Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: September 6, 2012

K. Lynn
Consultant
D. Sturek
Grid2Home
March 5, 2012

Extended Multicast DNS
draft-lynn-homenet-site-mdns-00

Abstract

Multicast DNS (mDNS) provides the ability to perform DNS-like operations on the local link in the absence of any conventional unicast DNS server. Extended mDNS (xmDNS) extends the specification of mDNS to site-local scope in order to support multi-hop LANs that forward multicast packets but do not provide a unicast DNS service.

Like mDNS, xmDNS designates a portion of the DNS namespace to apply to the site-local network and specifies rules for its use.

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 September 6, 2012.

Copyright Notice

Table of Contents

1. Introduction

Multicast DNS (mDNS) provides the ability to perform DNS-like operations on the local link in the absence of any conventional unicast DNS server. Extended mDNS (xmDNS) extends the specification of mDNS to site-local scope in order to support multi-hop LANs that forward multicast packets but do not provide a unicast DNS service.

Like mDNS, xmDNS designates a portion of the DNS namespace to apply to the site-local network and specifies rules for its use.

Extended mDNS implementations MUST support all of the features of Multicast DNS [I-D.cheshire-dnsext-multicastdns] in addition to the changes specified in this document. The organization of this document is identical to mDNS, with changes specified section by section below. It is important to note that xmDNS is not intended to replace wide-area DNS-Based Service Discovery (DNS-SD) [I-D.cheshire-dnsext-dns-sd], but rather to fill a gap between the link-local scope of mDNS and the highly scalable DNS-SD. In particular, the design target anticipates multi-hop residential LANs such as ethernet to wireless mesh.

2. Conventions and Terminology 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 "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

When this document uses the term "Multicast DNS", it should be taken to mean: "Clients performing DNS-like queries for DNS-like resource records by sending DNS-like UDP query and response packets on the local link over IP Multicast to UDP port 5353."

This document uses the term "Extended Multicast DNS" to indicate the distribution of mDNS queries and responses to all links that comprise the site-local area network. Exceptions to normal mDNS operation are specified in subsequent sections.

This document uses the term "host name" in the strict sense to mean a fully-qualified domain name that has an IPv4 or IPv6 address record. It does not use the term "host name" in the commonly used but incorrect sense to mean just the first DNS label of a host's fully qualified domain name.

A DNS (or mDNS) packet contains an IP TTL in the IP header, which is effectively a hop-count limit for the packet, to guard against

routing loops. Each Resource Record also contains a TTL, which is the number of seconds for which the Resource Record may be cached. This document uses the term "IP TTL" to refer to the IP header TTL (hop limit), and the term "RR TTL" or just "TTL" to refer to the Resource Record TTL (cache lifetime).

3. Extended Multicast DNS Names

Extended Multicast DNS specifies that the DNS top-level domain ".site." is a special domain with special semantics, namely that any fully-qualified domain name ending in ".site." is site-local, and names within this domain are meaningful only on the site-local area network where they originate. This is analogous to Unique Local IPv6 Unicast Address [RFC4193] prefixes, which are site-local and meaningful only on the site where they are defined.

Any DNS query for a name ending with ".site." MUST be sent to the xmDNS multicast address (FF05::FB or its IPv4 equivalent 239.255.255.TBD). Future versions of this document may specify a method for creating zones under the ".site." top-level domain and mapping these to alternative IPv6 multicast addresses.

Note that the ".site." and ".local." domains are functionally disjoint, both from a name space and address space perspective. Hosts wishing to register or discover names in both domains must do so separately.

4. Reverse Address Mapping

[RFC4193] recommends that queries for D.F.IPV6.ARPA be handled locally. [RFC6303] extends the recommendation to cover other well known IN-ADDR.ARPA and IP6.ARPA zones for which queries should not appear on the public Internet.

In the absence of a unicast DNS server in the LAN, any DNS query for a name within the reverse mapping domain ("d.f.ip6.arpa.") for Unique Local IPv6 Unicast addresses [RFC4193] SHOULD be sent to the xmDNS multicast address (FF05::FB or its IPv4 equivalent 239.255.255.TBD).

[TBD: See RFC 6303 for an expanded list of domains]

5. Querying

In cases where the desired scope of a query is the local link, Extended Multicast DNS queries MAY be sent with a link-local

[RFC4291] source address to FF05::FB.

Otherwise, Extended Multicast DNS queries SHOULD be sent with a Unique Local IPv6 Unicast (ULA) [RFC4193] source address.

Extended Multicast DNS queries SHOULD NOT be sent with a Global IPv6 Unicast [RFC4291] source address. The Source Address Check rules in Section 11 may not be able to determine whether the query was from an on-site host.

6. Responding

All Extended Multicast DNS responses (including responses sent via unicast) SHOULD be sent with IP TTL set to 255.

