

Network Working Group
Internet-Draft
Intended status: Informational
Expires: October 7, 2016

T. Lemon
Nominum, Inc.
R. Droms
Cisco, Inc.
April 5, 2016

Special Use TLD Problem Statement
draft-tldr-sutld-ps-00

Abstract

The Special-Use Domain Names registration policy in [RFC 6761](#) has been shown through experience to present unanticipated challenges. This memo explores current IETF and non-IETF publications relating on special-use TLDs, describes the history of domain names, and enumerates some problems that have come up in connection with the Special-Use Domain Names allocation process specifically as it related to top-level domains.

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 October 7, 2016.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

- [1. Introduction](#) [2](#)
- [2. The Status Quo in the IETF regarding SUTLDs](#) [3](#)
 - [2.1. Primary SUTLD Documents](#) [3](#)
 - [2.1.1. IAB Technical Comment on the Unique DNS Root](#) [3](#)
 - [2.1.2. Special-Use Domain Names](#) [3](#)
 - [2.2. Secondary documents relating to the SUTLD question](#) [4](#)
 - [2.2.1. The .onion Special-Use TLD](#) [4](#)
 - [2.2.2. Locally Served DNS Zones](#) [4](#)
 - [2.2.3. Name Collision in the DNS](#) [5](#)
 - [2.2.4. Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis](#) [5](#)
 - [2.2.5. Additional Reserved Top Level Domains](#) [5](#)
 - [2.3. Observations based on reading these documents](#) [5](#)
- [3. History](#) [7](#)
- [4. Problems with Special-Use TLDs and Their Allocation](#) [8](#)
 - [4.1. Purity of Domain Name Architecture Problem](#) [8](#)
 - [4.1.1. Are Domain Names DNS Names?](#) [8](#)
 - [4.1.2. Does Every Domain Name Have The Same Meaning Everywhere?](#) [9](#)
 - [4.1.3. Protocol Switch Problem](#) [9](#)
 - [4.2. Land Rush Problem](#) [10](#)
 - [4.3. IESG Getting Sued Problem](#) [10](#)
 - [4.4. Warranted Allocation Problem](#) [11](#)
 - [4.5. Experimental Squatting Problem](#) [11](#)
 - [4.6. Insecure Delegation Problem](#) [12](#)
- [5. Contributors](#) [12](#)
- [6. Informative References](#) [13](#)
- [Authors' Addresses](#) [14](#)

1. Introduction

At the time of this writing, the IETF has recently been asked to allocate several new special-use top-level domains (SUTLDs), and has experienced several different sorts of problems with the process as documented in Special-Use Domain Names [[RFC6761](#)]. Because of this, the IETF has decided to investigate the problem and decide if and how the [RFC 6761](#) process can be improved, or whether it should be deprecated.

In order to investigate this question, this document attempts first to summarize the status quo, then to briefly discuss the history of domain names, and finally to describe the set of problems that exist

Lemon & Droms

Expires October 7, 2016

[Page 2]

as reported by various IETF participants with experience in the various aspects of the problem.

2. The Status Quo in the IETF regarding SUTLDs

There are two primary and several secondary documents to consider when thinking about the SUTLD process.

2.1. Primary SUTLD Documents

2.1.1. IAB Technical Comment on the Unique DNS Root

At present there are two primary documents to be aware of that address the topic of SUTLDs. The first is the IAB Technical Comment on the Unique DNS Root [[RFC2826](#)]. This document is not an IETF consensus document, and appears to have been written to address a different problem than the SUTLD problem. However, it speaks directly to several of the key issues that must be considered, and of course coming as it does from the IAB, it is rightly treated with a great deal of authority despite not being an IETF consensus document. This document should be considered required reading for IETF participants who wish to express an informed opinion on the topic of SUTLDs.

