New gTLD Security and Stability Considerations

[Verisign Labs Technical Report #1130007 version 2.1]

Eric Osterweil
Verisign

Danny McPherson
Verisign

***Abstract*—The introduction of multitudes of new generic Top Level Domains (gTLDs) into the DNS (the Internet's de facto name mapping system) will have far-reaching effects. Any party concerned with the issues of privacy, trust, confidence, or the overall security of the DNS after the addition of new gTLDs (either from the consumer or the operator perspective) is implicitly depending on the Internet Corporation for Assigned Names and Numbers (ICANN) and the broader DNS community to appropriately address these issues before delegating any new gTLDs. The risk of a misstep during the process of introducing new gTLDs to the global DNS could have far-reaching and long-lasting residual implications.**

Many of the issues cataloged in this report focus on work that is currently *not* done, and should be completed before any new gTLDs can be deployed in a safe and secure manner. To both illustrate the concerns that exist and to clearly identify the rationale behind these concerns, the general areas of focus in this report are: Root Server System implications, Operational Readiness for gTLD Registries, and Risks of Name Collisions on the Internet, all of which will potentially have a considerable impact on the security and consumption of new gTLDs, as well as on the broader existing DNS ecosystem.

I. Introduction

The ecosystem surrounding the Domain Name System (DNS) [10] has been evolving since the late 1980's. During this time, it has become a multi-stakeholder system that is rich with technical, political, and operational policy complexity. Many understand the basic structure of the global name space of DNS is a simple tree structure; with a single root "zone" delegating to Top Level Domains (TLDs), and those TLDs delegating to Second Level Domains (SLDs), etc. However, the rich complexities of DNS' structure and interactions are not always deeply considered (especially in the face of those who wish to change them).

A great deal of prior work has outlined much of the nuanced (but extensive) complexities of DNS' ecosystem. Issues like Transitive Trust [16], [19], complexities that arise from adding cryptography to the DNS via the DNS Security Extensions (DNSSEC) [18], [15], the 3 R's of DNS (Registry, Registrar, and Registrant relationships) [18], the need to quantify DNS' (and DNSSEC's) operational status(es) [17], and many more. The list of these sorts of systemic complexities and dependencies betrays the relative caution that stewards of the DNS generally exercise when evolving the DNS. However, without proper precautions, processes, and safeguards being in place and fully specified, the advent of a multitude of new generic Top Level Domains (gTLDs) to the DNS root zone may give rise to security and stability issues.

Though the process of deciding whether to add new gTLDs, and the steps to outline how to add them, has been ongoing for many years, 2013 is the target year for the newest round (and largest addition of) new gTLDs to go live. In preparation, many concerns and prerequisites have been outlined. However, at the precipice of these new additions, many details, processes, questions, and concerns remain unaddressed. In this report, we outline many of the missing elements and some of the possible ramifications that could result from deploying new gTLDs, and the effects that could be felt throughout the Internet if proper caution is not exercised. Having been commissioned by the senior management at VeriSign, Inc., the goal of this *internal* report is to catalog many of the issues Verisign is currently facing related to new gTLDs and the surrounding processes, specifically from the perspective of a large registry and Internet infrastructure operator. Additionally, the aim is to raise the awareness of the various elements and subtleties of the rollout process, as well as the seriousness of issues that could prove to have significant consequences and perhaps even destabilize global operations of the DNS if not given due consideration.

The remainder of the work is organized as follows: we first motivate many of our concerns with real world examples of how unexpected complexities in the DNS ecosystem have already led to major operational challenges that were only visible after initial deployment (and which suggest changes be made slowly and cautiously), in Section II. Then, Section III discusses the non-obvious details of the operations of the DNS root zone, and how its holistic status is really the result of a federated effort, and provides discussion of the importance of measuring its status. Next, Section IV discusses the complexities that are involved in adding new gTLDs to the root zone. With these baselines, Section V introduces the complexities that arise from the DNS ecosystem's "interdisciplinary" nature. Lastly, we conclude in Section VI.

II. Systemic Complexities

The DNS is a deceptively complicated system: its namespace is a hierarchical tree structure, its data is often organized in simple zone files, its zones are served by simple name servers, and its clients recursively traverse the namespace tree, etc. However, the operational realities and complexities of today's Internet give rise to many non-obvious security and stability issues. To illustrate the nuanced, yet critical, issues that can arise from seemingly straightforward changes we look at two failure modes that were entirely surprising and arguably, unforeseeable before they occurred: i) "rollover and

die," and ii) the "poisoning" of recursive resolvers through DNS censorship.

In the early days of DNSSEC validation in 2010, an odd burst in traffic began to appear at various DNSSEC authoritative name servers. Increasing numbers of observations indicated that periodically, name servers would come under extreme load for `DNSKEY` queries from validating resolvers. Subsequent investigations identified what has come to be known as the "rollover and die" problem [9]. In this case, when validating resolvers found unverifiable `DNSKEY`s, they began iterating over the DNSSEC chain of trust (all the way up to the DNS root), until the chain was mended. As a result, even simple `DNSKEY` rollovers resulted in extremely large traffic volumes all the way up the delegation chain. This unintended behavior resulted in a way for descendant zones (or arbitrary depth in the DNS hierarchy) to prompt all validators to direct excessive query loads to all ancestors. While this behavior has been corrected in validating resolvers, it serves as a warning that the relationship between what data is being served, and what query behavior it may elicit is not always obvious.

Another example of unintended interactions was observed in conjunction with Internet censorship. For some time, it has been known that different political climates result in DNS censorship [2]. One unexpected consequence of this increasingly common behavior is that it can result in a Denial of Service (DoS) vector to name servers [1], [25], [8], [27]. What has been observed is that if a DNSSEC validating resolver is presented with a non-validatable response from a name server, it will try other name servers for the same zone, *and* it will blacklist that name server (considering it invalid). This can happen if, for example, a censoring country has middleboxes rewriting responses to queries from censored domains, and serving those to the querying resolvers. The rationale behind the resolvers' behavior (blacklisting the name servers) is to protect the DNSSEC cache from a source that seems to be sending invalid data. However, in the event that the entire set of name servers appear to be transmitting unverifiable (such as when censorship is being employed), resolvers may no longer be able to reach any information served by these authorities.

