

DNS-SD/mDNS Extensions
Internet-Draft
Intended status: Informational
Expires: April 25, 2014

K. Lynn, Ed.
Consultant
S. Cheshire
Apple, Inc.
October 22, 2013

Requirements for Scalable DNS-SD/mDNS Extensions
draft-lynn-dnssd-requirements-00

Abstract

DNS-SD/mDNS is widely used today for discovery and resolution of services and names on a local link, but there are use cases to extend DNS-SD/mDNS to enable service discovery beyond the local link. This document provides a problem statement and a list of requirements.

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 25, 2014.

Copyright Notice

Copyright (c) 2013 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. Problem Statement	3
3. Basic Use Cases	5
4. Internationalization Considerations	6
5. Namespace Considerations	6
6. Requirements	6
7. IANA Considerations	7
8. Security Considerations	7
9. Acknowledgments	8
10. References	9
Authors' Addresses	10

1. Introduction

DNS-Based Service Discovery [DNS-SD] in combination with its companion technology Multicast DNS [mDNS] is widely used today for discovery and resolution of services and names on a local link. However, as users move to multi-link home or campus networks they find that mDNS does not work across routers. DNS-SD can also be used in conjunction with conventional unicast DNS to enable wide-area service discovery, but this capability is not yet widely deployed. This disconnect between customer needs and current practice has led to calls for improvement, such as the Educause petition [EP].

In response to this and similar evidence of market demand, several products now enable service discovery beyond the local link using different ad-hoc techniques. However, it is unclear which approach represents the best long-term direction for DNS-based service discovery protocol development.

DNS-SD/mDNS in its present form is also not optimized for network technologies where multicast transmissions are relatively expensive. Wireless networks such as [IEEE.802.11] may be adversely affected by excessive mDNS traffic due to the higher network overhead of multicast transmissions. Wireless mesh networks such as 6LoWPAN [RFC4944] are effectively multi-link subnets where multicasts must be forwarded by intermediate nodes.

It is in the best interests of end users, network administrators, and vendors for all interested parties to cooperate within the context of the IETF to develop an efficient, scalable, and interoperable standards-based solution.

This document defines the problem statement and gathers requirements for Scalable DNS-SD/mDNS Extensions.

1.1. Requirements Language

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 "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

1.2. Terminology [TBD]

Discovery Scope

Zero Configuration

Incremental Deployment

2. Problem Statement

Service discovery beyond the local link is perhaps the most important feature currently missing from the DNS-SD/mDNS framework. The issues and requirements are summarized below.

2.1. Multilink Naming and Discovery

A list of desired DNS-SD/mDNS improvements from network administrators in the research and education community was issued in the form of the Educause petition [EP]. The following is a technical summary of the issues:

- o Products that advertise services such as printing and multimedia streaming via DNS-SD/mDNS are not currently discoverable by devices on other links. It is common practice for enterprises and institutions to use wireless links for client access and wired networks for server infrastructure, typically on different subnets. DNS-SD used with conventional unicast DNS does work when devices are on different links, but the resource records that describe the service must somehow be entered into the unicast DNS namespace.
- o Entering DNS-SD records manually into a unicast DNS zone file works, (as has been done for many years for the Terminal Room printers at IETF meetings) but requires the DNS administrator to know how to do that [static] and is fragile when IP addresses of devices may change, as is common when DHCP is used.

- o Automatically adding DNS-SD records using DNS Update works, but requires that the DNS server be configured to allow DNS Updates, and requires that devices be configured with the DNS Update credentials to permit such updates, which has proven to be onerous.
- o Therefore, a mechanism is desired that populates the DNS namespace with the appropriate DNS-SD records with less manual administration than typically needed for a unicast DNS server.

