root server letters over a long period of time.

In this paper, our contribution is two-fold. First, we leverage built-in RIPE Atlas [54] measurements to identify probes affected by the November 2021 route leak. We show that at least two months prior to be reported by Meta, the Chinese `k-root` instance had already been accessible from 32 autonomous systems (ASes) in 15 countries. Second, we set up a nine-month DNS measurement campaign and observe the same problem in the wild, although this time mostly over IPv6. More broadly, the DNS manipulation experienced by Mexican clients motivates us to study the extent of DNS response injection when contacting root servers. We reveal that even though less than 1% of queries are affected, they originate from 7% of probes located in 66 countries.

2 Manipulating Root DNS Traffic

2.1 Background on DNS Root Server System

DNS namespace is organized in a tree-like manner with the root zone at the very top. It is served by 13 root servers each referred to by letters “a” to “m” (`[a--m].root-servers.net.`). Despite being managed by 12 different organisations, each is contracted to provide an identical copy of the zone file, maintained by the Internet Assigned Numbers Authority (IANA). Every root server letter, in turn, can be accessed at its IPv4 and IPv6 addresses, both deployed using anycast (i.e., the same BGP prefix is announced from multiple anycast sites across

the globe). This ensures that the root DNS service remains highly available. As of October 2022, there are 1,575 anycast instances accessible either worldwide (global) or only within a limited range of networks (local) [1]. In the latter case, local network operators typically limit the propagation of BGP routes by using `NO_EXPORT` or `NOPEER` BGP community attributes [37]. DNS queries are then routed to the nearest anycast locations based on the routing tables. As BGP is latency agnostic, it may eventually map clients to instances from another continent, even when closer ones are available [43].

To assist in troubleshooting, root servers support DNS queries that identify individual anycast instances. This is achieved by one of the `CHAOS`-class `TXT` queries [18] (e.g., `id.server`, `hostname.bind`) or the `NSID` option [8]. The latter does not require issuing a separate query because the nameserver identifier provided in the `OPT` resource record is stored in the Additional section of a DNS response packet.

2.2 Previous Route Leaks

Root traffic manipulation has been previously reported twice. In both cases, a BGP route leak made China-located local instances globally available. Even though root servers themselves were legitimately run by their operators, the GFW or other interceptors were likely present in transit. In the first case in 2010, clients located in the USA and Chile had their queries for three domains (`twitter.com`, `facebook.com`, and `youtube.com`) answered with bogus IP addresses. Upon further investigation, it was found that original queries were sent towards the `i-root` instance in Beijing [59]. Similarly, in 2011 the clients in Europe and the USA were directed towards the `f-root` instance in Beijing [56], although no response injection was reported at that time.

2.3 November 2021: Mexico Event

In November 2021, `whatsapp.net` and `facebook.com` became inaccessible for some clients located in Mexico [15]—an event closely resembling those happening in 2010 and 2011. Figure 1 describes the course of events as reproduced by Meta engineers. A RIPE Atlas probe from Mexico was instructed to resolve the IP address of `d.ns.facebook.com` (one of the Meta nameservers) by contacting the `k-root` server directly. As it would not provide an authoritative answer for such a query, the injected response (`202.160.128.195`—an IP address belonging to *Twitter*) demonstrated that the DNS request was intercepted. To identify where the query was routed, Meta engineers configured the same probe in Mexico to send an `id.server` `CHAOS` `TXT` query. The response pointed to the Guangzhou instance in China (`ns1.cn-ggz.k.ripe.net`)—a *legitimate* `k-root` server as confirmed by its operator RIPE NCC. A traceroute measurement to `k-root`'s anycast IP address demonstrated that the seven penultimate hops went through AS4134—the Chinese telecommunication operator. As previously, the local root server instance should have been announced only to clients in China, but nevertheless leaked to the whole Internet.

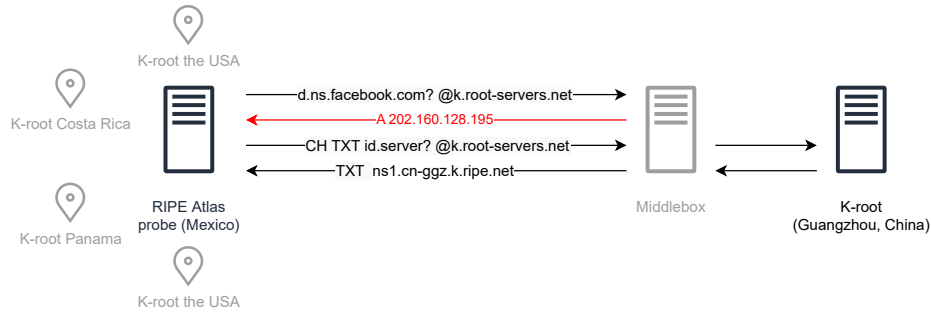


Fig. 1. DNS traffic seen on the RIPE Atlas probe from Mexico. A middlebox intercepts the DNS query and injects the bogus response for `d.ns.facebook.com`. The CHAOS-class TXT query for `id.server` confirms that the request was routed towards the Chinese instance of the `k-root`.

