

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 21, 2013

Y. Cui
P. Wu
J. Wu
Tsinghua University
T. Lemon
Nominum, Inc.
March 20, 2013

DHCPv4 over IPv6 Transport
draft-ietf-dhc-dhcpv4-over-ipv6-06

Abstract

In IPv6 networks, there remains a need to provide IPv4 service for some residual devices. This document describes a mechanism for allocating IPv4 addresses to such devices, using DHCPv4 with an IPv6 transport. It is done by putting a special relay agent function (Client Relay Agent) on the client side, as well as extending the behavior of the server; in the case where DHCP server only supports IPv4 transport, a relay agent is extended to support IPv6 transport (IPv6-Transport Relay Agent) and relay DHCP traffic for the server, with a new Relay Agent Information sub-option added to carry the IPv6 address of the Client Relay Agent.

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 September 21, 2013.

Copyright Notice

Copyright (c) 2013 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
2.	Requirements Language	3
3.	Terminology	3
4.	Protocol Summary	4
5.	Client Relay Agent IPv6 Address Sub-option	6
6.	Client Relay Agent Behavior	6
7.	IPv6-Transport Server Behavior	7
8.	IPv6-Transport Relay Agent Behavior	8
9.	Security Consideration	8
10.	IANA consideration	9
11.	Contributors	9
12.	References	9
12.1.	Normative References	9
12.2.	Informative References	10
Appendix A.	Motivation for selecting this particular solution	10
A.1.	Configuring IPv4 with DHCPv6	11
A.2.	Tunnel DHCPv4 over IPv6	11
A.3.	DHCPv4 relayed over IPv6	12
Appendix B.	Discussion on One Host Retrieving Multiple Addresses through One CRA	12
Authors' Addresses	13

1. Introduction

DHCPv4 [[RFC2131](#)] was not designed with IPv6 in mind: DHCPv4 cannot operate on an IPv6 network. However, as dual-stack networks become a reality, the need arises to allocate IPv4 addresses in an IPv6 environment. To meet this demand, this document extends the DHCPv4 protocol to allow the use of an IPv6 network for transport.

A typical scenario that probably requires this feature is IPv4-over-IPv6 hub and spoke tunnel [[RFC4925](#)]. In this scenario, IPv4-over-IPv6 tunnel is used to provide IPv4 connectivity to end users (hosts or end networks) across an IPv6 network. If the IPv4 addresses of the end users are provisioned by the concentrator side, then the provisioning process should be able to cross the IPv6 network. One such tunnel mechanism is demonstrated in [[I-D.ietf-softwire-public-4over6](#)]. DHCPv4 over IPv6 would be a generic solution for this scenario.

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 [[RFC2119](#)].

3. Terminology

This document makes use of the following terms:

- o DHCPv4: IPv4 Dynamic Host Configuration Protocol [[RFC2131](#)].
- o Client Relay Agent(CRA): a special DHCPv4 Relay Agent which relays between DHCPv4 client and DHCPv4 server using an IPv6 network. A CRA either sits on the same, IPv6-accessible host with the DHCPv4 client, or sits on the same link with the host running DHCPv4 client.
- o Host Client Relay Agent(HCRA): a CRA which sits on the same, IPv6-accessible host with the DHCPv4 client.
- o On-Link Client Relay Agent(LCRA): a CRA which sits on the same link with the host that runs DHCPv4 client.
- o IPv6-Transport Server(TSV): a DHCPv4 Server that supports IPv6 transport. TSV listens on IPv6 for incoming DHCPv4 messages, and sends DHCPv4 messages in IPv6 packets.

- o IPv6-Transport Relay Agent(TRA): a DHCPv4 Relay Agent that supports IPv6 transport. TRA sits on a machine which has both IPv6 and IPv4 connectivity, and relays DHCP messages between CRA and regular DHCPv4 server. Unlike CRA, TRA sits on the remote end of IPv6 network, and communicates with DHCPv4 server through IPv4.
- o Client Relay Agent IPv6 Address Sub-option (CRA6ADDR sub-option): a new sub-option of the DHCP Relay Agent Information Option [[RFC3046](#)], defined in this document, which is used to carry the IPv6 address of the CRA.

