

Softwire WG
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2014

T. Mrugalski
ISC
X. Deng

O. Troan
Cisco
C. Bao
Tsinghua University
W. Dec, Ed.
Cisco
L. Yeah
Freelancer Technologies
July 15, 2013

**DHCPv6 Options for configuration of Softwire Address and Port Mapped
Clients
draft-ietf-softwire-map-dhcp-04**

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, as specified by the IETF Softwires Working Group, including those where the IPv4 A+P parameters are optionally derived from the CE's IPv6 network prefix by algorithmic means.

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 January 16, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.	Softwire46 Service Information	3
4.	DHCPv6 Softwire46 Options	4
4.1.	Softwire46 Options Cardinality	4
4.2.	Softwire46 Container Option	4
4.3.	S46 Basic Rule Option	5
4.4.	S46 Forwarding Rule Option	6
4.5.	S46 Port Parameters Option	8
4.6.	Definition of Alternative Port Paramaters Option	8
5.	DHCPv6 Server Behavior	9
6.	DHCPv6 Client Behavior	9
6.1.	Processing of information by clients supporting MAP	10
7.	DHCPv6 Deployment Considerations	11
8.	S46 Options Examples	12
8.1.	Basic Rule with no embedded address bits and no IPv4 address sharing	12
8.2.	Basic Rule Option Example with shared addressing and MAP	13
8.3.	Forwarding Rule Option Example with MAP mesh mode	13
8.4.	Basic Rule Option Example with MAP and explicit PSID	14
8.5.	Forwarding Rule Option example for MAP-T DMR prefix	16
9.	IANA Considerations	16
10.	Security Considerations	16
11.	Acknowledgements	17
12.	References	17
12.1.	Normative References	17
12.2.	Informative References	17
	Authors' Addresses	18

1. Introduction

A number of architectural solution proposals formed in the IETF Softwires Working Group use Address and Port (A+P) [[RFC6346](#)] as their technology base in providing IPv4 connectivity service to end users using CE devices and doing so across a service provider's IPv6 network, while also allowing for shared or dedicated IPv4 addressing of the CEs. This document specifies DHCPv6 options for the provisioning of such Softwire A+P based Customer Edge (CE) devices, referred to collectively as Softwire46 options.

Examples of such proposals are those of Mapping of Address and Port (MAP) defined in [[I-D.ietf-softwire-map](#)], [[I-D.ietf-softwire-map-t](#)], and [[I-D.ietf-softwire-lw4over6](#)]. In all cases the solution architecture consists of one or more domain gateways (e.g. a MAP Border Relay or lw46 AFTR), responsible for forwarding between an IPv6 network domain and external public IPv4 network, and one or more Customer Edge (CE) routers, responsible for forwarding between a user's private IPv4 network and the IPv6 network domain. Collectively the CE and domain router form a domain when configured with common service parameters. These parameters consist primarily of the IPv4 address and the transport layer port-range(s) that each CE should use in the domain, as well as the parameters for the IPv6 network transport mode (e.g. Encapsulation or translation, address of the domain gateway, etc). Depending on the characteristics of the domain, the IPv4 address and port parameters can be optionally algorithmically derived from the CE's native end-user IPv6 prefix. Networks with numerous CEs require dynamic configuration of these parameters on CEs, which forms the requirement for a dynamic provisioning mechanism using DHCPv6 [[RFC3315](#)] options that are the subject of this document. The means of configuration of the domain's gateway is outside the 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. Softwire46 Service Information

The following presents the information parameters that are used to configure a Softwire46 CE and defined in this document:

- o A Basic Rule (BR). This rule governs the A+P configuration of the CE in terms of deriving the CE's IPv4 address and port range(s) as well as completing the CE's IPv6 address for the S46 service when

- needed. Key parameters of a BR include: i) The IPv4 Prefix - Used to derive the CE's IPv4 address; ii) The Embedded Address bit length - Used to derive how many, if any, of the CE's end-user IPv6 prefix are to be mapped to its IPv4 address and port-set. iii) The IPv6 prefix - used to scope the CE's IPv6 domain for the S46 service
- o A Forwarding Rule (FR). This rule governs the forwarding behaviour to destinations covered by the rule expressed in terms of an IPv4 prefix.
 - o Port Parameters: The specification allows flexibility in the level of automation a CE uses to derive its transport layer port-ranges, ranging from full derivation of this parameter from the CE's IPv6 prefix, to full parametrization of configuration independent of the CE's IPv6 prefix.

