

Root Cause Analysis - New gTLD Collisions

1. Introduction	3
2. Background	3
2.1. DNS Suffix Configuration	4
2.2. Controlled Interruption	4
2.3. Chrome Browser NXDOMAIN Probing	5
2.4. WPAD-Related Queries	5
2.5. ISATAP-Related Queries	5
3. Data Sets	5
3.1. Name Collisions Reports Submitted via ICANN's Web Form	5
3.2. DNSDB	6
3.2.1. Controlled Interruption Queries	6
3.2.2. Queries Post Controlled Interruption	7
3.3. DITL	7
3.4. Web Search Results	7
4. Name Collisions Report Analysis	8
4.1. TLD Statistics	8
4.2. Reporting Entity	10
4.3. Impact	11
4.4. Root Cause Identification	13
4.5. Other Observations	14
5. Web Search Results Analysis	14
5.1. TLD Statistics	14
5.2. Applications In Use	16
5.3. Name Collisions Root Causes	17
5.4. Other Observations	20
6. Leaked Suffix Identification	20
6.1. Suffix Identification via Chrome NXDOMAIN Probing	21
6.2. Suffix Identification Using WPAD and ISATAP DNS Queries	22
6.3. Results	23
6.3.1. Validation of Identification Methods	23
6.3.2. Distribution of Suffixes Across TLDs	24
6.4. Analysis of Chrome Identification Methodology	26
6.5. Validation of WPAD Identification Methodology	26
7. Controlled Interruption Analysis	26

8. Root Server Query Analysis	28
8.1. Data Set	29
8.2. Results	30
8.3. Qname Minimization Considerations	35
8.3.1. Summary of Recent Study of Qname Minimization	35
8.3.2. Application of Qname Minimization Data	36
9. Name Collisions Survey	41
9.1. Survey Content	41
9.2. Survey Recipients	41
9.3. General Survey Results	42
9.3.1. TLDs Used	42
9.3.2. Technical Issues Experienced	43
9.3.2.1. DNS Resolver Configuration	43
9.3.2.2. Discovery, Impact, and Resolution	44
9.3.2.3. Root Cause Identification	44
9.3.2.4. Other Observations	44
9.4. Targeted Survey Results	45
10. Discussion	45
10.1. Findings	45
10.2. Proposed Future Work	47
Appendix A - Name Collisions Report Form	49
Appendix B - Web Search Results for “127.0.53.53”	52
Appendix C - General Name Collisions Survey	55
Appendix D - Targeted Name Collisions Survey	61
Appendix E - General Email Sent to NANOG Subscribers	67
Appendix F - Targeted Email Sent to AS Contacts	69

1. Introduction

In 2013, the International Corporation for Assigned Names and Numbers (ICANN) began allowing new top-level domains (TLDs) to be introduced into the DNS root zone. Analysis showed that this new practice might adversely affect existing networks and systems, because of *name collisions*: the notion that a system uses a given DNS namespace in *private* and *relies* on it not resolving in the public DNS, but then, through delegation, that namespace becomes publicly resolvable. Because of the potential problems associated with name collisions, newly delegated TLDs were required to go through a period known as “controlled interruption,” beginning in August 2014. This practice, described in more detail hereafter, was intended to make users and administrators that *might* be affected by a TLD’s delegation aware of its delegation preemptively—before the problems became critical.

ICANN’s Security and Stability Advisory Committee (SSAC) commissioned the Name Collisions Analysis Project (NCAP) to “facilitate the development of policy on Collision Strings to mitigate potential harm to the stability and security of the DNS posed by delegation of such strings.”¹ This document is part of the NCAP effort. In particular, this study seeks to analyze various aspects of name collisions and controlled interruption since controlled interruption was instituted and to identify the root cause of related incidents reported by affected parties to ICANN. The analysis primarily takes into consideration TLDs delegated between August 2014 and June 2021. Three data sources are used in this analysis:

- collision reports submitted via ICANN’s name collisions Web submission form²;
- passive DNS from the 100 days of controlled interruption during the initial delegation of each TLD; and
- root query data from the 48-hour once-yearly day-in-the-life (DITL) collection from 2014 to 2021.

We begin with some technical background information related to our analysis and then briefly describe our data sets. We then perform an analysis of the name collision reports submitted to ICANN. Next we describe our methodology for quantifying the private use of newly delegated TLDs, and we share the results of our analysis of controlled interruption and leaked DNS queries intended for privately maintained namespace. We describe a survey that we commissioned to obtain more qualitative data associated with our analysis. Finally, we summarize our findings and propose future work.

2. Background

This section provides technical background related to our study.

¹ <https://community.icann.org/display/NCAP/>

² <https://www.icann.org/en/forms/report-name-collision>; [Appendix A](#).

2.1. DNS Suffix Configuration

The network configuration for most operating systems includes the option for a DNS “suffix” (e.g., `example.com`) to be specified for various purposes. The system’s stub resolver library, which is used by applications to resolve DNS names to addresses, might apply this domain to unqualified DNS names that are to be resolved (e.g., `foo` becomes `foo.example.com`). Or the domain might be used to identify certain network resources associated with the organization, such as the organization’s HTTP proxy server (see [Section 2.4](#)) or potential routers for IPv6-over-IPv4 tunneling (see [Section 2.5](#)).

This domain is configured in the “domain” and “search” entries of `/etc/resolv.conf` on UNIX and Linux systems. In macOS, the DNS configuration pane contains a “Search Domains” box to add this domain. On Windows, the “DNS suffix search list” is used.

Throughout this document, we use the term *DNS suffix* to refer to this domain, independent of the specific system on which it is configured.

2.2. Controlled Interruption

Some systems query the public DNS for names under a non-existent TLD, for a variety of possible reasons. *Prior* to the delegation of the TLD in the root zone, these names would not resolve but would rather result in an NXDOMAIN (name error)—or *negative response*. In some cases, a negative response from the public DNS was *relied on* to properly access a given resource (e.g., search list processing). In other cases, a negative response from the public DNS would simply *prevent* a system from accessing a given internal resource except from *within* the proper network for doing so (e.g., private namespace used within a corporate network). In all cases, negative responses played a role in *expected* application behavior.

Controlled interruption involves inserting *wildcard* records in the otherwise empty zone file associated with a previously undelegated TLD. The wildcard `A` (IPv4 address) record in the zone file maps to a non-routable address: `127.0.53.53`. Thus, *any* `A`-type query made to the public DNS for names under that TLD result in a *positive* response—as opposed to the negative response that would have resulted prior to controlled interruption. Note that there is no IPv6 equivalent for queries of type `AAAA` (IPv6 address).

Controlled interruption has been required of all TLDs delegated in the root zone since August 2014, for the first 100 days of its delegation. In cases where negative responses were required for expected behavior, it was expected that systems encountering controlled interruption would experience some sort of disruption to their “normal” behavior, a sort of signal that something had changed in the public DNS. Additionally, it was the hope that this disruption would be noticed by the affected parties, such that they would investigate and take action, by reporting the problem and/or changing their configuration.

2.3. Chrome Browser NXDOMAIN Probing

On startup, the Google Chrome Web browser historically issued three queries, appending the system DNS suffix (see [Section 2.1](#)) to three randomly-generated alphabetic strings. This is to detect infrastructure providing synthetic positive responses to DNS queries that would otherwise be classified as name errors (NXDOMAIN). During controlled interruption for a given TLD, queries under that TLD related to Chrome NXDOMAIN probing result in positive DNS responses.

2.4. WPAD-Related Queries

With the Web Proxy Auto Discovery Protocol (WPAD), browsers (e.g., Mozilla Firefox and Google Chrome) and operating systems (e.g., MacOS and Windows) auto-detect HTTP proxy settings using the DNS and HTTP. The specification designates that a WPAD client append the DNS suffix with which a system is configured (see [Section 2.1](#)) to the label `wpad`. If no answer is found for the newly-formed domain name, then the left-most label in the DNS suffix is stripped, and `wpad` is prepended to the resulting suffix. Thus, a browser on a system configured with DNS suffix `foo.example.com` would issue a DNS query for `wpad.foo.example.com` then (assuming the domain name did not resolve) `wpad.example.com`, etc. This process is repeated until an answer is found or the suffixes are exhausted. During the controlled interruption period for a given TLD, all WPAD-related queries under the TLD result in positive DNS responses.

2.5. ISATAP-Related Queries

The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) is used for creating a link-local IPv6 address from an IPv4 address and discovering a neighbor through which IPv6 traffic might be tunneled. As part of this process, a host discovers potential routers by performing a DNS lookup for the qname formed by appending the system's DNS suffix (see [Section 2.1](#)) to the string `isatap`. Thus, for a system configured with the DNS suffix `example.com`, the DNS lookup would consist of a lookup for `isatap.example.com`. During the controlled interruption period for a given TLD, all ISATAP-related queries under the TLD result in positive DNS responses.

3. Data Sets

In this section, we describe the data sets that were used as the basis for our analysis.

3.1. Name Collisions Reports Submitted via ICANN's Web Form

After ICANN began introducing new TLDs into the root zone, a Web form was created whereby users could submit reports of problems experienced, each potentially related to the delegation

of new TLDs³. Each report included, among other information, the date of the report, the TLD in question, a brief description of the problem, and contact information of the submitter. The entire form is included in [Appendix A](#). We use the data from these reports to better understand user and organization experience associated with the delegation of new TLDs in [Section 4](#).

3.2. DNSDB

DNSDB, operated by Farsight Security (part of DomainTools), is a DNS database populated by passive DNS sensors at operators world-wide. It contains historical domain-name-to-resource mappings going back more than 10 years. For example, it could show that `example.com` (A record type) resolved to 192.0.2.1 from March 2014 to October 2015 and to 192.0.2.2 from December 2015 to February 2019. It also supports historical response data for other record types, including NS, MX, and others. However, it only contains an entry where there is a legitimate mapping observed by a sensor. Thus, the database is limited to network locations where sensors are deployed. Also, if an observed query for a given name results in a negative response (i.e., no mapping), DNSDB will have no entry for that name.

We used DNSDB to create two data sets in this work: *query names* observed *during* the controlled interruption period; and *mappings* observed *since* controlled interruption. We describe each in the following sections.

3.2.1. Controlled Interruption Queries

We used ICANN's published list of delegated strings⁴ to obtain the list of TLDs delegated between August 2014 and June 2021, as well as the delegation date of each. August 2014 was when the requirement for controlled interruption began for newly delegated gTLDs. The following table shows the breakdown by year of each of the 885 domains delegated during the time period:

Year	TLDs Delegated	Year	TLDs Delegated
2014 (Aug - Dec)	131	2018	5
2015	390	2019	3
2016	340	2020	4
2017	12	2021 (Jan - Jun)	0
Total: 885			

For each of the new TLDs delegated, we issued a DNSDB query to solicit mappings observed during the dates of its control interruption period—i.e., the first 100 days of its delegation.

³ <https://www.icann.org/en/forms/report-name-collision>

⁴ <https://newgtlds.icann.org/en/program-status/delegated-strings>

Because controlled interruption results in a mapping (i.e., to 127.0.53.53) for *any* DNS queries under the TLD, the DNSDB queries effectively yielded every DNS name *queried* during the controlled interruption period—and observed by passive DNS sensors—for DNS names under the new TLD, along with a count of how many times it was queried. We refer to this data set as DNSDB-CI.

3.2.2. Queries Post Controlled Interruption

Requesting a complete history of *all* DNS mappings observed for every one of the 885 new TLDs delegated since their controlled interruption period ended would have been infeasible because the data sets would be so huge. However, for this analysis, we were interested in only the subset of namespace under each TLD that was associated with name collision activity. This namespace is identified in [Section 6](#) and refined in [Section 8](#), ultimately resulting in 2,266 subdomains associated with 166 of the new TLDs. We issued queries to DNSDB for all query names under each of the 2,266 subdomains (using a wildcard DNSDB query, such as `*.example.com` for the DNS suffix `example.com`), in each case requesting all mappings observed since the 100-day period of controlled interruption for the TLD associated with the subdomain. We refer to this data set as DNSDB-PostCI.

3.3. DITL

Various DNS root server operators contribute to a yearly collection of 48 hours of DNS queries observed at the root server system. This collection is known as the “Day in the Life” or DITL collection and is sponsored by the DNS Operations, Analysis, and Research Center (DNS-OARC). For this analysis, we extracted the query name and querying IP address for all queries associated with the 2,266 subdomains that we identify in [Section 8](#) for DITL collections between 2014 and 2021, inclusive, from root letters A, C, H, and J. This subset of four root letters was selected because each of these letters was available in each of the DITL years we were interested in (not all root letters are represented in all years). We refer to this data set as DITL-2014-2021. This is further described in [Section 8](#).

3.4. Web Search Results

We performed two Web searches using Google’s search engine. We searched for the following search terms between September 28 and October 4, 2022: “controlled interruption” and “127.0.0.53”. Each search term included the quotation marks. In each case, we looked at the first six pages of search results. For the search term “controlled interruption,” every result was either completely unrelated to the controlled interruption implemented by ICANN, or it involved documentation or announcements involving controlled interruption. For the search term “127.0.53.53,” we observed 17 results that described unique cases in which controlled interruption was experienced in such a way that the normal flow appeared to be disrupted. If the page in the result appeared to convey a matter of mere curiosity about behavior, rather than disruption, then we excluded the results. The full set of results are found in [Appendix B](#). An analysis of the results is found in [Section 5](#).

4. Name Collisions Report Analysis

We now analyze the reports submitted to ICANN via the Web form (see [Section 3.1](#)). We note that this data set has inherent bias in three ways. *First*, the submission of the report itself *implies* that a user or organization was impacted in some way by name collisions, so we cannot suggest that the reports herein are the *only* experience that was had; it is possible that some users of private DNS namespace were not impacted and that their story is not captured. *Second*, the submission implies that they found the online form. This means that the question of the *effectiveness* of the use of the controlled interruption IP address (127.0.53.53) in helping the user or administrator trace the problem to ICANN and the delegation of new TLDs cannot be evaluated; there is simply nothing in this data set to compare *against*. *Finally*, ICANN’s Web form invites users to submit a report only if they are “suffering demonstrably severe harm as a consequence of name collision.” Thus, some users impacted by name collisions—but not in such an extreme way as described by the form instructions—might have been dissuaded from submitting a report at all. Later in the paper (see [Section 9](#)) we describe a survey sent to a general audience of network administrators as well as a targeted audience of organizations potentially affected by the delegation of new TLDs—a study without those same biases.