2.1.2. Special-Use Domain Names

The second important document is Special-Use Domain Names [[RFC6761](#)]. [RFC 6761](#) represents the current IETF consensus on designating and recording SUTLDs. The IETF has experienced problems with the designation process described in [RFC 6761](#), which motivate this document. Again, familiarity with this document is a prerequisite for having an informed opinion on the topic of SUTLDs.

[RFC 6761](#) defines two aspects of SUTLD: designating a domain name to have a special purpose and registering that special use in the Special-Use Domain Names registry. The designation process is defined in a single sentence ([RFC 6761, section 4](#)):

If it is determined that special handling of a name is required in order to implement some desired new functionality, then an IETF "Standards Action" or "IESG Approval" specification [[RFC5226](#)] MUST be published describing the new functionality.

This sentence implies that the designation be subject to the same open review and consensus process as used to produce and publish all other IETF specifications.

The registration process is a purely mechanical process, in which the existence of the newly designated special use name is recorded, with a pointer to a section in the relevant specification document that defines the ways in which special handling is to be applied to the name.

[RFC 6761](#) provided the explicit designation and inclusion in the Special-Use Names registry for '.local' and for a set of names that had been previously used or defined to have special uses before the publication of [RFC 6761](#).

2.2. Secondary documents relating to the SUTLD question

In addition to these documents, there are several others with which participants in this discussion should be familiar.

2.2.1. The .onion Special-Use TLD

The first of these is The .onion Special Use TLD [[RFC7686](#)]. This document is important because it is the most recent IETF action on the topic of SUTLDs; although it does not set new policy, the mere fact of its publication is worth thinking about.

The two important points to consider about this document are that:

- o The IETF gained consensus to publish it
- o The situation was somewhat forced, both by the fact of its unilateral allocation by the TOR project without following the [RFC 6761](#) process, and because a deadline had been set by the CA/Browser Forum [[SDO-CABF-INT](#)] after which all .onion PKI certificates would expire and no new certificates would be issued, unless the .onion TLD were recognized by the IETF.

2.2.2. Locally Served DNS Zones

Locally Served DNS Zones [[RFC6303](#)] describes a particular use case for zones that exist by definition, and that are resolved using the DNS protocol, but that cannot have a global meaning, because the host IP addresses they reference are not unique. This applies to a variety of addresses, including Private IPv4 addresses [[RFC1918](#)], Unique Local IPv6 Unicast Addresses [[RFC4193](#)] (in which this practice was first described) and IANA-Reserved IPv4 Prefix for Shared Address Space [[RFC6598](#)].

This use case is distinct from the use-case for SUTLDs like '.local' and '.onion' in that the names are resolved using the DNS protocol, but shares the problem that such names cannot be assumed either to be

unique or to be functional in all contexts for all internet-connected hosts.

2.2.3. Name Collision in the DNS

Name Collision in the DNS [[SDO-ICANN-COLL](#)] is a study commissioned by ICANN that attempts to characterize the potential risk to the Internet of creating a registration process for new top-level domains that would allow corporations and persons to, for a fee, request a delegation for a particular top-level domain. The document concludes that such allocations do present some risks, and mentions some specific top-level domains that present sufficient risk that they must never be allocated by ICANN.

2.2.4. Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis

Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis [[RFC7050](#)] is an example of a document that successfully used the [RFC 6761](#) process to designate '.ipv4only.arpa' as a special-use name; in this case the process worked smoothly and without controversy.

2.2.5. Additional Reserved Top Level Domains

Additional Reserved Top Level Domains [[I-D.chapin-additional-reserved-tlds](#)] is an example of a document that attempted to reserve several TLDS identified by ICANN as particularly at risk for collision as special-use domain names with no documented use. This attempt failed, for better or for worse.

2.3. Observations based on reading these documents

Readers of this memo are strongly encouraged to familiarize themselves with the documents referenced above, and to draw their own conclusions. However, for the sake of beginning the discussion, some observations are presented here.

