How crazy is .internal/DOT

Wes Hardaker

April 27, 2026

Contents

1 TL;DR – aka Executive Summary

We used 5000 RIPE Atlas nodes to send queries to two domains under our control that exhibit various types of resolution success and failure scenarios. We recorded all result data from RIPE Atlas itself, traffic captures from the name server serving the zones, and traffic captures from the B-Root DNS root server. We performed 6 different sets of tests that included both regular requests of various existing common success and failure cases, and for a special case of a name server (NS) record with a target of ".".

None of the tests revealed a general increase in query results or abnormal responses when comparing the NS records with a target of "." to the other types of requests commonly seen on the internet today. Specifically, the introduction of a NS record with a target of "." exhibits similar traffic patterns to other common deployment errors seen today, such as a non-existent sub-zones, or sub-zones with unresolvable NS addresses.

Video presentation: if you'd prefer a video presentation of the results in this document, see the presentation I gave at IETF-125 IEPG meeting.

2 Inroduction

It has been suggested that we could create a DNS record for a child zone with a NS target of "." (redirecting back to the root). This could be a way to signal that a private name space exists below such a point which is not resolvable within the global DNS. In particular, this is being considered as one solution for how to (not) delegate the ".internal" (non) TLD to nowhere. This would allow it to be considered a DNSSEC insecure delegation, while at the same time not actually being delegated to a sub-zone.

But is this a good idea? Is it crazy enough that it just might work. Or will it cause endless resolver loops? The research described in this document aims to answer at least that last question.

3 Methodology

To study whether or not a NS record pointing back to the root (in other words, with the (only) NS target being the domain ".") causes real world problems, we created test records in two semantically similar zones under our control. We chose two similar domains under two different TLDs (.com and .games) to test if there was a behavioral difference between resolutions to more historic TLDs (.com) vs TLDs from the new gTLD round (.games). The registrant and operator of these zones is the author, providing the study with full control of the zones under test.

The test setup began with creating scenario records for under *sub-bbb.frostedaxe.com* and *sub-bbb.frostedaxe.games*, where the *sub-bbb* sub-domain in each case was reserved entirely for these tests. Only one NS record was used for each *sub-bbb* subdomain in each parent in order to capture all traffic to these domains while the test was run. We conducted 6 tests using 5000 RIPE Atlas probes sending queries to their configured (DHCP assigned) resolvers to these zones.

Note: The odd *-bbb* naming suffix was simply a distinguishing identifier to be used when processing data. Below we will generally refer to these domain names only by their *sub-bbb* prefix for brevity.

Note: though all zones are signed, the author failed to put DS records in for the sub-bbb children so they're technically insecure delegations. I may re-run it yet again.

3.1 DNS setup

The entire *sub-bbb.frostedaxe.com* domain looked like the following (without the corresponding DNSSEC material and related signatures):

```
;; hosting zone top records
sub-bbb.frostedaxe.com.          600  SOA  ns1.sub-bbb.frostedaxe.com. admin.frostedaxe.com. 202601
sub-bbb.frostedaxe.com.          600  A    107.220.113.177
sub-bbb.frostedaxe.com.          600  NS   ns1.sub-bbb.frostedaxe.com.
sub-bbb.frostedaxe.com.          600  TXT  "Frosted Axe Studios test zone by Wes Hardaker"

;; NS records for hosting zone
ns1.sub-bbb.frostedaxe.com.      600  A    107.220.113.177

;; sub-domain NS records
dnedot.sub-bbb.frostedaxe.com.   600  NS   .
dneext.sub-bbb.frostedaxe.com.   600  NS   dnens.frostedaxestudios.com.
dneint.sub-bbb.frostedaxe.com.   600  NS   dnens.frostedaxe.com.

;; real A records
exists.sub-bbb.frostedaxe.com.   600  A    107.220.113.177
```

Note: as we will see next in the testing descriptions, the 600 second TTLs value was chosen test response cache timings.

3.2 Test and RIPE Atlas setup

To study whether or not resolvers behave differently when using a NS record with a target of . vs other common situations, we conducted 6 different tests. Once the initial set of 5000 atlas nodes were selected during the first execution of the first test, the test harness requested the use of the same 5000 nodes for all subsequent tests.

In the following queries, a prefix of *probeid* indicates the use of RIPE atlas's ability to pre-prepend DNS requests with the probe's identifier as the first label or labels in a DNS request. This was done in all except the first test, which explicitly tested for an address (A) record that does exist.

1. query for an A record for a record that exists (*exists.sub-bbb*).
2. query for a record that doesn't exist (*probeid.sub-bbb*)
3. query for a record in a sub-domain that doesn't exist (*probeid.dnesub.sub-bbb*)
4. query for a record in a sub-domain with a broken in-bailiwick delegation (*probeid.dneint.sub-bbb*, where *dnsint.sub-bbb* is delegated to an NS of *dnens.frostedaxe.com* that doesn't itself exist. No glue is returned during this test either.)
5. query for a record in a sub-domain with a non-existent out-of-bailiwick delegation (*probeid.dneext.sub-bbb*, where *dneext.sub-bbb* delegates to an NS of *dnens.frostedaxestudios.com* that itself doesn't exist – this zone is also in control of the author. The nameserver for *frostedaxestudios.com* is actually the same name server, but no glue or NS records are returned in response to the *dneext.sub-bbb* query, so resolvers were forced to next resolve *frostedaxestudios.com* for its NS records.)
6. Finally, query for a record inside *probeid.dnedot.sub-bbb* where the *sub-bbb* zone contains an NS record for *dnedot.sub-bbb* with a . target.

Each of these tests was conducted 3 times in a row. Specifically, for each test, the following steps were taken:

1. a query was launched using RIPE Atlas for the domain in question
2. an identical query was launched 4 minutes later. This is within the 5 minute TTL of the *sub-bbb* records, and may have been retrievable from the resolvers cache.
3. a final query was sent at 10 minutes and 30 seconds after test 2. This is specifically after all caches should be expired.

In between each test group of three, a 10 minutes and 30 second delay was also introduced to ensure all caches were (hopefully) empty of associated data at the beginning of each new experiment. (Note that because the tests were in alternating TLDs too, resolvers already had a significant cache flushing opportunity while the tests under the alternate TLD were run.)

4 Results

Below are sub-sections containing results of:

1. Analyzing the collected PCAPs from the authoratative nameservers that host the *sub-bbb* zones.
2. Analyzing the data returned from the RIPE ATLAS probes themselves.
3. Analyzing the data gathered from the B-Root DNS root server.

In all of the following results tables and charts, the counts have been labeled with the following identifiers, which match the 6 tests discussed above plus an additional text hint identifier.

