

SAC 42

SSAC Comment on the Root Scaling Study Team Report and the TNO Report



A Comment from the ICANN
Security and Stability
Advisory Committee
(SSAC)
17 December 2009

Preface

This is a Comment by the Security and Stability Advisory Committee (SSAC) on the following reports: The Root Scaling Study Team report issued on 31 August 2009 entitled, *Scaling the Root: Report on the Impact on the DNS Root System of Increasing the Size and Volatility of the Root Zone* by the Root Scaling Study Team (RSST) and an accompanying report entitled, *Root Scaling Study: Description of the DNS Root Scaling Model* issued on 1 October 2009 by the Dutch organization TNO.

The SSAC advises the ICANN community and Board on matters relating to the security and integrity of the Internet's naming and address allocation systems. This includes operational matters (e.g., matters pertaining to the correct and reliable operation of the root name system), administrative matters (e.g., matters pertaining to address allocation and Internet number assignment), and registration matters (e.g., matters pertaining to registry and registrar services such as WHOIS). SSAC engages in ongoing threat assessment and risk analysis of the Internet naming and address allocation services to assess where the principal threats to stability and security lie, and advises the ICANN community accordingly. The SSAC has no official authority to regulate, enforce or adjudicate. Those functions belong to others, and the advice offered here should be evaluated on its merits.

The contributors to this Comment, reference to the committee members' biographies and statements of interest, and committee members' objections to the findings or recommendations in this Comment, are at end of this Comment.

Table of Contents

1. Introduction	4
2. Comments on the RSST Report.....	5
2.1 Gaps.....	5
2.1.1 Comments on Broad Gaps	5
2.1.2 Comments on Gaps Tied to Specific Sections:	8
2.2 Contradictions, Confusions, and Questions.....	11
2.2.1 General Comments	11
2.2.2 Comments by Section.....	13
2.2.3 Comments By Page Number.....	19
2.3 Small Errors.....	33
2.4 Irrelevant or Inappropriate Text.....	34
3. Comments on the TNO Report.....	38
4. Acknowledgments, Statements of Interests, and Objections and Withdrawals	45

1. Introduction

ICANN Board of Directors' Resolution 2009-02-03-04, February 3, 2009, asked the Root Server System Advisory Committee (RSSAC), the Security and Stability Advisory Committee (SSAC), and ICANN staff to study the potential impact on the stability of the root level of the Domain Name System (DNS) when Internet Protocol Version 6 (IPv6) address records, internationalized domain names (IDNs) and other new top level domains (TLD), and new resource records to support DNS security (DNSSEC) are added to the root zone. The Board resolution asked that the study consider both the technical and operational issues related to expanding the root zone, including the processes for generating, changing, and distributing the zone. From the study, the Board seeks to better understand the impact of each new addition as well as the aggregate effect of including and managing these simultaneously.

In response to the Board resolution, RSSAC, SSAC, and ICANN staff formed a Root Scaling Steering Group to define the parameters and deliverables for a focused study. These are described in a *Terms of Reference* published 5 May 2009 (the "ToR").¹ The Steering Group identified a team of experts to conduct the study. This Root Scaling Study Team (RSST) issued a report on 31 August 2009 entitled, *Scaling the Root: Report on the Impact on the DNS Root System of Increasing the Size and Volatility of the Root Zone* (the RSST Report).² An accompanying report entitled, *Root Scaling Study: Description of the DNS Root Scaling Model* issued on 1 October 2009 (the "TNO Report") provides a quantitative modeling of the root zone to simulate scenarios relevant to the root scaling study.³ These reports were provided as input to RSSAC, SSAC and ICANN staff for review. They were published in parallel for community consideration and to provide transparency to the process.

The SSAC welcomes the opportunity to comment on the RSST Report and the TNO Report. This Comment also includes responses provided by the RSST. This Comment is organized as follows:

1. Introduction: An explanation of the genesis of this Comment;
2. Comments on the RSST Report:
 - 2.1 Gaps: These are parts of the report that were expected but not provided;
 - 2.2 Contradictions, Confusions, and Questions: These are grouped together because what seems contradictory may be only a lack of clarity in the presentation;
 - 2.3 Small Errors: These are actual misstatements, usually specific and relatively small;

¹ Root Scaling Study Terms of Reference <<http://www.icann.org/en/committees/dns-root/root-scaling-study-tor-05may09-en.htm>>.

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

³ Root Scaling Study: Description of the DNS Root Scaling Model <<http://www.icann.org/en/committees/dns-root/root-scaling-model-description-29sep09-en.pdf>>.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

2.4 Irrelevant or inappropriate text: These are political comments or other judgments that are outside the scope and purpose of this report;

3. Comments on the TNO Report; and
4. Acknowledgements, statements of interests, and objections and withdrawals: In the interest of transparency, this is a standard section in our documents to provide the reader information on three aspects of our process.

2. Comments on the RSST Report

2.1 Gaps

These are parts of the report that were expected but not provided.

2.1.1 Comments on Broad Gaps

1. The ToR called for a baseline model of the existing system, complete with current numbers and current capacities. The RSST Report includes some of these numbers, but generally for illustration, not reference. We expected, and believe it is needed, authoritative information that can be verified and referenced.

RSST Response: On the provisioning side, the current root zone management system is a mostly manual process. At the beginning of the root scaling study we were told that a new system automating many of the manual steps—the Root Zone Workflow Automation System (RZ-WAS)—would be deployed by IANA [Internet Assigned Numbers Authority], NTIA [National Telecommunications and Information Administration], and VeriSign “by the time your study is completed.” We concluded that a model of “the existing system” would be obsolete almost as soon as it had been constructed, and decided to model instead the root system as it is expected to be after RZ-WAS has been deployed. Most of the current numbers and current capacities, which quantify the operation of the legacy manual system, are therefore not relevant. We obtained and incorporated data for those parts of the provisioning system that will not change under RZ-WAS, and for those parts of the root system that are not affected by the advent of RZ-WAS (such as the publication/distribution system).

2. The ToR called for explicit information on the errors that have occurred and a model of how the rate of errors might change as the root is scaled. The RSST Report does not include this information.

RSST Response: Our study found only one error that could be reliably documented—the 2005 incident in which a name server address change was not properly validated with all of the TLD registries that relied on the affected name server. None of the actors (IANA, NTIA, VeriSign, root server operators) systematically tracks and records “internal” errors which are detected and corrected before they affect the integrity of the root zone database, and none provided us with any information about the type, frequency, or handling of internal errors. We did not try to build an error model based on a single data point. Appendix B of the TNO report (“Explanation of the

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

error model”) describes the way in which errors in the root zone management process might be modeled, but it is not based on empirical data (which do not exist).

SSAC Response to RSST Comment: We noted that the August 2005 incident was a repeat of a similar incident in November 2004. We have heard, but we do not have the data, that one or two others also have occurred. In addition, the RSST comment on error rates seems to imply that people will do everything humanly possible to keep error rates as close to zero as possible, but the TNO report also makes it clear that error rates are intentionally being under-recorded (by humans). Does the RSST think something needs to change (if so what?) about how errors are defined, reported, and shared, to create reliable input data to an early warning system or does it believe that error handling and the early warning system are independent?

RSST Further Response: We don't read the TNO report as either stating or implying that errors are being intentionally under-recorded by human actors, although it does frequently refer to the "lack of data" concerning actual error rates. However, we think that in any case error detection and reporting is not an important issue for the early warning system, other than in the aggregate. Perhaps a better way to put it is that the propagation of a visible error (one that is consequential with respect to the observable behavior of the DNS) into the root zone would be important input to the early warning system, but no practical benefit to the early warning system would be realized from trying to define a more discrete error recognition and reporting regime within each of the root system components. Because the standard is "zero errors," and the tolerance for deviation from that standard is nil, we could waste a lot of effort trying to get everyone involved to agree on what constituted an "error," how "errors" should be measured, how and to whom they should be reported, etc. Our sense is that it will be more effective, in practice, to build the early warning system on the assumption that each actor is competent (perhaps in regular consultation with other actors) to assess the impact of "what lies ahead" on its own ability to maintain the zero-error standard. If nothing else, such an early warning system - which would be participatory, rather than regulatory or investigative - would require much less instrumentation (some of which would likely be considered invasive) and policing.

3. The ToR called for an analysis of the priming sequence. However, it appears that the RSST Report does not answer the following questions with respect to the priming sequence: 1) How big will the response be if a resolver turns on the DO and CD bits in a query for the list of the name servers for “.”? 2) Will this be the same or different for BIND, NSD, et al? 3) If the responses are too big to fit into a reasonable sized packet even with EDNS0, is there a way to break up the response by asking multiple questions?

RSST Response: The ToR refers to the priming sequence only in the Appendix (“Priming”), which does not ask any of the specific questions cited above. The study team performed an analysis of the effects of root scaling on the priming sequence, and its findings are included in the report.

SSAC Response to RSST Comment: See the following email exchange specifically raising the issue of the need for a parameterized model of the priming sequence.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

-----begin email exchange -----

From: bmannings@vacation.karoshi.com

Date: July 1, 2009 6:50:12 PM EDT

To: Steve Crocker <steve@shinkuro.com>

Subject: Re: Modeling the priming sequence

thanks!

On Wed, Jul 01, 2009 at 06:48:20PM -0400, Steve Crocker wrote:

Bill,

This is a follow up to our brief IM exchange about the priming sequence, which in turn was a follow up to an exchange Lyman and I had in Sydney.

The appendix to the ToR includes the following:

When a validating resolver is first started, it uses a hints file or other initial guess to find a root server, and then it asks that root server for the current list of root servers. The answer is the full list of thirteen root servers and their addresses. Until very recently, that answer fit within the 512 byte limit of a traditional IPv4 packet. With the inclusion of IPv6 addresses for root servers, the response is now longer. Fortunately, longer packets are routinely supported by most transport systems. See SSAC report 018, <http://www.icann.org/committees/security/sac018.pdf>.

However, when DNSSEC signatures are added to the root zone, the response to the priming query will increase yet again. Preliminary examination suggests the response cannot be accommodated within a single packet, so the primary query will necessarily become a priming sequence. Moreover, it appears that responses from NSD and BIND are different, so there is some work to be done to flesh out the details and make sure there is a feasible priming sequence for all of the implementations used across the thirteen root servers.

What I'd really like to have is a parameterized model of the priming sequence that shows the sequence of queries and responses needed to do a full priming for all combinations of the following:

- o Which name server software is responding, e.g. BIND, NSD, or something else
- o Number of root name servers. (This has been 13 for a very long time, but as long as we're making a model, let's include this as a parameter so we can see the effect of changing it.)
- o Number of IPv4 addresses for the root name servers
- o Number of IPv6 addresses for the root name servers
- o Number and length of KSKs
- o Number and length of ZSKs

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

For this exercise, let's assume the algorithms are the usual, RSA with SHA1. If you want to include SHA256, so much the better.

The length and content of each of the responses in the sequence is the information of central interest.

Thanks,

Steve

-----end email exchange -----

4. The load on the root servers is dependent on the number of queries generated in the entire Internet and the effectiveness of the caches throughout the net. Further, there is a sharp difference in the effectiveness of the caches in answering queries for existent versus non-existent TLDs. The RSST Report should have addressed what is known and what is not known about those numbers and how those numbers will change as the root is scaled.

RSST Response: We did not include the Internet fabric (including caches) that intervenes between resolvers and root name servers in our model.

5. The ToR called for an analysis of the impact of IDN TLDs. Although DNAMEs were not explicitly mentioned, they may play a role in handling variants. However, the RSST Report does not address the question of what is the status of DNAME implementation in the root.

RSST Response: The root does not currently support DNAME. We discussed the impact of IDN TLDs with experts involved in the IDNabis work in the IETF and others familiar with the alternative proposals for accommodating variant forms of IDN TLD U-labels, and found no clear consensus for or against the viability of DNAME as an approach to dealing with variants. As most of the disagreement seemed to be political, rather than technical, we did not include DNAME in our analysis of the impact of IDN TLDs.