Extending Meta’s initial analysis, we further investigated the issue by analyzing built-in `id.server` measurements from *all* the $\sim 13\text{K}$ RIPE Atlas probes active between September 1, 2021, and November 30, 2021. We identified 57 probes from 32 ASes in 15 countries reaching the local Guangzhou instance of the `k-root` at least once prior to the leak being reported on November 6, 2021. Consequently, the instance became reachable outside China for some time before the `k-root` BGP leak was noticed and fixed in November 2021. Interestingly, we saw one probe allegedly located in the USA reaching the Guangzhou instance between September 2021 and July 2022, thus after the leak was fixed. It is only since August 2022 that the probe would be routed to the `ns1.ar-bue.k.ripe.net` site in Argentina. We referred to the historical RIPE Atlas dataset with the metadata for all the probes and confirmed that the aforementioned probe was located in China until August 2022 and was then moved to the USA. Therefore, that case does not constitute a BGP route leak.

3 Characterizing DNS Manipulation in the Wild

The events described in §2.2 and §2.3 demonstrate the cases when queries directed to certain DNS root servers resulted in response injection. In this section, we set out to characterize the extent of this phenomenon for all the root letters. We identify the response injectors and factors influencing such manipulation. We analyze more than 1 billion DNS RIPE Atlas measurements issued between February and October 2022.

3.1 Measurement Setup

Figure 2 shows the series of queries we issue on each RIPE Atlas probe, directed to all the root servers on their IPv4 and IPv6 anycast addresses. We explicitly request them not to perform recursion by setting the Recursion Desired (RD) flag

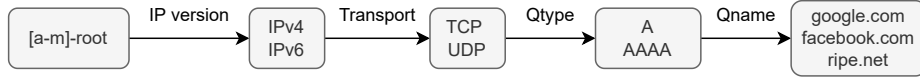


Fig. 2. DNS queries issued by each RIPE Atlas probe every 12 hours. We send queries to all root servers to resolve A/AAAA records of `google.com`, `facebook.com`, and `ripe.net` over two transport protocols (TCP/UDP) and both IP versions (IPv4/IPv6).

to false, even though correctly operating root servers would not do it anyways. Each root server is requested to resolve A and AAAA records of three domain names (`google.com`, `facebook.com`, and `ripe.net`) over TCP and UDP. The first two domains are known to trigger censorship middleboxes [48], while the third one (`ripe.net`) is a control domain. In addition, we request to include the NSID string in all the responses to learn which anycast instance (if any) answers our queries. In total, each available probe performs 312 DNS lookups every 12 hours.

3.2 The Guangzhou k-root Instance

We first check whether the local Guangzhou-located instance of the `k-root` was still reachable outside China after the BGP leak was fixed. We experimentally verified and confirmed with RIPE NCC that the server’s `id.server` string is identical to the NSID identifier (`ns1.cn-ggz.k.ripe.net`). We then extracted the latter from all the DNS responses received during the 9 months of our measurements. As expected, the instance was mainly accessible locally, but 12 probes outside mainland China (from Russia, Israel, Mexico, Denmark, and Hong Kong) would also reach it. We then verified that those probes had `k-root` sites in neighboring countries and most of the time would be routed there. For example, the probe from Hong Kong would be routed to the `ns1.vn-han.k.ripe.net` instance in Vietnam, while the probe from Russia would go to Russia itself (`ns1.ru-led.k.ripe.net`) and `ns1.kz-pwq.k.ripe.net` in Kazakhstan. Yet, during one day both probes were exceptionally reaching the Guangzhou-located instance of the `k-root`. We shared these findings with the `k-root` operator and they explained that sporadic route leaks do occasionally happen, but are then fixed quickly.

Interestingly, 11 probes out of 12 would reach the local China-located instance of the `k-root` over IPv6. Those probes received *bogus* responses for `facebook.com` queries containing, among others, IP addresses of Dropbox and Twitter. In IPv4, only one probe would reach the Chinese instance of the `k-root`. However, the event lasted a single day and no injected response was received. Overall, DNS injection is not persistent—the 12 aforementioned probes would occasionally reach the local `k-root` instance in Guangzhou without receiving rewritten responses even for sensitive domains.

Table 1. The number of injected responses per domain name and response type.

Domain	A	AAAA	URI	SOA	CNAME
ripe.net	726,993	486,111	0	0	0
google.com	3,573,083	2,156,859	13,512	0	4,536
facebook.com	2,730,047	1,456,150	29,065	6,687	0
Total:	7,030,123	4,099,120	42,577	6,687	4,536

3.3 Injected Responses

We now analyze all the DNS responses received on probes when sending queries to the 13 root servers. Recall that root servers do not directly answer queries for second-level domains, such as `example.com`. Instead, they point to authoritative DNS servers of top-level domains. Therefore, we refer to each measurement result as either *non-injected* (the answer section of the DNS response is empty) or *injected* (the answer section contains the response). In the collected dataset, over 9M responses (0.82%) were *injected* and contained more than 11M individual resource records of different types. Table 1 presents the response types received per domain name:

