The Internet Corporation for Assigned Names and Numbers

Summary of the Impact of Root Zone Scaling

Date of Publication: October 2010

Executive Summary

In February 2009, the ICANN Board requested a study be undertaken to examine the impact of the inclusion of a number of new technologies and the potential addition of significant numbers of new top-level domains to the root of the DNS. While some of these technologies had, by that time, already seen some deployment, some concerns were raised in the community that the stability of the DNS might be at risk if changes and additions were pursued without caution. As a result of the ICANN Board request, two studies were performed, one focusing on the impact of the new technologies and TLD additions on one root server, the other taking a wider view and looking at all processes associated with the management of the root system.

The new technologies of interest included IPv6 (both in terms of IPv6 addresses being associated with top-level domains and root servers as well as supporting IPv6 queries sent to the root servers), Internationalized Domain Names (IDNs), and security enhancements for the DNS (DNSSEC). However, since (and even in some cases, prior to) the ICANN Board resolution, all of these technologies have been deployed or implemented at the root, thus some empirical evidence exists which can be used in understanding the impact of these technologies.

To date, the deployment of IPv6, DNSSEC, and IDNs to the root system has had no significant harmful impact. While the deployment of these new technologies may have caused some minor degradation of service due to the lack of robust IPv6 infrastructure and/or the larger response size (due to the addition of IPv6 records or the DNSSSEC-signing of the root) causing that response to be dropped resulting in timeouts and retransmissions, no impacts were significant enough to have raised any concern among relevant communities.

Looking forward, with the assumption that estimates relating to a cap of less than 1000 new gTLDs per year being added to the root zone are accurate and assuming other parameters relating to the management of the DNS root are not altered substantively, it seems probable that normal operational upgrade cycles and resource allocations will be sufficient to ensure that scaling of the root, both in terms of new technologies as well as new content, will have no significant impact on the stability of the root system.

However, with the understanding that the management of the root of the DNS involves multiple parties and in the interest of the highest levels of care with respect to the stability of the root of the DNS, monitoring of root management system should be improved, particularly in the areas most sensitive to changes in

rate of growth or which require significant lead-time in which to change. In addition, clearer and more frequent communication between relevant root management partners and other stakeholders, including formal communications between ICANN staff and root server operators regarding projected numbers of approved applications, additional technologies that need to be deployed and in what timeframes, etc. would likely improve the confidence that changes to the root system won't negatively affect the stability of that system.

Introduction

Between 2004 and 2010, the root of the DNS has been undergoing significant change, both in terms of content as well as its support infrastructure. From the addition of Internationalized Domain Names (IDNs) in the root to the deployment of IPv6 and DNSSEC, it is safe to say that more change has occurred in the last 5 or 6 years than has occurred since the DNS was first deployed. With the imminent acceptance of applications for new generic Top-Level Domains (gTLDs), further substantive changes in the root of the DNS can be expected.

In keeping with ICANN's mission "to ensure the stable and secure operation of the Internet's unique identifier systems"1 ICANN's Board requested a study to be performed jointly by ICANN's Root Server System Advisory Committee (RSSAC) and ICANN's Security and Stability Advisory Committee (SSAC) with support by senior ICANN staff to investigate the impact of the proposed modifications to the DNS root system. However, both prior to and during the implementation of this study, many of the changes in the root system of interest to the Board were already implemented with no observable negative consequences.

This paper provides a summarization of the changes that have occurred to the DNS root and provides an analysis of those changes along with estimates as to the projected impact of future changes including the addition of new top-level domains.

Background

On 3 February 2009, the ICANN Board unanimously resolved in resolution 2009-02-03-04[2] that a joint RSSAC and SSAC study be conducted to analyze "*the impact to security*

[1] From "Article 1, Section 1. Mission" of ICANN's By Laws, see http://www.icann.org/en/general/bylaws.htm

[2] See http://www.icann.org/en/minutes/prelim-report-03feb09.htm

and stability within the DNS root server system of [the IPv6, IDN TLDs, DNSSEC, and new gTLDs] proposed implementations." The resolution stated that the joint study should:

- "*[A]ddress the implications of initial implementation of these changes occurring during a compressed time period."*
- "*[A]ddress the capacity and scaling of the root server system to address a stressing range of technical challenges and operational demands that might emerge as part of the implementation of proposed changes."*
- "*[D]evelop a terms of reference for the Study and appoint a steering committee to guide the effort by 28 February 2009."*
- "*[I]nvolve direct participation by senior ICANN technical staff involved with its planned implementations of these activities and to provide necessary support to implement aspects of this study under terms and with ultimate approval of the advisory committees."*
- Ensure *"the process for establishing the study terms, design and implementation will address the technical and operational concerns regarding expanding the DNS root zone that have been expressed on this topic."*
- Provide to the ICANN Board *"study findings and recommendations by 15 May 2009."*

As a result of this resolution, two efforts were undertaken, a study focused on the impact of scaling the root on one root server (the "L" root server operated by ICANN) and a more general study that aimed to model the processes in the root management system and analyze the results of scaling the system. An ad hoc study team known as the "Root Server Scaling Team" (RSST) was established comprised of members of RSSAC, SSAC, and outside experts to perform this second study.

The "L" Root Study

The "L" Root Study performed by the Domain Name System Operations and Research Center (DNS-OARC) under contract to ICANN focused specifically on the impact of different combinations of adding IPv6, DNSSEC, and new TLDs to a laboratory simulation of the "L" Root Server. The final report of this study, entitled "Root Zone Augmentation and Impact Analysis" was published on 17 September 2009 and is available at http://www.icann.org/en/topics/ssr/root-zone-augementation-analysis-17sep09-en.pdf.

The RSST Study

The RSST Study, which used the "L" Root Study as part of its input, outsourced the development of a simulation of root management processes, and conducted interviews with root server operators, IANA staff, VeriSign, NTIA, and others, was far more general, aiming to look at not only the impact on the root servers, but also on the provisioning systems that lead up to the root zone being propagated to the root servers. The final

report of this study, entitled “Scaling the Root” with a sub-title of “Report on the Impact on the DNS Root System of Increasing the Size and Volatility of the Root Zone” was published on 31 Aug 2009 and is available at http://www.icann.org/en/committees/dns-root/root-scaling-study-report-31aug09-en.pdf.

Root Scaling Events

Prior to and since the ICANN Board requested SSAC, RSSAC, and senior ICANN staff to undertake the study of the implications of scaling the root, many of the subjects of that study have already been implemented. The timeline associated with the introduction of new technologies to the root is provided in *Table 1*.

Date	Technology	Event
July 2004	IPv6	First IPv6 addresses added to the root zone for top-level domains (KR and JP).
November 2005	DNSSEC	First top-level domain (.SE) signed.
June 2007	DNSSEC	IANA DNSSEC-signed root test bed made available.
August 2007	IDNs	Test IDN top-level domains added to the root.
February 2008	IPv6, gTLDs	First IPv6 addresses added for root servers (A, F, J, K, L, and M). A limit of a maximum of less than 1000 new gTLDs per year is derived from estimates of gTLD processing times.
January 2010	DNSSEC	Deliberately Unvalidatable Root Zone (DURZ) published on first root server ("L").
May 2010	IDNs, DNSSEC	First production IDNs added to the root (for Egypt, Saudi Arabia, and United Arab Emirates). DURZ deployed on all 13 root servers.
June 2010	DNSSEC	First DS records are published in the root zone (for .UK and .BR).
July 2010	DNSSEC	Root is DNSSEC-signed and the root trust anchor is published.

Table 1 - Root Scaling Events

Impacts

During the period from July 2004 when the first IPv6 addresses were added to the root zone for TLD name servers until the root was DNSSEC-signed and DS records were inserted into the root in July 2010, root DNS service has continued with no reported or publicly visible degradation of service related to these events. This section examines the impact of each of the various changes to the DNS root.

IPv6

The inclusion of IPv6 in the root of the DNS has two components: adding IPv6 "glue" records[3] in the root zone for the authoritative name servers of TLDs and adding IPv6 "glue" records to the root servers. Each of these impacts will be examined in turn.

[3] Glue records are IPv4 ("A") and IPv6 ("AAAA") resource records associated with name servers that are in the zone being looked up. See RFC 1034 (http://www.ietf.org/rfc/rfc1034.txt) for the definition of glue records.

