

Softwire WG
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2014

T. Mrugalski
ISC
O. Troan
Cisco
I. Farrer
Deutsche Telekom AG
S. Perreault
Viagenie
W. Dec
Cisco
C. Bao
Tsinghua University
L. Yeh
CNNIC
X. Deng
Yingke Law Firm
March 13, 2014

**DHCPV6 Options for configuration of Softwire Address and Port Mapped
Clients
draft-ietf-softwire-map-dhcp-07**

Abstract

This document specifies DHCPV6 options, termed Softwire46 options, for the provisioning of Softwire46 Customer Edge (CE) devices. Softwire46 is a collective term used to refer to architectures based on the notion of IPv4 Address+Port (A+P) for providing IPv4 connectivity across an IPv6 network.

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 14, 2014.

Copyright Notice

Copyright (c) 2014 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 [2](#)
- [2.](#) Conventions [3](#)
- [3.](#) Softwire 46 Overview [3](#)
- [4.](#) Common Softwire46 DHCPv6 Options [4](#)
 - [4.1.](#) S46 Rule Option [5](#)
 - [4.2.](#) S46 BR Option [6](#)
 - [4.3.](#) S46 DMR Option [7](#)
 - [4.4.](#) S46 IPv4/IPv6 Address Binding Option [7](#)
 - [4.5.](#) S46 Port Parameters Option [8](#)
- [5.](#) Softwire 46 Container DHCPv6 Options [9](#)
 - [5.1.](#) Softwire46 MAP-E Container Option [9](#)
 - [5.2.](#) Softwire46 MAP-T Container Option [10](#)
 - [5.3.](#) Softwire46 LightWeight 46 Container Option [11](#)
- [6.](#) DHCPv6 Server Behavior [11](#)
- [7.](#) DHCPv6 Client Behavior [12](#)
- [8.](#) Security Considerations [12](#)
- [9.](#) IANA Considerations [13](#)
- [10.](#) Acknowledgements [13](#)
- [11.](#) References [13](#)
 - [11.1.](#) Normative References [13](#)
 - [11.2.](#) Informative References [13](#)
- Authors' Addresses [14](#)

1. Introduction

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

An example is Mapping of Address and Port (MAP) defined in [[I-D.ietf-softwire-map](#)]. 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 private 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 Softwire46 proposals.

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 will continue to use DHCPv4 [[RFC2131](#)].

This memo specifies a set of DHCPv6 [[RFC3315](#)] options to provision Softwire46 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 of 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. Softwire 46 Overview

This document describes a set of common DHCPv6 options for configuring the MAP-E [[I-D.ietf-softwire-map](#)], MAP-T [[I-D.ietf-softwire-map-t](#)] and Lightweight 4over6 [[I-D.ietf-softwire-lw4over6](#)] mechanisms.

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 [RFC2473](#) [[RFC2473](#)] IPv4 over IPv6 tunnelling,

while MAP-T uses NAT64 [[RFC6145](#)] based translation. 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 DHCP options described here tie the provisioning parameters, and hence the IPv4 service itself, to the End-user IPv6 prefix lifetime. The validity of a softwire'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 Option Request Option DHCPv6 [[RFC3315](#)] function 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 option. The provisioning parameters for that mechanism are expressed by embedding the common format options within the respective container.

This approach implies that all of the provisioning options MUST appear only within the container options. The client MUST NOT request any of the provisioning options directly within an ORO. Likewise, the server MUST NOT send the provisioning options directly within a DHCPv6 message, without encapsulation in the corresponding container option.

The document is organized with the common sub-options described first, followed by the three container options. Some sub-options are mandatory in some containers, some are optional and some are not permitted at all.

4. Common Softwire46 DHCPv6 Options

The DHCPv6 protocol is used for Softwire46 CE provisioning following regular DHCPv6 notions, with the CE assuming the role of a DHCPv6 client, and the DHCPv6 server providing options following typical DHCPv6 server side policies. The format and usage of the options are defined in the following sub-sections.

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, 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 OPTION_S46_PORTPARAMS to specify parameters and port sets for the port range algorithm.

4.1. S46 Rule Option

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

A server MAY send more than one S46 Rule Option in a container, if it is configured to do so. Clients MUST NOT send a S46 Rule Option.

