Network Working Group   O. Kolkman
Internet-Draft   NLnet Labs
Intended status: Informational   A. Sullivan
Expires: February 2, 2014   Dyn, Inc.
                                                             W. Kumari
                                                          Google, Inc.
                                                           Aug 1, 2013

              A Procedure for Cautious Delegation of a DNS Name
                    draft-kolkman-cautious-delegation-02

Abstract

   NOTE: The authors recognize that the statistical models that would
   inform the process are not well understood and that the possibilities
   to game the system might be unmountable. Unless we reach more
   insights on how to deal with this details this work is abandoned.

   Sometimes, a DNS name is known to be in use in the wild even though
   it was never properly delegated. This situation appears
   particularly, but not only, true in certain domains near the root of
   the tree: people have independently used those non-existent top-level
   domains as private namespaces. If those names are to be delegated in
   the public DNS, prudence dictates that collisions between the private
   uses and the public use be minimized. We outline a procedure to
   evaluate the harm of delegation.

Status of this Memo

Copyright Notice

Table of Contents

1. Terminology

   NXDOMAIN: an alternate expression for the "Name Error" RCODE as
        described in [[RFC1035] Section 4.1.1]. The two terms are used
        interchangeably in this document. (definition from [RFC2308])

   In this document we will be using the terms zone, domain and sub-
   domain. When envisioning the domain namespace as a tree, with nodes
   at the places where the dots separate the labels in a domain name,
   then:

   a 'domain' is an entire branch. e.g. The .org domain is the branch
        of the domain name tree for which all names end in .org.
   a 'sub-domain' is a subordinate namespace of a given domain. e.g.
        all names ending in example.org are in the domain example.org
        which is a sub-domain of .org
   a 'zone' is a piece of the domain space that is under administrative
        control of one party. e.g. the .org zone has delegated the
        example.org domain to the example.org maintainers.

2. Background and Introduction

   DNS names have always co-existed with other namespaces that are
   virtually indistinguishable from the DNS. The DNS was itself
   deployed alongside the host ### table. NetBIOS ### names, though
   only one label long, could always interact with the DNS search path
   mechanism to generate DNS names. Additionally, mDNS [RFC6762] names
   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 such sites. 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 where 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 (although not uncommon).

2.1. 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.

As an (hypothetical) example, Example Widgets 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." and so the users 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 to the incorrect server.

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.

3. Predelegation determination of use of a name

It is possible for the operator of a zone authoritative for some

domain name to tell whether a particular subordinate name has a
widespread use outside the DNS. In order to do this, the operator of
the zone monitors queries against the zone to learn the names for
which there are queries, ignoring those names that actually exist
i.e. those names the zone owner delegated or created resource records
for (in the remainder of this document we will not make the
distinction between entering data with a name or making a delegation;
within the context of this document the same considerations apply).
The operator then establishes a baseline "noise" level of queries for
non-existent subordinate names. Any name that is queried with
significantly greater frequency is to be treated as in widespread
private use, and it should not be released for delegation. The rest
of this section describes the mechanisms for such determination in
detail.

3.1. Predelegation testing is needed

In order for this procedure to be useful, it should be undertaken
before any subordinate names are delegated. Otherwise, it will be
difficult to tell whether a subordinate name is being queried because
it is already delegated or because it is in private use.

At the same time, it is possible that the operator of a zone may wish
to consider the private use of a descendant name, where some
intermediate namespace has been delegated. In that case, it is
necessary to ensure that the descendant name is not actually
delegated when evaluating queries against that name.

3.2. Determining the names of concern

[ED NOTE: This methodology needs to tested. First assessment of
data indicates that this approach may be far to trivial]

There are two modes of operation for determining names of concern.
The most usual is to examine names for which there is no intermediate
delegation. This is useful in case the operator of the zone is
deciding whether to permit delegation or addition of a particular
name. The second, more unusual mode, is to examine subordinate names
inside a sub-domain that has already been delegated. This mode is
useful only as part of a regime of contract enforcement with the
operator of the (already delegated) sub-domain. [WK Note: Are we
sure we even want to address/suggest this second "limited delegation"
option? If we are going to discuss it, referring to it as "limited
delegation" or similar may help clarify. Personally I think 'tis a
silly idea, but... There is talk of doing "test" delegations -
basically launch a TLD / domain with a low TTL. If nothing goes
"boom" then delegate for longer...]

