Softwire WG
Internet-Draft
Updates: RFC7568 (if approved)
Intended status: Standards Track
Expires: August 5, 2017

I. Farrer
Deutsche Telekom AG
Q. Sun
Y. Cui
L. Sun
Tsinghua University
February 1, 2017

DHCPv4 over DHCPv6 Source Address Option
draft-fsc-softwire-dhcp4o6-saddr-opt-07

Abstract

DHCPv4 over DHCPv6 [RFC7341] describes a mechanism for dynamically configuring IPv4 over an IPv6-only network. For DHCPv4 over DHCPv6 to function with some IPv4-over-IPv6 softwire mechanisms and deployment scenarios, the operator must learn the /128 IPv6 address that the client will use as the source of IPv4-in-IPv6 tunnel. This address, in conjunction with the IPv4 address and the Port Set ID allocated to the DHCP 4o6 client are used to create a binding table entry in the softwire tunnel concentrator. This memo defines two DHCPv6 options used to communicate the source tunnel IPv6 address between the DHCP 4o6 client and server. It also describes a method for configuring the client with the IPv6 address of the border router so that the softwire can be established. It is designed to work in conjunction with the IPv4 address allocation process.

Table of Contents

1. Introduction

Deterministic IPv4-over-IPv6 transition technologies require that elements are pre-configured with binding rules for routing traffic to clients. This places a constraint on the location of the client's tunnel endpoint: The tunnel endpoint has to be a pre-determined prefix which is usually be configured on the home gateway device. [RFC7597] describes a DHCPv6 based mechanism for provisioning such deterministic softwires.

A dynamic provisioning model, such as using DHCPv4 over DHCPv6 [RFC7341] allows much more flexibility in the location of the IPv4-over-IPv6 tunnel endpoint, as the IPv6 address is dynamically

signalled back to the service provider so that the corresponding tunnel configuration in the border router (BR) can be created. The DHCP 4o6 client and tunnel client could be run on end devices attached to any routable IPv6 prefix allocated to an end-user, located anywhere within an arbitrary home network topology. Dynamic allocation also helps to optimize IPv4 resource usage as only clients which are currently active are allocated IPv4 addresses.

This document describes a mechanism for dynamically provisioning softwires created using DHCPv4 over DHCPv6 (DHCP 4o6), including provisioning the client with the address of the softwire border router (BR) and informing the service provider of client's binding between the dynamically allocated IPv4 address and Port Set ID and the IPv6 address that the softwire Initiator will use for accessing IPv4-over-IPv6 services.

It is used with DHCP 4o6 message flows to communicate the binding over the IPv6-only network. The service provider can then use this binding information to provision other functional elements in their network accordingly, e.g. using the client's binding information to synchronise the binding table in the border router.

2. Applicability

The mechanism described in this document is only suitable for use for provisioning softwire clients via DHCP 4o6. The options described here are only applicable within the DHCP 4o6 message exchange process. Current softwire technologies suitable for extending to incorporate DHCPv4 over DHCPv6 with dynamic IPv4 address leasing include [RFC7597] and [RFC7596].

3. 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].

4. Solution Overview

The solution in this document is intended for the dynamic establishment of IPv4-over-IPv6 softwires. DHCP 4o6 [RFC7341] supports dynamically allocating (shared) IPv4 address. For a softwire to be successfully created, the IPv4 address has to be linked to the client's IPv6 tunnel source address. Within this process, the DHCP 4o6 client uses a DHCPv6 option to signal its tunnel source IPv6 address to the DHCP 4o6 server so that the operator's provisioning system can create the binding and configure the tunnel concentrator accordingly.

Two new DHCPv6 options are defined in this memo: OPTION_DHCP4O6_SADDR_HINT and OPTION_DHCP4O6_SADDR. They are intended to be used alongside the normal DHCPv4 IPv4 address allocation message flow in the context of DHCP 4o6. If a DHCP 4o6 client supports this mechanism, it MUST include the code of OPTION_DHCP4O6_SADDR_HINT in the Option Request Option (ORO) [RFC3315] when requesting IPv4 configuration through DHCP 4o6.

