NEW GTLD COLLISION OCCURRENCE MANAGEMENT

Proposal to manage the collision occurrences between new gTLDs
and existing private uses of the same strings

1. INTRODUCTION

ICANN's mission and core values call for ICANN to preserve and enhance the operational stability, reliability, security, and global interoperability of the Internet's system of unique identifiers (names, IP numbers and protocol parameters). In pursuing these goals and following the direction of its Board of Directors as well as taking into consideration the advice of the Security and Stability Advisory Committee, ICANN commissioned a study on the potential security impacts of the applied-for new-gTLD strings. The study was to consider whether name collisions might occur between applied-for new gTLD strings and domain names that may be in use in private namespaces ("non-delegated TLDs"). The study was also to review the possibility of name collision occurrences arising from the use of internal names for which X.509 digital certificates have been issued.

A name collision occurs when users unknowingly access a name that has been delegated in the public DNS when the user's intent was to access a resource identified by the same name in a private network. Circumstances like these, where the administrative boundaries of private and public namespaces overlap and name resolution yields unintended results, present concerns and should be avoided if possible. However, the collision occurrences themselves are not the concern, but whether such collisions cause unexpected behavior or harm, the nature of the unexpected behavior or harm and the severity of consequence.

On 5 August 2013, ICANN published and made available a name collision study http://www.icann.org/en/about/staff/security/ssr/name-collision-02aug13-en.pdf (the "Study") that identifies categories of strings according to the occurrences of queries, as observed in root server log samples obtained from the "Day in the Life of the Internet" (DITL) initiative from DNS-OARC. The Study used as input: 1) samples of DNS requests transmitted to root servers (from the DITL initiative), complemented with 2) information from Certificate Authorities regarding the issuance of internal name certificates (e.g., TLS/SSL certificates for non-delegated names). A full description of the methodology of the Study can be found in section 3.4 of the Study.

The Study also included options to mitigate the risks; however, it does not make specific recommendations for each of the categories. Based on the Study, ICANN staff published a proposal to manage the risk of name collision for public comment from 5 August to 17 September 2013 (http://www.icann.org/en/about/staff/security/ssr/new-gtld-collision-mitigation-05aug13-en.pdf).

This paper describes a revised proposal to manage name collision occurrences between new gTLDs and existing private uses of the same strings based on input received during the public comment period. Appendix I describes the main points made in the public comment forum and how these shaped the updated proposal.

2. HIGH-RISK (HOME, CORP)

The Study identifies two strings, ***home*** and ***corp***, that will likely cause problems if delegated[1], given their high frequency of occurrence in the 2012 and 2013 DITL data (an order of magnitude higher than the next most frequently occurring string). The Study identifies these strings as having a level of queries in the realm of heavily used TLDs. Both strings are also widely used in private namespaces within internal networks (for example, see Appendix G of RFC 6762, http://tools.ietf.org/html/rfc6762). Additionally, ***corp*** is identified as the string with the highest number of internal name certificates (see Appendix C of the Study).

Based on the analysis of frequency of occurrence and the perceived severity of impact, ICANN will defer delegating ***home*** and ***corp*** indefinitely.

ICANN will collaborate with the technical and security community to continue to study the issues presented by these strings.

3. PROPOSAL TO MANAGE COLLISION OCCURRENCES

3.1. COLLISION OCCURRENCE MANAGEMENT FRAMEWORK

ICANN will commission a study to develop a name collision occurrence management framework. The framework will include appropriate parameters and processes to assess both probability and severity of impact resulting from name collision occurrences. Examples of the parameters include number of DNS requests, type of DNS requests, type of queries, diversity of query source and appearances in internal name certificates.

The framework will specify a set of name collision occurrence assessments and corresponding mitigation measures if any, that ICANN or TLD applicants may need to implement per second level domain name (SLD) seen in the DITL and other relevant dataset (e.g., information from Certificate Authorities regarding the issuance of internal name certificates)[2]. The proposed name collision management framework will be made available for public comment.

[1] See section 6 of the Study.

