Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 30, 2012

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

March 29, 2012

DHCPv4 over IPv6 Transport
draft-ietf-dhc-dhcpv4-over-ipv6-02

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 extending DHCP client and server behavior, and by adding a new Relay Agent Information option to carry the IPv6 address of the client.

Status of This Memo

Copyright Notice

Table of Contents

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 DHCPv4 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, too. One such tunnel mechanism is demonstrated in [I-D.ietf-softwire-public-4over6]. DHCPv4 over IPv6 would be a generic solution for this scenario.

Three main flavours of solutions may be considered:

- 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.

- Use DHCPv4 and tunnel DHCPv4-in-IPv4 messages over IPv6. Unlike the previous approach where DHCPv6 is used for both IPv4 and IPv6 connectivity, this approach consists in preserving the separation between IPv4 and IPv6 connectivity information. It allows to maintain the IPv4 service without major modification of IPv6-related provisioning resources, and sustains DHCPv4 to be the IPv4-related information carrier. However, this approach enforces an IPv4-in-IPv6 tunnel on DHCP, and requires extra efforts to maintain tunnel endpoint information for encapsulation use.

- Use DHCPv4 and extend it to work over IPv6 transport. Instead of relying on IPv4-in-IPv6 tunnel, this flavour uses IPv6 directly for DHCP message transport, and it keeps the advantage of separation with IPv6 connectivity information. This document focuses on this flavour. The document will define the extensions of DHCPv4 protocol behavior, as well as a new suboption of the Relay Agent Information Option, to fully support DHCPv4 over IPv6.

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:

- DHCPv4: IPv4 Dynamic Host Configuration Protocol [RFC2131].

- Client Relay Agent(CRA): a special DHCPv4 Relay Agent that sits on the same, IPv6-accessable host with the DHCPv4 client. CRA works as a "bridge" between DHCPv4 client and the IPv6 network, to convert between IPv4 transport and IPv6 transport.

- On-link Client Relay Agent(LCRA): a CRA sits on the link of the host rather then inside the host.

- IPv6-Transport Server(TSV): a DHCPv4 Server that supports IPv6 transport. TSV can listen on IPv6 for incoming DHCPv4 messages, and send DHCPv4 messages in IPv6 packets.

- 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 normal DHCPv4 server.

- Client Relay Agent IPv6 Address Sub-option(CRA6ADDR suboption): a new suboption of DHCP Relay Agent Information Option [RFC3046] defined in this document. CRA6ADDR suboption is used by TRA to carry the IPv6 address of a 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 by default IPv4 UDP packets, either unicast or broadcast. To bridge this gap, both the client side and the server/relay side should enable DHCPv4 over IPv6 transport. More precisely, they should 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 client. 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.

```
        +-------------------------+
     +------+                     |
     |DHCPv4|                     |
     |Client|                 +-------+
     +------+                 |DHCPv4 |
        |      IPv6 Network   |Server/|
     +------+                 |Relay  |
     |DHCPv4|                 +-------+
     |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 will become 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, while receives DHCP message from the server in IPv4 and sends it to the CRA in IPv6. TRA has to use the DHCP Relay Agent Information Option(Option 82) to record the IPv6 address of a CRA, which will be used as forwarding destination when relaying a DHCP message from the server. Since Option 82 doesn't have an existing suboption that fits in the case, this document defines a new Client Relay Agent IPv6 Address Sub-option.

```
     +------+                +------+
     |client|  IPv6 network  |TSV   |
     |+CRA  |----------------|      |
     +------+                +------+
     (a)client--server case

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

```
     |+CRA  |----------------|      |--------------|      |
     +-------+                +-------+              +-------+
     (b)client--relay--server case
```

Figure 2 Protocol Summary

5. Client Relay Agent IPv6 Address Sub-option

This suboption MUST be added by a DHCPv4 TRA. It encodes the IPv6 address of the host from which a DHCPv4-in-IPv6 CRA-to-TRA message was received. It is intended for the TRA to relay DHCPv4 replies back to the proper CRA. To be more specific, the TRA uses the IPv6 address encoded in this suboption as the destination IPv6 address when relaying a DHCPv4 reply to IPv6 network.

The CRA IPv6 address MUST be unique in the IPv6 domain.

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

DHCP servers handles it following the standard option 82 procedure defined in [RFC3046]. DHCP servers MAY use this suboption to select parameters specific to particular hosts. Servers MAY parse this suboption and extract the semantic of IPv6 address.

```
              SubOpt    Len     Agent Remote ID
             +------+------+------+------+------+-     -+------+
             | tbd  |  16  |  a1  |  a2  |  a3  |  ...  |  a16 |
             +------+------+------+------+------+-     -+------+
```

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. CRA is a special type of relay agent, which relays DHCPv4 messages between regular client and TSV/TRA. The communication between CRA and the client happens within the machine using IPv4, and the communication between CRA and TSV/TRA happens on the IPv6 network using IPv6.