We use these two cases as just two illustrations of how unforeseen interactions must be *expected* when making changes to DNS, and that the types of complex interactions may often be *unknowable*, a priori.

III. The Root Server System

The ability of the root server system to introduce the multitudes of new gTLDs (rather than the very few that occasionally are added today), deal with the increased frequency of change, and identify any issues before they negatively impact the operation of the global Internet is paramount. The following two subsections, Section III-A (*System Monitoring and Reporting*) and Section III-B (*Automated TLD Additions to the Root*) catalog areas where action needs to be taken before any new gTLD delegations occur in order to minimize risk to the overall Internet ecosystem, as delegation of names in the root has far broader implications than simply updating the root zone file.

A. System Monitoring and Reporting

Practically since the beginning of the Domain Name System (DNS) [10] in the early 1980s, the root servers have been operated by different organizations. Today twelve organizations operate the thirteen root server IPv4 addresses, and only nine root servers have IPv6 addresses present in the root zone. Because of techniques enabled by IP anycasting, there are many more than 13 physical servers, but by convention, all instances of a given root server IP address run by the same operator are still referred to collectively as a "root server."

Each operator is responsible for only its own root server, and has complete autonomy in the design of its infrastructure, the type of hardware and software used, and the operating and monitoring procedures. The operators are not legally or contractually required to be accountable to anyone (with the exception of Verisign, which operates the "A" root server under a cooperative agreement with the U.S. Department of Commerce).

The root operators are independent organizations; there is no umbrella organization to which they all belong. They do coordinate loosely via various means, including a shared email list and in-person meetings three times per year. The operators have tried unsuccessfully multiple times to agree on even a baseline set of common principles. The only unifying and uncontentious precept is a pledge to publish the root zone administered by the Internet Assigned Numbers Authority (IANA) function.

This multi-operator arrangement makes administration of the root zone unusual compared to other critical DNS zones. We are not familiar with any other similar situation. Many DNS operators outsource responsibility for operation of some or all of their zone's authoritative servers to third-party providers. But in this case, a single operator still has oversight and control over the provider, usually through a contractual relationship. The operator establishes requirements and chooses a DNS provider accordingly. If the provider does not meet expectations, the operator can stop using that provider organization and replace them.

The dedication and technical expertise of the root operators has allowed the current unusual administrative model to largely succeed. There is one glaring exception: unified monitoring and reporting of the root server system to allow visibility into the system as a whole.

To varying degrees, each operator monitors its own root server. The operators communicate significant service changes to one another, such as planned or unplanned outages. They share information about unusual traffic seen at their servers, including suspected and confirmed attacks, and they usually report any resulting service degradation. Some operators even report the status and performance of their services publicly.

But this monitoring applies to individual root servers, not to the entire root server system. There is no mechanism currently in place to allow a detailed view of the entire

system, short of the annual "Day in the Life" (DITL) [6] data collection, which covers only a short time period and has been implemented for research purposes, not to provide overall monitoring and an operational view of system health. Such a view can only result from agreement of the operators: only individual operators have access to detailed information about their respective services, and a high-resolution picture of the system's performance can only be obtained by combining this detailed per-server data. As an example, while individual operators know how much traffic their individual servers receive, literally no one knows how much traffic the entire root server system receives.

Since at least 2009, there has been lengthy discussion within the ICANN community about the ability of the root zone to scale to accommodate new gTLDs. This discussion ultimately led to a request by the ICANN Board in late 2012 [14] affirming the need observed by the Security and Stability Advisory Committee (SSAC) for "actual measurement, monitoring, and data-sharing capability of root zone performance." [20] Separately, the Root Server System Advisory Committee (RSSAC) had earlier agreed in late 2010 [13] on the need to establish standard metrics to be collected and reported by all operators so that the entire root server system's performance could be monitored to observe any possible negative effect from the added TLDs.

Unfortunately, all the discussion about the need for coordinated monitoring and measurement of the root has not yet led to any implementation. As of this writing, RSSAC is still finalizing the specific metrics that each root operator will collect. Some, but not all, operators are collecting data based on the preliminary data collection metric specification still under discussion. A separate discussion on the RSSAC list about where the operators would send their respective data for analysis once collected did not reach a conclusion and showed that there is not universal agreement among the operators on the destination for collected metrics.

At this point it is unclear when (or even if) all the operators will be reporting metrics for analysis. This inability to view the root server system's performance as a whole presents a risk when combined with the impending delegation of the multitude of new gTLDs. There is no holistic view of the system's performance that would enable the technical community to determine the impact of added TLDs. Worse, if harmonized monitoring and reporting doesn't start sufficiently ahead of when the first new gTLDs are delegated, there will be no baseline establishing the current normal performance of the system. Without a baseline, determining whether or not future changes represent a negative trend becomes very difficult.

It would be inconceivable to begin a rapid expansion of any system currently in operation, much less one involving such critical infrastructure as the root zone, without a baseline of harmonized measurements from all root servers to establish the normal performance and behavior of the system. Then and only then can one begin to compare the baseline behavior to behavior of the root with the new gTLDs.

B. Automated TLD Additions to the Root

Three organizations cooperate to produce the root zone: ICANN, as the IANA Functions Operator, processes requests for changes to the zone. (Usually these changes come from existing TLD operators, but requests for delegation of new gTLDs will come from elsewhere within ICANN.) ICANN communicates change requests to U.S. Department of Commerce's (DoC's) National Telecommunications and Information Administration (NTIA) for authorization and Verisign, the Root Zone Maintainer, for implementation. Verisign maintains the database of record for the root zone from which the zone file is generated.

