

Internet Engineering Task Force (IETF)
Request for Comments: 7598
Category: Standards Track
ISSN: 2070-1721

T. Mrugalski
ISC
O. Troan
Cisco Systems
I. Farrer
Deutsche Telekom AG
S. Perreault
Jive Communications
W. Dec
Cisco Systems
C. Bao
Tsinghua University
L. Yeh
Freelancer Technologies
X. Deng
The University of New South Wales
July 2015

DHCPv6 Options for Configuration of Software Address
and Port-Mapped Clients

Abstract

This document specifies DHCPv6 options, termed Software46 options, for the provisioning of Software46 Customer Edge (CE) devices. Software46 is a collective term used to refer to architectures based on the notion of IPv4 Address plus Port (A+P) for providing IPv4 connectivity across an IPv6 network.

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 5741.

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/rfc7598>.

Copyright Notice

Copyright (c) 2015 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. Conventions	3
3. Software46 Overview	4
4. Common Software46 DHCPv6 Options	5
4.1. S46 Rule Option	5
4.2. S46 BR Option	7
4.3. S46 DMR Option	8
4.4. S46 IPv4/IPv6 Address Binding Option	9
4.5. S46 Port Parameters Option	10
5. Software46 Containers	11
5.1. S46 MAP-E Container Option	11
5.2. S46 MAP-T Container Option	12
5.3. S46 Lightweight 4over6 Container Option	13
6. Software46 Options Encapsulation	14
7. DHCPv6 Server Behavior	14
8. DHCPv6 Client Behavior	14
9. Security Considerations	15
10. IANA Considerations	16
11. References	16
11.1. Normative References	16
11.2. Informative References	17
Acknowledgements	18
Authors' Addresses	19

1. Introduction

A number of architectural solution proposals discussed in the IETF Software Working Group use Address plus Port (A+P) [RFC6346] as their technology base for providing IPv4 connectivity to end users using Customer Edge (CE) devices across a service provider's IPv6 network, while allowing for shared or dedicated IPv4 addressing of CEs.

An example is Mapping of Address and Port with Encapsulation (MAP-E) as defined in [RFC7597]. The MAP solution consists of one or more MAP Border Relay (BR) routers responsible for stateless forwarding between a MAP IPv6 domain and an IPv4 network, and one or more MAP Customer Edge (CE) routers responsible for forwarding between a user's IPv4 network and the MAP IPv6 network domain. Collectively, the MAP CE and BR form a domain when configured with common service parameters. This characteristic is common to all of the Software46 mechanisms.

To function in such a domain, a CE needs to be provisioned with the appropriate A+P service parameters for that domain. These consist primarily of the CE's IPv4 address and transport-layer port range(s). Furthermore, the IPv6 transport mode (i.e., encapsulation or translation) needs to be specified. Provisioning of other IPv4 configuration information not derived directly from the A+P service parameters is not covered in this document. It is expected that provisioning of other IPv4 configuration information will continue to use DHCPv4 [RFC2131].

This memo specifies a set of DHCPv6 [RFC3315] options to provision Software46 configuration information to CE routers. Although the focus is to deliver IPv4 service to an end-user network (such as a residential home network), it can equally be applied to an individual host acting as a CE. Configuration of the BR is out of scope for this document.

2. Conventions

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].

3. Software46 Overview

This document describes a set of common DHCPv6 options for configuring the Mapping of Address and Port with Encapsulation (MAP-E) [RFC7597], Mapping of Address and Port using Translation (MAP-T) [RFC7599], and Lightweight 4over6 [RFC7596] mechanisms. For definitions of the terminology used in this document, please see the relevant terminology sections in [RFC7597], [RFC7599], and [RFC7596].

