Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 26, 2021

T. Lemon
S. Cheshire
Apple Inc.

February 22, 2021

Multicast DNS Discovery Relay
draft-ietf-dnssd-mdns-relay-04

Abstract

This document complements the specification of the Discovery Proxy for Multicast DNS-Based Service Discovery. It describes a lightweight relay mechanism, a Discovery Relay, which, when present on a link, allows remote clients, not attached to that link, to perform mDNS discovery operations on that link.

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 https://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 August 26, 2021.

Table of Contents

1. Introduction

This document defines a Discovery Relay. A Discovery Relay is a companion technology that works in conjunction with Discovery Proxies, and other clients.

The Discovery Proxy for Multicast DNS-Based Service Discovery [RFC8766] is a mechanism for discovering services on a subnetted network through the use of Discovery Proxies. Discovery Proxies issue Multicast DNS (mDNS) requests [RFC6762] on various multicast links in the network on behalf of a remote host performing DNS-Based Service Discovery [RFC6763].

In the original Discovery Proxy specification, it was imagined that for every multicast link on which services will be discovered, a host will be present running a full Discovery Proxy. This document introduces a lightweight Discovery Relay that can be used in conjunction with a central Discovery Proxy to provide discovery services on a multicast link without requiring a full Discovery Proxy on every multicast link.

The primary purpose of a Discovery Relay is providing remote virtual interface functionality to Discovery Proxies, and this document is written with that usage in mind. However, in principle, a Discovery Relay could be used by any properly authorized client. In the context of this specification, a Discovery Proxy is a client to the Discovery Relay. This document uses the terms "Discovery Proxy" and "Client" somewhat interchangably; the term "Client" is used when we are talking about the communication between the Client and the Relay, and the term "Discovery Proxy" when we are referring specifically to a Discovery Relay Client that also happens to be a Discovery Proxy. One example of another kind of device that can be a client of a Discovery Relay is an Advertising Proxy [AdProx].

The Discovery Relay operates by listening for TCP connections from Clients. When a Client connects, the connection is authenticated and secured using TLS. The Client can then specify one or more multicast links from which it wishes to receive mDNS traffic. The Client can also send messages to be transmitted on its behalf on one or more of those multicast links. DNS Stateful Operations (DSO) [RFC8490] is used as a framework for conveying interface and IP header information associated with each message. DSO formats its messages using type-length-value (TLV) data structures. This document defines additional DSO TLV types, used to implement the Discovery Relay functionality.

The Discovery Relay functions essentially as a set of one or more remote virtual interfaces for the Client, one on each multicast link to which the Discovery Relay is connected. In a complex network, it

is possible that more than one Discovery Relay will be connected to the same multicast link; in this case, the Client ideally should only be using one such Relay Proxy per multicast link, since using more than one will generate duplicate traffic.

How such duplication is detected and avoided is out of scope for this document; in principle it could be detected using HNCP [RFC7788] or configured using some sort of orchestration software in conjunction with NETCONF [RFC6241] or CPE WAN Management Protocol [TR-069].

Use of a Discovery Relay can be considered similar to using Virtual LAN (VLAN) trunk ports to give a Discovery Proxy device a virtual presence on multiple links or broadcast domains. The difference is that while a VLAN trunk port operates at the link layer and delivers all link-layer traffic to the Discovery Proxy device, a Discovery Relay operates further up the network stack and selectively delivers only relevant Multicast DNS traffic. Also, VLAN trunk ports are generally only available within a single administrative domain and require link-layer configuration and connectivity, whereas the Discovery Relay protocol, which runs over TCP, can be used between any two devices with IP connectivity to each other.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here. These words may also appear in this document in lower case as plain English words, absent their normative meanings.

The following definitions may be of use:

Client A network service that uses a Discovery Relay to send and receive mDNS multicast traffic on a remote link, to enable it to communicate with mDNS Agents on that remote link.

mDNS Agent A host which sends and/or responds to mDNS queries directly on its local link(s). Examples include network cameras, networked printers, networked home electronics, etc.

Discovery Proxy A network service which receives well-formed questions using the DNS protocol, performs multicast DNS queries to find answers to those questions, and responds with those answers using the DNS protocol. A Discovery Proxy that can communicate with remote mDNS Agents, using the services of a Discovery Relay, is a Client of the Discovery Relay.

