ï»¿

Internet Engineering Task Force (IETF)
Request for Comments: 8987
Category: Standards Track
ISSN: 2070-1721

I. Farrer
Deutsche Telekom AG
N. Kottapalli
Benu Networks
M. Hunek
Technical University of Liberec
R. Patterson
Sky UK Ltd.

February 2021

DHCPv6 Prefix Delegating Relay Requirements

Abstract

This document describes operational problems that are known to occur when using DHCPv6 relays with prefix delegation. These problems can prevent successful delegation and result in routing failures. To address these problems, this document provides necessary functional requirements for operating DHCPv6 relays with prefix delegation.

It is recommended that any network operator using DHCPv6 prefix delegation with relays ensure that these requirements are followed on their networks.

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/rfc8987.

Table of Contents

1. Introduction

For Internet service providers that offer native IPv6 access with prefix delegation to their customers, a common deployment architecture is to have a DHCPv6 relay agent function located in the ISP's Layer 3 customer edge device and a separate, centralized DHCPv6 server infrastructure. [RFC8415] describes the functionality of a DHCPv6 relay, and Section 19.1.3 of [RFC8415] mentions this deployment scenario, but it does not provide details for all of the functional requirements that the relay needs to fulfill to operate deterministically in this deployment scenario.

A DHCPv6 relay agent for prefix delegation is a function commonly implemented in routing devices, but implementations vary in their functionality and client/server interworking. This can result in operational problems such as messages not being forwarded by the relay or unreachability of the delegated prefixes. This document provides a set of requirements for devices implementing a relay function for use with prefix delegation.

The mechanisms for a relay to inject routes (including aggregated ones) on its network-facing interface based on prefixes learned from a server via DHCP prefix delegation (DHCP-PD) are out of scope of the document.

Multi-hop DHCPv6 relaying is not affected. The requirements in this document are solely applicable to the DHCP relay agent co-located with the first-hop router to which the DHCPv6 client requesting the prefix is connected, so no changes to any subsequent relays in the path are needed.

2. Terminology

2.1. General

This document uses the terminology defined in [RFC8415]. However, when defining the functional elements for prefix delegation, [RFC8415], Section 4.2 defines the term "delegating router" as:

> The router that acts as a DHCP server and responds to requests for delegated prefixes.

This document is concerned with deployment scenarios in which the DHCPv6 relay and DHCPv6 server functions are separated, so the term "delegating router" is not used. Instead, a new term is introduced to describe the relaying function:

Delegating relay:
: A delegating relay acts as an intermediate device, forwarding DHCPv6 messages containing IA_PD and IAPREFIX options between the client and server. The delegating relay does not implement a DHCPv6 server function. The delegating relay is also responsible for routing traffic for the delegated prefixes.

Where the term "relay" is used on its own within this document, it should be understood to be a delegating relay unless specifically

stated otherwise.

In CableLabs DOCSIS environments, the Cable Modem Termination System (CMTS) would be considered a delegating relay with respect to Customer Premises Devices (CPEs) ([DOCSIS_3.1], Section 5.2.7.2). A Broadband Network Gateway (BNG) in a DSL-based access network may be a delegating relay if it does not implement a local DHCPv6 server function ([TR-092], Section 4.10).

[RFC8415] defines the "DHCP server" (or "server") as:

> A node that responds to requests from clients. It may or may not be on the same link as the client(s). Depending on its capabilities, if it supports prefix delegation it may also feature the functionality of a delegating router.

This document serves the deployment cases where a DHCPv6 server is not located on the same link as the client (necessitating the delegating relay). The server supports prefix delegation and is capable of leasing prefixes to clients, but it is not responsible for other functions required of a delegating router, such as managing routes for the delegated prefixes.

The term "requesting router" has previously been used to describe the DHCP client requesting prefixes for use. This document adopts the terminology of [RFC8415] and uses "DHCP client" or "client" interchangeably for this element.

2.2. Topology

The following diagram shows the deployment topology relevant to this document.