[3.2.1](). Mode 1: Prior to any delegation

The procedure starts with the name of a zone, which is called the "starting domain". In order to determine what subordinate names may be problematic, the starting domain zone operator captures all the names it receives in queries. The operator discards as irrelevant any sub-domain it has already delegated in its namespace. Every other queried name will result in a response of Name Error, RCODE=3 ###STD13 ("NXDOMAIN" ###Negative cache). We call the resulting list the "NX names". (See [Section 4]() for guidance on the sample size.)

The operator then takes the list of NX names, and builds a frequency of queries for each potential delegation point (in practice all immediately subordinate names). The operator proceeds in the fully-qualified domain name ("FQDN") label by label until the next label past the operator's namespace (in practice, these are the names at which delegation will potentially take place). We call these the "target names". The operator counts the number of queries for each target name.

The operator determines the mean and median number of queries over the set of target names. Any name that receives more queries than ###SIGMA -- needs xref to params### greater than the mean, or ###SIGMA2### greater than the median, should be regarded as in widespread private use on the Internet and therefore not a candidate for delegation.

It is possible that only a portion of a namespace subordinate to a target name is actually in private use. It is possible to measure this situation simply by treating the beginning of the namespace in question as the starting domain, and then repeating the procedure above. This could be useful in order to establish baseline restrictions on the operator of a subordinate namespace prior to delegation.

[3.2.2](). Mode 2: After delegation

This mode is more likely to be useful if the evaluation at the end of the previous section has already been performed. In this case, some sub-domain to the operator's zone is to be evaluated for possible private use, where that sub domain has already been delegated. The zone operator operates the "parent starting zone", and is evaluating use inside a starting domain already operated by someone else. The very same mechanisms as are outlined in [Section 3.2.1]() are used, but the evaluation must take into consideration the effects of negative TTLs ### for the starting domain. Because of the combining effects of multiple negative TTLs, it is inadvisable to attempt to perform this evaluation beyond the boundary of a single delegation.

4. Parameters for operation of this procedure

   This section ought to have some words about sane parameters to use
   for the procedure.?

4.1. Median or Mean

   In this section we would like to describe some likely distributions.
   Our assumption is that incoming queries will usually follow some
   dictionary pattern. The 'everybody wants to be Mr. Black'
   [ResevoirDogs] effect is that queries are much more likely for
   popular names than for labels filled with random content. Therefore
   distributions for non-existent names will have relatively little
   power in the long tail. However, the long tail is significant in the
   sense that the names in the long tail are most likely not to exist.

   The exact type of distribution and the statistical parameters that
   signify it is subject for a future version of the draft.

4.2. Discussion of Alternatives

   The above method is based on looking at names that the querying
   population perceives to exist. Alternatively one could count queries
   for a set of random name like "ao42hft3tofj4irsavc4owajhro.example".
   That type of measurement will set the baseline of _real_ non-existing
   names and set the noise level (likely zero queries within a
   reasonable timescale). However, using truly random names introduced
   the problem that any signal (e.g. a handful of queries used for
   probing of availability) will make the domain name unavailable.

5. Security considerations

   Applying this mechanism as the basis for decisions on whether or not
   to delegate domains introduces a motivation for gaming the system.
   The reception of a lot of queries for a particular domain may cause
   it to not be delegated, while the reception of many random queries
   (changing the properties of the query distribution) may cause a
   domain that is in common use to be delegated (by hiding the actual
   use of names in that domain in the noise). Careful analysis of data
   (for example, by studying root for queries, and taking into account
   historical trending) could, in case of suspicion of gaming, help to
   supplement decisions.

6. IANA Considerations

   This document makes no requests of the IANA.

7. Informative References

[RFC1035]   Mockapetris, P., "Domain names - implementation and
            specification", STD 13, RFC 1035, November 1987.

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

[RFC2308]   Andrews, M., "Negative Caching of DNS Queries (DNS
            NCACHE)", RFC 2308, March 1998.

[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>.

Appendix A. Document Editing Details

[To Be Removed before publication]

$Id: draft-kolkman-cautious-delegation.xml 3 2013-05-02 14:27:06Z
olaf $

A.1. version 00

Documenting the first rough outline based on hallway discussions with
the specific purpose to document the idea in the public domain.

$Id: draft-kolkman-cautious-delegation.xml 5 2013-06-11 21:49:28Z
warren $

A.2. version 01

   o Bunch 'o nits.
   o Added section on search-path processing.
   o

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