2.1.2 Comments on Gaps Tied to Specific Sections:

Page 12:

1. The RSST Report states, "Although as noted there have been instances in the past of errors in the root zone file..." It is not clear where this was noted. In particular, we did not see any concrete data in the RSST Report on the history of errors, so there is no baseline from which to work. Nor does the RSST Report include a model of errors developed, although the ToR specifically called for it.

RSST Response: See above, 3.1, number 2.

2. The RSST Report states, "The systems of error checking and correction and of capacity sharing appear to be working successfully." It is not clear what is the basis for this statement or what is its intended interpretation.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

Page 14:

The RSST Report states, “This is well within the computational capabilities of the servers because of the substantial over provisioning of the system.” It is not clear where is the data on the over provisioning.

RSST Response: In Section 4.1.4.

Page 33:

In Section 4.1, first paragraph, we think that the load on the servers should also have been mentioned.

RSST Response: Yes.

Page 38:

In Section 4.1.3, third paragraph the RSST Report states, “The server capacity and bandwidth resources required by the provisioning system... are negligible, and although increasing the frequency of change requests would place additional demands on both, the effect would be too small to be of any significance.” We think that the RSST Report should have provided the basis for these data and their limits.

RSST Response: A better way to state this would have been “...would be too small relative to the effect on other parts of the system that involve human intervention to be of any significance.” The scaling properties of the provisioning system are completely dominated by the steps in the process that require manual inspection. Every other factor is orders of magnitude less significant.

Page 48:

In the last paragraph on the page the RSST Report states, “It is possible that with an increase in the number of new delegations that some types of these queries, particularly those outside the Internet name space, will find themselves included in the Internet name space and therefore cache-able.” It is difficult us to understand precisely what this statement means. Also, a larger problem is that it was specifically requested that the study include consideration of how the load on the root servers might change. Yet, the RSST Report states that the load might change, but it does not offer a quantifiable model.

Page 56:

1. With respect to the discussion of diversity at the top of the page, we think it would have been useful if the RSST Report had included metrics on diversity, if these were available.

RSST Response: No. Diversity in the root server system arises from a belief on the part of most of the root server operators that it is “a good thing” (increasing the resilience of the system by diluting the effectiveness of any single attack vector). For most operators this means deliberately not coordinating or documenting the ways in which they are unlike other operators.

2. With respect to the “Root server adaptation matrix” we note that the RSST Report does not include this matrix.

RSST Response: Yes. The report “suggests the development of a matrix” but does not provide one.

Page 58:

1. With respect to Section 5.4, “The TNO modeling process” we note that the overriding message in this section is that the work was not accomplished.

RSST Response: Yes. The many obvious weaknesses of Section 5 are due, in part, to the limited time available to incorporate TNO’s quantitative modeling work into the study team’s analysis. The quantitative model described in the report (and in the separate report written by TNO, the publication of which was unfortunately delayed, contributing to considerable early confusion) should be considered a first step to demonstrate what could be done, rather than a finished piece of work.

2. With respect to Section 6, “Findings and Recommendations” we note that this section is written with heavy use of subjunctive mode, e.g. “If . . . could” and hence falls short of the explicitly requested baseline description. Some potentially useful numbers are given such as the IPv6 addresses will increase a TLD by a factor of 1.25 and DNSSEC will increase a TLD entry by a factor of 4, but these are not presented as authoritative nor is any reference or derivation shown. We expected that there would be real numbers, backed up with documentation. The comment at the end of Section 6.1.2 about changes due to key signing keys needed elaboration. For example, it would have been useful to know how large the effect is likely to be. In addition, Tables 2 and 3 in section 6.1.5 are filled with “X” instead of useful numbers. Also, it is unclear why the addition of IPv6 addresses changes the number of changes each TLD makes per year.

Page 66-67:

In Section 6.2, “Qualitative Effects” with respect to the discussion of the increase in the size of DNS messages we note that the discussion refers to 512 byte packets, but that this ignores the use of Extension Mechanisms for DNS (EDNS0). We believe that several members of the RSST are very, very familiar with DNSSEC and with EDNS0, so this is a surprising presentation. Further, this section creates inappropriate concern for the reader who may not be familiar with the details.

Pages 70-71:

In Section 6.6, “The effect of adding more TLDs to the root” we note that the secondary effect should have been modeled and measured.

Pages 71-72:

1. With respect to Section 6.7 “The Priming Sequence” we note that this section discusses that the priming sequence may be affected by DNSSEC, but the ToR asked the RSST to work out the effects and how the software used in the different root servers would respond. The last two paragraphs of this section discuss what needs to be done, but the work was not completed. Moreover, the discussion in the third paragraph about room for only two quad A (AAAA) records in a 512-byte response is very much out of date. RSSAC and SSAC spent quite a lot of time on this specific topic, resulting in SAC 018 in March 2007. Three of the RSST are members of one or the other of these committees and are presumably aware of that work. Thus, it is not clear why the RSST Report includes no reference to that work.

2. We note that Section 6.11.1 “Priming Query and Response” does not address the ToR, and it also ignores the work already done with EDNS0.

2.2 Contradictions, Confusions, and Questions

These are grouped together because what seems contradictory may be only a lack of clarity in the presentation.

2.2.1 General Comments

The RSST Report states that the minimum time from when a change enters the system at IANA and then appears in the root zone is two business days. Further, it emphasizes business days and describes them as 8:00 a.m. to 5:00 p.m. during a typical work week that also takes into account holidays. This means that the actual calendar time from submission to publication could be as long as six calendar days over certain holidays depending on exactly when the request came in. This presents a significant issue when DNSSEC is deployed.

Specifically, the RSST Report does not address the timing constraints that are characteristic of specific types of requests. For example, what if a TLD needs to execute an emergency key rollover? In this circumstance a TLD would probably want at least two things to happen. First the compromised key needs to be removed from the root zone and second the new key needs to be inserted in the root zone. We assert that, in this case, these actions need to be executed in hours, not days. The RSST Report should have identified this general issue and, given that it also recommends that DNSSEC be deployed sooner rather later, it should have called out this specific "problem" in the provisioning system, as it currently exists.

The RSST Report mentions why the root is different to those TLDs mentioned above, and why the issues for root zone provisioning are different. We think the RSST Report should have amplified that point. For example, there are several factors here that distinguish the root from TLDs:

1. As far as we are aware, the root is the only top-level registry that maintains no control over who or how name service for its zones is provided. Change requests involve multiple parties, where the organization that receives/vets the change request is different than the organization that authorizes the change which is different than the organization that edits and publishes the zone. The processes by which changes are made, some of which are not well known or understood, are very difficult to change.
2. What else differentiates the root in such a way as to make growing to the size of other TLDs infeasible? In particular, the RSST Report should have specified improvements to coordination among root zone system players. The recommendations presented in the RSST Report are predicated on the idea that the root management process in place today will need to be modified to meet the demands of an evolving root zone. The RSST Report states that in order for an early warning system to be effective it will require new agreements and channels for communication among the participants of the root management process (page 3). Thus, it would have been useful if the RSST Report would have specified what agreements construct a more comprehensive picture of risk amongst all players. Also, it would have been helpful if the RSST Report described the

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

techniques that exist (or are needed) to halt growth if the early warning system begins to show signs of weakness.

3. The RSST Report should have addressed risks and implications associated with any increase in size in the root zone. The RSST Report states that, "risks associated with doubling or tripling the size of the root zone can be mitigated without hitting a discontinuity in the scaling properties of any of the root system management functions (page 4)." However, this statement seems to contradict the report's overall recommendation that any change to the root zone should be introduced through a careful and coordinated process (a process not in place today) to minimize risk. As the statement stands, it has the potential to create a de facto growth threshold below which improved coordination among all players in the root zone management process would not be required. The existence of the RSST Report suggests that more coordination, policy, and processes are needed to ensure that *any* change to the root zone is conducted in a safe and carefully planned manner.
4. The RSST Report should have fully explored risks related to the introduction of a limited number of fast track Internationalized Domain Names. ICANN is continuing to move forward with the IDN fast track process (see Status Update: IDN ccTLD Fast Track Process Implementation, September 9, 2009, available at: <http://www.icann.org/en/announcements/announcement-2-09sep09-en.htm> <<http://www.icann.org/en/announcements/announcement-2-09sep09-en.htm>>). The RSST Report posits that introducing IDNs into the root zone before DNSSEC would cause undue risk, but it does not specifically comment on whether the IDNs targeted for fast track (a smaller number of IDNs) could be accommodated. Does the community anticipate fast track introduction to also be a problem if introduced concurrently with DNSSEC?

RSST Response: The IDN fast track program anticipates roughly 50 new TLDs over a period of 1-2 years (although most of those are expected to arrive during the first few months). Nothing in the study results suggests that adding such a small annualized number of TLDs would push any root system actor out of its normal operating range, even if it occurs concurrently with the introduction of DNSSEC.

SSAC Response to RSST Response: This comment appears to contradict a finding of the RSST on page 5 of its original report: "With aggressive re-planning (some of which is already underway), the system is capable of managing the risks associated with adding either (a) DNSSEC or (b) new TLDs, IDNs, and IPv6 addresses over a period of 12-24 months, but not both."

RSST Further Response: There's no contradiction, although we agree that the wording of the comment and the finding on page 5 of the original report make it easy to infer otherwise. The comment says that the root system can absorb "a small annualized number of TLDs" without pushing any of the actors out of its normal operating range. Adding "even if it occurs concurrently with the introduction of DNSSEC" implies that we could do both DNSSEC and fast-track IDNs without pushing any of the actors out of its normal operating range. In fact, just doing DNSSEC creates that push, as "[w]ith aggressive re-planning..." on page 5 of the original report suggests. What was meant in the comment is that also adding the

small number of fast-track iTLDs would not *in itself* cause any actor to hit that discontinuity - adding fast-track iTLDs to the list of "things we're going to do to the root" would have an insignificant additive effect on stability risk. This is admittedly an obscure point; if you're already re-planning to deal with the impact of DNSSEC, it's of small matter that you could also add a few iTLDs without making your problems measurably worse.

2.2.2 Comments by Section

Section 2.2:

1. With respect to the root zone file and the following sentences "the number of TLDs listed therein is 28"... "There is also a small number of resource records that relate to the root itself. This leads to a total number of 2,942 records". We found these sentences to be a little confusing.
2. With respect to the phrase, "...mutually update each other's BGP tables using the router-to-router pathways defined by the protocol...", we found this to be unclear. Perhaps the phrase "dynamically update each other's tables using the BGP protocol" would have been clearer.
3. Could the ".NET" have been replaced with ".com" or ".net"? This is largely introductory text and many people think of the Microsoft framework when seeing ".NET."

Section 2.3:

1. This section talks about the introduction of a "distribution master [...] operated by VeriSign, Inc." In fact there are multiple such distribution masters. Figure 1 on page 16 correctly shows "distribution masters" (plural).
2. The RSST Report refers to "The IANA" in various places (e.g. 2.4.2), which we casually note in passing might be contentious to those who insist that no such entity exists, and who might suggest that the correct terminology is something like "The IANA Functions Operator."

RSST Response: Anyone who reads the report closely enough to spot this in Table 1 deserves a finer label than "casual reader"!

Section: 2.4.5:

In this section, the RSST Report describes the RIPE NCC as "a cooperative body of European Internet Service Providers". We think it might have been informative to mention that they are also a regional Internet registry (RIR), in the interests of characterizing an additional dimension of community represented amongst the various operators.

RSST Response: That's a good point—the RIPE NCC is both the operator of the K root and one of the five regional Internet registries (RIRs).

Section 2.4.6:

The RSST Report uses the term "Iterative Mode Resolvers (IMRs)". This seems to be a new expansion of IMR. We are not convinced that anybody actually uses these terms in practice.

Section 4.0:

This section makes a “qualitative” case that the operational issues will prevent an explosion in TLDs (i.e. it would take drastic changes in operational modes at several bureaucracies), but also admits we have no way to model those operational parameters, and we only know their values will change in the future, in directions that promote expansion.

