Internet Engineering Task Force (IETF)
Request for Comments: 9928
Category: Standards Track
ISSN: 2070-1721

C. Porfiri
Ericsson
S. Krishnan
Cisco
J. Arkko
M. Khlewind
Ericsson

March 2026

DHCPv4 over DHCPv6 with Relay Agent Support

Abstract

This document describes a mechanism for networks with legacy IPv4-only clients to use services provided by DHCPv4 over DHCPv6 in a Relay Agent. RFC 7341 specifies the use of DHCPv4 over DHCPv6 in the client only. This document specifies an approach based on RFC 7341 that allows a Relay Agent to implement the DHCPv4-over-DHCPv6 encapsulation and decapsulation of DHCPv4 messages in DHCPv6 messages on behalf of a DHCPv4 client.

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/rfc9928.

Table of Contents

1. Introduction

[RFC7341] describes a transport mechanism for carrying DHCPv4 [RFC2131] messages using DHCPv6 [RFC9915] for dynamic provisioning of IPv4 addresses and other DHCPv4-specific configuration parameters across IPv6-only networks. The deployment of [RFC7341] requires support in DHCP clients and at the DHCPv6 server. However, if a client is embedded in a host that only supports IPv4 and cannot easily be replaced or updated (which could be due to any number of technical or business reasons), this approach does not work.

Similarly, the DHCPv6 Relay Agent specification defined in [RFC9915], which also refers to [RFC6221] for the Lightweight DHCPv6 Relay Agent (LDRA) behavior, does not provide any mechanism to handle legacy DHCPv4, except by requiring the client to implement the DHCPv4-over-DHCPv6 encapsulation and decapsulation.

This document specifies a solution based on [RFC7341] that can be implemented in intermediate nodes such as switches or routers, without putting any requirements on clients. No new protocols or extensions are needed; instead, this document specifies a new use case for [RFC7341] that allows a Relay Agent to perform the DHCPv4-over-DHCPv6 encapsulation and decapsulation instead of the client.

1.1. Applicability Scope

The mechanisms described in this document apply to the configuration phase of hosts that need to receive an IPv4 address when a DHCP server for IPv4 [RFC2131] is not reachable directly from the host. Furthermore, the host is unable to implement a DHCP client conformant to [RFC7341], as it is connected to an IPv4-only network. However, there is a DHCPv6 server that can provide IPv4 addresses by means of the mechanisms specified in [RFC7341].

2. Conventions and Definitions

The following terms and abbreviations are used in this document:

DHCP:
: Refers to DHCPv4 and/or DHCPv6 if not otherwise specified.

DHCP Relay Agent:
: Refers to a common concept in all of the following protocols, although the details differ between them: the Bootstrap Protocol (BOOTP) [RFC0951] [RFC1542], DHCPv4 [RFC2131] [RFC2132], and DHCPv6 [RFC9915].

DHCPv4:
: Refers to DHCP as defined in [RFC2131].

DHCPv4 over DHCPv6 (DHCP 4o6):
: Refers to the architecture, the procedures, and the protocols specified in the DHCPv4-over-DHCPv6 document [RFC7341].

DHCPv4-over-DHCPv6 Relay Agent (4o6RA):
: Refers to a Relay Agent that implements the DHCP 4o6 transport as specified in this document.

Layer 3 Relay Agent (L3RA):
   Refers to a DHCP Relay Agent as specified in [RFC9915] that is not a LDRA.

Lightweight DHCPv6 Relay Agent (LDRA):
   Refers to an extension of the original DHCPv6 Relay Agent specification, to allow Layer 2 (L2) only devices to perform a Relay Agent function [RFC6221].

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. DHCPv4-over-DHCPv6 Relay Agent (4o6RA)

This document assumes a network where IPv4-only hosts are connected to a network that supports IPv6 and limited IPv4 services.

To address such a network setup, this document extends DHCPv6 Relay Agents with DHCPv4 over DHCPv6, as shown in Figure 1.

