Internet Engineering Task Force (IETF)
Request for Comments: 8026
Category: Standards Track
ISSN: 2070-1721

M. Boucadair
Orange
I. Farrer
Deutsche Telekom AG

November 2016

Unified IPv4-in-IPv6 Softwire Customer Premises Equipment (CPE): A DHCPv6-Based Prioritization Mechanism

Abstract

In IPv6-only provider networks, transporting IPv4 packets encapsulated in IPv6 is a common solution to the problem of IPv4 service continuity. A number of differing functional approaches have been developed for this, each having their own specific characteristics. As these approaches share a similar functional architecture and use the same data plane mechanisms, this memo specifies a DHCPv6 option, whereby a single instance of Customer Premises Equipment (CPE) can interwork with all of the standardized and proposed approaches to providing encapsulated IPv4-in-IPv6 services by providing a prioritization mechanism.

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 http://www.rfc-editor.org/info/rfc8026.

Table of Contents

1. Introduction

IPv4 service continuity is one of the major technical challenges that must be considered during IPv6 migration. Over the past few years, a number of different approaches have been developed to assist with this problem (e.g., as described in [RFC6333], [RFC7596], and [RFC7597]). These approaches, referred to as "S46 mechanisms" in this document, exist in order to meet the particular deployment, scaling, addressing, and other requirements of different service providers' networks.

A common feature shared among all of the differing modes is the integration of softwire tunnel endpoint functionality into the Customer Premises Equipment (CPE) router. Due to this inherent data plane similarity, a single CPE may be capable of supporting several different approaches. Users may also wish to configure a specific mode of operation.

A service provider's network may also have more than one S46 mechanism enabled in order to support a diverse CPE population with differing client functionality, such as during a migration between mechanisms or where services require specific supporting softwire architectures.

For softwire-based services to be successfully established, it is essential that the customer's end node and the service provider's end node and provisioning systems are able to indicate their capabilities and preferred mode of operation.

A number of DHCPv6 options for the provisioning of softwires have been standardized:

RFC 6334 Defines DHCPv6 option 64 for configuring Basic Bridging BroadBand (B4) [RFC6333] elements with the IPv6 address of the Address Family Transition Router (AFTR) [RFC6333].

RFC 7341 Defines DHCPv6 option 88 for configuring the address of a DHCPv4-over-DHCPv6 server, which can then be used by a softwire client for obtaining further configuration.

RFC 7598 Defines DHCPv6 options 94, 95, and 96 for provisioning Mapping of Address and Port with Encapsulation (MAP-E) [RFC7597], Mapping of Address and Port using Translation (MAP-T) [RFC7599], and Lightweight 4over6 [RFC7596] respectively.

This document describes a DHCPv6-based prioritization method, whereby a CPE that supports several S46 mechanisms and receives configuration for more than one can prioritize which mechanism to use. The method requires no server-side logic to be implemented and only uses a simple S46 mechanism prioritization to be implemented in the CPE.

The prioritization method as described here does not provide redundancy between S46 mechanisms for the client. That is, if the highest priority S46 mechanism that has been provisioned to the client is not available for any reason, the means for identifying this and falling back to the S46 mechanism with the next highest priority is not in the scope of this document.

1.1. Terminology

This document makes use of the following terms:

- Address Family Transition Router (AFTR): The IPv4-in-IPv6 tunnel termination point and the Network Address Translator IPv4/IPv4 (NAT44) function deployed in the operator's network [RFC6333].

- Border Relay (BR): A MAP-enabled router managed by the service provider at the edge of a MAP domain. A BR has at least an IPv6-enabled interface and an IPv4 interface connected to the native IPv4 network [RFC7597].

- Customer Premises Equipment (CPE): Denotes the equipment at the customer edge that terminates the customer end of an IPv6 transitional tunnel. In some documents (e.g., [RFC7597]), this functional entity is called the Customer Edge (CE).

1.1.1. 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 [RFC2119].

1.2. Rationale

The following rationale has been adopted for this document:

(1) Simplified solution migration paths: Define unified CPE behavior, allowing for smooth migration between the different S46 mechanisms.

(2) Deterministic CPE coexistence behavior: Specify the behavior when several S46 mechanisms coexist in the CPE.

(3) Deterministic service provider coexistence behavior: Specify the behavior when several modes coexist in the service providers network.