A CRA is configured with one or more IPv6 addresses of TSV/TRA. This configuration is provided before DHCPv4 process, for example through DHCPv6 option, or by some other mechanisms depending on the

application scenarios.

A CRA listens for DHCP messages on IPv4 UDP port 67. When it receives from IPv4 any DHCP message with bootp op field = 1, it forwards the message using the standard DHCP relay agent format, but over UDPv6, with source port 67 and destination port 67. Here the CRA MUST NOT include an option 82 or modify the giaddr field of the DHCP message. The CRA forwards the message to each of the DHCP server or relay agent with which it is configured. The CRA MUST use a global IPv6 address if it has onbe.

A CRA also listens for DHCP messages on IPv6 UDP port 68. When it receives from IPv6 any DHCP message with bootp op field = 2, the CRA checks to see if the message contains option 82, and if so, it discards the message. Otherwise, it delivers the message to the DHCP client using IPv4.

A CRA can also sits on the link of the host (LCRA). The basic functionility of LCRA is the same with CRA. LCRA SHOULD NOT include a option 82 or modify the giaddr of DHCP message. If it has to, [I-D.ietf-dhc-dhcpv4-relay-encapsulation] can be used as the solution to the coexistence of LCRA and TRA. LCRA does not necessarily need an IPv4 address, though it may be configured with one.

A CRA MUST only serve the client inside the same host, while the LCRA MUST serve any client on the link. When the IPv6 address of TSV/TRA is provisioned, the DHCP client uses CRA; else the client uses LCRA.

7. IPv6-Transport Server Behavior

To support IPv6 transport, the behavior of DHCPv4 server should be 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. When it receives from IPv6 a DHCP message, it MUST record the IPv6 source address of that message and retain it as the return address of the message. That is to say, when sometime later the TSV responds to this message, it MUST send the reply message to the IPv6 return address retained earlier, with destination port 68. When filling in the server id option of DHCP replies, the TSV MUST choose an IPv4 address which is reachable from the client once the residual IPv4 serivice 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 can also listen on IPv4 UDP port 67 like a normal DHCPv4 server, and process normally when receives IPv4 DHCPv4 message.

This document places no new requirements on DHCPv4 servers that do not listen on UDPv6--in order to use an IPv4-only DHCPv4 server through an IPv6 connection, a TRA is required.

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 TRA and server uses IPv4. A TRA listens on IPv6 UDP port 67 for DHCP messages with bootp op field = 1, as well as IPv4 UDP port 68 for DHCP messages with bootp op field = 2.

When relaying a DHCP message from CRA to server, TRA MUST add an option 82 with a CRA6ADDR suboption. This suboption contains the IPv6 source address of the message (the CRA's IPv6 address) which is retained when the message is received in IPv6. The TRA MUST also store the IPv4 address of itself in the giaddr field of the DHCP message. The TRA MAY include a Link Selection Suboption [RFC3527] to indicate to the DHCP server which link to use when choosing an IP address.

When receiving a DHCP message from the DHCP server, if the option 82 in the message contains no CRA6ADDR suboption, the TRA MUST discard the message. Otherwise, it processes it as required by [RFC3046], and forwards it to the IPv6 address recorded in the CRA6ADDR suboption, with source port 67 and destination port number 68. TRA SHOULD drop DHCPv4-over-IPv6 traffic that is not originated from configured server address.

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 suboption code from the registry of DHCP Agent Sub-Option Codes maintained in http://www.iana.org/assignments/bootp-dhcp-parameters. This suboption code will be assigned to the Client Relay Agent IPv6 Address Sub-option.

11. Contributors

Francis Dupont and Tomek Mrugalski.

12. References

12.1. Normative References

[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.

[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.

[RFC4925] Li, X., Dawkins, S., Ward, D., and A. Durand, "Softwire Problem Statement", RFC 4925, July 2007.

12.2. Informative References

[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.

[I-D.ietf-softwire-public-4over6] Cui, Y., Wu, J., Wu, P., Metz, C., Vautrin, O., and Y. Lee, "Public IPv4 over IPv6 Access Network", draft-ietf-softwire-public-4over6-01 (work in progress), March 2012.

Authors' Addresses

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

Phone: +86-10-6260-3059
EMail: cuiyong@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 94063
USA

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