Label	Description
1-exists	records from querying a DNS record that actually exists (a real A record)
2-dnelabel	records observed from querying a domain name (left label = probeid) that didn't exist
3-dnesub	records observed from querying labels (probeid) under a non-existing sub domain
3-dnesub-NS	traffic seen querying for NS records of the subdomain label itself in the process of a resolver querying during the 3-dnesub domain
4-intbal	records observed from querying labels under a subdomain with non-existent internally named NS
4-intbal-NS	traffic seen querying for NS records of the subdomain label itself in the process of a resolver querying during the 4-intbal domain
5-extbal	records observed from querying labels under a subdomain with non-existent external named NS
5-extbal-NS	traffic seen querying for NS records of the subdomain label itself in the process of a resolver querying during the 5-extbal domain
6-DOT	records observed from querying labels under a NS record for .
6-DOT-NS	traffic seen querying for NS records of the subdomain label itself in the process of a resolver querying during the 6-DOT domain
all	records received for parent zones for all tests (not typically shown)

An important element to understand from these labels is that they all start with the test number (e.g. "2-"), followed by a simple keyword indicating the specific types of results extracted from test. The labels that end in "-NS" indicate that the query observed was for the nameserver itself of the child zone, not the record within the child zone. For example, when a query for *probeid.dnesub.frostedaxe.com.* was requested from RIPE Atlas nodes, the logged queries to the authoritative server for that exact record would be counted under *3-dnesub* but a NS query for *dnesub.frostedaxe.com.* itself would be counted separately under/3-dnesub-NS/.

Note that queries to "-NS" records were even seen as address record queries (A records) to the subdomain name even when the subdomain's NS records were different or unavailable. For example, even though the *dneint.sub-bbb.frostedaxe.com.* domain has an NS record with a target of *dnens.frostedaxe.com.*, resolvers even sent address (A) queries for the subdomain *dneint.sub-bbb.frostedaxe.com.* itself in addition to the address record it was supposedly trying to look up (such as *probeid.dneint.sub-bbb.frostedaxe.com.*) or an NS query for the subdomain (*dneint.sub-bbb.frostedaxe.com.*).

Finally, note that RIPE ATLAS probes frequently have multiple upstream resolvers (typically provided via DHCP), the packet counts can be higher than the total probe count even in near perfect conditions (IE, 5000 probes were requested, but more than 5000 queries are present in the data).

4.1 Authoritative server PCAP analysis

To get ground truth directly from the network running the authoratative nameserver supporting the test cases, we studied the PCAPs collected during the tests. We look for the address and nameserver queries collected from each test, the general total load during the test period, and the top talkers in terms of addresses and ASNs.

4.1.1 Address Requests received

The following table shows the count of address (A and AAAA records) queries received at the authoritative server for each type of expected request. The count is given in the first column, followed by the test label, and then the query type and aggregated query name (substituting *probeid* when a numerical probe identifier was actually observed in the data).

```
==== TLD: com

   39330 6-DOT             A probeid.dnedot.sub-bbb.frostedaxe.com.
   38585 4-intbal          A probeid.dneint.sub-bbb.frostedaxe.com.
   38310 5-extbal          A probeid.dneext.sub-bbb.frostedaxe.com.
   32123 2-dnelabel        A probeid.sub-bbb.frostedaxe.com.
   26454 3-dnesub          A probeid.dnesub.sub-bbb.frostedaxe.com.
   18987 5-extbal-NS       A dneext.sub-bbb.frostedaxe.com.
   16952 6-DOT-NS          A dnedot.sub-bbb.frostedaxe.com.
   16249 4-intbal-NS       A dneint.sub-bbb.frostedaxe.com.
   10232 1-exists          A exists.sub-bbb.frostedaxe.com.
    6677 3-dnesub-NS       A dnesub.sub-bbb.frostedaxe.com.
     723 4-intbal       AAAA probeid.dneint.sub-bbb.frostedaxe.com.
     340 5-extbal       AAAA probeid.dneext.sub-bbb.frostedaxe.com.
     338 2-dnelabel     AAAA probeid.sub-bbb.frostedaxe.com.
     301 6-DOT          AAAA probeid.dnedot.sub-bbb.frostedaxe.com.
     299 1-exists       AAAA exists.sub-bbb.frostedaxe.com.
     285 3-dnesub       AAAA probeid.dnesub.sub-bbb.frostedaxe.com.
       9 3-dnesub-NS    AAAA dnesub.sub-bbb.frostedaxe.com.
       9 4-intbal-NS    AAAA dneint.sub-bbb.frostedaxe.com.
       5 6-DOT-NS       AAAA dnedot.sub-bbb.frostedaxe.com.
       4 5-extbal-NS    AAAA dneext.sub-bbb.frostedaxe.com.

==== TLD: games

   51034 2-dnelabel        A probeid.sub-bbb.frostedaxe.games.
   45031 3-dnesub          A probeid.dnesub.sub-bbb.frostedaxe.games.
   25742 6-DOT             A probeid.dnedot.sub-bbb.frostedaxe.games.
   23175 5-extbal          A probeid.dneext.sub-bbb.frostedaxe.games.
   20156 4-intbal          A probeid.dneint.sub-bbb.frostedaxe.games.
   20153 4-intbal-NS       A dneint.sub-bbb.frostedaxe.games.
   19052 5-extbal-NS       A dneext.sub-bbb.frostedaxe.games.
   14794 6-DOT-NS          A dnedot.sub-bbb.frostedaxe.games.
   10063 1-exists          A exists.sub-bbb.frostedaxe.games.
    7344 3-dnesub-NS       A dnesub.sub-bbb.frostedaxe.games.
     450 4-intbal       AAAA probeid.dneint.sub-bbb.frostedaxe.games.
     380 2-dnelabel     AAAA probeid.sub-bbb.frostedaxe.games.
     380 5-extbal       AAAA probeid.dneext.sub-bbb.frostedaxe.games.
     376 6-DOT          AAAA probeid.dnedot.sub-bbb.frostedaxe.games.
     364 3-dnesub       AAAA probeid.dnesub.sub-bbb.frostedaxe.games.
      39 1-exists       AAAA exists.sub-bbb.frostedaxe.games.
       9 4-intbal-NS    AAAA dneint.sub-bbb.frostedaxe.games.
       6 3-dnesub-NS    AAAA dnesub.sub-bbb.frostedaxe.games.
       3 6-DOT-NS       AAAA dnedot.sub-bbb.frostedaxe.games.
       1 5-extbal-NS    AAAA dneext.sub-bbb.frostedaxe.games.
```

We display these also in the following bar charts, looking at the queries received by each test type for each TLD being used:

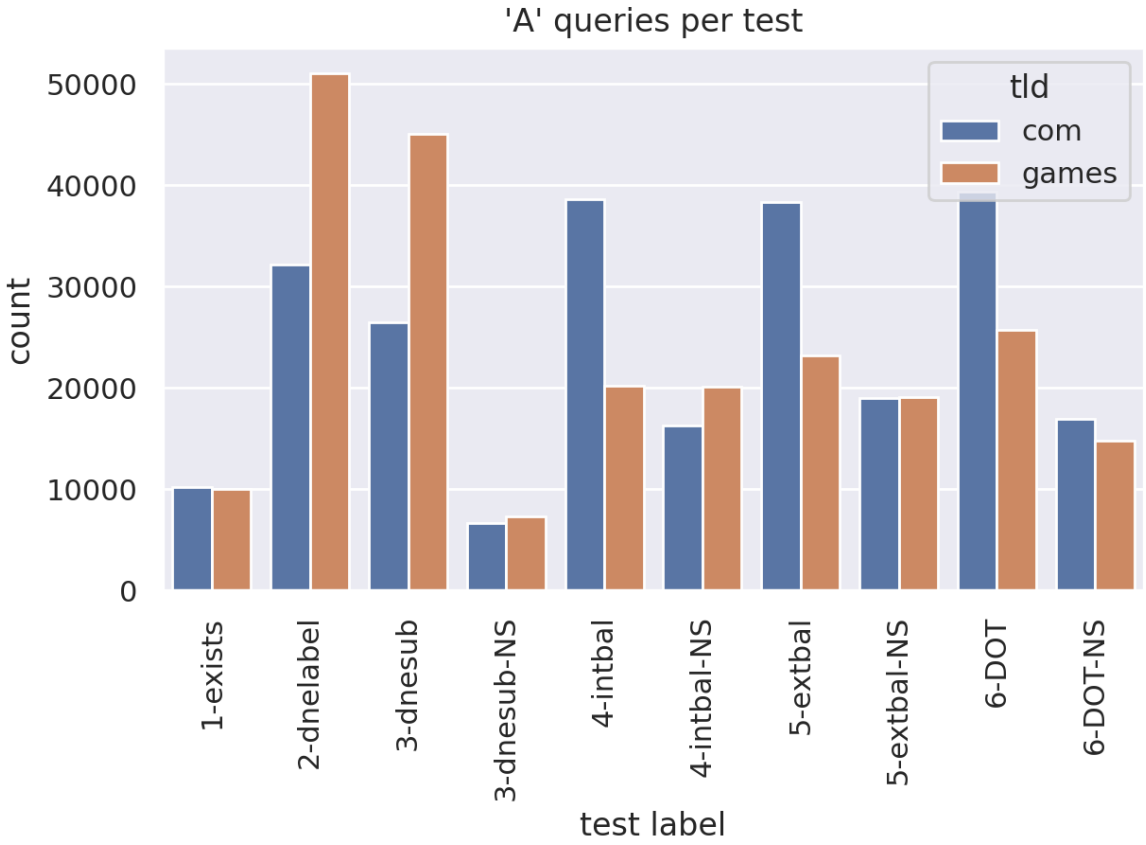

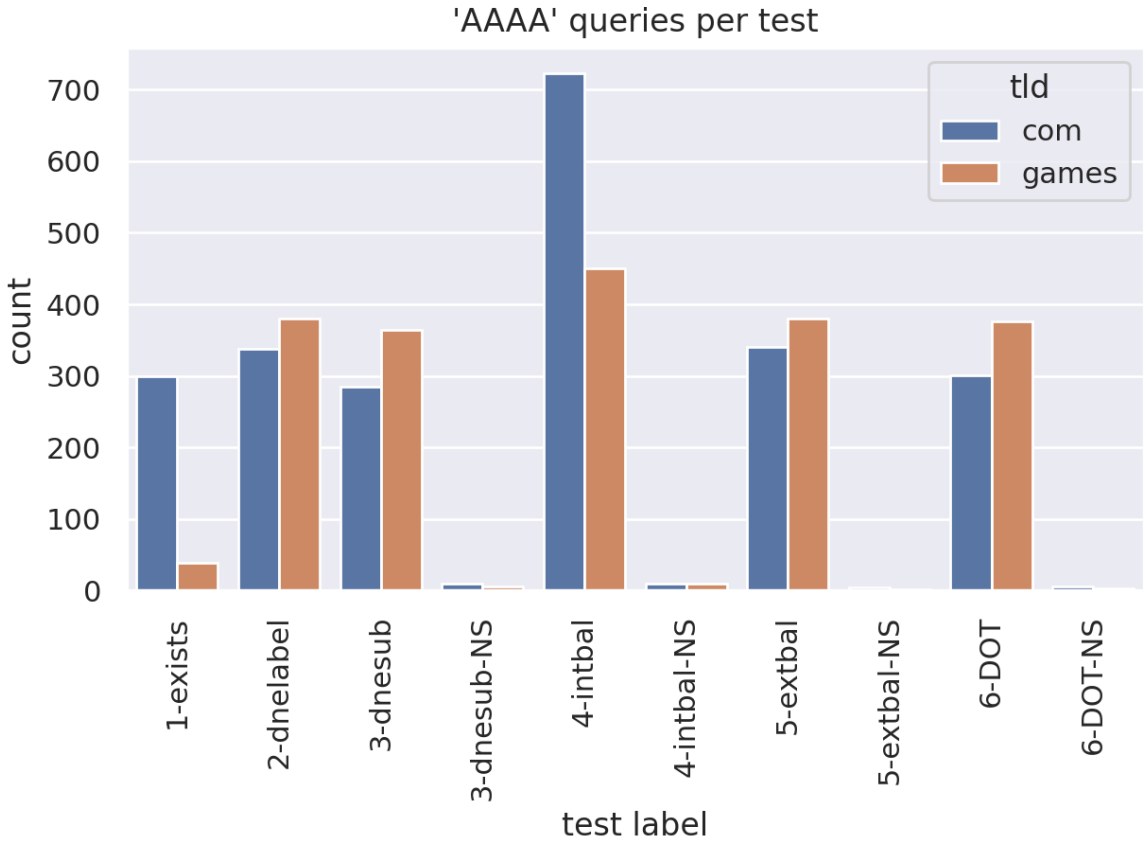

The test generating the least amount of address (A) record request traffic was for the domain name

that actually existed (*exists.sub-bbb*). The other tests generally generated more traffic, probably due to various fallback tests by the resolvers themselves and the probe's resolution process falling back to using other resolvers.

The takeaway: The test with an NS record target of "." did not cause more traffic than the other failure cases.