- **A:** maps a domain name to an IPv4 address—the most common response type received on 1,005 probes. As expected, the two sensitive domain names (`facebook.com` and `google.com`) triggered significantly more injected responses than `ripe.net`. The distribution of the returned IPv4 addresses is diverse—29 were private, 517 were from 120 globally routable ASes, and 1.8K belonged to Google and Facebook. Globally routable IPs are known to be injected by national censors [28] so that it complicates the detection of injection. Interestingly, 49% of `facebook.com` and 89.6% of `google.com` responses contained valid IP addresses of Facebook and Google, respectively. Therefore, the response injection did not prevent access to those domain names at the DNS level.
- **AAAA:** maps a domain name to an IPv6 address, received on 678 RIPE Atlas probes. Similarly to aforementioned **A** type responses, `google.com` and `facebook.com`—the two sensitive domains—experience significantly more injection than `ripe.net`. Overall, injected responses contained 3.2K unique IPv6 addresses—1.7K from reserved IP address ranges, 3 globally routable (belonging to Russian, German, and American telecommunication operators), and the remaining 1.4K addresses from Facebook and Google. The ratio of valid responses for sensitive domains was even higher than in IPv4—98.3% of `google.com` and 64.4% of `facebook.com` responses were correct. Once again, despite being injected, the majority of the DNS responses contained genuine IP addresses of the requested domains.
- **URI:** maps domain names to Uniform Resource Identifiers. We located 15 probes in Iran that received URI resource records in response to `google.com` and `facebook.com`. The responses contained 6-character strings (one for each domain name), but did not appear to be meaningful for external observers.

- **SOA**: contains administrative information about a DNS zone. One probe from the USA received SOA records for `facebook.com` queries as if they came from Facebook’s authoritative nameservers. However, the nameserver and maintainer names revealed the true originator—a DNS content filter [20]. As no valid IP address of `facebook.com` was returned, end users would not be able to access the domain name.
- **CNAME**: maps one domain name to another. We found 4,536 aliases that pointed `google.com` to `forcesafesearch.google.com`—the service [25] to exclude explicit content (e.g., pornography, violence) from search results. It is configured by adding a CNAME record to local DNS configurations. All the six affected probes (located in Spain, the USA, the Netherlands, and Russia) received corresponding A or AAAA resource records along with CNAMEs. Apart from one probe that received a bogus IP address, others would still access `google.com`, although some parts of search results would be filtered.

Overall, DNS injection impacted only 0.82% of all the queries issued during nine months in 2022. Figure 3 further demonstrates that the weekly ratio of response injection never exceeded 1%, yet proving that it is constantly present in the wild. Interestingly, response injection does not necessarily prevent access to requested domains at the DNS level - the majority of all the injected responses were not bogus. Therefore, if not coupled with other filtering techniques, such as HTTP(S) interception or destination IP address blocking, DNS manipulation would stay transparent to end users.

3.4 Identifying Injectors

The injected responses demonstrate that DNS queries originated from RIPE Atlas probes must have encountered middleboxes on the way to root servers. Such devices were shown to serve different purposes. Transparent forwarders [45] only relay incoming DNS requests to alternative resolvers, such as public or network’s internal DNS resolvers. Importantly, they do not inject spoofed responses, but rather let those alternative resolvers respond to end clients directly. More intrusive DNS interceptors, such as national censors, impersonate intended query destinations and actively inject bogus responses [38]. DNS interception is often accomplished at Customer Premises Equipment [51] and can be detected by issuing CHAOS-class or other DNS queries with the NSID option enabled.

We thus leverage the nameserver identifier option (NSID) to fingerprint services that provided responses to RIPE Atlas probes. We extracted and manually analyzed more than 12K unique NSID strings from over 1 billion measurements. We consulted the web pages of root server operators, online documentation, and issued additional DNS queries to validate our assumptions. We then generated regular expressions to match each identified service. Finally, we contacted root server operators and six of them (including Verisign that manages `a-root/j-root` with the same pattern) responded confirming the validity of our mappings.

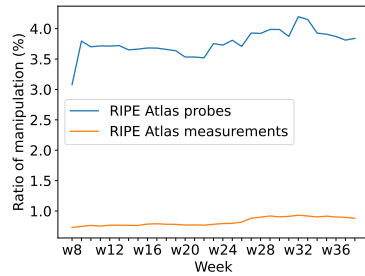


Fig. 3. Ratio of probes and measurements experiencing response injection per week.

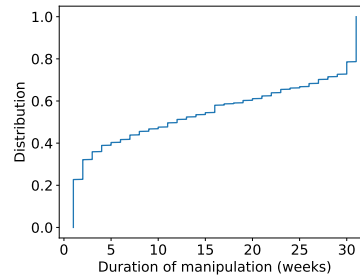


Fig. 4. Distribution of probes by the duration of DNS response manipulation in weeks.

Table 2 in Appendix provides the list of services and examples of corresponding nameserver identifiers. Apart from root server instances, we got responses from 4 public DNS providers (Cloudflare DNS, Google DNS, OpenDNS, Quad9), one DNS filtering service (CleanBrowsing), and other unknown entities. In most cases, we could deduce service names (e.g., `cmh1.groot`), airport codes (e.g., `a1.us-mia.root`), and city/country codes (e.g., `ua-kiv-aa`) from returned NSID strings. Other identifiers, e.g., `114m93` for Cloudflare DNS, could not be easily mapped to originators. We had to additionally issue queries directly to predefined addresses such as `1.1.1.1` (the IP address of Cloudflare public resolver) to prove the hypothesis.

As expected, none of the injected responses contained valid nameserver identifiers of root servers—78% of NSIDs were empty and the remaining ones included public resolvers, DNS filtering, and other undetermined services. Naturally, valid (i.e., empty) DNS responses were mostly returned by genuine root servers—they account for 95% of all the extracted NSIDs. As for the remaining 5%, we assume the presence of transparent forwarders. Additionally, we may have encountered middleboxes serving the root zone locally [21], thus not possessing the identifiers of root servers. Overall, our findings show that one can rely on nameserver identifiers to understand whether queries were answered by genuine root servers.

