Workgroup: DBOUND2
Internet-Draft:
draft-tjw-dbound2-problem-statement-01
Published: 10 July 2023
Intended Status: Informational
Expires: 11 January 2024
Authors: T. Wicinski

Domain Boundaries 2.0 Problem Statement

Abstract

Internet clients attempt to make inferences about the administrative relationship based on domain names. Currently it is not possible to confirm organizational boundaries in the DNS. Current mitigation strategies have there own issues. This memo attempts to outline these issues.

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 https://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 11 January 2024.

Copyright Notice

Copyright (c) 2023 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 (https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction and Motivation

Working off of the earlier problem statement [I-D.sullivan-dbound-problem-statement], which we still consider valid. Various Internet protocols and applications require some mechanism for determining whether two domain names have some relation.

The concept of an administrative boundary is by definition not present in the DNS. Relying on the DNS to divine administrative structure thus renders such solutions unreliable and unnecessarily constrained. For example, confirming or dismissing a relationship between two domain names based on the existence of a zone cut or common ancestry is often unfounded, and the notion of an upward "tree walk" as a search mechanism is, therefore, unacceptable.

Currently, the most well known solution in existence is the Public Suffix List (PSL). The PSL is maintained by and is kept current by volunteers on a best-effort basis. It contains a list of points in the hierarchical namespace at which registrations take place, and is used to identify the boundary between so-called "public" names (below which registrations can occur, such as ".com" or ".org.uk") and the private names (organizational names) that domain registrars create within them. When this list is inaccurate, it exposes a deviation from reality that degrades service to some and can be exploited by others. As the PSL is the de-facto resource, and as there is not a more comprehensive, alternative solution for relationship identification, the PSL has often been misused to accomplish things beyond its capabilities. For example, there is no way to confirm the relationship between two domain names -- the PSL may only signal that there is or is not a public boundary between the two. Additionally,

there are questions about the scalability, central management, and third-party management of the PSL as it currently exists.

Applications and organizations impose policies and procedures that create additional structure in their use of domain names. This creates many possible relationships that are not evident in the names themselves or in the operational, public representation of the names.

(This document is currently being edited at https://github.com/moonshiner/draft-tjw-dbound2-problem-statement)

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here. DNS terminology is as described in [RFC8499].

3. Simplifying the list of possible Use Cases

A main topic that immediately arises from this discussion is the replacement of the Public Suffix List (PSL). Currently, this document is not looking at the problem space with regards to it.

From the previous problem statement, the one use case which

3.1. HTTP State management cookies

*Cookies that have the same-site in browsers. For example:

allowing www.example.com to respond with a set-cookie for example.com so *.example.com will work.

not allowing example1.co.uk to respond with a set-cookie for example2.co.uk.

*CA wildcards: it's OK to sign a cert for *.example.co.uk or *.example.com but not for .co.uk or *.com.

Other applications and organizations impose policies and procedures that create additional structure to express many possible relationships, such as in first-party-sets or IDN-UA. These are not always evident in the names themselves, and any solutions developed here may or may not suit these existing policies and procedures.

3.2. Service Boundaries in shared cloud environments

A public cloud provider may have a large number of boundaries they need to publish. Taking an example resource within an example service, that resource might have a domain name that follows the below pattern (2LD may vary):

"resource-id.cluster-id.servicename.region-name.example.com"

With the current PSL, the need may exist to publish 1+ records per region per service. If there are many services across many more regions, these updates do not scale. Putting this information into DNS allows us to publish records, without manual updates.

3.3. Network resource boundaries in shared cloud environments

Network suffixes and resouces that are public, but not routable or resolvable outside of one's network (public-like-PSL, not public-like-pool). A good example is an internal zone within a virtual private clouds (VPCs).

VPCs are resources customers can create, which reserves them a logically-isolated portion of AWS's network. Within each VPC, there is a VPC-internal DNS zone which contains DNS records for resources within the VPC (suffix zone of "compute.internal" or equivalent), which is separated per-VPC. Every network interface in a VPC will have an associated per-interface record in this zone (for VPCs using a default configuration).

These customers may chose to peer their VPC with another customer, and these VPC-internal zone would now have multiple different customers operating within it.

*Linking domains together

In terms of specific use cases, within the realm of email there is a desire to link an arbitrary fully-qualified domain name (FQDN) to the organizational domain name (at some point in the namespace above it), in order to identify a deterministic location where some sort of statement of policy regarding that FQDN can be found.

There is another growing use case within organizations that need to identify relationships between different FQDNs, but also different top level domains. However, there is also desire to reliably identify relationships outside of the realm and constraints of the namespace tree.

4. Security Considerations

None at this time.

5. IANA Considerations

None at this time.

6. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

7. Informative References

[I-D.sullivan-dbound-problem-statement] Sullivan, A., Hodges, J., and J. R. Levine, "DBOUND: DNS Administrative Boundaries Problem Statement", Work in Progress, Internet-Draft, draft-sullivan-dbound-problem-statement-02, 18 February 2016, <https://datatracker.ietf.org/doc/html/draft-sullivan-dbound-problem-statement-02>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

Appendix A. Acknowledgements

The author leans heavily on the initial problem statement and thanks Andrew Sullivan, John Levine, Murray Kucherawy and Paul Vixie for comments and suggestions.

Appendix B. Appendix

Acknowledgements

Author's Address

Tim Wicinski
Elkins, WV 26241
United States of America

Email: tjw.ietf@gmail.com