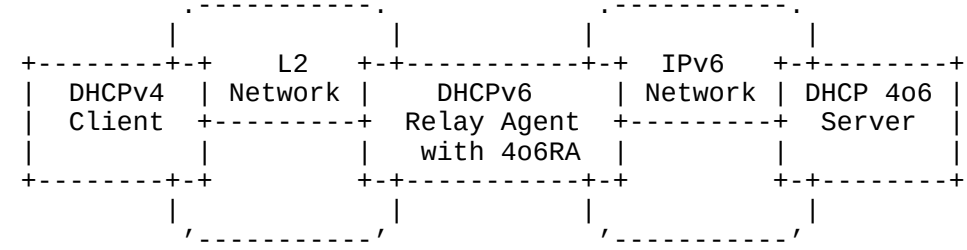

Figure 1: Architecture Example with Legacy DHCP Client

This document specifies the encapsulation and decapsulation specified in [RFC7341] to be performed in the Relay Agent without requiring any changes on the DHCPv4 client. In this case, it is up to the Relay Agent to provide the full DHCP 4o6 support, and the legacy DHCPv4 client is not aware that it is being served via a DHCP 4o6 service. As the 4o6RA acts as a DHCP 4o6 client, all prerequisites and configurations that apply to the DHCP client in Section 5 of [RFC7341] are also applied to the 4o6RA.

As the 4o6RA takes the role of the client in respect to [RFC7341], it is responsible for determining a suitable interface where it acts as a DHCPv6 client, and it is responsible for locating a suitable DHCPv6 server or Relay Agent and obtaining the necessary IPv6 configuration. As specified in [RFC7341], the 4o6RA, acting as DHCP 4o6 client, therefore has to request the DHCP 4o6 Server Address option from the server by sending the Option Request option as described in [RFC9915] before it can use the DHCP 4o6 transport.

To maintain interoperability with existing DHCPv6 relays and servers, the message format is unchanged from [RFC9915]. The 4o6RA implements the same message types as a DHCPv6 Relay Agent (see Section 6 of [RFC7341]).

However, in this specification, the 4o6RA, instead of the client, creates the DHCPV4-QUERY message and encapsulates the DHCP request message received from the legacy DHCPv4 client.

When the DHCPV4-RESPONSE message is received by the DHCP 4o6 Relay Agent, it looks for the DHCPv4 message option within this message. If this option is not found or the DHCPv4-RESPONSE message is not well-formed, it MUST be discarded. If the DHCPv4 message option is

present and correct, the 4o6RA MUST extract the DHCPv4 message and forward the encapsulated DHCPv4-RESPONSE to the requesting DHCPv4 client, given that the encapsulated DHCPv4-RESPONSE is correct and can be actually forwarded.

Layer 2 (L2) Relay Agents receiving DHCPV4-QUERY or DHCPV4-RESPONSE messages MUST handle them as specified in Section 6 of [RFC6221].

In any given environment, DHCPv6 servers to which DHCPV4-QUERY requests are routed are expected to be compliant with DHCP 4o6 according to [RFC7341]. No additional requirements on DHCPv6 servers are set by this specification.

3.1. Intermediate Relays

Intermediate relays shall behave according to Section 10 of [RFC7341].

3.2. 4o6RA and Topology Discovery

In some networks, the configuration of a host may depend on the topology. However, when a new host attaches to a network, it may be unaware of the topology and, consequently, how it has to be configured.

DHCPv4 [RFC2131] and DHCPv6 [RFC9915] specifications describe how addresses can typically be allocated to clients based on network topology information provided by a DHCP relay.

Address/prefix allocation decisions are integral to the allocation of addresses and prefixes in DHCP, as described in detail in [RFC7969]. This specification aims to guarantee that the 4o6RA does not break any legacy capability when used for topology discovery.

Topology discovery as described in [RFC7969] differs between IPv4 and IPv6 as follows:

* IPv4: When using DHCP on IPv4, only the first Relay Agent SHOULD set the giaddr field (Section 3.1 of [RFC7969]). Thus, in a network that has more than one Relay Agent, only part of the topology is transported via DHCPv4.

* IPv6: When using DHCPv6, all Relay Agents SHOULD send link-address and Interface-ID options that provide information about the complete path between the DHCPv6 client and the DHCPv6 server to the DHCPv6 server.

In Layer 2 networks, Lightweight DHCPv6 Relay Agents (LDRAs) [RFC6221] can be used.

When provided, the topology information is available at the DHCPv6 server in the form of a sequence of the link-address field and Interface-ID option.

Then, topology information for the given IP address can be obtained from the DHCPv6 server and used for configuration or other purposes.