Top-Level Domains

In July 2004, the .JP and .KR domains were the first TLDs to have IPv6 "glue" records added. As of 6 September 2010, there are 283 IPv6 "glue" records in the root zone covering 203 TLDs. One impact of the increased use of IPv6 "glue" records has been an increase in the number of resolutions using IPv6 transport. As of 6 September 2010, at least one root server (the "L" Root Server) is seeing approximately 1.3% of DNS queries over IPv6[4]. Due to the less robust IPv6 network infrastructure within the Internet today, IPv6 queries and/or responses may be lost more frequently than with IPv4, resulting in more timeouts and retransmissions that would have occurred without IPv6 support in the TLDs. However, this impact has minimal negative consequences and is expected to improve as IPv6 deployment moves forward.

Root Servers

When some of the root server operators added IPv6 addresses for their root name server records, the size of the "priming query" increased significantly. As discussed in report produced jointly by RSSAC and SSAC labeled SAC018 and entitled "Accommodating IP Version 6 Address Resource Records for the Root of the Domain Name System"[5], there were concerns due to the fact that the priming response was anticipated to grow to more than the "classic" DNS maximal non-truncated response of 512 bytes. If the resolver requesting the priming response did not provide a larger response buffer size via the EDNS0[6] extension, it was feared the root servers might indicate a truncated response causing the requesting resolver to retransmit the request over TCP. Since TCP-based DNS queries are significantly more resource-intensive than the normal UDP-based queries, there was some concern that the root servers could be overloaded resulting in degradation of service to all users that queried the root servers. In addition, there was some concern that the larger response from the root servers would be blocked or filtered by firewalls, NATs, and other "middlebox" devices that "knew" (incorrectly) that a DNS response could never be more than 512 bytes. In such cases, there was a risk that the requestors might never receive a response and thus be unable to obtain the addresses of the root servers.

After significant study and testing of this issue, IPv6 addresses were added to the root zone in February 2008. In practice, the DNS server implementations running on the root discarded non-essential ("Additional Section") information in preference to truncating responses to queries that did not specify a sufficiently large buffer via EDNS0 (or did not use EDNS0). This may have resulted in a slight uptick in the number of queries sent to the root servers as resolvers were required to issue additional queries for data that had

[4] Private communication with the operators of the "L" root server. Other root servers should see a similar percentage of queries.

[5] See http://www.icann.org/en/committees/security/sac018.pdf

[6] EDNS0 is defined in RFC 2671 (see http://www.ietf.org/rfc/rfc2671.txt).

previously been supplied in the Additional Section, however if so, the increase was not noticeable.

For those requestors that supplied a larger buffer size via the EDNS0 extension, there may have been an increase in the number of fragmented packets which could have resulted in dropped responses either due to the loss of a fragment or because middleboxes were configured to discard fragments. In addition, some security policies have suggested (erroneously) that TCP-based DNS should be blocked. In such cases, a priming query without the EDNS0 option (or in which the offered buffer was less than the size of the response) could result in an answer that was blocked. However, in the more than two and a half years since the first IPv6 "glue" records for the root servers were installed into the root, there have been no significant (if any) reports of negative consequences.

Looking at the processing side of the root management system, ICANN root management processes and system as well as VeriSign processes and systems required some modification to deal with the IPv6 "AAAA" resource records and to verify IPv6 reachability in "technical checks" performed by both parties. The impacts to both ICANN and VeriSign were minimal however and these processes and systems continue to operate today without incident.

Internationalized Domain Names (IDNs)

From the perspective of the DNS, aside from a slightly longer average label length, Internationalized Domain Names are essentially indistinguishable from any other domain name. The addition of IDNs to the root was thus no different to the DNS than adding any other non-IDN TLD to the root. As such, no impact at the DNS level was observed.

There was, however, some impact in ICANN root management processes and systems. In order to usefully display IDN information, IANA staff needed to revise processes to request U-labels in addition to A-labels and had to modify IANA systems such as the Whois server to support to display both A-labels and U-labels. More generally, the support of IDNs in backend systems, particularly in the display of registrant data, continues to be a topic of ongoing discussion in ICANN (and other, e.g., security-related) forums. It can be anticipated that the proper display of IDN information will be a non-trivial impact across (at least) registrars in the future.

