Network Working Group   W. Hardaker
Internet-Draft   USC/ISI
Intended status: Standards Track   November 02, 2020
Expires: May 6, 2021

DNS Private Namespace Options
draft-hardaker-dnsop-private-namespace-options-00

Abstract

This document discusses the trade-offs between various options about creating a private namespace within top level domains within the root zone.

Table of Contents

1. Introduction

HATS: The author is not wearing any hats while writing this document.

Deployed DNS clients within the Internet typically communicate with upstream resolvers using their own in-application stub resolver. These upstream resolvers may be run by ISPs, or may be a customer-premises equipment (CPE) that may or may not forward requests to its upstream ISP.

In an entirely singular Internet DNS there would be no name collisions as all data is uniquely named. However, the prevalence of local private name spaces within companies, organizations, governments, home LANs, etc have shown that existence of a single, unique naming system rarely exists. The deployment of Internet of Things (IoT) devices is only accelerating this trend for private namespaces by devices that bootstrap their names with the easy solution of "just make one up until the customer provides us with a better one", followed by the customer never providing one. This document makes no judgment on whether this is right or wrong, and takes this assumption as simply the state of the current world.

The for special use names is well spelled out in [RFC6761]. [RFC8244] provides additional insight into areas that are still under discussion and where work is needed. Recently ICANN's SSAC has issued [SAC113] entitled "SSAC Advisory on Private-Use TLDs", wherein it suggests the creation of a private-use DNS TLD.

This document considers the aspects associated with DNSSEC and the potential choices for a private-use TLD (also see [RFC8244] bullet 21). Specifically, we consider the case where somewhere in the resolution path DNSSEC validation is in use, potentially at an end-device (phone, laptop, etc), a CPE, or at an ISP's resolver.

1.1. Document state

This document is not a fully complete analysis, but rather a starting point for discussion and continued analysis by both the author and others that wish to contribute.

1.2. Requirements notation

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.

2. Analysis of choices

Note that this analysis is not (yet) exhaustive. It does describe some of the differences in the two approaches.

2.1. Assumptions

We make the following assumptions to begin:

1. A local environment needs to use both the Global Internet's DNS (GID), as well as at least one private name space as well.

2. A end-device, a CPE and/or a resolver may choose to validate DNS requests.

3. The validating resolver wishes to validate both responses from the GID as well as local names using DNSSEC.

4. The validating resolver will, thus, need Trust Anchors (TAs) for both the GID and all private namespaces, or will need a list of names which are assumed insecure and exemptions to the GID.

5. The device may (unfortunately) move to another network where private namespace resolution is not available, and thus queries to it will leak to the GID. This is extremely common today.

6. We take as accepted consensus that anything protocol needing a private name space that is not user visible can be properly housed under .arpa. This document assumes a private-namespace TLD is needed, as discussed in other documents ([SAC113, etc]) to aid in user presentation and understanding. This document does not make judgment on whether this or user-education may be the right approach to this problem.

2.2. TLD choices

Given these assumptions, we consider the cases where a private namespace TLD exists that is:

1. Is a special-use domain per [RFC6761], and does not (and will never) exist in the GID. In this document, we refer to this as ".internal" for discussion purposes only following conventions in [draft-wkumari-dnsop-internal].

2. Is an unsigned delegation within the (GID's) DNS root, with NS records likely pointing eventually to something like 127.0.53.53. In this document, we refer to this as ".zz" following convention in [draft-ietf-dnsop-private-use-tld]. We note that [draft-ietf-dnsop-alt-tld] also proposed a private namespace (".alt") that also fits into this category.

This document recognizes that .zz itself is actually not necessarily a normal special use domain, and [RFC6761] may not apply as its an ISO reversed name. However, in other aspects it will behave like a special-use registered domain and its under current consideration by dnsop so we leave it in here as the example name.

In summary:

- .internal is an unsigned TLD

- .zz is a special-use-like TLD that MUST never be assigned

2.2.1. Working state aside

The next two sections mix together DNSSEC validation at end-devices and resolvers; it would add significant more clarity to discuss them individually, which will be done in a future version.

