

# Staff Report of Public Comment Proceeding

## Continuous Data-driven Analysis of Root Server System Stability Draft Report

<b>Publication Date:</b>	9 February 2017
<b>Prepared By:</b>	Eleeza Agopian and CDAR consortium

Public Comment Proceeding	
Open Date:	27 October 2016
Close Date:	15 January 2017
Staff Report Due Date:	9 February 2017

Important Information Links
<a href="#">Announcement</a>
<a href="#">Public Comment Proceeding</a>
<a href="#">View Comments Submitted</a>

<b>Staff Contact:</b>	Eleeza Agopian	<b>Email:</b>	<a href="mailto:eleeza.agopian@icann.org">eleeza.agopian@icann.org</a>
-----------------------	----------------	---------------	------------------------------------------------------------------------

### Section I: General Overview and Next Steps

ICANN commissioned this study in response to a recommendation from the Governmental Advisory Committee to examine the technical impact of the delegation of new gTLDs on the security or stability of the root DNS system. The study, which was conducted by independent research organization TNO and its consortium partners, SIDN and NLnet Labs, will help the ICANN community determine if additional steps are required to safeguard the root DNS system’s security and stability. The comments received on the report – most of which supported the study’s findings and recommendations – will inform changes to a final report, which will be published in March 2017. Those findings and recommendations will then be presented to the ICANN Board for consideration.

### Section II: Contributors

*At the time this report was prepared, a total of [number] (n) community submissions had been posted to the forum. The contributors, both individuals and organizations/groups, are listed below in chronological order by posting date with initials noted. To the extent that quotations are used in the foregoing narrative (Section III), such citations will reference the contributor’s initials.*

#### Organizations and Groups:

Name	Submitted by	Initials
Registries Stakeholder Group	Paul Diaz	RySG
Root Server System Advisory Committee	Andrew McConachie	RSSAC
Internet Service Provider and Connectivity Providers Constituency	Chantelle Doerksen	ISPCP
NIC.br	Rubens Kuhl	NIC.br
Business Constituency	Steve DelBianco	BC
At-Large Advisory Committee	ICANN At-Large Staff	ALAC

#### Individuals:

Name	Affiliation (if provided)	Initials
John Poole	DomainMondo.com	JP

### Section III: Summary of Comments

*General Disclaimer: This section intends to summarize broadly and comprehensively the comments submitted to this public comment proceeding but does not address every specific position stated by each contributor. The preparer recommends that readers interested in specific aspects of any of the summarized comments, or the full context of others, refer directly to the specific contributions at the link referenced above (View Comments Submitted).*

ICANN compiled the following summary of comments.

The CDAR report received seven public comments. Four commenters broadly agreed with the report's findings and supported recommendations for continuing a gradual rate of delegation of new gTLDs to the root zone; as well as a recommendation to apply continuous monitoring of the root to ensure its security and stability. Two commenters raised additional issues for consideration by the report authors and asked questions regarding the findings. One commenter said the report was inadequate.

A. RSSAC presented a list of 16 questions for the report's authors to consider in the revised report. Some of these questions are requests for clarification and/or further information in the report, while other questions invite the authors for further interpretation of their findings. In addition, some comments are a suggestion for refining the recommendations related to monitoring the root DNS' security and stability. These individual questions and the report authors' responses are included below in the Analysis of Comments.

B. RySG, BC, ALAC and ISPCP all endorsed the report's key recommendations. All four agreed that the gradual delegation rate of new gTLDs was a prudent approach. RySG noted the recommendations for additional continuous monitoring and suggested those "should be used as guidance to whether a ceiling (on delegation rates) is in effect required and what the ceiling should be."

NIC.br also notes that while the report suggests a ceiling on the delegation rate it neither prescribes one nor provides a method for determining what that ceiling ought to be; noting that the previous guidance on delegating no more than 1,000 new gTLDs per year was an arbitrary figure. Rather than prescribe a delegation rate, NIC.br suggests that ICANN use its monitoring systems to determine what the rate of future delegations ought to be. NIC.br noted that it does not recommend additional research, but asked that root server stakeholders use the available data to make informed decisions.

