

Internet Engineering Task Force (IETF)  
Request for Comments: 7558  
Category: Informational  
ISSN: 2070-1721

K. Lynn  
Verizon Labs  
S. Cheshire  
Apple, Inc.  
M. Blanchet  
Viagenie  
D. Migault  
Ericsson  
July 2015

Requirements for Scalable DNS-Based Service  
Discovery (DNS-SD) / Multicast DNS (mDNS) Extensions

Abstract

DNS-based Service Discovery (DNS-SD) over Multicast DNS (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 for scalable DNS-SD.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7558>.

Copyright Notice

Copyright (c) 2015 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 . . . . . 4
- 3. Basic Use Cases . . . . . 6
- 4. Requirements . . . . . 7
- 5. Namespace Considerations . . . . . 9
- 6. Security Considerations . . . . . 9
- 7. References . . . . . 11
- Acknowledgements . . . . . 13
- Authors' Addresses . . . . . 14

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. As users move to multi-link home or campus networks, however, they find that mDNS (by design) 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. As of yet, no consensus has emerged regarding which approach represents the best long-term direction for DNS-based Service Discovery protocol development.

Multicast DNS in its present form is also not optimized for network technologies where multicast transmissions are relatively expensive. Wireless networks such as Wi-Fi [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 IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) [RFC4944] are effectively multi-link subnets [RFC4903] 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. Terminology and Acronyms

**Service:** A listening endpoint (host and port) for a given application protocol. Services are identified by Service Instance Names.

**DNS-SD:** DNS-based Service Discovery [DNS-SD] is a conventional application of DNS resource records and messages to facilitate the naming, discovery, and location of services. When used alone, the term generally refers to the wide-area unicast protocol.

**mDNS:** Multicast DNS [mDNS] is a mechanism that facilitates distributed DNS-like capabilities (including DNS-SD) on a local link without need of traditional DNS infrastructure.

**SSD:** Scalable Service Discovery (or Scalable DNS-SD) is a future extension of DNS-SD (and perhaps mDNS) that meets the requirements set forth in this document.

**Scope of Discovery:** A subset of a local or global namespace, e.g., a DNS subdomain, that is the target of a given SSD query.

**Zero Configuration:** A deployment of SSD that requires no administration (although some administration may be optional).

**Incremental Deployment:** An orderly transition, as a network installation evolves, from DNS-SD/mDNS to SSD.

## 2. Problem Statement

Service discovery beyond the local link is perhaps the most important feature currently missing from the DNS-SD-over-mDNS framework (also written as "DNS-SD over mDNS" or "DNS-SD/mDNS"). Other issues and requirements are summarized below.

### 2.1. Multi-link 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 summary of their technical issues:

- o It is common practice for enterprises and institutions to use wireless links for client access and wired links for server infrastructure; typically, they are on different subnets. Products that advertise services such as printing and multimedia streaming via DNS-SD over mDNS are not currently discoverable by client devices on other links. DNS-SD used with conventional unicast DNS does work when servers and clients are on different links, but the resource records that describe the services must somehow be entered into the unicast DNS namespace.
- o DNS-SD resource records may be entered manually into a unicast DNS zone file [STATIC], but this task must be performed by a DNS administrator. It is labor intensive and brittle when IP addresses of devices change dynamically, as is common when DHCP is used.
- o Automatically adding DNS-SD records using DNS Update works, but it requires that the DNS server be configured to allow DNS Updates and that devices be configured with the DNS Update credentials to permit such updates, which has proven to be onerous.

Therefore, a mechanism is desired that populates the DNS namespace with the appropriate DNS-SD records with less manual administration than is typically needed for a conventional unicast DNS server.

The following is a summary of technical requirements:

- o It must scale to a range of hundreds to thousands of DNS-SD/mDNS-enabled devices in a given environment.
- o It must simultaneously operate over a variety of network link technologies, such as wired and wireless networks.

