

Network Working Group
Internet-Draft
Intended status: Experimental Protocol
Expires: May 20, 2014

G. Huston
APNIC
O. Kolkman
NLnet Labs
A. Sullivan
Dyn, Inc.
W. Kumari
Google, Inc.
November 18, 2013

**Using Test Delegations from the Root Prior to Full Allocation and
Delegation
draft-kolkman-root-test-delegation-02**

Abstract

The delegation of certain strings as generic Top Level Domains (gTLDs) may cause stability and security issues if such strings have been used in private environments prior to their delegation. Test delegations can be used to enable empirical research on the extent of the potential for name collision. This document describes one such approach to an empirical testing framework for name collision, and considers the applicability of this approach to detect other forms of name collision.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 20, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | | |
|-----------------------------|---|--------------------|
| 1. | Introduction and Motivation | 2 |
| 1.1. | Scire est mensurare | 3 |
| 2. | Terms and Conventions Used in this Memo | 4 |
| 3. | Principle of Operation | 4 |
| 3.1. | Measurements Servers and Zones | 5 |
| 3.2. | Query Generation | 5 |
| 3.3. | Sampling | 6 |
| 4. | Evaluation | 6 |
| 5. | Name Resolution Considerations | 7 |
| 6. | Security Considerations | 9 |
| 7. | References | 9 |
| Appendix A. | Acknowledgements | 10 |
| | Authors' Addresses | 10 |

[1.](#) Introduction and Motivation

[[The authors are aware that this version of the document is not fully consistent. However they would value feedback on whether the idea is worth further study. A mail list to discuss this draft is collisions@lists.dns-oarc.net.]]

While certain names have been reserved for internal or private use [[RFC6761](#)], there is evidence [[SAC45](#)] that various sites connected to the Internet have used other names for internal purposes. In fact, the Multicast DNS specification [[RFC6762](#)] advises not to use .local for private use and observes: "the following top-level domains have been used on private internal networks without the problems caused by trying to reuse ".local." for this purpose:

- .intranet
- .internal.
- .private.
- .corp.
- .home.

.lan.

Huston, Kolkman, Sullivan Expires May 20, 2014

[Page 2]

In the event such names are delegated for use in the public DNS, there will be inevitable consequences for sites that have used those names. Some of those consequences may have security implications, with the potential for leakage of credentials and HTTP cookies ([RFC6265]). Responsible administration of the public namespace therefore requires careful consideration in permitting public delegation of any name when there are grounds to believe it is in widespread use as a private namespace, even though such private namespaces are (from the point of view of the DNS) irregular, even if common.

One form of name collision involves network domains that use selected names as local-use top level domains, as noted in [RFC6762]. In the case where the same label is delegated in the global DNS as a gTLD, then hosts in the local domain will be unable to resolve domain names in the context of the gTLD. This state of name occlusion is further compounded by a number of scenarios where the resolution of a name is performed across multiple name scope domains. This may happen with a mobile host (in the case, for example, when the host uses a statically defined "home page" on their local browser that is defined within a particular local scope), or even with applications, such as, for example, mail delivery (in the case where multiple MTAs who are listed as mail servers for a domain reside in different name scope domains, some of which have this name collision between the domain and locally defined pseudo-TLDs).

Name collision opens up the potential for misdirection, where the named remote point being contacted by the application may not necessarily be the intended service point for the transaction. When a host leaves the intranet environment, the host's applications may anticipate that the DNS names associated with a label return an RCODE 3 (NXDOMAIN) response, but may encounter an unanticipated response when the gTLD is deployed with a colliding name. Similarly, a host that has an association with a named service point within the gTLD may encounter unanticipated responses when the host is placed into an intranet environment where the same name exist as a locally-scoped pseudo-TLD.

There is a subtle form of interaction of names when the same name is placed on a local name search list. Certain name resolver libraries first query the original name, and if the query returns an NXDOMAIN, then they apply the local search list to the original name. When this process occurs in the context of a visible gTLD name colliding with the local name there is the possibility of the name resolving in the context of the gTLD, which then bypasses the application of the local search list.