The communication of parameters between the client and server is a two-way process: OPTION_DHCP4O6_SADDR_HINT is optionally used by the DHCP 4o6 server to indicate to the client a preferred IPv6 prefix for binding the received IPv4 configuration and sourcing tunnel traffic. This may be necessary if there are multiple IPv6 prefixes in use in the customer network (e.g. ULAs), or if the specific IPv4-over-IPv6 transition mechanism requires the use of a particular prefix for any reason. When the client has selected an IPv6 address to bind the IPv4 configuration to, it passes the address back to the DHCP 4o6 server using OPTION_DHCP4O6_SADDR.

4.1. Provisioning the BR Address

To configure a softwire, the initiator also requires the IPv6 address of the BR. Section 4.2 of [RFC7598] defines option 90 (OPTION_S46_BR) for this purpose, but mandates that the option can only be used when encapsulated within one of the softwire container options: OPTION_S46_CONT_MAPE, OPTION_S46_CONT_MAPT or OPTION_S46_CONT_LW. From Section 3 of [RFC7598]:

> "Softwire46 DHCPv6 clients that receive provisioning options that are not encapsulated in container options MUST silently ignore these options."

This document updates [RFC7598] to remove this restriction for DHCPv6 option 90 (OPTION_S46_BR) allowing it to appear directly within the list of options in the client's ORO request and directly within subsequent messages sent by the DHCPv6 server.

5. IPv6/IPv4 Binding Message Flow

The following diagram shows the client/server message flow and how the options defined in this document are used. In each step, the relevant DHCPv4 message is given above the arrow and the relevant options below the arrow. All the DHCPv4 messages here are encapsulated in DHCPv4-query or DHCPv4-response messages, and those options are included in the 'options' field of the DHCPv4-query or DHCPv4-response message.

