
Stream: Internet Engineering Task Force (IETF)
RFC: [8777](#)
Updates: [7450](#)
Category: Standards Track
Published: April 2020
ISSN: 2070-1721
Author: J. Holland
Akamai Technologies, Inc.

RFC 8777

DNS Reverse IP Automatic Multicast Tunneling (AMT) Discovery

Abstract

This document updates RFC 7450, "Automatic Multicast Tunneling" (or AMT), by modifying the relay discovery process. A new DNS resource record named AMTRELAY is defined for publishing AMT relays for source-specific multicast channels. The reverse IP DNS zone for a multicast sender's IP address is configured to use AMTRELAY resource records to advertise a set of AMT relays that can receive and forward multicast traffic from that sender over an AMT tunnel. Other extensions and clarifications to the relay discovery process are also defined.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8777>.

Copyright Notice

Copyright (c) 2020 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 (<https://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
 - 1.1. Background
 - 1.2. Terminology
 - 1.2.1. Relays and Gateways
 - 1.2.2. Definitions
 - 1.2.3. Requirements Language
2. Relay Discovery Overview
 - 2.1. Basic Mechanics
 - 2.2. Signaling and Discovery
 - 2.3. Example Deployments
 - 2.3.1. Example Receiving Networks
 - 2.3.2. Example Sending Networks
3. Relay Discovery Operation
 - 3.1. Optimal Relay Selection
 - 3.1.1. Overview
 - 3.1.2. Preference Ordering
 - 3.1.3. Connecting to Multiple Relays
 - 3.2. Happy Eyeballs
 - 3.2.1. Overview
 - 3.2.2. Algorithm Guidelines
 - 3.2.3. Connection Definition
 - 3.3. Guidelines for Restarting Discovery
 - 3.3.1. Overview

- 3.3.2. Updates to Restarting Events
 - 3.3.3. Tunnel Stability
 - 3.3.4. Traffic Health
 - 3.3.5. Relay Loaded or Shutting Down
 - 3.3.6. Relay Discovery Messages vs. Restarting Discovery
 - 3.3.7. Independent Discovery per Traffic Source
 - 3.4. DNS Configuration
 - 3.5. Waiting for DNS Resolution
 - 4. AMTRELAY Resource Record Definition
 - 4.1. AMTRELAY RRTYPE
 - 4.2. AMTRELAY RDATA Format
 - 4.2.1. RDATA Format - Precedence
 - 4.2.2. RDATA Format - Discovery Optional (D-bit)
 - 4.2.3. RDATA Format - Type
 - 4.2.4. RDATA Format - Relay
 - 4.3. AMTRELAY Record Presentation Format
 - 4.3.1. Representation of AMTRELAY RRs
 - 4.3.2. Examples
 - 5. IANA Considerations
 - 6. Security Considerations
 - 6.1. Use of AMT
 - 6.2. Record-Spoofing
 - 6.3. Congestion
 - 7. References
 - 7.1. Normative References
 - 7.2. Informative References
- Appendix A. Unknown RRTYPE Construction

[Acknowledgements](#)

[Author's Address](#)

1. Introduction

This document defines DNS Reverse IP AMT Discovery (DRIAD), a mechanism for AMT gateways to discover AMT relays that are capable of forwarding multicast traffic from a known source IP address.

AMT (Automatic Multicast Tunneling) is defined in [\[RFC7450\]](#) and provides a method to transport multicast traffic over a unicast tunnel in order to traverse network segments that are not multicast capable.

[Section 4.1.5](#) of [\[RFC7450\]](#) explains that the relay selection process for AMT is intended to be more flexible than the particular discovery method described in that document. That section further explains that the selection process might need to depend on the source of the multicast traffic in some deployments, since a relay must be able to receive multicast traffic from the desired source in order to forward it.

[Section 4.1.5](#) of [\[RFC7450\]](#) goes on to suggest DNS-based queries as a possible solution: DRIAD is DNS based. This solution also addresses the relay discovery issues in the "Disadvantages of this configuration" lists in [Sections 3.3](#) and [3.4](#) of [\[RFC8313\]](#).

The goal for DRIAD is to enable multicast connectivity between separate multicast-enabled networks without preconfiguring any peering arrangements between the networks when neither the sending nor the receiving network is connected to a multicast-enabled backbone.

This document extends the relay discovery procedure described in [Section 5.2.3.4](#) of [\[RFC7450\]](#).

1.1. Background

The reader is assumed to be familiar with the basic DNS concepts described in [\[RFC1034\]](#), [\[RFC1035\]](#), and the subsequent documents that update them, particularly [\[RFC2181\]](#).

The reader is also assumed to be familiar with the concepts and terminology regarding source-specific multicast as described in [\[RFC4607\]](#) and the use of Internet Group Management Protocol Version 3 (IGMPv3) [\[RFC3376\]](#) and Multicast Listener Discovery Version 2 (MLDv2) [\[RFC3810\]](#) for group management of source-specific multicast channels, as described in [\[RFC4604\]](#).

The reader should also be familiar with AMT, particularly the terminology listed in [Sections 3.2](#) and [3.3](#) of [\[RFC7450\]](#).

1.2. Terminology

1.2.1. Relays and Gateways

When reading this document, it's especially helpful to recall that once an AMT tunnel is established, the relay receives native multicast traffic and sends unicast tunnel-encapsulated traffic to the gateway. The gateway receives the tunnel-encapsulated packets, decapsulates them, and forwards them as native multicast packets, as illustrated in [Figure 1](#).

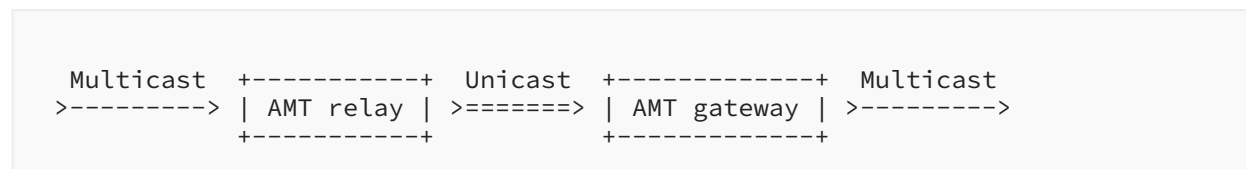


Figure 1: AMT Tunnel Illustration

1.2.2. Definitions

Term	Definition
(S,G)	A source-specific multicast channel, as described in [RFC4607] . A pair of IP addresses with a source host IP and destination group IP.
CMTS	Cable Modem Termination System
discovery broker	A broker or load balancer for AMT relay discovery, as mentioned in Section 4.2.1.1 of [RFC7450] .
downstream	Further from the source of traffic, as described in [RFC7450] .
FQDN	Fully Qualified Domain Name, as described in [RFC8499] .
gateway	An AMT gateway, as described in [RFC7450] .
L flag	The "Limit" flag described in Section 5.1.4.4 of [RFC7450] .
OLT	Optical Line Terminal
relay	An AMT relay, as described in [RFC7450] .
RPF	Reverse Path Forwarding, as described in [RFC5110] .
RR	A DNS Resource Record, as described in [RFC1034] .
RRType	A DNS Resource Record Type, as described in [RFC1034] .
SSM	Source-specific multicast, as described in [RFC4607] .

Term	Definition
upstream	Closer to the source of traffic, as described in [RFC7450].

Table 1: Definitions

1.2.3. Requirements Language

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.

2. Relay Discovery Overview

2.1. Basic Mechanics

The AMTRELAY resource record (RR) defined in this document is used to publish the IP address or domain name of a set of AMT relays or discovery brokers that can receive, encapsulate, and forward multicast traffic from a particular sender.

The sender is the owner of the RR and configures the zone so that it contains a set of RRs that provide the addresses or domain names of AMT relays (or discovery brokers that advertise relays) that can receive multicast IP traffic from that sender.