MAP-E, MAP-T, and Lightweight 4over6 are essentially providing the same functionality: IPv4 service to a CE router over an IPv6-only access network. MAP-E and MAP-T may embed parts of the IPv4 address in IPv6 prefixes, thereby supporting many clients with a fixed set of mapping rules and Mesh mode (direct CE-to-CE communication). MAP-E and MAP-T CEs may also be provisioned in hub-and-spoke mode and in 1:1 mode (with no embedded address bits). The difference between MAP-E and MAP-T is that they use different means to connect to the IPv6 domain. MAP-E uses IPv4-over-IPv6 tunneling [RFC2473], while MAP-T uses IPv4-to-IPv6 translation based on [RFC6145]. Lightweight 4over6 is a hub-and-spoke IPv4-over-IPv6 tunneling mechanism, with complete independence of IPv4 and IPv6 addressing (zero embedded address bits).

The DHCPv6 options described here tie the provisioning parameters, and hence the IPv4 service itself, to the End-user IPv6 prefix lifetime. The validity of a Software46's IPv4 address, prefix, or shared IPv4 address; port set; and any authorization and accounting are tied to the lifetime of its associated End-user IPv6 prefix.

To support more than one mechanism at a time and to allow for a possibility of transition between them, the DHCPv6 Option Request Option (ORO) [RFC3315] is used. Each mechanism has a corresponding DHCPv6 container option. A DHCPv6 client can request a particular mechanism by including the option code for a particular container option in its ORO. The provisioning parameters for that mechanism are expressed by embedding the common format options within the respective container option.

This approach implies that all of the provisioning options appear only within the container options. Software46 DHCPv6 clients that receive provisioning options that are not encapsulated in container options MUST silently ignore these options. DHCPv6 server administrators are advised to ensure that DHCPv6 servers are configured to send these options in the proper encapsulation.

This document is organized with the common encapsulated options described first (Section 4), followed by the three container options (Section 5). Some encapsulated options are mandatory in some containers, some are optional, and some are not permitted. This is shown in Table 1 (Section 6).

4. Common Software46 DHCPv6 Options

The DHCPv6 protocol is used for Software46 CE provisioning following regular DHCPv6 notions, with the CE assuming the role of a DHCPv6 client, and the DHCPv6 server providing options following DHCPv6 server-side policies. The format and usage of the options are defined in the following subsections.

Each CE needs to be provisioned with enough information to calculate its IPv4 address, IPv4 prefix, or shared IPv4 address. MAP-E and MAP-T use the OPTION_S46_RULE option, while Lightweight 4over6 uses the OPTION_S46_V4V6BIND option. A CE that needs to communicate outside of the A+P domain also needs the address or prefix of the BR. MAP-E and Lightweight 4over6 use the OPTION_S46_BR option to communicate the IPv6 address of the BR. MAP-T forms an IPv6 destination address by embedding an IPv4 destination address into the BR's IPv6 prefix conveyed via the OPTION_S46_DMR option. Optionally, all mechanisms can include the OPTION_S46_PORTPARAMS option to specify parameters and port sets for the port-range algorithm.

Software46 options use addresses rather than Fully Qualified Domain Names (FQDNs). For the rationale behind this design choice, see Section 8 of [RFC7227].

4.1. S46 Rule Option

Figure 1 shows the format of the S46 Rule option (OPTION_S46_RULE) used for conveying the Basic Mapping Rule (BMR) and Forwarding Mapping Rule (FMR).

This option follows behavior described in Sections 17.1.1 and 18.1.1 of [RFC3315]. Clients can send those options, encapsulated in their respective container options, with specific values as hints for the server. See Section 5 for details. Depending on the server configuration and policy, it may accept or ignore the hints. Clients MUST be able to process received values that are different than the hints it sent earlier.