Until recently, the process to update the root zone was largely manual: the three organizations communicated changes using human-readable forms via cryptographically signed email. In July 2011, ICANN and Verisign each deployed new systems that largely automated the change process for existing TLDs. At ICANN, TLD operators can submit changes via a web interface rather than filling out an email template; ICANN's system communicates change requests using predictable semantics to Verisign's system using Extensible Provisioning Protocol (EPP), the same protocol used between many TLD registrars and registries; and DoC NTIA personnel authorize changes via a web interface rather than a signed email. Not every aspect of the system is automated, however, which was an intentional design decision. At Verisign, for example, each pending change (and a corresponding report showing the results of various validation tests relating to the change) is approved by an engineer before it is implemented in the root zone. The process keeps the best part of the manual aspect - inspection by a human - while removing the risk involved with interpreting templates and transcribing changes from email to database.

Unfortunately, not every aspect of TLD administration ultimately planned for automation is actually automated at this point. ICANN's root zone management system only processes changes for existing TLDs. The process of creating TLDs has historically used a different system, and ICANN ruled the "add TLD" process out of scope for the initial automation implementation. As a result, ICANN's system cannot use the EPP interface with Verisign's system to add a TLD. Instead, the add TLD process still uses email templates among ICANN, Verisign and DoC. Without automating the wrought portions of the process, this approach re-introduces new points for unintentional errors (often called "fat fingers") that stem from incorrect transcription. After addition of a new TLD is authorized by DoC, Verisign engineers enter the new TLD parameters from the email template into a tool that communicates via EPP to Verisign's automation system, which has always supported the add TLD function. This process is undesirable for several reasons. First and foremost, it still requires an engineer to transcribe requests. It is also manually intensive and time-consuming, introduces scale issues, and risks errors if pushed to faster speed than would allow for careful checking. Additionally, it limits the speed at which ICANN can make

changes, etc.

Verisign has been requesting that ICANN automate the add TLD process since shortly after the initial automation system was deployed in July 2011. We have expressed concern that adding new TLDs without automation introduces unnecessary risk and is resource intensive. We have been particularly concerned that new gTLDs would be ready for delegation before automation is implemented.

ICANN has recently started a development effort to implement add TLD. As of this writing, ICANN is testing their systems in Verisign's test environment, but has not communicated a firm timeline for the test process, success criteria for exiting quality assurance, nor a date for moving to production. As a result, while Verisign remains committed to not being a gating factor in the process [24], we continue to be concerned that a significant number of new gTLDs will need to be added without the data integrity and validation benefits offered by automation.

Figure 1 illustrates the growth of the number and type of objects in the root zone itself over time. Because each new TLD requires multiple associated objects to implement the delegation, the overall root zone size has a multiplicative relationship relative to the number of new gTLDs. In other words, for every new TLD, multiple records must be added to the zone. One can see the growth of new TLDs (country code and generic) over time has been fairly modest, while associated resource records have been added at varying increasing frequencies. The combinatorial effects of introducing DNSSEC [3], [5], [4], IPv6 AAAA resource records [23], Internationalized Domain Names (IDNs) [7], and new gTLDs into the root zone all within a relatively short timeframe is still not well understood, particularly aspects relating to the impact on the broader DNS ecosystem, as well as the impact on external *dependent systems*, as discussed in later sections.

IV. Operational Readiness for gTLD Registries

Rolling out the multitudes of new gTLDs and implicitly adding the associated levels of churn into the relatively stagnant DNS root zone will require a lot of process and dependencies. The smaller scale evolution that the root has traditionally experienced likely has not illuminated the process related complexities and problems that, in our considered opinion, will almost certainly abound with larger scale changes. Without a well constructed and well reasoned process model, and at the scale of changes foreseen with the addition of the unprecedented rate of the new gTLDs being added, the entire DNS hierarchy faces the potential for issues at or near the root of the DNS tree, and the fallout from such a change could affect all delegations.

From our perspective as a large-scale registry operator, it is apparent that the new gTLD timeliness provide little accommodation for the operator aspects of the process. That is, given the information to date, it seems as though ICANN has taken a very ICANN-centric role in establishing these timeliness, and gave little consideration to the fact that registry operators would need to ingest requirements, implement, specifications, pilot internally, prepare to accommodate test cases, consider the security implications and change management functions required, etc.. It actually appears as though there is little to no time allotted for operators to adequately prepare, and the lack of stable fundamental specifications and test plans within days of publicly stated testing times illustrates a clear disconnect between aspirational timelines with the New gTLD program and operational realities with which various stakeholders are constrained. The following subsections capture some areas where timelines, sequence and dependencies have been and continue to be problematic for Verisign in this area: the *Trademark Clearing House (TMCH)* will determine which registrants are entitled to register within gTLDs without a trademark infringement, *Pre-Delegation Testing (PDT)* will ensure that new gTLDs are ready to go online, *Emergency Back End Registry Operator (EBERO) + Escrow + Zone File Access (ZFA)* all ensure the "who," "what," and "how" of handling cases where a gTLD suffers a failure and goes offline, *SLA Monitoring* tracks the security and status of new gTLDs, *WhoIs* provides the access to registrant information, and it remains unclear whether *ICANN's Government Advisory Council (GAC)* determines if ICANN is even *allowed* to deploy any given gTLD.

A. Trademark Clearinghouse (TMCH)

The Trademark Clearing House (TMCH) is an important rights protection mechanism for the introduction of New gTLD Registries. The TMCH is comprised of two primary functions:

1) A repository of marks submitted by trademark holders. The provider of this service receives and validates the marks. ICANN selected Deloitte Enterprise Risk Services as the trademark validation provider for the TMCH. On February 25, 2013, ICANN announced that this system would be available on March 26, 2013.
2) A centralized database of validated marks. This database includes provisioning relevant data to new gTLD registries and registrars that are required for Sunrise [26] and Trademark Claims services as part of new gTLD launches. This database is critical in order to avoid conflicts that may arise when multiple parties claim to have jurisdiction over the same gTLD identifier. ICANN is working with IBM and has announced that these systems are anticipated to be operational later in 2013.