This enables AMT gateways in remote networks to discover an AMT relay that is capable of forwarding traffic from the sender. This, in turn, enables those AMT gateways to receive the multicast traffic tunneled over a unicast AMT tunnel from those relays and then pass the multicast packets into networks or applications that are using the gateway to subscribe to traffic from that sender.

This mechanism only works for source-specific multicast (SSM) channels. The source address of the (S,G) is reversed and used as an index into one of the reverse mapping trees (in-addr.arpa for IPv4, as described in Section 3.5 of [RFC1035], or ip6.arpa for IPv6, as described in Section 2.5 of [RFC3596]).

This mechanism should be treated as an extension of the AMT relay discovery procedure described in Section 5.2.3.4 of [RFC7450]. A gateway that supports this method of AMT relay discovery **SHOULD** use this method whenever it's performing the relay discovery procedure, the source IP addresses for desired (S,G)s are known to the gateway, and conditions match the requirements outlined in Section 3.1.

Some detailed example use cases are provided in Section 2.3, and other applicable example topologies appear in Sections 3.3, 3.4, and 3.5 of [RFC8313].

2.2. Signaling and Discovery

This section describes a typical example of the end-to-end process for signaling a receiver's join of an SSM channel that relies on an AMTRELAY RR.

The example in [Figure 2](#) contains two multicast-enabled networks that are both connected to the internet with non-multicast-capable links and which have no direct association with each other.

A content provider operates a sender, which is a source of multicast traffic inside a multicast-capable network.

An end user who is a customer of the content provider has a multicast-capable Internet Service Provider (ISP), which operates a receiving network that uses an AMT gateway. The AMT gateway is DRIAD capable.

The content provider provides the user with a receiving application that tries to subscribe to at least one (S,G). This receiving application could, for example, be a file transfer system using File Delivery over Unidirectional Transport (FLUTE) [[RFC6726](#)], a live video stream using RTP [[RFC3550](#)], or any other application that might subscribe to an SSM channel.

2. The join propagates with RPF through the receiver's multicast-enabled network with PIM [RFC7761] or another multicast routing mechanism until the AMT gateway receives a signal to join the (S,G).
3. The AMT gateway performs a reverse DNS lookup for the AMTRELAY RRType by sending an AMTRELAY RRType query for the reverse IP domain name for the sender's source IP address (the S from the (S,G)).

The DNS resolver for the AMT gateway uses ordinary DNS recursive resolution until it has the authoritative result that the content provider configured, which informs the AMT gateway that the relay address is 2001:db8::c:f.

4. The AMT gateway performs AMT handshakes with the AMT relay as described in Section 4 of [RFC7450], then forwards a membership report to the relay, indicating subscription to the (S,G).
5. The relay propagates the join through its network toward the sender and then forwards the appropriate AMT-encapsulated traffic to the gateway, which decapsulates and forwards it as a native multicast through its downstream network to the end user.

In the case of an IPv4 (S,G), the only difference in the AMT relay discovery process is the use of the in-addr.arpa reverse IP domain name, as described in Section 3.5 of [RFC1035], instead of the in6.arpa domain name. For example, if the (S,G) is (198.51.100.12, 232.252.0.2), the reverse IP FQDN for the AMTRELAY query would be "12.100.51.198.in-addr.arpa."

Note that the address family of the AMT tunnel is independent of the address family for the multicast traffic.

2.3. Example Deployments

2.3.1. Example Receiving Networks

2.3.1.1. Internet Service Provider

One example of a receiving network is an Internet Service Provider (ISP) that offers multicast ingest services to its subscribers, illustrated in Figure 3.

In the example network below, subscribers can join (S,G)s with MLDv2 or IGMPv3 as described in [RFC4604], and the AMT gateway in this ISP can receive and forward multicast traffic from one of the example sending networks in Section 2.3.2 by discovering the appropriate AMT relays with a DNS lookup for the AMTRELAY RR with the reverse IP of the source in the (S,G).

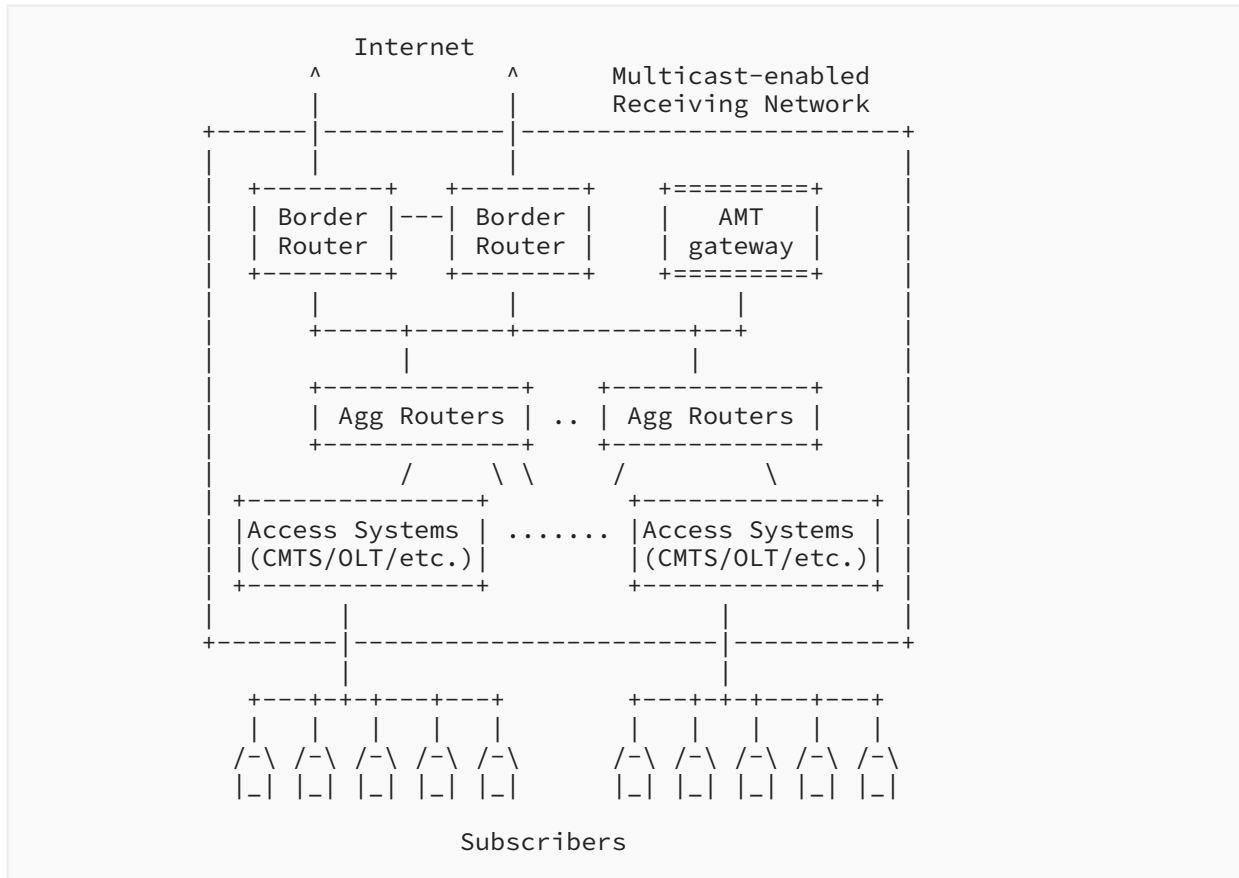


Figure 3: Receiving ISP Example

2.3.1.2. Small Office

Another example receiving network is a small branch office that regularly accesses some multicast content, illustrated in Figure 4.

This office has desktop devices that need to receive some multicast traffic, so an AMT gateway runs on a LAN with these devices to pull traffic in through a non-multicast next hop.

