Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 9, 2016

J. Abley
Dyn, Inc.
P. Koch
DENIC
A. Durand
ICANN
W. Kumari
Google

March 08, 2016

Problem Statement for the Reservation of Top-Level Domains in the Special-Use Domain Names Registry
draft-adpkja-dnsop-special-names-problem-01

Abstract

The dominant protocol for name resolution on the Internet is the Domain Name System (DNS). However, other protocols exist that are fundamentally different from the DNS, and may or may not share the same namespace.

When an end-user triggers resolution of a name on a system which supports multiple, different protocols (or resolution mechanisms) for name resolution, it is desirable that the protocol used is unambiguous, and that requests intended for one protocol are not inadvertently answered using another.

[RFC6761] introduced a framework by which, under certain circumstances, a particular domain name could be acknowledged as being special. This framework has been used twice to reserve top-level domains (.local and .onion) that should not be used within the DNS to avoid the possibility of namespace collisions in parallel use of non-DNS name resolution protocols.

Various challenges have become apparent with this application of the guidance provided in [RFC6761]. This document aims to document those challenges in the form of a problem statement, to facilitate further discussion of potential solutions.

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 9, 2016.

Table of Contents

1. Terminology

Clear and unambiguous use of terminology is important for the clear formulation of any problem statement. The DNS protocol suffers from imprecise and overloaded terminology (e.g. see RFC7719). The use of terms and concepts from other naming systems that are similar (but different) simply confuses matters further.

In the interests of clarity, the following terms used in this document are to be interpreted as follows:

> Registry (n): the Special-Use Domain Names Registry created by [RFC6761] and published at <https://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml>

[This section to be completed following review and refinement of the rest of the text.]

2. Introduction

A number of systems use the last label in a name to act as a switch to a different, non-DNS resolution process - examples of such switches include: .local (use mDNS) and .onion (use Tor). This switch practice is not explicitly documented anywhere, and the method for accomplishing this varies by implementation. As an interesting aside, the full semantics of domain names isn't really documented anywhere either, although [Ed Lewis domain-names draft] is a current attempt to rectify this.

This technique of using the last label as a switch has a number of properties which make it attractive to people implementing alternate name resolution systems, including:

- The names can follow the common DNS syntax of LDH labels, separated by dots. This means that these names can be entered in any application which takes exiting DNS names.

- The switch to the new resolution process can be implemented in a number of ways, such as custom application code, a shim in the normal DNS resolution process, or on the system's configured DNS servers.

- The names "look" like names to users.

At this point, one should note RFC6303, which already defines "locally served zones", with the important difference that per RFC6303 the names get registered for special treatment if they are already special - they are not declared special by the registration.

[RFC6761] defines ways to reserve domain names and could be read to augment the technical exemption made in [RFC2860] (IETF-ICANN MoU):

> "Note that (a) assignments of domain names for technical uses (such as domain names for inverse DNS lookup), (b) assignments of specialized address blocks (such as multicast or anycast blocks), and (c) experimental assignments are not considered to be policy issues, and shall remain subject to the provisions of this Section 4."

The framework in [RFC6761]RFC6761 has recently been used to reserve the .onion label, allowing it to be used as a switch to the tor resolution process[RFC7686]. By the .onion label in the "Special-Use Domain Names" registry [TODO: WK - Link], The Tor Project can be assured that there will not be a .onion TLD created in the IANA rooted DNS, and thus the possibility of collisions in the namespace will be avoided.

The discussions in the DNSOP WG and the IETF Last Call processes about the .onion registration in the Special Use Domain Names registry (1,200 messages) have made it apparent that clarity about if and how to treat this "protocol switching" practice would help a lot in deciding the merit of future similar applications.

One possible outcome of the discussion would be to decline to recognize such usage of domain names in the architecture, another one is to formalize it and better understand the issues that come with it.

An additional consideration is that names which follow the DNS syntax (including those which use alternate name resolutions processes to the DNS) are in the same namespace as names in the DNS. This means that currently both the IETF (through [RFC6761]) and ICANN are making allocations or reservations from a shared namespace. If this continues to be the case, in order to avoid conflict, close coordination is necessary.

