

IETF
Internet-Draft
Intended status: Informational
Expires: April 30, 2015

A. Sullivan
Dyn
October 27, 2014

On Interoperation of Labels Between mDNS and DNS
draft-sullivan-dnssd-mdns-dns-interop-01

Abstract

Despite its name, DNS-Based Service Discovery can use naming systems other than the Domain Name System when looking for services. Different name systems use different conventions for the characters allowed in any name. In order for DNS-SD to be used effectively in environments where multiple different name systems are in use, it is important to attend to differences in the underlying technology. This memo presents an outline of the requirements for selection of labels for mDNS and DNS when they are expected to interoperate in this manner.

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 April 30, 2015.

Copyright Notice

Copyright (c) 2014 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

Internet-Draft

mDNS and DNS Interop

October 2014

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.	Conventions and terms used in this document	3
2.	Requirements for a profile for label interoperation	3
3.	DNS-SD portions	4
3.1.	The <Instance> Portion of the Service Instance Name .	4
3.2.	The <Service> Portion of the Service Instance Name .	5
3.3.	The <Domain> Portion of the Service Instance Name . .	5
4.	Acknowledgements	6
5.	IANA Considerations	6
6.	Security Considerations	6
7.	Informative References	6
	Author's Address	7

[1.](#) Introduction

DNS-Based Service Discovery (DNS-SD, [\[RFC6763\]](#)) specifies a mechanism for discovering services using queries both to the Domain Name System (DNS, [\[RFC1034\]](#), [\[RFC1035\]](#)) and to Multicast DNS (mDNS, [\[RFC6762\]](#)). Conventional use of the DNS generally follows the host name rules [\[RFC0952\]](#) for labels -- the so-called LDH rule. That convention is the reason behind the development of Internationalized Domain Names for Applications (IDNA2008, [\[RFC5890\]](#), [\[RFC5891\]](#), [\[RFC5892\]](#), [\[RFC5893\]](#), [\[RFC5894\]](#), [\[RFC5895\]](#)). It is worth noting that the LDH rule is a convention, and not a strict rule of the DNS. It is assumed to be true widely enough, however, that in many circumstances names cannot be used unless they cleave to the LDH rule.

At the same time, mDNS requires that labels be encoded in UTF-8, and permits a range of characters in labels that are not permitted by IDNA2008 or the LDH rule. For example, mDNS encourages the use of spaces and punctuation in mDNS names (see [\[RFC6763\]](#), [section 4.1.3](#)). It does not restrict which Unicode code points may be used in those labels, so long as the code points are UTF-8 in Net-Unicode [\[RFC5198\]](#) format.

Users of applications are, of course, frequently unconcerned with

(not to say oblivious to) the name-resolution system(s) in service at any given moment, and are inclined simply to use the same names in different contexts. As a result, the same string might be tried as a name using different name resolution technologies. If DNS-SD is to be used in an environment where both mDNS and DNS are to be queried

for services, then some parts of the names to be queried will need to be compatible with the rules and conventions for both DNS and mDNS.

One approach to interoperability under these circumstances is to use a single operational convention for names under the different naming systems. This memo assumes such a use profile, and outlines what is necessary to make it work.

It is worth noting that users of DNS-SD do not use the service discovery names in the same way that users of other domain names might. Most domain names might as easily be typed in as direct user input as any other method. But the service discovery context generally assumes users are picking a service from a list. As a result, the sorts of application considerations that are appropriate to the general-purpose DNS name, and that resulted in the A-label/ U-label (see below) in IDNA2008, are not the right approach for DNS-SD.

1.1. Conventions and terms used in this document

Wherever appropriate, this memo uses the terminology defined in [Section 2 of \[RFC5890\]](#). In particular, the reader is assumed to be familiar with the terms "U-label", "LDH label", and "A-label" from that document. Similarly, the reader is assumed to be familiar with the U+NNNN notation for Unicode code points used in [\[RFC5890\]](#) and other documents dealing with Unicode code points. In the interests of brevity and consistency, the definitions are not repeated here.

