A Perspective Study of DNS Queries for Non-Existent Top-Level Domains

A Technical Report from the Name Collision Analysis Project (NCAP)

13 July 2022

Preface

This is a report to the ICANN Board, the ICANN organization (ICANN org), the ICANN community from the Name Collision Analysis Project.

Table of Contents

Executive Summary

As part of the goals and objectives of the Name Collision Analysis Project (NCAP) Study Two, the NCAP Discussion Group commissioned a study to better understand within the context of name collisions how representative DNS data is at various points of the DNS hierarchy. The study's main objective is to provide insights and guidance for future examinations of the DNS name collision data that will be used by ICANN for risk analysis and assessments of TLD string applications. This study, referred to as "A Perspective Study of DNS Queries for Non-Existent Top-Level Domains", focuses on two key measurements: (1) comparing traffic received at each root server identifier and (2) comparing traffic received at public recursive resolver(s) and the root server system. The former measurement provides insights into the ability of name collision DNS data to be collected and analyzed by using a single or subset of root servers, while the latter provides insights into the completeness of DNS measurements taken only at the root by examining DNS name collision traffic at the recursive layer of the DNS hierarchy.

Our analysis shows that no view at a single root server is comprehensive. However, when considering DNS clients that meet a defined query rate, a single root server observes query traffic from about two thirds of resolvers that are observed across the entire system. Additionally, there are notable differences in DNS traffic observed by recursive resolvers and at the root server system. These findings are significant in terms of how future guidance and advice may be applied to name collision risk assessments. Specifically, these perspective differences affect the effectiveness of top-N lists, particularly when they are generated from a single source.

The publication of top-N lists of non-existent TLDs can make applicants aware of strings that exhibit some risk associated with name collision. However, the effectiveness of such lists is limited. The very fact that these lists contain only the top N, ranked by some criteria, is constraining. This is particularly so when they are generated only from a single data source (e.g., root server queries or a single recursive resolver). Because there are multiple perspectives in the DNS ecosystem, the absence of a string on a top-N list does not provide any assurance the string is void or absent of name collision risks nor does the magnitude or ranking of a string that does show up in the list. For example, this analysis shows that non-existent TLDs observed at high volumes by some recursive resolvers are not seen in the same rankings by root servers.

Terminology

- Root Server Identifier (RSI)[1] - is the DNS name associated with a root server operator that appears in the root zone and root hints file. For example, c.root-servers.net is the root server identifier associated with the root server managed by Cogent at the time this document was published.
- Day-In-The-Life (DITL)[2] - a large-scale data collection project undertaken every year since 2006. This data has historically been the primary measurement asset for name collision studies.
- Delegation[3] - The introduction of a TLD into the Internet's authoritative database, known as the Root Zone.

Background

Preceding the round of new gTLDs in 2012, numerous studies were conducted by JAS Global Advisors, Interisle, ICANN, Verisign, and other researchers using various types of DNS data to measure and assess name collision risks.[4] The primary data used was root server DNS traffic data collected by DNS-OARC's DITL project. The DITL data provided the most complete view/collection of DNS traffic to the root servers despite being limited to a small number of days per year. The DITL data helped form the guidance issued by JAS Global Advisors to assess the risk of the applied-for TLDs based on query volume and other metrics observed at the root.

The next round of new gTLD applications will require name collision risk assessments by the applicants and ICANN. However, DITL and root data may not be adequate or even available to assure accurate and complete assessments due to anonymization efforts by root server operators and general changes within the DNS ecosystem that raise concerns about availability and accuracy. This study aims to understand the distribution of DNS name collision traffic throughout the DNS hierarchy and provide insights into where and how DNS data can be collected and assessed and data limitations within the context of name collisions.

---

[1] Root Server System Advisory Committee (RSSAC), "RSSAC 026 – RSSAC Lexicon" 14 March 2017, https://www.icann.org/en/system/files/files/rssac-026-14mar17-en.pdf.

[2] DNS Operations, Analysis, and Research Center (DNS-OARC), "Day In The Life of the Internet (DITL)," accessed 26 January 2022, https://www.dns-oarc.net/.

[3] ICANN, New Generic Top-Level Domains, "Delegated Strings," accessed 26 January 2022, https://newgtlds.icann.org/en/program-status/delegated-strings.

[4] ICANN, "Mitigating the Risk of DNS Namespace Collisions Final Report by JAS Global Advisors," ICANN Announcements, 30 November 2015, https://www.icann.org/en/announcements/details/mitigating-the-risk-of-dns-namespace-collisions-final-report-by-jas-global-advisors-30-11-2015-en.

Studies

“A Perspective Study of DNS Queries for Non-Existent Top-Level Domains” consists of two main studies: the comparison of traffic among RSIs and the comparison of name collision traffic observed by root and recursive resolvers. Together these two studies help provide insights into how risk assessments of name collisions should be evaluated based on the availability of DNS traffic data.

RSI Comparison: This study uses root server data collected by the 2020 DNS-OARC DITL to compare recursive resolver traffic received by each RSI. Using the source IP address and its number of queries issued, various measurements comparing the overlap and distribution of these sources to the various root server identifiers are calculated. Further analysis looking at A and J root server traffic data compares the top name collision strings based on two previously established critical diagnostic measurements, query volume and source diversity.

Recursive Resolver and Root Comparison: This study aims to examine the relatively opaque and widely inaccessible data for name collision analysis traffic to recursive resolvers. This study uses root server data at A, J, and L root servers, a public open recursive resolver, and commercial recursive resolver to compare the top name collision strings based on two critical diagnostic measurements established in the NCAP case study for .CORP, .HOME, and .MAIL: query volume and source IP address diversity.

Study 1: Root Server Identifier Comparison

Data

In order to compare RSIs, data was sourced from the DNS-OARC DITL 2020. At the time, the data for 2021 was not yet available. The 2020 DITL data was collected from May 5th to the 7th, 2020. The contributing root server identifiers were A, B, C, D, E, F, H, I, J, K, L, M. Note that B, E, and F data files are very "small" in terms of data stored in the 2020 DITL fileshare.

This study utilized a data set comprised of the following query characteristics extracted from UDP queries collected as part of DITL:

- IP Address
- Number of queries
- Number of priming queries (i.e., NS . queries)
- Root letter

Notable Limitations of the Data

Two of the root server identifiers, L and I, anonymize the source IP address. Unfortunately, this limits the ability to use those RSI's data. For example, the I-root data takes the source IP address and anonymizes all of them into the 10.0.0.0/8 IP address space. L-root anonymized the source IP address across the whole IPv4 range. IP anonymization impedes most of our measurements, thus both I and L RSI's were excluded from this study's measurements. Furthermore, the size and completeness of B, E, and F RSI's data was inadequate for this study's required measurements and these RSIs were also excluded. These exclusions reduced the original twelve RSIs down to seven.

The study "Case Study on Collisions Strings" identifies several "critical diagnostic measurements." The current analysis includes some of those metrics, including query volume and query origin diversity (both IP address and ASN). However, it does not include query name label diversity or query type diversity. The scope of the current analysis is thus limited to what can be assessed without those measurements.

Measurements

The following twelve measurements[5] were taken:

1. Query volume per RSI
2. Unique source IP address at each RSI

[5] Further investigation could explore the number of sites per RSI and other ratios related to queries and source IP addresses.

3. Distribution of query volume per source IP to all of the root server system
4. Identifying top talkers[6] that constitute a large percentage of overall traffic
5. Measuring overlap of top talkers at each RSI
6. Comparing the set of IPs at each RSI to the other RSIs
7. How many RSIs must be analyzed to reach 100% of the top talkers
8. How many RSIs does a typical top talker IP query
9. Are there any geospatial outliers within the top talker set of IPs
10. How evenly do top talkers distribute the query volume over RSIs
11. Is there a geographical bias for various countries to favor a subset of RSIs
12. What variation exists in the Top-N non-existent TLDs per RSI

Query Volume per RSI

The first baseline comparison of RSI traffic is the number of queries each receives. As shown in Figure 1 below, the number of queries received at each RSI varies; accordingly, this measurement provided insights into data collection issues with B, E, and F and why they were ultimately excluded from further analysis. The total query count for these root server letters (B, E, and F) combined was less than 0.8% of the entire 2020 DITL.

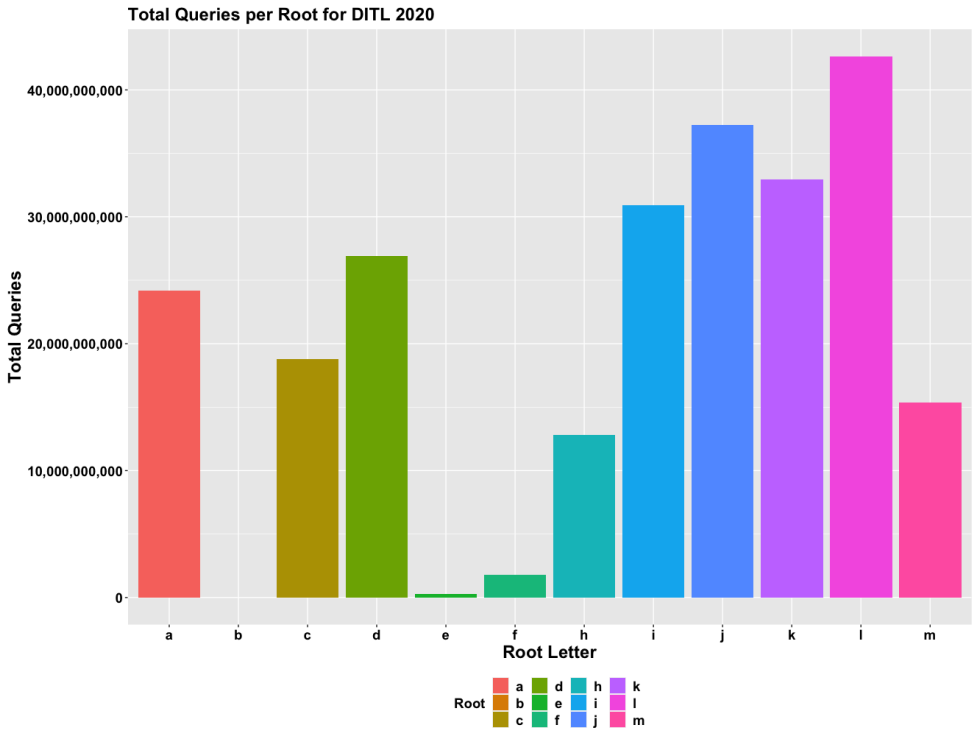

Figure 1 - Query Volume per RSI

[6] Top talkers are recursive resolvers that issue the largest amount of DNS queries to the RSS.

Top Talkers

A second fundamental measurement was to understand the number of unique IP addresses seen at each RSI. This is useful to understand if we should expect IP-RSI affinities, which would have a direct impact on any future name collision analysis that uses a subset of RSIs. Figure 2 below shows the number of unique IPv4 and IPv6 addresses seen at each RSI. That distribution, on the included RSIs, is relatively even (median = 8.52M, mean = 7.61M, standard deviation = 1.8M). In total, 15.51M unique IPv4 and 1.56M IPv6 addresses were observed. The complete set of client IP addresses observed in the 2020 DITL data are associated with 56,174 ASNs and 2.9M unique /24 IPv4 blocks and 1.2M unique /48 IPv6 blocks. The top 10 ASNs account for 25% of total 2020 DITL queries collected. Their distribution is as follows:

1. ASN 24560 - Bharti Airtel (4.089764%)
2. ASN 9498 - Bharti Airtel (3.599579%)
3. ASN 15169 - Google (2.717901%)
4. ASN 4837 - China Unicom (2.687349%)
5. ASN 7018 - AT&T (2.232134%)
6. ASN 12322 - Free SAS (2.226789%)
7. ASN 4134 - China Telecom (2.129250%)
8. ASN 7552 - Viettel Group (1.863532%)
9. ASN 16509 - Amazon (1.855771%)
10. ASN 38266 - Vodafone India (1.763266%)

The focus of this study is to understand how similar the sets of source IP addresses and queries for non-existent TLDs are across different RSIs. There are numerous similarity measurement approaches but a simple and often reliable measurement is the Jaccard index that is a statistic used in understanding the similarities between sample sets. Already from Figure 2 we can see that any type of set measure of unique IP sources will have significant variance. In order to support the focus of this study, a smaller number of IP addresses, that are representative of the entire RSS, is needed for the use of set similarity. To that end, understanding the profile of how many queries each IP sent is needed.

Figure 2 - Unique IP Addresses Observed per RSI

To understand the query volume distribution over the set of IP addresses observed in the 2020 DITL collection, a cumulative distribution measurement was made by ranking IP addresses in ascending order by the number of total queries that IP sent to the RSS. Figure 3 below depicts this distribution measurement relative to the total percentage of IP addresses observed during the 2020 DITL. A typical Power Law Distribution[7] was observed:

- 15% of IP addresses issued only 1 query.
- 27% of IP addresses issued 2 or fewer queries.
- 50% of IP addresses issued 10 or fewer queries.
- 98% of IP addresses issued 10,000 or fewer queries.

It is unclear as to what those source IP addresses that send so few queries actually are. A typical recursive resolver with even a minimal amount of users, clients or other stub systems behind it might be expected to generate more than 10 queries to the RSS over a period of 48 hours. It is possible these are odd pieces of software, spoofed source IPs, or a variety of other things. From Figure 4 we observe that the vast majority of the traffic collected during 2020 DITL comes from a small subset of IP addresses that send large amounts of queries. Figure 4 shows a distribution of the number of the top querying IP addresses relative to the total percentage of 2020 DITL queries collected. This measurement shows that 90% of the total 2020 DITL can be represented by the top 115K IPs. Likewise, 95% of the total 2020 DITL can be represented by the top 250K IPs. The top 115K set of client IP addresses observed in the 2020 DITL data are associated with 9,968 ASNs and 26,816 unique /24 IPv4 blocks and 5,022 unique /48 IPv6 blocks.