The IAB Technical note makes some important points:

- o The Internet requires a globally unique namespace
- o Private networks may operate private namespaces, but still require that names in the public namespace be globally unique.
- o The Domain Name System [[RFC1035](#)] is not the only protocol that may be used for resolving domain names.
- o Users cannot be assumed to know how to distinguish between symbols that have local meaning and symbols that have global meaning.

Users may therefore share symbols that incorporate domain names with no global meaning (for example, a URL of 'http://mysite.example.corp', where 'example.corp' is a domain allocated privately and informally as described in [[SDO-ICANN-COLL](#)]). Such symbols might refer to the object the user intends to share within that user's context, but either refer to some other object or objects in the recipients' context, or might not refer to any object at all in the recipients' context. The effect of this is that the user's intended communication will not be able to be understood by the recipients of the communication. This same problem can also occur simply because a single user copies a name from one context in which it has one meaning, into a different context in which it has a different meaning-- for example copying a '.onion' URL out of an email message, where it has no meaning, into a TOR browser, where it does.

There are several important conclusions that can be taken from this. First, domain names with unambiguous global meaning are preferable to domain names with local meaning which will be ambiguous. Second, at least at the time of the writing of this document the IAB was of the opinion that there might well be more than one name resolution protocol used to resolve domain names.

There are at least three important points to think of with respect to the [RFC6761](#):

- o A special-use name may be a name that should be resolved using the DNS protocol with no special handling.
- o A special-use name may be a name that should be resolved using the DNS protocol, but requires special handling in the resolver.
- o A special-use name may be a TLD like '.local' which acts as a signal to indicate that the local stub resolver should use a non-DNS protocol for name resolution.
- o The current IETF consensus (from a process perspective, not necessarily from the perspective of what would be consensus if the IETF were to attempt to produce a new consensus document) is that both of these applications for special-use names are valid.

It should be noted also that [RFC6761](#) does not limit special-use names to TLDs. However, at present, all special-use names registered in the IANA Special-Use Domain Names registry [[SDO-IANA-SUDR](#)] are either intended to be resolved using the DNS protocol, or are top-level domains, or both. That is, at present there is no protocol switch

required by [RFC6761](#) for names other than at the top level of the naming hierarchy.

However, the obvious corollary to this is that at present, [RFC6761](#) does require that stub resolvers for hosts that participate in mDNS [[RFC6762](#)] service discovery [[RFC6763](#)] or name resolution to use the presence of a special label, '.local', to indicate that mDNS, rather than DNS, be used to resolve names under that label.

Neither [RFC 6761](#) nor [RFC 7686](#) require any kind of delegation for '.onion' or '.local'. At least in the case of '.onion', it might be better if there were a delegation in order to limit the scope of leakage of '.onion' queries made to resolvers that do not know how to resolve '.onion' names. However, the DNS protocol does not currently provide a mechanism for signaling that a particular name is not intended to be resolved using the DNS protocol, and it might be better to address that problem before defining a delegation for '.onion'.

3. History

Newcomers to the problem of resolving domain names may be under the mistaken impression that the DNS sprang, like the greek legend of Athena, directly from Paul Mockapetris' forehead. This is not the case. At the time of the writing of the IAB technical document, memories would have been fresh of the evolutionary process that led to the DNS' dominance as a protocol for domain name resolution.

In fact, in the early days of the Internet, hostnames were resolved using a text file, HOSTS.TXT, which was maintained by a central authority, the Network Information Center, and distributed to all hosts on the internet using the File Transfer Protocol (FTP) [[RFC0959](#)]. The inefficiency of this process is cited as a reason for the development of the DNS [[RFC1034](#)] in 1983.

However, the transition from HOSTS.TXT to the DNS was not smooth. For example, Sun Microsystems's Network Information System [[CORP-SUN-NIS](#)], at the time known as Yellow Pages, was an active competitor to the DNS, although it failed to provide a complete solution to the global naming problem.

Most modern operating systems can still use the '/etc/hosts' file for name resolution. Many still have a name service switch that can be configured on the host to resolve some domains using NIS. Most have the capability to resolve names using mDNS by recognizing the special meaning of the '.local' SUTLD.

