Continuous Data-driven Analysis of Root Stability (CDAR)

Deliverable D2: Root Stability Report

Revision: Final

Delivery date: March 8, 2017

Responsible: TNO (on behalf of NLnet Labs, SIDN, TNO consortium)

Project Leader: Bart Gijsen (TNO)

Dissemination level: public

Executive summary

The New gTLD Program started in 2012 and resulted in the addition of more than 1,100 new gTLDs to the root zone since October 2013. Following previous advisories from the Governmental Advisory Committee and its Board of Directors, ICANN commissioned an empirical study into the technical impact of the New gTLD Program on the security and stability of the root DNS system, which was conducted by a consortium consisting of TNO, SIDN and NLnet Labs.

The study's primary research question was: *did the delegation of new gTLDs degrade the stability or security of the root DNS system?* And based on the analysis carried out for this research question, our second research question is: *Can we expect that the delegation of more new gTLDs will degrade the stability or security of the root DNS system in the future?*

Observed impact

To determine the technical impact on the security and stability of the root DNS system to date, we have analyzed large amounts of historical Internet measurement data sets, including RSSAC002, DNS-OARC's Day In The Life of the Internet (DITL) and RIPE NCC's Atlas measurement to the root DNS system.

Overall, our analysis shows that the root DNS system has been able to absorb the increase in root server traffic during the period in which new gTLDs were delegated without problems. We did not find any measurable degradation of the stability or security of the root DNS system in the period until the writing of this report (February 2017) that we could attribute to the new gTLDs.

In further detail, we found that:

- The total query rate to the root DNS system has been growing during the period in which new gTLDs were being delegated. More detailed analysis of the total query volume to the root DNS system indicates that we should distinguish between queries for valid and invalid (non delegated) TLDs, complemented with incidental peaks of query volumes during rare events. The fraction of queries for invalid TLD names increased over time and was around 64% in the DITL 2016 measurements. We found no indication that the delegation of new gTLDs contributed to this increase. Similarly, we found no indication that new gTLDs have contributed to incidental peaks of query volumes to the root DNS system (such as the unusually high query volumes on 30 November 2015 and 25 June 2016).
- Queries for new gTLDs did contribute to the *valid* query volume to the root DNS system although this contribution is very small. For example, in the DITL 2016 measurements (April 2016), the fraction of queries for new gTLDs constituted only 1.1% of the total *valid* query volume received at the root DNS system (and it constituted only 0.4% of the total query volume). Although this fraction has been increasing slightly every year since the start of the New gTLD Program, we conclude that the contribution of queries for new gTLDs so far form an insignificant fraction of the total query volumes at the root DNS system.
- In a more microscopic perspective, we can observe some specific effects of the delegation of new gTLDs. For example, when we zoom in on days around the delegation of a new gTLD, we see that the query rate the root DNS system receives for that TLD does fluctuate. But for each of these effects the root DNS system was capable of quickly converging to a new stable state. Moreover, we observed no significant impact of such fluctuations surrounding delegations on the Round Trip Time performance and the reachability of the

root DNS system from the user's perspective (using measurements from outside the root DNS system).

- The consistency of the data at the root is very high. To verify the potential impact of the delegation of new gTLDs on the consistency of DNS data we scanned for root zone file errors and DNSSEC validation errors. The results from these (simple) scans show no indication that there has been an impact.

Since our conclusion is limited to the results of the analyses that we designed and executed and is confined by the imperfections of the available measurement data, we have engaged with the technical DNS community at many events. The discussions during these events revealed no additional possible negative effects of the introduction of new gTLDs on the stability and security of the root DNS system, and confirmed our measurement-based analyses as an appropriate approach.

Possible Future Impact

From our analysis of the historic data sets we inferred time-invariant characteristics in root query rates and in the query-response behavior. A particular time-invariant correlation we identified is the ratio between the number of valid queries for a TLD at the root DNS system and the number of registered (second level) domain names in the TLD, which enables us to extrapolate the impact of new gTLDs on future growth in DNS traffic. If we presume that this correlation will remain time-invariant (as it has in past years), then the valid query volume for new gTLDs will only become significant if the number of domains in a new gTLD grows to .com-like proportions or if the number of new gTLDs grows in an unbounded way. When extrapolating the relatively steady growth of domain names in the new gTLDs, we do not see an indication that this is likely to happen in the near future.

More speculatively, we believe that the removal of new gTLDs from the root zone file could be a potential stability risk. This could lead to clients that leak or direct large invalid query volumes to the root DNS system. Another risk parameter is an increase in the amount of processing on root name servers, which would reduce the amount of traffic that they can handle. This might for instance occur if resolvers switch from using UDP to TCP on a massive scale. Similarly, the use of DNS Server Cookies and potentially other protocol extensions may increase the amount of server-side processing on root servers. However, we believe that this scenario is unlikely to materialize as a result of new gTLD delegations.

Conclusions

In conclusion, we did not observe a degradation of the security and stability of the root DNS system as a result of the delegation of new gTLDs. Moreover, presuming that the evolution of new gTLD delegations continues to exhibit the pattern we observed since the New gTLD Program's first delegations in October 2013, we see no signs that the delegation of more new gTLDs in itself will degrade the stability or security of the root DNS system in the near future.

Recommendations

The absence of an observed degradation of the security and stability of the root DNS system is no reason to be less cautious about possible future impact of the New gTLD Program. In particular, we recommend the New gTLD Program to continue enforcing a controlled rate of delegating new gTLDs, which is one of its current preventive root zone scaling measures. Further, we advise more frequent monitoring of the impact of new gTLD delegations. To enable monitoring in a more frequent manner an upgrade of current data collection efforts is recommended. Finally, we recommend including frequent monitoring of the identified risk parameters.

Authors

Name	Partner	Email
Bart Gijsen	TNO	bart.gijsen@tno.nl
Benno Overeinder	NLnet Labs	benno@NLnetLabs.nl
Cristian Hesselman	SIDN	cristian.hesselman@sidn.nl
Daniël Worm	TNO	daniel.worm@tno.nl
Giovane Moura	SIDN	giovane.moura@sidn.nl
Jaap Akkerhuis	NLnet Labs	jaap@NLnetLabs.nl

Table of Contents

1 Introduction

Generic Top-Level Domains (gTLDs) are one of the categories of top-level domains and include extensions such as .com, .net, and .org. The New gTLD Program was developed via the ICANN multi-stakeholder process to increase competition and choice in the domain name space by extending the number of gTLDs. The new gTLD application window opened in 2012, and ICANN received 1,930 applications for new gTLDs. As of September 2016, more than 1,100 new gTLDs have been delegated (i.e., added to the DNS root zone).

Figure 1 shows the growth of the root zone in terms of the total number of TLDs (ccTLDs, gTLDs and new gTLDs). The first new gTLD delegation was on 23 October 2013. As we can see in the figure, the curve shows sharp growth since the initial delegations of new gTLDs.

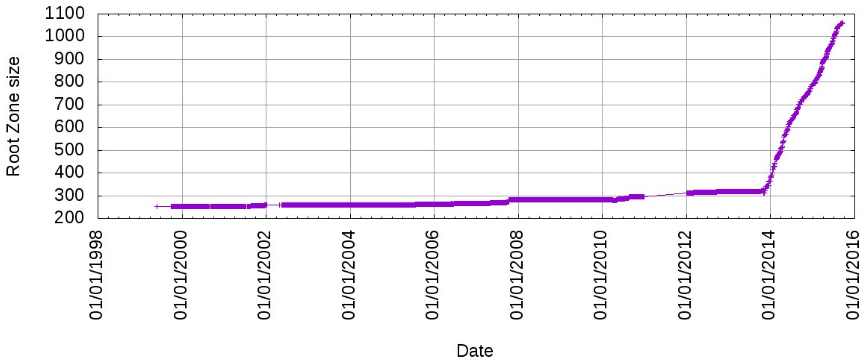

Figure 1: Growth of the root zone in terms of TLDs.
Data source: Zone File Repository, DNS-OARC

Given the critical role of the root DNS system for the operation of the DNS and the Internet at-large, ICANN issued a Request for Proposal (RFP) for a root stability study [1]. The main goal defined in the RFP was to determine whether this growth of the root zone could impact both security and stability of the root DNS system.

The Scaling the Root study of 2009 [2] addressed the possible impact of the introduction of DNSSEC, IDNs, IPv6 and new gTLDs on the root DNS system. In this report it was emphasized that the root DNS system is a highly dynamic system for which any change involves risk. Therefore it was recommended that changes to the root DNS system should be made gradually and that its effects should be monitored. Until the present there have not been widespread reports of performance and stability issues on the root DNS system due to growth of the root zone. Nevertheless, in the context of the substantive growth of the root zone size shown in Figure 1 it remains advisable to monitor the impact of the New gTLD Program, especially now that large-scale empirical measurement data has become available for more thorough, quantitative analysis.

With the goal of empirically determining whether the growth of the root zone has in any way impacted the stability or security of the root DNS system, ICANN commissioned TNO and its consortium partners, SIDN and NLnet Labs to conduct the Continuous Data-driven Analysis of Root Server System Stability (CDAR) study. (For consortium details see Appendix C.) The study started in the second half of 2015, and this report includes the project's final results.

1.1 Goal, Research Questions and Approach

The *primary goal* of this study is to determine the technical impact of the New gTLD Program on the security and stability of the root DNS system. To achieve this goal, we focused on the primary research question:

> *Has the introduction of new gTLDs degraded the stability or security of the root DNS system?*

And as a secondary question:

> *Can we expect that the introduction of many more gTLDs will degrade the stability or security of the root DNS system in the future?*

Our approach to answer the primary research question is to analyze large amounts of historical and longitudinal measurement data, with the aim of determining if there has been any correlation between root zone file growth and stability/security of the root DNS system. To provide an objective answer to this research question, we analyzed the data sets without any prejudice or particular expectations.

The approach to answer the secondary research question is to infer, from the analysis of the primary research question, time-invariant correlations between parameters that define the size of the root zone and the stability/security metrics of the root DNS system. Using these correlations we can investigate extrapolation of the findings, assuming that these correlations will remain constant in the near future.

1.2 Contributions

The key contribution of our study is that it is the first empirical study into the impact of the introduction of the new gTLDs on the security and stability of the root DNS system using large amounts of publicly available, historical measurement data.

The second contribution is the methodology we developed, which introduces root security and stability parameters such as query volume stability that are based on commonly referenced definitions of DNS security and stability. We analyzed and measured these parameters using passive data obtained from various sources such as RSSAC-002 data, and data collected by/through DNS-OARC, as well as through active measurements that we as well as others carried out. The methodology is supported by a set of tools that we developed to analyze our root stability parameters, which may be useful for future studies.

Another contribution of our study involves frequent outreach to the ICANN and DNS communities (e.g., DNS-OARC, IETF/IEPG and RIR meetings). This enabled us to share and discuss our approach and intermediate results with ICANN's multi-stakeholder community and refine our analyses based on their feedback.

1.3 Scope and Limitations

The CDAR study is focused on the *technical* impact of the introduction of new gTLDs on the security and stability of the root DNS system[1]. We emphasize that policy recommendations based on our technical observations are *not* part of this study. Instead, ICANN intends to use the results from this study, as well as those from other requested studies, as input for community discussions regarding the future expansion of the root zone[2].

It is important to understand that the findings in this report are focused on the technical impact of *new gTLDs* on the security and stability of *the root DNS system*. This implies that the study is not focused on other security and stability risks for the root DNS system, nor on the impact of the delegation of new gTLDs on parts of the DNS other than the root DNS system. For example, we have been requested to investigate the impact of the delegation of new gTLDs on resolvers, but such a study was out of scope because resolvers are not part of the root DNS system.

Additionally, investigating the impact of new gTLDs on the security and stability of the root DNS system is challenging given the complexity of the system and its constituents, as well as the incompleteness of available data. For example, the validity period for some of the study results is limited in time, because the root DNS system and its interaction with the global Internet infrastructure are continuously changing. Also, the complexity of the root DNS system and the limitations of the available data for this study prevent us from providing detailed explanations for all of the presented results. For example, we observe a growth in the total number of queries sent to the root DNS system (see Section 4.1.1), but were unable to deduce from the data why it is growing. In turn, the absence of such explanations reduces our options to extrapolate our observations for possible future DNS developments. Nevertheless we can still deduce relevant and meaningful extrapolations for some possible future developments (see Section 6).

1.4 Reading Guide

We presume that the reader is aware of the New gTLD Program and has a basic understanding of the operations of the root DNS system. To assist the reader interested in more detailed information about the program and the root DNS system, we have included relevant references throughout this report.

The following section provides background information about the root DNS system and a breakdown of its security and stability. We also include a brief summary of relevant results from other DNS studies. In Section 3 we present the CDAR methodology, including an overview of the interactions with ICANN's multi-stakeholder community. Results from the measurement-based analyses that are described in this methodology are presented in Section 4. We provide a comprehensive discussion of these results and their implications regarding our original research questions and present a conclusion in Section 5. In Section 6 we provide some reflections on possible future impact of the New gTLD Program on the security and stability of the root DNS system.

[1] In Section 2.1 background information about the root DNS system is provided. In general, the DNS terminology used in this report is aligned with RFC7719 (https://tools.ietf.org/html/rfc7719), "DNS Terminology," as much as possible.

[2] ICANN's Board of Directors commissioned the study following a recommendation from the Governmental Advisory Committee: https://archive.icann.org/en/topics/new-gtlds/board-notes-gac-scorecard-clean-15apr11-en.pdf

2 Background

In this section we provide some limited background information on the root DNS system, the New gTLD Program, as well as related work. We focus on the background information necessary to understand this report and present the reader with the relevant literature for a more detailed description of the root DNS system.

2.1 Root DNS System

The root DNS system is pivotal in the DNS resolving process. In principle every DNS query starts at the root DNS system and recursively follows delegations in the DNS tree down to the authoritative name server that provides the answer to the query. A graphical representation of this process is shown in Figure 2.

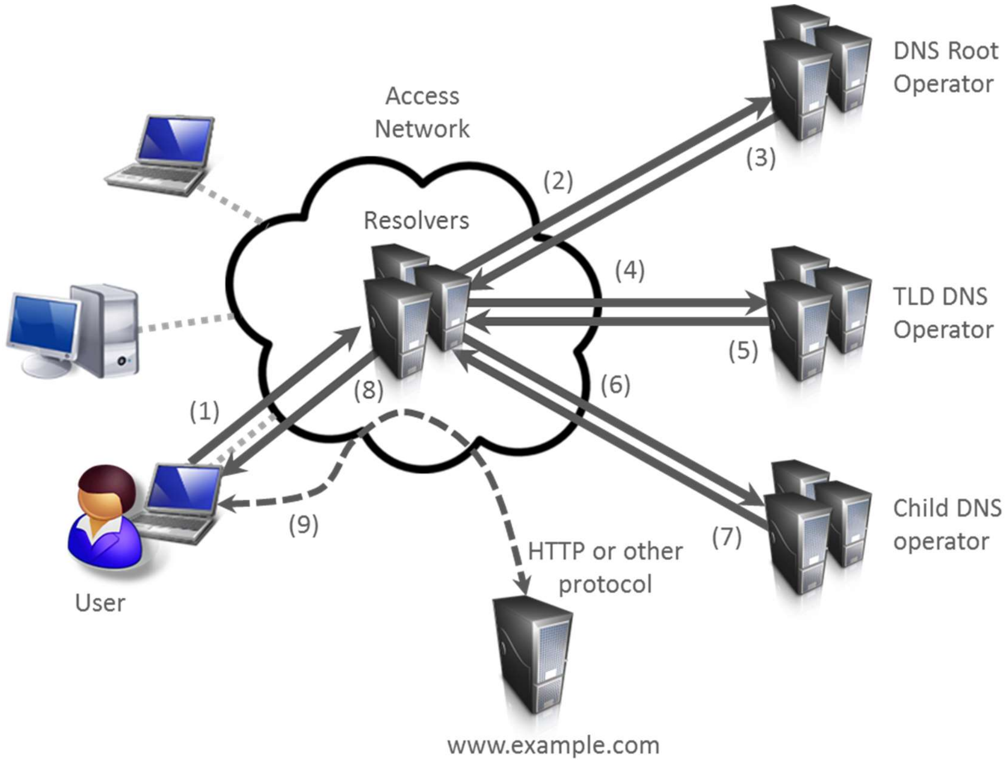

Figure 2: High-level DNS resolving process

In the resolution process, the root DNS system provides the records for all TLDs. Because of the importance of the root DNS system and to increase its resilience, there are multiple levels of redundancy employed: letter, site, and server.

Instead of having one root server operator, there are 12 root server operators responsible for the 13 root server letters (a–m), as shown in Table 1. These organizations are independent from each other, and have different budgets/setup/configurations for their respective letter [3], [4].

Table 1: Root Server Letters and Organizations. Sites reported by [3] on 16 February 2017 and 13 March 2008, respectively [5]

Letter	Organization	Sites (2017)	Sites (2008)
A	Verisign	5	1
B	ISI/USC	1 (unicast)	1
C	Cogent	8	4
D	University of Maryland	111	1
E	NASA	71	1
F	Internet Systems Consortium	58	43
G	DOD	6	1
H	U.S. Army Research Lab	2 (primary/backup)	1
I	Netnod	50	31
J	Verisign	127	37
K	RIPE NCC	50	17
L	ICANN	157	2
M	WIDE Project	8	6

As can been seen, the configurations of the letters vary significantly. To improve performance and availability, 11 of the 13 letters employ IP anycast, a technology that allows the same IP address of the DNS servers to be hosted in multiple places around the world. As a consequence, the stability of one root server letter is also related to its anycast deployment. There are significant differences across the letters. For example, B-ROOT has one site as of this writing, while L-ROOT has 154 (a site typically is a city in which these letters are hosted). Figure 3 shows the global distribution of root server letters.

In addition, there might be an extra layer of redundancy: on each anycast site multiple servers may be used. For example, K-ROOT, at its Tokyo, Japan site, has three different servers (as of December 2015). Ultimately, that adds a significant level of redundancy to the root DNS system, even in the event of a DDoS attack [4].

