

- 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 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 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:

```

      0 1 2 3 4 5 6 7
      +---+---+---+---+
      |Reserved      |F|
      +---+---+---+---+

```

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 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 rule 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 4 shows the format of the OPTION_S46_BR option.



Figure 3: S46 BR 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 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.

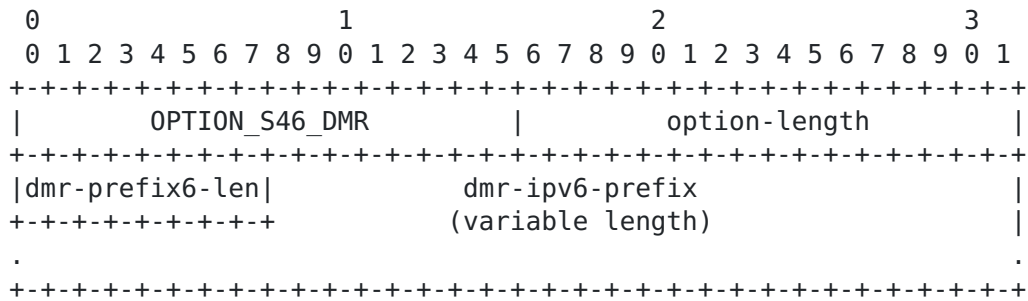


Figure 4: S46 DMR Option

- o option-code: OPTION_S46_DMR (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 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 (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.

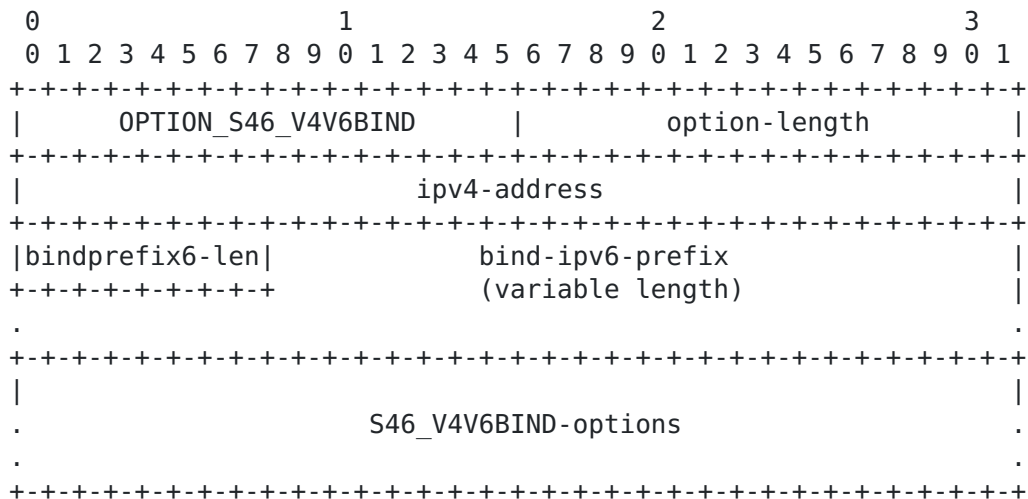


Figure 5: S46 IPv4/IPv6 Address Binding Option

- o option-code: OPTION_S46_V4V6BIND (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. This document specifies one such option, OPTION_S46_PORTPARAMS.

4.5. S46 Port Parameters Option

The Port Parameters Option (OPTION_S46_PORTPARAMS) specifies optional Port Set information that MAY be provided to CEs.

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. Default values for this field are specific to the softwire mechanism being implemented and are defined in the relevant specification document.

- 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 `OPTION_S46_PORTPARAMS` option with an explicit PSID, the client MUST use this explicit PSID in 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 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. Software46 MAP-E Container Option

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

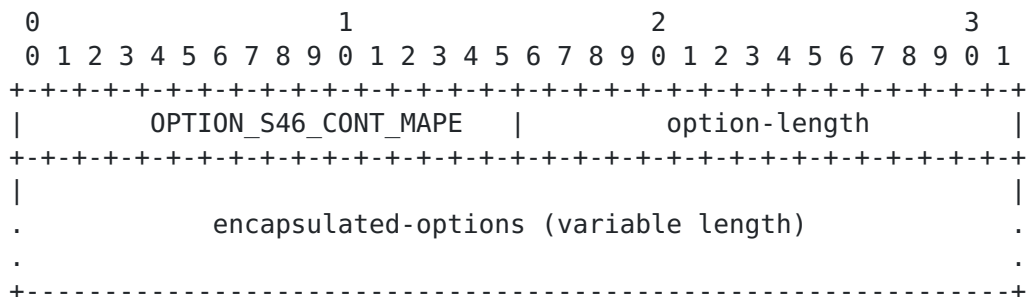


Figure 7: MAP-E Container Option

- o option-code: `OPTION_S46_CONT_MAPE` (TBD6)
- o option-length: Length of encapsulated options

- o encapsulated-options: options associated with this Software46 MAP-E domain.