C. Three commenters included in their suggestions recommendations to ICANN on how to implement the recommendations included in the report. These included:

- With regard to additional continuous monitoring, RySG supported the idea and recommended that improvements should be made and tested before implementation.
- ISPCP recommended that ICANN, in its Office of the Chief Technical Officer, should provide continuous monitoring, rather than via a third party.
- ISPCP also recommended greater cooperation between root server operators, ICANN and the IETF, with a specific recommendation for the Root Server Caucus to "act as ICANN's informal liaison to the IETF on DNS protocol issues."
- BC, ISPCP and ALAC all concurred that ICANN should provide greater public access to data related to this study. Commenters suggested that data related to this study should be made public and thus provide transparency into how conclusions were drawn. ISPCP recommended regular, public reports, in addition to publishing raw data from monitoring efforts.

D. Two commenters suggested new areas of research:

- ALAC suggested that ICANN may consider future research into the impact on stability of the root should new gTLDs be removed, noting concerns raised in the CDAR report regarding such scenarios.
- ALAC also wrote that the report notes growth in invalid queries to the root and marked this as an area of concern which merits further research, including referral to the SSAC and RSSAC.
- BC reiterated a suggestion from its comment on the CDAR study plan, that new gTLD impacts on addresses also should be assessed and monitored, and related reports and data be made public.

E. With regard to the report’s recommendations, NIC.br writes that the recommendations are not supported by evidence in the report. For example, it notes that while the removal of a popular new gTLD from the root zone is a risk, it is a very small one, and one that also applies to TLDs which existed prior to the New gTLD Program.

F. One commenter, JP, criticized the report as it does not “definitively answer the essential questions”. JP also suggested the authors note comments submitted in 2016 regarding the CDAR study plan, which states that the absence of instability is unlikely to be reliably predicted by this study.

**Section IV: Analysis of Comments**

*General Disclaimer: This section intends to provide an analysis and evaluation of the comments submitted along with explanations regarding the basis for any recommendations provided within the analysis.*

In response to the public comments, ICANN and the CDAR team thank everyone who responded to the public comment period of the CDAR draft report for their valuable feedback. In this section CDAR responds to the comments made and indicate the changes that will be made to the CDAR report, as well as other next steps that will be taken. For ease of reference the responses will use the same reference letters as the summarized comments in the previous section.

ICANN provides its response to comments and suggestions directed at the organization further in the response.

A. In the following table the list of questions raised by RSSAC is included, complemented with the CDAR team’s responses. The responses also include the modifications that the CDAR team is considering to the final CDAR report, which will be published in March 2017.

RSSAC comment	CDAR team response
<p>The first finding states that the total number of queries to the root grows over time. Can the authors of the report draw any conclusions about what drives the root server traffic to increase over time? It would not appear to be number of TLDs. What is the primary factor in traffic growth over time?</p>	<p>The analyses of the CDAR study were focused on the question regarding the impact of the introduction of new gTLDs on the security and stability of the root DNS system. In this context Finding 1 is a relevant observation, independently of the explanation behind this observation. As indicated in the “Scope and Limitations” subsection of the study “we cannot deduce from the data why the number of queries is growing.”</p>