4.1. TLD Statistics

The following table contains a summary of the reports submitted, based on factors such as the date of the report, the TLD and its delegation date, and the reporting entity.

Category	Count	Subcat. %	Total %
Total reports	47	100%	100%
do not include TLD	4	8.5%	8.5%
include TLD	43	91%	91%
delegated prior to new TLD program*	7	16%	15%
delegated as part of new TLD program	36	84%	77%
prior to controlled interruption (pre-Aug 2014)**	2	6%	4.3%
with controlled interruption (Aug 2014 or later)	34	94%	72%
report date is during controlled interruption	25	74%	53%
report date is post controlled interruption	9	26%	19%
reported by organization	24	71%	51%


reported by individual	9	26%	19%
reported origin unknown	1	3%	2.1%
Total TLDs reported	20	100%	100%
delegated prior to new TLD program*	1	5%	5%
delegated as part of new TLD program	19	95%	95%
prior to controlled interruption (pre-Aug 2014)**	2	11%	10%
with controlled interruption (Aug 2014 or later)	17	89%	85%

Each percentage in the “Subcategory %” column is taken from the “Count” in the “Parent” category or subcategory (i.e., the bolded count most immediately above). The percentages in the “Total %” column are taken from the “Count” in the “Total TLDs” or “Total Reports” category.

While the table captures the data of all reports, we pay particular focus to the subset of 34 (72%) reports that pertain to TLDs delegated after the controlled interruption period. Of the 20 TLDs mentioned, 17 (84%) fit this category. Other TLDs mentioned are `name`, `nyc`, and `kitchen`. The `name` TLD, delegated before the new TLD program (*), was associated with 7 reports. All 7 reports were associated with the delegation of `wpad.domain.name`, which allowed the HTTP traffic of affected parties to be monitored and intercepted by third parties. This is discussed in a separate report (ref report). The `nyc` and `kitchen` TLDs were delegated as part of the new TLD program prior to controlled interruption**.

The following plot shows the distribution of reports by TLD, including those that were not delegated as part of the new TLD program (*) and those that were delegated prior to controlled interruption (**). For the 17 reported TLDs that were delegated after controlled interruption was introduced, each bar in the plot is composed of the numbers of reports received during and after the controlled interruption period for the TLD.

Name Collisions Reports by Report Date

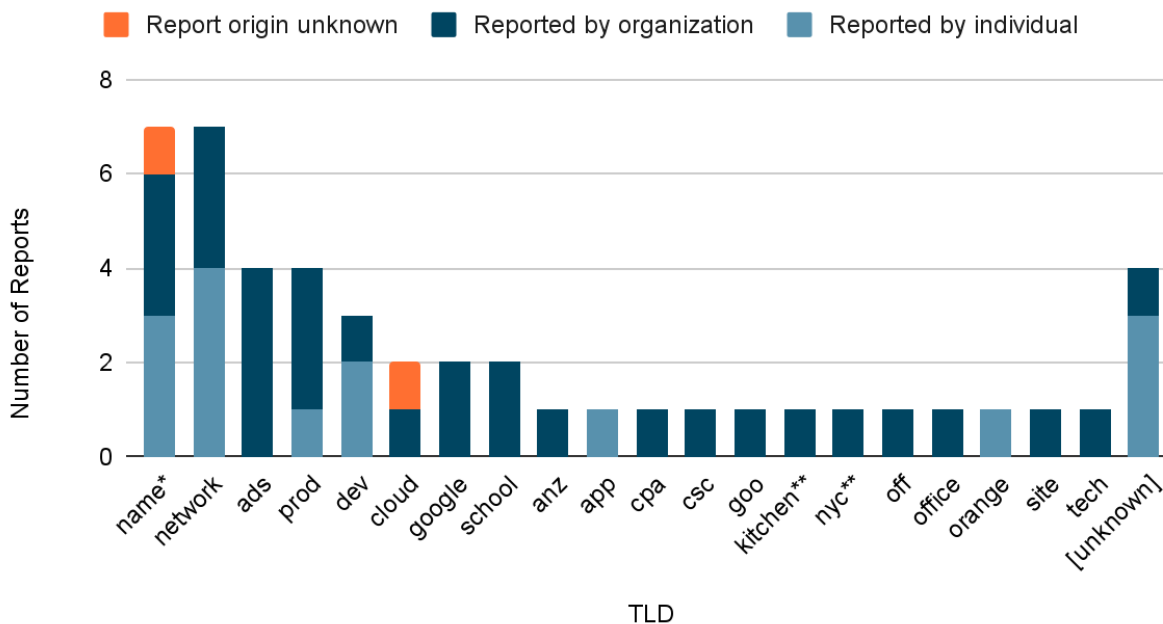


In most (74%) cases, report(s) were submitted during the controlled interruption period for the TLDs; in the remaining cases, the report was submitted after the controlled interruption period. For three TLDs, (`dev`, `app`, and `cpa`), *all* reports came *after* the controlled interruption period. With the exception of `cloud`, all TLDs for which reports were received after the controlled interruption period were observed using the controlled interruption IP address (127.0.53.53) beyond the designated time: an additional 693 days for `ads`, 1,042 days for `dev`, 593 days for `app`, and at least 644 days for `cpa`. (Domain names within the `cpa` TLD were still resolving at the time we retrieve the historical data.) See [Section 7](#) for more. In every one of these cases, the report date was prior to the date that the controlled interruption IP address was last observed for the TLD in question.

4.2. Reporting Entity

Reports were categorized as having been submitted on behalf of an organization, submitted by an individual, or for which the origin was unknown. Considering only the 34 reports for TLDs delegated after the introduction of controlled interruption, the counts were 24 (71%) by organization, 9 (26%) by individual, and 1 (3%) unknown. The breakdown is shown in the following plot, which includes TLDs that were not delegated as part of the new TLD program (*) and those that were delegated prior to controlled interruption (**).

Name Collisions Reports by Reporting Entity



Two standouts are `ads` and `school`, for which reports were made exclusively by organizations. The `ads` TLD (as well as `local` and `intern`) is reportedly (as indicated in one report, but not independently verified) used in books and training resources for creating Microsoft Active Directory domains. Other reports indicated that `office`, `off`, `school`, and `site` are used by organizations for Active Directory services. `school` is reportedly used by some school districts as a private DNS namespace and—at least in some cases—for Active Directory, as mentioned previously.

4.3. Impact

In addition to the quantitative analysis associated with the affected TLDs and their reporting organizations, we now use additional report details to add a qualitative analysis. We consider only the 34 reports associated with TLDs delegated after the introduction of controlled interruption.

We first categorize impact based on the self-reported description and size of organization, if reported. We group incidents into four categories based on what we could infer from the content of these fields:

- *severe*. A large number of users were affected, network access as a whole was affected, and/or the submitter described the impact as severe.
- *significant*. The number of affected users or systems was more moderate, and/or only specific network applications were impacted.
- *small-scale*. The number of affected users or systems is small, and/or impacts seem nominal.

- *unknown*. There is insufficient data in the report to justify assignment to one of the other categories.

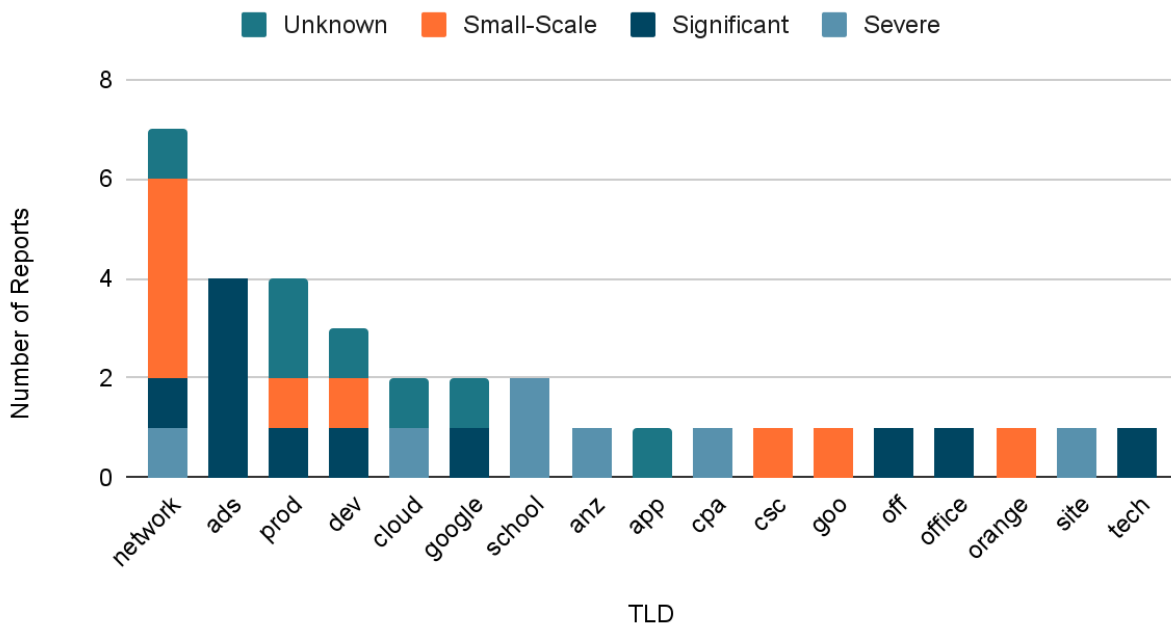
In the following table, we list the count for each category as well as sample comments from each report that led us to categorize them accordingly (except for unknown, for which details were too few to categorize otherwise):

Category	Count	Descriptions
Severe	7	<p>“more 30,000 employees in over 7 countries and these employees interact with one another and with the organization via an internal network.... employees had trouble accessing their internal network.”</p> <p>“Network down, no internet access”</p> <p>“this is causing all of our staff laptops to crash when off of our network... this is causing severe problems”</p> <p>“All clients are having problem and freeze during usage.”</p> <p>“This is affecting all users in the organisation at various times”</p> <p>“1400 servers in 800 schools”</p> <p>“The scale of the impact is fairly critical. All VPN tunneling to our network cannot resolve DNS.... it is affecting all of our external users needing to resolve anything internal via DNS name. 300 users affected. All systems that reside outside of the office...”</p>
Significant	10	<p>“CRM, MAIL and other Services provided by our Company do not work correctly”</p> <p>“Unable to send mail”</p> <p>“150 users”</p> <p>“No network shares access.”</p> <p>“Do not operate normally computers are connected to a domain controller”</p> <p>“VPN sessions with split tunnelling do not work as the DNS lookup fails.”</p> <p>“If our applications are started before the corporate VPN connection is up... we cannot use the app's anymore”</p> <p>“Unable to resolve internal Hostnames”</p> <p>“some Clients... not correct working with the DNS Suffix Searchlist”</p> <p>“Users cant loggon to local domain”</p>
Small-Scale	10	<p>“Internet browsing issues from LAN”</p> <p>“can't access to some servers”</p> <p>“home network disruption”</p> <p>“Having trouble connecting to some network resources”</p> <p>“i cant use my sub domain... any longer”</p>
Unknown	7	
Total	34	

Our analysis shows that only half (50%) were classified as either severe or significant. However, as noted previously, the text on the submission form suggests that reports are for systems “suffering demonstrably severe harm as a consequence of name collision” and that emergency response actions would be taken “only where there is a reasonable belief that the name collision presents a clear and present danger to human life.” Thus, either our classifications are inaccurate, the reports understate the magnitude of the problems experienced, and/or the reports were submitted notwithstanding the suggested criteria—perhaps in an effort to officially document the problem.

The following figure shows a plot of the severity of the 34 reports by TLD:

Name Collisions Reports by Report Severity



Of the reported TLDs, 14 (83%) included at least one report categorized as causing significant or severe impact. Thus, severity was not isolated.

4.4. Root Cause Identification

Clearly, all 34 reports were led to ICANN’s name collisions report page to submit the report. Of the 34 reports, 8 (24%) specifically either mentioned “127.0.53.53” or referred to “controlled interruption” by name. It is unclear from the other reports whether the controlled interruption IP address itself contributed to finding the ICANN form, but we can say that at least one quarter observed 127.0.53.53.

4.5. Other Observations

We here record two significant trends that we observed in our analysis of the reports.

First, 8 of the reports mentioned “remote users” or “VPN” (Virtual Private Network). These account for 33% of reports submitted by organizations and 17% of all reports. A VPN is typically used to connect the systems of these users to the corporate network. Once VPN-connected, the remote system typically uses the corporate DNS servers, but prior to connection, they must use a non-corporate (i.e., “public”) DNS resolver. A common configuration for organizations using private DNS namespaces is for the corporate DNS resolvers to be configured to answer authoritatively for the private DNS namespace. This “works” when corporate systems *only* ever issue queries to the corporate DNS resolver—not to the public DNS. However, as evidenced by the submitted reports analyzed in this section, observed leakage of DNS queries for private DNS namespace (see [Section 6](#)), and responses to our survey (see [Section 9](#)), this is not always the case.

Second, of the 24 reports submitted by organizations, 8 (33%) explicitly mentioned Active Directory services. One additional report did not mention Active Directory, but the associated TLD was `ads`, so it might be inferred. Three (37%) of the reports mentioning Active Directory *also* mentioned VPN usage, i.e., that it was the combination of the two that caused the disruption. This shows that the impact of name collisions on systems using Active Directory are not isolated.

5. Web Search Results Analysis

We now analyze the results of the Web search for “127.0.53.53” (see [Section 3.4](#)). Each of these results represents a circumstance in which the IP address 127.0.53.53 was unexpectedly observed in connection with resolving a given domain name ending in a TLD which has been recently introduced into the root zone (with one exception, which will be shown hereafter) as part of the new gTLD program. Thus, we cannot evaluate how often 127.0.53.53 was observed when name collisions were experienced, as this data set *only* includes experiences of name collisions where 127.0.53.53 was observed. However, we again refer the reader to [Section 9](#), where we describe a survey distributed to individuals and organizations *potentially* affected by the delegation of new TLDs, the results of which have no such bias.

