

Networking Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 30, 2018

N. Shen  
E. Chen  
Cisco Systems  
October 27, 2017

**Generalized UDP Source Port for DHCP Relay**  
**draft-ietf-dhc-relay-port-07**

Abstract

This document proposes an extension to the DHCP protocols that allows a relay agent to use any available source port for upstream communications, and to include a DHCP option that can be used to statelessly route responses back to the appropriate source port on downstream communications.

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 30, 2018.

Copyright Notice

Copyright (c) 2017 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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">1.1.</a>	<a href="#">Requirements Language</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Changes to DHCP Specifications</a>	<a href="#">3</a>
<a href="#">3.1.</a>	<a href="#">Changes to DHCPv4 in <a href="#">RFC 2131</a></a>	<a href="#">3</a>
<a href="#">3.2.</a>	<a href="#">Changes to DHCPv6 in <a href="#">RFC 3315</a></a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Relay Source Port Sub-option and Option</a>	<a href="#">4</a>
<a href="#">4.1.</a>	<a href="#">Source Port Sub-option for DHCPv4</a>	<a href="#">4</a>
<a href="#">4.2.</a>	<a href="#">Relay Source Port Option for DHCPv6</a>	<a href="#">5</a>
<a href="#">5.</a>	<a href="#">Relay Agent and Server Behavior</a>	<a href="#">6</a>
<a href="#">5.1.</a>	<a href="#">DHCPv4</a>	<a href="#">6</a>
<a href="#">5.2.</a>	<a href="#">DHCPv6</a>	<a href="#">6</a>
<a href="#">5.3.</a>	<a href="#">Compatibility</a>	<a href="#">7</a>
<a href="#">5.4.</a>	<a href="#">Deployment Considerations</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">An IPv6 Cascaded Relay Example</a>	<a href="#">7</a>
<a href="#">7.</a>	<a href="#">IANA Considerations</a>	<a href="#">9</a>
<a href="#">8.</a>	<a href="#">Security Considerations</a>	<a href="#">9</a>
<a href="#">9.</a>	<a href="#">Acknowledgments</a>	<a href="#">9</a>
<a href="#">10.</a>	<a href="#">Normative References</a>	<a href="#">9</a>
	<a href="#">Authors' Addresses</a>	<a href="#">10</a>

## [1.](#) Introduction

[RFC 2131](#) [[RFC2131](#)] and [RFC 3315](#) [[RFC3315](#)] specify the use of UDP as the transport protocol for DHCP. They also define both the server side and client side port numbers. The IPv4 server port is UDP number (67) and the client port is UDP number (68); for IPv6 the server port is (546) and the client port is (547).

The fixed UDP port combinations for the DHCP protocol scheme creates challenges in certain DHCP relay operations. For instance, in a large scale DHCP relay implementation on a single switch node, the DHCP relay functionality may be partitioned among multiple relay processes. All these DHCP relay processes may share the same IP address of the switch node. If the UDP source port has to be a fixed number as currently specified, the transport socket operation of DHCP packets would need to go through a central entity or process which would defeat the purpose of distributing DHCP relay functionality.

In some large-scale deployment, the decision to split the DHCP functionality into multiple processes on a node may not be purely based on DHCP relay computational load. But rather DHCP relay could just be one of the functions in a multi-process implementation.

Although assigning a different IPv4/IPv6 source address for each DHCP relay process can be a solution, it would introduce operational and



network management complexities, especially given the scarceness of the IPv4 addresses.

This document proposes an extension to relax the fixed UDP source port requirement for the DHCP relay agents. This extension requires a DHCP server to remember the inbound packet's UDP port number along with the IPv4/IPv6 address. The DHCP server when sending back replies MUST use the UDP port number that the incoming relay agent uses instead of the fixed DHCP port number. In the case of IPv6 cascaded relay agents [[RFC3315](#)], the upstream relay agent needs to use the "Relay Source Port Option" to record the downstream source port and it MUST use this recorded port number instead of the fixed DHCP port number when replaying the reply messages.

### **1.1. 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](#)].

## **2. Terminology**

**Downstream Device:** In the DHCP relay context, it refers to the next relay agent for forwarding Relay-reply Messages.

**Upstream Device:** In the DHCP relay context, it refers to the next relay agent or DHCP server for forwarding Relay-forward Messages.

**Relay Source Port:** This is the UDP port that a relay agent uses to receive Relay-forward Messages from an upstream device.

**Downstream Source Port:** This is the UDP port that the downstream device uses when forwarding Relay-forward Messages to this relay agent device. This UDP port is to be used by this relay agent device when forwarding the Relay-reply Messages to that downstream device.

**Non-DHCP UDP Port:** Any valid UDP port other than port 67 for DHCPv4 and port 547 for DHCPv6.

## **3. Changes to DHCP Specifications**

### **3.1. Changes to DHCPv4 in [RFC 2131](#)**

[Section 4.1 of RFC 2131](#) [[RFC2131](#)] specifies that:



DHCP uses UDP as its transport protocol. DHCP messages from a client to a server are sent to the 'DHCP server' port (67), and DHCP messages from a server to a client are sent to the 'DHCP client' port (68).