4.1.2 NS Requests received

The following table shows the count of NS record queries received:

```
==== TLD: com

   12031 2-dnelabel       NS probeid.sub-bbb.frostedaxe.com.
    7621 3-dnesub         NS probeid.dnesub.sub-bbb.frostedaxe.com.
    6311 5-extbal-NS      NS dneext.sub-bbb.frostedaxe.com.
    5842 3-dnesub-NS      NS dnesub.sub-bbb.frostedaxe.com.
    5659 6-DOT-NS         NS dnedot.sub-bbb.frostedaxe.com.
    5630 4-intbal-NS      NS dneint.sub-bbb.frostedaxe.com.
    3238 6-DOT            NS probeid.dnedot.sub-bbb.frostedaxe.com.
    3228 4-intbal         NS probeid.dneint.sub-bbb.frostedaxe.com.
    3000 5-extbal         NS probeid.dneext.sub-bbb.frostedaxe.com.
      88 1-exists         NS exists.sub-bbb.frostedaxe.com.

==== TLD: games

   11970 2-dnelabel       NS probeid.sub-bbb.frostedaxe.games.
    7724 3-dnesub         NS probeid.dnesub.sub-bbb.frostedaxe.games.
    6252 3-dnesub-NS      NS dnesub.sub-bbb.frostedaxe.games.
    5973 6-DOT-NS         NS dnedot.sub-bbb.frostedaxe.games.
    5939 4-intbal-NS      NS dneint.sub-bbb.frostedaxe.games.
    5756 5-extbal-NS      NS dneext.sub-bbb.frostedaxe.games.
    3778 6-DOT            NS probeid.dnedot.sub-bbb.frostedaxe.games.
    3644 5-extbal         NS probeid.dneext.sub-bbb.frostedaxe.games.
    3459 4-intbal         NS probeid.dneint.sub-bbb.frostedaxe.games.
      76 1-exists         NS exists.sub-bbb.frostedaxe.games.
```

Or graphically and sorted by test type:

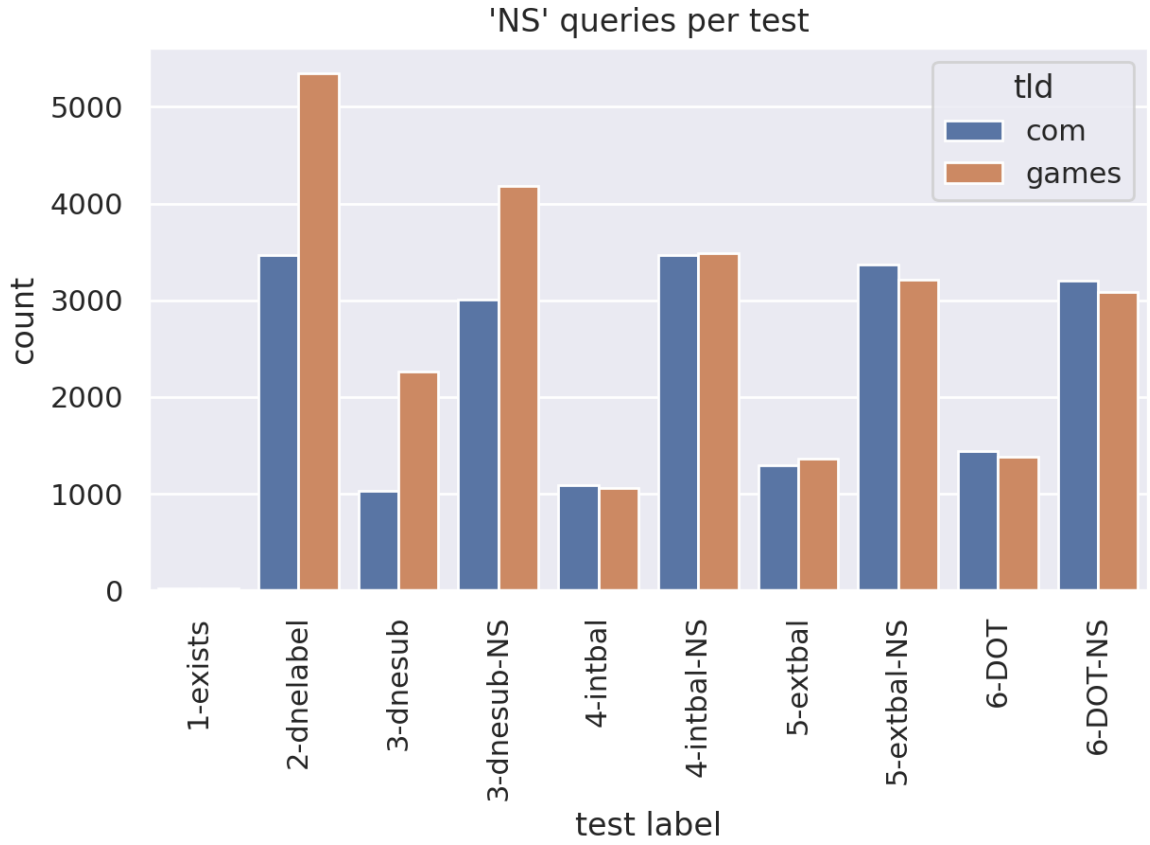

The takeaway: again we find the test for an NS with a target of "." did not generally generate more traffic than the other test cases.

4.1.3 Traffic loads over time

We also look at the traffic received over the lifetime of all of the tests, plotted as *queries per minute*:

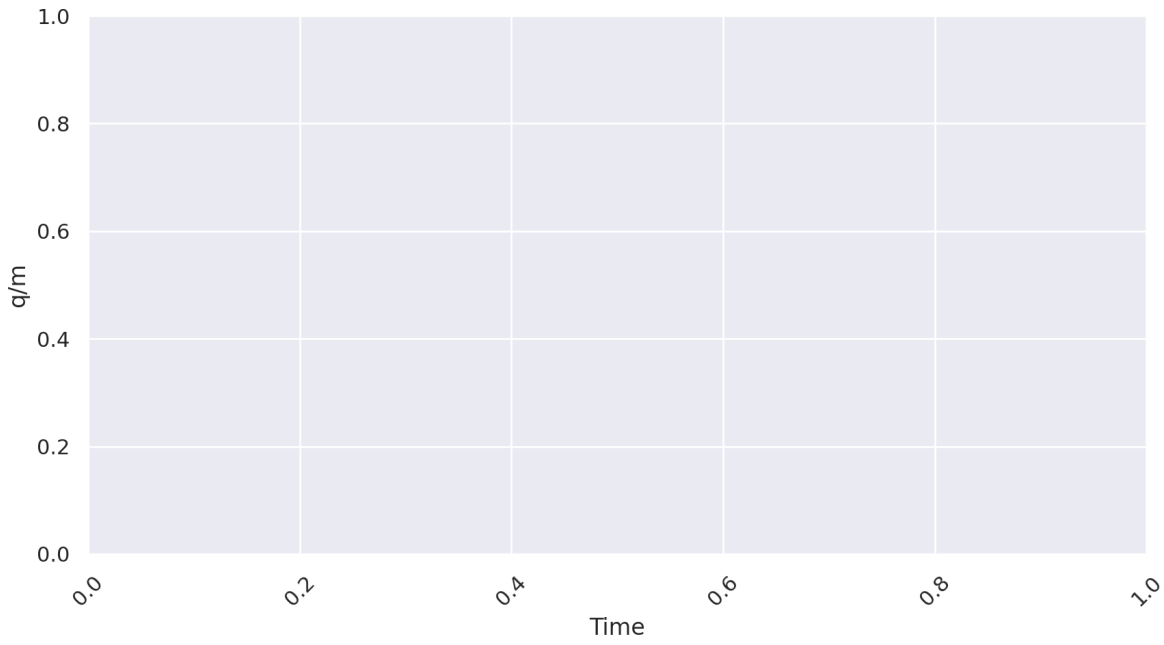

Though not labeled directly, the query rate does not spike unusually in any of the tests. And specifically, the final 6 peaks (which come from the testing of the NS target with .) do not show an increase in query traffic, alleviating our fears that a NS to "." would cause a recursion loop.

