

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 8, 2016

T. Hardie
March 07, 2016

Considerations for establishing resolution contexts for Internet Names draft-hardie-resolution-contexts-02

Abstract

If we model the system of Internet names as a set of directed graphs in an absolute naming context, following [RFC 819](#), an Internet name is not necessarily a name in the domain name system, but is simply a unique name associated with a particular directed graph. The resolution of the name, in other words, is independent from it being an "Internet name". The DNS is a common, but not the only, resolution context for Internet names. This document discusses the consequences of the need to select among multiple resolution contexts.

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 September 8, 2015.

Copyright Notice

Copyright (c) 2015 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 (<http://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.

1. Introduction

The history in [[I-D.lewis-domain-names](#)] and the usage in [[RFC3986](#)] both suggest that names registered in the domain name system are part of a larger set of Internet names. If we model the system of Internet names as a set of directed graphs in an absolute naming context, following [RFC 819](#) [[RFC0819](#)], an Internet name is not necessarily a name in the domain name system, but is simply a unique name associated with that particular directed graph. The resolution of the name, in other words, is independent from it being an "Internet name". The DNS is a common, but not the only, resolution context for Internet names.

2. Resolution Contexts

The Domain Name System [[RFC1034](#)] [[RFC1035](#)] provides the most common resolution system for Internet names by many orders of magnitude. It has not, however, met all resolution requirements. Multicast DNS [[RFC6762](#)] uses an alternative resolution service, as does TOR [[TOR](#)]. Tor's .onion names, in particular, appear to be effectively Internet names within a globally shared naming context; they simply happen to use an alternative resolution method.

The key practical question that follows from the existence of alternative resolution contexts is how you can determine what resolution context to use for a particular Internet name. Practically, this often means starting with the question of whether it is part of the Domain name set of Internet names, or part of a different set. The de facto signal we are using now is the top-most label of the Internet name. If it is within the known set of DNS top-most labels, we have a definite yes. If it is within an established set of non-DNS top-most labels, we have a definite no. For those with a definite no, there is an available registry set up by [[RFC6761](#)] to identify the alternative resolution context or to note that there is no resolution context (as is the case for example domains).

There are at least two unfortunate sets of potentially conflicting

cases, where people are using labels with the intent to use this signal but have not risen to the level of "established no". In the first case, their usage may be mistaken for non-fully qualified names within the domain name system, resulting in the construction of a new Internet name where one was not intended (e.g. `www.sld.allium`

becoming `www.sld.allium.corp.example.com`, rather than `.allium` being used as signal that this Internet name is not within the set of domain names). The second case, which may overlap, is one in which the growth of the set of names in the domain name system causes overlap (a new gTLD like `.allium` being assigned would conflict with the attempted use of `.allium` as a resolution context signal).

The risks of the two conflicting cases are pretty obvious, but despite that the use of a pseudo-TLD signal seems desirable to many setting up alternative resolution contexts. It seems likely that this is because the services within the alternative resolution contexts wish to use protocols defined for DNS names as if they were defined for their Internet names. The `.onion` example was driven, in other words, at least in part because its users wanted <https://identifier.onion/> to work. In order to share the HTTPS URI context, they needed to minimize the changes to the form of the URI. That meant using `https://` with a resolution trigger, rather than changing the URI (`tor-https://`, for example).

The implication for the universe of architecturally appropriate responses is that any means for signalling that a name is not within the DNS context but is still meant to be an Internet name must continue to allow those Internet names to be used in common protocol contexts. It also means that any Internet name must expect restrictions to achieve that (viz. it must be a unique name within a directed graph within the overall Internet name namespace).

3. Available Alternatives

Given that restriction, the universe of possible resolution context signals seems to be limited. One option is using a designated sub-tree of the Internet namespace for non-DNS resolutions, with labels within the tree indicating which resolution context is meant. [\[I-D.ietf-dnsop-alt-tld\]](#) describes one specific approach to this option. While the use of this sub-tree may be esthetically less pleasing than a pseudo-TLD, it avoids the ambiguities which may arise

during the development of alternative resolution context.

A second alternative is to fix either the set of top-level domains or the number of resolution contexts, so that ambiguity cannot occur. While a fixed set of top-level domains might have seemed practical when the number of TLDs was limited to country codes and a strictly limited set of generic top-level domains, this has ceased to be a practical alternative. Similarly, the creation of alternative resolution contexts cannot be effectively stifled, even were this desirable; those interested can implement and deploy them without registration of any kind. That these may not interoperate or conflict with other deployments is, of course, a risk.

A third alternative within the DNS context is to continue the current registration of pseudo-TLDs and accept the consequences of ambiguity. This will mean that conflicts between pseudo-TLDs marking alternative resolution contexts and potential future TLDs must be managed and that the operational impact must be addressed. A focus on deployment of mitigation strategies may reduce the operational consequences. As an example, the deployment of loopback root zones [[RFC7706](#)] will reduce the impact of queries for pseudo-TLDs leaking to the root DNS name servers. Similarly, policies for names registered as pseudo-TLDs may also limit potential conflict.