4. Protocol Summary

The scenario for DHCPv4 over IPv6 transport is shown in Figure 1. DHCPv4 clients and DHCPv4 server/relay are separated by an IPv6 network in the middle. DHCP messages between a client and the server/relay cannot naturally be forwarded to each other because they are IPv4 UDP packets, either unicast or broadcast. To bridge this gap, both the client side and the server/relay side must enable DHCPv4 over IPv6 transport. More precisely, they must support delivering and receiving DHCP messages in IPv6 UDP packets and thereby traverse the IPv6 network.

On the client side, a special relay agent called Client Relay Agent is placed on the same host with the client, or on the link of the host. CRA is used to relay DHCP messages from the client to the server, and from the server to the client. CRA sends DHCPv4 messages to the server through unicast IPv6 UDP, and receives unicast IPv6 UDP packets with the DHCPv4 messages from the server. By using CRA, no extension is required on the DHCP client.




Figure 1 Scenario of DHCPv4 over IPv6 Transport

The IPv6-Transport DHCPv4 server can receive DHCP messages delivered in IPv6 UDP from CRA, and send out DHCP messages to CRA using IPv6 UDP (figure 2(a)). TSV should send DHCP messages to the IPv6 address from which it receives relevant DHCP messages earlier.

When CRAs communicate with an IPv6-Transport Relay Agent rather than with a server directly, the situation becomes a little more complicated. Besides the IPv6 communication with CRA, TRA also communicates with a regular DHCPv4 server through IPv4. Therefore, when TRA relays DHCP messages between a CRA and the DHCPv4 server, it receives DHCP message from the CRA in IPv6 and sends it to the server in IPv4, as well as receives DHCP message from the server in IPv4 and sends it to the CRA in IPv6.

TRA sends the IPv6 address of the CRA to the DHCP server using the Client Relay Agent IPv6 Address suboption, defined in this document. The DHCP server returns this suboption to the TRA as required in [\[RFC3046\]](#). The TRA then uses the returned CRA6ADDR suboption to determine the destination address to which to relay the response.

```

+-----+           +-----+
|client| IPv6 network |TSV  |
|+HCRA |-----|      |
+-----+           +-----+

```

```

+-----+ +-----+           +-----+
|client| |LCRA | IPv6 network |TSV  |
|      |--|      |-----|      |
+-----+ +-----+           +-----+

```

(a)client--server case

```

+-----+           +-----+           +-----+
|client| IPv6 network |TRA  | IPv4 network |server|
|+HCRA |-----|      |-----|      |
+-----+           +-----+           +-----+

```

```

+-----+ +-----+           +-----+           +-----+
|client| |LCRA | IPv6 network |TRA  | IPv4 network |server|
|      |--|      |-----|      |-----|      |
+-----+ +-----+           +-----+           +-----+

```

(b)client--relay--server case

Figure 2 Protocol Summary

5. Client Relay Agent IPv6 Address Sub-option

The CRA6ADDR suboption is a suboption of the Relay Agent Information Option [RFC3046]. It encodes the IPv6 address of the machine from which a DHCPv4-in-IPv6 CRA-to-TRA message was received. It is used by the TRA to relay DHCPv4 replies back to the proper CRA. The TRA uses the IPv6 address encoded in this suboption as the destination IPv6 address when relaying a DHCPv4 message from the DHCP server to the CRA.

The CRA6ADDR sub-option has a fixed length of 18 octets. The SubOpt code is tbd by IANA, the length field is 16, and the following 16 octets contain the CRA IPv6 address.




Figure 3 Client Relay Agent IPv6 Address Sub-option format

6. Client Relay Agent Behavior

A Client Relay Agent sits on the same host with the DHCPv4 client (HCRA), or on the same link as the host (LCRA). CRA listens for DHCP packets on IPv4 on port 67, and also listens for DHCP packets on IPv6 on port 67.

A CRA is configured with one or more IPv6 addresses of TSV/TRA, using a DHCPv6 option or some other mechanism. The CRA cannot forward DHCPv4 messages before it is configured with an IPv6 address itself, so to function properly, the IPv6 address for the CRA SHOULD be configured before the DHCPv4 client starts. The CRA SHOULD use a global IPv6 address.

