Case Study of Collision Strings

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

The NCAP discussion group met over the course of approximately two years to evaluate and consider questions posed by the ICANN Board on the delegation of currently reserved TLDs such as .CORP, .HOME, and .MAIL. The group undertook a review of past studies and literature, and conducted its own analysis from two root server identities. The result of that review is a modern picture of the impact and potential harm due to name collisions with the undelegated names under study. The analysis provides a sufficient basis from which to draw a number of important findings. Among these include the observation that queries for these undelegated names are both increasing in volume and diversity. These facts suggest that challenges relating to impact and mitigation are also increasing. The group also identified a number of critical diagnostic measurements that help determine the scope, impact, and potential harm of name collisions.

1   Introduction

In resolutions (2017.11.02.29 - 2017.11.02.31)[1] the Internet Corporations for Assigned Names and Numbers (ICANN) Board requests the ICANN Security and Stability Advisory Committee (SSAC) conduct studies to present data, analysis and points of view, and provide advice to the Board on the topics around DNS name collisions. In response, SSAC formed the Name Collision Analysis Project. This project is organized into three studies. The first study,[2] which provided a primer on the topic of name collision and a list of datasets that either existed at the time of the study or would need to be generated to support further analysis, was finalized on 19 June 2020 and submitted to the Board.

The second study, with a somewhat revised scope as described in "SSAC2021-02: Revised Study Two Proposal for the Name Collision Analysis Project,"[3] focuses on four key areas:

- Perform a study of ICANN Collision Reports
- Perform an Impact and Data Sensitivity Analysis with respect to name collisions
- Respond to Board Questions relating to Study Two
- Produce a final report on Study Two

To aid the deliberations of the four key areas above, the NCAP proposes the following case studies be done by the project:

> "Using the similar data sources and methodologies by JAS Global Advisors and Interisle Consulting Group, perform updated case studies of the CORP, MAIL, HOME, and other strings. The study should highlight changes over time of the properties of DNS queries, and traffic alterations as a result of DNS evolution." (page 6 of NCAP Revised Study 2 Proposal)

Later part of the proposal asked for case studies of CORP, MAIL, HOME, and non-delegated strings that receive more than 100 million queries per day at the root. Using this threshold and DNS query data from A and J root servers results in six strings:.CORP, .HOME, .INTERNAL, .LAN, .LOCAL, and .MAIL.

This report, produced by the NCAP discussion group, documents the result of the case study of the six strings, and answers the two study objectives listed above. The rest of the report is organized as follows. In section 2, we review several important past work on name collision as a reference point for the current analysis. In section 3, we describe the data sources and data analysis methodology used in our study. In section 4, we describe the results of the case studies. In section 5, we analyze the result and discuss several key findings. In section 6, we summarize our findings.

---

[1] See https://www.icann.org/resources/board-material/resolutions-2017-11-02-en#2.a
[2] See
https://community.icann.org/display/NCAP/NCAP+Documents+and+Correspondence?preview=/79437474/153519703/ncap-study-1-report-19jun20-en.pdf
[3] See https://www.icann.org/resources/board-material/resolutions-2021-03-25-en#2.b

2   Background and Related Work

2.1   Related work on name collision studies

2.2.1 Interisle Study on Name Collisions

On August 2, 2013 ICANN released an ICANN report from Interisle Consulting Group titled, "Name Collision in the DNS: A study of the likelihood and potential consequences of collision between new public gTLD labels and existing private uses of the same strings, version 1.5."[4]

The Interisle Study analyzed data sets from the following sources:

1. the DNS request stream at the root servers that participated in the "Day in the Life of the Internet" (DITL) exercises organized by the DNS Operations, Analysis, and Research Center (DNS-OARC) in 2012 and 2013;[5]
2. the DNS request stream at servers operated by a global DNS resolver organization that contributed to the 2012 DITL exercise; and
3. data concerning the issuance of internal name certificates provided by organizations that operate Certificate (or Certification) Authorities that issue public key digital certificates, many of them members of the Certification Authority/Browser (CA/B) Forum.

With respect to the six strings studied in our current case study. Here are the results from the 2012 Interisle Study.

String	2012 DITL queries (thousands)	2013 DITL queries (thousands)	Query count change 2012-2013	2012 Prefix count[6] (thousands)	2013 Prefix count (thousands)
local		2,501,349			
home	595,024	952,944	60%	1,015	302
lan		362,914			
corp	122,794	144,507	18%	793	185

---

[4] Interisle Consulting Group, "Name Collision in the DNS: A study of the likelihood and potential consequences of collision between new public gTLD labels and existing private uses of the same strings, version 1.5," 2 August 2013. [Online]. Available: https://www.icann.org/en/system/files/files/name-collision-02aug13-en.pdf.

[5] The data captured during the 2012 DITL exercise consisted of full-stream packet captures from the A, C, E, F, H, I, J, K, L, and M root servers. The exercise was conducted during the 3-day period from 17 April to 19 April 2012. The data set amounted to 5.2 TB, comprising 230,000 compressed pcap files which contained a total of 55 billion DNS requests. The data captured during the 2013 DITL exercise consisted of full-stream packet captures from the A, C, D, E, F, H, I, J, K, L, and M root servers. The exercise was conducted during the 3-day period from 28 May to 30 May 2013. The data set amounted to 1.7 TB, comprising 290,000 compressed pcap files which contained a total of 39 billion DNS requests.

[6] For each proposed TLD, each distinct IP address prefix was determined and the number of distinct IP address prefixes that appeared for each proposed TLD was counted. For IPv4 the prefix is a /24; for IPv6 the prefix is a /32.

mail	1,505	2,383	58%	713	526
internal		508,937			

String	2013 DITL Query Count (thousands)	2013 DITL Rank	2012 Recursive Resolver Rank	2012 Recursive Resolver Count (thousands)
local	2,501,349	3		
home	595,024	5		15,308
internal	508,937	8		
lan	362,914	13		
corp	122,794	23		17, 963
mail	1,505	118		35,873

2.2.1 JAS Study on Name Collisions

JAS Global Advisors, using DITL 2012 and 2013 data, performed a study on name collision and provided a framework for risk mitigation.[7] The key recommendations of their reports are:

- The TLDs .corp, .home, and .mail be referred to the Internet Engineering Task Force (IETF) for potential RFC 1918-like protection/treatment.
- ICANN continue efforts to make technical information available in fora frequented by system operators (e.g., network operations groups, system administration-related conferences, etc.) regarding the introduction of new gTLDs and the issues surrounding DNS namespace collisions.
- Emergency response options are limited to situations where there is a reasonable belief that the DNS namespace collision presents a clear and present danger to human life.
- Root-level de-delegation of a production TLD is not considered as an emergency response mechanism under any circumstances.
- ICANN leverage the EBERO mechanisms and functionality to respond to DNS namespace-related issues. ICANN must have the following capabilities on a 24x7x365, emergency basis: 1) Analyze a specific report/incident to confirm a reasonable clear and present danger to human life; 2) Direct the registry on an emergency basis to alter, revert,

---

[7] See https://www.icann.org/en/system/files/files/name-collision-mitigation-final-28oct15-en.pdf

or suspend the problematic registrations as required by the specific situation; 3) Ensure that the registry complies in a timely manner; and 4) Evaluate and monitor the specific situation for additional required actions. Furthermore, we recommend that ICANN develop policies and procedures for emergency transition to an EBERO provider in the event the registry is unable and/or unwilling to comply. We recommend ICANN maintain this capability indefinitely.

- ICANN require new TLD registries to publish the controlled interruption zone immediately upon delegation in the root zone. After the 90-day period, there shall be no further collision-related restrictions on the registry.
- ICANN require registries that have elected the "alternative path to delegation" rather than a wildcard, instead publish appropriate A and SRV resource records for the labels in the ICANN 2LD Block List to the TLD's zone with the 127.0.53.53 address for a period of 90 days. After the 90-day period, there shall be no further collision-related restrictions on the registry.
- ICANN relieve the prohibition on wildcard records during the controlled interruption period.
- ICANN monitor the implementation of controlled interruption by each registry to ensure proper implementation and compliance.
- ICANN work with the IETF to identify a mechanism for IPv6 that provides similar functionality to that available in IPv4's "localhost" reserved prefix.
- ICANN, DNS-OARC, and the root operators explore a medium-latency, aggregated summary feed describing queries reaching the DNS root.
- ICANN, DNS-OARC, and the root operators explore establishment of a single, authoritative, and publicly available archive for historical data related to the root.
- ICANN explore collecting NXDOMAIN entries in DNS query logs from registry operators and contribute them to an independent data repository such as DNS-OARC for further analysis. To limit the potential for commercial gaming or use by malicious parties, we recommend that logs be provided six months in arrears.
- ICANN request that the appropriate bodies further explore issues relating to collisions in existing DNS namespace, the practice of "domain drop catching," and the associated data feeds that may be leveraged by attackers when attempting to exploit collisions.

Some of these recommendations were adopted by ICANN as a base to develop the Name Collision Occurrence Management Framework that has been in use by new gTLDs in the 2012 new gTLD round.

3 Study Methodology

3.1 Data Source

The best known source of historic query data for the root servers is the DITL (Day In The Life) of the Internet collection maintained by DNS-OARC. The Root Server System Advisory Committee (RSSAC) collects and publishes measurement data that define the desired service trends for the root server system. Such data and trends also helps to inform the current NCAP Study 2.

Verisign also has several years of A, J, and old J-root[8] query data. For this initial phase of the NCAP project, the A/J data was used due to its relative longitudinal completeness and availability to the NCAP discussion group. From here on out, this data is simply referred to as A/J root data.

As for the strings to be studied, the NCAP Revised Proposal asked for case studies of CORP, MAIL, HOME, and non-delegated strings that receive more than 100 million queries per day at the root. Using this threshold and DNS query data from A and J root servers results in six strings:.CORP, .HOME, .INTERNAL, .LAN, .LOCAL, and .MAIL.

The reliance on A/J root data is not without its limitations. While it covers a significant portion of anycast root instances on the Internet, the data set is an incomplete one. Resolver selection algorithms, BGP peering preferences to root server address prefixes, and even network policies of ISPs or governments means a significant proportion of root server queries will never be seen by A/J root. Nonetheless, for the purposes of this initial phase, A/J root query data is believed to be sufficiently representative of global root server query trends. The NCAP is performing *sensitivity analysis* to understand the limitations of root data sampling.

Current A/J root data is used for this NCAP Study 2 analysis phase right up until the time it was presented, which spans the first few months of 2021. Neither the raw A/J data nor the specific tools used to conduct the analysis is publicly available, with the exception of two days each year through the DITL project. While the sources and methods are not presented in great detail, for the purposes of this summary we outline the general methods applied to the data.

Where applicable, the analysis included data going back two, three, and even four or more years where trend lines were useful to convey important changes over time.

3.2   Data Analysis Methodology

Source IP addresses, originating autonomous systems (ASes), query name strings, and query types were the most common attributes of the data examined, but other characteristics such as QNAME minimization were also considered. The analysis largely focused on aggregate volume and query distribution of an attribute under study, such as a time series plot of aggregate query volume, or the percentage of queries for a particular label as seen by different originating ASNs. This led the analysis to focus primarily on the historical trends for an examined set of characteristics or the distribution of an examined set of query characteristics (e.g. the second-level label for a given TLD). While most historical trends were presented as time series plots, bar graphs, geolocation maps, venn diagrams, and even clustering algorithms were used to present varying views of the data. The sheer volume of the data and the limited resources available for analysis during this phase made deeper analysis prohibitive. For each of the labels under study, the following analysis methods were used across each data set:

---

[8] The original J-root IP address was 198.41.0.10. This root instance was renumbered in 2002 to 192.58.128.30 in order to facilitate the use of anycast. The old address still receives DNS query traffic, which Verisign continues to service and monitor. Also see https://j.root-servers.org/

- Query volume (daily)
- QTYPE distribution
- Unique daily query source IPv4 and IP6 addresses
- Geographic distribution
- ASN distribution
- Label (analysis) distribution
- SLD overlap between roots
- ASN overlap between roots

Some select analysis using clustering algorithms was also presented after each label was analyzed. This clustering was provided to help demonstrate the potential for more advanced analysis techniques to identify and understand the ramifications of collisions. A proposed set of data attributes from which to study collisions for future studies is summarized in the table below, many of which were utilized in this phase of the project.