New gTLD Registries cannot begin sunrise when there is no TMCH, thereby presenting a number of concerns:

1) As late as Q4 2012, the proposed date for the TMCH system to be live was January 2013. The February 25th notice does not define the date that registries and registrars can expect the systems to be available, including both the centralized database and the repository of marks.
2) The specifications for registries and registrars have not been finalized, resulting in a lack of information for these organizations to prepare their systems. The initial draft specifications (version 00) that define the mapping

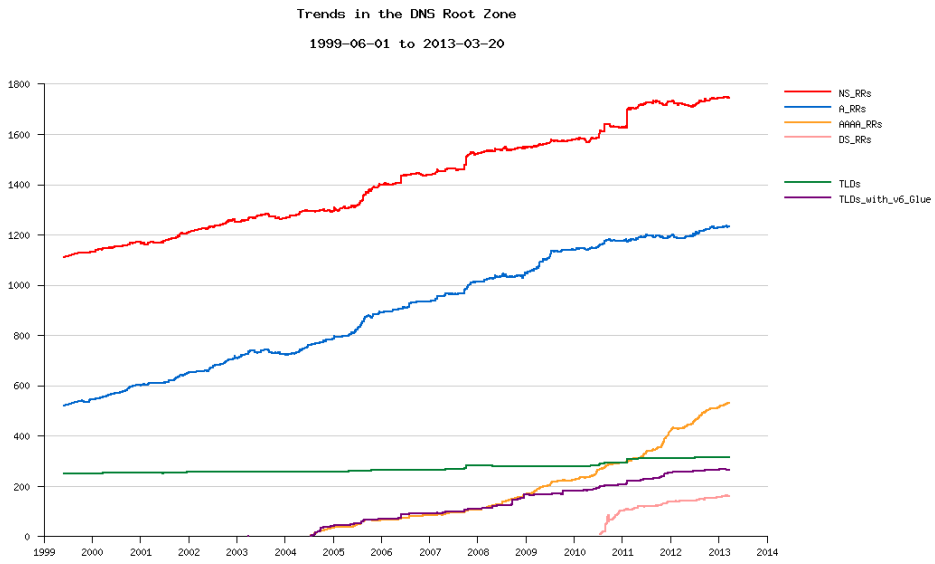

Fig. 1. Trends in Root Zone Growth

for marks and signed mark data were posted by ICANN only on February 25, 2013. As expected with any initial draft, this has generated discussion and many recommended modifications. Therefore, the current Internet draft does not represent a stable specification that can be built upon.

3) The absence of any clear commitment for when the specifications and the TMCH will be available for integration testing and when it will be live precludes organizations from proceeding with planning, scheduling, development and normal business planning and associated communications with customers. Further, a compressed testing time frame will jeopardize the stability and scalability of the TMCH system as ICANN will not have an opportunity to validate the system against real world issues like DDOS attacks and inconsistent global network connectivity.

B. Pre-Delegation Testing (PDT)

Pre-delegation testing is required prior to the introduction of all new gTLDs. ICANN issued an RFP to solicit for testing providers in November 2012 and selected Stiftelsen för Internetinfrastruktur (.se) registry in December. ICANN is currently developing a pilot for selected registries to participate in pilot testing during March 21 through April 5. There are a number of outstanding issues related to the PDT, as were outlined in detail by the Registry Stakeholder Group (RySG), as of late. These issues include, but are not limited to insufficient documentation or specifications from ICANN and the PDT provider, related to the following:

- Test Plan Documentation
- Load Capacity Testing Information
- Reachability Documentation
- TCP and DNSSEC Capabilities
- Testing Against Existing Infrastructure
- IDN Table Testing
- Trademark Clearinghouse

C. Emergency Back End Registry Operator (EBERO)

ICANN identified five critical functions required of all new gTLD registry operators, so that should a registry fail, an Emergency Back End Registry Operator (EBERO) provider would assume these functions. ICANN's original plan was to select EBERO providers in June 2012 with simulations and drills in January and February 2013 and providers prepared to be live in March. While ICANN issued an RFI in July 2012 and developed a short list in December 2012, there are currently no signed/tested EBERO providers, to our knowledge.

With no EBERO providers, the continuity risk of any registry failure is significantly increased. The original plan allotted EBERO providers with six months from selection to preparation for testing. Therefore, it is reasonable to conclude that from the date EBERO providers are selected that ICANN should expect that they will require six months to prepare for testing with another two months of testing prior to being prepared for a registry failure. Should any registries launch prior to EBERO provider readiness in the event of a registry failure or outage, the continuity of operations risk could be measured in days, weeks, or even months, rather than the hours specified in the EBERO service level requirements.

D. Escrow

All gTLD registry operators are required to enter into an escrow agreement and make incremental and full deposits of registry data on a prescribed schedule. The availability of escrow is an essential component for Emergency Back End

Registry Operators (EBRO) in the event of a registry failure in order to provide the critical functions and requirements defined in the Applicant Guidebook, noting that escrow is one of these critical functions. The Application Guidebook references draft escrow specifications expressly. There are currently two draft specifications (XML and CSV) that are being coordinated to merge into a single specification, although the specifications are not stable references at this time.

The lack of an approved, stable, final specification will require registry operators to either select a draft specification for implementation, or will result in delaying the launch of any gTLDs until a final specification is defined and implemented by the registry operators. Should various registries select from available or future draft specifications, an EBERO provider's timely and accurate data recovery during an emergency situation is significantly an increased risk, an observation that should certainly be considered as critical as new gTLD operations are being brought online.

E. Zone File Access

New gTLD registries are required to provide third-party access to zone file data. Registry operators are required to enter into an agreement that will be standardized, facilitated and administered by a Centralized Zone Data Access (CZDA) provider. The community has not been informed of any progress on soliciting for a CZDA provider or of any progress on creation of the standardized agreement or process for registry operators to provide users with access to zone file data.

Registry operators will not be able to comply with contractual requirements to provide zone file data in the manner specified and required in the gTLD Applicant Guidebook; therefore, third parties will not have access to zone file data for new gTLDs. Additionally, EBERO providers will not have access to zone file data because ICANN has not specified EBERO providers, nor the manner with which providers once designated will access zone files.

F. SLA Monitoring

Monitoring of registry performance specifications for new gTLDs relies on the implementation of external probes operated by ICANN. However, these systems, methodologies, and processes have not yet been defined.

