

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 2, 2012

O. Troan
cisco
S. Matsushima
SoftBank Telecom
T. Murakami
IP Infusion
X. Li
C. Bao
CERNET Center/Tsinghua
University
January 30, 2012

Mapping of Address and Port (MAP)
draft-mdt-softwire-mapping-address-and-port-03

Abstract

This document describes a generic mechanism for mapping between IPv4 addresses and IPv6 addresses and transport layer ports.

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 August 2, 2012.

Copyright Notice

Copyright (c) 2012 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	4
3.	Terminology	5
4.	Architecture	6
5.	Mapping Rules	7
5.1.	Port mapping algorithm	8
5.1.1.	Bit Representation of the Algorithm	9
5.1.2.	GMA examples	9
5.1.3.	GMA Provisioning Considerations	10
5.2.	Basic mapping rule (BMR)	10
5.3.	Forwarding mapping rule (FMR)	13
5.4.	Default mapping rule (DMR)	14
6.	The IPv6 Interface Identifier	15
7.	IANA Considerations	16
8.	Security Considerations	16
9.	Contributors	16
10.	Acknowledgements	16
11.	References	17
11.1.	Normative References	17
11.2.	Informative References	17
	Authors' Addresses	19

1. Introduction

The mechanism of mapping IPv4 addresses in IPv6 addresses has been described in numerous mechanisms dating back to 1996 [[RFC1933](#)]. The Automatic tunneling mechanism described in [RFC1933](#), assigned a globally unique IPv6 address to a host by combining the host's IPv4 address with a well-known IPv6 prefix. Given an IPv6 packet with a destination address with an embedded IPv4 address, a node could automatically tunnel this packet by extracting the IPv4 tunnel end-point address from the IPv6 destination address.

There are numerous variations of this idea, described in 6over4 [[RFC2529](#)], 6to4 [[RFC3056](#)], ISATAP [[RFC5214](#)], and 6rd [[RFC5969](#)]. The differences between these are the use of well-known IPv6 prefixes, or Service Provider assigned IPv6 prefixes, and the position of the embedded IPv4 bits in the IPv6 address. Teredo [[RFC4380](#)] added a twist to this to achieve NAT traversal by also encoding transport layer ports into the IPv6 address. 6rd, to achieve more efficient encoding, allowed for only the suffix of an IPv4 address to be embedded, with the IPv4 prefix being deduced from other provisioning mechanisms.

NAT-PT [[RFC2766](#)](deprecated) combined with a DNS ALG used address mapping to put NAT state, namely the IPv6 to IPv4 binding encoded in an IPv6 address. This characteristic has been inherited by NAT64 [[RFC6146](#)] and DNS64 [[RFC6147](#)] which rely on an address format defined in [[RFC6052](#)]. [[RFC6052](#)] specifies the algorithmic translation of an IPv6 address to IPv4 address. In particular, [[RFC6052](#)] specifies the address format to build IPv4-converted and IPv4-translatable IPv6 addresses. [RFC6052](#) discusses the transport of the port-set information in an IPv4-embedded IPv6 address but the conclusion was the following (excerpt from [[RFC6052](#)]):

"There have been proposals to complement stateless translation with a port range feature. Instead of mapping an IPv4 address to exactly one IPv6 prefix, the options would allow several IPv6 nodes to share an IPv4 address, with each node managing a different set of ports. If a port-set extension is needed, it could be defined later, using bits currently reserved as null in the suffix."

The commonalities of all these IPv6 over IPv4 mechanisms are:

- o Automatically provisions an IPv6 address for a host or an IPv6 prefix for a site
- o Algorithmic or implicit address resolution for tunneling or encapsulation. Given an IPv6 destination address, an IPv4 tunnel endpoint address can be calculated. Likewise for translation, an

IPv4 address can be calculated from an IPv6 destination address and vice versa.

- o Embedding of an IPv4 address or part thereof and optionally transport layer ports into an IPv6 address.

In phases of IPv4 to IPv6 migration, IPv6 only networks will be common, while there will still be a need for residual IPv4 deployment. This document describes a more generic mapping of IPv4 to IPv6 that can be used both for encapsulation (IPv4 over IPv6) and for translation between the two protocols.

Just as the IPv6 over IPv4 mechanisms referred to above, the residual IPv4 over IPv6 mechanisms must be capable of:

- o Provisioning an IPv4 prefix, an IPv4 address or a shared IPv4 address.
- o Algorithmically map between an IPv4 prefix, IPv4 address or a shared IPv4 address and an IPv6 address.

