Patterns in Queries Sent to Root Servers

ICANN Office of the Chief Technology Officer

Paul Hoffman

OCTO-020
15 December 2020

ICANN

TABLE OF CONTENTS

This document is part of ICANN's Office of the Chief Technical Officer (OCTO) document series. Please see the OCTO publication page for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

1 Introduction

Research into the traffic seen at root servers typically is done by using the “day in the life” (DITL) data[1] whose collection is facilitated by the DNS Operations Analysis and Research Center (DNS-OARC).[2] That data is collected once a year and is published for use by researchers and DNS operators.

While the DITL data has proven useful for some research, it only covers up to two days of interactions with the root servers which is not long enough to capture some interesting signals. For example, because the time to live (TTL) for most records in the root zone is two days, more than two consecutive days of data is needed to look at how resolvers cache answers received from the root zone. Root server operators (RSOs) can perform longer-term research on their own data.

This report looks at five consecutive days worth of data from the ICANN Managed Root Server (IMRS, also known as L-root) to see what patterns can be found in the interactions with the resolvers that send the largest number of queries during this period. These interactions can be considered “conversations” because recursive resolvers are expected to cache the results of queries, with pauses in the conversations caused by caching.

The research started with the hypothesis that the most active query sources would be caching resolvers, and would thus share some traits in common; if so, knowing these traits might be valuable when evaluating the root server system as a whole. We were not able, however, to find enough evidence to support this hypothesis, and the most active query sources have little in common; in fact, they might not be traditional resolvers at all. There are many reasons why they have so little in common; some of these reasons are covered at the end of this report.
The intended audiences for this report are DNS researchers and others in the DNS technical community.

This document supports ICANN’s strategic goal to improve assessment of, and responsiveness to, new technologies which impact the security, stability, and resiliency of the Internet’s unique identifier systems by greater engagement with relevant parties. It is part of ICANN’s strategic objective to evolve the unique identifier systems in coordination and collaboration with relevant parties to continue to serve the needs of the global Internet user base.

2 Test Setup

All DNS traffic to the IMRS is captured for analysis. The traffic at each instance is compressed and temporarily stored at each instance, and then retrieved by the IMRS operators at a central location. The research in this report looks at all IMRS traffic from 1 to 5 April 2020.

The raw data for the five days for all instances takes up about 2.4 terabytes for approximately 76 billion queries. The data was processed to only keep the following fields for each query: time stamp (accurate to 1 second), IPv4 or IPv6 address of the query source, QNAME, QTYPE, and

[1] See https://www.dns-oarc.net/oarc/data/ditl
[2] See https://www.dns-oarc.net/

the value of the RD bit. This data was then imported into a database (PostgreSQL, in this case) for analysis.

2.1 Distribution of Queries by Address

The approximately 76 billion queries received came from approximately 11 million unique IP addresses. Of these addresses, about 0.15 percent were from the IPv4 and IPv6 private use address space and, thus, were clearly not from query sources that could be responded to.

The addresses were ordered by query volume, and then split into two categories, “top” and “rest,” with each category having about 38 billion queries. The top category had 1,611 addresses, which is a tiny percentage of the 11 million addresses seen during the five-day period. That is, this tiny portion of the 11 million addresses sends about half of the queries processed by the IMRS.

Figure 1 shows the request counts for the top query sources, which are the subject of the rest of this analysis, and Figure 2 shows the distribution of the rest. Note that the y-axes in both figures are in a logarithmic scale.