Without monitoring in place, ICANN will have no ability to detect, understand if, or evaluate and identify if a registry encroaches on any emergency thresholds for critical functions, nor will they have visibility into performance against specific service levels. ICANN has had the right to perform similar monitoring under the Cross Network Nameserver Performance (CNNP) specifications in existing gTLDs since the first new gTLDs were introduced more than a decade ago, yet ICANN has never established this monitoring and measurement capability. ICANN has engaged in significant discussions on the Continuity of Operations Instrument yet has not established that the operational mechanisms themselves are ready for the launch of new gTLDs.

G. WhoIs Change Requirements

New standards and or requirements may need to be implemented for models such as RESTful interfaces to WhoIs, even though specifications are still unstable.

H. Government Advisory Council (GAC) Advice

The requirement for GAC advice delegations, with possible separate EU advice, could delay timelines and/or impact applications. Or worse, it is unclear if it would be possible that a TLD could launch, then receive GAC advice – would ICANN withdraw the delegation? Or might some countries begin to block certain TLDs causing fragmentation of the Internet? The uncertainty in this area is of considerable concern to applicants, operators, and consumers alike.

V. Risks of Name Collisions on the Internet

Since the inception of the Internet DNS it has been known that interactive complexities can result in large scale security and stability issues as well as hard to diagnose corner cases where consumer expectations are unaddressed or users are provided an unsafe or otherwise less than desirable experience. Express interdisciplinary studies are paramount to proactive mitigation of issues with which the DNS interacts. In this section we catalog several such issues, provide general discussion related to usability of new gTLDs, as well as consider the current status of some requests that have been made of ICANN to conduct broader studies in this area.

A. Internal Names Certificates

On March 15th 2013, the ICANN Security and Stability Advisory Committee (SSAC) published an advisory titled SSAC Advisory on Internal Name Certificates (SAC057) [21]. The advisory identified a Certificate Authority (CA) practice that, if widely exploited, "could pose a significant risk to the privacy and integrity of secure Internet communications." As conveyed from ICANN to the "New gTLD Communications" email distribution list on March 15, 2013, just after the advisory was made public, it is believed that this CA practice could impact the New gTLD Program. Appendix A of the advisory provides a timeline that suggests the issue was brought to the attention of ICANN via the SSAC mid-November 2012 at the annual SSAC workshop, and upon recognizing the seriousness of this issue, the SSAC formed a work party to develop advice for ICANN and the community. That effort led to discussions with the ICANN Security Team, the Certification Authority/Browser (CA/B) Forum, and working party representatives from SSAC, the product of which is reflected in the Advisory.

The following 5 Findings were captured in [21] and provide an overview of the Internal Names Certificates issue that has given reason for security and stability concerns, particularly when intersecting the implications with the delegation and operation of new gTLDs:

> Finding 1: The SSL observatory data shows that at least 157 CAs have issued internal name certificates.
>
> Finding 2: The exact number of internal name certificates

> that end in an applied for new gTLD cannot be known unless CAs voluntarily disclose the list.
>
> Finding 3: Enterprises use internal name certificates for a variety of reasons.
>
> Finding 4: The practice for issuing internal name certificates allows a person, not related to an applied for TLD, to obtain a certificate for the TLD with little or no validation, and launch a man-in-the-middle attack more effectively.
>
> Finding 5: The CA / Browser (CA/B) forum is aware of this issue and requests its members to stop this practice by October 2016. The vulnerability window to new gTLDs is at least 3 years.

The empirical data provided in the report was from 2010 datasets, and as indicated in the report and discussed in Finding 2, only represents a lower bound observation of Internal Name Certificates. That is, given that these certificates are presumably allocated for internal-only use (i.e., not visible on Internet-facing servers), it is reasonable to assume that the number of issued certificates is much larger, and without a complete corpus across all Certification Authorities, it would be impossible to ascertain the potential impact with the delegation of any new gTLD.

The Recommendations section of [21] follows. [Note: numbered bullets inserted, and original footnote pointers removed]:

> Recommendation: The ICANN Security Team should immediately develop and execute a risk mitigation plan.
> The mitigation plan should include at least:
>
> 1) Outreach to the CA/B forum and CAs, requesting that they treat applied for new gTLDs as if they were delegated TLDs as soon as possible, as well as discussing the broader implications and mitigation steps. In doing so, ICANN should seek to create trust relationships between ICANN and CA/B Forum and CAs. Because of the potential for collateral harm to users if disclosure is made public before mitigation is effected, the SSAC believes it is important to conduct correspondence confidentially.
> 2) A Disclosure Policy as informed by industry best practices for vulnerability disclosure (e.g. CERT / CC vulnerability disclosure). Such a policy should take into consideration that once the disclosure is public, it is trivial to exploit the vulnerability.
> 3) A communication plan on informing affected parties as determined by the disclosure policy.
> 4) A contingency plan to be executed if the vulnerability is leaked to the public prematurely, as well as a proactive vulnerability disclosure plan.

Regarding the first recommendation, as a result of the SSAC efforts and subsequent dialog between the ICANN Security Team and the CA/B Forum, efforts to remedy at least some of this vulnerability have led to some marked improvements in minimizing the potential exposure window of Internal Name Certificates. Specifically, as conveyed in the detailed Findings information in [21], the initial vulnerability window for new gTLDs may have persisted for three years or more, via the product of new "Baseline Requirements for the Issuance and Management of Publicly Trusted Certificates" from the CA/B Forum that went into effect July 1, 2012, and is as follows:

> As of the Effective Date [1 July 2012] of these Requirements, prior to the issuance of a Certificate with a subjectAlternativeName extension or Subject commonName field containing a Reserved IP Address or Internal Server Name, the CA SHALL notify the Applicant that the use of such Certificates has been deprecated by the CA / Browser Forum and that the practice will be eliminated by October 2016. Also as of the Effective Date [1 July 2012], the CA SHALL NOT issue a certificate with an Expiry Date later than 1 November 2015 with a subjectAlternativeName extension or Subject commonName field containing a Reserved IP Address or Internal Server Name. Effective 1 October 2016, CAs SHALL revoke all unexpired Certificates whose subjectAlternativeName extension or Subject commonName field contains a Reserved IP Address or Internal Server Name.