The office also hosts some mobile devices that have AMT gateway instances embedded inside apps in order to receive multicast traffic over their non-multicast wireless LAN. (Note that the "Legacy Router" is a simplification that's meant to describe a variety of possible conditions; for example, it could be a device providing a split-tunnel VPN as described in [RFC7359], deliberately excluding multicast traffic for a VPN tunnel, rather than a device that is incapable of multicast forwarding.)

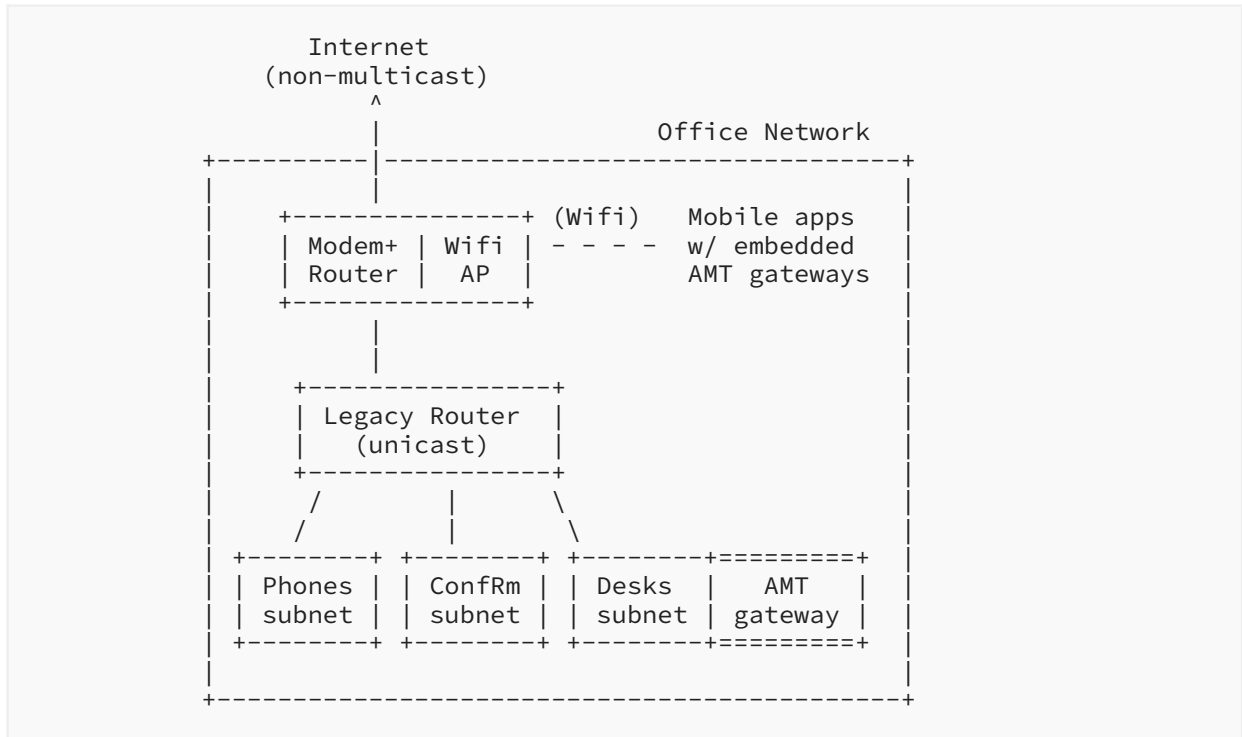


Figure 4: Small Office (No Multicast Up)

By adding an AMT relay to this office network as in [Figure 5](#), it's possible to make use of multicast services from the example multicast-capable ISP in [Section 2.3.1.1](#).

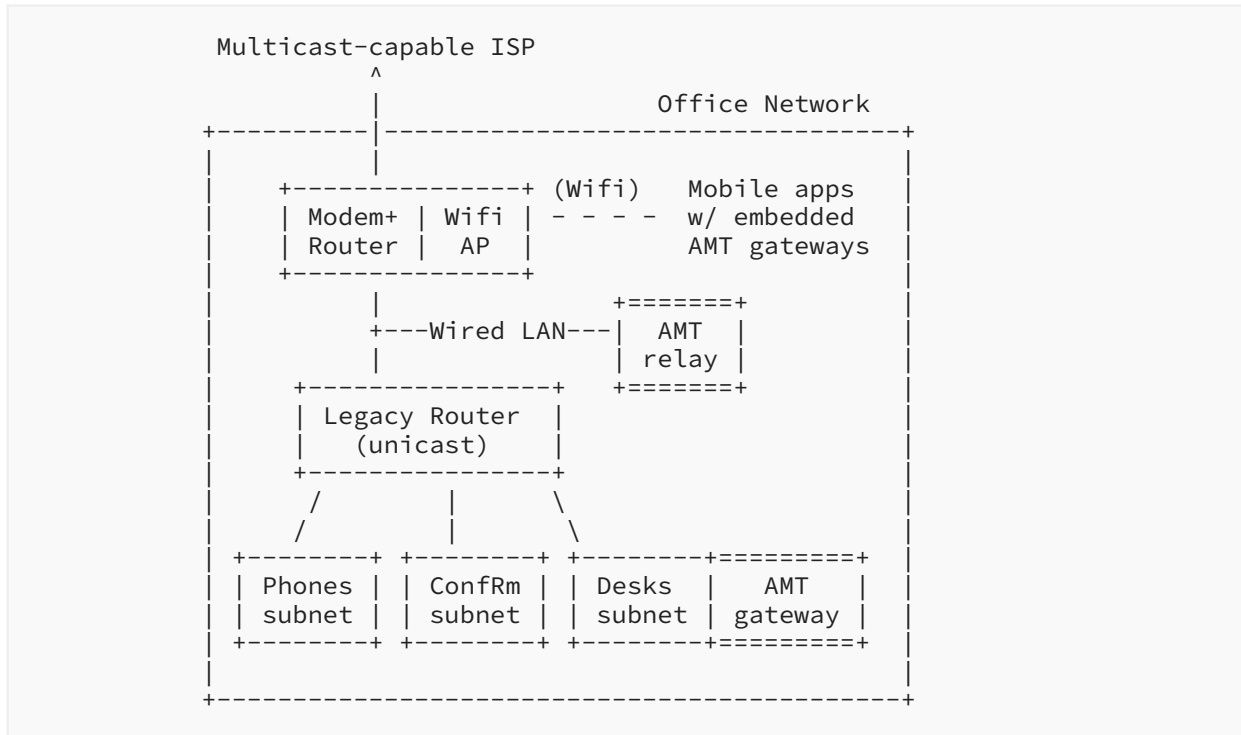


Figure 5: Small Office Example

When multicast-capable networks are chained like this, with a network like the one in [Figure 5](#) receiving Internet services from a multicast-capable network like the one in [Figure 3](#), it's important for AMT gateways to reach the more local AMT relay in order to avoid accidentally tunneling multicast traffic from a more distant AMT relay with unicast and failing to utilize the multicast transport capabilities of the network in [Figure 3](#).

2.3.2. Example Sending Networks

2.3.2.1. Sender-Controlled Relays

When a sender network is also operating AMT relays to distribute multicast traffic, as in [Figure 6](#), each address could appear as an AMTRELAY RR for the reverse IP of the sender. Alternately, one or more domain names could appear in AMTRELAY RRs, and the AMT relay addresses can be discovered by finding A or AAAA records from those domain names.




Figure 6: Small Office Example

2.3.2.2. Provider-Controlled Relays

When an ISP offers a service to transmit outbound multicast traffic through a forwarding network, it might also offer AMT relays in order to reach receivers without multicast connectivity to the forwarding network, as in [Figure 7](#). In this case, it's recommended that the ISP also provide at least one domain name for the AMT relays for use with the AMTRELAY RR.

