

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.

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/>](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>.

- [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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: February 7, 2020

E. Lewis
ICANN
August 6, 2019

RFC Origins of Domain Names
draft-lewis-domain-names-13

Abstract

Is the concept of Domain Names owned by the DNS protocol or does the DNS protocol exist to support the concept of Domain Names? This question has become pertinent in light of proposals to use Domain Names in protocols in ways incompatible with the DNS protocol and the operational environment built to run the protocol.

This document is intended to help answer this question by presenting a look into the recorded history of relevant Requests for Comments. This document comprises the research and views of the author and has benefited from review and input from many IETF experts, but it does not represent the consensus opinion of the IETF.

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 February 7, 2020.

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
1.1. Goal	4
1.2. Terminology	5
2. Early RFCs	5
3. Emergence of Domain Names	6
3.1. The Term "Domain Name" Itself	7
3.2. The Term "Resolve"	8
3.3. Where Does It Start?	9
4. Dialects, So To Speak, of Domain Names	10
4.1. Domain Names in the DNS	10
4.2. Host Names	11
4.3. URI Authority and Domain Names	12
4.4. Internet Protocol Address Literals	12
4.5. Internationalized Domain Names in Applications	12
4.6. Restricted for DNS Registration	13
4.7. Tor Network Names	13
4.8. X.509	13
4.9. Multicast DNS	14
4.10. /etc/hosts	14
4.11. Other Protocols	14
4.12. Other Others	15
5. Interoperability Considerations	15
6. IANA Considerations	16
7. Security Considerations	16
8. Acknowledgements	16
9. Informational References	17
Author's Address	22

1. Introduction

Which came first, the concept of Domain Names or the DNS? This question is at the heart of whether or how Domain Names are put to use in ways avoiding the DNS protocol.

The discussion leading to "The '.onion' Special-Use Domain Name" [RFC7686], a document designating "onion" as a top-level domain in the Special Use Domain Names registry (see "Special Use Domain Names" [RFC6761]), opened the question of how to treat Domain Names that were designed to be used external to the DNS. The history of Domain

Names and the DNS had become intertwined over time to the point that what is essentially a case of permission-less innovation led to contentious discussions on the IETF's DNS Operations (DNSOP) working group mail list and an interim meeting of the DNSOP working group [DNSOP].

A portion of the discussion centered around a seeming conflict among processes to register Domain Names, such as the process launched from "Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority" [RFC2860], for registering a name in the global, public DNS and the process for registering a name in the Special Use Domain Names registry. The latter process is documented in "Special Use Domain Names" cited in the previous paragraph.

To help establish a way forward, a look backward is thought to be a good start. A document search, sticking to RFC documents, reveals evidence of discussions on Domain Names prior to the DNS, with the DNS protocol's base documents indicating that the DNS is based on some simplifying assumptions, implying there is a larger concept in play.

To help bolster the idea that Domain Names came first, a look at how other protocols have treated identifying names, how Domain Names are put to use, how what a name is further restricted for the protocol's needs. From this it has become apparent that the concept of Domain Names has drifted over time, which leads to some uncertainty when it comes to looking forward.

During reviews of this document, documented studies of other difficulties resulting from the uncertainty have surfaced. "IAB Thoughts on Encodings for Internationalized Domain Names" [RFC6055] documents issues related to converting human-readable forms of Domain Names in forms useful to automated applications when there is no clear architecture or precise definition of how to handle Domain Names. "Issues in Identifier Comparison for Security Purposes" [RFC6943] documents issues related to the same conversion as related to evaluating security policies. The presence of these studies suggests a need to examine the architecture of naming and identifiers.

The most glaring omission discovered by the document survey is a definitive foundation for Domain Names. There are abstract descriptions of the concept that come close to qualifying as a definition. The descriptions though are too loose to be something that can be tested objectively, frustrating discussions when it comes to innovations in the use of Domain Names.

In reviews of this document, an important thought has been expressed. During the era when the early RFC documents were prepared, many considerations now deemed important were not considered, discussed, examined. This lack of attention should be taken simply as the limits of the problem space perceived at the time and not as an intentional non-statement about questions now under consideration. The fact that the history is presented here does not imply that history will contain the answers and guide the way forward, the history as presented is only a starting point.

This document is a literature search covering the RFC series and makes a case for clarifications to be made. There are obvious continuations to this work, such as the earlier Internet Engineering Notes series (IEN), other published works, and interviews with participants from the early days. This document is intended to help answer this question of whether the concept of Domain Names is owned by the DNS protocol or does the DNS protocol exist to support the concept of Domain Names. It does this by presenting a look into the recorded history of relevant Requests for Comments. This document comprises the research and views of the author and has benefited from review and input from many IETF experts, but it does not represent the consensus opinion of the IETF.

1.1. Goal

To establish a solid foundation accommodating an installed base and permission-less innovation, having a clear definition of Domain Names would be great. This document, however, does not attempt to achieve a definition. This document's goal has settled into compiling a narrative on the history, within perhaps artificial bounds (the RFC series), and declaring that there is a need to clarify Domain Names.

In this document are criteria for performing a clarification, recognizing from experience in preparing "The Role of Wildcards in the Domain Name System" [RFC4592] and "DNS Zone Transfer Protocol (AXFR)" [RFC5936] that clarifications may have adverse impacts on deployed software, thus entering into a clarifications activity is not to be taken without considerations.

There are a few deviations from the strict rule of relying on the RFC series. First is the research into the term "resolve" and then further additions during late reviews of this document. The experience of these deviations illustrates the need to expand the literature search beyond the RFC series and to include other publications and recollections.

1.2. Terminology

Throughout this document (except in document quotations) the term "Domain Names" is capitalized to emphasize the concept of the names and "the DNS" is used to describe the protocol and algorithms described in STD 13, including any applicable updates, related standards track documents and experimental track documents. The words "DNS domain names" refers to the definition of Domain Names within the DNS (as well as, for example, "Tor domain names" referring to the definition of Domain Names within the Tor system).

The term "domain" is a generic term, having many dictionary definitions. There are many naming systems in existence, many unrelated to the Internet. The use of the term Domain Names in this document refers to a roughly-defined set of protocols defined in IETF RFC documents and their applications' use of a somewhat common, interoperating, naming structure. Lacking a formal, documented definition for Domain Names, which is why this document exists, it is hard to avoid a hand-waving reference.

2. Early RFCs

Two or three decades into the history of Domain Names, a popular notion has taken hold that Domains Names were defined and specified in the definition of the Domain Name System (DNS). There are two documents that form the basic definition of the DNS, "Domain names - concepts and facilities" [RFC1034] and "Domain names - implementation and specification" [RFC1035] referred to as RFC 1034 and RFC 1035, respectively. (Note that there is another pair of Request for Comments documents with the same titles [RFC0882] [RFC0883] that precede RFC 1034 and RFC 1035, declared obsolete in favor of the newer documents.) Together RFC 1034 and RFC 1035 form STD 13, a full standard cataloged by the RFC Editor. The definitions of the DNS' version of Domain Names within RFC 1034 and RFC 1035 have become the apparently-authoritative source for discussions on what is a Domain Name.