When the CRA receives any DHCP message on IPv4 with BOOTP op field set to 1, it forwards the message over UDP on IPv6 using a standard DHCP message format, with source port 67 and destination port 67. The CRA forwards the message to each TSV or TRA address with which it is configured.

When the CRA receives any message on IPv6 with BOOTP op field set to 2, the CRA checks to see if the message contains option 82. If it does, the CRA silently discards the message. Otherwise, it relays the message to the DHCP client using IPv4.

When the CRA receives any message on IPv6 with BOOTP op field set to 4, it decapsulates the message as specified in DHCPv4 Relay Agent Encapsulation [[I-D.ietf-dhc-dhcpv4-relay-encapsulation](#)]. If the CRA does not support encapsulation, it MUST silently discard the message.

The LCRA or HCRA MUST NOT use the Relay Agent Information Option [[RFC3046](#)]. If either type of CRA needs to send relay agent options, it MUST use relay agent encapsulation as defined in [[I-D.ietf-dhc-dhcpv4-relay-encapsulation](#)].

An HCRA MUST only serve the client inside the same host, while the LCRA SHOULD serve any client on the link. When the IPv6 address of TSV/TRA is provisioned to the host running the DHCP client, it uses HCRA; else the client depends on LCRA. A HCRA serves only one link; the multiple link case MUST be handled by multiple HCRA instances. A LCRA does not necessarily need an IPv4 address, though it may be configured with one.

In HCRA case, the DHCPv6 client (or other IPv6 configuration processes), DHCPv4 client and CRA runs on the same physical interface. If possible, the host running the DHCPv4 client and CRA SHOULD defer the operation of the DHCPv4 client until an IPv6 address of the interface has been acquired, as well as the TSV/TRA address information. If this is not done, the DHCPv4 client may send several messages that the CRA cannot relay, and this could result in long delays before the DHCPv4 client actually gets an IPv4 address.

7. IPv6-Transport Server Behavior

To support IPv6 transport, the behavior of DHCPv4 server is extended. The IPv6-Transport Server can listen on IPv6 port 67 for DHCPv4 messages, and send DHCPv4 messages through IPv6.

A TSV listens for DHCP messages on IPv6 UDP port 67 and IPv4 UDP port 67. When it receives a DHCP message on IPv6, it MUST retain the IPv6 source address of that message until it has sent a response. When it sends a response, it MUST send the response to this IPv6 address, with destination port 67.

The TSV MUST send a server identifier option [[RFC2132](#)] containing an IPv4 address which will be reachable from the client once the residual IPv4 service is set up. This follows the server id option requirement in [[RFC2131](#)].

The rest of TSV DHCP process is the same with normal DHCPv4 server. A TSV MUST also listen on IPv4 UDP port 67 like a normal DHCPv4 server, and process IPv4 DHCPv4 messages normally. This requirement

exists because when a DHCPv4 client renews, it sends its renewal messages directly to the server, rather than broadcasting them.

Because the CRA may use relay agent encapsulation [[I-D.ietf-dhc-dhcpv4-relay-encapsulation](#)], the TSV SHOULD support it. A TSV that does not support it will not interoperate with a CRA that sends relay agent options.

8. IPv6-Transport Relay Agent Behavior

An IPv6-Transport Relay Agent sits between IPv6 network and IPv4 network, and relays DHCPv4 message between CRAs and IPv4-only DHCPv4 server. The communication between CRAs and the TRA uses IPv6, while the communication between the TRA and the server uses IPv4. A TRA listens on IPv6 UDP port 67 for DHCP messages with BOOTP op field set to 1 or 3, as well as IPv4 UDP port 67 for DHCP messages with BOOTP op field set to 2 or 4.

When relaying a DHCP message from CRA to server, TRA MUST add a CRA6ADDR suboption. The TRA sets the contents of this suboption to the IPv6 source address of the message. The TRA MUST also store one its own IPv4 addresses in the giaddr field of the DHCP message. The TRA MAY include a Link Selection sub-option [[RFC3527](#)] to indicate to the DHCP server which link to use when choosing an IP address. If the received message is a RELAYFORWARD message, the TRA MUST encapsulate the message in a new RELAYFORWARD message and store the CRA6ADDR in the new relay segment. If it is some other message, the TRA SHOULD append a Relay Agent Information Option as described in [[RFC3046](#)], but MAY encapsulate it in the same way as RELAYFORWARD message instead.