- o It must not significantly increase network traffic (wired or wireless).
- o It must be cost-effective to manage at up to enterprise scale.

## 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-medium Ethernet networks, multicast frames place little additional demand on the shared network medium compared to 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 use access points that transmit multicast frames using a series of link-layer unicast frames.

Wireless links may be orders of magnitude less reliable than their wired counterparts. To improve transmission reliability, the IEEE 802.11 Medium Access Control (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 higher loss rates of multicast frames on wireless network technologies than on 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 the Routing Protocol for LLNs (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 wireless mesh network representing a single logical subnet may often extend to multiple hops [RFC4903]; therefore, a larger multicast scope is required to span it [RFC7346]. Multicast DNS was intentionally not specified for greater than link-local scope because of the additional off-link multicast traffic that it would generate.

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 characteristics to help motivate, distinguish, and classify the target requirements. They cover a spectrum of increasing deployment and administrative complexity.

(A) Personal Area Networks (PANs): The simplest example of a network may consist of a single client and server, e.g., one laptop and one printer, on a common link. PANs that do not contain a router may use Zero Configuration Networking [ZC] to self-assign link-local addresses [RFC3927] [RFC4862] and Multicast DNS [mDNS] to provide naming and service discovery, as is currently implemented and deployed in Mac OS X, iOS, Windows [B4W], and Android [NSD].

(B) Classic home or 'hotspot' networks, with the following properties:

- \* Single exit router: The network may have one or more upstream providers or networks, but all outgoing and incoming traffic goes through a single router.
- \* One-level depth: A single physical link, or multiple physical links bridged to form a single logical link, that is connected to the default router. The single logical link provides a single broadcast domain, facilitating use of link-local Multicast DNS, and also ARP, which enables the home or 'hotspot' network to consist of just a single IPv4 subnet.
- \* Single administrative domain: All nodes under the same administrative authority. Note that this does not necessarily imply a network administrator.

(C) Advanced home and small business networks [RFC7368]:

Like B, but consists of multiple wired and/or wireless links, connected by routers, generally behind a single exit router. However, the forwarding nodes are largely self-configuring and do not require routing protocol administration. Such networks should also not require DNS administration.

(D) Enterprise networks:

Consists of arbitrary network diameter under a single administrative authority. A large majority of the forwarding and security devices are configured, and multiple exit routers are

more common. Large-scale conference-style networks, which are predominantly wireless access, e.g., as available at IETF meetings, also fall within this category.

(E) Higher-Education networks:

Like D, but the core network may be under a central administrative authority while leaf networks are under local administrative authorities.

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

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

#### 4. Requirements

Any successful SSD solution(s) will have to strike the proper balance between competing goals such as scalability, deployability, and usability. With that in mind, none of the requirements listed below should be considered in isolation.

REQ1: For use cases A, B, and C, there should be a Zero Configuration mode of operation. This implies that servers and clients should be able to automatically determine a default scope of discovery in which to advertise and discover services, respectively.

REQ2: For use cases C, D, and E, there should be a way to configure scopes of discovery that support a range of topologically independent zones (e.g., from department to campus wide). This capability must exist in the protocol; individual operators are not required to use this capability in all cases -- in particular, use case C should support Zero Configuration operation where that is desired. If multiple scopes are available, there must be a way to enumerate the choices from which a selection can be made. In use case C, either Zero Configuration (one flat list of resources) or configured (e.g., resources sorted by room) modes of operation should be available.

REQ3: As stated in REQ2 above, the discovery scope need not be aligned to network topology. For example, it may instead be aligned to physical proximity (e.g., building) or organizational structure (e.g., "Sales" vs. "Engineering").