Data Attributes When Evaluating Collision Strings

Traffic Properties:	Qname and Labels:	Other Attributes:
- Network diversity - Number of unique ASNs, /24s, etc. - Distribution of traffic (e.g. heavily weighted in a few ASNs) - Geographical diversity - Qtype distribution - Query volume - Longitudinal trends	- Distinct SLDs - Distribution of traffic over SLDs - Amount of "noise" (e.g. Chromium) - SLDs appear to be delegated TLDs - First label features - DNS-SD - Common protocols - Qname Minimization effect	- The string's context - OSINT of string being used - Data sensitivity and catchment of data collector

The analysis concluded by examining an array of potential vulnerabilities that may arise due to collisions. This includes a number of well-known, specific cases such as those with WPAD[9], ISATAP, ZEROCONF[10], and others.

3.3   Limitations

The majority of the NCAP analysis conducted focused on the aggregate view of DNS traffic. However there may be regional differences that may be more or less important when it comes to collisions and their mitigation. A breakdown by root instance for example may help uncover

---

[9] https://ieeexplore.ieee.org/abstract/document/7546529
[10] https://dl.acm.org/doi/10.1145/3133956.3134084

interesting topological patterns. These concerns will be further examined in the data sensitivity analysis effort conducted by the NCAP discussion group.

In addition to the A and J root data sets, historical data from the DNS-OARC Day-In-The-Life (DITL) of the Internet data sets, that include query traffic over the course of 48 hours from many of the root servers and some other DNS systems going back many years, was included. Examining the difference in queries for names in newer TLDs may provide some clues and insight into what future delegations may expect to see.

There are however a few questions that were difficult or impossible to answer. For example, the use of QNAME minimization limits the ability for passive measurement at the root servers to uncover the valuable insights from label analysis. Recent research has shown that 6% of resolvers and 40% of queries associated with authoritative servers for a TLD exhibited qname minimization.[11] Practically all studies conducted thus far have also lacked an aggregate view of alternative transports such as DNS over TLS and DNS over HTTPS. While those technologies are still relatively new, they are beginning to see significant deployment that appears to be on the rise and likely to grow.[12] The NCAP analysis lacks what effect these technologies currently and are projected to have in the coming years. Additionally, this analysis is limited to data available at the root server system and lacks any insights into what recursive resolvers observe in terms of name collision queries.

4 Results

4.0 Overall DNS Traffic Evolution

Over the past few years the analysis highlighted a number of DNS root server traffic trends, some well known, others less so. DNS traffic volume has shown a steady year-over-year increase up until the end of 2020. Comparable to global Internet traffic trends seen in most networks, traffic volume doubles roughly every few years. However, at the end of 2020, two phenomena seem to have disrupted this once expected growth trend.

DNS traffic volume is heavily influenced by modern web browsers. In December of 2020 the Chromium browser altered their method used to detect the existence of captive portals. This method utilized random query labels and accounted for a significant amount of DNS traffic that made its way to the root servers to test if NXDomain interception was happening locally. Upon the root server operators' request, a code change to chromium was made. Once this change was made, the reduction in query traffic was noticeable.[13]

Another feature being deployed in many modern web browsers is DNS over HTTPS (DoH). When browsers utilize this feature they often bypass the typical resolution path, which may include a query to the root server system. While DoH does not eliminate the need to query the root server system, it may reduce the volume of traffic to the root servers. This is because DoH

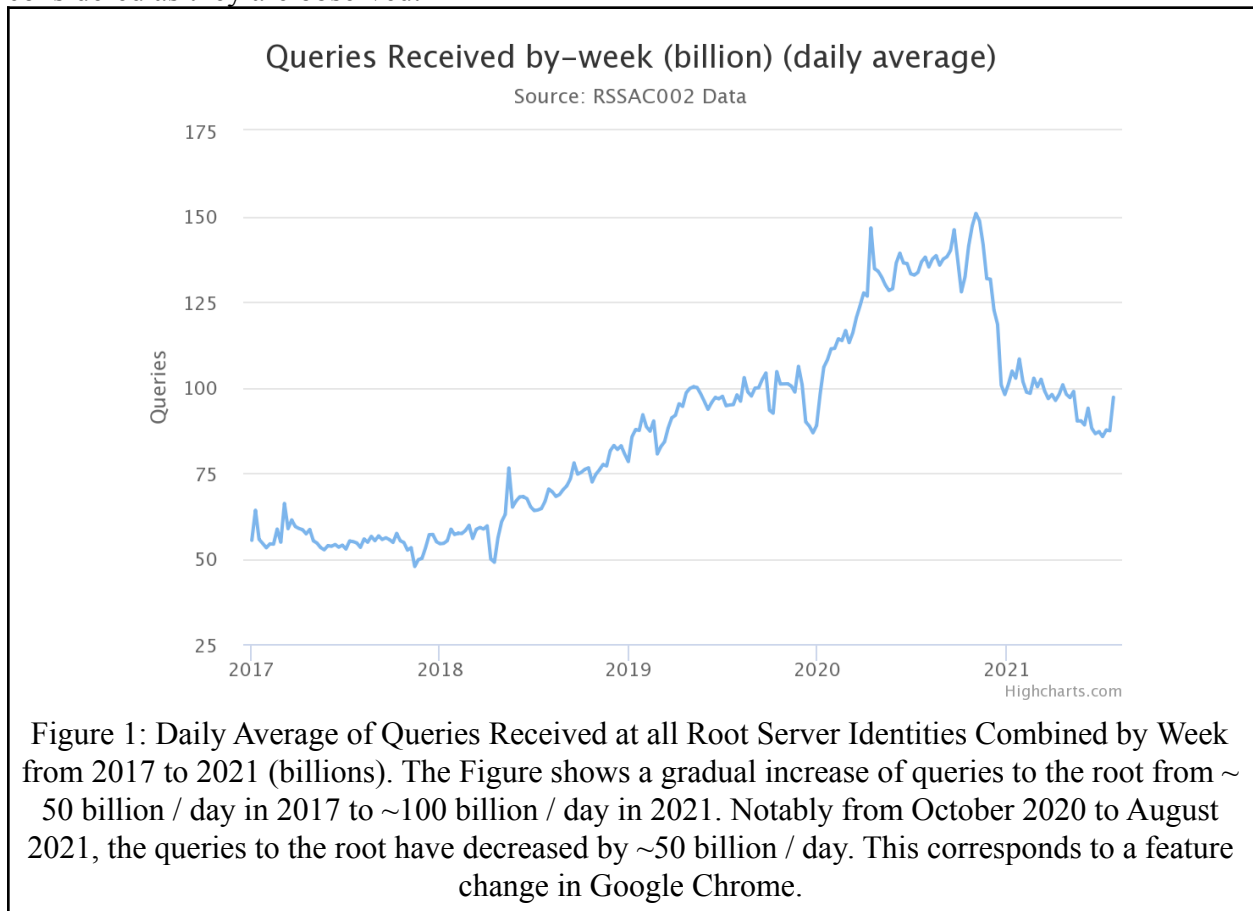[11] See https://link.springer.com/content/pdf/10.1007%2F978-3-030-15986-3_10.pdf.
[12] See https://dl.acm.org/doi/abs/10.1145/3359989.3365435.
[13] https://blog.verisign.com/domain-names/chromiums-reduction-of-root-dns-traffic/

servers, which are often global resolvers, can concentrate the DNS caching or run local instances of root, both of which reduce the frequency to query the root server system. At this time of this report, though, the overall effect of DoH--and to a greater extent, these new features appearing in web browsers--isn't yet entirely clear.

A better known trend is the increasing use of QNAME minimization. Most modern resolver implementations now support this feature, and it is often enabled by default.[14] This development limits the ability of top-level server operators to only see a subset of the query name label string. While this feature enhances client query privacy, it prohibits diagnostic query analysis from those servers that return referrals.

Other important DNS traffic evolutionary trends are occurring (e.g. DNSSEC deployment, aggressive negative caching, NXDomain cut, and local root servers), but they are not considered as fundamental to the issues of this name collision study. However, what is important is that DNS traffic is evolving, and it is reasonable to assume that as the Internet continues to evolve, so will DNS traffic. The effect of these evolutionary changes on name collisions will need to be considered as they are observed.

Figure 1: Daily Average of Queries Received at all Root Server Identities Combined by Week from 2017 to 2021 (billions). The Figure shows a gradual increase of queries to the root from ~ 50 billion / day in 2017 to ~100 billion / day in 2021. Notably from October 2020 to August 2021, the queries to the root have decreased by ~50 billion / day. This corresponds to a feature change in Google Chrome.

---

[14] At the time of writing, the latest versions of BIND, unbound, and knot resolvers enable qname minimization by default. See http://ftp.isc.org/isc/bind9/9.17.16/doc/arm/html/reference.html, https://nlnetlabs.nl/documentation/unbound/unbound.conf/, and https://knot-resolver.readthedocs.io/_/downloads/en/latest/pdf/