3. RFC6761

In Section 5, [RFC6761] describes seven questions to be answered to justify how and why a particular domain name is special. These seven questions can be broadly categorized as follows:

1. impact on end-users;

2. impact on applications;

3. impact on name resolution APIs and libraries;

4. impact on recursive resolvers;

5. impact on authoritative DNS servers;

6. impact on DNS server operators;

7. impact on DNS registries and registrars.

The intent of those seven questions was originally to serve as the justifications for *why* the special-use registration should be granted, demonstrating that it (a) provides a result that the community judges to be good, and (b) the aforementioned good result cannot reasonably be achieved in another way. The rough consensus from significant discussion was that .onion did satisfy both (a) and (b), but this was not clearly demonstrated by the answers to the "seven questions". Furthermore, it is unclear if and how these questions could reliably and unambiguously be used to make the determination, leading to the conclusion that they are generally inadequate for making the determination whether a particular domain name qualifies as requiring special/different treatment. Applications which follow the [RFC6761] process are likely to devolve into a "beauty contest". More over, the answers to the seven questions are not available in a machine readable form to applications that want to follow [RFC6761].

So the answers to these seven questions can better be seen as providing guidance to the corresponding seven audiences on how to handle a special-use domain name once it has been reserved by inclusion in the Registry, and not as entrance filters for inclusion in the registry.

They specify desired behavior in the internet for handling a particular domain name, not the basis for deciding whether the effort to implement special behavior across all of those audiences is worth the cost. This indifference to costs is not necessarily scalable.

The justification in [RFC6761] is concerned with the rationale of reserving a domain name that precludes its subsequent use as a generic top level domain name. However, the document fails to offer such a rationale, and instead requires the justification of the reserved name to include the provision of guidance to a number of audiences (users, application developers, DNS resolver applications, DNS resolution service operators, and name registries and registrars) as to how to handle names that are listed in this registry. But this guidance is not, in and of itself, an adequate rationale for the selection of a particular name value to be reserved in this registry.

What is missing in [RFC6761] is the consideration of the name itself. If one were to contrast the procedures relating to the admission of a name to the IETF Special Use Name registry to the processes associated with the New gTLD Program operated by ICANN, then it is evident that the IETF process does not admit many considerations which appear to be areas of evaluation in the new gTLD program. More on this in a subsequent section.

This memo proposes to categorize considerations related to the usage of RFC6761 registry for protocol switches in 3 categories: Architectural, Technical and Organizational. This memo then lists a number of questions to drive the discussion. The list of issues discussed here is non-exhaustive.

However, some voices have noted that [RFC6761] describes other alternative special handling aside from protocol switches. That alternative special handling must be considered carefully at the time of publication of the defining RFC, regardless of the nature of the special use.

4. Architectural considerations

The first thing to consider in this discussion is that not all names (or domain names) are part of the Domain Name System. See [ID-lewis-domain-names] for an in-depth discussion on this topic.

At the time of writing, two top-level domain names reserved by inclusion in the Registry are used by name resolution protocols other than the DNS and went through the [RFC6761] process:

.local is used by the Multicast DNS protocol specified in [RFC6762] which is similar in some respects to the DNS, but which uses a different well-known port number and is limited to a particular multicast scope;

ONION is used to construct names that designate services reachable via the Tor network using onion routing.

The two name resolution protocols described above are, to varying degrees, different from the DNS, and the namespaces used in each naming scheme are also different (albeit similar, in the .local case). The top-level label is effectively being used as a name resolution protocol identifier. At the core of the issue is that different "strings" that look like "domain names" (i.e. are within the same name space) but are not DNS names are used interchangeably in the URI (or URN). In particular, DNS imposes constraints on name syntax. An example of such constraints is the 64 octet limit per label. Strings used in the ONION domain do not have that constraint. It could be argued that in the absence of a more elegant alternative, a pragmatic choice to embed protocol selectors as namespace tokens has effectively already been made. The running code and effective consensus in how it should be used by significant user bases should not be discounted. Although the reservation of names in the DNS namespace can be made at any level, the two examples above demonstrate use-cases for reservation at the top-level, and hence that case must be considered.