Section: 4.1.4

In this section the RSST Report notes that root server operators "observe the prudent network engineering rule of thumb that critical systems should maintain at least 10% overhead in capacity [...]". We don't know many commercial network operators who think that running networks or systems at 90 percent is prudent. However, that aside, the following paragraph notes that "as a whole the system has tended to keep between 50% and 100% overhead in critical resources available at all times," which we believe is quite an understatement. Based on our experience with a few root servers, we would have conservatively estimated headroom in terms of query-answering and associated network capacity of the order of 100x the steady state load (so, 10,000 percent). In addition, the headroom maintained in the current system seems like it has direct and relevant implications for the rate at which it is safe to scale the root zone. This seems like an important number to get right, and hence we think some justification for the number is warranted. For example, Figure 6 on page 76 would look very different if the headroom was 100x the steady state traffic levels. In that diagram it's more like 1x.

Section 5.4:

The RSST Report references "O(100)". However, O(N) suggests a computational complexity notation, which could mean "in the order of 100 times," which made very little sense as compared to the data. In addition, there is a problem with the notion of "scale by O(100)" without context. For example, if the current 80,000-byte file grows by O(100), it is unclear what that means in terms of ratios of new/IDN TLDs to DNSSEC RRs to AAAA records or in terms of numbers of changes. More specifically, based on the discussion of the workflow between IANA, VeriSign and NTIA in Section 4, we find it hard to understand the notion that a hundred-fold increase in adds/drops/changes is within the normal operating region of all parties involved. Also, we are uncomfortable with this number given the complexity and highly human-dependent nature of the RZM. Moreover, it is not clear why the threshold is O(100) and O(1000), as opposed to O(10) and O(100) or O(1000) and O(10,000). If the bottlenecks the RSST Report identified are primarily human checking functions, then the RSST Report appears to be saying that those manual checks scale by 100 but not 1000.

RSST Response: Yes—we should have used a scaling convention and nomenclature other than “O(X),” which has been widely misinterpreted. In the report, “O(100)” doesn't mean “100 times as many” or “hundredfold”; it means “an increase on the order of hundreds” (of new entries, or of additional change requests). Starting at the current size of the root, an “O(100)” increase would represent a 2x or 3x (roughly) increase, not a 100x increase.

The “O(100)” and “O(1000)” ranges are orders of magnitude, not thresholds (in the sense of crisp boundaries). They are based on the data we collected from IANA and NTIA people about (a) the manpower resources currently allocated to their root zone management tasks, and (b) how they manage the staffing of those tasks within their respective organizations. With normal HR [human resource] and organizational management processes, they can scale at O(100); if they

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

had to scale at $O(1000)$, they could do so, but they would have to make significant changes to the way in which they operate (and they could not necessarily make those changes quickly). This finding was not constructed mathematically; it represents the analytical judgment of the members of the study team, who spent months reviewing the information collected from all of the root zone management system actors.

A fully quantitative model would be useful only as a continuously updated component of an early warning system. It makes no sense to expect a quantitative model to produce meaningful simulation results beyond the very near term, because the root system is highly decentralized and adapts dynamically to change in ways that can't realistically be modeled. For example, IANA might respond to increased load (scale) by continuing to operate the way it currently does, hiring more people to do the additional work. But it might also respond by changing the way it operates—perhaps reducing or otherwise changing its dependence on the trust relationship that is currently established and maintained with each administrative contact for all TLDs, or automating a process that is currently performed manually. NTIA might respond to increased load by adding staff to perform a larger number of the same checks it performs today. But it might also respond by modifying its approval policy so that different types of change request were subjected to different levels of scrutiny. There is no meaningful quantitative answer to the question “what would happen to the root system if we added a million new TLDs?” because the system would start to evolve in both predictable and unpredictable ways as soon as the scaling process began, quickly invalidating whatever baseline had been used to construct the model. The root server operators, for example, are already re-planning in expectation of a signed root by the end of the year (and they have been in that re-planning region since at least June 3, when the formal announcement was made).

SSAC Response to RSST Comment: With respect to the comment above, we assume that is why the RSST Report recommends that, “Root system oversight should focus on early warning rather than threshold prediction.”

RSST Further Response (1): Yes, precisely.

SSAC Response to RSST Comment (2): Could the RSST comment on exactly what it thinks could and/or should be usefully measured/shared as part of this early warning system which it deems a reasonable substitute for modeling and is this something the RSST recommends being in place before X, for some value of X, and if so what?

RSST Further Response (2): Choosing metrics and measurement protocols (using “protocol” in the sense of a defined policy-governed process, rather than the narrower sense in which it appears in (for example) “transmission control protocol”) for an early warning system is of course a “topic for further study” (section 8.1 of the report). However, during the study we captured at least the following observations:

1. The goal of an early warning system would be to detect signs of stress at critical points in the root system, so that plans could be made and executed to deal with the stress before it caused problems. The metrics and measurement protocols best suited to detecting early signs of stress may or may not be the same as metrics and measurement protocols developed to serve other goals, such as the more general query rate, query type, IPv6

support, and other statistics relevant (for example) to the Day in the Life of the Internet (DITL)⁴ goal of compiling a database that comprehensively profiles Internet traffic under realistic operating conditions. For example, the rate at which queries arrive at a root server is a useful measurement in the context of an effort to capture data that describe how traffic flows in the real-world Internet. It says almost nothing, however, about whether or not a root server is approaching a discontinuity in its ability to process queries.

2. Our discussions with root server operators suggest that at least one relevant measure of stress is the amount of time by which the state of the root as expressed in responses to queries at a particular root server instance lags the instantaneous authoritative state of the root as expressed by the contents of the root zone database maintained by VeriSign. (Actually, if we model root zone latency using the contents of the distribution masters, rather than the contents of the root zone database, as the authoritative state, we can factor out the "sawtooth" component introduced by the periodic (rather than immediate) update of the DMs, which produces a more intelligible graph of the latency function.) Informally, this measures the distance (in time) between the "real" root and the root that is being served by a given root server instance. Most of the important distribution-side parameters contribute to this measurement - how long it takes a root server to obtain a new zone file from a Distribution Master (DM) after receiving a NOTIFY; how long it takes a root server with anycast instances to propagate a new zone file to all of its instances; and how long it takes a root server to begin responding to queries with data that reflect the state of the root captured in the newly received zone file.
3. Although it would be natural to assume that "error rate" would be an important metric for early warning, our study suggests that this is not the case for the root system, in which the pressure to prevent the propagation of even a single error is intense. At least analytically (as we have reported elsewhere, we were able to document only one actual error, which effectively forestalled the development of a data-based error model), we see no clear relationship between potential errors and scaling stress. To put that in more concrete terms, we believe that each actor in the provisioning system would not under any circumstances allow the risk of error to increase by even a small amount, and would take whatever steps were necessary to ensure that. [Process throughput, or process throughput variance, might be shadow metrics that track what would otherwise be "error rate."]
4. We do not recommend that we pursue the "early warning" goal by instrumenting the root system in such a way that an independent party would collect all of the various "early warning" signals and decide when to raise appropriate alerts. A more effective (and also more readily implementable) approach to constructing an early warning system would engage the existing root system actors, the measurements they already (in almost all cases) make internally, and a neutral fabric through which the self-assessed "stress" status of each actor could be effectively communicated and integrated.

⁴ See <<http://www.caida.org/projects/ditl/>>.

SSAC Response to RSST Comment (3): With respect to your goal of an early warning system, isn't this what the Domain Name System Operations Analysis and Research Center (DNS-OARC) was created for, and most if not all rootops (and ICANN, and research groups) are members?

RSST Further Response (3): If one looks at the "DNS-OARC Governance Update" presented recently it's hard not to think of [Operations, Analysis, and Research Center (OARC)] OARC as primarily a forum for root operators and researchers - it comes across with a strong flavor of IETF/dnsops combined with a DNS-specific [North American Network Operators Group (NANOG)] NANOG (or *NOG), invigorated somewhat beyond those analogies by the research component. That doesn't mean, of course, that OARC couldn't take on the additional role of designing the early warning system and at least participating in its operation. Operating the early warning system in such a way that it has a meaningful effect on policy almost certainly will require some sort of "inter-agency" arrangement anyway. The other obvious candidate for an early warning system role is RSSAC. We have been somewhat surprised to hear almost no discussion of this, either in the context of the current "organizational review" of RSSAC or elsewhere (although we are on the RSSAC list, and may easily have missed not just something but everything).

Section 5.7:

The parameters noted in Section 5.7 are measurable, yet the RSST Report does not include the following:

1. Number of TLDs (affecting both provisioning and publication);
2. Bandwidth variations in the publication infrastructure (affecting only publication);
3. Processing time of various steps (related parameters: working hours in a day, working days in the week);
4. The mix of requests for social data or name server changes; and
5. Time waiting on actions from actors outside the system (such as delays in getting confirmation from zone managers for change requests).

Section 6.2:

1. With respect to new TLDs and IDNs these have primarily quantitative effects on the root system, but the RSST Report asserts that these are more of the same. We think this is a baseless assertion.
2. The RSST Report states, "The primary effect of adding new TLDs to the root will be felt in the distribution of the zone data to the root servers themselves, but a secondary expected effect is increased traffic to the root name servers, due to the reshaping of the name space. New TLDs will have a much smaller effect on the size of any response to a query since, as delegation records, they are not signed in a DNSSEC enabled zone." While it seems reasonable to anticipate more traffic to the root due to flattening the namespace, it is unclear how this traffic would compare to bogus traffic now getting to root servers. Moreover, this effect would depend on the extent to which name lookup is moved to a new TLD rather than just replicated at a TLD while significant lookup traffic still flows to the (slightly more)

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

hierarchical domain name. We believe that the traffic conclusion is much weaker than other conclusions in the RSST Report. In addition, while delegation records do not have signatures, the RSST Report should have indicated that there is a little growth in the record size for the designated signer (DS) record(s).

3. In Section 6.6 there is some discussion of the potential impact on poorly connected regions of the world that currently enjoy the benefits of a local instance of a root sever. We think the RSST Report could have made the point more forcefully that the ongoing distribution of root server instances to tenuously connected corners of the Internet is laudable goal but may be difficult as the size of the root zone expands. Careful study of the interactions between the distribution process and the size of the root zone should be carried out.

Section 6.8:

A representative quote: "Not only must the entire block [of 40,000 famous marks as TLDs] be added, but once begun the process cannot be reversed, nor can it be stopped or even slowed for a significant period of time." This assumption appears to be unsupported. If there is a clear indication of strain in the root system (and we believe the strain would be seen much earlier in the provisioning side than in the operational side), then there is clear and supportable justification to change behavior.

Section 6.11.2:

The RSST Report states, "Root server operators and DNS management companies report that normal zone transfer to any site (not just poorly connected 'distant' sites) becomes infeasible when the number of records in a zone file reaches approximately 20,000,000. The operators of large ccTLDs give similar reports. The experience of these operators suggests that managing a zone with 1,000,000 records is readily accomplished with today's technologies and root server architecture, and that at some point after a zone grows to 10,000,000 records it becomes unmanageable without significant change to the server architecture."

We would like to make two points with respect to this statement:

1. As has been presented in public forums before, Afilias' name server platform (those nodes they run themselves and those nodes that are supplied and operated by PCH⁵, as a conscious effort by Afilias to promote diversity) use BIND9⁶ and Name Server Daemon (NSD)⁷

⁵ PCH = Packet Clearing House, <http://www.pch.net/home/index.php>.

⁶ "BIND (pronounced /'baɪnd/, for *Berkeley Internet Name Domain*, or *named* (/ 'neɪm.di:/, "name D"), is the most commonly used Domain Name System (DNS) server on the Internet^[1] On Unix-like systems it is the *de facto* standard...A new version of BIND (BIND 9) was written from scratch in part to address the architectural difficulties with auditing the earlier BIND code bases, and also to support DNSSEC (DNS Security Extensions). Other important features of BIND 9 include: TSIG, DNS notify, nsupdate, IPv6, rndc flush (remote name daemon control), views, multiprocessor support, and an improved portability architecture. rndc uses a shared secret to provide encryption for local and remote terminals during each session." Wikipedia < <http://en.wikipedia.org/wiki/BIND>>.