Figure 1: The top query sources

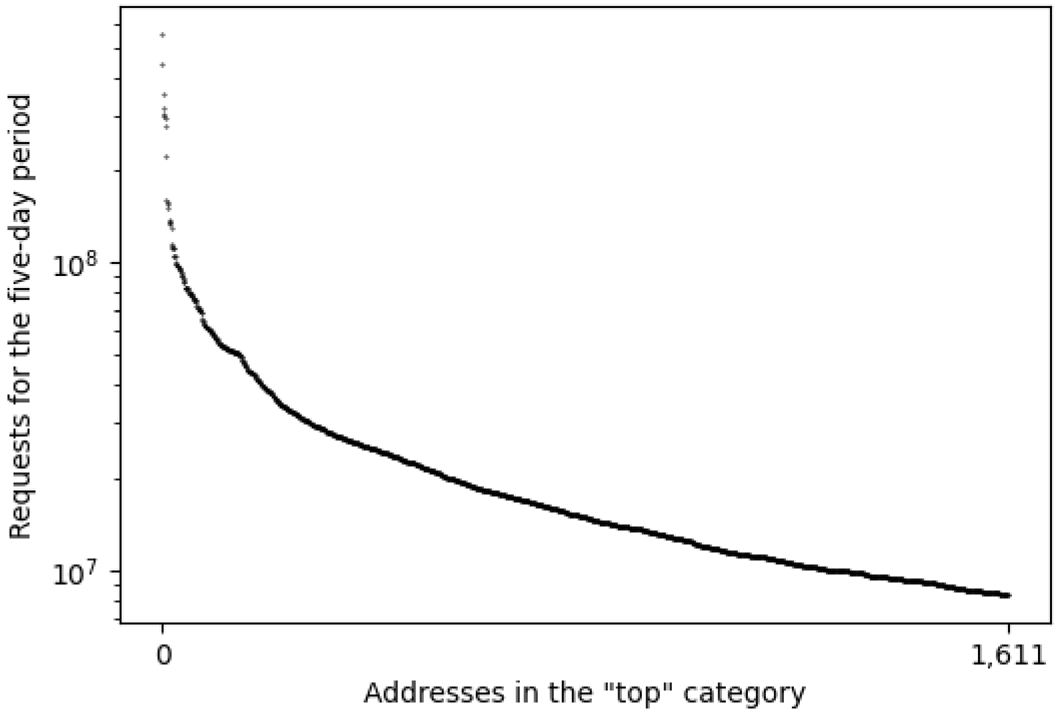

Figure 2: The rest of the query sources

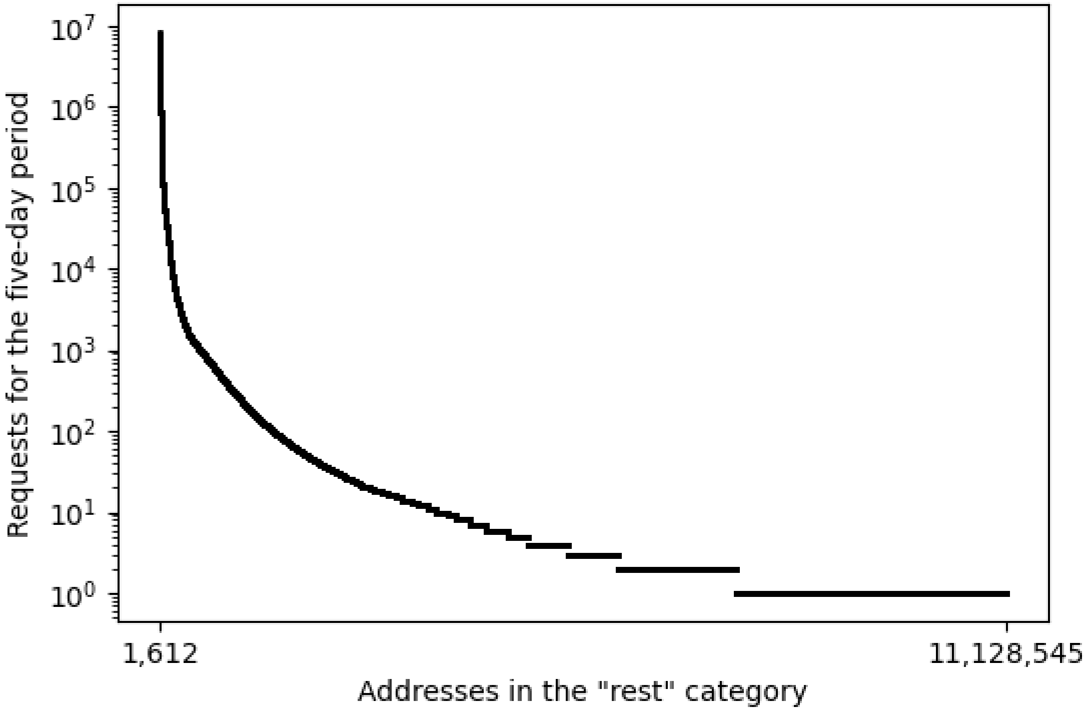

The query volumes in the top category range from 552,962,054 to 8,326,702 queries during the five-day period, with a median of 14,889,157. The volumes in the second category (the “rest”) range from 8,325,640 queries to just 1 during the period, with a median of just 3.

3 Analysis of Top Query Sources

As seen in Figure 1, the 1,611 addresses in the top category have a very wide range of query volumes over the five-day period. Of these 1,611 addresses, 200 were selected at random for the purpose of this analysis. The result detailed below shows that there was little commonality found in this sample of query sources.