[RFC7341] enables the client to use DHCPv6 for topology discovery even within a DHCPv4 context, as the DHCPv6 Relay Agent knows the interface where the encapsulated DHCP request is received. However, as shown in Figure 2, the introduction of DHCP 4o6 at the edge of the IPv6 network hides the Layer 2 network from the DHCPv6 RA. As such, moving DHCP 4o6 to an intermediate node rather than performing it at the client breaks the topology propagation, as 4o6RA-only solutions do not provide any interface information in the encapsulated message.

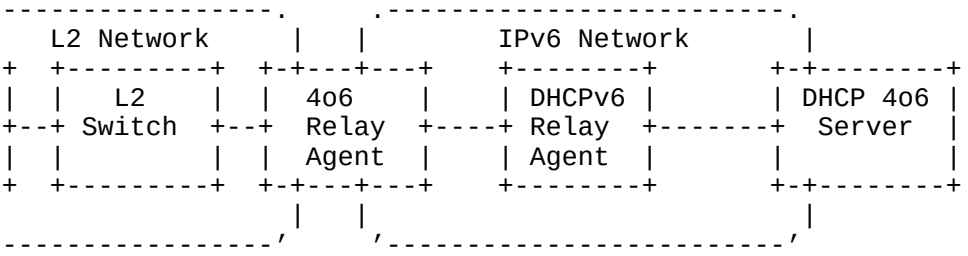

Figure 2: Broken Topology Information

In order to provide full topology information, it is RECOMMENDED that any implementation of 4o6RA be combined with an LDRA implementation [RFC6221] in a back-to-back structure and that the LDRA implementation includes a mechanism to obtain interface information that can be used to provide the Interface-ID option to outgoing DHCPV4-QUERY messages, as specified in Section 5.3.2 of [RFC6221].

The internal mechanisms to exchange interface information, their format, and whether the interface information contains an indication that a 4o6RA is involved, are out of the scope for this document.

The resulting architecture is shown in Figure 3 where the Relay Agent is implementing 4o6RA and LDRA and has an internal interface to propagate topology information from 4o6RA to LDRA.

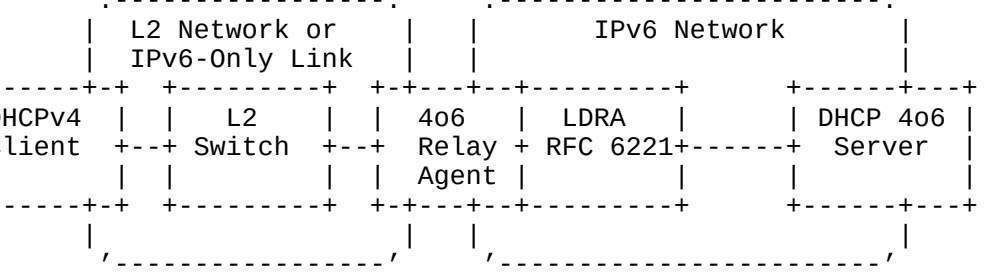

Figure 3: Topology Information Preserved with LDRA

In a simple case, where the same node hosts the 4o6RA and the DHCP 4o6 server, it might be enough to only use 4o6RA, as shown in Figure 4, where CPE stands for "Customer Premises Equipment".

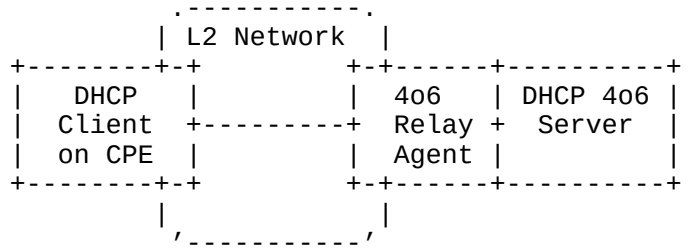

Figure 4: Topology Information Preserved by 4o6 Relay Agent in DHCP Server

4. Deployment Considerations

As clients are unaware of the presence of 4o6RA, the network deployment needs to ensure that all DHCPv4 broadcast and unicast messages to and from clients are steered via a 4o6RA. This can be achieved by placing the 4o6RA in a central position that can intercept all traffic from the clients or by using Network Address Translation (NAT) with the 4o6RA address for unicast messages.

