

# Security and Stability of the DNS Topics

The following are a set of general topic areas to underpin the study of this question. The question is predicated by scoping this topic to matters that fall with ICANN's remit. The broader consideration of the security and stability of the DNS ecosystem, its architecture, technologies, applicable standards, vendors, operating practices and actors is a topic that intersect with ICANN's mission in a number of areas, but the general topics are substantially larger matters and will not be considered as part of this review[1].

The focal point here is those activities that are directed or coordinated by ICANN or the PTI, or where ICANN has a substantive presence by virtue of hosting the community-driven policy process or through existing activities.[2]

## 1. Root Zone Management Practices

ICANN coordinates the contents of the root zone of the public DNS system, so consideration of the security aspects of the way in which this responsibility is ~~discharged~~<sup>fulfilled</sup> is relevant to this topic. Sub-topics include:

### TLD Label management (what labels go in the root zone)[3]

- What guidelines and constraints govern the labels that are placed into the root zone of the DNS? For example, are single character domains, either in Ascii or in their Unicode equivalent permitted? Are two letter codes other than those defined through ISO3166 permitted?  
This report should either enumerate these constraints, or preferably reference such a document that the community use as a reference document as to the permissible labels in the DNS. If such a document does not exist then perhaps it should and we might recommend as such.
- How are these constraints managed by ICANN, and how are the constraints communicated to the community?  
For example, are the procedures used by ISO3166 in the management of two letter country codes binding on the DNS? What procedures are in place to ensure that retirement of a two letter code from the ISO 3166 registry is coordinated with the retirement of the corresponding two letter code in the DNS. If these procedures are not synchronised what measures have been taken by ICANN to ensure that the two name spaces do not diverge?
  - We may wish to pass this as a query to ICANN staff, and record our response to their answers to this question.
- How is change control exercised over these constraints? Presumably, one could envisage a scenario where some group is wanting to change these constraints. What process would ICANN follow to examine such a request for a change in these constraints?  
For example, is ICANN bound to exclusively follow the specifications as described in current RFCs? Under what circumstances, if any, might ICANN reach a conclusion to deviate from the RFCs? If it were to do so how would ICANN codify its adopted specifications?
- What extent impacts and behavior relate to labels that will be, or are, in the Root zone, how can this be longitudinally considered?  
This issue is exemplified by the name collisions work [Verisign TR, S&P publication, US-CERT TA]  
The domain name space is a space of all possible labels that could be described by the syntax roles of a "valid" domain name. This space forms a far larger set than the set of delegated names. As names are added to the root zone of the DNS, such names are being drawn from a notional 'unused' domain name pool. At the time time there are other parties also drawing names from the same name pool for their own use. Such actions include the registration of labels in the IETF's Special Use Names registry, the use of names in contexts that are not the DNS (such as multicast DNS, or DNS resolution based on the ToR network, and similar), or the use of names in private network realms?  
[We should perhaps reference some recent SDSAC studies on this topic]
- If a proposed TLD contains non-ascii unicode characters (IDN) what procedures are followed to ensure that the label meets these criteria?

Historically, the contents of the TLDs listed in the Root Zone have been constrained by technical specifications developed by the IETF and published as RFCs. RFC1034 described the "LDH" rule, of ascii characters that are permissible in labels in the root zone, and the equivalence of upper and lower case. Punycode is an encoding of any unicode character into a LDH-constrained label. The implication is that using a very strict literal interpretation of this LDH rule could allow any Unicode character to be permitted into a DNS label, including dashes and characters that display as period characters, as well as punctuation characters and diacritics.

What approach is used to ensure that the Unicode characters encoded in DNS labels also conform to some unicode equivalent of the LDH rule? Or is the LDH rule only relevant for the ascii character set?

What studies and reference material is ICANN using to provide appropriate guidance to the community on the use of unicode characters in the DNS?

This question refers to the matters raised in SSAC 084 and the EPSRP?

## NS and DS record management in the Root Zone

When a name is delegated in the root zone, the delegation is reflected by the presence of NS records in the root zone, and the DNSSEC security binding is reflected by the presence of DS records in the root zone.

- Are appropriate security practices used to ensure that changes are duly authorised by the correct party prior to inclusion onto the root zone?

• ICANN (or perhaps PTI) staff should be in a position to provide an answer to this question which this review can then comment

- Are the NS and DS records validated by ICANN (or PTI) prior to inclusion in the root zone? What steps are taken if validation fails?

If PTI is not in a position to validate a registered DNSSEC crypto algorithm, and hence unable to validate a DS record signed with the algorithm, should it accept this DS record without validation or not?

• So all NS and DS records be validated by PTI all of the time? What procedures are used to audit the root zone and what steps are taken if any of the entries fail this audit?

[4]

- Are these records, and the associated glue records, regularly audited to ensure their continuing accuracy?

## Respective roles of RSSAC and ICANN

The Root Server operators serve an authoritative copy of the current root zone contents from their secondary servers. While the content of the root zone is the same across all root servers, these server operators have considerable latitude as to how the root zone is served in terms of finer level of granularity of the technical aspects of the service.

- How consistent are the contents across all Root servers (letters and instances)?
- How much lag is there in attaining consistency after a change?
- Are there any instances of sustained inconsistency or variation, and how would this be detected?
- Is it appropriate that such variation exists across the Root Servers?

The finer level of detail of how each root server responds to queries are largely undocumented, and it is unclear if this diversity enhances or compromises the security and resilience of the root service.

