

dhc Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 3, 2016

S. Gandhewar
Juniper Networks, Inc.
October 1, 2015

DHCPv6 Relay Initiated Release
draft-gandhewar-dhc-v6-relay-initiated-release-01

Abstract

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) is initiated by a DHCPv6 client. A DHCPv6 server can force DHCPv6 client to send RENEW or INFORMATION-REQUEST by sending a RECONFIGURE message. There may be multiple DHCPv6 network devices connected in between a DHCPv6 client and a server, each one reserving resources for the DHCPv6 client. There are no DHCPv6 messages that a relay can initiate in order to control the client binding.

A DHCPv6 client may not always send a RELEASE message when it no longer needs the IPv6 address or prefix and network resources for the associated services it is using. This document specifies a way to request release to be initiated by an intermediate DHCPv6 network device, e.g. DHCPv6 relay, on behalf of DHCPv6 client. This helps to relinquish network resources sooner than the lease expiration time.

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 <http://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 April 3, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://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	3
1.1.	Problem Description	3
1.2.	Relay Initiated Release	4
1.3.	Applicability	5
2.	Requirements Language	7
3.	Protocol Details	7
3.1.	Message Definitions	7
3.1.1.	RELEASE-REQUEST	7
3.1.2.	RELEASE-REQUEST-REPLY	8
3.2.	Message Validation	8
3.2.1.	RELEASE-REQUEST	8
3.2.2.	RELEASE-REQUEST-REPLY	9
4.	Functionality	9
4.1.	First DHCPv6 Network Device Behavior	9
4.1.1.	Generation and Transmission of RELEASE-REQUEST Message	10
4.1.2.	Receipt of RELEASE-REQUEST Message	11
4.2.	Intermediate DHCPv6 Network Device Behavior	11
4.3.	DHCPv6 Server Behavior	11
4.4.	Receipt of RELEASE-REQUEST-REPLY	12
5.	Security Considerations	12
6.	IANA Considerations	12
7.	Acknowledgements	13
8.	References	13
8.1.	Normative References	13
8.2.	Informative References	13
	Author's Address	14

1. Introduction

DHCPv6 [RFC3315] and [RFC3633] provides a framework for configuring clients with network addresses, address prefixes and other network parameters. It includes a relay agent capability where DHCPv6 server may not be directly connected to the DHCPv6 client. A relay agent is an intermediate node that passes DHCPv6 messages between DHCPv6 clients and DHCPv6 servers. As per [RFC3315], a relay agent cannot generate a message on its own which can control the client binding. Figure 1 below shows a typical network with multiple DHCPv6 devices.



Figure 1: Typical DHCPv6 Network

1.1. Problem Description

While providing an IPv6 address or IPv6 Prefix to the DHCPv6 Client, Service Providers (e.g. Broadband Service Providers), creates a logical interface per client, programs various routes (e.g. access routes, framed routes) for the client to access the network and services, attaches services (e.g. voice, video, data), maintains policy, applies QoS. Along with these resources there is a need for memory and bandwidth per client. Since all these resources are limited on a network device (e.g. Broadband Network Gateway), it defines the scaling capacity of the device. Since the availability of the IPv6 addresses is large, subscription rate for the Service Providers is thus limited by the availability of the resources on their network device.

A DHCPv6 client may be connected to the DHCPv6 server through multiple DHCPv6 network devices, e.g. multiple DHCPv6 relays. These network resources remain reserved for the client at all the DHCPv6 network devices until the lease expires.

In some situations, there might be need to clear the client binding administratively. The process of administratively clearing the client binding is very cumbersome. The administrator needs to access every single DHCPv6 network device (relay, relay-proxy) and also the DHCPv6 server, and clear the DHCPv6 client binding at each of these devices manually.

In some situations when the DHCPv6 client is replaced (e.g. replacing the set-top-box) due to the device failure or upgrade, the older DHCPv6 client might not have sent the RELEASE message on its failure.

In this case, the previously assigned IPv6 address or prefix and network resources for the older (stale) client will stay reserved and unused until the lease expires.