This specification adds the following extension to the above paragraph.

DHCP messages from a relay agent to a server are sent to the 'DHCP server' port (67), and the UDP source port it uses can be any valid UDP port available in the relay system, including the DHCP port 67. The default port number is 67 if there is no explicit configuration for the generalized source UDP port extension for DHCP relay.

### **3.2. Changes to DHCPv6 in [RFC 3315](#)**

[Section 5.2 of RFC 3315](#) [[RFC3315](#)] specifies that:

Clients listen for DHCP messages on UDP port 546. Servers and relay agents listen for DHCP messages on UDP port 547.

This specification adds the following extension to the above paragraph.

A DHCP relay agent can listen for DHCP messages from a server or another upstream relay agent device on any valid UDP port available in the relay system including the DHCP UDP port 547. The default UDP port is 547 if there is no explicit configuration for the generalized UDP source port extension for DHCP relay.

## **4. Relay Source Port Sub-option and Option**

Relay agents do not maintain state. To return a message to its source, the relay agent must include all the required information in the Relay-Forward message. When a relay in a sequence of cascaded relays does not use the standard source port, that source port must be included along with the source address. This option allows the relay agent to do so.

### **4.1. Source Port Sub-option for DHCPv4**

The Relay Agent "Source Port Sub-option" is a new option, and it is part of the relay-agent-information option for DHCPv4 [[RFC3046](#)].