[7] Wikipedia contributors, "Power law," *Wikipedia, The Free Encyclopedia*, accessed January 26, 2022, https://en.wikipedia.org/w/index.php?title=Power_law&oldid=1065551786.

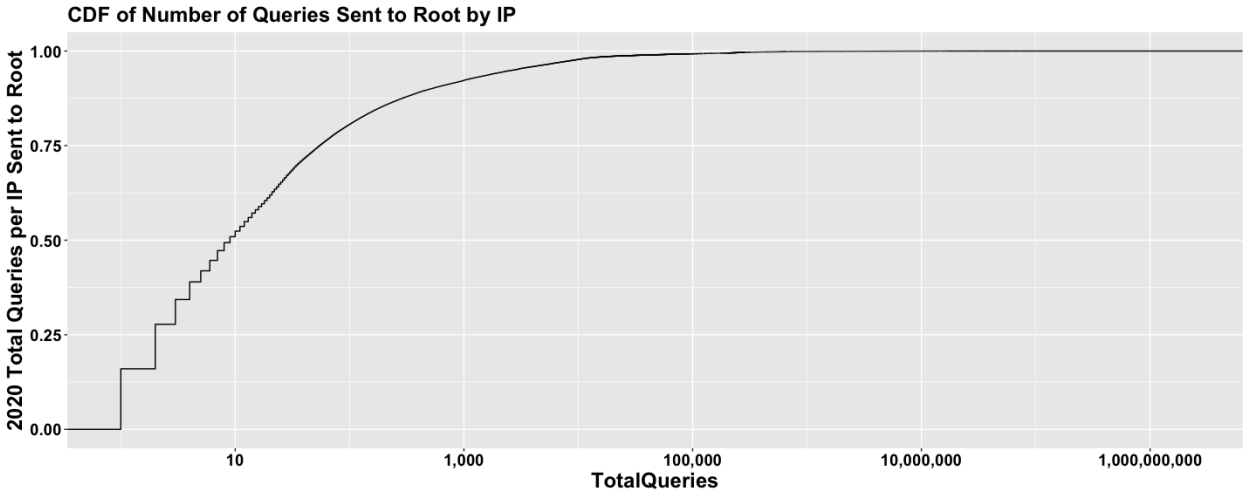

Figure 3 - Cumulative Distribution of the Number of Queries Sent to RSS by IP Address

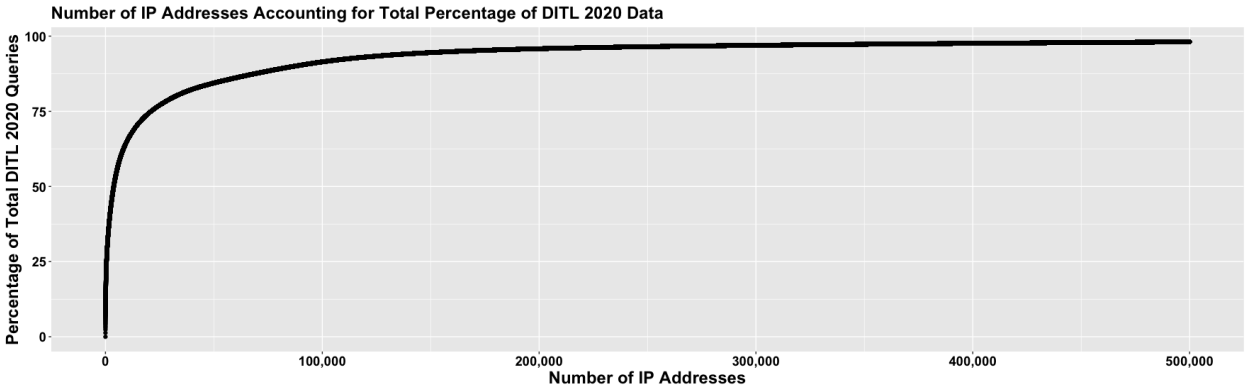

Figure 4 - Number of IP Addresses Accounting for Total Percentage of DITL 2020

As seen in Figure 3, nearly 92% of all source IP addresses seen during 2020 DITL sent less than 1,000 queries during the 48 hour collection interval. Additional analysis was conducted on this long-tail of source IP addresses with low query volume to better understand the domain names these source IP addresses requested.

Nearly 15% (2.7 million of the total 15.5 million) of IP addresses sent a single query during 2020 DITL. Using the single query set of IP addresses, Table 1 examines the top queried names from those IPs at A and J root servers during the 2020 DITL. A and J roots observed 732,711 of the 2.7 million single querying IP addresses. The most frequently observed name was for a root priming query (i.e., '.') from 160,740 unique source IPs, which account for 22% of the total single querying IPs. Other notable names include RFC 8145 trust anchor signals. Two other interesting findings appear in the examination of these top single query names: 1.) Most of the names are delegated in the global DNS and would result in a successful response of NOERROR from the RSIs. This is not expected based on RSSAC002 data, which shows most traffic at the RSS results in NXDOMAIN responses. 2.) Many of these top names appear to be multi-labeled domains under delegated namespaces such as qq.com and in-addr.arpa.

Rank	Qname	Number IPs	Rank	Qname	Number IPs
1	.	160740	31	_xmpp-client._tcp.wallapop.com	1985
2	sy.eu.angsrvr.com	20494	32	ns2.dnsimple.com	1933
3	im.mielse.com	14582	33	ns4.dnsimple.com	1932
4	ad.afy11.net	13421	34	api.surfeasy.com	1929
5	chat.grindr.com	11664	35	ns1.dnsimple.com	1900
6	com	10813	36	ns3.dnsimple.com	1865
7	gsm1.g4c5j.com	10747	37	dion.ns.cloudflare.com	1785
8	www.google.com	9119	38	lb._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1767
9	216.58.202.4.in-addr.arpa	9000	39	dr._dns-sd._udp.0\\162z\\001\\248\\025\\022\\001	1713
10	150.109.167.160m\\183{\\128\\156\\023\\001\\233\\172\\186\\158\\244vah\\026\\001	7368	40	a1-14.akam.net	1666
11	xmpp-prod.monksoftware.it	6845	41	hostname.bind	1598
12	net	5788	42	r._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1566
13	_ta-4f66	5341	43	reachit.lenovo.com	1561
14	_sip._udp.smart.0038.net	4652	44	sso.cloudsoftphone.com	1517
15	ns-635.awsdns-15.net	3997	45	fire-base.com	1512
16	ns-1891.awsdns-44.co.uk	3904	46	ns3.msft.net	1508
17	ns-1192.awsdns-21.org	3884	47	-.-	1481
18	ns-399.awsdns-49.com	3878	48	tv	1461
19	providers.cloudsoftphone.com	3701	49	m.root-servers.net	1383
20	_ta-4a5c-4f66	3634	50	p.qpic.cn	1379
21	cx-messenger.imbeepro.es	3292	51	livem.l.qq.com	1379
22	local	3192	52	tv.l.qq.com	1377
23	glborigintest.canarytest.net	2652	53	qq.m.cn.miaozhen.com	1364
24	moiawsorigin.clo.footprintdns.com	2647	54	vgdt.gtimg.cn	1363
25	qq.com	2545	55	mtcls.qq.com	1362
26	teredo.ipv6.microsoft.com	2513	56	dns.weixin.qq.com	1353
27	support0.bigo.sg	2465	57	appmedia.qq.com	1353
28	49.51.82.122\\023r\\170j\\002\\180\\128\\156\\023\\001\\233\\172\\186\\158,vwukq\\00	2256	58	cm.l.qq.com	1352
29	home	2051	59	nadia.ns.cloudflare.com	1352
30	49.51.82.122\\023r\\170j\\002\\180(~\\024\\001\\233\\172\\186\\158${=rq\\016	2002	60	la.gtimg.com	1341

Table 1 - Top Names Queried from source IPs only with one request in 2020 DITL

Table 2 examines the ranking of the most popular aggregated second level domain names queried by single query IP addresses. Again, we see that the top querying names all would return a response of NOERROR and that they account for almost 70% of all single query sources at A and J roots.

Rank	Qname (SLD)	Number IPs
1	qq.com.	198905
2	.	160740
3	in-addr.arpa.	31893
4	gtimg.com.	25111
5	angsrvr.com.	20523
6	google.com.	16991
7	mielse.com.	14582
8	co.uk.	13496
9	afy11.net.	13422
10	grindr.com.	11665

Table 2 - Top second level names queried from source IPs only with one request

With such a large percentage of single query source IPs sending queries for resolvable names as opposed to the overall norm of non-existent names being sent to the RSS, additional analysis investigating the top NXDomain names was examined in Table 3. Here we observe a significant amount of those names are for RFC 8145 trust anchors. Many of the other names contain DNS service discovery labels (e.g. _dns-sd._udp) and some type of encoded\escaped non-existent label.

Some previous research[8] looking at RFC 8145 queries to the RSS identified certain pieces of software that linked to DNS software libraries would cause inadvertent queries into the public DNS. Based on the names seen in Tables 1, 2, and 3 these names might suggest a similar root cause in which a piece of software (e.g. a QQ mobile app) might inadvertently send DNS queries.

Rank	Qname	Number IPs
1	150.109.167.160m\\183{\\128\\156\\023\\001\\233\\172\\186\\158\\244vah\\026\\001	7368
2	_ta-4f66	5341
3	_ta-4a5c-4f66	3634
4	local	3192
5	49.51.82.122\\023r\\170j\\002\\180\\128\\156\\023\\001\\233\\172\\186\\158,vwukq\\001	2256
6	home	2051
7	49.51.82.122\\023r\\170j\\002\\180(~\\024\\001\\233\\172\\186\\158$[=rq\\016	2002
8	lb._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1767
9	dr._dns-sd._udp.0\\162z\\001\\248\\025\\022\\001	1713
10	r._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1566
11	-.-	1481
12	perforce	1295
13	db._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1164
14	b._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1102
15	150.109.181.132m\\178\\162\\128\\156\\023\\001\\233\\172\\186\\158\\244vah	1089
16	b._dns-sd._udp.0\\162z\\001\\248\\025\\022\\001	1077
17	r._dns-sd._udp.(\\146z\\001\\248\\025\\022\\001\\192\\168\\001m	1031
18	dr._dns-sd._udp.\\168\\147z\\001\\248\\025\\022\\001\\192\\168\\001m	1000
19	r._dns-sd._udp.\\184k^\\001\\248\\025\\022\\001	965
20	b._dns-sd._udp.(\\146z\\001\\248\\025\\022\\001\\192\\168\\001m	939
21	db._dns-sd._udp.0\\162z\\001\\248\\025\\022\\001	924
22	ns.zyxel-usg	919
23	db._dns-sd._udp.@}^\\001\\248\\025\\022\\001\\192\\168\\001m	908

Table 3 - Top NXDomains queried from source IPs only with one request

[8] Roll, Roll, Roll your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover

Rank	Qname	Number IPs	Rank	Qname	Number IPs
1	.	27166942	21	m.root-servers.net	457922
2	net	4022360	22	ns7.cloudflare.com	456293
3	com	3672953	23	ns6.cloudflare.com	454773
4	gsm1.g4c5j.com	2678614	24	pdns196.ultradns.biz	448244
5	ns3.msft.net	1655606	25	wpad.zyxel-usg	447540
6	ns1.msft.net	1630519	26	biz	433063
7	local	1348070	27	ns2.afrinic.net	424529
8	org	970844	28	www.google.com	414777
9	au	779287	29	ari.gamma.aridns.net.au	413913
10	www.microsoft.com	768918	30	dns10.ovh.net	407591
11	manus.authdns.ripe.net	765465	31	ns10.ovh.net	407491
12	uk	707717	32	ari.delta.aridns.net.au	405893
13	_ta-4f66	666081	33	u1.amazonaws.com	388165
14	ns-2027.awsdns-61.co.uk	590959	34	ari.alpha.aridns.net.au	385307
15	ns-1384.awsdns-45.org	587544	35	pixels.change.me	375795
16	\\000	582250	36	u2.amazonaws.com	363989
17	ns-749.awsdns-29.net	568703	37	ari.beta.aridns.net.au	358587
18	localdomain	541666	38	https://app-measurement.com/sdk-exp	357873
19	info	516986	39	e.root-servers.net	343246
20	arpa	510894	40	de	341469

Table 4 - Top names queried from source IPs only with 2 to 1000 requests

An examination of names being queried by source IPs that sent between 2 and 1,000 queries during 2020 DITL are listed in Table 4. These names reaffirm the previously observed behavior of single query source IP addresses, that the majority of names being requested are for existent names and RFC 8145 trust anchors.

The next examination of these low querying source IP addresses focuses on the abnormal proportion of NOERROR (rcode:0) to NXDOMAIN (rcode:3). RSSAC002 data shows that A and J NXDOMAIN rates to be around 55% of total queries. That rate of NXDOMAIN responses was not being expressed in the low query volume source IP addresses. In order to better understand the ratio of NOERROR to NXDOMAIN, a measurement was created to calculate the return code percentages based on the number of queries an IP sent - e.g., For all source IP addresses that sent one query what percentage were NXDOMAIN. For all source IP addresses that sent two queries what percentage were NXDOMAIN, etc.

Figure 5 illustrates the NXDOMAIN percentage based on source IP total query volume during 2020 DITL. A very obvious dichotomy of NXDOMAIN percentage rates are observed in source IP addresses that issue smaller amounts of queries (i.e. from 1 to ~1,000 which is marked with a vertical line) than those source IP addresses issuing larger DNS query volumes. Based on NXDOMAIN rates, these lower querying source IPs behave differently than higher query volume sources. This supports the exclusion of these low querying source IP addresses for RSS similarity measurements as they are not material in terms of the traffic related to this study.