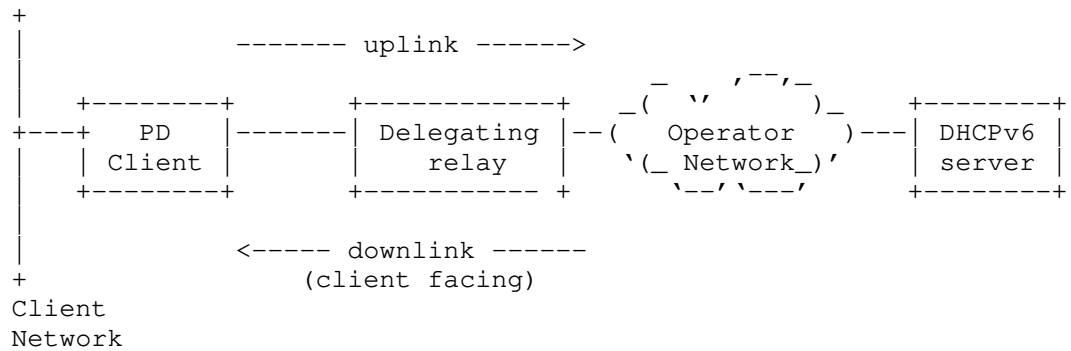

Figure 1: Topology Overview

The client requests prefixes via the downlink interface of the delegating relay. The resulting prefixes will be used for addressing the client network. The delegating relay is responsible for forwarding DHCP messages, including prefix delegation requests and responses between the client and server. Messages are forwarded from the delegating relay to the server using multicast or unicast via the operator uplink interface.

The delegating relay provides the operator's Layer 3 edge towards the client and is responsible for routing traffic to and from clients connected to the client network using addresses from the delegated prefixes.

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.

3. Problems Observed with Existing Delegating Relay Implementations

The following sections of the document describe problems that have been observed with delegating relay implementations in commercially available devices.

3.1. DHCP Messages Not Being Forwarded by the Delegating Relay

Delegating relay implementations have been observed not to forward messages between the client and server. This generally occurs if a client sends a message that is unexpected by the delegating relay. For example, the delegating relay already has an active PD lease entry for an existing client on a port. A new client is connected to this port and sends a Solicit message. The delegating relay then drops the Solicit messages until either it receives a DHCP Release message from the original client or the existing lease times out. This causes a particular problem when a client device needs to be replaced due to a failure.

In addition to dropping messages, in some cases, the delegating relay will generate error messages and send them to the client, e.g., "NoBinding" messages being sent in the event that the delegating relay does not have an active delegated prefix lease.

3.2. Delegating Relay Loss of State on Reboot

For proper routing of client traffic, the delegating relay requires a corresponding routing table entry for each active prefix delegated to a connected client. A delegating relay that does not store this state persistently across reboots will not be able to forward traffic to the client's delegated leases until the state is reestablished through new DHCP messages.

3.3. Multiple Delegated Prefixes for a Single Client

DHCPv6 [RFC8415] allows a client to include more than one instance of OPTION_IA_PD in messages in order to request multiple prefix delegations by the server. If configured for this, the server supplies one (or more) instance of OPTION_IAPREFIX for each received instance of OPTION_IA_PD, each containing information for a different delegated prefix.

In some delegating relay implementations, only a single delegated prefix per DHCP Unique Identifier (DUID) is supported. In those cases, only one IPv6 route for one of the delegated prefixes is installed, meaning that other prefixes delegated to a client are unreachable.

3.4. Dropping Messages from Devices with Duplicate MAC Addresses and DUIDs

It is an operational reality that client devices with duplicate Media Access Control (MAC) addresses and/or DUIDs exist and have been deployed. In some networks, the operational costs of locating and swapping out such devices are prohibitive.

Delegating relays have been observed to restrict forwarding client messages originating from one client DUID to a single interface. In this case, if the same client DUID appears from a second client on another interface while there is already an active lease, messages originating from the second client are dropped, causing the second client to be unable to obtain a prefix delegation.

It should be noted that in some access networks, the MAC address and/or DUID are used as part of device identification and authentication. In such networks, enforcing uniqueness of the MAC address and/or DUID is a necessary function and is not considered a problem.