⁷ "In Internet computing, NSD (for "name server daemon") is an [open-source](#) server program for the Domain Name System. It was developed by NLnet Labs of Amsterdam in cooperation with the RIPE NCC, from scratch as an authoritative name server (i.e., not implementing the recursive caching function by design). The intention of this development is to add variance to the "gene pool" of DNS implementations used by higher level name servers and thus increasing the resilience of DNS against software flaws or exploits." Wikipedia <<http://en.wikipedia.org/wiki/NSD>>.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

version 3, and all distribution of zone data is done using Asynchronous Full Transfer Zone (AXFR), Incremental Zone Transfer (IXFR), and NOTIFY.

2. The .org zone is already substantially bigger than the 10,000,000 records that the RSST Report cites as being the point where "it becomes unmanageable". AXFR/IXFR appears to be working for Afiliacs in conditions that far exceed what we might expect for the root zone (e.g. many registry changes per second, registry changes propagated to name servers globally in under a minute, many remote nodes). Moreover, we think that this particular topic would have benefited from the presentation of real-world data, if not results from a balanced experiment.

2.2.3 Comments By Page Number

Page 3:

1. We do not think it is appropriate for the RSST Report to assert that, "Adding support for DNSSEC, for example, has a negative impact on root system stability" without context in the Executive Summary.

RSST Response: The context is stated in the bold-face first sentence of that paragraph: "Any increase in the size or volatility of the root zone involves risk." The statement for which the Executive Summary provides no context is the continuation of the sentence you quote: "...but a positive impact on DNS and Internet security"—that's accepted as given (it's not a finding of the study).

2. The RSST Report includes the following phrase "...early warning system than threshold protection," which implies some sort of early warning system, but the RSST Report does not discuss what that might mean, who measures what, or what actions are taken on what triggers.
3. The RSST Report states, "On the provisioning side, the ability to scale the root is completely dominated by the steps that involve human intervention." We are not sure how to interpret the verb tense since there continue to be overwhelming economic, political, and technical pressures to automate whatever can be automated.

RSST Response: The "human intervention" steps that remain after RZ-WAS are there not because a good way to automate them has not yet been developed, but because all three provisioning actors (IANA, NTIA, and VeriSign) have made deliberate policy decisions to have a "human in the loop" on every change request. All three told us that they have no plans to revisit or change these policies.

SSAC Response to RSST: Do you have confidence that "having no plans to revisit or change" means it will not happen? NTIA's letter dated 18 December 2008⁸ suggests a number of steps ICANN must take to prepare for new TLDs prior to implementing them, so it is unclear whether the NTIA is in the process of planning for them.

⁸ See National Telecommunications and Information Administration letter dated 18 December 2008 at <http://74.125.47.132/search?q=cache:S_G10ecbZKAJ:www.ntia.doc.gov/comments/2008/ICANN_081218.pdf+N TIA+letter+new+TLDs+October+2008&cd=1&hl=en&ct=clnk&gl=us&client=firefox-a>

RSST Further Response: We don't know how to calibrate the "no plans to revisit" assertions. Each actor is convinced that it can't maintain a zero error rate (even when we agree that "zero" actually means "really, really, really close to zero") without human inspection. To the extent that this is in fact true, the "zero" error rate mandate is clearly at odds with the expectation that each actor will scale up to handle a much larger and more volatile root.

Page 4:

1. The summary on this page considers how the human inspection rate limits root zone changes, but the RSST Report makes no comment regarding the increased potential for human error. Since one popularly expressed concern is the increase in the number and frequency of changes, we think that anticipating and compensating for human error in work flows merit consideration. (We do not intend any disrespect to the parties who manually inspect root zone information.)

RSST Response: This is a good point, and the additional human resources required to maintain a target "zero error" rate as load increases probably should be given greater attention as a driver of the super-linear growth in resource requirements described in Sections 4.1.2 and 4.1.3.

2. The discussion in the RSST Report on "A much larger root zone" is not specific. In particular, it is not clear what the RSST Report means by the term "much". For example, does the RSST Report mean $O(100)$ or $O(1,000,000)$? In addition, is it only size that matters and not frequency of publication? Using the quoted ~80K size for the current root, the RSST Report should have included an example such as "A root zone of $O(1000)$ the current size that changes hourly requires connections capable of transporting 80 Megabytes in less than one hour to assure global root zone accuracy." Such an example would have revealed several dimensions to the problem not easily gleaned from the Report text.

RSST Response: The non-specific usage is appropriate in the Executive Summary; the details are in Section 6.9.1 and 6.11.2. The report talks about several approaches that root server operators might take to distributing a larger root zone, including a shift from AXFR to IXFR updates and a complete change in the way in which the zone file is distributed (e.g. shifting away from DNS protocol XFRs to some other mechanism to move the file reliably to remote sites).

3. It is unclear whether the claim "[$O(100)$] can be managed without changing any actor's current arrangements" is corroborated in the context of human inspection and error. For example, it would have been useful to include how many changes will be manually inspected per work hour in this scenario.

RSST Response: The human inspection steps in the processes at IANA and NTIA are not like assembly-line operations, in which individual workers are dedicated to a specific task and the speed with which each worker can perform that task can be accurately quantified. Neither IANA nor NTIA is amenable to a classical time-and-motion study. What we determined, based on analysis of the processes that IANA and NTIA follow, the distribution of processing times for completing their checks, and those organizations' self-descriptions of the way in which they are currently prepared to manage growth, is that each could absorb an annual increase in the number of entries in the root zone or the number of changes to the zone on the order of "hundreds"—hundreds of new root zone entries, or hundreds of additional change requests, per year—without being forced outside their normal operating region. That means that with growth at $O(100)$, their

normal processes for adjusting work assignments, bringing in new hires, training people to perform different types of work, etc. would enable them to “keep up.” That does not mean that if growth exceeded $O(100)$ they would be “unable to keep up”—it means that growth beyond $O(100)$ would force them out of their “normal operating region” into what we have called a “re-planning region,” in which significant changes would have to be planned, justified, approved by senior management, and provisioned.

4. The RSST Report states, “These bottlenecks govern the rate at which root zone changes can be processed; by comparison, every other factor is ‘in the noise.’” While we realize this is in the subsection “On the provisioning side,” it would have been useful for the RSST Report to reassert this with the following text: “These bottlenecks govern the rate at which root zone changes can be processed; by comparison, every other provisioning factor is ‘in the noise.’”
5. With respect to “On the publication side...” the text makes the statement that, “Scaling the root is likely to place additional demands on those operators who use Internet Protocol (IP) anycast⁹ to deploy root servers in economically less-developed parts of the world.” While we understand the point being made here, we would observe that not using IP anycast deploying root servers in resource constrained environments would introduce additional demands on scaling the root. That is, there is nothing particularly special about IP anycast here, and scaling the root surely involves expanded footprints and availability of services in environments.

Page 5:

1. The RSST Report states, "If a choice must be made, DNSSEC should come first." The basis for this decision seems to be "prove the resiliency of the root zone system by introducing the most stressful change first."

RSST Response: It isn't necessary to invoke “seems to be” to discover “the basis for this decision”—it’s right there in the same Executive Summary paragraph: “Because the step-function impact of signing the root will be proportionally greater the larger the root becomes, deploying DNSSEC before the other changes have increased the size of the root would significantly lower the risk it presents to DNS stability.” That has nothing to do with “proving the resiliency of the root zone system.”

2. In contrast, one could argue that adding IPv6 resource records will become a necessary condition when IPv4 space is exhausted. It seems the broader conclusion, which we support, is that adding DNSSEC and IPv6 resource records to the root should take precedence over new TLDs and IDN TLDs.

RSST Response: IPv6 resource records have already been added to the root for roughly 20 percent of TLDs, and given the importance of IPv6 to the Internet’s future, it’s very hard to imagine a scenario in which IANA would reject a request from a TLD registry to add IPv6 address information to its RRset.

⁹ “Anycast is a network addressing and routing scheme whereby data is routed to the “nearest” or “best” destination as viewed by the routing topology.” Wikipedia <<http://en.wikipedia.org/wiki/Anycast>>.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

With respect to RSST Report discussion on “The risks associated with DNSSEC...” the conclusion that it is either (a) DNSSEC OR (b) new TLDs, IDNs, and IPv6 addresses seems to be rather loose and unqualified in the text. Furthermore, the RSST Report seems to completely ignore the fact that 222 AAAA records are already present in the root zone (7 for root servers, 215 for TLDs) – with approximately 20 percent IPv6 AAAA penetration already.

Page 11:

The RSST Report states, "it is not modified in any way by anyone in the provisioning or publication process." It would have been helpful to provide a footnote that the RSST Report is describing the publication process associated with what is popularly called "the authoritative root zone file." In practice, private organizations and alternative root operators routinely modify this file, see the SSAC publication “SAC009: Alternative TLD Name Systems and Roots: Conflict, Control and Consequences.”¹⁰

RSST Response: Yes, but as the reference occurs immediately after the URL of the authoritative root zone file is cited, the context is reasonably clear (although there would certainly be no harm in reinforcing the point).

1. After stating, "However, relatively simple error detection and correction procedures can prevent such errors" the RSST Report identifies an error detection method but does not explain how this error is corrected in practice. We think this is too important a detail to omit.
2. With respect to the statement "there has been none in recent times" we think the RSST Report should have changed "has" to "have" and quantified "recent" (i.e., since 2007?)

RSST Response: We used “recent” because it’s not clear precisely when, other than “not recently,” the last instance of a significant error occurred. As noted above, we were able to reliably document only one error that resulted in incorrect data appearing in the authoritative root zone database.

3. The RSST Report states, "However, relatively simple error detection and correction procedures can prevent such errors." This identifies an error detection method, but the RSST Report does not explain how this error is corrected in practice. We think this is too important a detail to omit.

Page 14:

1. With respect to the statistics on root servers it is unclear why the RSST Report does not use OARC DITL data.

Page 15:

1. The RSST Report statement "...unexpected failure of a computer or communication link" depicts single points of failure that we hope are not present in root operations. Surely no single personal computer or communications link will cause a root name server to become non-operational? Do root operators really become non-operational during maintenance cycles?

¹⁰ See < <http://www.icann.org/en/committees/security/alt-tlds-roots-report-31mar06.pdf>>.

RSST Response: This should have been stated more clearly; the reference to “one or more of the root name servers” is to individual server instances, not to a “root name server” operation as a whole. It’s easy to deduce this from the remainder of the paragraph, but it would have been better to use a more precise term at the outset.

2. With respect to “However, *sustained* problems...” we note that another reason would be a sustained malicious attack on the system as a whole by a very large number of computers – or a very effective attack vector.
3. The following statement appears to contradict the finding on page 5 of the RSST Report: “Because the step-function impact of signing the root will be proportionally greater the larger the root becomes, deploying DNSSEC before the other changes have increased the size of the root would significantly lower the risk it presents to DNS stability. That has nothing to do with proving the resiliency of the root zone system.”

RSST Response: We don't see how this contradicts either our comments with respect to Section 4. Contradictions, Confusions/Questions, subsection 4.1 General Comment (above) or the finding on page 5 of the original report. It just says that the smaller the root when DNSSEC is deployed, the less stability risk to the system. One could combine this with the our comment on subsection 4.1 to conclude that "the increase in the size of the root that would occur with the delegation of a small number of fast-track iTLDs is too small to have a significant effect on the cost/risk of deploying DNSSEC in the root."

SSAC Response to RSST Comment: If you are already re-planning to deal with the impact of DNSSEC, it is a small matter to add a few IDN TLDs without making your problems measurably worse. So, is this why the RSST Report recommends adding a few (fast-track 50-ish IDNs), but not more new TLDs/IDNs in the 12-24 months after introducing DNSSEC?

