Staff Report of Public Comment Proceeding

Continuous Data-driven Analysis of Root Server System Stability (CDAR) Study Plan	
Publication Date:	17 February 2016
Prepared By:	Eleeza Agopian

Public Comment Proceeding	
Open Date:	2 December 2015
Close Date:	3 February 2015
Staff Report Due Date:	17 February 2016

Important Information Links
Announcement
Public Comment Proceeding
View Comments Submitted

Staff Contact:	Eleeza Agopian	**Email:**	Eleeza.agopian@icann.org

Section I: General Overview and Next Steps

A review of the New gTLD Program (the Program) for security and stability impact is a previous commitment based on advice from ICANN's Governmental Advisory Committee and other discussions. Specifically, ICANN committed to review the effects of the Program on the operations of the DNS root system, and to postpone delegations in a future round until it is determined that the delegations in the 2012 round have not jeopardized the root system's security or stability.

In October 2015, ICANN contracted with TNO and its partners, SIDN and NLnet Labs, to conduct the study: Continuous Data-driven Analysis of Root Server System Stability (CDAR). The consortium presented its methodology at a workshop session at ICANN 54 and published the study plan for public comment on 2 December.

The consortium partners have begun to collect data and conduct preliminary analyses. In addition, they have presented early findings and discussed methodology with the RSSAC and SSAC. A public session is planned at ICANN 55 to bring the community up to date on the study's progress, as well as address any changes in the study plan, in response to the public comments as outlined below.

In response to the public comments, the CDAR team thanks all who responded for their valuable feedback. After analyzing the comments, the CDAR team feels an update to the study plan is not required. Rather, the team has indicated that the suggestions will be taken into account during the execution of the study and details its plans below.

A public session is planned at ICANN 55 in Marrakech on 8 March. Representatives of the CDAR team will discuss the comments received to date, as well present preliminary findings. More detailed technical findings will be presented by the CDAR team at DNS-OARC 24 from March 31-1 April 2016. The study's timeline remains on track to publish a draft report in the third quarter of 2016, with a final report expected in April 2017.

Section II: Contributors

At the time this report was prepared, a total of 2 (two) usable 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
Business Constituency	Steve Del Bianco	BC

Individuals:

Name	Affiliation (if provided)	Initials
Daniel Karrenberg	RIPE NCC	DK

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).

Two commenters submitted their input for the CDAR study. Both emphasized acknowledging the limitations of the study. BC noted this may be a particular problem when measuring past performance of the root, prior to the delegation of new gTLDs. DK writes that modeling may be too simplified for analyzing a complex system like the root server system. In particular, DK writes, a model "cannot predict the *absence* of instabilities in complex real-world system (sic) with any useful level of confidence." BC suggests the CDAR plan carefully note what may be achieved by the study.

DK notes that keeping expectations in check regarding the study's predictive possibilities requires ICANN to prepare "proper contingency plans for the unpredictable cases where root zone expansion causes or contributes to instabilities in the DNS root server system." Similarly, the BC recommends (comments 1, 2 and 3) identifying the right parameters for the study to ensure possible risks to the systems are identified early on and mitigated.

BC also emphasizes (comment 4) outreach to the broader DNS community to validate the methodology, measurements and models, with a particular focus on reaching out to those who may be experts in those topics.

More specifically, the BC provided several comments on the parameters:

- In WP-2, the BC recommends (comment 5) evaluating "the value of examining stability from the new gTLD registry perspective. It also notes that only scalability can be analyzed from the measurements identified there and recommends expanding them to include security and stability.
- The BC suggests (comment 6) the parameters ought to be clearly defined with regard to gTLDs, ccTLDs, new gTLDs and others, including geographic boundaries and outlined the following

questions:

- Will this study gather measurements across all TLDs, or only new gTLDs?
- Will it be able to differentiate between first round new gTLDs and subsequent round new gTLDs?
- Will it be able to differentiate between gTLDs and ccTLDs?
- Will it be able to break results down geographically?

Moving forward, the BC recommends (comment 8) working with the community to review results, both preliminary results and more final findings. It suggests (comment 9) that the CDAR team should allow for additional work should there be a need for more measurements and analysis. It also notes that this may require extending the timeline and cautions that the approach should be realistic in estimating the timeline to completion.