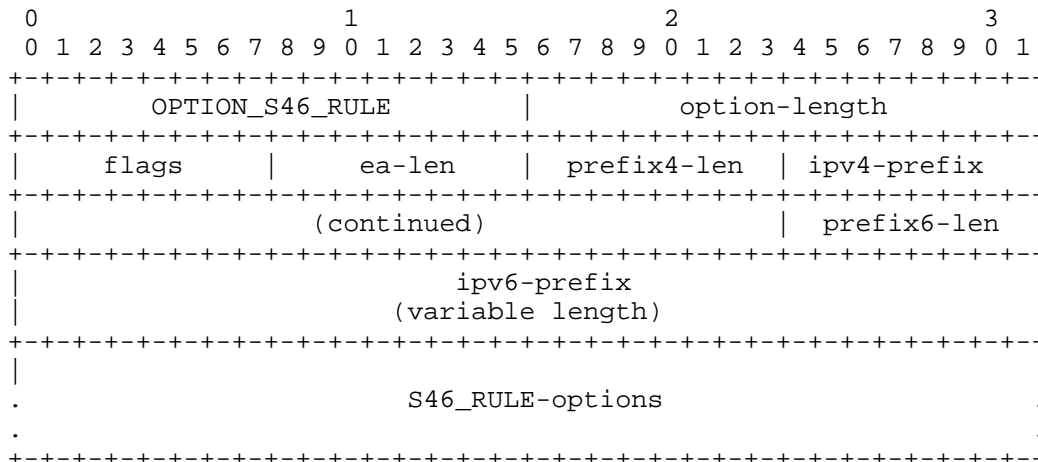


Figure 1: S46 Rule Option

- o option-code: OPTION_S46_RULE (89)
- o option-length: length of the option, excluding option-code and option-length fields, including length of all encapsulated options; expressed in octets.
- o flags: 8 bits long; carries flags applicable to the rule. The meanings of the specific bits are explained in Figure 2.
- o ea-len: 8 bits long; specifies the Embedded Address (EA) bit length. Allowed values range from 0 to 48.
- o prefix4-len: 8 bits long; expresses the prefix length of the Rule IPv4 prefix specified in the ipv4-prefix field. Allowed values range from 0 to 32.
- o ipv4-prefix: a fixed-length 32-bit field that specifies the IPv4 prefix for the S46 rule. The bits in the prefix after prefix4-len number of bits are reserved and MUST be initialized to zero by the sender and ignored by the receiver.
- o prefix6-len: 8 bits long; expresses the length of the Rule IPv6 prefix specified in the ipv6-prefix field. Allowed values range from 0 to 128.
- o ipv6-prefix: a variable-length field that specifies the IPv6 domain prefix for the S46 rule. The field is padded on the right with zero bits up to the nearest octet boundary when prefix6-len is not evenly divisible by 8.

- o S46_RULE-options: a variable-length field that may contain zero or more options that specify additional parameters for this S46 rule. This document specifies one such option: OPTION_S46_PORTPARAMS.

The format of the S46 Rule Flags field is:




Figure 2: S46 Rule Flags

- o Reserved: 7 bits; reserved for future use as flags.
- o F-flag: 1-bit field that specifies whether the rule is to be used for forwarding (FMR). If set, this rule is used as an FMR; if not set, this rule is a BMR only and MUST NOT be used for forwarding. Note: A BMR can also be used as an FMR for forwarding if the F-flag is set. The BMR is determined by a longest-prefix match of the Rule IPv6 prefix against the End-user IPv6 prefix(es).

It is expected that in a typical mesh deployment scenario there will be a single BMR, which could also be designated as an FMR using the F-flag.

4.2. S46 BR Option

The S46 BR option (OPTION_S46_BR) is used to convey the IPv6 address of the Border Relay. Figure 3 shows the format of the OPTION_S46_BR option.