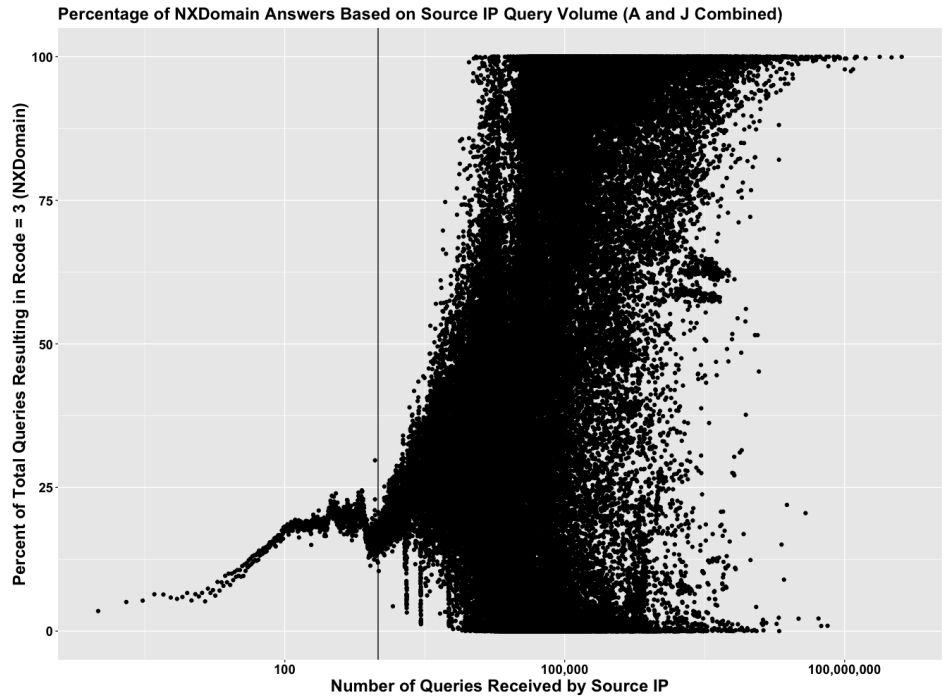

Figure 5 - Percentage of NXDOMAIN answers per source IP total query volume.

In Figure 6, A and J roots are separated to ensure this trend is observed at both of the RSIs. Due to processing and time limitations, this measurement was not feasible for the remaining 2020 DITL RSIs.

The vertical line in Figures 5 and 6 marks a threshold of source IP addresses that sent 1,000 queries during 2020 DITL. Within the context of name collision analysis, capturing queries that result in NXDOMAIN is critical for risk assessment purposes. Based on the clear behavioral change depicted in Figure 5, a threshold for IPs to send a minimum of 1,000 queries was selected. This threshold, which includes all IPs to the right of the vertical line in Figure 5, results in 98% of total 2020 DITL NXDOMAIN responses to be represented in our similarity measurements.

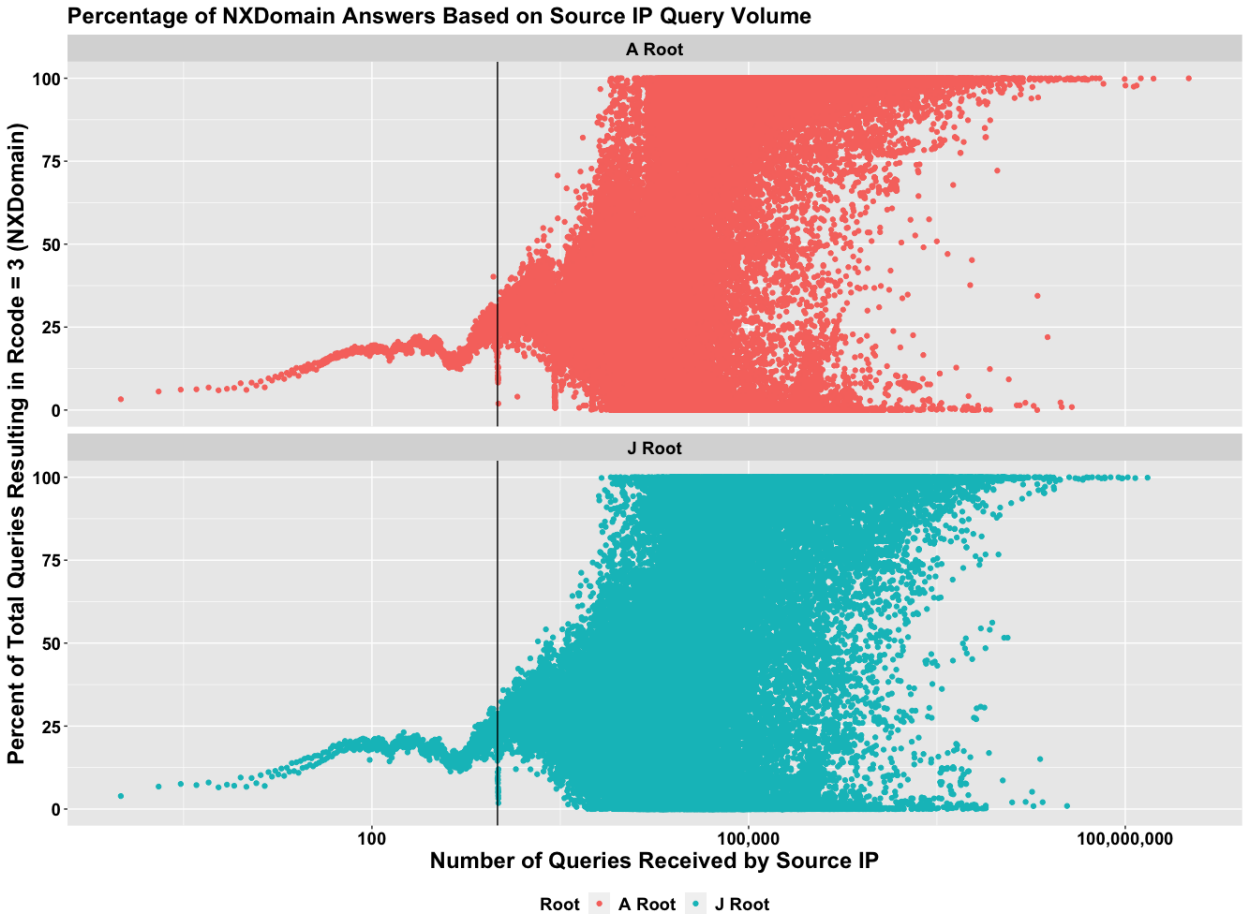

Figure 6 - Pct. of NXDOMAIN answers per source IP total query volume by RSI.

Understanding this behavioral profile of IP addresses helps inform traffic comparisons across RSIs. Any measurement of similarity will likely be very skewed by the nature of having so many IP addresses that account for negligible amounts of RSS traffic. Using a threshold limit to measure similarity does come at the cost of disregarding the longtail of low querying source IPs, but it will facilitate the intended measurement of this study. However, as seen in the previous Tables and Figures, those low querying source IP addresses and their associated queried names are not representative or relevant to the general context of name collisions and this study.

Using a threshold of 1,000 queries or more for a source IP address to be included in similarity measurements results in 1.3M unique source IP addresses. Those 1.3M IP addresses account for 99.5% of the queries observed in the 2020 DITL data. The threshold subset spanned 38,361ASNs and 0.56M unique /24 IPv4 blocks and 0.08M unique /48 IPv6 blocks.

From this point on, our analysis will be performed on the source IP addresses with total query volume greater than 1,000. This distribution shown in Figure 7 indicates 66.1% of these IP addresses are seen by all RSIs and 78.1% are seen at 6 or more RSIs- indicators that any RSI may be representative of the general RSS. This analysis was also conducted on the top 115K IPs and found that 89% of those IPs are seen by all seven of the RSIs.

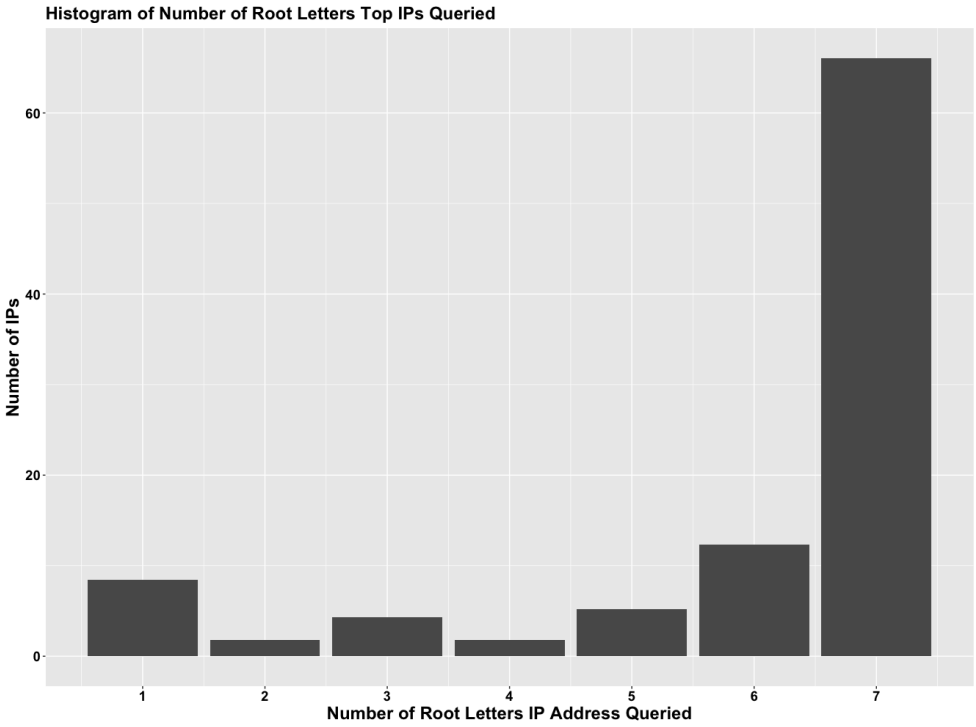

Figure 7 - Histogram of Number of Root Letters IPs Queried

A more detailed measurement of how the source IPs are observed at any two RSIs is depicted in Figure 8. The figure shows one-half of a similarity matrix that utilizes the Jaccard index, a similarity measurement that is further clarified in the Appendix, to measure the amount of overlap between two RSIs and the IPs that sent queries to them. From a source diversity perspective, any root letter, in general, sees a very high percentage of IPs compared to any other root. On average 86% of the IP addresses are observed at any two roots. This analysis was also conducted on the top 115K source IP addresses and found that 96% of the top 115K IPs are observed at any two roots. The IPs are widely seen at all root letters. Data from any combination of three RSIs will include 98.4% of IPs, though all RSIs must be included to reach 100% of IPs.

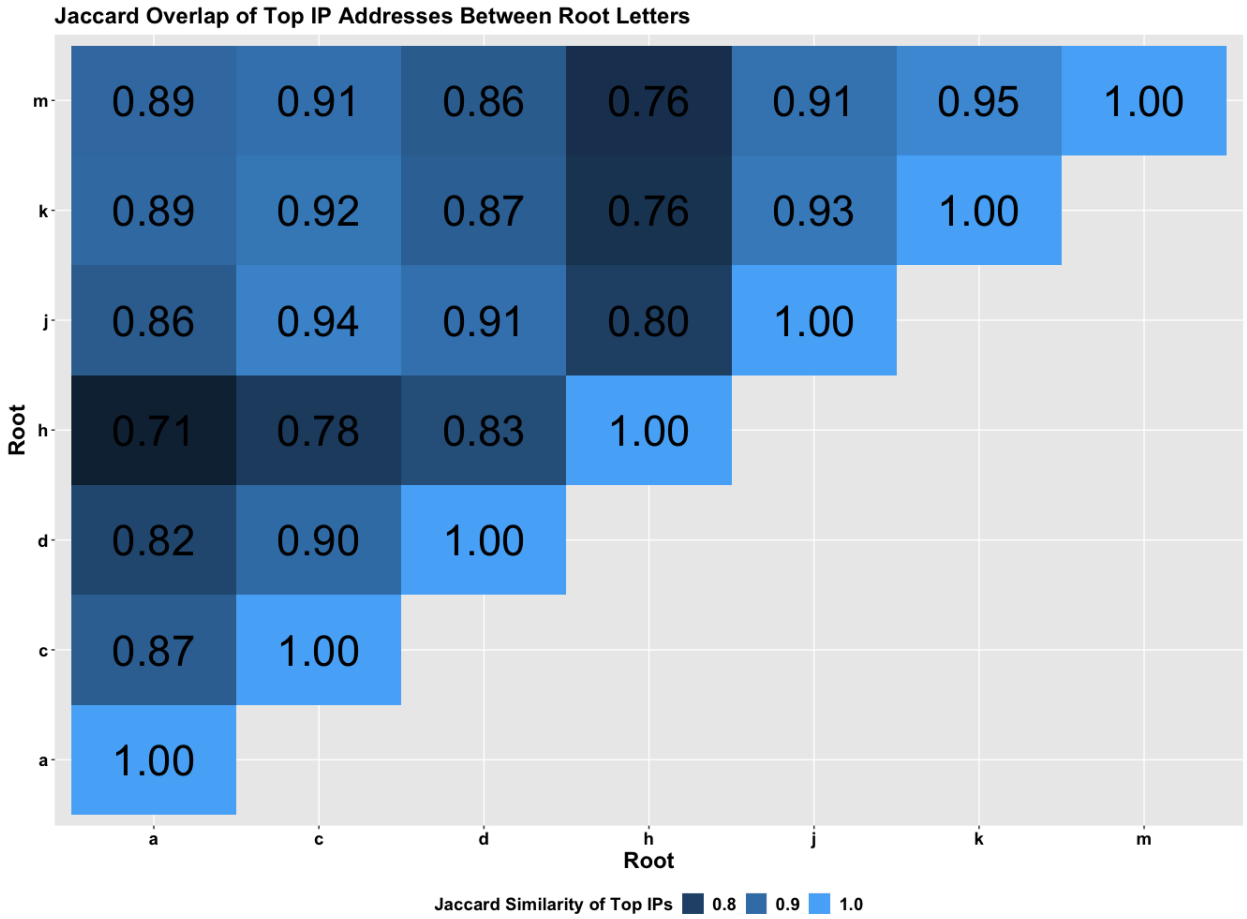

Figure 8 - Jaccard Overlap of IP Addresses Between RSIs

Geographic Relevance

The preceding measurements provided insights into the distribution of IPs over the RSIs. The following measurements continue to compare the distribution of those IPs from spatial and geographic means. Spatial representation of the IPv4 space is achieved via the use of a tool called IPv4 Heatmap.