Discovery Relay A network service which relays mDNS messages received on a local link to a Client, and on behalf of that Client can transmit mDNS messages on a local link.

multicast link A maximal set of network connection points, such that any host connected to any connection point in the set may send a packet with a link-local multicast destination address (specifically the mDNS link-local multicast destination address [RFC6762]) that will be received by all hosts connected to all other connection points in the set. Note that it is becoming increasingly common for a multicast link to be smaller than its corresponding unicast link. For example it is becoming common to have multiple Wi-Fi access points on a shared Ethernet backbone, where the multiple Wi-Fi access points and their shared Ethernet backbone form a single unicast link (a single IPv4 subnet, or single IPv6 prefix) but not a single multicast link. Unicast packets sent directly between two hosts on that IPv4 subnet or IPv6 prefix, without passing through an intervening IP-layer router, are correctly delivered, but multicast packets are not forwarded between the various Wi-Fi access points. Given the slowness of Wi-Fi multicast [I-D.ietf-mboned-ieee802-mcast-problems], having a packet that may be of interest to only one or two end systems transmitted to hundreds of devices, across multiple Wi-Fi access points, is especially wasteful. Hence the common configuration decision to not forward multicast packets between Wi-Fi access points is very reasonable. This further motivates the need for technologies like Discovery Proxy and Discovery Relay to facilitate discovery on these networks.

allow-list A list of one or more IP addresses from which a Discovery Relay may accept connections.

silently discard When a message that is not supported or not permitted is received, and the required response to that message is to "silently discard" it, that means that no response is sent by the service that is discarding the message to the service that sent it. The service receiving the message may log the event, and may also count such events: "silently" does not preclude such behavior.

Take care when reading this document not to confuse the terms "Discovery Proxy" and "Discovery Relay". A Discovery Proxy [RFC8766] provides Multicast DNS discovery service to remote clients. A Discovery Relay is a simple software entity that provides virtual link connectivity to one or more Discovery Proxies or other Discovery Relay clients.

3. Protocol Overview

This document describes a way for a Client to communicate with mDNS agents on remote multicast links to which the client is not directly connected, using a Discovery Relay. As such, there are two parts to the protocol: connections between Clients and Discovery Relays, and communications between Discovery Relays and mDNS agents.

3.1. Connections between Clients and Relays (overview)

Discovery Relays listen for incoming connection requests. Connections between Clients and Discovery Relays are established by Clients. Connections are authenticated and encrypted using TLS, with both client and server certificates. Connections are long-lived: a Client is expected to send many queries over a single connection, and Discovery Relays will forward all mDNS traffic from subscribed interfaces over the connection.

The stream encapsulated in TLS will carry DNS frames as in the DNS TCP protocol [RFC1035] Section 4.2.2. However, all messages will be DSO messages [RFC8490]. There will be four types of such messages between Discovery Relays and Clients:

- Control messages from Client to Relay
- Link status messages from Relay to Client
- Encapsulated mDNS messages from Client to Relay
- Encapsulated mDNS messages from Relay to Client

Clients can send four different control messages to Relays: Link State Request, Link State Discontinue, Link Data Request and Link Data Discontinue. The first two are used by the Client to request that the Relay report on the set of links that can be requested, and to request that it discontinue such reporting. The second two are used by the Client to indicate to the Discovery Relay that mDNS messages from one or more specified multicast links are to be relayed to the Client, and to subsequently stop such relaying.

Link Status messages from a Discovery Relay to the Client inform the Client that a link has become available, or that a formerly-available link is no longer available.

Encapsulated mDNS messages from a Discovery Relay to a Client are sent whenever an mDNS message is received on a multicast link to which the Discovery Relay has subscribed.

Encapsulated mDNS messages from a Client to a Discovery Relay cause the Discovery Relay to transmit the mDNS message on the specified multicast link to which the Discovery Relay host is directly attached.

During periods with no traffic flowing, Clients are responsible for generating any necessary keepalive traffic, as stated in the DSO specification [RFC8490].

3.2. mDNS Messages On Multicast Links

Discovery Relays listen for mDNS traffic on all configured multicast links that have at least one active subscription from a Client. When an mDNS message is received on a multicast link, it is forwarded on every open Client connection that is subscribed to mDNS traffic on that multicast link. In the event of congestion, where a particular Client connection has no buffer space for an mDNS message that would otherwise be forwarded to it, the mDNS message is not forwarded to it. Normal mDNS retry behavior is used to recover from this sort of packet loss. Discovery Relays are not expected to buffer more than a few mDNS packets. Excess mDNS packets are silently discarded. In practice this is not expected to be a issue. Particularly on networks like Wi-Fi, multicast packets are transmitted at rates ten or even a hundred times slower than unicast packets. This means that even at peak multicast packets rates, it is likely that a unicast TCP connection will able to carry those packets with ease.

Clients send encapsulated mDNS messages they wish to have sent on their behalf on remote multicast link(s) on which the Client has an active subscription. A Discovery Relay will not transmit mDNS packets on any multicast link on which the Client does not have an active subscription, since it makes no sense for a Client to ask to have a query sent on its behalf if it’s not able to receive the responses to that query.

4. Connections between Clients and Relays (details)

When a Discovery Relay starts, it opens a passive TCP listener to receive incoming connection requests from Clients. This listener may be bound to one or more source IP addresses, or to the wildcard address, depending on the implementation. When a connection is received, the relay must first validate that it is a connection to an IP address to which connections are allowed. For example, it may be that only connections to ULAs are allowed, or to the IP addresses configured on certain interfaces. If the listener is bound to a specific IP address, this check is unnecessary.