When the sender wishes to use the relays provided by the ISP for forwarding multicast traffic, an AMTRELAY RR should be configured to use the domain name provided by the ISP to allow for address reassignment of the relays without forcing the sender to reconfigure the corresponding AMTRELAY RRs.




Figure 7: Sending ISP Example

3. Relay Discovery Operation

3.1. Optimal Relay Selection

3.1.1. Overview

The reverse source IP DNS query of an AMTRELAY RR is a good way for a gateway to discover a relay that is known to the sender.

However, it is *not* necessarily a good way to discover the best relay for that gateway to use, because the RR will only provide information about relays known to the source.

If there is an upstream relay in a network that is topologically closer to the gateway and is able to receive and forward multicast traffic from the sender, that relay is better for the gateway to use since more of the network path uses native multicast, allowing more chances for packet replication. But since that relay is not known to the sender, it won't be advertised in the sender's reverse IP DNS record. An example network that illustrates this scenario is outlined in [Figure 5](#) from [Section 2.3.1.2](#).

It's only appropriate for an AMT gateway to discover an AMT relay by querying an AMTRELAY RR owned by a sender when all of these conditions are met:

1. The gateway needs to propagate a join of an (S,G) over AMT because in the gateway's network, no RPF next hop toward the source can propagate a native multicast join of the (S,G);

2. The gateway is not already connected to a relay that forwards multicast traffic from the source of the (S,G);
3. The gateway is not configured to use a particular IP address for AMT discovery, or a relay discovered with that IP is not able to forward traffic from the source of the (S,G);
4. The gateway is not able to find an upstream AMT relay with DNS-based Service Discovery (DNS-SD) [RFC6763] using "_amt._udp" as the Service section of the queries, or a relay discovered this way is not able to forward traffic from the source of the (S,G) (as described in Section 3.3.4.1 and 3.3.5); and
5. The gateway is not able to find an upstream AMT relay with the well-known anycast addresses from Section 7 of [RFC7450].

When all of the above conditions are met, the gateway has no path within its local network that can receive multicast traffic from the source IP of the (S,G).

In this situation, the best way to find a relay that can forward the required traffic is to use information that comes from the operator of the sender. When the sender has configured an AMTRELAY RR, gateways can use the DRIAD mechanism defined in this document to discover the relay information provided by the sender.

Note that the above conditions are designed to prefer the use of a local AMT relay if one can be discovered. However, note also that the network upstream of the locally discovered relay would still need to receive traffic from the sender of the (S,G) in order to forward it. Therefore, unless the upstream network contains the sender or has a multicast-capable peering with a network that can forward traffic from the sender, the upstream network might still use AMT to ingest the traffic from a network that can receive traffic from the sender. If this is the case, the upstream AMT gateway could still rely on the AMTRELAY RR provided by the sender, even though the AMTRELAY RR is not directly used by gateways topologically closer to the receivers. For a concrete example of such a situation, consider the network in Figure 5 connected as one of the customers to the network in Figure 3.

3.1.2. Preference Ordering

This section defines a preference ordering for relay addresses during the relay discovery process. Gateways are encouraged to implement a Happy Eyeballs [RFC8305] algorithm to try candidate relays concurrently (see Section 3.2), but even gateways that do not implement a Happy Eyeballs algorithm **SHOULD** use this ordering, except as noted.

When establishing an AMT tunnel to forward multicast data, it's very important for the discovery process to prioritize network topology considerations ahead of address selection considerations in order to gain the packet replication benefits from using multicast instead of unicast tunneling in the multicast-capable portions of the network path.

The intent of the advice and requirements in this section is to describe how a gateway should make use of the concurrency provided by a Happy Eyeballs algorithm to reduce the join latency while still prioritizing network efficiency considerations over address selection considerations.

Section 4 of [RFC8305] requires a Happy Eyeballs algorithm to sort the addresses with the Destination Address Selection defined in Section 6 of [RFC6724], but for the above reasons, that requirement is superseded in the AMT discovery use case by the following considerations:

- Prefer Local Relays

Figure 5 and Section 2.3.1.2 provide a motivating example to prefer DNS-SD [RFC6763] for discovery strictly ahead of using the AMTRELAY RR controlled by the sender for AMT discovery.

For this reason, it's **RECOMMENDED** that AMT gateways by default perform service discovery using DNS Service Discovery (DNS-SD) [RFC6763] for `_amt._udp.<domain>` (with `<domain>` chosen as described in Section 11 of [RFC6763]) and use the AMT relays discovered that way in preference to AMT relays discoverable via the mechanism defined in this document (DRIAD).

- Prefer Relays Managed by the Containing Network

When no local relay is discoverable with DNS-SD, it still may be the case that a relay local to the receiver is operated by the network providing transit services to the receiver.

In this case, when the network cannot make the relay discoverable via DNS-SD, the network **SHOULD** use the well-known anycast addresses from Section 7 of [RFC7450] to route discovery traffic to the relay most appropriate to the receiver's gateway.

Accordingly, the gateway **SHOULD** by default discover a relay with the well-known AMT anycast addresses as the next-best option after DNS-SD when searching for a local relay.

- Let Sender Manage Relay Provisioning

A related motivating example is provided by considering a sender whose traffic can be forwarded by relays in a sender-controlled network like Figure 6 in Section 2.3.2.1 and by relays in a provider-controlled network like Figure 7 in Section 2.3.2.2, with different cost and scalability profiles for the different options.

In this example about the sending-side network, the precedence field described in Section 4.2.1 is a critical method of control so that senders can provide the appropriate guidance to gateways during the discovery process in order to manage load and failover scenarios in a manner that operates well with the sender's provisioning strategy for horizontal scaling of AMT relays.

Therefore, after DNS-SD, the precedence from the RR **MUST** be used for sorting preference ahead of the Destination Address Selection ordering from Section 6 of [RFC6724] so that only relay IPs with the same precedence are directly compared according to the Destination Address Selection ordering.

Accordingly, AMT gateways **SHOULD** by default prefer relays in this order:

1. DNS-SD

2. Anycast addresses from [Section 7](#) of [\[RFC7450\]](#)
3. DRIAD

This default behavior **MAY** be overridden by administrative configuration where other behavior is more appropriate for the gateway within its network.

Among relay addresses that have an equivalent preference as described above, a Happy Eyeballs algorithm for AMT **SHOULD** use the Destination Address Selection defined in [Section 6](#) of [\[RFC6724\]](#).

Among relay addresses that still have an equivalent preference after the above orderings, a gateway **SHOULD** make a non-deterministic choice (such as a pseudorandom selection) for relay preference ordering in order to support load balancing by DNS configurations that provide many relay options.

The gateway **MAY** introduce a bias in the non-deterministic choice according to information that indicates expected benefits from selecting some relays in preference to others. Details about the structure and collection of this information are out of scope for this document but could, for example, be obtained by out-of-band methods or from a historical record about network topology, timing information, or the response to a probing mechanism. A gateway in possession of such information **MAY** use it to prefer topologically closer relays.

Within the above constraints, gateways **MAY** make use of other considerations from [Section 4](#) of [\[RFC8305\]](#), such as the address family interleaving strategies, to produce a final ordering of candidate relay addresses.

Note also that certain relay addresses might be excluded from consideration by the hold-down timers described in [Section 3.3.4.1](#) or [3.3.5](#). These relays constitute "unusable destinations" under Rule 1 of the Destination Address Selection and are also not part of the superseding considerations described above.

The discovery and connection process for the relay addresses in the above described ordering **MAY** operate in parallel, subject to delays prescribed by the Happy Eyeballs requirements described in [Section 5](#) of [\[RFC8305\]](#) for successively launched concurrent connection attempts.

3.1.3. Connecting to Multiple Relays

In some deployments, it may be useful for a gateway to connect to multiple upstream relays and subscribe to the same traffic in order to support an active/active failover model. A gateway **SHOULD NOT** be configured to do so without guaranteeing that adequate bandwidth is available.

A gateway configured to do this **SHOULD** still use the same preference-ordering logic from [Section 3.1.2](#) for each connection. (Note that this ordering allows for overriding by explicit administrative configuration where required.)

3.2. Happy Eyeballs

3.2.1. Overview

Often, multiple choices of relay will exist for a gateway using DRIAD for relay discovery. Happy Eyeballs [RFC8305] provides a widely deployed and generalizable strategy for probing multiple possible connections in parallel. Therefore, it is **RECOMMENDED** that DRIAD-capable gateways implement a Happy Eyeballs algorithm to support fast discovery of the most preferred available relay by probing multiple relays concurrently.

The parallel discovery logic of a Happy Eyeballs algorithm serves to reduce join latency for the initial join of an SSM channel. This section and the preference ordering of relays defined in [Section 3.1.2](#) together provide guidance on use of a Happy Eyeballs algorithm for the case of establishing AMT connections.

Note that according to the definition in [Section 3.2.3](#) of this document, establishing the connection occurs before sending a membership report. As described in [Section 5](#) of [RFC8305], only one of the successful connections will be used, and the others are all canceled or ignored. In the context of an AMT connection, this means the gateway will send the membership reports that subscribe to traffic only for the chosen connection after the Happy Eyeballs algorithm resolves.

3.2.2. Algorithm Guidelines

During the "Initiation of asynchronous DNS queries" phase described in [Section 3](#) of [RFC8305], a gateway attempts to resolve the domain names listed in [Section 3.1](#). This consists of resolving the SRV queries for DNS-SD domains for the AMT service, as well as the AMTRELAY query for the reverse IP domain defined in this document.

Each of the SRV and AMTRELAY responses might contain:

- one or more IP addresses (as with type 1 or type 2 AMTRELAY responses or when the SRV Additional Data section of the SRV response contains the address records for the target, as urged by [RFC2782]), or
- only domain names (as with type 3 responses from [Section 4.2.3](#) or an SRV response without an additional data section).

When present, IP addresses in the initial response provide resolved destination address candidates for the "Sorting of resolved destination addresses" phase described in [Section 4](#) of [RFC8305]), whereas domain names without IP addresses in the initial response result in another set of queries for AAAA and A records, whose responses provide the candidate resolved destination addresses.

Since the SRV or AMTRELAY responses don't have a bound on the count of queries that might be generated aside from the bounds imposed by the DNS resolver, it's important for the gateway to provide a rate limit on the DNS queries. The DNS query functionality is expected to follow ordinary standards and best practices for DNS clients. A gateway **MAY** use an existing DNS client implementation that does so and **MAY** rely on that client's rate-limiting logic to avoid issuing

excessive queries. Otherwise, a gateway **MUST** provide a rate limit for the DNS queries, and its default settings **SHOULD NOT** permit more than 10 queries for any 100-millisecond period (though this **MAY** be overridable by the administrative configuration).

As the resolved IP addresses arrive, the Happy Eyeballs algorithm sorts them according to the requirements and recommendations given in [Section 3.1.2](#) and attempts connections with the corresponding relays under the algorithm restrictions and guidelines given in [\[RFC8305\]](#) for the "Establishment of one connection, which cancels all other attempts" phase. As described in [Section 3](#) of [\[RFC8305\]](#), DNS resolution is treated as asynchronous, and connection initiation does not wait for lagging DNS responses.

3.2.3. Connection Definition

[Section 5](#) of [\[RFC8305\]](#) non-normatively describes a successful connection attempt as "generally when the TCP handshake completes".

There is no normative definition of a connection in the AMT specification [\[RFC7450\]](#), and there is no TCP connection involved in an AMT tunnel.

However, the concept of an AMT connection in the context of a Happy Eyeballs algorithm is a useful one, and so this section provides the following normative definition:

- An AMT connection is established successfully when the gateway receives from a newly discovered relay a valid Membership Query message ([Section 5.1.4](#) of [\[RFC7450\]](#)) that does not have the L flag set.

See [Section 3.3.5](#) of this document for further information about the relevance of the L flag to the establishment of a Happy Eyeballs connection. See [Section 3.3.4](#) for an overview of how to respond if the connection does not provide multicast connectivity to the source.

To "cancel" this kind of AMT connection for the Happy Eyeballs algorithm, a gateway that has not sent a membership report with a subscription would simply stop sending AMT packets for that connection. A gateway only sends a membership report to a connection it has chosen as the most preferred available connection.

3.3. Guidelines for Restarting Discovery

3.3.1. Overview

It's expected that gateways deployed in different environments will use a variety of heuristics to decide when it's appropriate to restart the relay discovery process in order to meet different performance goals (for example, to fulfill different kinds of service level agreements).

In general, restarting the discovery process is always safe for the gateway and relay during any of the events listed in this section but may cause a disruption in the forwarded traffic if the discovery process results in choosing a different relay because this changes the RPF forwarding tree for the multicast traffic upstream of the gateway. This is likely to result in some dropped or duplicated packets from channels actively being tunneled from the old relay to the gateway.

The degree of impact on the traffic from choosing a different relay may depend on network conditions between the gateway and the new relay, as well as the network conditions and topology between the sender and the new relay, as this may cause the relay to propagate a new RPF join toward the sender.

Balancing the expected impact on the tunneled traffic against likely or observed problems with an existing connection to the relay is the goal of the heuristics that gateways use to determine when to restart the discovery process.

The non-normative advice in this section should be treated as guidelines to operators and implementors working with AMT systems that can use DRIAD as part of the relay discovery process.

3.3.2. Updates to Restarting Events

[Section 5.2.3.4.1](#) of [\[RFC7450\]](#) lists several events that may cause a gateway to start or restart the discovery procedure.

This document provides some updates and recommendations regarding the handling of these and similar events. The first five events are copied here and numbered for easier reference, and the remaining four events are newly added for consideration in this document:

1. When a gateway pseudo-interface is started (enabled).
2. When the gateway wishes to report a group subscription when none currently exists.
3. Before sending the next Request message in a membership update cycle.
4. After the gateway fails to receive a response to a Request message.
5. After the gateway receives a Membership Query message with the L flag set to 1.
6. When the gateway wishes to report an (S,G) subscription with a source address that does not currently have other group subscriptions.
7. When there is a network change detected; for example, when a gateway is operating inside an end user device or application and the device joins a different network or when the domain portion of a DNS-SD domain name changes in response to a DHCP message or administrative configuration.
8. When substantial loss, persistent congestion, or network overload is detected in the stream of AMT packets from a relay.
9. When the gateway has reported one or more (S,G) subscriptions but no traffic is received from the source for some timeout (see [Section 3.3.4.1](#)).

This list is not exhaustive, nor are any of the listed events strictly required to always force a restart of the discovery process.

Note that during event #1, a gateway may use DNS-SD but does not have sufficient information to use DRIAD, since no source is known.

3.3.3. Tunnel Stability

In general, subscribers to active traffic flows that are being forwarded by an AMT gateway are less likely to experience a degradation in service (for example, from missing or duplicated packets) when the gateway continues using the same relay as long as the relay is not overloaded and the network conditions remain stable.

Therefore, gateways **SHOULD** avoid performing a full restart of the discovery process during routine cases of event #3 (sending a new Request message), since it occurs frequently in normal operation.

However, see Sections 3.3.4, 3.3.6, and 3.3.4.3 for more information about exceptional cases when it may be appropriate to use event #3.

3.3.4. Traffic Health

3.3.4.1. Absence of Traffic

If a gateway indicates one or more (S,G) subscriptions in a Membership Update message but no traffic for any of the (S,G)s is received in a reasonable time, it's appropriate for the gateway to restart the discovery process.

If the gateway restarts the discovery process multiple times consecutively for this reason, the timeout period **SHOULD** be adjusted to provide a random exponential back-off.

The **RECOMMENDED** timeout is a random value in the range [initial_timeout, MIN(initial_timeout * 2^retry_count, maximum_timeout)], with a **RECOMMENDED** initial_timeout of 4 seconds and a **RECOMMENDED** maximum_timeout of 120 seconds (which is the recommended minimum NAT mapping timeout described in [RFC4787]).

Note that the recommended initial_timeout is larger than the initial timeout recommended in the similar algorithm from Section 5.2.3.4.3 of [RFC7450]. This is to provide time for RPF Join propagation in the sending network. Although the timeout values may be administratively adjusted to support performance requirements, operators are advised to consider the possibility of join propagation delays between the sender and the relay when choosing an appropriate timeout value.

Gateways restarting the discovery process because of an absence of traffic **MUST** use a hold-down timer that removes this relay from consideration during subsequent rounds of discovery while active. The hold-down **SHOULD** last for no less than 3 minutes and no more than 10 minutes.

3.3.4.2. Loss and Congestion

In some gateway deployments, it is also feasible to monitor the health of traffic flows through the gateway -- for example, by detecting the rate of packet loss by communicating out of band with receivers or monitoring the packets of known protocols with sequence numbers. Where feasible, it's encouraged for gateways to use such traffic health information to trigger a restart of the discovery process during event #3 (before sending a new Request message).

However, if a transient network event that affects the tunneled multicast stream -- as opposed to an event that affects the tunnel connection between the relay and gateway -- occurs, poor health detection could be triggered for many gateways simultaneously. In this situation, adding a random delay to avoid synchronized rediscovery by many gateways is recommended.

The span of the random portion of the delay should be no less than 10 seconds by default but may be administratively configured to support different performance requirements.

3.3.4.3. Ancient Discovery Information

In most cases, a gateway actively receiving healthy traffic from a relay that has not indicated load with the L flag should prefer to remain connected to the same relay, as described in [Section 3.3.3](#).

However, a relay that appears healthy but has been forwarding traffic for days or weeks may have an increased chance of becoming unstable. Gateways may benefit from restarting the discovery process during event #3 (before sending a Request message) after the expiration of a long-term timeout on the order of multiple hours or even days in some deployments.

It may be beneficial for such timers to consider the amount of traffic currently being forwarded and to give a higher probability of restarting discovery during periods with an unusually low data rate to reduce the impact on active traffic while still avoiding relying on the results of a very old discovery.

Other issues may also be worth considering as part of this heuristic; for example, if the DNS expiry time of the record that was used to discover the current relay has not passed, the long-term timer might be restarted without restarting the discovery process.

3.3.5. Relay Loaded or Shutting Down

The L flag (see [Section 5.1.4.4](#) of [\[RFC7450\]](#)) is the preferred mechanism for a relay to signal overloading or a graceful shutdown to gateways.

A gateway that supports handling of the L flag should generally restart the discovery process when it processes a Membership Query packet with the L flag set. If an L flag is received while a concurrent Happy Eyeballs discovery process is underway for multiple candidate relays ([Section 3.2](#)), the relay sending the L flag **SHOULD NOT** be considered for the relay selection.

It is also **RECOMMENDED** that gateways avoid choosing a relay that has recently sent an L flag, with approximately a 10-minute hold-down. Gateways **SHOULD** treat this hold-down timer in the same way as the hold-down in [Section 3.3.4.1](#) so that the relay is removed from consideration for subsequent short-term rounds of discovery.

3.3.6. Relay Discovery Messages vs. Restarting Discovery

All AMT relays are required by [\[RFC7450\]](#) to support handling of Relay Discovery messages (e.g., in [Section 5.3.3.2](#) of [\[RFC7450\]](#)).

So a gateway with an existing connection to a relay can send a Relay Discovery message to the unicast address of that AMT relay. Under stable conditions with an unloaded relay, it's expected that the relay will return its own unicast address in the Relay Advertisement in response to such

a Relay Discovery message. Since this will not result in the gateway changing to another relay unless the relay directs the gateway away, this is a reasonable exception to the advice against handling event #3 described in [Section 3.3.3](#).

This behavior is discouraged for gateways that do support the L flag to avoid sending unnecessary packets over the network.

However, gateways that do not support the L flag may be able to avoid a disruption in the forwarded traffic by sending such Relay Discovery messages regularly. When a relay is under load or has started a graceful shutdown, it may respond with a different relay address, which the gateway can use to connect to a different relay. This kind of coordinated handoff will likely result in a smaller disruption to the traffic than if the relay simply stops responding to Request messages and stops forwarding traffic.

This style of Relay Discovery message (one sent to the unicast address of a relay that's already forwarding traffic to this gateway) **SHOULD NOT** be considered a full restart of the relay discovery process. It is **RECOMMENDED** that gateways support the L flag, but for gateways that do not support the L flag, sending this message during event #3 may help mitigate service degradation when relays become unstable.

3.3.7. Independent Discovery per Traffic Source

Relays discovered via the AMTRELAY RR are source-specific relay addresses and may use different pseudo-interfaces from each other and from relays discovered via DNS-SD or via a non-source-specific address, as described in [Section 4.1.2.1](#) of [\[RFC7450\]](#).

Restarting the discovery process for one pseudo-interface does not require restarting the discovery process for other pseudo-interfaces. Gateway heuristics about restarting the discovery process should operate independently for different tunnels to relays when responding to events that are specific to the different tunnels.

3.4. DNS Configuration

Often, an AMT gateway will only have access to the source and group IP addresses of the desired traffic and will not know any other name for the source of the traffic. Because of this, typically, the best way of looking up AMTRELAY RRs will be by using the source IP address as an index into one of the reverse mapping trees (in-addr.arpa for IPv4, as described in [Section 3.5](#) of [\[RFC1035\]](#), or ip6.arpa for IPv6, as described in [Section 2.5](#) of [\[RFC3596\]](#)).

Therefore, it is **RECOMMENDED** that AMTRELAY RRs be added to reverse IP zones as appropriate. AMTRELAY records **MAY** also appear in other zones, since this may be necessary to perform delegation from the reverse zones (see, for example, [Section 5.2](#) of [\[RFC2317\]](#)), but the use case enabled by this document requires a reverse IP mapping for the source from an (S,G) in order to be useful to most AMT gateways. This document does not define semantics for the use of AMTRELAY records obtained in a way other than by a reverse IP lookup.

When performing the AMTRELAY RR lookup, any CNAMEs or DNAMEs found **MUST** be followed. This is necessary to support zone delegation. Some examples outlining this need are described in [\[RFC2317\]](#).

See Sections 4 and 4.3 for a detailed explanation of the contents of a DNS zone file.

3.5. Waiting for DNS Resolution

DNS query functionality is expected to follow ordinary standards and best practices for DNS clients. A gateway **MAY** use an existing DNS client implementation that does so and **MAY** rely on that client's retry logic to determine the timeouts between retries.

Otherwise, a gateway **MAY** resend a DNS query if it does not receive an appropriate DNS response within some timeout period. If the gateway retries multiple times, the timeout period **SHOULD** be adjusted to provide a random exponential back-off.

As with the waiting process for the Relay Advertisement message from [Section 5.2.3.4.3 of \[RFC7450\]](#), the **RECOMMENDED** timeout is a random value in the range [initial_timeout, MIN(initial_timeout * 2^retry_count, maximum_timeout)], with a **RECOMMENDED** initial_timeout of 1 second and a **RECOMMENDED** maximum_timeout of 120 seconds.

4. AMTRELAY Resource Record Definition

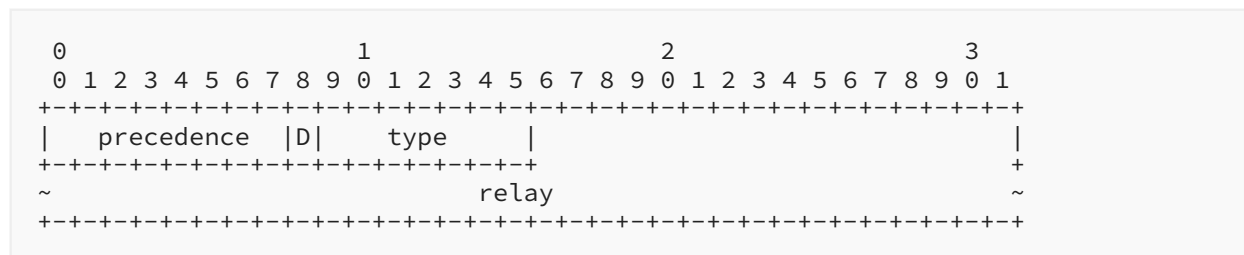
4.1. AMTRELAY RRType

The AMTRELAY RRType has the mnemonic AMTRELAY and type code 260 (decimal).

The AMTRELAY RR is class independent.

4.2. AMTRELAY RData Format

The AMTRELAY RData consists of an 8-bit precedence field, a 1-bit "Discovery Optional" field, a 7-bit type field, and a variable length relay field.



4.2.1. RData Format - Precedence

This is an 8-bit precedence for this record. It is interpreted in the same way as the PREFERENCE field described in [Section 3.3.9 of \[RFC1035\]](#).

Relays listed in AMTRELAY records with a lower value for precedence are to be attempted first.

4.2.2. RData Format - Discovery Optional (D-bit)

The D-bit is a "Discovery Optional" flag.

If the D-bit is set to 0, a gateway using this RR **MUST** perform AMT relay discovery as described in [Section 4.2.1.1](#) of [\[RFC7450\]](#) rather than directly sending an AMT Request message to the relay.

That is, the gateway **MUST** receive an AMT Relay Advertisement message ([Section 5.1.2](#) of [\[RFC7450\]](#)) for an address before sending an AMT Request message ([Section 5.1.3](#) of [\[RFC7450\]](#)) to that address. Before receiving the Relay Advertisement message, this record has only indicated that the address can be used for AMT relay discovery, not for a Request message. This is necessary for devices that are not fully functional AMT relays but rather load balancers or brokers, as mentioned in [Section 4.2.1.1](#) of [\[RFC7450\]](#).

If the D-bit is set to 1, the gateway **MAY** send an AMT Request message directly to the discovered relay address without first sending an AMT Discovery message.

This bit should be set according to advice from the AMT relay operator. The D-bit **MUST** be set to zero when no information is available from the AMT relay operator about its suitability.

4.2.3. RData Format - Type

The type field indicates the format of the information that is stored in the relay field.

The following values are defined:

- type = 0: The relay field is empty (0 bytes).
- type = 1: The relay field contains a 4-octet IPv4 address.
- type = 2: The relay field contains a 16-octet IPv6 address.
- type = 3: The relay field contains a wire-encoded domain name. The wire-encoded format is self-describing, so the length is implicit. The domain name **MUST NOT** be compressed (see [Section 3.3](#) of [\[RFC1035\]](#) and [Section 4](#) of [\[RFC3597\]](#)).

RRs with an undefined value in the Type field **SHOULD NOT** be considered by receiving gateways for AMT relay discovery.

4.2.4. RData Format - Relay

The relay field is the address or domain name of the AMT relay. It is formatted according to the type field.

When the type field is 0, the length of the relay field is 0, and it indicates that no AMT relay should be used for multicast traffic from this source.

When the type field is 1, the length of the relay field is 4 octets, and a 32-bit IPv4 address is present. This is an IPv4 address as described in [Section 3.4.1](#) of [\[RFC1035\]](#). This is a 32-bit number in network byte order.

When the type field is 2, the length of the relay field is 16 octets, and a 128-bit IPv6 address is present. This is an IPv6 address as described in [Section 2.2](#) of [RFC3596]. This is a 128-bit number in network byte order.

When the type field is 3, the relay field is a normal wire-encoded domain name, as described in [Section 3.3](#) of [RFC1035]. For the reasons given in [Section 4](#) of [RFC3597], compression **MUST NOT** be used.

For a type 3 record, the D-bit and preference fields carry over to all A or AAAA records for the domain name. There is no difference in the result of the discovery process when it's obtained by type 1 or type 2 AMTRELAY records with identical D-bit and preference fields vs. when the result is obtained by a type 3 AMTRELAY record that resolves to the same set of IPv4 and IPv6 addresses via A and AAAA lookups.

4.3. AMTRELAY Record Presentation Format

4.3.1. Representation of AMTRELAY RRs

AMTRELAY RRs may appear in a zone data master file. The precedence, D-bit, relay type, and relay fields are **REQUIRED**.

If the relay type field is 0, the relay field **MUST** be ".".

The presentation for the record is as follows:

```
IN AMTRELAY precedence D-bit type relay
```

4.3.2. Examples

In a DNS authoritative nameserver that understands the AMTRELAY type, the zone might contain a set of entries like this:

```
$ORIGIN 100.51.198.in-addr.arpa.  
12      IN AMTRELAY 10 0 1 203.0.113.15  
12      IN AMTRELAY 10 0 2 2001:db8::15  
12      IN AMTRELAY 128 1 3 amtrelays.example.com.
```

This configuration advertises an IPv4 discovery address, an IPv6 discovery address, and a domain name for AMT relays that can receive traffic from the source 198.51.100.12. The IPv4 and IPv6 addresses are configured with a D-bit of 0 (meaning discovery is mandatory, as described in [Section 4.2.2](#)) and a precedence 10 (meaning they're preferred ahead of the last entry, which has precedence 128).

For zone files in name servers that don't support the AMTRELAY RRType natively, it's possible to use the format for unknown RR types, as described in [RFC3597]. This approach would replace the AMTRELAY entries in the example above with the entries below:

```

10  IN TYPE260  \# (
      6  ; length
      0a ; precedence=10
      01 ; D=0, relay type=1, an IPv4 address
      cb00710f ) ; 203.0.113.15
10  IN TYPE260  \# (
      18 ; length
      0a ; precedence=10
      02 ; D=0, relay type=2, an IPv6 address
      20010db800000000000000000000000f ) ; 2001:db8::15
10  IN TYPE260  \# (
      24 ; length
      80 ; precedence=128
      83 ; D=1, relay type=3, a wire-encoded domain name
      09616d74726556c617973076578616d706c6503636f6d ) ; domain name

```

See [Appendix A](#) for more details.

5. IANA Considerations

This document updates the DNS "Resource Record (RR) TYPES" registry by assigning type 260 to the AMTRELAY record.

This document creates a new registry named "AMTRELAY Resource Record Parameters" with a subregistry for the "Relay Type Field". The initial values in the subregistry are:

Value	Description
0	No relay is present
1	A 4-byte IPv4 address is present
2	A 16-byte IPv6 address is present
3	A wire-encoded domain name is present
4-255	Unassigned

Table 2: Initial Contents of the "Relay Type Field" Registry

Values 0, 1, 2, and 3 are further explained in Sections [4.2.3](#) and [4.2.4](#). Relay type numbers 4 through 255 can be assigned with a policy of Specification Required (as described in [RFC8126]).

6. Security Considerations

6.1. Use of AMT

This document defines a mechanism that enables a more widespread and automated use of AMT, even without access to a multicast backbone. Operators of networks and applications that include a DRIAD-capable AMT gateway are advised to carefully consider the security considerations in [Section 6](#) of [\[RFC7450\]](#).