3.1 Observed Top Level Domain Caching

For the measurements of top-level domain (TLD) caching, the data for query sources is presented in two ways. Over the five days, a single resolver should make queries two or three times, with a two-day gap between queries as per the root zone TTL. Common resolvers such as BIND and Unbound are known to allow maximum TTLs for cache entries lower than two days, but hopefully no less than one day. In extreme cases, the maximum TTL of six hours might be enforced. Thus, the data is shown both as “total queries over five days” and “number of six-hour periods with no queries.”

Resolvers may set a maximum TTL for their caches, and that maximum might be below the two days that is the TTL for all TLD records in the root zone. To test this, an analysis of queries with

QNAMEs in the ".org" zone was made. The .org TLD was selected because it is one of the TLDs for which the most queries are seen.

Figure 3 shows the number of times an address sent a query with .org as the TLD in the QNAME. Note the extremely wide dispersion of values and the log scale; a caching resolver using the TTLs from the root zone should have fewer than ten queries.

Figure 3: The total number of .org queries in a five-day period, by address

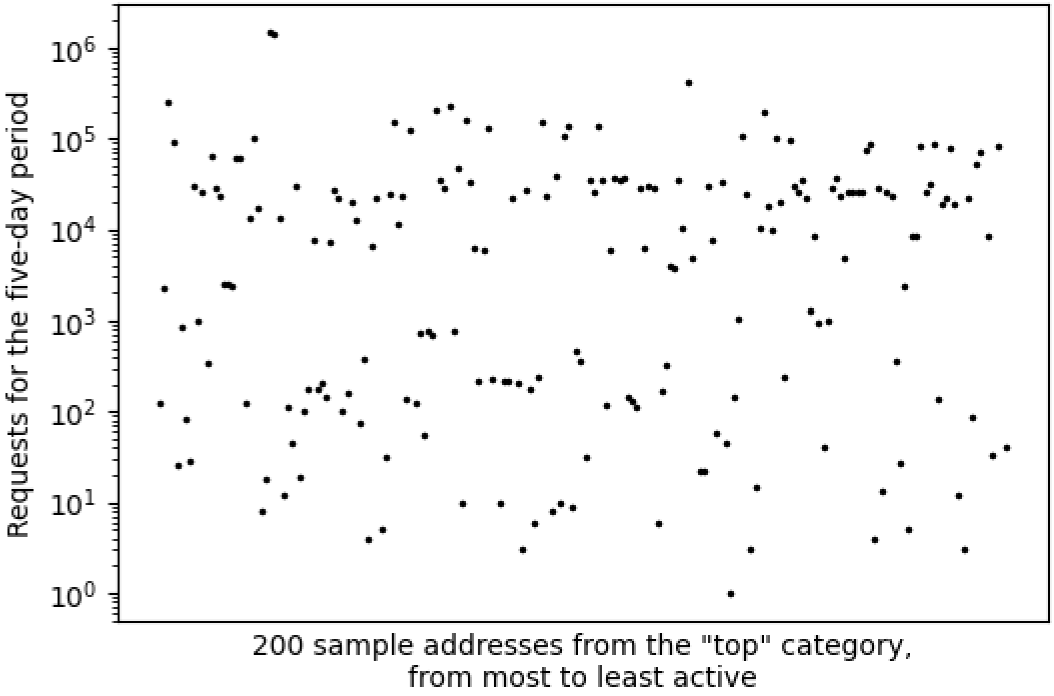

Figure 4 shows the number of six-hour periods that have no .org queries during the five-day period. A typical caching resolver using full TTLs should only ask two times during a five-day period, and thus have 18 periods with no .org queries.

Figure 4: The number of periods with no .org queries in a five-day period, by address