This memo refers to names in the DNS as though the LDH rule and IDNA2008 are strict requirements. They are not. DNS labels are, in principle, just collections of octets, and therefore in principle the LDH rule is not a constraint. In practice, applications often intercept labels that do not conform to the LDH rule and apply IDNA and other transformations.

The term "owner name" (common to the DNS vernacular) is used here to

apply not just to the names to be looked up in the DNS, but to any name that might be looked up either in the DNS or using mDNS.

2. Requirements for a profile for label interoperation

Any interoperability between mDNS and DNS will require interoperability across some of the portions of a DNS-SD Service Instance Name (see [Section 3](#)) that are implicated in regular mDNS and DNS lookups. Only some portions are implicated. In any case, if a given portion is implicated, the profile will need to apply to all labels in that portion.

In addition, because DNS-SD Service Instance Names can be used in a domain name slot, care must be taken by DNS-SD resolvers to undertake the special processing outlined here, so that DNS-SD portions that do not use IDNA2008 will not be treated as U-labels and will not undergo IDNA processing.

Because the profile will need to apply to names that might need to interoperate with names in the DNS, and because mDNS permits labels that IDNA does not, the profile might reduce the labels that could be used with mDNS. Consequently, some recommendations from [\[RFC6763\]](#) will not really be possible to implement using names subject to the profile. In particular, [\[RFC6763\], section 4.1.3](#) recommends that labels always be stored and communicated as UTF-8, even in the DNS. Because IDNA2008 libraries will treat any Unicode-encoded labels as candidate U-labels and attempt to perform resolution in A-label form, the advice to store and transmit labels as UTF-8 in the DNS is likely to encounter problems. In particular, the <Domain> part of a Service Instance Name is unlikely to be found in its UTF-8 form in the public DNS tree for zones that are using IDNA2008. By contrast, mDNS normally uses UTF-8.

U-labels cannot contain upper case letters. That restriction extends to ASCII-range upper case letters that work fine in LDH-labels. It may be confusing that the character "A" works in the DNS when none of the characters in the label has a diacritic, but does not work when there is such a diacritic in the label. Labels in mDNS names may contain upper case characters, so the profile will need either to restrict the use of upper case or come up with a reliable and predictable (to users) convention for case folding even in the

presence of diacritics.

3. DNS-SD portions

DNS-SD specifies three portions of the owner name for a DNS-SD resource record. These are the <Instance> portion, the <Service> portion, and the <Domain>. The owner name made of these three parts is called the Service Instance Name. It is worth observing that a portion may be more than one label long. See [\[RFC6763\], section 4.1.](#)

3.1. The <Instance> Portion of the Service Instance Name

[RFC6763] is clear that the <Instance> portion of the Service Instance Name is intended for presentation to users, and therefore virtually any character is permitted in it. There are two ways that a profile might address this portion.

The first way would be to treat this portion as likely to be intercepted by system-wide IDNA-aware resolvers. In this case, the

portion needs to be made subject to the profile, thereby curtailing what characters may appear in this portion. This approach permits DNS-SD to use any standard system resolver but presents inconsistencies with the DNS-SD specification and with DNS-SD that is exclusively mDNS-based. Therefore, this strategy is rejected.

Instead, DNS-SD implementations can intercept the <Instance> portion of a Service Instance Name and ensure that those labels are never handed to IDNA-aware resolvers that might attempt to convert these labels into A-labels. Under this approach, the DNS-SD <Instance> portion works as it always does, but at the cost of using special resolution code built into the DNS-SD system.

3.2. The <Service> Portion of the Service Instance Name

DNS-SD includes a <Service> component in the Service Instance Name. This component is not really user-facing data, but is instead control data embedded in the Service Instance Name. This component includes so-called "underscore labels", which are labels prepended with U+005F (_). The underscore label convention was established by DNS SRV ([\[RFC2782\]](#)) for identifying metadata inside DNS names. A system-wide resolver (or DNS middlebox) that cannot handle underscore labels will

not work with DNS-SD at all, so it is safe to suppose that such resolvers will not attempt to do special processing on these labels. Therefore, the <Service> portion of the Service Instance Name will not be subject to the profile.