RSST Further Response: Yes - partly because the rules for the fast track ensure that it will be limited to a relatively small number of new TLDs, but also because at this point (particularly after the Board resolution in Seoul) [at the ICANN annual meeting in Seoul, South Korea on 30 October 2009¹¹] it would be very difficult for ICANN to renege on (or even substantially delay) the fast track for IDNs. As we've said before, the question is not "can the system handle X?" but "what would the system have to do in order to handle X?" - and then you decide whether or not all of the various actors are willing and prepared to deal with what the system would "have to do." So when I say "but also because..." above, we are not suggesting that an aura of inevitability is a justifiable reason to go ahead and do something that would otherwise be considered unacceptably risky; we mean that people and organizations generally agree to accept risk in some (ideally, rational :-)) proportion to the expected reward. To the extent that ICANN and its decisions represent "what the community wants," the community has decided that the benefit of IDN ccTLDs is great enough to justify the risk of expecting the root system to absorb whatever increased load those IDN ccTLDs will represent. The question, again, is not "can the root system handle ~50 new iTLDs?" but "what will the root system have to do to handle ~50 new iTLDs?" We can (and during the study did) directly observe the effort that each of the root system actors is already investing to prepare for DNSSEC, and ask what each of them would have to do, in

¹¹ See Approval of Final Implementation Plan for IDN ccTLD Fast Track Process at <<http://www.icann.org/en/minutes/resolutions-30oct09-en.htm#2>>.

addition, to also support the fast track's limited introduction of new iTLDs. What we heard from the people we talked with during the study is that the additional capacity (of various kinds, depending on the role of each actor) required to support the introduction of a limited number of iTLDs (where the word "limited" is really important) is so small relative to what they are doing to prepare for DNSSEC that it would be almost unnoticeable.

As a practical matter, we can probably agree that the following will happen, regardless (or almost regardless) of what we do or don't say at this point about root scaling:

- a. We'll continue to steadily add AAAA records and IPv6 glue to the root. Considering the rapidly increasing pressure on the Internet as a whole to move to the v6 address space, it's essentially impossible to imagine that IANA would reject a request from a TLD registry to put v6 addresses for its name servers into the root.
- b. The root will be signed more or less in accordance with the timetable that Matt [Larson] (and others) have described to a variety of audiences over the past few months.
- c. The fast track will add roughly 50 IDN ccTLDs to the root over the next 12 months. I'm sure that there are arguments to be made about the details, and perhaps about how these three things should be ranked in order of certainty, but I have no doubt that the root system will be expected to absorb all three of them roughly during CY 2010.

Obviously, that means we can't have "first do DNSSEC, and wait before doing anything else." We will have to deal with some chunk of "everything else" alongside the DNSSEC rollout. Fortunately, we learned during the root scaling study that the effects of (a) and (c) on the root system actors are not great enough - particularly relative to the effect of (b) - to create risks that the actors can't manage without major re-planning (beyond what they're already doing for (b)).

Page 16:

With respect to the statement "IANA tries to know each of the actual persons directly; the total number of these individuals currently known by IANA is between 400 and 500" we wonder whether the RSST Report is saying that personal relationships IANA relies on cannot scale. For example, using current figures, if the number of TLDs increase by a factor of 10, then IANA must personally know 4,000 to 5,000 individuals as well as keep track of human adds, drops, and changes. Yet it is unclear what dependencies are affected when TLD contacts can no longer be someone you know.

Page 17:

The RSST Report states, "The root server system has evolved beyond that point long ago," which refers to system evolution but make no mention whether the responsibilities have been enumerated and documented since 2000. We believe that the community has no visibility into the responsibilities (and accountability) of the root operators and some will argue that this is a scaling issue. At some point, root operators can decide their gift to all users of the Internet is too costly. For completeness, we think the RSST Report should have considered this an issue.

RSST Response: We talk about "institutional commitment" as a critical resource (with respect to ability to scale) for the root server operators in Section 4.1.4, but we don't take up the issue of "what would happen if the root server operators decided to quit?" explicitly in this study.

Page 18:

1. With respect to the statement "It sends a so called "priming query"" we thought that although this is the typical case it is not the definitive case for all resolvers.
2. B-root at University of Southern California Information Sciences Institute (ISI) is listed as being IPv6-enabled in Table 1, however, the current root zone does not contain an AAAA record for B.

Page 20:

1. The RSST Report states, "The role of root server operators... Each server is expected to have the capacity to respond to many times the rate of queries it receives and must increase its capacity at least as fast as the query rate increase." It is unclear what the RSST Report means by "many times" and where it sets this expectation.
2. The RSST Report states, "The role of root server operators... but some of the ideas that they describe were outdated already when the documents were finally published, and the root server system has evolved beyond that point long ago." If this is indeed the case, the RSST Report should have captured what it has evolved to, or what its current state is relative to terms aligned with those of RFC 2870.

Page 21:

The RSST report states, "As virtually anyone on the Internet can create a DNS resolver at any time, there is no way to precisely determine how many DNS resolvers are 'out there,' where they are, what software they are running, or other details of their configuration." The RSST Report should not make it sound like these things are impossible to study. They are not.

Page 22:

With respect to the statement, "growth is virtually open-ended" we think this insight, apparently still controversial, deserved more attention. There seems a cloudy debate over the need to find the appropriate number (N) to keep the world safe, but there seems less emphasis on the more likely scenario, which is no politically or economically feasible way to keep N finite, i.e. there will be no way for IANA to 'know each TLD manager personally' which seems to be perceived as a fundamental lever of trust in the system.

Page 24:

This is the third time the RSST Report includes the stabilizing properties of the root zone, but it is not clear that this is relevant in this case.

Page 27:

Instead of the statement "Anycast started as early as 2000" we think this should be "Root system use of anycast started as early as 2000" given that at least one commercial ISP (mica) was using it for DNS and other services (e.g. Network Time Protocol (NTP)) as early as late 1995.

Page 31:

We find the statement "allowing the size of the root zone to increase may frustrate future innovations" to be confusing. In particular, the root zone increases *now*, albeit slowly, and this growth has not frustrated innovations (to our knowledge).

RSST Response: This was not a point we (the study team) made; it's from a public comment (<http://forum.icann.org/lists/scaling/msg00001.html>).

Page 32:

1. With respect to the statement "The analytical (or qualitative) model of the root system is constructed from a set of six primary actors" we expected a discussion of these actors to follow.

RSST Response: It does—4.1.1 through 4.1.4. The actors on either end are covered in 2.4, but not in Section 4 (as they are not part of the analytical model).

2. In particular, we assume the six are: TLD operators, IANA, NTIA, VeriSign, Root operators, and resolver operators. If we are right, the RSST Report should have listed them and used "operators" throughout the discussion of actors and in the diagram (rather than registries and resolvers).

RSST Response: The only "operators" that appear in the model are the root server operators.

3. The model illustrates that TLD operators are the only place that change requests might originate. Yet it is unclear how requests for readdressing, new records, etc., are accommodated for the root servers themselves in this system.

Page 33:

With respect to the phrase "...will have to rebuild it in flight" this is not the only end case, and in fact, the less likely end case. We also point out that Operation, Administration, Maintenance, and Provisioning (OAM&P) antedates and is not exclusive to information technology. More importantly, the way the RSST Report presents this model/argument is convincing in the abstract but less so in practice. The RSST Report speculates that there is a risk that the rate will become steep without sufficient notification or that no/insufficient rate limiting will be present or enforced. However, some will argue that the TLD acceptance process and cost is itself a sufficient constraint. Others will say that rate limiting is inherently present in a system that involves multiple parties manually inspecting and approving changes so the end state the Report discusses is not significant. Thus, it would have been helpful if the RSST Report had included some historical data demonstrating how it calculated and maintained headroom, with an extrapolation of how these might apply to future scenarios.

Page 34:

1. The RSST Report states, "This number of additional contacts would place substantial additional demand on IANA's ability to scale the 'personal knowledge' aspect of the trust relationship." We believe this understates the problem. In particular, the current number of IANA (full time equivalents) FTEs with personal knowledge of authorized TLD contacts is on the order of three to four people. Increasing new TLDs tenfold or accepting server changes directly from DNS providers makes this model intractable. Even if many staff are added there will not be the same degree of shared knowledge of contacts or the ability to keep up with changes in contacts across so many organizations. This increases the risk of impersonation and the RSST Report should have recommended a different authentication or integrity check.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

2. With respect to “Restructuring” it is unclear to us why the RSST Report is discussing funding models. In particular, we don’t see them in the “Root System Model” diagram and we certainly wouldn’t consider them in scope of the operational root server system, yet we do see discussion of them scattered throughout the Report.

Page 35:

1. With respect to the statement "As for IANA, the demand for this resource increases linearly" we do not believe it is accurate to say that demand for manpower was or is a linear function in IANA or NTIA (and footnote 50 is not convincing on this point).

RSST Response: For the most part, the manpower resources allocated to the performance of root zone management tasks at IANA and NTIA are not full FTEs—that is, most of the people who perform these tasks do other things too. The quantum is some fraction of an FTE, not a whole FTE. Obviously at some magnification the graph of the function is not continuously linear, but the “steps” are much smaller than they would be if full FTEs were involved.

2. With respect to the statement "NTIA does not plan to change the way in which it exercises its oversight role" the RSST Report identified the rate limiting function, but based on this statement one could conclude that the root will only grow as fast as NTIA will approve its growth. Yet, NTIA manpower lies outside ICANN/root operator scope, oversight and budget.

RSST Response: Yes.

3. In Footnote 47 the RSST Report states, “The size of the root zone does not significantly affect the demand on this resource.” This footnote is a comment on the statement “... the demand for [IANA’s manpower] increases linearly with the frequency of arrival of change requests.” Yet elsewhere the RSST Report estimates the number of transactions at one per year per TLD.

Page 37:

1. In Section 4.1.2 the RSST Report states, “As far as this study was able to determine, NTIA does not plan to change the way it exercises its oversight role in order to accommodate an increase in the frequency of arrival of root zone change requests.” We think that this statement should not have been included in the RSST Report, as it presupposes things on behalf of NTIA that may not be accurate.
2. At the end of second paragraph the RSST Report states, “This number of additional contacts would place substantial additional demand on IANA’s ability to scale the ‘personal knowledge’ aspect of the trust relationship.” We note that the number of large DNS providers is only 20, which seems pretty modest. But even the total of 1,500 smaller providers is not too large, particularly if the IANA group is presumably ready to handle many, many more TLDs. However, we note that the facts also are wrong. A short time ago SSAC ran the numbers. At the time of the survey, there were 1,561 name server records associated with the TLDs. There were only 1,047 distinct name servers. Of these, 895 served exactly one TLD. That means the number of name servers shared by two or more TLDs was 152. Of these, 69 were shared by exactly two TLDs, 21 were shared by exactly three TLDs, and another 21 were shared by exactly 4 TLDs. This accounts for 111 of the 152, leaving 41 name servers shared by 5 or more TLDs. These are not large numbers, and,

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

as noted above, if IANA is presumably prepared to handle a very large number of TLDs, this additional load should not be of major concern.

Page 38:

1. We think what the RSST Report is saying here is that the engineering principles are "Build Capacity +10% and threshold/warn when utilization is Capacity +1%".

RSST Response: No. The "prudent network engineering rule of thumb" is to determine the load that you have to handle, and then ensure that the system you build can handle at least 10% more than that "maximum" load; and most operators run with a threshold warning set to trigger an alarm when load reaches 1% of capacity (not "capacity + 1%").

2. The RSST Report states, "During the root scaling study, many operators told the study team that they were prepared to do 'whatever it takes' to operate their root server." We appreciate the intended candor and commitment, but we imagine that the root operators intended for this to be taken in the context of money, manpower and machines, yet the phrase "whatever it takes" may imply adding root operators, oversight, etc. to some of the target audience.

RSST Response: The statement should be read as an expression of the general attitude of many root server operators, not a specific commitment or promise on the part of any root server operator or of root server operators as a group.

3. In the paragraph beginning "In order to achieve..." the RSST Report also observes "the prudent network engineering rule of thumb that critical systems should maintain at least a 10% overhead in capacity to absorb either attacks or unanticipated increases in legitimate demand." However, 10 percent seems like a less-than-prudent percentage overhead in capacity. In particular, it was not clear how well quantified this 10 percent value is across the root operators.