3.5 Participating Probes

We leveraged 14,335 RIPE Atlas probes from 177 countries and 4,132 ASes. A great majority of them did not experience DNS manipulation, but a smaller fraction (1,010 or 7.05%) received injected responses—a substantial increase since 2016 when less than 1% of RIPE Atlas probes were reported to be intercepted when contacting DNS root servers [44]. We compute the fraction of affected probes per country in Table 3 and plot the results on Figure 6 in Appendix. Overall, the manipulation ratio remains low—113 countries do not host a single probe experiencing DNS injection. On the contrary, some of the other 66 countries have a significant ratio of manipulated probes - 97.1% in Iran, 83.15%

in China, 66.67% in Palestine, 50% in Yemen, and 50% in Saint Barthélemy. As for the remaining countries (see Table 3 in Appendix), the ratio of manipulated probes does not usually exceed 30%. The autonomous system distribution is much more diverse, but more than half of the ASes host only a single probe. Overall, 5.61% of ASes only host probes experiencing manipulation, 88.92% only host probes that do not, and the remaining 5.47% have probes of both types.

RIPE Atlas probes are not constantly available and may get occasionally disconnected, which makes it non-trivial to run longitudinal measurements. However, Figure 3 shows that the proportion of probes experiencing manipulation to all the participating probes per week remains stable (the corresponding proportion of measurements exhibits the same behavior). Figure 4 additionally shows that roughly 20% of probes (mostly located in Iran, China, the USA, and Russia) experienced response manipulation during all the weeks of the experiment.

We emphasize that DNS interception and injection may happen anywhere in transit between RIPE Atlas probes and DNS root servers. Therefore, we refer to countries and networks as those *hosting* probes that experience injection. We do not assume that those entities are necessarily responsible for manipulating with DNS traffic of their clients.

3.6 Factors Driving DNS Manipulation

We leverage a generalized linear mixed-effects model (see Appendix for more details) to quantify whether the following variables make certain DNS requests more prone to response injection: query type (A, AAAA), the use of a sensitive domain name (google.com, facebook.com), transport protocol (UDP, TCP), IP version (IPv4, IPv6), and the queried root server. Figure 5 in Appendix shows that three variables significantly affect the probability of getting an injected response. As expected, the queried domain name is the factor that has the highest impact. Domain names such as facebook.com and google.com are 5.99 and 4.49 times, respectively, more likely to be manipulated than ripe.net. The second factor that impacts most the probability of DNS response injection is the transport protocol used in the DNS request. Requests over UDP are 3.50 times more likely of getting manipulated than requests over TCP. Similarly, AAAA requests are 0.70 times less likely to get manipulated than A requests. We also analyze the odds of DNS response injection when querying different root servers. While some root servers present marginally statistically significant increase of the probability of getting injected response compared to the a-root, only queries to the k-root were 0.72 times less likely of getting manipulated.

3.7 Limitations

We acknowledge that the presented measurement study has certain limitations. Using two sensitive domains would not trigger all the existing censorship middleboxes. Yet, domains from other categories (e.g., gambling or adult content) would potentially put the owners of RIPE Atlas probes in danger as those could break local laws. We thus consider this limitation acceptable and suggest the

reader to interpret the reported results as a lower-bound estimation of the problem. We also note that we only study the extent of interception and injection when sending DNS queries to root servers. Other filtering mechanisms, e.g., IP-blocking, may eventually block access to domain names at the network level.

4 Countermeasures

Some techniques presented below can effectively reduce the risks of DNS injection. However, a deliberate interceptor, especially when located close to the query source, is capable of monitoring all the client activity and reacting accordingly:

- **BGP Communities:** Li et al. [36] proposed to encode geographic coordinates of anycast sites in BGP announcements so that routers could choose the closest location. If deployed by DNS root operators, the client from Mexico could have detected nearby `k-root` instances and privilege those routes over the one from China. This would eventually avoid DNS middleboxes that injected bogus responses. While it is relatively easy to include coordinates in BGP announcements, many routers worldwide would require additional configuration to parse those geographic hints.
- **DNS Query Name Minimisation:** recursive resolvers usually include a full query name in requests to authoritative nameservers [26]. While in some cases it could reduce the total number of packets sent, query names may trigger keyword-based filtering. The client from Mexico could have issued a minimal necessary request (`.net` instead of `whatsapp.net`) [14] to the `k-root` server located in China—it would be still intercepted, but it might have not triggered keyword-based response injection. Queries to root servers can be avoided altogether when serving the root zone locally [34]. At some point, the resolver will be forced to issue a DNS request with a full query name, but if the path to the authoritative nameserver is not under interception, the response injection will not happen.
- **Encrypted DNS:** making DNS exchanges private effectively hides the communication from on-path observers. Such techniques (DNS-over-TLS [30], DNS-over-QUIC [31], and DNS-over-HTTPS [29]) are getting slowly but steadily adopted [40,6] between end clients and recursive resolvers. Yet, securing the communication link between resolvers and authoritative nameservers is still a work in progress [24]. Once standardized, it will effectively prevent middleboxes from sniffing plain-text domain names and injecting responses, as it happened for the client in Mexico. While DNS encryption is a promising technique, note that unsolicited TCP connections can be trivially torn down with injected RST packets [12], provided encrypted DNS is using a dedicated well-known port (e.g., port 853 for DoT and DoQ).
- **DNS Security Extensions (DNSSEC):** originally designed to fight cache poisoning attacks, DNSSEC [55] helps ensure that received responses are genuine. While it is an effective mechanism to detect bogus responses (e.g., a Twitter IP address returned in response to `facebook.com` query), it requires

the domain names in question to be cryptographically signed and recursive resolvers to be able to perform validation—the two criteria not met for Mexican clients. Generally, DNSSEC signing and validation are far from being deployed universally [17,10].