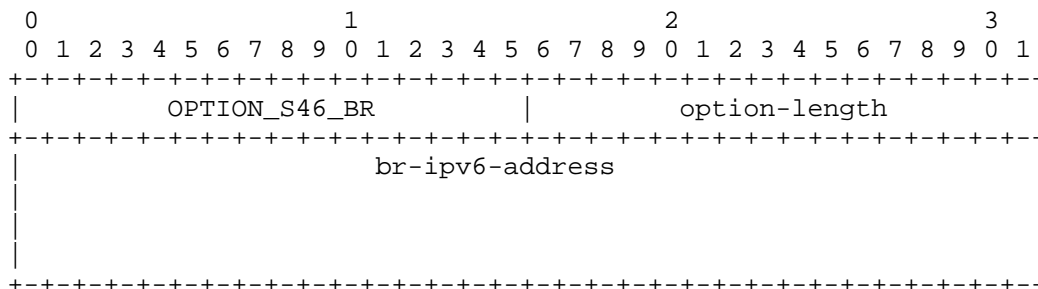


Figure 3: S46 BR Option

- o option-code: OPTION_S46_BR (90)
- o option-length: 16

- o br-ipv6-address: a fixed-length field of 16 octets that specifies the IPv6 address for the S46 BR.

BR redundancy can be implemented by using an anycast address for the BR IPv6 address. Multiple OPTION_S46_BR options MAY be included in the container; this document does not further explore the use of multiple BR IPv6 addresses.

4.3. S46 DMR Option

The S46 DMR option (OPTION_S46_DMR) is used to convey values for the Default Mapping Rule (DMR). Figure 4 shows the format of the OPTION_S46_DMR option used for conveying a DMR.

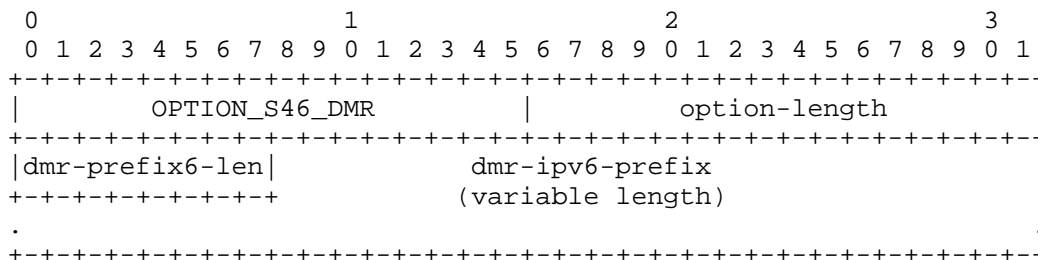


Figure 4: S46 DMR Option

- o option-code: OPTION_S46_DMR (91)
- o option-length: 1 + length of dmr-ipv6-prefix specified in octets.
- o dmr-prefix6-len: 8 bits long; expresses the bitmask length of the IPv6 prefix specified in the dmr-ipv6-prefix field. Allowed values range from 0 to 128.
- o dmr-ipv6-prefix: a variable-length field specifying the IPv6 prefix or address for the BR. This field is right-padded with zeros to the nearest octet boundary when dmr-prefix6-len is not divisible by 8.

4.4. S46 IPv4/IPv6 Address Binding Option

The S46 IPv4/IPv6 Address Binding option (OPTION_S46_V4V6BIND) MAY be used to specify the full or shared IPv4 address of the CE. The IPv6 prefix field is used by the CE to identify the correct prefix to use for the tunnel source.