Interpretation note: if you look closely at the graphs, there is always a peak followed later by a second peak – these are the first two tests where the second should be hopefully "in-cache". This double peak is then followed by a singular peak, which is the third test run 10.5 minutes later, after the test data should have cleared from the cache again. It is unfortunate to note how high the peaks are for data that should be in cache but is queried for anyway.

4.1.4 Network owners sending the most traffic

We also studied which networks sent the most traffic. To do this we looked at the top 20 networks sending traffic, as determined by the *ip2asn* database.

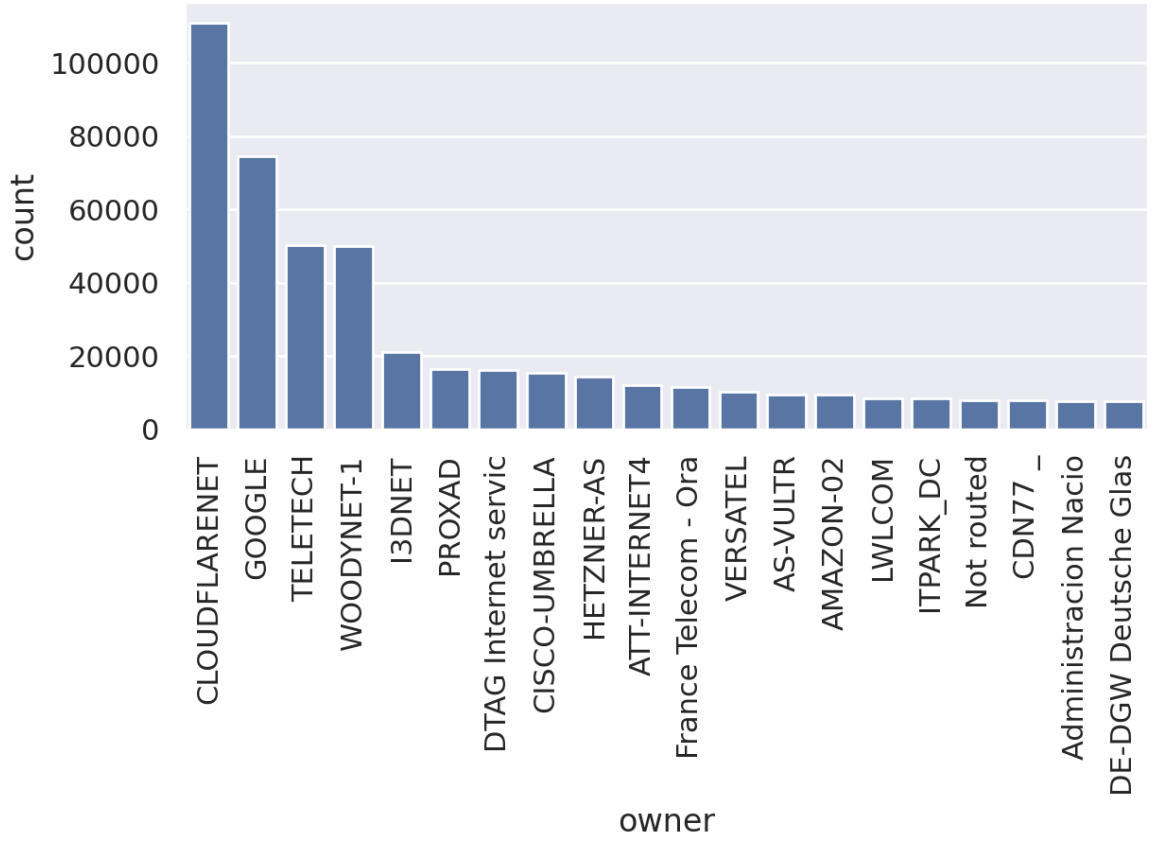

It is not surprising that the top few are all large network providers, some of which run large open resolvers that the probes may be using. It is unclear, however, why TELETECH is so large.

4.1.5 Network owners graphed over time

The following graph shows the traffic from the top 10 networks and the traffic they sent over the lifetime of the tests.

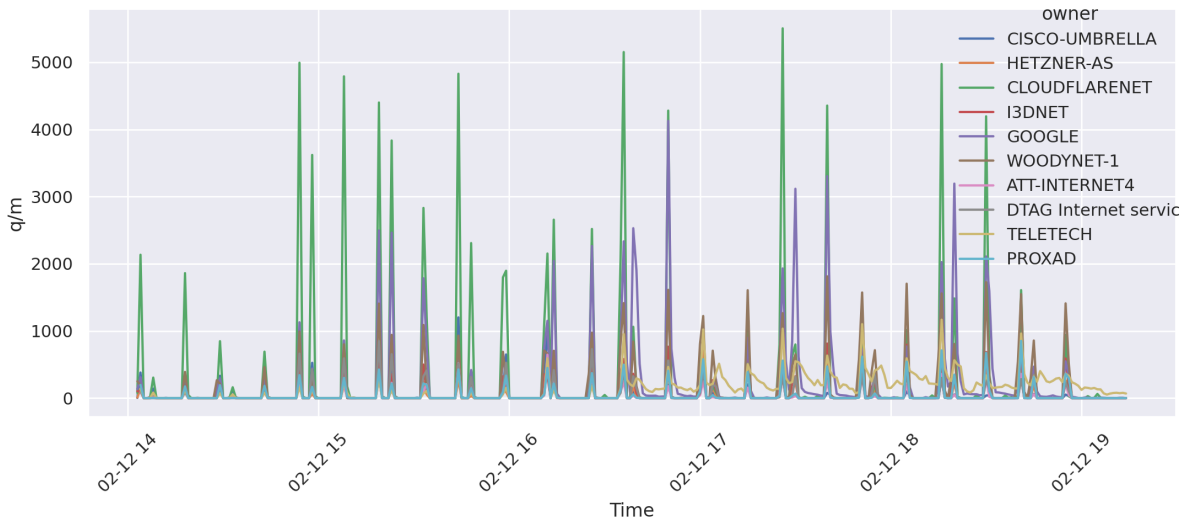

The traffic per source owner looks similar in effect to the previous time-aggregated bar graph. It is also similar to the above graph of total traffic, showing that the largest number of sources did not significantly increase in traffic sent during the NS to . testing.