Same is the situation where clients move frequently without sending RELEASE e.g. in the case of mobile networks, network resources stay reserved and unused. Similarly, network resources stay reserved and unused where DHCPv6 clients login and logout frequently without sending RELEASE e.g. Wi-Fi access centers.

As per DHCPv6 protocol it is not mandatory for the DHCPv6 client to send a RELEASE message while disconnecting. As per the statistics from Service Providers, 95% of the cases DHCPv6 client does not send RELEASE message when it no longer needs the service. It is also possible that the UDP datagram carrying a RELEASE message may get dropped due to network issues.

All the resources including the IPv6 address or prefix remains reserved for the client at all the DHCPv6 network devices until the lease expires. Service Providers needs to take into account such situations and are forced to lower the subscription rate. Thus it reduces the scaling per network device. Also it causes errors for the time based billing.

1.2. Relay Initiated Release

It is possible for the first DHCPv6 network device, i.e. "DHCPv6 Relay 1" in Figure 1 which is closest to the DHCPv6 client, to detect that the DHCPv6 client is replaced, moved or is no longer present on the network. In this scenario, the relay agent doesn't have any mechanism to inform the server to release the client's binding and subsequently relinquish network resources.

With the relay initiated release message, when a DHCPv6 relay detects client's unavailability or needs to clear the client binding administratively, it can generate the release message on behalf of the client and send it to the server. Thus, all the DHCPv6 network devices along the path will be in synchronization with respect to the client's binding information and network resources can be relinquished earlier than the lease expiry. The server MAY choose to integrate some mechanism to confirm with the client, e.g. generate RECONFIGURE message before sending reply to the relay. It is outside the scope of this document.

Generation of the relay initiated release SHOULD be a configurable behavior at the first relay. The configuration at Relay SHOULD be further granular to indicate the situation under which relay should

initiate the release e.g. administratively clearing DHCPv6 binding, client replaced, client moved, client unavailable, etc.

Forwarding of the relay initiated release related messages SHOULD be a configurable behavior at the intermediate DHCPv6 network devices.

Acceptance of relay initiated release SHOULD also be a configurable behavior at the server.

The purpose of such configurable behavior is explained in Section 1.3.

1.3. Applicability

As per the statistics from Service Providers, 95% of the cases DHCPv6 client does not send RELEASE message when it no longer needs the service. This functionality is useful in order to relinquish network resources sooner than the lease expiry. This allows Service Providers for higher subscription rate and accurate time based billing.

This functionality described in Section 1.2 is useful for clearing the client binding administratively, client replacement, frequent client login and logout without sending RELEASE (e.g. at Wi-Fi centers) or where client moves frequently without sending RELEASE (e.g. mobile networks). All these situations can be detected by the first DHCPv6 network device. Thus this functionality is applicable to all these situations without any problems.

This functionality is also useful where client unavailability can be detected. Client unavailability could be because of multiple reasons. Client may become unavailable due to powered-off, disconnect from the network or problems in the network itself. Since it is difficult to identify the cause of client's absence, precaution must be taken in such situations. With this functionality described in Section 1.2, the state of the binding is cleared and network resources are relinquished at DHCPv6 Relay, DHCPv6 Server and all the intermediate network devices. However it is possible that the binding is still not cleared at the DHCPv6 client. There may be a situation where client remembers the IPv6 address or prefix as well as the lease it received and continue to use when network comes back. This situation may happen when the network between Relay and client becomes unavailable and Relay may assume that the client is unavailable.

When such a situation happens where all the DHCPv6 network devices cleared the binding but client still remembers and tries to use the address or prefix, at that moment there is no way to clear the

binding at the client. The client's binding will get cleared at the client at the time of Renew or Rebind or when the lease expires or when client restarts DHCPv6 process.

This may not be a problem in case of DSL based networks where DHCPv6 is over PPP session. The failed PPP session will cause the DHCPv6 client to bring up the PPP session and restart the DHCPv6 discovery process. However it may be a problem with an Ethernet based access network since there is no trigger event to the CPE (client) to restart the DHCPv6 binding process.