[2] Note that measures taken by ICANN or TLD applicants are attempts to mitigate unintended consequences or harm by preventing a name collision from occurring. These measures do not mitigate the causes of collision occurrences. Mitigating causes is a matter for users, private network operators, software developers, or equipment manufacturers to address.

3.2. COLLISION OCCURRENCE ASSESSMENT

ICANN will apply the final name collision occurrence framework, using DITL and other relevant data as an input, to each applied-for TLD and will deliver a name collision occurrence assessment to each applicant. The assessments will be published.

The assessment for each applied-for TLD will include a list of SLDs, an associated name collision occurrence assessment, and suggested mitigation measures; for example,

- Block the SLD indefinitely. In this proposal "block" means that the SLD must not be made available for registration, must not be delegated or otherwise activated in the TLD zone file (i.e., the SLD must not resolve and must return the same DNS results (NXDOMAIN) that the public DNS returns today, i.e., prior to the delegation of the new gTLD), and must not be used in any way by the registry operator,
- Block the SLD temporarily, i.e., until analysis or evidence that the cause of collision occurrence has been mitigated or data are available to demonstrate the collision occurrences are substantially reduced (e.g., demonstrably "negligible),
- Conduct a trial delegation of some form,
- Make the SLD available to the single entity that is the sole originator of name collisions for that SLD, or
- Other mitigation measures that may be identified during the course of the collision occurrence assessment or other studies.

ICANN will proceed with its established processes and procedures for delegating each applied-for gTLD. The registry operator will either (a) implement the mitigation measures described in its SLD collision occurrence assessment before activating any SLD, or (b), the registry operator can block those SLDs for which the mitigation plan has not been implemented, and proceed with delegating SLDs that are not listed in the report. The implementation of the mitigation measures may allow the release of blocked SLDs at a later time, based on analysis or evidence that the cause of collision occurrence has been mitigated.

Additionally, registry operators will implement a "wait" period of no less than 120 days from the date that a registry agreement is signed before it may activate any names under the applied-for TLD in the DNS. The length of this period is based on the Baseline Requirement 11.1.4 for Certification Authorities (CAs)[3]. Impact on TLD launch should be minimal in most cases because a set of activities must be completed between contracting and launch that account for a significant part of the 120 days (see figure 1). This measure will help mitigate the risks related to the internal name certificates issue as described in the Study report and SAC 057, SSAC Advisory on Internal Name Certificates located at http://www.icann.org/en/groups/ssac/documents/sac-057-en.pdf.

Registry operators, if they choose and if otherwise allowed by their registry agreement, may accept registrations during this period, but they will not be permitted to activate them in the DNS. If a registry operator chooses to register names during this 120-day period, the operator must clearly inform the registrants (through the registrars) about the inability to activate names until the period ends.

[3] https://www.cabforum.org/Baseline_Requirements_V1_1_6.pdf

It is possible that name collision occurrences of some second-level labels that did not appear in the study dataset might occur after the applied-for gTLD begins operation. To mitigate the risk that name collisions not observed in the study dataset occur and cause severe impact, ICANN and the registry operator shall implement a process to enable an affected party(ies) to report and request the blocking of a domain name (SLD) that causes demonstrably severe harm as a consequence of name collision occurrences. Such reports must be processed through an ICANN point of contact, which will coordinate the notification with registry operators and ensure that the report is acted upon in an expedited manner. The process will allow the deactivation (SLD removal from the TLD zone) of the name for a period of up to two (2) years in order to allow the affected party to effect changes to its network to eliminate the DNS request leakage that causes collisions, or mitigate the harmful impact. The process will be in effect only for the first two years after delegation.

Figure 1 shows the timeline of the activation of SLDs considering the processes introduced by this paper.