The truth is, the documents comprising STD 13 do not define Domain Names, the documents define only how Domain Names are used and processed in the DNS. However, the way in which those RFC documents read seem to lend to the confusion.

RFC 1034, section 2 begins with this text:

"This RFC introduces domain style names, their use for Internet mail and host address support, and the protocols and servers used to implement domain name facilities."

That text seems to indicate that RFC 1034 is the origin of Domain Names. Immediately following is section 2.1, entitled "The history of domain names" which includes the following text. (The text differs from the original presentation only in wrapping of text to fit current formatting rules.)

Continuing the quote from RFC 1034:

"The result was several ideas about name spaces and their management [IEN-116, RFC-799, RFC-819, RFC-830]. The proposals varied, but a common thread was the idea of a hierarchical name space, with the hierarchy roughly corresponding to organizational structure, and names using "." as the character to mark the boundary between hierarchy levels. A design using a distributed database and generalized resources was described in [RFC-882, RFC-883]. Based on experience with several implementations, the system evolved into the scheme described in this memo."

The only reference included in that text not otherwise mentioned in this document is to "INTERNET NAME SERVER", identified as IEN-116.[IEN116]

The DNS as it is known today did not invent Domain Names. Work on the Simple Mail Transfer Protocol (SMTP) preceding work on the DNS mentions Domain Names, and even SMTP too was not the origin of the concept. The DNS is not even the first attempt at an Internet naming system, see "The Domain Naming Convention for Internet User Applications" [RFC0819] and "A Distributed System for Internet Name Service" [RFC0830].

One important phrase to keep in mind is:

"To simplify implementations,"

which appears in both of the RFC 1034 and RFC 1035 documents, as well as their predecessor pair RFC 882 and RFC 883. This gives credence to the notion that Domain Names exist beyond the DNS, in that the text following the phrase is meant to limit an existing definition or concept as opposed to introducing a new idea.

3. Emergence of Domain Names

The first effort taken, in preparation for writing this document, was to scan for the earliest use of the term "domain name" or "name domain". This work is detailed in the following section, but, as noted in private email by reviews of early versions of the document, gave the impression that Domain Names were somehow a by-product of the effort to develop electronic mail. To challenge the notion that

email begat Domain Names, a search through RFC documents for the use of the term resolve as it applies to Domain Names was also done.

3.1. The Term "Domain Name" Itself

Domain Names emerged from the need to build a hierarchy around the growing number of identified hosts exchanging email. "SIMPLE MAIL TRANSFER PROTOCOL" [RFC0788], explains, in its section 3.7:

"At some not too distant future time it might be necessary to expand the mailbox format to include a region or name domain identifier. There is quite a bit of discussion on this at present, and is likely that SMTP will be revised in the future to take into account naming domains."

Knowing the origins of a concept helps setting the correct boundaries for discussion. The past isn't meant to restrict the future but meant to help provide a context, include forgotten ideas, and help identify rational for scope creep.

"Internet Name Domains" [RFC0799] has (arguably) the first formation of what is a Domain Name:

"In its most general form, a standard internet mailbox name has the syntax

`<user>.<host>@<domain>` ,

where `<user>` is the name of a user known at the host `<host>` in the name domain `<domain>`."

Prior to this, the term "domain" referred to principally an administrative domain, such as the initial organizations involved in networks at the time.

"NCP/TCP TRANSITION PLAN" [RFC0801] contains this, indicating the passage from the host tables:

"It might be advantageous to do away with the host name table and use a Name Server instead, or to keep a relatively small table as a cache of recently used host names."

"Computer Mail Meeting Notes" [RFC0805] contains this:

"The conclusion in this area was that the current "user@host" mailbox identifier should be extended to 'user@host.domain' where 'domain' could be a hierarchy of domains."

"The Domain Naming Convention for Internet User Applications" [RFC0819] contains this:

"A decision has recently been reached to replace the simple name field, "<host>", by a composite name field, '<domain>' "

A Domain Name began to take on its current form:

"Internet Convention: Fred@F.ISI.ARPA"

In addition, "simple name" is defined as what we now call a label, and a "complete (fully qualified) name" is defined as "concatenation of the simple names of the domain structure tree nodes starting with its own name and ending with the top level node name". Noticeably absent is a terminating dot or any mention or representation of a root.

"The Domain Naming Convention for Internet User Applications" (RFC 819) also defines ARPA as a top-level name (as opposed to top-level domain name). This is an early mention of the role of top-level names. Additionally, the use of "." [RFC0020][ANSIX34] as a separating character is mentioned.

This walk thru history relies solely on the record left behind inside RFC documents. The precise chain of events is likely slightly different and nuanced. The point of the exercise is to show that Domain Names are a concept that emerged over time, spawned the DNS with its Domain Names, a definition of host names derived from the host tables, and was heavily influenced by SMTP as the driving application. The definition of the FTP protocol, originally defined in "FILE TRANSFER PROTOCOL" [RFC0959], never mentions hosts, domains or host names. No formal definition of Domain Names has been written and recorded.

Note: Concurrent with the writing of this document, the Domain Name Systems Operations working group is documenting a definition for "Domain Names". The first edition of "DNS Terminology" [RFC7719] has a recitation of the original definition from STD 13, the successor edition (still in preparation) has a new, further reaching definition.

3.2. The Term "Resolve"

In looking for other early mentions of Domain Names, a look for the use of the term "resolve" or "resolution" was conducted, reading through early (arbitrarily defined as pre-1000) RFC documents. The term "resolve" appears numerous times, but in many different

contexts. "Resolve" has many meanings, consulting a dictionary, such as Merriam Webster's dictionary [MWDICT], none which seem to match the use associated with domain names. For example, a committee can resolve to solve a certain question. This use of "resolve" occurred numerous times in early RFC documents unrelated to Domain Names.

In "Proposed Official Standard for the Format of ARPA Network Messages" [RFC0724] the term resolve was used in the sense of mapping an identifier into an address or something actionable. A section on Semantics (C), Address fields (1.), General (a.), bullet 1 states:

"<path>s are used to refer to a location, on the ARPANET, containing a stored address list. The <phrase> should contain text which the referenced host can resolve to a file. This standard is not a protocol and so does not prescribe HOW data is to be retrieved from the file."

Private email to the (reachable) authors of the document pointed to the use of "resolve" stemming from work on programming languages and compiler theory. In that field of work, variables are associated with machine addresses when linking code. There are formal papers including "A Theory of Name Resolution" [TONR15] using the term and the term resolution is used in the field of "Automated Reasoning" [WIKIAR].

The exercise of determining how the term "resolve" came to be part of Domain Names and DNS shows that there are influences, topics, terms and concepts from technologies preceding Domain Name and DNS that can be researched to help establish a foundation from which to build. There is more work to do here.

3.3. Where Does It Start?

Without a definitive introduction to Domain Names it is hard to know how far back in documented history to search for references to the concept. Chasing "domain" and variations has not necessarily found the beginning, chasing "resolution" and variations also has not necessarily found the beginning. During later reviews of this document, a significant early document has been identified for inclusion, an IEN entitled "A Note On Inter-Network Naming, Addressing, and Routing" by John F. Shoch. That document is tagged as IEN-19 [IEN019].