5 Ethical Considerations

Measurement research must be designed extremely carefully so that it minimizes any risk for involved parties but maximizes the probable benefits, as outlined in The Menlo Report [9]. This is especially important in censorship studies, which usually involve actively generating traffic to trigger censors. A rich body of research [57,50,58,47,48] performed experiments similar to ours, in particular using RIPE Atlas infrastructure [13,2], and reported that no evidence suggested that any harm was caused. We further received a formal approval from the institutional review board (IRB) of our institution. They judged our research as the one complying with all the ethical requirements.

Our choice of the measurement platform was dictated by several reasons. RIPE Atlas is an opt-in service where all the participants accept the Terms and Conditions [52]. In particular, probe hosts agree that i) the permission to install probes was obtained (§5.1), ii) other users can perform measurements on probes, in particular for research (§5.4, §4.5), iii) probes may be disabled one month after a written request is received (§8.4), and iv) measurement results will be made public either fully or in the aggregated form (§4.2). Our measurements comply with these terms and in this paper, we do not expose sensitive information about individual probes, even when allowed (§4.3).

All the RIPE Atlas probes regularly perform a set of built-in measurements [53] for different protocols. More than half of 242 recurring DNS measurements (running every 4 minutes to 12 hours) are destined to root servers. Moreover, each probe is also requesting A records of popular domain names every 10 minutes. Consequently, the traffic we generate for this experiment does not stand out from the normal operation of RIPE Atlas probes. Two out of three domain names that we query, namely `google.com` and `facebook.com`, are the first and the third most popular worldwide respectively [49]. Thus, queries to such domains are challenging to link to particular end hosts when observed in the wild.

6 Related Work

DNS middleboxes were previously known to interfere with root DNS traffic. In 2013, Fan et al. [21] reported that 1.75% of 64K vantage points worldwide would encounter DNS proxies, rogue root servers, and other unusual behavior when sending requests to the `f-root`. Moreover, some of the queries for `www.facebook.com` would be answered directly. Jones et al. [32] further formalized the phenomenon as *DNS root manipulation*. Measurements towards the `b-root` from 8K RIPE Atlas probes revealed 10 DNS proxies and one root server replica in China. Moura et al. [44] identified 74 RIPE Atlas probes (less than 1%

of all the 9K probes at that time) that would have root server queries answered by third parties. More recently, Li et al. [35] found that some of the queries originated inside China to **k-root** root server instances located inside the country would also result in hijacking. In our paper, we present a nine-month longitudinal study that characterizes the extent of DNS response injection across all the root server letters. We show that the ratio of affected probes has significantly risen compared to previous studies.

More broadly, various types of DNS manipulation have been extensively studied in the literature. Censors [57,50,58,47,48,7,23,27,46], transparent forwarders [45,33], rogue DNS servers [19], and middleboxes [51,38,61,60,16]—all interfere with the normal DNS resolution process. Particular attention has been paid to the GFW of China [3,4,5,11,22,28,39], known to intercept DNS traffic and inject bogus responses. Hoang et al. [28] provided the most complete picture of DNS manipulation by the GFW to date. The authors issued queries for 534M domain names from outside China to controlled servers inside the country. Their measurements triggered the GFW, suggesting that it indeed operates on the traffic coming from outside. As witnessed during the November 2021 route leak, middleboxes were shown to inject globally routable IP addresses in their bogus responses. These findings are in line with the previous study of the anonymous researcher [3], showing that the GFW is acting on the traffic that is barely traversing Chinese ASes. Overall, the existing research suggests that clients from Mexico were affected by the operation of the GFW.

7 Conclusions

In this paper, we explored the November 2021 BGP route leak that resulted in DNS response injection as a side effect. We identified 32 ASes worldwide that would reach the Guangzhou instance of the **k-root** and potentially encounter injecting middleboxes on the way. While this particular problem was quickly fixed, our longitudinal measurements revealed that DNS injection is omnipresent. Queries to DNS root servers are constantly getting intercepted and may result in injected responses, especially when involving sensitive domain names.

We also revealed that the Guangzhou **k-root** instance became reachable outside mainland China several months before it was reported. As it is crucial to identify such events early enough before the impact on end users becomes apparent, RIPE NCC deployed BGP community attributes identifying each **k-root** server instance. Therefore, such leaks are now detectable.

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).

