



Impact on Root Server Operations and Provisioning Due to New YF>6e

Joe Abley & Kim Davies
© 2012

Table of Contents

Management Summary	4
Introduction	5
The Root Server System.....	5
Prior Root Zone Deployment Experience	5
Risk to Root Server Operations	7
Rate of change is slow	7
Operations affected by number of queries, not number of TLDs	8
Limited delegation rate ensures slow Root Zone growth.....	9
Conclusions of studies	9
Root Server Operators indicate no anticipated risk	9
Propagation rate not affected by expected growth rate.....	9
Other Factors	10
Diversity provides protections.....	10
Coordination among Root Server Operators provides protections.....	11
Impact of failure in diverse system is minimal	11
Risk to Root Zone Provisioning	11
Root Zone Management Operations.....	13
Current activity level.....	13
Expected level.....	14
Automation System	14
Staffing Plan.....	14
Distribution from Root Zone Maintainer to Root Server Operator	15
Distribution within Root Server Operators’ Infrastructure.....	15
Monitoring	15
Facets of monitoring.....	15
Slow process	16
Formal communication channel with root operators	16
Mitigation	16
Appendix A — Previous Studies and Analyses.....	18
Appendix B — Root Server System Overview.....	20
Roles and Responsibilities	20
Root Server Diversity	20
Coordination of Root Server Operators.....	21
Root Server Capacity	21
Impact of Failure.....	21
Service Availability	22
Root Zone Provisioning.....	22

Appendix C — DNS Records in the Root Zone	23
Start of Authority (“SOA”) Record	23
Nameserver (“NS”) Record	23
IPv4 Address (“A”) Record	24
IPv6 Address (“AAAA”) Record	24
Delegation Signer (“DS”) Record	24
DNSSEC Key (“DNSKEY”) Record	25
Resource Record Signature (“RRSIG”) Record	25
Next Domain (“NSEC”) Record	25
Appendix D — Analysis of Root Zone Decay.....	27
Methodology	27
Analysis	28
Appendix E — Quantitative Analysis of Root Zone Growth	30
Executive Summary	30
Introduction.....	30
Data Collection	30
Root Zone Data	30
Analysis.....	31
Root Zone Size	31
Delegations.....	32
Root Zone Size per Delegation.....	33
Resource Records in the Root Zone.....	34
Projections.....	36
Zone Size per Delegation	36
Zone Size.....	37
On the Appreciation of Size	38
Appendix F — Case Study: Unusual Traffic Received by L-Root, June 2011	44
Executive Summary	44
Summary of Events.....	44
Impact on L-Root	44
Impact on Other Root Servers	46

Management Summary

The Root Server System has served Internet users well, due in no small measure to a robustness that is due to both its spare capacity and diversity.¹ That robustness exists in both the ability of root servers to serve the query load from the Internet and in provisioning updates to the data they serve.

This document considers the impact of the allowed growth due to new gTLD delegations from the Root Zone on all parts of the Root Server System:

- the impact on “Root Server Operations”, which is the ability of the Root Server System to answer queries from end users; and
- the impact on “Root Zone Provisioning”, comprised of:
 - the ability of the “Root Zone Management” process (the IANA process) to receive, evaluate and implement changes to the Root Zone received from TLD managers;
 - the ability of the Root Zone Maintainer (Verisign) to distribute updates to the Root Zone to individual Root Server Operators; and
 - the ability of individual Root Server Operators to distribute update within their infrastructure.

The Root Zone is comprised of resource records, a small set for each top-level domain. The number and size of records in the Root Zone has successfully grown over time, in part to accommodate new developments like the introduction of IDN ccTLDs², the first two rounds of new gTLDs³, the introduction of IPv6 glue records, and the deployment of DNSSEC in the Root Zone. There is also a natural tendency for the number of name servers per delegation to increase as TLD name server infrastructure matures over time⁴.

ICANN’s New gTLD Program⁵ is expected to substantially increase the number of records in the Root Zone. ICANN has limited⁶ growth due to New gTLDs to a maximum of 1,000 new delegations per year. This document describes evidence to support the assertion that growth in the size of the Root Zone due to new gTLDs will be accommodated comfortably by existing system capacity, and that processes exist to support such growth with the planned deployment of additional resources.

This document also reviews the impact of a failure in any of the identified processes in the event that growth does exceed system capacity, and finds that in all cases the impact on end users is negligible, and, in any case, that the time available for mitigation of any failure is substantial.

We conclude that in relation to scaling the Root Zone, it is safe to launch new gTLDs, as:

- there is very low risk that the Root Server System will suffer degraded performance due to the delegation of new gTLDs at the indicated maximum rate of delegation;
- the impact of any degradation in performance that might occur would be very low; and

¹ See **Appendix B** — “Root Server System Overview”

² <http://www.icann.org/en/resources/idn/fast-track>

³ ICANN has conducted two rounds of new gTLD additions prior to the current New gTLD Program — the “proof-of-concept” round in 2000, and the “sponsored” round in 2003.

⁴ Top-level domain operators deploy multiple name servers for redundancy and resiliency purposes, just as the Root Server System has done. As the list of name servers grow, it must be reflected in the contents of the Root Zone. A fuller analysis of these trends is provided in Appendix C — “Quantitative Analysis of Root Zone Growth.”

⁵ <http://newgtlds.icann.org/>

⁶ Section 1.1.2.5, “New gTLD Applicant Guidebook” (11 January 2012), <http://newgtlds.icann.org/en/applicants/agb/guidebook-full-11jan12-en.pdf>

- even if a degradation of performance were to occur, negative trends will be observable well in advance; the slow rate at which Root Server System performance would be impacted provides ample time to mitigate any degradation.

Introduction

When ICANN was formed in 1998 as a not-for-profit, multistakeholder organization dedicated to coordinating the Internet’s addressing system, its primary purpose was to promote competition in the domain name system (“DNS”) marketplace while ensuring internet security and stability. ICANN’s Bylaws and other foundational documents articulate that the promotion of competition in the registration of domain names is one of ICANN’s core missions.⁷

One part of this mission is fostering competition by allowing additional Top-Level Domains⁸ (“TLDs”) to be created. ICANN began this process with the “proof of concept” round for a limited number of new generic TLDs (“gTLDs”) in 2000, and then permitted a limited number of additional “sponsored” TLDs in 2004-2005. These additions to the Root Zone demonstrated that TLDs could be added without adversely affecting the security and stability of the domain name system.

After an extensive policy development process, in August 2007, the ICANN Generic Name Supporting Organization (GNSO) issued a lengthy report in which it recommended that ICANN permit a significant expansion in the number of new gTLDs. The report recognized that the introduction of new gTLDs would require the expansion of the top-level DNS zone in the DNS hierarchy known as the DNS Root Zone (“Root Zone”). This expansion of the Root Zone, along with ICANN’s recent and concurrent implementation of other changes to the root of the DNS, caused some members of the community to ask ICANN to review how the expansion of the Root Zone could impact its stability.⁹

Between 2004 and 2010, the root of the DNS underwent significant changes, both in content as well as support infrastructure. These changes included the addition of Internationalized Domain Names (“IDNs”) to the Root Zone, the deployment of IPv6, and the implementation of Domain Name System Security Extensions (“DNSSEC”). The broad scope of these changes was unprecedented. Now with new gTLDs on the horizon, further substantive changes in the root of the DNS are expected.

In response to comments from members of the community, ICANN commissioned a number of studies to address the capacity and scaling of the Root Server System with the goal of ensuring the stable and secure addition of new gTLDs. The studies improved ICANN’s understanding of the scalability of the Root Zone as it pertains to new gTLDs, and they reinforced confidence in the technical capability and stability of the Root Zone at the projected expansion rates. The studies also helped to inform and improve ICANN’s approach to monitoring the scalability and stability of the Root Zone.

The Root Server System

This report assumes familiarity with the role of the Root Zone and the Root Server System in supporting the operation of top-level domains. An overview of the Root Server System, including the various components and roles, is provided as **Appendix B** to this document — “Root Server System Overview”.

Prior Root Zone Deployment Experience

In order to determine the stability of the Root Zone with the implementation of the New gTLD Program, we first review the impact of the significant changes that had already been implemented or were in the process of being implemented into the Root Zone.

⁷ ICANN Bylaws, Article 1, Section 2.6
<http://www.icann.org/en/about/governance/bylaws>

⁸ A list of top-level domains that are currently in existence can be found at <http://www.iana.org/domains/root/db>.

⁹ <http://gns0.icann.org/issues/new-gtlds/pdp-dec05-fr-part08aug07.htm>

This section concludes that the fact that the Root Zone has absorbed significant changes without significant impact is an indicator of its resiliency and one predictor that the delegation of new gTLDs will not affect Root Zone stability.

Since February 2008, there have been significant additions to the Root Zone with the adoption and implementation of IDNs, IPv6 and DNSSEC. During the period between July 2004 when the first IPv6 addresses were added to the Root Zone for name servers, until July 2010 when the root was DNSSEC-signed and Delegation Signer Records¹⁰ were inserted, the root DNS service continued with no reported or publicly visible degradation of service. We evaluated the impact of each individual addition to the Root Zone to date, and determined that the addition of IPv6 to the root system, IDN TLDs and the deployment of DNSSEC had no significant harmful effects that were observed or reported. **Figure 1** shows a timeline of the various additions to the Root Zone since July 2004.

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	First IPv6 addresses added for root servers (A, F, J, K, L and M).
	gTLDs	A limit of a maximum of 1,000 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	First production IDNs added to the root (for Egypt, Saudi Arabia and United Arab Emirates).
	DNSSEC	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.

Figure 1 - Notable events impacting the size of the DNS Root Zone.

The deployment of new technologies continues without any significant impact to Root Zone stability. Deployment of IPv6 in the Root Zone, which began in 2004, caused no significant harmful effects. Insertion of IDNs into the Root Zone in 2007 similarly did not affect stability of the DNS, and deployment of DNSSEC in the Root Zone starting in January 2010 resulted in no observable or reported negative consequences. The empirical data drawn from the deployment of these new technologies can be used to validate the observations. Comparison of this data with the continued stability of the Root Zone throughout the implementation of these programs, indicates that the introduction of the new gTLD program at the proposed maximum rate of 1,000 applications per year would similarly not impact the stability of the Root Zone.

¹⁰ Delegation Signer (DS) Records are used as part of the DNSSEC technology, but are independent of DNSSEC-signing the Root Zone.

Risk to Root Server Operations

The primary consideration in relation to scaling the Root Zone is the ability of the root servers to reliably continue to publish the contents of the Root Zone. Without this availability, stability of the DNS could be compromised because the contents of the Root Zone are required to facilitate the DNS resolution process.