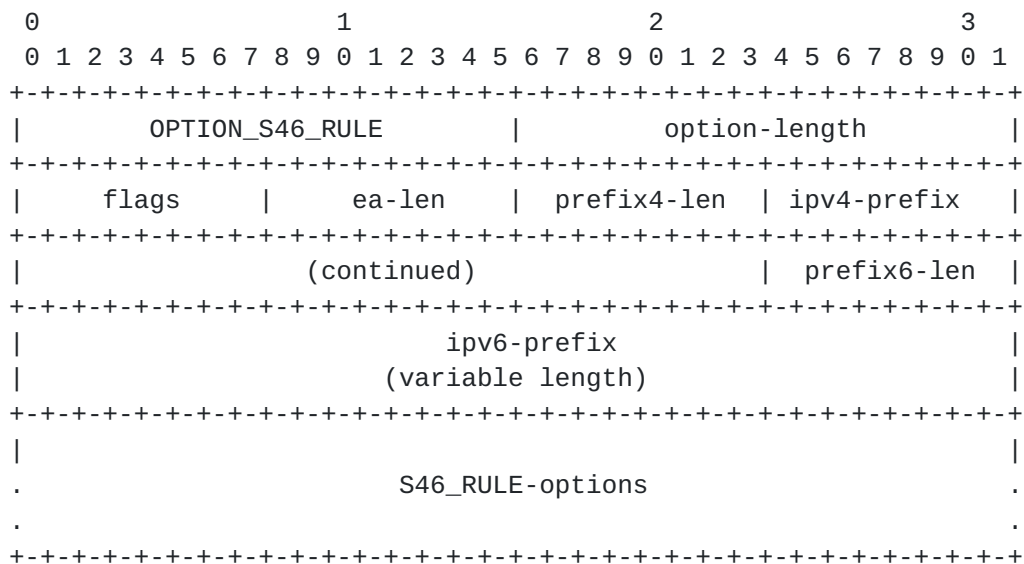


Figure 1: S46 Rule Option

- o option-code: OPTION_S46_RULE (TBD1)
- o option-length: length of the option, excluding option-code and option-length fields, including length of all encapsulated options, expressed in bytes.
- o flags: 8 bits long field carrying flags applicable to the rule. The meaning of specific bits are explained in Figure 2.
- o ea-len: 8 bits long field that specifies the Embedded-Address (EA) bit length. Allowed values range from 0 to 48.
- o prefix4-len: 8 bits long field expressing the prefix length of the IPv4 prefix specified in the rule-ipv4-prefix field. Valid values 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 field expressing the length of the IPv6 prefix specified in the rule-ipv6-prefix field.

- o ipv6-prefix: a variable length field that specifies the IPv6 domain prefix for the S46 rule. The field is padded with follow up zero bits up to the nearest octet boundary when prefix6-len is not divisible by 8.
- o S46_RULE-options: a variable field that may contain zero or more options that specify additional parameters for this S46 rule, e.g. a Port Parameter Option.

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 a FMR, if not set this rule is a BMR. Note: A BMR rule can also be an FMR rule by setting the F flag. The BMR rule is determined by a match of the Rule-IPv6-prefix against the CPE's 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 is used to convey the IPv6 address of the Border Relay. Figure 4 shows the format of the BR option.



Figure 3: S46 DMR Option

- o option-code: OPTION_S46_BR (TBD2)
- 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 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 is used to convey values for the Default Mapping Rule (DMR). Figure 4 shows the format of the MAP Rule option used for conveying a DMR.