Figure 2: Response Code Distribution for the Queries Received at the Root from 2018.1.1 - 2021.4.12. 71.6% are for non-existent names where a NXDomain Response is returned, 28.3% are for legitimate domains where a NoError is returned.
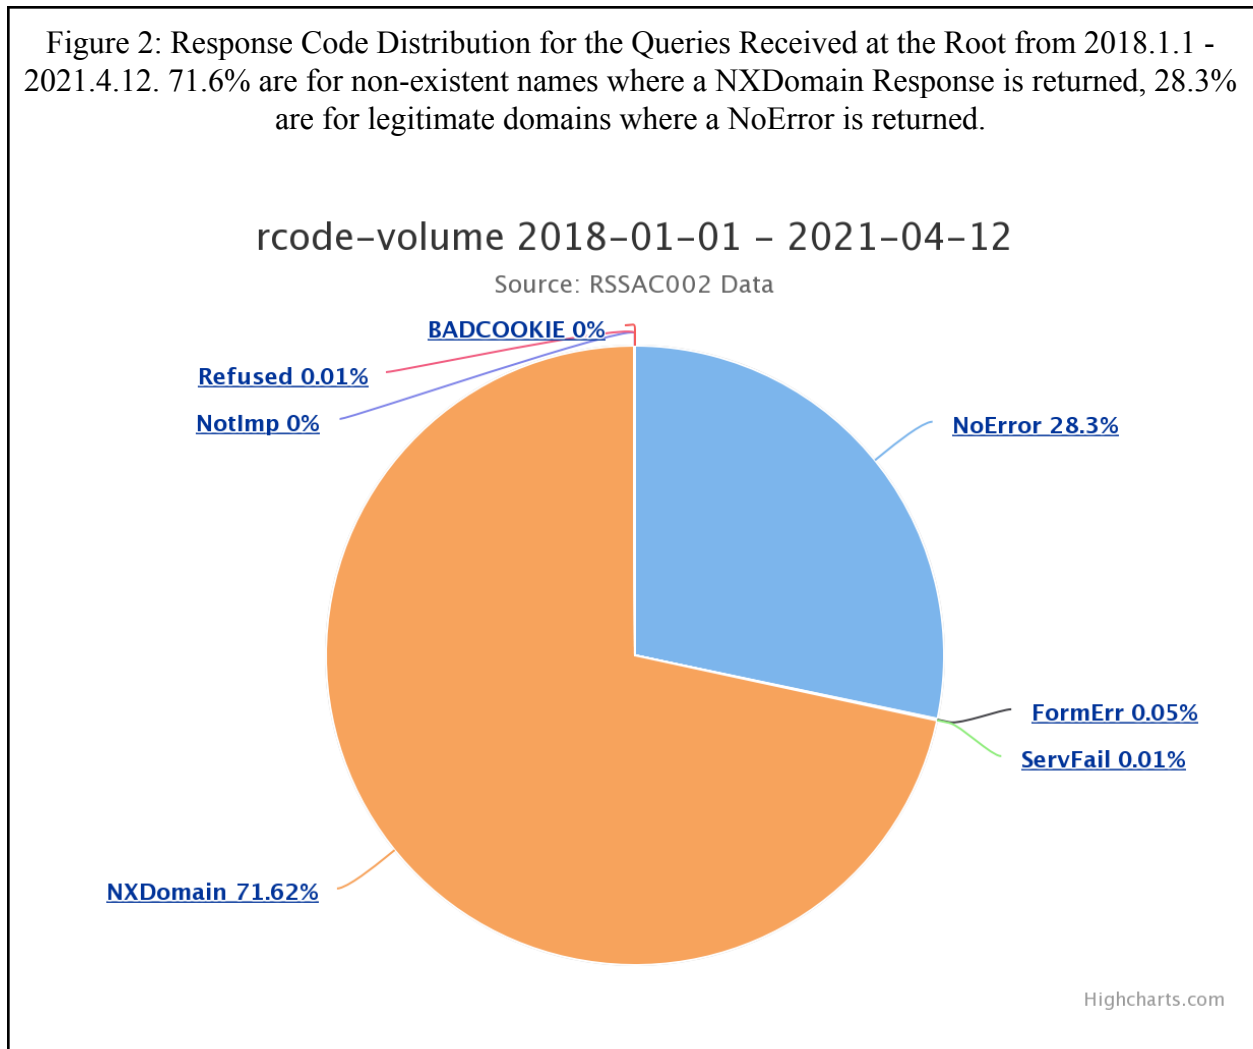
Figure 2 showed that for all the traffic to the DNS root server system, 71.6% are queries for non-existent domains, that's roughly 71 billions queries / day as of March 2021. This remains a substantial percentage throughout the measurement period. This data is significant not because the root server system cannot handle the load. On the contrary, the root server system has been very resilient and is capable of handling such load and many more. The data is significant as it reflects the amount of queries of non-existent domains, which could be a significant source of name collisions should some of these TLDs be delegated. In our subsequent analysis, we focus on understanding the name collisions impact to end users and system, and how to better measure them. We will not consider its impact to DNS root name servers further, as the root server system has been resilient and has no issues handling the increased query volume of non-delegated TLD strings.

Figure 3: Unique IPv4 and IPv6 Sources Querying the Root Servers Per Day (million). The number of unique IPv4 and IPv6 sources remain stable until 2020 in which the number of unique IP addresses querying the root grows significantly from ~50 million per day to ~75 million per day.
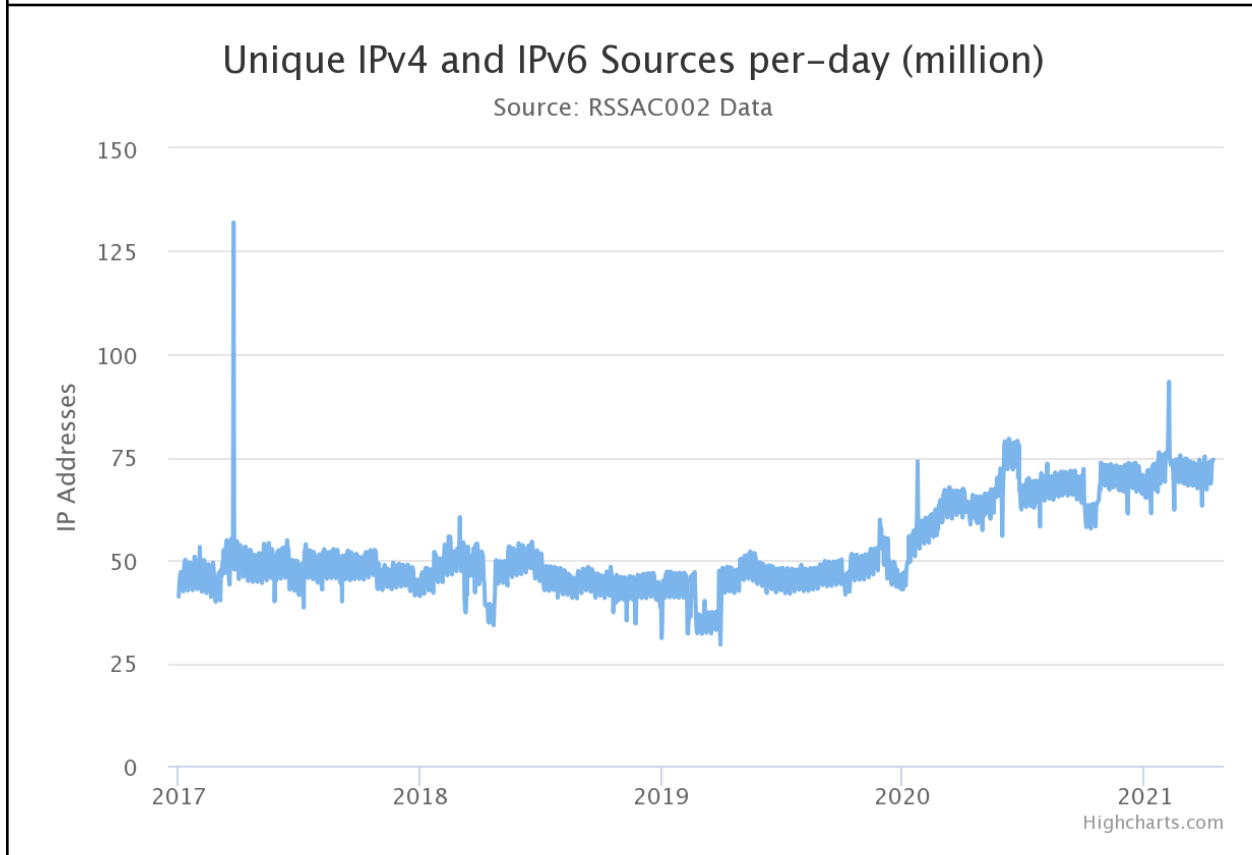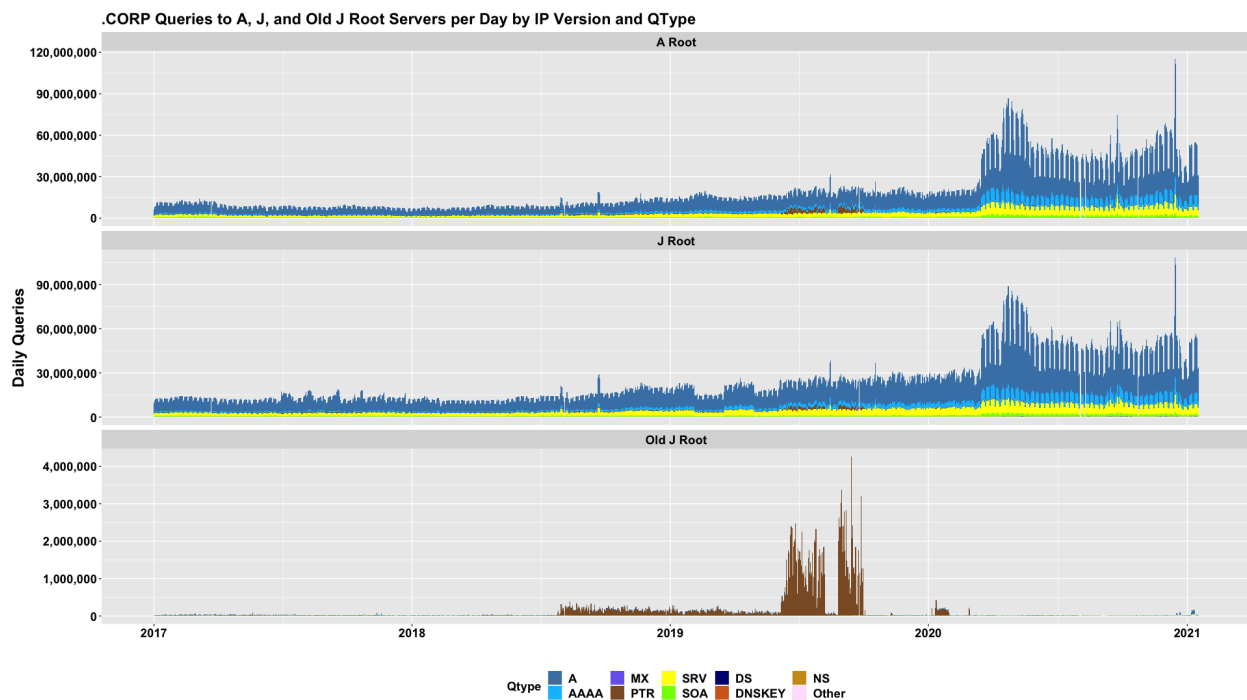
4.1 **Traffic Analysis**

4.1.1 Query Volume Analysis

By sheer volume, .HOME accounts for one or two orders of magnitude more of DNS query volume than does either .CORP or .MAIL. At the end of 2020 A/J root data shows .HOME volume approaching 400 million queries/day per root server. Whereas .CORP has reached roughly 60 million queries/day and .MAIL approximately 2 million queries/day. To put this in perspective, the A/J root servers receive tens of billion queries per day per letter.[15]

Both .CORP and .HOME volume rose significantly in March 2020, coinciding with the COVID-19 outbreak. The reason for this change in traffic volume is largely believed to be associated with a large global shift in working at home, where queries for these names were previously captured by employer-run resolvers.[16]