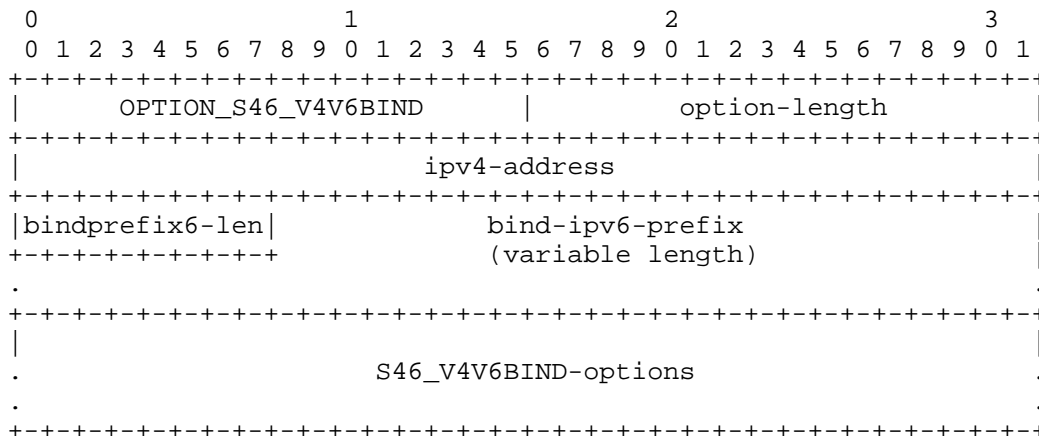


Figure 5: S46 IPv4/IPv6 Address Binding Option

- o option-code: OPTION_S46_V4V6BIND (92)
- o option-length: length of the option, excluding option-code and option-length fields, including length of all encapsulated options; expressed in octets.
- o ipv4-address: a fixed-length field of 4 octets specifying an IPv4 address.
- o bindprefix6-len: 8 bits long; expresses the bitmask length of the IPv6 prefix specified in the bind-ipv6-prefix field. Allowed values range from 0 to 128.
- o bind-ipv6-prefix: a variable-length field specifying the IPv6 prefix or address for the S46 CE. This field is right-padded with zeros to the nearest octet boundary when bindprefix6-len is not divisible by 8.
- o S46_V4V6BIND-options: a variable-length field that may contain zero or more options that specify additional parameters. This document specifies one such option: OPTION_S46_PORTPARAMS.

4.5. S46 Port Parameters Option

The S46 Port Parameters option (OPTION_S46_PORTPARAMS) specifies optional port set information that MAY be provided to CEs.

See Section 5.1 of [RFC7597] for a description of the MAP algorithm and detailed explanation of all of the parameters.




Figure 6: S46 Port Parameters Option

- o option-code: OPTION_S46_PORTPARAMS (93)
- o option-length: 4
- o offset: Port Set Identifier (PSID) offset. 8 bits long; specifies the numeric value for the S46 algorithm's excluded port range/offset bits (a-bits), as per Section 5.1 of [RFC7597]. Allowed values are between 0 and 15. Default values for this field are specific to the software mechanism being implemented and are defined in the relevant specification document.
- o PSID-len: 8 bits long; specifies the number of significant bits in the PSID field (also known as 'k'). When set to 0, the PSID field is to be ignored. After the first 'a' bits, there are k bits in the port number representing the value of the PSID. Consequently, the address-sharing ratio would be 2^k.
- o PSID: 16 bits long. The PSID value algorithmically identifies a set of ports assigned to a CE. The first k bits on the left of this field contain the PSID binary value. The remaining (16 - k) bits on the right are padding zeros.

When receiving the OPTION_S46_PORTPARAMS option with an explicit PSID, the client MUST use this explicit PSID when configuring its software interface. The OPTION_S46_PORTPARAMS option with an explicit PSID MUST be discarded if the S46 CE isn't configured with a full IPv4 address (e.g., IPv4 prefix).

The OPTION_S46_PORTPARAMS option is contained within an OPTION_S46_RULE option or an OPTION_S46_V4V6BIND option.

5. Software46 Containers

5.1. S46 MAP-E Container Option

The S46 MAP-E Container option (OPTION_S46_CONT_MAPE) specifies the container used to group all rules and optional port parameters for a specified domain.