DNSSEC

The addition of DNSSEC to the root had significant impact, both in terms of the size of the root zone, size of responses to root queries, as well as the implications deploying DNSSEC has had to ICANN, VeriSign, and NTIA, the parties involved in root zone management. In terms of root zone size, as of 6 September 2010, the signed root zone

(as transmitted over the wire in a full zone transfer) was 222,246 bytes. When all DNSSEC-related records, namely DNSKEY, NSEC, DS, and RRSIG resource records, were stripped from that zone, the resulting zone size was 122,657 bytes. However, based on data from the "L" Root Study, it was anticipated that the additional data load on any reasonably configured name server imposed by DNSSEC would be inconsequential and in practice, this was borne out: there were no reports of any difficulties experienced by any of the root server operators loading and serving the DNSSEC-signed zone during the deployment of the "Deliberately Unvalidatable Root Zone (DURZ)", the staged deployment of DNSSEC in the root prior to publishing the root trust anchor.

Potentially more significantly, the size of the majority of responses from the root servers grew by a non-trivial amount, e.g., a query for the root name servers went from 492 bytes to 829 bytes when a DNSSEC-signed response was requested. As opposed to zone data size, a doubling of the size of a DNS response was of concern due to the 512-byte limit discussed previously in the context of IPv6. The DNSSEC specifications addressed this limit by requiring the use of EDNS0 to signal the resolver was equipped to handle responses that included DNSSEC-related resource records. However, as it turns out, most resolvers on the Internet, at least those querying the root servers, by default use EDNS0 and set a bit in DNS queries (the "DNSSEC OK" bit) to indicate the resolver understands responses that include DNSSEC-related resource records (regardless of whether or not the resolver will make use of those resource records). As a result, between 50% and 80% of the queries hitting the root server prior to the root being signed had the "DNSSEC OK" bit set and thus, when the signed root was served from all the root servers, those servers immediately started returning an aggregate of at least 50,000 DNSSEC-related resource records per second[7].

Prior to the root being signed, significant concerns existed regarding the impact of the larger DNSSEC-signed responses being returned to clients who may not be expecting them. In particular, there were concerns that middleboxes would, like in the case of IPv6 mentioned earlier, discard responses larger than 512 bytes. As a result, ICANN, VeriSign, and NTIA agreed upon a phased deployment of the signed root zone (the "DURZ") that also included substantial instrumentation of root servers to observe any change in query patterns. However, after deploying the signed root zone to all 13 root servers over the course of 6 months, no reports of negative consequences were received by any of the parties involved in signing the root.

In terms of process changes, deployment of DNSSEC at the root resulting in the creation of elaborate new processes along with new physical facilities that are necessary to securely manage the root key-signing key by ICANN and the root zone-signing key by VeriSign. New processes were also established to allow TLD administrators to securely provide "delegation signer" (DS) information to ICANN (and to allow ICANN to submit DS

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

information to VeriSign for inclusion in the root zone) to enable the creation of a "chain of trust" from the root to signed child zones. To date, these new processes have operated without incident.

Summary

Summarizing the impacts to date of the addition of IPv6 to the root system, IDN top-level domains, and the deployment of DNSSEC, no significant harmful effects have either been observed by or reported to ICANN.

However, with that said, one point that has been raised in the context of discussions regarding root scaling is the need for improved communications between stakeholders involved in the management of the root system. In some cases, the introduction of new technologies could likely have been improved with more formal communication of requirements from all parties that may have been impacted, discussion of those requirements and impacts, documented plans with timelines, etc. The communications, documentation, and discussions surrounding the deployment of the signed root have been suggested as an example of movement in the right direction in this regard.

Projections

The root system continues to undergo changes, albeit now more in terms of continued deployments of existing technologies than in structural changes such as the introduction of new technologies. This section examines some projections of likely changes, making the assumption that parameters such as zone refresh times, DNS record Time-To-Live (TTL) values, rates of root zone changes, and the length and complexity of administrative processes do not vary wildly or unexpectedly from historical values.

IPv6

It is highly likely that in the future, additional top-level domains will add IPv6 address records for their name servers. As of 6 September 2010, the root zone contains 283 IPv6 "glue" records corresponding to 203 out of 294 top-level domains having at least one IPv6 address record for their name servers. As IPv6 becomes more fully deployed, it is safe to assume more TLDs will be adding IPv6 support, eventually to cover all TLDs, and that the average number of IPv6-supporting name servers for those TLDs will go up. Until the Internet's IPv6 infrastructure improves to be on par with the IPv4 infrastructure, end users may experience some negative consequences in the form of delays resulting from queries sent to IPv6 name servers timing out.

In the case of the root, SAC018 documents that the size of the priming query response when all root servers have deployed IPv6 should be 811 bytes. While the root server operators that have not yet deployed IPv6 have not provided dates when they plan on enabling IPv6 on their root servers, they have all indicated they do intend to do so[8].