(4) Reusability: Maximize the reuse of existing functional blocks including tunnel endpoints, the port-restricted Network Address Port Translator IPv4/IPv4 (NAPT44), forwarding behavior, etc.

(5) Solution agnostic: Adopt neutral terminology and avoid (as far as possible) overloading the document with solution-specific terms.

(6) Flexibility: Allow operators to compile CPE software only for the mode(s) necessary for their chosen deployment context(s).

(7) Simplicity: Provide a model that allows operators to only implement the specific mode(s) that they require without the additional complexity of unneeded modes.

1.3. DHCPv6 S46 Priority Option

The S46 Priority Option is used to convey a priority order of IPv4 service continuity mechanisms. Figure 1 shows the format of the S46 Priority Option.

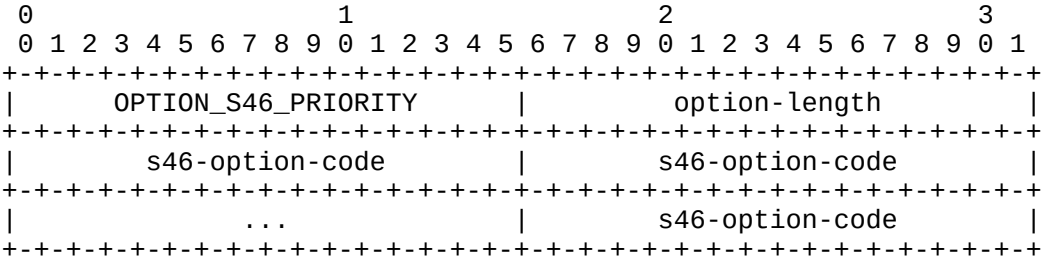

Figure 1: S46 Priority Option

- option-code: OPTION_S46_PRIORITY (111)
- option-length: >=2 and a multiple of 2, in octets.
- s46-option-code: 16-bit IANA-registered option code of the DHCPv6 option that is used to identify the softwire mechanism. S46 mechanisms are prioritized in the appearance order in the S46 Priority Option.

Codes in OPTION_S46_PRIORITY are processed in order; if a client receives more than one s46-option-code with a particular value, it should consider this case to be invalid. DHCP servers MAY validate the list of s46-option-code values to detect invalid values and duplicates. The option MUST contain at least one s46-option-code.

1.4. DHCPv6 Client Behavior

Clients MAY request the OPTION_S46_PRIORITY option, as defined in [RFC3315], Sections 17.1.1, 18.1.1, 18.1.3, 18.1.4, 18.1.5, and 22.7. As a convenience to the reader, we mention here that the client includes requested option codes in the Option Request Option.

Upon receipt of a DHCPv6 Advertise message from the server containing OPTION_S46_PRIORITY, the client performs the following steps:

1. Check the contents of the DHCPv6 message for options containing valid S46 mechanism configuration. A candidate list of possible S46 mechanisms is created from these option codes.

2. Check the contents of OPTION_S46_PRIORITY for the DHCPv6 option codes contained in the included s46-option-code fields. From this, an S46 mechanism priority list is created, ordered from highest to lowest following the appearance order.

3. Sequentially check the priority list against the candidate list until a match is found.

4. When a match is found, the client MUST configure the resulting S46 mechanism.

In the event that no match is found between the priority list and the candidate list, the client MAY proceed with configuring one or more of the provisioned S46 softwire mechanism(s). In this case, which mechanism(s) are chosen by the client is implementation specific and not defined here.

If an invalid OPTION_S46_PRIORITY option is received, the client MAY proceed with configuring the provisioned S46 mechanisms as if OPTION_S46_PRIORITY had not been received.

If an unknown option code is received in the OPTION_S46_PRIORITY option, the client MUST skip it and continue processing other listed option codes if they exist. The initial option codes that are allowed to be included in an OPTION_S46_PRIORITY option are listed in Section 4.1.

1.5. DHCPv6 Server Behavior

Sections 17.2.2 and 18.2 of [RFC3315] govern server operation in regard to option assignment. As a convenience to the reader, we mention here that the server will send a particular option code only if configured with specific values for that option code and if the client requested it.

Option OPTION_S46_PRIORITY is a singleton. Servers MUST NOT send more than one instance of the OPTION_S46_PRIORITY option.

2. Operator Deployment Considerations for Deploying Multiple Softwire Mechanisms

The following subsections describe some considerations for operators who are planning on implementing multiple softwire mechanisms in their network (e.g., during a migration between mechanisms).