In some provider networks, DHCPv6 Relay has liveness detection. When the network between DHCPv6 Relay and client becomes unavailable, DHCPv6 Relay may initiate Release, whereas client is completely unaware. It is not possible to differentiate between network unavailable and client unavailable. This will very likely be the case with cable network configurations. If the link between Cable Modem and the CMTS goes down, the Relay running on CMTS may initiate release for the Cable Modem as well as the devices behind the Cable Modem unless Cable Modem runs the DHCPv6 Relay. The granular configuration to initiate Release on client unavailability should be turned off in such networks.

However, there are some Service Provider networks where DHCPv6 client runs the liveness detection e.g. BFD on the provider facing interface. Such DHCPv6 clients can identify the network unavailability and may restart the DHCPv6 binding process.

In some Service Provider networks, Relay takes up longer lease from the Server but gives out very small lease to the DHCPv6 client. This forces DHCPv6 client to frequently renew the lease. Thus recovery from problematic state of the DHCPv6 client will be much faster in such network configurations.

For some of the Service Provider's configurations, DHCPv6 Relay adds access routes per subscriber (DHCPv6 client) and remove these routes on clearing the binding on receiving the REPLY for RELEASE or the RELEASE-REQUEST-REPLY. Thus the Relay can restrict DHCPv6 client's network traffic based on the source or the destination address and thus restrict the harm and protects from two devices accessing the network with the same IPv6 address.

This functionality SHOULD be a configurable behavior since there is no clear way to distinguish between DHCPv6 client unavailable and network unavailable. Having configurable behavior equips administrator to enable this granular knob (send Relay Initiated Release on DHCPv6 client's unavailability) at Relay only if it is certain that such a situation will not occur or client will clear the

binding state and reestablish or the risk of such situation is being accounted.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. Protocol Details

3.1. Message Definitions

This document specifies 2 new DHCPv6 message types:

- o RELEASE-REQUEST
- o RELEASE-REQUEST-REPLY

The RELEASE-REQUEST and RELEASE-REQUEST-REPLY messages use the Client/Server Message Formats described in Section 6 of [RFC3315], similar to the LEASEQUERY and LEASEQUERY-REPLY in [RFC5007].

3.1.1. RELEASE-REQUEST

This is the relay initiated release request message.

The RELEASE-REQUEST message MAY be generated by the first DHCPv6 network device ("DHCPv6 Relay 1" in Figure 1), on behalf of the DHCPv6 client. The RELEASE-REQUEST message MUST contain one or more Client Data Options as described in Section 4.1.2.2 of [RFC5007], requesting release for one or more clients.

The RELEASE-REQUEST message MUST contain the Server Identifier Option. It MAY contain Interface-Id Option indicating common values for all the clients requesting the release. This reduces the redundant data when there are multiple clients with common information.

Each Client Data Option MUST include the Client Identifier Option OPTION_CLIENTID. It MUST also include options containing the IAs - OPTION_IAADDR, OPTION_IAPREFIX, etc. - for the addresses or prefixes it is releasing. If the Interface-Id option is different from the one included directly under RELEASE-REQUEST message then it MUST be included here.

3.1.2. RELEASE-REQUEST-REPLY

This is the reply for the RELEASE-REQUEST message.

The message RELEASE-REQUEST-REPLY will be generated by the DHCPv6 Server to communicate the status of the request. The server conveys the success or failure of the RELEASE-REQUEST by including Status Code Option at different levels:

- o Status Code Option directly inside RELEASE-REQUEST: Indicates success or failure of the complete RELEASE-REQUEST message it received.
- o Status Code Option inside Client Data Option: Indicates success or failure to release all the addresses or prefixes for a particular client. Client Data Option MUST include the Client-Id Option.
- o Status Code Option inside IA Option: Indicates success or failure to release a particular address or prefix for a particular client. Client Data Option MUST include the Client-Id Option and the IA option.

The RELEASE-REQUEST-REPLY message MAY contain one or more Client Data Options, described in Section 4.1.2.2 of [RFC5007], responding to the request to release for each of the clients.

The RELEASE-REQUEST-REPLY message SHOULD contain the Interface-Id option if it was included in RELEASE-REQUEST message.

3.2. Message Validation

3.2.1. RELEASE-REQUEST

Clients MUST silently discard any received RELEASE-REQUEST messages.

Relay MAY accept or discard any received RELEASE-REQUEST messages depending upon the configuration as explained in Section 4.1.2.