This section concludes that the introduction of new gTLDs does not pose a threat to root server operations because:

- Due to the slow rate of change to the Root Zone, the impacts of failures in the process are not felt for long periods (days or longer), but can be detected and reacted to quickly (in minutes, or hours).
- Root zone performance is primarily predicated on the number of queries, not the number of records (i.e. TLDs) in the Root Zone.
- The limited delegation rate ensures slow Root Zone growth.
- Other studies have been conducted and reach the same conclusion.
- A survey of Root Zone operators and RSSAC indicates that maximum delegation rates can be safely accommodated.
- Other characteristics of the Root Server System such as system diversity and root server operator coordination.

Rate of change is slow

One of the risks to successful publication of the Root Zone data is that the data cannot be updated in a timely manner, and therefore the data being served is out of date.

The Root Zone consists of a set of data that is largely static. It is comprised of two sources of data — the output of the Root Zone Management process, made up of technical delegation information largely provided by TLD Managers; and the DNSSEC-signing related data applied by the Root Zone Maintainer¹¹ prior to dissemination to the Root Server Operators. A description of the different data formats and how they are used is provided in **Appendix B** — “DNS Records in the Root Zone”.

We’ve reviewed several years of historical Root Zone data, in order to analyze the impact of old data¹² on the published Root Zone. With respect to the data that is supplied by TLD Managers to be published in the Root Zone, we find that TLDs would generally have high resiliency with aged data. Even if the Root Zone could not be updated for a very long period (say, a year), the vast majority of TLDs would continue to function effectively. The details of this analysis are provided in **Appendix D** — “Analysis of Root Zone Decay”.

Concerning the DNSSEC signing data¹³, the cryptographic signatures applied to the Root Zone have specific validity periods¹⁴. An inability to update these signatures can result in stale DNSSEC signature data. These stale signatures can cause validation failures, which will impact those who rely on DNSSEC

¹¹ This is through the application of the DNSSEC signing process using a “Zone Signing Key” (ZSK). Currently, Verisign conducts this process as part of its responsibilities to the US Government.

¹² For example, if updates supplied by TLD managers could not propagate, because the Root Zone Management process could not be executed.

¹³ This aspect is not influenced by the number of TLDs, and would apply equally no matter how many TLDs the Root Zone contained.

¹⁴ These intervals are stipulated in the DNSSEC policy statement, issued by the Root Zone Maintainer. See “DNSSEC Practice Statement for the Root Zone ZSK Operator”, Section 6.6, http://www.verisigninc.com/assets/root_zone_ZSK_operator.pdf

validation¹⁵. According to the current schedule, signatures are valid for 15 calendar days, and are replaced within 10 days. This means, in the worst-case scenario, the window to update the Root Zone is 5 days before the signatures are considered stale.

Therefore, due to the slow rate of change to the Root Zone, the impacts of failures in the process are not felt for long periods (days, months, or even years), but can be detected quickly (in a matter of minutes or hours). This provides comfort that delays can be identified before they introduce material harm by serving old Root Zone data.

Operations affected by number of queries, not number of TLDs

Another risk to consider is the inability of the Root Server System to sustain the load of queries being directed at it. To understand this risk, one must consider what the trigger is for additional load to the Root Server System.

The traffic to the Root Servers is driven by the number of times computer applications need to perform a DNS lookup. A DNS lookup is performed when a person on the Internet performs an action such as visiting a web page, sending an email, or views an online video¹⁶. The volume of DNS lookups, therefore, is driven by utilization of applications that use the Internet, such as web browsers and other Internet-connected software. The main operational requirement for a Root Server Operator is to successfully respond to this load of queries in a timely way, including additional capacity to cater for peaks in demand¹⁷. It is the inability to respond to the queries being transmitted to the Root Server System that would represent a deterioration of the Root Server System's capabilities.

There is no evidence to support the notion that overall Internet usage (see **Figure 2**) is related to the number of TLDs. Having twice as many TLDs does not mean that the average Internet user visits twice as many web pages, or writes twice as many emails. Rather, Internet usage is driven by growth in overall Internet adoption. Having more TLDs available does not directly incur increased usage of the DNS; rather it will exchange a subset of their query load from existing TLDs to new TLDs.




Figure 2 - Internet usage compared to number of top-level domains.
(Traffic data: Cisco Visual Networking Index.)

¹⁵ While not desirable, it is worth considering in a worst-case scenario of a known DNSSEC-related fault that impacted the signatures; DNSSEC could be disabled in such a way as to restore DNS service temporarily until proper DNSSEC signing are restored.

¹⁶ This is not an exhaustive list. The DNS has performance optimizations (i.e. caching) that mean that a lookup does not need to be performed each and every time.

¹⁷ This includes not just natural peaks in demand, but those generated by denial of service attacks and other anomalous behavior.

In **Figure 2**, the number of TLDs has experienced a small amount of growth due to expansion caused by the New gTLD Programs in 2000 and 2004, as well as the IDN Fast Track process, whereas Internet usage as a whole has grown exponentially.

Limited delegation rate ensures slow Root Zone growth

ICANN has committed to introducing no more than 1,000 new generic top-level domains into the Root Zone per year. The current Root Zone is around 150,000 bytes, and previous studies have identified that the maximum growth rate should grow the actual Root Zone file by no more than about half a megabyte per year. This is a very manageable number, and contextually represents a small file. The laboratory tests conducted showed that there was no measurable impact to root server performance metrics until millions of new TLDs were added, and the Root Zone was in the range of hundreds of millions of bytes.

As part of this assessment, we have reviewed years of historical root zone files in order to ascertain the impacts on the growth of the root zone, and to project how the root zone can potentially grow. This analysis is summarized in **Appendix E** — “Quantitative Analysis of Root Zone Growth”. At the maximum possible growth rate of 1,000 new generic top-level domains per year, the size of root zone would not grow in the foreseeable future to a magnitude that would trigger potential scaling issues using the existing root server infrastructure.

Conclusions of studies

Over recent years, a number of studies on root scaling have been conducted¹⁸. None of these studies have identified risks to safe and stable Root Zone operation that would be impacted by growing the Root Zone by 1,000 generic top-level domains per year. Further, it is identified that any potential risks will take a long period to emerge, providing ample time for remediation.

Root Server Operators indicate no anticipated risk

Root Server Operators have variously made undertakings to ICANN in relation to their ability to commit any required resources to the ongoing management of their systems to support publication of the Root Zone:

- Letter from Autonomica AB, operator of I-Root, to ICANN¹⁹, May 2009
- Letter from RIPE NCC, operator of K-Root, to ICANN²⁰, May 2009
- Letter from The WIDE Project, operator of M-Root, to ICANN²¹, May 2009
- Mutual Responsibilities Agreement between ISC, operator of F-Root, and ICANN²², December 2007

In response to ICANN’s work on studying Root Scaling, the RSSAC also wrote to ICANN in November 2010²³. The communication concluded “in the case of the proposed gradual expansion of no more than 1000 entries per year for the next several years, RSSAC expects the system to remain stable and robust.”

Propagation rate not affected by expected growth rate

Part of the operation of root servers involves ensuring that each Root Server Operator disseminates any revised data to all those servers in a timely fashion. This ensures that when DNS queries are performed against the Root Servers, they provide the most recent authoritative data.

¹⁸ See Appendix A — “Previous Studies and Analyses.”

¹⁹ <http://www.icann.org/en/correspondence/lindqvist-to-twomey-08may09-en.pdf>

²⁰ <http://www.icann.org/en/correspondence/pawlik-to-twomey-06may09-en.pdf>

²¹ <http://www.icann.org/en/correspondence/murai-to-twomey-06may09-en.pdf>

²² <http://archive.icann.org/en/froot/ICANN-ISC-MRA-26dec07.pdf>

²³ <http://www.icann.org/en/correspondence/murai-to-board-25nov10-en.pdf>

Historically, and today, the Root Zone is updated twice per day. At a minimum, this change involves altering a serial number that is contained within the zone.

As this update schedule is not predicated on either the number of TLDs, or the number of changes made to the Root Zone on any given day, the growth of the number of TLDs will not impact the propagation rate.

Even though the Root Zone update schedule is conservative and the new TLD program does not require this schedule to be altered, it has been shown by TLD operators at the second level (e.g. the .COM registry managed by Verisign), that the DNS can support and propagate highly dynamic zones with over a hundreds million domains, that alter many times per day²⁴. There is no possibility that the Root Zone can reach this volume of transactions based on the maximum growth rate, but deployment strategies that have been used in more dynamic zones could be developed and used, if necessary.

Other Factors

Diversity provides protections

The Root Server System features significant diversity, a conscious architectural approach intended to reduce the prevalence of single points of failure. For example, a newly identified software defect in the DNS software BIND might have operational consequences for those root servers which use it, but other root servers that use different software might reasonably expect to continue normal operations.

Examples of diversity in the system are:

- **Organizational diversity:** twelve different organizations²⁵ contribute to the Root Server System, and each organization makes its own choices in the architecture and operational procedures used for the root server components it is responsible for. The impact of operator error or architectural failure is hence limited to redundant parts of the system.
- **Jurisdictional diversity:** root server operators are variously headquartered in the Netherlands, Sweden, the USA and Japan. Those operators located in the USA are a mixture of for-profit corporations, not-for-profit public benefit corporations, educational institutions and government agencies.
- **Software diversity:** a variety of operating systems, DNS software, routing software and measurement/management software is in use across the system, limiting the impact of a defect or exploit in any particular software package.
- **Platform diversity:** a variety of server and network platforms are in use across the system, limiting the impact of a defect or exploit in any particular platform.
- **Topological diversity:** root server instances are located in a mixture of carrier networks and exchange points, and each instance in general operates under a unique routing policy. The effect of even multi-point topological or routing errors on the system as a whole is hence minimal.

Individual Root Servers by design incorporate significant redundancy. For example, at the time of writing L-Root service is provided by over 200 individual servers deployed in over 100 locations.²⁶

²⁴ Typically, TLD registries will deploy a dynamic zone update and distribution platform that would see changes made to a given domain reflected in the DNS within a few minutes. If we conservatively assume updates happen on average every 5 minutes, this would mean around 300 revised zones per day, with each revision potentially containing many different changes.

²⁵ A list of the organizations responsible for the architecture and operation of individual root servers can be found at <http://www.root-servers.org/>.

²⁶ A list of locations at which the various Root Server Operators maintain servers is available at www.root-servers.org.