3.5. Forwarding Loops between Client and Relay

If the client loses information about an active prefix lease it has been delegated while the lease entry and associated route are still active in the delegating relay, then the relay will forward traffic

to the client. The client will return this traffic to the relay, which is the client's default gateway (learned via a Router Advertisement (RA)). The loop will continue until either the client is successfully reprovisioned via DHCP or the lease ages out in the relay.

4. Requirements for Delegating Relays

To resolve the problems described in Section 3 and to preempt other undesirable behavior, the following section of the document describes a set of functional requirements for the delegating relay.

In addition, relay implementers are reminded that [RFC8415] makes it clear that relays MUST forward packets that either contain message codes it may not understand (Section 19 of [RFC8415]) or options that it does not understand (Section 16 of [RFC8415]).

4.1. General Requirements

G-1: The delegating relay MUST forward messages bidirectionally between the client and server without changing the contents of the message.

G-2: The relay MUST allow for multiple prefixes to be delegated for the same client IA_PD. These delegations may have different lifetimes.

G-3: The relay MUST allow for multiple prefixes (with or without separate IA_PDs) to be delegated to a single client connected to a single interface, identified by its DHCPv6 Client Identifier (DUID).

G-4: A delegating relay may have one or more interfaces on which it acts as a relay, as well as one or more interfaces on which it does not (for example, in an ISP, it might act as a relay on all southbound interfaces but not on the northbound interfaces). The relay SHOULD allow the same client identifier (DUID) to have active delegated prefix leases on more than one interface simultaneously unless client DUID uniqueness is necessary for the functioning or security of the network. This is to allow client devices with duplicate DUIDs to function on separate broadcast domains.

G-5: The maximum number of simultaneous prefixes delegated to a single client MUST be configurable.

G-6: The relay MUST implement a mechanism to limit the maximum number of active prefix delegations on a single port for all client identifiers and IA_PDs. This value MUST be configurable.

G-7: It is RECOMMENDED that delegating relays support at least 8 active delegated leases per client device and use this as the default limit.

G-8: The delegating relay MUST update the lease lifetimes based on the client's reply messages it forwards to the client and only expire the delegated prefixes when the valid lifetime has elapsed.

G-9: On receipt of a Release message from the client, the delegating relay MUST expire the active leases for each of the IA_PDs in the message.

4.2. Routing Requirements

R-1: The relay MUST maintain a local routing table that is dynamically updated with leases and the associated next hops as they are delegated to clients. When a delegated prefix is released or expires, the associated route MUST be removed from the relay's routing table.

R-2: The delegating relay's routing entry MUST use the same prefix length for the delegated prefix as given in the IA_PD.

R-3: The relay MUST provide a mechanism to dynamically update ingress filters permitting ingress traffic sourced from client delegated leases and blocking packets from invalid source prefixes. This is to implement anti-spoofing as described in [BCP38]. The delegating relay's ingress filter entry MUST use the same prefix length for the delegated prefix as given in the IA_PD.

R-4: The relay MAY provide a mechanism to dynamically advertise delegated leases into a routing protocol as they are learned. If such a mechanism is implemented, when a delegated lease is released or expires, the delegated route MUST be withdrawn from the routing protocol. The mechanism by which the routes are inserted and deleted is out of the scope of this document.

R-5: To prevent routing loops, the relay SHOULD implement a configurable policy to drop potential looping packets received on any DHCP-PD client-facing interfaces.

The policy SHOULD be configurable on a per-client or per-destination basis.

Looping packets are those with a destination address in a prefix delegated to a client connected to that interface, as follows:

* For point-to-point links, when the packet's ingress and egress interfaces match.

* For multi-access links, when the packet's ingress and egress interface match, and the source link-layer and next-hop link-layer addresses match.

An ICMPv6 Type 1, Code 6 (Destination Unreachable, reject route to destination) error message MAY be sent as per [RFC4443], Section 3.1. The ICMP policy SHOULD be configurable.

4.3. Service Continuity Requirements

S-1: To preserve active client prefix delegations across relay restarts, the relay SHOULD implement at least one of the following:

* Implement DHCPv6 Bulk Leasequery as defined in [RFC5460].