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

- REQ5: SSD should leverage and build upon current link scope DNS-SD/mDNS protocols and deployments.
- REQ6: SSD must not adversely affect or break any other current protocols or deployments.
- REQ7: SSD must be capable of operating across networks that are not limited to a single link or network technology, including clients and services on non-adjacent links.
- REQ8: It is desirable that a user or device be able to discover services within the sites or networks to which the user or device is connected.
- REQ9: SSD should operate efficiently on common link layers and link types.
- REQ10: SSD should be considerate of networks where power consumption is a critical factor; for example, nodes may be in a low-power or sleeping state.
- REQ11: SSD must be scalable to thousands of nodes with minimal configuration and without degrading network performance. A possible figure of merit is that, as the number of services increases, the amount of traffic due to SSD on a given link remains relatively constant.
- REQ12: SSD should enable a way to provide a consistent user experience whether local or remote services are being discovered.
- REQ13: The information presented by SSD should closely reflect the current state of discoverable services on the network. That is, new information should be available within a few seconds and stale information should not persist indefinitely. In networking, all information is necessarily somewhat out of date by the time it reaches the receiver, even if only by a few microseconds or less. Thus, timeliness is always an engineering trade-off against efficiency. The engineering decisions for SSD should appropriately balance timeliness against network efficiency.
- REQ14: SSD should operate over existing networks (as described by use cases A through F above) without requiring changes to the network at the physical, link, or internetworking layers.



REQ15: The administrator of an advertised service should be able to control whether the service is advertised beyond the local link.

## 5. Namespace Considerations

The traditional unicast DNS namespace contains, for the most part, globally unique names. Multicast DNS provides every link with its own separate link-local namespace, where names are unique only within the context of that link. Clients discovering services may need to differentiate between local and global names and may need to determine when names in different namespaces identify the same service.

Devices on different links may have the same mDNS name (perhaps due to vendor defaults) because link-local mDNS names are only guaranteed to be unique on a per-link basis. This may lead to a local label disambiguation problem when results are aggregated (e.g., for presentation).

SSD should support rich internationalized labels within Service Instance Names, as DNS-SD/mDNS does today. SSD must not negatively impact the global DNS namespace or infrastructure.

The problem of publishing local services in the global DNS namespace may be generally viewed as exporting local resource records and their associated labels into some DNS zone. The issues related to defining labels that are interoperable between local and global namespaces are discussed in a separate document [INTEROP-LABELS].

## 6. Security Considerations

Insofar as SSD may automatically gather DNS-SD resource records and publish them over a wide area, the security issues are likely to include the union of those discussed in the Multicast DNS [mDNS] and DNS-based Service Discovery [DNS-SD] specifications. The following sections highlight potential threats that are posed by deploying DNS-SD over multiple links or by automating DNS-SD administration.

### 6.1. Scope of Discovery

In some scenarios, the owner of the advertised service may not have a clear indication of the scope of its advertisement.

For example, since mDNS is currently restricted to a single link, the scope of the advertisement is limited, by design, to the shared link between client and server. If the advertisement propagates to a larger set of links than expected, this may result in unauthorized

clients (from the perspective of the owner) discovering and then potentially attempting to connect to the advertised service. It also discloses information (about the host and service) to a larger set of potential attackers.

Note that discovery of a service does not necessarily imply that the service is reachable by, or can be connected to, or can be used by, a given client. Specific access-control mechanisms are out of scope of this document.

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

## 6.2. Multiple Namespaces

There is a possibility of conflicts between the local and global DNS namespaces. Without adequate feedback, a discovering client may not know if the advertised service is the correct one, therefore enabling potential attacks.

## 6.3. Authorization

DNSSEC can assert the validity but not the accuracy 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.

An impostor may register on the local link and appear as a legitimate service. Such "rogue" services may then be automatically registered in unicast DNS-SD.

## 6.4. Authentication

Up to now, the "plug-and-play" nature of mDNS devices has relied only on physical connectivity. If a device is visible via mDNS, then it is assumed to be trusted. This is not likely to be the case in foreign networks.

If there is a risk that clients may be fooled by the deployment of rogue services, then application-layer authentication should be considered as part of any security solution. Authentication of any particular service is outside the scope of this document.