References

1. Root Server Technical Operations Association. <https://root-servers.org> (2022)
2. Anderson, C., Winter, P., Ensafi, R.: Global Censorship Detection over the RIPE Atlas Network. In: USENIX FOCI (2014)
3. Anonymous: The Collateral Damage of Internet Censorship by DNS Injection. SIGCOMM Comput. Commun. Rev. **42**(3) (Jun 2012)
4. Anonymous: Towards a Comprehensive Picture of the Great Firewall’s DNS Censorship. In: USENIX FOCI (2014)
5. Anonymous, Niaki, A.A., Hoang, N.P., Gill, P., Houmansadr, A.: Triplet Censors: Demystifying Great Firewall’s DNS Censorship Behavior. In: USENIX FOCI (2020)
6. APNIC: Encrypted DNS World Map. <https://stats.labs.apnic.net/edns> (Jan 2023)
7. Arturo Filastò and Jacob Appelbaum: OONI: Open Observatory of Network Interference. In: USENIX FOCI (2012)
8. Austein, R.: DNS Name Server Identifier (NSID) Option. RFC 5001 (2007)
9. Bailey, M., Kenneally, E., Maughan, D., Dittrich, D.: The Menlo Report. IEEE Security & Privacy **10**(02) (Mar 2012)
10. Bayer, J., Nosyk, Y., Hureau, O., Fernandez, S., Paulovics, I., Duda, A., Korczyński, M.: Study on Domain Name System (DNS) abuse : technical report. Appendix 1. Publications Office of the European Union (2022). <https://doi.org/doi/10.2759/473317>
11. Bhaskar, A., Pearce, P.: Many Roads Lead To Rome: How Packet Headers Influence DNS Censorship Measurement. In: USENIX Security (2022)
12. Bock, K., Alaraj, A., Fax, Y., Hurley, K., Wustrow, E., Levin, D.: Weaponizing Middleboxes for TCP Reflected Amplification. In: USENIX Security (2021)
13. Bortzmeyer, S.: DNS Censorship (DNS Lies) As Seen By RIPE Atlas. https://labs.ripe.net/author/stephane_bortzmeyer/dns-censorship-dns-lies-as-seen-by-ripe-atlas/ (Dec 2015)
14. Bortzmeyer, S., Dolmans, R., Hoffman, P.E.: DNS Query Name Minimisation to Improve Privacy. RFC 9156 (2021)
15. Bretelle, M.: [dns-operations] K-root in CN leaking outside of CN. <https://lists.dns-oarc.net/pipermail/dns-operations/2021-November/021437.html> (Nov 2021)
16. Chung, T., Choffnes, D., Mislove, A.: Tunneling for Transparency: A Large-Scale Analysis of End-to-End Violations in the Internet. In: IMC (2016)
17. Chung, T., van Rijswijk-Deij, R., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C.: A Longitudinal, End-to-End View of the DNSSEC Ecosystem. In: USENIX Security (2017)
18. Conrad, D.R., Woolf, S.: Requirements for a Mechanism Identifying a Name Server Instance. RFC 4892 (2007)
19. Dagon, D., Lee, C., Lee, W., Provos, N.: Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority. In: NDSS (2008)
20. DNSFilter: DNS Threat Protection. <https://www.dnsfilter.com> (2022)
21. Fan, X., Heidemann, J., Govindan, R.: Evaluating Anycast in the Domain Name System. In: IEEE INFOCOM (2013)
22. Farnan, O., Darer, A., Wright, J.: Poisoning the Well: Exploring the Great Firewall’s Poisoned DNS Responses. In: WPES (2016)

23. Gill, P., Crete-Nishihata, M., Dalek, J., Goldberg, S., Senft, A., Wiseman, G.: Characterizing Web Censorship Worldwide: Another Look at the OpenNet Initiative Data. *ACM Trans. Web* **9**(1) (Jan 2015)
24. Gillmor, D.K., Salazar, J., Hoffman, P.E.: Unilateral Opportunistic Deployment of Encrypted Recursive-to-Authoritative DNS. Internet-Draft draft-ietf-dprive-unilateral-probing-02, Internet Engineering Task Force (Sep 2022), work in Progress
25. Google: SafeSearch. <https://safety.google/products/#search> (2022)
26. Hilton, A., Deccio, C., Davis, J.: Fourteen Years in the Life: A Root Server's Perspective on DNS Resolver Security. In: *USENIX Security* (2023)
27. Hoang, N.P., Doreen, S., Polychronakis, M.: Measuring I2P Censorship at a Global Scale. In: *USENIX FOCI* (2019)
28. Hoang, N.P., Niaki, A.A., Dalek, J., Knockel, J., Lin, P., Marczak, B., Crete-Nishihata, M., Gill, P., Polychronakis, M.: How Great is the Great Firewall? Measuring China's DNS Censorship. In: *USENIX Security* (2021)
29. Hoffman, P.E., McManus, P.: DNS Queries over HTTPS (DoH). RFC 8484 (2018)
30. Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., Hoffman, P.E.: Specification for DNS over Transport Layer Security (TLS). RFC 7858 (2016)
31. Huitema, C., Dickinson, S., Mankin, A.: DNS over Dedicated QUIC Connections. RFC 9250 (2022)
32. Jones, B., Feamster, N., Paxson, V., Weaver, N., Allman, M.: Detecting DNS root manipulation. In: *PAM* (2016)
33. Kühner, M., Hupperich, T., Rossow, C., Holz, T.: Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In: *USENIX Security* (2014)
34. Kumari, W.A., Hoffman, P.E.: Running a Root Server Local to a Resolver. RFC 8806 (2020)
35. Li, C., Cheng, Y., Men, H., Zhang, Z., Li, N.: Performance Analysis of Root Anycast Nodes Based on Active Measurement. *Electronics* **11**(8), 1194 (2022)
36. Li, Z., Levin, D., Spring, N., Bhattacharjee, B.: Internet Anycast: Performance, Problems, & Potential. *SIGCOMM* (2018)
37. Lindqvist, K.E., Abley, J.: Operation of Anycast Services. RFC 4786 (2006)
38. Liu, B., Lu, C., Duan, H., Liu, Y., Li, Z., Hao, S., Yang, M.: Who Is Answering My Queries: Understanding and Characterizing Interception of the DNS Resolution Path. In: *USENIX Security* (2018)
39. Lowe, G., Winters, P., Marcus, M.L.: The Great DNS Wall of China. Tech. rep., New York University (2007)
40. Lu, C., Liu, B., Li, Z., Hao, S., Duan, H., Zhang, M., Leng, C., Liu, Y., Zhang, Z., Wu, J.: An End-to-End, Large-Scale Measurement of DNS-over-Encryption: How Far Have We Come? In: *IMC* (2019)
41. Mockapetris, P.: Domain names - concepts and facilities. RFC 1034 (1987)
42. Mockapetris, P.: Domain names - implementation and specification. RFC 1035 (1987)
43. Moura, G.C.M., Heidemann, J., Hardaker, W., Charnsethikul, P., Bulten, J., Ceron, J.a.M., Hesselman, C.: Old but Gold: Prospecting TCP to Engineer and Live Monitor DNS Anycast. In: *PAM* (2022)
44. Moura, G.C.M., de O. Schmidt, R., Heidemann, J., de Vries, W.B., Müller, M., Wei, L., Hesselman, C.: Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In: *IMC* (2016)
45. Nawrocki, M., Koch, M., Schmidt, T.C., Wählisch, M.: Transparent Forwarders: An Unnoticed Component of the Open DNS Infrastructure. In: *CoNEXT* (2021)