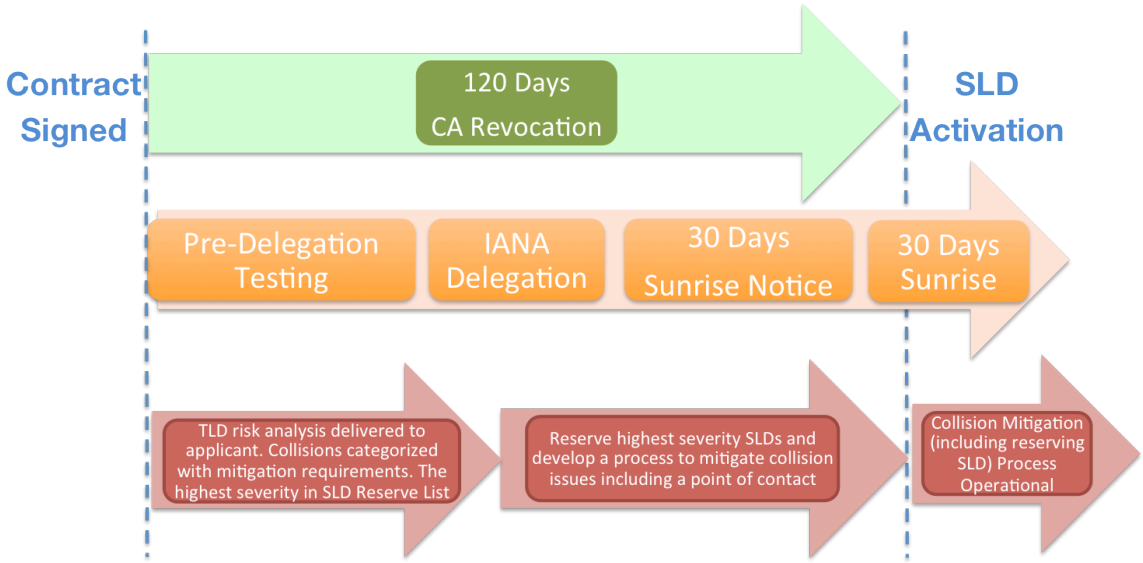

Figure 1 – Timeline to activate SLDs

3.3. Alternate Path to Delegation

A registry operator may elect to proceed to delegation (subject to established processes and procedures) prior to receiving its corresponding SLD collision occurrence assessment report. If the registry operator so chooses, it must implement a conservative collision mitigation measure and initially block *all* SLDs that appear in the DITL and other relevant dataset while the assessment is conducted. ICANN will develop a list of labels to be blocked at the second level under the TLD, and then determine whether the proposed TLD is eligible for this option to delegation. This list will be made publicly available and will consist of all the second-level labels that appeared in DNS requests to the applied-for TLD in the DITL and other relevant dataset. Blocking all second level labels (and thus preventing these labels from resolving in the newly delegated TLD) ensures that corresponding DNS requests for such labels in the newly delegated TLD will return the same DNS results (NXDOMAIN) that the public DNS returns today, i.e., prior to the delegation of the new gTLD.

The registry operator will have the option to (1) request its corresponding SLD collision occurrence assessment in order to implement the mitigation measures or (2) leave the SLD blocking in place. The registry operator will still be required to participate with ICANN in the process that enables affected party(ies) to report and request the blocking of a domain name (SLD) that causes demonstrably severe harm as a consequence of name collision occurrences.

3.4. OUTREACH CAMPAIGN

ICANN will develop an outreach campaign to

a) Make the public as well as private network operators aware of the possibility of name collision occurrences as new TLDs are delegated (e.g., raise general awareness of the problem space using multiple communications media, technical briefs, or social media),

b) Advise users and private network operators of the measures that ICANN and new TLD registries are able to and will take to minimize the potential for unintended consequences or harm, (e.g., measures that manage collision occurrences by maintaining the same (NXDOMAIN) responses for queries that appear in the public DNS),

c) Assist users, private network operators, and software or equipment manufacturers with the identification of causes (origins) of name collisions.

ICANN will invite and collaborate with other parties and members of the community that share a common interest in identifying strategies for eliminating or managing name collision causes from their networks.

4. CONCLUSION

ICANN's mission and core values call to preserve and enhance the operational stability, reliability, security, and global interoperability of the Internet's system of unique identifiers (names, IP numbers and protocol parameters). ICANN is fully committed to the delegation of new gTLDs in accordance with its mission and core values. ICANN appreciates the community's involvement in the process and look forward to further collaboration on the remaining work.

