

Network Working Group
Internet-Draft
Intended status: Informational
Expires: October 29, 2019

J. Levine
Taughannock Networks
April 27, 2019

Publishing Organization Boundaries in the DNS
draft-levine-dbound-dns-03

Abstract

The organization that manages a subtree in the DNS is often different from the one that manages the tree above it. We describe an architecture to publish in the DNS the boundaries between organizations that can be adapted to various policy models and can be queried with a small number of DNS lookups.

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 October 29, 2019.

Copyright Notice

Copyright (c) 2019 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 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	2
2.	Design Issues	2
3.	TXT record format	3
4.	Lookup Process	4
5.	DNS Records	5
6.	Application scenarios	7
6.1.	DMARC	7
6.2.	Cookies	7
6.3.	SSL Certificates	7
7.	Discussion	7
8.	ABNF syntax of bound records	8
9.	Security Considerations	8
10.	Variations	8
11.	IANA considerations	9
12.	References	10
12.1.	Normative References	10
12.2.	Informative References	10
Appendix A.	Change Log	10
	Author's Address	11

[1.](#) Introduction

Often, the organization that manages a subtree in the DNS is different from the one that manages the tree above it. Many applications use information about such boundaries to implement security policies. For example, web browsers use them to limit the names where web cookies can be set, and Secure Socket Layer (SSL) certificate services use them to determine the party responsible for the domain in a signing request. Mail security applications such as Domain-based Messaging Authentication, Reporting and Conformance (DMARC) use them to locate an organization's policy records in the DNS. This specification is intended to provide boundaries usable for DMARC, and possibly for other applications.

[[Please direct discussion of this draft to the dbound working group at dbound@ietf.org.]]

[2.](#) Design Issues

Organization boundaries can be assigned on what one could call an opt-in or opt-out basis. "Opt-in" means that two names are only managed by the same organization if both actively assert that they are related. "Opt-out" means that if there is any boundary information at all for a DNS subtree, each name is assumed to be under the same management as its parent unless there is a boundary assertion to the contrary. This design describes an opt-out model.

Levine

Expires October 29, 2019

[Page 2]

Within the opt-out model, this design can adapt to a variety of scenarios:

- o Policies can be published by the domains themselves, or by a third party. In the former case, each domain might assert its own boundary policies. In the latter case, the third party makes the assertions, which may or may not agree with what the domains themselves would want.
- o Multiple levels of delegation may be implemented, which is different from irregular boundaries. For example, "ca", "on.ca", and "toronto.on.ca" are irregular boundaries, because they're all handled by the Canadian Internet Registration Authority (CIRA). CentralNIC's "uk.com" would be a second level of delegation below Verisign's com.
- o Different sets of boundary rules can be published for different applications. For example, the boundaries for SSL certificates might be different from the boundaries for e-mail policies, or for web cookie setting policies.

In the lookup process below, the boundary point data is stored in the DNS tree in a TXT record. The boundary is considered to be directly below the name that the process returns, similarly to the names in the PSL [[PSL](#)]. If the process returned "abc.example", then "foo.abc.example" and "bar.abc.example" are separated by the boundary, but "foo.abc.example" and "bar.foo.abc.example" are not.

Each domain can publish its own policies within its own domain name space, or a separate authority can publish a global set of policies in a separate name space.

3. TXT record format

```
*.toronto.on._bound.ca IN TXT "bound=1 . . toronto.on.ca"
```

The "bound=1" tag is to prevent confusion when a domain publishes a wildcard such as *.example.com that could match a _bound name.

The bound record contains a tag, two keyword strings and a domain name.

Each keyword string is a series of comma separated keywords. If the string would otherwise be empty, it is a single dot.

The first keyword string expresses policy options. It can include NOLOWER which means that no lower level boundaries can exist below

this one, and NOBOUND which means that this name is not a boundary for this application.

The second keyword string identifies the application(s) to which this boundary applies. The strings DMARC, COOKIE, and CERT mean that the applications re DMARC, HTTP cookies, and SSL certificate signing respectively; a dot means it is a default for any applications not otherwise specified.

The domain name is an absolute domain name, without the final dot. The first label in the name may be "*" to indicate that the boundary is at any name one label deeper than the rest of the name. That is, the asterisk matches a single label, not the usual DNS sense of matching any string of labels.

4. Lookup Process

In general, the lookup process takes as input a domain name and application. It returns the name of the boundary node in the DNS. This may be the domain itself or a parent. If there is no policy for the domain the lookup fails; there are no defaults, and the DNS root is not within any organization boundary. (Applications may apply defaults of their own, but that is beyond the scope of this specification.)

Names of boundary information records use the tag "_bound" which is intended to be unique.

For the first lookup, the client extracts the top level component (i.e., the rightmost label, as "label" is defined in [Section 3 of \[RFC1034\]](#)) of the domain name from the subcomponents, if any, and inserts the prefix in front of that component, after other components if any. For example, if the domain to be checked is "example.com" or "www.example.com", the client issues a DNS query for "example._bound.com" or "www.example._bound.com". If the domain is a dotless one such as "example", the client looks up "_bound.example".

