

# Summary of the Impact of Root Zone Scaling

## Summary of Comments and Analysis

---

ICANN conducted a public comment period on the Summary of the Impact of Root Zone Scaling from 6 October 2010 to 5 November 2010. Two comments were received in the forum.

### Summary of Comments

ICANN received comments from two individuals, one of which explicitly noted the comments were provided in an individual capacity. An analysis of these comments is provided below. The comment forum can be viewed at <http://forum.icann.org/lists/root-zone-scaling-impact/>.

### Detailed Analysis

#### Stakeholder Comments

Comments in this category included input from Eric Brunner-Williams and “k claffy” both provided via the ICANN comment forum.

#### Eric Brunner-Williams

Brunner-Williams notes that it may not be universally felt that “more change has occurred [to the DNS] in the last 5 or 6 years than has occurred since the DNS was deployed.” This comment in the Summary of the Impact of Root Zone Scaling was derived from various sources, e.g., the DNSSEC Notice of Inquiry undertaken by the U.S. Dept. of Commerce, NTIA stated “A DNSSEC signed root zone would represent one of most significant changes to the DNS infrastructure since it was created.”<sup>1</sup> As such, the combination of the deployment of DNSSEC, IDNs, IPv6, and new gTLDs would represent a combination of changes unprecedented in the deployment of the DNS. However, it is acknowledged that such an evaluation is without defined metrics, thus can be viewed as subjective.

Brunner-Williams takes exception to the use of the term “senior ICANN staff.” This term is used 3 times in the “Summary of the Impact of Root Zone Scaling” document and is derived from wording used in the ICANN Board Resolution which explicitly stated:

*“... involve direct participation by senior ICANN technical staff ...”*

Brunner-Williams notes that the estimate of a maximum addition of “less than 1000 new gTLDs per year” is an estimate of the maximum throughput of the entire processing system, including administrative processing such as the number of contracts ICANN counsel estimates it can process, not simply the IANA root zone

---

<sup>1</sup> <http://www.ntia.doc.gov/dns/DNSSECNOI.doc>

change management process. He feels this explanation of the estimation should be made more explicit.

Brunner-Williams observes that the Draft Application Guide makes IPv6 support mandatory for new gTLDs, however the Summary of the Impact of Root Zone Scaling notes that the IPv6 infrastructure does not (yet) match the IPv4 infrastructure and suggests support of IPv6 in new gTLD registries be made optional.

Brunner-Williams believes the discussion of the impact of IDNs on IANA and registrar systems is inappropriate for inclusion in a report on scaling. The section in question is summarizing the impacts the recent changes to the DNS have had on various systems, thus it is likely a summary that does not include the impacts to IANA and registrar systems resulting from the deployment of IDNs would be viewed as incomplete.

Brunner-Williams takes exception to the comparison between the processing of AAAA records by IANA and VeriSign, believing that ICANN “does bookkeeping and makes requests” while VeriSign “operationalizes every change requested, generates the signed zones, and publishes the signed zones.” While this observation is true, it may minimize the often time-consuming and labor-intensive communication processes IANA staff must undertake to obtain acknowledgement from TLD administrators that the requests are valid and appropriate while potentially over-estimating the impact of the largely automated processes performed by VeriSign. In terms of processing time (as measured from the point the initial request is received by IANA staff to the point that request is forwarded to the U.S. Department of Commerce, NTIA) is typically far greater than the amount of time spent by VeriSign in implementing the request in the root zone.

Brunner-Williams believes there are two aspects of the Summary of the Impact of Root Scaling that are not provided:

1. The impact of time critical consequences when an incorrect change is published
2. How will coordination of all TLD administrators be performed as was useful in the response to the “C” variant of the Conficker worm.

With regards to the first aspect, it seems likely that any incorrect change being published will result in a decrease in the throughput of the root management system as a whole as various root management systems (IANA, NTIA, and VeriSign) are pre-empted to process corrections, thus reducing the potential number of new entrants to the root. Another way of viewing this concern would be to reasonably infer that with a constant error rate, the increase in the number of TLDs will result in an increase in the number of corrections needing to be processed in a time critical fashion. The impact of errors is discussed to some extent in section 2.2 of the “Scaling the Root” report<sup>2</sup> with estimates of response time being between 5 minutes

---

<sup>2</sup> <http://www.icann.org/en/committees/dns-root/root-scaling-study-report-31aug09-en.pdf>

and 36 hours, however this is not mentioned within the Summary of the Impact of Root Scaling.

The second aspect, that of coordination of TLD administrators for response to events such as Conficker, is indeed not taken into account in any of the root scaling reports nor is it addressed in the Impacts section of the “Summary of Impacts of Root Scaling” report as the situation hasn’t arisen in which coordination has been required across a significantly increased number of TLDs. A separate analysis and proposal for mechanisms by which TLD coordination can be assured may be warranted.