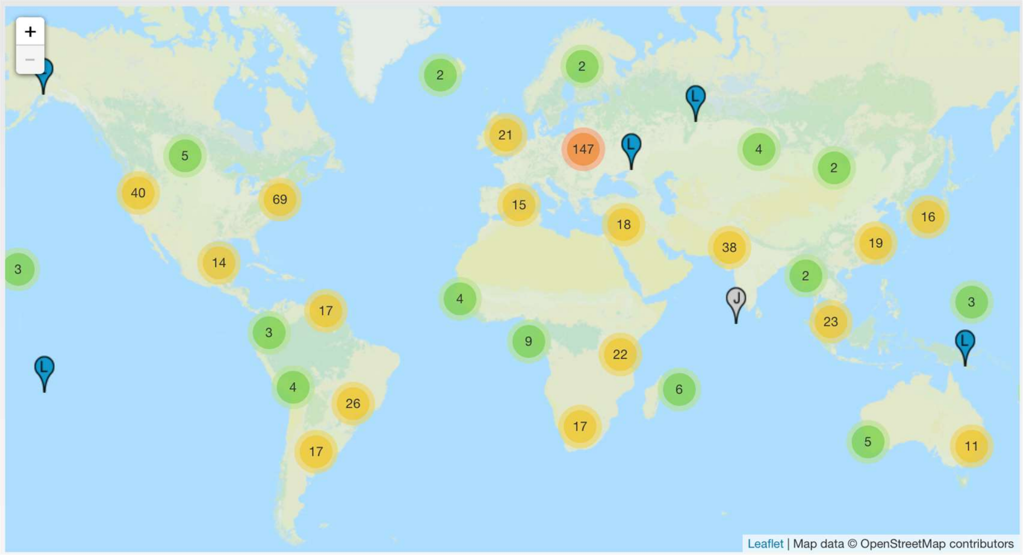

Figure 3: Distribution of root server letters (16 February 2017), source: [3]

Root DNS Evolution

The root DNS system is also upgraded over time: more servers/sites/links have been and will be added to its infrastructure. As an illustration, we show in the right most column of Table 1 the configuration of each root server letter in 2008.

2.2 New gTLD Program

One of ICANN's key responsibilities is introducing and promoting competition in the registration of domain names, while ensuring the security and stability of the DNS[3]. From this responsibility ICANN's Generic Names Supporting Organization (GNSO) initiated the process for the delegation of new gTLDs in 2005. After extensive development, the New gTLD Program [6] was defined and the application window for new gTLDs opened on 12 January 2012. After passing the evaluation procedures, the applied-for new gTLDs continue to be delegated to the root zone. On 23 October 2013 ICANN announced the first gTLD delegations[4]. Since then new gTLD delegations continue to be gradually delegated to the root zone[5].

Given the development period of the New gTLD Program our primary focus is on the time period from 2012 until present. Moreover, our focus is on the changes to the root zone that have taken place as a result of the New gTLD Program, rather than changes in the root zone that have taken place in the same period due to other developments such as the IDN ccTLD Fast Track Process. In order to restrict our focus to the impact of new gTLDs we used relevant data provided by ICANN [6].

2.3 Security and Stability

Given the objective of our study the terms of security and stability of the root DNS system are at the core of this study. In general, the definitions of the security and stability of the root DNS system have been discussed frequently and there is no consensus about them (cf. [7], [8]). In [9] high-level definitions for DNS security and stability are presented:

[3] Quoted from [6]

[4] https://www.icann.org/resources/press-material/release-2013-10-23-en

[5] https://newgtlds.icann.org/en/program-status/statistics

- *DNS security*: “The ability of the components of the DNS to protect the integrity of DNS information and critical DNS system resources.”
- *DNS stability*: “The ability of the entire name resolution system and its component parts to be able to respond to DNS queries.”

These definitions are applicable to the wider scope of the DNS, and so we use and specialize them for the security and stability of the root DNS system in this report. From the perspective of the CDAR study the breakdown of the security and stability of the root DNS system into metrics is more relevant than the phrasing of their definitions.

2.3.1 Metrics

In DNS threat analysis literature several breakdowns of security and stability aspects can be observed. For example, in [10] a threat tree is presented in which the highest level branch distinguishes between denial of service and data corruption (apart from privacy threats). In [9] and [11] this categorization is acknowledged. Inspired by these analyses we derived the breakdown of root DNS system security and stability into the high-level metrics presented in Figure 4.

We distinguish two main categories of metrics: operational stability and DNS data consistency. The threat analysis literature provides a further subdivision of these two categories, although these sub-categories are not consistent between the different analyses. For our study we subdivide operational stability into: query rate stability at the root and query/response stability as perceived by users. This is shown in Figure 4. Data consistency covers both the correctness of the root zone data responded by the root DNS system (e.g. zero-error level and accuracy of the root zone data), as well as the consistency between the root zone data and the TLD data. The latter is in particular relevant in the context of DNSSEC, where a chain of trust is derived linking cryptographic data in the root to data in a TLD.

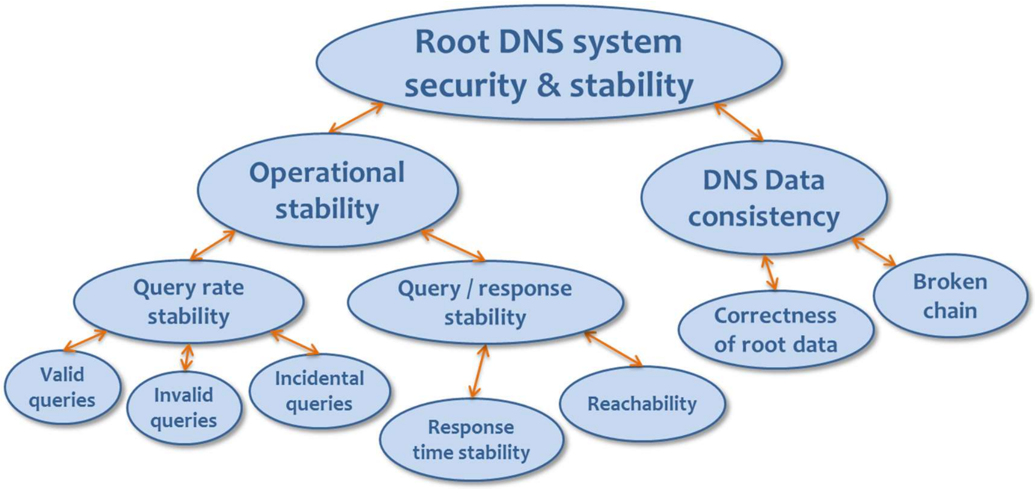

Figure 4: Breakdown of root DNS system security and stability into metrics

In following sections this high-level breakdown will be used to refine our research questions and to identify appropriate measurement data to analyze the impact of new gTLDs on the security and stability of the root DNS system.

2.3.2 Granularity

To understand the data collection and analysis results presented in this report, there is another aspect of DNS security and stability that we need to point out. This regards the different scales at which “root DNS system security and stability” can be observed. For example, we can observe from a macroscopic view where the stability of the root DNS system as a whole is considered, or we can zoom in into a single root server letter or even on a microscopic root DNS system level such as a site or an individual name server. As indicated in [12] a disruption at a single root server letter would not cause immediate stability issues for the root server system.

A second scale on which stability levels can be distinguished is the degree to which the system is able to respond to DNS queries. In one extreme case the system is able to respond to all queries whereas in another extreme case it is able to respond to none of the queries. The latter occurs if the system is disrupted. It is also possible that the system operates in a degraded state that is “in between” these two extreme cases. For example, a highly loaded system in the root DNS system may be able to respond to queries, but only at slow response times. Or it may respond to only a part of the queries that are received.

Figure 5 presents a (simplified) illustration of both scales of different granularity levels of root DNS system stability. Both axes, the root DNS system level and the state of operations, have a “continuous” scale. As an example the operational state at root DNS system, root server letter and node level are shown for the high query rate event of 30 November and 1 December 2015 [13]. During this event the incident traffic saturated network connections near some locations of root name servers, which resulted in timeouts for queries. Other root name servers were continuously reachable for the entire duration of the incident. Because the DNS protocol is designed to cope with partial reachability among a set of name servers at the same delegation level, the impact was barely perceptible by Internet users. As such the root DNS system as a whole was operating as intended.

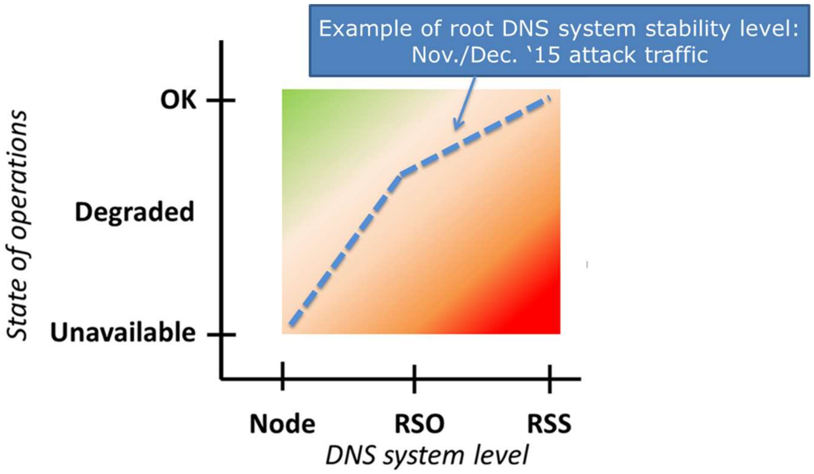
Example of root DNS system stability level:
Nov./Dec. ‘15 attack traffic
OK
Degraded
Unavailable
State of operations
Node
RSO
RSS
DNS system level

Figure 5: Granularity levels of root DNS system stability

For this study our primary focus is on the macroscopic view, but we also present some observations of more microscopic behavior to show some effects more clearly.

2.4 Related Work

The CDAR study extends related root stability studies. Early research presented in [14], [15] and [16] provided first, measurement-based insight in the DNS query and response volumes and characteristics at a single root server letter.

Obtaining insight into the behavior at root DNS system scale required more extensive data collection that would be coordinated between multiple root server letters. Such investigation was actually started after the first large scale root DNS system data collection effort in the day in the life of the Internet (DITL) project[6]. Since 2006 more and more root server operators (RSO) (and other organizations) have been contributing raw DNS data from their root servers for coordinated two-day DITL periods, once a year. Since the start of the DITL measurements valuable insights into the behavior of the root DNS system have been published. For example, [17] discusses the evolution of several root DNS system traffic characteristics, based on the DITL data collection experiments from 2006 until 2009.

Active monitoring systems are also being used to monitor the root DNS system from the outside. A prominent example is the RIPE Atlas infrastructure, which consists of thousands of measurement probes distributed around the globe. The subset of publicly available measurements particularly useful for the CDAR study is DNSMON [18]. Some of the DNSMON data traces back to 2012, the year prior to the delegation of new gTLDs. The wide variety of active measurements are, for example, used to analyze the effect of the high query rate event of 30 November and 1 December 2015 on the root DNS system [4].

There have also been data-driven studies that focused on investigating specific root security and stability events. For example, in the L-ROOT scaling study in 2009 [19] an analysis is provided to estimate the possible impact of IPv6, DNSSEC, and new gTLDs prior to their introduction to ICANN's L-ROOT servers. Another study focused on namespace collisions in the Global Internet DNS [20].

Some of the analyses in the CDAR study are inspired by this previous work and their results are complemented by the results from this study. A distinguishing factor of the CDAR study is that it is based on a wide range of root DNS system measurements (including DITL and RIPE Atlas), as explained in Section 3. Furthermore, driven by the CDAR study objective, more focus is put on the impact of the delegation of new gTLDs on the security and stability of the root DNS system. Zooming in on TLD specific characteristics adds a new dimension to the studies of the root DNS system, and this report presents new insights that have resulted from this perspective.

[6] https://www.caida.org/projects/ditl/

3 Methodology

3.1 Overview

The methodology we apply to answer the research questions raised in Section 1.1 is to analyze empirical measurement data collected from the root DNS system. The measurements used for this study are collected and made publicly available by other parties, including RSOs and DNS-OARC.

These measurements can be distinguished into two main categories: active and passive measurements. Figure 6 shows their relation to the root DNS system. Passive measurements refer to measurements collected, in our case, in the root DNS servers themselves. For example, the RSSAC002 [21] reports are based on passive measurements generated by processing the incoming traffic (denoted by “n” in Figure 6) on the servers.

Active measurements, on the other hand, refer to measurements collected by measurement devices outside the root DNS system. Active measurement devices send DNS queries (“q” in Figure 6) and receive responses (“r”), and these query response pairs can be used to measure properties such as reachability and performance. The measurements collected by RIPE Atlas are an example of this category.

Both passive and active measurements have their advantages and disadvantages: passive measurements provide information with regard to the query traffic received by the root DNS system. However, they do not allow measurement of the user experience when querying the root DNS system. Active measurements, on the other hand, cannot infer the total traffic received by the root server letters, but they provide an approximation for user behavior and experience.

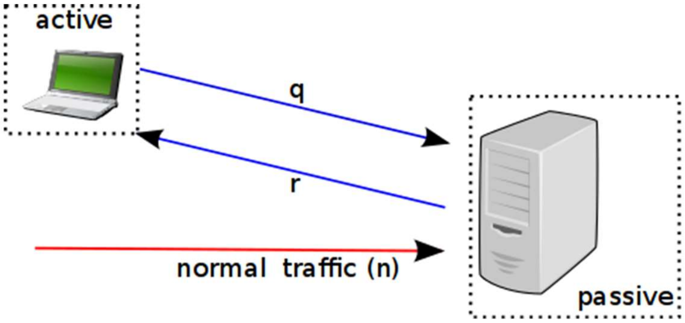

Figure 6: Active and passive measurements on the root DNS system

For measuring the metrics indicated in Figure 4 we apply both passive and active measurements. For example, to analyze the stability of the query rate at the root DNS system we can use the RSSAC002 passive measurements, while complementary active measurement data from RIPE Atlas are better suited to analyze response time stability. In the following subsections we explain which measurement data we used for our analyses at high level, including their use for analyzing the metrics. More detailed descriptions of the data sets are included in Appendix A.

3.2 Passive Measurement Data Sets

3.2.1 RSSAC002

In 2014, RSSAC published RSSAC002 [21]: an advisory on daily aggregated measurements of the Root Server System. This document contains descriptions of an initial set of parameters that according to RSSAC would be useful to monitor and establish a baseline trend of the root server system. These parameters include, among others, the number of queries and responses (traffic-volume) and response type distribution (rcode-volume), divided into sub-metrics.

The RSSAC002 measurements that contain the longest period of historical data trace back to 2013 (A-, I- and J-ROOT). As of February 2017 a total of eleven root servers are daily collecting and publishing RSSAC002 measurements, where the date of initial published measurement varies per root server (see Appendix A for more details).

The data is publicly available via www.root-servers.org, stored in YAML-files. Furthermore, DNS-OARC has taken up the responsibility to act as a long-term repository for this data[7]. For this study, we collected the measurement data via www.root-servers.org and we mainly used RSSAC002 measurements for analyses regarding day-to-day, root-letter aggregated query volumes over a longer period of time.

The timeframe is limited: there are no RSSAC002 measurements available covering a period *before* the delegation of new gTLDs. Not all suggested metrics are measured by the participating root servers and some root servers deviate from the proposed YAML-format. Furthermore, the data (such as daily number of queries) is not split out per TLD, which would have been particularly useful for this study. To some extent this restricts the scope of the analyses that we could do. Regarding the accuracy of the data, it should be noted that some data is missing (missing spans of 24h series)[8]. In addition we observed some inconsistencies that indicate errors or misinterpretation in the aggregation of data (see Appendix A). In our analyses we had to take these limitations into account (see remarks in Section 4).

3.2.2 DITL data

The DNS Operations, Analysis, and Research Center (DNS-OARC) organizes a yearly two-day data collection effort [22], also known as the Day In The Life of the Internet (DITL) data traces. The data consists of the queries received by the root and other DNS servers. The results are stored in files in an “industry standard” pcap format[9]. The data itself can be accessed by DNS-OARC members[10] and contains a lot of information although some of the data is obfuscated because of privacy requirements.

DITL data has been collected since 2006, although not all root server operators contribute for the complete two-day period or contribute every year. Also, root server operators do not always measure the full two-day period on all of their individual servers.

[7] https://www.dns-oarc.net/oarc/rssac002

[8] https://www.dns-oarc.net/oarc/rssac002

[9] http://www.tcpdump.org

[10] Subject to DNS-OARC’s data sharing agreement.

There are quite some tools available to process the data. The CDAR team used some of these tools to decide what type of data would be of interest for the study. These tools are very general and flexible in nature and often require quite some post-processing and data aggregation to present answers to questions such as "how many queries per second are received for these delegated domains".

After these requirements were established one of the CDAR partners wrote a set of targeted programs that collected and aggregated the data needed. This sped up the analysis of the DITL data and also simplified later (repeated) data processing and aggregation by the usual UNIX tools.

3.2.3 Root Zone Files Archive

DNS-OARC assembled a historical archive[11] of the DNS root zone files. The covered timeframe runs from 1999 till now, with typically one root zone file per day.

The Root Zone Files Archive data is available only to DNS-OARC members, either as raw zone files or as a Subversion repository. As DNS-OARC members, we downloaded and parsed the raw root zone files for this study.

There are gaps (missing days), especially before 2006, but also more recent: in the period 30 December 2012 until 21 September 2015, 44 days are not included in the archive. Dates suggested by the root zone file names do not (always) correspond to the actual date when the root zone file was published. However, the actual publish date can always be found in the serial number contained in the SOA record of the file. For example, the file with filename 'root-20130105.021102' refers to the root zone file with serial number 2013010301. Some files in the archive are duplicated under different names.

3.2.4 ICANN Registry Reports

For one of our analyses we used historic data regarding the number of domain names registered in various TLDs. For the delegated (new) gTLDs we used ICANN's monthly registry reports[12]. In particular, in the monthly transactions reports the data item "total-domains" denotes the total number of registered domains in the gTLD. ccTLDs do not report such data to ICANN, but most of them publish aggregate data such as the number of registered domain names themselves.