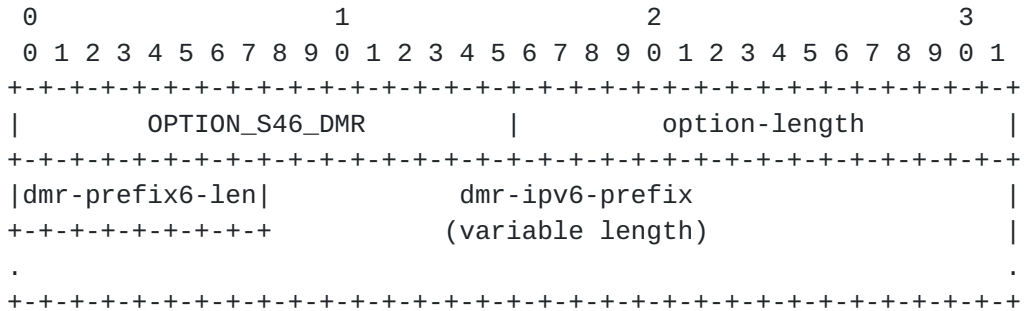



Figure 4: S46 DMR Option

- o option-code: OPTION_S46_DMR (TBD3)
- o option-length: 1 + length of dmr-ipv6-prefix specified in bytes.
- o dmr-prefix6-len: 8 bits long field expressing the bit mask length of the IPv6 prefix specified in the dmr-ipv6-prefix field.
- o dmr-ipv6-prefix: a variable length field specifying the IPv6 prefix or address for the S46 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 IPv4 address Option 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.



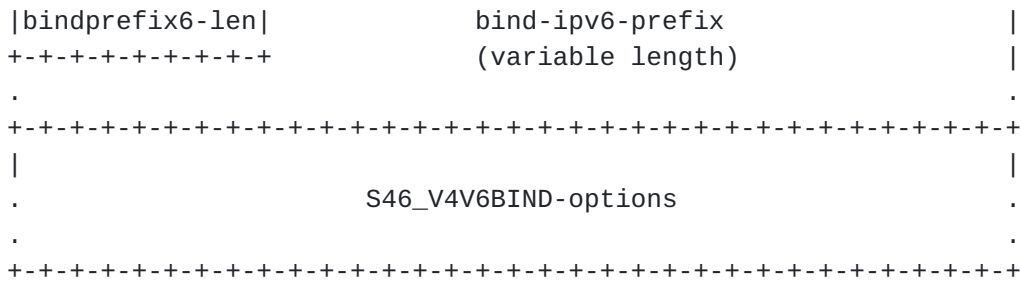


Figure 5: S46 IPv4/IPv6 Address Binding Option

- o option-code: OPTION_S46_V4V6BIND (TBD4)
- o option-length: 4
- o ipv4-address: A fixed field of 4 octets specifying an IPv4 address.
- o bindprefix6-len: 8 bits long field expressing the bit mask length of the IPv6 prefix specified in the bind-ipv6-prefix field.
- o bind-ipv6-prefix: a variable length field specifying the IPv6 prefix or address for the S46. 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 field that may contain zero or more options that specify additional parameters e.g. a Port Parameters Option.

4.5. S46 Port Parameters Option

The Port Parameters Option specifies optional Rule Port Parameters that MAY be provided as part of the Mapping Rule for CEs using the MAP algorithm.

See [[I-D.ietf-softwire-map](#)], Section 5.1 for a description of MAP algorithm, explaining all of the parameters in detail.




Figure 6: S46 Port Parameters Option

- o option-code: OPTION_S46_PORTPARAMS (TBD5)
- o option-length: 4
- o offset: (PSID offset) 8 bits long field that specifies the numeric value for the S46 algorithm's excluded port range/offset bits

(A-bits), as per section 5.1.1 of [[I-D.ietf-softwire-map](#)]. Allowed values are between 0 and 15, with the default value being 6.

- o PSID-len: Bit length value of 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 Port Set Identifier (PSID). Consequently, the address sharing ratio would be 2^k.
- o PSID: Explicit 16-bit (unsigned word) PSID value. 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 value. The remaining (16-k) bits on the right are padding zeros.

When receiving the Port Parameters option with an explicit PSID, the client MUST use this explicit PSID in configuring its MAP interface. If the conveyed IPv4 address is not 32 bit-long, the option MUST be discarded. The formula for this check is "prefix4-len + ea-len = 32" and serves to ensure that the explicit PSID is only applied to configurations with a completely formed IPv4 address.

The OPTION_S46_PORTPARAMS option MUST be encapsulated in a OPTION_S46_RULE option or an OPTION_S46_V4V6BIND option. It MUST NOT appear directly within a container option.

