Network Working Group
Internet-Draft
Intended status: Experimental Protocol
Expires: March 22, 2014

O. Kolkman
NLnet Labs
A. Sullivan
Dyn, Inc.
W. Kumari
Google, Inc.

September 20, 2013

Using Test Delegations from the Root Prior to Full Allocation and Delegation
draft-kolkman-root-test-delegation-00

Abstract

The delegation of certain strings as TLDs will cause stability and security issues if such strings have been used in private environments prior to their delegation. It is recommended that test delegations be used to enable empirical research on the extent of the possible disruption prior to actual allocation and delegation of any label in the root zone.

Table of Contents

1. Introduction and Motivation

[[The authors are aware that this first version of the document does is not fully consistent. However they would value feedback on whether the idea is worth further pursuance.]]

[Editor not: An appropriate mailinglist for discussion of this draft has not yet been identified]

DNS names have always co-existed with other namespaces that are virtually indistinguishable from the DNS. The DNS was itself deployed alongside the host [RFC0822] table. NetBIOS [NETBIOS] names, though only one label long, could always interact with the DNS search path mechanism to generate DNS names. Additionally, mDNS [RFC6762] names often look just like DNS names. Because different naming systems are usually linked together in the user interface, from an end user's point of view these name spaces are all one -- even though they function differently on the Internet.

While [RFC6761] reserved certain special names for internal or private use, there is evidence [SAC45] that various sites connected to the Internet have used other names for internal purposes. In fact, [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.

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 have implications for security, with the potential for leakage of credentials and HTTP cookies ([RFC6265]). Responsible administration of the public namespace therefore requires great care in permitting public delegation of any name when there is good reason to suppose 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.

1.1. Search-path interaction.

In many cases a string appears to be used as an "undelegated TLD" (being used as the rightmost label in an name), but this is simply an artifact of domain search list processing.

For example, suppose the Example Widgets corporation uses a sub-domain (.corp) of their primary domain (example.com) to name their employee workstations, servers, printers and similar. They have an "intranet" server named intranet.corp.example.com. In order to allow their employees to simply type "intranet.corp" to access this server, the users' workstations are configured (probably using [RFC3379]) with the search-list set to "corp.example.com, example.com".

When a user enters "intranet.corp", their workstation will try and resolve the name. RFC1535 [RFC1535] specifies that "in any event where a "." exists in a specified name it should be assumed to be a fully qualified domain name (FQDN) and SHOULD be tried as a rooted name first." So, the user's workstation will first try and resolve "intranet.corp.". As there is (currently) no .corp TLD this will result in an NXDOMAIN response. The workstation will then append entries in the search-list until it is able to resolve the (now fully-qualified) name.

If the .corp label were to be delegated as a TLD and the sub-domain "intranet" created within .corp, the first lookup ("intranet.corp") would no longer generate an NXDOMAIN response. This would stop the

search-list processing, and direct the user somewhere other than where the user intended to go -- the "wrong server", in the eyes of the user, even if right according to the DNS.

It is worth noting that a researcher analyzing DNS queries hitting the root servers would see queries before search-list processing expands them. While this may not change whether or not it is safe to delegate these names, having an understanding of the cause is valuable.

1.2. Scire est mensurare

The local use of undelegated top-level domain names is troublesome because it produces different experience depending on the search path and location of a given device. That is a normal effect of the search path mechanism or the roaming of users, but with the advent of new generic top-level domains (gTLDs), the problem gets more acute, because many TLDs are intended to be mnemonics that will be intuitive to humans. Since names higher in the DNS tree are likely also to use those same intuitive labels, there is potential for user confusion and information leakage.

At the same time, it is not clear that the DNS protocol was designed around a static list of top-level domains (TLDs), and therefore it seems reasonable to plan for the possible addition of new TLDs whose use might conflict with deployed search path settings. Yet prudent operation of the root zone requires that deployment of new names in the root should not cause widespread untoward effects for users of the DNS, particularly when those users are relying on features that have always been part of the protocol.

What is needed is a mechanism to test whether a particular delegation from the root zone presents a serious conflict with widespread use. This memo presents a methodology for making such a determination.

The methodology 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 tradeofs are likely to include other factors than the effects measured herewith.

3. Principle of Operation

In order to assess whether there is significant use of a given candidate string (CandidateTLD), probes will send out sets of queries from a large number of random locations. The queries are answered by dedicated servers that collect statistics. In this section we describe the query generation, data-collection, and what the dedicated servers should answer to not interfere with Internet traffic.

The goal of the experiment is to assess whether there is significant existing use of a single CandidateTLD.

3.1. Measurements Servers and Zones

In addition to the candidate string ("CandidateTLD") the methodology uses a specific sting "TestName". During an experiment the Proposed TLD under evaluation ("CandidateTLD") and a control TLD ("TestName") are delegated from the root zone to a special DNS name server, the experiment's server. Further, a second control name (TestName.ExistingTLD) is delegated from a 'common' existing TLD (ExistingTLD) to the experiment's server.