IPv4 Heatmap is a program[9] that generates a map of IPv4 address data using a space-filling Hilbert Curve. Each pixel in the image represents a single /24 network and is assigned one of 256 colors. Pixel colors range from blue (1 host) to red (256 hosts), while black represents no data (0 hosts). Figure 9 below is an example of how an IPv4 spatial distribution can be visualized.

[9] Duane Wessels, ipv4-heatmap, last updated 1 March 2021, https://github.com/measurement-factory/ipv4-heatmap.

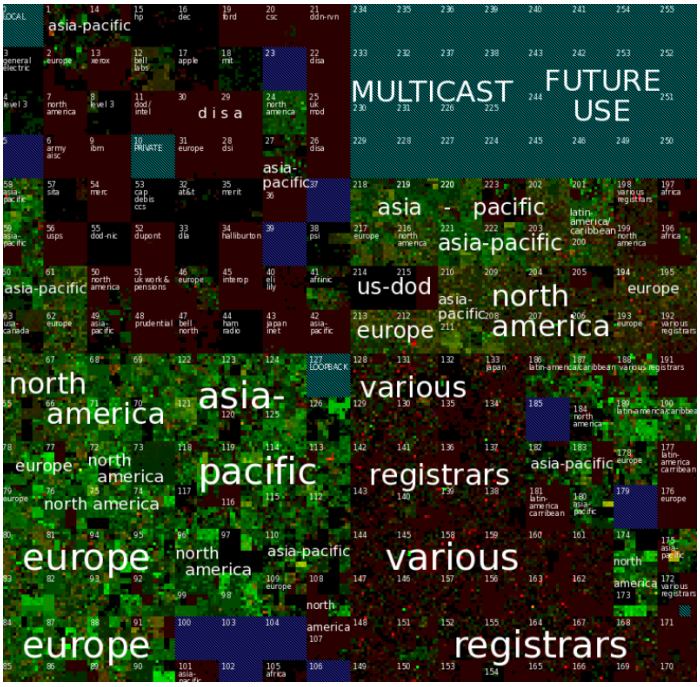

Figure 9 - Example of Hilbert Curve IPv4 Visualization

After exclusion of various RSIs due to IP anonymization or minimal data, seven RSIs were bucketed into the 256 color range by increments of 256/7. Each of those colors was then assigned in increasing order to represent the number of RSIs an individual IP address queried during the 2020 DITL. Blue dots are IP addresses that only queried 1 root while red dots represent top talker IP addresses that queried all 7 RSIs. As seen in Figure 10 below, there is a heterogeneous distribution across IPv4 address space. There are some notable exceptions in which several netblocks have a concentration. A few interesting groups of IP addresses, which queried only one or a few RSIs, appear in a small number of netblocks (e.g. 178.0.0.0, 172.0.0.0, etc.). It remains unclear as to what those resolvers are or their purpose without a more thorough analysis of their specific queries. Overall, these measurements indicate no large biases of source IP addresses showing specific RSI affinity.

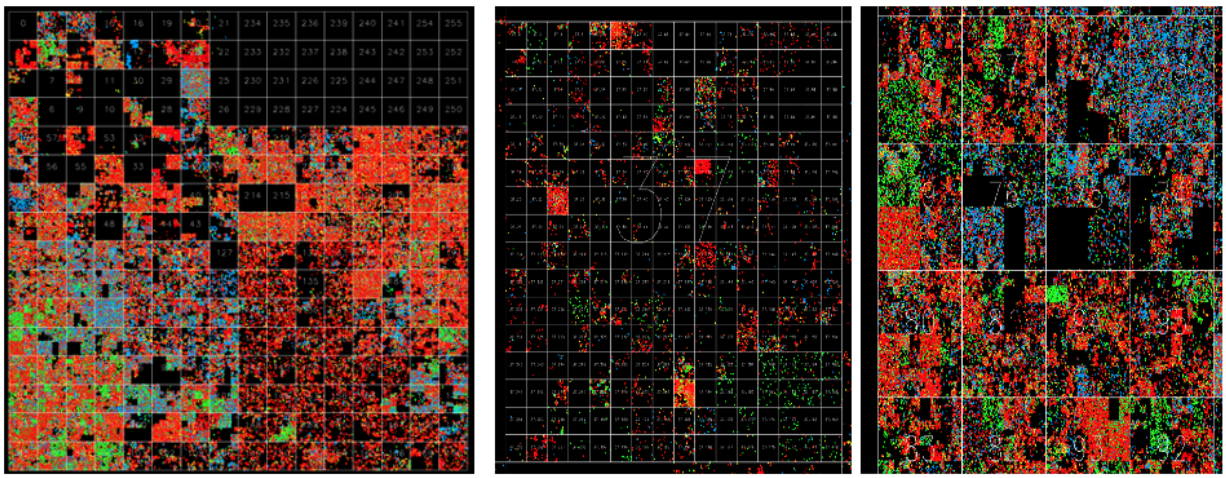

Figure 10 - IPv4 Spatial Distribution of Top Talker IP Addresses

Expanding into geospatial measurements, we next used the Gini coefficient to measure how much inequality an IP has for the distribution of root letters. We also geolocated the IP addresses to determine country-root inequality. The Gini coefficient measures the inequality among values of a frequency distribution (for example, levels of income). Gini values are bound between 0 and 1, in which a value of 0 would indicate the values are evenly distributed and 1 would indicate complete inequality.

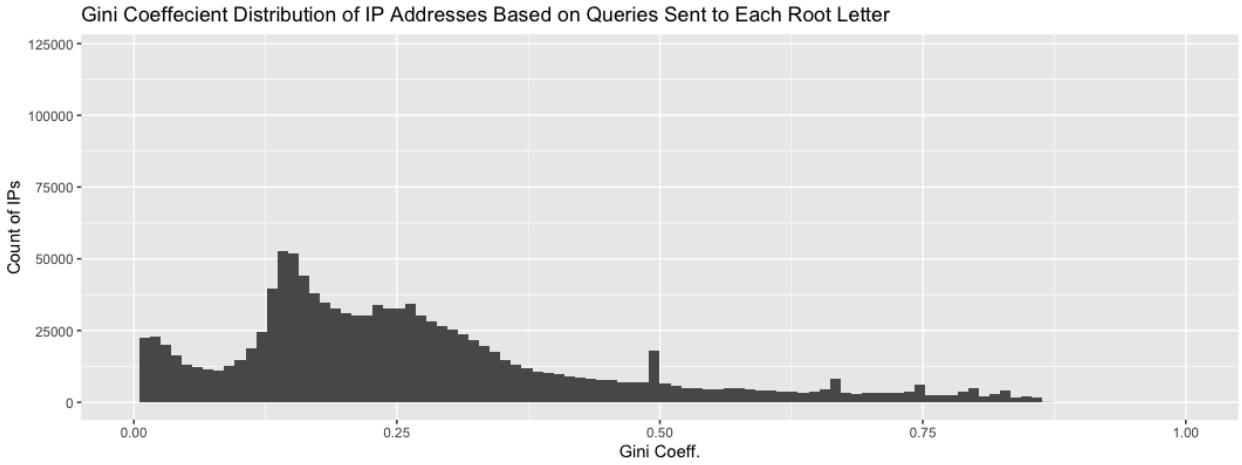

Figure 11 - Gini Coefficient Histogram of IPs

A Gini coefficient was calculated for each source IP address based on the number of queries the IP sent to each of the seven RSIs. Figure 11 shows the distribution of those Gini coefficients. While it appears to be multi-modal, the majority of the IP addresses resulted in values nearer to zero, indicating that these source IPs are distributing their query load over all of the participating RSIs. Likewise, the IPs were mapped to

countries using the Maxmind GeoIP database and the country traffic for each RSI was calculated. Figure 12 shows a geographical plot coloring in which the shading of the country is based on its Gini coefficient.

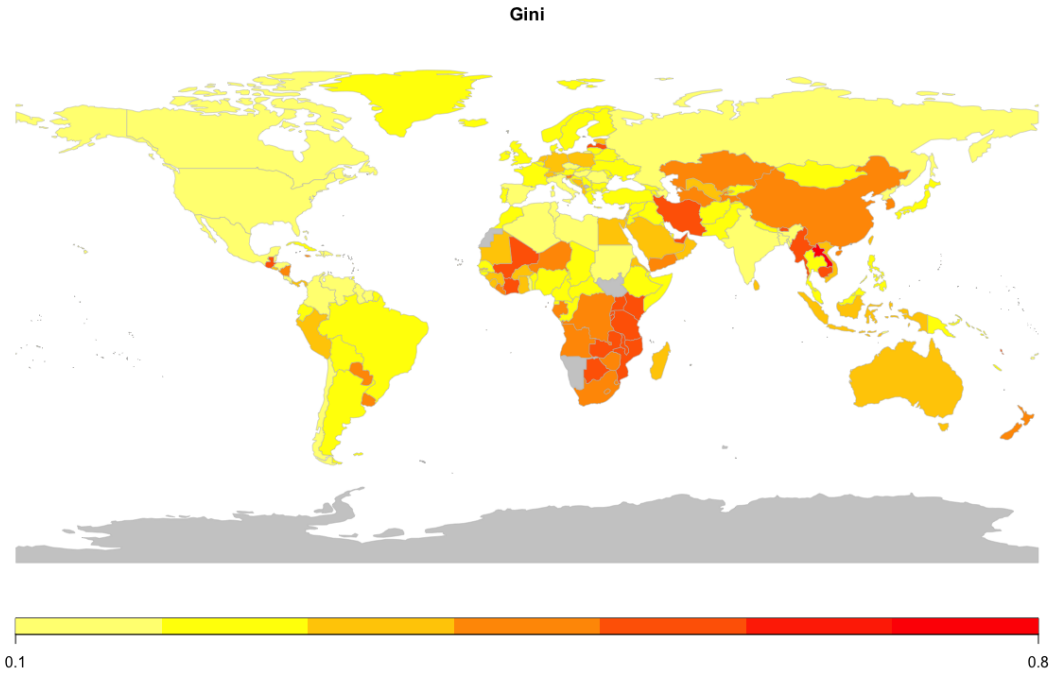

Figure 12 - Country to RSI Traffic Distribution

The overall per country Gini was an average of 0.304. Certain regions of Africa, Asia, and island countries have elevated Gini values and stronger affinities to certain root letters (it is expected this is likely due to placement/peering). An example of this bias/affinity can be seen in Figure 13 below, in which 2.8% of K-root traffic originated from Iran while other RSIs observed rates closer to 0.3%. Overall, this measurement helps confirm there is no large geographical bias of RSIs.

```
> country_root[CC == "IR",]
   CC Root CountryRootTotal   RootTotal CountryRootPercent
1: IR    a        114110956 23911477254         0.47722253
2: IR    c         70774822 18630303950         0.37989086
3: IR    d        127623076 26796152646         0.47627388
4: IR    h          4464366 12736101199         0.03505285
5: IR    j         61807444 37047068935         0.16683491
6: IR    k        919097084 32678054518         2.81258201
7: IR    m         54315964 15211977077         0.35706052
```

Figure 13 - Gini Coefficient for RSIs in Iran

Top N Non-Existent TLDs

One of the primary comments the NCAP DG discussed was that there is a need for clarity and predictability for applicants in future rounds of TLD applications and delegation decisions. To that end, having published lists of known name collision strings would benefit applicants. The publication of such lists may help inform applicants about name collision risks prior to their application; however, the absence of a string on a top-N list does not provide any assurance the string is void or absent of name collision risks

The previous analysis shows any RSI will see queries from roughly two thirds of what is seen more generally for a given 48-hour period. This is important because it provides some confidence that future name collision measurements could be taken by any RSI without requiring an RSS-wide collection. However, the data also suggests there can be no assurances that name collision measurements conducted by a single or a subset of RSIs to be complete.

In addition to looking at how representative traffic is received from querying recursive resolvers to RSIs, the following measurements will look at the similarity of the names. Specifically, the following figures and tables will examine what variation exists in the top N non-existent TLDs using both critical diagnostic measurements of query volume and network diversity.

In order to understand how top non-existent TLDs compare at each RSI, the entire set of non-existent TLDs were compared at A and J RSIs using the 2020 DITL data. The TLD strings were also required to match the regular expression [a-z0-9]{3,63}[10]. This resulted in 13.9 billion unique non-existent TLDs. This large number was caused by Chromium probing queries[11]. In order to more easily process these non-existent TLDs and to remove the garrulous noise of Chromium queries, a minimum of five queries was required for the TLD to be further analyzed, resulting in 4.84 million non-existent TLDs. For each TLD, the number of total queries, the number of distinct IPs, the number of distinct /24 and /48 netblocks, and the number of distinct ASNs was calculated.