The distribution of subsequent revisions of the Root Zone by the Root Zone Maintainer to individual Root Server Operators also incorporates significant geographic and topological diversity. Consequently, observed propagation times for new Root Zones are low, and the data being served by all root servers is highly consistent. Even in the event that individual Root Servers fail to obtain a timely revision of the Root Zone, the resulting inconsistency in answers from different servers has very low impact, with this inconsistency being accommodated by the DNS protocol itself²⁷.

Coordination among Root Server Operators provides protections

Representatives from Root Server Operators maintain frequent contact through e-mail and scheduled in-person and telephone meetings. Technical staff has access to a complete set of contact details for their peers in other organizations, and facilities exist to establish emergency teleconference bridges with all involved technical staff at all organizations. The contact methods and emergency facilities are exercised regularly to ensure that they are functional and effective.

Diversity in the Root Server System is maintained through regular coordination between Root Server Operators, and architectural decisions incorporate the goal of sustained diversity. Data collection and analysis of abnormal conditions are coordinated, and results of analysis are shared in near real-time to allow system-wide events to be characterized effectively.

Impact of failure in diverse system is minimal

The redundancy in design and operations of individual Root Servers has been shown to provide a high-fidelity and stable service despite unavailability of individual components, either due to planned maintenance or unplanned failures.

The Root Server System as a whole has been observed to provide continued, effective service to the DNS even in the event that a single root server is unavailable, although such events are rare and unusual²⁸. The DNS protocol itself allows for robust DNS service despite such outages, and testing in captive lab environments using a comprehensive and representative set of DNS software confirms that the system as a whole continues to function effectively even if 12 of the 13 Root Servers are unavailable, a situation that has never occurred on the production Internet.

Risk to Root Zone Provisioning

In addition to the role of the Root Server System to publish data contained in the Root Zone in a robust manner, one must consider the impact of growth of TLDs upon the ability for changes to be promulgated to the root servers. This “provisioning process” is responsible for the life-cycle of a change in the Root Zone data: (1) the request (typically by a TLD manager) for a change in Root Zone data; (2) the authentication and approval of that request; (3) the introduction of the change into the Root Zone System; and (4) the publication of that change in Root Zone data on all the root servers. Requests for changes can be for: delegations of new TLDs, redelegations of TLDs to a new operator, and changes to the data contained in a TLD’s records such as replacement of an authoritative name server.

Root Provisioning can be considered as three components:

- Root Zone Management Operations — The role of accepting change requests to the Root Zone, verifying and authenticating the request, and having fully-qualified changes

²⁷ The DNS protocol is built upon the notion of “loose coherence.” This means that at any given moment of time, different servers may return different answers to the same query, due to caching and other artifacts of the way DNS data is propagated. The DNS protocol is designed with resiliency to this in mind, such that this is acceptable and does not usually compromise the ability of domain names to function.

²⁸ For example, H-Root service was unavailable to much of the Internet for a prolonged period due to a network failure in October 2010. There was no reported degradation in performance or other failure in resolution due to this outage for the DNS as a whole.

implemented in the master Root Zone. This is the process currently conducted jointly among ICANN (as the IANA functions operator), the U.S. Government National Telecommunications and Information Administration (NTIA), and Verisign (as the Root Zone Maintainer);

- Distribution from the Root Zone Maintainer to Root Server Operator — After the Root Zone is modified, the Root Zone Maintainer distributes the updated data to the twelve Root Server Operators;
- Distribution within Root Server Operator’s Infrastructure — Each of the twelve Root Server Operators must distribute up-to-date copies of the Root Zone throughout their network of servers. This includes “anycast” instances that are be distributed across the globe, and in any single instance, ensuring a cluster of computing resources are all updated with the latest version.

This section concludes that the introduction of new gTLDs does not pose a threat to Root Zone provisioning because:

- Root Zone Management operations can address the increase in workload. The task is well-understood, there has been significant continual improvement over time. Manpower planning is largely complete but ICANN and its partners in the Root Zone Management chain will know precise workload requirements eight months in advance of the need.²⁹ Manpower needs are mitigated by the introduction of automation.
- The largest, immediate task – evaluation of gTLD applications – is currently and appropriately staffed. ICANN has retained qualified firms to perform each evaluation review. In most cases (i.e., for each type of evaluation). ICANN has retained multiple firms to process applications. The redundancy is intended to provide high capacity and address potential for conflicts. The firms are well established, reputable and generally well known. Each of the firms is prepared to take on the full load individually. Each firm has access to global resources to address cultural, language and other geographic issues. The evaluation categories and firms are:
 - Financial and technical evaluation: Ernst & Young, KPMG and JAS Advisors
 - Geographic names: The Economist Intelligence Unit, Interconnect Communications (working with the University College London)
 - String similarity: Interconnect Communications (working with the University College London), Interisle
 - DNS stability: Interisle³⁰
 - Community priority evaluation: The Economist Intelligence Unit; Interconnect Communications (working with the University College London)
- Distribution from Root Zone Maintainer to Root Server Operator and distribution within Root Server Operators’ infrastructure is a well-established exercise. New versions of the Root Zone are published twice daily. That frequency is expected to continue regardless of the number of TLDs. Propagation times are minimal compared to the frequency of publication and will not be affected by the forecasted changes. Sufficient geographical diversity in the publication locations already exists.

²⁹ “New gTLD Applicant Guidebook,” Section 1.1.3, <http://newgtlds.icann.org/en/applicants/agb>

³⁰ This is the only evaluation performed by one entity. Interisle is employing a large number of independent contractors that can individually backstop Interisle in case of service interruption by Interisle.

- Monitoring of these three mechanisms already exists. There is monitoring by independent third parties who publish results. Importantly, the rate of change and the rate of degradation are slow. Identification and mitigation of anomalies can easily occur before effects in Root Zone operations and provisioning would be apparent to users.
- Because detected anomalies have a delayed impact, ICANN can delay or halt delegation of additional TLDs if a delay in delegations is required. ICANN is formalizing communications with Root Server Operators to provide additional safeguards.

Root Zone Management Operations

The role of the Root Zone Management process is to accept change requests to the Root Zone, review these changes to ensure they are authorized by appropriate parties, ensure that the changes meet a certain number of technical requirements; and then implement those changes to the Root Zone once they have met all those tests.

In performing the process, ICANN takes primary responsibility in accepting change requests and shepherding through the Root Zone Management workflow. As part of the process, requests are transmitted to the NTIA, which has the role of authorizing all changes to the Root Zone. Verisign, as the Root Zone Maintainer, is responsible for executing changes that have been transmitted by ICANN, and authorized by the NTIA, in the master version of the Root Zone itself.

Current activity level

As of this writing, there are 313 top-level domains³¹, and the managers of each are responsible for communicating with ICANN any changes they wish to have effected in the Root Zone. ICANN publishes statistics on the number of such changes it receives on the ICANN Dashboard³². Over the past several years, on average, the number of requests received by ICANN is approximately one per TLD per year.




Figure 3 - Rate of Root Zone Change Requests to ICANN

ICANN is responsible for receiving and processing all requests to change the Root Zone through the IANA department. The rate of routine requests — day-to-day maintenance updates to an existing TLD — average approximately one per TLD per year. A small number of requests relate to delegations or redelegations, which refer to a material change in the party that operates the TLD, which involves a different set of procedures.

³¹ <http://data.iana.org/TLD/tlds-alpha-by-domain.txt>

³² <https://charts.icann.org/public/>

Expected level

ICANN has projected that the number of Root Zone change requests for new top-level domains will be comparable to the historical activity seen for the existing top-level domains, that is, averaging around one change per year per TLD. This means that, over the next 20 months, ICANN should plan for request rates increasing from approximately one per day to approximately four per day. In addition, ICANN will perform one-time delegation reviews for the new generic TLDs that successfully emerge from the evaluation process.

It is important to consider that the ICANN staff members who are responsible for Root Zone Management processes will not be responsible for the assessment and evaluation relating to new gTLD applications. Under the process of the New gTLD Program, applicants only commence Root Zone operations after they have completed the application evaluation process separately, and have executed a registry agreement with ICANN. This evaluation function is separately staffed to address the maximum anticipated delegation rates. This is distinct from ccTLD delegation and redelegation requests³³, where staff in the IANA Department performs the primary analysis and review.

Automation System

While ICANN can comfortably handle the requirements of processing the volume of anticipated requests through an increase in staffing, its deployment of a workflow automation system in 2011 has greatly increased the capacity of the function to scale without additional staff. As identified, the vast majority of requests to change the Root Zone data in a TLD record are categorized as routine changes. These routine changes can be automated through this system. Intervention by ICANN Staff in such changes is minimized, and therefore significant increases in change requests will not significantly impact resource requirements.

Staffing Plan

As part of the budgeting involved in the New gTLD Program, components of the ICANN budget are allocated to support funding additional staff on the IANA Department to address the additional workload caused by the new TLDs.

There are two key components to this staffing function. Firstly, a one-time IANA function review is required as part of the initial delegation of a new gTLD. This involves performing the appropriate staff review of the proposed delegation, and submitting a report to the NTIA to obtain consent to delegate the domain. While the new gTLD application evaluation is performed by ICANN's new gTLD evaluation function, the IANA function will play a role in checking the delegation request meets the requirements of the IANA Functions Contract.

Secondly, there is an ongoing administrative responsibility, as is currently required for the over 300 existing top-level domains. ICANN is responsible for managing the ongoing operational relationships with the TLD managers, including receiving updates to the Root Zone Database from qualified TLD managers, authenticating those requests, and processing them for implementation. These requests include redelegations, and updates to the name server records, the DS records, and the contact persons for a domain.

Experience from the existing community of TLDs, and the process of delegating new generic top-level domains from the previous two rounds, has given ICANN confidence to safely project the amount of resources required for both of these responsibilities. Furthermore, the introduction of automation has substantially reduced the staffing levels required, as much of the ongoing administration of domains will now be driven by the automated workflow.

³³ <http://www.iana.org/domains/root/delegation-guide>

The number and timing of potential new gTLD delegations will be known over eight months in advance of the IANA function delegation task. This provides ample time for retention and training of required staff.

Distribution from Root Zone Maintainer to Root Server Operator

In its role as Root Zone Maintainer, Verisign is also responsible for regularly distributing revisions of the Root Zone to individual Root Servers. This is done by publishing the Root Zone in a “stealth master” configuration using distributed master servers deployed in several geographically- and topologically-dispersed locations. This configuration provides private access to the Root Zone file to the Root Server Operators. Using authenticated access to the stealth masters, Root Server Operators use the DNS protocol to obtain the latest version of the Root Zone. The DNS protocol is also used to proactively notify Root Server Operators of updates to the Root Zone to trigger to update process.