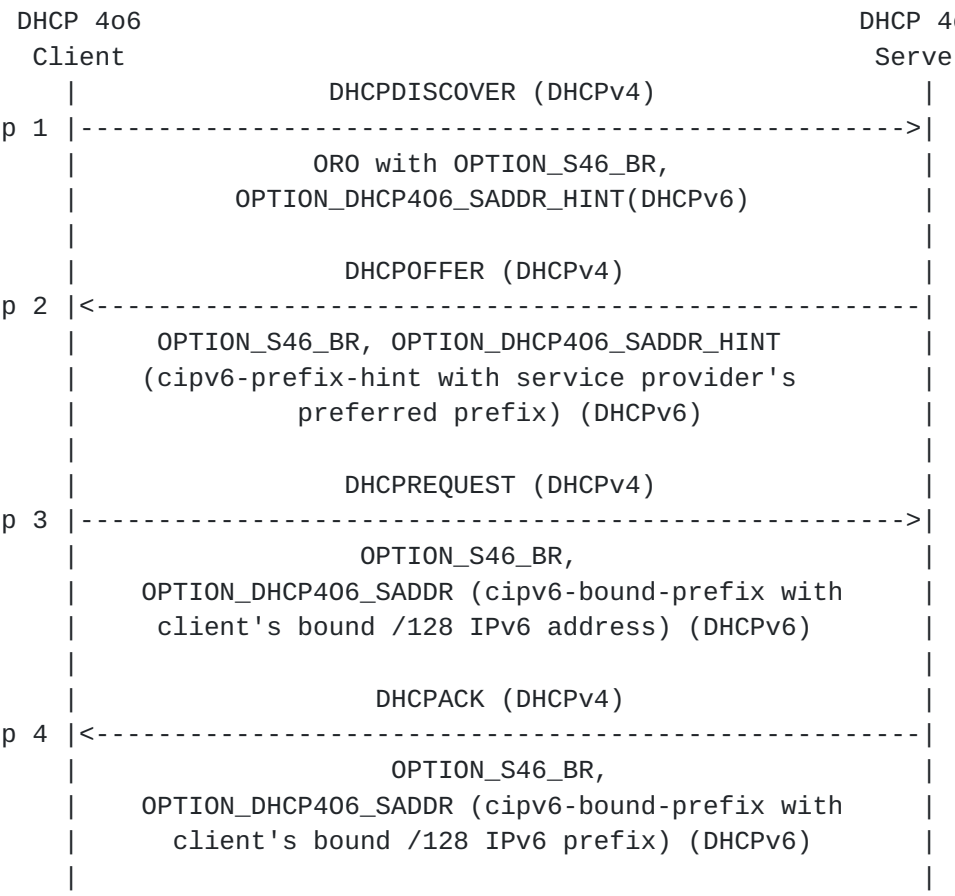

IPv6/IPv4 Binding Message Flow

A client attempting dynamic softwire configuration includes the option code for OPTION_BR_PREFIX, OPTION_DHCP4O6_SADDR_HINT in the DHCPv6 ORO in all DHCPv4-query messages it sends.

When a DHCP 4o6 Server replies with a DHCPOFFER message, it SHOULD include OPTION_S46_BR. It MAY also include OPTION_DHCP4O6_SADDR_HINT, which is used to indicate a preferred prefix that the client should use to bind IPv4 configuration to. If this option is received, the client MUST perform a longest prefix match between cipv6-prefix-hint and all prefixes/addresses in use on the device. If a match is found, the selected prefix MUST then be used to bind the received IPv4 configuration to and source the tunnel from. If no match is found, or the client doesn't receive OPTION_DHCP4O6_SADDR_HINT the client MAY select any valid IPv6 address to use as the tunnel source.

Once the client has selected which prefix it will use, it MAY use either an existing IPv6 address that is already configured on an interface, or create a new address specifically for use as the

softwire source address (e.g. using an Interface Identifier constructed as per Section 6 of [RFC7597]). If a new address is being created, the client MUST complete configuration of the new address, performing duplicate address detection (if required) before proceeding to Step 3.

OPTION_DHCP4O6_SADDR is used by the client to inform the DHCP 4o6 Server which IPv6 address the IPv4 configuration has been bound to. The client MUST put the selected IPv6 softwire source address into this option and include it in the DHCPv4-response message when it sends the DHCPREQUEST message.

6. DHCPv6 Options

6.1. DHCPv4 over DHCPv6 Source Address Hint Option

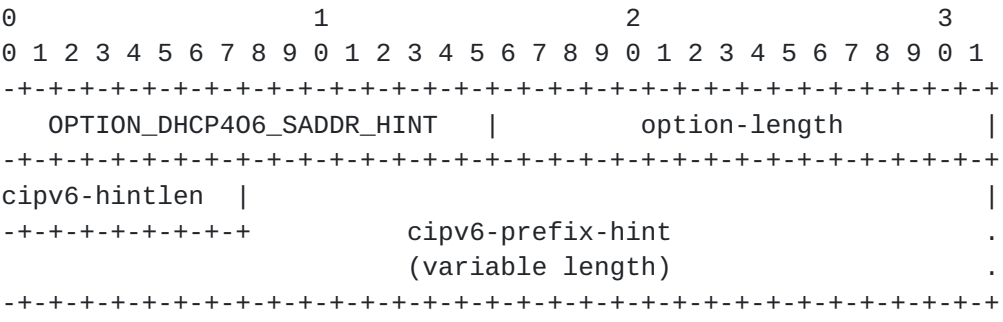

- option-code: OPTION_DHCP4O6_SADDR_HINT (TBA1)
- option-length: 1 + length of cipv6-prefix-hint, specified in bytes.
- cipv6-hintlen: 8-bit field expressing the bit mask length of the IPv6 prefix specified in cipv6-prefix-hint. Valid values are 0 to 128.
- cipv6-prefix-hint: The IPv6 prefix indicating the preferred prefix for the client to bind the received IPv4 configuration to. The length is (cipv6-hintlen + 7) / 8. The field is padded on the right with zero bits up to the nearest octet boundary when cipv6-prefix-hint is not evenly divisible by 8.