If the relay is using an IP address allow-list, the next step is for the relay to verify that that the source IP address of the connection is on its allow-list. If the connection is not permitted either because of the source address or the destination address, the Discovery Relay closes the connection. If possible, before closing the connection, the Discovery Relay first sends a TLS user_canceled alert ([RFC8446] Section 6.1). Discovery Relays SHOULD refuse to accept TCP connections to invalid destination addresses, rather than accepting and then closing the connection, if this is possible.

Otherwise, the Discovery Relay will attempt to complete a TLS handshake with the Client. Clients are required to send the post_handshake_auth extension ([RFC8446] Section 4.2.5). If a Discovery Relay receives a ClientHello message with no post_handshake_auth extension, the Discovery Relay rejects the connection with a certificate_required alert ([RFC8446] Section 6.2).

Once the TLS handshake is complete, the Discovery Relay MUST request post-handshake authentication ([RFC8446] Section 4.6.2). If the Client refuses to send a certificate, or the key presented does not match the key associated with the IP address from which the connection originated, or the CertificateVerify does not validate, the connection is dropped with the TLS access_denied alert ([RFC8446] Section 6.2).

Clients MUST validate server certificates. If the client is configured with a server IP address and certificate, it can validate the server by comparing the certificate offered by the server to the certificate that was provided: they should be the same. If the certificate includes a Distinguished Name that is a fully-qualified domain name, the client SHOULD present that domain name to the server in an SNI request.

Rather than being configured with an IP address and a certificate, the client may be configured with the server's FQDN. In this case, the client uses the server's FQDN as a Authentication Domain Name

[RFC8310] Section 7.1, and uses the authentication method described in [RFC8310] section 8.1, if the certificate is signed by a root authority the client trusts, or the method described in section 8.2 of the same document if not. If neither method is available, then a locally-configured copy of the server certificate can be used, as in the previous paragraph.

Once the connection is established and authenticated, it is treated as a DNS TCP connection [RFC7766].

Aliveness of connections between Clients and Relays is maintained as described in Section 4 of the DSO specification [RFC8490]. Clients must also honor the 'Retry Delay' TLV (section 5 of [RFC8490]) if sent by the Discovery Relay.

Clients SHOULD avoid establishing more than one connection to a specific Discovery Relay. However, there may be situations where multiple connections to the same Discovery Relay are unavoidable, so Discovery Relays MUST be willing to accept multiple connections from the same Client.

In order to know what links to request, the Client can be configured with a list of links supported by the Relay. However, in some networking contexts, dynamic changes in the availability of links are likely; therefore Clients may also use the Report Link Changes TLV to request that the Relay report on the availability of its links. In some contexts, for example when debugging, a Client may operate with no information about the set of links supported by a relay, simply relying on the relay to provide one.

5. Traffic from Relays to Clients

The mere act of connecting to a Discovery Relay does not result in any mDNS traffic being forwarded. In order to request that mDNS traffic from a particular multicast link be forwarded on a particular connection, the Client must send one or more DSO messages, each containing a single mDNS Link Data Request TLV (Section 8.1) indicating the multicast link from which traffic is requested.

When an mDNS Link Data Request message is received, the Discovery Relay validates that it recognizes the link identifier, and that forwarding is enabled for that link. If both checks are successful, it MUST send a response with RCODE=0 (NOERROR). If the link identifier is not recognized, it sends a response with RCODE=3 (NXDOMAIN/Name Error). If forwarding from that link to the Client is not enabled, it sends a response with RCODE=5 (REFUSED). If the relay cannot satisfy the request for some other reason, for example resource exhaustion, it sends a response with RCODE=2 (SERVFAIL).

If the requested link is valid, the Relay begins forwarding all mDNS messages from that link to the Client. Delivery is not guaranteed: if there is no buffer space, packets will be dropped. It is expected that regular mDNS retry processing will take care of retransmission of lost packets. The amount of buffer space is implementation dependent, but generally should not be more than the bandwidth delay product of the TCP connection [RFC7323]. The Discovery Relay should use the TCP_NOTSENT_LOWAT mechanism [NOTSENT][PRIO] or equivalent, to avoid building up a backlog of data in excess of the amount necessary to have in flight to fill the bandwidth delay product of the TCP connection.

Encapsulated mDNS messages from Relays to Clients are framed within DSO messages. Each DSO message can contain multiple TLVs, but only a single encapsulated mDNS message is conveyed per DSO message. Each forwarded mDNS message is sent in an Encapsulated mDNS Message TLV (Section 8.4). The source IP address and port of the message MUST be encoded in an IP Source TLV (Section 8.5). The multicast link on which the message was received MUST be encoded in a Link Identifier TLV (Section 8.3). As described in the DSO specification [RFC8490], a Client MUST silently ignore unrecognized Additional TLVs in mDNS messages, and MUST NOT discard mDNS messages that include unrecognized Additional TLVs.