1.1. Scire est mensurare

Huston, Kolkman, Sullivan Expires May 20, 2014

[Page 3]

The local use of undelegated top-level domain names is troublesome because it may produce different user experiences depending on the locally used name, the names placed in a local search list and the location of a given host, and the host's name resolution behaviour.

Prudent operation of the root zone requires that deployment of new names in the root should not necessarily cause widespread untoward effects for users of the DNS, particularly when those users are relying on name resolution outcomes that have always been part of the name resolution behaviour up unto this point.

What is useful in this context is a mechanism to test whether a particular delegation from the root zone presents a conflict with widespread local use. This memo presents a methodology for making such a determination.

The methodology considered here depends on temporary delegation of the top-level domains in question, and the use of a domain under an existing TLD in order to capture and compare queries generated by a large number of querying sources under the control of the experiment.

2. Terms and Conventions Used in this Memo

The mechanism outlined here is intended to complement the analysis already performed in "Name Collision in the DNS" [[namecollision](#)]. We therefore use the terms defined in section 1.1 of [[namecollision](#)] whenever appropriate.

Note that the evaluation methodology outlined here is intended to be complementary input to a risk analysis e.g. as found in [[namecollision](#)]; risk tradeoffs are likely to include other factors than the effects measured herewith.

3. Principle of Operation

The goal of the experiment is to assess whether there is significant existing use of a given candidate string ("CandidateTLD").

We propose the use of a software test that is executed by a large number of end hosts drawn from across the entire Internet. The execution of this test will cause the end host to attempt to retrieve a small set of URLs. This will trigger a set of DNS queries to resolve the domain name part of each URL, and subsequent HTTP queries to retrieve the object in the case that the DNS name is successfully resolved to an IP address. Both the DNS queries and the HTTP requests are answered by dedicated servers that analyse the received responses and match them to the original set of queries that were used by the end host. This will allow us to infer whether the host is located in an context where there is name collision with the CandidateTLD. In this section we describe the query generation, data-

collection, and analysis.

This methodology is based on earlier work by APNIC [[Method](#)].

3.1. Measurements Servers and Zones

In addition to the use of CandidateTLD, the methodology uses an additional name, delegated from a 'common' existing TLD, ("TestName.ExistingTLD") to the experiment's server.

The experiment's name server is authoritative for CandidateTLD and TestName.ExistingTLD. The name server will respond to an A and AAAA query for any name within "TestName.CandidateTLD" with the IPV4 or IPV6 address of the experiment's HTTP server. The name server will respond to queries for any other name within CandidateTLD with RCODE 3 (Name Error or NXDOMAIN). The name server will respond to A and AAAA queries in TestName.ExistingTLD with the IPV4 or IPV6 address of the experiment's HTTP server.

The experiment's HTTP server will respond with a "200 OK" for a request for the object "1x1.png" in TestName.CandidateTLD and in TestName.ExistingTLD. The server will respond with "404 Not Found" for any other object name.

3.2. Query Generation

The TestName is a synthetic name with no intentional semantic meaning, that is generated in such a way to reduce the likelihood of collision with any existing delegated name. It is suggested that it be generated by using the hex encoding of a randomly selected integer value between 1,000,000,000 and 2,000,000,000. The name must not be already delegated from the root or in the ExistingTLD.

Each query set constitutes one "measurement". A "measurement" is identified by a measurement identifier (<uniqueid>, syntactically a valid hostname) that is uniquely generated for each instance of a measurement. This ensures that when the domain name is resolved, and when the named object is retrieved there is no occlusion of the interaction with the experiment's services because of local name or web object caches. The set uses the following URLs:

A: `http://<unique_id>-a.TestName.CandidateTLD/1x1.png?<uniqueid>-a`

B: `http://<unique_id>-a.TestName.ExistingTLD/1x1.png?<uniqueid>-b`

C: `http://results.TestName.ExistingTLD/1x1.png?<uniqueid>?za=<a_result>&zb=<b_result>`

The A URL is intended to test if CandidateTLD is a locally used name. In other words, if local use of CandidateTLD occludes visibility of CandidateTLD as a gTLD. The DNS query for the A Fully Qualified Domain Name (FQDN) will only be received by the authoritative name server for this name if there is no local name resolution function that uses the CandidateTLD name as a locally defined pseudo-top level domain.