When receiving a DHCP message from the DHCP server, if the message contains no CRA6ADDR suboption, the TRA MUST discard the message. Otherwise, it processes it as required by [[RFC3046](#)] and [[I-D.ietf-dhc-dhcpv4-relay-encapsulation](#)], and forwards it to the IPv6 address recorded in the CRA6ADDR sub-option, with source port 67 and destination port number 67.

9. Security Consideration

This mechanism may rise a new form of DHCP protocol attack. A malicious attacker in IPv6 can interference with the DHCPv4 process by inject fake DHCPv4-in-IPv6 messages which will be handled by TSV or TRA. However, the damage is the same with the known DHCPv4 attack happened in IPv4. The only difference is the attacker and the victim could locate in different address families.

Another impact is DHCP filtering. There are firewalls today capable of filtering DHCP traffic (DHCPv4 over IPv4 and DHCPv6 over IPv6 packages). The DHCP messages with the new, DHCPv4-in-IPv6 style may bypass these firewalls. Nevertheless it is not difficult for them to make some slight modification and adapt to the new DHCPv4 message pattern.

10. IANA consideration

IANA is requested to assign one new sub-option code from the registry of DHCP Agent Sub-Option Codes maintained in <http://www.iana.org/assignments/bootp-dhcp-parameters>. This sub-option code will be assigned to the Client Relay Agent IPv6 Address Sub-option.

11. Contributors

The following gentlemen also contributed to the effort:

Francis Dupont
Internet Systems Consortium, Inc.

Email: fdupont@isc.org

Tomasz Mrugalski
Internet Systems Consortium, Inc.

Email: tomasz.mrugalski@gmail.com

Dmitry Anipko
Microsoft Corporation

Email: danipko@microsoft.com

12. References

12.1. Normative References

[I-D.ietf-dhc-client-id]
Swamy, N., Halwasia, G., and S. Unit, "Client Identifier Option in DHCP Server Replies",

[draft-ietf-dhc-client-id-07](#) (work in progress),
November 2012.

- [I-D.ietf-dhc-dhcpv4-relay-encapsulation]
Lemon, T., Deng, H., and L. Huang, "Relay Agent
Encapsulation for DHCPv4",
[draft-ietf-dhc-dhcpv4-relay-encapsulation-01](#) (work in
progress), July 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol",
[RFC 2131](#), March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor
Extensions", [RFC 2132](#), March 1997.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option",
[RFC 3046](#), January 2001.
- [RFC3527] Kinnear, K., Stapp, M., Johnson, R., and J. Kumarasamy,
"Link Selection sub-option for the Relay Agent Information
Option for DHCPv4", [RFC 3527](#), April 2003.
- [RFC4361] Lemon, T. and B. Sommerfeld, "Node-specific Client
Identifiers for Dynamic Host Configuration Protocol
Version Four (DHCPv4)", [RFC 4361](#), February 2006.
- [RFC4925] Li, X., Dawkins, S., Ward, D., and A. Durand, "Softwire
Problem Statement", [RFC 4925](#), July 2007.

[12.2.](#) Informative References

- [I-D.ietf-softwire-public-4over6]
Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public
IPv4 over IPv6 Access Network",
[draft-ietf-softwire-public-4over6-05](#) (work in progress),
February 2013.

[Appendix A.](#) Motivation for selecting this particular solution

We considered three possible solutions to the problem of configuring
IPv4 addresses on an IPv6 network.

A.1. Configuring IPv4 with DHCPv6

Use DHCPv6 instead of DHCPv4, to provision IPv4-related connectivity. In DHCPv6, the provisioned IPv4 address can be embedded into IPv6 address, or carried within a new option. Along with that, dedicated options are needed to convey IPv4-related information, such as the IPv4 address of DNS server, NTP server, etc. Therefore it will put a certain amount of IPv6-unrelated information into DHCPv6 protocol.

This solution was rejected for two reasons. First, the DHCPv6 protocol does not currently provide a mechanism for recording bindings between IPv4 addresses and DHCPv6 clients. Extending DHCPv6 to provide this functionality would be a substantial change to the existing protocol.