2.2.2. Analysis of an unsigned TLD (eg .internal)

An unsigned TLD such as .internal will:

- o Exist within the DNS root

- o have NS records pointing to something.arpa with on A/AAAA resolution

2.2.2.1. non-validating end-devices querying within .internal will:

- o inside the private network the client will:

  * Believe the upstream resolver's responses

- o outside the private network the client will:

  * Believe the upstream resolver's NXDOMAIN responses for anything deeper than .internal itself (IE, api.example.internal/A will return NXDOMAIN)

2.2.2.2. validating end-devices querying within .internal will:

- o inside the private network the client will:

  * must be configured with a private TA to enable DNSSEC within the private network (creating an island of trust)

  * If unconfigured, it will believe the upstream resolver's responses because its delegated insecure, and therefore has no basis to distrust the answers

- o outside the private network the client will:

  * if not configured with a TA, all answers to .internal will either be NXDOMAIN or spoofable

  * if configured with a TA, all answers will be detected as BOGUS

2.2.3. Analysis of a special-use TLD (eg .zz)

A special-use TLD will:

- o Not exist within the DNS root

- o Proven by the root's NSEC chain

2.2.3.1. non-validating end-devices querying within .zz will:

o inside the private network the client will:

* Believe the upstream resolver's responses

o outside the private network the client will:

* Believe the upstream resolver's NXDOMAIN or spoofed answers for all data within the .zz domain.

2.2.3.2. validating end-devices querying within .zz will:

o inside the private network the client will:

* with an upstream resolver

* self-resolving:

+ needs a configured TA or a configured negative trust anchor

+ possibly automatically obtained configuration with a bootstrapping mechanism, or-preconfigured in a ROM image

o outside the private network the client will:

* if not configured with a TA, all answers to .internal will either be NXDOMAIN or spoofable

* if configured with a TA, all answers will be detected as BOGUS

3. Other considerations

3.1. a unsigned delegated domain - .internal

o configuration of new TAs

o requires collaboration between the IETF and ICANN , since the TLD will exist and falls outside the scope of [RFC6761]. This process can be slow.

3.2. a special-use domain - .zz

o May require invoking [RFC6761] (depending on .zz or not .zz)

o may require more configuration per-device

4. Deployment considerations

During initial deployment of either of these, there is a fundamental difference for validating resolvers.

Specifically, until all validating resolvers are updated with a new TA for specific instances under a special-use TLD (e.g. .zz), the resolvers will fail to validate any names underneath as .zz is provable insecure. This could take a while to update all deployed validating resolvers.

On the other hand, deploying a newly allocated, unsigned TLD will take a long time in process both within the IETF and within ICANN.

And each may have impacts on what error processing results, based on the differing resolution characteristics (Section 2.2.2, Section 2.2.3).

5. Recommendation

This author recommends that the IETF take on both tracks simultaneously, and:

1. starts the process of communicating with ICANN and ISO about the use of .zz, or selects another name to use under [RFC6761] as a special-use name.

2. Issues a request to the ICANN board via the IAB to follow the guidance of [SAC113] and reserve a string or set of strings for use as a private-namespace(s) as an unsigned TLD. The ICANN board can not act on their own, based on ICANN bilaws, but can take requests from the IETF via the IAB to act.

This leaves vendors the freedom to chose that path that best meets their specific requirements. Recommendations about how to best select one given their situation is hinted above, but should be more formally written down in this document or others.

5.1. Selecting good TLD names

Unfortunately, here be dragons. Selecting a good name has been discussed multiple times in the IETF, and has always resulted in a lack of consensus. In part, this is because the IETF doesn't have the skillsets needed to hold a discussion about what language element(s) would be best for universal adoption and usage.

Instead, this author recommends that we direct ICANN to select the names that should be used for both of these cases. ICANN has

significant more skill breadth in the area of selecting names best suited to be understood by end-users. This discussion will not be faster, however, within ICANN but this author believes a resolution in that SDO will be more likely successful.

Thus, the IETF can make the technical recommendation and ICANN can implement these two choices.

6. Security Considerations

TBD

(though much of this draft is a security considerations itself)

7. IANA Considerations

TBD

8. References

8.1. 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>.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <https://www.rfc-editor.org/info/rfc6761>.

8.2. Informative References

[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>.

[RFC8244] Lemon, T., Droms, R., and W. Kumari, "Special-Use Domain Names Problem Statement", RFC 8244, DOI 10.17487/RFC8244, October 2017, <https://www.rfc-editor.org/info/rfc8244>.

Appendix A. Acknowledgments

Large portions of the technical analysis in this document derives from a discussion with Roy Arends and Warren Kumari (back when we could stand in front of a whiteboard together).

Author's Address

Wes Hardaker
USC/ISI

Email: ietf@hardakers.net