## 6.5. Access Control

Access Control refers to the ability to restrict which users are able to use a particular service that might be advertised via DNS-SD. In this case, "use" of a service is different from the ability to "discover" or "reach" a service.

While controlling access to an advertised service is outside the scope of DNS-SD, we note that access control today often is provided by existing site infrastructure (e.g., router access-control lists, firewalls) and/or by service-specific mechanisms (e.g., user authentication to the service). For example, networked printers can control access via a user ID and password. Apple's software supports such access control for USB printers shared via Mac OS X Printer Sharing, as do many networked printers themselves. So the reliance on existing service-specific security mechanisms (i.e., outside the scope of DNS-SD) does not create new security considerations.

## 6.6. Privacy Considerations

Mobile devices such as smart phones or laptops 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 ("opt-in") of their users. However, it should be possible to configure one or more "safe" zones in which mobile devices may automatically register their services.

## 7. References

### 7.1. Normative References

- [DNS-SD] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<http://www.rfc-editor.org/info/rfc6763>>.
- [mDNS] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<http://www.rfc-editor.org/info/rfc6762>>.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, DOI 10.17487/RFC3927, May 2005, <<http://www.rfc-editor.org/info/rfc3927>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.

- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<http://www.rfc-editor.org/info/rfc4903>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., 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, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [RFC7346] Droms, R., "IPv6 Multicast Address Scopes", RFC 7346, DOI 10.17487/RFC7346, August 2014, <<http://www.rfc-editor.org/info/rfc7346>>.
- [RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, DOI 10.17487/RFC7368, October 2014, <<http://www.rfc-editor.org/info/rfc7368>>.

## 7.2. Informative References

- [B4W] "Bonjour (software)", <[http://en.wikipedia.org/wiki/Bonjour\\_\(software\)](http://en.wikipedia.org/wiki/Bonjour_(software))>.
- [EP] Badman, L., "Petitioning Apple: From Educause Higher Ed Wireless Networking Admin Group", July 2012, <<https://www.change.org/p/from-educause-higher-ed-wireless-networking-admin-group>>.
- [IEEE.802.11] IEEE Computer Society, "IEEE Standard for 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, <<http://standards.ieee.org/about/get/802/802.11.html>>.

- [INTEROP-LABELS] Sullivan, A., "On Interoperation of Labels Between mDNS and DNS", Work in Progress, draft-sullivan-dnssd-mdns-dns-interop-01, October 2014.
- [NSD] Android, "NsdManager", <<http://developer.android.com/reference/android/net/nsd/NsdManager.html>>.
- [STATIC] "Manually Adding DNS-SD Service Discovery Records to an Existing Name Server", July 2013, <<http://www.dns-sd.org/ServerStaticSetup.html>>.
- [ZC] Cheshire, S. and D. Steinberg, "Zero Configuration Networking: The Definitive Guide", O'Reilly Media, Inc., ISBN 0-596-10100-7, December 2005.

#### Acknowledgements

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

## Authors' Addresses

Kerry Lynn  
Verizon Labs  
50 Sylvan Road  
Waltham, MA 95014  
United States

Phone: +1 781 296 9722  
Email: [kerry.lynn@verizon.com](mailto:kerry.lynn@verizon.com)

Stuart Cheshire  
Apple, Inc.  
1 Infinite Loop  
Cupertino, CA 95014  
United States

Phone: +1 408 974 3207  
Email: [cheshire@apple.com](mailto:cheshire@apple.com)

Marc Blanchet  
Viagenie  
246 Aberdeen  
Quebec, QC G1R 2E1  
Canada

Email: [Marc.Blanchet@viagenie.ca](mailto:Marc.Blanchet@viagenie.ca)  
URI: <http://viagenie.ca>

Daniel Migault  
Ericsson  
8400 Boulevard Decarie  
Montreal, QC H4P 2N2  
Canada

Phone: +1 514 452 2160  
Email: [daniel.migault@ericsson.com](mailto:daniel.migault@ericsson.com)