A Client may discontinue listening for mDNS messages on a particular multicast link by sending a DSO message containing an mDNS Link Data Discontinue TLV (Section 8.2). The Discovery Relay MUST discontinue forwarding mDNS messages when the Link Data Discontinue request is received. However, messages from that link that had previously been

queued may arrive after the Client has discontinued its listening. The Client should silently discard such messages. The Discovery Relay does not respond to the Link Data Discontinue message other than to discontinue forwarding mDNS messages from the specified links.

6. Traffic from Clients to Relays

Like mDNS traffic from relays, each mDNS message sent by a Client to a Discovery Relay is communicated in an Encapsulated mDNS Message TLV (Section 8.4) within a DSO message. Each message MUST contain exactly one Link Identifier TLV (Section 8.3). The Discovery Relay will transmit the mDNS message to the mDNS port and multicast address on the link specified in the message using the specified IP address family.

Although the communication between Clients and Relays uses the DNS stream protocol and DNS Stateless Operations, there is no case in which a Client would legitimately send a DNS query (or anything else other than a DSO message) to a Relay. Therefore, if a Relay receives any message other than a DSO message, it MUST immediately abort that DSO session with a TCP reset (RST).

When defining this behavior, the working group considered making it possible to specify more than one link identifier in an mDNSMessage TLV. A superficial evaluation of this suggested that this might be a useful optimization, since when a query is issued, it will often be issued to all links. However, on many link types, like Wi-Fi, multicast traffic is expensive [I-D.ietf-mboned-ieee802-mcast-problems] and should be generated frugally, so providing convenient ways to generate additional multicast traffic was determined to be an unwise optimization. In addition, because of the way mDNS handles retries, it will almost never be the case that the exact same message will be sent on more than one link. Therefore, the complexity that this optimization adds is not justified by the potential benefit, and this idea has been abandoned.

7. Discovery Proxy Behavior

Discovery Proxies treat multicast links for which Discovery Relay service is being used as if they were virtual interfaces; in other words, a Discovery Proxy serving multiple remote multicast links using multiple remote Discovery Relays behaves the same as a Discovery Proxy serving multiple local multicast links using multiple local physical network interfaces. In this section we refer to multicast links served directly by the Discovery Proxy as locally-connected links, and multicast links served through the Discovery Relay as relay-connected links. A relay-connected link can be thought of as similar to a link that a Discovery Proxy connects to using a USB Ethernet interface, just with a very long USB cable (that runs over TCP).

When a Discovery Proxy receives a DNS query from a DNS client via unicast, it will generate corresponding mDNS query messages on the relevant multicast link(s) for which it is acting as a proxy. For locally-connected link(s), those query messages will be sent directly. For relay-connected link(s), the query messages will be sent through the Discovery Relay that is being used to serve that multicast link.

Responses from devices on locally-connected links are processed normally. Responses from devices on relay-connected links are received by the Discovery Relay, encapsulated, and forwarded to the Client; the Client then processes these messages using the link-identifying information included in the encapsulation.

In principle it could be the case that some device is capable of performing service discovery using Multicast DNS, but not using traditional unicast DNS. Responding to mDNS queries received from the Discovery Relay could address this use case. However, continued reliance on multicast is counter to the goals of the current work in service discovery, and to benefit from wide-area service discovery such client devices should be updated to support service discovery using unicast queries.

8. DSO TLVs

This document defines a modest number of new DSO TLVs.

8.1. mDNS Link Data Request

The mDNS Link Data Request TLV conveys a link identifier from which a Client is requesting that a Discovery Relay forward mDNS traffic. The link identifier comes from the provisioning configuration (see Section 9). The DSO-TYPE for this TLV is TBD-R. DSO-LENGTH is always 5. DSO-DATA is the 8-bit address family followed by the link identifier, a 32-bit unsigned integer in network (big endian) byte order, as described in Section 9. An address family value of 1 indicates IPv4 and 2 indicates IPv6, as recorded in the IANA Registry of Address Family Numbers [AdFam].

The mDNS Link Data Request TLV can only be used as a primary TLV, and requires an acknowledgement.

At most one mDNS Link Data Request TLV may appear in a DSO message. To request multiple link subscriptions, multiple separate DSO messages are sent, each containing a single mDNS Link Data Request TLV.

A Client MUST NOT request a link if it already has an active subscription to that link on the same DSO connection. If a Discovery Relay receives a duplicate link subscription request, it MUST immediately abort that DSO session with a TCP reset (RST).

8.2. mDNS Link Data Discontinue

The mDNS Link Data Discontinue TLV is used by Clients to unsubscribe to mDNS messages on the specified multicast link. DSO-TYPE is TBD-D. DSO-LENGTH is always 5. DSO-DATA is the 8-bit address family followed by the 32-bit link identifier, a 32-bit unsigned integer in network (big endian) byte order, as described in Section 9.

The mDNS Link Data Discontinue TLV can only be used as a DSO unidirectional message TLV, and is not acknowledged.