The B URL is intended to function as the control test for the experiment, and the use of ExistingTLD in B is intended to operate as a name that does not collide with a local use context.

As the experiment uses the absence of a fetch of the A URL to infer the name resolution behaviour of the location where the measurement is being performed, it is necessary to ensure that the measurement code has run to completion. The measurement code starts a timer at the start of its execution. Upon expiration of the timer, or when both the A and B objects have been successfully retrieved, the code will schedule the retrieval of the C URL. The arguments to the C URL include the client-side measurement of the elapsed time to retrieve the A and B URLs.

3.3. Sampling

One way to perform this measurement is to embed the measurement in web content, using a scripting language. When the web content is loaded the script is activated, and the measurement sequence is performed.

One way to distribute this content to clients to perform the test is via an online (ad) campaign. If the measurement script is enclosed within the ad itself, then there is no reason for the campaign actually to cause users to click though in order to perform the test. Behavior of this sort is trivially achievable with a number of available online advertising systems.

It is also necessary to spread the delivery of the ad to a very broad spectrum of clients, use the ad should be presented across all time zones, across all language bases, and across all geographic regions.

4. Evaluation

To evaluate the results, we take those measurements that return the C URL. The use of the C URL ensures that we use measurement results where the ExistingTLD name is not being locally occluded. We count the number of experiments of each of the possible combinations of retrieving the A and B URLs. These combinations are:

Not A and Not B: This result contributes to experimental uncertainty. (We know that ExistingTLD is not locally occluded, so the failure to retrieve B is due to other factors that are not being examined in the context of this measurement.)

A and Not B: This result indicates that the client is able to resolve names in the CandidateTLD in the context of the global DNS, but the inability to retrieve the B URL contributes to experimental uncertainty. (The same reasoning about the ExistingTLD and local occlusion applies to this case).

Not A and B: This result is an indicator that the client's use of CandidateTLD is probably being occluded by some form of local use.

A and B: This result indicates that the client is able to resolve names in the CandidateTLD in the context of the global DNS.

If the CandidateTLD is in widespread private use then we would see the count of "Not A and B" be far in excess of the level of experimental uncertainty, then we can conclude that there are locales where the CandidateTLD is being used in local context. Analysis of the source IP addresses of the clients that fetch "Not A and B", and the BGP Origin AS of these addresses and their geolocation may indicate if such local use is clustered in a particular network or group or networks, or clustered in a particular geography or language region.

5. Name Resolution Considerations

Earlier versions of this memo proposed to use this experimental technique to detect name search list considerations. This section describes the name search list collision considerations, and describes some further investigation that has lead to the conclusion that this technique would not necessarily be applicable in that context.

The basic algorithm used in name resolution when search lists are present appears to be consistent across a number of implementations: various permutations of using the base name and appending individual values from the name search list are used as DNS queries in order to find a name that can be resolved by the local DNS resolver. The search process stops when the DNS query returns other than an NXDOMAIN response.

However the exact order of generating these candidate names has been observed to vary across implementations. To describe these observations it is first necessary to introduce some basic

terminology. There are four generic ways that name resolution libraries apply a search list to a "base name" in order to construct a set of FQDN that are used in DNS queries:

none the search list is not applied to the base name.

pre the search list is applied to the base name, then the base name alone is used.

post the base name alone is used, then the search list is applied to the base name.

always the search list is applied to the base name, and the base name alone is not used.

The form of name collision with search lists, as discribed in the introduction section of this memo, occurs in the "post" case, where the unexpected resolution of the base name causes the search list not to be applied to the base name, and the global name context is applied to the base name, rather applying a local name context, as defined by the search list.