Second, a deliberate choice was made when the DHCPv6 protocol was defined to avoid simply copying existing functionality from DHCPv4. While it is possible, using DHCPv6, to deliver IPv4 addresses as IPv6-encoded IPv4 addresses, it might be necessary to add additional DHCPv6 options simply to support IPv4. These options would then remain in the protocol, long after the need for IPv4 has gone.

By comparison, any extensions to DHCPv4 will naturally be forgotten when DHCPv4 is no longer needed. This means that whatever extensions we make to DHCPv4 to solve the problem, we can stop maintaining as soon as IPv4 is no longer needed.

A.2. Tunnel DHCPv4 over IPv6

Use DHCPv4 for configuration, and tunnel DHCPv4-in-IPv6 messages over IPv6. Unlike the previous approach where DHCPv6 is used for both IPv4 and IPv6 connectivity, this approach preserves the separation between IPv4 and IPv6 connectivity information. It maintains the IPv4 service without major modifications to IPv6-related provisioning resources, and sustains DHCPv4 to be the IPv4-related information carrier.

This approach was not chosen because it adds a requirement for DHCPv4 to operate over an IPv4-in-IPv6 tunnel. DHCPv4 clients generally operate on broadcast networks, not on tunnels. To make DHCPv4 operate over a tunnel would require substantial changes to the DHCPv4 client, as well as maintaining a tunnel over which to deliver DHCPv4 traffic.

This also creates a chicken-and-egg problem: how do we set up an IPv4 tunnel when we do not know our IPv4 address? Solutions to these problems were proposed, but they require significant changes to the DHCP client and significant additional work to make a tunnel that

could carry the DHCP packets.

A.3. DHCPv4 relayed over IPv6

Use DHCPv4 for configuration, and extend it to use an IPv6 transport for relayed messages. Essentially this involves a single change to the protocol, to allow DHCPv4 servers or relay agents to send and receive packets using an IPv6 transport. No changes are required on the client.

The working group chose this third solution because, of the three, it required the fewest changes to the DHCP protocol, so that it was easiest to specify and easiest to implement.

Appendix B. Discussion on One Host Retrieving Multiple Addresses through One CRA

This document is written with the intention of supporting a use case where a single DHCP client is configuring a single tunnel endpoint per physical link. The technique described in this document could be used by a host needing to configure more than one tunnel endpoint on the same physical link, i.e., to retrieve multiple addresses through the same CRA. However, the following additional behavior is REQUIRED to support this case.

DHCP server implementing this specification MUST implement Client Identifier Option in DHCP server replies [[I-D.ietf-dhc-client-id](#)].

In general this specification is intended not to require modification of DHCP clients. However, DHCP clients being used to configure multiple tunnel endpoints have to be modified; otherwise there is no way for such DHCP clients to differentiate between DHCP responses. Therefore, in such case, the DHCP client using this specification MUST use a different client identifier for each tunnel endpoint being configured. Such DHCP clients MUST examine the response from the DHCP server and use the client identifier to differentiate between the DHCP client state machines for each tunnel endpoint.

In order to satisfy the requirement that client identifiers be unique, DHCP clients configuring multiple tunnel endpoints MUST implement Node-specific Client Identifiers for DHCPv4 [[RFC4361](#)]. Such clients MUST use a different IAID for each tunnel endpoint.

It is assumed here that every client state machine on a multiple-tunnel-endpoint link can hear all the DHCP messages (and subsequently accept the messages intended for it). How this is accomplished is left to the implementor. However, implementations MUST follow this

requirement; otherwise, it will be impossible for multiple tunnel endpoints to be successfully configured. The easiest way to accomplish this is to have a single DHCP client process with multiple DHCP state machines, and to dispatch each DHCP message to the correct DHCP client state machine using the client identifier. However, this is not REQUIRED; any mechanism that results in client state machines receiving the messages that are intended for them will suffice.

Authors' Addresses

Yong Cui
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6260-3059
Email: yong@csnet1.cs.tsinghua.edu.cn

Peng Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6278-5822
Email: pengwu.thu@gmail.com

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6278-5983
Email: jianping@cernet.edu.cn

Ted Lemon
Nominum, Inc.
2000 Seaport Blvd
Redwood City, CA 94063
USA

Phone: +1-650-381-6000
Email: mellon@nominum.com