4. DHCPv6 Software46 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 typical DHCPv6 server side policies. The format and usage of the options is defined in the following sub-sections. Examples of use of the option as are presented in [Section 8](#).

4.1. Software46 Options Cardinality

When requested by clients by means of an DHCPv6 Option Request Option a server SHOULD return one Software46 Container Option for each Software46 domain that a given CPE is determined to belong to. Typically a CE will belong to only one domain, multiple CEs can also be part of a single domain and thus share their Software46 service configuration. The means of determining the domain of a given client is by following typical DHCPv6 server policies and not specified by this document.

A returned Software46 Container Option MAY include one or more Rule Options.

4.2. Software46 Container Option

This S46 Container Option specifies the container used to group all rules and optional port parameters for a specified domain.

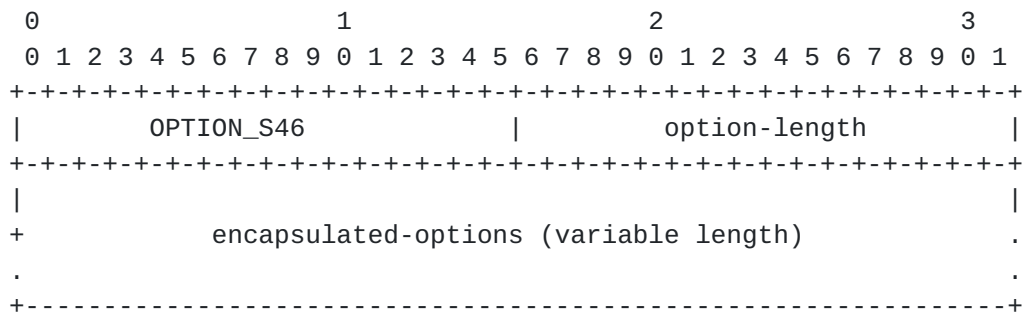


Figure 1: MAP Container Option

- o option-code: OPTION_S46 (TBD1)
- o option-length: Length of encapsulated options
- o encapsulated-options: options associated with this Software46 domain.

Currently there are two options that MAY appear encapsulated as sub-options in the Container Option: OPTION_S46_BRULE and OPTION_S46_FRULE. Other options suitable for a domain may be defined in the future. A DHCP message MAY include multiple S46 Container Options (representing multiple domains), but typically it will have only one.

4.3. S46 Basic Rule Option

Figure 2 shows the format of the S46 Rule option used for conveying the Basic Rule that is used for determining the key A+P characteristics intended for the client.