Distribution within Root Server Operators’ Infrastructure

Each of the twelve Root Server Operators has its own mechanisms to distribute the Root Zone data among their own constellation of root servers. There is no specific common approach, and this diversity of approaches provides resiliency should any specific approach become problematic.

For the L-Root server, operated by ICANN, new revisions of the Root Zone are obtained from the Root Zone Maintainer for publication in two locations: Los Angeles, USA and Prague, Czech Republic. Each redistribution point is equipped with redundant server and network hardware, and hence individual instances of L-Root are able to transfer the Root Zone from a variety of locations. Each transfer of a new Root Zone revision is authenticated cryptographically, and multiple notifications of the availability of a new revision are dispatched to ensure that new data is distributed promptly. Each new distribution of the Root Zone, typically done twice per day, is propagated across the entire L-Root system and is generally completed in less than 60 seconds.

While ICANN cannot speak authoritatively about zone distribution mechanisms in use by other Root Server operators, the distribution of Root Zone data across the system as a whole is observed to be prompt and accurate.

Monitoring

While all projections suggest that the growth of the Root Zone resulting from the New gTLD Program will have no discernable impact on function of the Root Server System, it is prudent to continue to measure and monitor service delivery to identify negative trends and remedy them.

Facets of monitoring

Monitoring impact on Root Zone Management — ICANN tracks the entire life cycle of all changes to the Root Zone. Tracking is triggered when requests are submitted, and follows the various steps as the business processes are conducted. This data is compiled in monthly reports that are transmitted to NTIA, and reviewed internally for trends. Publication of more comprehensive data for external review is planned. Individually, all TLD managers — as customers of the process — are provided with a detailed timeline that describes when the key steps of their request were processed, and are provided with real-time web access to review current status and processing times for their TLD.

The trend data from these various products of the Root Zone Management process provide clear and immediate indication of the ability of the process to sustain changes in the volume of requests.

If one looks at the average processing time for requests today, versus requests from 10 years ago, there is a considerable and demonstrable improvement. This is the result of many areas of process improvement, including process optimization and implementation of workflow automation systems.

Monitoring distribution of the Root Zone — As public data, the propagation of the DNS Root Zone to the root servers can be easily monitored by any party. This tracking is facilitated by the fact that the Root Zone is time-coded with a serial number that is updated at least twice a day. Any propagation delays can be identified quickly. ICANN monitors the propagation of the Root Zone to all of the root server operators, and proactively monitors and researches the cause of any undue delay in new Root Zones appearing. Independent third parties also monitor this. Historical experience has shown that any anomalies in propagation of the Root Zone has quickly been identified by both ICANN and third parties, and reported to the relevant Root Server Operator.

Monitoring distribution within root server operators — Individual Root Server Operators operate monitoring and measuring systems overseeing the distribution of new Root Zone data across their infrastructure. These systems seek to ensure the fidelity of data in the newly distributed zones and to quickly identify and cure any delays or anomalies.

For L-Root, observed delays or anomalies in Root Zone propagation are escalated using automated systems to on-call technical staff that are empowered to diagnose and mediate any observed problems. Any changes to production systems required for remediation are made using a documented industry standard change management process³⁴.

Third-party measurement systems³⁵ also track zone propagation based on active measurement of zone serial numbers and publicly report results.

Monitoring root server operations — The performance characteristics of individual Root Servers are monitored by individual Root Server Operators as well as third parties. Third parties publicly report performance. For L-Root, comprehensive query analysis and service element monitoring is performed using multiple tools³⁶; threshold alarms and other events are escalated automatically to technical staff for investigation.

Slow process

A key, well-understood and well-accepted finding of the various studies and analyses of the Root Server System is that the growth of the Root Zone will be a slow process compared to the ability to track and respond to changes or anomalies. This feature enables deleterious effects to be observed well in advance. This means that remedial measures will be able to be identified and implemented well in advance of the growth's impact on the Root Server System.

Formal communication channel with root operators

ICANN has taken steps to establish a formal communications channel between the Root Zone Management (IANA) function and the Root Server Operators. This provides even faster feedback in cases where delays or anomalies are reported.

Contact details for the IANA liaison are distributed to all Root Server Operators, and contact details for all Root Server operators are shared with ICANN's IANA liaison. The accuracy of contact details is tested regularly.

Mitigation

Based on the foregoing discussion there is little risk that there will be a noticeable impact of Root Server System performance caused by the introduction of new gTLDs. Among other factors, this low risk is

³⁴ The process uses the Information Technology Infrastructure Library (ITIL) methodology.

³⁵ For example, the RIPE NCC's "dnsmon" and "ATLAS" platforms, and those offered by Team Cymru at <http://www.team-cymru.org/Monitoring/DNS/>.

³⁶ Such as DSC, Observium, Intermapper, and Nagios.

assured by the pre-set maximum delegation rate for new gTLDs into the Root Zone, and the successful operation of DNS zones that are many orders of magnitude in size larger.

Nonetheless, ICANN is prepared for this unlikely probability. If monitoring should detect a negative effect caused by the growth of the number of new gTLDs, ICANN's response to the ICANN Governmental Advisory Committee (GAC) published in May 2011 made undertakings to address this issue³⁷:

- ICANN will implement a process with a clearly established chain of command that enables the delegation of TLDs to be slowed or stopped at any time in the event that anomalies occur and are not timely cured; since delegations are approved by ICANN, this procedure will have no external dependencies that might jeopardize Root Zone stability;
- ICANN commits to review the effects of the New gTLD Program on the operations of the Root Zone system, and to postpone delegations in a second round until it is determined that the delegations in the first round have not jeopardized the Root Zone system's security or stability;
- ICANN commits to ensuring that the operation of the IANA functions and ICANN's coordination of the Root Zone system will not be negatively affected.

³⁷

<http://www.icann.org/en/topics/new-gtlds/gac-comments-new-gtlds-26may11-en.pdf>

Appendix A — Previous Studies and Analyses

A number of studies and analyses have been compiled on the topic of scaling of the DNS Root Zone in the last three years:

- **Root Zone Augmentation and Impact Analysis**³⁸ — This study, also known as the “L-Root Study”, examined the impacts of multifaceted growth of the size of the Root Zone on the performance of “l.root-servers.net”, the root server that is managed by ICANN. This analysis considered the implications of IPv6 addresses, DNSSEC, as well as new TLDs in a laboratory simulation. The work was conducted by the independent DNS Operations and Research Center (DNS-OARC). This study was published in September 2009, and its conclusions include that root zone servers’ requirements for memory grow linearly with the number of top-level domains, that the then-deployed software was capable of handling at least 100,000 top-level domains before there was potential degradation in response times. Other software in use had higher thresholds (i.e. over a million TLDs) before the size of the root zone became a factor.
- **Report on the Impact on the DNS Root System of Increasing the Size and Volatility of the Root Zone**³⁹ — This report was developed by a specially convened “Root Server Scaling Team” (RSST), comprised of experts from the RSSAC, SSAC as well as experts from outside the ICANN community.
- **Summary of the Impact of Root Zone Scaling**⁴⁰ — An analysis of issues relating to the impact of Root Zone Scaling was prepared by ICANN and published the document in October 2010. This document considers the findings of the DNS-OARC and RSST reports, and the various root scaling events to date. It identifies the impacts from IPv6 deployment, TLD growth, DNSSEC deployment and other factors, and concludes the maximum growth to the Root Zone caused by the New gTLD Program is unlikely to cause any disruption.
- **Report of the Security and Stability Advisory Committee on Root Scaling**⁴¹ — ICANN’s Security and Stability Advisory Committee reviewed the original research questions relating to Root Zone Scaling, and provided recommendations relating to processes to handle the increase in the number of top-level domains.
- **Explanatory memorandum on Root Zone Scaling**⁴² — As part of a number of briefing papers associated with the New gTLD Program following dialogue between the ICANN Board and the Governmental Advisory Committee, ICANN published a memorandum concerning Root Scaling in April 2011.
- **Board response to the GAC on Root Zone Scaling**⁴³ — ICANN published an additional response to the issues raised on dialogue between the ICANN Board and the Governmental Advisory Committee, by providing further detail on how ICANN

³⁸ <http://www.icann.org/en/topics/ssr/root-zone-augmentation-analysis-17sep09-en.pdf>

³⁹ <http://www.icann.org/en/committees/dns-root/root-scaling-study-report-31aug09-en.pdf>

⁴⁰ <http://archive.icann.org/en/topics/new-gtlds/summary-of-impact-root-zone-scaling-06oct10-en.pdf>

⁴¹ <http://www.icann.org/en/groups/ssac/documents/sac-046-en.pdf>

⁴² <http://archive.icann.org/en/topics/new-gtlds/root-zone-scaling-15apr11-en.pdf>

⁴³ <http://archive.icann.org/en/topics/new-gtlds/root-scaling-30may11-en.pdf>

undertakes to address the ICANN community's and the GAC's concerns regarding root scaling. This response was published in May 2011.

Appendix B — Root Server System Overview

The Root Server System is a key component of the global DNS infrastructure. The vast majority of Internet applications make use of the DNS, both at the explicit direction of a user who uses a domain name to locate a specific resource or service and in lower-level protocols.

Applications, operating systems and intermediate DNS servers generally cache responses from the DNS. Data served by the DNS is consequently made available from a massively diverse and very widely distributed set of sources. Caching of DNS data enhances the performance and reliability of the DNS as a whole.

The Root Server System comprises the data contained within the Root Zone and the name server infrastructure (the Root Servers themselves) that makes that data available to DNS clients on the Internet.

A DNS client with no cached information (or one in which a particular set of cached data needs to be refreshed) consults authority servers. The fundamental authority servers in the DNS are those which serve the DNS Root Zone; the Root Servers which make that zone available provide referrals to other servers based on query name. Availability and accuracy of Root Zone data is hence important; the inability to receive a referral from any root server or the receipt of a response that is not accurate has the potential to degrade the effectiveness of the DNS globally.

Management of the data in the Root Zone is carried out under multi-party scrutiny, following well-defined processes. Distribution of the Root Zone data and publication of the Root Zone itself through Root Servers involves a great deal of technical and organizational diversity with the goal of making the system extremely stable and resilient, and as a result the system as a whole has exhibited very high availability.

Roles and Responsibilities

The data that is ultimately published in the Root Zone is maintained by ICANN as an IANA function, primarily under the direction of managers of Top-Level Domains (TLDs). Processes for Root Zone management ensure that the technical and organizational impact of any proposed change is carefully assessed prior to implementation.