Extended Multicast DNS Responders MUST return all available AAAA records with scope equal to or greater than the scope of the source address of the query. Extended Multicast DNS Responders SHOULD NOT include link-local AAAA records unless the source of the query is on the local link.

7. Traffic Reduction

Provisions of Multicast DNS Traffic Reduction, namely, Known Answer Suppression, Multi-Packet Known Answer Suppression, Duplicate Question Suppression, and Duplicate Answer Suppression SHALL be supported in Extended Multicast DNS with the following exceptions:

An Extended Multicast DNS Responder seeing a Multicast DNS Query with the TC (truncated) bit set SHALL defer its response for 1 second and then respond within a randomly selected time interval between 0 and 200 ms.

If the xmDNS Responder receives additional Known-Answer packets with the TC bit set, it SHOULD extend the delay as necessary to ensure a pause of 1 second (plus a random delay between 0 and 200 ms) after the last such packet before it sends its answer.

In multi-hop LAN deployments where a single Multicast DNS Query is propagated for longer than 1 second, the xmDNS Responder SHOULD extend the time it defers its response to at least 1 second longer than the maximum propagation time of a single Multicast DNS Query.

8. Probing and Announcing on Startup

Provisions of Multicast DNS Probing and Announcing SHALL be supported in Extended Multicast DNS.

9. Conflict Resolution

Provisions of Multicast DNS Conflict Resolution SHALL be supported in Extended Multicast DNS. When creating address records (i.e. host names) or resource records where uniqueness (or maintenance of some other defined constraint) is desired, xmDNS Responders SHOULD append some relatively unique string (i.e. low order bits of an EUI-64) to the name in order to minimize name conflict resolution traffic.

10. Resource Record TTL Values and Cache Coherency

Provisions of Resource Record TTL Values and Cache Coherency, namely, Goodbye Packets, Announcements to Flush Outdated Cache Entries, Cache Flush on Topology Change, Cache Flush on Failure indication and Passive Observation of Failure SHALL be supported in Extended Multicast DNS with the following exceptions:

Let TimeActive be the time duration that a single multicast request or response is active in a multi-hop LAN deployment instance (in seconds, rounded up to the next integer value).

Queriers of Extended Multicast DNS receiving a response with TTL of zero SHOULD set the TTL to 1 plus TimeActive and delete the record 1 second plus TimeActive later.

For Announcements to Flush Outdated Cache Entries, all timing values stated as "one second" SHOULD be read as "one second plus TimeActive" to address the propagation of multicast packets in a multi-hop LAN instance.

11. Source Address Check

Source address check must ensure that queries originate from on-site prefixes. All other queries must be silently dropped.

12. Special Characteristics of Extended Multicast DNS Domains

[TBD]

13. Enabling and Disabling Multicast DNS

[TBD]

14. Considerations for Multiple Interfaces

[TBD]

15. Considerations for Multiple Responders on the Same Machine

[TBD]

16. Multicast DNS Character Set

[Same as mDNS]

17. Multicast DNS Message Size

[Same as mDNS]

18. Multicast DNS Message Format

[Same as mDNS]

19. Summary of Differences Between Multicast DNS and Unicast DNS

[Same as mDNS]

20. IPv6 Considerations

An IPv4-only host and an IPv6-only host behave as "ships that pass in the night". Even if they are on the same Ethernet, neither is aware of the other's traffic. For this reason, each physical link may have *two* unrelated ".site." zones, one for IPv4 and one for IPv6. Since for practical purposes, a group of IPv4-only hosts and a group of IPv6-only hosts on the same Ethernet act as if they were on two entirely separate Ethernet segments, it is unsurprising that their use of the ".site." zone should occur exactly as it would if they really were on two entirely separate Ethernet segments.

A dual-stack (v4/v6) host can participate in both ".site." zones, and

should register its name(s) and perform its lookups using both IPv4 and IPv6. This enables it to reach, and be reached by, both IPv4-only and IPv6-only hosts. In effect this acts like a multi-homed host, with one connection to the logical "IPv4 Ethernet segment", and a connection to the logical "IPv6 Ethernet segment". When such a host generates NSEC records, if it is using the same host name for its IPv4 addresses and its IPv6 addresses on that network interface, its NSEC records should indicate that the host name has both A and AAAA records.

21. Security Considerations

[TBD]

22. IANA Considerations

IANA has allocated the IPv6 multicast address set FF0X::FB for Multicast DNS [mcast6]. The use of FF02::FB (Link-Local Scope) is described in [I-D.cheshire-dnsext-multicastdns] and the use of address FF05::FB (Site-Local Scope) is defined in this document.

When this document is published, IANA should designate a list of domains which are deemed to have only site-local significance, as described in Section 12 of this document ("Special Characteristics of Extended Multicast DNS Domains") [I-D.cheshire-dnsext-special-names].

Specifically, the designated site-local domains are:

site.
d.f.ip6.arpa.