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

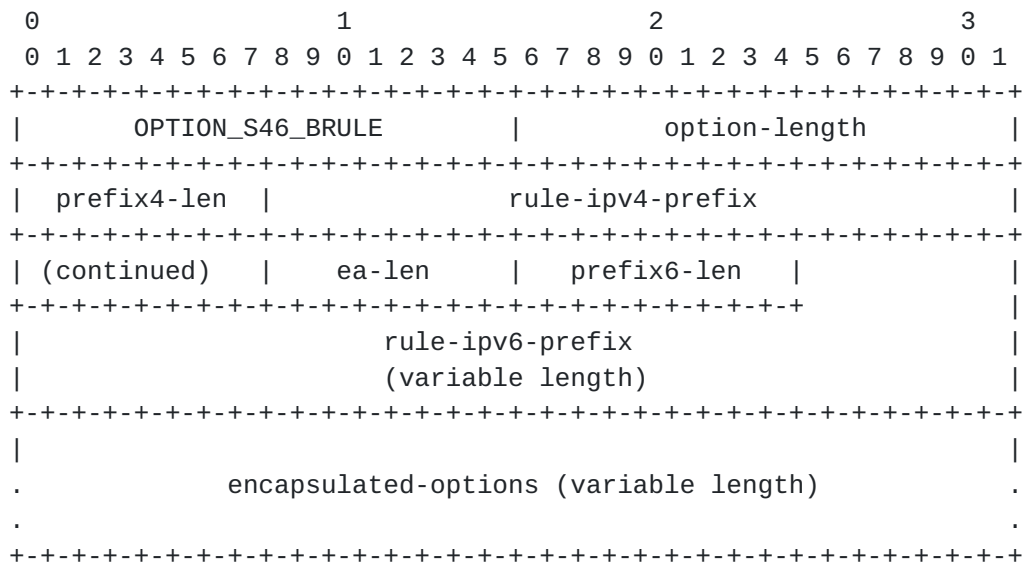


Figure 2: S46 Basic Rule Option

- o option-code: OPTION_S46_BRULE (TBD2)
- o option-length: length of the option, excluding option-code and option-length fields, including length of all encapsulated options, expressed in bytes.
- o prefix4-len: 8 bits long field expressing the bit mask length of the IPv4 prefix specified in the rule-ipv4-prefix field. Allowed values are 0 to 32.
- o rule-ipv4-prefix: a fixed length 32 bit field that specifies the IPv4 prefix for the S46 rule.
- o ea-len: 8 bits long field that specifies the Embedded-Address (EA) bit length. Values allowed range from 0 to 48. A value of 0 signifies no embedded address.
- o prefix6-len: 8 bits long field expressing the bit mask length of the IPv6 prefix specified in the rule-ipv6-prefix field. Allowed values are 0 to 128.
- o rule-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 encapsulated 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.

4.4. S46 Forwarding Rule Option

Figure 3 shows the format of the S46 Forwarding Rule option used for conveying a Forwarding Rule that governs the IPv4 forwarding behaviour for destinations covered by the rule.

- o option-code: OPTION_S46_FRULE (TBD2)
- o option-length: length of the option, excluding option-code and option-length fields, including length of all encapsulated options, expressed in bytes.
- o prefix4-len: 8 bits long field expressing the bit mask length of the IPv4 prefix specified in the rule-ipv4-prefix field. Valid values 0 to 32.
- o rule-ipv4-prefix: a fixed length 32 bit field that specifies the IPv4 prefix for the S46 rule.
- o ea-len: 8 bits long field that specifies the Embedded-Address (EA) bit length. Values allowed range from 0 to 48. A value of 0 signifies no embedded address.
- o prefix6-len: 8 bits long field expressing the bit mask length of the IPv6 prefix specified in the rule-ipv6-prefix field. Allowed values are 0 to 128.
- o rule-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 encapsulated 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.

4.5. S46 Port Parameters Option

The S46 Port parameters Option conveys information used for determining the transport port-range(s) to be used by the client. The port-range is generally represented by an index, the Port Set Identifier (PSID), that pertains to a given port-indexing reversible algorithm.

When used, the option MUST appear as encapsulated option in S46 Basic or Forwarding Rule options and it MUST NOT appear directly in a message.




Figure 4: S46 Port Parameters Option

- o option-code: OPTION_S46__PORTPARAMS (TBD4)
 - 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 in [[I-D.ietf-softwire-map](#)]. Allowed values are between 0 and 15. 0 means no exclusion.
 - o PSID-len: 8 bits long field expressing the bit length value of the number of significant bits in the PSID field (also known as 'k' bits). When set to 0, the PSID field is to be ignored.
 - o PSID: Explicit 16-bit (unsigned word) PSID value. The first k-bits on the left of this 2-octets field is the PSID value. The remaining (16-k) bits on the right are padded with zeros.
- The numeric sum of offset and PSID-len MUST NOT exceed 16.