The note introduces the difference between names, addresses, and routes. The term "domain" is used to scope a name space, giving examples from telephony and networking. But there still is no formal definition of Domain Names nor any solution path towards Domain Names as they are commonly known today.

A relatively more modern document (15 years later), entitled "On the Naming and Binding of Network Destinations" [RFC1498], refers to IEN-19, extending the discussion on naming to divide into four categories of objects. This document illustrates the continuing conceptual work covering naming as opposed to further developing the solution known as Domain Names.

4. Dialects, So To Speak, of Domain Names

Subtypes of Domain Names have come to be defined for different protocols, evolving and sometimes building on previous definitions.

4.1. Domain Names in the DNS

The DNS protocol defines a subset of Domain Names that referred to as DNS domain names. The DNS places size restrictions on Domain Names and defines rules for matching DNS domain names, treating sets of DNS domain names as equivalent to each other. (This matching refers to treating upper case and lowercase ASCII letters as equivalent.) The DNS defines the format used to transmit the names across the network as well as rules for displaying them inside text zone files. The DNS creates the notion that names are assigned by an authority per zone.

Placing size restrictions on a DNS domain name is significant in reducing the overall population of names that can be represented in the DNS. The matching rules have the effect of creating (to use a term from graph theory) cliques, distorting the tree-nature of the Domain Name graph. A clique is a completely connected sub-graph implying cyclic paths, a tree is a graph that is acyclic. In sum, the treatment of ASCII (and only ASCII) cases as equivalent is a distortion of the DNS domain name hierarchy.

The DNS defines two representational formats for DNS domain names. One format is the "on-the-wire" format used inside messages, a flags-and-length octet followed by some count of octets for each label with the final length of 0 representing the root. The other format is a version that can be rendered in printable ASCII characters, complete with a means to represent other characters via an escape sequence. This does not alter the Domain Name concept but has implications when it comes to interoperating with other protocol definitions of their domain name use.

The DNS assumes that there is, in concept, a central authority creating names within the DNS management structure (called a zone). Although the DNS does not define how a central authority is implemented nor how it manages names, the names have to come from a single point to appear in a zone. There are other means for claiming names, an example will be mentioned later.

DNS domain names allows for names to appear as address literals, such as "192.0.2.1" or "0:0:0:0:0:FFFF::192.0.2.1". But such Domain Names are not used in the DNS for two reasons. Applications expecting a Domain Name (as a comment line parameter as an example) could opt to treat the string as an address literal and therefore not look for the string in the DNS domain name space. And, if addresses were stored using this representation, there would be no means to aggregate managed address ranges into zones.

By reversing the order of the address components, DNS domain names can be aggregated (as in routing) into the same zone. E.g., the network address 192.0.2.1 would be represented by a DNS domain name as "1.2.0.192.in-addr.arpa." as described in RFC 1035. For IPv6, the convention used is documented in "DNS Extensions to Support IP Version 6" [RFC3596], section 2.5.

See also "Issues in Identifier Comparison for Security Purposes" [RFC6943] section 3.1, "Host Names", in particular, section 3.1.1 and 3.1.2 on address literals, and section 4.1, "Conflation."

DNS domain names have become the dominant definition of Domain Names due to the success (scale) of the DNS on the public Internet. Many protocols interact with the DNS but instead of supporting the complete definition of DNS domain names the protocols rely on a subset more commonly called host names.

4.2. Host Names

Work on the definition of a host name began well before the issuance of the STD 13 documents defining the DNS. The rules for the Preferred Syntax in RFC 1034 conform to the host name rules outlined in "DoD Internet host table specification" [RFC0952]. The host name definition was presented again in "Requirements for Internet Hosts -- Application and Support" [RFC1123] (which is part of STD 3). In section 2.1 of RFC 1123, one (of two mentions) definition of host name is presented, noting that the definition is a relaxation of what is in RFC 952.

Host names are subsets of DNS domain names in the sense that the character set is limited. In particular, only "let" (i.e., presumably letters a-z), "digits" and "hyphen" can be used, with hyphen only internal to a label. (This description is meant to be illustrative, not normative. See the grammar presented on page 5 of "DoD Internet host table specification" for specifics.) "Hypertext Transfer Protocol -- HTTP/1.0" [RFC1945], Section 3.2.2 "http URL" specifically references section 2.1 of RFC 1123. The reference is explicit.

"Simple Mail Transfer Protocol" [RFC5321] refers to RFC 1035 for a definition of Domain Names but includes text close to what is in the previous paragraph, noting that Domain Names as used in SMTP refer to both hosts and to other entities. RFC 5321 updates RFC 1123, but does not cite the latter for a definition of host names. RFC 5321 additionally requires brackets to surround address literals, referring to the use case as an "alternative to a domain name."

See also "IAB Thoughts on Encodings for Internationalized Domain Names" [RFC6055], particularly section 3 entitled "Use of Non-ASCII in DNS" for more thoughts on host names.

4.3. URI Authority and Domain Names

In "Uniform Resource Identifier (URI): Generic Syntax" [RFC3986], also known as STD 66, mentions in its section 3.2.2 (page 20) that the host subcomponent of the URI Authority (section 3.2) "should conform to the DNS syntax". This comes after discussion that the host subcomponent is not strongly tied to the DNS, i.e., names can be managed via a concept other than the DNS. There's no discussion on the rationale but this enables the reuse of code parsing and marshaling the host subcomponent between different Domain Name environments.

This reinforces the notion that there's a need to understand how Domain Names interoperate amongst protocols and applications. And reinforces the need to derive or make explicit a way for client software to know how to resolve a name, that is, convert a name into a network address.

4.4. Internet Protocol Address Literals

The above definition includes address literals such as 192.0.2.1 for IPv4 and even IPv6 literals such as ::ffff:192.0.2.1. Yes, these qualify as Domain Names. The addresses might be encased in square brackets "[" and "]" (SMTP mentioned already). In the DNS, as previously described in section 3.1, they are represented per appropriate conventions.

4.5. Internationalized Domain Names in Applications

The original uses of Domain Names (such as DNS domain names and host names) assumed the ASCII character set. Specifically, making the labels case insensitive prohibited a straightforward use of any method of representation of non-ASCII characters.

"Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework" [RFC5890], with associated other documents,

defines IDNA2008 as a convention for handling non-ASCII characters in DNS domain names. In figure 1 of that document, the sets of legal DNS domain name formats are defined. Noted in the footnotes of the figure, applications unaware of IDNA2008 cannot distinguish the subsets defined by the document meaning this definition is not an alteration of Domain Names, but, like host names, yet another subset of DNS domain names.

4.6. Restricted for DNS Registration

"Suggested Practices for Registration of Internationalized Domain Names (IDN)" [RFC4290] presents reasons why DNS domain name registration is restricted in the context of IDN. (That RFC refers to a obsoleted version of IDNA but the concepts still apply.) This is yet another convention related to DNS domain names, which excludes names that fit the syntax but would lead to undesirable outcomes in applications.