Proposed changes in the Root Zone are implemented using systems and processes that involve external organizations. Changes are communicated to the US Department of Commerce, National Telecommunications and Information Administration (NTIA) for authorization, and to Verisign, Inc., acting as Root Zone Maintainer, for technical implementation and publication.

Authorized changes, once fully implemented, are published using standard DNS protocols and distributed by the Root Zone Maintainer to Root Server Operators.

Twelve Root Server Operators operate thirteen root server instances⁴⁴, known as A.ROOT-SERVERS.NET, B.ROOT-SERVERS.NET, through M.ROOT-SERVERS.NET. The Root Server operated by ICANN is known as L.ROOT-SERVERS.NET (or, colloquially, as L-Root).

Root Server Diversity

The twelve independent, autonomous organizations that operate the Root Servers use a variety of geographic locations, network connectivity providers, network and server hardware and DNS software. Service components are managed according to individual operators' best practices, ensuring massive diversity in almost every aspect of service delivery. This diversity makes it possible for the system as a

⁴⁴ A complete list of the organizations that carry out the role of Root Server Operator can be found at <http://www.root-servers.org/>, together with other related technical information.

whole to survive even unlikely coincidences of operator errors, software defects, network congestion or failure and natural disasters.

Contributing further to the diversity described above, many of the Root Server Operators have incorporated additional diversity into the deployment of their individual Root Servers. L-Root, for example, is deployed in a large number of locations using anycast⁴⁵. At the time of writing, the number of service delivery points for the Root Server System as a whole exceeds 260.

In security terms, the significant redundancy and diversity of locations, software, hardware and management of the Root Servers provides a substantial challenge to any potential attacker, since denial of service to the system as a whole would require a coordinated attack on an enormous array of software and infrastructure in multiple legal jurisdictions.

Coordination of Root Server Operators

Root Server Operators, whilst maintaining autonomy and operational independence, communicate regularly and maintain close inter-personal familiarity in order to ensure that functional coordination is possible both in normal operations and in crisis management.

Technical representatives from Root Server Operators meet in person at least three times per year, coincident with IETF meetings. Communication by teleconference, e-mail and instant messaging are also regularly exercised.

Root Server Operators have participated in multiple table-top exercises during which fictional crises were managed in an attempt to model real-world responses to events which might threaten the stability and availability of the Root Server System.

Root Server Capacity

In 2000, the IETF has provided guidance⁴⁶ for capacity planning for Root Server Operators. That guidance suggested that individual Root Servers be provisioned with sufficient capacity to accommodate a sustained query load three times greater than the normal measured steady-state load.

In current practice, the increased performance of network and server components and the prevalence of service distribution using *anycast* since that guidance was provided has resulted in infrastructure that far exceeds those requirements. For example, L-Root's baseline traffic is less than 20,000 queries/second; aggregate capacity of L-Root servers (distributed between more than 50 locations) exceeds 10,000,000 queries/second, more than 500 times the steady-state load.

Impact of Failure

DNS data is frequently cached, as a natural consequence of the way that the DNS protocols are designed. At the time of writing, referrals from the Root Zone are provided with a time-to-live (TTL) of 172,800 seconds, or 48 hours. A DNS client with a cached referral hence has no requirement to send more than one query every 48 hours for each TLD.

By consequence, a simultaneous, catastrophic failure of all Root Servers would have no immediate global consequence. Failures in the DNS would be sporadic and widely-distributed, and impact to end users would increase in prevalence as records expire from caches.

⁴⁵ RFC 4786, "Operation of Anycast Services", J. Abley, K. Lindqvist, December 2006.
<http://tools.ietf.org/html/rfc4786>

⁴⁶ RFC 2870, "Root Name Server Operational Requirements", R. Bush, D. Karrenberg, M. Koster, R. Plzak, June 2000.
<http://tools.ietf.org/html/rfc2870>

Isolated failures in individual Root Servers have been observed, with no discernable impact on end users. This is a consequence of the DNS protocol; the non-availability of individual Root Servers does not prevent name resolution so long as at least one other Root Server is available to provide responses.

Service Availability

There are no recorded examples of catastrophic failure of the Root Server System as a whole (i.e. global, simultaneous non-availability of all Root Servers). Observed availability of the Root Server System since its inception is hence 100%.

Root Zone Provisioning

The ICANN Root Zone Management staff, as part of the IANA functions contract, coordinates the contents of the Root Zone. This process involves receiving requests from top-level domain managers to make changes, reviewing the changes for accuracy and ensuring they are agreed by authorizing parties from that TLD manager; and then submitting them for authorization and implementation by the NTIA and Verisign, respectively.

While the process has successfully functioned for many years, in 2011 the parties concluded a multi-year roll-out of a new Root Zone Workflow Automation system. This system automates, to the greatest extent possible, the processing of Root Zone change requests. The automated process still provides for human review to ensure that there is a safety check to prevent errant changes reaching the Root Zone. Importantly, the process has been designed in such a way that the parties can always agree to revert to fully manual processing in the event of a disaster, or if the automation system was in any way compromised.

The automation system has been designed to cope with volumes of change requests far in excess of requirement, and well beyond the highest projections of new TLD growth. Experience has shown typically TLD managers make, on average, 1-2 Root Zone changes per annum. Even with thousands or tens of thousands of TLDs, the system can comfortably support this number of changes. Both ICANN and Verisign have experience running transactional systems with higher transaction rates, and are confident the system can be managed for growth well in advance of requirements.

In an emergency scenario, it is unlikely an emergency Root Zone change is required to sustain operation of a TLD. The DNS is expressly designed to utilize multiple redundant authoritative name servers, such that even if some had a problem, ongoing function of a given TLD would be uncompromised. ICANN conducts technical reviews in order to assess whether each TLD's name server setup is sufficiently diverse to continue operating in such an event. This review occurs in detail during a delegation or redelegation request, and is also conducted with every Root Zone change as part of the standard set of technical checks.

There has never been a case where a catastrophic failure of a TLD's name servers has necessitated wholesale replacement of their Root Zone records to restore function of the TLD. The only scenario where a TLD has required a change of this magnitude was due to business failure of a TLD, rather than as a result of technical issues.

In the rare event such an emergency scenario needs to be enacted, ICANN maintains a 24x7 emergency response service. Launched in 2006, this service has only once been enacted for an emergency. After liaising with the reporting party, that emergency was deemed not to require a Root Zone change. ICANN regularly tests the service to ensure it is functioning correctly.

Appendix C — DNS Records in the Root Zone

The DNS Root Zone is comprised of a set of data records, published as a single file. Those data records are comprised of several types of data, each of which are used in different parts of the DNS resolution process. This appendix briefly summarizes the purpose of each of these record types, where the data they contain originates, and actual samples of these records to illustrate the nature of the data that appears in the Root Zone.

Start of Authority (“SOA”) Record

The SOA record⁴⁷ is provided as the first record in the Root Zone, and contains important metadata relating to the operation of the entire zone. Most notably, it contains a serial number which is used as a signal to identify when the contents of the zone has changed. By convention, the serial number for the Root Zone is of the form YYYYMMDDSS, whereby the first four digits represent the year of publication, the next two digits represent the month, the next two digits the day, and the last two digits represent the edition on a given day (the first being “00”, the second “01”, and so on.)

The source of the data contained within the SOA record is the Root Zone Maintainer. Apart from changes to the serial number, the other values of the SOA record are generally not subject to change.

A sample SOA record is:

```
. 86400 IN SOA a.root-servers.net. nstld.verisign-grs.com.  
2012031201 1800 900 604800 86400
```

Serial number


Nameserver (“NS”) Record

The NS record⁴⁸ is used to list the names of the individual nameservers that are authoritative for a particular domain. For the root zone, this includes both the list of the root servers themselves, and the list of nameservers for every top-level domain.

The source of the data contained within a set of NS records for a given TLD, is the manager of that specific top-level domain. When a top-level domain administrator wishes to change the set, firstly they change a matching set configured in their servers, and then they request a root zone change from IANA through the IANA functions.

A sample set of NS records is:

```
mw. 172800 IN NS mw.cctld.authdns.ripe.net.  
mw. 172800 IN NS rip.psg.com.  
mw. 172800 IN NS sec3.apnic.net.  
mw. 172800 IN NS domwe.sdn.mw.  
mw. 172800 IN NS chambo.sdn.org.mw.
```

Authoritative name servers for .mw

⁴⁷ RFC 1035, “Domain Implementation and Specification”, P. Mockpetris, Section 3.3.13.
<http://tools.ietf.org/html/rfc1035>

⁴⁸ Ibid., Section 3.3.11.

IPv4 Address (“A”) Record

The A record⁴⁹ is used to signal the IPv4 address for a given name server. These records, as they appear in the root zone, are known as “glue” records and are used to facilitate discovery of the network location of name servers.

```
ns1.hkirc.net.hk. 172800 IN A 203.119.2.18
                        └──────────┘
                        IPv4 address
```

IPv6 Address (“AAAA”) Record

The AAAA record⁵⁰ is used to signal the IPv6 address for a given name server. As with A records, they represent “glue” records that facilitate discovery of the network location of name servers.

The nameserver operator initially provides the IPv6 address to the operator of the top-level domain. The top-level domain operator then transmits the contents to ICANN in order to commence the Root Zone change process.

An example AAAA record is:

```
ph.cctld.authdns.ripe.net. 172800 IN AAAA 2001:67c:e0:0:0:0:104
                                           └──────────┘
                                           IPv6 address
```

Delegation Signer (“DS”) Record

The DS record⁵¹ is a summary of the key used to sign a specific top-level domain. When a top-level domain implements DNSSEC, they provide this summary (known as a “digest”) to be listed in the DNS root zone to implement the “chain of trust” between the signed root zone, and the signed top-level domain.

The origin of this data is the top-level domain manager, which provides changes to the listed DS records in the root zone to ICANN via the IANA functions. The digest itself is machine-generated by the TLD manager’s software using algorithms based on the top-level domain’s DNSKEYs that the TLD manager maintains.

An example DS record is:

```
xn--jxalpdlp. 86400 IN DS 56231 8 1 C686FC34C432A82BD0F7A8569C32BBA8152B2D0D
                                           └──────────┘
                                           key digest
```

⁴⁹ Ibid., Section 3.4.1

⁵⁰ RFC 3596, “DNS Extensions to Support IP Version 6”, S. Thomson et.al.
<http://tools.ietf.org/html/rfc3596>

⁵¹ RFC 4034, “Resource records for the DNS Security Extensions”, R. Arends et.al., Section 5.
<http://tools.ietf.org/html/rfc4034>

DNSSEC Key (“DNSKEY”) Record

The DNSKEY record⁵² represents the public component of the key(s) that is used to sign the Root Zone. In the DNS Root Zone, each individual record either represents a “Key Signing Key”, or a “Zone Signing Key”. A “Key Signing Key” is a key that was generated in a key ceremony conducted by ICANN, involving trusted community representatives. A “Zone Signing Key” is a key generated by Verisign as the Root Zone Maintainer, and signed using a Key Signing Key by ICANN. Each of these keys is used in DNSSEC as part of the signing and verification process.