The client does a DNS lookup of TXT records at that name, which will return zero or more TXT records. A failure such as NXDOMAIN is considered to return zero records. A lookup can return multiple records if different applications have different boundaries or policy options. The lookup process discards any records that do not start with "bound=1".

If a relevant policy record is returned, and the record does not contain the NOBOUND keyword, the domain name in the record is the policy boundary. A policy record is relevant if it lists the desired application, or it is a default policy and there is no record with

the application's keyword. For example, a check for a boundary above "example.com" would be issued at "example._bound.com", and the expected TXT record could be "bound=1 . . com".

If there are no boundaries below the queried point, the policy record contains "bound=1 NOBOUND . ." indicating the root. For example, if all subdomains of the "example" top-level domain (TLD) are under the same management as the TLD itself, checks for "_bound.example" or "www._bound.example" would return "bound=1 NOBOUND . .".

If the relevant record has the NOLOWER keyword set, the process stops. Otherwise, the client inserts the prefix tag into the name just below (i.e., to the left of) the name at the largest matching boundary indicated in the lookup result, and repeats the lookup. For example:

- o When evaluating "www.foo.example.com", the first query would be to "www.foo.example._bound.com". If the reply to this is "bound=1 . . com", then the second query would go to "www.foo._bound.example.com".
- o When evaluating "www.example.on.ca", the first query would be to "www.example.on._bound.ca". If the reply to this is "bound=1 . . on.ca", the next lookup would be to "www._bound.example.on.ca".

This process repeats until a DNS lookup returns a relevant record with the NOLOWER keyword, or a lookup returns no relevant records, at which point the boundary is the domain name in the last retrieved relevant record.

If an otherwise relevant record has the NOBOUND keyword, the process continues if the NOLOWER keyword is not present, but there is no boundary at the name with the NOBOUND keyword. These NOBOUND keyword enables a name in the hierarchy to be a boundary for some applications but not for others.

5. DNS Records

The publishing entity uses wildcards and prefixed names that parallel the regular names under a TLD to cover the domain's name space.

If there is a boundary at a given name, an entry in the TLD record covers the names below it. For example, if there is a boundary at ".TEST", a suitable record would be:

```
*._bound.test IN TXT "bound= . . test"
```


If the boundary is above the TEST domain, i.e., TEST is under the same management as FOO.TEST, the record would indicate no boundaries, and an additional non-wildcard record is needed to cover TEST itself:

```
*._bound.test IN TXT "bound=1 . . ."  
_bound.test   IN TXT "bound=1 . . ."
```

In domains with irregular policy boundaries, multiple records in the record describe the boundary points. For example, in the CA (Canada) TLD, for national organizations there might be a boundary directly below the national TLD; for provincial organizations there might be a boundary below a provincial subdomain such as "on.ca"; and for local (e.g., municipal) organizations, a boundary below a municipal subdomain such as "toronto.on.ca" might exist. A suitable set of records covers this structure. The closest encloser rule in [RFC 4592](#) [[RFC4592](#)] makes the wildcards match the appropriate names.

```
*._bound.ca           IN TXT "bound=1 . . ca"  
*.on._bound.ca        IN TXT "bound=1 . . on.ca"  
*.toronto.on._bound.ca IN TXT "bound=1 . . toronto.on.ca"
```

In some cases, a domain may assert that every name below a given name is a boundary, or that every name other than a specific set of exceptions is a boundary. For example (adapted from the Mozilla PSL) every name below kobe.jp is a boundary other than city.kobe.jp. This could be expressed as:

```
*.kobe._bound.jp      IN TXT "bound=1 . . *.kobe.jp"  
city.kobe._bound.jp   IN TXT "bound=1 NOBOUND . city.kobe.jp"  
*.city.kobe._bound.jp IN TXT "bound=1 NOBOUND . city.kobe.jp"
```

For any set of policy boundaries in a tree of DNS names, a suitable set of policy records can describe the boundaries, so a client can find the boundary for any name in the tree with a single policy lookup per level of delegation.

Since the delegation structure is unlikely to change frequently, long time-to-live (TTL) values in the TXT records are appropriate.

If different applications have different boundaries or policy options, the policy records for each application are put at the appropriate names for the boundaries. Due to the way DNS wildcards work, each name with any policy records MUST have records for all policies, with the NOBOUND bit for policies for which the name is not in fact a boundary.

6. Application scenarios

Here are some ways that DMARC and potentially other applications can use BOUND data.

6.1. DMARC

If a DMARC lookup for the domain in a message's From: header fails, the client would do a boundary check for the domain name using the "DMARC" application. The organizational domain is the immediate subdomain of the boundary domain. (Note that the boundary will always be the one looked up or an ancestor.)

6.2. Cookies

If an http request attempts to set a cookie for a domain other than the request's own domain, the client would do boundary check for a "cookie" application for both the request's domain and the cookie domain. If they are not separated by a boundary, the request is allowed.