Page 39:

In Section 4.1.4 at the top of page the RSST Report seems to say root server operators operate at two-thirds of capacity ("50% overhead") or one-half of capacity ("100% overhead"), and that if the load increases by 10 percent, it triggers planning for an increase in capacity. Further, the discussion in this paragraph suggests it takes 90 to 180 days to actually install the increased capacity. However, it was not clear to us how much of an increase is included in this change. Moreover, it was difficult to tell how the picture presented in this paragraph matches with the information in the immediately previous paragraph on the prior page.

Page 40:

With respect to the statement "... they are in the same category with respect to the scope of our study as query/response load on the root name servers" it is not clear to us whether the RSST Report ruled these out of scope and notes that the load on the root name servers was supposed to be in scope for this study.

Page 41:

1. With respect to the Section labeled "Automated root zone management" we note that there are two distinct parts to the automation, and it would be helpful to treat them separately. One

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

part of the automation is the interaction between TLD operators and IANA. The other part is between IANA, NTIA and VeriSign.

2. With respect to the Section labeled “Initiating change requests” it was not clear what is the bottom line on how this process will scale and whether new TLDs be required to use automation.

Page 42:

With respect to the first paragraph, beginning “IANA will not act on a request (howsoever received) until they confirm it...” we question whether this is the controlling factor no matter how much automation exists.

Page 45:

We note that it would have been helpful if the RSST Report explained what constitutes a “stealth delegation” and how it is detected.

Page 54:

1. We take issue with the RSST Report’s apparent dual excuse that 1) certain information was not provided because the system was too complex and 2) that there was not sufficient time to produce certain information. In particular, we note that the RSST agreed to the ToR, but that the RSST Report appears to argue that it is impossible to complete the ToR while admitting there was not enough time to accomplish these requirements.
2. The RSST Report states, “the limitations of the model developed for this study, which was constructed from incomplete information over a relatively short period of time.” With respect to this statement we wonder what information is missing and why and note that a key requirement for the root scaling study is a solid description of the existing system and a clear path for incrementally improving the model.
3. In Section 5.3.1 the RSST Report states, “The often-heard answer was: ‘we can adapt to that, given the time; the details of how we would adapt would depend on the specific circumstances.’” We wonder whether the RSST Report is providing this as a primary message from the RSOs and a conclusion of the Study Team.

Page 55:

1. We note that an earlier comment from an RSO claiming “whatever it takes” seems to be contradicted by the cautionary statements made in the paragraphs in “Architectural and design choices.” Also, it is unclear how the RSST intended the audience to benefit from the discussion of RSO funding and motivations. In particular, we think that some of this material seems orthogonal to the purpose of the RSST Report, and could mislead the community. For example, we are concerned that statements such as “The RSOs are independent actors, and they guard that status fervently” may antagonize an already sensitive matter and distract from the main issue.
2. With respect to the statement “...these sometime include cost-sharing or even revenues beyond simple cost recovery” we do not understand what the RSST Report means by the phrase “cost-sharing” and suggests elaboration would have been useful.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

3. In the phrase "...do as the RSS grows by orders of magnitude..." we assume "RSS" here is root server system, but we do not think we have seen this term defined and it is not completely intuitive here.

RSST Response: Yes—this acronym should have been expanded here, where it first appears.

Page 57:

In the Section on "Complexity of combinations" we think the implication in this section is incorrect in as much as there are, indeed, interactions among the effects, but the interactions are actually fairly small and simple.

RSST Response: This is not what we observed or concluded.

Page 59:

The RSST Report states, "...but many detailed parameters are scattered across different parts of the program. This makes the implementation and evaluation of scenarios a non-trivial task." With respect to this statement we think it would be far better for the RSST Report to have presented a readable, consistent and clear model.

Page 61:

1. The RSST states, "The reader should be even more careful than usual when interpreting current results of the modeling and basing any conclusions on these results. Doing so would be premature." Yet, we note that the RSST Report does not present any results of the modeling.

RSST Response: Yes—unfortunately the TNO quantitative model, including documentation and simulation results, was not available to the Steering Group at the same time as the study team's report.

2. In Section 5.9 concerning discussion of the business day we wonder whether this discussion was important enough to include in the RSST Report. In addition, we think the discussion of representing the "... zone file size with '.1', an undefined number" is unclear.
3. The RSST Report states, "Work is progressing on the validation of the model. However, there is a problem that not all data are presently known or available in a form suitable to put in the model. We have just begun to..." Concerning this statement we wonder where are the data and why the RSST Report does not list what data are needed. In particular, we think the RSST should have determined where the model requires more precise data about how much time certain events take or within what range. Similarly, the RSST Report states, "For example, the time needed for validation checks was available only at the last moment and is consequently implemented in the model with an exponential distribution; such distributions may need to be refined." In this case we think the RSST Report should have included a plan for real validation. In both of these cases we note that it was unsettling to keep reading that the important work has yet to be done.
4. The RSST Report states, "By looking at the block where this number is input, it is possible to determine what this represents. It would have been better to have this defined differently, which we anticipate will be done in a future version of the model." We think this statement suggests that the RSST is does not find the current model to be sufficient.

Page 62:

1. The RSST Report states, "...by 4,480 TLDs, change output starts to slow down, and the whole simulation system becomes fully overloaded...with 8K the model becomes fully overloaded." With respect to this statement it is not clear whether the reference is to the simulation software, or a prediction of the software that some aspect of the system becomes overloaded. Moreover, we note that a model is a mental construct that cannot be overloaded by a computer so it is unclear what was overloaded and what it tells us about the system.
2. In Section 5.10 at end of the first paragraph the RSST Report states, "TNO has developed and run two validation scenarios, in which they have changed only two different variables: the number of TLDs and the size of the zone file..." We think it would have been useful to know how different are these two variables.
3. The RSST Report states, "In this case, it takes about 2 days for a change to become effective (that is, to be present across the root server instances." It is unclear what part of the process is being measured and we note that this figure is at variance from reports we have heard. Moreover, it is unclear whether the IANA, NTIA and VeriSign processes are included.

Page 71:

Much of the priming discussion was stated earlier. In particular, we note that the paragraph about message sizes, including the historical reference to 1996 and the explanation of "how we filled a UDP datagram" seems unnecessary in a "Findings" section.

Page 72:

It may have been worth mentioning in Section 6.8 that "famous marks" may also be early adopters of DNSSEC in section 6.8 as this would have underscored the practical matter of dealing with this scenario.

RSST Response: Yes—this is a good point.

Page 73:

In Section 6.9.1 the discussion of the publication headroom is unclear. Also, there is a recommendation that "[s]tagging has provided a valuable backstop in the past and should not be abandoned without good cause." We do not understand the intent of this recommendation. In particular, as shown in Figure 6, 11 of the 13 RSOs do *not* use staging so we wonder whether the RSST Report is recommending staging. Moreover, we suggest that whether or not to use staging is within the purview of VeriSign and the root server operators to decide. Finally, it is unclear where is the line between making this recommendation and making comparable recommendations for how to modify the structure or operation of other parts of the system.

Page 74:

The RSST Report states, "... the current distribution system could absorb a growth of 20K to 40K in the size of the root zone file without noticeable effect on propagation delay or jitter and without pushing any root server operator out of its 'normal operations' region." This suggests that there is only room to expand the root zone by 50 percent, which is enough to add about 140 new TLDs without IPv6. In fact, at the end of the paragraph the RSST Report suggests 60 to 120, which is close enough – or about 56 zones if IPv6 addresses are added for all TLDs. (IPv6

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

adds 25 percent.) However, this suggests that there is no possibility of accommodating DNSSEC, since that would cause the root zone to be four times as large. We wonder if this is the intended message and, if so, whether this contradicts the claim from all of the RSOs that they are ready to take the signed zone and serve it.

Page 75:

1. The RSST Report states, "This suggests that adding DNSSEC to the existing root zone would cause enough direct and indirect impact to the root system and the Internet at large that would preclude combinatorially adding any other feature until the operations have returned to normal." We think this should have been stated as either a finding or a recommendation. As written, there is strong potential for the target audience to choose the label that fits its interest, i.e., a staunch new TLD advocate will say "well doing DNSSEC first is only a *finding*..."
2. The RSST Report states, "The latest test results from VeriSign suggest that the delay will be higher than these extrapolations." Yet, we note that the RSST Report does not appear to provide the test results from VeriSign.

Page 77:

1. In the first paragraph the RSST Report argues that the distribution process can handle 1,000,000 records. This is inconsistent with the recommendation that an "O(100)" increase is okay but an "O(1000)" increase is not, and is further inconsistent with the discussion at the top of page 74.

<p>RSST Response: The "O(100)" and "O(1000)" findings refer to the root system as a whole, not just to the distribution components.</p>
--

2. In Section 6.13.3 the RSST Report states, "Early simulation results suggest that the human-inspection bottlenecks at IANA and NTIA break down when the root zone contains entries for between 3,000 and 8,000 TLDs, depending on other variables. The current requirements for multiple human intervention steps in the root zone management process therefore appear to limit the growth to O(100) entries." However, we note that the RSST Report does not include the details of the processes at NTIA and IANA, does not provide the simulation results, does not provide a validation, and that the conclusion of "O(100)" seems unrelated to a bottleneck of "3,000 to 8,000." (A simple rule of thumb from queuing theory: it is generally okay to run at 70 percent of capacity. That would suggest 2,000 TLDs would be okay. We are not suggesting this is a valid figure to use, but it's a better conclusion from a simulation result of "3,000 to 8,000" than "O(100).")

Page 78:

With respect to the phrase "...distribution master is hidden..." we wonder whether this should have reflected the plural, "multiple DMs", or the "DM cluster", or whatever it is called.

2.3 Small Errors

These are actual misstatements, usually specific and relatively small.

Page 20-21:

The RSST Report states, “The root server operators have no formal relationship to each other, to ICANN, or to NTIA (the exception is VeriSign...,” but we note that the L-root is operated by ICANN. The E, G and H-roots are operated by various parts of the U.S. Government. Thus, there may be interagency relationships within the U.S. Government regarding these operations.

Page 22:

In Section 2.5.1 the RSST Report states, “Whereas newer TLDs tend to be 3-6 character names.” We note that this accurately describes the handful of newer TLDs in the past, but is far less likely to be the case for IDN TLDs and perhaps ASCII gTLDs in the future. In any case, this plays the smallest role in any quantitative analysis of the size of the root.

Page 48:

In the last paragraph the RSST Report states, “Unfortunately, none of the answers to these queries is cache-able because the answer is ‘NXDOMAIN’ – ‘not in this domain.’” Our understanding is that negative answers do, in fact, get cached, and that a large fraction of the queries that get negative answers do get handled out of caches. However, the fraction is not as high as for positive answers.

Page 50:

The data that appears in Figure 6 seem to be new. In particular, it is interesting that only two operators, B and I, stage the updates. B runs a single instance, and I runs 34 instances. This means 158 machines ($191 - 35 = 156$ instances plus the two staging machines for B and I) are all drawing from the distribution master(s) at the same time. However, after subsequent discussion at the SSAC Retreat on 01 October, it was noted that the information in Figure 6 appears to be inaccurate as several SSAC members confirmed that more than two operators stage updates, depending on the definition of staging. Consequently, we think the dichotomy of staging versus not staging is not a precise enough description and does not convey a clear enough picture about what are the issues.

RSST Response: Two figures in the report are labeled “figure 6”—“Figure 6: Distribution Architecture” in Section 4.3.5, and “Figure 6—Root signing effect on root server operator headroom” in Section 6.10. This comment refers to the first of these.

Page 56:

The RSST Report states, “There are two excellent historical examples. The first, is the transition from unicast to anycast for root zone file distribution. The second, now underway, is the IANA transition from its largely manual style of work to the RZ-WAS automated web tool described in Section 4.2.3.” We note that the first was primarily a change to the services run by the RSOs and it entailed *no* changes to the distribution. In fact, it arguably made distribution worse. The second does not affect the root servers at all.

2.4 Irrelevant or Inappropriate Text

These are political comments or other judgments that are outside the scope and purpose of this report.

Page 12:

The paragraph "Entry of a top-level domain into the root zone file has become a subject of substantial economic and social importance. Consequently, who controls entry into the root, and by what means, have become controversial subjects" is interesting but seems out of place here.

RSST Response: We agree. Neither of the points made in these two sentences is germane to the study or the report.