2.1. Client Address Planning

As an operator's available IPv4 resources are likely to be limited, it may be desirable to use a common range of IPv4 addresses across all of the active softwire mechanisms. However, this is likely to result in difficulties in routing ingress IPv4 traffic to the correct Border Relay (BR) / AFTR instance, which is actively serving a given CE. For example, a client that is configured to use MAP-E may send its traffic to the MAP-E BR; however, on the return path, the ingress IP traffic gets routed to a MAP-T BR. The resulting translated packet that gets forwarded to the MAP-E client will be dropped.

Therefore, operators are advised to use separate IPv4 pools for each of the different mechanisms to simplify planning and IPv4 routing.

For IPv6 planning, there is less of a constraint as the BR/AFTR elements for the different mechanisms can contain configuration for overlapping the client's IPv6 addresses, provided that one mechanism is actively serving a given client at a time. However, the IPv6 address that is used as the tunnel concentrator's endpoint (BR/AFTR address) needs to be different for each mechanism to ensure correct operation.

2.2. Backwards Compatability with Existing Softwire Clients

Deployed clients that can support multiple softwire mechanisms, but do not implement the prioritization mechanism described here may require additional planning. In this scenario, the CPE would request configuration for all of the supported softwire mechanisms in its DHCPv6 Option Request Option (ORO), but would not request

OPTION_S46_PRIORITY. By default, the DHCPv6 server will respond with configuration for all of the requested mechanisms, which could result in unpredictable and unwanted client configuration.

In this scenario, it may be necessary for the operator to implement logic within the DHCPv6 server to identify such clients and only provision them with configuration for a single softwire mechanism. It should be noted that this can lead to complexity and reduced scalability in the DHCPv6 server implementation due to the additional DHCPv6 message processing overhead.

3. Security Considerations

Security considerations discussed in [RFC6334] and [RFC7598] apply for this document.

Misbehaving intermediate nodes may alter the content of the S46 Priority Option. This may lead to setting a different IPv4 service continuity mechanism than the one initially preferred by the network side. Also, a misbehaving node may alter the content of the S46 Priority Option and other DHCPv6 options (e.g., DHCPv6 Option 64 or 90) so that the traffic is intercepted by an illegitimate node. Those attacks are not unique to the S46 Priority Option but are applicable to any DHCPv6 option that can be altered by a misbehaving intermediate node.

4. IANA Considerations

IANA has allocated the following DHCPv6 option code:

111 OPTION_S46_PRIORITY

All values should be added to the DHCPv6 option code space defined in Section 24.3 of [RFC3315].

4.1. S46 Mechanisms and Their Identifying Option Codes

IANA has created a new registry titled "Option Codes permitted in the S46 Priority Option". This registry enumerates the set of DHCPv6 option codes that can be included in the OPTION_S46_PRIORITY option. Options may be added to this list using the IETF Review process described in Section 4.1 of [RFC5226].

The following table shows the option codes that are currently defined and the S46 mechanisms that they represent. The contents of this table shows the format and the initial values for the new registry. Option codes that have not been requested to be added according to the stated procedure should not be mentioned at all in the table, and

they should not be listed as "reserved" or "unassigned". The valid range of values for the registry is the range of DHCPv6 option codes (1-65535).

Option Code	S46 Mechanism	Reference
64	DS-Lite	[RFC6334]
88	DHCPv4 over DHCPv6	[RFC7341]
94	MAP-E	[RFC7598]
95	MAP-T	[RFC7598]
96	Lightweight 4over6	[RFC7598]

Table 1: DHCPv6 Option to S46 Mechanism Mappings

5. References

5.1. 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>.

[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>.

[RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", RFC 6334, DOI 10.17487/RFC6334, August 2011, <http://www.rfc-editor.org/info/rfc6334>.

[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>.

[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>.

5.2. Informative References

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <http://www.rfc-editor.org/info/rfc6333>.

[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>.

[RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <http://www.rfc-editor.org/info/rfc7599>.

Acknowledgements

Many thanks to O. Troan, S. Barth, A. Yourtchenko, B. Volz, T. Mrugalski, J. Scudder, P. Kyzivat, F. Baker, and B. Campbell for their input and suggestions.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes
France

Email: mohamed.boucadair@orange.com

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

Email: ian.farrer@telekom.de