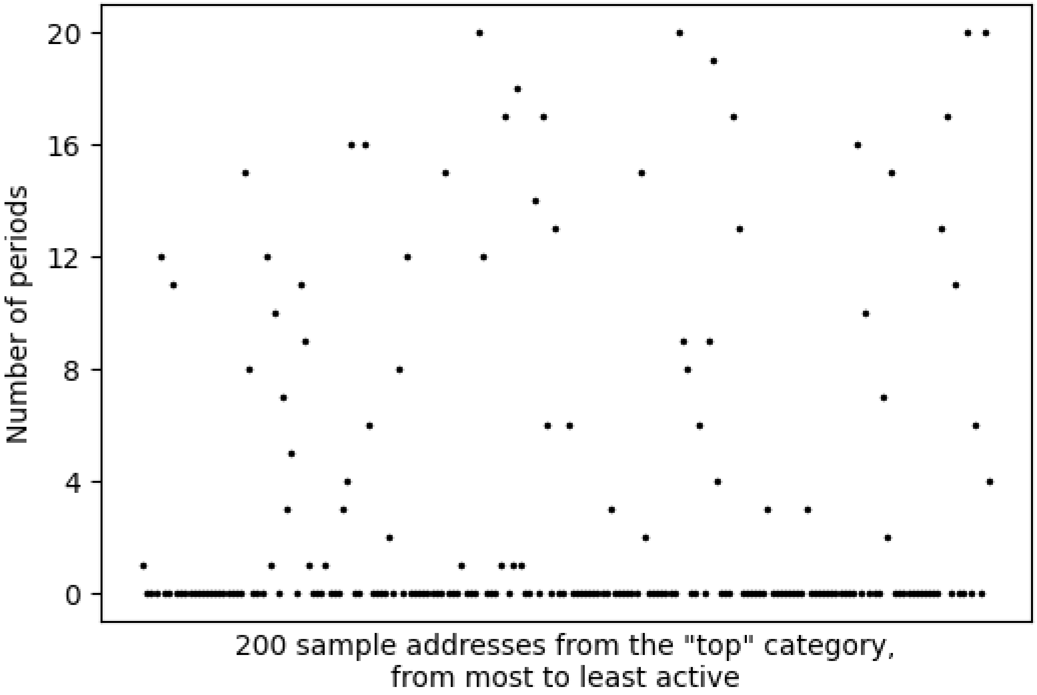

3.2 Chromium-based Queries

Reports from many researchers show that approximately half of all queries that go to the root server system originate from Google Chrome and other browsers based on Google’s Chromium browser source code.[3] These queries are single-label names whose length is between 7 and 15 letters, inclusive.

Some people have interpreted these reports as implying that about half the queries from typical query sources are Chromium-based queries. Figure 5 shows that, and while this simple assumption might be statistically valid, the variation between query sources is quite high, and seems uncorrelated with the number of queries sent.

[3] See https://blog.verisign.com/domain-names/chromiums-impact-on-root-dns-traffic for a good example of such a report