46. Niaki, A.A., Cho, S., Weinberg, Z., Hoang, N.P., Razaghpanah, A., Christin, N., Gill, P.: ICLab: A Global, Longitudinal Internet Censorship Measurement Platform. In: IEEE S&P (2020)
47. Pearce, P., Ensafi, R., Li, F., Feamster, N., Paxson, V.: Towards Continual Measurement of Global Network-Level Censorship. In: IEEE S&P (2018)
48. Pearce, P., Jones, B., Li, F., Ensafi, R., Feamster, N., Weaver, N., Paxson, V.: Global Measurement of DNS Manipulation. In: USENIX Security (2017)
49. Pochat, V., Van Goethem, T., Tajalizadehkhoob, S., Korczyński, M., Joosen, W.: Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In: NDSS (2019)
50. Raman, R.S., Stoll, A., Dalek, J., Ramesh, R., Scott, W., Ensafi, R.: Measuring the Deployment of Network Censorship Filters at Global Scale. In: NDSS (2020)
51. Randall, A., Liu, E., Padmanabhan, R., Akiwate, G., Voelker, G.M., Savage, S., Schulman, A.: Home is Where the Hijacking is: Understanding DNS Interception by Residential Routers. In: IMC (2021)
52. RIPE Atlas: Legal. <https://atlas.ripe.net/legal/terms-conditions/> (2020)
53. RIPE Atlas: Built-in Measurements. <https://atlas.ripe.net/docs/built-in-measurements/> (2022)
54. RIPE NCC: RIPE Atlas. <https://atlas.ripe.net> (2022)
55. Rose, S., Larson, M., Massey, D., Austein, R., Arends, R.: DNS Security Introduction and Requirements. RFC 4033 (2005)
56. Snabb, J.: F.ROOT-SERVERS.NET moved to Beijing? <https://seclists.org/nanog/2011/Oct/12> (Oct 2011)
57. Sundara Raman, R., Shenoy, P., Kohls, K., Ensafi, R.: Censored Planet: An Internet-Wide, Longitudinal Censorship Observatory. In: CCS (2020)
58. VanderSloot, B., McDonald, A., Scott, W., Halderman, J.A., Ensafi, R.: Quack: Scalable Remote Measurement of Application-Layer Censorship. In: USENIX Security (2018)
59. Vergara Erece, M.: [dns-operations] Odd behaviour on one node in I root-server (facebook, youtube & twitter). <https://lists.dns-oarc.net/pipermail/dns-operations/2010-March/005263.html> (Mar 2010)
60. Weaver, N., Kreibich, C., Nechaev, B., Paxson, V.: Implications of Netalyzrs DNS Measurements. In: SATIN (2011)
61. Weaver, N., Kreibich, C., Paxson, V.: Redirecting DNS for Ads and Profit. In: USENIX FOCI (2011)

Appendix

Generalized Linear Mixed-effects Model

To determine which factors make DNS queries more susceptible to manipulation, we fit a generalized linear mixed-effects model (GLMM) assuming a binomial distribution with logit link function (logistic regression) while accounting for the country-level random effects, i.e., with the response variable as a logit transformation of DNS response state: 1 (manipulated) or 0 (not manipulated). With the inclusion of country effects, we account for observable and unobservable factors specific to queries executed within a country such as state-level filtering factors, which potentially influence DNS response manipulation. We describe the

results in odds ratios, indicating the change in the odds of DNS queries getting manipulated. The modeling results are presented in Figure 5 and discussed in detail in §3.6.