4.1.6 Graphing the top src IPs observed

We next look at the top 10 individual srcips that sent queries to the test zone's name server. We also plot the top 3 just for better visibility.

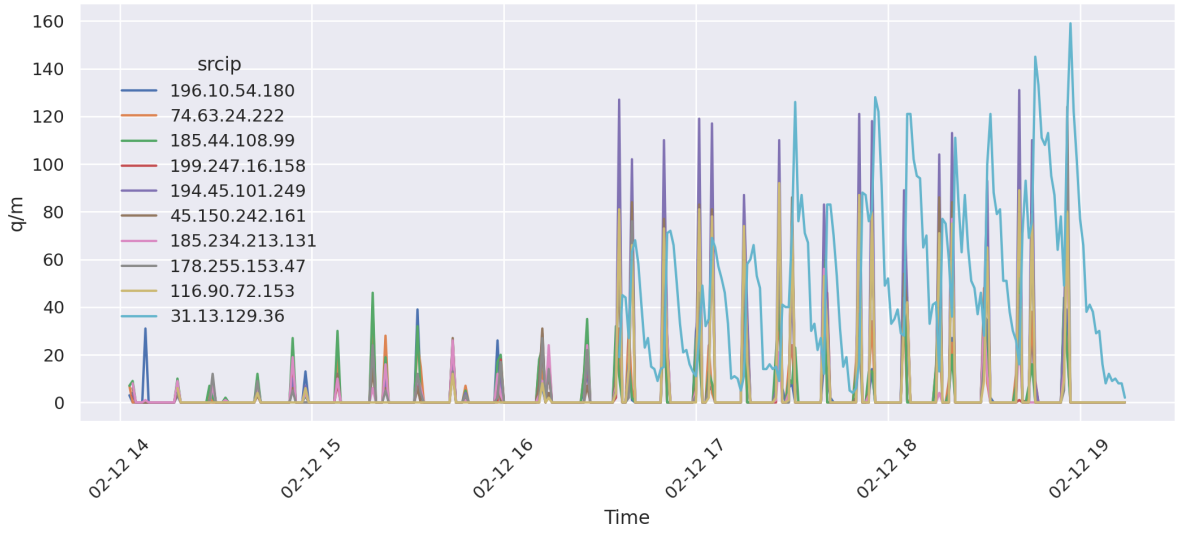

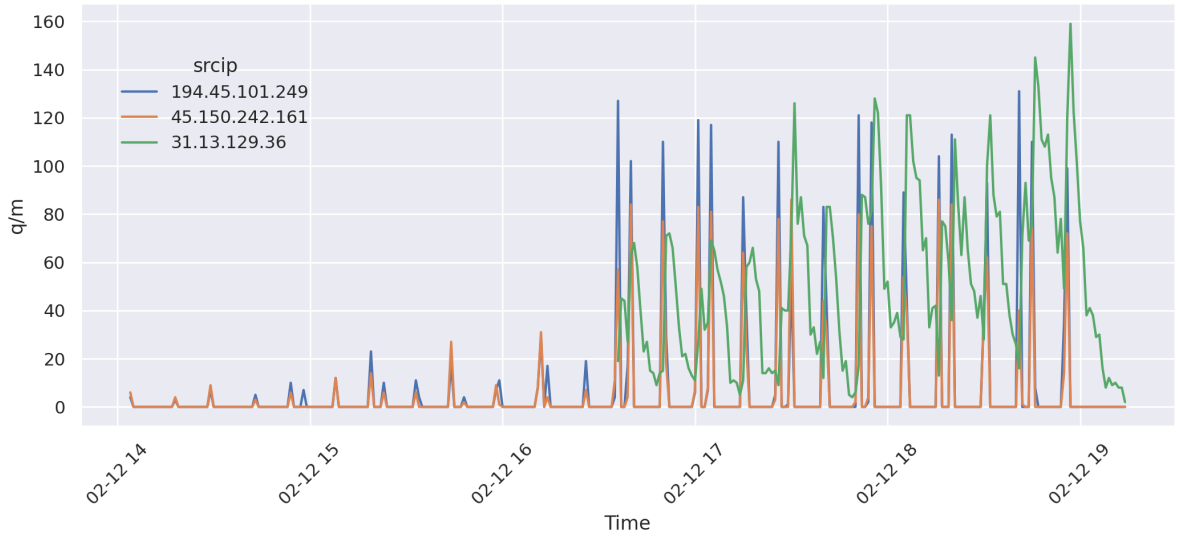

Takeaway: One of the top hosts sending queries to the zones seem to be excess queries that might be concerning, but still falls off after a short period. After running this experiment 4 times, this final run was the only time that showed one probe's resolver behaving questionably. Note that the resolver's address exhibits issues for multiple tests, not just the NS to ".". The probe and resolver in question is in Russia:

```
Address: 31.13.129.36
  Numeric ip: 520978724
         ASN: 197765
       Owner: ITPARK_DC
     Country: RU
    ip_range: 31.13.128.0 - 31.13.135.255
```

4.2 ATLAS result analysis

To analyze the results reported by RIPE Atlas directly, we break the ATLAS data into parts and analyze each set of responses for each of the 3 runs for each of the 6 test cases (36 datasets in all). Our analysis simply determined how many RIPE probes are generally answering correctly in conjunction with their associated resolvers.

To study these response, we graph the types of responses returned for each different test: NOERROR, NXDOMAIN, SERVFAIL, or a timeout.