At most one mDNS Link Data Discontinue TLV may appear in a DSO message. To unsubscribe from multiple links, multiple separate DSO messages are sent, each containing a single mDNS Link Data Discontinue TLV.

8.3. Link Identifier

This option is used both in DSO messages from Discovery Relays to Clients that contain received mDNS messages, and from Clients to Discovery Relays that contain mDNS messages to be transmitted on the multicast link. In the former case, it indicates the multicast link on which the message was received; in the latter case, it indicates the multicast link on which the message should be transmitted. DSO-TYPE is TBD-L. DSO-LENGTH is always 5. DSO-DATA is the 8-bit address family followed by the link identifier, a 32-bit unsigned integer in network (big endian) byte order, as described in Section 9.

The Link Identifier TLV can only be used as an additional TLV. The Link Identifier TLV can only appear at most once in a Discovery Relay DSO message.

8.4. Encapsulated mDNS Message

The Encapsulated mDNS Message TLV is used to communicate an mDNS message that a Relay is forwarding from a multicast link to a Client, or that a Client is sending to a Relay for transmission on a multicast link. Only the application-layer payload of the mDNS message is carried in the DSO "Encapsulated mDNS Message" TLV, i.e., just the DNS message itself, beginning with the DNS Message ID, not the IP or UDP headers. The DSO-TYPE for this TLV is TBD-M. DSO-LENGTH is the length of the encapsulated mDNS message. DSO-DATA is the content of the encapsulated mDNS message.

The Encapsulated mDNS Message TLV can only be used as a DSO unidirectional message TLV, and is not acknowledged.

8.5. IP Source

The IP Source TLV is used to report the IP source address and port from which an mDNS message was received. This TLV is present in DSO messages from Discovery Relays to Clients that contain encapsulated mDNS messages. DSO-TYPE is TBD-S. DSO-LENGTH is either 6, for an IPv4 address, or 18, for an IPv6 address. DSO-DATA is the two-byte source port, followed by the 4- or 16-byte IP Address. Both port and address are in the canonical byte order (i.e., the same representation as used in the UDP and IP packet headers, with no byte swapping).

The IP Source TLV can only be used as an additional TLV. The IP Source TLV can only appear at most once in a Discovery Relay DSO message.

8.6. Link State Request

The Link State Request TLV requests that the Discovery Relay report link changes. When the relay is reporting link changes and a new link becomes available, it sends a Link Available message to the Client. When a link becomes unavailable, it sends a Link Unavailable message to the Client. If there are links available when the request is received, then for each such link the relay immediately sends a Link Available Message to the Client. DSO-TYPE is TBD-P. DSO-LENGTH is 0.

The mDNS Link State Request TLV can only be used as a primary TLV, and requires an acknowledgement. The acknowledgment does not contain a Link Available TLV: it is just a response to the Link State Request message.

8.7. Link State Discontinue

The Link State Discontinue TLV requests that the Discovery Relay stop reporting on the availability of links supported by the relay. This cancels the effect of a Link State Request TLV. DSO-TYPE is TBD-Q. DSO-LENGTH is 0.

The mDNS Link State Discontinue TLV can only be used as a DSO unidirectional message TLV, and is not acknowledged.

8.8. Link Available

The Link Available TLV is used by Discovery Relays to indicate to Clients that a new link has become available. The format is the same as the Link Identifier TLV. DSO-TYPE is TBD-V. The Link Available TLV may be accompanied by one or more Link Prefix TLVs which indicate IP prefixes the Relay knows to be present on the link.

The mDNS Link Available TLV can only be used as a DSO unidirectional message TLV, and is not acknowledged.

8.9. Link Unavailable

The Link Unavailable TLV is used by Discovery Relays to indicate to Clients that an existing link has become unavailable. The format is the same as the Link Identifier TLV. DSO-TYPE is TBD-U.

The mDNS Link Unavailable TLV can only be used as a DSO unidirectional message TLV, and is not acknowledged.

8.10. Link Prefix

The Link Prefix TLV represents an IP address or prefix configured on a link. The length is 17 for an IPv6 address or prefix, and 5 for an IPv4 address or prefix. The TLV consists of a prefix length, between 0 and 32 for IPv4 or between 0 and 128 for IPv6, represented as a single byte. This is followed by the IP address, either four or sixteen bytes. DSO-TYPE is TBD-K.

The Link Prefix TLV can only be used as a secondary TLV.

9. Provisioning

In order for a Discovery Proxy to use Discovery Relays, it must be configured with sufficient information to identify multicast links on which service discovery is to be supported and, if it is not running on a host that is directly connected to those multicast links, connect to Discovery Relays supporting those multicast links.

A Discovery Relay must be configured both with a set of multicast links to which the host on which it is running is connected, on which mDNS relay service is to be provided, and also with a list of one or more Clients authorized to use it.

On a network supporting DNS Service Discovery using Discovery Relays, more than one different Discovery Relay implementation may be present. While it may be that only a single Discovery Proxy is present, that implementation will need to be able to be configured to interoperate with all of the Discovery Relays that are present. Consequently, it is necessary that a standard set of configuration parameters be defined for both Discovery Proxies and Discovery Relays.