The unified mapping scheme described here supports translation mode, encapsulation mode, in both mesh and hub and spoke topologies.

This document describes delivery of IPv4 unicast service across an IPv6 infrastructure. IPv4 multicast is not considered further in this document.

The A+P (Address and Port) architecture of sharing an IPv4 address by distributing the port space is described in [\[RFC6346\]](#). Specifically [section 4 of \[RFC6346\]](#) covers stateless mapping. The corresponding stateful solution DS-lite is described in [\[RFC6333\]](#). The motivation for work is described in [\[I-D.ietf-softwire-stateless-4v6-motivation\]](#).

A companion document defines a DHCPv6 option for provisioning of MAP [\[I-D.mdt-softwire-map-dhcp-option\]](#). Deployment considerations are described in [\[I-D.mdt-softwire-map-deployment\]](#).

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. Terminology

MAP domain:	A set of MAP CEs and BRs connected to the same virtual link. A service provider may deploy a single MAP domain, or may utilize multiple MAP domains.
MAP Rule	A set of parameters describing the mapping between an IPv4 prefix, IPv4 address or shared IPv4 address and an IPv6 prefix or address. Each MAP node in the domain has the same set of rules.
MAP node	A device that implements MAP.
MAP Border Relay (BR):	A MAP enabled router managed by the service provider at the edge of a MAP domain. A Border Relay router has at least an IPv6-enabled interface and an IPv4 interface connected to the native IPv4 network. A MAP BR may also be referred to simply as a "BR" within the context of MAP.
MAP Customer Edge (CE):	A device functioning as a Customer Edge router in a MAP deployment. A typical MAP CE adopting MAP rules will serve a residential site with one WAN side interface, and one or more LAN side interfaces. A MAP CE may also be referred to simply as a "CE" within the context of MAP.
Port-set:	Each node has a separate part of the transport layer port space; denoted as a port-set.
Port-set ID (PSID):	Algorithmically identifies a set of ports exclusively assigned to the CE.
Shared IPv4 address:	An IPv4 address that is shared among multiple CEs. Only ports that belong to the assigned port-set can be used for communication. Also known as a Port-Restricted IPv4 address.
End-user IPv6 prefix:	The IPv6 prefix assigned to an End-user CE by other means than MAP itself. E.g. provisioned using DHCPv6 PD [RFC3633] or configured manually. It is unique for each CE.

MAP IPv6 address:	The IPv6 address used to reach the MAP function of a CE from other CE's and from BR's.
Rule IPv6 prefix:	An IPv6 prefix assigned by a Service Provider for a mapping rule.
Rule IPv4 prefix:	An IPv4 prefix assigned by a Service Provider for a mapping rule.
IPv4 Embedded Address (EA) bits:	The IPv4 EA-bits in the IPv6 address identify an IPv4 prefix/address (or part thereof) or a shared IPv4 address (or part thereof) and a port-set identifier.
MRT:	MAP Rule table. Address and Port aware datastructure, supporting longest match lookups. The MRT is used by the MAP forwarding function.

4. Architecture

A full IPv4 address or IPv4 prefix can be used like today, e.g. for identifying an interface or as a DHCP pool. A shared IPv4 address on the other hand, MUST NOT be used to identify an interface. While it is theoretically possible to make host stacks and applications port-aware, that is considered a too drastic change to the IP model [[RFC6250](#)].

The MAP architecture described here, restricts the use of the shared IPv4 address to only be used as the global address (outside) of the NAPT [[RFC2663](#)] running on the CE. The NAPT MUST in turn be connected to a MAP aware forwarding function, that does encapsulation/decapsulation or translation to IPv6.

For packets outbound from the private IPv4 network, the CE NAPT MUST translate transport identifiers (e.g. TCP and UDP port numbers) so that they fall within the assigned CE's port-range.

The forwarding function uses the MRT to make forwarding decisions. The table consist of the mapping rules. An entry in the table consists of an IPv4 prefix and PSID. The normal best matching prefix algorithm is used. With a maximum key length of 48 (32 + 16). E.g. with a sharing ratio of 64 (6 bit PSID length) a host route for this CE would be a /38 (32 + 6).

5. Mapping Rules

A MAP node is provisioned with one or more mapping rules.