The monthly registry reports are provided by the TLD registry operators and are then published by ICANN, although their publication is withheld for three months due to contractual reasons. The reports trace back for many years (depending on the date of initial delegation). For example, for .com, registry reports are available from January 2001. We collected the domain name registration statistics for ccTLDs from the publication channels that are used by the ccTLDs. For these statistics we have a sufficiently long history of data for the purpose of our analysis.

The registry reports are published per month, which is granular enough for most TLDs because the number of registered domain names mostly fluctuates only a few percent in the time period of one month. For new gTLDs the number of registered domain names per month can fluctuate more. However, for the purpose of our analysis we consider the monthly registry reports to be granular enough.

[11] https://www.dns-oarc.net/oarc/data/zfr/root

[12] https://www.icann.org/resources/pages/registry-reports

3.3 Active Measurement Data Sets

3.3.1 RIPE Atlas data

The RIPE Atlas active measurement network from RIPE NCC consists of more than 9000 measurement probes that provide vantage points distributed around the globe. It is the largest measurement network for which data sets are publicly available.

All probes continuously measure all root DNS letters. RIPE uses different measurement IDs to identify each root server letter [23]. A subset of the probes are used for DNSMON [18] measurements. This study takes into account all probes and not only those employed on DNSMON. We use the RIPE Atlas data to analyze RTT performance and reachability (or the lack of reachability measured as queries that do not get responded to) of root DNS letters on a continuous time scale since the delegation of new gTLDs.

RIPE measurements are carried every 4 minutes[13]. The DNS requests sent by the probes are in the form of CHAOS queries that return the name of the server that responds to the queries.

We map all observations into a time series of 10 minutes. In each time bin we identify the root server letter and the response (either the anycast site or the error code). Each time represents 2.5 Atlas probing intervals, a similar approach used in a related study [2].

Most of the RIPE Atlas probes are located in Europe and North America. This, however, does not interfere with our study because we disregard probes that fail independently. Further, we remark that RIPE Atlas probes measure the RTT to the different root server letters, which is not the same as resolving a DNS query. A resolver resolving a query typically shows a more complex behavior, including selecting the root server letter for best performance.

3.3.2 DNSSEC Validation Tool

The new gTLD registries are contractually required to have at minimum the TLD itself protected by running the DNS protocol with the Security Extensions (DNSSEC).

Plain DNS is rather tolerant against errors and misconfigurations in the sense that "broken" setups are often not fatal. In general, the rigorous response validation done for DNSSEC makes the setup more brittle. Small data configuration errors will have more impact and can cause DNSSEC validation failures. These cause the resolver to decide that the whole domain is "bogus" or even doesn't exist. This can result in unstable DNS behavior.

In 2012 NLnet Labs started to monitor the status of all DNSSEC-signed TLDs, every twelve hours. The methodology used is primitive but rather effective. At first a list is generated of all TLDs that have DS (Delegation Signer) records in the root zone. For all the TLDs on this list we validate whether or not the DNSSEC chain of trust is established with a standard available program (unbound-hosts). If the validation fails, a notification is sent to the observer so the problem can be analyzed in due course. More details were presented at ICANN55 [24].

[13] With the exception of A-Root that was measured with 30 minute intervals until mid-2015.

3.4 Community Interaction

3.4.1 Dissemination Events

In Section 1.2 we stated that we regard it important to reach out to the ICANN and DNS communities (i.e., DNS-OARC, IETF/IEPG, RIR meetings) for this study. Therefore our consortium presented the study and results at a number of community interaction activities. Table 2 presents an overview of most of these interactions by the CDAR consortium.

Table 2: CDAR dissemination events

Event	Presented material	Date
ICANN 54 (Dublin)	Study plan for public comment[14]	20 October 2015
RSSAC conference call	Study plan and preliminary findings	7 January 2016
SSAC conference call	Study plan and preliminary findings	14 January 2016
ICANN 55 (Marrakech)	Study progress and preliminary findings[15]	8 March 2016
DNS-OARC 24 (Buenos Aires)	Intermediate findings and preliminary conclusion[16]	31 March 2016
IEPG (Buenos Aires)	Intermediate findings and preliminary conclusion[17]	3 April 2016
ENOG (Moscow)	Additional findings[18]	8 June 2016
SSAC retreat (Washington D.C.)	Study progress and additional findings	September 2016
RSSAC conference call	Study progress and summary of draft conclusion	6 October 2016
RIPE 73 (Madrid)	Impact of New gTLD on the Root System – Preliminary Results[19]	27 October 2016
ICANN 57 (Hyderabad)	Draft report for public comment	8 November 2016

3.4.2 Community Feedback

During the community interaction activities we received questions and feedback on the assumptions made in this study, and facilitated constructive discussion from ICANN's multi-stakeholder community.

In response to the call for public comments regarding the CDAR study plan[20], two commenters submitted their feedback. Both commenters emphasized acknowledging the limitations of the study, e.g. in terms of measuring performance of the root prior to the delegation of new gTLDs. Also, community feedback suggested that the study team should be careful with extrapolations of the results towards future scenarios. One of the commenters requested the identification of possible risk

[14] https://meetings.icann.org/en/dublin54/schedule/tue-root-stability
[15] https://meetings.icann.org/en/marrakech55/schedule/tue-root-stability-study
[16] https://indico.dns-oarc.net/event/22/session/1/contribution/28
[17] http://iepg.org/2016-04-03-ietf95/cdar-iepg.pdf
[18] https://www.enog.org/presentations/enog-11/167-cdar-enog.pdf
[19] https://ripe73.ripe.net/programme/meeting-plan/dns-wg/
[20] https://www.icann.org/en/system/files/files/report-comments-cdar-study-plan-17feb16-en.pdf

parameters. Further, the commenters emphasized the need for outreach to the broader DNS community to validate the methodology, the measurements and models. A response to these comments was published in [25]. We made an effort to process these remarks in the analyses and formulations in the draft report.

In the subsequent round for public comments on the CDAR draft report seven public comments were received. These comments included a number of questions for clarification, for example about how to perform the recommended continuous monitoring[21] of the security and stability of the root DNS system. Also, remarks were made about the relation between the objective results from the data analyses and the more speculative identification of risk parameters in Section 6. The comments on the draft report and the responses from ICANN and the CDAR team were published in [26]. These comments inspired several clarifications and modifications in this final report[22].

[21] In this final report we replaced the term '(more) continuous monitoring' used in the draft report with '(more) frequent monitoring', since that better fits our intended meaning.

[22] In addition to the modifications made in response to the public comments, minor modifications to this report also have been made based on additional analyses on more recent root DNS system data that has become available after publication of the draft report for public comments.

4 Results

This section presents our findings on how the stability and security of the root DNS system evolved from the beginning of the New gTLD Program. In Section 4.1, we discuss our analysis of the passive measurement data sets available from the root server operators and DNS-OARC (see Section 3.2) to get an “inside view” on the DNS queries that the root receives. Section 4.2 complements this analysis and focuses on end-user perception of the root behavior (query-response pairs) using the active measurement data sets of RIPE Atlas (see Section 3.3). In Section 4.3 we also use RIPE Atlas measurements to analyze the evolution of the size of the root DNS system in terms of the number of name servers that it contains. Finally, in Section 4.4 we will look at root DNS data consistency via the DNSSEC validation tool and by analyzing root zone files.

In this section we will briefly describe the analyses that we performed on the data sets, as well as the findings from these analyses. In the next section we discuss how these findings provide answers to our research questions.

4.1 Queries to the Root DNS System

We characterize the incoming queries at the root in terms of the total volume of queries (Section 4.1.1), as well as the breakdown into valid and invalid query volumes and the contribution of new gTLDs to these (Section 4.1.2). Further, we analyze the affinity to local DNS name servers of the queries for geographic new gTLDs[23] (Section 4.1.3) and we compare how query volumes to new gTLDs and other TLDs are distributed over RR types and protocols (Section 4.1.4).

4.1.1 Total Query Volume

We use the available RSSAC002 measurement data in order to analyze the dynamics in the total query volume over time:

Finding 1: The total number of queries to the root grows over time.

Analysis supporting Finding 1:
In Figure 7, we plot the daily number of queries per root server reported in the RSSAC002 measurements, by summing the different categories contained within the traffic-volume metric. For reference we also indicated the dates of the DITL periods. The observant reader can see a weekly recurring pattern in the query volumes, which is statistically significant in the traffic to the root DNS system.

[23] By geographic new gTLDs we mean a new gTLD denoting a geographical, geopolitical, ethnic, social or cultural representation, such as .tirol or .tokyo.

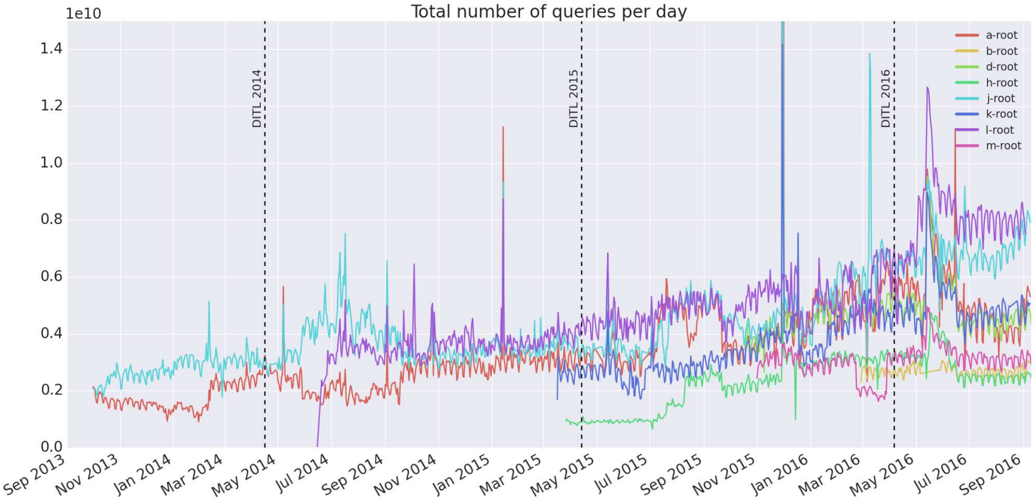

Figure 7: Daily number of queries from the RSSAC002 measurements

The daily number of queries shows an upward trend for most of the root servers. The number of queries to the root servers A and J (the only two root servers with RSSAC002 data available from November 2013 onwards) have tripled over three years. Note that the upward trends visible in Figure 7 may not only be due to an actual increase in the query volumes, but also due to the fact that some root servers have increased the coverage of their measurements over time. Also note that the query volumes for some root server letters show a significant decrease for some time periods. The RSSAC002 data does not allow us to find a decisive explanation for these observations.

In Figure 7 we also observe several peaks of query volumes. One of these peaks coincides with the extremely high, incidental query volumes on 30 November 2015 (reported in [13] and [4], amongst others). As far as we can deduce from reports about such incidental peaks, these are caused by other effects[24] than the existence of new gTLDs. For the purpose of this study we therefore distinguish the total query volume in 'regular' and 'incidental' traffic.

Moreover, we can *further distinguish the 'regular' traffic* in 'valid queries' (that we define as queries for names that have been delegated to the root zone) and 'invalid queries' (queries to names that have *not* been delegated to the root zone, which would be responded to with an NXDOMAIN by the root DNS system), as we will clarify in the following subsection. As a consequence, we characterize the total query volume by the equation: total query volume = valid (regular) queries + invalid (regular) queries + incidental queries.

4.1.2 Valid and Invalid Queries

To analyze the valid and invalid queries to the root, including the contribution of queries directed to new gTLDs, we make use of the DITL sets as well as other pcap data made available to us (see Section 3.2). These data sets allow us to distinguish

[24] An example of another effect is a DNS amplification attack whereby root DNS name servers can be used as reflection points.

the queries per TLD, and thereby extract the influence of specific TLDs on the query rate to the root. Even though DITL data sets are “snapshots” (they are collected only for two consecutive days per year), it is the only available data set that we can use to characterize the distribution of queries to the root servers. In absence of more detailed data we are left to presume that the results obtained on these data sets are representative for other dates.

Note that the *total number of queries*[25] to the different root server letters (split out per TLD) is difficult to extrapolate from the DITL sets with high accuracy, because

- the exact actual period being covered by measurements may differ per root server letter and per year;
- DITL periods cover different weekdays in different years, meaning that the weekly patterns can have an impact on the query volumes;
- individual root servers may capture pcap data at only part of their sites.

Therefore, in most of the following analyses we focus on the distribution (the fraction) of queries to different categories of TLDs, since the fractions are expected to be less sensitive to the fluctuations mentioned above.

For each DITL data set, we divide the queries into the following seven categories, depending on the name being queried.

The first three categories correspond to *valid queries*:

1. *Delegated new gTLDs*: Queries to new gTLDs that are delegated at the (end) time of the DITL set.
2. *.com:* Queries to .com.
3. *Other TLDs:* Queries to the delegated TLDs at the (end) time of the DITL set that remain when excluding .com and excluding the delegated new gTLDs.

The other four categories correspond to *invalid queries*:

4. *.home*: Queries to the non-delegated domain name home.
5. *.corp*: Queries to the non-delegated domain name corp.
6. *Non-delegated (potential) new gTLDs*: This category consists of a list of potential new gTLDs that are not yet delegated. It has been constructed as follows: We start with a large list of domains, consisting of all applied-for new gTLDs and all TLDs that are or have ever been part of the root zone. We then remove from this list all TLDs that have been delegated at the time of the specific DITL set, as well as home and corp.
7. *Invalid (remainder):* Queries to remaining names that do not correspond to delegated TLDs, excluding names that appear in the categories above.

[25] See Figure 18 in Appendix A for fluctuations in total number of queries between root server letters and DITL sets.

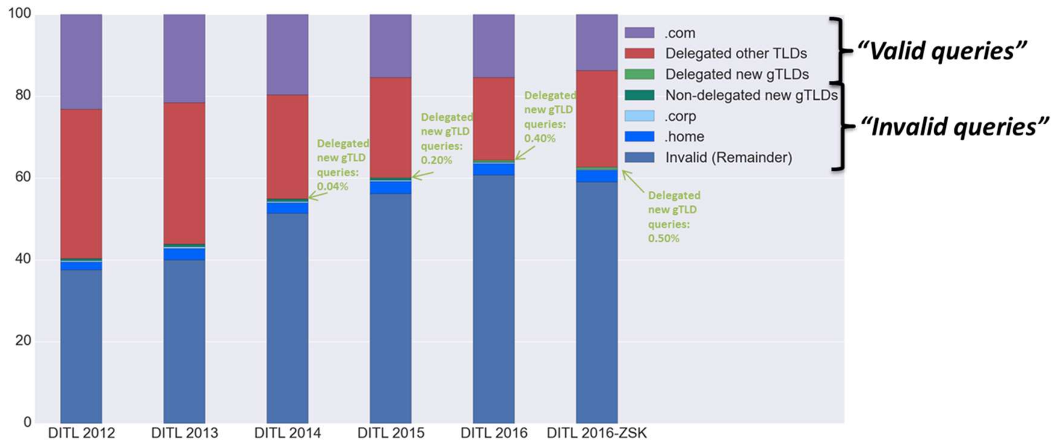

Figure 8: Evolution of the fraction of queries to the root distinguished in valid and invalid (new g)TLDs

Finding 2: The fraction of invalid queries (queries to invalid TLDs) increases significantly over time.

Analysis supporting Finding 2:
The fraction of invalid queries increases from 40.32% in the DITL period of 2012 to 63.93% in the DITL period of 2016, with a minor decrease to 62.09% in the additional ZSK-DITL round in October 2016 (see Figure 8). While this increase coincides with the delegation period of new gTLDs, we can see from historic DNS root studies that this trend seems to be[26] visible long before the delegation of new gTLDs: the DITL set of March 2009 indicated the percentage of queries to invalid TLD to be at around 30% [17][27], while in 2001-2002 this percentage for F-ROOT was around 20% [15], [16].

To investigate the impact of the delegation of new gTLDs on the evolution of valid and invalid queries in more detail we distinguish the analyses of valid and invalid queries in the remainder of this subsection.

Valid queries

Finding 3: Delegated new gTLDs attracted only a small number of DNS queries, which resulted in a negligible increase in the valid query rate to the root.

Analysis supporting Finding 3:
In this analysis we counted the number of queries in these DITL sets going to delegated new gTLDs and compared that to the number of queries going to all delegated TLDs (the 'valid' queries).

Figure 8 illustrates that while the fraction of queries to new gTLDs increases over time, it remains negligible compared to the total query rate (0.40% of total query rate in DITL 2016) and the valid query rate (1.10% of valid query rate in DITL 2016). The

[26] We have to be cautious here, because the DITL data collection data sets are not exactly comparable over the years, nor are the definitions of the valid and invalid queries.

[27] The percentage was derived from the "query validity" figure 6 in this paper by summing the query categories 'invalid TLD', 'A-for-A' and 'Unused query class'. This percentage is merely a rough estimate.

new gTLD with the highest query rate in DITL 2016 (.xyz) consists of 0.027% of total query rate in DITL 2016.

As a possible input for extrapolations to the future, we investigated potential relations between the number of domains within a TLD and the number of queries to that TLD:

Finding 4: The valid query rate for any TLD that is received by the root DNS system is 'bound' by the number of domains in that TLD. In general, this 'bound' appears to be lower for new gTLDs than for other TLDs.

Analysis supporting Finding 4:
As can be observed from previous findings, the rate at which TLD-related queries are sent to the root DNS system varies strongly among the TLDs. Statistics such as the "Most Popular TLDs Queried" graphs published by, for example, K-ROOT[28] support this observation. These statistics are collected using DSC, a system for collecting and exploring statistics from DNS servers (now developed by DNS-OARC).[29] The .com TLD is the "most popular" in terms of query rate to the root. This may not be surprising in the sense that this TLD is also "most popular" in terms of the number of domain names that are registered for this TLD.

Moreover, the "most popular" rankings in terms of the query rate to the root DNS system and the number of registered domain names appear to be similar. In Table 3 we present the number of queries received by K-ROOT during the DITL 2016 data collection period, for several different TLDs. In the third column the corresponding number of domain names registered for these TLDs is shown, at the time of the DITL 2016 period. The last column presents the ratios between these two values (in scientific notation).

Table 3: Number of queries (K-ROOT, DITL 2016) and number of domains

TLD	Nr. of queries/ TLD	Nr. of domains/ TLD	Query/domain ratio
.com	1 514 419 879	129 378 123	1,17 E+01
.net	777 776 811	16 221 406	4,79 E+01
.org	143 556 952	11 392 417	1,26 E+01
.cn	154 090 887	18 507 125	8,33 E+00
.br	49 465 477	3 788 150	1,31 E+01
.xyz	2 883 292	2 710 459	1,06 E+00
.top	1 497 090	1 818 175	8,23 E-01
.london	66 652	64 080	1,04 E+00

The 'query rate / domains' ratio in this table illustrates an interesting observation: the ratio appears to be relatively constant, regardless of how 'popular' the TLD is. While the number of queries (and the number of domains) per TLD can vary by a factor of thousands, the ratio only varies by a factor of tens.

In absence of a precise explanation for this observation, we verified this observation statistically with historic measurement data. We investigated this ratio for the DITL periods between 2013 and 2016, for all TLDs (that were delegated during the DITL periods) and all root-letters. It appears from these thousands of historic cases that the 'query rate / domains ratio' remains in the same order of magnitude[30] over time

[28] https://www.ripe.net/analyse/dns/k-root/statistics?type=ROOT&increment=daily&

[29] https://www.dns-oarc.net/tools/dsc

[30] i.e. the largest query rate / domains ratio is less than ten times the smallest ratio.

(the ratio is time-invariant) for almost all combinations of TLD and root-letter[31]. As such we can consider this to be a 'bound' on the query volume for each TLD that is dependent on the number of domains in the TLD. Moreover, up till now this 'bound' appears to be lower for new gTLDs than for other TLDs.

The relevance of this finding is that, presuming that this historic bound remains invariant in the near future, the number of (valid) queries for new gTLDs that are sent to the root DNS system is bound. More specifically, the query volume grows proportionally (within bound) with the size of the TLD, e.g. if it reaches the size of .com, the TLD will generate the same order of traffic at the root. In other words, *Finding 3* would remain valid.

The DITL data provides insight into the query rates to the root DNS system for a sequence of two-day sample periods in the relevant years for the delegation of new gTLDs (2012-now). In order to be able to analyze possible impact of the delegation of new gTLDs on a more continuous time-scale, for instance in the weeks after delegation of a new gTLD, we complemented the DITL analyses with analysis of a data set for a recent period of nine weeks (Beginning of November 2015 – Beginning of January 2016). This data set was collected by H-ROOT and therefore only includes queries to that root server letter.

This continuous data set enables us to analyze the impact of initial delegation of new gTLDs on the query rates to H-ROOT.

Finding 5: The impact of a delegation of a new gTLD to the root on the query rate to the root is microscopic in the period immediately after delegation. In most cases the number of queries for new gTLDs that were sent to the root even decreases after delegation.

Analysis supporting Finding 5:
We investigated the query rate fluctuations around delegation for all new gTLDs delegated in the H-ROOT renumbering period. We observed a limited number of typical patterns, which are illustrated in the four plots in Figure 9. These four TLDs were chosen because their query rates illustrate the different query rate patterns. To emphasize that the patterns apply to multiple new gTLDs, we anonymized these plots.

[31] See Figure 19 in Appendix B for an illustration.

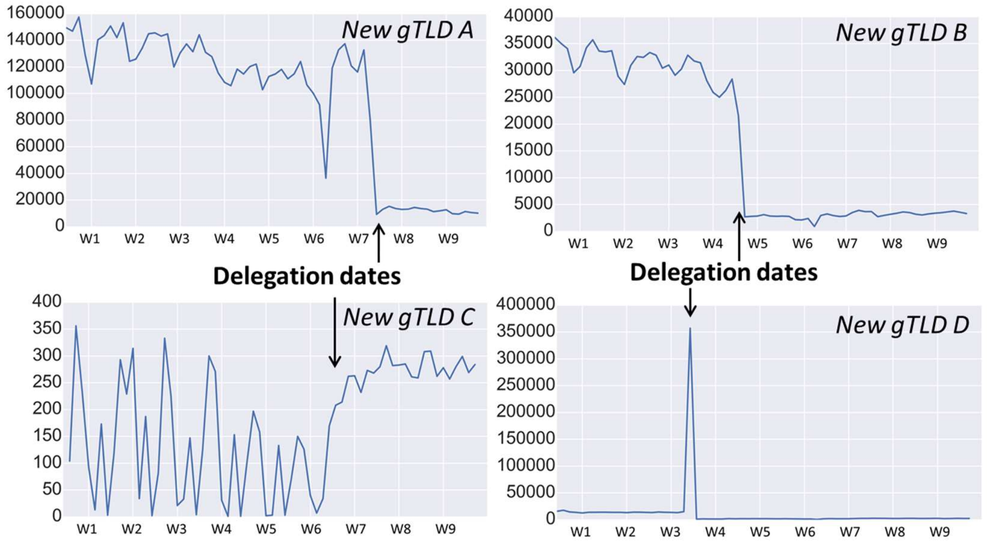

Figure 9: Daily query rates to H-ROOT for four selected new gTLDs, around their initial delegation dates, during a period of nine weeks

These plots illustrate that the volume of root traffic for a new gTLD may significantly decrease after delegation (gTLDs A and B). This is probably due to the effect of caching in recursive name servers. For some new gTLDs the query rate to H-ROOT increased (gTLD C) or increased temporarily (gTLD D) around delegation. But any increase is minor in comparison with the total query rate.

In total, 89 new gTLDs have been delegated in the period between 4 November 2015 and 7 January 2016 (the time period for which H-ROOT renumbering data was made available to us). When considering daily average query rates before, during and after delegation date for each of these new gTLDs, we observe the following:

- 64 of these 89 new gTLDs show a decrease in daily query rate after delegation. On average the query rate after delegation is 57% of the query rate before delegation.
- 25 of these 89 new gTLDs show an increase in daily query rate after delegation. On average the query rate after delegation is 240% of the query rate before delegation.
- In most of these cases the query rate converges to a new ‘steady state’; in general the delegations did not lead to heavy oscillations of query rates to H-ROOT *after* delegation.

In summary, the majority of new gTLDs delegated in this period show a decrease in daily query rates after delegation. Also, in general, it appears that the new gTLDs that do show an increase in daily query rate after delegation are new gTLDs with a relatively low daily query rate *before* delegation compared to those of the other new gTLDs (663 versus 4490).

Invalid queries

Analogous to the valid queries, an analysis of the DITL sets indicates that the addition of new gTLDs has a very slight influence on the invalid query rate to the root:

Finding 6: Applied-for but not-yet-delegated new gTLDs attracted a small number of DNS queries to the root, which resulted in a negligible increase in the invalid query rate to the root.

Analysis supporting Finding 6:
Based on Finding 5 one might conjecture that the New gTLD Program leads to many queries to not-yet delegated new gTLDs. This would then cause an increase in the invalid query rate. We can investigate this using the *Non-delegated (potential) new gTLDs* category, including *.home* and *.corp*, the two most queried-for non-delegated names. While *.home* and *.corp* have been applied for as new gTLDs, they will not be delegated in the current round because they have been classified as strings with a high-risk for name collisions[32].

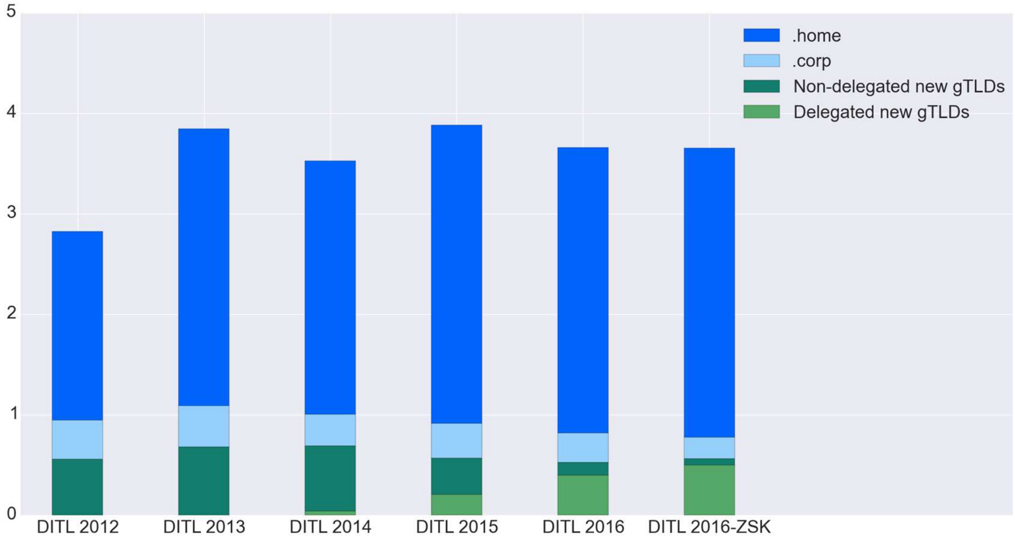

Figure 10: Percentage of queries for new gTLDs to the root DNS system distinguished in delegated and non-delegated new gTLDs

The total of queries to (potential) new gTLDs stays below 4% of all queries for each DITL set (see Figure 10). Also, no increasing trend is visible in these five years. When excluding *.home* and *.corp*, the remainder of queries appear to be quite constant over time: between 0.53% and 0.69% (of all queries), with a slight increase between 2012 and 2014, and a slight decrease between 2014 and 2016. Furthermore, note that, by Finding 2, the number of invalid queries is increasing, so these potential new gTLDs actually decrease relative to all invalid queries. These numbers indicate that the addition of new gTLDs to the root has a negligible impact on the total invalid query rate to the root.

4.1.3 Geographical Affinity

A portion of the applied-for new gTLDs are so-called GeoTLDs: a TLD denoting geographical, geopolitical, ethnic, social or cultural representation[33]. Such GeoTLDs might generate geographically uneven distributed query volumes to the root DNS system. Partly inspired by a remark from the Business Constituency during the public

[32] https://www.icann.org/en/system/files/files/new-gtld-collision-mitigation-05aug13-en.pdf
[33] https://icannwiki.com/GeoTLD

comment period for the CDAR study plan[34], we investigated whether there is any geographic (localized) affinity in terms of traffic for geographic new gTLDs.

Finding 7: Although there is geographic affinity for geographic new gTLDs, at each root name server node the fraction of traffic to such geographic new gTLDs is insignificant.

Analysis supporting Finding 7:
Based on the list of new gTLDs that were identified from the list of geographic applications on the ICANN wiki, we selected the following delegated TLDs for analyzing the impact of geographic new gTLDs on the geographic distribution of DNS queries to root server nodes (note that all continents are covered in this selection):

Table 4: Selection of geographic new gTLDs

TLD	Delegation date	Area type	Country
bayern	3 May 2014	Region	Germany
capetown	19 June 2014	City	South Africa
doha	25 March 2015	City	Qatar
london	22 March 2014	City	United Kingdom
melbourne	10 July 2014	City	Australia
moscow	24 April 2014	City	Russia
nyc	20 March 2014	City	USA
rio	22 May 2014	City	Brazil
sydney	5 November 2014	City	Australia
tirol	4 June 2014	Region	Austria
tokyo	29 January 2014	City	Japan
vlaanderen	18 June 2014	Region	Belgium
xn--80adxhks	24 April 2014	City	Russia

We considered F-ROOT and L-ROOT. These two are a subset of root server letters whose individual server nodes could be mapped to specific geographic locations (cities). Furthermore, we focused on DITL 2015 and DITL 2016 since all TLDs in Table 4 were delegated to the root zone before DITL 2015. For each of the TLDs in the table above we want to know for a given server node location whether the fraction of queries arriving in this location that query the TLD is more or less than the fraction of queries for the whole root server that query the TLD. We do this by dividing the former by the latter.

As an example, Table 5 indicates a factor of 2.14 for TLD *.bayern* and server node location Frankfurt. This means that the fraction of queries in Frankfurt that query *.bayern* is 2.14 times bigger than the fraction of all F-ROOT queries that query *.bayern*.

[34] https://www.icann.org/en/system/files/files/report-comments-cdar-study-plan-17feb16-en.pdf

Table 5: F-ROOT, DITL 2016: How much more relative localized traffic to TLDs compared to the average relative traffic to TLDs.

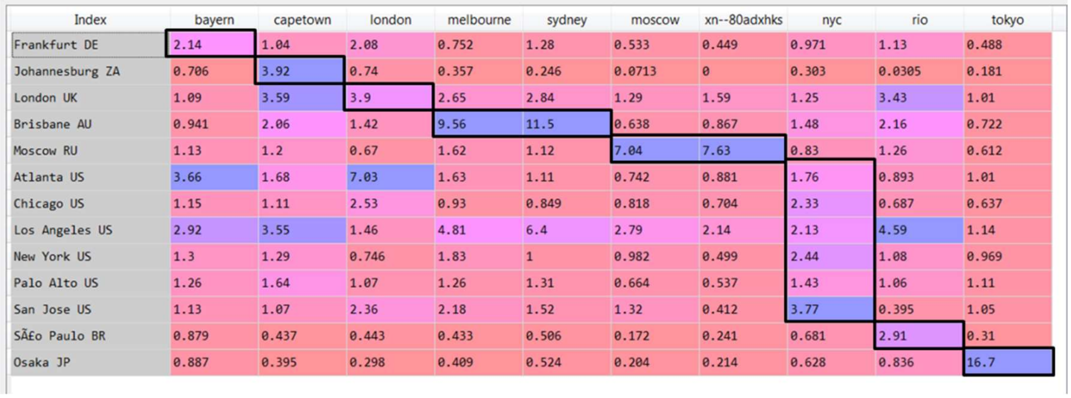

Index	bayern	capetown	london	melbourne	sydney	moscow	xn--80adxhks	nyc	rio	tokyo
Frankfurt DE	2.14	1.04	2.08	0.752	1.28	0.533	0.449	0.971	1.13	0.488
Johannesburg ZA	0.706	3.92	0.74	0.357	0.246	0.0713	0	0.303	0.0305	0.181
London UK	1.09	3.59	3.9	2.65	2.84	1.29	1.59	1.25	3.43	1.01
Brisbane AU	0.941	2.06	1.42	9.56	11.5	0.638	0.867	1.48	2.16	0.722
Moscow RU	1.13	1.2	0.67	1.62	1.12	7.04	7.63	0.83	1.26	0.612
Atlanta US	3.66	1.68	7.03	1.63	1.11	0.742	0.881	1.76	0.893	1.01
Chicago US	1.15	1.11	2.53	0.93	0.849	0.818	0.704	2.33	0.687	0.637
Los Angeles US	2.92	3.55	1.46	4.81	6.4	2.79	2.14	2.13	4.59	1.14
New York US	1.3	1.29	0.746	1.83	1	0.982	0.499	2.44	1.08	0.969
Palo Alto US	1.26	1.64	1.07	1.26	1.31	0.664	0.537	1.43	1.06	1.11
San Jose US	1.13	1.07	2.36	2.18	1.52	1.32	0.412	3.77	0.395	1.05
SÃ£o Paulo BR	0.879	0.437	0.443	0.433	0.506	0.172	0.241	0.681	2.91	0.31
Osaka JP	0.887	0.395	0.298	0.409	0.524	0.204	0.214	0.628	0.836	16.7

In the table above we choose only those F-ROOT server locations that coincide with one of the countries in Table 4. As can be seen in the delineated rectangles, there tends to be a geographic affinity for the geographic TLDs: for each of the TLDs, the fraction of queries to these TLDs in server node locations within the countries indicated by the TLDs is higher than average. Sometimes this increase may be relatively large: for .tokyo the fraction of queries in Osaka is 16.7 times the average fraction of queries to .tokyo. Actually, in the data for F-ROOT, DITL 2016 we see that the highest fraction to any of the chosen TLDs occurs indeed for .tokyo in Osaka: 0.015%. Note that this remains an insignificant portion of the total traffic in Osaka. An additional analysis presented in Appendix B indicates that in general the fraction of traffic to such geographic new gTLDs remains an insignificant part of the total traffic to individual server nodes.

4.1.4 Query Type Distribution

Below we show some findings on the distribution of query types. One of the reasons why we are interested in this is that the query type may have an influence on the load and response time performance of a root name server. For example, [27] (section 4.2) and [28] (section 5) provide considerations about increased latency, higher overhead and server memory demand in case TCP is used, relative to UDP[35]. Investigation of such query type dependent performance effects is beyond the scope of this study. We restrict ourselves to analyzing the evolution of the distribution over query types and, in particular, if this distribution is distinct for new gTLDs versus other TLDs.

Finding 8: The fraction of queries that use the TCP-protocol increases over time, but remains a small portion of the total number of queries to the root.

Analysis supporting Finding 8:

Using the DITL data, we see the following percentage of queries using the TCP-protocol over time: 0.34% (2012); 0.26% (2013); 0.34% (2014); 1.03% (2015); 1.91% (2016). From Figure 1 in [17] it can be derived that historically the percentage of queries using the TCP-protocol was even lower: 0.018% (2007); 0.026% (2008); 0.037% (2009).

[35] The latency impact of using TCP versus UDP is for example monitored by the RIPE Atlas monitoring framework: https://atlas.ripe.net/results/maps/root-server-performance/

A significant increase is visible from 2013 onwards, which coincides with the delegation of new gTLDs. However, the following finding indicates that this increase is not *because* of the new gTLDs.

Finding 9: The fraction of queries towards delegated new gTLDs that use the TCP-protocol is lower than that fraction of queries towards other delegated TLDs.

Analysis supporting Finding 9:
Figure 11 indicates the fraction of queries corresponding to UDP/TCP queries and IPv4/IPv6. Three different categories are considered here:

- *New gTLDs:* Queries to new gTLDs that are delegated at the time of the DITL set.
- *Other TLDs:* Queries to the delegated TLDs at the time of the DITL set that remain when excluding the delegated new gTLDs.
- *Invalid:* Remaining queries (to non-delegated TLDs).

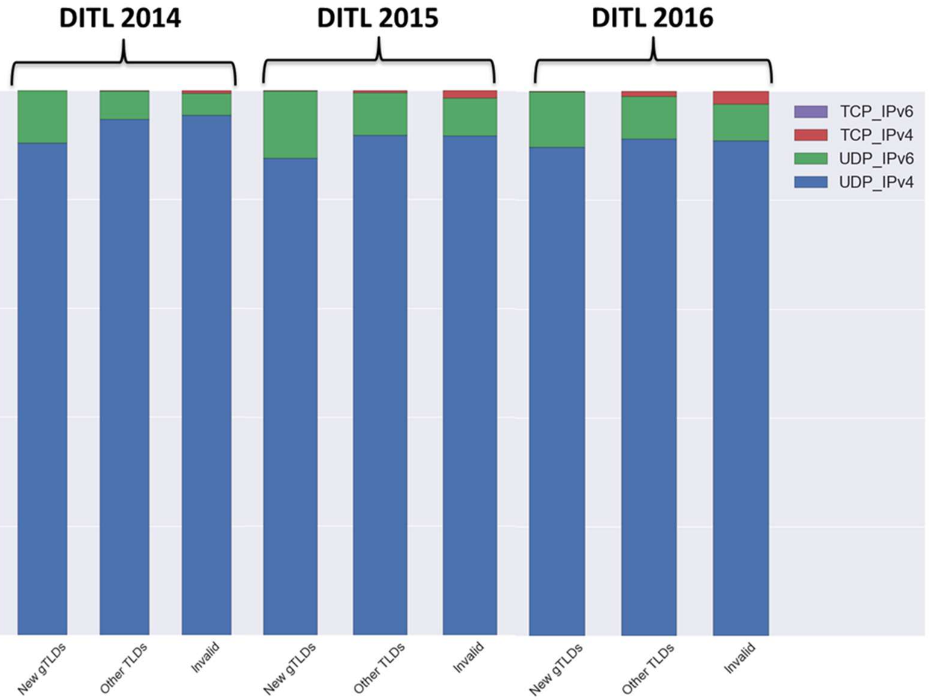

Figure 11: UDP/TCP and IPv4/IPv6 distribution

For all indicated years, less than 0.5% of the queries to new gTLDs use the TCP-protocol.

Finding 10: Some differences between the type distribution of delegated new gTLDs and the type distribution of other delegated TLDs are visible, but the type distributions seem to converge over time.

Analysis supporting Finding 10:
From Figure 12 below, which shows a breakdown of query types, some observations can be made:

- Initially (DITL 2014) the fraction of queries to new gTLDs of A and AAAA type is much lower than the fraction of queries to other TLDs of A and AAAA type.

This fraction increases for queries to new gTLDs over time. A possible explanation for this observation is the low initial number of domains within the new gTLDs.

- Initially (DITL 2014) a much larger fraction of queries for new gTLDs have type DS and DNSKEY compared to queries to other delegated TLDs. These fractions seem to converge to the distribution for other TLDs over the DITL periods. The fact that DNSSEC-related records such as DS and DNSKEY are required for new gTLDs, while more and more of the other TLDs are being signed, could be one explanation of this observation. An additional explanation will be the previous observation: the fact that the fraction and volume of queries for A and AAAA records is increasing, decreases the relative portion of queries for other records.

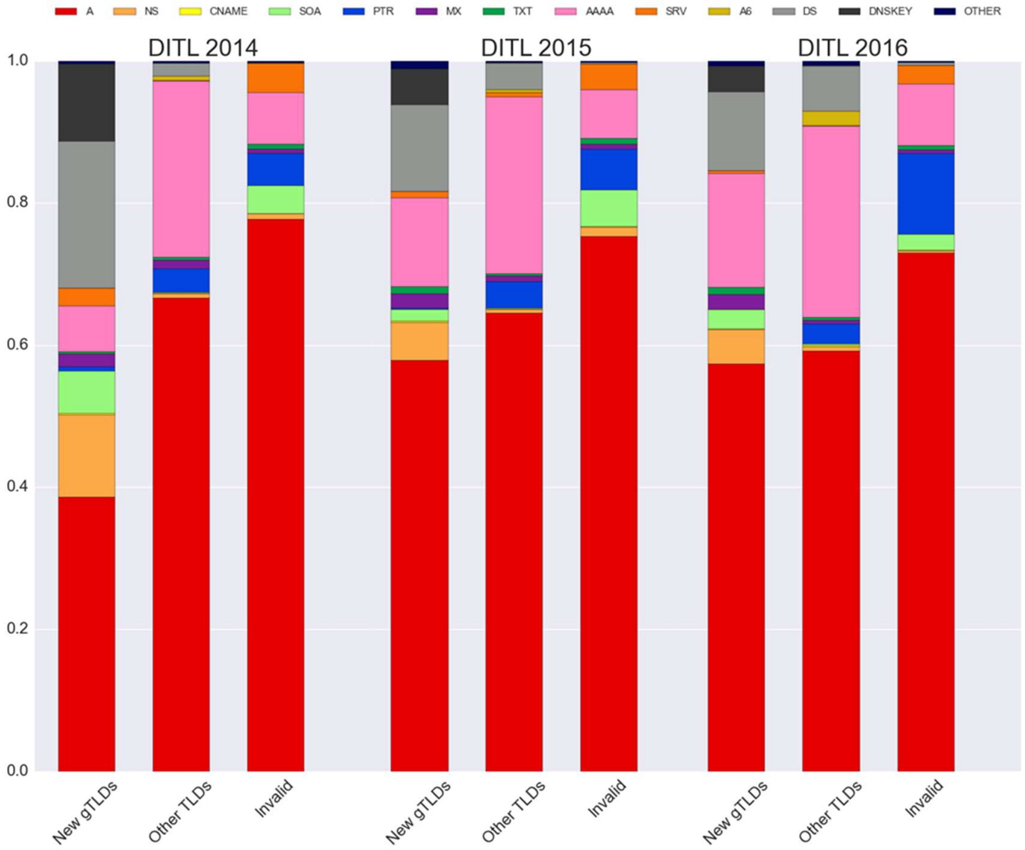

Figure 12: Breakdown into query types

4.2 Query-Response of the Root DNS System

For an analysis of the impact of new gTLDs on performance and reachability of the root DNS system as experienced by users, we will use a subset of the available RIPE Atlas network measurement data (see Section 3.3). These are measurements that are carried out not at the root servers themselves, but via Atlas probes towards the root servers, behaving like potential users/resolvers.

As we discussed in Section 2.1, the root DNS system is constantly changing. More anycast sites are added, more network links, more or higher capacity servers. Since most of these details are not public, we cannot determine precisely when and how these changes occur in the root DNS system, which is likely to affect results in reachability and RTT performance as observed by RIPE Atlas data. To avoid this issue, we used the following approach:

- Choose a date in which seven or more new gTLDs have been delegated.
- Analyze RTT performance and reachability two days before and two days after the delegations, and observe any significant changes.
- Repeat this procedure for many dates.

By choosing a time window of four days (-2 days, +2 days) around each delegation, we narrow down the time in which the root DNS system may have changed and assume it as stable, and therefore present our analysis.

Metrics:
We devise two main metrics for the analyses:

- *RTT Performance*: defined as the round-trip time (RTT) which is the time difference between the DNS request and the response, measured at the sender.
- *Reachability*: measured in the fraction of RIPE Atlas probes that obtained an answer from each letter from the root DNS system.

These metrics are important for the following: if by any means the delegation of new gTLDs leads to any sort of instability on any root server letter, by any unforeseen way, we would be able to notice that in terms of RTT and reachability. If a root server letter becomes overloaded, RIPE Atlas probes would experience changes in their RTT. If a root server letter becomes unresponsive, its reachability would be significantly affected.

This behavior is not only theoretical, it has been in fact observed in the root DNS system when they were target of a 100 times the average load DDoS attack: some root server letters became unreachable; some had severe performance degradation [4].

4.2.1 RTT Performance

Finding 11: The overall RTT (performance) is not significantly influenced by the delegation of new gTLDs to the root zone.

Analysis supporting Finding 11:
Using public RIPE Atlas data, statistics were gathered regarding the mean RTT per hour, the median RTT per hour and the 90th percentile RTT per hour of queries to different root server letters in periods of four days surrounding the delegation of at least seven new gTLDs.

For some periods and/or root server letters we did not manage to obtain sufficient RIPE Atlas data per hour. We only select those periods where for at least some root server letters we obtained at least 50 000 measurements for each hour, and show those in the table below. In total, we analyzed data comprising 22 different date points in which new gTLDs were delegated. Our hypothesis is that if this increase would lead to any sort of instability, it would be noticed in multiple dates shown in Table 6.

Table 6: Selected periods for performance and reachability analysis

Date (number of TLDs delegated)	Day	Period
2013/12/17 (+ 17)	Tue	2013/12/15 - 2013/12/18
2013/12/28 (+ 19)	Sat	2013/12/26 - 2013/12/29
2014/01/23 (+ 10)	Thu	2014/01/21 - 2014/01/24
2014/02/04 (+ 10)	Tue	2014/02/02 - 2014/02/05
2014/03/31 (+ 12)	Mon	2014/03/29 - 2014/04/01
2014/04/11 (+ 15)	Fri	2014/04/09 - 2014/04/12
2014/04/23 (+ 15)	Wed	2014/04/21 - 2014/04/24
2014/05/15 (+ 9)	Thu	2014/05/13 - 2014/05/16
2014/05/22 (+ 8)	Thu	2014/05/20 - 2014/05/23
2014/05/31 (+ 10)	Sat	2014/05/29 - 2014/06/01
2014/06/19 (+ 7)	Thu	2014/06/17 - 2014/06/20
2014/07/18 (+ 9)	Fri	2014/07/16 - 2014/07/19
2014/08/16 (+ 10)	Sat	2014/08/14 - 2014/08/17
2014/08/30 (+ 10)	Sat	2014/08/28 - 2014/08/31
2014/09/15 (+ 9)	Mon	2014/09/13 - 2014/09/16
2014/10/15 (+ 8)	Wed	2014/10/13 - 2014/10/16
2014/12/13 (+ 7)	Sat	2014/12/11 - 2014/12/14
2015/01/24 (+ 11)	Sat	2015/01/22 - 2015/01/25
2015/06/22 (+ 7)	Mon	2015/06/20 - 2015/06/23
2015/06/26 (+ 9)	Fri	2015/06/24 - 2015/06/27
2015/07/08 (+ 7)	Wed	2015/07/06 - 2015/07/09
2015/11/25 (+ 7)	Wed	2015/11/23 - 2015/11/26

As an illustration we plot the median RTT per hour for the root server letters for two representative dates:

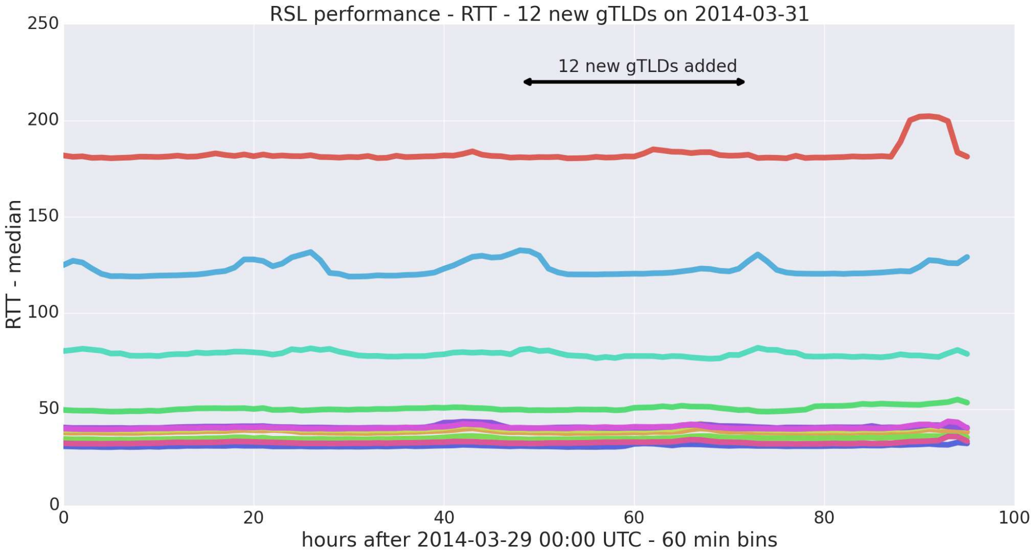

Figure 13: Median RTT per hour per root server letter: 12 new gTLDs delegated on 31 March 2014

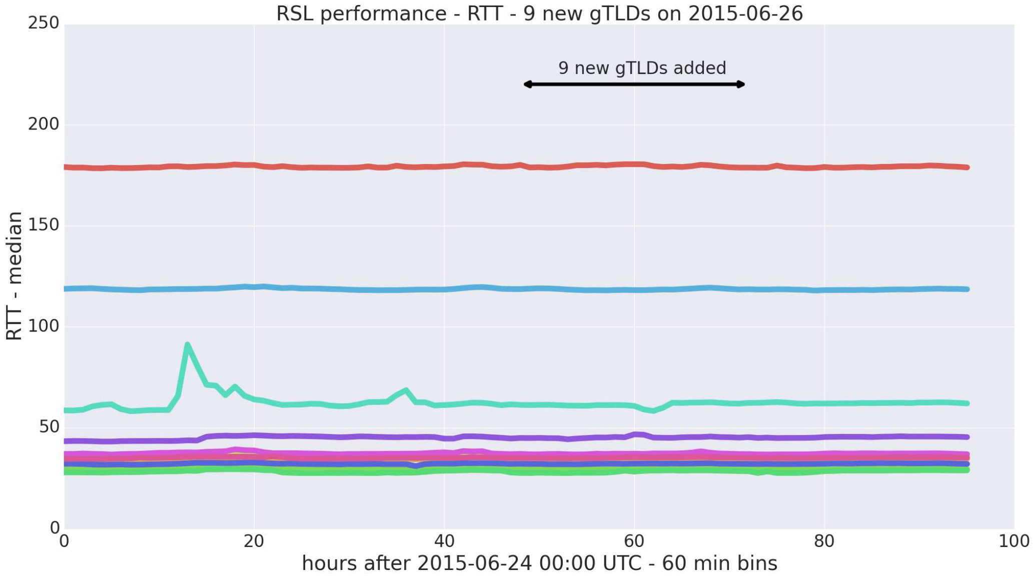

Figure 14: Median RTT per hour per root server letter: 9 new gTLDs delegated on 26 June 2015

Figure 13 and Figure 14 indicate that the (median) RTT is not significantly affected by the delegation of new gTLDs. We support this observation by the following statistical analysis.

When comparing hourly statistics, we do not want the number of measurements per hour to fluctuate too much. We therefore only consider those combinations of period and root server letter where the number of measurements per hour does not differ from the maximum number of hourly measurements by more than 5%. This excludes some combinations of period and root server, including all of the 2014/01/23-period. 208 different combinations of period and root server letter remain in our analysis.

- Statistically, using both the Welch's t-test and the Student t-test we observe that in 71% of the cases the mean RTT per hour does not change significantly between the 48-hour period before delegation and the 48-hour period after delegation.
- Of the 29% of the cases where there *is* a *statistically* significant change of mean RTT per hour, 61% concerns a *decrease* of mean RTT and 39% an *increase.* So no clear trend of an increase or decrease of RTT is visible.
- Furthermore, in those 29% of the cases of statistically significant change, the change in mean RTT is in general still very small: *less than 5% in the vast majority of cases.*
- The median and 90th percentile RTT may both increase and decrease after delegation with almost equal probability, and the size of this increase or decrease is in general very small.

4.2.2 Reachability

Finding 12: The fraction of answered queries (reachability) is not significantly affected by the delegation of new gTLDs to the root zone.

Analysis supporting Finding 12:

We perform a similar analysis on the processed RIPE Atlas data as in the analysis supporting Finding 11, but then on the fraction of RIPE Atlas probes that obtained an answer from the root DNS system (reachability). Figure 15 and Figure 16 illustrate the reachability per hour for the root server letters for two of the chosen dates:

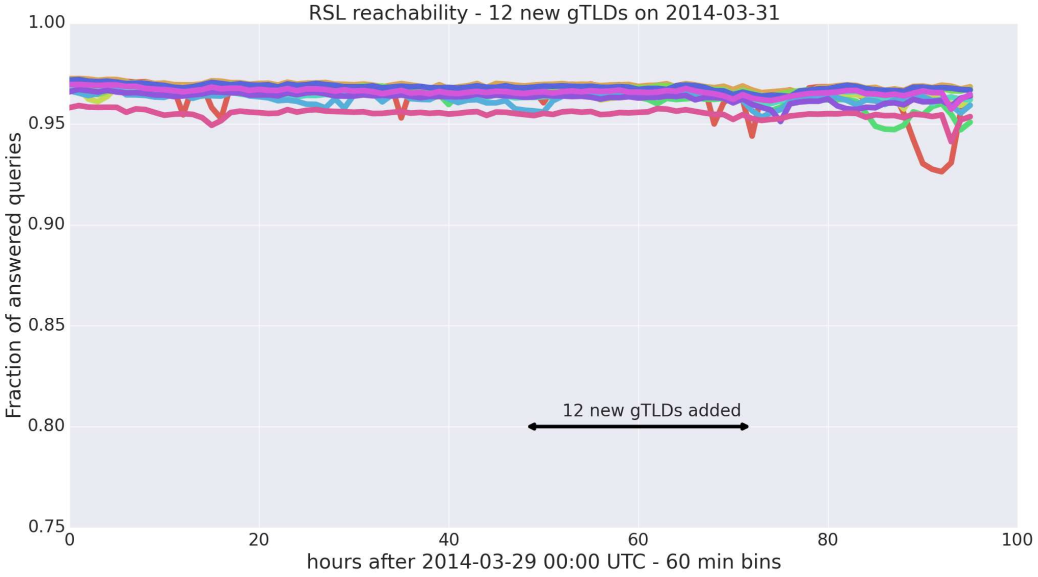

Figure 15: Reachability per hour per root server letter: 12 new gTLDs delegated on 31 March 2014

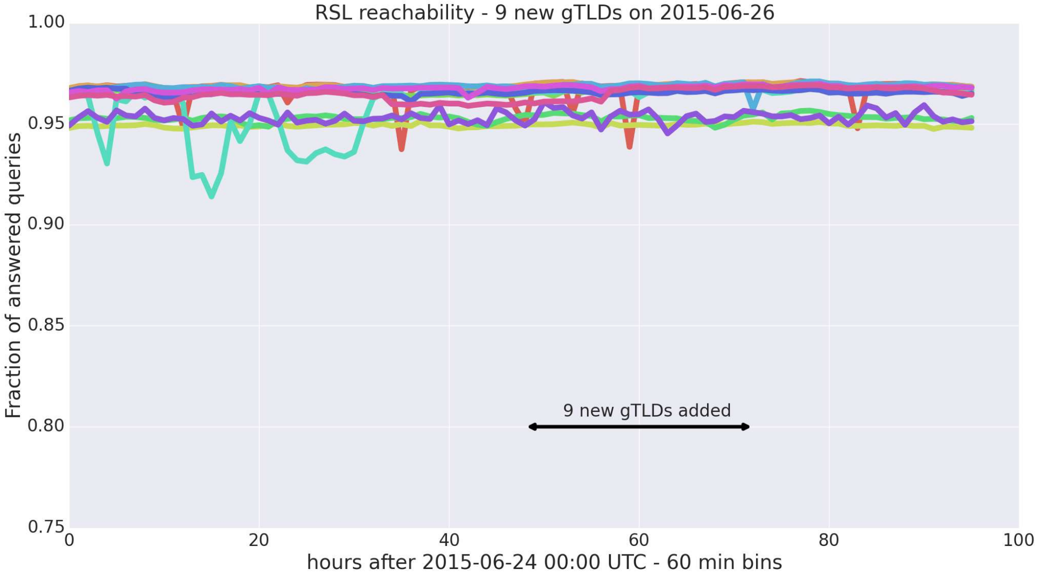

Figure 16: Reachability per hour per root server letter: 9 new gTLDs delegated on 26 June 2015

The figures above indicate that the reachability does not change after delegation of new gTLDs. This is supported by considering all cases of chosen dates and root servers:

- In 51% of the cases, the fraction of answered queries per hour averaged over the 48 hours after delegation decreases compared to the average over the 48 hours just before delegation. In the other 49% of the cases the average fraction of answered queries increases.
- In almost all cases, the difference in the fraction of answered queries remains minor.

Since the reachability changes both up and down with almost equal probability, it seems that changes in reachability are due to factors other than the delegation of new gTLDs.

Therefore, we can conclude that, judging from the data points we chose to carry out these analyses, the delegation of new gTLDs does not have any impact as observed by users/resolvers.

4.3 Size of the Root DNS system

As already indicated in Section 2.1, the root DNS system is continuously evolving. As input for reflections on extrapolation of this evolution (in Section 6) we provide a rough indication for the growth of the root DNS system.

Finding 13: The size of the root DNS system (in terms of number of anycast sites) increases over time.

Analysis supporting Finding 13:

While most of the internals of the root DNS systems are kept private, some statistics are public [3]. Even though we cannot make statements about the evolution of more accurate capacity indicators of the root server letters (e.g. in terms of bandwidth and processing power), we can see the number of public anycast sites as a rough indication of how the root DNS system evolves.

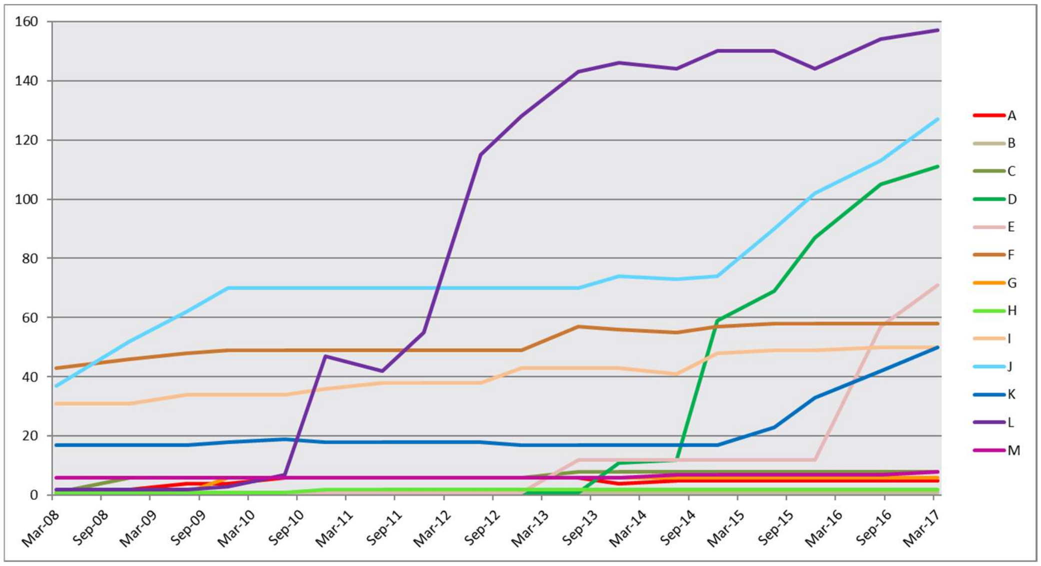

Figure 17: Number of anycast sites reported on root-servers.org

Figure 17 shows the growth over the last eight years, obtained from [3] (using the Wayback Machine[36] for historic counts). As can be seen, several of the letters have shown a considerable increase in anycast sites, most notably in the last two years.

4.4 Root DNS Data Consistency

4.4.1 DNSSEC Broken Chain Analysis

Finding 14: The delegation of new gTLDs has not contributed to a significant increase in DNSSEC validation errors between the root zone and the Top Level Delegations.

Analysis supporting Finding 14:
Using the DNSSEC validation tool described in Section 3.3.2 we noticed about 550 validation failures since 2012. This number is dominated by more than 400 repeated notifications for the same TLD.

Most other problems were resolved within 24 hours. The errors followed similar patterns. Often we observed that signatures were expired and this repeated sometimes on a regular basis. Other errors were caused by non-matching algorithms between the parent (root) and child (TLD) zones. This points to an improper roll over of signing algorithms. Also, signatures sometimes disappeared, probably due to an improper update of the TLD zones.

When we classify the errors by nature of TLDs, we see that there are a lot of repeated failures with ccTLDs. We suspect that this is due to the fact that the operation of (small) ccTLDs is often done by hand and thus caused by human errors. For "non-new" gTLDs most errors were resolved very quickly and might just as well be caused by incidental network errors during the probing. Errors with new gTLDs often happened only once when they just got delegated, which might be due to problems with the delegation starting up.

More details can be found in the presentation [24] on this subject in the DNSSEC workshop at ICANN 55.

In the last year we noticed that the DNSSEC validation errors became less frequent. This is probably due to improved monitoring by the TLD operators and improvements of tools available to automate the processes needed to run DNSSEC signed zones.

4.4.2 Root Zone Files Correctness

In [9] and [10], for example, it is indicated that corruption of DNS data can have impact on DNS security as well as DNS stability. Based on this identified concern we analyzed if the delegation of any new gTLD has had an impact on the correctness of data in the root zone file.

Finding 15: The delegation of new gTLDs has not contributed to errors in the root zone files.

[36] http://web.archive.org/web/*/www.root-servers.org

Analysis supporting Finding 15:
We parsed all root zone files stored in the root zone file archive for the period from 2012 until 2015. We encountered no syntax errors, apart from some differences with respect to delimiters in different files (tabs versus spaces). Although this verification is not conclusive in itself, it does indicate that the delegation of new gTLDs has not contributed to errors in root zone files.

This result is expected, because the process of preparing and distributing each root zone file includes validation steps by the root zone maintainer[37]. Such validation steps are aimed to eliminate errors in the root zone files, before they are published. As Finding 14 and Finding 15 indicate, these validations appear to work well, and the delegation of new gTLDs does not seem to affect that.

[37] See for example http://yazvs.verisignlabs.com/

5 Conclusion

In Section 4 we presented the findings that resulted from the individual experiments, measurements and analysis of empirical data, covering the period of September 2013 through September 2016. In this section we combine these findings and relate them to the security and stability metrics we formulated in Section 1 and Section 3. We revisit our security and stability metrics (Section 5.1) and discuss their main categories: query rate stability (Section 5.2), query-response stability (Section 5.3), and data consistency (Section 5.4). Finally, in Section 5.5 we summarize our overall conclusion.

5.1 Security and Stability Metrics Revisited

Our analysis focused on answering the first main research question of the CDAR study: *Has the introduction of new gTLDs degraded the stability or security of the root DNS system?*

Our approach was to answer this research question using the root DNS security and stability metrics we presented in Figure 4 (repeated below), using the methodology we outlined in Section 3 and taking into consideration the scope and limitations of our study as discussed in Section 1.3.

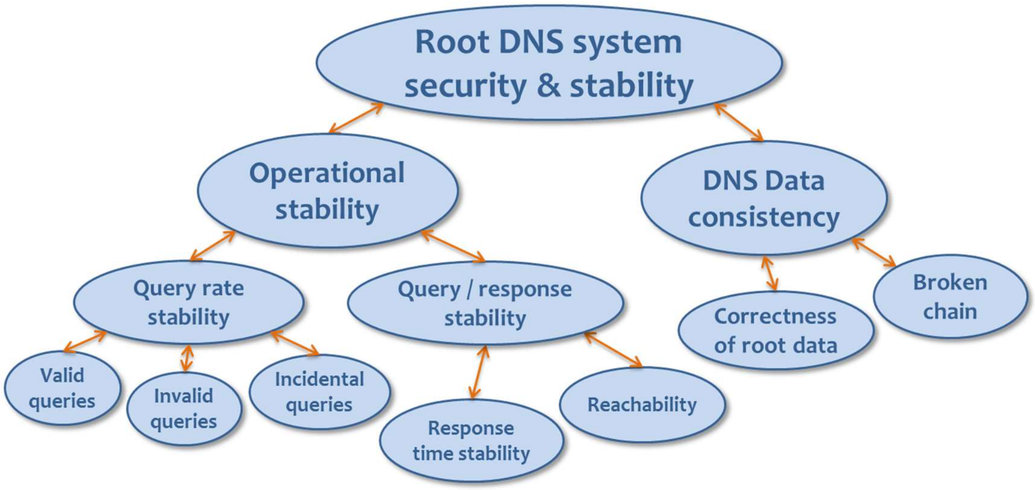

Repeated Figure 4: root DNS system stability & security metrics

We have used this breakdown into metrics to refine the primary research question into more detailed research questions. In particular, whether the introduction of new gTLDs has led to:

- an increase of the query rate towards the root DNS system?
- a degradation of query / response stability?
- an increase of incorrect DNS data, broken (DNSSEC) chain events or other data inconsistencies due to root DNS data?

In the following subsections we relate the findings presented in Section 4 to these three more detailed research questions.

5.2 Query Rate Stability

Summarizing we can state that the total number of queries to the root DNS system has been growing at a gradual rate (Finding 1). At the same time the number of anycast sites deployed by the root server letters (as a rough indication of the capacity of the root DNS system to respond to those queries) has also increased (Finding 13).

More detailed analysis of the total query volume to the root DNS system requires a distinction between valid and invalid queries, complemented with incidental peaks of query volumes during rare events. In the total query volume the fraction of queries for invalid TLD names has increased over time (Finding 2) and forms the majority in 2016.

Regarding the valid query volume the contribution of queries to new gTLDs has been insignificant (Finding 3) in comparison to the total number of valid queries. Moreover, the valid query volume seems to be bound by the number of registered domain names in the TLD (Finding 4). Although our empirical analyses did not explain why this statistical bound exists, it may be a useful result as an indicator for the future growth in DNS traffic (see Section 6).

We have found no empirical evidence of an impact of the delegation of new gTLDs on incidental or invalid query volumes (Finding 5 and Finding 6, respectively). Although Finding 5 allows us to exclude the existence of a significant impact to incidental query volumes surrounding delegations, this does not mean that we can exclude *any* impact. Also, the contribution of newly delegated gTLDs to invalid queries is somewhat ambiguous: while we observe that the query volume for ‘applied-for, but non-delegated new gTLDs’ is actually higher compared to the query volume for delegated new gTLDs (though still small compared to the total query volume), it is dominated by queries for the invalid gTLDs .home and .corp. These query volumes have been large for years before the New gTLD Program was introduced, and are therefore most likely unrelated to the applied-for new gTLDs.

Apart from the query volume, the type of queries and the protocols that are used (e.g. UDP/TCP) may also have an impact on the root DNS system’s ability to respond. To this end we investigated if the distribution of queries for new gTLDs is significantly distinct from the distributions of queries for other TLDs. We investigated the distribution of queries over geographic regions (Finding 7), over query types (Finding 10) and the protocols used (Finding 8 and Finding 9). These findings show that the distributions for new gTLDs do not significantly differ from the distributions of other TLDs, or at least they do not contribute to higher load on the root DNS system.

In summary, we have found no evidence that the introduction of new gTLDs has incurred a significant increase in traffic that impacts the root. Also, we did not find a significant change in the distribution of different types of queries compared to existing TLDs. The conclusion is that if upcoming new gTLDs exhibit a similar query pattern after delegation as the currently delegated new gTLDs, we expect the root DNS system to continue to operate without any problems.

5.3 Query-Response Stability

Complementary to the analysis of queries towards the root DNS system we analyzed the query-response behavior. By measuring the query-response behavior from the outside of the root DNS system we can get an indication of the impact of the introduction of new gTLDs on internet user experience.

The key user experience indicators that we analyzed are the RTT (Round Trip Time) performance and the reachability (i.e. the fraction of queries that result in a response) of the root DNS system.

Analysis of the active measurements shows that the overall RTT performance is not significantly influenced during the introduction period of new gTLDs. Even when we zoom in on specific days where relatively many new gTLDs are delegated to the root zone (22 days in total), we found no evidence of significant changes to the RTT performance (Finding 11).

The same observation holds for the fraction of queries that do not result in a response from the root DNS system: the reachability is not significantly degraded by the delegation of new gTLDs to the root zone (Finding 12).

Thus there is no evidence that, from the user's point of view, the introduction of new gTLDs has caused any impact on the performance and reachability of the root DNS system. Even during periods of strongly increased query volumes it is observed that the system as a whole remains stable, although some parts of the system suffer from degraded performance and reachability (as concluded in reports from other studies, such as [4]). Such rare events demonstrate that there is a correlation between query rate stability and query / response stability, and therefore it is not surprising that our findings and conclusions for these metrics coincide.

5.4 DNS Data Consistency

Where the findings and conclusions in the previous subsections are focused on the operational stability of the root DNS system, the third category of metrics is focused on consistency of root DNS data. These metrics have a stronger relation to security of the root DNS system.

Data inconsistencies can potentially be caused by many factors. Given the scope of this study we restricted our analyses to the possible impact of the introduction of new gTLDs on data consistency. In particular, we analyzed two specific aspects of new gTLDs. First, new gTLDs appear to cause a slightly higher rate of changes in the root zone file, which can potentially result in data errors in the zone file. Secondly, the mandatory enabling of DNSSEC requires coordination between cryptographic keys in the root zone file and the matching key data in the new gTLD zone files, in order to avoid validation errors. Our analyses indicate that none of these two potential effects have occurred. In fact, DNSSEC validation measurements show that new gTLDs show no more validation errors than other TLDs (Finding 14). A scan of all root zone files that were published since the first delegation of new gTLDs also indicates that changes to new gTLD data have not led to any data errors (Finding 15).

We recognize that these simple verification analyses do not enable us to conclude that the introduction of new gTLDs does not have an impact on DNS data consistency. However, our findings are expected (based on the fact that root zone files are validated before they are published) and confirmed by the lack of published issues with data consistencies that can be attributed to the introduction of new gTLDs. Therefore, we render more detailed investigation of this metric less useful.

5.5 Overall Conclusion and Recommendations

Our analysis of the available large amounts of historical and longitudinal measurement data shows that the root has been able to handle the increase in root server traffic over the past years, including the period in which new gTLDs have been delegated. In this period no significant degradation of the stability or security of the root DNS system can be attributed to the new gTLDs.

Moreover, from this analysis we inferred time-invariant correlations between root zone size parameters and stability/security metrics of the root DNS system. Presuming that these correlations will remain time-invariant for the near future we do not see obvious signals that the delegation of more new gTLDs in itself will degrade the stability or security of the root DNS system in the near future.

However, the absence of an observed degradation of the security and stability of the root DNS system is no reason to be less cautious for possible future impact of the New gTLD Program. In particular, the preventive root zone scaling measure in the New gTLD Program to limit the rate of delegations of new gTLDs may have contributed to the absence of degradation of the security and stability of the root DNS system. We advise the New gTLD Program to retain a controlled rate of delegating new gTLDs.

Further, we advise more frequent monitoring of the impact of new gTLD delegations, in order to obtain more detailed insight and to identify and respond to events impacting root DNS system stability on a short time scale[38]. The data-driven methodology developed in this study can be used and automated for this purpose. To enable monitoring in a more frequent manner an upgrade of the current set of data collection methods is recommended. In particular, for the analyses that require aggregation of DNS query volumes per name (TLD) that is queried for, we had to resort to using raw, snapshot (DITL) data. In principle, the aggregated, daily RSSAC002 measurements would be a more appropriate method, if the RSSAC002 measurement parameters would be complemented with daily aggregated traffic volumes per (most popular) TLDs[39]. We recommend the technical DNS community to consider this extension of the RSSAC002 measurements.

[38] For example, monitoring on a daily basis allows analysis of query rate fluctuations at the days around initial delegation of new gTLDs, such as illustrated in Figure 9. Further, short time-scale monitoring would allow the analysis of the impact of new gTLDs on incidental traffic peaks, if any (see the remark about incidental peaks in Section 4.1.1).

[39] Note that some RSOs already provide such DNS Statistics Collector data (DSC), but these are not part of the specified RSSAC002 metrics.