DNS Service Discovery generally operates within a constrained set of links, not across the entire internet. This section assumes that what will be configured will be a limited set of links operated by a single entity or small set of cooperating entities, among which services present on each link should be available to users on that link and every other link. This could be, for example, a home network, a small office network, or even a network covering an entire building or small set of buildings. The set of Discovery Proxies and Discovery Relays within such a network will be referred to in this section as a 'Discovery Domain'.

Depending on the context, several different candidates for configuration of Discovery Proxies and Discovery Relays may be applicable. The simplest such mechanism is a manual configuration file, but regardless of provisioning mechanism, certain configuration information needs to be communicated to the devices, as outlined below.

In the example we provide here, we only refer to configuring of IP addresses, private keys and certificates. It is also possible to use FQDNs to identify servers; this then allows for the use of DANE ([RFC8310] Section 8.2) or PKIX authentication [RFC6125]. Which method is used is to some extent up to the implementation, but at a minimum, it should be possible to associate an IP address with a self-signed certificate, and it should be possible to validate both

self-signed and PKIX-authenticated certificates, with PKIX, DANE or a pre-configured trust anchor.

9.1. Provisioned Objects

Three types of objects must be described in order for Discovery Proxies and Discovery Relays to be provisioned: Discovery Proxies, Multicast Links, and Discovery Relays. "Human-readable" below means actual words or proper names that will make sense to an untrained human being. "Machine-readable" means a name that will be used by machines to identify the entity to which the name refers. Each entity must have a machine-readable name and may have a human-readable name. No two entities can have the same human-readable name. Similarly, no two entities can have the same machine-readable name.

9.1.1. Multicast Link

The description of a multicast link consists of:

link-identifier A 32-bit identifier that uniquely identifies that link within the Discovery Domain. Each link MUST have exactly one such identifier. Link Identifiers do not have any special semantics, and are not intended to be human-readable.

ldh-name A fully-qualified domain name for the multicast link that is used to form an LDH domain name as described in section 5.3 of the Discovery Proxy specification [RFC8766]. This name is used to identify the link during provisioning, and must be present.

hr-name A human-readable user-friendly fully-qualified domain name for the multicast link. This name MUST be unique within the Discovery Domain. Each multicast link MUST have exactly one such name. The hr-name MAY be the same as the ldh-name. (The hr-name is allowed to contain spaces, punctuation and rich text, but it is not required to do so.)

The ldh-name and hr-name can be used to form the LDH and human-readable domain names as described in [RFC8766], section 5.3.

Note that the ldh-name and hr-name can be used in two different ways.

On a small home network with little or no human administrative configuration, link names may be directly visible to the user. For example, a search in 'home.arpa' on a small home network may discover services on both ethernet.home.arpa and wi-fi.home.arpa. In the case of a home user who has one Ethernet-connected printer and one Wi-Fi-connected printer, discovering that they have one printer on ethernet.home.arpa and another on wi-fi.home.arpa is understandable and meaningful.

On a large corporate network with hundreds of Wi-Fi access points, the individual link names of the hundreds of multicast links are less likely to be useful to end users. In these cases, Discovery Broker functionality [I-D.sctl-discovery-broker] may be used to translate the many link names to something more meaningful to users. For example, in a building with 50 Wi-Fi access points, each with their own link names, services on all the different physical links may be presented to the user as appearing in 'headquarters.example.com'. In this case, the individual link names can be thought of similar to MAC addresses or IPv6 addresses. They are used internally by the software as unique identifiers, but generally are not exposed to end users.

9.1.2. Discovery Proxy

The description of a Discovery Proxy consists of:

name a machine-readable name used to reference this Discovery Proxy in provisioning.

hr-name an optional human-readable name which can appear in provisioning, monitoring and debugging systems. Must be unique within a Discovery Domain.

certificate a certificate that identifies the Discovery Proxy. This certificate can be shared across services on the Discovery Proxy Host. The public key in the certificate is used both to uniquely identify the Discovery Proxy and to authenticate connections from it. The certificate should be signed by its own private key.

private-key the private key corresponding to the public key in the certificate.

source-ip-addresses a list of IP addresses that may be used by the Discovery Proxy when connecting to Discovery Relays. These addresses should be addresses that are configured on the Discovery Proxy Host. They should not be temporary addresses. All such addresses must be reachable within the Discovery Domain.

public-ip-addresses a list of IP addresses that a Discovery Proxy listens on to receive requests from clients. This is not used for interoperation with Discovery Relays, but is mentioned here for completeness: the list of addresses listened on for incoming client requests may differ from the 'source-ip-addresses' list of addresses used for issuing outbound connection requests to Discovery Relays. If any of these addresses are reachable from outside of the Discovery Domain, services in that domain will be discoverable outside of the domain.

multicast links a list of multicast links on which this Discovery Proxy is expected to provide service

The private key should never be distributed to other hosts; all of the other information describing a Discovery Proxy can be safely shared with Discovery Relays.