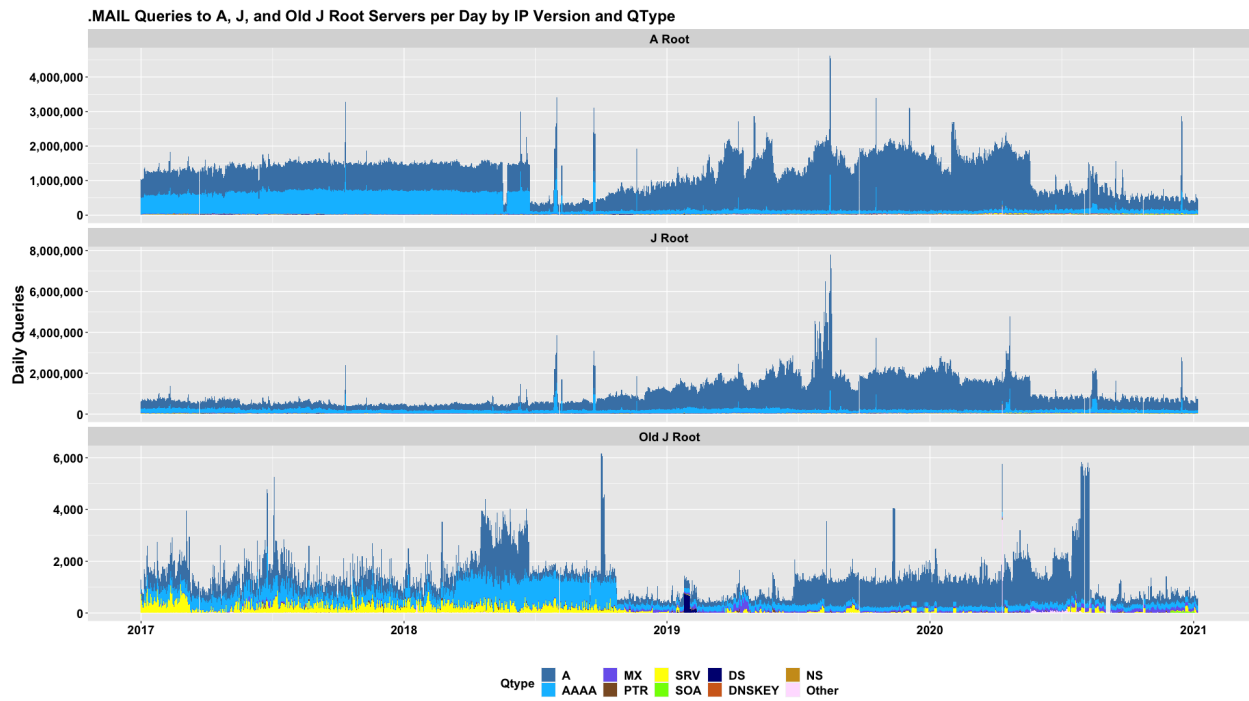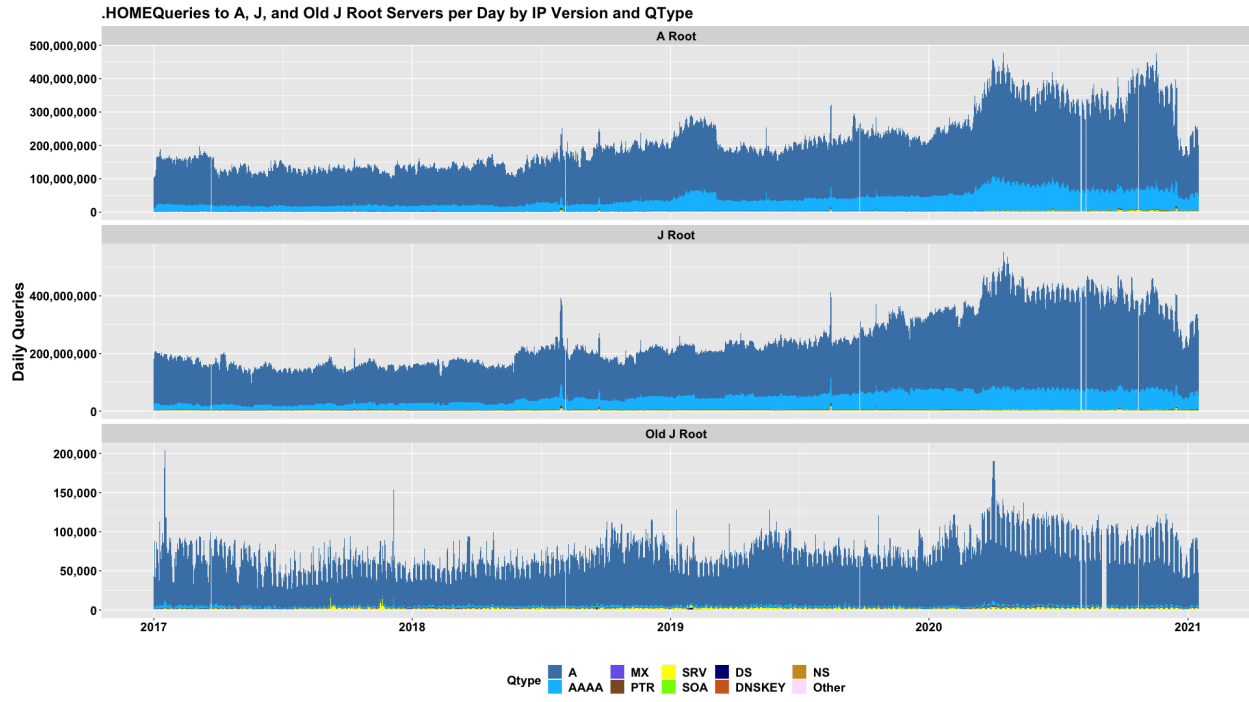### 4.1.2 Query Type Distribution Analysis

Figure below shows queries to A, J and old J Root server per day by IP version and query type for the six strings. The query type (QTYPE) distribution by label varies, sometimes significantly. For example, .HOME is dominated by A and AAAA queries, but .MAIL sees more diversity and for a largely different set of QYTPEs with MX, SRV, and DNSKEY amongst the most common.

.CORP Queries to A, J, and Old J Root Servers per Day by IP Version and QType

---

[15] https://blog.verisign.com/domain-names/chromiums-reduction-of-root-dns-traffic/
[16] See https://www.icann.org/en/system/files/files/octo-008-15apr20-en.pdf.

Case Study of Collision Strings

.HOMEQueries to A, J, and Old J Root Servers per Day by IP Version and QType

.MAIL Queries to A, J, and Old J Root Servers per Day by IP Version and QType

.INTERNAL Queries to A, J, and Old J Root Servers per Day by IP Version and QType

.LAN Queries to A, J, and Old J Root Servers per Day by IP Version and QType

.LOCAL Queries to A, J, and Old J Root Servers per Day by IP Version and QType

4.1.3 Unique Daily Source IP Addresses

The number of distinct source IP addresses issuing queries grew significantly in 2020 across all the labels examined. While there was a noticeable drop at the end of 2020, most likely due to the change in behavior by Chromium, the trend appears to continue to be up and to the right. The analysis did no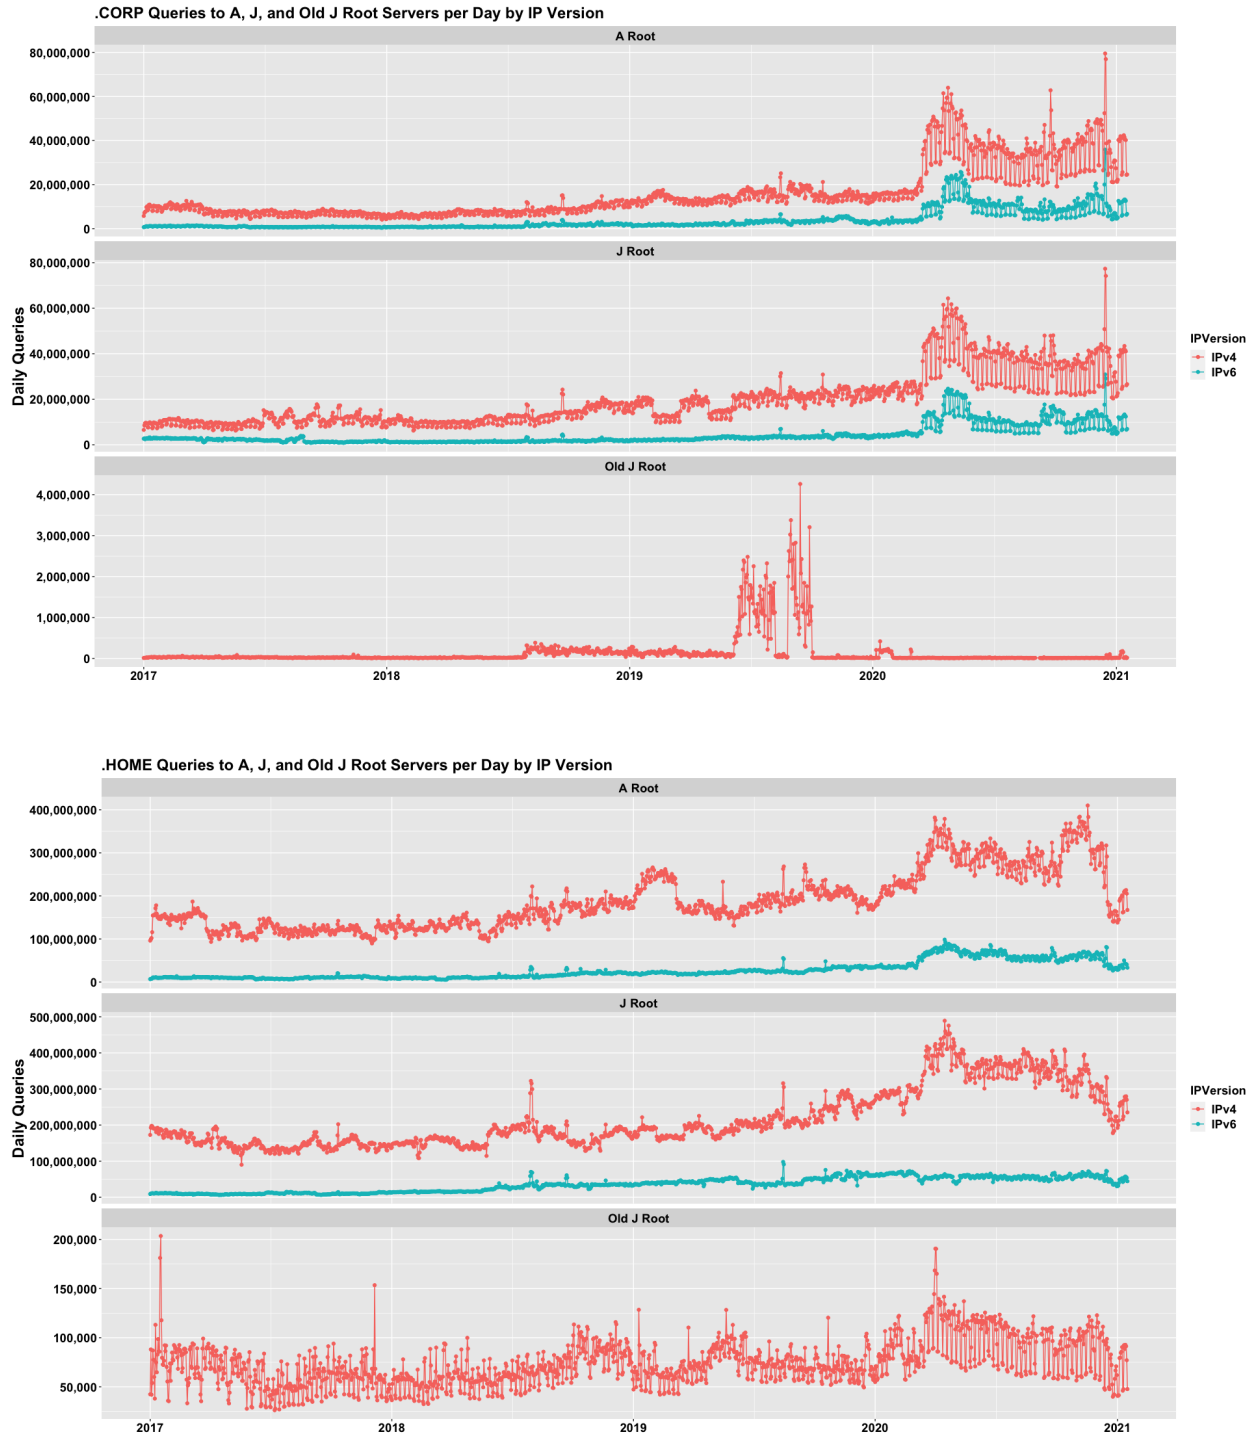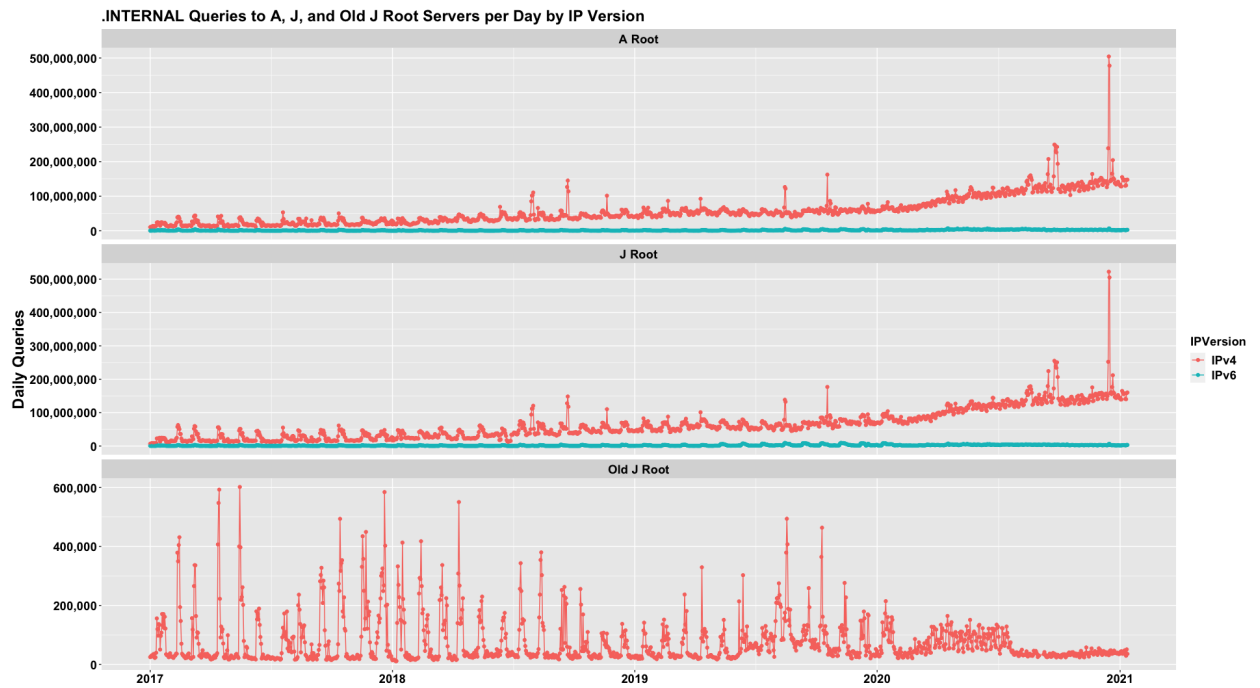t provide enough detail to explain this phenomenon, but there are theories. One has been the ongoing shift in working from home around the globe. Data in the coming years may help shed light on this hypothesis. Another possible explanation is the growing use of IPv6 and address privacy, but this seems less likely since resolvers that talk to the root would seem to be less likely to use address privacy than originating clients themselves. The plots below depict the common shape of unique daily source IP address volume for each of the six strings studied.