As highlighted in [21], while that was indeed welcome news, it was problematic because ICANN plans to delegate new TLDs in 2013, introducing vulnerability for potential new gTLDs until October 2016. Further dialog with the CA/B Forum led to CA/B Forum advanced Ballot 96 on new gTLDs. The ballot called for CAs to stop issuing certificates that end in an applied-for-gTLD string within 30 days of ICANN signing the contract with the registry operator, and to revoke any existing certificates within 120 days of ICANN signing the contract with the registry operator. The full text of the ballot is included as Appendix B of [21]. On February 20, 2013, the CA/B Forum passed Ballot 96 (Wildcard certificates and new gTLDs). March 12, 2013, the SSAC finalized its advisory based on the mitigations and additional input provided by the Certificate Authority Security Council, and published [21] shortly thereafter.

This effort shortened the exposure window considerably, but residual vulnerabilities persist with new gTLDs for up to 120 days after ICANN signs a contract with a registry operator, and even then, only where revocation capabilities are enabled. Additionally, residual vulnerabilities may exist even longer in closed environments (e.g., air gapped) or with other instantiations where relying party software cannot or is otherwise provisioned to not support Certificate Revocation List (CRL) and/or Online Certificate Status Protocol (OCSP) certificate validation status checking functions. Unfortunately, one could imagine scenarios where attackers could leverage the 30-day issuance window to initiate and exploit certificate issuance in order to attack relying parties for up to 120 days after the contract is signed, or perhaps even to specifically target specific relying party software or organizations that do not download CRLs or support OCSP (e.g., browsers that disable checking in order to minimize latency, enterprises that block downloads or filter OCSP transactions, etc.).

ICANN also needs to develop a process and capability to notify CAs as soon as these contracts are signed, as that infrastructure does not exist yet, CAs can't be subscribed to such a notification capability as of yet. Once this *contract registry* has been established and the format specification made public, CAs will need to automate ingestion and certificate revocation functions from ICANN based on the registry contract signing dates provided therein, and notify their own customers under this framework that those corresponding certificates must be and/or have been revoked. It would seem natural that CAs would not revoke any certificates until the final day of the 120-day window, as under the ballot provisions provided they very well may have just issued the certificate (and collected associated monies), or intend to issue it within the upcoming 30 day window. Note that if the above contract registry exists and is made publicly available, it should be expected that miscreants will surely exploit the information as well, and the current model could actually lead to "runs" on certificates related to newly delegated namespace.

An astute reader might surmise that it may well be advised to delay delegation of a new gTLD until at least 120 days has passed since ICANN has signed the contract and compliant CAs have revoked any issued certificates that relate to that namespace, although as noted above, this would only address a portion of the residual vulnerability.

Another consideration is the current wording the CA/B Forum ballot referenced in the [21] advisory:

> If a wildcard would fall within the label immediately to the left of a registry-controlled or public suffix, CAs MUST refuse issuance unless the applicant proves its rightful control of the entire Domain Namespace. (e.g. CAs MUST NOT issue "*.co.uk" or "*.local", but MAY issue "*.example.com" to Example Co.).

Specifically, the text leaves some ambiguities for CAs to issue certificates for new gTLD holders themselves, who could then exploit the vulnerabilities for child domains within the gTLD. This could potentially result in security and privacy issues for Internet users, and fosters the sort of interdisciplinary concerns that have been raised for several years. "Rightful control" can be hard to prove because there is no general accepted way of communicating which specific entities have rights to which portions of the domain namespace, similar to the problem with "zone cuts" described next.

Regarding recommendations 2 through 4 of SAC057, on March 11, 2013 ICANN published the first Coordinated Vulnerability Disclosure Reporting at ICANN practice statement, describing the basic principles of Coordinated Vulnerability Disclosure Reporting as practiced by ICANN.

B. Suffix Lists, Cookies, and Administrative Boundaries in the DNS

The in-bailiwick scope of cookies is already an acknowledged threat to Internet browsing, consumer privacy, and Internet security in general. A cookie for example.com could be used to affect subzone.example.com. The way this is addressed currently is an annotated *trusted* list (or repository) of zone cuts (administrative boundaries) [12] that browsers fetch in order to disambiguate where zone cuts exist, and therefore, cookie authority ends. With the addition of new gTLDs, a number of problems arise.

The process of mapping DNS' evolving zone cuts into a standard file that is then fetched and ingested by browsers is itself architecturally ill advised because, much like the issues "certificate pinning" introduces in browsers, it codifies DNS' dynamic structure in an offline medium. The result is that the representation of DNS contained there will (likely) always tend towards being out of date. No matter how fast the DNS is polled and then flattened into a file, zone cuts may well have just changed.

In addition to the primitives associated with authority, trust, and creation of the zone cut file itself, there exists the problem of browsers fetching and maintaining it. The primary mechanisms that are in use today have not been cataloged or measured, and as such, it is not clear how well they are working. Interestingly, at the time of this writing, we observe that the Mozilla repository is currently unprotected (i.e., it lacks both HTTPS and certification credentials from a CA). However, by inspection we can infer that as the mapping file grows and churn increases at the highest levels of the namespace hierarchy (i.e., new gTLDs) it becomes more likely to be stale, and the process of ingesting it into all browsers and other applications inherits this staleness, and adds its own. The larger size file will cause greater resources to be needed on the server side, and slower synchronization on the browser side. While this further enshrines the heightened complexity of the browser and broader Internet application ecosystem, it also opens a potential vulnerability and variability window.

Also, with the fan-out that new gTLDs will prompt, many existing zones are likely to create doppelgangers under new gTLDs. That is, example.com will likely seek example.newgtld, etc. This will lead to a greater corpus for browsers to poll and serialize, and will tend to increase the likelihood of zone cuts being unsynchronized with browser mapping files. Some work [22], albeit under considerable debate, has been proposed in various forums to address one aspects of this problem, but there are many other interacting elements that need to be considered and operationally integrated in order to mitigate the risks from the delegation of any new gTLDs.