In some configurations it may make sense for the Discovery Relay not to have a list of links, but simply to support the set of all links available on relays to which the Discovery Proxy is configured to communicate.

9.1.3. Discovery Relay

The description of a Discovery Relay consists of:

name a required machine-readable identifier used to reference the relay

hr-name an optional human-readable name which can appear in provisioning, monitoring and debugging systems. Must be unique within a Discovery Domain.

certificate a certificate that identifies the Discovery Relay. This certificate can be shared across services on the Discovery Relay Host. Indeed, if a Discovery Proxy and Discovery Relay are running on the same host, the same certificate can be used for both. The public key in the certificate uniquely identifies the Discovery Relay and is used by a Discovery Relay Client (e.g., a Discovery Proxy) to verify that it is talking to the intended Discovery Relay after a TLS connection has been established. The certificate must either be signed by its own key, or have a signature chain that can be validated using PKIX authentication [RFC6125].

private-key the private key corresponding to the public key in the certificate.

listen-tuple a list of IP address/port tuples that may be used to connect to the Discovery Relay. The relay may be configured to listen on all addresses on a single port, but this is not required, so the port as well as the address must be specified.

multicast links a list of multicast links to which this relay is physically connected.

The private key should never be distributed to other hosts; all of the other information describing a Discovery Relay can be safely shared with Discovery Proxies.

In some cases a Relay may not be configured with a static list of links, but may simply discover links by monitoring the set of available interfaces on the host on which the Relay is running. In that case, the relay could be configured to identify links based on the names of network interfaces, or based on the set of available prefixes seen on those interfaces. The details of this sort of configuration are not specified in this document.

9.2. Configuration Files

For this discussion, we assume the simplest possible means of configuring Discovery Proxies and Discovery Relays: the configuration file. Any environment where changes will happen on a regular basis will either require some automatic means of generating these configuration files as the network topology changes, or will need to use a more automatic method for configuration, such as HNCP [RFC7788].

There are many different ways to organize configuration files. This discussion assumes that multicast links, relays and proxies will be specified as objects, as described above, perhaps in a master file, and then the specific configuration of each proxy or relay will reference the set of objects in the master file, referencing objects by name. This approach is not required, but is simply shown as an example. In addition, the private keys for each proxy or relay must appear only in that proxy or relay's configuration file.

The master file contains a list of Discovery Relays, Discovery Proxies and Multicast Links. Each object has a name and all the other data associated with it. We do not formally specify the format of the file, but it might look something like this:

```
Relay upstairs
  certificate xxx
  listen-tuple 192.0.2.1 1917
  listen-tuple fd00::1 1917
  link upstairs-wifi
  link upstairs-wired
  client-allow-list main

Relay downstairs
  certificate yyy
  listen-tuple 192.51.100.1 2088
  listen-tuple fd00::2 2088
  link downstairs-wifi
  link downstairs-wired
  client-allow-list main

Proxy main
  certificate zzz
  address 203.1.113.1

Link upstairs-wifi
  id 1
  hr-name Upstairs Wifi

Link upstairs-wired
  id 2
  hr-name Upstairs Wired

Link downstairs-wifi
  id 3
  hr-name Downstairs Wifi

Link downstairs-wired
  id 4
  hr-name Downstairs Wired
```

9.3. Discovery Proxy Private Configuration

The Discovery Proxy configuration contains enough information to identify which Discovery Proxy is being configured, enumerate the list of multicast links it is intended to serve, and provide keying information it can use to authenticate to Discovery Relays. It may also contain custom information about the port and/or IP address(es) on which it will respond to DNS queries.

An example configuration, following the convention used in this section, might look something like this:

```
Proxy main
  private-key zzz
  subscribe upstairs-wifi
  subscribe downstairs-wifi
  subscribe upstairs-wired
  subscribe downstairs-wired
```

When combined with the master file, this configuration is sufficient for the Discovery Proxy to identify and connect to the Discovery Relays that serve the links it is configured to support.

9.4. Discovery Relay Private Configuration

The Discovery Relay configuration just needs to tell the Discovery Relay what name to use to find its configuration in the master file, and what the private key is corresponding to its certificate (public key) in the master file. For example:

```
Relay Downstairs
  private-key yyy
```

