        **Using Labels With DNS-Based Service Discovery, mDNS, and DNS**
                 **draft-sullivan-dnssd-label-miprofile-00**

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 follow a common set of conventions for naming.  This
   memo presents a convention for maximizing such interoperability.

Table of Contents

## 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 the names to be queried will need to be compatible
   with the rules and conventions for both DNS and mDNS.  This memo
   provides advice on how to do that.  For the sake of brevity, in what
   follows the use of labels that work reliably with both mDNS and DNS
   is called the "maximally inter-operative profile", or "MI profile".

   It is important to emphasize that this profile is maximally
   interoperable in the sense that it encourages the most
   interoperability between DNS and mDNS environments; but it does not
   guarantee it.  IDNA2008 does not constrain DNS operators from putting
   any labels they want (including those from outside the IDNA2008-
   permissible repertoire) in zones.  Rather, this profile is intended
   to reduce the scope for variability between systems so that a minimal
   (but predictable) subset of possible behaviour is available to
   everyone.

## 1.1.  Conventions and terms used in this document

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this

document are to be interpreted as described in RFC 2119 [RFC2119].

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.

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.  The MI Profile

In the following, we make recommendations for how to use DNS-SD where that use needs to work seamlessly across DNS and mDNS.  The recommendations involve limitations on what labels should (and should not) be used for names used with DNS-SD.  The MI profile applies to labels, not names.

The MI profile has three rules:

1.  If the label is made entirely of LDH code points, then the label MUST be an LDH label.

2.  All LDH code points MUST be folded to lower case.

3.  If the label contains any other code point, then the label MUST be a well-formed U-label.

[[anchor3: Rule 1 is a tautology.  I've wondered whether it's needed, but it makes rule 2 clearer. --ajs@anvilwalrusden.com]]

### 2.1.  The MI Profile and DNS-SD

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.

To be effective, the MI profile is either applied to every label in a Service Instance Name portion, or it is not applied to that portion at all.  The reason the profile might not be applied to a portion is because different portions have different functions within the

Service Instance Name: some of them function as control data, and therefore have special handling applied.  Those portions are not intended for user display.

Because the MI profile is to apply to names that might need to interoperate with names in the DNS, the profile reduces the scope for labels to be used with DNS or mDNS.  Consequently, some recommendations from [RFC6763] cannot really be implemented using names subject to the MI profile.  In particular, [RFC6763], section 4.1.3 recommends that rich text, human-readable labels be used, and includes punctuation and space characters in the examples.  Such uses are incompatible with the MI profile, because spaces and most punctuation are permitted neither in U-labels nor in LDH labels.  In addition, the same section 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 and is NOT RECOMMENDED.  Naturally, because mDNS always uses UTF-8, mDNS labels SHOULD be transmitted as UTF-8 unless there is strong reason to suppose that some mDNS responder is using A-labels. The subset of allowable characters under the MI profile remains the same, however, so some characters that would be available in mDNS without the MI profile are not available when the MI profile is in use.

The reason for rule 2 is merely to reduce potential user confusion. U-labels cannot contain upper case letters.  That restriction extends to ASCII-range upper case letters that work fine in LDH-labels.  It is 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.  Therefore, MI profile requires folding to lower case even traditional DNS labels, in the interests of maximizing interoperability.

### 2.1.1.  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.  Because the <Instance> portion may actually be part of the QNAME submitted for DNS resolution, and because such names are subject to being intercepted by a system-wide resolver that is IDNA2008-aware, use of the MI profile on the <Instance> portion of the Service Instance Name is RECOMMENDED.  This will probably reduce some of the utility of the <Instance> portion, but it provides the benefit that the entire name can be looked up and used with DNS-SD when using the DNS. [[anchor4: I am torn by this recommendation.  This version is conditioned by a

mental model where a resolution system (more than a DNS resolver, but
including IDNA for instance) looks at a label.  If the label has
Unicode characters in it, then the resolver attempts an IDNA2008
transformation on the label; otherwise, it attempts to use the label
in stock DNS operation.  It's possible, however, that some systems
pick out things like underscore labels first, and thereby identify
"control" labels that purport to represent particular pieces of
functionality.  In that case, the resolver could treat the whole name
differently, and pull off the Instance portion prior to the Service
portion.  If it could do that, it could use straight UTF-8, spaces,
punctuation, and everything else.  I'm sceptical of the reliability
of this, though, so it seems to me it'd be better to apply the
profile to anything that wasn't a control label.
--ajs@anvilwalrusden.com]]

## 2.1.2.  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.
Note that underscore labels do not meet the requirements of the MI
profile, so the MI profile MUST NOT be applied to the <Service>
portion of the Service Instance Name.

## 2.1.3.  The MI Profile and 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.  Therefore,
use of the MI profile on such names is RECOMMENDED, unless there is
strong evidence that no resolvers in the resolution chain will
attempt to perform a U-label to A-label transformation during lookup,
and that the actual DNS server will have U-labels rather than
A-labels stored.  In practice, these restrictions will permit plain
UTF-8 lookups in special conditions (e.g. on a local network with a
DNS server and careful administration) only.

## 3.  Acknowledgements

The author gratefully acknowledges the insights of Kerry Lynn.

## 4.  IANA Considerations

This memo makes no requests of IANA.

## 5.  Security Considerations

This memo recommends a subset of available characters for use in DNS-
SD-related queries, consistent with the rules of mDNS and IDNA2008.
The security considerations of those protocols apply broadly to this
memo, but this memo introduces no additional security considerations
on its own.

## 6.  References

### 6.1.  Normative References

[RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
            Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC6763]   Cheshire, S. and M. Krochmal, "DNS-Based Service
            Discovery", RFC 6763, February 2013.

### 6.2.  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.

     [RFC6762]   Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762,
                   February 2013.


Author's Address

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

     Email: asullivan@dyn.com