Page 13-14:

In the first paragraph the RSST Report states, "There are 13 publicly accessible well-known IPv4 address on the Internet from which such service can be obtained." On the next page the RSST Report states, "Technical constraints make it difficult to increase the number of root name server IP addresses beyond 13." This is not entirely accurate. As written, and with no reference to the SAC AAAA reports and EDNS0, this suggests to the layperson that there is no alternative but to keep the number at 13. Also, if the RSST Report is going to enumerate the rationale and conflict surrounding root zone operators, we believe it is arguably unfair to characterize every incentive to add a root zone operator as a nation/regional desire to have its own root. Moreover, we note that for some nations and regions, it is less about having one's own root but more about putting the decision to create more root operators in the hands of a clear, discernable, publicly accountable body responsible for oversight and governance rather than the current system.

RSST Response: We shouldn't have raised the "political acrimony" point at all — it unnecessarily stirs up an old controversy without any offsetting benefit to the arguments being made.

Page 17:

1. The RSST Report states, "The IANA is currently operated by ICANN under a contract from the U.S. Department of Commerce..." implicitly suggests this may be likely to change. We note that while the statement is factually true, the emphasis is inappropriate unless balanced with comparable phrasing for the other parties, e.g. oversight of the root system is currently vested in the NTIA, editing and distribution of the root zone is currently done by VeriSign, and operation of the thirteen lettered root operations is currently done by VeriSign, USCI/ISI, Cogent, the University of Maryland, NASA, Internet Systems Consortium, etc.
2. In Section 2.4.4 the RSST Report states, "It is important to distinguish the separate roles played by VeriSign... or by any one organization." Although we find the first part of the paragraph, up to "different groups" to be informative, the remainder of the paragraph goes further and suggests these functions can or perhaps should be moved. While true, this is out of scope and unnecessarily provocative.

Page 18:

In Section 2.4.5 the RSST Report states, “The current root name server operators were not selected through a formal evaluation and qualification process, although they play a fundamental role in ensuring the availability and reliability of the root.” It is not clear what is the purpose of this sentence. For example, was the RSST Report trying to say the operation of the root is at risk because the operators do not meet any explicit criteria? If so, the RSST Report should have been more clear on this point. Alternatively, this sentence is out of place because it directs the reader to think about the political process – or lack of one – behind the vagaries of the selection process quite some time ago. This statement should have been deleted or rewritten.

Page 20:

1. The RSST Report states, “Historically, the operators of the root servers have not charged fees for resolving Internet address queries.” We wonder whether the RSST Report is suggesting the root server operators are unstable or reaching capacity because they do not have enough money and, if so, whether there is any evidence to support that assertion. This statement seems out of place.
2. In the next paragraph the RSST Report states, “Nevertheless, it is a valuable service, whose provision is a little-known and little-appreciated gift in kind to all users of the Internet.” This statement may or may not be true. We suspect there is actually broad awareness, understanding and appreciation among the technical and operations community. In any case, we don’t see the relevance for the RSST Report.

Page 41:

In the paragraph beginning “Operational changes...” we find the discussion about communications between NTIA and ICANN about changes in the process qualitatively different from a discussion how the system works and scales. This is not relevant and should not have been included.

Page 47:

The RSST Report states, “...and take no further action with the aborted change request.” We find this hard to believe in practice. On the contrary, surely it is queued for a later response from NTIA by IANA and resubmitted, and IANA is notified of the abort such that they can at least be prepared to deal with the change requestor.

Page 54:

1. The RSST Report states, "...the limitations of the model developed for this study, which was constructed from incomplete information over a relatively short period of time..." We think the RSST Report should have explained why the information was incomplete and whether additional time would have allowed the RSST to gather the necessary information.
2. Here and elsewhere the informal responses the RSST Report quotes of RSOs create the impression that the RSOs treated the RSST’s inquiries cavalierly, evasively, or dismissively. The alternative conclusion is that the RSST did not pursue the RSOs until it received a suitable answer, but instead conceded that the RSST would not get the information it needed. We suggest that this approach only reinforces the opinions of those who are suspicious over the way the root zone system is managed.

Page 55:

1. The RSST Report states, “However, in any event, the DNS does not support charging for queries made to root server instances.” By this statement the RSST appears to be responding to the question of how to provide revenue to root server operators or perhaps how to throttle the load on root servers. However, the RSST was not asked to provide answers to these questions so this statement seems inappropriate and out of scope.
2. In the next paragraph the RSST Report states, “Each root server operates within an organization that is willing to fund their root server for reasons, collateral and otherwise, and where the costs of doing so are a reasonably small fraction of overall organizational costs. Some of their motivations for running an RSO relate to marketing and stature, others consider a strategic value for their business, still others see it within their role as being members of the Internet technical community.” We think this argument is speculative and out of place.
3. At the end of the next paragraph the RSST Report states, “Providing external funding may not be as simple as it may sound. The RSOs are independent actors, and they guard that status fervently.” We find this statement to be inappropriate.

Page 57:

The RSST Report states, “Determining parameters... These parameters will include, for example, the capacity of internal distribution links used by root server operators.” It is unclear what is an “internal distribution link” and how it fits into the system with root server operators. Specifically, earlier text suggested that standard DNS techniques and Internet connections were used to update/distribute the root zone to the RSOs.

Page 61:

1. We find the text under “Limitations of the model” to be unclear.
2. We find the statement, “... which we anticipate will be done in a future version of the model” to be misleading and wonder whether there is support or commitment to “a future version of the model.”

Page 63:

A visual depiction of the outcomes would have been useful here. In particular, we think the RSST Report should have included some discussion of the resources that were input to the model (e.g. IANA staff). For example, the RSST Report describes the model and the outcomes, but it does not discuss any of the input parameters. We think it is appropriate to highlight inputs as well as outcomes here.

Pages 64-70:

Since the RSST Report indicates that DNSSEC takes precedence, should this not be discussed first, e.g. as subsection 6.1.1? In fact, it would have been helpful to discuss "adding support" in order of precedence. Moreover, the level of detail presented with respect to effects is heavily imbalanced towards describing the effects of adding DNSSEC. Some of this detail might have been better presented earlier in the report. We suggest that certain readers may feel that more attention was paid to DNSSEC than v6 and new TLDs simply due to the fact DNSSEC was discussed in more detail.

Page 66:

With respect to Table 3 our interpretation of the previous text would lead us to believe that at least another 12 squares should be populated with an ‘X’ – specifically Row 1, Column 5; Row 2, Column 4, 5 and 6; Row 3, Column 2,3,5 and 6; Row 4, Column 3 and 4; Row 5, Column 3 and 7. It would have been useful for the RSST Report to include an explanation of how these are *not* affected.

Page 67:

The RSST Report states, “If TCP transactions become more prevalent, the anycast architecture for root zone distribution may require changes.” We note that technically anycast is not used for “distribution” if one follows the terminology the RSST Report outlined at the beginning of S.4. Also, it would follow then that anycast only has implications on the “resolver” side of the model. In addition, it would have been helpful if the RSST Report had described what changes one might employ to get away from the constraint of non short-lived transactions besides moving away for geographically dispersed anycast nodes.

Page 69:

In Section 6.4 we were surprised the RSST Report did not mention the 222 AAAA records in the current root zone or how that represents approximately 20 percent parity with A records in the zone already. Also, one might extrapolate from this argument that in the future no IPv4/A record will be necessary.

RSST Response: Section 2.5.3: “IPv6 records are already present in the root zone, are being added at a steady pace, and currently represent 15% of all IP addresses in the root zone.”
Reiterating this in Section 6.4 would have been a good idea.

Pages 78-80:

Section 7, “Frequently Asked Questions” appears to be based on a suggestion by the Root Scaling Steering Group as a way to handle questions that came up during the Study that fell outside the scope and that might be pursued later. However, we think that the actual questions and answers the RSST Report includes in this section are not very useful and some are inappropriate. Questions 1-4 and 11 and 12 are informative and would have been better added as footnotes in the appropriate places. Questions 5-10 are political in nature and do not belong here. Questions 13 and 14 should have been included as part of the RSST Report introduction. Finally, question 15 assumes there is an active request to add 40,000 famous marks, but there is not. This issue would have been more appropriately addressed at the beginning of the RSST Report, in the same status as the signing of the root.

3. Comments on the TNO Report

The following are the SSAC comments on the TNO Report organized by page number and section, including a response from the RSST.

Page 6:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "The hierarchical modeling concept enables the creation of an initial model in which for example IANA is modeled as a black-box, which can be worked out in more detail during the modeling process without having to model the interaction between IANA and other systems in the provisioning process again." It is not clear what is the basis for this assertion.
2. "In line with this modeling approach, we chose to use the simulation model SW package ExtendSim (or in fact, Extend OR, version 6.0.8) to enable fast simulation SW development." The explanation for why the software was chosen is inadequate and it is unclear whether anyone else uses it for verification.
3. "If time permits, the dynamical behavior will be visualized by means of graphical animation." This statement does not seem to belong in the Report as it refers to activities that do not appear to have taken place.
4. "Further by performing sensitivity analysis on the input parameters." It is unclear how the sensitivity analysis was performed or what was learned from it.

Page 8:

General comment: The acronym "EPP" is used several times on this page without an explanation of what it means to what it refers.

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "At the same time it begins its technical check, meaning that this process is conducted in parallel with the DoC check." It is unclear what are these "technical checks" and it would have been helpful if the Report had included a table with a list of steps.
2. "For each arriving request the (configurable) Success Rate attribute is set to 1, indicating that the probability that the request content contains an error is zero. See appendix B for details about the error model." This statement seems to suggest that since no relevant parties collect any information on their own errors the assumption is that there are no errors. This seems to be an inadvisable approach to modeling anything, much less critical infrastructure.
3. "This is included in the model to simulate the case where IANA performs multiple subsequent authorization checks towards multiple stakeholders." An example would

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

have been useful in this case to clarify the meaning of this statement.

4. "However, we consider it likely that at least several authorization checks, executed by authorized persons, will remain in the process." It is unclear why this is considered unlikely and whether this statement was based on information gathered from people involved in the authorization checks.

Page 11:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. In the last paragraph on page 11 TNO Report statements seem to suggest that the Department of Commerce (DoC) and IANA try to avoid revealing any problems with a root zone request from VeriSign by having IANA, rather than the DoC, ask VeriSign to withdraw a request. These statements, while providing a helpful description of the process, seem to suggest that the process is less than transparent.
2. The TNO Report states, "In order to keep the number of input parameters limited we assumed a rather even distribution of processing time over the steps." It is not clear from the tables which data are actual and which are estimates. It would have been helpful if the Report had provided the source of the data.

Page 12:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "The holding time of the information is limited to 10 working days (336 hours), after which the request is discarded by VeriSign. However, in practice the situation where this 10 working day timer actually expires does not occur. Therefore, we chose not to include this timer in the model." This suggests a level of accuracy in the model that is not reflected in actual practice.
2. "If the results of the DoC authorization check are though such that the DoC does send a "YES" to VeriSign within the holding period, the DoC will notify IANA via PGP-encrypted email. VeriSign receives a PGP-encrypted email message from IANA requesting the EPP transaction to be aborted. In this case VeriSign does not preserve any state with respect to an expunged request. This means that if the request is re-submitted at a later point in time, VeriSign will consider it as a new request." This paragraph, as noted above, suggests that some aspects of the process are less than transparent.
3. The assumptions about root zone file size seem wrong in a number of ways. They appear to ignore the impacts of DNSSEC, IPv6, and new alphabets, all of which increase the size of the zone.
4. In footnote 7 the TNO Report states, "As far as we know this is the case and in the model the zone file production time is assumed to be independent from changes being updated in the DM database." It is unclear why this information was not verified.