The following is a technical summary of the requirements:

- o It must scale to a range of hundreds or thousands of DNS-SD/mDNS enabled devices in a given environment.
- o It must work with wired and wireless networks from different vendors.
- o It must not significantly increase network traffic (wired or wireless).
- o It must be easily managed at an enterprise scale.
- o It must be provided at a reasonable cost. [CapEx + OpEx. KEL]

2.2. IEEE 802.11 Wireless LANs

Multicast DNS was originally designed to run on Ethernet - the dominant link-layer at the time. In shared Ethernet networks, multicast frames place little additional demand on the shared network medium above unicast frames. In IEEE 802.11 networks however, multicast frames are transmitted at a low data rate supported by all receivers. In practice, this data rate leads to a larger fraction of airtime being devoted to multicast transmission. Some network administrators block multicast traffic or convert it to a series of link-layer unicast frames.

Wired links may be orders of magnitude less reliable than their wired counterparts. To improve transmission reliability, the IEEE 802.11 MAC requires positive acknowledgement of unicast frames. It does not, however, support positive acknowledgement of multicast frames. As a result, it is common to observe much higher loss of multicast frames on wireless as compared to wired network technologies.

Enabling service discovery on IEEE 802.11 networks requires that the number of multicast frames be restricted to a suitably low value, or replaced with unicast frames to use the MAC's reliability features.

2.3. Low Power and Lossy Networks (LLNs)

Emerging wireless mesh networking technologies such as RPL [RFC6550] and 6LoWPAN present several challenges for the current DNS-SD/mDNS design. First, Link-Local multicast scope [RFC4291] is defined as a single-hop neighborhood. A single subnet prefix in a wireless mesh network may often span multiple links, therefore a larger multicast scope is required to span it [I-D.ietf-6man-multicast-scopes].

Additionally, low-power nodes may be offline for significant periods either because they are "sleeping" or due to connectivity problems. In such cases LLN nodes might fail to respond to queries or defend their names using the current design.

3. Basic Use Cases

The following use cases are defined with different constraints to help distinguish and classify the target requirements.

(A) Personal Area networks; e.g., one laptop and one printer. This is the simplest example of a DNS-SD/mDNS network.

(B) Classic home networks, consisting of:

- * Single exit router: the network may have multiple upstream providers or networks, but all outgoing and incoming traffic goes through a single router.
- * One level depth: all links on the network are connected to the same default router.
- * Single administrative domain: all nodes under the same admin entity.

(C) Advanced residential and small business networks [I-D.ietf-homenet-arch]:

Like B but consist of two or more wired and/or wireless links, connected by routers, behind the single exit router. However, the forwarding nodes are largely self-configuring and do not require routing protocol administration.

(D) Enterprise networks:

Like C but consist of arbitrary diameter under a single administrative domain. A large majority of the forwarding and security devices are configured.

(E) Higher Education networks:

Like D but core network may be under a central administrative domain while leaf networks are under local administrative domains.

(F) Mesh networks such as RPL/6LoWPAN:

Multi-link subnets with prefixes defined by one or more border routers. May comprise network B and any part of networks C, D, or E.

4. Internationalization Considerations

The solution should support rich international text, as do DNS-SD and mDNS today. Users will not accept a solution that does not allow the richness of service naming that they currently have with mDNS, manual zone files, and DNS Update today.

5. Namespace Considerations

The unicast DNS namespace contains globally unique names. Naming services over a local scope contain locally unique names. Clients discovering services need to be able to differentiate global names from local names.

6. Requirements

[This is a strawman proposal. MB]

REQ1: The scope of the discovery should be either automatically found by the discovering devices and/or configured.

REQ2: For use cases A, B, and C, there should be a zero configuration mode of operation.

REQ3: For use cases D and E, there should be a way to configure the scope of the discovery and also support both smaller (ex: department) and larger (ex: campus-wide) discovery scopes.

REQ4: For use cases D and E, there should be an incremental way to deploy the solution.

REQ5: The new solution should integrate or at least should not break any current link scope DNS-SD/mDNS protocols and deployments.

REQ6: The new solution MUST be capable of spanning multiple links (hops) and network technologies.