4.6. Definition of Alternative Port Paramaters Option

Given that all the S46 options in this document are regular DHCPv6 options (with global option codes), alternative port parameters options relating to alternative port-indexing methods and intended for use along with the S46 options can be defined following DHCPv6 extension procedures [[I-D.ietf-dhc-option-guidelines](#)].

5. DHCPv6 Server Behavior

[RFC 3315 Section 17.2.2](#) [[RFC3315](#)] describes how a DHCPv6 client and server negotiate configuration values using the ORO. As a convenience to the reader, we mention here that a server will by default not reply with a S46 Rule Option if the client has not explicitly enumerated it on its Option Request Option.

A Server following this specification MUST allow the configuration of one or more S46 Rule Options, and optional Port Parameters Option and SHOULD send such options grouped under a single S46 Container Option. The server SHOULD allow the S46 Rule Options to be part of a wider set of options returned to a client.

The DHCPv6 Server MUST use a S46 Container Option, which encapsulates all S46 Basic and Forwarding Rules, as well as Port parameters Options, when sending responses to clients that requested S46 configuration.

A DHCPv6 server that is enabled to provision requesting CE's with Softwire46 options MUST include exactly one OPTION_S46 option in a REPLY message for each domain the CE is determined to be part of - by default this is one domain.

6. DHCPv6 Client Behavior

A CE acting as DHCPv6 client will request S46 configuration to be assigned from a DHCPv6 server located in the IPv6 network. Such a client SHOULD request OPTION_S46 option in its ORO in SOLICIT, REQUEST, RENEW, REBIND and INFORMATION-REQUEST messages. The other S46 options MAY also be included as options in the ORO.

When processing received S46 options the following behaviour is expected:

- o A client MUST support processing multiple received OPTION_S46_FRULE options in a container OPTION_S46 option.
- o A client MUST use the contents of the OPTION_S46_BRULE rule-ipv6-prefix parameter to determine the IPv6 interface prefix to select for the S46 service, for example by performing a longest match between the rule-ipv6-prefix and the locally configured interfaces. A null rule-ipv6-prefix (rule6-prefix-len = 0) communicates to the client that any IPv6 prefix can be used from the interface as used by the DHCPv6 client receiving this option.
- o The contents of the OPTION_S46_FRULE convey information about the mode that the client should use for transporting IPv4 traffic across the IPv6 network. A client receiving an OPTION_S46_FRULE with an 0/0 IPv4 prefix (i.e. prefix4-len parameter = 0) and

prefix6-len length of 128 SHOULD assume encapsulation mode [RFC2473] of IPv6 transport. Conversely, when the received option has a prefix length shorter than 128, the client SHOULD assume a stateless NAT46 [RFC6145] translation mode of IPv6 transport. Clients not supporting the intended mode of transport SHOULD discard the option and log an "DHCP S46 FRule Option received requires form of IPv6 transport not supported by this device" message, or similar.

- o The presence of a OPTION_S46_BRULE option indicates to the CE that NAPT44 mode of operation is to be used, and the conveyed A+P parameters applied. Conversely, the absence of an OPTION_S46_BRULE option indicates that NAPT44 mode is not required and SHOULD be disabled. This allows the client to be compatible for use with [RFC6333] architectures*.
- o The processed contents of the OPTION_S46_FRULE and OPTION_S46_FRULE override any previously acquired rules for the domain.

*Note: An operator intending to use the client in a [RFC6877] architecture, can provision all clients with the same OPTION_S46_BRULE passing in the rule-ipv4-prefix any suitable RFC1918 IPv4 address. This would be sufficient to allow a NAT64 capable client (e.g. a MAP-T CE) to be used in stateful core NAT64 deployments, besides stateless NAT64.