The format of the "Source Port Sub-option" is shown below:



```

+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| SubOpt Code  |      Len      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Where:

SubOpt Code: SUBOPT\_RELAY\_PORT. 8 bits value, to be assigned by IANA.

Len: 8 bits value to be set to 0.

#### [4.2.](#) Relay Source Port Option for DHCPv6

The "Relay Source Port Option" is a new DHCPv6 option. It MUST be used either by a DHCPv6 relay agent that uses a non-DHCP UDP port (not 547) communicating with the IPv6 server and the upstream relay agent, or by a IPv6 relay agent that detects the use of a non-DHCP UDP port (not 547) by a downstream relay agent.

The format of the "Relay Source Port Option" is shown below:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  OPTION_RELAY_RELAY_PORT  |      Option-Len      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Downstream Source Port    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Where:

Option-Code: OPTION\_RELAY\_RELAY\_PORT. 16 bits value, to be assigned by IANA.

Option-Len: 16 bits value to be set to 2.

Downstream Source Port: 16 bits value. To be set by the IPv6 relay either to the downstream relay agent's UDP source port used for the UDP packet, or to zero if only the local relay agent uses the non-DHCP UDP port (not 547).





## **5. Relay Agent and Server Behavior**

### **5.1. DHCPv4**

When a relay agent uses a non-DHCP UDP port (not 67) communicating with the DHCP server, it MUST include the "Source Port Sub-option" in Relay-forward messages to indicate that.

When an IPv4 server receives a message from a relay agent with the "Source Port Sub-option", it MUST remember the UDP source port of the message and use that port number as the UDP destination port when sending the reply message to the same relay agent.

### **5.2. DHCPv6**

The IPv6 relay agent MUST include the "Relay Source Port Option" when it uses a non-DHCP UDP port (not 547) to communicate to a DHCPv6 server or an upstream IPv6 relay agent. Also when an IPv6 relay agent detects that a downstream relay agent uses a non-DHCP UDP port in the packet, it MUST record the port number in the "Downstream Source Port" field of this option. If this option is included to indicate only the local non-DHCP UDP port usage and there is no downstream relay agent's non-DHCP UDP port usage, the field Downstream Source Port field MUST be set to zero.

The IPv6 relay agent MUST include this option in the following three cases:

- 1) The local relay agent uses a non-DHCP UDP port (not 547).
- 2) the downstream relay agent uses a non-DHCP UDP port (not 547).
- 3) the local relay agent and the downstream relay agent both use non-DHCP UDP ports (not 547).

In the first case, the value of the "Downstream Source Port" field is set to zero. In the other two cases, the value of the field is set to the UDP port number that the downstream relay agent uses.

When an IPv6 server receives a Relay-forward message with the "Relay Source Port Option", it MUST copy the option when constructing the Relay-reply chain in response to the Relay-forward message. This option MUST NOT appear in any message other than a Relay-forward or Relay-reply message. Additionally, the IPv6 server MUST check and use the UDP source port from the UDP packet of the Relay-forward message in replying to the relay agent.



When a relay agent receives a Relay-reply message with the "Relay Source Port Option" from a server or from an upstream relay agent, if the "Downstream Source Port" field in the option is non-zero, it MUST use this UDP port number to forward the Relay-reply message to the downstream relay agent.

### **5.3. Compatibility**

Sites that need for relay agents to specify a source port will need to install new DHCP server and DHCP relay agent software with this feature. If a site installs only DHCP relay agent software with this feature, there is no possibility that the DHCP server will be able to communicate to the relay agent.

### **5.4. Deployment Considerations**

During deployment, it may be advisable the operator and/or user of the new DHCP relay port implementation upgrade the DHCP server first when possible, before the relay implementations are deployed. This would ensure that the erroneous case noted in [Section 5.3](#) is not encountered.


When the DHCP relay port implementation is deployed, it is recommended that the configuration is setup to allow for the mode of operation where a non-DHCP port can be used for the DHCP relay agents. The recommended configuration then permits the relay agent to utilize the default DHCP UDP port, or a non-DHCP UDP port as desired.

Although if the network uses firewall to block or allow DHCP packets with both static UDP source and destination port numbers, this may no longer match the packets from new DHCP relay agent and server software with this extension. The firewall rules need to be modified only to match the DHCP server side of the UDP port number, and if necessary, IP addresses and other attributes.

## **6. An IPv6 Cascaded Relay Example**

An example of IPv6 cascaded relay agents with the "Relay Source Port Option" is shown below.





In the above diagram, all the IPv6 devices support this generalized UDP source port extension except for Relay3. Relay1 is the only relay agent device uses a non-DHCP UDP port (not 547). Relay2 is the upstream device of Relay1.

Both Relay1 and Relay2 include the "Relay Source Port Option" in Relay-forward message. Relay1 sets the "Downstream Source Port" field in the option to zero. Relay2 notices the "Relay Source Port Option" is included in the message from Relay1, and it determines that the UDP source port used by Relay1 is 1000. Relay2 will include the "Relay Source Port Option" and it sets the "Downstream Source Port" field in the option to 1000. The IPv6 server copies the "Relay Source Port Option" when replying with the Relay-reply message.

When Relay2 receives the Relay-reply message with the "Relay Source Port Option", it finds the "Downstream Source Port" field has the value of 1000. Relay2 then uses this port number in the UDP packet when sending the Relay-reply message to Relay1.

When Relay1 receives the Relay-reply message with the "Relay Source Port Option", it finds that the "Downstream Source Port" field has the value of zero. Relay1 then uses the normal IPv6 port 547 in the packet sending the Relay-reply message to its downstream relay agent or uses UDP port 546 to an IPv6 client.

This DHCP extension works with any combination of IPv6 cascaded relay agents, as long as the relay agent which uses a non-DHCP UDP port (not 547) and its upstream relay device support this generalized UDP source port extension.

Similar to the above example, now assume that Relay2 uses the UDP source port of 2000 instead of 547 as in the diagram. The Relay3 device needs to support this DHCP extension and it will set 2000 in its "Downstream Source Port" field of the option in the Relay-forward message. When DHCP server sends the DHCP Relay-reply to Relay3, Relay3 finds its own relay option has this "Downstream Source Port" with the value of 2000. Relay3 will use this UDP port when sending the Relay-reply message to Relay2.



## **7. IANA Considerations**

A new sub-option, DHCPv4 Relay Source Port, is defined in this document within the IPv4 Relay Agent Information Option. It needs to be assigned by IANA in the "DHCP Relay Agent Sub-Option Codes" registry, <http://www.iana.org/assignments/bootp-dhcp-parameters> as specified in [RFC3046].

A new option, DHCPv6 Relay Source Port, is defined in this document for DHCPv6 and it needs to be assigned by IANA for the DHCPv6 option code, in the "Option Codes" registry for DHCPv6, <http://www.iana.org/assignments/dhcpv6-parameters> as specified in [RFC3315].

## **8. Security Considerations**

[RFC3118] and [RFC3315] described many of the threats in using DHCP. This extension does not raise addition security issues.

## **9. Acknowledgments**

The authors would like to thank Peter Arberg, Luyuan Fang, Bhanu Gopalasetty, Scott Kelly, Andre Kostur, Victor Kuarsingh, Ted Lemon, Kishore Seshadri and Jackelyn Shen for their review and comments of this document.

The authors would like to thank Bernie Volz for discussions that led to the definition of The Relay Source Port sub-option and DHCPv6 Relay Source Port Option.

The RFC text was produced using Marshall Rose's xml2rfc tool.

## **10. 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>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", [RFC 3046](#), DOI 10.17487/RFC3046, January 2001, <<https://www.rfc-editor.org/info/rfc3046>>.





- [RFC3118] Droms, R., Ed. and W. Arbaugh, Ed., "Authentication for DHCP Messages", [RFC 3118](#), DOI 10.17487/RFC3118, June 2001, <<https://www.rfc-editor.org/info/rfc3118>>.
- [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, <<https://www.rfc-editor.org/info/rfc3315>>.

#### Authors' Addresses

Naiming Shen  
Cisco Systems  
560 McCarthy Blvd.  
Milpitas, CA 95035  
US

Email: [naiming@cisco.com](mailto:naiming@cisco.com)

Enke Chen  
Cisco Systems  
560 McCarthy Blvd.  
Milpitas, CA 95035  
US

Email: [enkechen@cisco.com](mailto:enkechen@cisco.com)