Mapping rules are used differently depending on their function. Every MAP node must be provisioned with a Basic mapping rule. This is used by the node to configure itself with an IPv4 address, IPv4 prefix or shared IPv4 address from an End-user IPv6 prefix. This same basic rule can also be used for forwarding, where an IPv4 destination address and optionally a destination port is mapped into an IPv6 address or prefix. Additional mapping rules can be specified to allow for e.g. multiple different IPv4 subnets to exist within the domain. Additional mapping rules are recognized by having a Rule IPv6 prefix different from the base End-user IPv6 prefix.

Traffic outside of the domain (IPv4 address not matching (using longest matching prefix) any Rule IPv4 prefix in the Rules database) will be forward using the Default mapping rule. The Default mapping rule maps outside destinations to the BR's IPv6 address or prefix.

There are three types of mapping rules:

1. Basic Mapping Rule - used for IPv4 prefix, address or port set assignment. There can only be one Basic Mapping Rule per End-user IPv6 prefix. The Basic Mapping Rule is used to configure the MAP IPv6 address or prefix.
 - * Rule IPv6 prefix (including prefix length)
 - * Rule IPv4 prefix (including prefix length)
 - * Rule EA-bits length (in bits)
 - * Rule Port Parameters (optional)
2. Forwarding Mapping Rule - used for forwarding. The Basic Mapping Rule is also a Forwarding Mapping Rule. Each Forwarding Mapping Rule will result in an entry in the MRT for the Rule IPv4 prefix.
 - * Rule IPv6 prefix (including prefix length)
 - * Rule IPv4 prefix (including prefix length)
 - * Rule EA-bits length (in bits)
 - * Rule Port Parameters (optional)

3. Default Mapping Rule - used for destinations outside the MAP domain. A 0.0.0.0/0 entry is installed in the MRT for this rule.

- * Rule IPv6 prefix (including prefix length)

- * Rule BR IPv4 address

A MAP node finds its Basic Mapping Rule by doing a longest match between the End-user IPv6 prefix and the Rule IPv6 prefix in the Mapping Rule database. The rule is then used for IPv4 prefix, address or shared address assignment.

A MAP IPv6 address (or prefix) is formed from the BMR Rule IPv6 prefix. This address MUST be assigned to an interface of the MAP node and is used to terminate all MAP traffic being sent or received to the node.

Port-aware IPv4 entries in the MRT are installed for all the Forwarding Mapping Rules and an IPv4 default route for the Default Mapping Rule.

In hub and spoke mode, all traffic MUST be forwarded using the Default Mapping Rule.

5.1. Port mapping algorithm

Different Port-Set Identifiers (PSID) MUST have non-overlapping port-sets. The two extreme cases are: (1) the port numbers are not contiguous for each PSID, but uniformly distributed across the port range (0-65535); (2) the port numbers are contiguous in a single range for each PSID. The port mapping algorithm proposed here is called the Generalized Modulus Algorithm (GMA) and supports both these cases.

For a given sharing ratio (R) and the maximum number of contiguous ports (M), the GMA algorithm is defined as:

1. The port number (P) of a given PSID (K) is composed of:

$$P = R * M * j + M * K + i$$

Where:

- * PSID: $K = 0$ to $R - 1$

- * Port range index: $j = (4096 / M) / R$ to $((65536 / M) / R) - 1$, if the port numbers (0 - 4095) are excluded.

* Contiguous Port index: $i = 0$ to $M - 1$

2. The PSID (K) of a given port number (P) is determined by:

$$K = (\text{floor}(P/M)) \% R$$

Where:

* % is the modulus operator

* floor(arg) is a function that returns the largest integer not greater than arg.

5.1.1. Bit Representation of the Algorithm

Given a sharing ratio ($R=2^k$), the maximum number of contiguous ports ($M=2^m$), for any PSID (K) and available ports (P) can be represented as:




Figure 1: Bit representation

Where j and i are the same indexes defined in the port mapping algorithm.

For any port number, the PSID can be obtained by bit mask operation.

For $a > 0$, j MUST be larger than 0. This ensures that the algorithm excludes the system ports ([\[I-D.ietf-tsvwg-iana-ports\]](#)). For $a = 0$, j MAY be 0 to allow for the provisioning of the system ports.