The underlying discussion here is the tussle between the applications and the network. Application architects see using special name tags (a la .onion) as an easy way to get new features deployed. They consider the hurdles of deploying new URI schemes such as http:/onion/onion-name as too onerous and too slow to deploy for their needs. Network architects worry of overloading the semantics of DNS names and/or creating a name space that is larger than the DNS namespace. They refer to bad precedents such as .uucp and .bitnet.

The fundamental point to consider here is the unicity (or multiplicity) of the name space. Are we talking about one namespace with different resolution protocols or independent name spaces?

It might it be helpful to point out that the property of interest here is the assurance of uniqueness of a name, and another way of thinking about the question is whether it applies across domain names as people expect or need it to? None of this would matter if people didn't expect names constructed according to whatever rules they're following to be unique across a set of names that spans multiple operating environments and resolution protocols.

In [RFC2826] the IAB noted that

> "To remain a global network, the Internet requires the existence of a globally unique public name space. The DNS name space is a hierarchical name space derived from a single, globally unique root."

> "Maintaining a globally-unique public namespace that supports different name resolution protocols is hence an architectural requirement, and some facility for reservation of top-level domains in the DNS is necessary."

If we were to accept the notion that the last label of a domain name is actually a protocol switch, we are actually building a catalog of all top level domains and what resolution protocol each one invokes. Note that such a catalog does not formally exist today, as [RFC6761] is an exception list to the general case which is supposed to use regular DNS as resolution protocol. Such a catalog may remain a concept to guide this discussion or be implemented as an actual IANA registry. In effect, it would associate TLDs with indications on how applications and resolvers should treat them. However, such an approach would leave open the question of not-yet-defined TLDs. No resolution mechanism could be associated with those.

It should also be noted that there are choices for a protocol switch other than reserving labels. In particular, a proposal to move those protocol switches under a specific top level domain has been discussed (.ALT). If that architecture choice is made, some of the questions listed in the sections below would become moot.

Note: [RFC6761] mentions the reserved names could be any label in any random string, not just the rightmost one (or ones). However, this creates a number of complications and has not seen much support in the community as of now.

5. Technical considerations

Each of the seven questions posed by [RFC6761] has the potential to describe why special handling of the requested name(s) in applications by a particular audience may be necessary. However, aside from reserving the name, it is not entirely clear what any of those audiences might further expect as a result of a successful request to add a top-level domain to the Registry.

For example, reservation of a top-level domain by the IETF does not guarantee that DNS queries for names within a reserved domain will not be sent over the Internet. The requirements of the operators of recursive resolvers in the DNS cannot be relied upon to be implemented; the impact on the operators of DNS authoritative servers hence cannot be reliably assumed to be zero. In the case of [I-D.ietf-dnsop-onion-tld], leakage of .onion queries on the Internet might lead to disclosure of private information that, in some cases, might pose a risk to the personal safety of end-users.

At the time of writing, the [RFC6761] registry does not include direct guidance for any of the seven audiences, relying instead upon a reference for each entry in the Registry to the document that requested its insertion. Such documents might well be opaque to many readers; ([RFC6762] is a seventy-page protocol specification, for example, which is arguably not the most effective way to set expectations of non-technical end-users).

Useful reservations of top-level domains should be accompanied by documentation of realistic expectations of each of the seven audiences, and the evaluation of particular requests should consider the practical likelihood of those expectations being met and the implications if they are not.

Here is a non-exhaustive list of additional questions that have surfaced in discussion of requests for names to be added to the Special Use Names registry:

What does it mean to have a "non-DNS" entry in the registry described above?

Are applications supposed to check that registry to know what to do?

Can/Should applications do this check dynamically?

What if an application makes this dynamic check and realizes the name contains a switch it does not know how to treat?

Similar questions applies to resolvers (DNS and non-DNS); what is the expected behavior?

One particular avenue of investigation would be to see if such considerations could be encoded in machine understandable code in an extension of the current [RFC6761] registry.

6. Organizational considerations

Organizational considerations can be broken down in two categories, internal and external.