The elements of the record itself are generally unintelligible, but act as inputs into cryptographic functions that DNSSEC-enabled software uses to verify the validity of a domain.

A sample DNSKEY record is:

```
. 172800 IN DNSKEY 257 3 8 AwEAAgAIK1VZrpC6Ia7gEzah0R+9W29euxhJhVVLOyQbSEW008g
cCjFFVQUTf6v58fLjwBd0YI0EzrAcQqBGCzh/RStIo08g0NfnfL2MTJRkxoXbfDaUeVPQuYEhg37NZW
AJQ9VnMVDxP/VHL496M/QZxkjf5/Efucp2gaDX6RS6CXpoY68LsvPVjR0ZSwzz1apAzvN9dlzEheX7I
CJBBtuA6G3LQpzW5h0A2hzCTMjJPJ8LbqF6dsV6DoBQzgu10sGIcGOY170yQdXfZ57re1SQageu+ipAd
TTJ25AsRTAoub80NGcLmqrAmRLKBP1dfwhYB4N7knNnulqQxA+Uk1ihz0=
```

Resource Record Signature (“RRSIG”) Record

The RRSIG record⁵³ represents the signature for a specific set of resource records of a given type and domain. This signature is verified using the DNSSEC protocol in order to ascertain whether the resource records being signed have been modified in transit.

As with the DNSKEY record, the elements of the record are not intelligible, but rather act as inputs into the cryptographic verification functions of DNSSEC. The source of the RRSIG record is the zone signing software used by Verisign when it creates a new edition of the Root Zone. The exact ordering and contents is fully dictated by the formal requirements of the DNSSEC protocol.

A sample RRSIG record is:

```
BIZ. 86400 IN RRSIG DS 8 1 86400 20120609000000 20120601230000 56158 . qyAqJIRE
r8g9qAF0b1DEdAVBva0ugspZi0oS1W6Ba5u2HEPz0RtZYJYX3IShZmDaZdGEG/EYQPTCItHUMVpoTAi
hu2K9kf7drz3tdM2v4N8JNnrOfJXv6f6yEFJWSBzoZcn79oWBG+LM5SNWDOEjOXkw10wprYucIUE5iw
0AACM=
```

Next Domain (“NSEC”) Record

The NSEC record⁵⁴ is used to denote the next secure record contained within a zone. Such records are used by DNSSEC validators to identify which records verifiably exist in a zone, and by extension, which records verifiably do not exist in a zone. It is used by DNSSEC to prove non-existent domains do not exist.

The source of the NSEC record is the zone signing software used by Verisign when it creates a new edition of the Root Zone. The exact ordering and contents is fully dictated by the formal requirements of the DNSSEC protocol.

⁵² Ibid., Section 2.

⁵³ Ibid., Section 3.

⁵⁴ Ibid., Section 4.

A sample NSEC record is:

```
et. 86400 IN NSEC eu. NS RRSIG NSEC
```

Appendix D — Analysis of Root Zone Decay

To assess the impact of stale root zone data, an analysis was performed using historical data to identify which elements of the root zone had changed over time and to calculate the operational impact of those changes. This gives a good indication of the operational impact of events that may delay propagation of Root Zone changes.

Methodology

To conduct this assessment, a set of historical root zones were compared against a contemporary root zone⁵⁵, and compared to identify which specific resource records had changed, and what the impact of those changes would be.

To build a picture of decay over time, a daily snapshot was taken for every day in the two years prior to the contemporary root zone. By comparing the older snapshot and the current snapshot of root zone data, the following elements were determined for each top-level domain:

- **No errors.** Are all the list name servers and DS records for a given TLD matching between the contemporary root zone and the old data set? If so, there is no impact to that TLD by using the old root zone data.
- **Nameserver mismatch.** Are any of the name servers⁵⁶ for the TLD listed in the old data set not listed in the contemporary data set? If so, the TLD is considered “partly lame”, because queries using the old data would be directed to name servers that no longer are authoritative⁵⁷ for that TLD.
- **No common nameservers.** Are all of the name servers for the TLD in the contemporary data set not listed in the old data set? If so, the TLD is considered “lame” and non-functional. This is because using the old data, you could not reach any of the name servers for the given TLD.
- **DS record missing.** Are there no DS records in the old data set, when there are DS records in the contemporary data set? If so, the TLD would function correctly, but would not be DNSSEC-signed and therefore be considered “insecure”.
- **DS record mismatch.** Are there DS records different in the old data set compared to the contemporary data set? If so, the TLD is considered “provably insecure” because a resolution would fail due to a DNSSEC validation error. This is considered a hard failure.

⁵⁵ The root zone as published on 31 May 2012 (serial 2012053101) was used as the contemporary reference point for this analysis.

⁵⁶ Comparisons were performed against the IP addresses, rather than the hostname. This is because it is a change to the IP address, rather than the hostname, that affects the ability of a given host to answer authoritatively for a domain name.

⁵⁷ This approach assumes that once a nameserver is removed from the zone, it is no longer answering authoritatively for a domain, which is the worst case scenario. In reality, this isn't always true, and therefore it could be expected that the actual impact would be lower than the results of this analysis.

Analysis

The analysis as performed, graphed in **Figure 4**, shows how the age of the root zone data impacts the operations of TLDs.




Figure 4 - Impact of age of root data on successful TLD operation.

The red line indicates the number of TLDs that would not function due to old Root Zone data. It would take weeks or even months of being unable to update the Root Zone before there was a significant noticeable impact on TLD delegations.

Some of the insights that can be gleaned from this include:

- Using week old data, all top-level domains would function correctly.
- Using one month old data, all top-level domains would function correctly, with the exception of 1 TLD which is provably insecure due to a DNSSEC key rollover. One other TLDs would be deemed insecure as their DS record is not listed, but would still function. Thirteen TLDs are partially lame, due to some name server changes, but would continue to function because they still have operational name servers.
- Using three month old data, two top-level domains would not function correctly as there were no common nameservers in the root zone. Thirty-one TLDs would be partially lame, but continue to function. Three TLDs would be provably insecure.
- Using one year old data, eleven top-level domains would not function correctly as there were no common nameservers in the root zone. 142 TLDs would be partially lame, but continue to function. Thirteen would be provably insecure.

The most common case of negative affects that may present is where one or more name servers is no longer functioning, but others are still functioning. In such cases, known as having a “partially lame” delegation, clients will retry their DNS queries until they reach the functioning name servers, resulting in a successful lookup. The ability of the DNS to be resilient to partial availability issues such as this is one of its strengths.

In conclusion, with respect to the data that is supplied by TLD Managers to be published in the Root Zone, we find that TLDs would generally have high resiliency with aged data. Even if the Root Zone could

not be updated for a very long period (say, a year), the vast majority of the 314 TLDs in existence would continue to function effectively from the perspective of end users, even with the stale Root Zone data.

Appendix E — Quantitative Analysis of Root Zone Growth

Management Summary

Copies of root zones published over a twelve-year period were compared in order to investigate the characteristics of root zone growth.

The introduction of IPv6 glue (AAAA records) and DNSSEC (DNSKEY, NSEC, RRSIG and DS) resource records appears to have had a visible, positive impact on the growth rate. The impact of DNSSEC has been more significant than that of IPv6 glue.

The contribution to the compiled size of the root zone by individual delegations has increased from under 200 bytes in 2000 to over 600 bytes in 2012. The trend since 2010⁵⁸ is for this size contribution per delegation to increase, and that growth now being driven primarily by IPv6 glue and DNSSEC resource records.

A crude model is provided for extrapolating root zone size in the future, incorporating a proposed linear increase in root zone size per delegation together with a projected 1,000 new gTLDs added per year from the beginning of 2013. Under this model, the size of the root zone is expected to increase to a size of over 2.5 megabytes by 2016 from its current size of around 200 kilobytes.

Introduction

The root zone of the Domain Name System⁵⁹ (DNS) has grown over time due to various contributory factors, such as the delegation of new top-level domains (TLDs), the addition of DNS resource records to existing delegations and the deployment of DNSSEC in the root zone.

This document provides some quantitative insight into the effect of these various contributory factors, and presents a crude model of root zone size based in part on observed behavior, and in part on a projected 1,000 new delegations per year added as part of ICANN's New gTLD Program.

Data Collection

Root Zone Data

The collection of root zone data used for this analysis is that retained by the Domain Name System Operations Analysis and Research Center (DNS-OARC) as part of its Zone File Repository project⁶⁰. This collection contains root zone snapshots from between May 1999 and April 2012, and although incomplete, it includes over 4,000 distinct root zones that provide a representative set of data for the purposes of this analysis⁶¹.

⁵⁸ Insofar as trends are reasonable to infer, given the relatively short sample period since that time.

⁵⁹ RFC 1034, "Domain Names — Concepts and Facilities", RFC 1034, P. Mockapetris
<http://tools.ietf.org/html/rfc1034>

⁶⁰ <https://www.dns-oarc.net/oarc/data/zfr>

⁶¹ It should be noted that the collection does not include root zones distributed with DNSSEC resource records before the production deployment of DNSSEC; that is, the Deliberately Unvalidatable Root Zones that were used as part of the trial deployment of DNSSEC.

Each root zone in the collection was examined, and the following parameters⁶² were extracted:

Figure 5 - Parameters

<i>serial</i>	The serial number field from the root zone Start of Authority (SOA) RDATA.
<i>size</i>	The compiled (wire-format) size of the root zone, in bytes.
<i>count_tld</i>	The number of delegations present in the root zone (and hence, the number of top-level domains in the DNS).
<i>length_tld</i>	The length of all top-level domain labels (A-Labels, in the case of IDN TLDs) concatenated together, in bytes.
<i>length_glue</i>	The length of all glue labels (that is, the owner names of all A and AAAA resource record sets in the root zone) concatenated together, in bytes.
<i>count_rr</i>	The number of resource records in the root zone.
<i>count_ns</i>	The number of NS records in the root zone.
<i>count_a</i>	The number of A records in the root zone.
<i>count_aaaa</i>	The number of AAA records in the root zone.
<i>count_dnskey</i>	The number of DNSKEY records in the root zone.
<i>count_ds</i>	The number of DS records in the root zone.
<i>count_rrsig</i>	The number of RRSIG records in the root zone.
<i>count_nsec</i>	The number of NSEC records in the root zone.