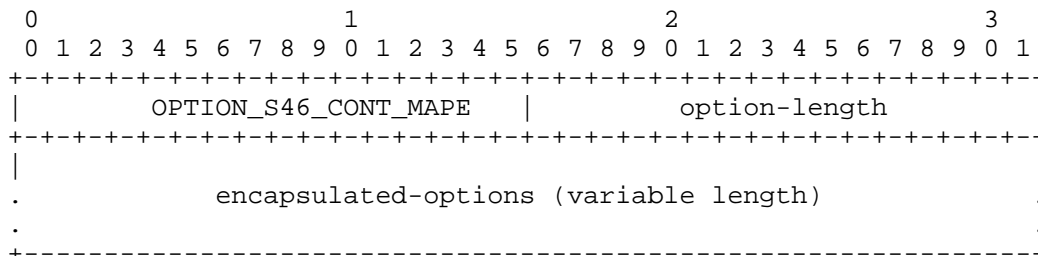


Figure 7: S46 MAP-E Container Option

- o option-code: OPTION_S46_CONT_MAPE (94)
- o option-length: length of encapsulated options, expressed in octets.
- o encapsulated-options: options associated with this Software46 MAP-E domain.

The encapsulated-options field conveys options specific to the OPTION_S46_CONT_MAPE option. Currently, there are two encapsulated options specified: OPTION_S46_RULE and OPTION_S46_BR. There MUST be at least one OPTION_S46_RULE option and at least one OPTION_S46_BR option.

Other options applicable to a domain may be defined in the future. A DHCPv6 message MAY include multiple OPTION_S46_CONT_MAPE options (representing multiple domains).

5.2. S46 MAP-T Container Option

The S46 MAP-T Container option (OPTION_S46_CONT_MAPT) specifies the container used to group all rules and optional port parameters for a specified domain.

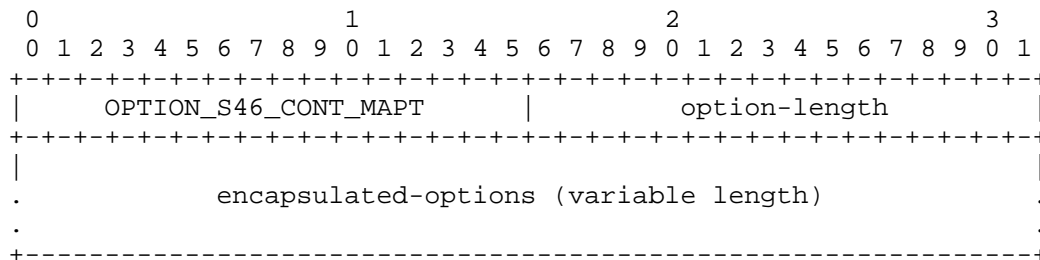


Figure 8: S46 MAP-T Container Option

- o option-code: OPTION_S46_CONT_MAPT (95)
- o option-length: length of encapsulated options, expressed in octets.
- o encapsulated-options: options associated with this Software46 MAP-T domain.

The encapsulated-options field conveys options specific to the OPTION_S46_CONT_MAPT option. Currently, there are two options specified: the OPTION_S46_RULE and OPTION_S46_DMR options. There MUST be at least one OPTION_S46_RULE option and exactly one OPTION_S46_DMR option.

5.3. S46 Lightweight 4over6 Container Option

The S46 Lightweight 4over6 Container option (OPTION_S46_CONT_LW) specifies the container used to group all rules and optional port parameters for a specified domain.

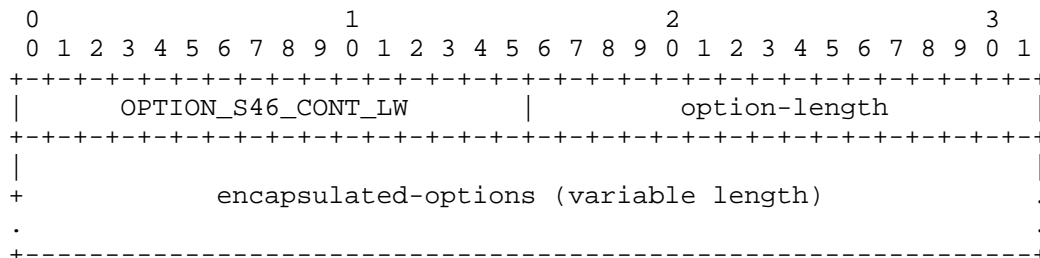


Figure 9: S46 Lightweight 4over6 Container Option

- o option-code: OPTION_S46_CONT_LW (96)
- o option-length: length of encapsulated options, expressed in octets.
- o encapsulated-options: options associated with this Software46 Lightweight 4over6 domain.