5.1. TLD Statistics


The following table contains a summary of the search results, based on factors such as the date of the report, the TLD and its delegation date, and the reporting entity.

Category	Count	Subcat. %	Total %
Total search results	17	100%	100%

do not include TLD	3	18%	18%
include TLD	14	82%	82%
delegated prior to new TLD program*	1	7.1%	5.9%
delegated as part of new TLD program	13	93%	76%
prior to controlled interruption (pre-Aug 2014)**	2	15%	12%
with controlled interruption (Aug 2014 or later)	11	85%	65%
result date is during controlled interruption	5	45%	29%
result date is post controlled interruption	6	55%	35%
Total TLDs in search results	11	100%	100%
delegated prior to new TLD program*	1	9.0%	9.0%
delegated as part of new TLD program	10	91%	91%
prior to controlled interruption (pre-Aug 2014)**	2	20%	18%
with controlled interruption (Aug 2014 or later)	8	80%	73%

The following plot shows the distribution of search results by TLD, including those that were not delegated as part of the new gTLD program (*) and those that were delegated prior to controlled interruption (**). For those results that were associated with the 8 TLDs that were delegated after controlled interruption was introduced, each bar in the plot is composed of the numbers of reports received during and after the controlled interruption period for the TLD. As noted previously, the mappings to “127.0.53.53” were observed for the `bar` and `dental` TLDs (both marked with **) even though they are labeled “Controlled Interruption N/A” because they were delegated prior to the start of controlled interruption.

"127.0.53.53" Search Results by Result Date



Of the search results corresponding to TLDs delegated as part of the new gTLD program (i.e., excluding `int`), only 38% were dated during the controlled interruption period for the TLD. These correspond to 45% when only considering the TLDs that were delegated after controlled interruption (i.e., excluding `int`, `bar`, and `dental`). These fractions are comparatively lower than the 74% observed in our analysis of the reports submitted to ICANN (see [Section 4.1](#)). However, we note that `dev`, `box`, `cisco`, and `cpa` all continued exhibiting controlled interruption behavior (i.e., returning 127.0.53.53 for non-existent domain names) for 1,042 days, 78 days, 193 days, and (at least) 644 days, respectively, [according to DNSDB](#) (see [Section 7](#)). The dates for search results for `dev` and `cpa` were prior to the date that the controlled interruption IP address was last observed. However, the dates of the search results for `box` and `cisco` were past the dates for which the controlled interruption IP address was last observed. Among the possible explanations for the discrepancy are the following. The passive sensors contributing to the historical DNSDB mappings did not have sufficient network placement to observe the controlled interruption experienced by those that posted the report found in the search results. Or it could be that the report (i.e., associated with the Web search result) was made long after controlled interruption was experienced.

The only inexplicable instance of controlled interruption is the one search result corresponding to the `int` TLD. The `int` TLD was delegated in 1988, and we have no data to suggest that it implemented controlled interruption, other than the search result itself.

5.2. Applications In Use

Because the search results often contained more detail than the name collision reports, we were able to glean more about each incident. In 12 (63%) of the 19 reports, a primary application was identified associated with the incident. In cases where the main application was unclear, we categorized it “unknown.” This included cases where we inferred that the application might simply be a diagnostic test but that the main application was something else. The resulting categorization was imperfect but still provided some insight into the use case leading to the collision.

Category	Count	Subcat. %	Total %
Total search results	17	100%	100%
no application identified	7	41%	41%
application identified	10	59%	59%
Web browser	2	20%	12%
ping	2	20%	12%
Apache Kafka (unit testing)	1	10%	5.9%
gitlab-ci-multi-runner	1	10%	5.9%
php, tnsping	1	10%	5.9%
RDP	1	10%	5.9%
SSH	1	10%	5.9%
valet	1	10%	5.9%

First, we note that these results show that there is a variety of applications with which users have experienced name collisions. Additionally, of the search results for which applications were inferred, Web browsers accounted for only 20%.

5.3. Name Collisions Root Causes

The detail in the search results also allows us to better understand the root cause of name collisions affecting applications and end users. We begin with discussion of configurations that contribute to name collisions and then present our findings. Note that these configurations include—but are not limited to—the scenarios described in section 2.3.3 of the NCAP study 1 RFP⁵ and section 2.2 of the NCAP study 1 report⁶.

⁵ <https://www.icann.org/en/system/files/files/rfp-ncap-study-1-09jul19-en.pdf>

⁶ <https://www.icann.org/en/system/files/files/ncap-study-1-report-12feb20-en.pdf>

Private and Non-private. Much of this document refers to the private use of TLD namespace as the primary cause of name collisions. This is the case in which systems use a presumably non-existent TLD to name resources that they wish to access. If queries for domain names under that TLD reach the public DNS authoritative servers, then there is a name collision. While the private use of TLD namespace seems to be the most prevalent use case for name collisions, there are situations in which name collisions do not involve the private use of domains. We refer to such use as non-private. Name collisions involving non-private use of domain names are typically associated with the use of multi-label, unqualified domain names (discussed hereafter).

Single- and Multi-Label Unqualified. Unqualified names are those that are not intended to be resolved without the application of a DNS suffix (see [Section 2.1](#)). There are two variants to those names: single-label (e.g., `foo`) and multi-label (e.g., `foo.bar`). Single-label names traditionally do not resolve to an IP address (exceptions are described later in this section), making them a clear candidate for application of a DNS suffix for proper resolution. In contrast, multi-label, unqualified names have the appearance of being fully qualified, simply because they have more than one label. Yet multi-label, unqualified domain names are known to be used in practice. In the case where the right-most label of a multi-label, unqualified name corresponds to a TLD which has (relatively recently) been delegated is used as the unqualified domain name, the search suffix logic might result in the name being resolved without qualification—ending in a name collision. For example, if `foo.bar` is used as an unqualified domain name, and `bar` is delegated, then `foo.bar` might resolve as if it were fully qualified, regardless of which DNS suffixes are available to be applied.

DNS Suffix Devolution. Some systems use a technique referred to as *DNS suffix devolution* to resolve an unqualified domain name. Given the DNS suffix `foo.bar.com`, suffix devolution involves attempting to resolve the unqualified name `www` first with `www.foo.bar.com` then with `www.bar.com`, etc. An observed variant of this is the following⁷. Given the DNS suffix `bar.local` and the unqualified name `www`, the system attempts to resolve `www.bar` if `www.bar.local` does not resolve. If `bar` corresponds to a TLD that is newly delegated, then there is a name collision.

Deliberately Unresolvable. One of the causes of name collisions involving unqualified names is that a system or user expects an unqualified name to ultimately resolve in a certain way. In order for this ultimate resolution to work as expected, certain *intermediate* iterations of suffix application (or not) should not resolve. However, in some cases, the user or system uses a name with the expectation that *ultimately* it will not resolve. We refer to these names as *deliberately unresolvable*.

Single-Label Resolution. While the DNS protocol does not prohibit domain names with only a single label (e.g. “dotless domains”) from resolving to an IP address, new gTLDs are

⁷ https://www.reddit.com/r/sysadmin/comments/2jcdso/workstations_resolving_domainlocal_to_12705353/

administratively prohibited from allowing this type of resolution⁸. Nonetheless, *single-label resolution* has been enabled for at least some gTLDs, and applications take advantage of this functionality. Whether explicitly or inadvertently, this behavior has resulted in name collisions of various types.

Web Search Term. Many Web browsers use a single input area for users to enter either a search string, a domain name, or a URL—any of which must eventually be converted to a URL. Behavior across browsers varies as to the handling of such an input to make this determination. Some browsers attempt to resolve a single “word” (i.e., no spaces) as a domain name, only using it as a search term after it has been shown to not resolve. In such cases, a word intended as a search term that corresponds to a TLD that has been delegated and configured for single-label resolution, results in a name collision.

VPN. As discussed in [Section 4.5](#), name collisions are often manifest when a VPN is in use. In such cases, the system is potentially operating under two network environments, and what might otherwise be *controlled* use of private namespace can be exposed to public authoritative DNS, resulting in name collisions.

We categorize the use cases according to the following:

- *Private Namespace.* Was the user’s system using the TLD in a private context?
- *Qualification.*
 - Was an unqualified single- or multi-label name the target of resolution? For unqualified, single-label names, was some form of suffix devolution used for search list processing? For unqualified, multi-label names, was the name non-private?
 - Was the name fully qualified? If so, was the name intended to be deliberately unresolvable?
 - Was a single label being resolved to an address? Was the intention for the domain name to be used as a Web search term?
- *VPN.* Was a VPN involved? If so, was the name private?

Category	Count	Subcat. %	Total %
Total	17	100%	100%
Private Namespace			
Private	11	65%	65%
Non-Private	2	12%	12%

8

<https://www.icann.org/en/announcements/details/new-gtld-dotless-domain-names-prohibited-30-8-2013-e>
n

N/A	1	5.9%	5.9%
Unknown	3	18%	18%
Qualification			
Unqualified	6	35%	35%
Single-Label	4	67%	24%
Suffix Devolution	1	25%	5.9%
Multi-Label	2	33%	12%
Non-Private	2	100%	12%
Fully-Qualified	10	59%	59%
Deliberately Unresolvable	3	30%	18%
Single-Label Resolution	1	5.9%	5.9%
Search Term	1	100%	5.9%
Unknown	2	12%	12%
VPN	2	12%	12%
Private	1	50%	5.9%
Unknown	1	50%	5.9%

Perhaps the most interesting aspect of the analysis is that the causes are so diverse, particularly for such a relatively small dataset. Nearly every conceivable use case is represented. As mentioned previously, use of private namespace accounts for the majority (65%) of search results. The use of unqualified names with search list processing accounted for only 35% of cases, with two thirds of those involving single-label, unqualified names and the rest involving multi-label, unqualified names. Of the nearly 60% of cases that involved fully-qualified domain names, 30% were cases where the fully-qualified name was ultimately not intended to resolve. In 12% of cases VPN usage was mentioned—compared to 17% reported in the name collision reports submitted to ICANN. There was one case of nuanced DNS suffix devolution. Finally, there was one case where single-label resolution was at play, and it corresponded to the use of a label as a search term.

5.4. Other Observations

Among the other observations were the following. First, while all 17 search results contained a reference to the controlled interruption IP address, 127.0.53.53, 13 (76%) of those additionally included a reference to ICANN and controlled interruption; only 4 (24%) did not reference

ICANN. Thus, there was a relatively high success rate in associating the IP address 127.0.53.53 to ICANN and controlled interruption—for those that observed the IP address.

Second, we note the sentiment expressed in each of the scenarios gleaned from search results was generally neutral (16 results or 94%). That is to say that the public commentary accompanying the situations in which users encountered name collisions was neither positive nor negative towards controlled interruption. In only one instance (6%) did the language convey anger—which was towards both ICANN and Google, the registry for the TLD in question.

6. Leaked Suffix Identification

The queries in DNSDB-CI provide a look into the quantity and nature of controlled interruption queries being issued. This is enlightening because it corresponds to DNS queries leaked—whether intentionally or unintentionally—to the public DNS. These are queries which, prior to controlled interruption for the given TLD, would have resulted in an NXDOMAIN response from the root servers. Finding a meaningful way to systematically measure these queries is the next important step in our analysis.

Typical metrics for quantifying the DNS query activity associated with a given TLD include query count, IP address distribution, ASN distribution, second-level domain (SLD) distribution, and query name (qname) distribution. Unfortunately, of all these metrics, only one is feasible *and* useful: the query count—both per-qname and per-TLD. While IP address and origin ASN *would* be useful, neither is available with DNSDB. This is because DNSDB only provides a mapping of domain name to a resource and a query count associated with each mapping—no query source information. The diversity of SLDs and query names is only an effective measure inasmuch as there is additional context to understand how to categorize those SLD and qnames. For example, consider the qnames `foo1.bar.baz.com` and `foo2.bar.baz.com`. These are certainly distinct qnames and can be counted as such. But when considering the organizational diversity of these names, the question might be asked: do they originate from the same organization? This is difficult to know with only the qnames themselves, but if we had additional contextual data indicating that the DNS suffix (i.e., the right-most set of labels) `bar.baz.com` is common for a given organization, then that increases confidence that they do in fact originate from the same organization. Similarly, qnames `foo.bar1.baz.com` and `foo.bar2.baz.com` are clearly from the same SLD, but there is insufficient data in the names themselves to assert that they are from the same organization. For example the domains `state.ut.us` and `k12.ut.us` are delegated to two different entities, even if they have a common SLD.

Rather than using qnames or SLDs, we identify *DNS suffixes* to apply our query metrics (see [Section 2.1](#)). This allows us to more effectively measure the nature and diversity of DNS queries because each query can be associated with a given network configuration setting that would be expected to be applied consistently to systems in the administering organization.

Our analysis applies three heuristic techniques to identify these DNS suffixes, given a set of queries: Chrome NXDOMAIN probing, WPAD lookups, and ISATAP preferred router lookups. In all three cases, we use the DNSDB-CI data set to provide the queries.

6.1. Suffix Identification via Chrome NXDOMAIN Probing

The first method of DNS suffix identification involves inferring Chrome NXDOMAIN probing from DNS queries observed in the DNSDB-CI data set. Any such activity would indicate Chrome browser usage, suggesting it originated from end-user application usage. Additionally it would identify the DNS suffix in use by the respective systems and users.

We note that queries associated with Chrome NXDOMAIN probing would not normally be found with DNSDB queries because, by definition, there is no mapping associated with NXDOMAIN responses. However, during the controlled interruption period for a TLD, *all* queries for qnames under that TLD result in an answer. Such is the case with the DNSDB-CI data set.

We now explain the procedure we employed to identify NXDOMAIN probing behavior. Chrome sends three DNS queries, all with the same DNS suffix, each with a randomly-generated first label, and all in rapid succession. Therefore, we look for DNS mappings (i.e., associated with DNS queries) exhibiting that pattern. We use DNSDB's "first seen" timestamp to group mappings first observed at a given timestamp. We then considered all mappings observed at each timestamp, according to the following criteria:

- **First label.** Only mappings for which the first label of the domain name had a length of between 7 and 15 characters consisting of all alphabet letters were considered.
- **Query type.** Only mappings for which the query type was `A` were considered.
- **Qname observed only once.** Because the first label of the qnames related to Chrome NXDOMAIN probing are randomly generated, it is probabilistically unlikely—though not impossible—that the same qname would be observed more than once in a mapping. Thus, we only considered mappings for which the "first seen" timestamp equals the "last seen" timestamp, i.e., it was only observed once.