5. Softwire 46 Container DHCPv6 Options

Option	MAP-E	MAP-T	Lightweight 4over6
OPTION_S46_RULE	M	M	N/A
OPTION_S46_BR	M	N/A	M
OPTION_S46_PORTPARAMS	0	0	0
OPTION_S46_DMR	N/A	M	N/A
OPTION_S46_V4V6BIND	N/A	N/A	0

M - Mandatory, 0 - Optional, N/A - Not Applicable

Table 1: Option to Container Mappings

5.1. Softwire46 MAP-E Container Option

The MAP-E Container Option specifies the container used to group all rules and optional port parameters for a specified domain.






Figure 7: MAP-E Container Option

- o option-code: OPTION_S46_CONT_MAPE (TBD6)
- o option-length: Length of encapsulated options
- o encapsulated-options: options associated with this Softwire46 MAP-E domain.

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

Other options applicable to a domain may be defined in the future. A DHCP message MAY include multiple S46 MAPE Container Options (representing multiple domains).

5.2. Softwire46 MAP-T Container Option

The MAP-T Container Option specifies the container used to group all rules and optional port parameters for a specified domain.

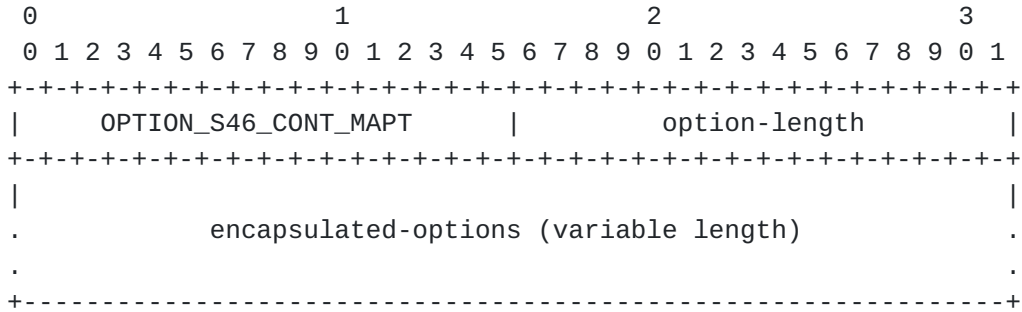


Figure 8: MAP-E Container Option

- o option-code: OPTION_S46_CONT_MAPT (TBD7)
- o option-length: Length of encapsulated options
- o encapsulated-options: options associated with this Softwire46 MAP-T domain.

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

5.3. Softwire46 Lightweight 46 Container Option

The LW46 Container Option specifies the container used to group all rules and optional port parameters for a specified domain.

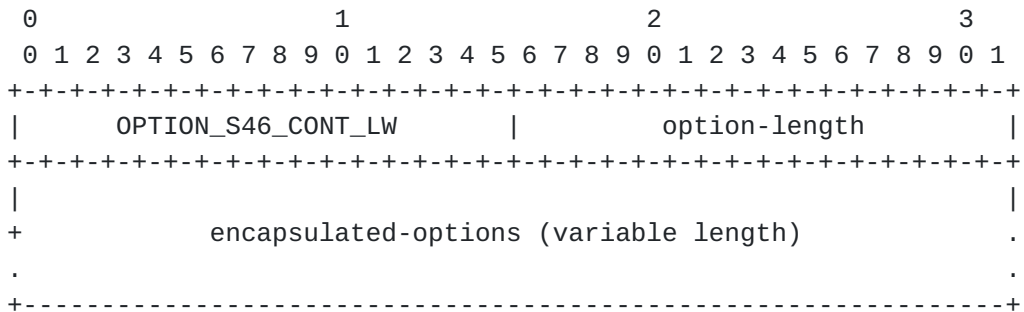


Figure 9: LW46 Container Option

- o option-code: OPTION_S46_CONT_LW (TBD8)
- o option-length: Length of encapsulated options
- o encapsulated-options: options associated with this Softwire46 domain.

The encapsulated options field conveys options specific to this Lightweight 4over6 Option. Currently there are two options specified for the OPTION_S46_CONT_LW option, 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. DHCPv6 Server Behavior

[RFC3315] Section 17.2.2 describes how a DHCPv6 client and server negotiate configuration values using the ORO. As a convenience to the reader, we mention here that by default, a server will not reply with a Softwire 46 Container Option if the client has not explicitly enumerated one in its Option Request Option.