The encapsulated-options field conveys options specific to the OPTION_S46_CONT_LW option. Currently, there are two options specified: OPTION_S46_V4V6BIND and OPTION_S46_BR. There MUST be at most one OPTION_S46_V4V6BIND option and at least one OPTION_S46_BR option.

6. Software46 Options Encapsulation

The table below shows which encapsulated options are mandatory, optional, or not permitted for each defined container option.

Option	MAP-E	MAP-T	Lightweight 4over6
OPTION_S46_RULE	M	M	N/P
OPTION_S46_BR	M	N/P	M
OPTION_S46_PORTPARAMS	O	O	O
OPTION_S46_DMR	N/P	M	N/P
OPTION_S46_V4V6BIND	N/P	N/P	O

M - Mandatory, O - Optional, N/P - Not Permitted

Table 1: Options for Container Mappings

Software46 DHCPv6 clients that receive container options that violate any of the above rules MUST silently ignore such container options.

7. DHCPv6 Server Behavior

Section 17.2.2 of [RFC3315] describes how a DHCPv6 client and server negotiate configuration values using the ORO. As a convenience for the reader, we mention here that by default a server will not reply with a Software46 container option if the client has not explicitly enumerated one in its ORO.

A CE router may support several (or all) of the mechanisms mentioned here. In the case where a client requests multiple mechanisms in its ORO, the server will reply with the corresponding Software46 container options for which it has configuration information.

8. DHCPv6 Client Behavior

An S46 CE acting as a DHCPv6 client will request S46 configuration parameters from the DHCPv6 server located in the IPv6 network. Such a client MUST request the S46 container option(s) that it is configured for in its ORO in SOLICIT, REQUEST, RENEW, REBIND, and INFORMATION-REQUEST messages.

When processing received S46 container options, the following behavior is expected:

- o A client MUST support processing multiple received OPTION_S46_RULE options in a container OPTION_S46_CONT_MAPE or OPTION_S46_CONT_MAPT option.
- o A client receiving an unsupported S46 option or an invalid parameter value SHOULD discard that S46 container option and log the event.

The behavior of a client that supports multiple Software46 mechanisms is out of scope for this document. [Unified-v4-in-v6] describes client behavior for the prioritization and handling of multiple mechanisms simultaneously.

Note that a system implementing CE functionality may have multiple network interfaces, and these interfaces may be configured differently; some may be connected to networks using a Software46 mechanism, and some may be connected to networks that are using normal dual-stack or other means. The CE should approach this specification on an interface-by-interface basis. For example, if the CE system is MAP-E capable and is attached to multiple networks that provide the OPTION_S46_CONT_MAPE option, then the CE MUST configure MAP-E for each interface separately.

Failure modes are out of scope for this document. Failure recovery mechanisms may be defined in the future. See Section 5 of [RFC7597] for a discussion of valid MAP Rule combinations. See Section 11 of [RFC7227] and Sections 18.1.3, 18.1.4, and 19.1 of [RFC3315] for parameter-update mechanisms in DHCPv6 that can be leveraged to update configuration after a failure.

9. Security Considerations

Section 23 of [RFC3315] discusses DHCPv6-related security issues.

As with all DHCPv6-derived configuration states, it is possible that configuration is actually being delivered by a third party (Man in the Middle). As such, there is no basis on which access over MAP or Lightweight 4over6 can be trusted. Therefore, softwires should not bypass any security mechanisms such as IP firewalls.