5. Security Considerations

This document specifies the applicability of DHCP 4o6 in a scenario where legacy IPv4 clients are connected to DHCP 4o6 Relay Agents that perform the encapsulation and decapsulation. This document does not change anything else in the DHCP 4o6 specification [RFC7341]; therefore, the security considerations of that document still apply. Specifically, since the legacy IPv4 client is not aware of the encapsulation and decapsulation, 4o6RA has to provide the protections that are specified in the security considerations in Section 12 of [RFC7341].

The mechanisms defined here differ from [RFC7341] as they allow the DHCP client to send and receive DHCPv4 messages, whereas in [RFC7341], the client only sends DHCPv6 messages. This makes it possible that in improperly configured networks where the client is located on the same Layer 2 scope of a DHCPv4 server, DHCPv4 messages could reach a DHCPv4 server without using the 4o6RA. While this can cause erroneous state in both clients and servers and potentially even lead to misconfigurations that impact reachability, this is seen as a deployment error rather than a security concern. Further, even though this mechanism may be used for attacks from within the network, this is not a new concern introduced by this specification.

More generally, legacy IPv4 clients are not aware of this mechanism; however, even when DHCP 4o6 is used, the client does not have any control about the information provided by the Relay Agent. As such, this change does not raise any additional security concerns.

6. IANA Considerations

This document has no IANA actions.

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>.

[RFC6221] Miles, D., Ed., Ooghe, S., Dec, W., Krishnan, S., and A. Kavanagh, "Lightweight DHCPv6 Relay Agent", RFC 6221, DOI 10.17487/RFC6221, May 2011, <https://www.rfc-editor.org/info/rfc6221>.

[RFC7341] Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4-over-DHCPv6 (DHCP 4o6) Transport", RFC 7341, DOI 10.17487/RFC7341, August 2014, <https://www.rfc-editor.org/info/rfc7341>.

[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>.

[RFC9915] Mrugalski, T., Volz, B., Richardson, M., Jiang, S., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", STD 102, RFC 9915, DOI 10.17487/RFC9915, January 2026, <https://www.rfc-editor.org/info/rfc9915>.

7.2. Informative References

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

[RFC1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, DOI 10.17487/RFC1542, October 1993, <https://www.rfc-editor.org/info/rfc1542>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <https://www.rfc-editor.org/info/rfc2132>.

[RFC7969] Lemon, T. and T. Mrugalski, "Customizing DHCP Configuration on the Basis of Network Topology", RFC 7969, DOI 10.17487/RFC7969, October 2016, <https://www.rfc-editor.org/info/rfc7969>.

Appendix A. Example Use Case: Topology Discovery for IPv4-Only Radio Unit in 3GPP RAN with Switched Fronthaul

In 3GPP mobile network architecture, the User Equipment (UE) is connected via a Radio Access Network (RAN). RAN is built up with Baseband Units (BBUs) and Radio Units (RUs). A radio Fronthaul (FH) network connects RUs and BBUs. Each RU and BBU is an IP host, and they may support IPv4 only, IPv6 only, or both, depending on the vendor and the model. Each RU is unique as it is tied to a set of antennas, and each antenna is serving a specific Cell and Sector. Each RU is configured by the BBU depending on the Cell and Sectors it serves. However, that dependency is only specified by the cabling between RUs and antennas. BBUs can be cabled to RUs directly or via a Layer 2 switched network.

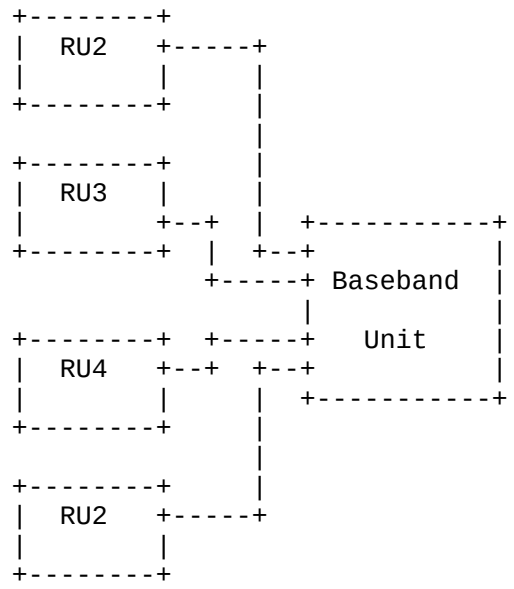

Figure 5: 3GPP RAN Where RUs Are Cabled Directly to BBUs

In Figure 5, the BBU is directly cabled to a set of RUs, and the BBU can recognize the relationship between RUs and Cell/Sectors based on the cabling between the RUs and antennas.

