

Internet Engineering Task Force (IETF)
Request for Comments: 8357
Category: Standards Track
ISSN: 2070-1721

N. Shen
E. Chen
Cisco Systems
March 2018

Generalized UDP Source Port for DHCP Relay

Abstract

This document defines an extension to the DHCP protocols that allows a relay agent to use any available source port for upstream communications. The extension also allows inclusion of 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 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/rfc8357>.

Copyright Notice

Copyright (c) 2018 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 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	2
2. Terminology	3
2.1. Requirements Language	3
2.2. Definitions	3
3. Changes to DHCP Specifications	4
3.1. Additions to DHCPv4 in RFC 2131	4
3.2. Additions to DHCPv6 in RFC 3315	4
4. Relay Source Port Sub-option and Option	4
4.1. Source Port Sub-option for DHCPv4	5
4.2. Relay Source Port Option for DHCPv6	5
5. Relay Agent and Server Behavior	6
5.1. DHCPv4	6
5.2. DHCPv6	6
5.3. Compatibility	7
5.4. Deployment Considerations	7
6. Example of an IPv6-Cascaded Relay	7
7. IANA Considerations	9
8. Security Considerations	9
9. Normative References	9
Acknowledgments	10
Authors' Addresses	10

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- 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 (547) and the client port is (546).

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 of 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 deployments, the decision to split the DHCP functionality into multiple processes on a node may not be purely based on DHCP relay computational load. 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, this would introduce operational and network management complexities, especially given the scarceness of IPv4 addresses.

This document defines 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. When sending back replies, the DHCP server 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.

2. Terminology

2.1. 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.

2.2. Definitions

Downstream Device: In the DHCP relay context, this refers to the next relay agent that forwards Relay-reply messages.

Upstream Device: In the DHCP relay context, this refers to the next relay agent or DHCP server that forwards 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 and non-zero UDP port other than port 67 for DHCPv4 and port 547 for DHCPv6.

3. Changes to DHCP Specifications

3.1. Additions 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).

Relay agents implementing this specification may be configured instead to 1) use a source port number other than 67 when relaying messages toward servers and 2) receive responses toward clients on that same port. This will only work when the DHCP server or relay agent to which such a relay agent is forwarding messages is upgraded to support this extension.

3.2. Additions 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.

Relay agents implementing this specification may be configured instead to 1) use a source port number other than 547 when relaying messages toward servers and 2) receive responses toward clients on that same port. This will only work when the DHCP server or relay agent to which such a relay agent is forwarding messages is upgraded to support this extension.

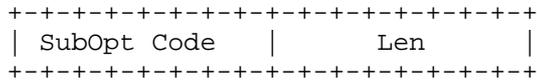
4. Relay Source Port Sub-option and Option

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:



Where:

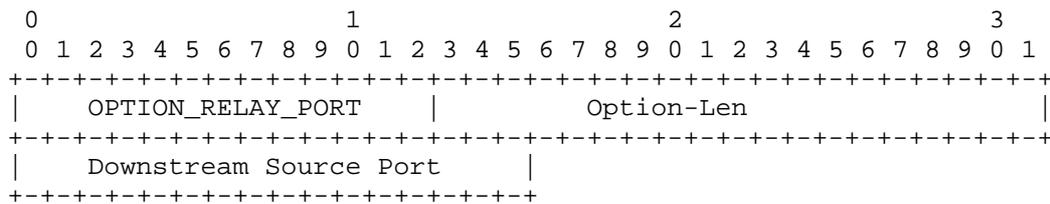
SubOpt Code: SUBOPT_RELAY_PORT. 8-bit value, 19.

Len: 8-bit 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 by either 1) a DHCPv6 relay agent that uses a non-DHCP UDP port (not 547) communicating with the IPv6 server and the upstream relay agent or 2) an 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:



Where:

Option-Code: OPTION_RELAY_PORT. 16-bit value, 135.

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

Downstream Source Port: 16-bit 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) to communicate 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 the port is not being used by a downstream relay agent, the 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 relay agents to specify a source port 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, the operator and/or user of the new DHCP relay port implementation should upgrade the DHCP server before the relay implementations are deployed. This would ensure that the erroneous case noted in Section 5.3 is not encountered. If the upstream relay agent or server does not support this extension, this DHCP relay port feature needs to be disabled.

When the DHCP relay port implementation is deployed, the default relay agent behavior should use the DHCP UDP port, and it is recommended that the configuration be set up to allow for the mode of operation where a non-DHCP port can be used for the DHCP relay agents.

If the network uses a 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 to match only the DHCP server side of the UDP port number and, if necessary, IP addresses and other attributes.

6. Example of an IPv6-Cascaded Relay

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

```

      (forward)           (forward)           (forward)
Relay1 -----> Relay2 -----> Relay3 -----> Server
      (1000)              (547)              (547)
      (reply)            (reply)            (reply)
      <-----           <-----           <-----

```

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 that 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 the Relay-forward message. Relay1 sets the "Downstream Source Port" field in the option to zero. Relay2 notices that 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 that 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 that 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, but 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 the DHCP server sends the DHCP Relay-reply to Relay3, Relay3 finds that 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. Relay2 finds its own relay option also has this "Downstream Source Port" with the value of 1000. Relay2 will use this UDP port when sending the Relay-reply message to Relay1.

7. IANA Considerations

IANA has registered "DHCPv4 Relay Source Port Sub-option" (19) for the IPv4 "Relay Agent Information" option within the "DHCP Relay Agent Sub-Option Codes" registry <<https://www.iana.org/assignments/bootp-dhcp-parameters>> [RFC3046].

IANA has registered "OPTION_RELAY_PORT" (135) for the DHCPv6 "Relay Source Port Option" within the DHCPv6 "Option Codes" registry <<http://www.iana.org/assignments/dhcpv6-parameters>> [RFC3315].

8. Security Considerations

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

9. 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>>.
- [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>>.

Acknowledgments

The authors would like to thank Peter Arberg, Luyuan Fang, Bhanu Gopalasetty, Scott Kelly, Andre Kostur, Victor Kuarsingh, Ted Lemon, Adam Roach, 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".

Authors' Addresses

Naiming Shen
Cisco Systems
560 McCarthy Blvd.
Milpitas, CA 95035
United States of America

Email: naiming@cisco.com

Enke Chen
Cisco Systems
560 McCarthy Blvd.
Milpitas, CA 95035
United States of America

Email: enkechen@cisco.com