APPENDIX I – RESPONSES TO PUBLIC COMMENTARY

During the public comment period for the proposal to manage name collision risks, 75 comments were received[4]: 35 in favor of moving forward with the new gTLD delegation in the current projected timeframe in one way or another, 31 against rolling out new gTLDs in the current projected timeframe without first doing additional studies, and 9 making neutral proposals. The public comment report summarizing the comments, and the full comments can be found at http://forum.icann.org/lists/comments-name-collision-05aug13/.

Various commenters stated that the proposed categorization of strings as high risk, low risk and uncalculated risk was not correct, that frequency of occurrence be differentiated from severity of impact, and impact be properly studied. Verisign and others proposed that basing decisions only on frequency of DNS request was not enough and that ICANN should consider a series of parameters (e.g., request frequency, type of query, type of request, etc.) to define the risk level. The ISPCP constituency and others commented that further study was needed before moving forward with delegating any new gTLD. Radix and others commented that ICANN used arbitrary methods for dividing strings into categories.

ICANN acknowledges that a collision occurrence assessment is comprised of two components, namely frequency of occurrence and severity of impact. ICANN agrees that other parameters, besides request frequency, should be considered in assessing the threat, particularly the potential for harm caused by name collisions. ICANN will adopt the advice regarding the use of the other proposed parameters when developing a collision occurrence management framework.

NTAG and others suggested the idea to block Second Level Domain names (SLDs) that are being queried to eliminate the name collision risks. Blocking SLDs that appeared in the DITL data (i.e., by prohibiting these names from resolving in the newly delegated TLD) so that DNS queries that currently leak into the public DNS will continue to return "name error (NXDOMAIN)" responses, avoids the possibility of harm while the blocking is in effect. ICANN will adopt the idea by NTAG and others to block Second Level Domain names (SLDs) that are being queried.

The ANA and others requested that the public comment period be extended in order to have more time to study the risks in their networks posed by the name collision. With the adoption of the SLD blocking measure, the potentially affected parties will have more time to study the issues in their networks without being affected by new gTLD delegations. Additionally, to address effects by SLDs that were not blocked but that generate significant harm, ICANN will enable an affected party to report and request the suspension of a domain name that by virtue of name collisions is causing severe harm.

Daniel Karrenberg and others commented that ICANN should neither mandate nor recommend that registry operators notify the point of contacts of IP addresses that issue DNS requests for an non-delegated TLD or names under it, on the basis that such notifications will not be effective and pose a significant risk for abuse. ICANN acknowledges the issue and has removed the 30-day notification measure from the proposal. ICANN looks forward to learning of new and better ways to notify parties potentially affected by name collision with new gTLDs and seeks comment on the elements of the outreach campaign described in this proposal.

[4] There were 80 comments in total, however, 5 were duplicates

There were a number of comments regarding the 120-day period of no activation of names, mostly by new gTLD applicants. Some requested there be a waiver of the period for certain cases, others that it be shortened, or at least started immediately without having to wait for the signing of the registry agreement. ICANN remains committed to mitigate the internal name certificate risk and will not waive the requirement at this time. The period is aimed at all new gTLDs, because there is no way to obtain the entire set of internal name certificates that have been issued for names under a new gTLD. The data provided by the CA/B Forum for the Study represented only a subset of CAs (only those CAs that were willing to provide it). ICANN acknowledges the request to shorten the period or start it early and will consider liaising with the CA/B Forum to explore alternatives.

Warren Kumari and Danny McPherson provided an explanation of the search list processing in operating systems as one of the culprits of queries in the public DNS for non-delegated TLDs. ICANN appreciates the description of the issue and acknowledges that search list processing requires further consideration as part of identifying means to mitigate causes of internal name queries that are submitted to the public DNS (see Section 3.4 of this Proposal).

The Association of the German Internet Industry, Google, and NTAG provided evidence that is, in general terms, consistent with the findings in the Study regarding resolver data. The use of root server data, as opposed to resolver data, seems to overestimate the number of collisions for most non-delegated TLDs as a fraction of the overall query traffic overall. ICANN acknowledges the differences in behavior that are observed from root or resolver and will take this into consideration in developing the collision occurrence management framework.