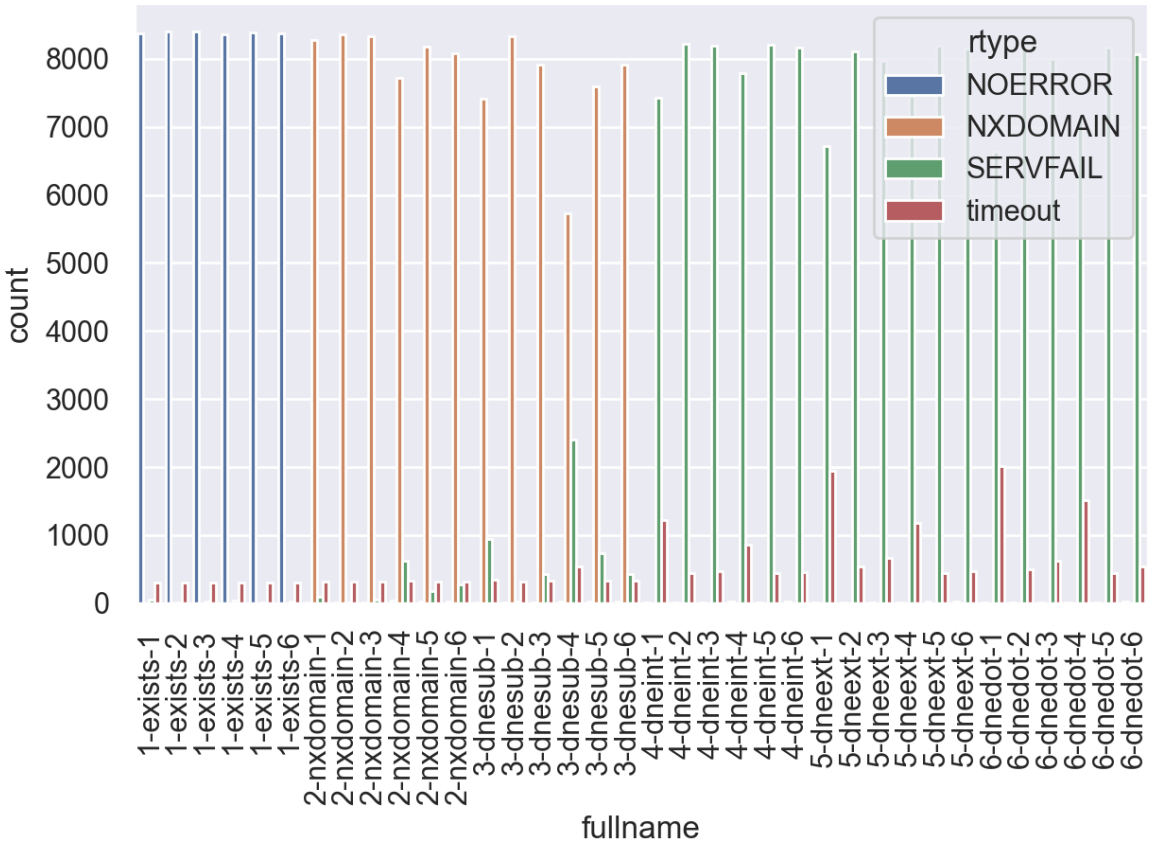

We also show this in a heatmap format:

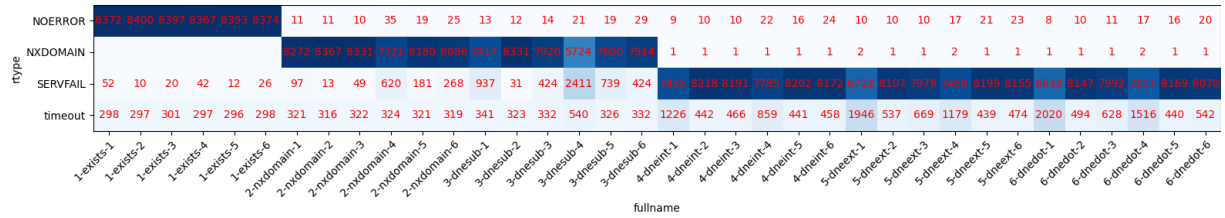

Because these are hard to visualize with 36 tests, we also aggregated the different answer types together for each of the six major test types:

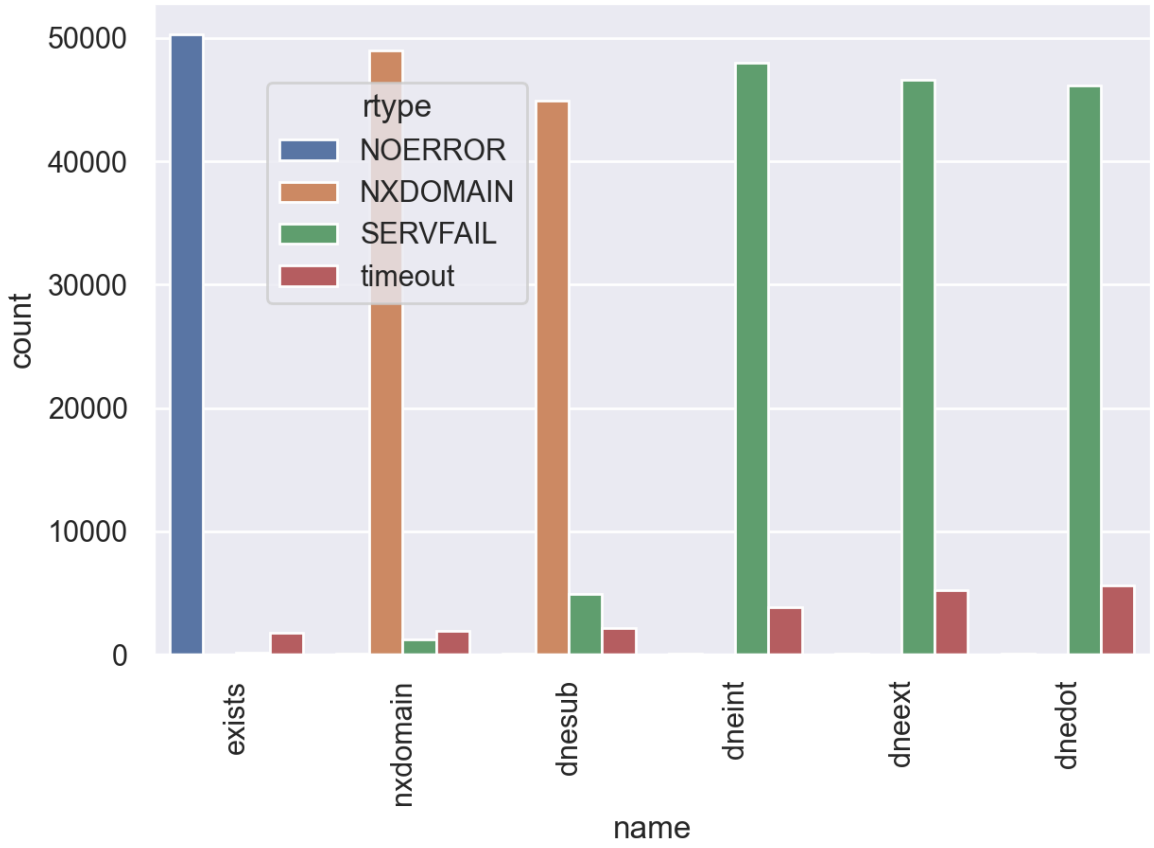

Takeaway: the tests for the NS target of "." did not show a significant increase in the number of concerning responses (timeout) than was shown for the other missing or invalid NS server tests.

4.3 Traffic analysis as seen at b.root-servers.net

In partnership with the b.root-servers.net (aka B-Root) operational team, we studied all the related traffic that was received at B-Root during the testing period, looking for queries to the associated *sub-bbb* zones. This analysis shows similar results to the pcap data collected directly at the authoritative server. None of the top talkers exhibited a statistically significant increase in sent queries than during the other tests.