6.3. SSL Certificates

The client would do a boundary check for the domain name in a normal certificate, or the name after the "*" in a wildcard certificate for a "cert" application. If the boundary is above the name, the name is allowed.

7. Discussion

The total number of DNS lookups is the number of levels of boundary delegation, plus one if the last boundary doesn't have the NOLOWER keyword. That is unlikely to be more than 2 or 3 in realistic scenarios, and depends on the number of boundaries, not the number of components in the names that are looked up.

Some domains have very irregular boundaries. This may require a relatively large number of records to describe all the boundaries, perhaps several hundred, but it doesn't seem like a number that would challenge modern DNS servers, or need unduly complex scripts to create them.

The wildcard lookup means that each time an application looks up the boundaries for a hostname, the lookup results use DNS cache entries that will not be reused other than for subsequent lookups for the identical hostname. This might cause cache churn, but it seems at worst no more than we already tolerate from DNSBL lookups.

8. ABNF syntax of bound records

The syntax of bound records is something like this:

```
BOUND = "bound=1" WSP BFLAGS WSP BKWDS WSP DOMAIN
```

```
BFLAGS = ( BFLAG *("," BFLAG) ) / "."
```

```
BFLAG = "NOLOWER" / "NOBOUND"
```

```
BKWDS = ( BKWD *("," BKWD) ) / "."
```

```
BKWD = "DMARC" / "COOKIE" / "CERT"
```

9. Security Considerations

The purpose of publishing organization boundaries is to provide advice to third parties that wish to know whether two names are managed by the same organization, allowing those names to be treated "as the same" in some sense. Clients that rely on published boundaries are outsourcing some part of their own security policy to the publisher, so their own security depends on the publisher's boundaries being accurate.

Although in some sense domains are always in control of their subdomains, there are many situations in which parent domains are not expected to influence subdomains. For example, second level domains in global TLDs (gTLDs) operated by registries with contracts with the Internet Corporation for Assigned Names and Numbers (ICANN) Since there is no technical bar to a parent publishing records that shadow part or all of the boundary record namespace for delegated subdomains, correct operation depends on the parent and subdomains agreeing about who publishes what.

The DNS is subject to a variety of attacks. DBOUND records are subject to the same ones as any other bit of the DNS, and the same countermeasures, such as using DNSSEC, apply.

10. Variations

Some boundary schemes distinguish between public and private subtrees. If that were useful, a PUBLIC flag keyword could indicate that the subtrees below a boundary were public rather than the default of private.

Since nothing but the boundary records should be published at names with `_bound` components, one could get the same effect with a new

DBOUND RRTYPE, which would avoid the problem of confusion with other TXT wildcards.

If third parties want to publish boundary information, they could do it in their own subtree of the DNS. For example, if policy.example were publishing boundary information about boundaries, the records for the test domain described above would be:

```
*._bound.test.policy.exaple IN TXT "bound=1 . . ."
_bound.test.policy.example  IN TXT "bound=1 . . ."
```

11. IANA considerations

This document defines a new _bound prefix keyword.

This document requests that IANA create a registry of dbound Flag keywords. Its registration policy is IETF Review. Its initial contents are as follows. [[NOTE: new flags are likely to change the lookup algorithm]]

Keyword	Reference	Description
NOLOWER	(this document)	No lower level policies
NOBOUND	(this document)	No boundary at this name

Table 1: BOUND Flag Keywords Initial Values

This document requests that IANA create a registry of BOUND Application keywords. Its registration policy is First Come First Served. Its initial contents are as follows. [[Note: New applications don't affect the lookup process, and shouldn't affect existing applications.]]

Value	Reference	Description
.	(this document)	Any application without a specific boundary record
DMARC	(this document)	DMARC organizational domains
COOKIE	(this document)	HTTP cookies
CERT	(this document)	Owner of certificate requests

Table 2: BOUND Applications Initial Values

12. References

12.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](https://www.rfc-editor.org/info/rfc1034), DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC4592] Lewis, E., "The Role of Wildcards in the Domain Name System", [RFC 4592](https://www.rfc-editor.org/info/rfc4592), DOI 10.17487/RFC4592, July 2006, <<https://www.rfc-editor.org/info/rfc4592>>.

12.2. Informative References

- [PSL] Mozilla Foundation, "Public Suffix List", Nov 2015.

Appendix A. Change Log

NOTE TO RFC EDITOR: This section may be removed upon publication of this document as an RFC.

02 to -03 Add wildcard labels like in the PSL.

01 to -02 Make TXT record the proposal, new RR as alternative.

-00 to -01 Editorial changes to limit standard use to DMARC.

non-WG to -00 Add NOBOUND record to make wildcard matches do the right thing

Rename to match WG name

Author's Address

John Levine
Taughannock Networks
PO Box 727
Trumansburg, NY 14886

Phone: +1 646 481 7726
Email: standards@taugh.com
URI: <http://jl.ly>