At this point, we grouped the mappings observed within a timestamp by common suffix of the qname—defined as everything to the right of the first (i.e., left-most) label. We then applied the following additional criteria:

- **Qnames with common suffix found in groups of three.** Only suffixes found in groups of three were considered, i.e., corresponding to the number of probing queries issued by Chrome.
- **Qname group only seen once.** Only groups of qnames observed exactly once were considered because of the improbability of observing two groups of randomly-generated qnames that were exactly the same.

The list that resulted consisted of the suffixes (i.e., everything after the first label) for every qname group that met the criteria above.

As an example, suppose the following queries were observed:

First seen	Last seen	Query (qname/type)	Reason for Disqualification
1649687014	1649687014	sujenbfd.foo.example.com/A	
1649687014	1649687014	pwfiksd.foo.example.com/A	
1649687014	1649687014	nmzuhes.foo.example.com/A	
1649687014	1649687017	lkaubqq.foo.example.com/A	More than 1 second
1649687020	1649687020	polkuhadev.bar.example.com/ A	Group of 2 qnames
1649687020	1649687020	fvqiyjas.bar.example.com/A	Group of 2 qnames
1649687020	1649687020	hnsjmirc.baz.example.com/A	Group of 1 qname

This query data would result in the following DNS suffix: `foo.example.com`. Other potential DNS suffixes above (e.g., `bar.example.com`, `baz.example.com`, `example.com`) are not part of the resulting set because they do not meet all of the aforementioned criteria.

An analysis of the Chrome identification methodology is found in [Section 6.4](#).

Even with the measures we took, there still might be room for false positives. In [Section 8](#), we further filter the suffixes to increase confidence in the data set used for our later analysis.

6.2. Suffix Identification Using WPAD and ISATAP DNS Queries

To identify suffixes using DNS queries related to WPAD and ISATAP, we identified all qnames with first label was “wpad” or “isatap”, respectively. The suffix list was built by extracting the suffix (i.e., everything after the first label) from every qname beginning with “wpad” or “isatap.”

We validate our methodology related to DNS suffix identification in [Section 6.5](#).

6.3. Results

6.3.1. Validation of Identification Methods

The total number of DNS suffixes identified in the DNSDB-CI data set was 2,762. The following table shows the counts and percentages of DNS suffixes identified using different combinations of the methods:

Identification Method(s)	Suffixes Identified		
	Count	Subcategory %	Total %

Chrome, WPAD, or ISATAP - Any	2,762	100%	100%
Chrome, WPAD, and ISATAP - All	1,064	39%	39%
Chrome	1429	52%	52%
Chrome only	197	14%	7%
Chrome and WPAD or ISATAP	1,232	86%	45%
WPAD	2,084	75%	75%
WPAD only	360	17%	13%
WPAD and ISATAP or Chrome	1,724	83%	62%
ISATAP	2,065	75%	75%
ISATAP only	453	22%	16%
ISATAP and Chrome or WPAD	1,612	78%	58%

Each percentage in the “Subcategory %” column is taken from the “Count” in the “parent” category or subcategory (i.e., the bolded count most immediately above). The percentages in the “Total %” column are taken from the “Count” in the “Chrome, WPAD or ISATAP - Any” category.

Each method resulted in the identification of between 52% (Chrome) and 75% (WPAD and ISATAP) of all 2,762 suffixes. These percentages show that each identification method contributed to the set of DNS suffixes. To further validate the suffixes identified with each method, we further analyze the contributions of each subsequently.

The subcategories whose label includes “and” (e.g., “Chrome and WPAD or ISATAP”) show how many of the suffixes identified by *one* method (e.g., Chrome) were identified by at least one *other* method (e.g., WPAD or ISATAP). Higher values indicate more confidence in the method, i.e., because multiple applications were used in the environment exposing this DNS suffix. For all three methods, the percentage of suffixes identified by at least one other method was at least 45%.

The subcategories labeled “only” (e.g., “WPAD only”) identify the *individual contributions* of each method—that is, how many of the suffixes were identified *only* because the listed method was employed. Larger numbers are a possible indicator that the suffix identification method was inaccurate, finding many suffixes that were not found by any other methodology. However, we also would not expect a zero value because of the diversity of application deployment within network environments. In every case, these figures are under 20% of the total. The ISATAP methodology was the single largest contributor, from which 16% of the suffixes were identified. The Chrome NXDOMAIN probing had the lowest individual contribution, yet without it, 7% of DNS suffixes would not have been identified.

6.3.2. Distribution of Suffixes Across TLDs

While at least one suffix was found in 498 (56%) of the 885 new delegated TLDs, the distribution of suffixes across TLDs was such that most of the suffixes were concentrated within a relative few. The following table shows a per-TLD statistical breakdown of the suffixes, both overall and by individual identification method:

	Number of Suffixes per TLD			
	Median	90th percentile	99th percentile	Max
WPAD	0	3	37	223
ISATAP	0	3	40	240
Chrome	0	2	27	145
Combined	1	3	52	297

Thus, half of TLDs were associated with at most one suffix, and fewer than 10% of TLDs were associated with more than three suffixes. Particularly interesting is the disproportionately high number of DNS suffixes identified in newly delegated TLDs and their inclusion in reports submitted via ICANN's Web form. The following table lists each reported TLD, in order of rank, along with the numbers of DNS suffixes identified in each. Only the 17 TLDs delegated after controlled interruption (August 2014) are included, as they are the only ones for which we have suffix data from the DNSDB-CI data set *because* of controlled interruption. Numbers that are underlined indicate a value above the 90th percentile.

TLD	ICANN Reports	DNS Suffixes Identified Using Method			Total DNS Suffixes Identified
		Chrome	WPAD	ISATAP	
network*	7	<u>60</u>	<u>86</u>	<u>115</u>	<u>134</u>
ads*	4	<u>139</u>	<u>233</u>	<u>234</u>	<u>247</u>
prod*	4	<u>32</u>	<u>64</u>	<u>66</u>	<u>71</u>
dev*	3	<u>62</u>	<u>100</u>	<u>98</u>	<u>113</u>
cloud*	2	<u>10</u>	<u>14</u>	<u>12</u>	<u>14</u>
google**	2	1	<u>6</u>	3	3

school*	2	<u>29</u>	<u>37</u>	<u>40</u>	<u>47</u>
anz	1	0	2	0	2
app*	1	<u>3</u>	3	<u>5</u>	<u>6</u>
cpa*	1	2	<u>6</u>	3	<u>4</u>
csc	1	2	2	2	3
goo	1	0	1	1	1
off*	1	<u>7</u>	<u>15</u>	<u>14</u>	<u>14</u>
office*	1	<u>145</u>	<u>216</u>	<u>240</u>	<u>264</u>
orange*	1	<u>3</u>	<u>5</u>	<u>4</u>	<u>5</u>
site*	1	<u>18</u>	<u>23</u>	<u>33</u>	<u>50</u>
tech*	1	<u>18</u>	<u>25</u>	<u>30</u>	<u>33</u>

* All DNS suffix counts were in the 90th percentile.

** At least one DNS suffix count was in the 90th percentile—but not all counts were.

At least one DNS suffix was identified for every TLD for which problems were reported, and all reported TLDs except one (goo) had suffix counts greater than the median. In 13 (76%) of the 17 TLDs for which problems were reported, the number of DNS suffixes were in the 90th percentile. In only 3 (18%) of the 17 TLDs for which reports were submitted were all suffix counts below the 90th percentile. Further, the 4 (24%) TLDs with the most reports (i.e., the four highest ranking) had suffix counts within 99th percentile.

The trends here are clear. There are disproportionately high counts of DNS suffixes amongst the 17 reported TLDs, with 76% having DNS suffix counts in the 90th percentile. The trend clearly suggests that reports for a given TLD are more prevalent where the DNS suffix count is higher.

6.4. Analysis of Chrome Identification Methodology

We previously identified rules for detecting DNS suffixes by recognizing qnames associated with Chrome browser NXDOMAIN probing behavior (see [Section 6.1](#)). Two of the criteria for considering potential suffixes using the Chrome method were that they showed up in groups of three at a given timestamp and that the exact group of three qnames does not show up at any other timestamp. We now mention some statistics with regard to those which were “rejected” from candidacy because of failure to meet that criteria. A total of 21,768 *potential* suffixes were identified before considering the number of mappings with a given suffix at a given timestamp and uniqueness of groups of qnames. Of those 20,336 (93%) were eliminated because they

were not in groups of three, and an additional 3 were eliminated because the same qnames were found at a different timestamp. This is a fairly high percentage, and we suspect that some of these are false negatives. However, the intent was to reduce false *positives*. As mentioned previously, 86% of the DNS suffixes identified with Chrome browser identification were also identified by either the ISATAP or the WPAD methodology, and only 14% were found exclusively using the Chrome technique. This percentage is comparable to those of the WPAD and ISATAP methods, which were 17% and 22%, respectively. These numbers provide confidence in the methodology. While a more rigorous validation of this and other methods is possible, it is beyond the scope of the work.

6.5. Validation of WPAD Identification Methodology

As mentioned previously ([Section 6.2](#)), there was some question about false positives produced when using the WPAD identification methodology. Specifically, there was some concern that “ancestor” names of a legitimate DNS suffix might be falsely identified as DNS suffixes because of the iteration performed by WPAD clients. We evaluated our results to look for evidence of such behaviors.

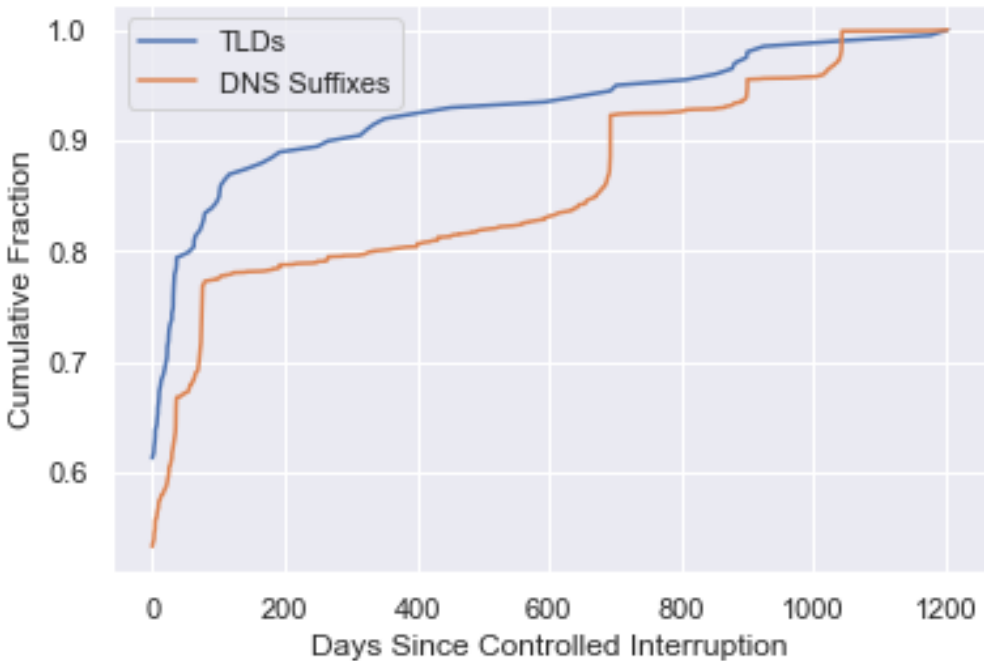
Of the DNS suffixes using the WPAD identification methodology, 1,728 suffixes were composed of two or more labels. For only those cases, only 153 (8.9%) was the “parent” DNS name also identified as a suffix using the WPAD methodology. In 91 (59%) of those cases, the parent name was identified independently as a DNS suffix using one of the other methodologies. Thus, in only 62 (3.6%) of cases was a parent name identified *exclusively* by our WPAD methodology as a DNS suffix. It is possible that every one of these “parent” suffixes is a legitimate DNS suffix, but even if not, the low percentage shows that this is not a pervasive behavior.

7. Controlled Interruption Analysis

We use the DNSDB-PostCI data to learn more about the use of controlled interruption and the use of the observed DNS suffixes identified as being in conflict with new TLDs being delegated. By considering only DNS suffixes that had two or more labels (see also [Section 8](#)), we reduced the number of DNS suffixes to 2,300, within 200 TLDs—instead of the full set of 2,762 suffixes within 498 TLDs. With this reduced data set we looked at the mappings observed since the first 100 days of delegation for each DNS suffix. Note that this filtered set of DNS suffixes included 16 (94%) of the 17 TLDs reported to ICANN; only the `goo` TLD (associated with a single ICANN report) was excluded.


As mentioned previously ([Section 2.2](#)), the IP address 127.0.53.53 is returned for all names under a TLD during the first 100 days of its delegation, i.e., the controlled interruption period. By analyzing the mappings in DNSDB-PostCI, we were able to determine how long controlled interruption was observed for each TLD and at what point non-controlled interruption addresses (i.e., other than 127.0.53.53) were observed in relation to the controlled interruption period.

The following plot shows the cumulative distribution of the number of days after the controlled interruption period for which the controlled interruption address was observed—on a per-TLD basis and a per-suffix basis:



For about 53% of DNS suffixes and 62% of TLDs, the controlled interruption address was not observed after the controlled interruption period, i.e., the first 100 days of delegation. However, the controlled interruption IP address was observed for a year or more after the controlled operation period for about 10% of TLDs and for 20% of DNS suffixes.

While a glimpse of how long controlled interruption was maintained beyond the prescribed period, perhaps more interesting and useful is an understanding of how soon after the controlled interruption period non-controlled interruption addresses were introduced for suffixes known to be used in conjunction with private DNS namespaces. The following plot shows the cumulative distribution of days since controlled interruption representing those mappings:



For about 72% of TLDs and 80% of DNS suffixes, no mappings were observed for known DNS suffixes. However, for the remaining 28% and 20% of TLDs and suffixes, respectively, non-controlled interruption mappings were observed at some point after the controlled interruption period ended. In both cases, those mappings were observed immediately after; for 10% of suffixes and 20% of TLDs mappings were observed within 500 days (about 16 months).