Figure 5: The percent of Chromium-based queries in a five-day period, by address

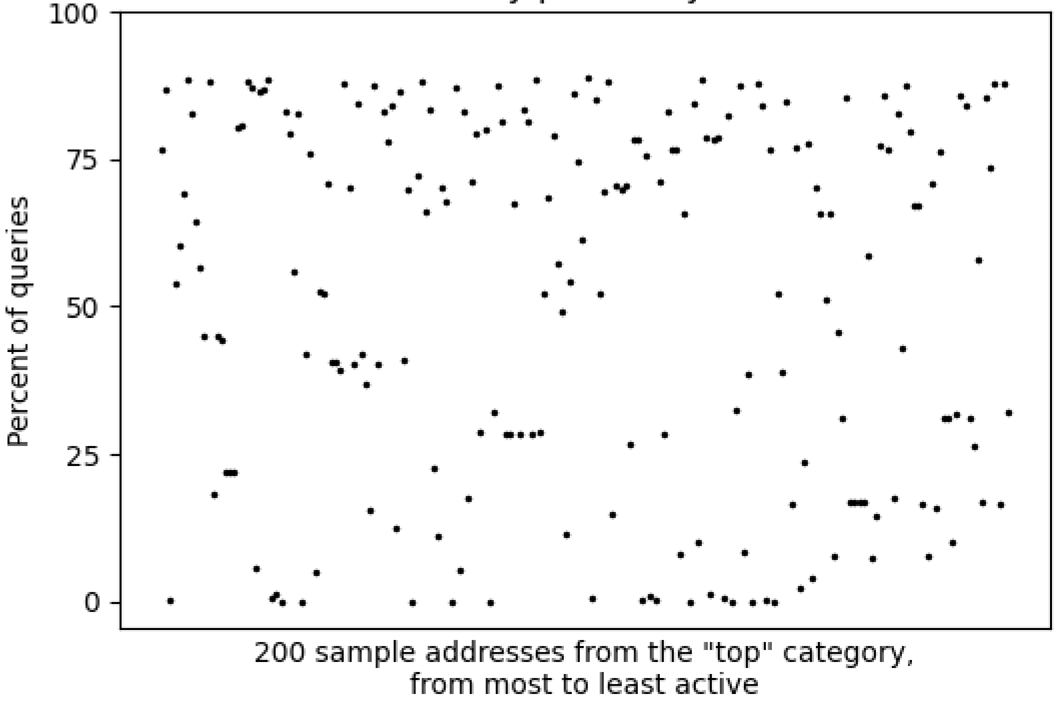

3.3 Stickiness

There is no expectation that query sources will necessarily stick with one root server instance such as the IMRS for more than a few hours when instances of different RSOs are close together. Typical caching resolvers based on open source software will check for the fastest connection to the root server identifiers every few hours and will change preferred root server identifiers based on these results. For many resolvers, however, only one root server will always clearly be the fastest. Additionally, some common resolver software will randomly pick root server identifiers within the set of fastest results. However, there is a reasonable expectation that the query sources sending the most traffic to a root server instance might use that instance all the time.

Figure 6 shows that the latter expectation is borne out for the IMRS in the five-day period under study: only 4 of the 200 sampled query sources had any six-hour periods with no queries. It is assumed that the queries from those query sources were sent to other root server identifiers.

Figure 6: The number of periods with no queries in a five-day period, by address

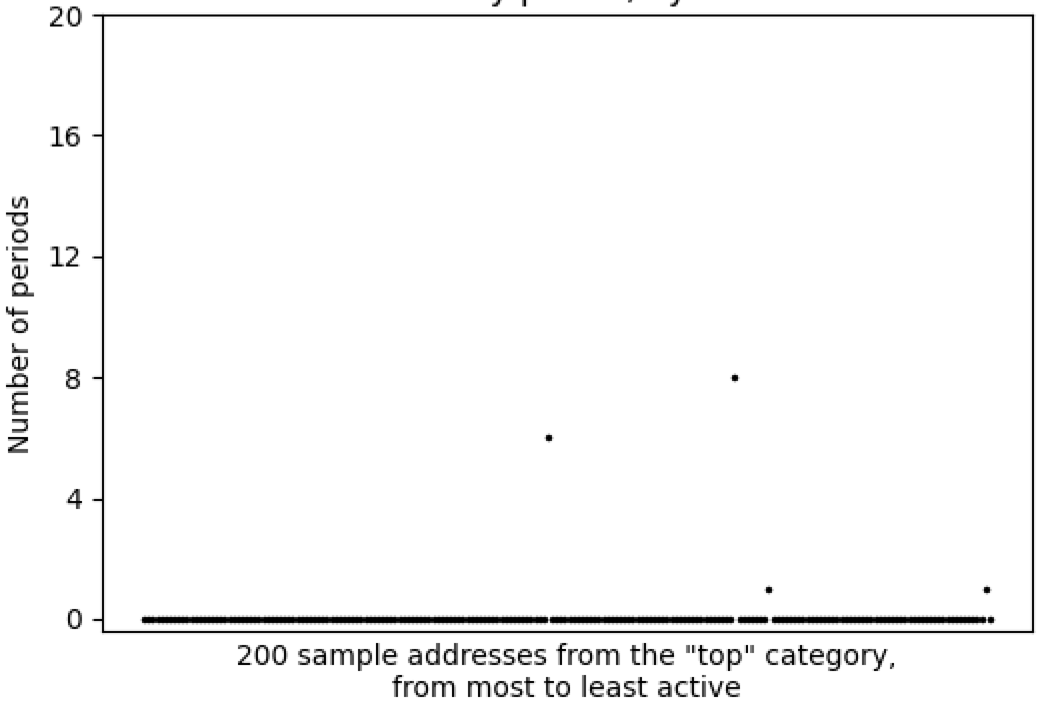

4 Conclusion

This report shows the lack of patterns in the interactions of the top query sources with a root server. There are many reasons why this could be, most notably that only a small percentage of the top query sources are in fact traditional caching resolvers.

This behavior could be explained by the top query sources being DNS forwarders, or by systems acting as front ends for large farms or independent resolvers. However, neither of these scenarios would explain why query sources are making tens of thousands of queries per day for the same TLD. Another theory is that a large number of very active caching resolvers have their caches badly configured.

Given the unpredictability of the results, the data in this report does not lead to any greater understanding of how to make the root server system more stable or reliable. It also implies that root server system designers should not assume conversations but instead assume individual transactions. Further research could examine the entire dataset looking for query sources that act like caching resolvers that follow the given TTLs; separating these out may lead to a better understanding of the different behaviors.