4.7. Tor Network Names

The Tor network is an activity organized by the Tor Project, Inc., described on its main web page "<https://www.torproject.org/index.html.en>".

One component of the Tor network name space are Domain Names ending in ".onion". (There are other suffixes in use, but it isn't very clear how they are used, defined or whether they are active.)

The way in which Domain Names are used in Tor is described in two web documents "Tor Rendezvous Specification" [RENDEV] and "Special Hostnames in Tor" [OHOST] available from the project's website.

Syntactically, a Tor domain name fits within the DNS domain name definition but the manner of assignment is different in a manner incompatible with the DNS. (Not better or worse, still significantly different.) Tor domain names are derived from cryptographic keys and organized by distributed hash tables, instead of assigned by a central authority per zone.

4.8. X.509

"Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile" [RFC5280], section 4.2.1.6 "Subject Alternative Name" a dNSName is defined to be a host name, with the further restriction that the name " " cannot be used.

4.9. Multicast DNS

Multicast DNS uses a name space ending with ".local." as described in "Multicast DNS" [RFC6762]. The rules for Multicast DNS domain names differ from DNS domain names. Multicast DNS domain names are encoded as Net-Unicode as defined in "Unicode Format for Network Interchange" [RFC5198] with the DNS domain name tradition of case folding the ASCII letters when matching names. Appendix F of RFC 6762 gives an explanation of why the punycode algorithm, defined in "Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)" [RFC3492], is not used.

4.10. /etc/hosts

The precursor to the DNS, host tables, still exists in remnants in many operating systems. There are library functions, used by applications to resolve Domain Names, that can return names of arbitrary length (meaning, for example, longer than what DNS domain names are defined to be).

"Basic Socket Interface Extensions for IPv6" [RFC3493], addresses this in Section 6, further documentation can be found as part of "The Open Group Base Specifications Issue 7" [IEEE1003] and "Microsoft Winsock Functions" [WINSOCK].

4.11. Other Protocols

This section is used to list (some) other protocols that use Domain Names but in general do not impose any other restrictions that what has been mentioned above.

SSH, documented in "The Secure Shell (SSH) Protocol Architecture" [RFC4251] uses host names, using the name when storing public keys of hosts. SSH clients, not necessarily the protocol, illustrate how applications juggle the different forms of Domain Names. SSH can be invoked to open a secure shell with a host via its DNS domain name/ host name or it can be used to open a secure shell with a host via its Multicast DNS domain name. Or, many others, including name of a purely local, per-user scope. (Note that SSH does not distinguish between DNS domain names and Multicast DNS domain names in the protocol definition, the difference is handled in resolution libraries belonging to the computing platform.)

FTP, defined in "FILE TRANSFER PROTOCOL (FTP)" [RFC0959], is silent on Domain Names but client implementations of the protocol behave as SSH clients, being unaware the differences between definitions of Domain Names.

DHCP, defined in "Dynamic Host Configuration Protocol" [RFC2131], includes Domain Names in many DHCP options. The use is described in many documents, starting with "DHCP Options and BOOTP Vendor Extensions" [RFC2132]. In "Dynamic Host Configuration Protocol (DHCP) Domain Search Option" [RFC3397] the encoding of Domain Names uses the on-the-wire format of DNS domain names. In "The DHCP Client FQDN Option" [RFC4702] the same format is used. "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)" [RFC8415] contains definitions related to DHCP designed for IPv6. The significance of the DHCP protocols implementation of Domain Names is that, while most other protocols represent DNS domain names or host names in a human readable form, DHCP is using the machine-friendly format.

4.12. Other Others

If there is a use of Domain Names not listed here it is merely an omission. The goal in this document is to provide a survey that is sufficient to avoid hand-waving arguments, recognizing the diminishing return building a complete roster of uses of Domain Names.

5. Interoperability Considerations

Any single protocol can define a format for a conceptual Domain Name. Examples given above show that many protocols have done so. From the examples, it is clear that the way in which protocols have interpreted Domain Names has varied, leading to, at least, user interfaces having to have built-in intelligence when handling names and, at worst, a growing confusion over how the Domain Name space is to be managed.

When protocols having different formats and rules for Domain Names interact, software implementing the protocols translate one protocol's domain name format to another's format. Even when the translation is straightforward, it is predictable that software will fail to handle this situation well.

Often the clash of definitions impacts the design of a new protocol and/or an extension of a protocol. For example, adding non-ASCII domain names has to be done with backwards compatibility with an installed base of ASCII-assuming code. This clash can inhibit new uses of Domain Names.

Search lists are a Domain Name mechanism studied in "SSAC Advisory on DNS 'Search List' Processing" [SSAC064]. (Note that the advisory's title labels search lists as a DNS mechanism although the idea of a search list spans many different naming schemes.) One of the particular use cases related to this topic is the issuance of search

lists via DHCP and then used by any user-client protocol implementation. This emphasizes an interoperability consideration for how Domain Names are treated in different protocols, not just among implementations of one protocol.

The detection and handling of Fully Qualified Domain Names is an interoperability issue as well. At issue is the significance of the terminating separation character in a printed version of a Domain Name. Many clients, when passed a Domain Name as an identifier will add a dot at the end of the argument if the argument does not already end in a dot. [TRAILDOT1] Some do this only after applying the aforementioned search list. As mentioned in the SSAC document in the previous paragraph, inconsistency leads to unpredictable results.

The Special Use Domain Names registry lists Domain Names that are to be treated in a manner inconsistent with the DNS normal processing rules. This registry contains Domain Names regardless of whether the name is a DNS domain name and regardless whether the name is a top-level (domain) name or is positioned elsewhere in the tree structure.

These are reasons this document is needed. The reason for the confusion over what's a legal domain name stems from application-defined restrictions. For example, using a one-label domain name ("dotless") for sending email is not a problem with the DNS nor the name in concept, but is a problem for mail implementations that expect more than one label. (One-label names may be assumed to be in ARPA host table format.) The "IAB Statement: Dotless Domains Considered Harmful" [IABSTMT] elaborates.

6. IANA Considerations

None.

7. Security Considerations

Nothing direct. This document proposes a definition of the term "Domain Name" and surveys how it has been variously applied. In some sense, loosely defined terms give rise to security hazards. Beyond that, there is no impact of "security."

8. Acknowledgements

Comments or contributions from Andrew Sullivan, Paul Hoffman, George Michaelson, Kevin Darcy, Joe Abley, Jim Reid, Tony Finch, Robert Edmonds, hellekin, Stephane Bortzmeyer, Ray Bellis, Bob Harold, Alec Muffett, Stuart Cheshire, Dave Thaler, Niall O'Reilly, John Klensin, Dave Crocker, Ken Pogran, John Vittal, Lixia Zhang, Ralph Droms and a

growing list of others I am losing track of. Not to imply endorsement.