<p>If new TLDs do not appear to be the driving factor, do the authors have any concerns that the historical or forecasted rates of growth pose a threat to root server stability?</p>	<p>Based on the CDAR data analyses the historical rates of growth of the number of queries have not posed a threat, or at least we found no indication that the introduction of new gTLDs poses a threat. Presuming that the evolution of new gTLD delegations continues to exhibit the observed pattern we see no signs that the delegation of more new gTLDs will degrade the stability or security of the root DNS system in the near future.</p>
<p>The second finding states that the fraction of invalid queries (queries to invalid TLDs) increases significantly over time. Could the authors of the report say whether they believe the increase in fraction of invalid queries to be a threat to root server stability?</p>	<p>Similar to the previous answer the increase of the fraction of invalid queries is an observation, for which we did not find a conclusive explanation. As indicated in the “Scope and Limitations” subsection “the absence of such explanations reduces our options to extrapolate [this] observation for possible future DNS developments.”</p> <p>We can add here that a repetition of this CDAR analysis on the ZSK roll-over data (which became available at the end of 2016) shows no sign that the fraction of invalid queries is increasing at a concerning rate.</p> <p>Results from this analysis are included in the final CDAR report, to be published in March 2017.</p>
<p>In section 4.1.4 (page 32) the report states:</p> <p>“One of the reasons why we are interested in this is that the query type may have an influence on the load of a root name server.”</p> <p>Could the authors please clarify or back up this assertion?</p> <p>Do the authors mean DNS query types such as A, AAAA, MX, NS? Or do they mean other characteristics of a query such as transport (UDP/TCP), inclusion of DNSSEC data, and other protocol features? Is there any citable research that could quantify the additional load imposed by either different query types or protocol features?</p>	<p>In particular transport characteristics (UDP/TCP) are meant, but also larger responses (DNSSEC and other protocol features). This will be formulated more clearly in the final CDAR report.</p> <p>In general, the difference between UDP and TCP transport protocol are known to have a different effect on the load of servers and their response time performance. For example, RFC 1035 (section 4.2) and RFC 7858 (section 5) provide considerations about increased latency, higher overhead and server memory demand (due to additional state information) in case TCP is used, relative to UDP. The latency impact of using TCP / UDP is for example monitored by the RIPE Atlas monitoring framework. See: <a href="https://atlas.ripe.net/results/maps/root-server-performance/">https://atlas.ripe.net/results/maps/root-server-performance/</a>.</p>

	<p>In general, no remark about the exact impact of each of the different protocols on root name servers can be made, because this depends on many aspects including server configuration details. Nevertheless, the differences typically do have an impact and that is the motivation for analyzing the distribution of DNS query types in the CDAR study.</p> <p>In the final CDAR report citations will be included to reflect these remarks.</p>
<p>In section 6.1 the report makes the following recommendation:</p> <p>“We recommend monitoring and analyzing DNS traffic across all root server letters on a continuous basis to detect new .com-like gTLDs early on and to continue to enforce a gradual rate of delegation of new gTLDs to the root zone.”</p> <p>Rather than burden the root servers with additional monitoring requirements, could this instead be achieved by examining the zone files stored in the Centralized Zone Data Service (CZDS)?</p>	<p>In Section 6.1 the risk factor of potential rapid increase of “large” (.com-like) TLDs is identified, which are both large in terms of domain names as well as in terms of queries to the root DNS system. As such, zone files are a usable data source indeed. Moreover, using Finding 4, examination of the zone files (and registry reports) also provides a rough indication for the query rate.</p> <p>Nevertheless, we are convinced that more continuous monitoring of traffic rates for (most-popular) TLDs on a daily basis (such as the DSC, DNS Statistics Collector, most popular TLD metrics that are already measured and published by several RSOs) is essential to provide further insight in the impact of new gTLDs on the stability of the root DNS systems. In addition, such data would enable further analysis of the possible impact of the observed invalid queries (see RSSAC question in the second row of this table).</p>
<p>“.home”-like gTLDs</p> <p>In section 6.2 the report makes the following recommendation:</p> <p>“We recommend analyzing the levels of invalid queries across all root server letters on a continuous basis to detect “.home”-like gTLDs early on.”</p> <p>Do the authors believe that root server system stability depends on how invalid queries are distributed among different invalid TLDs? For example, consider two contrived scenarios: (1) a million invalid TLDs each with one query per second, vs (2) two invalid TLDs each with 500,000 per second.</p>	<p>Section 6 contains “speculations on risk parameters that we believe are worth monitoring ...” As such, this cited text is intended as a suggestion for consideration, rather than a recommendation based on the CDAR data analyses.</p> <p>For the stability of the root DNS system only the total query load is relevant, independent of the distribution among TLDs. Nevertheless, monitoring the distribution among TLDs is relevant for investigating the impact of new gTLDs and invalid queries on the stability of the root DNS system.</p>