REQ7: The new solution MUST be scalable to thousands of servers with minimal configuration and without degrading network performance.

REQ8: The new solution MUST provide a consistent user experience whether local or global services are being discovered.

7. IANA Considerations

This document currently makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

8. Security Considerations

[Not complete - initial ideas. MB/KEL]

If the scope of the discovery is not properly setup or constrained, then information leaks will happen outside the appropriate network.

Visiting nodes on a network may discover more services than desired by the network policies, if filtering of discovery packets was not properly setup. [Is this a NAC or DNS problem? KL]

Depending on the chosen solution, there is a possibility of name space conflicts between the DNS tree and this solution. In this case, a node may not know if the target node or service is the right one, therefore enabling ground for various attacks.

The DNS-SD/mDNS framework security considerations also apply.

DNSSEC can assert the validity but not the veracity of records in a zone file. The trust model of the global DNS relies on the fact that human administrators either a) manually enter resource records into a zone file, or b) configure the DNS server to authenticate a trusted device (e.g., a DHCP server) that can automatically maintain such records.

By contrast, the "plug-and-play" nature of mDNS devices has up to now depended only on physical connectivity. If a device is visible via mDNS then it is assumed to be trusted. This is no longer likely to be the case in larger networks. Still, the new solution SHOULD leverage existing security solutions and not invent new ones.

Mobile devices such as smart phones that can expose the location of their owners by registering services in arbitrary zones pose a risk to privacy. Such devices MUST NOT register their services in arbitrary zones without the approval of their operators. However, it

SHOULD be possible to configure one or more "home" zones, e.g., based on subnet prefix, in which mobile devices may automatically register their services.

9. Acknowledgments

We gratefully acknowledge contributions and review comments made by RJ Atkinson, Marc Blanchet, Tim Chown, Ralph Droms, Educause, David Farmer, Matthew Gast, Peter Van Der Stok, and Thomas Narten.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007.
- [RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.
- [mDNS] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013.
- [DNS-SD] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, February 2013.

10.2. Informative References

- [I-D.ietf-6man-multicast-scopes] Droms, R., "IPv6 Multicast Address Scopes", draft-ietf-6man-multicast-scopes-00 (work in progress), August 2013.
- [I-D.ietf-homenet-arch] Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "Home Networking Architecture for IPv6", draft-ietf-homenet-arch-10 (work in progress), August 2013.
- [EP] "Educause Petition", <https://www.change.org/petitions/from-educause-higher-ed-wireless-networking-admin-group>, July 2012.
- [IEEE.802.11] "Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications ", IEEE Std 802.11-2012, 2012, <<http://standards.ieee.org/getieee802/download/802.11-2012.pdf>>.

[static] "Manually Adding DNS-SD Service Discovery Records to an Existing Name Server", July 2013, <<http://www.dns-sd.org/ServerStaticSetup.html>>.

Authors' Addresses

Kerry Lynn (editor)
Consultant

Phone: +1 978 460 4253
Email: kerlyn@ieee.org

Stuart Cheshire
Apple, Inc.
1 Infinite Loop
Cupertino , California 95014
USA

Phone: +1 408 974 3207
Email: cheshire@apple.com

IETF
Internet-Draft
Intended status: Informational
Expires: April 21, 2014

A. Sullivan
Dyn
October 18, 2013

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.

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 21, 2014.

Copyright Notice

Copyright (c) 2013 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	3
1.1. Conventions and terms used in this document	3
2. The MI Profile	4
2.1. The MI Profile and DNS-SD	4
2.1.1. The <Instance> Portion of the Service Instance Name . .	5
2.1.2. The <Service> Portion of the Service Instance Name . .	6
2.1.3. The MI Profile and the <Domain> Portion of the Service Instance Name	6
3. Acknowledgements	7
4. IANA Considerations	7
5. Security Considerations	7
6. References	7
6.1. Normative References	7
6.2. Informative References	7
Author's Address	8

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.

2.1.1.1. 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