Analysis

Root Zone Size

The compiled size of the root zone over the entire period for which data was available is shown in **Figure 6**. The growth is observed to be positive and linear over three distinct time intervals:

- Before 20 July 2004, when the first AAAA record was published in the root zone;
- Between 20 July 2004 and 15 July 2010, when the first DNSSEC records were published in the root zone; and
- Following 15 July 2010.

There is no obvious change in the characteristics of root zone growth due to other notable effects during the data collection period, such as the introduction of test internationalized top-level domains (IDN TLDs) in August 2007⁶³, and the introduction of production IDN TLDs in May 2010⁶⁴.

⁶² The script and resulting data used for this report is not included in their report, but can be provided upon request. Data collection was performed on DNS-OARC systems, in accordance with DNS-OARC's Participation Agreement at <https://www.dns-oarc.net/files/agreements/oarc-participation-201102.pdf>

⁶³ <http://www.iana.org/domains/idn-test/>

⁶⁴ <http://www.icann.org/en/news/announcements/announcement-05may10-en.htm>




Figure 6 - Size of Root Zone File (1999-2012).

The deployment of DNSSEC in the root zone involved the publication of a substantial number of new resource records, many of which are comparatively large. For example, the DNSKEY and RRSIG resource record sets contain cryptographic keys and signatures generated using them. Each additional resource record set is published with an accompanying RRSIG record; each additional owner name is published with an accompanying NSEC; and each additional DS resource record set, used to signal a chain of trust to a signed child zone, increases the data associated with its respective delegation. The deployment of DNSSEC, therefore, has provided an additional, and more pronounced, amplification to any natural growth evident during the period observed.

Delegations

The number of delegations in the root zone over time can be seen in **Figure 7**. The number of delegations is observed to have grown over time, although the total number of delegations remains relatively small⁶⁵.

Significant events that are apparent include:

- the delegation of eleven test IDN TLDs in August 2007; and
- the deployment following the delegation of the first production IDN TLDs in May 2010.

The number of delegations is observed to decline occasionally, but the general trend over the observed period is for the number of delegations to increase. The observed rate of growth has increased since the introduction of IDN TLDs.

⁶⁵ There were 314 delegations from the root zone as at 18 April 2012.




Figure 7 - Number of delegations (TLDs) in the Root Zone (1999-2012).

Root Zone Size per Delegation

The impact on root zone size of additional delegations (simply the compiled size of the root zone divided by the number of delegations present in the zone) is shown in **Figure 8**.

The increase in impact on zone size for individual delegations over time is consistent with the observed tendency for TLD infrastructure to mature, e.g. a trend towards more nameservers per TLD, reduced reliance on individual nameservers which serve multiple TLDs. For more discussion, see observations of the use of particular resource records by TLDs in the following section.

Delegations have a greater impact on zone size following the deployment of DNSSEC, consistent with the discussion on DNSSEC impact on overall root zone size.




Figure 8 - Root Zone Size per delegation (1999-2012). This graph represents the size of the DNS root zone, divided by the number of delegations.

Resource Records in the Root Zone

The number of various types of resource record present in the root zone per delegation is shown in **Figure 9**.

It is observed that the number of non-DNSSEC resource record types per delegation (NS, A, AAAA) increased steadily from the beginning of the dataset (or, in the case of AAAA, from the time that AAAA records were first published in 2004). However, the final two years of data suggest that the trend for increased nameservers and accompanying IPv4 glue records per delegation has peaked. The number of AAAA records per delegation continues to increase, which is consistent with general industry trends to deploy IPv6, an effort it is clear is far from over.




Figure 9 - Average number of records per delegation, by type. (1999-2012)

Of the DNSSEC resource records shown:

- DNSKEY growth per delegation is flat; DNSKEY resource records are placed at the apex of the root zone and vary according to key rollover schedules, which are independent of the number of delegations in the zone;
- NSEC growth per delegation is flat; since NSEC records are attached to owner names, and every new delegation necessarily involves a new owner name, it follows that the ratio of delegations to NSEC resource records should be one;
- DS growth per delegation is increasing slowly as more TLDs are signed. The rate (and amount) of DNSSEC adoption remains fairly modest, consistent with the small number of DS records per delegation and the observed low but positive growth; and
- RRSIG growth per delegation is fairly static; since RRSIG records are attached to resource record sets for which the root zone is authoritative (apex SOA, NS and DNSKEY resource records, NSEC resource records and DS resource records). Since the number of NSEC records in the zone increases linearly with owner names (and hence delegations) and the number of apex resource record sets is constant, growth in RRSIG records per delegation is primarily driven by DS resource record provisioning.

Of the resource records surveyed, it appears that only AAAA and DS records are experiencing growth across all delegations, and that NS and A records may be declining slightly. The net effect of these contributing factors is shown in **Figure 10**.




Figure 10 - Total number of resource records per delegation (1999-2012). The distinct change in July 2010 reflects the commencement of publishing a DNSSEC-signed root zone, which introduced an immediate and significant increase of resource records per delegation.

While the ongoing deployment of DNSSEC provides continued (yet modest) growth in the number of resource records per delegation, the trend observed since 2000 of the number of non-DNSSEC resource records per delegation to increase appears to have reached a steady state with low growth, perhaps indicative of maturity in the DNS service market.

Simple Projection

The following is a simple extrapolation of observed trends in root zone size per delegation, together with an anticipated additional growth in the number of delegations due to ICANN's New gTLD Program. Confidence in this projection is derived from the observed linear relationship between zone size and the number of delegations. This model does not expose individual parameters, however, and assumes that the underlying factors driving growth in zone size per delegation will continue unchanged.

Zone Size per Delegation

Extrapolating from the post-DNSSEC zone size per delegation relationship observed in **Figure 8** using parameters from the least-square derivation of the linear relationship shown, and assuming continued linear growth, we obtain the projection shown in **Figure 11**.




Figure 11 - Projected root zone size per delegation (2012-2016).

Zone Size

The New gTLD Program⁶⁶ at ICANN has the potential to result in an additional 1,000⁶⁷ new delegations to the root zone every year. Natural growth in delegations (without new gTLDs) is observed to be positive but small (see **Figure 7**) and for the purposes of this simple model will be assumed to be zero⁶⁸. The model assumes 1,000 new gTLDs are added to the root zone evenly throughout the year, starting on 1 January 2013, with each contributing towards the size of the root zone according to the projections shown in **Figure 11**.

The projected future size of the root zone implied by this model can be seen in **Figure 12**. The size of the root zone is shown to increase to over 2.5 megabytes in size by 2016 from its current size of around 200 kilobytes.

⁶⁶ <http://newgtlds.icann.org/>

⁶⁷ The number 1,000 is a limit that has been set by policy, and represents an upper bound on the number of delegations that could be added under this program, as currently defined.

⁶⁸ Growth by other factors is predominantly caused by the addition and removal of new country-code top-level domains, which is usually predicated on the creation of new countries and the dissolution of former countries.




Figure 12 - Projected size of root zone file (2012-2016)

Numerical Model

The simple projection presented in **Figure 12**, as noted, does not expose any of the individual parameters in the data set, and instead assumes that observed patterns in growth will continue as the number of delegations continues to increase. The following model is provided to allow estimates of root zone size to be made based on different patterns.

Interdependence of Parameters

The Pearson product-moment correlation coefficient between all pairs of parameters was calculated in order to provide some measure of whether particular pairs of parameters exhibit a linear relationship. Three data sets were considered, as in earlier analysis: root zones prior to the introduction of AAAA glue records (see **Figure 13**), root zones with AAAA records but before the deployment of DNSSEC (**Figure 14**), and root zones after the deployment of DNSSEC (**Figure 15**).

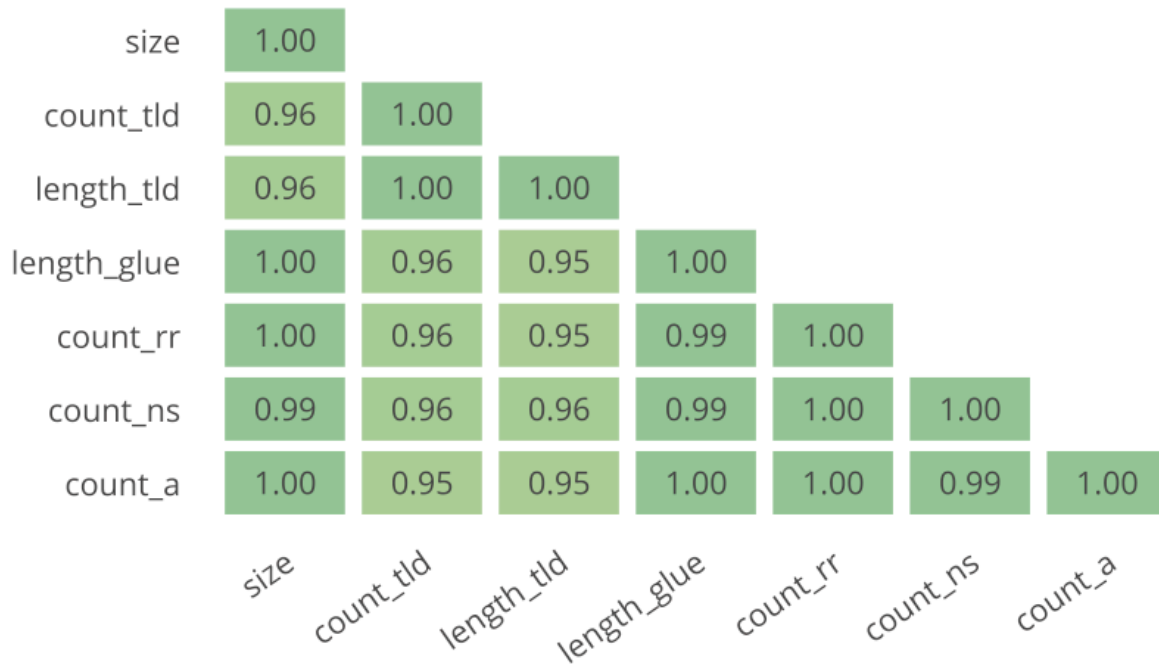


Figure 13. Correlation Coefficients of Parameters (Root Zone prior to the introduction of AAAA records).




Figure 14 – Correlation Coefficients of Parameters (Root Zone with AAA records, prior to DNSSEC).




Figure 15 - Correlation Coefficients of Parameters (Root Zone with AAAA records and DNSSEC).