6.1. Processing of information by clients supporting MAP

Clients using the MAP algorithm MUST be capable of applying the received S46 option parameters for the configuration of the local MAP service: The MAP Base Mapping Rule (BMR) and Forward Mapping Rule (FMR) parameters are to be derived from the OPTION_S46_BRULE and OPTION_S46_FRULE Options respectively and MUST be used to configure the CE's MAP service characteristics as per Section 5 of [I-D.ietf-softwire-map], including the CE's MAP interface address.

When processing OPTION_S46_FRULE containing an parameter of prefix4-len = 0, a MAP CE client SHOULD determine the MAP BR's address or prefix [I-D.ietf-softwire-map-t], from the OPTION_S46_FRULE as follows:

- o If the prefix6-len parameter is 128, the BR's address is contained in the rule-ipv6-prefix parameter. This forms the MAP BR address for MAP-E.
- o If the prefix6-len parameter is <128, the BR's prefix is contained in the rule-ipv6-prefix parameter. This forms the Default Mapping Rule (DMR) for MAP-T.
- o In either of the above cases, the determined parameter MUST override the MAP BR's address or prefix parameter that was determined using other** means.

****Note:** The MAP CE's BR's address or prefix, as well as implicit encapsulating or translating mode of operation may be derived using other means, such as [\[RFC6334\]](#) or [\[I-D.ietf-behave-nat64-discovery-heuristic\]](#). This allows for the deployment of MAP CEs in existing DS-Lite or 464xlat deployments. If the operator prefers to utilize these other means, then the Forwarding Rule Option with the 0/0 IPv4 parameter is simply not distributed.

The MAP algorithm, as per [Section 5.1 \[I-D.ietf-software-map\]](#), allows the port-indexing algorithm to be parameterized in terms of an excluded port range (a-bits) and an explicit Port Set ID (PSID) value, which can be conveyed via the S46 Port Parameters Option. The following behaviour is expected when processing this option by MAP clients: A client receiving a S46 Port parameters Option in an OPTION_S46_BRULE or OPTION_S46_FRULE whose derived IPv4 address is not 32 bit-long MUST discard the option. The formula for this check is "prefix4-len + ea-len == 32". This is to ensure that the explicit PSID is only applied to configurations with a completely formed IPv4 address.

When received in an OPTION_S46_BRULE, the MAP CE client MUST use the contents of S46 Port Option for configuring its MAP Basic Mapping Rule (BMR), and similarly, it MUST use the contents for configuring its MAP Forward Mapping Rule (FMR) when received in an OPTION_S46_FRULE.

7. DHCPv6 Deployment Considerations

The usage of the S46 Option(s) is intended for stateless server operation, whereby for any given client, the server establishes no dynamic state regarding the S46 options that are sent to the client. For cases that require per client configuration, the DHCPv6 server needs to have means to access a per-client option configuration store. This is commonly realized and utilized DHCPv6 server platforms servers using back-end mechanisms such as [\[RFC4510\]](#) for example. This type of DHCP server and configuration store deployment is recommended when the intent is to utilize full IPv4-IPv6 address independence of any given CE, by issuing unique options to each CE.

In typical MAP deployments, many clients will belong to one of a few MAP domains, with the same MAP domain level configuration applicable to of its clients. The DHCPv6 server is then expected to provide the same set of MAP Option(s) to all clients in a given domain, governed by a regular DHCP server policy.

Domains that require more extensive IPv4 client configuration

options, i.e. More extensive than those possible by means of DHCPv6 in conjunction with the options in this document, should consider the use of DHCPv4 over IPv6 as specified in [\[I-D.ietf-dhc-dhcpv4-over-ipv6\]](#).

8. S46 Options Examples

This section presents some examples of the use of the S46 Options.