<p>Is one of these better or worse than the other in terms of root server stability?</p> <p>Could the authors please explain why detection of “.home”-like gTLDs requires continuous monitoring? Why is periodic monitoring insufficient? Similarly, why is monitoring at all root server letters recommended? Would it not be possible to detect such gTLDs at a subset of root server letters?</p>	<p>To avoid confusion the text in section 6.2 (in particular the text box starting with “We recommend ...”) will be reformulated: “Analyzing the levels of invalid queries across all root server letters on a continuous basis enables timely detection of the risk parameter of leaking ‘.home’-like gTLDs, even when the gTLD lifecycle becomes more dynamic.” The CDAR team may further reformulate this text in the final report.</p> <p>The detection of “.home”-like TLDs requires more frequent data collection than is currently the case (for the CDAR study we had to build on yearly DITL data collection). So, doing such monitoring on a quarterly or monthly basis would already be a strong improvement. Monitoring a DSC / most-popular TLD-like metric on a daily basis would immediately resolve this recommendation.</p> <p>In principle, we can already use such DSC / most-popular TLD metrics, which are being reported by some root server letter operators, which is very useful. However, this provides a partial view on the root DNS system, which only allows for partial answers to questions about the stability of the root DNS system.</p>
<p>In section 6.3 the report makes the following recommendation:</p> <p>“We recommend continuously analyzing the use of non-UDP transports and new DNS extensions across root letters to detect trends in server-side processing.”</p> <p>While tracking deployment of new protocol features may be very interesting, does it accomplish the goal of providing insight into server processing limitations?</p>	<p>We agree with the suggestion inherent in this question. No, tracking the deployment status of new protocols does not provide direct insight into server processing limitations.</p> <p>In the final CDAR report this text will be rephrased.</p>

B. *CDAR Response:* Some of the comments refer to statements in the report concerning the “gradual rate of delegation”, which in our view is misinterpreted as a “ceiling on the rate of delegation” in some of the comments. In particular, the term “ceiling” is not mentioned anywhere in the report.

We re-iterate that the CDAR analysis is based on historic measurement data and we cannot exclude that the measured data has been influenced by the delegation rate measure adopted in the New gTLD Program. Or, as stated in Section 5.5: “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.”

Further, based on the observations from the data analysis it is stated in the report that *IF* “the evolution of new gTLD delegations continues to exhibit the pattern we observed since the New gTLD Program’s first delegations in October 2013”, *THEN* “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”.

The intention of the “gradual rate of delegation” recommendation is that by retaining control over the rate of delegation ICANN can influence at least one aspect of the precondition in the previous paragraph. Thereby having a control (but no full control) regarding the stability of the root DNS system. Or to present an extreme, uncontrolled-rate-of-delegation example: we would *not* have advised ICANN to organize a “New gTLD Launch Day”, where all 1,930 new gTLDs would have been delegated at a single point in time (as the extreme example of an uncontrolled rate of delegation), since our data (figure 9-D) suggests that such an event could theoretically have led to a query load peak comparable to the .com daily query load. Such an extreme event probably would not have been disastrous, but it is an event that one would like to prevent (in particular if a potential next round of new gTLDs would result in a higher number of applications than the recent round).

Therefore, we believe that it should be possible to control the rate of delegation, in case signs are being observed that the stability or security of the root DNS system might be harmed. We agree with the commenters that monitoring should be used to detect such signs. Although the results of the analyses presented in the CDAR report provide suggestions for specific metrics, a design of such a monitoring system is out of the scope of the CDAR study.

Further, we emphasize that a controlled rate of delegation is not necessarily the same as enforcing a “ceiling on the delegation rate”. A controlled rate of delegation might also be temporarily postponing further delegations in case potentially harmful signs are detected, or other measures. It is out of the scope of the CDAR study to propose an extensive list of such possible mitigating measures.

To make this clearer, we will rephrase our recommendation in Section 5.5 in the final report, and in other places, by replacing “gradual rate” with “controlled rate” and added some clarifying text.

- C. *Public Comment*: “With regard to additional continuous monitoring, RySG supported the idea and recommended that improvements should be made and tested before implementation.”

*CDAR Response*: We agree that any improvements of data collection and monitoring should be “tested before being implemented.” Moreover, in section 3 we point out several specific issues regarding missing data and inconsistencies that we encountered in the public data sets that were used during the study.

*Public Comment*: “ISPCP recommended that ICANN, in its Office of the Chief Technical Officer, should provide continuous monitoring, rather than via third party.”

*ICANN Response*: The data referenced in the report is publicly available. ICANN’s DNS Engineering team continuously measure the performance of the L root servers. The Office of the Chief Technical Officer intends to perform regular analysis of the root-server traffic and intends to publish the results.

*Public Comment:* “ISPCP also recommended greater cooperation between root server operators, ICANN and the IETF, with a specific recommendation for the Root Server Caucus to “act as ICANN’s informal liaison to the IETF on DNS protocol issues.” “

*ICANN Response:* Both the IETF and RSSAC have a liaison on the ICANN board. ICANN organization participates in IETF meetings on DNS and other protocol developments in various levels. ICANN organization operates one of 13 root server letters. ICANN organization and many IETF participants are active members of the RSSAC Caucus.

*Public Comment:* “BC, ISPCP and ALAC all concurred that ICANN should provide greater public access to data related to this study. Commenters suggested that data related to this study should be made public and thus provide transparency into how conclusions were drawn. ISPCP recommended regular, public reports, in addition to publishing raw data from monitoring efforts.”

*CDAR Response:* For the analyses in the CDAR draft study we have made use of publicly available data sources, as much as possible. Links to where this data is available are provided in section 3 of the report and appendix A. Part of the data used, for example DITL data collection data and the Root Zone File Archive, is available via DNS-OARC. Access to that data is privileged to DNS-OARC members (as mentioned in the report), due to data sharing agreements.

*ICANN Response:* ICANN is launching an Open Data Initiative that could provide some of this data. However, because some of the data is not owned by ICANN, ICANN would need to work with the data owners to make it public.

- D. *ICANN Response:* This is an active area of research. With regard to a study on addresses, this report did not mention addresses and their impact on gTLDs, so we must assume that none were seen. With regard to the ALAC comment on future research into the impact on root stability if new gTLDs are removed, ICANN organization agrees to consider this topic as a possible subject for future research.
- E. *CDAR Response:* The suggestions and speculations referred to in this public comment can be found in Section 6.2. Section 6 is the part of the report that contains “speculations on risk parameters that we believe are worth monitoring ...” and is therefore of a highly speculative nature as indicated in the report. This ‘speculative’ section has been included to respond to requests from the community to identify potential risk parameters. These speculations towards future developments are indeed not ‘proven’ by the CDAR data analyses and are therefore deliberately placed in a section after the conclusions and recommendations related to the CDAR data analyses that are summarized in Section 5. Apparently, this has not been formulated clear enough in the report and the “We recommend ...” stated in several text boxes might have fed this confusion. In the final CDAR report these text boxes and other parts of Section 6 are reformulated.

The remark that the risk parameter “also applies to TLDs which existed prior to the New gTLD Program” is correct, but the strong increase in the number of delegated new gTLDs makes this risk more prominent.

- F. *CDAR Response:* The CDAR team has extensively analyzed the available data to answer the study’s research question regarding the impact of the New gTLD Program on the security and stability of the root DNS system. In addition, the team put much effort in reaching out to the technical community to obtain additional insights. These efforts resulted in the conclusions in the draft report.

The authors feel that the comment from the previous public comment that states: “absence of instability is unlikely to be reliably predicted by this study” has not been neglected in the draft CDAR report. In the conclusion section and in the management summary, we note: “The absence of observed [...] degradation [...] is no reason to be less cautious [...]”.