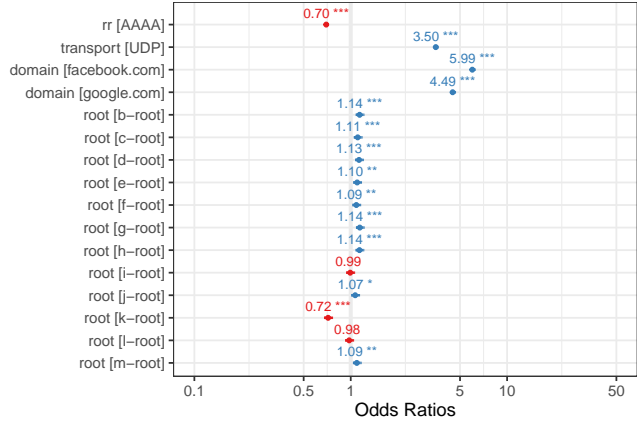


Fig. 5. Odds ratios of DNS injection survival. Values above 1 (in blue) indicate that the corresponding variables increase the chances of DNS injection, while ratios below 1 (in red) decrease the chances of DNS injection. The 95% confidence limits are delimited by horizontal lines. Those that do not cross the zero line correspond to variables that affect DNS injection more significantly.

Table 2. Services identified from 12,150 unique NSID strings.

Name	Example	Name	Example
A/J-root	a1.us-mia.root	Empty string	-
B-root	b4-iad	Unclassified	-
C-root	jfk1b.c.root-servers.org	Cloudflare DNS	114m93
D-root	jbsa4.droot.maxgigapop.net	Google DNS	gpdns-waw
E-root	p01.atlc.eroot	CleanBrowsing	CleanBrowsing v1.6a [...]
F-root	abq1f.f.root-servers.org	OpenDNS	r2.fra
G-root	cmh1.groot	Quad9	res760.qfra3.rrdns.pch.net
H-root	001.hkg.h.root-servers.org		
I-root	s1.pnh		
K-root	ns2.gb-lon.k.ripe.net		
L-root	ua-kiv-aa		
M-root	M-NRT-DIXIE-1		

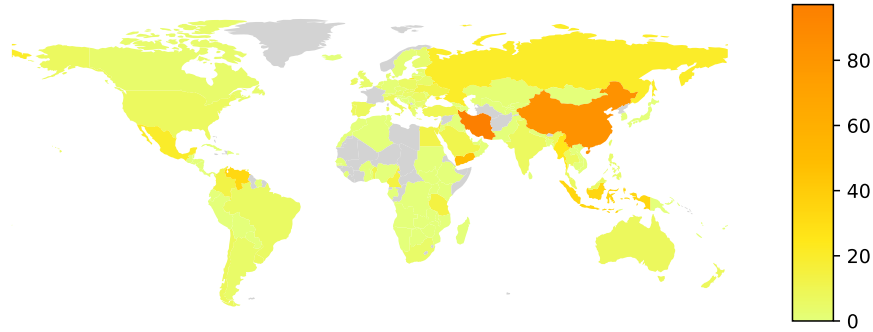


Fig. 6. The ratio (in %) of probes that experienced response injection to all the probes participating in our measurements. We did not receive any results for countries highlighted in grey.

Table 3. The ratio of RIPE Atlas probes experiencing manipulation to all those hosted in a particular country. This table only includes countries with at least one probe experiencing manipulation.

Rank	Country	Ratio (%)	Rank	Country	Ratio (%)
1.	Iran	97.1%	34.	Malaysia	6.67%
2.	China	83.15%	35.	Ireland	6.56%
3.	Palestine	66.67%	36.	Uruguay	6.25%
4.	Yemen	50.0%	37.	Poland	6.01%
5.	Saint Barthélemy	50.0%	38.	Brazil	5.77%
6.	Indonesia	34.18%	39.	Slovakia	5.45%
7.	Venezuela	33.33%	40.	Romania	5.38%
8.	Myanmar	25.0%	41.	Armenia	5.0%
9.	Russia	20.57%	42.	South Africa	4.85%
10.	Mexico	20.41%	43.	Argentina	4.76%
11.	Israel	15.38%	44.	Belgium	4.74%
12.	Albania	14.81%	45.	Lithuania	4.44%
13.	Tanzania	14.29%	46.	Hungary	4.35%
14.	Cameroon	14.29%	47.	Austria	4.17%
15.	Colombia	13.04%	48.	Canada	4.11%
16.	Egypt	12.5%	49.	The UK	4.07%
17.	Ukraine	12.0%	50.	The Philippines	3.85%
18.	Saudi Arabia	11.11%	51.	Belarus	3.85%
19.	Benin	11.11%	52.	Taiwan	3.57%
20.	Türkiye	9.59%	53.	Finland	3.38%
21.	France	8.25%	54.	Bulgaria	3.03%
22.	Hong Kong	8.16%	55.	Denmark	2.9%
23.	Portugal	7.79%	56.	Estonia	2.7%
24.	Italy	7.77%	57.	Greece	2.68%
25.	New Zealand	7.5%	58.	Panama	2.22%
26.	Chile	7.5%	59.	Luxembourg	2.08%
27.	Spain	7.46%	60.	Czechia	1.93%
28.	Australia	7.22%	61.	The Netherlands	1.91%
29.	Thailand	7.14%	62.	Singapore	1.57%
30.	Cyprus	7.14%	63.	Kazakhstan	1.39%
31.	India	7.1%	64.	Sweden	1.32%
32.	Norway	6.96%	65.	Germany	1.24%
33.	The USA	6.88%	66.	Switzerland	1.04%