[8] Private communications with the co-chair of RSSAC and with the operator of the "L" root server.

However, since larger than 512 byte responses have already been encountered, the additional 100+ bytes in a priming query response is unlikely to have noticeable impact.

DNSSEC

As of 15 July 2010, the root zone has been signed and is being distributed to all instances of all 13 root servers. As such, further impact to the root zone from DNSSEC is likely to be limited to the addition, modification, and deletion of Delegation Signer (DS) resource records, the potential for changes in key algorithms, key lengths, or number of keys, and key rollover events.

Since DS resource records can vary in size based on the hashing algorithm used, the exact increase in size the addition of DS records will have in the future is difficult to accurately predict. However, given the structure of DS resource records, it can be argued that a pessimistic estimate of DS record size would be 64 bytes. As of 6 September 2010, there are 49 DS records for 29 TLDs (including the 11 test IDN TLDs still in the root). Assuming, as the ”L” Root study does, that full deployment of DS records by TLDs will result in a total of 1440 DS RRs for 1000 zones, the total number of bytes DS records will add would be less than 100 Kbytes. The actual number will likely be significantly less as it is tied to the number of TLDs and, as discussed in the subsequent section, this number is expected to be significantly less than the 1000 new TLDs assumed in the “L” Root study.

With regards to changes in key algorithms, key lengths, and number of keys, it is possible that the most significant change will be to move to Elliptic Curve Cryptography, which will result in significantly smaller keys at the same cryptographic strength.

Finally, while it is more of an operational issue than a root scaling issue, key rollover events occur with some regularity with all DNSSEC-signed zones. In the normal course of events, key rollovers of key-signing keys will require updated DS records to be provided to the parent zone administrator. In the case of the root zone, rolling the root key-signing key will require updating the root trust anchor in all resolvers configured for validation. It is hoped that RFC 5011-based mechanisms will enable much of the root key-signing key rollover to be automated, but it can be anticipated that some disruption will occur when the root key-signing key is changed and thus, rolling the root key-signing key should be done with some care.

Top-Level Domains

In the analysis done in the draft document “Delegation Rate Scenarios for new gTLDs”[9], ICANN staff estimates that the expected rate of new TLDs entering the root will be on the order of 200 to 300, even with higher than anticipated application rates. The same paper infers that regardless of the number of applications, there will be a process-

[9] See http://www.icann.org/en/topics/new-gtlds/anticipated-delegation-rate-model-25feb10-en.pdf

imposed limit in the addition of new TLDs of less than a maximum of 1000 new gTLDs per year[10]. For the purposes of this analysis, a fixed number of 1000 per year additional new TLDs will be assumed.

Based on work done in the "L" Root study, the anticipated size of the DNSSEC-signed root zone with IPv6 and full DS deployment and with 1000 new top-level domains is 624,791 bytes. Based on input received from root server operators, it is unlikely this amount of zone data will stress any of the root servers. In addition, this root zone must be distributed to each instance of all 13 root servers. For the sake of this analysis, making the assumption that the effective minimum bandwidth (taking into consideration line noise, interrupted communications, etc.) to the worst connected instance of all root servers is 300 bits per second, it would take approximately 4 and a half hours to transfer the entire zone, well within the current 12-hour root zone regeneration period[11].

Looking forward 10 years, and still assuming a maximum of 1000 new TLDs per year, the "L" Root study projects the root zone will have grown to 7,471,784 bytes. Again, based on input from root server operators, it is unlikely this amount of zone data will stress any of the root servers. With regards to bandwidth, the minimum bandwidth necessary to transfer the zone of this size in the 12-hour window would be approximately 1400 bits per second.

Another potential future impact of the addition of new TLDs is related to root query "splay". That is, the dispersion of queries across an increased number of TLDs may have some impact on the operation of individual caching servers. While it is not certain that an increased number of TLDs will result in an increased number of queries or that query patterns will change drastically, taken to an extreme, if a resolver sends a query to each TLD in the root, the cache of that resolver will end up holding the NS records for each TLD (along with IPv4 and IPv6 "glue" records and DNSSEC-related records if they exist) for the duration of the Time-To-Live (TTL) of those records. Compared to the limited number of TLDs today, this would increase the amount of memory consumed by the caching name server and, depending on caching name server's memory management techniques, could increase the likelihood that the caching name server could run out of memory. However, caching name servers already must cope with these sorts of memory management challenges since there are already sufficient domain names that can be queried (at all levels) to overflow pretty much any memory configuration if queries are asked quickly enough (that is, within the TTLs of the records such that more new records are added than records are expired). As such, the impact associated with a higher degree of "splay" within the root zone is not expected to result in significant impact on caching servers.