AMT gateway operators also are encouraged to take appropriate steps to ensure the integrity of the data received via AMT, for example, by the opportunistic use of IPsec [\[RFC4301\]](#) to secure traffic received from AMT relays when IPSECKEY records [\[RFC4025\]](#) are available or when a trust relationship with the AMT relays can be otherwise established and secured.

Note that AMT does not itself provide any integrity protection for Multicast Data packets ([Section 5.1.6](#) of [\[RFC7450\]](#)), so absent protections like those mentioned above, even an off-path attacker who discovers the gateway IP, the relay IP, and the relay source port for an active AMT connection can inject multicast data packets for a joined (S,G) into the data stream if he can get data packets delivered to the gateway IP that spoof the relay as the source.

6.2. Record-Spoofing

The AMTRELAY resource record contains information that **SHOULD** be communicated to the DNS client without being modified. The method used to ensure the result was unmodified is up to the client.

There must be a trust relationship between the end consumer of this resource record and the DNS server. This relationship may be end-to-end DNSSEC validation or a secure connection to a trusted DNS server that provides end-to-end safety to prevent record-spoofing of the response from the trusted server. The connection to the trusted server can use any secure channel, such as with a TSIG [\[RFC2845\]](#) or SIG(0) [\[RFC2931\]](#) channel, a secure local channel on the host, DNS over TLS [\[RFC7858\]](#), DNS over HTTPS [\[RFC8484\]](#), or some other mechanism that provides authentication of the RR.

If an AMT gateway accepts a maliciously crafted AMTRELAY record, the result could be a Denial of Service or receivers processing multicast traffic from a source under the attacker's control.

6.3. Congestion

Multicast traffic, particularly interdomain multicast traffic, carries some congestion risks, as described in [Section 4](#) of [\[RFC8085\]](#).

Application implementors and network operators that use AMT gateways are advised to take precautions, including monitoring of application traffic behavior, traffic authentication at ingest, rate-limiting of multicast traffic, and the use of circuit-breaker techniques such as those

described in [Section 3.1.10](#) of [\[RFC8085\]](#) and similar protections at the network level in order to ensure network health in the event of misconfiguration, poorly written applications that don't follow UDP congestion control principles, or a deliberate attack.

[Section 4.1.4.2](#) of [\[RFC7450\]](#) and [Section 6.1](#) of [\[RFC7450\]](#) provide some further considerations and advice about mitigating congestion risk.

7. References

7.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [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>>.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/info/rfc2181>>.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, DOI 10.17487/RFC2782, February 2000, <<https://www.rfc-editor.org/info/rfc2782>>.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", STD 88, RFC 3596, DOI 10.17487/RFC3596, October 2003, <<https://www.rfc-editor.org/info/rfc3596>>.
- [RFC3597] Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", RFC 3597, DOI 10.17487/RFC3597, September 2003, <<https://www.rfc-editor.org/info/rfc3597>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, DOI 10.17487/RFC4604, August 2006, <<https://www.rfc-editor.org/info/rfc4604>>.

- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <<https://www.rfc-editor.org/info/rfc4607>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.
- [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>>.
- [RFC7450] Bumgardner, G., "Automatic Multicast Tunneling", RFC 7450, DOI 10.17487/RFC7450, February 2015, <<https://www.rfc-editor.org/info/rfc7450>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [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>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.

7.2. Informative References

- [RFC2317] Eidnes, H., de Groot, G., and P. Vixie, "Classless IN-ADDR.ARPA delegation", BCP 20, RFC 2317, DOI 10.17487/RFC2317, March 1998, <<https://www.rfc-editor.org/info/rfc2317>>.
- [RFC2845] Vixie, P., Gudmundsson, O., Eastlake 3rd, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, DOI 10.17487/RFC2845, May 2000, <<https://www.rfc-editor.org/info/rfc2845>>.
- [RFC2931] Eastlake 3rd, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, DOI 10.17487/RFC2931, September 2000, <<https://www.rfc-editor.org/info/rfc2931>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<https://www.rfc-editor.org/info/rfc3550>>.
- [RFC4025] Richardson, M., "A Method for Storing IPsec Keying Material in DNS", RFC 4025, DOI 10.17487/RFC4025, March 2005, <<https://www.rfc-editor.org/info/rfc4025>>.

- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4787] Audet, F., Ed. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <<https://www.rfc-editor.org/info/rfc4787>>.
- [RFC5110] Savola, P., "Overview of the Internet Multicast Routing Architecture", RFC 5110, DOI 10.17487/RFC5110, January 2008, <<https://www.rfc-editor.org/info/rfc5110>>.
- [RFC6726] Paila, T., Walsh, R., Luby, M., Roca, V., and R. Lehtonen, "FLUTE - File Delivery over Unidirectional Transport", RFC 6726, DOI 10.17487/RFC6726, November 2012, <<https://www.rfc-editor.org/info/rfc6726>>.
- [RFC7359] Gont, F., "Layer 3 Virtual Private Network (VPN) Tunnel Traffic Leakages in Dual-Stack Hosts/Networks", RFC 7359, DOI 10.17487/RFC7359, August 2014, <<https://www.rfc-editor.org/info/rfc7359>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8313] Tarapore, P., Ed., Sayko, R., Shepherd, G., Eckert, T., Ed., and R. Krishnan, "Use of Multicast across Inter-domain Peering Points", BCP 213, RFC 8313, DOI 10.17487/RFC8313, January 2018, <<https://www.rfc-editor.org/info/rfc8313>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.

Appendix A. Unknown RRType Construction

In a DNS resolver that understands the AMTRELAY type, the zone file might contain this line:

```
IN AMTRELAY 128 0 3 amtrelays.example.com.
```

In order to translate this example to appear as an unknown RRType as defined in [RFC3597], one could run the following program:

```
<CODE BEGINS>
$ cat translate.py
#!/usr/bin/env python3
import sys
name=sys.argv[1]
wire=''
for dn in name.split('.'):
    if len(dn) > 0:
        wire += ('%02x' % len(dn))
        wire += (''.join('%02x'%ord(x) for x in dn))
print(len(wire)//2) + 2
print(wire)

$ ./translate.py amtrelays.example.com
24
09616d7472656c617973076578616d706c6503636f6d

<CODE ENDS>
```

The length of the RData and the hex string for the domain name "amtrelays.example.com" are the outputs of this program.

The length of the wire-encoded domain name is 22, so 2 was added to that value (1 for the precedence field and 1 for the combined D-bit and relay type fields) to get the full length 24 of the RData. For the 2 octets ahead of the domain name, we encode the precedence, D-bit, and relay type fields, as described in [Section 4](#).

This results in a zone file entry like this:

```
IN TYPE260 \# ( 24 ; length
           80 ; precedence = 128
           03 ; D-bit=0, relay type=3 (wire-encoded domain name)
           09616d7472656c617973076578616d706c6503636f6d ) ; domain name
```

Acknowledgements

This specification was inspired by the previous work of Doug Nortz, Robert Sayko, David Segelstein, and Percy Tarapore, presented in the MBONED Working Group at IETF 93.

Thanks to Jeff Goldsmith, Toerless Eckert, Mikael Abrahamsson, Lenny Giuliano, Mark Andrews, Sandy Zheng, Kyle Rose, Ben Kaduk, Bill Atwood, Tim Chown, Christian Worm Mortensen, Warren Kumari, Dan Romanescu, Bernard Aboba, Carlos Pignataro, Niclas Comstedt, Mirja Kühlewind, Henning Rogge, Eric Vyncke, Barry Lieba, Roman Danyliw, Alissa Cooper, Suresh Krishnan, Adam Roach, and Daniel Franke for their very helpful reviews and comments.

Author's Address

Jake Holland

Akamai Technologies, Inc.

150 Broadway

Cambridge, MA 02144

United States of America

Email: jakeholland.net@gmail.com