9. Informational References

- [ANSIX34] American National Standards Institute (formerly United States of America Standards Institute), "USA Code for Information Interchange, ANSI X3.4-1968", 1968.
- [DNSOP] Woolf, S., "Interim DNSOP WG meeting on Special Use Names: some reading material", 2015, <https://mailarchive.ietf.org/arch/msg/dnsop/VnSjuXYiwd89K_mt05-CJLa0lbs>.
- [IABSTMT] Board, I. A., "IAB Statement: Dotless Domains Considered Harmful", 2013, <<https://www.iab.org/documents/correspondence-reports-documents/2013-2/iab-statement-dotless-domains-considered-harmful/>>.
- [IEEE1003] Group, T. I. A. T. O., "The Open Group Base Specifications Issue 7, IEEE Std 1003.1, 2013 Edition, Copyright 2001-2013 The IEEE and The Open Group", 2013, <<http://pubs.opengroup.org/onlinepubs/9699919799/functions/freeaddrinfo.html>>.
- [IEN019] Shoch, J., "A note on Inter-Network Naming, Addressing, and Routing", IEN 19, January 1973, <<https://www.rfc-editor.org/ien/ien19.txt>>.
- [IEN116] Postel, J., "INTERNET NAME SERVER", IEN 116, August 1979, <<https://www.rfc-editor.org/ien/ien116.txt>>.
- [MWDICT] Merriam-Webster, Incorporated, "Merriam-Webster's Online Dictionary, 11th Edition (Merriam-Webster's Collegiate Dictionary)", 2003, <<https://www.merriam-webster.com/>>.
- [OHOST] Mathewson, N., "Special Hostnames in Tor", undated, <<https://gitweb.torproject.org/torspec.git/tree/address-spec.txt>>.
- [RENDEV] Anonymous, "Tor Rendezvous Specification", undated, <<https://gitweb.torproject.org/torspec.git/tree/>>.
- [RFC0020] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <<http://www.rfc-editor.org/info/rfc20>>.

- [RFC0724] Crocker, D., Pogran, K., Vittal, J., and D. Henderson, "Proposed official standard for the format of ARPA Network messages", RFC 724, DOI 10.17487/RFC0724, May 1977, <<http://www.rfc-editor.org/info/rfc724>>.
- [RFC0788] Postel, J., "Simple Mail Transfer Protocol", RFC 788, DOI 10.17487/RFC0788, November 1981, <<http://www.rfc-editor.org/info/rfc788>>.
- [RFC0799] Mills, D., "Internet name domains", RFC 799, DOI 10.17487/RFC0799, September 1981, <<http://www.rfc-editor.org/info/rfc799>>.
- [RFC0801] Postel, J., "NCP/TCP transition plan", RFC 801, DOI 10.17487/RFC0801, November 1981, <<http://www.rfc-editor.org/info/rfc801>>.
- [RFC0805] Postel, J., "Computer mail meeting notes", RFC 805, DOI 10.17487/RFC0805, February 1982, <<http://www.rfc-editor.org/info/rfc805>>.
- [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>>.
- [RFC0830] Su, Z., "Distributed system for Internet name service", RFC 830, DOI 10.17487/RFC0830, October 1982, <<http://www.rfc-editor.org/info/rfc830>>.
- [RFC0882] Mockapetris, P., "Domain names: Concepts and facilities", RFC 882, DOI 10.17487/RFC0882, November 1983, <<http://www.rfc-editor.org/info/rfc882>>.
- [RFC0883] Mockapetris, P., "Domain names: Implementation specification", RFC 883, DOI 10.17487/RFC0883, November 1983, <<http://www.rfc-editor.org/info/rfc883>>.
- [RFC0952] Harrenstien, K., Stahl, M., and E. Feinler, "DoD Internet host table specification", RFC 952, DOI 10.17487/RFC0952, October 1985, <<http://www.rfc-editor.org/info/rfc952>>.
- [RFC0959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, DOI 10.17487/RFC0959, October 1985, <<http://www.rfc-editor.org/info/rfc959>>.

- [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>>.
- [RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, DOI 10.17487/RFC1123, October 1989, <<http://www.rfc-editor.org/info/rfc1123>>.
- [RFC1498] Saltzer, J., "On the Naming and Binding of Network Destinations", RFC 1498, DOI 10.17487/RFC1498, August 1993, <<https://www.rfc-editor.org/info/rfc1498>>.
- [RFC1945] Berners-Lee, T., Fielding, R., and H. Frystyk, "Hypertext Transfer Protocol -- HTTP/1.0", RFC 1945, DOI 10.17487/RFC1945, May 1996, <<http://www.rfc-editor.org/info/rfc1945>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<http://www.rfc-editor.org/info/rfc2131>>.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <<https://www.rfc-editor.org/info/rfc2132>>.
- [RFC2860] Carpenter, B., Baker, F., and M. Roberts, "Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority", RFC 2860, DOI 10.17487/RFC2860, June 2000, <<http://www.rfc-editor.org/info/rfc2860>>.
- [RFC3397] Aboba, B. and S. Cheshire, "Dynamic Host Configuration Protocol (DHCP) Domain Search Option", RFC 3397, DOI 10.17487/RFC3397, November 2002, <<http://www.rfc-editor.org/info/rfc3397>>.
- [RFC3492] Costello, A., "Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)", RFC 3492, DOI 10.17487/RFC3492, March 2003, <<http://www.rfc-editor.org/info/rfc3492>>.

- [RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 3493, DOI 10.17487/RFC3493, February 2003, <<http://www.rfc-editor.org/info/rfc3493>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, DOI 10.17487/RFC3596, October 2003, <<http://www.rfc-editor.org/info/rfc3596>>.
- [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>>.
- [RFC4251] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Protocol Architecture", RFC 4251, DOI 10.17487/RFC4251, January 2006, <<http://www.rfc-editor.org/info/rfc4251>>.
- [RFC4290] Klensin, J., "Suggested Practices for Registration of Internationalized Domain Names (IDN)", RFC 4290, DOI 10.17487/RFC4290, December 2005, <<http://www.rfc-editor.org/info/rfc4290>>.
- [RFC4592] Lewis, E., "The Role of Wildcards in the Domain Name System", RFC 4592, DOI 10.17487/RFC4592, July 2006, <<https://www.rfc-editor.org/info/rfc4592>>.
- [RFC4702] Stapp, M., Volz, B., and Y. Rekhter, "The Dynamic Host Configuration Protocol (DHCP) Client Fully Qualified Domain Name (FQDN) Option", RFC 4702, DOI 10.17487/RFC4702, October 2006, <<http://www.rfc-editor.org/info/rfc4702>>.
- [RFC5198] Klensin, J. and M. Padlipsky, "Unicode Format for Network Interchange", RFC 5198, DOI 10.17487/RFC5198, March 2008, <<http://www.rfc-editor.org/info/rfc5198>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<http://www.rfc-editor.org/info/rfc5280>>.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <<http://www.rfc-editor.org/info/rfc5321>>.

- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <<http://www.rfc-editor.org/info/rfc5890>>.
- [RFC5936] Lewis, E. and A. Hoenes, Ed., "DNS Zone Transfer Protocol (AXFR)", RFC 5936, DOI 10.17487/RFC5936, June 2010, <<https://www.rfc-editor.org/info/rfc5936>>.
- [RFC6055] Thaler, D., Klensin, J., and S. Cheshire, "IAB Thoughts on Encodings for Internationalized Domain Names", RFC 6055, DOI 10.17487/RFC6055, February 2011, <<http://www.rfc-editor.org/info/rfc6055>>.
- [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>>.
- [RFC6943] Thaler, D., Ed., "Issues in Identifier Comparison for Security Purposes", RFC 6943, DOI 10.17487/RFC6943, May 2013, <<http://www.rfc-editor.org/info/rfc6943>>.
- [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>>.
- [RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <<https://www.rfc-editor.org/info/rfc7719>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.
- [SSAC064] Anonymous, "SSAC Advisory on DNS "Search List" Processing", 2014, <<https://www.icann.org/en/system/files/files/sac-064-en.pdf>>.
- [TONR15] Wachsmuth, P. N. A. P. T. E. V. G., "A Theory of Name Resolution", last seen 2015, <[https://msdn.microsoft.com/en-us/library/windows/desktop/ms738520\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/ms738520(v=vs.85).aspx)>.

- [TRAILDOT1] by), S. C. (. M., "Trailing Dots in Domain Names",
Undated,
<<http://www.dns-sd.org/trailingdotsindomainnames.html>>.
- [WIKIAR] Anonymous, "Automated Reasoning", last edit 2016,
<https://en.wikipedia.org/wiki/Automated_reasoning>.
- [WINSOCK] Microsoft, "getaddrinfo function", last seen 2017,
<[https://msdn.microsoft.com/en-us/library/windows/desktop/ms738520\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/ms738520(v=vs.85).aspx)>.

Author's Address

Edward Lewis
ICANN

Email: edward.lewis@icann.org

Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 24, 2018

B. Trammell
ETH Zurich
September 20, 2017

Properties of an Ideal Naming Service
draft-trammell-inip-pins-04

Abstract

This document specifies a set of necessary functions and desirable properties of an ideal system for resolving names to addresses and associated information for establishing communication associations in the Internet. For each property, it briefly explains the rationale behind it, and how the property is or could be met with the present Domain Name System. It is intended to start a discussion within the IAB's Names and Identifiers program about gaps between the present reality of DNS and the naming service the Internet needs by returning to first principles.

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 March 24, 2018.

Copyright Notice

Copyright (c) 2017 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	3
2. Terminology	3
3. Query Interface	4
3.1. Name to Address	4
3.2. Address to Name	4
3.3. Name to Name	4
3.4. Name to Auxiliary Information	5
3.5. Name/Address to Auxiliary Information	5
4. Authority Interface	5
5. Properties	5
5.1. Semantics	5
5.1.1. Meaningfulness	5
5.1.2. Distinguishability	6
5.1.3. Minimal Structure	6
5.2. Authority	6
5.2.1. Federation of Authority	6
5.2.2. Uniqueness of Authority	6
5.2.3. Transparency of Authority	7
5.2.4. Revocability of Authority	7
5.2.5. Consensus on Root of Authority	7
5.3. Authenticity	7
5.3.1. Authenticity of Delegation	7
5.3.2. Authenticity of Response	8
5.3.3. Authenticity of Negative Response	8
5.4. Consistency	8
5.4.1. Dynamic Consistency	8
5.4.2. Explicit Inconsistency	8
5.4.3. Global Invariance	9
5.5. Performance Properties	9
5.5.1. Availability	9
5.5.2. Lookup Latency	10
5.5.3. Bandwidth Efficiency	10
5.5.4. Query Linkability	10
5.5.5. Explicit Tradeoff	10
5.6. Trust in Infrastructure	11
6. Observations	11
6.1. Delegation and redirection are separate operations	11
6.2. Queries and assertion contexts are presently implicit	11
6.3. Unicode alone may not be sufficient for distinguishable names	12
6.4. Implicit inconsistency makes global invariance	

challenging to verify	12
7. IANA Considerations	12
8. Security Considerations	12
9. Acknowledgments	13
10. Informative References	13
Author's Address	14

1. Introduction

The Internet's Domain Name System (DNS) [RFC1035] is an excellent illustration of the advantages of the decentralized architecture that have made the Internet able to scale to its present size. However, the choices made in the evolution of the DNS since its initial design are only one path through the design space of Internet-scale naming services. Many other naming services have been proposed, though none has been remotely as successful for general-purpose use in the Internet.

This document returns to first principles, to determine the dimensions of the design space of desirable properties of an Internet-scale naming service. It is a work in progress, intended to start a discussion within the IAB's Names and Identifiers program about gaps between the present reality of DNS and the naming service the Internet needs.

Section 3 and Section 4 define the set of operations a naming service should provide for queriers and authorities, Section 5 defines a set of desirable properties of the provision of this service, and Section 6 examines implications of these properties.

2. Terminology

The following capitalized terms are defined and used in this document:

- o Subject: A name, address, or name-address pair about which the naming service can answer queries
- o Association: A mapping between a Subject and information about that Subject
- o Authority: An entity that has the right to determine which Associations exist within its namespace
- o Delegation: An Association that indicates that an Authority has given the right to make assertions about the Associations within the part of a namespace identified by a Subject to a subordinate Authority.

3. Query Interface

At its core, a naming service must provide a few basic functions for queriers, associating a Subject of a query with information about that subject. The information available from a naming service is that which is necessary for a querier to establish a connection with some other entity in the Internet, given a name identifying it.

3.1. Name to Address

Given a Subject name, the naming service returns a set of addresses associated with that name, if such an association exists, where the association is determined by the authority for that name. Names may be associated with addresses in one or more address families (e.g. IP version 4, IP version 6). A querier may specify which address families it is interested in receiving addresses for, and the naming system treats all address families equally.

This mapping is implemented in the DNS protocol via the A and AAAA RRTYPES.

3.2. Address to Name

Given an Subject address, the naming service returns a set of names associated with that address, if such an association exists, where the association is determined by the authority for that address.

This mapping is implemented in the DNS protocol via the PTR RRTYPE. IPv4 mappings exist within the in-addr.arpa. zone, and IPv6 mappings in the ip6.arpa. zone. This mechanism has the disadvantage that delegations in IPv4 only happen on octet (8-bit) boundaries, and in IPv6 only happen on hex digit (4-bit) boundaries, which make delegations on other prefixes operationally difficult.

3.3. Name to Name

Given a Subject name, the naming service returns a set of object names associated with that name, if such an association exists, where the association is determined by the authority for the subject name.

This mapping is implemented in the DNS protocol via the CNAME RRTYPE. CNAME does not allow the association of multiple object names with a single subject, and CNAME may not combine with other RRTYPES (e.g. NS, MX) arbitrarily.

3.4. Name to Auxiliary Information

Given a Subject name, the naming service returns other auxiliary information associated with that name that is useful for establishing communication over the Internet with the entities associated with that name.

Most of the other RRTYPES in the DNS protocol implement these sort of mappings.

3.5. Name/Address to Auxiliary Information

As a name might be associated with more than one address, auxiliary information as above may be associated with a name/address pair, as opposed to just with a name.

This mapping is not presently supported by the DNS protocol.