OPTION_DHCP4O6_SADDR_HINT is a singleton. Servers MUST NOT send more than one instance of the OPTION_DHCP4O6_SADDR_HINT option.

6.1.1. Client Option Validation Behavior

On receipt of the OPTION_DHCP4O6_SADDR_HINT option, the client makes the following validation checks:

- The received cipv6-hintlen value is not larger than 128.

- The received cipv6-hintlen value is not larger than the number of bytes sent in the cipv6-prefix-hint field. (e.g. the cipv6-hintlen is 128 but the cipv6-prefix-hint has only 8 bytes).

For either of these cases the receiver MAY either discard the option and proceed to attempt configuration as if the option had not been received, or attempt to use the received values for the long prefix match anyway.

The receiver MUST only use bits the cipv6-prefix-hint field up to the value specified in the cipv6-hintlen when performing the longest prefix match. cipv6-prefix-hint bits beyond this value MUST be ignored.

6.2. DHCPv4 over DHCPv6 Source Address Option

The format of DHCPv4 over DHCPv6 Source address option is defined as follows:

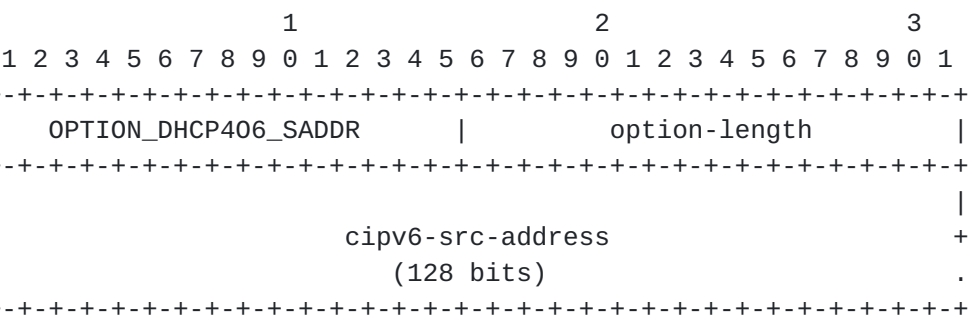

- option-code: OPTION_DHCP4O6_SADDR (TBA2)
- option-length: 16.
- cipv6-src-address: 16 bytes long; The IPv6 address that the client has bound the allocated IPv4 configuration to.

7. Security Considerations

Security considerations which are applicable to [RFC7341] are also applicable here.

8. IANA Considerations

IANA is requested to allocate a DHCPv6 option code for OPTION_DHCP4O6_SADDR_HINT and a DHCPv4 option code for OPTION_DHCP4O6_SADDR.

9. Acknowledgements

The authors would like to thank Ted Lemon, Lishan Li and Tatuya Jinmei for their contributions and comments.

10. References

10.1. 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>.

[RFC7341] Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4-over-DHCPv6 (DHCP 4o6) Transport", RFC 7341, DOI 10.17487/RFC7341, August 2014, <http://www.rfc-editor.org/info/rfc7341>.

10.2. Informative References

[I-D.farrer-softwire-br-multiendpoints] Farrer, I. and Q. Sun, "Multiple BR Tunnel Endpoint Addresses", draft-farrer-softwire-br-multiendpoints-01 (work in progress), July 2015.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <http://www.rfc-editor.org/info/rfc2131>.

[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>.

[RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <http://www.rfc-editor.org/info/rfc7596>.

[RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <http://www.rfc-editor.org/info/rfc7597>.

[RFC7598] Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for Configuration of Softwire Address and Port-Mapped Clients", RFC 7598, DOI 10.17487/RFC7598, July 2015, <http://www.rfc-editor.org/info/rfc7598>.

Authors' Addresses

Ian Farrer
Deutsche Telekom AG
CTO-ATI, Landgrabenweg 151
Bonn, NRW 53227
Germany

Email: ian.farrer@telekom.de

Qi Sun
Tsinghua University
Beijing 100084
P.R. China

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

Yong Cui
Tsinghua University
Beijing 100084
P.R. China

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

Linhui Sun
Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-5822
Email: lh.sunlinh@gmail.com