Finally, the BC also suggests (comment 10) that this or a future study ought to consider the impact of new gTLDs' impact on addresses, as well as names, such as the growth in reverse DNS queries, resolving IP addresses to new gTLD domain names. The BC writes that a future study building on the same methodology may explore this topic.

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.

TNO, the CDAR consortium leader, provides the following analysis of the comments.

DK positions the development of a Root Server System (RSS) model as the core deliverable of the CDAR project. Although the model will be an important output, the team emphasizes CDAR's data analysis new gTLDs' impact on the RSS stability to be equally valuable. The CDAR project will provide RSS stability measurement results (as well as the measurement methodology) for the period in which new gTLDs were introduced in the root zone as of October 2013. These measurements will be useful as a point of reference for future measurements.

In addition, the CDAR team will endeavor to measure stability-related metrics of studies that predate the New gTLD Program, insofar as their results are publicly available. In this way, the CDAR project can be linked to measurements in the period before the launch of the New gTLD Program. We acknowledge that any such correlation is constrained by the data and published results that are available to the CDAR team. In the draft study report we will carefully document what data was used for our conclusions and how the data was analyzed.

Furthermore, the CDAR project will investigate relatively invariant stability-related correlations, such as the correlation between the growth of new gTLDs (e.g. in terms of the number of registered domain names) and query rates to the RSS. We are well aware of the fact that such invariants, if any, can only be interpreted with some margins of uncertainty and that such correlation will only reflect "operations on a typical day" (i.e. excluding abnormal events such as the extremely high root query rates observed on 30 November 2015). We do not pretend to present a model with uncertainty margins within a few percent, but rather a model that can be used to assess the order of magnitude, and certainly not a model that will be capable of predicting future RSS instability. Nevertheless, we do see an added value in a RSS stability model that enables us to extrapolate current measurements with regard to potential future growth of the root zone, under ceteris paribus

conditions (i.e. under the precondition that other influences on RSS stability remain invariant).

We underline that we will only draw conclusions based on the data that we will have at our disposal. In the draft report we will carefully document what data was used for our conclusions and how the data was analyzed.

BC's first comment is a request for a clearer indication of what the study can and cannot accomplish and the second comment is about the limitations of the study. We think these are valid comments and they are in line with DK's feedback and reference our response above.

In the third and fourth comment, BC suggests seeking broader community input for the parameters to be analyzed and/or include contributions to the study (for example from those who rely upon the DNS). We emphasize that the CDAR team is already making as many efforts as reasonably possible to collect such input, in particular via public comments, open sessions at ICANN meetings, sessions with ICANN advisory committees, planned sessions with the larger, technical DNS community, and direct engagements with representatives from a broad range of stakeholders. We therefore feel that we are actively reaching out to a broad community and the CDAR study team will of course remain open for constructive feedback throughout our study.

Comments five and seven are requests for clarifying the measurements that we will conduct, beyond those described in the study plan. The CDAR team plans to expand on this further in future public sessions, such as at ICANN 55 and DNS-OARC 24 / IETF 95. In comment six, BC asks whether we will distinguish different parameters for different TLD categories. This will indeed be the case, since this is in line with the nature of study's goals. The CDAR team will collect and perform measurements that will distinguish between TLD categories, and will report on these.

Moreover, in comment six BC suggests a geographic breakdown of results. This suggestion was not explicitly described in our plan, but we will take this into account in the execution of our study.

In comment seven, it is requested the study will data to examine potential contributing factors to metrics such as latency in order to enable risk mitigation. Such an investigation is not in the primary scope of the study.

Comment eight suggests including an additional work package for reviewing preliminary results. We emphasize that the project plan includes several dissemination activities to present preliminary results and to request feedback from the community. The CDAR team may consider further publication of its results via other media, such as relevant journals. The study is, however, constrained in budget and time, which means that we need to be selective in facilitating dissemination events on preliminary results.

Comment nine concerns the timelines for collection measurements for a sufficiently long period of time. The CDAR team agrees that this is a point of attention and will address this during the study. However, we note that we are also collecting data that traces back in time as much as possible, as referenced in our response to DK's feedback.

The additional note (comment ten) suggests expanding the scope of the study to include the impact of numbers in addition to addresses. Numbers are outside the scope of this study.