6 Possible Future Developments

In this section, we speculate[40] on risk parameters that we believe are worth monitoring more frequently if the ICANN community decides to expand the root zone with additional gTLDs after the current round. We selected these parameters based on our interactions with the community, who regularly asked us to speculate on the future of the DNS in general.

We focus on the risk parameters that we considered the most relevant which are: (i) multiple .com-sized gTLDs (Section 6.1), (ii) post-retirement and pre-delegation traffic (Section 6.2), and (iii) an increase in server-side processing on the root DNS system (Section 6.3). This list may not be exhaustive.

6.1 Multiple .com-sized TLDs

We speculate that an increase in the number of “large” gTLDs in a relatively short timeframe might form a stability risk. By “large” we mean .com-like gTLDs, which are both large in terms of domain names under management as well as in terms of queries on the root DNS system. Such gTLDs might form a stability risk because our analysis shows that .com is responsible for a significant portion of the total number of DNS queries that the root DNS system handles (see Figure 8)

We estimate that the probability of this scenario to unfold is low because we believe it is unlikely that a gTLD will grow to a .com-like size within 12-24 months, which is the time frame root server operators need to significantly update their infrastructure [2]. However, we also cannot exclude it and we therefore advocate monitoring and analyzing DNS traffic across all root server letters on a more frequent basis to detect new .com-like gTLDs early on. Our analysis suggests that there exists a relatively stable relation between the size of a gTLD and the number of DNS queries that the root DNS system receives for this gTLD (see Finding 4). This implies that the extrapolated growth in domain name registrations for a gTLD may be one of the indicators for a “large” new gTLD to emerge.

For this same reason, we also recommend continuing to enforce a controlled rate of delegation of new gTLDs to the root zone so that root server operators have sufficient time to further increase the capacity of their infrastructure. This is in line with the recommendations of the 2009 Scaling the Root study [2], which indicate that root server operators need about 18 months to significantly update their infrastructure and about 3 months for regular upgrades. For such a monitoring and control cycle to be effective, the monitoring intervals should be frequent enough to provide a trend view that can be used for control decisions (e.g. infrastructure updates or postponing new delegations). Typically, the monitoring intervals should be in the order of a quarter of a year, a month or preferably even shorter. At least it ought to be shorter than the yearly DITL data collection that we had to resort to for this study.

Monitoring and analyzing DNS traffic across all root server letters more frequently allows the detection of new .com-like gTLDs early on that can be used for guiding a controlled rate of delegation of new gTLDs to the root zone.

[40] We speak of “speculate” here, because our data analyses do not provide conclusive evidence for the risk parameters mentioned in this section. Nevertheless some of the risk parameters were inspired by speculative extrapolation of our data analyses.

Based on our analysis (see Section 4), we also speculate that the root will likely be able to handle the additional traffic that additional "normal" gTLDs will introduce. By "normal" we mean gTLDs that attract traffic loads similar to the ones that have been delegated as of the beginning of the New gTLD Program (see Finding 3 and its accompanying analysis). Our rationale is that the root DNS system is a highly diverse, flexible, and distributed system, which the root server operators are able to grow dynamically as the demand for capacity increases. Our analysis supports this as the root DNS system was able to handle an increase in DNS traffic in the period of Sep 2013-Sep 2016 (see Figure 7) using a growing number of root server sites (see
Figure 17). This included the relatively small amount of DNS traffic for currently delegated gTLDs (see Figure 8) as well as new gTLD DNS traffic on root servers that are "local" to geographic gTLDs (see Table 5).

6.2 Post Retirement and Pre-delegation Traffic

Another potential stability risk is gTLDs that were removed from the root zone file. This is because clients may have such retired TLDs hardcoded in their software and as a result they continue to put a load on the root DNS system. The .home and .corp queries that we see being "leaked" to the root (see Figure 8) are of a similar nature, although these TLDs have never actually been delegated.

A scenario like this might unfold in the "Internet of Things" when several "large" gTLDs have been withdrawn from the root and significant numbers of abandoned devices with outdated firmware continue to query domain names within these TLDs. This would further increase the number of invalid queries, in addition to the steady increase of invalid queries that we have observed since 2012 (see Figure 8). While the root has been able to handle this increase so far, a large number of retired gTLDs that attract a significant amount of leaked traffic may form a stability risk in the long run.

The number of retired gTLDs might for instance increase as a result of the lifecycle of a gTLD becoming more dynamic, similar to that of a domain name. For example, if applicants are able to apply for a new gTLD and can get it delegated in a matter of days at relatively low costs, then they might also retire it more easily. A scenario like this also poses a challenge for root server operators, who may need to enhance their capabilities to upgrade their infrastructure more quickly (see Section 6.1).

Analyzing the levels of invalid queries across root server letters on a more frequent basis enables the timely detection of retired gTLDs that attract a significant amount of "leaked" traffic, even when the gTLD lifecycle becomes more dynamic.

Conversely, we have seen that the root may already receive DNS traffic for applied-for gTLDs that have not yet been delegated (see Figure 8 and Figure 9). However, the query rates we found in our analysis were relatively low (around 150,000 queries per day max, see Figure 9A), which forms a limited risk for the stability of the root. The relatively low query rate may be the result of a new gTLD not being used widely yet around the time of delegation.

6.3 More Server-side Processing

Our final risk parameter is an increase in the amount of processing on root name servers, which would reduce the amount of peak traffic they can handle. A development like this would likely be the result of changes in the DNS protocol or changes of resolver behavior and would be independent of the New gTLD Program.

We believe that changes in the DNS protocol or in the way it is being used might even have a larger operational impact than the addition of more TLDs.

6.3.1 DNS over non-UDP Transports

An example of a risk parameter is that resolvers switch from UDP to TCP on a massive scale. This would require root servers to handle many more stateful connections than today for the same amount of traffic, which increases the risk that they run out of resources. We believe this scenario is unlikely to materialize in the near future because by far most DNS transactions use UDP as their transport protocol (see Figure 11).

Ongoing discussions in the IETF suggest that TCP transport might however become more popular in the future. Similar considerations can be made regarding the DNS-over-HTTPS proposals, which are also being discussed within the IETF. This also holds for the proposed Datagram Transport Layer Security (DTLS) protocol, which will likely require more resources as well.

6.3.2 DNS Cookies

DNS Server Cookies [27] may also increase the amount of server-side processing on root servers. Server Cookies are pseudorandom numbers that name servers generate to loosely authenticate incoming DNS queries. The amount of processing that the cookie mechanism requires depends on the particular algorithm that the server uses to generate cookies.

The advantage of using DNS server cookies is that they provide servers with a lightweight mechanism to treat validated requests in a different way than non-validated requests, for instance by severely rate limiting the latter. This enables servers to reduce the impact of reflection and denial of service attacks to their response rate limit, which contributes to increasing the stability of the DNS.

6.3.3 QNAME Minimization

Another example is QNAME minimization [29], which is an experimental DNS extension that increases the privacy of the DNS protocol by reducing the amount of information that resolvers send to name servers[41]. For example, to resolve "domain.example", a resolver with QNAME minimization sends a query for ".example" to the root and a subsequent query for "domain.example" to the name servers of .example. Without QNAME minimization, the resolver would send a query for "domain.example" to both.

A side effect of QNAME minimization is that it reduces the number of queries on the root for non-existing TLDs (invalid queries). This is because with QNAME minimization a resolver only needs to send a query for ".example" to the root to learn that .example does not exist and can authoritatively answer NXDOMAIN for subsequent queries to second-level domains. Without QNAME minimization, it would only learn that individual domains with the .example TLD do not exist (such as "domain.example") and as a result would send a query to the root for each .example second-level domain request.

A continuous awareness is needed regarding the possible disruptive impact of server-side processing at the root DNS system caused by new DNS extensions and new ways of using the DNS (e.g. non-UDP transports).

[41] QNAME minimization is not widely deployed yet, although a resolver like Unbound supports it.

Acknowledgements

For this study we have gratefully used root DNS data sets that have been collected by other parties. In particular, we would like to acknowledge DNS-OARC as a provider of DITL and other data to this study and RIPE NCC as a provider of RIPE Atlas data.

In addition, we have been given the opportunity to discuss our study approach and preliminary results with stakeholders of the DNS community. We thank the organizations that have arranged meetings for these interactions.

Bibliography

[1] ICANN, "Root Stability Study RFP," 2015. https://www.icann.org/news/announcement-2-2015-06-05-en.

[2] J. Akkerhuis, L. Chapin, P. Fältström, G. Kowack, L.-J. Liman and B. Manning, "Scaling the Root: Report on the Impact on the DNS Root System of Increasing the Size and Volatility of the Root Zone," 2009.

[3] "root-servers.org," http://root-servers.org.

[4] G. Moura, R. d. Schmidt, J. Heidemann, W. d. Vries, M. Müller, L. Wie and C. Hesselman, "Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event," 2016. http://www.isi.edu/%7ejohnh/PAPERS/Moura16a.pdf.

[5] "Internet archive - wayback machine," https://web.archive.org/web/20080313095445/http://www.root-servers.org/.

[6] ICANN, "New gTLDs Program," ICANN, https://newgtlds.icann.org/en/about/program.

[7] ICANN, "Measuring the Health of the Domain Name System," 2010. https://www.icann.org/en/system/files/files/dns-ssr-symposium-report-1-03feb10-en.pdf.

[8] ICANN, "DNS Stability, Security and Resiliency," 10 February 2012. https://www.gcsec.org/keyportal/uploads/dns_ssr3_report_20120210_001.pdf.

[9] D. Conrad, "www.internetsociety.org," 2012. http://www.internetsociety.org/sites/default/files/bp-dnsresiliency-201201-en_0.pdf.

[10] O. Kolkman, M. Santcroos, "DNS Threat Analysis," 3 May 2007. https://nlnetlabs.nl/downloads/se-consult.pdf.

[11] S. Baranowski, "How Secure are the Root DNS Servers?," March 2003. https://www.sans.org/reading-room/whitepapers/dns/secure-root-dns-servers-991.

[12] RSSAC, "RSSAC Statement Concerning The Impact of the Unavailability of a Single Root Server," 2016. https://www.icann.org/en/system/files/files/rssac-unavailability-single-root-server-08sep16-en.pdf.

[13] Root Server Operators, "Events of 2015-11-30," 4 December 2015. www.root-servers.org/news/events-of-20151130.txt.

[14] P. Danzig, K. Obraczka and A. Kumar, "An Analysis of Wide-Area Name Server Traffic," 1992. http://nms.lcs.mit.edu/6829-papers/danzig92analysis.pdf.

[15] N. Brownlee, K. Klaffy and E. Nemeth, "DNS Measurement at a Root Server: an early investigation of DNS traffic at the F-root servers," 2001. https://www.caida.org/publications/papers/2001/DNSMeasRoot/dmr.pdf.

[16] D. Wessels and M. Fomenkova, "Wow that's a lot of packets: an early investigation of DNS traffic at the root," 2003. https://www.caida.org/publications/papers/2003/dnspackets/wessels-pam2003.pdf.

[17] S. Castro, M. Zhang, W. John, D. Wessels and K. Klaffy, "Understanding and Preparing for DNS Evolution," 2010. https://www.caida.org/publications/papers/2010/understanding_dns_evolution/understanding_dns_evolution.pdf.

[18] RIPE NCC, "RIPE NCC DNS Monitoring Service," https://atlas.ripe.net/dnsmon.

[19] D. Wessels and G. Sisson, "L-Root Zone Augmentation Analysis," 17

September 2009. https://www.icann.org/news/announcement-2009-09-17-en.

[20] Global Advisors, "Mitigating the Risk of DNS Namespace Collisions," 28 October 2015. www.icann.org/en/system/files/files/name-collision-mitigation-final-28oct15-en.pdf.

[21] RSSAC, "RSSAC002 - RSSAC Advisory on Measurements of the Root Server System," https://www.icann.org/en/system/files/files/rssac-002-measurements-root-20nov14-en.pdf, 2014.

[22] DNS-OARC, "Description of annual DITL data collection," DNS-OARC, https://www.dns-oarc.net/oarc/data/ditl.

[23] RIPE NCC, "RIPE Atlas Measurements," https://atlas.ripe.net/measurements/ .

[24] J. Akkerhuis, "DNSSEC Monitoring - On a Shoe String," 7 March 2016. https://meetings.icann.org/en/presentation-dnssec-monitoring-07mar16-en.

[25] ICANN, "Staff Report of Public Comment Proceeding - CDAR Study Plan," 17 February 2016. https://www.icann.org/en/system/files/files/report-comments-cdar-study-plan-17feb16-en.pdf.

[26] ICANN, "Staff report of Public Comment Proceeding - CDAR Draft Report," February 2017. https://www.icann.org/en/system/files/files/report-comments-cdar-draft-09feb17-en.pdf.

[27] IETF, "RFC 1035," November 1987. https://tools.ietf.org/html/rfc1035.

[28] IETF, "RFC 7858," May 2016. https://tools.ietf.org/html/rfc7858.

[29] IETF, "RFC 7816," March 2016. https://tools.ietf.org/html/rfc7816.

A Appendix: Details about Data Sets

A.1 RSSAC002

Data description

In 2014 RSSAC published RSSAC002[42]: an advisory on daily aggregated measurements of the Root Server System. This document contains descriptions of an initial set of parameters that according to RSSAC would be useful to monitor and establish a baseline trend of the root server system. These parameters are as follows (divided into sub-metrics): Daily measurements of:

- Latency in publishing available data (load-time)
- The size of the overall root zone (zone-size)
- The number of queries and responses (traffic-volume)
- The query and response size distribution (traffic-size)
- The response type distribution (rcode-volume)
- The number of sets seen (unique-sources)

Timeframe

As of 28 June 2016 eight of the RSOs have started to collect and publish RSSAC002 measurements. The timeframe varies per root server letter:

- *Oct 2013 – Now*: A, J
- *Jun 2014 – Now:* L
- *Jan 2015 – Now:* C
- *Mar 2015 – Now:* H, K
- *Oct 2015 – Now:* D
- *Nov 2015 – Now:* M
- *Dec 2015 – Now:* B

After 28 June 2016 RSSAC002 E-, G- and I-ROOT also published RSAC002 measurements (tracing back to August 2016, July 2016, respectively April 2013).

Data availability and collection

The data is publicly available via www.root-servers.org. This site provides links to each root server letter and (if available) the RSSAC002-measurements, stored in YAML-format. Within the CDAR study we used Python scripts to read and aggregate the online YAML files for further analysis.

DNS-OARC has undertaken to act as a long-term repository for this data, see https://www.dns-oarc.net/node/348. admin@dns-oarc.net may be contacted if one is interested in the analysis of this data.

Data limitations and issues

- The timeframe is limited: only two of the root servers reach back as far as October 2013 and there are no RSSAC002 measurements *before* the delegation of new gTLDs (only I-ROOT has a few months of measurements available prior to October 2013).
- The data (such as daily number of queries) is not split out per TLD.

[42] https://www.icann.org/en/system/files/files/rssac-002-measurements-root-20nov14-en.pdf

- DNS-OARC lists some issues regarding this data[43]. While this document is outdated, some of these issues are still valid:
    - Data not found in the right place (some days' data found in next month's directory tree)
    - Data missing (missing spans of 24h series) for some days across a number of sources.
    - Some root servers do not record zone-size and load-time.
- Some root servers deviate from the proposed YAML-format
- There are some inconsistencies in the data. e.g., the total number of queries per root server letter per day can be aggregated in two different ways from RSSAC002 data:
    - Adding the counts in the different subcategories related to queries within the traffic-volume metric.
    - Adding the counts in the different query size bins in the traffic-size metric.

  These two aggregations are not equal for all root server letters – for some there are minor differences (<1%) but in some periods of time the differences between aggregations become large; often we see then that the traffic-volume aggregations are higher and therefore probably more accurate.

A.2 DITL data

Data description

DNS-OARC annually organizes a two-day data collection effort [22], also known as the Day In The Life of the Internet (DITL) data traces. The data consists of the queries received by the root and other DNS servers. The results are stored in files in an "industry standard" pcap format[44] and contains a lot of information although some of the data is obfuscated because of privacy requirements.

Timeframe

DITL data has been collected since 2006 onwards, with a consecutive two-day period per year (usually in March, April or May).

Data availability and collection

The data itself can be accessed by DNS-OARC members[45]. There are quite some tools available to process the data. The CDAR team used some of these tools to decide what type of data would be of interest for the study. These tools are very general and flexible in nature and often require quite some post-processing and data aggregation to present answers to questions such as "how many queries per second are received for these delegated domains".

After these requirements were established one of the CDAR partners wrote a set of targeted programs that collected and aggregated the data needed. This sped up the analysis of the DITL data and also simplified later (repeated) data processing and aggregation by the usual tools such as awk and others.

Data limitations

Not all root server operators contribute the complete two-day period or contribute every year. Also, root server operators do not always measure the full two-day period on all of their individual servers. See in Figure 18 the amount of queries as

[43] https://indico.dns-oarc.net/event/21/contribution/32/material/slides/0.pdf
[44] http://www.tcpdump.org
[45] Subject to the DNS-OARC's data sharing agreement.

stored in the DITL sets for eight of the root server letters, all measured between 2012 and 2016. Note that fluctuations are visible, both between root server letters and between years.