5.1.2. GMA examples

For example, for $R = 1024$, PSID offset: $a = 4$ and PSID length: $k = 10$ bits

	Port-set-1	Port-set-2
PSID=0	4096, 4097, 4098, 4099,	8192, 8193, 8194, 8195, ...
PSID=1	4100, 4101, 4102, 4103,	8196, 8197, 8198, 8199, ...
PSID=2	4104, 4105, 4106, 4107,	8200, 8201, 8202, 8203, ...
PSID=3	4108, 4109, 4110, 4111,	8204, 8205, 8206, 8207, ...
...		
PSID=1023	8188, 8189, 8190, 8191,	12284, 12285, 12286, 12287, ...

Example 1: with offset = 4 ($a = 4$)

For example, for $R = 64$, $a = 0$ (PSID offset = 0 and PSID length = 6 bits):

	Port-set
PSID=0	[0 - 1023]
PSID=1	[1024 - 2047]
PSID=2	[2048 - 3071]
PSID=3	[3072 - 4095]
...	
PSID=63	[64512 - 65535]

Example 2: with offset = 0 ($a = 0$)

5.1.3. GMA Provisioning Considerations

The number of offset bits (a) and excluded ports are optionally provisioned via the "Rule Port Mapping Parameters" in the Basic Mapping Rule.

The defaults are:

- o Excluded ports : 0-4095
- o Offset bits (a) : 4

To simplify the GMA port mapping algorithm the defaults are chosen so that the PSID field starts on a nibble boundary and the excluded port range (0-1023) is extended to 0-4095.

5.2. Basic mapping rule (BMR)

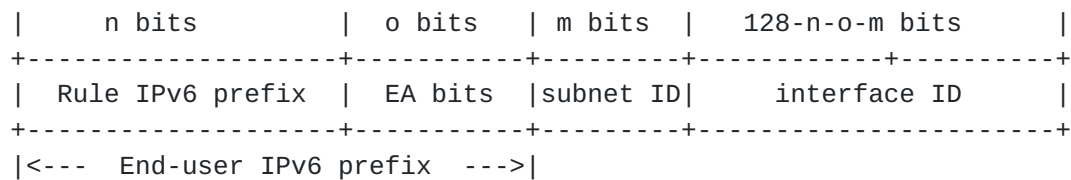


Figure 2: IPv6 address format

The Embedded Address bits (EA bits) are unique per end user within a Rule IPv6 prefix. The Rule IPv6 prefix is the part of the End-user IPv6 prefix that is common among all CEs using the same Basic Mapping Rule within the MAP domain. The EA bits encode the CE specific IPv4 address and port information. The EA bits can contain a full or part of an IPv4 prefix or address, and in the shared IPv4 address case contains a Port-Set Identifier (PSID).

The MAP IPv6 address is created by concatenating the End-user IPv6 prefix with the MAP subnet-id and the interface-id as specified in [Section 6](#).

The MAP subnet ID is defined to be the first subnet (all bits set to zero). A MAP node MUST reserve the first IPv6 prefix in a End-user IPv6 prefix for the purpose of MAP.

Shared IPv4 address:




Figure 3: Shared IPv4 address

Complete IPv4 address:




Figure 4: Complete IPv4 address

IPv4 prefix:




Figure 5: IPv4 prefix

The length of r MAY be zero, in which case the complete IPv4 address or prefix is encoded in the EA bits. If only a part of the IPv4 address/prefix is encoded in the EA bits, the Rule IPv4 prefix is provisioned to the CE by other means (e.g. a DHCPv6 option). To create a complete IPv4 address (or prefix), the IPv4 address suffix (p) from the EA bits, are concatenated with the Rule IPv4 prefix (r bits).

The offset of the EA bits field in the IPv6 address is equal to the BMR Rule IPv6 prefix length. The length of the EA bits field (o) is given by the BMR Rule EA-bits length. The sum of the Rule IPv6 Prefix length and the Rule EA-bits length MUST be less or equal than the End-user IPv6 prefix length.

If $o + r < 32$ (length of the IPv4 address in bits), then an IPv4 prefix is assigned.

If $o + r$ is equal to 32, then a full IPv4 address is to be assigned. The address is created by concatenating the Rule IPv4 prefix and the EA-bits.

If $o + r$ is > 32 , then a shared IPv4 address is to be assigned. The number of IPv4 address suffix bits (p) in the EA bits is given by $32 - r$ bits. The PSID bits are used to create a port-set. The length of the PSID bit field within EA bits is: $o - p$.

In the following examples, only the suffix (last 8 bits) of the IPv4 address is embedded in the EA bits ($r = 24$), while the IPv4 prefix (first 24 bits) is given in the BMR Rule IPv4 prefix.