8.1. Basic Rule with no embedded address bits and no IPv4 address sharing

Example 1 - Basic Rule:

This example demonstrates the use of Basic Rule Option to convey to a client a non-shared IPv4 address 192.0.2.1/32. It is applicable to any S46 architecture. The end-user ipv6 prefix is assumed to be 2001:db8:0012:3400::/56.

DHCPv6_S46_OPTION:

option-length: 18

DHCPv6_S46_BRULE_OPTION:

option-length: 14

prefix4-len: 32

rule-ipv4-prefix: 192.0.2.1

ea-len: 0

prefix6-len: 56

rule-ipv6-prefix: 2001:db8:0012:34

8.2. Basic Rule Option Example with shared addressing and MAP

Example 2 - Basic Rule with MAP:

Given an IPv6 prefix assigned to the end-user:
2001:db8:0012:3400::/56, the following DHCPv6 S46_BRULE_OPTION
allows a client to derive its basic A+P configuration
for a shared address domain. The shared IPv4 subnet is
192.0.2.0/24.

```
DHCPv6_S46_OPTION:
    option-length: 14
DHCPv6_S46_BRULE_OPTION:
    option-length: 10
    prefix4-len: 24
    rule-ipv4-prefix: 192.0.2.0
    ea-len: 16
    prefix6-len: 40
    rule-ipv6-prefix: 2001:db8:00
```

Using the above a MAP CE node will determine the IPv4 address
and port-set as shown below:

```
EA bits offset: 40
IPv4 suffix bits (p) Length of IPv4 address (32) - IPv4 prefix
length (24) = 8
IPv4 address 192.0.2.18 (0xc0000212)
PSID start: 40 + p = 40 + 8 = 48
PSID length (q): o - p = (End-user prefix len -
rule IPv6 prefix len) - p = (56 - 40) - 8 = 8
PSID derived from IPv6 end-user-prefix: 0x34
```

Available ports (63 ranges): 1232-1235, 2256-2259, ,
63696-63699, 64720-64723

IPv6 address of MAP CE: 2001:db8:0012:3400:0000:c000:0212:0034

8.3. Forwarding Rule Option Example with MAP mesh mode

Example 3 - Mesh mode Forwarding Rule:

Given an IPv6 prefix assigned to the end-user:
2001:db8:0012:3400::/56, the following
DHCPv6 S46_FRULE_OPTION allows the client to communicate
directly with other CEs in the domain addressed by
the 192.0.2.0/24 prefix.

```
DHCPv6_S46_OPTION:
    option-length: 14
DHCPv6_S46_FRULE_OPTION:
    option-length: 10
    prefix4-len: 24
    rule-ipv4-prefix: 192.0.2.0
    ea-len: 16
    prefix6-len: 40
    rule-ipv6-prefix: 2001:db8:00
```

8.4. Basic Rule Option Example with MAP and explicit PSID

Example 4 - Basic Rule with PSID:

Given an IPv6 prefix assigned to the end-user:
2001:db8:0012:3400::/56, the following DHCPv6 S46_BRULE_OPTION
allows a client to derive its basic A+P configuration
for a shared MAP address domain with no relation to the
IPv6 end-user prefix.

```
DHCPv6_S46_OPTION:
    option-length: 34
DHCPv6_S46_BRULE_OPTION:
    option-length: 30
    prefix4-len: 32
    rule-ipv4-prefix: 192.0.2.1
    ea-len: 0
    prefix6-len: 56
    rule-ipv6-prefix: 2001:db8:0000:00
DHCPv6_S46_PORTPARAMS_OPTION:
    option-length: 4
    rsv: 0
    offset: 6
    rsv: 0
    PSID-len: 8
    PSID: 0x20
```

The Basic Rule information allows a MAP CE also to determine
its full IPv6 address by combining the IPv6 prefix with the MAP
interface identifier (that embeds the IPv4 address and PSID).

IPv6 address of MAP CE: 2001:db8:0012:3400:0000:c000:0212:0034

Note that the IPv4 address and PSID is not derived from the IPv6
prefix assigned to the CE, but provisioned separately using for
example MAP options in DHCPv6.