Figure 14 below is a Venn diagram showing the overlap of the 4.84 million non-existent TLDs between A and J. Only 564K (12%) of the non-existent TLDs were seen at both RSIs. To better understand which TLDs did overlap at both RSIs, rank correlation plots for all three TLD measurements were evaluated.

---

[10] This regular expression was used because it will loosely match some of the ASCII label technical and policy string requirements set out in section 2.2.1.3.2 (“String Requirements”) of the evaluation procedures in the gTLD Applicant Guidebook (v. 2012-06-04) for DNS Stability. Many of the top non-existent TLD strings contain non-alphanumeric characters that do not match this requirement.

[11] https://blog.verisign.com/domain-names/chromiums-impact-on-root-dns-traffic/

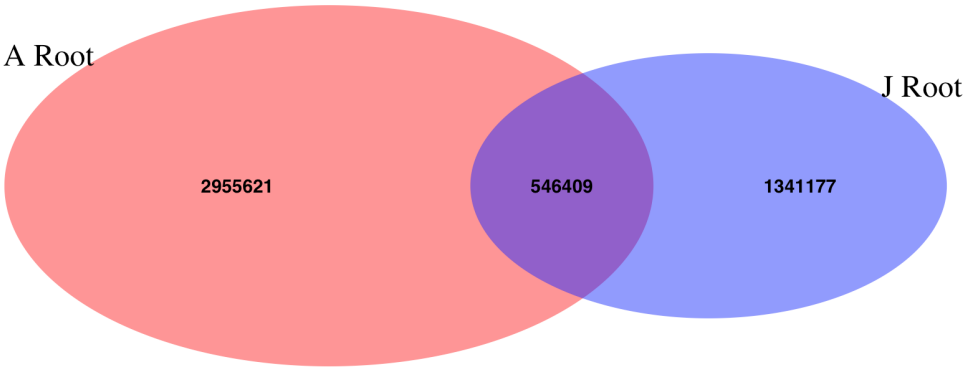

Figure 14 - Venn Diagram of TLD Overlap between A and J RSIs

If a TLD was observed at one RSI but not at the other RSI, a rank value of zero was assigned to that TLD. Thus any TLD depicted in Figures 15, 16, and 17 in which the dot is at x=0 or y=0 means that particular TLD was not seen by the other RSI.

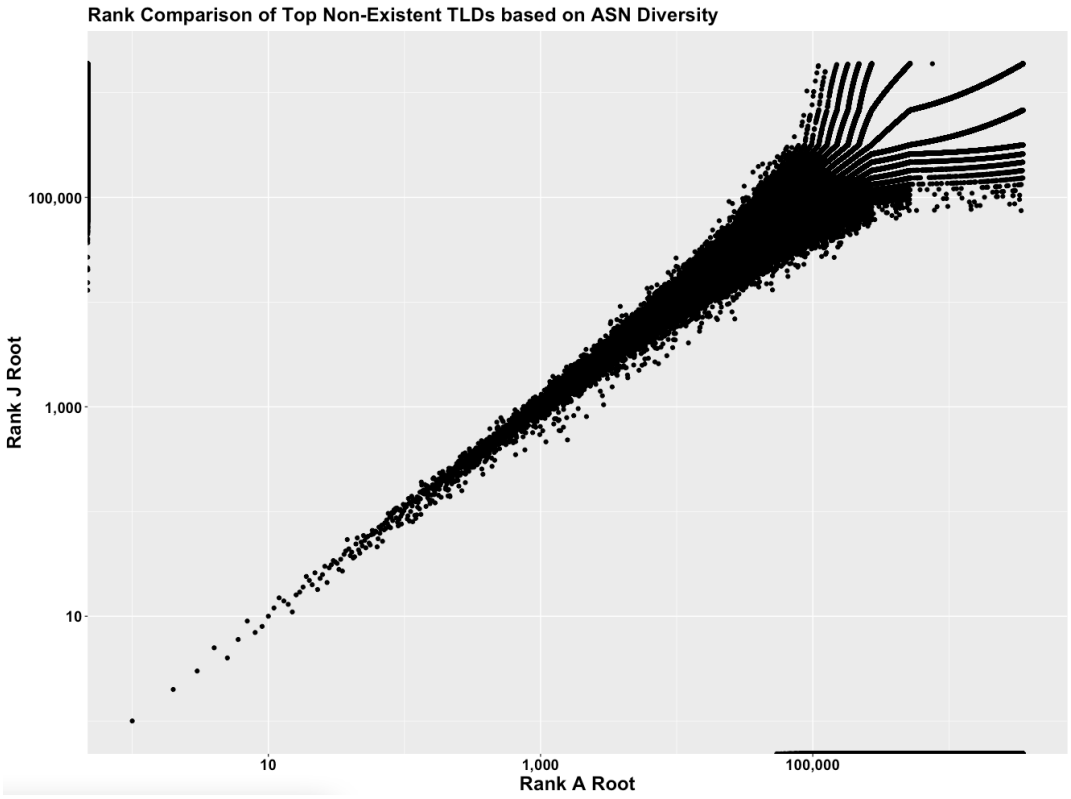

Figure 15 - Rank Correlation Between Top Non-Existent TLDs at A and J Based on ASN Diversity

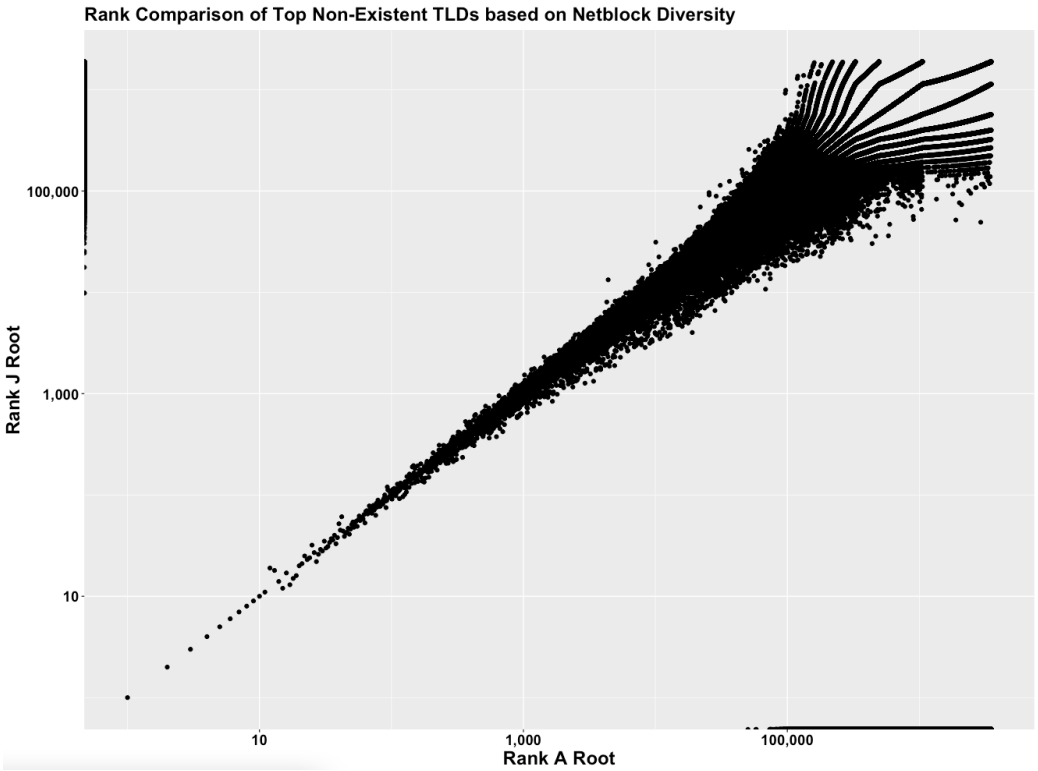

Figure 16 - Rank Correlation Between Top Non-Existent TLDs at A and J Based on Netblock Diversity

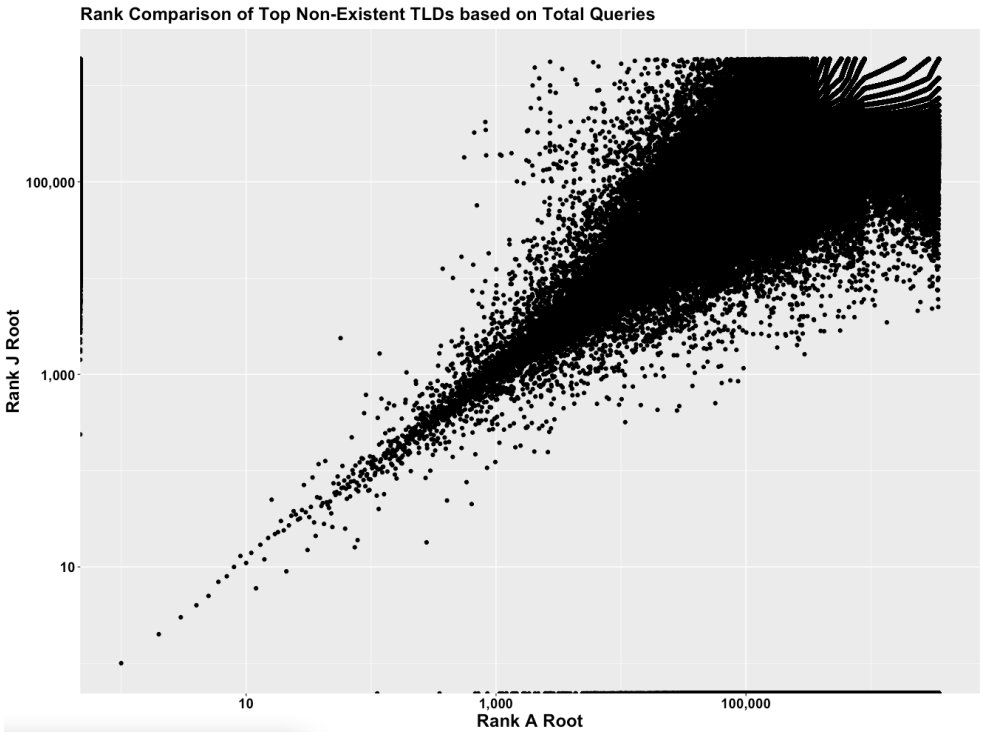

Figure 17 - Rank Correlation Between Top Non-Existent TLDs at A and J Based on Query Volume

Both network diversity measurements of ASNs and netblocks shown in Figures 15 and 16 exhibit a very strong rank correlation for the non-existent TLDs up to approximately rank 10K. Non-existent TLD strings observed at only one RSI became more frequent at rank levels above 100K. Meanwhile, query volume shown in Figure 17 also displays a strong correlation for the top non-existent TLDs up to rank 1,000 and non-existent TLD strings only observed at one RSI become more common after that level.

To understand how similar top-N lists are at various rank depths, Figure 18 shows the Jaccard value of the set similarity between A and J using the three critical diagnostic measurement ranking functions. Both network diversity measurements of netblock and ASNs show roughly 90% overlap until rank level 10K in which the overlap begins to degrade due to the TLDs being observed by just one of the RSI. Likewise, the query volume measurements show roughly 70% overlap until rank level 1K. These data suggest appropriate levels of N to be used for the publication of top-N lists. It also suggests that publishers of top-N lists should consider providing multiple lists based on the critical diagnostic measurement ranking methods and issue guidance to users that top-N data has limitations as to how complete it may be.

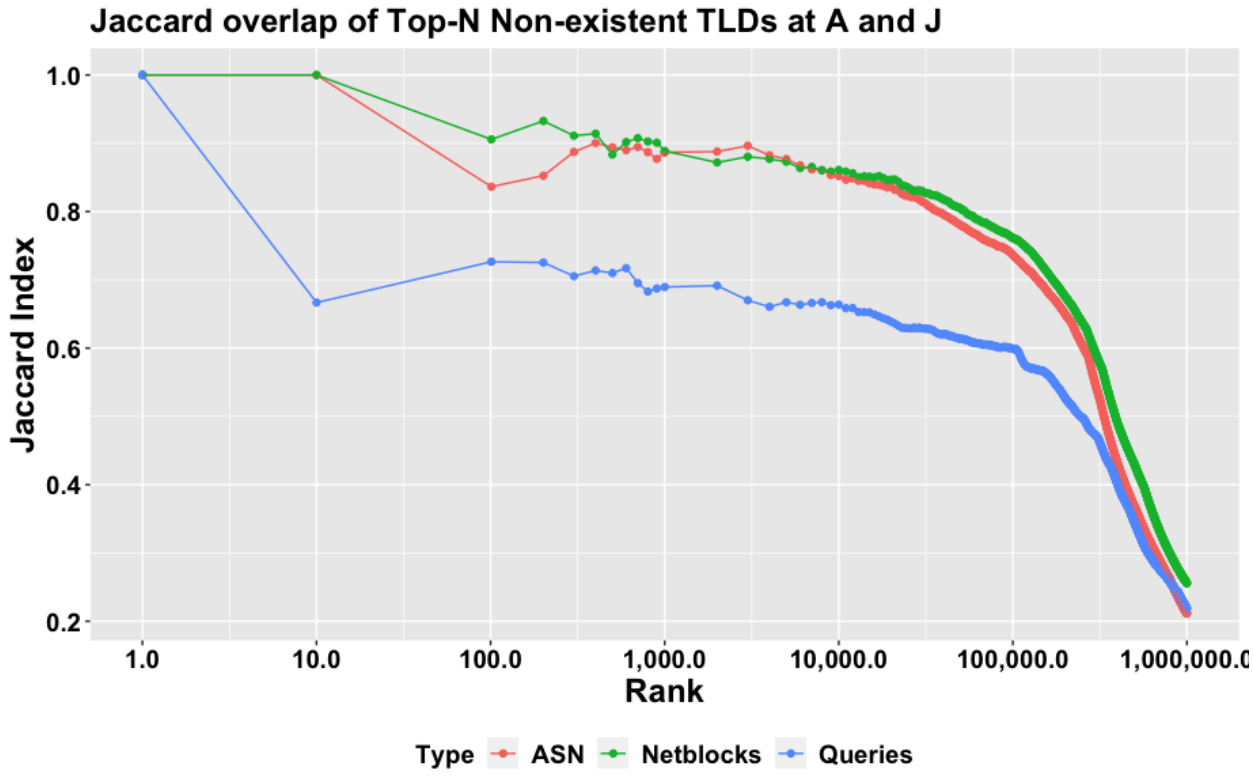

Figure 18 -Jaccard overlap of Top-N non-existents TLDs at A and J using different ranking criteria

Study 1 Key Observations:

In the analysis that we have conducted, we have observed the following:

- Ninety percent of RSS queries are sent from a relatively small set of IP addresses (115K).
- Low query volume source IP addresses express a different query behavior than the source IP addresses that constitute 99.5% of all queries received in 2020 DITL. They issue queries primarily for existent domains, RFC 8145 trust anchor signals, and oddly encoded DNS service discovery names.
- Source IPs that issued more than 1,000 queries are broadly seen at all root letters. Queries from 66% of those IP addresses are observed at all root servers. On average 86% of the IP addresses are observed at any two roots.
- Some geographic affinity/preference to certain root letters does occur.
- Initial research using two RSIs show top non-existent TLD strings between letters appear to generally correlate for the top 1K strings based on query volume and that lower volume non-existent strings appear to be more root letter dependent. Likewise, the data shows top non-existent TLD strings between letters appear to generally correlate for the top 10K strings based on network diversity measurements and that lower network diversity non-existent strings appear to be more root letter dependent.

Study 2: Recursive Resolver and Root Comparison

Since the 2012 round of TLD delegations, several new technologies and recommended best practices within the DNS ecosystem now have a significant impact on the volume and fidelity of DNS queries observed at name servers in the DNS hierarchy. The emergence of popular open recursive resolvers has also dramatically shaped the DNS ecosystem since the new gTLD delegations. These recursive services may provide a richer and more complete understanding of name collisions if they can be utilized for analysis. Therefore Study 2 was designed to investigate the differences of name collision strings at the RSS level as well as the recursive resolver level.

Data

In order to understand how DNS traffic compares at various layers of the DNS hierarchy, query data sent to several root server identifiers and one public recursive resolver and one commercial recursive resolver were collected and measured in such a way that would facilitate the examination of top non-existent TLD observed in queries. The data was measured using two sorting functions that reflect the importance of our critical diagnostic measurements: (1) Query Volume and (2) IP Address diversity. Two lists of the top 1000 non-existent TLDs matching the regular expression [a-z0-9]{3,63} were generated based on the two sorting functions. The resulting aggregated data was used to measure how recursive and root server query volume compare by examining rank ordering as well as general TLD string overlap.

Notable Limitations of the Data

While efforts were made to obtain recursive resolver data from numerous sources, only two recursive resolver operators provided the data. The limiting factor appears to be data privacy concerns. To that end, the recursive resolvers that did provide the data will not be identified and herein simply referred to as the "public recursive resolver" (PRR) and the "recursive resolver" (RR). Without obtaining data from other public recursive resolvers, it is unclear how each recursive resolver compares to another. It is likely due to their underlying user-base, deployment size, and internal DNS protocol optimizations, that each recursive resolver represents a unique vantage point of the DNS; however, without additional data this will remain only a hypothesis. The measurements presented in this study, while only looking at the PRR and the RR, do provide a novel and previously unknown understanding of name collisions via passive DNS telemetry data used for quantifying and assessing name collision risks at multiple collection points within the DNS hierarchy.

The patterns of DNS queries observed at recursive resolvers are expected to be very different from those observed at the root server system because of the very nature of their roles. Two specific characteristics that differentiate the two are caching and client aggregation. These characteristics complicate comparisons between the two systems. Caching makes the comparison of query volume between the two sources unreliable because recursive resolvers presumably receive many more queries for a given domain name than will ever show up at any authoritative server, including the root servers. This, in turn, makes

ranking based on query count prone to error. Yet the data obtained from the PRR and RR contains only the non-existent TLDs most queried for. Client aggregation refers to the fact that a recursive resolver can serve stub resolvers (i.e., end-user systems) and/or other recursive resolvers. The diversity of query sources can mean different things between PRR and RR and between either of these and the root servers system. Based on the challenges inherent in comparing these very different data sets, we recognize that this analysis is merely a heuristic and should not be taken as definitive.

PRR and RR Source Diversity and Query Volume vs. A and J RSIs

Comparing PRR and RR data with RSI data is problematic due to their unique positions within the DNS hierarchy. For the purpose of this study, efforts were made to compare the top 1K PRR and RR non-existent TLDs collected at two sources to all non-existent TLDs seen at A and J based on both query volume and source diversity. The findings reveal some interesting insights. Figure 19 shows that both using query volume and distinct IP addresses, the first 100 top non-existent TLD strings roughly correlate between the PRR/RR and the RSI. However, higher ranking non-existent TLDs exhibit huge discrepancies (several orders of magnitude) between the PRR/RR and RSI ranking. From a name collision perspective, this suggests that even if a non-existent TLD has a very high rank based on RSI data, that measurement may not reflect the entire name collision impact posed by that string. RSI measurements quantify global recursive resolver behavior, while the PRR and RR data is measuring impacted clients. In the case where PRRs and RRs have strong concentration/localization/preference to a set of clients, a low rank in the PRR/RR and a very high rank in RSI would be expressed – which again suggests that a complete and accurate picture of name collision issues cannot be assessed solely on either PRR, RR, or RSI data.

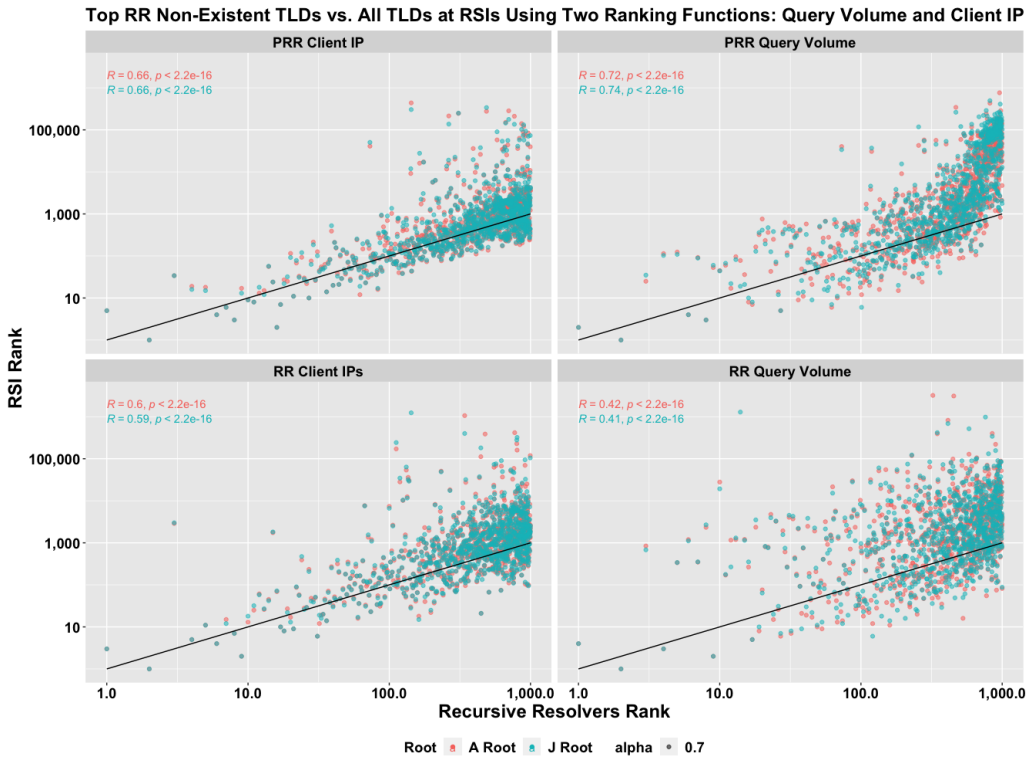

Figure 19 - Rank comparison of PRR and RR non-existent TLDs based on source diversity and volume

Figure 20 shows the distribution of the ratio of rank at the PRR and RR to the rank at RSI. Essentially, think of it as a distribution of x-rank divided by y-rank, in which an equal ranking would result in the value of 1. Most TLDs exhibit a +/- 1 magnitude difference in ranking but there is a subset of the top 1-K PRR and RR non-existent TLDs that exhibit differences of more than 3+ orders of magnitude. Showing that the top-N at a given PRR or RR can be significantly different than how an RSI may quantify that string.

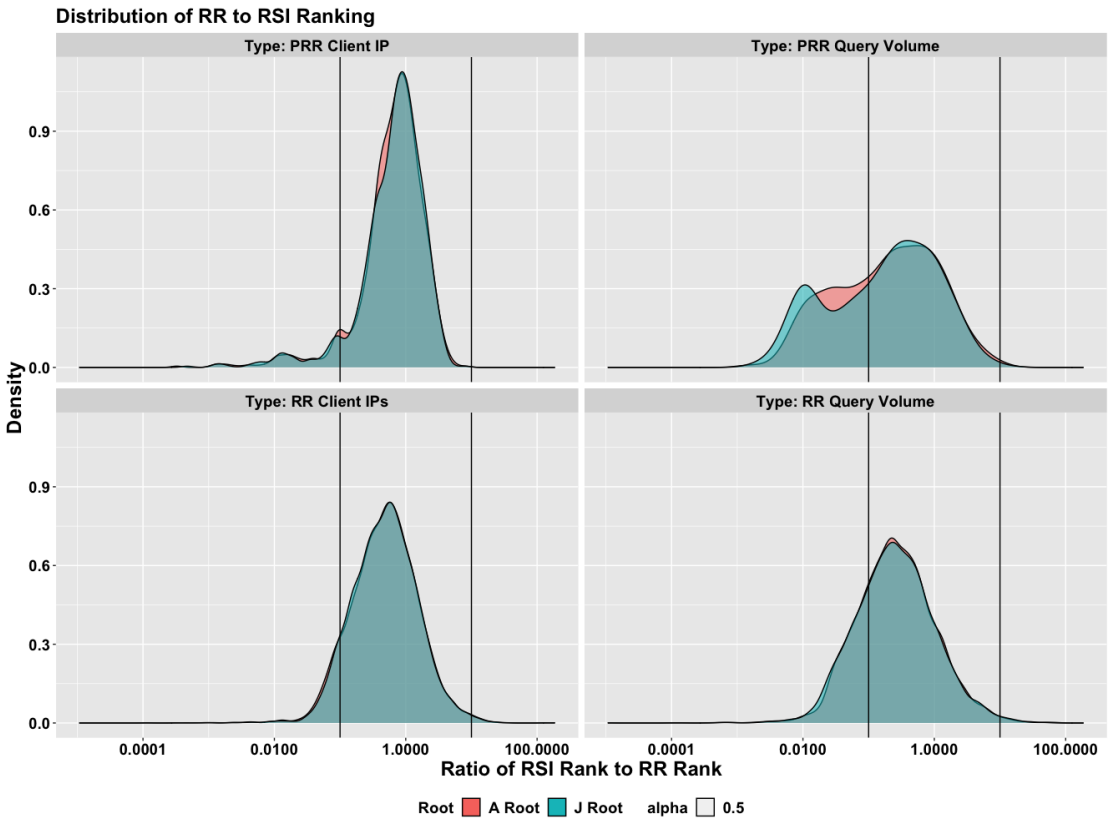

Figure 20 -Distribution of ratio of RSI rank to PRR and RR based on diversity and volume

Figures 21 through 24 show Venn diagrams depicting the various overlap of the top-N strings from the PRR, RR, and A and J RSIs. The PRR provided exactly the top 1K non-existent TLDs based on client diversity while the RR provided the intersection of the top 10K based on query volume and client diversity - resulting in 8020 strings. Figure 21 compares the overlap of the PRR and RR and shows that 46% of the PRR top non-existent TLDs are not observed by the RR. This suggests that each recursive resolver will have a unique observation space based on the end clients using the resolver.

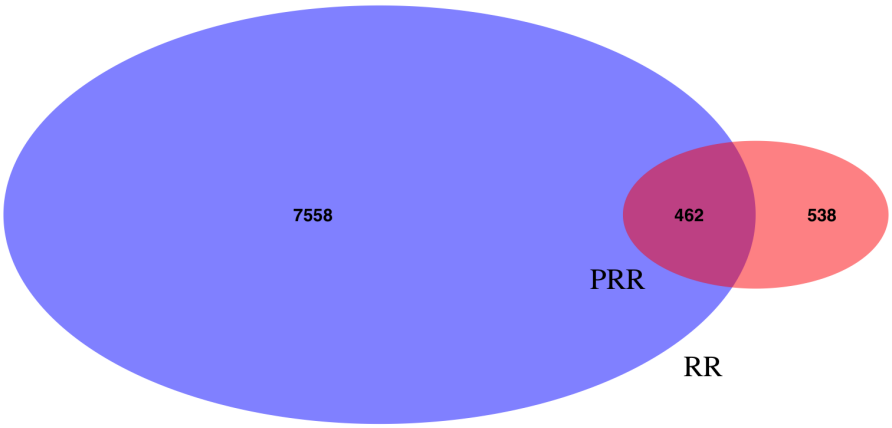

Figure 21 -Venn diagram showing overlap between PRR and RR using source diversity

Figure 22 and Figure 23 both compare the overlap of the PRR and RR to A and J RSIs. Here we see that except for a small number of top non-existent TLDs seen at the RR, that nearly all of the strings are observed by either A or J RSIs.