6.1. Non-exhaustive list of external organizational considerations

The policy surrounding the implementation and management of top-level domains in the DNS has been developed using a multi-stakeholder process convened by ICANN according to the MoU between ICANN and IETF [RFC2860]. It is out of scope for this document to revisit that MoU.

Whilst discussing the particular attributes that make a domain name special, [RFC6761] notes that "the act of defining such a special name creates a higher-level protocol rule, above ICANN's management of allocatable names on the public Internet."

[RFC2860] draws a line between what is policy and what is technical. A variety of opinions have been expressed regarding whether [RFC6761] blurs this line. In particular, see http://www.circleid.com/posts/20151222_whats_in_a_name/ for a certain viewpoint on the topic. As noted earlier, it is out of scope for this document to analyse this issue beyond noting that such a variety of views exist.

Taking a different perspective, it has been argued that [RFC6761] specifically extends the DNS protocol to include special treatment for names in the registry, and that there's nothing in 2860 at all that limits the IETF's authority to change the protocol.

However, it should be noted that, if the IETF were to formalize the concept of protocol/name switch in the Internet architecture, coordination would be require between ICANN and IETF on such names. Using the analogy described above of a catalog/registry of such switches, care must be taken to make sure we do not end up with 2 process streams allowed to create entries without any synchronization.

6.2. IETF Internal considerations

6.2.1. Process

[RFC6761] specifies the way in which "an IETF 'Standards Action' or 'IESG Approval' document" should present answers to the questions described above (see Section 2), but does not describe the process by which the answers to those questions should be evaluated.

For example, it is not clear who is responsible for carrying out an evaluation. A document which requests additions to the Registry might be performed by the IESG, by the IAB, by the DNSOP working group, by an ad-hoc working group, by expert review or any combination of those approaches. [RFC6761] provides no direction.

As an illustration of the inconsistency that has been observed already, [RFC6762] was published as an AD-sponsored individual submission in the INT area, and the IESG evaluation record does not reveal any discussion of the reservation of the .local top-level domain in the DNS. [I-D.ietf-dnsop-onion-tld], however, was published as a working group document through DNSOP, and an extensive discussion by both the participants of DNSOP and the IESG on the

merits of the request took place. The evaluation process, in the absence of clear direction, is demonstrably inconsistent.

We should point to RFC 5226 and explicitly quote the definition of "Standards Action" or "IESG Approval":

> IESG Approval is not intended to be used often or as a "common case"; indeed, it has seldom been used in practice during the period RFC 2434 was in effect. Rather, it is intended to be available in conjunction with other policies as a fall-back mechanism in the case where one of the other allowable approval mechanisms cannot be employed in a timely fashion or for some other compelling reason. IESG Approval is not intended to circumvent the public review processes implied by other policies that could have been employed for a particular assignment. IESG Approval would be appropriate, however, in cases where expediency is desired and there is strong consensus for making the assignment (e.g., WG consensus).

So, while it is very interesting to note that [RFC6761] was an AD sponsored individual submission in spite of two active DNS related WGs, 6762 is probably clean: it defines the protocol and is itself on standards track.

RFC 7686 however, while on standards track, does not define the TOR protocol, so it was used to fulfill the 'standards action' requirement by the letter. It contains normative references to non-IETF protocols, which is noteworthy.

A comparison of the two '7 question forms' reveals that at least the responses to questions 2, 3, and 4, differ significantly while there is no defined way to communicate the difference to the affected software entities.

An alternate view has been expressed with regard to the protocol evaluation. It states that the authority belongs to the IESG to seek whatever support it likes, within the established process, in making standards decisions, including delegating evaluation of a specific registry change proposal to a WG or a directorate. The IESG might have varied what guidance it sought, but that does not constitute "inconsistency" under the process. That being said, more complete evaluation guidance would be helpful to the IESG and the community.

6.2.2. Technical criteria

Regardless of the actual name being proposed as protocol and/or namespace switch, it is also not clear what technical criteria the evaluation body should use to examine the merit of an application for

such a reserved name/protocol switch. For example, is large scale prior deployment an acceptable criteria? A number of voices have clearly answered "no" to that question as it would only encourage "squatting" on names.