[3.3.](#) The <Domain> Portion of the Service Instance Name

The <Domain> portion of the Service Instance Name forms an integral part of the QNAME submitted for DNS resolution, and a system-wide resolver that is IDNA2008-aware is likely to interpret labels with UTF-8 in the QNAME as candidates for IDNA2008 processing. Operators of Internationalized Domain Names will almost certainly publish them in the DNS as A-labels. Therefore, these labels will need to be subject to the profile. DNS-SD implementations ought to identify the <Domain> portion of the Service Instance Name and treat it subject to IDNA2008 in case the domain is to be queried from the global DNS. This is different to the rule for resolution published in [[RFC6763](#)].

One might argue against this restriction on either of two grounds:

1. It is possible the names may be in the DNS in UTF-8, and [RFC 6763](#) already specifies a fallback strategy of progressively attempting first the U-label lookup and then the A-label lookup.

2. Zone administrators that wish to support DNS-SD can publish a UTF-8 version of the zone along side the A-label version of the zone.

The first of these is rejected because it represents a potentially significant increase in DNS lookup traffic for no value. It is possible for a DNS-SD application to identify the <Domain> portion of the Service Instance Name. The standard way to publish IDNs on the Internet uses IDNA. Therefore, additional lookups should not be encouraged. When [[RFC6763](#)], the bulk of IDNs were lower in the tree, but now that there are internationalized labels in the root zone, it seems reasonable to use only the single lookup strategy.

The second reason depends on the idea that it is possible to maintain two names in sync with one another. This is not strictly speaking

true, although in this case the domain operator could simply create a DNAME record [[RFC6672](#)] from the UTF-8 name to the IDNA2008 zone. This still, however, relies on being able to reach the (UTF-8) name in question, and it is unlikely that the UTF-8 version of the zone will be delegated from anywhere. Moreover, in many organizations the support for DNS-SD and the support for domain name delegations are not performed by the same department, and depending on a coordination between the two will make the system more fragile or slower or both.

[4.](#) Acknowledgements

The author gratefully acknowledges the insights of Stuart Cheshire and Kerry Lynn.

[5.](#) IANA Considerations

This memo makes no requests of IANA.

[6.](#) Security Considerations

This memo presents some requirements for future development, but does not specify anything. Therefore, it has no implications for security.

[7.](#) Informative References

- [RFC0952] Harrenstien, K., Stahl, M., and E. Feinler, "DoD Internet host table specification", [RFC 952](#), October 1985.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), February 2000.
- [RFC5198] Klensin, J. and M. Padlipsky, "Unicode Format for Network

Interchange", [RFC 5198](#), March 2008.

- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", [RFC 5890](#), August 2010.
- [RFC5891] Klensin, J., "Internationalized Domain Names in Applications (IDNA): Protocol", [RFC 5891](#), August 2010.
- [RFC5892] Faltstrom, P., "The Unicode Code Points and Internationalized Domain Names for Applications (IDNA)", [RFC 5892](#), August 2010.
- [RFC5893] Alvestrand, H. and C. Karp, "Right-to-Left Scripts for Internationalized Domain Names for Applications (IDNA)", [RFC 5893](#), August 2010.
- [RFC5894] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Background, Explanation, and Rationale", [RFC 5894](#), August 2010.
- [RFC5895] Resnick, P. and P. Hoffman, "Mapping Characters for Internationalized Domain Names in Applications (IDNA) 2008", [RFC 5895](#), September 2010.
- [RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", [RFC 6672](#), June 2012.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", [RFC 6762](#), February 2013.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", [RFC 6763](#), February 2013.

Author's Address

Sullivan

Expires April 30, 2015

[Page 7]

Internet-Draft

mDNS and DNS Interop

October 2014

Andrew Sullivan

Dyn
150 Dow St.
Manchester, NH 03101
U.S.A.

Email: asullivan@dyn.com