Page 13:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "This formula reflects: a) the model assumption that the root zone file size will grow linear with the number of TLDs (at least with AXFR this will be the case; for iXFR this will be different);" It is not clear when or how when or how it will be different for iXFR.
2. "b) a multiplier factor for introducing DNSSEC (without DNSSEC the multiplier is 1; with DNSSEC the root zone file size is expected to triple) and;" However, since the TNO Report models this with a value of 1 this appears to be incomplete.
3. "c) one additional input parameter enabling representation of other influences on the root zone file size." It is unclear where this is reflected in the table or model and what values were used.
4. "Theoretically, the DoC approval process for each change request is limited to 10 working days." It is unclear why this is theoretical and whether this is an upper or lower limit.
5. "In the study we were not able to obtain specific information about the checks performed by DoC, or the time that those checks require." Since the DoC urged that there should be a study of security and stability before advancing with new gTLDs (http://www.ntia.doc.gov/comments/2008/ICANN_081218.pdf) it would have been helpful that if the RSST could not get specific information from the DoC it could have followed up with a formal request for that information.
6. "Further, it is estimated that the actual amount of work needed to verify the change request is around 2 hours." It is not clear what is the basis for this statement.

Page 14:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. The attempt to model root zone change requests coming in during evening and weekend hours seems to assume that the stakeholders are working on evenings and weekends. It is not clear whether this is consistent with what we know about when change requests currently happen or whether these data are available.
2. "However, if the results of the DoC authorization check are such that the DoC does not send a "YES" to VeriSign within the holding period, it will notify IANA via PGP-encrypted email, which in turn will notify VeriSign to abort the EPP transaction. As mentioned before this will rarely happen." It is unclear why the statement says that this will "rarely happen" and this appears to suggest that copies of the out-of-band PGP-encrypted email are not saved.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

3. “With respect to the likelihood of introducing errors into the content of a change request we assume that DoC will never make actual changes to the content of a change request. DoC will approve, or not. It will not modify. Therefore the error rate at the DoC process steps is assumed to be zero.” The statement that errors are “assumed to be zero” appears to contradict other statements that suggest that the process does not report errors in a transparent manner.
4. “An RSO responds with an acknowledgement to the DM. If the DM does not receive an acknowledgement within a certain period, a DM may send more than one notifies (RFC 1996, Section 3.6). Details are implementation-dependent but the interval between retransmissions, and the total number of retransmissions, should be operational parameters specifiable by the name server administrator, perhaps on a per-zone basis. Reasonable defaults are a multiple of the RTT (or timeout if using TCP), and a maximum of 5 attempts (for UDP). It is considered reasonable to use additive or exponential back off for the retry interval.” We would have assumed that there would be two opportunities each day to measure this interval, yet it doesn’t appear to have been measured. In addition, although the TNO Report appears to claim that the details of when to retransmit are “implementation-dependent” this appears to be a VeriSign parameter, not an RSO parameter, so it is not clear why the implementation would not be measurable. It also is not clear to which name-server administrator this statement refers and who is specifying the Round Trip Time (RTT) of the root zone on a “per-zone” basis.
5. “Once again, note that a “freshness” older than one week may cause an RSO to “go dark.” It is not clear whether this statement suggests that stale zone files are significantly different from unavailable zone files and what are the relationships between these files.

Page 16:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. “RSOs differ in the manner that they retrieve the root-zone file, e.g., in a staged or non-staged manner. Further, the RSO can be a cluster at a single location or it can have multiple locations using anycast.” It would have been helpful if these other factors could have been modeled or if the TNO had provided an explanation of why they were not.
2. “A query for a root-zone file entry is first handled by a load balancer which forwards a query to one of the name servers.” It is unclear whether this is always true, particularly with respect to anycast nodes, which are first handled by border gateway protocol (BGP), a load balancer.
3. “In the root-scalability study we do not consider the query-response side of the RSO, we focus on the publication side.” This appears to be a gap in the study despite the fact that it was not considered in the scope of the project. Moreover, the RSO modeling section seems incomplete. In particular, it is not clear how the TNO Report integrated the four separate models of 13 root name servers into the simulation.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

4. "In order to keep the simulation model insightful, we decided not to model each interaction between each DM and RSO reattempt separately, nor to model it on a packet level." It would have been helpful to understand what whether there were any negative affects from the decision to simplify the model.
5. "We do not have information concerning the DM to Name server transport protocol settings, nor do we precisely know where the specific bottlenecks are." It seems that this information could have been obtained at a meeting of the RSSAC or the OARC.
6. "Third, this enables the option to do pre-processing of meta-model input parameters, into the ExtendSim model parameters." This appears to make it easier to do proprietary simulation the results of which may be impossible for anyone to verify.
7. "Then the average transfer time is $\text{avg}(T) = s/R$, in the model the transfer time is randomly sampled from a statistical distribution with mean value $\text{avg}(T)$, e.g., distributed according to the Normal distribution." It would appear that this information would be easy to measure directly, that there would be two opportunities each day to measure it, and that someone should have been asked to perform this measurement.

Page 21:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "In the ExtendSim model the internal trigger is not modeled." It is not clear why the internal trigger was not modeled.
2. "In ExtendSim we do not incorporate the time-varying quality of the connection between the DMs and the RSOs." It is unclear why this was not incorporated.

Section 2.2:

This section appears to suggest that all bottlenecks are human, and will remain so, without supporting data. It would have been helpful to see the explicit request for data and the response.

Section 3:

General comment: The topic of this section is validation, but the section does not appear to include information on validation.

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. The TNO Report states, "The parameters are set according to their default values as described throughout section 2." It would have been helpful if the TNO Report had included a table of all parameters here, including the source of data to justify the choice of value.
2. Footnote 11 states, "The values shown in this table are those that are used in the simulation runs. In the input text file Input_ZoneFile.txt the Base Root zone file size is kept constant at 0.1 MBytes, respectively 3.0 MBytes." It is unclear to what the term "respectively" refers.

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

3. "When reading this input ExtendSim executes calculation (1), which results in the presented file size." It is unclear to what the "(1)" refers.
4. "For the modeling of the DNS notify and the SOA request/response, which are performed by UDP, we use the packet loss ratio and RTT." The source of these data is unclear.
5. "Therefore this is implemented in the model by a XFR success probability and a goodput during the transmission. For the numerical results we used the parameter settings for the good connection presented in Tables 7 and 8." It would have been helpful if the Report had provided the source of the parameter value for each cell in each table.
6. "The curves in Figure 7 match with the expectation, based on the applied model input parameters." It would appear that the curves in Figure 7 that are based on the model input parameters are the expectation.

Page 27:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "The production of the zone file is independent of the provisioning process," It would have been helpful to have a citation for this statement.
2. "Since we ran scenarios with good connectivity, the distribution typically takes far less than 1 hour. (see also right hand side of Figure 4)." The reference to the right hand side of Figure 4 is unclear.
3. "*The model results should not be used to draw conclusions from those scenarios.*" It would have been helpful if the TNO Report included a reference to the section and table with the results on which these conclusions are based.
4. "We emphasize once more that these cases and numerical results are only included for the purpose of illustration." It would have been more accurate to label this section "Illustration."
5. "Compared to the near-zero error rate in practice these error rates are too large. Apparently our assumption of 1% error rate per manual action is too large." This appears to contradict previous statements that claim a zero error rate.
6. "Recall that the study team was not able to obtain any numbers of error rates in the operational process. Therefore, this 1% value was purely an assumption." This does not appear to be an accurate statement since the information to which it refers is not stated previously in the text, but in the appendix.

Page 30:

On this page the TNO report makes the following statements. SSAC comments are included after each statement.

1. "Due to the fact that sound data regarding error causes and probabilities in the provisioning and distribution process are lacking the resulting output does not match the current near zero-failure practice. However, the error model is adaptable, such that obtaining valid error predictions merely requires configuration of input parameters with

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

error data from current practice.” This paragraph appears to imply correct error data should be obtained before using any results of this model.

2. “Further, the modeling exercise itself has shown that there remain questions regarding several details of its current operation and that it was quite hard to find quantitative data as input for the model. For example, there was little information obtained to model the request handling by DoC. Also information and quantitative data concerning interaction between IANA, DoC and VeriSign, especially in exception cases (e.g. an authorization check that turns out to be not OK) is hardly available. On the root zone file publication side the key, anticipated risk is to load a large zone file over a ‘bad link,’ but there is no clear characterization of a ‘bad link.’” It is not clear how to address these issues before making disruptive changes to the roots, or how often the “bad link” transfer has occurred and what was the impact. It would seem that if errors are not accurately reported this type of modeling would not be possible.

References:

It is unclear why the TNO Report did not include DNS-related references, except for a link to the RSST Report. It also is not clear how the TNO Report used the data on Savannah River Side Nonreactor Facilities (<http://www.osti.gov/bridge/servlets/purl/10102668-ndfaFm/webviewable/10102668.pdf>) and the personality/mood impacts on quality control tasks. It also is unclear whether the Report authors have expertise in DNS or the criteria by which they were selected.

Appendix:

The following statements appear in the Appendix. The SSAC comments are included after each statement.

1. “Then, twice a day the zone file will be produced. Potentially this may also introduce errors in the zone file, such as missing records or an empty zone file.” However, according to retained records it appears that this has never occurred.
2. “Secondly, an increase in the number of TLDs may lead to higher name server sharing factor. Further, it is assumed that an increase in the number of change requests leads to higher workload for human resources. We can easily assume the opposite in both of these cases, as automation is even more assured than new TLDs.” It would have been helpful for the Report to consider alternative scenarios.
3. The Appendix includes the following paragraphs. The SSAC comments are included following the three paragraphs.
 - a. “These failure causes and their relation with the model flowcharts and input parameters lead to a so-called “reward model,” that is an integral part of the root scalability model. The way it works is as follows.
 - b. “The purpose of this (human) error model is to estimate the quantitative relationship between the model input parameters and the resulting increase or decrease of error probabilities. In particular, it is not aimed at assessing the

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

(business) impact of errors, nor on detection or recovery actions for errors that get published in zone files.

- c. “Further, we will focus on the relative relationship of the model input parameters and the error rates; not on the prediction of absolute error rate values. This approach makes the error model is less sensitive to the availability and accuracy of data concerning errors in the current provisioning and publication process.”

These three paragraphs seem to imply that the TNO Report compensated for not having adequate data, but in the main text, the Report appears to emphasize that the results cannot be safely interpreted without additional data. This apparent contradiction is problematic.

RSST Response to SSAC Comments on the TNO Report: The errors as cited in the TNO Report appear to fall into two categories: 1) data-entry or transcription errors and 2) propagation or transmission errors. Only the first has the ability to disrupt the whole system. These have been so rare that there have been less than one per year, as reported by IANA. The second category of errors of significant magnitude to disable an instance of an authoritative root are almost as rare, with only a half dozen reported cases where a node was shutdown due to infrastructure instability. We think you suggest the correct way forward: to construct a framework of test and measurement points across the entire root zone system in order to gain a better understanding of the impact of change. It is true that DNSSEC is the most disruptive change to the system. We hope we can focus our energies on taking the root system architecture that has been defined and use that to help us construct an instrumentation framework around it.

4. Acknowledgments, Statements of Interests, and Objections and Withdrawals

In the interest of greater transparency, we have added these sections to our documents to provide the reader information on three aspects of our process. The Acknowledgments section lists the members who contributed to this particular document. The Biographies and Statements of Interest section points to the biographies of the Committee members and any conflicts of interest, real, apparent or potential, that may bear on the material in this document. The Objections and Withdrawals section provides a place for individual members to disagree with the content of this document or the process for preparing it.

Acknowledgments

The committee wishes to thank the following SSAC members and invited guests and members of the Root Scaling Study Team for their time, contributions, and review in producing this Comment.

Harald Alvestrand
Jaap Akkerhuis

SSAC Comment on the Root Scaling Study Team Report and the TNO Report

Lyman Chapin
KC Claffy
Steve Crocker
Patrik Fältström
Jim Gavin
Jeremy Hitchcock
Glenn Kowack
Warren Kumari
Matt Larson
Lars-Johan Liman
Bill Manning
Danny McPherson
Ram Mohan
Russ Mundy
Dave Piscitello
Barbara Roseman
John Schnizlein
Bruce Tonkin
Suzanne Woolf

Statements of Interest

SSAC member biographical information and Statements of Interest are available at:
<http://www.icann.org/en/committees/security/biographies.htm>.

Objections and Withdrawals

There are no objections or withdrawals.