The Sun Microsystems model of having private domains within a corporate site, while supporting the global domain name system for off-site persisted even after the NIS protocol fell into disuse. Microsoft used to recommend that site administrators allocate a "private" top-level domain for internal use, and this practice was very much a part of the zeitgeist at the time. This attitude is the root of the widespread practice of simply picking an unused top-level domain and using it for experimental purposes, which persists even at the time of writing of this memo.

4. Problems with Special-Use TLDs and Their Allocation

[RFC 6761](#) serves two purposes. First, it enumerates a set of domain names that require special treatment either by stub resolvers or by authority and caching name servers. Second, it establishes a registry for future such names, and in [section 5](#) it describes considerations for deciding when such names should be allocated. This has been a success in the sense that this registry now exists and has been used, but the stated considerations don't work for all use cases, and there are questions about whether this is even the right approach. This section attempts to enumerate these problems in no deliberate order.

4.1. Purity of Domain Name Architecture Problem

There is some question as to whether special-use TLDs may violate the Domain Name System Architecture. Much discussion has centered around three aspects of this question. There is, of course, no document titled "The Domain Name System Architecture," so we have to rely on the various documents referenced earlier to consider this question. [RFC 2826](#) seems to speak most directly to it.

4.1.1. Are Domain Names DNS Names?

The first of the three aspects of the purity question is whether domain names are specific to the DNS protocol. [RFC 2826](#) appears to state unequivocally that they are not. However, many people have expressed the opinion that they are. This is a question that may need to be explored further through the IETF process.

At present, there are three obvious exceptions to the use of the DNS protocol to resolve domain names: '.onion', '.local' and any name that appears in /etc/hosts or the equivalent. If indeed DNS is the only protocol to be used for resolving domain names, these use cases would have to be deprecated.

4.1.2. Does Every Domain Name Have The Same Meaning Everywhere?

The second aspect of the purity question is the question of whether at any particular time, every domain name refers to the same entity regardless of where the query comes from.

It has been stated by various participants in this discussion that in fact they do. [RFC 2862](#) seems to make a very strong case that to the extent possible they should be, but does not go so far as to insist that in all cases they must be. '.onion' actually should exhibit this property. However, '.local' does not.

There are other examples of domain names that cannot be assumed to have the same meaning everywhere. For example, it is not unreasonable for a site that uses [RFC 1918](#) addresses to populate the DNS reverse lookup tree for their copy of the address space. Fortunately, names of this sort are rarely seen by users who do not understand their limited scope of applicability, making this exception a violation in fact but not in spirit of this principle.

There is the additional case of names that refer to different IP addresses depending on the location of the host that originates the query. Although at first glance this appears to be an example of the problem, it really is not, because the entity to which the name refers is the same entity, cached in different locations.

Something similar can be said about domain names that refer to objects that may be harmful to devices that access them. A caching resolver may refuse to answer queries for such names, or may return answers that differ from the answer that the name server authoritative for that name is returning. Again, although this seems like an exception, it probably is not, because in this case the user who is attempting to share a reference to the entity is either malicious or has been duped, and in either case a failure to successfully share the entity is, presumably, an intended outcome from the perspective of the recipient of the reference.

If it were to be determined that the principle expressed in [RFC 2862](#) that domain names should always refer to the same object is too weak, this would necessitate changing the operational model of Multicast DNS, and would render the use of names derived from non-unique IP addresses at least in principle out of bounds.

4.1.3. Protocol Switch Problem

The third of the three aspects of the purity question is the problem that if we do not resolve all domains using the same protocol, there has to be a mechanism for determining which protocol to use to

resolve any particular name. If the names that must be resolved using different protocols all exist directly under the root of the namespace, this determination is relatively simple, but clearly using the DNS protocol for all queries would be simpler.