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 option, the server SHOULD reply with all the corresponding Softwire 46 Container options, enumerated in the Option Request Option, the the server is configured for.

7. DHCPv6 Client Behavior

An S46 CE acting as DHCPv6 client will request S46 configuration parameters from the DHCPv6 server located in the IPv6 network. Such a client SHOULD include 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 behaviour 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 supporting multiple Softwire 46 mechanisms, is out of scope of this document. [[I-D.ietf-softwire-unified-cpe](#)] describes client behaviour for the prioritization and handling of multiple mechanisms simultaneously.

Note that system implementing CE functionality may have multiple network interfaces, and these interfaces may be configured differently; some may be connected to networks that call for MAP, and some may be connected to networks that are using normal dual stack or other means. The CE system should approach this specification on an interface-by-interface basis. For example, if the CE system is MAP capable and is attached to multiple networks that provide the MAP Mapping Rule Option, then the CE system MUST configure a MAP service (i.e. a translation or encapsulation) for each interface separately as each MAP provides IPv4 connectivity for each distinct interface. The means to bind a MAP configuration to a given interface in a multiple interfaces device are out of scope of this document.

8. Security Considerations

Implementation of this document does not present any new security issues, but as with all DHCPv6-derived configuration state, 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 lw4o6 can be trusted. Therefore, softwires should not bypass any security mechanisms such as IP firewalls.

Readers concerned with security of MAP provisioning over DHCPv6 are encouraged to read [[I-D.ietf-dhc-secure-dhcpv6](#)].

Section 11 of [[I-D.ietf-software-map](#)] discusses security issues of the MAP mechanism.

[Section 23 of \[RFC3315\]](#) discusses DHCPv6-related security issues.

9. IANA Considerations

IANA is kindly requested to allocate the following DHCPv6 option codes:

TBD1 for OPTION_S46_RULE
TBD2 for OPTION_S4_BR
TBD3 for OPTION_S46_DMR
TBD4 for OPTION_S46_V4V6BIND
TBD5 for OPTION_S46_PORTPARAMS
TBD6 for OPTION_S46_CONT_MAPE
TBD7 for OPTION_S46_CONT_MAPT
TBD8 for OPTION_S46_CONT_LW

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

10. Acknowledgements

This document was created as a product of a MAP design team. 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, Tomasz 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.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.

11.2. Informative References

[I-D.ietf-dhc-secure-dhcpv6]

Jiang, S. and S. Shen, "Secure DHCPv6 Using CGAs", [draft-ietf-dhc-secure-dhcpv6-07](#) (work in progress), September 2012.

[I-D.ietf-softwire-lw4over6]

Cui, Y., Qiong, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", [draft-ietf-softwire-lw4over6-03](#) (work in progress), November 2013.

[I-D.ietf-softwire-map-t]

Li, X., Bao, C., Dec, W., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", [draft-ietf-softwire-map-t-04](#) (work in progress), September 2013.

[I-D.ietf-softwire-map]

Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, "Mapping of Address and Port with Encapsulation (MAP)", [draft-ietf-softwire-map-08](#) (work in progress), August 2013.

[I-D.ietf-softwire-unified-cpe]

Boucadair, M., Farrer, I., Perreault, S., and S. Sivakumar, "Unified IPv4-in-IPv6 Softwire CPE", [draft-ietf-softwire-unified-cpe-01](#) (work in progress), May 2013.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", [RFC 2473](#), December 1998.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", [RFC 6145](#), April 2011.

Authors' Addresses

Tomasz Mrugalski
Internet Systems Consortium, Inc.
950 Charter Street
Redwood City, CA 94063
USA

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

Ole Troan
Cisco Systems, Inc.
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
Viagenie
246 Aberdeen
Quebec, QC G1R 2E1
Canada

Phone: +1 418 656 9254
Email: simon.perreault@viagenie.ca

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
CN

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

Leaf Y. Yeh
CNNIC
4, South 4th Street, Zhong_Guan_Cun
Beijing 100190
P. R. China

Email: leaf.yeh.sdo@gmail.com

Xiaohong Deng
Yingke Law Firm
6 Floor, C Block, DaCheng International Center Chaoyang District
Beijing 100124
China

Phone: +61 3858 3128

Email: dxhbupt@gmail.com