The presence of non-controlled interruption does not pose an immediate threat in and of itself; it all depends on the existence of a mapping for a qname within a DNS suffix and, of course, the nature of the application or service relying on the resolution. However, it does indicate the *potential* for third-party interception of traffic, whether intentionally or inadvertently. While we have not carried out a general search of qname mappings, we did search for two prominent qname patterns, which, if present, could have a significant impact on systems relying on the non-resolution of certain DNS qnames used for private use: `wpad` and `isatap` (see [Section 2.4](#) and [Section 2.5](#)). Fortunately, we found no mappings for such qnames in the DNSDB-Post-CI data.

8. Root Server Query Analysis

The DNS suffixes identified in [Section 6](#) provide a unit of measurement for quantifying the usage of newly-delegated TLDs, prior to and after their delegation, and to identify organizations from which their associated queries originated. In this section we describe our measurement methodology.

8.1. Data Set

We used the DITL data from 2014 through 2021 (see [Section 3.3](#)) to observe queries at the root servers related to the DNS suffixes associated with leaked DNS queries, i.e., those identified previously. Extracting query information from the DNS root servers requires resources related to both computation and storage. For this reason, we reduced the computational resources required by limiting the suffixes against which we compared DITL queries in two ways.

Eliminate TLDs. First, we reduced the suffixes by eliminating those that were themselves TLDs. For example, `office` is a TLD, but it was *also* identified as a DNS suffix through one or more of the identification methods. Thus, DNS queries associated with the suffix `office` because it was a TLD. The rationale behind excluding TLDs was two-fold. First, by including a TLD, our filter would include *all* queries ending with that TLD. Many of those queries would be false positives, and we have no way to reliably exclude false positives from the data set when the suffix is a TLD. Additionally, as mentioned previously, one of the objectives of this analysis is to identify organizations from which the DNS suffix originated, as part of root cause, by using the suffix itself. For example, the suffix `acme.network` originating from a network with name “ACME” would support an association between the network and the DNS suffix. However, a single label is typically too generic to help us associate suffixes to organizations in that way.

Further Filtered TLDs. Second, we further limited our analysis to suffixes with TLDs meeting one or more of the following criteria:

- The number of DNS suffixes identified from ISATAP-related queries was at least one;
- The number of DNS suffixes identified from WPAD-related queries was at least one; or
- The number of total DNS suffixes identified as at least two.

This effectively eliminated DNS suffixes for TLDs that were *only* part of the data set because of a single suffix identified with our Chrome NXDOMAIN probing technique. While all three of our suffix identification techniques were merely heuristics, Chrome NXDOMAIN probing was the most susceptible to false positives. This filter eliminated some of the weaker contributors in the data set.

	Suffixes	TLDs
All DNS Suffixes	2,762	498
DNS Suffixes - no TLDs	2,300	200
DNS Suffixes - no TLDs and further filtered (TLD has at least one WPAD suffix, one ISATAP suffix, or more than 1 total suffix)	2,266	166


Note that this filtered set of DNS suffixes included 16 (94%) of the 17 TLDs that were the subject of reports submitted to ICANN via their Web submission form (see [Section 4](#)). The only TLD that was excluded was `goo`.

Having our updated DNS suffix list in hand, we utilized a two-step process to actually extract the DNS queries from the DITL: 1) we filtered all DITL queries, keeping only those with a query name under one of the newly-delegated TLDs; then 2) we tested each of the resulting queries to see if the query name was under one of the 2,266 DNS suffixes we identified previously.

8.2. Results

We first consider the number of DNS suffixes observed in root queries during each DITL collection period between 2014 and 2021. The following plot shows: 1) the total number of DNS suffixes for which their TLD was delegated during the time of the DITL collection for the corresponding year (i.e., all 2,266 were delegated by the time of the 2021 DITL collection); 2) the total number of DNS suffixes for which DNS queries were observed at the root servers, out of the 2,266 total suffixes; 3) the subset of observed DNS suffixes that were the subject of ICANN reports (see [Section 4](#)); and 4) The number of DNS suffixes for which DNS queries were observed and for which non-CI mappings (i.e., other than 127.0.53.53) were identified after the CI period for the respective TLD (i.e., after the first 100 days).

DNS Suffixes Observed in DNS Queries to Root Servers




While over 1,900 (84%) of the 2,266 DNS suffixes were observed as early as 2014, the number of suffixes observed in DNS queries has consistently decreased over time, as new TLDs have been delegated, such that in 2021 1,434 (63%) suffixes were observed. Nearly half of those DNS suffixes are associated with the reported TLDs, specifically between a low of 43% (2018) and a high of 51% (2014). This disproportionately high contribution of observed DNS suffixes again emphasizes the significance of the name collisions reports submitted to ICANN.

We note that *all* of these suffixes were observed during the controlled interruption period for their respective TLDs and have thus been associated with leakage of “private” DNS queries colliding with public DNS namespace. However, we cannot know from these query observations alone whether the queries at the root were associated with previous, private use of the TLD (i.e., prior to its delegation) or use of the TLD in connection with its delegation. The latter is certainly the case in 2014 because none of the new TLDs or their suffixes were delegated by the time of the 2014 DITL collection, but for 2015 and beyond, it is not known. See [Section 6](#) for more.

Between 2015 and 2021, there is a steadily increasing number of DNS suffixes observed in query data for which non-CI mappings exist (see [Section 7](#)). In 2021, queries were observed for 336 suffixes that had a non-CI mapping. That accounts for 23% of all DNS suffixes observed in queries at the DNS root and 15% of all 2,266 DNS suffixes. As mentioned, it is difficult to tell with current data whether the queries associated with these suffixes were in connection with private use or not, but it does raise some concerns.

We now consider the same data, but with respect to TLD. The following plot shows: 1) the total number of TLDs delegated during the time of the DITL collection for the corresponding year (i.e., a total of 885 delegated TLDs by the time of the 2021 DITL collection); 2) the total number of *filtered* TLDs delegated at the time of DITL data collection (i.e., a total of 166 TLDs by the time of the 2021 DITL collection); and 3) the total number of TLDs having DNS suffixes for which DNS queries were observed at the root servers, out of the 166 filtered TLDs. In other words, this plot shows the number of TLDs experiencing some sort of name collision behavior over time.

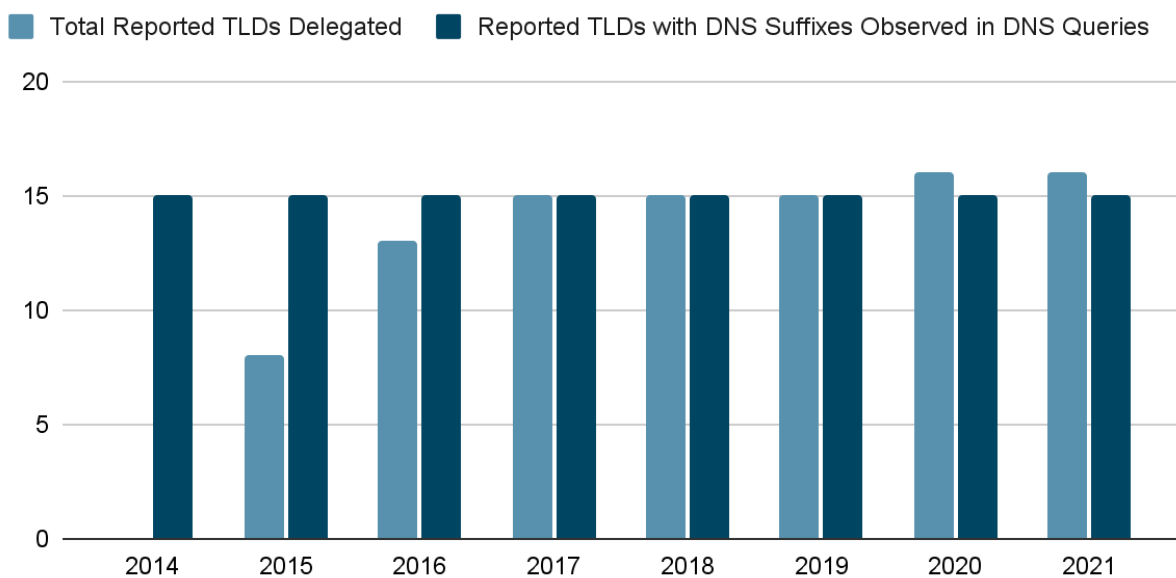
TLDs with DNS Suffixes Observed in DNS Queries to Root Servers



The number of TLDs experiencing name collisions, by observation, has remained relatively steady from 2014, when queries for DNS suffixes associated with 146 TLDs (88% of filtered, 16% of all TLDs) were observed, through 2021, when 133 TLDs exhibited name collision behavior (80% of filtered, 15% of all TLDs). The peak was in 2016 when 154 TLDs (93% of filtered, 17% of all TLDs) exhibited name collision behavior.

When we consider only the 16 TLDs that were the subject of reports and part of the filtered set of DNS suffixes, the following plot is the result:


Reported TLDs with DNS Suffixes Observed in DNS Queries to Root Servers




This shows that in every DITL collection between 2014 and 2021, queries for DNS suffixes within 15 (94%) of the 16 reported TLDs, after filtering, were consistently observed. Only the TLD `google` was not observed. While the general trend was mostly consistent, this trend was completely consistent.

We now consider several other metrics to help us quantify name collision behavior between 2014 and 2021. Specifically, for DNS suffixes experiencing queries each year, we consider the number of queries, unique qnames, querying IP addresses, and origin ASes of queries. The median and 75th percentile values are shown in the following two plots:

Median Counts for DNS Suffixes for which Queries were Observed




75th Percentile Counts for DNS Suffixes for which Queries were Observed



In the next figure, we show overall counts for DNS queries associated with identified DNS suffixes, as a fraction of those observed in 2014:

Total Counts for Observed Queries Associated with DNS Suffixes



The plot is normalized because of the significant difference in scale between the different categories. For reference, the following table shows the raw counts:

Year	Queries	Qnames	IP Addresses	ASNs
2014	62,305,672	25,937,776	112,374	12,296
2015	21,358,020	6,504,348	98,555	10,287
2016	16,061,683	6,349,761	97,640	10,356
2017	4,586,613	754,204	75,294	10,050
2018	4,126,353	729,336	73,658	8,854
2019	1,846,412	268,356	77,951	9,469
2020	5,855,426	695,784	88,393	8,944
2021	3,636,318	531,233	66,304	6,472

In all plots, a clear trend of decreasing per-suffix and overall usage metrics is evident. However, the cause of this trend is unknown. One possible cause might be actual administrative changes eliminating the use of those suffixes in configurations, possibly because of the effects of controlled interruption. However, it could also be due to reduced DNS query data at the root servers associated with local root deployments⁹ or qname minimization¹⁰, which we explain in the following paragraphs.

⁹ [RFC 8806](#)

¹⁰ [RFC 7816](#)

The local root specification was first published in November 2015 and updated in June 2020. It provides guidance for serving a copy of the root zone on a recursive resolver. This keeps the resolver from having to issue any queries to the root servers because it has all the answers it needs locally. It thus achieves benefits of both privacy and performance. There are currently no research studies to provide insight into the prevalence of local root deployment. However, the publication date of the original specification for local root deployments was *after* the prominent decrease in per-suffix and total counts related to name collisions, which was first observed in April 2015.

With qname minimization, a recursive resolver only reveals the necessary parts of the name it is attempting to resolve in the queries it issues to authoritative DNS servers. For example, when a resolver is resolving `www.example.com`, it might have historically sent the entire name, `www.example.com`, to a root server. However, qname-minimizing resolvers take advantage of the fact that the only *required* component is `com`, i.e., to elicit a referral. They use various techniques to conceal more specific query information from authoritative servers. Recent studies suggest that qname minimization affects 12% of Internet resolvers and 40–48% of queries as of 2018¹¹. We consider the effects of qname minimization in [Section 8.3](#).

To gain additional insight into the causes of the query behaviors we observed, we supplement our quantitative measurements with a qualitative study, which we discuss in [Section 9](#).

8.3. Qname Minimization Considerations

The data that has been presented thus far has been compiled independent of qname minimization. However, because qname minimization has seen an increase in deployment, and its effects might contribute to some of the downward trends in our analysis, we now perform additional analysis that takes qname minimization into account.


8.3.1. Summary of Recent Study of Qname Minimization

We first summarize recent work analyzing the deployment of qname minimization on resolvers that queried A-root during between the years of 2008 and 2021, using the yearly DITL collection as its data source¹². In that work, the process for determining qname minimization behavior was as follows. A resolver was evaluated for qname minimization by testing for the following two query behaviors during the collection period: 1) the resolver issued a minimum of five queries for qnames other than the root name; and 2) the resolver issued no query with a qname having more than one label. If a resolver met both requirements, then it was considered to be qname-minimizing. If it met only the first, then it was considered to be non-qname-minimizing. If it met neither requirement, then no assessment could be made.

¹¹ <https://www.nlnetlabs.nl/downloads/publications/devries2019.pdf>

¹² “Fourteen Years in the Life: A Root Server’s Perspective on DNS Resolver Security” by Alden Hilton, Casey Deccio, and Jacob Davis. To appear in *Proceedings of USENIX Security ’23*.

We include below the plot from that work that shows the percentage of IP addresses (of the subset that could be evaluated, based on the five-query minimum) that exhibited qname-minimizing behavior:



Also shown are the percentage of ASes for which at least one qname-minimizing resolver was observed and the percentage of queries corresponding to the qname-minimizing resolvers. The labeled vertical lines represent (A) the submission of the initial qname minimization Internet Draft, (B) its adoption by the unbound resolver, (C) its adoption by Knot resolver and its publication as an RFC, and (D) its adoption by BIND resolver. From the vantage point of A-root, the percentage of resolvers that use qname minimization has increased from 1% to 12% between 2018 and 2021. The percentage of overall queries that come from qname-minimizing resolvers has risen from 1% to 14% between 2014 and 2021, with it reaching as high as 27% in 2019.

Notably, the upward trend in deployment of qname minimization does not correlate with the downward trend associated with the name collision queries observed at the root servers. While significant uptick of qname-minimizing resolvers did not occur until 2019 with its inclusion in BIND, the significant decrease in per-suffix name collision queries occurred in 2015, which was the first DITL collection after controlled interruption was instituted.

8.3.2. Application of Qname Minimization Data

We next sought to isolate the resolvers identified as non-qname-minimizing and run our analysis again on *only* those, so we could compare the trends observed in this latest analysis with those resulting from the analysis that did not consider qname minimization (i.e., from [Section 7.2](#)). However, there were three challenges with this. *First*, the IP addresses observed at A-root constituted only 40% of all IP addresses seen at the collective root servers (except I-root and L-root, which anonymize their IP address data) during the 2021 DITL collection. Even so, these IP addresses represented 95% of ASes from which queries were received by the collective root servers. *Second*, only 36% of the IP addresses querying A-root met the criteria for qname minimization evaluation in 2021, corresponding to 15% of the total IP addresses observed in 2021. Of those, 88% of IP addresses exhibited behavior characteristic of non-qname-minimizing resolvers. Thus, the percentage of 2021 IP addresses that are used for

our analysis is 13%. *Finally*, the set of IP addresses observed in DITL collections prior to 2021 is not the same as the set observed in 2021; various factors over time contributed to the variance between those sets.


We used the IP addresses of the non-qname-minimizing resolvers identified in the 2021 DITL collection as the basis for carrying out the analysis in the previous years. We did this under the assumption that if a resolver was not using qname minimization in 2021, then it was likely not using qname minimization before 2021. This assumption greatly simplified the data set we were working with and its analysis. A summary of the numbers of IP addresses comprising the analysis for each year since 2018 is found in the following table. In each case, the percentage reflects the percentage of all IP addresses observed in the given DITL collection year:

DITL Year	IP Addresses (all root servers)	IP Addresses (only A-root 2021)	Qname Min. Evaluated	Non-Qname Min.
2018	17,017,222	7,047,980 (41%)	1,205,290 (7%)	1,121,513 (7%)
2019	12,651,567	7,071,314 (56%)	1,511,110 (12%)	1,395,088 (11%)
2020	17,343,285	8,718,048 (50%)	2,089,481 (12%)	1,893,877 (11%)
2021	26,463,953	10,612,429 (50%)	3,845,577 (15%)	3,380,341 (13%)

Thus, the sample of data from which we take our analysis ranges from 7% (2018) to 13% (2021). Sample data prior to 2018 is not currently available.

The median per-suffix counts for queries, unique qnames, IP addresses, and ASNs is shown in the following figure:

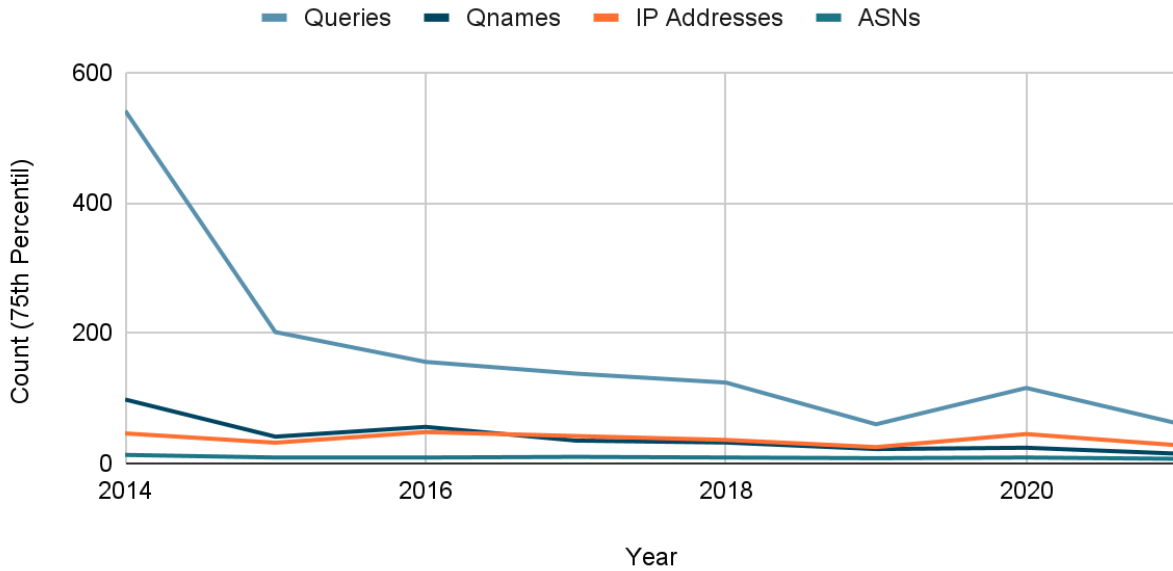
Median Counts for DNS Suffixes for which Queries were Observed



The raw numbers are, expectedly, much lower each year in this plot than they are in the previous plot, which considers the entire set of querying IP addresses; this is due to the very fact that we are working with only a subset of the data. However, the trends in this plot match those in the previous plot, especially in the following ways: 1) both plots show a significant decrease in median counts between 2014 and 2015; both plots show relatively little change between 2015 and 2021; and 3) both plots show a slight increase in median counts in 2020.


The trends associated with the per-suffix 75th percentile counts for non-qname-minimizing resolvers also match those of the plots that consider all resolvers:

75th Percentile Counts for DNS Suffixes for which Queries were Observed



Finally, we consider the total counts associated with name collision queries, across all DNS suffixes, each shown as a percentage of the respective 2014 value:

Total Counts for Observed Queries Associated with DNS Suffixes




The corresponding raw numbers are shown in the following table:

Year	Queries	Qnames	IP Addresses	ASNs
2014	5,439,148	1,315,878	172,559	45,634
2015	5,330,047	1,344,184	181,541	36,869
2016	7,639,202	3,423,237	210,865	32,692
2017	2,650,592	491,464	193,517	33,514
2018	2,807,826	569,824	188,971	31,571
2019	1,263,280	181,099	146,644	32,020
2020	5,314,238	601,583	219,667	38,490
2021	3,562,760	526,523	187,830	28,524

In this case, a consistent trend is hard to observe within the data itself, and it differs significantly from its counterpart, which includes queries from all IP addresses, rather than just non-qname-minimizing IP addresses. Because the median is so consistent, these relatively high and inconsistent counts are likely related to outliers—suffixes that receive *many* more queries than the median or 75th percentile, within the non-qname-minimizing IP addresses. To test this, we plot the maximum count values across all DNS suffixes, for each of the years:

Maximum Counts for DNS Suffixes for which Queries were Observed



There are features in the plot that clearly demonstrate outlier behavior, the most prominent of which is the relatively high number of queries and unique qnames queried for name collisions in 2016. The cause of these outliers requires further research, but it is outside the scope of this

work. It is sufficient to indicate that outlier behavior is at play with the plot showing the total counts associated with name collisions queries.

Based on the analysis presented herein, we conclude that the trends related to name collision DNS queries observed at the root servers from DITL collection data are not affected by qname minimization behaviors.

9. Name Collisions Survey

To better understand the metrics we presented in the previous section, we conducted a survey to solicit experiences related to name collisions. The survey was given to two different target audiences: a general audience of network operators and a targeted audience consisting of organizations presumably affected by name collisions related to the delegation of new TLDs.

9.1. Survey Content

The questions were common to both surveys, with some slight variants in wording. They solicited the following information:

- What DNS suffixes under newly delegated TLDs are in use by organizations.
- Which newly delegated TLDs are associated with DNS suffixes in use.
- What DNS configuration is being used in the organization in connection with suffix use.
- Whether or not problems were experienced with the use of the DNS suffixes since the delegation of the TLD.
- What the effects of suffixes were, in terms of time to detection, number of users affected, and time to resolution.
- What was the role of controlled interruption IP address (127.0.53.53) in diagnosing the problem.

The complete set of survey questions for the general and targeted audiences are found in Appendixes [C](#) and [D](#), respectively.

9.2. Survey Recipients

The general version of the survey was sent to the North American Network Operators Group (NANOG) mailing list on March 29, 2022, with a reminder email sent on April 4, 2022. The text of the message is in [Appendix E](#).

The recipients for the targeted version of the survey consisted of network administrators for which the autonomous system (AS) description matched DNS suffixes corresponding to queries originating from that AS number (ASN). We created this list using the following methodology:

- Create suffix-ASN mappings from queries observed at root servers, based on DITL data (see [Section 8](#)).

- Filter suffix-ASN mappings to include only suffixes for which at least 10 unique qnames (implies at least 10 queries) were observed for the suffix for any collection year. This filter was used to establish additional confidence in the sample set of suffixes that would be used for targeted reach-out.
- Further filter suffix-ASN mappings to include only ASNs that included a single suffix. This filter is applied to exclude ASNs that likely provide a DNS resolver service for other organizations.
- For each suffix-ASN mapping, perform a WHOIS lookup of the ASN, and compare the organization information provided by WHOIS with the DNS suffix itself (typically the left-most label). Include only mappings for which a positive match was made.

This process resulted in a list of 28 mappings in 18 TLDs for which we could associate ASN technical contact information. These included 7 (44%) of the set of 16 reported TLDs (after filtering). However, there was no selection bias based directly on report TLDs; we selected all mappings from the sample for which we were able to positively identify a match between DNS suffix and ASN.

The targeted messages sent to ASN contacts contained not only a link to the survey, but also the DNS suffix associated with the mapping—that is, the one for which DNS queries were observed as having originated from the ASN. The text of the message is in [Appendix F](#).

One known limitation of our methodology is that the mappings consist of DNS suffixes that match the ASN descriptions; however, one of the observations made in [Section 4](#) is that a significant contributor to name collisions is systems querying the public DNS from *outside* their corporate network (in which the DNS resolvers might be configured to answer authoritatively). Thus, the targeted survey results have some bias related to the symptoms and possibly the network configuration causing the issues. The targeted surveys might also represent a community with a private query leakage caused by something different than the remote user/VPN configuration noticed in [Section 4](#). However, as will be noted, this bias has little impact on our findings because the response rate was so low.

9.3. General Survey Results

The survey sent to the NANOG mailing list generated 31 responses. Of those 31, 21 (68%) indicated that their organization did not employ any DNS suffixes that were associated with newly delegated TLDs. We focus the remainder of this analysis on the 10 (32%) respondents that indicated that they *did* use DNS suffixes under new TLDs.

9.3.1. TLDs Used

The following tables lists the TLDs associated with survey responses, representing DNS suffixes in use by organizations:

Delegated Before Controlled Interruption	Delegated After Controlled Interruption	Not Delegated
--	---	---------------

audio	dev*	corp
foo	group	example
media	llc	internal
pro	network*	test
	office*	
	tech*	

* Included in name collisions reports submitted to ICANN.

Most pertinent to this root cause analysis are the TLDs in the middle column, which represent the TLDs that have been delegated since controlled interruption (i.e., since August 2014). Four of those (marked with *) were also the subject of reports submitted to ICANN via their Web form.

9.3.2. Technical Issues Experienced

Of the 10 reports in which DNS suffix use was indicated, 7 (70%) reported experiencing technical problems after delegation of the TLDs. We focus our analysis on just those 7 reports for the remainder of this section.

9.3.2.1. DNS Resolver Configuration

In three (43%) of the cases experiencing technical issues, the response indicated that the organization's configuration was such that the DNS resolvers were configured to answer authoritatively for the DNS suffixes in question; in two (29%) cases, that was *not* the configuration. Two respondents did not know details related to this configuration. There seems to be no strong correlation between the DNS resolver configuration and the presence of technical issues with the DNS suffix. Across the 10 responses confirming use of DNS suffixes within newly delegated TLDs and the 1 response confirming use from the targeted survey (**), we saw the following combinations:

DNS Resolver Authoritative	Issues Experienced	Count
No	No	2
No	Yes	2
Yes	Yes	2
Yes*	Yes	1
Yes*	No	1

Yes	No	1**
-----	----	-----

* Resolvers were changed to answer authoritatively at some point.

** Included from the targeted survey response.

9.3.2.2. Discovery, Impact, and Resolution

Three (43%) organizations discovered the problems within days of the delegation; one (14%) within weeks of the delegation; and three (43%) within months of the delegation. In terms of impact, three (43%) reported that only a few systems were affected, but two (29%) reported that many were affected, and two (29%) reported that nearly all systems were affected. Two (29%) reported that they were able to resolve the issue within days or weeks of its discovery. However, two (29%) reported that it took years to resolve, and two (29%) reported that it has not yet been resolved.

9.3.2.3. Root Cause Identification

With respect to the identifying the root cause of the problem, five (71%) respondents indicated that they knew the problems were related to the delegation of new TLDs before the problem was resolved, and two (29%) only discovered that the problems were related to delegation of new TLDs after the problem was resolved. In only one (14%) case was the controlled interruption IP address, 127.0.53.53, observed and helpful in leading the organization to ICANN and the delegation of the new TLD. One (14%) respondent reported that 127.0.53.53 was observed, but its meaning was unclear and was not helpful in identifying the problem. In the five (71%) remaining cases, 127.0.53.53 was not observed at all.

9.3.2.4. Other Observations

Some of the free-form comments received from respondents shed additional light on the experiences of those who were impacted by new TLD delegations.

One respondent indicated that their DNS resolvers were not configured as authoritative for their DNS suffix, but rather for the entire TLD (`dev`). The problems then came when `dev` was delegated. In this specific case, they reported that the problem was discovered within days of its delegation, affected “many” users or systems of an organization with fewer than 1,000 systems, and took weeks to fix. The fix involved changing the DNS suffix they were using internally (e.g., as opposed to changing the way their DNS resolvers were configured). In this case, 127.0.53.53 was not observed.

Another respondent commented:

“This was very expensive and disruptive. In addition, employees cannot reach websites in the network domain.”

This response indicated that “nearly all” systems or users were affected by the change, in an organization consisting of between 1,000 and 10,000 systems. Although the problem was discovered within days of the delegation, it reportedly took years to fix. In this case, 127.0.53.53 was observed, but its meaning was unclear or unhelpful in identifying the problem.