An additional consideration about this aspect of the question is whether future resolver implementations should be required to engage in any behavior to determine whether a particular TLD for which they have no special information is a SUTLD or a normal TLD. Also, since there can be privacy implications to special-use TLD queries, does it make sense to require that all resolvers take some action to test the special-use-ness of each TLD for which they are asked to resolve queries before sending the complete query out over the wire using the DNS Protocol.

4.2. Land Rush Problem

The land rush problem refers to the concern that has been expressed that if the IETF is seen to be able to allocate TLDs outside of the ICANN process, persons and entities who might otherwise apply to ICANN for the use of a particular TLD might instead apply to the IETF. The concern is that this could turn into a DoS attack on the IESG, particularly since one option for SUTLD allocation is "IESG Approval."

It is therefore important to clarify what sorts of use cases make sense for both the "Standards Action" and "IESG Approval." It seems clear that if there are good use cases for SUTLDs, they would be in the production of standards that the IETF considers important, not for the purpose of assigning domains for normal delegation in the DNS.

It is not the case that all special-use TLDs can be expected to be non-DNS names; for example, the Homenet Working Group is likely to propose the use of a special-use TLD for use on the homenet in cases where the homenet does not have its own globally unique name allocated. This would nevertheless be resolved in the same way as an [RFC1918](#) reverse query, by sending a query to the local resolver using the DNS protocol.

4.3. IESG Getting Sued Problem

An additional serious problem in allocating SUTLDs is that it exposes the IETF to the risk of lawsuits from interested parties. While this is obviously a real risk, it is not special to the SUTLD problem--the IETF is always potentially at risk of being sued. However, ICANN has developed a process for evaluating the issues that might arise with respect to the allocation of a TLD. The IETF could in principle take

advantage of this evaluation process through the agency of the IANA, which at present is a function managed by ICANN.

4.4. Warranted Allocation Problem

The issue here is similar to the Land Rush problem: the IETF is being asked to evaluate whether a particular protocol use of a SUTLD justifies the permanent sacrifice of that name in the registry, and the appearance of that name in the name service protocol switches of various stub resolver implementations.

Several things could be considered here. One is that there exist processes for other registries where names are temporarily reserved; if the protocol for which they were reserved turns out to be a success [[RFC5218](#)], then the reservation can be made permanent. If not, it's deleted from the registry.

A corollary to this is that if a protocol is not a success, it likely will not be implemented in the protocol switches of mainstream host operating systems. This mitigates this aspect of the damage that would be caused by a mistake here.

The final point to be made about this concern is that it is not at all unique to SUTLDs. The IETF frequently produces protocols that do not turn out to be successes by the criteria set forth in [RFC5218](#). Each of these protocol efforts began with the IETF chartering process, the working group consensus process, and the IETF and IESG review processes. These processes do not prevent unsuccessful protocols from being produced, but they do very much limit the rate at which they are produced, and, it is to be hoped, improve the ratio of failures to successes. There is no reason to think that this particular IANA registry is special in this regard, nor to treat it specially simply because we are aware of this problem.

4.5. Experimental Squatting Problem

One of the problems with SUTLDs is that various protocol development efforts find it convenient to use SUTLDs as a name service protocol switch precisely because such switches exist in most popular operating systems. So a project will allocate a name without going through the [RFC 6761](#) process, and then later on will want that allocation to be made permanent. Several organizations that have actually done this have indicated that they looked for a process they could use to get a name, found [RFC6761](#), realized it wouldn't work for them, found nothing else, and then just unilaterally allocated the name.

Precisely because of the principle stated in [RFC 2826](#), that names that do not consistently refer to the same object create problems for end users, this creates pressure once use of the protocol becomes somewhat widespread, yet not widespread enough or strategic enough to justify allocation by the IETF, for the organization that squatted on the name to try to get the IETF to formalize the allocation.

Solutions have been proposed for this problem. There already exists a TLD, '.test', under which such an allocation could in principle be made, but no RFC exists recommending this practice. Furthermore, [RFC 6761](#) recommends resolving .test using the DNS, so this isn't exactly the right thing for a name service protocol switch.