[TBD] This document also requests an IPv4 Scope Relative multicast address in the Local Scope range (239.255.255.0/24) [RFC2365] in order to differentiate xmDNS queries from normal mDNS queries and to facilitate modified xmDNS source address check rules.

23. Domain Name Reservation Considerations

The two domains listed in Section 22 above and any names falling within those domains (e.g. "MyServer.someZone.site.", "b.a.9.8.7.6.5.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.b.d.0.1.0.d.f.ip6.arpa.", "www._http._tcp.site.") are special DNS names [I-D.cheshire-dnsext-special-names] in the following ways:

1. Users may use these names as they would other DNS names, entering them anywhere that they would otherwise enter a conventional DNS name, or a dotted decimal IPv4 address, or a literal IPv6 address.

   Since there is no central authority responsible for assigning dot-site names, and all devices on the site-local network are equally entitled to claim any dot-site name, users SHOULD be aware of this and SHOULD exercise appropriate caution. In an untrusted or unfamiliar network environment, users SHOULD be aware that using a name like "www.site" may not actually connect them to the web site they expected, and could easily connect them to a different web page, or even a fake or spoof of their intended web site, designed to trick them into revealing confidential information. As always with networking, end-to-end cryptographic security can be a useful tool. For example, when connecting with ssh, the ssh host key verification process will inform the user if it detects that the identity of the entity they are communicating with has changed since the last time they connected to that name.

2. Application software may use these names as they would other similar DNS names, and is not required to recognize the names and treat them specially. Due to the relative ease of spoofing dot-site names, end-to-end cryptographic security remains important when communicating across a local network, just as it is when communicating across the global Internet.

3. Name resolution APIs and libraries SHOULD recognize these names as special and SHOULD NOT send queries for these names to their configured (unicast) caching DNS server(s). This is to avoid unnecessary load on the root name servers and other name servers, caused by queries for which those name servers do not have useful non-negative answers to give, and will not ever have useful nonnegative answers to give.

4. Caching DNS servers SHOULD recognize these names as special and SHOULD NOT attempt to look up NS records for them, or otherwise query authoritative DNS servers in an attempt to resolve these names. Instead, caching DNS servers SHOULD generate immediate NXDOMAIN responses for all such queries they may receive (from misbehaving name resolver libraries). This is to avoid unnecessary load on the root name servers and other name servers.

5. Authoritative DNS servers SHOULD NOT by default be configurable to answer queries for these names, and, like caching DNS servers, SHOULD generate immediate NXDOMAIN responses for all such queries they may receive. DNS server software MAY provide a

configuration option to override this default, for testing purposes or other specialized uses.

6. DNS server operators SHOULD NOT attempt to configure authoritative DNS servers to act as authoritative for any of these names. Configuring an authoritative DNS server to act as authoritative for any of these names may not, in many cases, yield the expected result, since name resolver libraries and caching DNS servers SHOULD NOT send queries for those names (see 3 and 4 above), so such queries SHOULD be suppressed before they even reach the authoritative DNS server in question, and consequently it will not even get an opportunity to answer them.

7. DNS Registrars MUST NOT allow any of these names to be registered in the normal way to any person or entity. These names are reserved protocol identifiers with special meaning and fall outside the set of names available for allocation by registrars. Attempting to allocate one of these names as if it were a normal DNS domain name will probably not work as desired, for reasons 3, 4, and 6 above.

24. Acknowledgments

We wish to thank the authors of [I-D.cheshire-dnsext-multicastdns] on whose work this document is heavily based. Reviews and comments were provided by Tom Herbst and Ralph Droms.

25. References

25.1. Normative References

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

[RFC2365] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC6303] Andrews, M., "Locally Served DNS Zones", BCP 163, RFC 6303, July 2011.

[mcast6] "IPv6 Multicast Address Space Registry", <http://www.iana.org/assignments/ipv6-multicast-addresses>.

25.2. Informative References

[I-D.cheshire-dnsext-dns-sd]
Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", draft-cheshire-dnsext-dns-sd-11 (work in progress), December 2011.

[I-D.cheshire-dnsext-multicastdns]
Cheshire, S. and M. Krochmal, "Multicast DNS", draft-cheshire-dnsext-multicastdns-15 (work in progress), December 2011.

[I-D.cheshire-dnsext-special-names]
Cheshire, S. and M. Krochmal, "Special-Use Domain Names", draft-cheshire-dnsext-special-names-02 (work in progress), December 2011.

[I-D.ietf-dnsop-default-local-zones]
Andrews, M., "Locally-served DNS Zones", draft-ietf-dnsop-default-local-zones-15 (work in progress), March 2011.

Authors' Addresses

Kerry Lynn
Consultant

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

Don Sturek
Grid2Home
404 Camino Del Rio South, Suite 600
San Diego, CA
USA

Phone: +1 619 504 3615
Email: d.sturek@att.net