The impact related to the introduction of new gTLDs may well result in insecurity for widely deployed web security and cookie mechanisms, and potentially contribute to the loss of consumer confidence and erosion of consumer privacy and trust in the Internet, as well as broader security and stability implications yet unforeseen. It is critical that work on interactive complexity and interactions of the proposed delegation and operation of new gTLDs with mainstream fixed and mobile devices, and applications such as web browsers and smart-phone applications, security software and hardware, and broader infrastructure, be examined in short order.

C. Interactive Complexity and Interdisciplinary Studies

The implications of new gTLDs on external systems has been a concern of the ICANN SSAC and much more broadly in the community since new gTLDs were first proposed. As a matter of fact, the introduction of .info over a decade ago highlighted just what sort of obvious and nuanced interdependencies may exist as new gTLDs are delegated and made available on the Internet while applications and other systems are ill-prepared. As provided in the previous two sections, Internal Name Certificates and Administrative Boundaries pinned to DNS zone cuts provide opportunity for Internet usability, safety and security issues to arise.

The proper functioning and expectations of new gTLD holders, and the consumer confidence, trust, and security associated with services residing under those new gTLDs is clearly paramount to the success of the New gTLD Program. The following quoted text contains an excerpt from a Security Week article [11] authored by Ram Mohan in August 2012, and captures the essence of some issues to be considered with the introduction of new gTLDs:

> **Getting to Universal Acceptance of All Domains**
>
> Ensuring that Internet software and sites understand all domains – not just the old three-letter domains like .COM and .NET – is called universal acceptance. For example, when is the last time you looked at your company's online contact forms? If you haven't revisited them for a while, you might discover that they are hard coded for certain domains like .NET or .ORG and may reject email addresses that use, say, a four-or-more-character domain like .INFO or .MOBI. (Full disclosure: .INFO and .MOBI are both domains managed by my company, Afilias).
>
> Or have you seen some TLDs that don't work in your browser? Some browsers, including mobile ones, screen out addresses as either "right" or "wrong," and many modern TLDs simply don't resolve because the browser doesn't understand how to handle the TLD.
>
> A real-life example: as late as 2007, you could not email an article from the New York Times website to anybody with a .INFO email address, which was actually fun for some of my colleagues because they would try to send me articles and say, "Oh, you didn't see it? Maybe you should get a .COM address."
>
> Some software and websites contain restraints that limit the scope of what is considered a valid domain name and, in the process, impose artificial – and many times, unintended – boundaries on the emails and websites that will (or won't) be accepted. Some of these universal acceptance issues are caused by improper logic in software for checking valid domains or older software that requires an upgrade.
>
> Other universal acceptance issues can be caused by a lack of support for Internationalized Domain Names (IDNs), which are domain labels that incorporate accent marks or non-ASCII characters like Chinese, Hindi and Hebrew. IDNs are six percent of all the new gTLDs currently applied for and, given that they haven't been widely used until now, it's likely that many websites and Internet applications won't recognize them as top-level domains. ICANN has been discussing the issue of universal acceptance for years. And at my own company, Afilias, we had the "fun" in 2001 of launching the first four-letter TLD – .INFO – that wasn't accepted practically anywhere.
>
> From that experience, I developed my three "rules" of TLD acceptance:
>
> 1) **An old TLD will be accepted more often than a new TLD.**
> 2) **An ASCII-only TLD will be accepted more than an IDN TLD.**
> 3) **A three-letter gTLD will be accepted more often than a longer string, even if it's a gTLD.**
>
> Universal TLD acceptance must overcome these three rules. And with the upcoming addition of hundreds of new TLDs and IDNs, the problems that result due to a lack of universal TLD acceptance have quickly moved to the front burner as a global issue.

Related to the notion of universal acceptance, the SSAC made 5 recommendations [20] to the ICANN Board related to root scaling and new gTLDs:

> Recommendation(1): Formalize and publicly document the interactions between ICANN and the root server operators with respect to root zone scaling. ICANN and the root server operators may choose to utilize RSSAC to facilitate this interaction.
>
> Recommendation(2): ICANN, National Telecommunications and Information Administration (NTIA), and VeriSign should publish statements, or a joint statement, that they are materially prepared for the proposed changes.
>
> Recommendation(3): ICANN should publish estimates of expected and maximum growth rates of TLDs, including IDNs and their variants, and solicit public feedback on these estimates, with the end goal of being as transparent as possible about the justification for these estimates.
>
> Recommendation(4): ICANN should update its "Plan for Enhancing Internet Security, Stability, and Resiliency," to include actual measurement, monitoring, and data-sharing capability of root zone performance, in cooperation with RSSAC and other root zone management participants to define the specific measurements, monitoring, and data sharing framework.
>
> Recommendation(5): ICANN should commission and incent interdisciplinary studies of security and stability implications from expanding the root zone more than an order of magnitude, particularly for enterprises and other user communities who may implement strong assumptions about the number of TLDs or use local TLDs that may conflict with future allocations.

As discussed in the Root Server System section earlier in the document, these five recommendations were provided to the

ICANN Board in 2010. The first and fourth recommendation have not been accommodated as of yet, whereas Recommendations 2 and 3 have been accommodated but with varying levels of fanfare. Recommendation 5, various aspects of which are reflected in some of the text earlier in this section, and some of which are reflected in the advisory released just last week, has resulted in numerous letters exchanged over the past three years between the ICANN Board and the SSAC, albeit with little to no actual action by the ICANN Board to ameliorate the issues provided therein, much to the frustration of some members of the SSAC. To further illustrate this point yet another letter is being drafted by an SSAC work party addressed to the ICANN Board as of this writing, the crux of which is that the SSAC believes that the community would benefit from further inquiry into lingering issues related to the expansion of the root zone as a consequence of the new gTLD program. Furthermore, issues that previous public comment periods have suggested were inadequately explored as well as issues related to cross-functional interactions of the changes brought about by root zone growth should be investigated, and that additional perspective regarding stubbornly unresolved concerns about the longer-term management of the expanded root zone and related systems would be prudent.