The experiment's server has two functions:

First, with one exception detailed below in Section 3.4, if any request for any name inside the CandidateTLD or the TestName namespace reaches the name server, the name server MUST respond with RCODE 3 (Name Error or NXDOMAIN) (Note that from a DNS client perspective the ultimate RCODE 3 response is indistinguishable from what is returned without the test delegation in place.)

Second, the experiment's server tracks every name in any request it receives, the time at which the request arrived, the RRTYPE and CLASS in the request, and the source network for the request.

3.2. Query Generation

Software probes will need to be deployed throughout the Internet (also see Section 3.3) these probes will send, in parallel, sets of queries.

Each set comprises one measurement and consists of three queries (or even 4, see Section 6). A measurement is identified by a unique measurement identifier (<uniqueid>, syntactically a valid hostname); the set will include the following:

the QNAME is <uniqueid>-a.TestName.

the QNAME is <uniqueid>-b.TestName.CandidateTLD

the QNAME is <uniqueid>-c.TestName.ExistingTLD

The TestName MUST be a randomly chosen name that remains constant for the duration of the experiment, MUST be a syntactically valid label, SHOULD be semantic nonsense, an MUST NOT be delegated from the root or in the ExistingTLD already.

Depending on the environment in which the probe is located the query that is send by a stub resolver is handled differently. We distinguish 4 possibilities.

Local CandidateTLD use without Search List: The probe is located within a network that locally resolves the candidateTLD and there are no searchlists being used that append CandidateTLD. The query with a QNAME of <uniqueid>-b.TestName.CandidateTLD will not exit the local network while the queries with a QNAME of <uniqueid>-c.TestName.ExistingTLD. and QNAME of <uniqueid>-a.TestName. will arrive at the authoritative servers for the respective domains.

Local CandidateTLD and CandidateTLD appened Search List: The probe is located within a network that locally resolves the candidateTLD and there are searchlists being used that append CandidateTLD. The query with a QNAME of <uniqueid>-b.TestName.CandidateTLD will not exit the local network while the queries with a QNAME of lt;uniqueid>-c.TestName.ExistingTLD. will arrive at the authoritative server for the domain. The query with a QNAME of <uniqueid>-a.TestName. is subject to the search algortithm, the query name will effectively be substitude to <uniqueid>-a.TestName.CandidateTLD and be resolved localy. The query for <uniqueid>-a.TestName. will therefore not be seen at the authoritative server.

No use of CandidateTLD and no use of Search List: The probe is located within a network that does not resolve the candidate TLD and no searchlists being used that append CandidateTLD. All queries will arrive at the authoritative servers for the respective domains.

No use of CandidateTLD and CandidateTLD appended by Search List: The probe is located within a network that does not resolve the candidate TLD but search list processing appends CandidateTLD. In this case the queries for Lt;uniqueid>-a.TestName. get rewritten to lt;uniqueid>-a.TestName.CandidateTLD and arrive, together with lt;uniqueid>-b.TestName.CandidateTLD at the authoritative server for CandidateTLD, while queries for the QNAME <uniqueid>-c.TestName.ExistingTLD arrive at the server authoritative for TestName.ExistingTLD.

As a result, by comparing what arrives at the authoritative servers one can establish the prevalence of the various scenarios. Under the

assumption of a broad unbiased sample exclusively observing the 3rd option (all queries hitting their respective servers) would be a strong indication that a candidate TLD is not in use.

3.3. Sampling

To perform the evaluation, a names of the form <uniqueid>.TestName.CandidateTLD and <uniqueid>.controlname.ExistingTLD are embedded in content that is placed around the web. As people visit web sites, the content is processed, yielding attempts at resolution of the names.

The easiest way to provide the content that causes the relevant DNS lookup is to use an online (ad) campaign. There is no reason for the campaign actually to cause users to click though, so it should be as boring as possible. However, the campaign must result in the independent resolution of both the control and test names. Behavior of this sort is trivially achievable with several available online advertising systems.

It is critical that the sampling be as representative of the Internet population as possible. This sort of sample already has the significant problem that it can only measure users of the web. And there may be sampling effects that might prevent measurements from taking place in those environments that need to be reached. For instance, add-blocking or different web surfing behavior in corporate environements.

The measurements should be unbiased with respect to temporal behavior like sleep-wake and work-rest cycles.

[[The sampling biases probably deserve their own section with much more elaboration and more possible biases]]

3.4. The Name Server

This procedure rests on a name server that is configured and instrumented in particular ways. First, the name server must be configurable so that it authoritative for all requests inside the Candidate TLD. Normally, it will always return RCODE 3 (NXDOMAIN) for all queries inside that Candidate TLD, except that the name server must also be configurable so that it is the authoritative name server for the test name. All names underneath the test name, however, also return RCODE 3. A summary of the behavior is in Table 1.