However, in the case of .local and .onion, those particular domain names were already in use by a substantial population of end-users at the time they were requested to be added to the Registry. Rightly or not, the practical cost of a transition away from the requested strings was argued as a justification for their inclusion in the registry.

6.2.3. Name evaluation

With regard to the actual choice of name, [RFC6761] is silent. The answers to the seven questions are expected to tell how a name, presumably already chosen outside of the process, might be handled if it is determined to be a "special use" name. However, it is silent on how to choose a name or how to evaluate a specific proposed name.

6.2.4. The ICANN process to evaluate names

Section 4.3 of [RFC2860] says:

> Two particular assigned spaces present policy issues in addition to the technical considerations specified by the IETF: the assignment of domain names, and the assignment of IP address blocks.

This remains as true today as when it was written (2000). Domain names have a number of considerations that have complex policy issues that ICANN deals with and which the IETF may not be well equipped to handle.

The ICANN process applicant have to go through to get a name is described in the applicant guide book https://newgtlds.icann.org/en/applicants/agb/guidebook-full-04jun12-en.pdf which is a 338 page document. It should however be noted that the current round of gTLD application is closed and rules may differ in the next round if and when it happens.

Considerations include, but are not limited to:

Geographical During the most recent round of new gTLD applications, there were a number of applications for so call "geographic" terms. These included applications for .amazon and .patagonia. The .amazon application in particular was controversial - the governments of Brazil and Peru requested that ICANN's Governmental

Advisory Committee (GAC) to issue a warning that granting .amazon to Amazon would "prevent the use of this domain for purposes of public interest related to the protection, promotion, and awareness raising on issues related to the Amazon biome." The IETF is not well suited to evaluating this sort of issue.

Brands / Trademark law If Wile E. Coyote approached the IETF requesting that the IETF reserve .acme, a trademark held by a large corporation making anvils and giant slingslots, the IETF could become embroiled in trademark lawsuit - and even if the IETF were not, we have enough armchair lawyers that the discussions would be extremely annoying :-). Closely related to this issue is "protected designation of origin (PDO)" - for example, Champagne.

String similarity ICANN has an entire process for evaluating the string similarity / confusability between applied for (and current) strings - for example, under what conditions would the IETF be able to make a determination if someone attempted to use RFC6761 to reserve .c0m?

International Organization Names Certain names and organizations get additional protection under trademark law - well known examples of this are the RedCross/RedCrescent and the International Olympic Committee (IOC). Whether or not this should be the case is well outside anything that the IETF should have an opinion on but, undoubtedly, there are many within the community who will have an opinion (and will want to argue it ad nauseam :-))

Offensive Terms There are a huge range of these, from the obscure / archaic (waesucks, gadsbudlikins) to the more obvious and current ([xml2rfc-error], [xml2rfc-error] and [xml2rfc-error]. Certain terms are sufficiently offensive that the IETF would have a hard time coming to any useful consensus (other then "Eeeew!")

Going back to the IETF process used for the evaluation of .local and .onion, one might ask the following questions:

For example, what consideration have there been in the intellectual property rights in the reservation of a name in this Special Use Name registry, and what procedures should be followed in the case of a dispute over the rights to use a name in this manner? Also, to what extent could such a reservation of a name in this Special Use Names registry be used to block competing interests and/or competing technologies? What are the competition and consumer issues that need to be considered if the reservation of a name in this registry causes some form of exclusive access and reduced competitive access, or where there is no ability for

consumers to exercise choice in a situation where providers compete in the offering of services?

A related consideration is that the current process of admission to the Special Use Name registry appears to admit no formal assessment of environmental impact. Is the name that is proposed to be entered into this registry already being used in local contexts, with or without an association with DNS name resolution, such that its use as a reserved name through an entry in this registry, and its continued use in local contexts could cause harm to users? To what extent can this impact be assessed, and what level of impact is considered acceptable?

While the "seven questions" relate to altered behaviours by specific audiences and users of names there is no explicit consideration of the security in this process. Is the registration of such a name a "safe" action for the IETF to take? To what extent could the use of this reserved name be used in a hostile or malicious manner? What measures have been taken to mitigate or otherwise address such potential vulnerabilities?