.CORP Queries to A, J, and Old J Root Servers per Day by IP Version

.HOME Queries to A, J, and Old J Root Servers per Day by IP Version

.MAIL Queries to A, J, and Old J Root Servers per Day by IP Version

.INTERNAL Queries to A, J, and Old J Root Servers per Day by IP Version

Case Study of Collision Strings

.LAN Queries to A, J, and Old J Root Servers per Day by IP Version

.LOCAL Queries to A, J, and Old J Root Servers per Day by IP Version
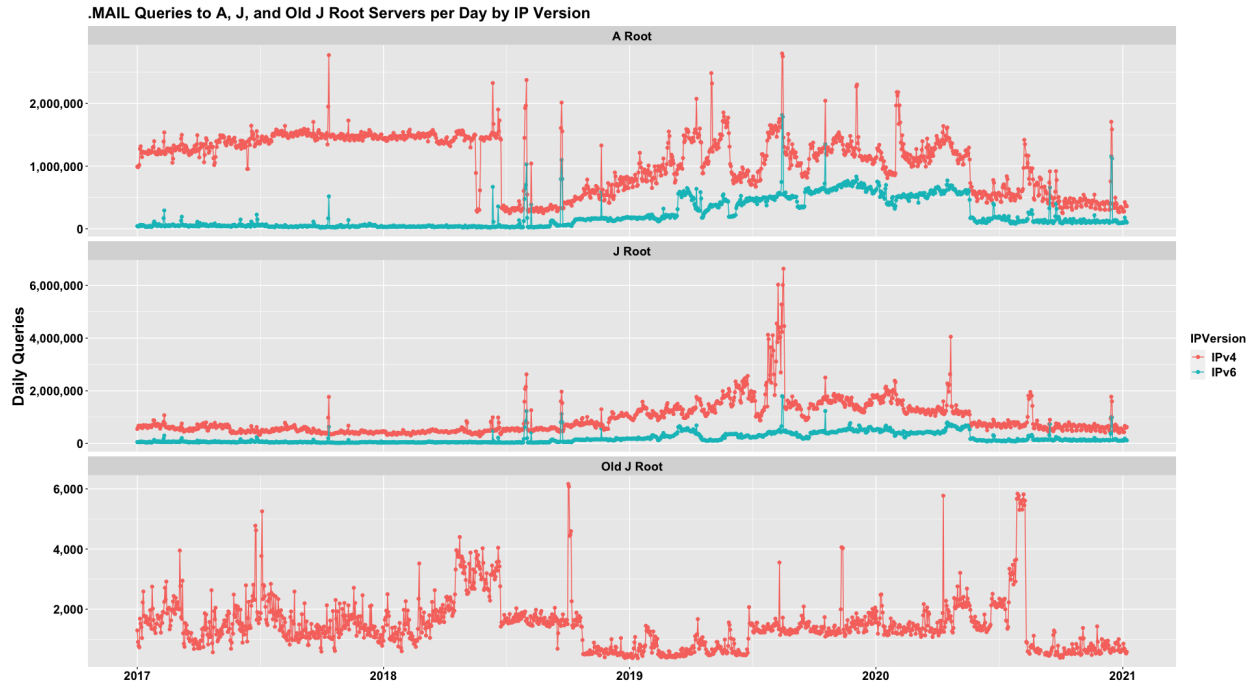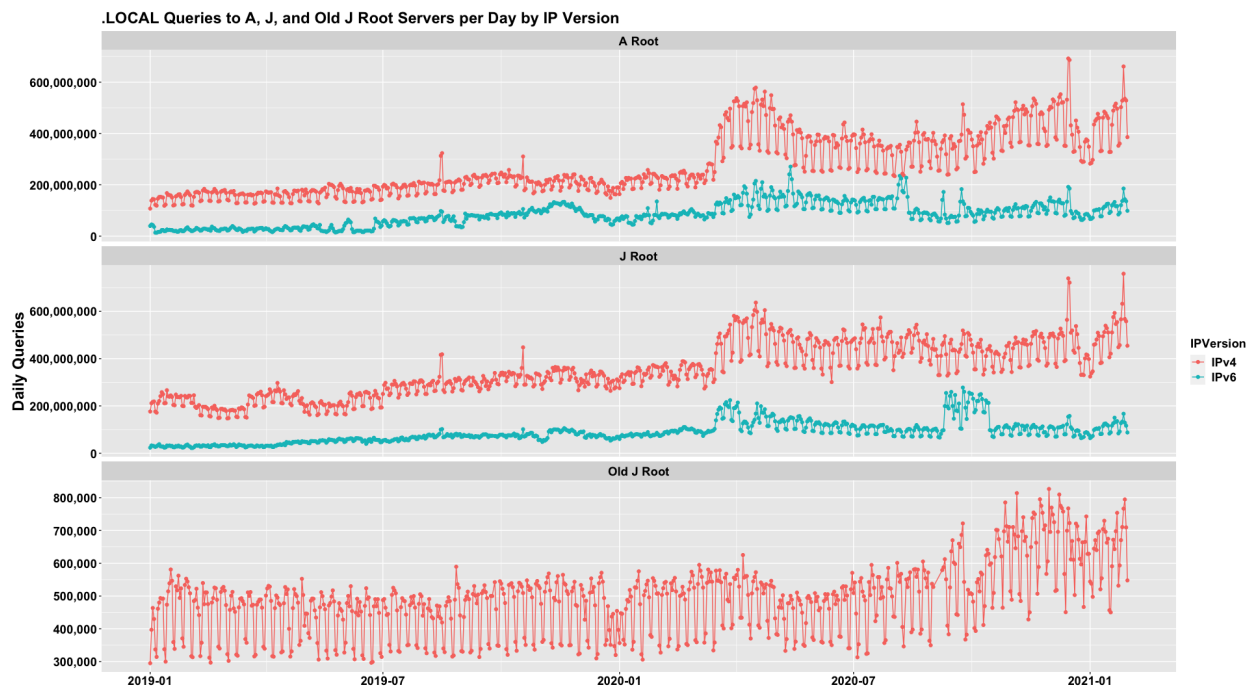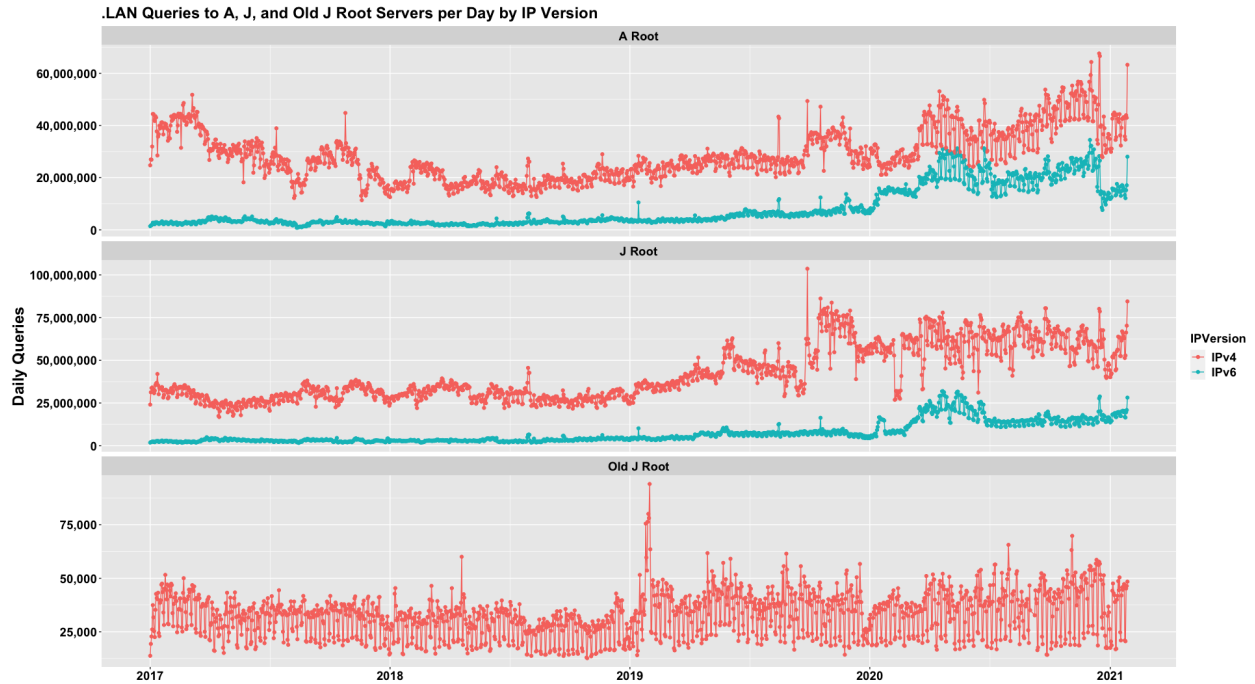
Number of Distinct IPs per Day at A and J Root Servers for .CORP Queries
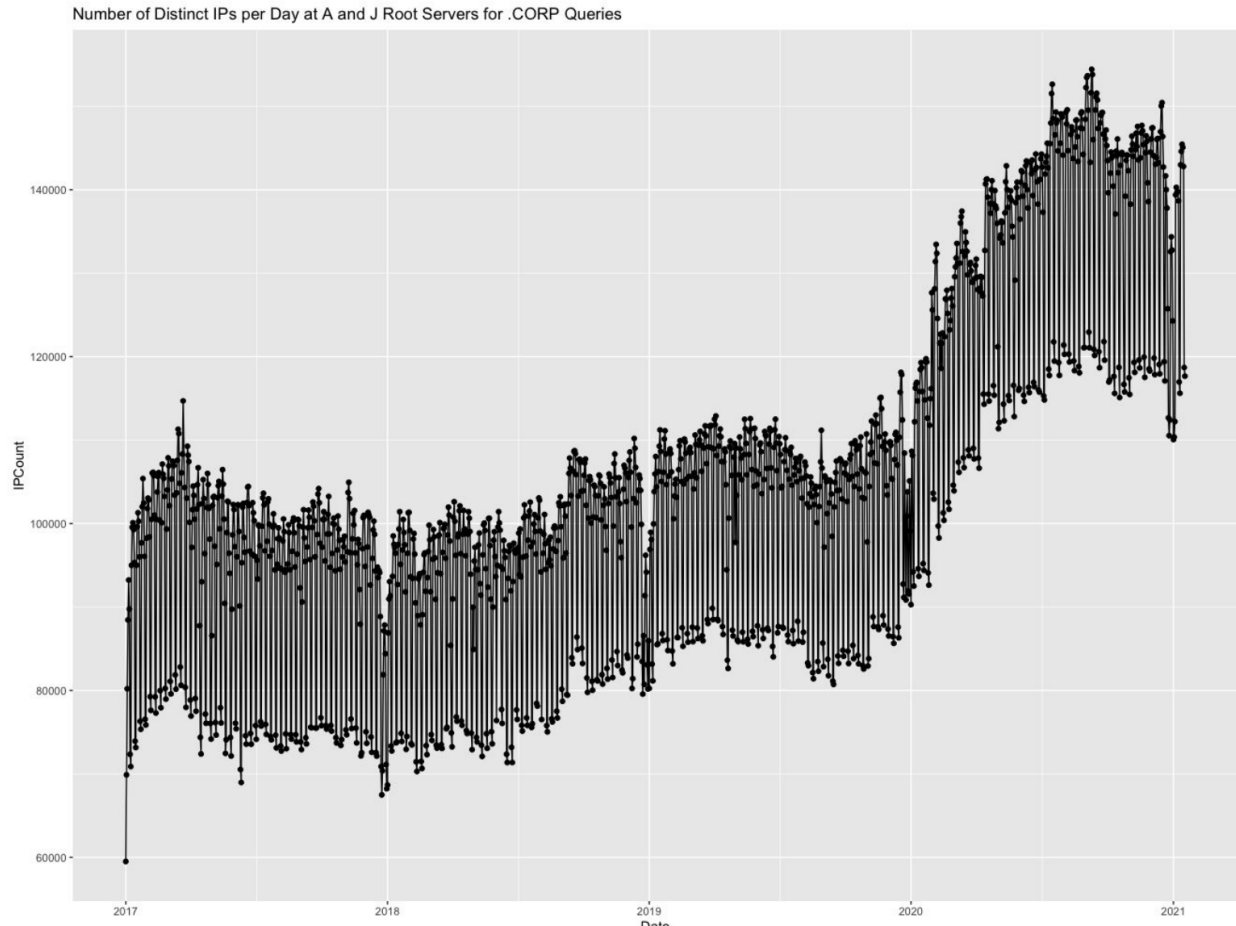
4.1.4 Geographical Distribution