A draft exists to allocate a '.alt' TLD to address this need. The draft as currently written attempts to solve the whole SUTLD problem, but fails to do so. This draft could easily be restricted to simply defining a SUTLD under which new non-DNS resolution strategies could be tested, and this would likely resolve the bulk of the experimental squatting problems that might come up in the future.

[4.6. Insecure Delegation Problem](#)

At present [RFC 6761](#) doesn't give specific advice about what delegation, if any, should appear for newly registered SUTLDs. It may be worthwhile to make a default recommendation for this. This may tie in with the idea that SUTLDs should have some signal in the DNS that tells the resolver not to try to use DNS to resolve names under that TLD, even if they don't have an alternative protocol to try.

[5. Contributors](#)

This document came about as a result of conversations that occurred the weekend before IETF 95 in the lobby. Stuart Cheshire, Mark Andrews, David Conrad, Paul Ebersman and Aaron Falk all made helpful and insightful observations or patiently answered questions. This should not be taken as an indication that any of these folks actually agree with what the document says, but their generosity with time and thought are appreciated in any case. Ralph started out as an innocent bystander, but discussion with him was the key motivating factor in the writing of this document, and he agreed to co-author it without too much arm-twisting. And this document owes a great deal to Ed Lewis' excellent work on what a "domain name" is [[I-D.lewis-domain-names](#)].

6. Informative References

- [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>>.
- [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>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), DOI 10.17487/RFC4193, October 2005, <<http://www.rfc-editor.org/info/rfc4193>>.
- [RFC5218] Thaler, D. and B. Aboba, "What Makes For a Successful Protocol?", [RFC 5218](#), DOI 10.17487/RFC5218, July 2008, <<http://www.rfc-editor.org/info/rfc5218>>.
- [RFC6303] Andrews, M., "Locally Served DNS Zones", [BCP 163](#), [RFC 6303](#), DOI 10.17487/RFC6303, July 2011, <<http://www.rfc-editor.org/info/rfc6303>>.
- [RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", [BCP 153](#), [RFC 6598](#), DOI 10.17487/RFC6598, April 2012, <<http://www.rfc-editor.org/info/rfc6598>>.
- [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>>.

- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", [RFC 6763](#), DOI 10.17487/RFC6763, February 2013, <<http://www.rfc-editor.org/info/rfc6763>>.
- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", [RFC 7050](#), DOI 10.17487/RFC7050, November 2013, <<http://www.rfc-editor.org/info/rfc7050>>.
- [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.chapin-additional-reserved-tlds]
Chapin, L. and M. McFadden, "Additional Reserved Top Level Domains", [draft-chapin-additional-reserved-tlds-02](#) (work in progress), March 2015.
- [I-D.lewis-domain-names]
Lewis, E., "Domain Names", [draft-lewis-domain-names-02](#) (work in progress), January 2016.
- [SDO-CABF-INT]
CA/Browser Forum, "Guidance on the Deprecation of Internal Server Names and Reserved IP Addresses", June 2012, <<https://www.digicert.com/internal-names.htm>>.
- [SDO-ICANN-COLL]
Interisle Consulting Group, LLC, "Name Collisions in the DNS", August 2013, <<https://www.icann.org/en/system/files/files/name-collision-02aug13-en.pdf>>.
- [SDO-IANA-SUDR]
Internet Assigned Numbers Authority, "Special-Use Domain Names registry", October 2015, <<http://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml>>.
- [CORP-SUN-NIS]
Sun Microsystems, "System and Network Administration", March 1990.

Authors' Addresses

Ted Lemon
Nominum, Inc.
800 Bridge Parkway
Redwood City, California 94065
United States of America

Phone: +1 650 381 6000
Email: ted.lemon@nominum.com

Ralph Droms
Cisco, Inc.

Email: rdroms.ietf@gmail.com