It is also unclear how this question fits into the SSR2 study. Is this a RSSAC Review question? A topic for SSAC study? Is this a detail of a more generic question for SSR2 study about the nature of the relationship between RSSAC and ICANN and the ACs and SOs?

It may be worth some comment at a generic level with respect to this, and the next topic about the difference between devolved and centralised systems. The risks of a devolved system is that the components may not provide uniform responses from every element of the devolved system, and this may cause issues for users. On the other hand the devolved system allows each component to use different approaches to some

common challenges, and these different approaches may ensure that a persistent client can get the answer they are seeking if they are persistent in querying across all elements of the system.

The other area of comment here is that the modes of community engagement differ between the two bodies. RSSAC tend to operate with little in the way of public consultation and the soliciting of advice and comment on aspects of their operation, which appears to be quite distinct to the mode of operation undertaken by ICANN.

## Respective roles of ICANN, PTI and Verisign over root zone contents

Each iteration of the root zone is produced as the outcome of a multi-party process, where the zone file is the result of records managed by PTI (NS, DS and glue records) and records provided by Verisign DNSSEC RRSIG records (generated through the use of the ZSK).

- **Is this separation of roles appropriate?**  
Does this multi-step process introduce vulnerabilities, or do they remove potential single points of failure by having multiple parties with oversight on the root zone as it is generated?

The case in point here is the inordinate delay in the signing of the zone used to name the root server instances.

[5]

## 2. Change Management

With respect to the root zone of the DNS, the list of delegated labels is not considered to be a static list, and the current mode of management of the root zone is to periodically open the zone for the inclusion of new top level labels. The general motivation behind this is that this expanded set of top level labels promotes diversity and competition in the domain name space, and this competition works to the benefit of the consumer in terms of reduced prices and improved focus on customer service for holders of second level name registrations. It is unclear if these changes and the increased diversity of top level names and top level name registries promote or detract from the overall security and integrity of the domain name system.

- Introduction of new TLDs  
  
Is the phased introduction of TLDs into the root zone, resulting in a zone that is in a state of constant flux better or worse than a single introduction of a set of new top level labels in a single event, or is the mode of introduction of new labels neutral to the security and stability of the root zone management function?
- Aside from ccTLDs (below) is there any consideration of the retirement of TLDs from the root zone?
- Coordination with ISO3166 or both introduction and retirement of ccTLDs  
  
What is the nature of the interaction between ISO 3166 and the root zone? Are all two letter TLDs reserved? What about Exceptionally reserved, traditionally reserved and indeterminately reserved names? Does ISO3166 provide for CC name retirement in an acceptable manner? [6]
- Coordination with IETF over Special Use Names Registry  
  
There is none at the moment! See SSAC 090
- Coordination between IETF and Unicode Consortium over IDNA standards and practice  
  
See SSAC 095 for the specific case with emojis, but the general observation holds true as well.
- Evolution of the Root Service  
  
Does the future security of the DNS root service rely solely on the current root server operators and the infrastructure that they operate? Should the model of distribution of root zone data evolve to include consideration of the opportunities offered by DNSSEC, such as local root secondary servers and recursive resolvers DNSSEC NSEC

caching. How can ICANN assist in these measures to assist in scaling the root and making it more resilient to known attack vectors?[7]

- Consistency of the Identifiers

The DNS resolution protocol is not the only protocol that performs a mapping from a domain name to a IP address (or other 'attached' attribute. To what extent should the ICANN community take steps to ensure that a domain name has a consistent meaning irrespective of the method of name resolution? i.e. is ICANN's remit solely concerned with domain names as resolved by the DNS protocol. Or does it extend to domain names as resolved by any protocol in the context of the public Internet? Or domain names irrespective of the name's manner and context of use?[8]

### 3. Roles and Responsibilities

There are many bodies who have an interest in the DNS and its operation. ICANN and its SOs and ACs bring many of these interests together, but not all. The broader DNS ecosystem includes aspects of applicable technical standards and their evolution, security and threat analysis, and modes of use of identifiers by applications and services. The observation is that this is not a static space and changes are anticipated. The consideration from security and stability is to ensure that such changes are considered carefully and the motivations that are driving such change are balanced against what is prudent and safe in terms of operational practice and use of available technology.

- Has ICANN achieved an effective balance relating to community policies, applicable standards, and SSR concerns?
- Are there checks and balances in the process, and do they have a voice?

### 4. Abuse and Threats

What are ICANN's responsibilities in this space?

*<more material needed here!>*

#### *Notes about DNSSEC.*

DNSSEC is the chosen mechanism to support security in the DNS, in so far that the resolution process itself is not protected, but the result can be validated to assure the client that the resolution response is genuine and complete. Being such a central part of DNS Security in general, should it be a topic of study in this sub-topic?

DNSSEC is largely under the control of the IETF as a piece of technology. If a party wants to alter the operation of DNSSEC or any other aspect of the way in which its operation is defined then the IETF is the place to undertake such a conversation. This implies that such technical aspects are beyond ICANN's remit.

DNSSEC, as seen by ICANN, is firstly a set of DS delegation records, similar in almost every respect to NS records. There are some operational questions in this process, noted in section 1.2. Secondly, DNSSEC is a root zone signing operation. This area includes the management of the KSK key which appears to fall under the subtopic of ICANN operations as it is not a generic DNS topic. The question of the choice of protocol and key length and the scheduling of periodic KSK rolls appears to also be an ICANN operational topic rather than a generic DNS stability issue.