Each of the three primary TLD labels under study exhibit a similar geographic distribution of queries to A/J servers. The U.S. dominates in all three cases with .CORP, .HOME, and .MAIL. The percentage of traffic from the U.S. for .HOME and .MAIL queries is roughly the same at approximately 30%. For .CORP however, the percentage rises to 50%.

4.1.5 ASN Distribution

A relatively small number of origin ASNs account for the vast majority of query traffic for .CORP, .HOME, and .MAIL. In all cases roughly 200 ASNs make up nearly 90% of the volume. This suggests that remediation efforts focused on a subset of the highest ISP contributors by volume could significantly reduce damage and risk of colliding names. See the plot depicting the ASN distribution for .MAIL below.

One interesting finding is the ASN distribution for .INTERNAL. Nearly 85% of the DNS queries observed originated from two ASNs, both of which are owned by the same large cloud provider. This distribution significantly differs from the other strings under study, in such that a large

percentage of traffic is directly associated with a single entity.
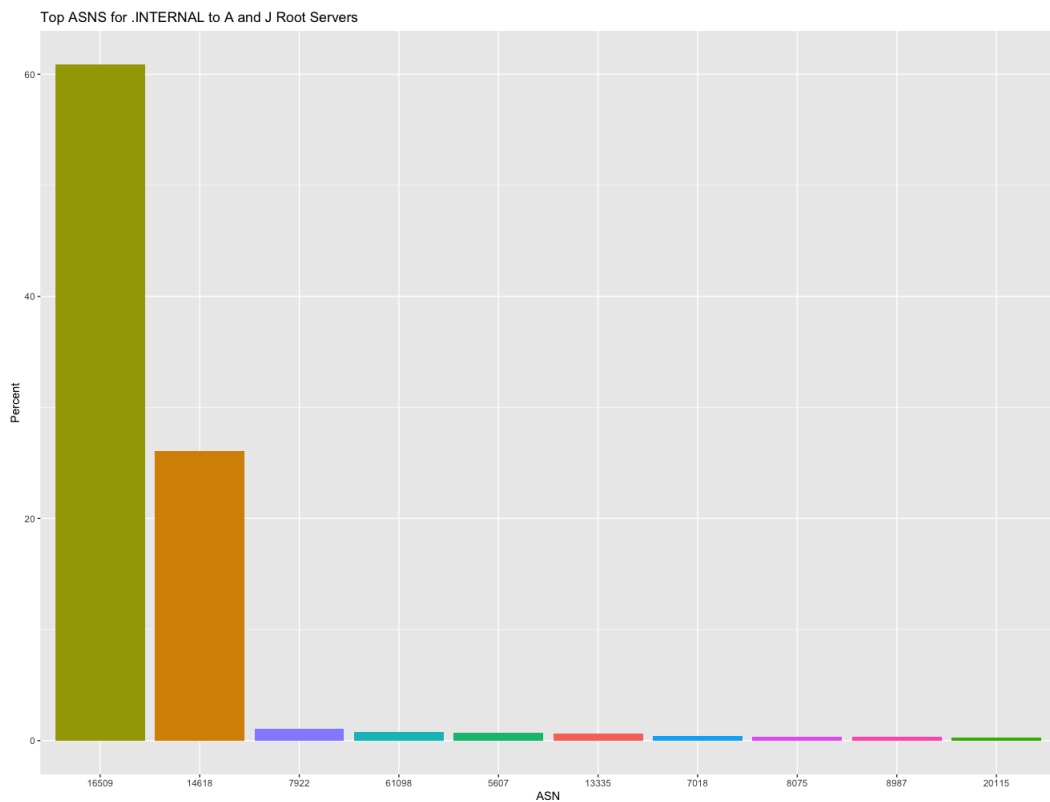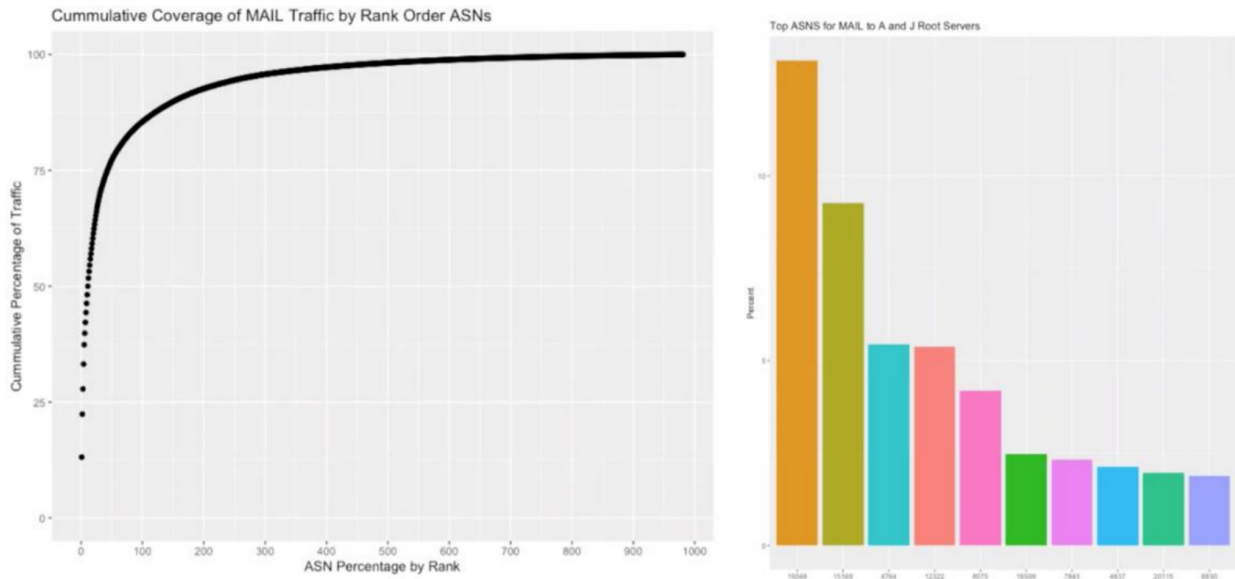
4.2 Label Analysis

4.2.1 SLD Analysis

The label analysis proves very insightful across all three names of interest. In the case of .CORP we can identify the most prevalent second-level names by volume with ease. The top four labels

account for greater than 20% of all query volume, whereas all other labels account for less than 1.6%. A similar trend can be seen with .HOME. In both cases, due to the long tail of unique names, this evidence may complicate the remediation process even if the number of ISPs originating these queries are of limited number. .MAIL shows a different pattern however, especially when examining the third label of the name. Here a preponderance of the names are associated with a smaller set of applications or services. For example, WPAD makes up almost 20% of all third label names. Likewise, roughly the same percentage of volume is associated with Windows email provisioning systems.

Label analysis provides a unique observational context into the underlying systems, networks, and protocols inducing leakage of DNS queries to the global DNS ecosystem. Understanding the diversity of labels can help provide a sense of how broadly disseminated the leakage is throughout the DNS. Labels can also provide clear indicators of who and or where the leaking queries originate from. Finally, specific labels can be identified and associated with known protocols that can cause harm in name collision scenarios.

It is worth reiterating yet again, that with the deployment and standardization of the DNS privacy enhancing Qname Minimization technique, label analysis at the upper-level of the DNS hierarchy will become completely impaired. So while label analysis plays an important role in this case study, future measurements may not be able to quantify label attributes and contexts.

