The undersigned thanks the opportunity to comment on the "Continuous Data-Driven Analysis of Root Server System Stability" report and some of the comments already posted.

First, we would like to point out that the data presented in the report demonstrates that query load on the root server system is not connected to growth in root zone. Despite an increase in the total number of TLDs in the root zone of over three times between January 2014 and January 2016, by 2016 queries for newly-delegated gTLDs accounted for only .4% of the total queries received by the root servers. By contrast, overall traffic to the root servers, likely attributable to the growth of Internet penetration and an increasingly digital life, increased by approximately three times in the overall period. In fact, making the root zone larger may actually decrease overall traffic to the root servers because of differences in behavior between positive and negative caching. First, positive answers have a larger TTL (Time To Live) than answers of non-existence, and second, as demonstrated in previous analysis of queries to recursive servers versus root servers, negative caching is applicable only to specific host names whereas positive caching applies to an entire TLD.

On the report itself, we note that it makes two recommendations not backed by any evidence contained or referred to in the report:

1. It suggests the retirement of popular TLDs as a risk factor. The general discussion around this risk is quite speculative and not tied to any specific observations (such as the recent retirement of .doosan). While we agree that this incredibly improbable situation (after all, why would a popular TLD be retired?) could be a risk, we note that this risk is not unique to new gTLDs established TLDs, and in fact is most applicable to existing TLDs which already comprise a much larger fraction of traffic to the root servers. The report also speculates about possible risks relating to "more dynamic" management of the root zone or traffic increases related to not-yet-delegated TLDs; these concerns seem unrelated to any specific data in the report.
2. It suggests keeping a ceiling on the delegation rate, but does not prescribe one or a method to determine one, and in doing so preserves a current guidance that was also not based on evidence. This "chicken soup" type of caution is common sense, but the previous guidance (1000 per year) was an arbitrary figure thought to be extremely conservative. Given that weekly delegation rates have already exceeded the theoretical cap of twenty per week without incident, we suggest that rather than continuing to impose an arbitrary cap that ICANN use the root server system instrumentation to continuously verify whether the integrity and performance goals are being met in order to determine whether the rate of future delegations needs to be throttled at all.

We would to like to highlight that we are not suggesting further research, which seems to be a common reaction to research reports with no definite guidance, but only to use data already available to the root server stakeholders to do informed decisions.

On behalf of NIC.br and other organisations that will provide their sign-in of this comment later in the process.