9.4. Targeted Survey Results

Of the 28 targeted surveys, two recipients (7%) filled out the survey. Of those, only one recipient confirmed use of the suffix provided in the email message; the other was symptomatic of false positive match between DNS suffix and ASN.

The admin that confirmed usage of the provided DNS suffix provided the following information with regard to its use:

- The suffix is associated with the `win` TLD.
- Use of the DNS suffix predated the delegation of the TLD, and the DNS suffix continues to be used by the organization.
- The organization’s DNS resolvers are configured to be authoritative for the DNS suffix, such that queries within those suffixes, when issued to their resolvers, are presumably not leaked to the public Internet.
- No known technical issues were experienced with the suffix after the delegation of its TLD.

10. Discussion

This work attempts to analyze several data sources consisting of mostly passive traffic data and couple that analysis with qualitative data from both a targeted and a general survey. We report here some of the key findings from the analysis, impact inferred from both quantitative and qualitative measurements, known and suspected limitations of this analysis, and proposed future work.

10.1. Findings

Private use of DNS suffixes is widespread. It is clear from the data that private use of DNS suffixes is not isolated. Apparently private use of DNS suffixes is exhibited within over half of newly delegated TLDs, even though a few TLDs are responsible for more usage than others. **Evidences.** Over half of the 885 TLDs delegated since August 2014 are being used as part of at least one configured DNS suffix for organizations, according to our measurements. Yet the use of DNS suffixes is not uniformly distributed across affected TLDs. Rather, 90% of TLDs are associated with three or fewer private-use DNS suffixes, but 1% have more than 52, reaching upwards of 297 (maximum).

Name collision reports are supported strongly by measured data. The TLDs appearing in name collision reports submitted to ICANN via their Web form rank disproportionately high in terms of the number of identified suffixes and DNS queries observed at the root servers. This bolsters the concerns associated with the reports and also indicates that there are likely others that experienced problems but did not submit reports. **Evidences.** About *two thirds* (66%) of reported TLDs were in the *90th percentile* of all TLDs for which DNS suffixes were identified, in terms of DNS suffix count. Additionally, TLDs associated with reports accounted for around half (between 43% and 51%) of the identified DNS suffixes that were observed in queries to the root servers, despite them comprising only 10% of the TLDs that were being watched for in the root server query data (i.e., the filtered set). Finally, while the observation rate of the *entire* filtered subset of TLDs ranged from 84% (2014) to 63% (2016), the fraction of *reported TLDs* for which DNS suffixes were observed in queries to the root servers was consistently 97%.

Usage of private DNS suffixes colliding with newly delegated TLDs has decreased over time. Various metrics related to DNS queries for DNS suffixes presumed to be used privately were measured over time and shown to be consistently decreasing since 2014. The reasons are unclear, but two considerations are 1) decreased DNS suffix usage and/or 2) reduced visibility at the root zones. **Evidences.** Both the median and 75th percentile counts of individual DNS queries, unique query names, querying IP addresses, and origin ASNs decreased sharply between 2014 and 2015, and have decreased more gradually since then. Some anecdotal data submitted by survey respondents supports the evidence of that decrease. We also reference outside studies that show some uptake of qname minimization, which reduces the query context available at root servers (see [Section 8](#)).

Controlled interruption is effective at disruption, but not at root cause identification. Controlled interruption has shown to be good at disruption, but not at helping affected users identify the cause of the problem—at least not in the way that was intended. **Evidences.** Of the survey respondents that indicated that they used of TLDs, 70% reported having experienced technical issues related to their suffix. Of those, 43% experienced the problems within days of delegation of the TLD. Over two-thirds (71%) of organizations experiencing technical problems indicated that they knew that the issues were related to TLD delegation before the problem was resolved. It appears that most of the ineffectiveness was due to the controlled interruption IP address not even being observed, which occurred in 71% of cases, according to the survey. However, when the controlled interruption IP address was observed, the success rate in identifying ICANN and controlled interruption as the cause was between 50% and 76%, according to the survey results and the Web search results analysis, respectively.

Configuring DNS resolvers as authoritative for DNS suffixes is not a panacea. DNS resolvers that respond authoritatively for private DNS suffixes do not prevent query leakage to the public DNS or name collision problems. **Evidences.** We have one confirmed account of DNS suffix usage where the queries were leaked to the public DNS: the targeted survey respondent confirmed usage of the DNS suffix, and we observed the queries within that suffix in the DITL query data. Additionally, the survey responses show no clear correlation between DNS resolvers thus configured and technical problems related to name collisions. In contrast, they

show all combinations of issues experienced and resolver authoritative configuration. Further, 8 (33%) of the 24 ICANN reports submitted by organizations explicitly mentioned remote users or VPN usage.

The impact of TLD delegation ranged from no impact to severe impact. The only data we have quantifying impact related to delegation of new TLDs is from the name collision reports and the survey responses. With the limited responses we received, it is hard to generalize impact. However, what we *can* say from the data is that: 1) there is a range of impact reported, from no impact to major impact; and 2) there was evidence of both severe and significant impact amongst affected parties. **Evidences.** On one side of the spectrum, the one targeted survey respondent that confirmed DNS suffix usage indicated no technical issues. Seven respondents of the general survey indicated that they had experienced technical issues, with one describing it as “expensive and disruptive,” impacting almost all users or systems of an organization with between 1,000 and 10,000 systems. The remaining survey responses reported impact somewhere between no impact and extensive impact, based on both number of systems affected and total number of systems. In the name collision reports, half (17 or 50%) of the reports imply severe or significant impact to the reporting entities.

The public response to controlled interruption was overall neutral. Name collisions and controlled interruption certainly impacted various individuals and organizations. Nonetheless in forums where users or administrators publicly posted questions or experiences with controlled interruption, the overall sentiment was neither positive nor negative, but neutral. **Evidences.** A sentiment analysis of the Web search results revealed that in 94% of cases, neither positive nor negative feelings were expressed towards controlled interruption. In only one case (6%) was negative sentiment expressed.

Name collisions were diverse, both in terms of the application involved and their root causes. Multiple applications were involved with name collisions, some with which users interface directly and some which are more process-driven. Name collisions were caused by the use of both private and non-private namespace. They were caused by the use of domain names that were fully-qualified and unqualified, including unqualified names with single and those with multiple labels. **Evidences.** Eight different applications were responsible for the 10 Web search results that revealed an application affected by name collisions. No single application was responsible for more than 20%, including Web browsers. While nearly two thirds (61%) of collisions identified in the Web search results were caused by the use of private use of TLD namespace, 10% involved the use of namespaces that were non-private. The Web search results also showed that name collisions were encountered in cases where a name was fully-qualified (59%), unqualified (35%), and even where a single-label was used (5.9%). Additionally, the use of unqualified domain names involved both single-level (67%) and multi-label (33%) unqualified domain names.

10.2. Proposed Future Work

This work has provided many insights into the impact of the delegation of new TLDs since 2014. However, it also leaves many unanswered questions—along with some paths to answer them. Some of the trends in the measured data are clear: private DNS suffix usage appears to be declining; and the reports submitted to ICANN are supported by the measured data. However, the amount of qualitative survey data is far from adequate. It provides enough of a picture to see that experience has varied widely, ranging from no impact to high impact. Yet it is insufficient to complement and interpret the measurement data.

To fill the knowledge gap on the experiences of organizations, we propose additional work, targeting *analysis* and *reach-out* related to the suffix-ASN mappings. The goal in both of these is to better understand how DNS suffixes are being used and to further our understanding of organizational impact with TLD delegation. In performing the manual inspection and alignment of identified DNS suffixes and ASNs for a *small* sample, we gained experience and insight into the effort that might be applied to carry out the same work, more efficiently and effectively on a large sample. The key observation is that there are a variety of different suffix-ASN mappings, which are suffix-dependent, ASN-dependent, and network configuration dependent. We provide several examples below:

1. **Even statically configured systems are mobile.** While DNS suffixes are applied by an organization to its systems, some of those systems are mobile. Evidence of mobile devices was observed in both root server queries and from name collision reports submitted to ICANN. Even when a DNS suffix can be associated with a given organization and its ASN, queries for that suffix will appear from other ASNs, as mobile systems travel. Further investigating the use of private DNS suffixes on mobile devices will not only help us better understand the configuration trends of mobile devices but might also help us more accurately determine the cause(s) of decreasing DNS queries for private-use suffixes over time.
2. **DNS queries might never leak from their origin ASN.** Because of corporate DNS configurations in which DNS resolvers answer authoritatively to queries in private namespace, the leakage associated with the configuration of one ASN might *only* appear to originate from other ASNs.
3. **Many ASNs are ISPs.** These exhibit the characteristics that 1) they are more ephemeral in terms of suffixes observed; and 2) there are potentially larger numbers of DNS suffixes mapped to ISP ASNs because of mobile systems. These can be identified by name (e.g., “comcast”, “cox”, or “sprint”), but also by keyword (e.g., “mobile”, “wireless”, “telecom”, “cable”, or “broadband”).
4. **Generic suffixes are in use.** Generic DNS suffixes like `local.site` and `modem.local` are, by their very nature, not specific to any organization. Thus, the organization which is using it in its configuration is more difficult to identify.
5. **Regional subdomain suffixes are in use.** Some organizations have deployed suffixes globally, with region-specific subdomains. For example `corp.sap`, `homeaway.live`, `hsbc`, with labels like the following prepended: `emea`, `mos`, `de`, `aus1`.

6. **Some TLDs are commonly used for Active Directory services.** This includes `school`, `ads`, `site`, `prod`, and possibly others. And some books and trainings for Microsoft Active Directory direct administrators to use a private suffix, including some of the aforementioned TLDs.

We believe that using knowledge gained in this analysis, including the findings noted above, a more automated workflow could be developed to better match DNS suffixes to their origin organization. It is our hope that this will both enrich our understanding of the use of private DNS suffixes, create more opportunity for reach-out, and ultimately better understand past and future impact of delegation of new TLDs.

Appendix A - Name Collisions Report Form

Report a name collision

A name collision occurs when an attempt to resolve a name used in a private name space (e.g. under a non-delegated Top-Level Domain, or a short, unqualified name) results in a query to the public Domain Name (Domain Name) System (DNS (Domain Name System)). When the administrative boundaries of private and public namespaces overlap, name resolution may yield unintended or harmful results.

Name collisions are not new. The introduction of any new Domain Name (Domain Name) into the DNS (Domain Name System), whether a generic TLD (Top Level Domain), country code TLD (Top Level Domain) or Second-Level Domain name (SLD (Second-level domain of the DNS)) creates the potential for name collision. A secure, stable and resilient Internet is ICANN (Internet Corporation for Assigned Names and Numbers)'s number one priority. Therefore, we've made a commitment to the Internet community to launch a substantial effort to mitigate and manage collision occurrence.

If your system is suffering demonstrably severe harm as a consequence of name collision, please fill in the form below to report the incident.

ICANN (Internet Corporation for Assigned Names and Numbers) will initiate an emergency response for name collision reports only where there is a reasonable belief that the name collision presents a clear and present danger to human life.

The emergency response could include temporarily removing the effected SLD (Second-level domain of the DNS) or the entire TLD (Top Level Domain) from the DNS (Domain Name System). ICANN (Internet Corporation for Assigned Names and Numbers) will serve as the initial reporting point, and if necessary will coordinate with registry operators to ensure that the report is acted upon in an expedited manner. ICANN (Internet Corporation for Assigned Names and Numbers)'s contracted Registry Operators are required to act on requests from ICANN (Internet Corporation for Assigned Names and Numbers) within 2 hours of receipt of the request from ICANN (Internet Corporation for Assigned Names and Numbers).

If you believe your name collision meets the criteria above (i.e. your system is suffering demonstrably severe harm as a consequence of name collision or you have a reasonable belief that the name collision presents a clear and present danger to human life), please use the form below to submit your report to ICANN (Internet Corporation for Assigned Names and Numbers).

After submitting the report, please review the [Guide to Name Collision Identification and Mitigation for IT Professionals \(https://www.icann.org/en/system/files/files/name-collision-mitigation-01aug14-en.pdf\)](https://www.icann.org/en/system/files/files/name-collision-mitigation-01aug14-en.pdf) [PDF, 476 KB] for more information.

All fields marked with asterisk (***) are required.

* Domain Name (Domain Name) Causing Harm Related to Name Collision (e.g., foo.bar.example):

* Requestor's Name:

* Requestor's Email:

* Requestor's Phone Number:

* Requestor's Address:

0/250 characters

Organization Name:

* Start Date of Domain Name (Domain Name) Usage:

Day	▼	Month	▼	Year	▼
-----	---	-------	---	------	---

* When did you learn about the issue?

Day	Month	Year
-----	-------	------

* Is the issue causing clear and present danger to human life?