ICANN has created an entire set of groups, organizations, committees, processes and procedures to deal with the evaluation of applied for new TLDs, complete with a cadre of lawyers and policy people. Unless the IETF were willing to do the same, it would have a hard time performing evaluation of the strings themselves, distinct from the evaluation of the technology behind the name resolution system.

An alternate view has been expressed, that such a process is not necessary because the IESG is the body that makes the decision on a specific name reserved by RFC6761, and the IETF has a workable appeal process to deal with any potential issues. However, looking at the level of contention created in the ICANN process around the choice of certain names, serious doubts have been expressed to the scalability and ultimate viability of such an appeal process.

7. Security Considerations

This document aims to provide a problem statement that will inform future work. Whilst security and privacy are fundamental considerations, this document expects that future work will include such analysis, and hence no attempt is made to do so here. See among other places SAC-057 [https://www.icann.org/en/system/files/files/sac-057-en.pdf]

Reserving names has been presented as a way to prevent leakage into the DNS. However, instructing resolvers to not forward the queries (and/or by instructing authoritative servers not to respond) will

garantee that such leakage will not happen. The security (or privacy) of an application MUST NOT rely on names not being exposed to the Internet DNS resolution system.

8. IANA Considerations

This document has no IANA actions.

9. Acknowledgements

Your name here, etc.

10. References

10.1. Normative References

[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>.

[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>.

[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>.

[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>.

10.2. Informative References

[I-D.ietf-dnsop-dns-terminology] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", draft-ietf-dnsop-dns-terminology-05 (work in progress), September 2015.

[I-D.ietf-dnsop-onion-tld]
Appelbaum, J. and A. Muffett, "The .onion Special-Use Domain Name", draft-ietf-dnsop-onion-tld-01 (work in progress), September 2015.

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

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC2826] Internet Architecture Board, "IAB Technical Comment on the Unique DNS Root", RFC 2826, DOI 10.17487/RFC2826, May 2000, <http://www.rfc-editor.org/info/rfc2826>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <http://www.rfc-editor.org/info/rfc6762>.

Appendix A. Editorial Notes

This section (and sub-sections) to be removed prior to publication.

A.1. Venue

An appropriate forum for discussion of this draft is for now the dnsop working group.

A.2. Pithy Quotes from History

> The question has arisen as to how the toplevel naming authority decides who gets a toplevel name and who must get by with a non-toplevel name. The suggestion was made by MOCKAPETRIS@USC-ISIF that perhaps the existing toplevel nameholders might vote on whether the applicant for a new toplevel name should be granted, with a majority needed for approval. It seems to me this might produce a clique whereby whoever initially gains power will hold it and prevent its "enemies" from getting in too. This will make the toplevel rather less than universal.

(E-mail from Robert Elton Maas to the namedroppers mailing list on 9 November 1983)

> My basic point is that as a world-wide network evolves it is ridiculous to force people to name resources in terms of one

> static hierarchy which very closely resembles the current internetwork topology (as the current scheme does). What we are eventually going to require is a distributed expert for making sense out of a name someone hands it. There will be no simple algorithm to be written on one page of an RFC that will suffice to resolve a name. Rather, a number of heuristics will let a resolver make sense out of a given name by querying other experts which it suspects may be more knowledgeable about the name than it is, or by forwarding a piece of mail to an expert which is at least one level closer to the destination in some hierarchy.

(E-mail from Peter Karp to the namedroppers mailing list on 8 February 1984)

A.3. Change History

A.3.1. draft-adpkja-special-names-problem-00

Initial draft circulated for comment.

Appendix B. Change history

[RFC Editor: Please remove this section before publication]

-00 to -01:

- Significant readability changes.
- [WK: Stopped at end of Sec 3]

-00:

- Initial draft circulated for comment.

Authors' Addresses

Joe Abley
Dyn, Inc.
103-186 Albert Street
London, ON N6A 1M1
Canada

Phone: +1 519 670 9327
Email: jabley@dyn.com

Peter Koch
DENIC

Email: pk@denic.de

Alain Durand
ICANN

Email: alain.durand@icann.org

Warren
Google

Email: warren@kumari.net