[10] 924 new TLDs per year to be specific.

[11] 300 bits per second is, of course, an unrealistically low number, however a more realistic number would allow for the zone to be transferred more quickly thus the use of 300 bits per second could be considered a worst case.

As discussed in the RSST report, the addition of new top-level domains will likely have impacts related to processes and back end systems in use by ICANN (in performance of the IANA function), VeriSign, and NTIA. For example, the quantities of data maintained in the database used to maintain contact information for TLD administrators is likely to increase significantly and the processes used to vet requests at each of the organizations involved in root management will likely need to change to cope with the increased load associated with day-to-day root zone modifications. However, all of the organizations involved in root management have indicated that they will adjust their resources to meet demand. The primary consideration thus becomes detecting the increased loads prior to them becoming an issue and to facilitate the adjustment of resources. As such, monitoring of root management systems at points in those system where bottlenecks may arise as well as defining thresholds that signal areas of concern is an area in which additional efforts are required.

Summary

Predicting the future is known to be somewhat challenging, however in the case of projecting the impact of scaling the root, it seems likely that if we assume historical patterns don't change in unanticipated ways, anticipated growth is well within the capacity of the system to adjust to that growth.

In the case of IPv6, nearly 70% of top-level domains have already deployed IPv6 as has 8 of the 13 root servers. It is unlikely that moving to 100% of both of these will have any negative consequences (modulo possible delays to end users resulting from timeouts due to the IPv6 infrastructure not yet being on par with the IPv4 infrastructure).

With DNSSEC, while there will be additions of new DS records as more TLDs sign their zones, it is unlikely this will cause any noticeable change in the root other than the root zone getting larger at a rate that will be (at most) tied to the number of new TLDs.

Finally, the addition of new TLDs has the potential for the greatest impact, however given the projected limit of less than 1000 new TLDs per year; it is unlikely the impact of this growth will cause any disruption as long as systems and processes are adjusted as part of normal operational upgrades.

Conclusion

As the DNS continues to grow and evolve to meet new requirements, ensuring that those changes do not negatively impact the stability of the DNS is of critical importance. As a result of ICANN Board resolution 2009-02-03-04, two studies were undertaken to analyze the impact of the addition of IPv6, DNSSEC, IDNs, and new gTLDs to the root of the DNS. In the "L" Root study, it was shown that at least one root server could easily handle both the deployment of the new technologies as well as several orders of magnitude more new TLDs than are anticipated to even be possible to be processed by ICANN for the foreseeable future. The RSST study suggested that absolute numbers weren't particularly relevant, rather it was the rate of change and how various root

management processes and back end systems are modified to deal with the changes that is important.

However, in the time between when resolution 2009-02-03-04 was issued and today, deployment of new technologies has continued, thus empirical data can be used to validate the observations of both studies. Deployment of IPv6 in the root, which began in 2004, has caused no significant harmful effects. Insertion of IDNs into the root in 2007 similarly was a non-event from the perspective of stability of the DNS, and deployment of DNSSEC in the root starting in January 2010 resulted in no observable or reported negative consequences.

Looking forward, further additions of IPv6, DNSSEC, and IDNs are unlikely to have any negative impact on the stability of the DNS, albeit the roll of the root key-signing key will need to be managed carefully to ensure validating resolvers have the new root trust anchor configured before the old trust anchor becomes invalid. The only remaining wildcard is related to the number of new TLDs inserted into the root.

One clear observation from the studies performed in response to ICANN Board Resolution 2009-02-03-04 and discussions related to those studies was that both monitoring of root management systems as well as communications between the various stakeholders involved in root management should be improved. While modifications to the root have, to date, not resulted in noticeable negative impact, it can be argued that without additional monitoring and improved communications, scaling of the root could pass a critical threshold without notice, resulting in scalability problems that could affect the stability of the DNS as a whole. With the assumptions that less than 1000 new TLDs will be added per year and that monitoring and communications among relevant stakeholders is improved, it seems clear that the root system should remain stable as it changes to meet new demands.