

Dynamic Host Configuration
Internet-Draft
Intended status: Standards Track
Expires: 15 February 2026

C. Porfiri
Ericsson
S. Krishnan
Cisco
J. Arkko
M. K_端hlewind
Ericsson
14 August 2025

DHCPv4-over-DHCPv6 with Relay Agent Support
draft-ietf-dhc-dhcpv4-over-dhcpv6-ra-05

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. RFC7341 specifies use of DHCPv4-over-DHCPv6 in the client only. This document specifies a RFC7341-based approach that allows a Relay Agent to implement the DHCP 4o6 encapsulation and decapsulation of DHCPv4 messages in DHCPv6 messages on behalf of a DHCPv4 client.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at
<https://datatracker.ietf.org/doc/draft-ietf-dhc-dhcpv4-over-dhcpv6-ra/>.

Source for this draft and an issue tracker can be found at
<https://github.com/mirjak/draft-dhc-dhcpv4-over-dhcpv6-ra>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 15 February 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Applicability Scope	3
2. Conventions and Definitions	3
3. DHCPv4 over DHCPv6 Relay Agent (4o6RA)	4
3.1. Intermediate relays	5
3.2. 4o6RA and Topology Discovery	5
4. Deployment Considerations	7
5. Security Considerations	8
6. IANA Considerations	8
7. References	8
7.1. Normative References	8
7.2. Informative References	9
Appendix A. Example Use Case: Topology Discovery for IPv4-only Radio Unit in 3GPP RAN with Switched Fronthaul	9
Acknowledgments	11
Authors' Addresses	11

1. Introduction

[RFC7341] describes a transport mechanism for carrying DHCPv4 [RFC2131] messages using DHCPv6 [draft-ietf-dhc-rfc8415bis] 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 specifications for DHCPv6 Relay Agents such as Lightweight DHCPv6 Relay Agent (LDRA) [RFC6221] or DHCPv6 Relay Agent (L3RA) [draft-ietf-dhc-rfc8415bis] do not foresee the possibility to handle legacy DHCPv4, other than implementing DHCP 4o6 in the client.

This document specifies an [RFC7341] based solution 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 an amendment to [RFC7341] that allows a Relay Agent to perform the DHCP 4o6 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 but 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. But 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 acronyms are used in this document:

- * DHCP: If not otherwise specified, DHCP refers to DHCPv4 and/or DHCPv6.
- * DHCPv4: DHCP as defined in [RFC2131].
- * DHCPv4 over DHCPv6 (or 4o6): The architecture, the procedures, and the protocols specified in the DHCPv4-over-DHCPv6 document [RFC7341].
- * DHCP Relay Agent: This is a concept in all of the following protocols, although the details differ between them: BOOTP [RFC951] [RFC1542], DHCPv4 [RFC2131] [RFC2132], and DHCPv6 [draft-ietf-dhc-rfc8415bis].
- * Lightweight DHCPv6 Relay Agent (or LDRA): This is an extension of the original DHCPv6 Relay Agent specification, to allow layer-2-only devices to perform a Relay Agent function [RFC6221].
- * DHCPv4 over DHCPv6 Relay Agent (or 4o6RA): Refers to a Relay Agent that implements the 4o6 transport as specified in this document.

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 configuration 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 obtain the necessary IPv6 configuration.. As specified in [RFC7341], the 4o6RA, acting as 4o6 client, therefore has to request the DHCP 4o6 Server Address option from the server by sending the Option Request option as described in [draft-ietf-dhc-rfc8415bis] before it can use the 4o6 transport.

To maintain interoperability with existing DHCPv6 relays and servers, the message format is unchanged from [draft-ietf-dhc-rfc8415bis]. The 4o6RA implements the same message types as a DHCPv6 Relay Agent 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 DHCPV4-RESPONSE Message is received by the 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, 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 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 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 the new host attaches to a network, it may be unaware of the topology and respectively how it has to be configured.

DHCPv4 [RFC2131] and DHCPv6 [draft-ietf-dhc-rfc8415bis] specifications describe how addresses can be allocated to clients based on network topology information provided by a DHCP relay, typically.

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:

- * 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 [RFC6221] can be used.

When provided, the topology information is available at the DHCPv6 server in form of 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 an DHCPv4 context, as the DHCPv6 Relay Agent knows the interface where the encapsulated DHCP request is received. As shown in Figure 2, the introduction of 4o6 at the edge of the IPv6 network, however, hides the Layer-2 network from the DHCPv6 RA. As such, moving 4o6 in an intermediate node rather than performing it at the client, breaks the topology propagation as 4o6RA-only does 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 implementation of an LDRA [RFC6221] in a back-to-back structure, and that the LDRA implementation has a mechanism to get 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 DHCP4o6 server, it might be enough to only use 4o6RA, as shown in Figure 4.




Figure 4: Topology information preserved by 4o6 Relay Agent in DHCP server

4. Deployment Considerations

As clients are not aware of the presence of 4o6RA, the network deployment needs to ensure that all DHCPv4 broadcast and unicast messages from and to clients are steered via a 4o6RA. This can be achieved by placing the 4o6RA in a central position that can observe all traffic from the clients or use Network Address Translation (NAT) with the 4o6RA address for unicast messages.

5. Security Considerations

This document specifies the applicability of 4o6 DHCP in a scenario where legacy IPv4 clients are connected to 4o6 DHCP Relay Agents that perform the encapsulation and decapsulation. This document does not change anything else in the 4o6 DHCP specification and therefore the security considerations of [RFC7341] still apply.

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

- [draft-ietf-dhc-rfc8415bis]
"Dynamic Host Configuration Protocol for IPv6 (DHCPv6)",
June 2025, <<https://datatracker.ietf.org/doc/draft-ietf-dhc-rfc8415bis/>>.
- [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/rfc/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/rfc/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/rfc/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/rfc/rfc8174>>.

7.2. Informative References

- [RFC1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, DOI 10.17487/RFC1542, October 1993, <<https://www.rfc-editor.org/rfc/rfc1542>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/rfc/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/rfc/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/rfc/rfc7969>>.
- [RFC951] Croft, W. and J. Gilmore, "Bootstrap Protocol", RFC 951, DOI 10.17487/RFC0951, September 1985, <<https://www.rfc-editor.org/rfc/rfc951>>.

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 Equipments (UE) are connected via Radio Access Network (RAN). RAN is built up with Baseband Units (BB) and Radio Units (RU). Radio Fronthaul Network (FH) connects RU and BB, each of RU and BB is an IP host, 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 BB depending on the Cell and Sectors it serves. However, that dependency is only specified by the cabling between RU and antennas. BB can be cabled to RU directly or via a Layer-2 switched network.




Figure 5: 3GPP RAN where RU are cabled directly to BB

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

When BBs and RUs are connected via a Layer-2 switched network, the added level of complexity requires the BBs 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 RU 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 RU are capable of IPv4 only but implement a 4o6 client according to [RFC7341], the same topology discovery mechanisms are applicable.

If RU are capable of IPV4 only and cannot implement a 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 also 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 Eric Vyncke, Mohamed Boucadair, David Lamparter, Michael Richardson, Alan DeKok and Dale Worley.

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 Kuehlewind
Ericsson
Email: mirja.kuehlewind@ericsson.com