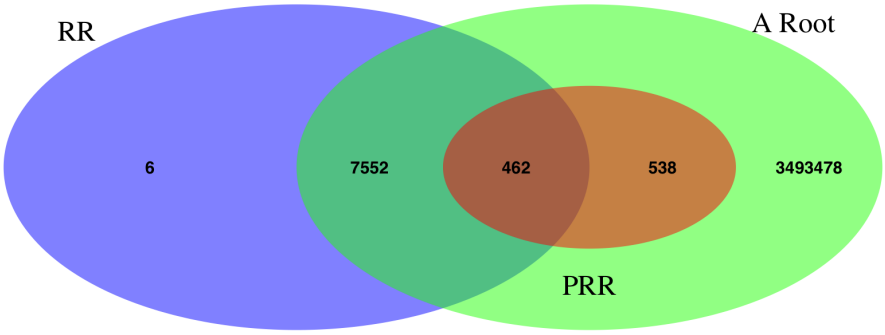

Figure 22 -Venn diagram showing overlap between PRR, RR, and A root using source diversity

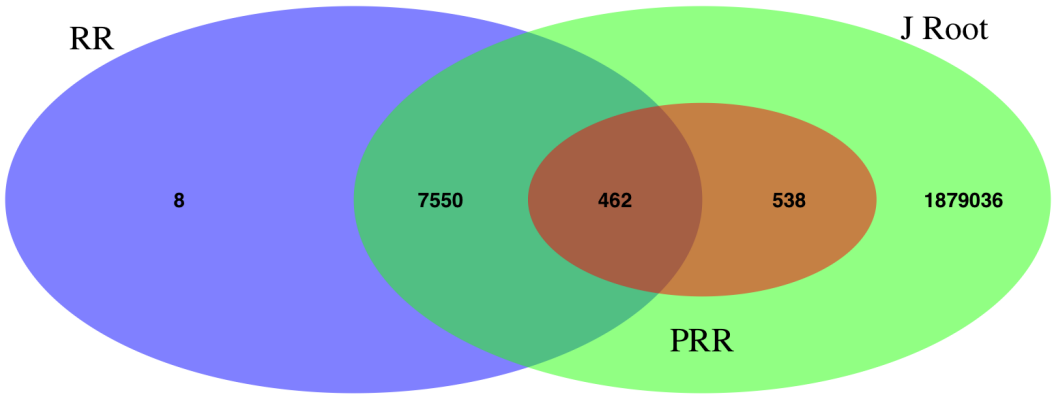

Figure 23 -Venn diagram showing overlap between PRR, RR, and J root using source diversity

Figure 24 shows a Venn diagram overlap of all four data sources: the PRR, the RR, A and J RSIs. Here we see that all top non-existent TLDs from the PRR and RR are seen by the RSIs. Additional supporting measurements looking at the PRR and RSI instances are in Annex 2.

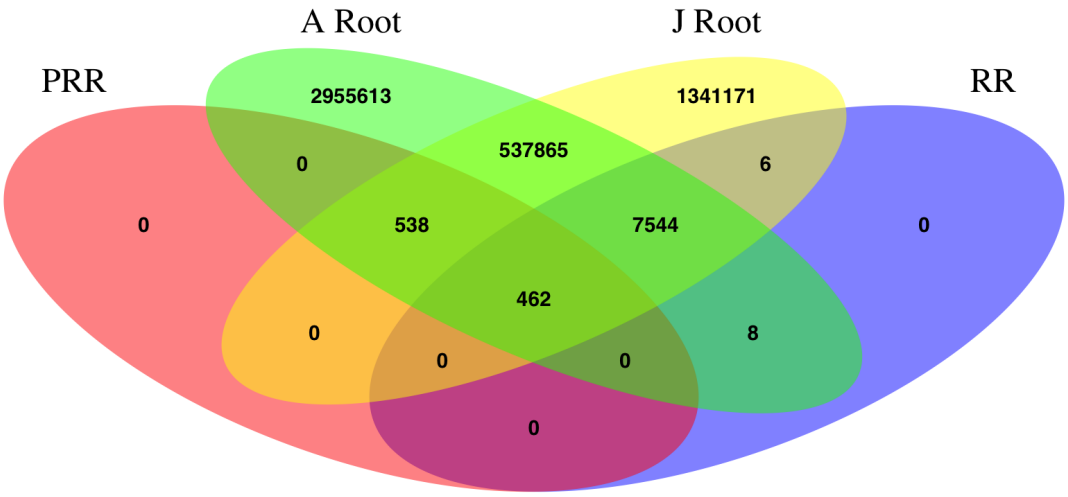

Figure 23 -Venn diagram showing overlap between PRR, RR, and A & J roots using source diversity

Study 2 Key Observations:

In the analysis that we have conducted, we have observed the following:

- Initial results from one PRR and one RR indicate there is a difference in top non-existent TLDs using either query volume or source diversity measurements. The top 1K non-existent TLDs seen at either the PRR or RR can be upwards of 3+ orders of magnitude different from RSI rankings.
- Significant differences between the PRR and the RR non-existent TLDs and those at the RSIs suggest that name collision strings cannot be measured or assessed for their presence within the context of a top-N list with an arbitrary N cutoff reliably using data only from either the RSS or the PRR/RR..

Key Findings

The two studies in this analysis provide two key findings that will help the NCAP provide guidance and advice to ICANN as to how future risk assessments of name collision strings should be evaluated. The risks of name collision strings have been cataloged in the NCAP Study 1 report.

Finding 1: The IP addresses meeting a defined threshold of root server queries are 66% likely to query any given root server in the course of two days, and queries from those IP addresses for non-existent TLDs are likely to be found with similar prevalence on different RSIs. However, the data also indicates this similarity degrades as the query volume and network diversity decrease for non-existent TLDs, in such a way that measurements from a subset of RSIs will not ensure a complete and accurate assessment of a name collision string.

Implications:

- DITL-like collections are extremely valuable for both name collision assessments and comparison across root server letter perspectives. Nevertheless, when such a comprehensive data set is not available, data collected from a subset of root servers can represent a significant fraction of queries observed by the whole system. Even a single root server observes traffic from roughly two thirds of DNS clients that meet a defined threshold of query activity.
- PRR and RR data further indicates that there is a very different view of the top non-existent TLDs based both on query volume and source diversity within the context of a top-N list.
- Appropriate advice should be issued that the presence of a string on top-N lists suggests existing name collision issues; however, the absence of a string on a top-N list does not provide any assurance the string does not have name collision risk.

Finding 2: Name collision traffic observed at the root within the context of top-N is not sufficiently representative of traffic received at recursive resolvers to guarantee a complete and or accurate representation of a string's potential name collision risks and impacts.

Implications:

- Name collision traffic observed via root server telemetry data should be considered the minimal recorded value.
- Given the current and likely future state of the DNS ecosystem, a complete and accurate risk assessment of a string's name collision potential cannot be determined prior to the string's temporary delegation.

Annex 1: Statistical Methods

Jaccard Index

The Jaccard index, also known as the Jaccard similarity coefficient, is a statistic used for gauging the similarity and diversity of sample sets.

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|}.$$

Jaccard index measurements are bound between 0 (identical sets) and 1 (completely distinct sets).
Note: This measurement is only accounting for set presence. It does not consider the magnitude/volume of queries sent - it is only if the IP appears in both sets.

Gini Coefficient

$$G = \frac{2\sum_{i=1}^{n} i y_i}{n\sum_{i=1}^{n} y_i} - \frac{n+1}{n}.$$

The Gini coefficient measures the inequality among values of a frequency distribution (for example, levels of income). Gini coefficient measurements are bound between 0 (even distribution) and 1 (completely uneven, e.g., one member receives all traffic).

Annex 2: Additional Recursive Resolver Measurements

Measurements

The following five measurements were conducted against the data:

1. Query volume distribution of RSIs and the PRR
2. Rank correlation between RSI and PRR based on query volume
3. String overlap between RSI and PRR based on query volume
4. Rank correlation between RSI and PRR based on source diversity
5. String overlap between RSI and PRR based on source diversity

Total Query Volume per TLD Distribution

A baseline measurement comparing query volume of the top 1,000 non-existent TLDs at two RSIs, A and J roots, and the PRR is depicted in Figure 16 below. The distributions appear similar in nature, forming a power-law distribution in which the top non-existent TLDs express query volumes that are several magnitudes higher than the other TLDs. All three distributions seem to "flatten out" into the long tail distribution after the top 50 TLDs.

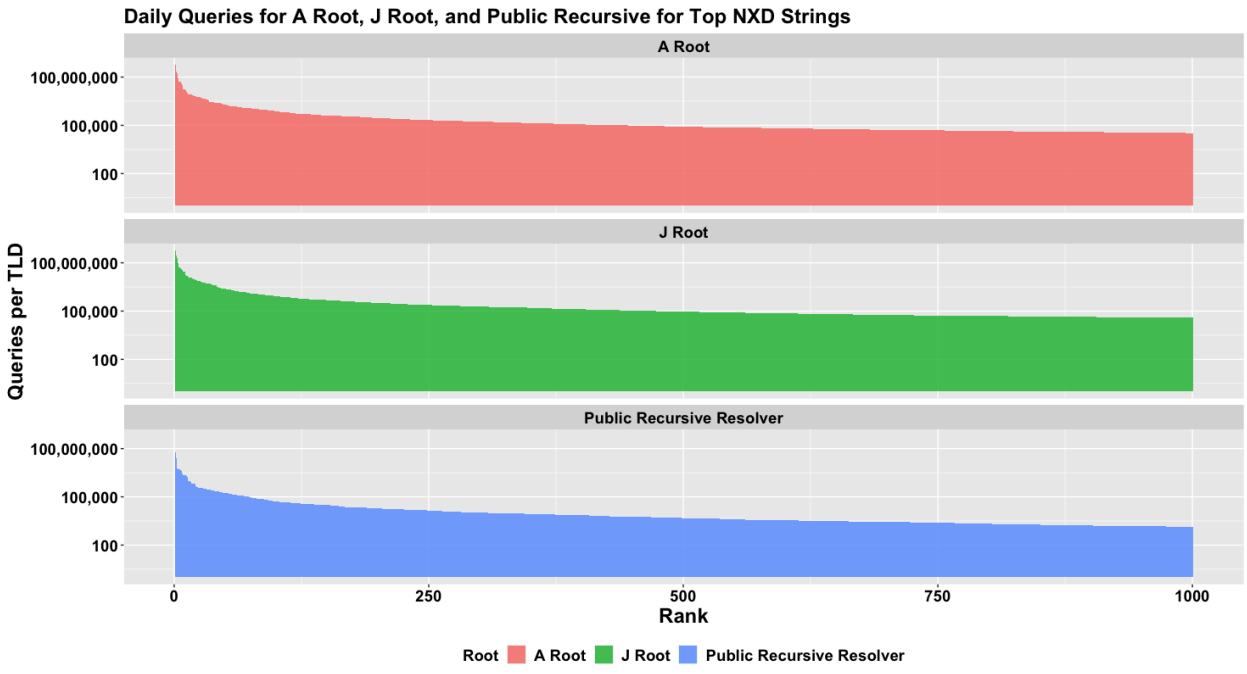

Annex 2 Figure 1 - Daily Queries for A and J RSIs and the PRR for Top NXD TLDs

A and J Root Servers Compared to a PRR Using Total Query Volume per TLD Ranking as a Function

While the initial query volume distribution shown in Figure 16 may have shown some similarities, no other strong similarities were found between the RSIs and the PRR data. Figure 17 below shows a simple scatter plot of the top RSI TLD rankings vs. those of the PRR. Unlike the rank scatter plots comparing top RSI TLD rankings relative to another RSI, the RSI to PRR plot shows no correlation between the two DNS data sets (e.g., there is no "diagonal" line with a slope of ~1).

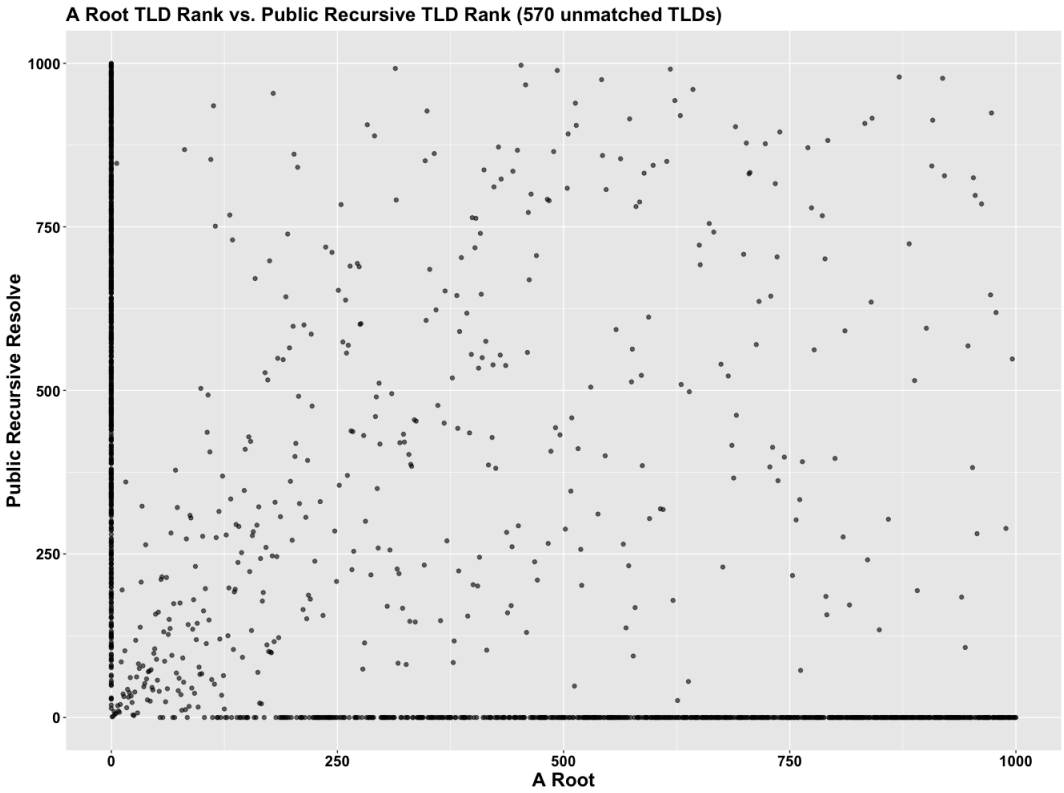

Annex 2 Figure 2 - Rank Correlation of Top TLDs at A Root and Public Resolver based on Query Vol.