Servers MUST discard any received RELEASE-REQUEST messages that meet any of the following conditions:

- o The message does not include a Relay Id Option.
- o The message does not include a Client Data Option.
- o The Client Data Option does not include a Client Identifier Option.

- o The message does not include a Server Identifier option.
- o The message includes a Server Identifier Option but the contents of the Server Identifier Option do not match the server's identifier.

3.2.2. RELEASE-REQUEST-REPLY

Clients MUST silently discard any received RELEASE-REQUEST-REPLY messages.

Servers MUST silently discard any received RELEASE-REQUEST-REPLY messages.

Relay MUST discard any received RELEASE-REQUEST-REPLY messages that meet any of the following conditions:

- o The "transaction-id" field in the message does not match the value used in the RELEASE-REQUEST message.
- o The message does not include a Status Code Option.

4. Functionality

The generation of a RELEASE-REQUEST message SHOULD be a configurable behavior at DHCPv6 network device. Similarly, taking action to release the binding SHOULD also be a configurable behavior at the DHCPv6 server and intermediate DHCPv6 network devices.

4.1. First DHCPv6 Network Device Behavior

Devices MAY be configured to generate the newly defined RELEASE-REQUEST message.

The first DHCPv6 network device ("DHCPv6 Relay 1" in Figure 1) can be configured such that when it detects the client is no longer available on the network or is replaced or the binding information needs to be deleted administratively, the device can generate the RELEASE-REQUEST message.

In order to generate the RELEASE-REQUEST message this network device needs to store the information related to the client, e.g. the client identifier and the server identifier used while obtaining the client lease.

4.1.1. Generation and Transmission of RELEASE-REQUEST Message

Set the "msg-type" field to RELEASE-REQUEST.

Generate a transaction ID and insert it in the "transaction-id" field.

MUST include Server-Id Option.

MUST include Relay-Id option [RFC5460].

MAY add Interface-Id option [RFC3315].

MUST include one or more Client Data Options each one:

- o MUST include Client Identifier and MUST be same as what was used when client obtained the lease.
- o MAY include options containing the IAs (IA_NA, IA_TA, IA_PD) for the addresses or prefixes it is requesting to be released. Absence of this option indicates release of all the addresses and prefixes associated with this Client Identifier.
- o MAY include Interface-Id option [RFC3315] if it is different from the one included outside of the Client Data Option

Because RELEASE-REQUEST messages MAY be lost, the message SHOULD be retransmitted if no RELEASE-REQUEST-REPLY message is received. The client transmits the message according to Section 14 of [RFC3315], using the following parameters:

- o IRT REL_TIMEOUT
- o MRT 0
- o MRC REL_MAX_RC
- o MRD 0

If RELEASE-REQUEST-REPLY from a DHCPv6 server is lost, then the RELEASE-REQUEST will be retransmitted, and the server MAY respond with a RELEASE-REQUEST-REPLY indicating a status as NoBinding. Therefore, in this message exchange, the relay SHOULD NOT treat a RELEASE-REQUEST-REPLY message with a status of NoBinding as an error.

4.1.2. Receipt of RELEASE-REQUEST Message

In order to protect against spoofed RELEASE-REQUEST messages attempting to disconnect the clients, the first DHCPv6 network device SHOULD drop any received RELEASE-REQUEST messages. It MUST be a configurable behavior if these messages are from the trusted sources and needs to be forwarded to the server.

4.2. Intermediate DHCPv6 Network Device Behavior

The behavior of the intermediate DHCPv6 network device can be configurable to either accept or reject these messages. On accepting, it can forward the messages as specified in Section 20.1 and 20.2 of [RFC3315].

4.3. DHCPv6 Server Behavior

DHCPv6 server ("DHCPv6 Server" in Figure 1) SHOULD be configurable to either accept or reject the relay initiated release message RELEASE-REQUEST. Upon receipt of a RELEASE-REQUEST message, the server MUST confirm the validity of the message.

If server does not support the new message type then it MAY simply drop the packet.

If the server is not configured to accept this relay initiated RELEASE-REQUEST message then it MAY simply drop the packet or send RELEASE-REQUEST-REPLY with status as NotConfigured.