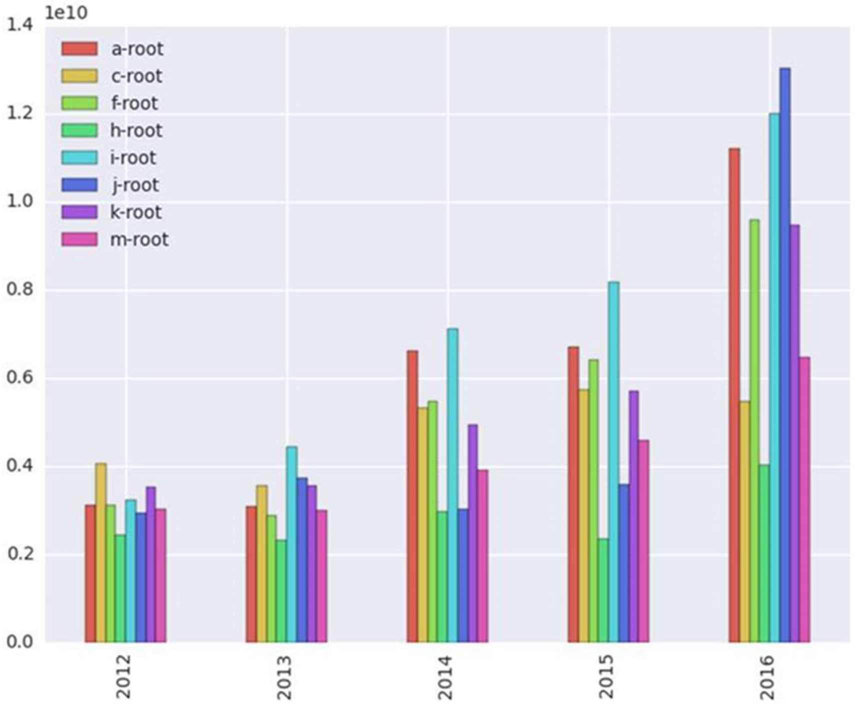

Figure 18: Total number of queries in DITL sets for eight of the root server letters

A.3 Root Zone File Archive

Data description

DNS-OARC has assembled a historical archive[46] of the DNS root zone files dating back to 1999. The data contains the raw zone files.

Timeframe

The covered timeframe runs from 1999 to now, with typically one root zone file per day. There are gaps (missing days), especially before 2006, but also in recent years. In the period 30 December 2012 to 21 September 2015, 44 days are not included in the archive.

Data availability and collection

The Root Zone Archive data is available to DNS-OARC members (subject to DNS-OARC's data sharing agreement), either as raw zone files or as a Subversion repository.

Data issues

- Dates suggested by the root zone file names do not (always) correspond to the actual date when the root zone file was published. The actual publish

[46] https://www.dns-oarc.net/oarc/data/zfr/root

date can always be found in the serial number contained in the SOA record of the file, e.g., the file with filename ‘root-20130105.021102’ refers to the root zone file with serial number 2013010301.

- Some files in the archive have different names but actually refer to the same root zone file. E.g., files ‘root-20130105.021102’ and ‘root-20130104.021102’ both refer to the root zone file with serial number 2013010301.
- No syntax errors found, nor were there empty zone files.

A.4 ICANN registry reports

Data description

ICANN publishes monthly registry reports on the generic TLDs. In particular, in the monthly transactions reports the data item “total-domains” denotes the total number of registered domains in a gTLD, which is valuable to the CDAR study. The monthly registry reports are provided by the TLD registry operators and are then published by ICANN, although their publication is withheld for three months due to contractual reasons.

Timeframe

The reports trace back for many years (depending on the date of initial delegation). For example, registry reports for .com are available from January 2001. The new gTLDs have reports starting from their delegation dates. The registry reports are published per month.

Data availability and collection

The reports are publicly made available via ICANN’s website[47]. For the purpose of the CDAR study we used the monthly transaction reports, most of which are available in CSV-format. Using a Python script, we read and parsed these reports to extract the total number of second-level domains per month per TLD.

Data limitations

The ccTLDs are not covered by the registry reports. The most recent three months are not available due to contractual reasons.

Data issues

The transaction reports stored in CSV-files are not all formatted equally. Furthermore, the transaction reports for some TLDs were stored in PDF, before switching to CSV, making it difficult to parse these reports.

A.5 DSC data

Data description

DSC (DNS Statistics Collector) data is an aggregated form of data based on collected DNS queries and responses entering DNS name servers. A DSC system consists of data collectors (run on or near DNS name servers) and data presenters that can display aggregated DSC data (e.g. to generate graphical plots). The DSC software code is currently maintained by DNS-OARC[48] and can be applied by operator of DNS name servers.

[47] https://www.icann.org/resources/pages/registry-reports

[48] https://www.dns-oarc.net/tools/dsc

The DSC system is flexible with respect to the type of DNS data that can be collected, aggregated and displayed. In published DSC output metrics are usually included, such as: number of queries received per time unit (day, week, month) and queries counts per time unit for the most queried domain name extensions. In particular, such DSC metrics that aggregate measurements per TLD are interesting for the CDAR study.

Timeframe

Typically, DSC data that is publicly available is represented as graphs on several periodic scales (daily, weekly, etc.) Such graphs are overwritten dynamically. DSC data archived by DNS-OARC traces back for several years.

Data availability and collection

Availability of public DSC data from RSOs is limited. Some of the RSOs publish DSC data for their root server letter infrastructure[49]. DNS-OARC archives DSC data that is contributed by root-letters C and K, which is available for its members. For the purpose of our study we also received some DSC data sets from individual RSOs.

Data limitations

While DSC data is potentially useful for the CDAR study its applicability turned out to be limited so far. This was due to the limited availability (in terms of number of RSOs that could provide DSC data) and the fact that our analyses required specific data aggregations that are not reported in default DSC statistics. From the CDAR perspective it was easier to use raw DNS data and process the aggregated metrics of interest.

Data issues

For the DSC data sets that we received we verified simple DSC counters, such as the number of queries received in a day, to the same counters from other data sets (RSSAC002 and raw PCAP data). It appears that significant deviations occur in these numbers. Investigation of one of these deviations appeared to be caused by underperforming data collectors that could not keep up with the querying speed of the name server.

A.6 RIPE Atlas data

Data description

The RIPE Atlas active measurement network from RIPE NCC consists of more than 9,000 measurement probes that provide vantage points distributed around the globe. It is the largest measurement network for which data sets are publicly available. All probes continuously measure all root DNS letters. RIPE uses different measurement IDs to identify each root server letter [23].

RIPE measurements are carried every 4 minutes[50]. The DNS requests sent by the probes are in the form of CHAOS queries that return the name of the server that responds to the queries.

Timeframe

The specific time period covered varies per measurement. Some go back as far as 2012. Out of caution, we disregard measurements from Atlas probes that had a firmware version before 4570, which was released in 2013.

[49] https://www.ripe.net/analyse/dns/k-root/statistics

[50] With the exception of A-Root that was measured with 30 minute intervals until mid-2015.

Data availability and collection

RIPE Atlas measurement data is publicly available via RIPE NCC [51], see https://atlas.ripe.net/measurements/. After cleaning, we map all observations into a time series of 10 minutes. In each time bin we identify the root server letter and the response (either the anycast site or the error code). Each time represents 2.5 Atlas probing intervals, a similar approach used in a related study [2].

Data limitations

While RIPE Atlas is a worldwide distributed measurement network, most of its probes are located in Europe and North America. This, however, does not interfere with our study, because we disregard probes that fail independently. Further, we remark that RIPE Atlas probes measure the RTT to the different root server letters, which is not the same as resolving a DNS query. A resolver resolving a query typically shows a more complex behavior, including selecting the root server letter for best performance.

[51] https://atlas.ripe.net/measurements/

B Appendix: Supplementary Analyses

Query rate / domain ratio

The figure below is the result of one of the analyses supporting Finding 4, and it illustrates that the query rate/domain ratio appears to be relatively constant, regardless of how ‘popular’ the TLD is.

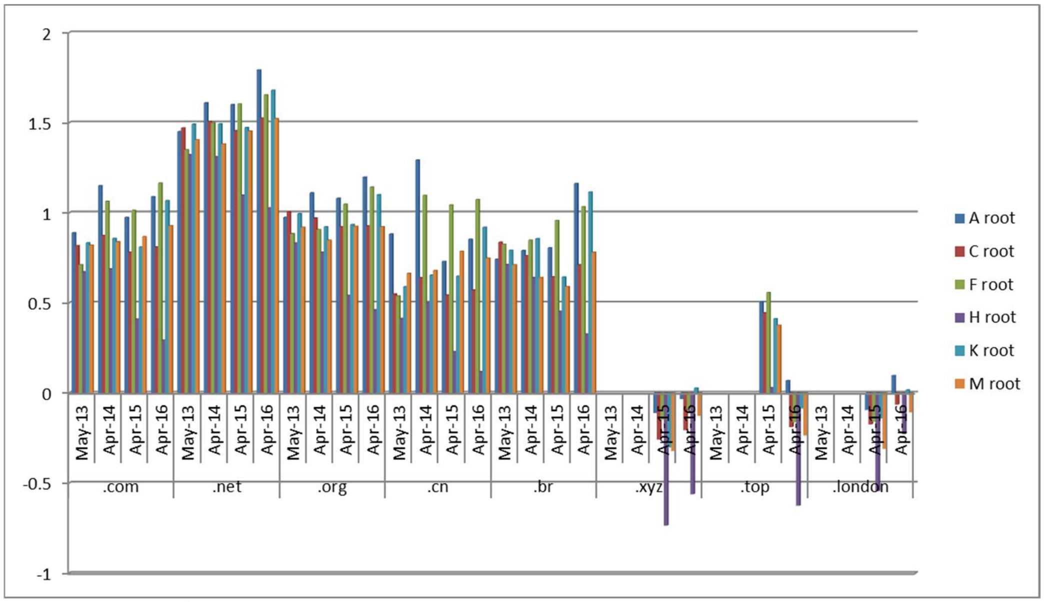

Figure 19: ^{10}log of ratio between the number of queries and the number of domains for different DITL sets and TLDs

Geographical affinity

In the table below we show for both F-ROOT and L-ROOT the maximum factors (localized fraction of queries to TLD divided by average fraction of queries to TLD) per root server letter and DITL set. We take only those locations into account at which at least 0.1% of the total queries arrive, to ensure the number of queries to new gTLDs are high enough to allow for an accurate estimation of the fractions. The highest factor visible is 56.8, while most others are much lower. This supports the finding that the fraction of traffic to such geographic new gTLDs remains insignificant.

TLD	F-ROOT DITL 2015	F-ROOT DITL 2016	L-ROOT DITL 2015	L-ROOT DITL 2016
bayern	3.7	6.8	7.6	3.7
capetown	21.9	15.1	10.0	23.5
com	2.2	1.8	1.8	2.3
doha	31.5	19.8	8.5	56.8
london	7.0	15.6	4.4	5.0
melbourne	15.4	9.6	19.5	23.0
moscow	7.0	12.6	9.0	10.4
nyc	7.2	10.2	5.2	5.1
rio	24.1	14.8	7.2	23.0
sydney	20.7	12.2	17.4	22.3
tirol	29.5	12.6	9.3	23.5
tokyo	16.7	23.2	4.2	8.1
vlaanderen	13.1	9.9	6.4	18.2
xn--80adxhks	9.9	10.2	10.9	11.1

C Appendix: Consortium

NLnet Labs (www.nlnetlabs.nl)
NLnet Labs is a non-profit research and development company that focuses on developments in Internet technology bridging the gap between theoretical insights and practical deployments; engineering and standardization, where public interest is often more pressing than commercial interest. It is NLnet Labs' goal to play an active and important role in the development of open source software, participation in development of open standards, and dissemination of knowledge through training, consultancy, and evangineering. NLnet Lab's software is an important component of the Internet infrastructure. NLnet Lab plays a significant role in standards development. Dissemination of knowledge is realized through education and collaboration. NLnet Labs has a staff of nine software developers and experts.

NLnet Labs is recognized for her expertise in Internet system technology, security and architecture, in particular in DNS, DNSSEC, inter-domain routing and addressing. With the development of authoritative name servers and recursive resolvers, NLnet Labs has deep knowledge of the DNS system and its protocols. Complementary to this, NLnet Labs has a strong track record in providing expertise to security and stability analysis of critical infrastructures like scaling the root study, SSAC, ENISA study of the routing infrastructure and is member of the ENISA Internet Infrastructure Security and Resilience Reverence Group. With this, NLnet Labs is strongly involved in the ICANN, DNS and Internet infrastructure community.

NLnet Labs team member: Benno Overeinder
Benno Overeinder is managing director of NLnet Labs in the Netherlands. He is active in the RIPE and IETF community, focusing on Internet infrastructure security and stability, both DNS and routing related. He is the chair of the RIPE Programme Committee and co-chair of the RIPE Best Current Operational Practices Taskforce. Overeinder has contributed to ENISA commissioned studies on Internet routing infrastructure security and stability, and is member of the ENISA Internet Infrastructure Security and Resilience Reverence Group.

NLnet Labs team member: Jaap Akkerhuis
Jaap Akkerhuis is a senior research engineer at NLnet Labs. He has been instrumental in the development of the Internet in the Netherlands and in Europe in the early 1980s. After some year in the US, he returned to the Netherlands where he joined the first independent ISP. Later he worked as a Technical Advisor for SIDN, the registry of the .NL TLD. Jaap Akkerhuis has served in the SSAC since its inception and is co-chair of the RIPE DNS working group and served as a co-chair for the IETF ProvReg WG. He is a regular consultant to ICANN and their member of the ISO 3166 Maintenance Agency.

SIDN (www.sidn.nl)
SIDN manages the Internet extension of the Netherlands, .nl. As the Dutch national domain name registry, we enable Internet users to safely use and register .nl domain names anytime and anywhere. We operate the .nl zone of the Domain Name System (DNS) and handle over a billion DNS queries every day for more than 5.6 million registered .nl domain names. Over 2.5 million of those are secured with DNSSEC, making .nl the largest secured Internet extension in the world. We also provide the backend services for the new

gTLDs .amsterdam and .politie (“Police” in Dutch) as well as for the country code .aw (Aruba).

SIDN has been actively involved in the ICANN community since its inception. We actively contributed to the cross-community working groups on the IANA Stewardship Transition and ICANN Accountability, we led the working group “Secure Email Communication for ccTLD Incident Response” (SECIR), and we are currently chairing the TLD-OPS Standing Committee. In the past, we had staff on the ccNSO Council and led the ccNSO working group “Strategic and Operational Planning” (SOP). In addition, SIDN is a long-time member of DNS-OARC and had one of our staff on the DNS-OARC board from 2012 until 2014.

SIDN Labs (www.sidnlabs.nl) is SIDN’s research team, which develops, prototypes, and evaluates new technologies and systems that further enhance the security and stability of .nl, the DNS, and the Internet at large. An example is ENTRADA (ENhanced Top-level domain Resilience through Advanced Data Analysis)[52], an experimental system that we have developed to capture, store, and analyze the DNS traffic we handle on our production systems. The goal of the platform is to develop new services and applications to discover anomalies and threats in the DNS traffic and use that information to enable SIDN as well as others to further increase the security and stability of the Internet. ENTRADA comes with a Privacy framework[53] to protect the privacy of Internet users.

SIDN team member: Cristian Hesselman

Cristian Hesselman is the head of SIDN Labs, which he also set up. Cristian was previously with *Telematica Instituut*, a Dutch public-private research facility, where he led and developed large national and international research projects. He also worked as a senior researcher on topics such as sensor systems, adaptive multimodal user interfaces, and service platforms. Before that, he was a software engineer at Lucent Technologies. Cristian holds a Ph.D. (2005) and an M.Sc. (1996) in computer science, both from the University of Twente, the Netherlands.

Cristian is a member of SIDN’s management team and serves on the board of NLnet Labs. He was previously on the board of Abuse Information Exchange, where he oversaw the development of AbuseHUB, a service for Dutch access providers to share information on botnet infections.

SIDN team member: Giovane Moura

Giovane Moura is a Data Scientist with SIDN Labs, the research arm of the top-level domain registry of the Netherlands (.nl). His interests include security and performance of computer networks, with emphasis on Internet measurements. Prior to SIDN, he worked as a postdoctoral researcher at TU Delft, and obtained his Ph.D. from the University of Twente, both in the Netherlands. He has also a Master’s degree in computer science from the Federal University of Rio Grande do Sul, Brazil.

TNO (www.tno.nl)

TNO[54] is one of the major internationally oriented contract research and technology organizations (RTO) in Europe. With a staff of approximately 3000 and an annual turnover

[52] https://www.sidnlabs.nl/uploads/tx_sidnpublications/NCSC-presentatie-BIG-data-pub.pdf

[53] https://www.sidn.nl/downloads/whitepapers/SIDN_Labs_Privacy_Framework_Position_Paper_V1.3_EN.pdf

[54] TNO is a not-for-profit organisation, whose acronym is an abbreviation of “Nederlandse Organisatie voor Toegepast-Natuurwetenschappelijk Onderzoek”

of 586 million Dollars, TNO carries out technological and life science research aimed at boosting innovation and achieving societal impact. By translating scientific knowledge into practical applications, TNO contributes to strengthening the innovation capacity of businesses and government. TNO is involved in many international projects (about 30% of the market turnover), including the Scaling the Root study commissioned by the ICANN board in 2009.

In TNO's innovation area of Information Society applied research is carried out along three lines:

- Technical Robust Infrastructures (Security, Stability & Quality)
- Information Creation (Media & Content Delivery; Big Data Evolution)
- Information Influence (Privacy & e-Identity; Strategic Use of Information)

TNO's Performance of Networks and Systems expertise (contributing to the Robust Infrastructures research line) was recognized as 'internationally leading' by an external knowledge auditing committee led by prof. W. Jonker. In the proposed study team TNO contributes quantitative modelling & analysis experts, who were involved in the Scaling the Root study team. Moreover, their research on the Global DNS reference model[55] was awarded the best paper award at the international DNS Easy conference in 2011.

TNO team member: Bart Gijsen

Bart Gijsen started working for KPN Research in 1997 as a technical performance analyst and joined TNO in 2003. Currently, Bart is a senior researcher and consultant in the department Cyber Security and Robustness. In 2009 Bart led TNO's contribution to the Scaling the Root study team and has been an active contributor to numerous DNS stability research initiatives. He presented his work at, amongst others, ICANN and DNS OARC meetings and DNS Health symposia. Bart has also led a study investigating the plans of Dutch multinationals regarding brand name TLDs, in cooperation with SIDN.

TNO team member: Daniël Worm

Daniël Worm has been working at TNO since 2011 as mathematical researcher and consultant. He has extensive experience with mathematical modelling and analysis, with a primary focus on stochastic modelling including statistics. He has participated in a variety of projects in the domains of ICT and energy. His work includes development of new models and performing stochastic analysis for telecom operators in order to estimate performance KPIs in their networks, applying optimization techniques and performing resilience and anomaly detection techniques.

[55] www.gc-sec.org/sites/default/files/files/dnseasy2011.pdf#page=6