ICANN needs to seriously consider these issues before new gTLDs are operationalized, as the safety and security of Internet users, and the infrastructure itself, is at risk.

VI. Conclusion

The recent disclosure of issues with Internal Name Certificates, and the broader implications of new gTLDs to parties that rely on the Internet DNS, will be far-reaching if these issues are not addressed by ICANN in a timely manner. Addressing these issues doesn't simply mean publishing a specification and expecting the community to have immediately implemented it and be capable of responding to all operational and security corner cases conveyed therein. It means working with the community in attempts to identify these issues before problems arise in operational systems. It also means that adequate buffers should exist in ICANN published timelines that account for implementation, internal testing, security auditing and vulnerability testing, pilots and early field trials, and deliberate transition to operations; it's apparent little consideration has been given to this in the current timelines published by ICANN. In order to ensure a successful implementation of each new gTLD, it is essential that proper planning be conducted in advance. This entails the development of a project plan (to include: a detailed schedule, communications plan, risk management plan, and deployment plan) for each new gTLD to be implemented. These plans should align with ICANNs timelines, thus minimizing impacts to current registry operations, as well as the overall DNS and broader Internet ecosystem.

Any party concerned with consumer and operator privacy, trust, confidence, and overall security of new gTLDs and the broader Internet would be well served by the ICANN Board addressing these issues as appropriate before delegating any new gTLDs, as the risk of a misstep in this direction could have far-reaching and long-lasting residual implications.

References

[1] 2010 REPORT TO CONGRESS, China Economic and Security Review Commission, November 2010. http://origin.www.uscc.gov/sites/default/files/annual_reports/2010-Report-to-Congress.pdf.

[2] Anonymous. The collateral damage of internet censorship by dns injection. *SIGCOMM Comput. Commun. Rev.*, 42(3):21–27, June 2012.

[3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirement. RFC 4033, March 2005.

[4] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Protocol Modifications for the DNS Security Extensions. RFC 4035, March 2005.

[5] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Resource Records for the DNS Security Extensions. RFC 4034, March 2005.

[6] The Cooperative Association for Internet Data Analysis (CAIDA). A Day in the Life of the Internet (DITL). http://www.caida.org/projects/ditl/.

[7] J. Klensin. Internationalized Domain Names in Applications (IDNA): Protocol. RFC 5891, August 2010.

[8] Danny McPherson. Uprooting the DNS Root, May 2008. http://www.circleid.com/posts/852211_uprooting_the_dns_root/.

[9] George Michaelson, Patrik Wallstrm, Roy Arends, and Geoff Huston. Roll over and die? http://www.potaroo.net/ispcol/2010-02/rollover.html, February 2010.

[10] P. Mockapetris and K. J. Dunlap. Development of the domain name system. In *SIGCOMM '88*, 1988.

[11] Ram Mohan. Will New Top Level Domains (TLDs) Work Everywhere? In *Security Week*, August 2012. http://www.securityweek.com/will-new-top-level-domains-tlds-work-everywhere.

[12] Mozilla. Mozilla Zone Cuts / Cookie File. http://mxr.mozilla.org/mozilla-central/source/netwerk/dns/effective_tld_names.dat?raw=1.

[13] Jun Murai. RSSAC Comments on root scaling, November 2010. http://www.icann.org/en/news/correspondence/murai-to-board-25nov10-en.pdf.

[14] ICANN Board of Directors. Preliminary Report — Regular Meeting of the ICANN Board, September 2012. http://www.icann.org/en/groups/board/documents/prelim-report-13sep12-en.htm.

[15] Eric Osterweil, Daniel Massey, and Lixia Zhang. Observations from the DNSSEC Deployment. In *The 3rd workshop on Secure Network Protocols (NPSec)*, 2007.

[16] Eric Osterweil, Danny McPherson, and Lixia Zhang. Operational Implications of the DNS Control Plane. *IEEE Reliability Society Newsletter*, May 2011.

[17] Eric Osterweil, Michael Ryan, Dan Massey, and Lixia Zhang. Quantifying the Operational Status of the DNSSEC Deployment. In *IMC '08*, 2008.

[18] Eric Osterweil and Lixia Zhang. Interadministrative challenges in managing dnskeys. *IEEE Security and Privacy*, 7(5):44–51, 2009.

[19] Venugopalan Ramasubramanian and Emin Gün Sirer. Perils of transitive trust in the domain name system. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, IMC '05, pages 35–35, Berkeley, CA, USA, 2005. USENIX Association.

[20] ICANN Security and Stability Advisory Committee (SSAC). SAC 046 Report of the Security and Stability Advisory Committee on Root Scaling, December 2010. http://www.icann.org/en/groups/ssac/documents/sac-046-en.pdf.

[21] ICANN Security and Stability Advisory Committee (SSAC). SSAC Advisory on Internal Name Certificates, March 2013. http://www.icann.org/en/groups/ssac/documents/sac-057-en.pdf.

[22] A. Sullivan. Asserting DNS Administrative Boundaries Within DNS Zones, institution = IETF, year = 2012, type = Draft, month = October. Technical report.

[23] S. Thomson, C. Huitema, V. Ksinant, and M. Souissi. DNS Extensions to Support IP Version 6. RFC 3596, October 2003.

[24] Verisign, NTIA, and ICANN. Joint Statement from the Root Zone Partners, November 2012. http://www.icann.org/en/news/correspondence/icann-et-al-to-icann-board-ssac-05nov12-en.pdf.

[25] Duane Wessels. How do validating caches deal with spoofed responses?, October 2010. https://www.dns-oarc.net/files/workshop-201010/Duane-OARC.pdf.

[26] ICANN Wiki. Sunrise Period, December 2011. http://icannwiki.com/index.php/Sunrise_Period.

[27] Earl Zmijewski. Identity Theft Hits the Root Name Servers, May 2008. http://www.renesys.com/blog/2008/05/identity_theft_hits_the_root_n_1.shtml/.