When BBUs and RUs are connected via a Layer 2 switched network, the added level of complexity requires the BBUs to have a deeper knowledge of the topology in order to properly configure the RUs, involving knowledge of all the cabling in the switched network.

Examples for switched networks are shown in Section 3 of [RFC7969]

and demonstrate the different levels of complexity. An example of a FH is depicted in Figure 6.

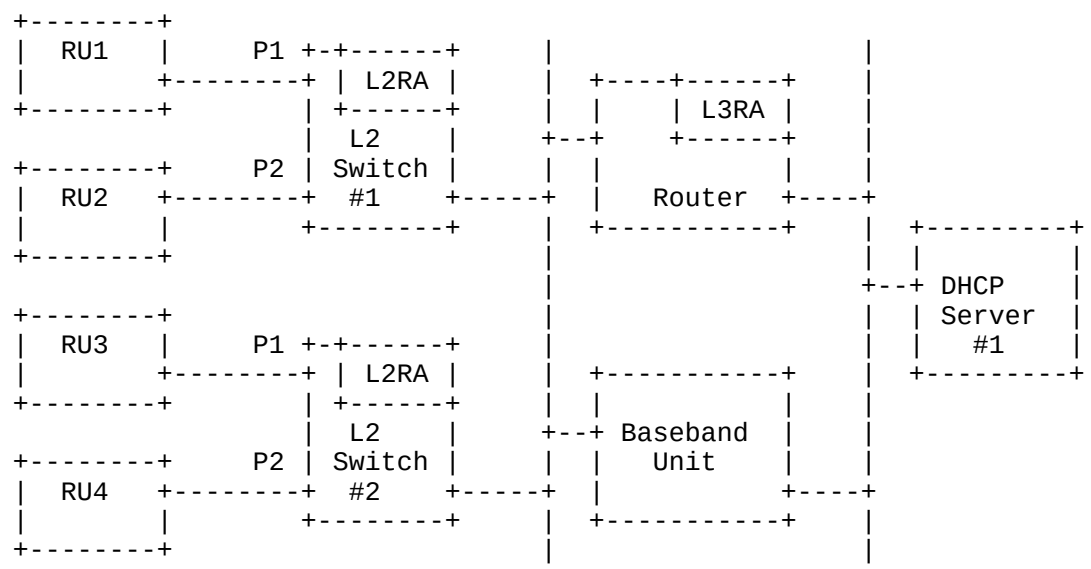

Figure 6: 3GPP RAN with Layer 2 Switched Fronthaul Example

If IPv6 is used and all RUs are capable of DHCPv6 in Figure 6, DHCP topology knowledge can be used for solving the RU configuration problem. Such solution would use the topology discovery mechanisms described in Section 3.2 of [RFC7969].

If RUs are capable of IPv4 only but implement a DHCP 4o6 client according to [RFC7341], the same topology discovery mechanisms are applicable.

If RUs are capable of IPv4 only and cannot implement a DHCP 4o6 client according to [RFC7341], the topology discovery mechanisms described in Section 3.2 of [RFC7969] can be used by introducing 4o6RA in the switches as described in this document.

Acknowledgments

The authors would like to acknowledge interesting discussions in this problem space with Sarah Gannon, Ines Ramadza, and Siddharth Sharma, as well as reviews and comments provided by ric Vyncke, Mohamed Boucadair, David Lamparter, Michael Richardson, Alan DeKok, Dale Worley, Paul Wouters, Deb Cooley, Erik Kline, Ketan Talaulikar, Mike Bishop, and Roman Danyliw.

Authors’ Addresses

Claudio Porfiri
Ericsson
Email: claudio.porfiri@ericsson.com

Suresh Krishnan
Cisco
Email: suresh.krishnan@gmail.com

Jari Arkko
Ericsson
Email: jari.arkko@ericsson.com

Mirja Khlewind

Ericsson
Email: mirja.kuehlewind@ericsson.com