10. Security Considerations

Part of the purpose of the Multicast DNS Discovery Relay protocol is to place a simple relay, analogous to a BOOTP relay, into routers and similar devices that may not be updated frequently. The BOOTP [RFC0951] protocol has been around since 1985, and continues to be useful today. The BOOTP protocol uses no encryption, and in many enterprise networks this is considered acceptable. In contrast, the Discovery Relay protocol requires TLS 1.3. A concern is that after 20 or 30 years, TLS 1.3, or some of the encryption algorithms it uses, may become obsolete, rendering devices that require it unusable. Our assessment is that TLS 1.3 probably will be around for many years to come. TLS 1.0 [RFC2246] was used for about a decade, and similarly TLS 1.2 [RFC5246] was also used for about a decade. We expect TLS 1.3 [RFC8446] to have at least that lifespan. In addition, recent IETF efforts are pushing for better software update practices for devices like routers, for other security reasons, making it likely that in ten years time it will be less common to be using routers that haven't had a software update for ten years. However, authors of encryption specifications and libraries should be aware of the potential backwards compatibility issues if an encryption algorithm becomes deprecated. This specification RECOMMENDS that if an encryption algorithm becomes deprecated, then rather than remove that encryption algorithm entirely, encryption libraries should disable that encryption algorithm by default, but leave the code present with an option for client software to enable it in special cases, such as a recent Client talking to an ancient Discovery Relay. Using no encryption, like BOOTP, would eliminate this backwards compatibility concern, but we feel that in such a future hypothetical scenario, using even a weak encryption algorithm still makes passive eavesdropping and tampering harder, and is preferable to using no encryption at all.

11. IANA Considerations

The IANA is kindly requested to update the DSO Type Codes Registry [RFC8490] by allocating codes for each of the TBD type codes listed in the following table, and by updating this document, here and in Section 8. Each type code should list this document as its reference document.

DSO-TYPE	Status	Name
TBD-R	Standard	Link Data Request
TBD-D	Standard	Link Data Discontinue
TBD-L	Standard	Link Identifier
TBD-M	Standard	Encapsulated mDNS Message
TBD-S	Standard	IP Source
TBD-P	Standard	Link State Request
TBD-Q	Standard	Link State Discontinue
TBD-V	Standard	Link Available
TBD-U	Standard	Link Unavailable
TBD-K	Standard	Link Prefix

DSO Type Codes to be allocated

12. Acknowledgments

Thanks to Derek Atkins for the secdir early review.

13. References

13.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <https://www.rfc-editor.org/info/rfc6125>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <https://www.rfc-editor.org/info/rfc6241>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <https://www.rfc-editor.org/info/rfc6762>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <https://www.rfc-editor.org/info/rfc6763>.

[RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, Ed., "TCP Extensions for High Performance", RFC 7323, DOI 10.17487/RFC7323, September 2014, <https://www.rfc-editor.org/info/rfc7323>.

[RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and D. Wessels, "DNS Transport over TCP - Implementation Requirements", RFC 7766, DOI 10.17487/RFC7766, March 2016, <https://www.rfc-editor.org/info/rfc7766>.

[RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <https://www.rfc-editor.org/info/rfc7788>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC8490] Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S., Lemon, T., and T. Pusateri, "DNS Stateful Operations", RFC 8490, DOI 10.17487/RFC8490, March 2019, <https://www.rfc-editor.org/info/rfc8490>.

[RFC8766] Cheshire, S., "Discovery Proxy for Multicast DNS-Based Service Discovery", RFC 8766, DOI 10.17487/RFC8766, June 2020, <https://www.rfc-editor.org/info/rfc8766>.

13.2. Informative References

[AdFam] "IANA Address Family Numbers Registry", <https://www.iana.org/assignments/address-family-numbers/>.

[AdProx] Cheshire, S. and T. Lemon, "Advertising Proxy for DNS-SD Service Registration Protocol", draft-sctl-advertising-proxy-00 (work in progress), July 2020.

[I-D.ietf-mboned-ieee802-mcast-problems] Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-12 (work in progress), October 2020.

[I-D.sctl-discovery-broker] Cheshire, S. and T. Lemon, "Service Discovery Broker", draft-sctl-discovery-broker-00 (work in progress), July 2017.

[NOTSENT] Dumazet, E., "TCP_NOTSENT_LOWAT socket option", July 2013, <https://lwn.net/Articles/560082/>.

[PRIO] Chan, W., "Prioritization Only Works When There's Pending Data to Prioritize", January 2014, <https://insouciant.org/tech/prioritization-only-works-when-theres-pending-data-to-prioritize/>.

[RFC0951] Croft, W. and J. Gilmore, "Bootstrap Protocol", RFC 951, DOI 10.17487/RFC0951, September 1985, <https://www.rfc-editor.org/info/rfc951>.

[RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, DOI 10.17487/RFC2246, January 1999, <https://www.rfc-editor.org/info/rfc2246>.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <https://www.rfc-editor.org/info/rfc5246>.

[TR-069] Broadband Forum, "CPE WAN Management Protocol", November 2013, <https://www.broadband-forum.org/technical/download/TR-069_Amendment-5.pdf>.

Authors' Addresses

Ted Lemon
Apple Inc.
One Apple Park Way
Cupertino, California 95014
United States of America

Phone: +1 (408) 996-1010
Email: elemon@apple.com

Stuart Cheshire
Apple Inc.
One Apple Park Way
Cupertino, California 95014
United States of America

Phone: +1 (408) 996-1010
Email: cheshire@apple.com