* Store active prefix delegations in persistent storage so they can be reread after the reboot.

S-2: If a client's next-hop link-local address becomes unreachable (e.g., due to a link-down event on the relevant physical interface), routes for the client's delegated prefixes MUST be retained by the delegating relay unless they are released or removed due to expiring DHCP timers. This is to reestablish routing for the delegated prefix if the client next hop becomes reachable without the delegated prefixes needing to be relearned.

S-3: The relay SHOULD implement DHCPv6 Active Leasequery as defined in [RFC7653] to keep the local lease database in sync with the DHCPv6 server.

4.4. Operational Requirements

O-1: The relay SHOULD implement an interface allowing the operator to view the active delegated prefixes. This SHOULD provide information about the delegated lease and client details such as the client identifier, next-hop address, connected

interface, and remaining lifetimes.

O-2: The relay SHOULD provide a method for the operator to clear active bindings for an individual lease, client, or all bindings on a port.

O-3: To facilitate troubleshooting of operational problems between the delegating relay and other elements, it is RECOMMENDED that a time synchronization protocol be used by the delegating relays and DHCP servers.

5. IANA Considerations

This document has no IANA actions.

6. Security Considerations

This document does not add any new security considerations beyond those mentioned in Section 4 of [RFC8213] and Section 22 of [RFC8415].

If the delegating relay implements [BCP38] filtering, then the filtering rules will need to be dynamically updated as delegated prefixes are leased.

[RFC8213] describes a method for securing traffic between the relay agent and server by sending DHCP messages over an IPsec tunnel. It is RECOMMENDED that this be implemented by the delegating relay.

Failure to implement requirement G-6 may have specific security implications, such as a resource depletion attack on the relay.

The operational requirements in Section 4.4 may introduce additional security considerations. It is RECOMMENDED that the operational security practices described in [RFC4778] be implemented.

7. References

7.1. Normative References

[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>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[RFC4778] Kaeo, M., "Operational Security Current Practices in Internet Service Provider Environments", RFC 4778, DOI 10.17487/RFC4778, January 2007, <https://www.rfc-editor.org/info/rfc4778>.

[RFC5460] Stapp, M., "DHCPv6 Bulk Leasequery", RFC 5460, DOI 10.17487/RFC5460, February 2009, <https://www.rfc-editor.org/info/rfc5460>.

[RFC7653] Raghuvanshi, D., Kinnear, K., and D. Kukrety, "DHCPv6 Active Leasequery", RFC 7653, DOI 10.17487/RFC7653, October 2015, <https://www.rfc-editor.org/info/rfc7653>.

[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>.

[RFC8213] Volz, B. and Y. Pal, "Security of Messages Exchanged between Servers and Relay Agents", RFC 8213, DOI 10.17487/RFC8213, August 2017,

<https://www.rfc-editor.org/info/rfc8213>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

7.2. Informative References

[BCP38] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

<https://www.rfc-editor.org/info/bcp38>

[DOCSIS_3.1] CableLabs, "MAC and Upper Layer Protocols Interface Specification", Version 10, DOCSIS 3.1, January 2017, <https://www.cablelabs.com/specification/CM-SP-MULPIv3.1>.

[TR-092] Broadband Forum, "Broadband Remote Access Server (BRAS) Requirements Document", Technical Report TR-092, August 2004, <https://www.broadband-forum.org/download/TR-092.pdf>.

Acknowledgements

The authors of this document would like to thank Bernie Volz, Ted Lemon, and Michael Richardson for their valuable comments.

Authors' Addresses

Ian Farrer
Deutsche Telekom AG
Landgrabenweg 151
53227 Bonn
Germany

Email: ian.farrer@telekom.de

Naveen Kottapalli
Benu Networks
WeWork Galaxy, 43 Residency Road
Bangalore 560025
Karnataka
India

Email: nkottapalli@benunets.com

Martin Hunek
Technical University of Liberec
Studentska 1402/2
46017 Liberec
Czech Republic

Email: martin.hunek@tul.cz

Richard Patterson
Sky UK Ltd.
1 Brick Lane
London
E1 6PU
United Kingdom

Email: richard.patterson@sky.uk