4.2.2 Labels Associated with Known Protocols that Could Cause Harm

A number of network protocols make use of top-level domains, including those under study, that if delegated may cause harm. WPAD and ISATAP are two that are well known. The former is discovering the web proxy auto configuration mechanism designated by a network; it is supported primarily by Microsoft browsers. WPAD prepends a label of the same name to the client's default domain. If an answer is not found, the client will "devolve"[17] issuing WPAD prepended labels up the namespace hierarchy until an answer is received or the process fails at a query for the WPAD top-level domain. ISATAP is a label used for a legacy IPv6 transition mechanism of the same name. ISATAP clients will prepend the label to their default domain, issue an A query, and if an answer is returned attempt to use the A answer as their IPv6 over IPv4 gateway. For both WPAD and ISATAP the harm may come in failed, or worse, hijacked network traffic.

In addition to the WPAD and ISATAP services, the name collision problem affects a broader set of internal network services. The following table summarizes a number of known protocols that can cause harm via various attack techniques that can be exploited when under name collision conditions.[18] These DNS-based service discovery protocols have potential security implications including MitM attacks, call spoofing, and information leakage.

---

[17] https://docs.microsoft.com/en-us/security-updates/SecurityAdvisories/2009/971888?redirectedfrom=MSDN

[18] Chen, Qi Alfred, et al. "Client-side name collision vulnerability in the new gtld era: A systematic study." *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 2017.

Exposed service functionality	Exposed service name	Potential security implications	Exposed service functionality	Exposed service name	Potential security implications
Proxy/tunnel config.	wpad① (N), isatap② (N), proxy② (N)	MitM attack	Remote access to computers/file systems	afs3-vlserver④, adisk④, smb④, afpovertcp④, ftp④, sftp-ssh④, rfb④, webdav⑤, odisk⑤, eppc⑤, telnet⑤	Phishing attack, info. leakage
Time config.	ntp③	Time shifting attack			
Software activation	vlmcs② (N)	DoS	System management	kpasswd② (N), airport③, servermgr⑤	System config. info leakage
Directory service (help a client locate a server of the requested service)	ns*① (N), alt*① (N), lb① (N), db① (N), dns-sd①, dr① (N), tracker② (N), dns-llq⑤, dns-update⑤	Server spoofing, service info. leakage	Mail	autodiscover① (N), outlook① (N), mail*① (N), pop3②, smtp②	Email spoofing, phishing
Web service	www*① (N), api① (N), static① (N), cf① (N), share① (N), http②, https③	Web-based phishing attack, malicious script execution	VoIP	sipinternaltls① (N), sip① sipinternal① (N), sipexternal① (N), sips③	Call spoofing, phishing
Server config. retrieval	stun④	Config. info. spoofing			
Multimedia file access	ptp③, dpap④	Phishing attack	Messaging	xmpp-server③, xmpp-client③	Msg. spoofing, phishing
Authentication service	kerberos①	DoS	Printer	printer③, pdl-datastream③, riousbprint③, ipp③	Internal/personal document leakage
Coding library retrieval	rubygems⑤	Malicious code injection			
Database service (organization data, calendar, contacts, etc.)	gc① (N), ldap① (N), carddav④, ldaps④, caldav④, caldavs④, carddavs④	Phishing attack, organization data leakage	Scanner/camera	scanner③, ica-networking⑤	Phishing attack
			Distributed computing	xgrid④	Malicious code execution
			System monitoring	syslog⑤	Organization info. leakage

4.3   Sensitivity Analysis

The NCAP is performing *sensitivity analysis* to understand the limitations of root data sampling and will hopefully elucidate these initial measurements provided from A and J to the broader root server system as a whole.

4.3.1 SLD Overlap Analysis

Each root server sees a unique set of queries. When comparing the intersection of query names or query sources each root server sees, it might seem likely that over time the union between the two might be relatively close to the set of either one. However, there is significant diversity in what a root sees. For one-time unique query names (e.g., random strings generated by Chromium) this may be perfectly logical and relatively irrelevant. The .CORP SLDs seen at both A and J (approximately 16 thousand) is almost equal to those seen at A-root alone, but J-root sees over 30,000 .CORP SLDs that A-root does not see.

4.3.2 Catchment Overlap Analysis

Originating ASN overlap may be slightly more interesting than SLD overlap. If the overlap is weak, this suggests a very biased view of the world, perhaps deriving from network policies and BGP peering relationships. Across all names studied, while A and J saw much in common, there was a non-negligible amount of uniqueness to each view. For example, A and J each saw queries from the same 5717 originating ASNs, but J saw 2477 ASNs that A didn't see and A saw 901 that didn't see.

5   Analysis and Discussion

5.1   Critical Diagnostic Measurements

The NCAP Study 2 uncovered a number of properties that help determine the scope, impact, and potential harm of name collisions. We term these *Critical Diagnostic Measurements (CDM)* to underscore their value and importance in name collision analysis and risk assessment process. No one measurement alone is generally going to provide sufficient quantitative or qualitative indications to thoroughly assess the name collision risks expressed by a string. For example,

query volume--one of the four major classes of measurements--is an important factor, but a single source that could be easily mitigated with a simple configuration may be responsible for high query of a name. Conversely, if not only query volume was high, but query origin diversity (i.e., from many networks and many systems) and query type diversity were also extremely high, this would suggest collision impact may be greater. This is because the expectation of negative responses is high, and the mitigation across multiple services, networks, and users is increasingly complex to perform. We briefly discuss the critical diagnostic measurements that ought to factor into name collisions risk assessments. The original "Data Attributes When Evaluating Collision Strings" table is reproduced below. This is followed by our revised and updated classification of these attributes into a grouping of Critical Diagnostic Measurement in the sections to follow.

      a. Query Volume
          i. DNS query count
      b. Query Origin Diversity
          i. IP distribution
          ii. Network diversity: ASN distribution
      c. Query Type distribution
      d. Label Diversity
      e. Other characteristics
          i. OSINT of string being used

Data Attributes When Evaluating Collision Strings

Traffic Properties:
- Network diversity
  - Number of unique ASNs, /24s, etc.
  - Distribution of traffic (e.g. heavily weighted in a few ASNs)
- Geographical diversity
- Qtype distribution
- Query volume
- Longitudinal trends

Qname and Labels:
- Distinct SLDs
  - Distribution of traffic over SLDs
- Amount of "noise" (e.g. Chromium)
- SLDs appear to be delegated TLDs
- First label features
  - DNS-SD
  - Common protocols
- Qname Minimization effect

Other Attributes:
- The string's context
- OSINT of string being used
- Data sensitivity and catchment of data collector

5.2   CDM: Query Volume

Our analysis of the query volume shows that the strings under study receive substantial amounts. By sheer volume, .HOME accounts for one or two orders of magnitude of DNS query volume more than does either .CORP or .MAIL. At the end of 2020 A/J root data shows .HOME volume approaching 400 million queries/day per root server. Whereas .CORP has reached roughly 60 million queries/day and .MAIL approximately 2 million queries/day.

Query volume also continues to grow year-over-year, and shows no signs of abating. A more intensive and thorough analysis would include other root server vantage points to minimize potential bias in the A and J catchments. Additional measurement from large recursive resolvers would also help elucidate any behaviors masked by negative caching and the population of stub resolvers.

The volume of query traffic for the undelegated names under study is substantial, particularly since this view comes from not only a small subset of root server systems, but also misses all the query traffic aggregated behind resolvers that are never seen in this study.

When interpreting the query volume metric, it is important to distinguish between root server system load concerns and the CDM of query volume and source diversity load concerns related to name collisions. The root server system has been resilient and has no issues handling the increased query volume of non-delegated TLD strings.

5.3   CDM: Query Origin Diversity

Not only is query volume growing over time, but the number of unique query source IP addresses (resolvers) is increasing as well. The current strings exhibit a diversity of origins, not only from an IP address perspective, but also with ASNs, application, and global resolvers. While some growth seems congruent with overall query volume growth, in most cases the rise in distinct sources was significantly more than expected. Regardless of the nature of the change, distinct IP addresses are originating from more networks and in more places around the globe.

This finding highlights the challenge associated with mitigation since diversity complicates mitigation coordination across an increasing number of parties (i.e. networks, vendors, applications, and users). Every additional distinct operator, region, and configuration contributing collision traffic may require their own unique mitigation response. In the best case, distinct sources share a common deployment that can be centrally altered, such as in the Chromium case. However, there are undoubtedly different policies, software implementations, default configurations, and operator skill sets widely dispersed around the globe. As the collision traffic diversity grows, particularly as a result of different operational practices, so too will the mitigation response resources.

A noteworthy finding that speaks to the importance of the origin diversity CDM is highlighted in .INTERNAL. While query volume for that string was voluminous, the diversity of the IP

addresses were mainly concentrated within two ASNs. This highlights the importance of understanding origin diversity as it can have direct implications to mitigating those leaking queries.

5.4   CDM: Query Type Diversity

Another unique attribute to consider when evaluating name collisions is the type of query (i.e. resource record type) being requested. The type of query is often associated with specific types of applications or used by other protocols. Therefore, an analysis of the query type diversity can reflect the potential number of different uses of the leaking string. This query type diversity CDM is highlighted in the finding of .MAIL, in which the string expresses a wider diversity of types including MX, SRV, and DNSKEY query types. This finding also highlights the challenge associated with mitigation since diversity complicates mitigation coordination across more systems, applications, etc.

It is also worth noting that also with the deployment and standardization of the DNS privacy enhancing Qname Minimization technique, query type diversity measurements will be impaired at the upper-level of the DNS hierarchy.

5.5   CDM: Label Diversity

The diversity of labels under a name collision string is another critical diagnostic measurement to consider. This measure helps provide insights into the variety of systems, devices, networks, or entities that are leaking DNS queries. This information may also help construct mitigation strategies. For example, an interesting label diversity analysis of .INTERNAL revealed that 91% of all the leaking queries used three common second level labels. This is in stark contrast to other strings, including .CORP, .HOME, and .MAIL, that exhibit a very large and diverse set of second level labels.

In addition to label diversity, measuring key labels that are associated with various network service discovery protocols, as enumerated in section 4.2.2, may also provide a critical diagnostic measure into known name collision exploitatable scenarios. Again, it is worth noting that Qname Minimization, when more ubiquitously deployed, will impair label analysis at the upper-level of the DNS hierarchy.

5.6   Comparison with 2012 Analysis

The Interisle and JAS collision analysis that began roughly a decade ago and concluded in 2015 is still sound. Nothing the NCAP analysis uncovered contradicts those earlier studies. While there are notable differences in data sets and anomalies, both the measured potential impact and projected harm essentially agree between the earlier studies and today. Those earlier studies examined the issue of collisions for .CORP, .HOME, and .MAIL, and the other applied for strings. They found that name collisions are practically unavoidable, but that the three names under consideration within the NCAP discussion group were particularly problematic.

This case study expands on the previous analysis by providing a longitudinal measurement that also includes insights to several new *critical diagnostic measurements*. These new findings support the previous risk assessment and also highlights new additional measurements and contextual insights into current name collision risks for .CORP, .HOME, and .MAIL.

5.8. Discussion on Potential Mitigations

In terms of name collision mitigation, we have seen progress in some cases[19] and significant uncertainty in others. Where there is progress for example, is when a significant proportion of DNS collision traffic under study can be attributed to a small set of large organizations and equipment. However, for popular undelegated names such as .CORP, .HOME, and .MAIL, there are many cases of incredibly high dispersion rates, where usage of these names is widespread and difficult to attribute to any one particular class of application or origin network. This diversity poses the greatest challenge for mitigation since there are potentially countless ramifications to be identified and resolved.

One successful mitigation was observed in the recent past with the NCAP study. The Chromium software changed its behavior in testing connectivity through DNS queries. This resulted in a significant and noticeable decrease in query names seen in the monitored period for collision domains under study. While not all collisions can be avoided through a single point of control, this demonstrated the far reaching effects of identifying and changing behavior in just one part of the ecosystem.