An alternative to signals within the DNS is making alternative signals easier. URI registrations have gotten significantly easier [[RFC7595](#)] over time, but it might be possible to lower the bar further by creating a convention for using alternative resolution contexts.

As an example, we could set aside a string delimiter for this purpose as we set aside xn- to single out the ACE encoding for Internationalized Domain Names [[RFC5891](#)]. That string delimiter could then be used to construct faceted URI schemes, one aspect of which contained the usual protocol indicator and the other the resolution context. The ABNF for scheme is:

```
scheme = ALPHA *( ALPHA / DIGIT / "+" / "-" / "." )
```

Setting aside a string delimiter such as ++ would allow something like [https://identifier.onion/](#) to become [https++.tor//identifier/](#). This would require updates to URI parsing libraries that intended to

handle alternative resolution contexts, but the use of a common delimiter would lower the amount of code needed both to identify the core protocol and the alternative resolution contexts. It might remain esthetically less pleasing, however, and it would prevent the use of IDNA-permitted characters as resolution context identifiers, something which the DNS-based solutions do allow.

[4.](#) Conclusions

There are clearly trade-offs among the available alternatives, as each has its own drawbacks as an indicator of resolution context. Given, however, that the existence of multiple signals could generate even further interoperability issues and operational concerns, the creation of multiple signals is undesirable. Any system which allows Internet names from alternate resolution contexts to be used in common protocol systems can likely be made to work, provided its drawbacks are accounted for and mitigated appropriately.

Hardie

Expires September 8, 2015

[Page 4]

Internet-Draft Resolution-Contexts-for-Internet-Names

March 2015

[5.](#) Security Considerations

This document describes a number of potential method for establishing a resolution context for an Internet name. Should the resolution context to be used with a name not be sufficiently clear, it may be possible to provide alternative information in a different context. That alternative information could provide an avenue for an attacker to stand up services which would mimic those present elsewhere, allowing the attacker to subvert the connection, steal credentials,

[6.](#) IANA Considerations

This document currently has no actions for IANA.

[7.](#) Acknowledgements

Thanks to Ed Lewis, Suzanne Wolff, and Andrew Sullivan for conversations leading up to this document; all errors of fact and judgement are, however, the author's.

[8.](#) Informative References

- [TOR] The Tor Project, "Tor", 2013,
<<https://www.torproject.org/>>.
- [RFC0819] Su, Z. and J. Postel, "The Domain Naming Convention for Internet User Applications", [RFC 819](#),
DOI 10.17487/RFC0819, August 1982,
<<http://www.rfc-editor.org/info/rfc819>>.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities",
STD 13, [RFC 1034](#), DOI 10.17487/RFC1034, November 1987,
<<http://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification",
STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66,
[RFC 3986](#), DOI 10.17487/RFC3986, January 2005,
<<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC5891] Klensin, J., "Internationalized Domain Names in Applications (IDNA): Protocol", [RFC 5891](#),
DOI 10.17487/RFC5891, August 2010,
<<http://www.rfc-editor.org/info/rfc5891>>.

Hardie

Expires September 8, 2015

[Page 5]

Internet-Draft Resolution-Contexts-for-Internet-Names

March 2015

- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names",
[RFC 6761](#), DOI 10.17487/RFC6761, February 2013,
<<http://www.rfc-editor.org/info/rfc6761>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", [RFC 6762](#),
DOI 10.17487/RFC6762, February 2013,
<<http://www.rfc-editor.org/info/rfc6762>>.
- [RFC7706] Kumari, W. and P. Hoffman, "Decreasing Access Time to Root Servers by Running One on Loopback", [RFC 7706](#),
DOI 10.17487/RFC7706, November 2015,
<<http://www.rfc-editor.org/info/rfc7706>>.
- [RFC7595] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", [BCP 35](#),

[RFC 7595](#), DOI 10.17487/RFC7595, June 2015,
<<http://www.rfc-editor.org/info/rfc7595>>.

[RFC7686] Appelbaum, J. and A. Muffett, "The ".onion" Special-Use Domain Name", [RFC 7686](#), DOI 10.17487/RFC7686, October 2015, <<http://www.rfc-editor.org/info/rfc7686>>.

[I-D.ietf-dnsop-alt-tld]
Kumari, W. and A. Sullivan, "The ALT Special Use Top Level Domain", [draft-ietf-dnsop-alt-tld-03](#) (work in progress), September 2015.

[I-D.lewis-domain-names]
Lewis, E., "Domain Names", [draft-lewis-domain-names-02](#) (work in progress), January 2016.

Author's Address

Ted Hardie

Email: ted.ietf@gmail.com