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.

Table of Contents

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:

- Subject: A name, address, or name-address pair about which the naming service can answer queries

- Association: A mapping between a Subject and information about that Subject

- Authority: An entity that has the right to determine which Associations exist within its namespace

- 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:

- Reduced latency for reduced dynamic consistency
- Increased dynamic consistency for increased latency
- Reduced request linkability for increased latency and/or reduced dynamic consistency
- 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/>.

[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>.

[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