Domain Name	RRTYPE queried	RCODE returned	Actions if any
CandidateTLD	anything except NS	3	
CandidateTLD	NS	0	return server in answer section RDATA
TestName	anything except NS	3	
TestName	NS	0	return server in answer section RDATA
TestName.CandidateTLD	NS	0	return server in answer section RDATA
TestName.CandidateTLD	anything except NS	3	
*.TestName.CandidateTLD	ANY	3	Queries to be measured
*.controlname.ExistingTLD	ANY	3	Queries to be measured
*.TestName	ANY	3	Queries to be measured

4. Evaluation

[TODO align with above]

To evaluate the results, it is only necessary to compare the number of resolution attempts of the test names against the resolution attempts of the control names. If the test name is not in wide private use, then a lookups for the name unique identifier in each name space will arrive nearly as often and at the same time (modulo the difference in the recursive nameserver following the delegation chain) at the instrumented name server.

If the test name is in widespread private use without a search list, or is otherwise resolved locally, then we should expect that lookups inside the test namespace will happen less often than lookups inside the control namespace. If there is no search list in use, then the test QNAMEs of the "b" form will arrive less often than those of "a" form and "c" form. If there is a search list in use, then the "a" form will also not arrive at the authoritative server. If the CandidateTLD is in a search list, we can expect to see duplicated queries of the "b" form on the authoritative server (because the "a" form gets the search list appended).

5. A Basis for Acceptable Behaviour

We assume that there will always be some "stray" queries to the DNS: queries for names that have a TLD-label that does not exist in the root-zone, and which were not intended to be sent to the global DNS. Therefore, it is necessary to establish some baseline level of these "noise" queries, and then use that to evaluate whether some proposed new name for the DNS presents a problem.

Because of the historic prominence of the .com TLD, it may be supposed that .com is, like the root itself, a special zone in which unusual behaviour might be expected. Therefore, names inside the .com name space are a poor guide for "normal" behaviour, and it should not be used for making these sorts of evaluations. The best guide will likely come from using TLDs that are themselves statistically normal.

In addition, overwhelming conservatism suggests that using comparisons with the TLD that is queried least often provides a great margin of safety. As of this writing, a string that is queried less often than that least-queried TLD seems likely not to be in widespread real use, and therefore comparisons with that least-queried TLD are a good conservative choice when evaluating.

6. Possible Experiment Extension

A bias to the measurement is introduced if in certain environments lists of existing TLDs are used in access lists of, for instance, firewalls. In that case queries for the QNAME <uniqueid>-b.TestName.CandidateTLD might be blocked. To calibrate that behavior an additional non-existent TLD could be delegated for the duration of the experiment:

the QNAME is <uniqueid>-d.TestName.RandomTLD

Whereby RandomTLD is a short TLD with the same properties as the candidate TLD. e.g. if the CandidateTLD is U-Label then the RandomTLD is a U-Label from the same script.

If the measurement servers receive queries where the QNAME is neither <uniqueid>-d.TestName.RandomTLD nor <uniqueid>-b.TestName.CandidateTLD then it is likely that all non-delegated domains are blocked. An alternative way of interpreting this is that the queries where the QNAME is <uniqueid>-d.TestName.RandomTLD that arrive at the measurement servers provide a baseline for the transparency of the querying environment for non-delegated TLDs.

7. Security Considerations

The delegation of the Proposed TLD (CandidateTLD) and control TLD (TestName) comes with some risk of interference with existing deployments. The risks for name collision for the TestName is under the control of the experimenter and can be minimized by taking random strings without semantic value.

The risk of name collision with the CandidateTLD is minimized reduced by having the experiment's server returning RCODE=3. Under the assumption of regular DNS implementations that response is indistinguishable from a direct root-response for applications that receive such answer from a stub resolver. The authors have no reason to believe that there are DNS implementations that would hand the applications a different response if a delegation is part of the resolution process.

The authors would advise against signing the various delegated domains, as the introduction of DNSSEC is likely to bias the experiment. At the root domain and ExistingTLDs, regular signing practices, including the inclusion of an NSEC or NSEC3 RR proving the non-existence of a DS record should continue.

The experiment can be biased by 3rd parties through sending queries that have properties like the ones specified herein. The experimentor SHOULD carefully control and record the <uniqueid> values used in the experiment and discard non-expected and non-unique queries that arrive at the nameserver.

8. References

[NETBIOS] IBM Corporation, "LAN Technical Reference: 802.2 and NetBIOS APIs ", Doc. number SC30-3587-01, April 1996.

[RFC0822] Crocker, D.H., "Standard for the format of ARPA Internet text messages", STD 11, RFC 822, August 1982.

[RFC1535] Gavron, E., "A Security Problem and Proposed Correction With Widely Deployed DNS Software", RFC 1535, October 1993.

[RFC3379] Pinkas, D. and R. Housley, "Delegated Path Validation and

Delegated Path Discovery Protocol Requirements", RFC 3379,
September 2002.

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

[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, an borrows heavily from, our earlier (abbandonded) work on "A Procedure for Cautious Delegation of a DNS Names". Discussion of that document in various hallways lead to inspiration for this document and we want to thank those that gave us feed-back.

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

Authors' Addresses

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