A correlation coefficient close to 1.00 is indicative of a compelling linear relationship (a coefficient of 1.00 indicates perfect correlation). The three matrices indicate that the correlation between the parameters changed between each of the three periods, most significantly due to DNSSEC.

Selection of Key Parameters

The relationship of various parameters and the degree to which they are predictable (unpredictable parameters being good candidates for exposure in any model) is discussed in more detail below.

length_tld, length_glue

TLD labels (i.e. the names of the top-level domains) must be encoded into the root zone, and hence have a direct impact on zone size.

The average length of a TLD label in the root zone has not varied significantly over the measurement period, which is to be expected given that the sample population is heavily weighted towards two-character country-code TLDs. However, the 1409 unique strings applied for in the first round of the ICANN New gTLD Program⁶⁹ do not suggest that the average TLD label size will increase: quite the contrary; if all 1409 strings were added to the root zone the average TLD label size would decrease from 9.10 characters to 7.68. Given no reason to expect dramatic increases in the contribution to the size of the root zone due to changes in TLD registration policy, this parameter will be excluded from the proposed model.

It is not clear that the average length of a glue record owner name in the root zone is likely to increase. There is technical motivation for TLD operators to keep the referral response for a TLD from a root nameserver at a manageable size whilst maximizing nameserver diversity⁷⁰, and a dramatic increase in the length of nameserver names would be in conflict with this. This parameter will be excluded from the proposed model.

⁶⁹ <http://newgtlds.icann.org/en/program-status/application-results/strings-1200utc-13jun12-en>

⁷⁰ For discussion, see <http://tools.ietf.org/html/ietf-dnsop-respsize>

count_ns, count_a, count_aaaa

It has been observed that the number of NS and A records associated with a delegation have reached a steady state, whilst the number of AAAA records in the zone has had a visible impact on root zone growth, an effect which appears to be increasing.

Whilst the growth in NS and A records appears predictable, the growth in AAAA records seems likely to increase with uptake of IPv6 on the Internet in general. Since it seems unlikely that IPv4 support for TLDs will be removed in the useful horizon of this model (and hence the expected limit on the number of AAAA records is the corresponding number of A records) the ratio of AAAA:A records will be exposed in the proposed model.

count_dnskey, count_ds, count_rrsig, count_nsec

As has been noted, the number of DNSKEY records in the root zone is related to published DNSSEC key rollover schedules, and has no significant impact on the growth of the root zone. The number of NSEC records is exactly equal to the number of owner names. RRSIG records are attached to apex records, NSEC records and DS records; of these, DS records are the most unpredictable quantity.

Increased DNSSEC deployment in TLDs will result in additional DS records in the root zone. Since future trends in DNSSEC deployment are difficult to predict, the number of DS records per TLD will be exposed in the proposed model.

Proposed Model

The size of the root zone in bytes, S , is expressed as a function of the number of TLDs, T , the number of AAAA records, A and the number of DS records, D :

$$S = a + bT + cA + dD$$

For simplicity, since the component of the root zone which are invariant with growth in TLDs is small, we set $a = 0$. Multiple linear regression across the independent variables T , A and D yields:

$$S = 496.225 * T + 70.1988 * A + 125.387 * D$$

with $R^2 = 0.999996$ indicating an excellent linear correlation. We represent A in terms of the average number of AAAA records per TLD, α :

$$a = A/T$$

and D in terms of the average number of DS records per TLD, δ :

$$\delta = D/T$$

Substituting, we obtain:

$$S = 496.225 * T + 70.1988 * aT + 125.387 * \delta T$$

$$S = T(496.225 + 70.1988 * a + 125.387 * \delta)$$

We observe that, over the sample period, $\alpha = 1.12852$ and $\delta = 0.35872$. Comparing with the simple projection in **Figure 11**, which considered a root zone with 3142 delegations, we obtain:

$$S = 3142 * 496.225 + 70.1988 * 1.12852 * 3142 + 125.387 * 0.35872 * 3142 = 1949374$$

i.e. an estimated root zone size of around 1.9MB. This is comparable with the projection made in **Figure 11**, and as expected is lower since the linear model shown here does not accommodate the growth in AAAA and DS records extrapolated from the sample period in the simple projection.

To obtain an upper bound, we note that the average number of A records per TLD in the sample period (which we have noted appears to have reached a steady state) is 3.9, and set $\alpha = 3.9$ to simulate equal deployment of IPv4 and IPv6 glue records in the root zone. Similarly, we note that the root zone at the time of writing contains 142 DS records corresponding to 87 signed TLDs, and set $\delta = 142 / 87 = 1.63$ to simulate equivalent DNSSEC deployment across all TLDs. We obtain:

$$S = 3142\,496.225 + 70.1988 * 3.9 + 125.387 * 1.63 = 3061505$$

i.e. an estimated root zone size of around 2.9MB.

On the Appreciation of Size

The models presented in this document predict substantial relative growth in the size of the root zone, but it should be noted that the absolute size is still modest. A comparison between the projected root zone size in 2016 and various other downloadable objects is shown in **Figure 16**. Successive revisions of the root zone are distributed only twice per day.

Data	Size (kilobytes)
English Wikipedia Main Page ⁷¹	61
Root Zone on 2012-04-18 (measured)	199
Root Zone on 2016-01-01 (estimated)	
Numerical Model, lower bound	1,900
Simple Projection	2,600
Numerical Model, upper bound	3,000
ICANN New gTLD Applicant Guidebook ⁷²	3,885
This document ⁷³	6,495
MP3 audio file ⁷⁴	7,075
iPhone game ⁷⁵	12,940
Television episode ⁷⁶	500,643
Linux operating system install image ⁷⁷	715,700

Figure 16. Comparison of file sizes.

It is useful to consider the comparatively small size of the root zone, both current and estimated, versus the size of files that modern computers are expected to accommodate.

Conclusion

The size of the root zone appears substantially linear with growth in the number of TLDs. Based on the data available, and considering the possibility of dramatic advances in the deployment of IPv6 and DNSSEC, the size of the root zone will increase but not dramatically so, in absolute terms. Further, the rate of growth is slow, providing ample opportunity (on the scale of years, for the estimates presented here) for considered policy and operational responses to any unexpected infrastructural challenges that might arise.

⁷¹ http://en.wikipedia.org/wiki/Main_Page, transferred on 2012-04-25

⁷² PDF format, <http://newgtlds.icann.org/en/applicants/agb/guidebook-full-11jan12-en.pdf>

⁷³ Microsoft Word for Mac default file format, suffix .docx

⁷⁴ Donna Summers, "Love to Love you Baby" (4:59), MPEG-1, Layer 3, 44.1kHz, 192kbps, joint stereo

⁷⁵ Angry Birds version 2.1.0, install archive

⁷⁶ House, Season 8, Episode 17, "We Need the Eggs" (43:16), MPEG-4, low complexity, H.264, 640x352

⁷⁷ Ubuntu 12.04-beta2 desktop install image, "ubuntu-12.04-beta2-desktop-amd64.iso"

Appendix F — Case Study: Unusual Traffic Received by L-Root, June 2011

[This appendix represents a briefing that was produced during an event in June 2011, to illustrate an example of a response to a specific root zone related issue. Details in this case study may no longer be applicable.]

Executive Summary

This is an executive briefing on events observed at L-Root between 2011-06-28 and 2011-06-30 intended to provide timely information regarding an operational event.

L-Root observed a substantial increase in its usual query load between approximately 2011-06-28 1630 UTC and 2011-06-30 1200 UTC.

The unusual traffic sent to L-Root was isolated in an L-Single node located in Sydney, Australia. No other nodes appear to have received any traffic. The additional traffic was well within the capabilities of the Sydney node to handle, and there was no impact on L-Root service due to the event.

DNS Operations staff has retained full packet captures of queries received by all L-Root nodes for a significant period of the attack. Analysis will be performed on this data to attempt to better characterize the traffic, with the ultimate goal of explaining why it occurred.

Similar events are understood to have been observed by other Root Server Operators.

Summary of Events

On 2011-06-28 at 1630 UTC, L-Root began to receive an unusually high rate of queries. The queries per second (qps) received at L-Root increased by around 40,000. The normal steady-state query rate for L-Root is around 15,000 qps.

Early analysis showed that the additional queries were characterized by QTYPE A, QCLASS IN and various QNAMEs under the domain 91WW.COM. In simple terms, the unusual queries were all seeking to obtain the IPv4 addresses for various names under 91WW.COM.

It appears that all the unusual traffic was received and answered by the L-Root Sydney node, an L-Single (i.e. single-host) node deployed as part of the L-Single field trial. No other nodes appear to have received any of the queries.

Various source addresses were observed, and the majority appears to be assigned to organizations in China.

On 2011-06-30 at 0200 UTC, a full packet capture of all queries was initiated for all L-Root nodes. This data will be retained for further analysis.

ICANN DNS Operations staff monitored the performance of L-Root (in Sydney and elsewhere) throughout the event and did not observe any performance degradation. DNS Operations staff also provided early briefings of the situation to the CEO, to SSR and to the Communications team.

The unusual traffic persisted until around 2011-06-30 1200 UTC. More normal traffic levels have since resumed. DNS Operations staff continues to monitor.

Impact on L-Root

Any impact of the unusual query traffic was restricted to the Sydney L-Root node. No other nodes appear to have received any traffic, and hence clients served by those other nodes would have been entirely unaffected.

The Sydney node is capable of serving around 120,000 qps, and the additional 40,000 qps did not threaten normal production service levels for that node.




Figure 17. Graph of query rate to various L-Root locations.

The RIPE DNS Monitoring Service⁷⁸ (DNSMON) indicates that there was no substantial change in the observed ability of DNSMON probes to obtain service from L-Root during the event. The red line on the graph from probe 45 (Bangladesh) appears to be a result of local network conditions, and not an indication of L-Root service degradation.

⁷⁸ <http://dnsmon.ripe.net/dns-servmon/>




Figure 18. DNSMON Graph of unanswered queries to L-Root.

Impact on Other Root Servers

ICANN is aware that other root server operators saw substantially similar traffic levels due to similar queries during the same period that the traffic was seen on L-Root.

The RIPE NCC made a public announcement about the situation⁷⁹. It appears that this announcement was the basis of at least one news article in the technical media⁸⁰.

⁷⁹ "Increased Query Load on Root Name Servers", W. Nagele, RIPE NCC, 29 June 2011
<http://labs.ripe.net/Members/wnagele/increased-query-load-on-root-name-servers>

⁸⁰ "Key internet address server sees spike in traffic", D. Goodin, The Register, 29 June 2011
http://www.theregister.co.uk/2011/06/29/k_root_traffic_spike/



<http://www.icann.org/>