This lack of correlation shown in Figure 17 is better explained by looking at the Venn diagram that examines the set overlap of the top 1,000 non-existent TLDs. Only 430 strings were both observed at the RSI and the PRR. This is significantly different from the overlap previously seen between RSIs in which ~800 of the strings overlap.

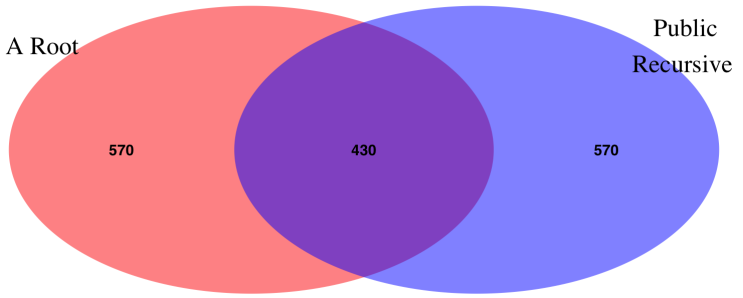

Annex 2 Figure 3 - Venn Diagram showing TLD overlap of A Root and PRR based on Query Vol.

Figure 19 below is another examination of a ranking scatter plot at a second RSI. Again no correlation is observed between the RSI and the PRR. This is again reconfirmed by the Venn diagram in Figure 20, in which only 417 of the top non-existent TLDs were observed by both the RSI and the PRR.

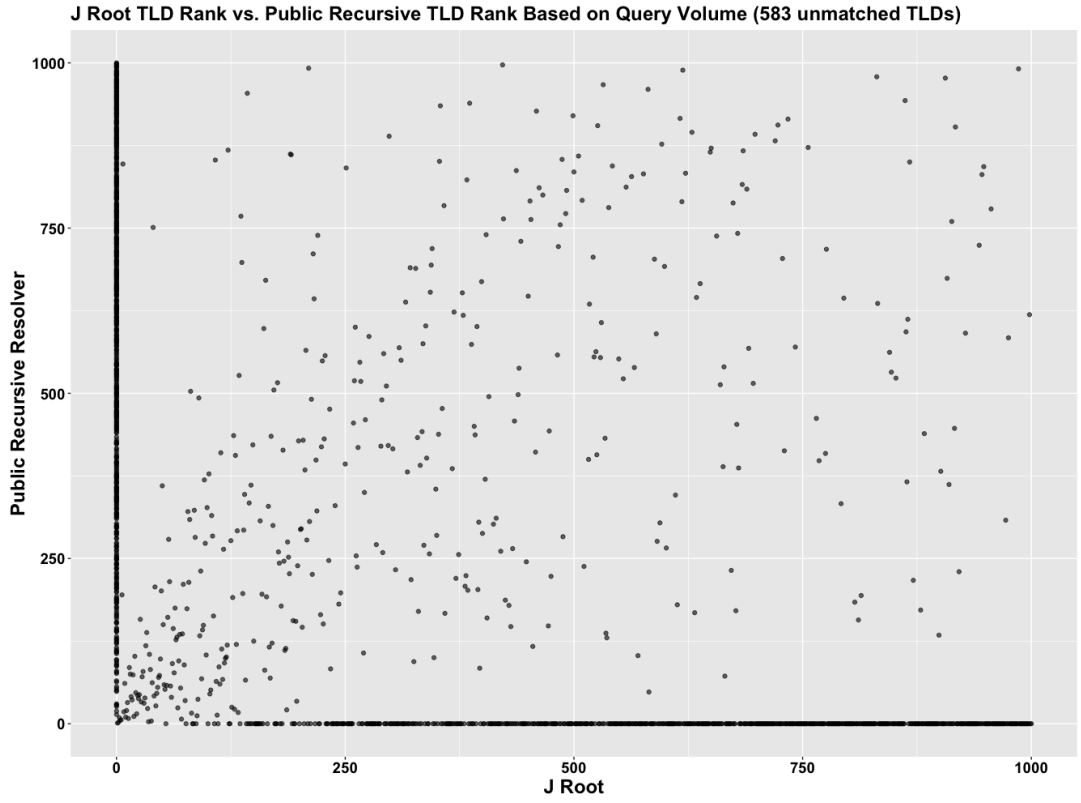

Annex 2 Figure 4 - Rank Correlation of Top TLDs at A Root and the PRR based on Query Volume

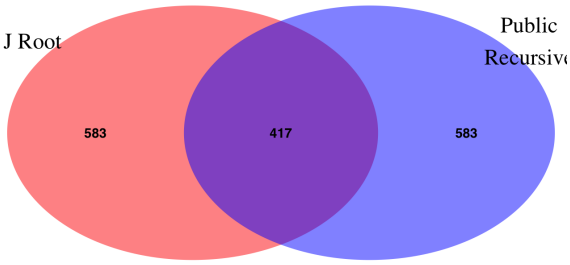

Annex 2 Figure 5 - Venn Diagram showing TLD overlap of the J RSI and PRR based on Query Vol.

These initial comparisons of top strings based on query volume observed at RSIs and the PRR reveal there is a significant difference in DNS queries. The small overlap of top strings between the two data sources further suggests that an accurate and complete picture and risk assessment of collision strings is not possible from RSS data alone. Figure 21 below show the top 50 non-existent TLD strings observed at A root and the PRR.

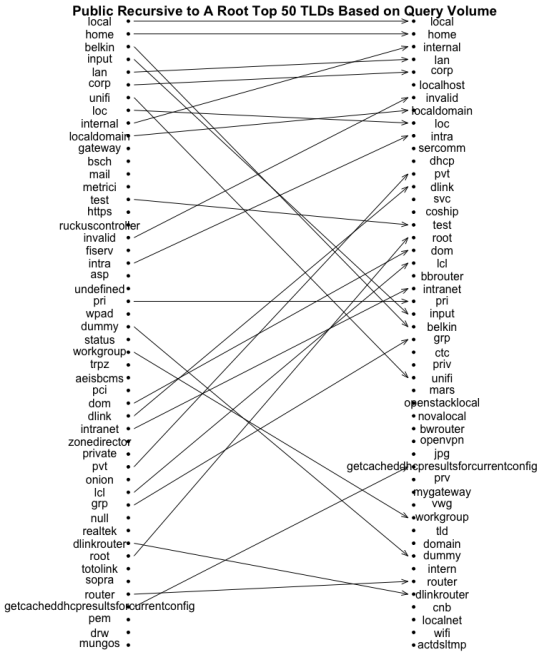

Annex 2 Figure6 - PRR to A Root Top 50 TLDs Based on Query Volume

A, J, and L Root Servers Compared To Public Recursive Using Distinct Source IPs per TLD Ranking Function

The top 1,000 non-existent TLDs were identified at each of the RSIs and the PRR based on the number of unique IP addresses observed per TLD. An initial measurement looking at TLD string overlap via a Venn diagram is shown in Figure 22 below. The PRR still observed 311 strings which none of the RSIs observed in their top 1,000. This measurement shows greater overlap between RSIs and the PRR than top strings by query volume. However, the significant dissimilarity between the PRR TLDs with the greatest source IP diversity and those of the RSIs means that name collision strings cannot be measured or assessed properly based on only using data from the RSS.

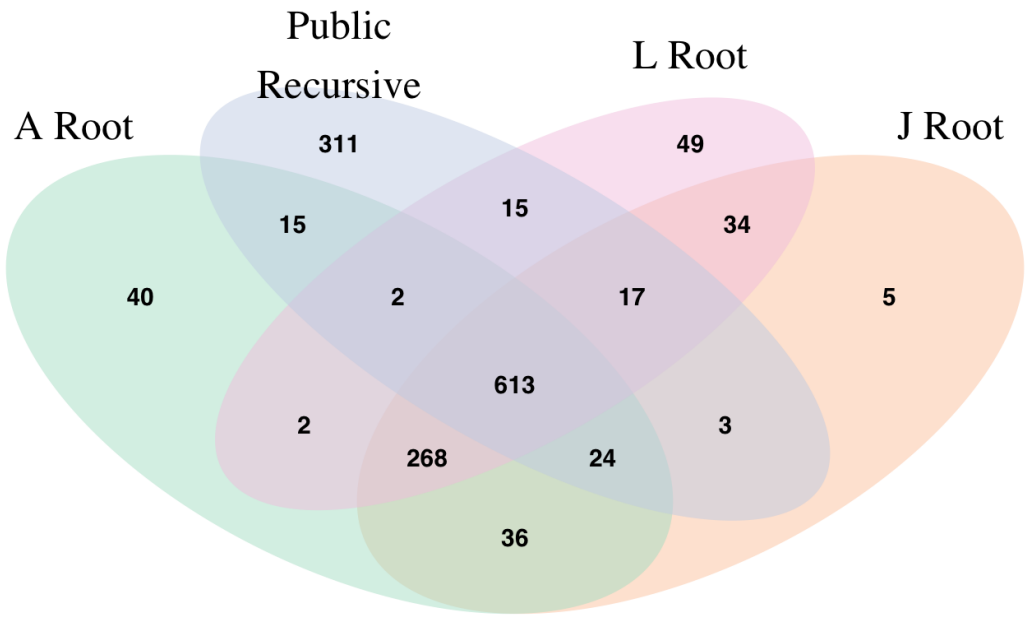

Annex 2 Figure 6 - Venn Diagram Comparing Overlap of Top TLDs at A, J, and L RSIs and PRR based on Source IP Address Diversity

Examining a rank scatter plot between an RSI and the PRR does indicate a slightly better correlation of TLD rankings; however, this correlation appears slightly weak (correlation coefficient of 0.36), at best, and mainly for the top-ranking strings that had large source diversity measurements (i.e., TLD rankings under 100).

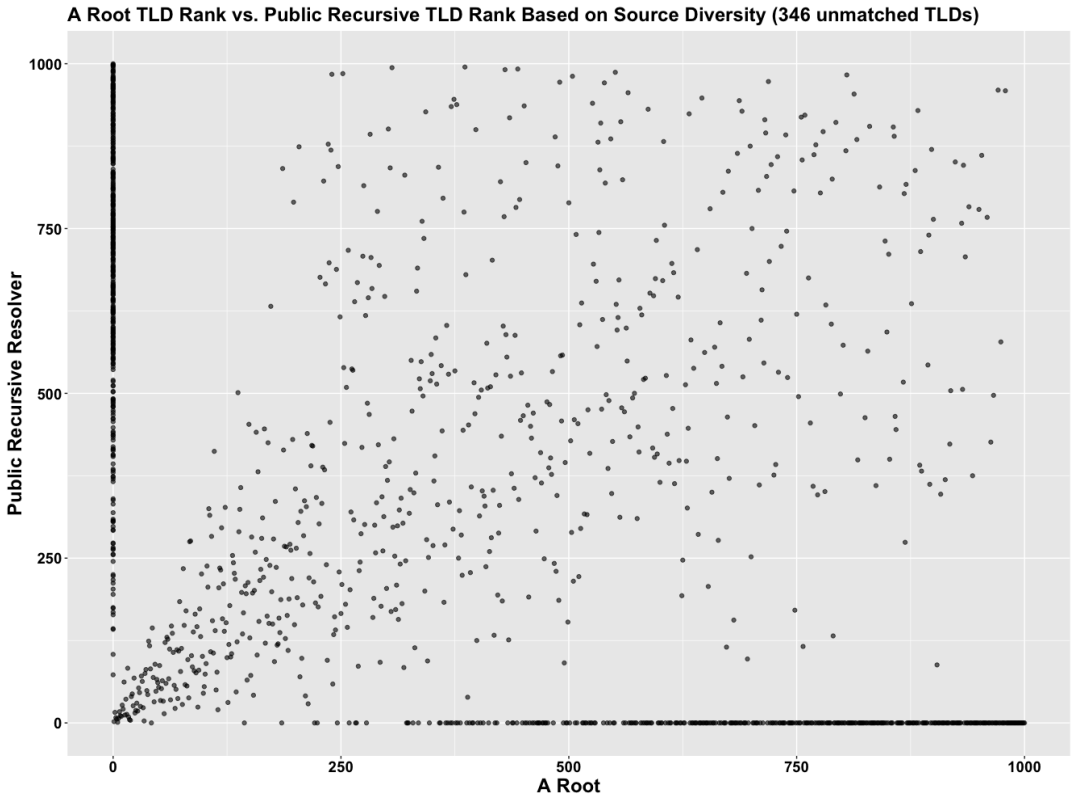

Annex 2 Figure 7 - Rank Correlation of Top TLDs at A Root and PRR based on IP Diversity