In IPv6-only networks that lack IPv4 firewalls, a device that supports MAP could be tricked into enabling its IPv4 stack and directing IPv4 traffic to the attacker, thus exposing itself to previously infeasible IPv4 attack vectors.

Section 10 of [RFC7597] discusses security issues related to the MAP-E mechanism, Section 9 of [RFC7596] discusses security issues related to the Lightweight 4over6 mechanism, and Section 13 of [RFC7599] discusses security issues related to the MAP-T mechanism.

Readers concerned with the security of Software46 provisioning over DHCPv6 are encouraged to read [Secure-DHCPv6].

10. IANA Considerations

IANA has allocated the following DHCPv6 option codes:

- 89 for OPTION_S46_RULE
- 90 for OPTION_S46_BR
- 91 for OPTION_S46_DMR
- 92 for OPTION_S46_V4V6BIND
- 93 for OPTION_S46_PORTPARAMS
- 94 for OPTION_S46_CONT_MAPE
- 95 for OPTION_S46_CONT_MAPT
- 96 for OPTION_S46_CONT_LW

All values have been added to the "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)" option code space defined in Section 24.3 of [RFC3315].

11. References

11.1. 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>>.

11.2. Informative References

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<http://www.rfc-editor.org/info/rfc2131>>.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <<http://www.rfc-editor.org/info/rfc2473>>.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <<http://www.rfc-editor.org/info/rfc6145>>.
- [RFC6346] Bush, R., Ed., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", RFC 6346, DOI 10.17487/RFC6346, August 2011, <<http://www.rfc-editor.org/info/rfc6346>>.
- [RFC7227] Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", BCP 187, RFC 7227, DOI 10.17487/RFC7227, May 2014, <<http://www.rfc-editor.org/info/rfc7227>>.
- [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>>.

[Secure-DHCPv6]

Jiang, S., Ed., Shen, S., Zhang, D., and T. Jinmei,
"Secure DHCPv6", Work in Progress,
draft-ietf-dhc-sedhcpv6-08, June 2015.

[Unified-v4-in-v6]

Boucadair, M., Farrer, I., Perreault, S., Ed., and S.
Sivakumar, Ed., "Unified IPv4-in-IPv6 Software CPE", Work
in Progress, draft-ietf-software-unified-cpe-01, May 2013.

Acknowledgements

This document was created as a product of a MAP design team. The following people were members of that team: Congxiao Bao, Mohamed Boucadair, Gang Chen, Maoke Chen, Wojciech Dec, Xiaohong Deng, Jouni Korhonen, Xing Li, Satoru Matsushima, Tomek Mrugalski, Tetsuya Murakami, Jacni Qin, Necj Scoberne, Qiong Sun, Tina Tsou, Dan Wing, Leaf Yeh, and Jan Zorz.

The authors would like to thank Bernie Volz and Tom Taylor for their insightful comments and suggestions.

Authors' Addresses

Tomek Mrugalski
Internet Systems Consortium, Inc.
950 Charter Street
Redwood City, CA 94063
United States

Phone: +1 650 423 1345
Email: tomasz.mrugalski@gmail.com
URI: <http://www.isc.org/>

Ole Troan
Cisco Systems
Philip Pedersens vei 1
Lysaker 1366
Norway

Email: ot@cisco.com

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

Email: ian.farrer@telekom.de

Simon Perreault
Jive Communications
Quebec, QC
Canada

Email: sperreault@jive.com

Wojciech Dec
Cisco Systems, Inc.
The Netherlands

Email: wdec@cisco.com
URI: <http://cisco.com>

Congxiao Bao
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
China

Phone: +86 10-62785983
Email: congxiao@cernet.edu.cn

Leaf Y. Yeh
Freelancer Technologies
China

Email: leaf.y.yeh@hotmail.com

Xiaohong Deng
The University of New South Wales
Sydney NSW 2052
Australia

Email: dxhbupt@gmail.com
URI: <https://www.unsw.edu.au/>