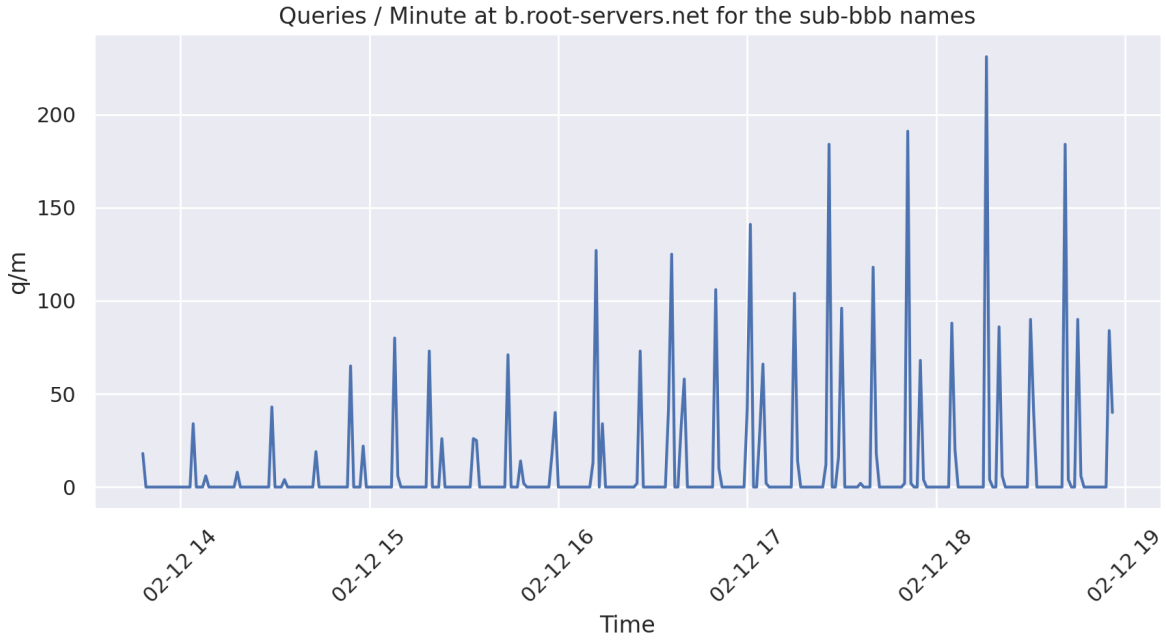

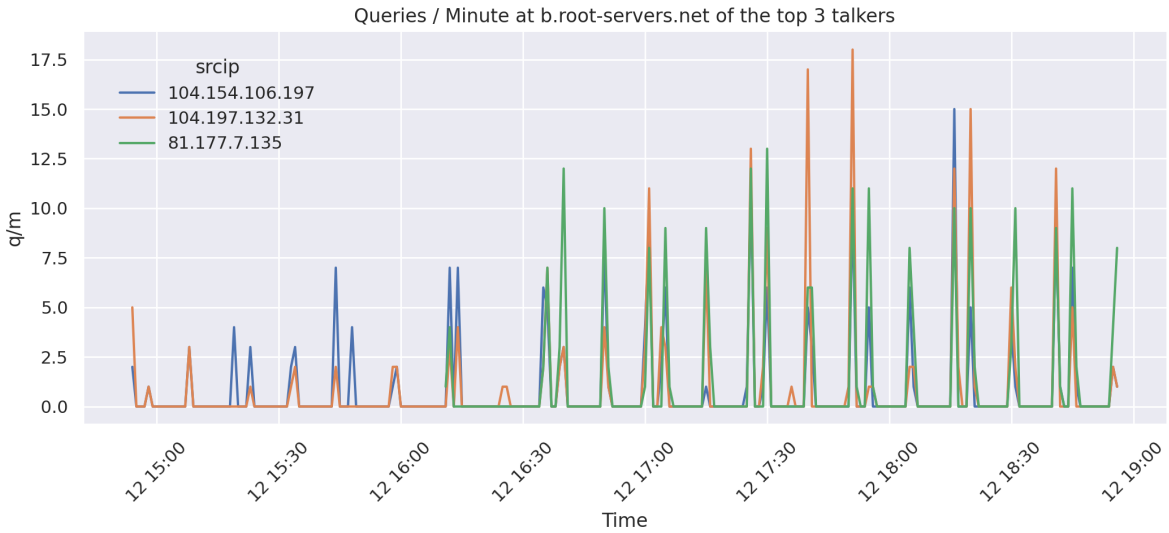

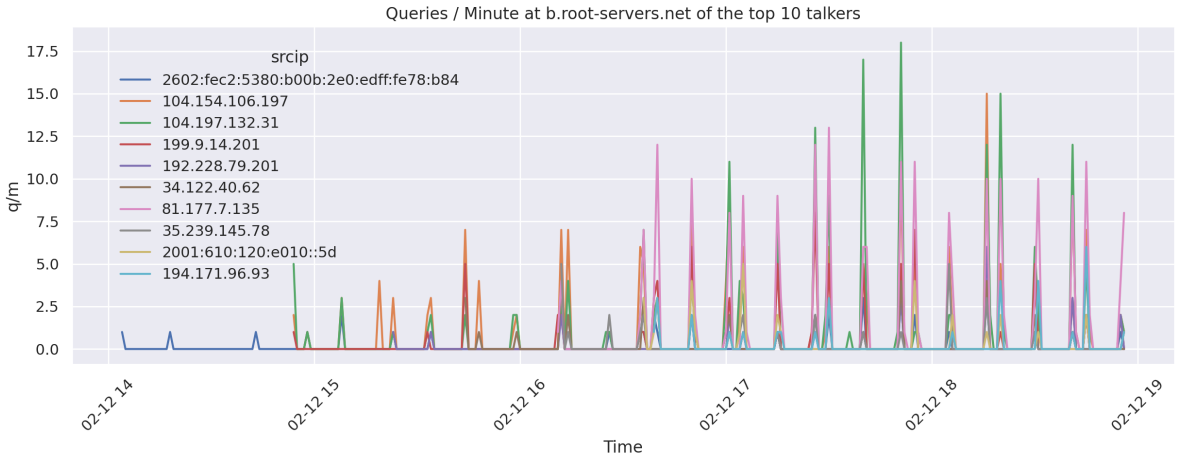

Takeaway: the tests for the NS target of "." did not show a significant increase in the number of concerning queries sent to B-Root as a result of these tests.

5 Conclusion

Based on the above three studies of different data sources from the many tests executed from 5000 RIPE Atlas nodes, it does not appear that these tests "broke the internet".

6 And now for something completely different. . . .

We also got a number of, um, more interesting queries. Here is a small set of sample queries for your entertainment:

QName	QType	Count
ns1.sub-bbb.frostedaxe.games.	A6	12
ns1.sub-bbb.frostedaxe.com.	A6	10
exists.sub-bbb.frostedaxe.games.	CNAME	4
ns1.sub-bbb.frostedaxe.games.	MX	4
ns1.sub-bbb.frostedaxe.games.	TXT	2