- Yes
- No

* Description of the issue:

0/1000 characters

* Describe the impact caused by the issue:

0/1000 characters

Describe any solutions, workarounds, or mitigation measures put in place to manage the issue:

0/1000 characters

Other domain names involved, e.g., another.example (if applicable):

0/500 characters

Other pertinent contact information (if necessary):

0/500 characters

Appendix B - Web Search Results for “127.0.53.53”

Date	Sentiment	gTLD (Delegated)	App	Root Cause Symptoms	ICANN Identified	Other
Sep 2014	Neutral	prod (Aug 2014)	SSH	Unqualified (suffix search list), non-private	Y	
https://serverfault.com/questions/626612/dns-just-started-resolving-my-server-prod-addresses-to-127-0-53-53						
Aug 2015	Neutral	drive (Jun 2015)	Web Browser	Single label resolution	Y	Google search intended
https://superuser.com/questions/958758/why-pinging-drive-gets-replies-from-127-0-53-53						
Oct 2016	Neutral	[Unknown]	[Unknown]	[Unknown]	Y	Firewall logs
https://community.helpsystems.com/forums/intermapper/general-network-questions/3c736b35-b09b-e611-80d8-0050568473e2						
Oct 2014	Neutral	dental (Apr 2014)	[Unknown]	Unqualified (suffix search list, WinXP-style), private	Y	
https://www.reddit.com/r/sysadmin/comments/2jcdso/workstations_resolving_domainlocal_to_12705353/						
Aug 2016	Neutral	dev (Dec 2014)	valet	FQDN, private	Y	Not intended to resolve
https://github.com/laravel/valet/issues/115						
Jan 2016	Neutral	cisco (May 2015)	ping	FQDN and suffix search list, private	N	
https://community.spiceworks.com/topic/1381179-host-name-pinging-to-127-0-53-53						

Feb 2020	Neutral	cpa (Sep 2019)	[Unknown]	FQDN, VPN, private	N	
https://community.meraki.com/t5/Security-SD-WAN/Receiving-127-0-53-53-when-connected-to-the-Client-VPN-FQDN-s/m-p/75929						
Apr 2017	Neutral	[unknown]	RDP	Unqualified (suffix search list), VPN	Maybe (arpa)	
https://community.logmein.com/t5/LogMeIn-Hamachi-Discussions/FQDN-for-hamachi-hosts-127-0-53-53/td-p/139663						
Jun 2015	Neutral	windows (Jun 2015)	[Unknown]	FQDN, private	Y	
https://social.technet.microsoft.com/Forums/en-US/63ac3e27-7e95-47d2-a969-4044737aec0a/dns-collisions-with-windows-tld?forum=winserveripamdhcpdns						
Sep 2014	Angry	prod (Aug 2014)	[Unknown]	Unqualified multi-label, non-private	Y	Google (registry) also known
https://domainincite.com/17278-victims-of-first-confirmed-new-gtld-collision-respond-fuck-google						
Feb 2017	Neutral	bar (Feb 2014)	Apache Kafka (unit testing)	FQDN, private	N	Not intended to resolve
https://issues.apache.org/jira/browse/KAFKA-4765						
Oct 2014	Neutral	[Unknown]	[Unknown]	[Unknown]	Y	
https://blog.51cto.com/u_8378022/1560434						
Aug 2017	Neutral	dev (Dec 2014)	Web browser	FQDN, private	Y	Dev environment
https://apple.stackexchange.com/questions/296588/cant-connect-to-server-app-local-sites						
May 2015	Neutral	int (??) (Nov 1988)	ping	FQDN, ??	Y	
https://blog.manton.im/2015/05/12705353-dns-name-collision.html?m=1						

Oct 2014	Neutral	world (Sep 2014)	php, tnspring	FQDN, private	Y	Access to DB backend; Not intended to resolve
https://crumblybits.com/?p=316						
Dec 2017	Neutral	dev (Dec 2014)	gitlab-ci-multi-runner	FQDN, private	N	
https://gitlab.com/gitlab-org/gitlab-foss/-/issues/41072						
Apr 2017	Neutral	box (Nov 2016)	[Unknown]	FQDN or unqualified, private	Y	Access to pi-hole on LAN
https://discourse.pi-hole.net/t/pi-hole-server-lose-awareness-of-it-self/2715/15						

Appendix C - General Name Collisions Survey

DNS Suffix Usage and new gTLD Delegation

This survey has been commissioned to better help ICANN understand the impact of delegating new generic top-level domains (gTLDs). Your responses will remain anonymous.

This survey uses the term "DNS suffix" to refer to a domain name used in the DNS resolver search list of a device, e.g., the "domain" and "search" entries in `/etc/resolv.conf` on UNIX/Linux, "Search Domains" in the macOS DNS configuration pane, and "DNS suffix search list" on Windows.

To communicate with us about this research or this survey, please contact Casey Deccio <casey.deccio@icann.org>.

1. Has your organization ever used a DNS suffix associated with a new gTLD for its internal configuration? See <https://newgtlds.icann.org/en/program-status/delegated-strings> for more information.

Mark only one oval.

- Yes
- No *Skip to question 15*

DNS Suffix Information

2. Which DNS suffix(es) have been used by your organization?

Your response will be kept anonymous.

3. Which gTLD(s) correspond to the DNS suffix(es) used by your organization?

For example: the gTLD for the DNS suffix "foo.network" would be "network". This is helpful in the case you chose not to provide the actual DNS suffix.

4. Did your organization ever use the DNS suffix(es) *before* the date the gTLD(s) was/were delegated? See <https://newgtlds.icann.org/en/program-status/delegated-strings> for more information, including delegation dates.

Mark only one oval.

- Yes
 No
 Not sure

5. Are the DNS suffix(es) *still* in use by your organization?

Mark only one oval.

- Yes
 No
 Not sure

6. Are your organization's DNS resolvers configured to answer authoritatively for the DNS suffix(es) (i.e., without querying servers on the Internet)?

Mark only one oval.

- Yes, its DNS resolvers have always been configured this way
 Yes, its DNS resolvers are currently configured this way, but it has not always been this way
 No
 Not sure

7. Has your organization experienced technical problems with the use of the DNS suffix(es) *since* the delegation of the gTLD(s)?

Mark only one oval.

- Yes
- No *Skip to question 15*
- Not sure

Technical Issues with DNS Suffixes

8. When did you become aware of the technical issues regarding the DNS suffix(es)?

Mark only one oval.

- Within days after the delegation date
- Within weeks after the delegation date
- Within months after the delegation date

9. How many individuals or computer systems were affected by the technical issues associated with the DNS suffix(es)?

Mark only one oval.

- Only a few individuals or systems were affected
- Many individuals or systems were affected
- Nearly all individuals or systems were affected

10. How many computer systems are in your organization?

Mark only one oval.

- Fewer than 1,000
- Between 1,000 and 10,000
- More than 10,000

11. How long did it take for your organization to *resolve* the technical problem(s) related to the DNS suffix(es) after they were discovered?

Mark only one oval.

- Days after they were discovered
- Weeks after they were discovered
- Months after they were discovered
- Years after they were discovered
- They have not been resolved

12. When did your organization learn that the technical issues might be related to the delegation of a new gTLD in the DNS?

Mark only one oval.

- Some time before the problem was resolved
- Some time after the problem was resolved
- Not until now

13. What role did the IP address 127.0.53.53 have in identifying the cause of the problem?

Mark only one oval.

- 127.0.53.53 was not observed
- 127.0.53.53 was observed, but its meaning and origin were unclear or not helpful
- 127.0.53.53 led us to ICANN and the delegation of the new gTLD

14. What more are you willing to share with regard to your organization's experience? For example, you might include additional details about your organization's system configuration, the problems experienced, other DNS suffixes affected, how you solved the issue, what other entities got involved, what it cost your organization in terms of time, effort, and money, and more.

Skip to question 16

Additional Information

15. What more are you willing to share with regard to your organization's DNS experience with gTLDs? For example, you might include additional details about your organization's system configuration, any other suffixes that might be in use, any questions you have, etc.

Skip to question 16

Contact Information

16. If you wish to communicate further, please share your email address. It will only be used for communications related to this survey. Any record of your email address will be deleted at the conclusion of the research project, which is expected to end in mid-2022.
-

This content is neither created nor endorsed by Google.

Google Forms

Appendix D - Targeted Name Collisions Survey

DNS Suffix Usage and new gTLD Delegation

This survey has been commissioned to better help ICANN understand the impact of delegating new generic top-level domains (gTLDs). Your responses will remain anonymous.

This survey uses the term "DNS suffix" to refer to a domain name used in the DNS resolver search list of a device, e.g., the "domain" and "search" entries in `/etc/resolv.conf` on UNIX/Linux, "Search Domains" in the macOS DNS configuration pane, and "DNS suffix search list" on Windows.

To communicate with us about this research or this survey, please contact Casey Deccio <casey.deccio@icann.org>.

1. Which DNS suffix was referenced in the email?

For example, "foo.network". Your response will be kept anonymous.

2. Which gTLD corresponds to the DNS suffix referenced in the email?

For example: the gTLD for the DNS suffix "foo.network" would be "network". This is helpful in the case you chose not to provide the actual DNS suffix.

3. Has your organization ever used the DNS suffix referenced in the email?

Mark only one oval.

- Yes
- No *Skip to question 15*
- Not sure

DNS Suffix Use

4. Did your organization ever use the DNS suffix **before** the date indicated?

Mark only one oval.

- Yes
 No
 Not sure

5. Is the DNS suffix **still** in use by your organization?

Mark only one oval.

- Yes
 No
 Not sure

6. Are your organization's DNS resolvers configured to answer authoritatively for the DNS suffix (i.e., without querying servers on the Internet)?

Mark only one oval.

- Yes, its DNS resolvers have always been configured this way.
 Yes, its DNS resolvers are currently configured this way, but it has not always been this way.
 No.
 Not sure.

7. Has your organization experienced technical problems with the use of the DNS suffix *since* the date indicated in the email?

Mark only one oval.

- Yes
- No *Skip to question 15*
- Not sure

Technical Issues with DNS Suffixes

8. When did you become aware of the technical issues regarding the DNS suffix?

Mark only one oval.

- Within days after the date listed
- Within weeks after the date listed
- Within months after the date listed

9. How many individuals or computer systems were affected by the technical issues associated with the DNS suffix?

Mark only one oval.

- Only a few individuals or systems were affected.
- Many individuals or systems were affected.
- Nearly all individuals or systems were affected.

10. How many computer systems are in your organization?

Mark only one oval.

- Fewer than 1,000
- Between 1,000 and 10,000
- More than 10,000

11. How long did it take for your organization to *resolve* the technical problem(s) related to the DNS suffix after they were discovered?

Mark only one oval.

- Days after they were discovered
- Weeks after they were discovered
- Months after they were discovered
- Years after they were discovered
- They have not been resolved

12. When did your organization learn that the technical issues might be related to the delegation of a new gTLD in the DNS?

Mark only one oval.

- Some time before the problem was resolved
- Some time after the problem was resolved
- Not until now

13. What role did the IP address 127.0.53.53 have in identifying the cause of the problem?

Mark only one oval.

- 127.0.53.53 was not observed.
- 127.0.53.53 was observed, but its meaning and origin were unclear or not helpful.
- 127.0.53.53 led us to ICANN and the delegation of the new gTLD.

14. What more are you willing to share with regard to your organization's experience? For example, you might include additional details about your organization's system configuration, the problems experienced, other DNS suffixes affected, how you solved the issue, what other entities got involved, what it cost your organization in terms of time, effort, and money, and more.

Skip to question 16

**Additional
Information**

What more are you willing to share with regard to your organization's experience? For example, you might include additional details about your organization's system configuration, any questions you have, etc.

15. Additional Information

Skip to question 16

Contact Information

16. If you wish to communicate further, please share your email address. It will only be used for communications related to this survey.

This content is neither created nor endorsed by Google.

Google Forms

Appendix E - General Email Sent to NANOG Subscribers

Dear colleagues,

tl;dr: Please take our survey on DNS suffix usage here: <https://forms.gle/ntvsn6eqzYH9YcTN6>

The Internet Corporation for Assigned Names and Numbers (ICANN) is researching the technical impact of delegating new generic top-level domains (gTLDs). This research is part of the Name Collision Analysis Project (NCAP). More information about NCAP can be found at <https://community.icann.org/display/NCAP>.

Since 2013 hundreds of new gTLDs have been introduced into the public DNS (<https://newgtlds.icann.org/en/program-status/delegated-strings>). In some cases those gTLDs might have been used as part of a DNS suffix by one or more organizations around the Internet, prior to their introduction. (By “DNS suffix” we mean a domain name used in the DNS resolver search list of a device, e.g., the “domain” and “search” entries in `/etc/resolv.conf` on UNIX/Linux, “Search Domains” in the macOS DNS configuration pane, and “DNS suffix search list” on Windows.) As a result, the behavior of systems or devices in these organizations might have changed because of a “name collision”. A name collision occurs when a name used in one context (in the organization's network) is interpreted in another context (in this case, in the public DNS after the corresponding gTLD went live).

We are researching the causes and impact of name collisions. We are seeking qualitative data based on experiences of those organizations potentially affected. We expect that this additional data will greatly enhance our understanding of name collisions that resulted from adding new gTLDs.

If you suspect that your organization has been impacted by the delegation of any new gTLDs, we invite you to please fill out the following brief survey regarding your experience. We would be grateful for your input!

<https://forms.gle/ntvsn6eqzYH9YcTN6>

Your responses will remain anonymous, and any personal information will be discarded after the research has concluded.

If you have any questions, please reply to this email.

Thank you for your help!

Sincerely,

18 January 2023

Casey Deccio
ICANN Name Collisions Analysis Project

Appendix F - Targeted Email Sent to AS Contacts

Dear network administrator,

The Internet Corporation for Assigned Names and Numbers (ICANN) is researching the technical impact of delegating new generic top-level domains (gTLDs). This research is part of the Name Collision Analysis Project (NCAP). More information about NCAP can be found at <https://community.icann.org/display/NCAP>.

Based on our research, we believe systems or devices in your organization might have been using the DNS suffix “«DNSSuffix»” when the top-level domain “«gTLD»” was added to the DNS root zone on «Date». (By “DNS suffix” we mean a domain name used in the DNS resolver search list of a device, e.g., the “domain” and “search” entries in /etc/resolv.conf on UNIX/Linux, “Search Domains” in the macOS DNS configuration pane, and “DNS suffix search list” on Windows.) We inferred possible use of this DNS suffix by analyzing several years of DNS queries captured at the DNS root servers as part of the annual Day In the Life (DITL) collection (<https://www.dns-oarc.net/oarc/data/ditl>). We used publicly available WHOIS information for your autonomous system to find your contact information and send this email.

After the TLD «gTLD» went live, the behavior of systems or devices in your organization might have changed because of a “name collision”. A name collision occurs when a name used in one context (in this case, inside your organization) is interpreted in another context (in this case, in the public DNS after «gTLD» went live).

We are researching the causes and impact of name collisions. We are seeking qualitative data based on experiences of those organizations potentially affected. We expect that this additional data will greatly enhance our understanding of name collisions that resulted from adding new gTLDs.

Would you be willing to please fill out the following brief survey regarding your experience? We would be grateful for your input!

<https://forms.gle/1kj6VtEK1M5ANq8JA>

Your responses will remain anonymous, and all personal information will be discarded after the research has concluded.

If you have any questions or would like to opt out of future communications related to this topic, please reply to this email.

Thank you for your help!

Sincerely,

18 January 2023

Casey Deccio
ICANN's Name Collisions Analysis Project

18 January 2023