Andrew Sullivan, O. Kolkman, and W. Kumari offered for consideration an Internet Draft (available at http://tools.ietf.org/html/draft-kolkman-root-test-delegation) that sketches a methodology that could be helpful to make some determinations regarding the possible disruption of name collisions prior to actual allocation and delegation. ICANN appreciates the suggested approach and encourages the authors to seek community input on the Internet Draft. ICANN will consider the approach as it develops the collision occurrence management framework.

NTAG and others commented that existing TLDs have been delegated in the past without incident; some of these TLDs were delegated having parts per million queries beyond those seen for most applied-for strings. ICANN acknowledges the observation, though notes that the current new gTLD Program is a first in terms of the number of TLDs that will be delegated and, therefore, requires cautious steps forward.

DotGreen requested that strings in the uncalculated-risk category be allowed to proceed to contracting. Similarly, other commenters complained about ICANN not allowing these strings to proceed to contracting when the public comment period for the proposal is still open. ICANN understands the interest of applicants to see their strings move as fast as possible through the new gTLD process and will remove that restriction. The adoption of the blocking of SLDs makes this restriction unnecessary.

Microsoft, Verisign, and Yahoo! requested that ICANN implement SAC045, SAC046, SAC057 and SAC059 recommendations before moving forward with new gTLDs. ICANN appreciates the requests and agrees on the potential value of having said recommendations implemented.

Certain of the SAC045 recommendations have been implemented, e.g., the advisory included in the new gTLD Applicant Guidebook. The remaining recommendations have been or are in the process of

being implemented as part of the collision occurrence management framework; for example, recommendations from SAC 045 regarding awareness or outreach.

Regarding SAC046 and SAC059 recommendations, ICANN notes that most of these recommendations, while valuable, are targeted to the root zone scaling issue and therefore not directly related to the name collision occurrences. Nevertheless, all of the recommendations are either being implemented, or have already been implemented. A more detailed report status on the implementation of individual recommendations can be found at http://www.icann.org/en/news/correspondence/moss-to-falstrom-30apr13-en.

ICANN notes that SAC 057 has been fully implemented.

Several commenters requested that ICANN implement an outreach campaign to educate potentially affected parties on the name collision occurrences. ICANN has identified elements of an outreach campaign in this proposal.

DotHome, Radix Registry, Donuts, and NTAG proposed that high-risk strings be allowed to continue the contracting and delegation process pending implementation of mitigation measures. ICANN considers that the Study presents sufficient evidence to classify ***home*** and ***corp*** as high-risk strings, including evidence on the severity of consequences beyond query frequency. Given the risk level presented by these strings, ICANN will not delegate either one pending further study of the issues presented by these strings and alternatives to address them.

ICANN notes the comments raised in the letter from Verisign dated 27 August 2013, where Verisign expressed concerns that "*the risks arising from name collisions (and other security and stability risks) should be mitigated by ICANN, and not applicants, and should be completed prior to delegation of any new gTLDs.*" As established in the ICANN Bylaws, ICANN's role is to "*coordinate, at the overall level, the global Internet's systems of unique identifiers* …" The proposed mitigation plan to address the collision risks between new gTLDs and existing private uses of the same string would serve to provide overall coordination of the issue. Based on the public comments received, ICANN has proposed additional revisions to the proposal to strengthen the coordination plan to mitigate the name collision risks.

General Electric and others questioned whether the DITL data is good enough to make inferences given its limitations, as it did not observe weekend, month-end, or other periodic or differentiated patterns in the Internet. ICANN observes that the DITL data comprises 8 different ~48-hour sampling periods from root servers during 8 years. The DITL data is deemed the best available data set because it is neutral, relatively broad, available for crosschecking, and historical. To complement the use of DITL data as the basis for proceeding with applied-for TLDs and for managing SLDs, ICANN is proposing an ongoing collision occurrence management process whereby parties are able to report harmful collisions to ICANN so that SLDs with demonstrable risk can be blocked or suspended.