4. Authority Interface

The query interface is not the only interface to the naming service: the interface a naming service presents to an Authority allows updates to the set of Associations and Delegations in that Authority's namespace. Updates consist of additions of, changes to, and deletions of Associations and Delegations. In the present DNS, this interface consists of the publication of a new zone file with an incremented version number, but other authority interfaces are possible.

5. Properties

The following properties are desirable in a naming service providing the functions in Section 3 and Section 4.

5.1. Semantics

Since the point of a naming service is to replace network-layer identifiers with more useful identifiers for humans (whether end users, software developers, or network administrators), the Subject names the naming service can provide must meet two semantic criteria:

5.1.1. Meaningfulness

A naming service must provide the ability to name objects that its human users find more meaningful than the objects themselves.

5.1.2. Distinguishability

A naming service must make it possible to guarantee that two different names are easily distinguishable from each other by its human users.

5.1.3. Minimal Structure

A naming service should impose as little structure on the names it supports as practical in order to be universally applicable. Naming services that impose a given organizational structure on the names expressible using the service will not translate well to societies where that organizational structure is not prevalent.

5.2. Authority

Every Association among names, addresses, and auxiliary data is subject to some Authority: an entity which has the right to determine which Associations and Subjects exist in its namespace. The following are properties of Authorities in our ideal naming service:

5.2.1. Federation of Authority

An Authority can delegate some part of its namespace to some other subordinate Authority. This property allows the naming service to scale to the size of the Internet, and leads to a tree-structured namespace, where each Delegation is itself identified with a Subject at a given level in the namespace.

In the DNS protocol, this federation of authority is implemented through delegation using the NS RRTYPE, redirecting queries to subordinate authorities recursively to the final authority. When DNSSEC is used, the DS RRTYPE is used to verify this delegation.

5.2.2. Uniqueness of Authority

For a given Subject, there is a single Authority that has the right to determine the Associations and/or Delegations for that subject. The unitary authority for the root of the namespace tree may be special, though; see Section 5.2.5.

In the DNS protocol as deployed, unitary authority is approximated by the entity identified by the SOA RRTYPE. The existence of registrars, which use the Extensible Provisioning Protocol (EPP) [RFC5730] to modify entries in the zones under the authority of a top-level domain registry, complicates this somewhat.

5.2.3. Transparency of Authority

A querier can determine the identity of the Authority for a given Association. An Authority cannot delegate its rights or responsibilities with respect to a subject without that Delegation being exposed to the querier.

In DNS, the authoritative name server(s) to which a query is delegated via the NS RRTYPE are known. However, we note that in the case of authorities which delegate the ability to write to the zone to other entities (i.e., the registry-registrar relationship), the current DNS provides no facility for a querier to understand on whose behalf an authoritative assertion is being made; this information is instead available via WHOIS. To our knowledge, no present DNS name servers use WHOIS information retrieved out of band to make policy decisions.

5.2.4. Revocability of Authority

An ideal naming service allows the revocation and replacement of an authority at any level in the namespace, and supports the revocation and replacement of authorities with minimal operational disruption.

The current DNS allows the replacement of any level of delegation except the root through changes to the appropriate NS and DS records. Authority revocation in this case is as consistent as any other change to the DNS.

5.2.5. Consensus on Root of Authority

Authority at the top level of the namespace tree is delegated according to a process such that there is universal agreement throughout the Internet as to the subordinates of those Delegations.

5.3. Authenticity

A querier must be able to verify that the answers that it gets from the naming service are authentic.

5.3.1. Authenticity of Delegation

Given a Delegation from a superordinate to a subordinate Authority, a querier can verify that the superordinate Authority authorized the Delegation.

Authenticity of delegation in DNS is provided by DNSSEC [RFC4033].

5.3.2. Authenticity of Response

The authenticity of every answer is verifiable by the querier. The querier can confirm that the Association returned in the answer is correct according to the Authority for the Subject of the query.

Authenticity of response in DNS is provided by DNSSEC.

5.3.3. Authenticity of Negative Response

Some queries will yield no answer, because no such Association exists. In this case, the querier can confirm that the Authority for the Subject of the query asserts this lack of Association.

Authenticity of negative response in DNS is provided by DNSSEC.

5.4. Consistency

Consistency in a naming service is important. The naming service should provide the most globally consistent view possible of the set of associations that exist at a given point in time, within the limits of latency and bandwidth tradeoffs.

5.4.1. Dynamic Consistency

When an Authority makes changes to an Association, every query for a given Subject returns either the new valid result or a previously valid result, with known and/or predictable bounds on "how previously". Given that additions of, changes to, and deletions of associations may have different operational causes, different bounds may apply to different operations.

The time-to-live (TTL) on a resource record in DNS provides a mechanism for expiring old resource records. We note that this mechanism makes additions to the system propagate faster than changes and deletions, which may not be a desirable property. However, as no context information is explicitly available in DNS, the DNS cannot be said to be dynamically consistent, as different implicitly inconsistent views of an association may be persistent.

5.4.2. Explicit Inconsistency

Some techniques require giving different answers to different queries, even in the absence of changes: the stable state of the namespace is not globally consistent. This inconsistency should be explicit: a querier can know that an answer might be dependent on its identity, network location, or other factors.

One example of such desirable inconsistency is the common practice of "split horizon" DNS, where an organization makes internal names available on its own network, but only the names of externally-visible subjects available to the Internet at large.

Another is the common practice of DNS-based content distribution, in which an authoritative name server gives different answers for the same query depending on the network location from which the query was received, or depending on the subnet in which the end client originating a query is located (via the EDNS Client Subnet extension {RFC7871}). Such inconsistency based on client identity or network address may increase query linkability (see Section 5.5.4).

These forms of inconsistency are implicit, not explicit, in the current DNS. We note that while DNS can be deployed to allow essentially unlimited kinds of inconsistency in its responses, there is no protocol support for a query to express the kind of consistency it desires, or for a response to explicitly note that it is inconsistent. [RFC7871] does allow a querier to note that it would specifically like the view of the state of the namespace offered to a certain part of the network, and as such can be seen as inchoate support for this property.

5.4.3. Global Invariance

An Association which is not intended to be explicitly inconsistent by the Authority issuing it must return the same result for every Query for it, regardless of the identity or location of the querier.

This property is not provided by DNS, as it depends on the robust support on the Explicit Inconsistency property above. Examples of global invariance failures include geofencing and DNS-based censorship ordered by a local jurisdiction.

5.5. Performance Properties

A naming service must provide appropriate performance guarantees to its clients. As these properties deal with the operational parameters of the naming service, interesting tradeoffs are available among them, both at design time as well as at run time (on which see Section 5.5.5).

5.5.1. Availability

The naming service as a whole is resilient to failures of individual nodes providing the naming service, as well as to failures of links among them. Intentional prevention of successful, authenticated query by an adversary should be as hard as practical.

The DNS protocol was designed to be highly available through the use of secondary nameservers. Operational practices (e.g. anycast deployment) also increase the availability of DNS as currently deployed.