The encapsulated options field conveys options specific to the `OPTION_S46_CONT_MAPE`. Currently there are two sub-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 DHCP message **MAY** include multiple `OPTION_S46_CONT_MAPE` options (representing multiple domains).

5.2. Software46 MAP-T Container Option

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

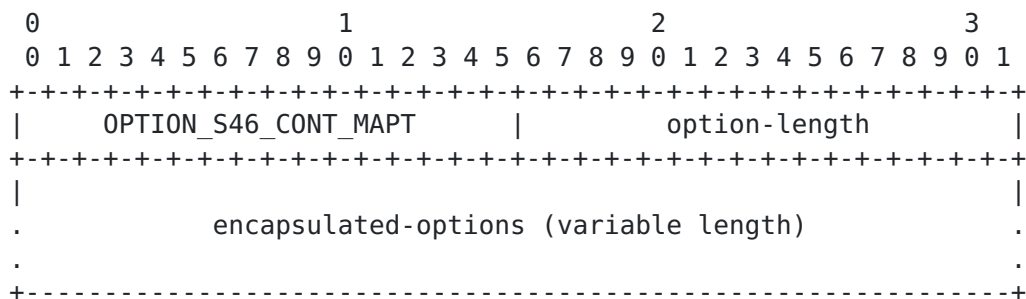


Figure 8: MAP-E Container Option

- o option-code: `OPTION_S46_CONT_MAPT` (TBD7)
- o option-length: Length of encapsulated options
- o encapsulated-options: options associated with this 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. Software46 Lightweight 46 Container Option

The LW46 Container option (`OPTION_S46_CONT_LW`) specifies the container used to group all rules and optional port parameters for a specified domain.

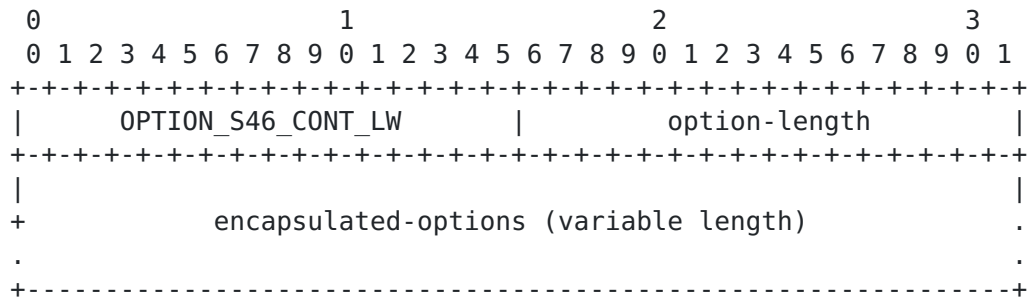


Figure 9: LW46 Container Option

- 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 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. Softwire46 Options Formatting

The below table shows which sub-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/A
OPTION_S46_BR	M	N/A	M
OPTION_S46_PORTPARAMS	O	O	O
OPTION_S46_DMR	N/A	M	N/A
OPTION_S46_V4V6BIND	N/A	N/A	O

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

Table 1: Option to Container Mappings

MAP-DHCP clients that receive container options that violate any of the above rules MUST silently ignore such container options.

7. DHCPv6 Server Behavior

[RFC3315] [Section 17.2.2](#) describes how a DHCPv6 client and server negotiate configuration values using the OR0. As a convenience to 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 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 OR0 option, 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 DHCPv6 client will request S46 configuration parameters from the DHCPv6 server located in the IPv6 network. Such a client MUST include the S46 Container option(s) that it is configured for in its OR0 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 Software46 mechanisms, is out of scope of this document. [[I-D.ietf-software-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 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 [\[I-D.ietf-softwire-map\]](#) for discussion on valid MAP rule combinations. See [Section 11 of \[RFC7227\]](#), Sections [18.1.3](#), [18.1.4](#) and [19.1 of \[RFC3315\]](#) for parameters 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 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.

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

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

Readers concerned with security of MAP provisioning over DHCPv6 are encouraged to read [\[I-D.ietf-dhc-sedhcpv6\]](#).

10. IANA Considerations

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

TBD1 for `OPTION_S46_RULE`

TBD2 for `OPTION_S46_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 \[RFC3315\]](#).

11. 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.

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.
- [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-dhc-sedhcpv6]
Jiang, S., Shen, S., Zhang, D., and T. Jinmei, "Secure DHCPv6", [draft-ietf-dhc-sedhcpv6-04](#) (work in progress), September 2014.
- [I-D.ietf-software-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-software-lw4over6-12](#) (work in progress), November 2014.
- [I-D.ietf-software-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-software-map-11](#) (work in progress), October 2014.

[I-D.ietf-software-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-software-map-t-06](#) (work in progress), October 2014.

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

Boucadair, M., Farrer, I., Perreault, S., and S. Sivakumar, "Unified IPv4-in-IPv6 Software CPE", [draft-ietf-software-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.

[RFC7227] Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", [BCP 187](#), [RFC 7227](#), May 2014.

Authors' Addresses

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

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
CT0-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
6 Floor, C Block, DaCheng International Center Chaoyang District
Beijing 100124
China

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