8.5. Forwarding Rule Option example for MAP-T DMR prefix

Example 5 - Forwarding Rule for MAP-T:

Given an IPv6 prefix assigned to the end-user:
2001:db8:0012:3400::/56, the following DHCPv6 S46_FRULE_OPTION
allows a client to derive its MAP-T's DMR prefix.

```
DHCPv6_S46_OPTION:
  option-length: 19
  DHCPv6_S46_FRULE_OPTION:
    option-length: 15
    prefix4-len: 0
    rule-ipv4-prefix: 0.0.0.0
    ea-len: 0
    prefix6-len: 64
    rule-ipv6-prefix: 2001:db8:ffff:0000
```

Using the above a MAP CE node will install a DMR using
prefix 2001:db8:ffff/64.

9. IANA Considerations

IANA is kindly requested to allocate the following DHCPv6 option codes: TBD1 for OPTION_S46, TBD2 for OPTION_S4_BRULE, TBD3 for OPTION_S46_FRULE, and TBD4 for OPTION_S46_PORTPARAMS. All values should be added to the DHCPv6 option code space defined in [Section 24.3 of \[RFC3315\]](#).

10. Security Considerations

Implementation of this document does not present any new security issues, but as with all DHCPv6-derived configuration state, it is completely possible that the configuration is being delivered by a third party (Man In The Middle). As such, there is no basis to trust that the access over the S46 service can be trusted, and it should not therefore bypass any security mechanisms such as IP firewalls.

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

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

11. Acknowledgements

This document was created as a product of a Software WG and 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.

Former MAP design team members are: Remi Despres.

Authors would like to thank Bernie Volz for his insightful comments and suggestions.

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

12.2. Informative References

- [I-D.ietf-behave-nat64-discovery-heuristic]
Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", [draft-ietf-behave-nat64-discovery-heuristic-17](#) (work in progress), April 2013.
- [I-D.ietf-dhc-dhcpv4-over-ipv6]
Cui, Y., Wu, P., Wu, J., and T. Lemon, "DHCPv4 over IPv6 Transport", [draft-ietf-dhc-dhcpv4-over-ipv6-06](#) (work in progress), March 2013.
- [I-D.ietf-dhc-option-guidelines]
Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", [draft-ietf-dhc-option-guidelines-12](#) (work in progress), June 2013.

- [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., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I.
Farrer, "Lightweight 4over6: An Extension to the DS-Lite
Architecture", [draft-ietf-softwire-lw4over6-00](#) (work in
progress), April 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-07](#)
(work in progress), May 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-03](#) (work
in progress), July 2013.
- [RFC4510] Zeilenga, K., "Lightweight Directory Access Protocol
(LDAP): Technical Specification Road Map", [RFC 4510](#),
June 2006.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation
Algorithm", [RFC 6145](#), April 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-
Stack Lite Broadband Deployments Following IPv4
Exhaustion", [RFC 6333](#), August 2011.
- [RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration
Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite",
[RFC 6334](#), August 2011.
- [RFC6346] Bush, R., "The Address plus Port (A+P) Approach to the
IPv4 Address Shortage", [RFC 6346](#), August 2011.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT:
Combination of Stateful and Stateless Translation",
[RFC 6877](#), April 2013.

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/>

Xiaohong Deng
6 Cordelia St.
South Brisbane QLD 4101
Australia

Phone: +61 3858 3128
Email: dxhbupt@gmail.com

Ole Troan
Cisco Systems, Inc.
Telemarksvingen 20
Oslo N-0655
Norway

Email: ot@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

Wojciech Dec (editor)
Cisco Systems, Inc.
The Netherlands

Phone:
Fax:
Email: wdec@cisco.com
URI:

Leaf Yeh
Freelancer Technologies
Shenzhen,
P. R. China

Phone:
Fax:
Email: leaf.yeh.sdo@gmail.com
URI:

