IETF
Internet-Draft
Intended status: Informational
Expires: August 18, 2014

A. Sullivan
Dyn

February 14, 2014

By Any Other Name: Considerations on DNS, Other Naming Protocols, and the Hierarchical Domain Name Space

draft-sullivan-draft-sullivan-namespaces-and-dns-00

Abstract

You should probably not read this. It's not done.

Not every domain name is intended to appear in the global DNS. It is also possible that not everything that looks like a domain name is intended to be one. Regardless of whether a given name is intended to appear in the DNS, such names often turn up in domain name slots. When choosing a naming scheme that is not intended to be part of the global DNS, it is necessary to understand the architectural implications of using domain names or a domain-name-like syntax.

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 August 18, 2014.

Table of Contents

1. Introduction

Not every domain name appears in the global DNS, and many domain names are not intended for use in the global DNS. At the very least, as [RFC6590] acknowledges, it is only pragmatic to recognize the existence of split-horizon DNS deployments; these often include so-called "private names".

At the same time, as [RFC2826] is at pains to say, the global DNS requires a globally unique public name space. This means that the set of labels near the root of the tree need to be shared by everyone, or the root itself becomes suspect. Private namespaces are fine, but not if they conflict with the public namespace.

Some recent developments have revealed the deep tension in this approach to the namespace. Three factors in particular seem to be important.

First, while historically the DNS root zone was mostly stable and changed quite slowly when it did change, the decision to expand the root zone dramatically does away with that historic stability. More than a thousand new TLDs are expected in the root zone as of this writing. In some cases, the changes make old assumptions about what is or is not a top-level domain obsolete; but in any case, the new root zone management policy makes keeping static lists of TLDs even worse than it used to be.

Second, the arrival of IDNA in the root zone means that protocols that rely on plain UTF-8 labels in the DNS may lead to ambiguous results, if the IDNA version of a zone and the "raw UTF-8" version of a zone become somehow unsynchronized. While from the DNS point of

view, these zones are unrelated to one another, from any user's point of view the zones should be the same thing.

Finally, and perhaps most importantly, a number of protocols have emerged that use either domain names, or else strings that appear to be just like domain names in most contexts, without relying on the public DNS. (There is an argument to be made that in at least some cases these names are not domain names, because they do not actually have the same restrictions. They may not, for example, restrict length to 63 octets per label. For the purposes of the present discussion, this distinction makes no difference.) For the present purpose, the most interesting of these cases are the ones which use so-called pseudo-top-level-domains (pseudoTLDs) to call out that a namespace has been shifted out of the DNS space and into some other protocol.

These different threads suggest that some careful thinking about name spaces is in order. This text, alas, is not that; indeed, at the moment it's little more than a jumble of ill-organized thoughts around this topic, and groping towards some coherence. But I hope to initiate some thoughts about whether domain names, and the domain name system, provide the right foundation for future name space use.

2. PseudoTLDs, Real Names

The idea of a pseudoTLD is fetching. With such a top level domain (TLD), one uses the right-most label of a domain name as a sort of protocol shifter, to indicate that while the namespace is still the same domain name space, the name is not to be looked up in the DNS. This strategy is not novel. For instance, Multicast DNS [RFC6762] uses the "local" TLD as such an indicator. Prior to the registration of local as a Special-Use Domain Name [RFC6761], local was a pseudoTLD.

In the era of a dynamic DNS root zone, the idea that pseudoTLDs can be used safely without co-ordination with the public root zone is a delusion. If a given pseudo use for a TLD takes off, it seems likely that there will be commercial pressures in favour of registration of the pseudoTLD in the root zone (i.e. as a "real" TLD) in an effort to gain automatic traffic. At the same time, some of the pseudoTLDs appear to be attempts to undermine the operation of the existing root either with alternative resolution systems riding atop the DNS (sometimes with claims about additional security or privacy as a concomitant benefit).

None of the issues with pseudoTLDs would matter, except that the expansion of the root zone has itself been fraught with political disagreements, and disputes about the costs of registrations. There

have also been the sorts of commercial tensions apparent whenever a scarce resource is divided up by commercial means.

At least some of the pseudoTLDs could as easily be accommodated in a tree elsewhere in the DNS. These cases would still be candidates for the Special-Use Domain Name registration; it is not clear why these cases need a top-level domain. This issue may be a red herring, however.

3. Is a Common Namespace a Good Idea?

The basic issue that we are facing is not collisions with the DNS root namespace. It is obvious that we could reserve chunks of the DNS namespace for private use in exactly the way number spaces do this (e.g., [RFC6890]). The deeper architectural question is why, if there is such desire to do away with the limitations of the DNS, such systems start from the premise that the DNS and its namespace are the right foundation. Oddly, the design of the DNS would appear to suggest another approach.

Conceptually [RFC1034], the DNS name space is divided up not only by name, but also by class. As a practical matter, on the Internet only the IN class is used. But in principle, the DNS design permits a given namespace to be classed such that the same name could have completely different RDATA at every owner name, for the same RRTYPE, in each of two different classes.

Unfortunately, because of the way CNAME is defined, the DNS classes do not really work. CNAME processing does not restart the processing of a given name in a given class, but rather restarts the processing of a name alone. For that reason, two DNS classes are not really completely separate name spaces.

The resort to a pseudoTLD as a kind of shift bit to indicate a new resolution protocol signifies that what is really wanted is a new resolution class. We have been down this road before. The use of so-called underscore labels in DNS names as a mechanism for subdividing space under a TXT RRTYPE was in effect an attempt to lift the burden of deploying a new RRTYPE. In that case, however, the desire was to publish data in the existing global DNS, without the hassle of making the global DNS work the way it was supposed to (by making it easy to deploy new RRTYPEs).

The present case, of alternative name resolution systems, is different. In this case, new resolution libraries all use the pseudoTLD as a trigger to spring into action. The pseudoTLD isn't there for compatibility with the DNS; some proposals are in fact inimical to the DNS. Instead, the DNS name space pattern is used

because it is familiar. This seems a poor reason to use an old-fashioned naming system that does not even support its own entire feature set. And the fact that the pseudoTLDs are a flag that new resolution systems need to be used suggests that in fact new code to be deployed across systems is acceptable in these cases. If true, that means that the traditional argument in favour of retaining the DNS -- its existing universal deployment -- is lost. After all, if the new naming system does not work except for those who have installed the new system, what reason is there to saddle the new system with the compromises (and long, crufty history) of the DNS?

The above suggests that a better goal would be to undertake the design of an improved naming system, incorporating lessons from the DNS as well as ideas from the new naming technologies and proposal.

4. IANA Considerations

This memo makes no requests of IANA.

5. Security Considerations

The security implications of the foregoing are as yet unknown.

6. Informative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC2826] Internet Architecture Board, "IAB Technical Comment on the Unique DNS Root", RFC 2826, May 2000.

[RFC6590] Falk, J. and M. Kucherawy, "Redaction of Potentially Sensitive Data from Mail Abuse Reports", RFC 6590, April 2012.

[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.

[RFC6890] Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, April 2013.

Author's Address

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

Email: asullivan@dyn.com