Table 1 provides a summary of the behaviour of various operating systems and their local name resolver library behaviour when resolving base names that contain a single label, and names that contain two labels. As can be seen, only Windows XP and Unix-based libraries perform the "post" form of search name application that would be susceptible for this form of name collision.

| System | Single Label | Multi-Label |
|---------------|--------------|-------------|
| MAC OSX 10.9 | always | never |
| Windows XP | always | post |
| Windows Vista | always | never |
| Windows 7 | always | never |
| Windows 8.1 | always | never |
| FreeBSD 9.1 | pre | post |
| Ubuntu 13.04 | pre | post |

The experimental approach described here does not necessarily use the operating system's name resolution libraries. The experimental technique forms a name query within the browser, so it is more relevant to examine the behaviour of the browsers when given single and multi-label names to lookup. Table 2 shows the behaviour of a number of browsers on two operating system platforms. (It should be noted that these results in Table 2 were obtained by using Javascript to feed names to the browser. The interactive data entry procedures in current browsers are a dual purpose URL and search engine term data entry, and the variations on behaviour between browsers in the

way in which entered data is interpreted is more due to the differences in the browser's input parser than it is due to any differences in the browser's name resolution library.)

| System | Single Label | Multi-Label |
|---------------------------|--------------|-------------|
| MacOS OSX 10.9 | | |
| Chrome (31.0.1650.39) | always | post |
| Opera (12.16) | always | never |
| Firefox (25.0) | always | never |
| Safari (7.0 9537.71) | always | never |
| Windows 8.1 | | |
| Chrome (30.0.1599.101) | always | never |
| Opera (17.0) | always | never |
| Firefox (25.0) | always | never |
| Safari (5.1.7 7534.57.2) | always | never |
| Explorer (11.0.900.16384) | always | never |

Only one browser / Operating System combination tested shows the "post" form of search name use, namely Chrome on the Mac OSX platform. In all other cases a single label name always has the local search list appended, and a multi-label name never applies the local search list.

6. Security Considerations

The delegation of the Proposed TLD (CandidateTLD) comes with some risk of interference with existing deployments. In the case where a local system queries a name, and that query returns a NXDOMAIN response, then local system then queries further name forms where each entry on a local name search list is appended to the original name in turn, searching for a name response that is not NXDOMAIN. The delegation of CandidateTLD for this experiment may interfere this behaviour.

However, two observations mitigate this concern. The first is that this situation of potential collision arises in the case where the local system is querying for the CandidateTLD name as a "dotless" name (as the only delegated subdomain in the CandidateTLD zone is TestName, which is intended to have no semantic meaning in any language). The second observation is that for such "dotless" names, the currently widely deployed name resolver libraries do not initially query the "dotless" domain name then apply the search list is the first query results in an RCODE 3 response. Many name resolver libraries do not query for "dotless" domain names at all, while those libraries that have been observed to perform such queries (Windows XP, Linux, FreeBSD) perform them after using the local search name list, rather than before.

7. References

[Method] APNIC, "APNIC Labs IPv6 Measurement System ", May 2013.

[RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), April 2011.

Huston, Kolkman, Sullivan Expires May 20, 2014

[Page 9]

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", [RFC 6761](#), February 2013.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", [RFC 6762](#), February 2013.

[SAC45] ICANN Security and Stability Advisory Committee, "Invalid Top Level Domain Queries at the Root Level of the Domain Name System", 11 2010, <<http://www.icann.org/en/groups/ssac/documents/sac-045-en.pdf>>.

[namecollision]

Interisle Consulting Group, "Name Collision in the DNS", August 2013.

[Appendix A](#). Acknowledgements

This draft is a follow-up of, and borrows heavily from, our earlier (abandoned) work on "A Procedure for Cautious Delegation of a DNS Names". Discussion of that document in various hallways led to inspiration for this document and we want to thank those that gave us feedback.

The idea of using different names to trigger events in a DNS server is due to Geoff Huston and George Michaelson.

The approach described here of using code embedded in ads delivered by online advertisement networks to generate a large volume of URL-based experiments performed by end users' browsers was developed by George Michaelson, Byron Ellacot and Geoff Huston.

Authors' Addresses

Geoff Huston
APNIC
6 Cordelia St
South Brisbane, QLD 4101
Australia

Email: gih@apnic.net

Olaf Kolkman
NLnet Labs
Science Park 400
Amsterdam, 1098 XH
The Netherlands

Email: olaf@NLnetLabs.nl

Andrew Sullivan
Dyn, Inc.
150 Dow St
Manchester, NH 03101
U.S.A.

Email: asullivan@dyn.com

Warren Kumari
Google, Inc.
1600 Amphitheatre Pkwy
Mountain View, CA 94043
U.S.A.

Email: warren@kumari.net