5.5.2. Lookup Latency

The time for the entire process of looking up a name and other necessary associated data from the point of view of the querier, amortized over all queries for all connections, should not significantly impact connection setup or resumption latency.

5.5.3. Bandwidth Efficiency

The bandwidth cost for looking up a name and other associated data necessary for establishing communication with a given Subject, from the point of view of the querier, amortized over all queries for all connections, should not significantly impact total bandwidth demand for an application.

5.5.4. Query Linkability

It should be costly for an adversary to monitor the infrastructure in order to link specific queries to specific queriers.

DNS over TLS [RFC7858] and DNS over DTLS [RFC8094] provide this property between a querier and a recursive resolver; mixing by the recursive helps with mitigating upstream linkability.

5.5.5. Explicit Tradeoff

A querier should be able to indicate the desire for a benefit with respect to one performance property by accepting a tradeoff in another, including:

- o Reduced latency for reduced dynamic consistency
- o Increased dynamic consistency for increased latency
- o Reduced request linkability for increased latency and/or reduced dynamic consistency
- o Reduced aggregate bandwidth use for increased latency and/or reduced dynamic consistency

There is no support for explicit tradeoffs in performance properties available to clients in the present DNS.

5.6. Trust in Infrastructure

A querier should not need to trust any entity other than the authority as to the correctness of association information provided by the naming service. Specifically, the querier should not need to trust any intermediary of infrastructure between itself and the authority, other than that under its own control.

DNS provides this property with DNSSEC. However, the lack of mandatory DNSSEC, and the lack of a viable transition strategy to mandatory DNSSEC, means that trust in infrastructure will remain necessary for DNS even with large scale DNSSEC deployment.

6. Observations

On a cursory examination, many of the properties of our ideal name service can be met, or could be met, by the present DNS protocol or extensions thereto. We note that there are further possibilities for the future evolution of naming services meeting these properties. This section contains random observations that might inform future work.

6.1. Delegation and redirection are separate operations

Any system which can provide the authenticity properties in Section 5.3 is freed from one of the design characteristics of the present domain name system: the requirement to bind a zone of authority to a specific set of authoritative servers. Since the authenticity of delegation must be a protected by a chain of signatures back to the root of authority, the location within the infrastructure where an authoritative mapping "lives" is no longer bound to a specific name server. While the present design of DNS does have its own scalability advantages, this implication allows a much larger design space to be explored for future name service work, as a Delegation need not always be implemented via redirection to another name server.

6.2. Queries and assertion contexts are presently implicit

Much of the difficulty with explicit inconsistency (Section 5.4.2) derives from the fact that assertions and queries about subjects exist within a context: .local names on the local network (whether link or site local), split-DNS names within the context of the "inside" side of the recursive resolver, DNS geographic load balancing within the geographic context of the client. Because DNS provides no protocol-level support for expressing these contexts, they remain implicit.

We note that protocol-level support for this context explicit could point toward solutions for a variety of problems in currently deployed naming services, from generalized solutions with privacy/efficiency tradeoffs ({RFC7871}} aside), to explicit redirection to alternate naming resolution for "special" names [RFC6761].

6.3. Unicode alone may not be sufficient for distinguishable names

Allowing names to be encoded in Unicode goes a long way toward meeting the meaningfulness property (see Section 5.1.1) for the majority of speakers of human languages. However, as noted by the Internet Architecture Board (see [IAB-UNICODE7]) and discussed at the Locale-free Unicode Identifiers (LUCID) BoF at IETF 92 in Dallas in March 2015 (see [LUCID]), it is not in the general case sufficient for distinguishability (see Section 5.1.2). An ideal naming service may therefore have to supplement Unicode by providing runtime support for disambiguation of queries and assertions where the results may be indistinguishable.

6.4. Implicit inconsistency makes global invariance challenging to verify

DNS does not provide a generalized form of explicit inconsistency, so efforts to verify global invariance, or rather, to discover Associations for which global invariance does not hold, are necessarily effort-intensive and dynamic. For example, the Open Observatory of Network Interference performs DNS consistency checking from multiple volunteer vantage points for a set of targeted (i.e., likely to be globally variant) domain names; see <https://ooni.torproject.org/nettest/dns-consistency/>

7. IANA Considerations

This document has no actions for IANA.

8. Security Considerations

Protocols implementing name resolution systems that meet these ideal properties will have to consider tradeoffs, especially with respect to privacy (Section 5.5.4) versus performance, as in Section 5.5.5. Many properties are security and privacy relevant. All the properties in Section 5.3 must hold for a client to be able to trust that assertions about a name are as intended by the authority for that name. Section 5.1.2 specifies a property which, when it does not hold, may be exploitable for phishing attacks, and Section 5.2.3 specifies a property which may ease operational defense against malware abuse of the naming system.

9. Acknowledgments

This document is, in part, an output of design work on naming services at the Network Security Group at ETH Zurich. Thanks to the group, including Daniele Asoni, Steve Matsumoto, and Stephen Shirley, for discussions leading to this document. Thanks as well to Ted Hardie, Wendy Selzter, Andrew Sullivan, and Suzanne Woolf for input and feedback.

10. Informative References

- [I-D.ietf-dprive-dns-over-tls]
Zi, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D.,
and P. Hoffman, "Specification for DNS over TLS", draft-
ietf-dprive-dns-over-tls-09 (work in progress), March
2016.
- [I-D.ietf-dprive-dnsodtls]
Reddy, T., Wing, D., and P. Patil, "Specification for DNS
over Datagram Transport Layer Security (DTLS)", draft-
ietf-dprive-dnsodtls-15 (work in progress), December 2016.
- [IAB-UNICODE7]
IAB, ., "IAB Statement on Identifiers and Unicode 7.0.0",
n.d., <[https://www.iab.org/documents/
correspondence-reports-documents/2015-2/
iab-statement-on-identifiers-and-unicode-7-0-0/](https://www.iab.org/documents/correspondence-reports-documents/2015-2/iab-statement-on-identifiers-and-unicode-7-0-0/)>.
- [LUCID] Freytag, A. and A. Sullivan, "LUCID problem (slides, IETF
92 LUCID BoF)", n.d.,
<[https://www.ietf.org/proceedings/92/slides/
slides-92-lucid-0.pdf](https://www.ietf.org/proceedings/92/slides/slides-92-lucid-0.pdf)>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and
specification", STD 13, RFC 1035, DOI 10.17487/RFC1035,
November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S.
Rose, "DNS Security Introduction and Requirements",
RFC 4033, DOI 10.17487/RFC4033, March 2005,
<<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC5730] Hollenbeck, S., "Extensible Provisioning Protocol (EPP)",
STD 69, RFC 5730, DOI 10.17487/RFC5730, August 2009,
<<https://www.rfc-editor.org/info/rfc5730>>.

- [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>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", RFC 7871, DOI 10.17487/RFC7871, May 2016, <<https://www.rfc-editor.org/info/rfc7871>>.
- [RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/RFC8094, February 2017, <<https://www.rfc-editor.org/info/rfc8094>>.

Author's Address

Brian Trammell
ETH Zurich
Universitaetstrasse 6
Zurich 8092
Switzerland

Email: ietf@trammell.ch