5.9. Discussion on Impact

An ongoing question whose answer continues to be particularly elusive has been "What is the harm if this name were delegated?" The study has shown very clear evidence of "impact", e.g., a tremendous amount of query traffic would be affected by a delegation of .CORP, .HOME, and .MAIL. However, one of the weaknesses of this study was the inability to truly measure the harm that might manifest as a result of a delegation. This led to at least two considerations that merit further investigation to better address this question:

- Study new delegations that have occurred for collision impact. The study of the name collision reports for the delegated strings in the 2012 round of TLDs would shed light on this topic.
- Conduct controlled experiments on .CORP, .HOME, or .MAIL to measure impact and potential harm. A controlled experiment doesn't have to be done on the entire Internet, but could be done on a subset population (e.g. cooperating ISP, region, users).

Up until now, speculation and educated guesses as to what harm may come from delegations has ruled the day. Recall, it was only after Verisign deployed SiteFinder[20] when the full ramifications of this change constituted. Instead of performing the actual delegation, however, a set of

---

[19]Verisign Outreach Program Remediates Billions of Name Collision Queries
https://blog.verisign.com/domain-names/verisign-outreach-program-remediates-billions-of-name-collision-queries/
[20] https://www.icann.org/en/system/files/files/report-redirection-com-net-09jul04-en.pdf

controlled experiments in portions of the network or by networks that are willing and able to conduct such exercises can prove vastly helpful in better understanding the full effect of what a new delegation means for real impact.

Finally, when considering the impact, it is important to note that we focus the impact of name collisions on end users and systems. The root server system has been resilient and has no issues handling the query volume of non-delegated TLD strings.

6   Conclusions

The recent NCAP study has helped provide insight into the current trends of DNS traffic on undelegated names and with potential collisions in some of the most notable potential TLDs such as .CORP, .HOME, .MAIL, .LAN, .LOCAL, and .INTERNAL. The analysis illuminates the significant impact delegations would mean, and presents some insight into the potential harm that may result. A precise accounting of harm is difficult to quantify and this recent study provides only hints at what harm may actually occur if delegations were to occur. Additional analysis, particularly in other parts of the DNS would help to better measure impact from different vantage points in the system.

Controlled experiments could help to further measure the real world harm that may or may not occur if such delegations were to take place in the face of potential collisions. Furthermore, future analysis should also take into account the recent interest and use of new technologies such as QNAME minimization and DNS over HTTPS.

Notwithstanding the limitations of our analysis, there would be a clear and significant impact if .CORP, .HOME, or .MAIL were delegated. The specific harm is difficult to quantify due to the sheer volume and diversity of query traffic. It is reasonable to believe that the significant diversity across the multiple critical diagnostic measurements would require enormous additional effort to analyze the impact, identify real and potential harm, and devise mitigation strategies.

7   Acknowledgments, Statements of Interest, and Dissents, Alternative Views and Withdrawals

In the interest of transparency, these sections provide the reader with information about aspects of the SSAC process. The Acknowledgments section lists the SSAC members, outside experts, and ICANN staff who contributed directly to this particular document. The Statements of Interest section points to the biographies of all SSAC members, which disclose any interests that might represent a conflict—real, apparent, or potential—with a member's participation in the preparation of this Report. The Dissents and Alternative Views section provides a place for individual members to describe any disagreement with, or alternative view of, the content of this document or the process for preparing it. The Withdrawals section identifies individual members who have recused themselves from discussion of the topic with which this report is concerned. Except for members listed in either the Dissents and Alternative Views or Withdrawals sections, this document has the consensus approval of all of the members of SSAC.

7.1   Acknowledgments

The committee wishes to thank the following NCAP Discussion Group members for their time, contributions, and review in producing this report.

NCAP Discussion Group members (* indicates SSAC member)
James Galvin - co-chair*
Patrik Faltstrom - co-chair*
Matthew Thomas - co-chair*
Jaap Akkerhuis*
Thomas Barrett
Dmitry Belyavsky
Justine Chew
Steve Crocker*
Julie Hammer*
Merike Kaeo*
Rubens Kühl
Warren Kumari*
Barry Leiba*
Danny McPherson*
Brantly Millegan
Ram Mohan*
Russ Mundy*
Jeff Neuman
Eric Osterweil
Chris Roosenraad*
Rod Rasmussen*
Anne Aikman-Scalese
Jeff Schmidt
Greg Shatan
Suzanne Woolf*

ICANN staff
John Kristoff (ICANN Research Fellow)
Kathy Schnitt
Kimberley Carlson
Matt Larson
Steve Sheng (editor)

7.2   Statements of Interest

SSAC member biographical information and Statements of Interest are available at:
https://www.icann.org/resources/pages/ssac-biographies-2020-07-02-en

NCAP Discussion Group member Disclosure of Interest are available at:

Case Study of Collision Strings

https://community.icann.org/display/NCAP/NCAP+Discussion+Group

7.3   Dissents and Alternative Views

Appendix A: DNS Evolution from 2012 to 2021

Since the last 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 nameservers in the DNS hierarchy. These technologies include running a Root Server Local to a Resolver (RFC 8806), Aggressive Use of DNSSEC-Validated Cache (RFC 8198), DNS Query Name Minimization (RFC 7816), DNS Queries over HTTPS (RFC 8484). In this section, we provide some background on these technologies, to better understand how these standards and technology changes will influence data collection capabilities as well as their impacts to data analysis of DNS traffic.

A1   Running a Root Server Local to a Resolver

Some resolver operators, usually due to either performance or privacy concerns, prefer to limit over the network DNS query traffic between themselves and authoritative server operators. Sometimes an instance of the authoritative zone can be located near the resolver, but this is not practical for the entire namespace. Another option is for the resolver to obtain and serve local copies of the authoritative data itself. IETF RFC 8806 describes this approach for serving the root zone from a resolver.

For resovers that implement this technique, the resolution process will short circuit the query to root name servers, instead of answering queries directly and immediately. It is unknown how prevalent this practice is, but it is believed to be deployed in some parts of the Internet for reasons cited above. It may also be done to avoid not only performance and privacy limitations, but to work around active interception and traffic manipulation, a form of avoiding network interference in the otherwise normal query path.

If widely deployed, RFC8806 technology reduces the volume and fidelity of DNS queries observed at the root.

A2   Aggressive Use of DNSSEC-Validated Cache

The DNS relies upon caching to scale; however, the cache lookup generally requires an exact match. RFC 8198 specifies the use of two DNSSEC related resource records (i.e., NSEC/NSEC3) to securely handle non-existent names in the DNS. It allows DNSSEC-validating resolvers to generate negative answers within a range and positive answers from wildcards. This increases performance, decreases latency, decreases resource utilization on both authoritative and recursive servers, and increases privacy.[21]

A3   DNS Query Name Minimization

RFC7816 defines the "DNS Query Name Minimisation to Improve Privacy." Prior to RFC7816, when a resolver received the query "what is the AAAA record for www.example.com?", it sent the full query name to a root server.

---

[21] See RFC 8198, Aggressive Use of DNSSEC-Validated Cache

RFC7816 minimises the amount of data sent from the DNS resolver to the authoritative name server. Instead of sending the full QNAME (www.example.com) and the original QTYPE (AAAA) upstream, a resolver that implements QNAME minimisation and does not already have the answer in its cache sends a request to the name server authoritative for the closest known ancestor of the original QNAME. In the example above, sending "What are the NS records for .com?" would be sent to the root (assuming the resolver does not already have the answer in the cache).

Implementation of query minimization reduces the visibility of queries at the root servers. If RFC7816 is fully adopted, the root server would have no visibility for strings other than the TLD level and no visibility of the actual query type. This loss would significantly limit the analysis possible when evaluating name collisions and considering a mitigation strategy.

A4   Evolution of DNS Resolution

In the original DNS protocol, DNS queries and responses are traditionally transported in clear text (unencrypted) over the underlying UDP or TCP protocol.[22] In the traditional model of DNS resolution, a DNS library is included in operating systems. While the resolver used by this library is sometimes configured by the end user, it is more often configured by the service provider through the use of the Dynamic Host Configuration Protocol (DHCP). The configured resolver is mostly a system-wide setting and generally not application-specific. As identified in SAC109, over the last decade, we see a confluence of factors disrupting the traditional model. These are:

- New technical standards and implementations have been developed to convey DNS queries and responses over alternative transport protocols, examples of which include HTTPS,[23] TLS,[24] DTLS,[25] and QUIC.[26] Such efforts move away from unencrypted UDP and TCP (the traditional protocols used to transport DNS traffic) and appear to be driven by privacy, confidentiality, security, and robustness considerations, as well as a desire for increased levels of control over the retrieval of DNS information for client applications.

- As of March 2020, open public resolver services not operated by ISPs handle roughly 16% of all DNS resolution on the Internet.[27] Examples of open public resolver operators include; Google,[28] OpenDNS[29] and 114DNS.[30] Such efforts consolidate DNS-based user behavior and substantially change the characteristics of DNS query traffic at root servers.

---

[22] See RFC 1035
[23] See RFC 8484
[24] See RFC 7858
[25] See RFC 8310
[26] See Huitema, C., Shore, M., Mankin, A., Dickinson, S., Iyengar, J., "Specification of DNS over Dedicated QUIC Connections", draft-huitema-quic-dnsoquic-07, September 2019, https://datatracker.ietf.org/doc/draft-huitema-quic-dnsoquic/
[27] See APNIC, Use of DNS Resolvers for World, https://stats.labs.apnic.net/rvrs/XA?hc=XA&hl=1&hs=1&ht=0&w=30&t=0&s=0
[28] See Google Public DNS, https://developers.google.com/speed/public-dns/
[29] See Cisco OpenDNS, https://www.opendns.com/
[30] See 114DNS, https://www.114dns.com/

- Vendors of browsers and other applications have incentives to embed addresses of resolvers directly into their applications, thereby bypassing the traditional model of using the resolver(s) configured in the operating system, and instead creating application specific resolution behaviors. An example of this is Mozilla Firefox, which at the time of publication, is sometimes pre-configured with a DoH capable resolver compliant with Mozilla's Trusted Recursive Resolver Program[31] (TRR).

- Mechanisms that use the DNS as a control point are continuing to be developed and implemented. These include DNS-aware firewalls and monitoring tools, as well as some forms of IPv6 transition technologies. These may be local to end-users or centralised, and are often deployed at the direct request of the end-user or to enforce a policy of the local network.

Among these new technologies, DoH appears to be making the most inroad. At the time of this publication, major operating systems, recursive DNS resolvers offer native support. Web browsers are including the technology in recent software releases and turning them on by default. These are documented below.

	DoH Support[32]
Operating Systems	IOS14, MacOS11, Windows 10 (upcoming)
Recursive DNS Resolvers	BIND 9.17.10 added native support for DOH, PowerDNS 1.4.0, Unbound 1.12.0
Web Browsers	Google Chrome 83, Microsoft Edge, Mozilla Firefox (with Cloudflare), Opera (with Cloudflare)

As voiced by the SSAC in SAC109, the main concern was that DNS resolution was fundamentally changing with the advent of DoT and especially DoH. Resolution from the application all the way up to the authoritative and back again was fundamentally changing how names get resolved on the Internet. In addition to changes in the resolution process, changes impacting the fundamental principle of a single shared namespace could also come about. Applications using DNS were increasingly performing the full stub resolution themselves, without using an operating system wide resolver library. Thus, different applications running on the same computer are more commonly getting different answers from the DNS.

---

[31] See Mozilla's Trusted Resolver Program, https://wiki.mozilla.org/Trusted_Recursive_Resolver
[32] See https://en.wikipedia.org/wiki/DNS_over_HTTPS