#### k claffy

Claffy contends that the “Summary of the Impact of Root Scaling” report makes a number of assertions about empirical evidence yet does not provide supporting citations either for informed estimations or assertions based on data. In several cases, data was derived from experiences of the “L” root server (since “L” root server monitoring data is publicly available at <http://stats.l.root-servers.org>) and extrapolated to the other root servers, making the assumption that all the root servers receive roughly the same distribution (albeit not number) of queries. In other cases, the analysis of impact is based on observed lack of consequence of the deployment of new technologies by ICANN staff either in press reports, Internet news sources such as blogs or Twitter feeds, or in network operations forums such as network operations mailing lists. As such, the use of tentative terms such as “may” and “no significant (if any)” was necessitated due to the lack of comprehensive monitoring mechanisms for the DNS. The “Summary of the Impact of Root Scaling” report explicitly notes that monitoring of root management systems should be improved.

Claffy asks a series of questions:

- “how will icann or others improve monitoring to be able to answer these questions more definitively?” Proposing improvements to the monitoring of root management systems was deemed out of scope for a document intended to summarize the impact of recent changes to the root system.
- “can icann point to supporting documents reporting e.g., how many IPv6 addresses were added in 2010, and how does it compare to what is expected in 5,10,15 years?” IANA staff reports<sup>3</sup> the following historical data:

End of Year	IPv6 Addresses
2007	107
2008	161
2009	226
2010	292

---

<sup>3</sup> Private e-mail with Kim Davies, Manager of Root Zone Services at ICANN

A linear best fit of these data results in the equation  $y = 62x + 41.5$  with an  $R^2$  of 0.9979. Extrapolating this linear best fit results in the following projections of IPv6 addresses associated with existing TLDs:

End of Year	IPv6 Addresses
2011	352
2012	414
2013	476
2014	538
2015	600
2016	662
2017	724
2018	786
2019	848
2020	910
2021	972
2022	1034
2023	1096
2024	1158
2025	1220
2026	1282

Since all new gTLDs are required to support IPv6, the total number of IPv6 addresses in the root zone will be dependent on the number of new TLDs and the number of IPv6 addresses associated with each of the name servers for those TLDs that are entered into the root zone.

- “how often was a signed response requested? what fraction of requests were they?” Data at the “L” root server<sup>4</sup> indicates that between 60% and 80% of queries had “DNSSEC OK” enabled.
- “what fraction of (byte) traffic were they, in both directions?” These data are not publicly available, however from the perspective of scaling the root server infrastructure, what matters is the response size (since responses are invariably larger than queries); regardless of whether the response is DNSSEC-signed. As can be seen from graphs of response size at the “L” root server<sup>5</sup>, the distribution is multi-modal with no response size going over 2.5% of all queries with the vast majority of responses being less than 800 bytes.
- “which root servers are being reported on [for EDNS0-enabled queries]?” The “L” root server was used as a representative of all root servers.

---

<sup>4</sup> [http://stats.l.root-servers.org/cgi-bin/dsc-grapher.pl?binsize=60&window=604800&plot=do\\_bit&server=L-root&yaxis=percent](http://stats.l.root-servers.org/cgi-bin/dsc-grapher.pl?binsize=60&window=604800&plot=do_bit&server=L-root&yaxis=percent)

<sup>5</sup> e.g., [http://stats.l.root-servers.org/cgi-bin/dsc-grapher.pl?binsize=60&window=604800&plot=rcode\\_vs\\_replylen&server=L-root&yaxis=percent](http://stats.l.root-servers.org/cgi-bin/dsc-grapher.pl?binsize=60&window=604800&plot=rcode_vs_replylen&server=L-root&yaxis=percent)

Claffy notes that studies of DITL data show that most EDNS-capable queries can be considered pollution and that at a client level the transmission of pollution has been decreasing over time. However, from the root server infrastructure perspective, it doesn't matter if a "DNSSEC OK" query is pollution or not: the root server must answer regardless. Current measurements at the "L" root server indicate a relatively constant ~70% of queries with "DNSSEC OK" enabled. The fact that a root server should actually only have received 30% of those queries is irrelevant since the root server must respond and the response will consume root server resources. If query pollution is reduced, it will mean increased headroom in the root server infrastructure.