If the server decides not to accept the RELEASE-REQUEST from a particular relay, it MAY simply drop the packet or send RELEASE-REQUEST-REPLY with status as NotAllowed.

The server SHOULD iterate through each of the Client Data Options and examine the Client-Id and the addresses in the IAs for validity. If the addresses or prefixes in the IAs have been assigned by the server, the server deletes the binding of these addresses and prefixes and makes them available for assignment to other clients. Server keeps note of these addresses and prefixes in the IAs for generating the RELEASE-REQUEST-REPLY.

After all of the clients have been processed, the server generates a RELEASE-REQUEST-REPLY message and includes a Status Code Option with value Success. It also includes Server Identifier option.

For each of the clients where there is a failure in releasing addresses or prefixes, server MUST include Client Data Option. In the Client Data Option, it MUST include the Client Identifier option

from the RELEASE-REQUEST message. It MUST also include Status Code Option for each of the failed IAs from the RELEASE-REQUEST message. For the clients or IAs for which the server has no binding information, correspondingly, the server MUST include a Status Code Option with the value NoBinding. No other options are included in the IA option.

4.4. Receipt of RELEASE-REQUEST-REPLY

The first DHCPv6 network device ("DHCPv6 Relay 1" in Figure 1), upon receipt of a valid RELEASE-REQUEST-REPLY message, considers the completion of RELEASE-REQUEST event. The action at this device is based on the status. For all of the IAs or clients where the Status Code is not Success or NoBinding, addresses and prefixes remain unchanged until the lease expires. For all other clients and IAs, bindings MUST be cleared.

5. Security Considerations

The RELEASE-REQUEST message provides a mechanism for releasing the client binding, it can be the cause of security threat. The DHCPv6 server SHOULD have some mechanism for determining that the relay agent is a trusted entity. DHCPv6 servers and relay agents MAY implement relay message authentication as described in Section 21.1 of [RFC3315]. DHCPv6 servers MAY also implement a control policy based on the content of a received Relay Identifier Option [RFC5460]. Administrators MAY configure one of these security mechanisms.

In an environment where the network connecting the relay agent to the DHCPv6 server is physically secure and does not contain devices not controlled by the server administrator, it MAY be sufficient to trust the Relay Agent Identifier provided by the relay agent. In networks where the security of the machines with access to the data path is not under the control of the server administrator, IPsec [RFC4301] is necessary to prevent spoofing of messages.

DHCPv6 servers MUST silently discard RELEASE-REQUEST messages originating from unknown or untrusted relay agents or reject the RELEASE-REQUEST. Section 4.3 specifies the error code to return when the server is configured to reject RELEASE-REQUEST messages.

6. IANA Considerations

We request IANA to assign following new message types from the registry of Message Types maintained in:
<http://www.iana.org/assignments/dhcpv6-parameters/>

- o RELEASE-REQUEST

- o RELEASE-REQUEST-REPLY

7. Acknowledgements

We would like to acknowledge Utae Kim (Smart GiGA Network Project, Korea Telekom), Dan Seibel (Sr. Engineer, TELUS), Ian Farrer (Network Architect, Deutsche Telekom) and Chris Topazi (Access Engineering, Cox Communications) for their valuable contributions, suggestions and support for this document.

We would like to thank Bernie Volz, Ted Lemon, Andrew Sullivan, Ole Troan and Shrivinas Joshi for their valuable comments and suggestions for improving the document.

Many thanks to Tomek Mrugalski, Bernie Volz and Jaya Bhawtankar (Lead Engineer, Coriant) for their support.

We would like to acknowledge Anand Vijayvergiya, Jeff Haas and Ross Callon for their guidance and tirelessly reviewing the document multiple times.

8. References

8.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.

8.2. Informative References

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.

[RFC5007] Brzozowski, J., Kinnear, K., Volz, B., and S. Zeng,
"DHCPv6 Leasequery", RFC 5007, DOI 10.17487/RFC5007,
September 2007, <<http://www.rfc-editor.org/info/rfc5007>>.

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

Author's Address

Sunil M. Gandhewar
Juniper Networks, Inc.

Email: sgandhewar@juniper.net