Claffy asks "which servers are meant" when the "Summary of the Impact of Root Scaling" report states ""when the signed root was served from all the root servers, those servers immediately started returning an aggregate of at least 50,000 DNSSEC-related resource records per second". As the footnote associated with that sentence in the report states:

*"Assuming a back-of-the-envelope estimate of an average of 8000 queries per second per root server cluster over 13 root server clusters and with the "DNSSEC OK" bit set in half the queries."*

The intent was to derive a conservative number without access to all data (since that data is not consistently or made publicly available by all the root server operators) to demonstrate the entire root infrastructure was being subjected to a non-trivial number of DNSSEC-enabled queries yet no indications of problems were received by ICANN (or other root zone signing partners). The "8000 queries per second" was chosen to be a rough lower-bound estimate of the aggregate average queries per second across all root server clusters taking into account the more heavily loaded servers (A, F, J) are rumored to receive significantly more than this (e.g., over 20000 qps) and the less loaded servers (D, G, H) are rumored to receive significantly less (e.g., less than 2000 qps). As stated, the estimate for the number of queries with "DNSSEC OK" was 50%, which was the lower bound of such queries received at the "L" root server since recording with the new monitoring system implemented at "L" began in 2007. As such, the derivation of the "at least" estimate is  $13 * 8000 * .5 = 52000$ , which is rounded down to 50,000. The actual average "DNSSEC OK"-enabled queries-per-second is certain to be significantly higher than the estimate used in this "back-of-the-envelope" calculation however higher numbers would imply more chances for something to go wrong that someone would notify the root zone signing partners of.

Claffy notes that "200 is not a rate" in the statement "... ICANN staff estimates that the expected rate of new TLDs entering the root will be on the order of 200 to 300, ...". The intended statement was that the expected rate of new TLDs would be 200 to 300 TLDs per year.

Claffy points out that RSSAC and NTIA have not committed (presumably publicly) to the statement that they "will adjust their resources to meet [root scaling] demand". Given RSSAC is an advisory committee with no means to impose behavior, it isn't

clear a commitment by RSSAC would imply anything useful with regards to root system scalability. With respect to NTIA, SSAC has recommended<sup>6</sup> that the root management partners (which includes NTIA) make statements (individually or jointly) that they are “materially prepared for the proposed changes.” To date, it is unclear whether this recommendation has been formally forwarded by ICANN to appropriate parties for action to be taken.

Claffy asks “what activities are being undertaken to address” the outstanding issue of detecting the increased loads on the root management system prior to those loads becoming an issue, contending the verb “detect” is incorrect for something that hasn’t happened yet and suggesting that if the verb “predict” was intended, ICANN would need to have a predictive model or an “uncontroversially conservative (openly justified) ‘concern threshold’ that would trigger a fresh inquiry.”

As the “Summary of the Impact of Root Scaling” was intended to make observations about events that had occurred since the ICANN Board asked for an analysis of the impact of root scaling activities, it was deemed out of scope to detail what activities should be undertaken to address this outstanding issue. However, as the root management system is an on-going operation with existing operational metrics, some of which are reported at <https://charts.icann.org/public/index-iana-main.html>, detection of increased load could be based on significant deviations from past behavior (that is, the predictive model is that tomorrow will be much like yesterday). Since deviations from past behavior occur with some frequency, the definition of a “concern threshold” is warranted; however how that “concern threshold” is defined is outside of the scope of the “Summary of the Impact of Root Scaling” report.

Claffy asks “is the 1000 TLDs per year considered uncontroversially conservative by the wider community?” The “Delegation Rate Scenarios For New gTLDs” report was posted for public comment from 6 October 2010 to 5 November 2010. During that period, a total of 17 comments were received. While (as of this writing) ICANN has yet to publish an analysis/summary of the comments, it appears that only one comment substantively addressed the projected estimates and that one comment felt the estimates provided by the report to be “20% to 40% higher than we will experience in practice.”

Claffy asks “has anyone else been consulted regarding possible problems with the widespread lack of negative caching if the dictionary at the top level of DNS gets much richer each year?” ICANN staff is unaware of any studies regarding this topic, however given zones exist that are orders of magnitude larger than the root zone and no ill effects have been reported regarding lack of negative caching in those zones, it is unclear as to why the root zone would be special in this regard.

### **Registry Operators, TLD Associations and Internet Organizations Comments**

No comments received.

---

<sup>6</sup> <http://www.icann.org/en/committees/security/sac046.pdf>

## Business Community Comments

No comments received.

## Next Steps

ICANN should review and implement as appropriate the SSAC Recommendations as documented in SAC046.

The comments provided have raised several points in which clarification or expansion of the Summary of the Impact of Root Scaling may be helpful. Four particular areas in which additional study may be warranted were identified:

- How will coordination of a large number of TLD administrators be effectuated for responses similar to the response towards the “C” variant of the Conficker worm?
- How can monitoring of root management systems be improved?
- How can a “concern threshold” be derived?
- What are the potential impacts of the growth of the root zone with respect to negative caching in resolvers?

## Comments Received

---

<i>Individual</i>	<i>Comment URL</i>
Eric Brunner-Williams	<a href="http://forum.icann.org/lists/root-zone-scaling-impact/msg00000.html">http://forum.icann.org/lists/root-zone-scaling-impact/msg00000.html</a>
k claffy	<a href="http://forum.icann.org/lists/root-zone-scaling-impact/msg00001.html">http://forum.icann.org/lists/root-zone-scaling-impact/msg00001.html</a>

---