Example:

Given:

End-user IPv6 prefix: 2001:db8:0012:3400::/56
Basic Mapping Rule: {2001:db8:0000::/40 (Rule IPv6 prefix),
192.0.2.0/24 (Rule IPv4 prefix),
16 (Rule EA-bits length)}
Sharing ratio: 256 ($16 - (32 - 24) = 8$. $2^8 = 256$)
PSID offset: 4

We get IPv4 address and port-set:

EA bits offset: 40
IPv4 suffix bits (p): Length of IPv4 address (32) -
IPv4 prefix length (24) = 8
IPv4 address: 192.0.2.18

PSID start: $40 + p = 40 + 8 = 48$
PSID length: $o - p = 16 (56 - 40) - 8 = 8$
PSID: 0x34
Port-set-1: 4928, 4929, 4930, 4931, 4932, 4933, 4934, 4935,
4936, 4937, 4938, 4939, 4940, 4941, 4942, 4943
Port-set-2: 9024, 9025, 9026, 9027, 9028, 9029, 9030, 9031,
9032, 9033, 9034, 9035, 9036, 9037, 9038, 9039
...
Port-set-15: 62272, 62273, 62274, 62275,
62276, 62277, 62278, 62279,
62280, 62281, 62282, 62283,
62284, 62285, 62286, 62287,

5.3. Forwarding mapping rule (FMR)

On adding an FMR rule, an IPv4 route is installed in the AP RIB for the Rule IPv4 prefix.

On forwarding an IPv4 packet, a best matching prefix lookup is done in the IPv4 routing table and the correct FMR is chosen.

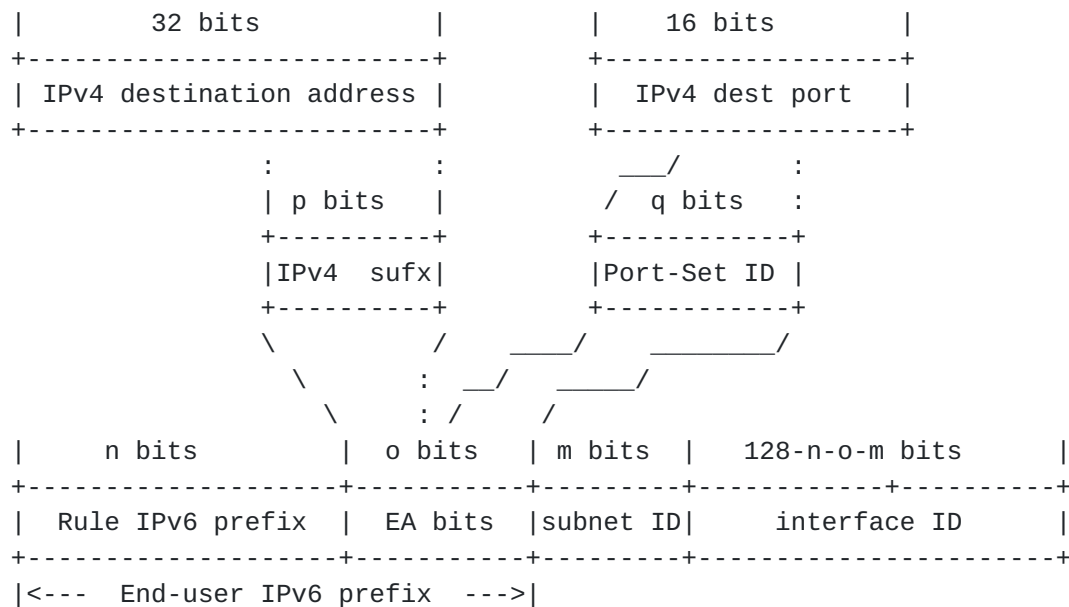


Figure 6: Deriving of MAP IPv6 address

Example:

Given:

IPv4 destination address: 192.0.2.18

IPv4 destination port: 9030

Forwarding Mapping Rule: {2001:db8:0000::/40 (Rule IPv6 prefix),
192.0.2.0/24 (Rule IPv4 prefix),
16 (Rule EA-bits length)}

PSID offset: 4

We get IPv6 address:

IPv4 suffix bits (p): 32 - 24 = 8 (18 (0x12))

PSID length: 8

PSID: 0x34 (9030 (0x2346))

EA bits: 0x1234

MAP IPv6 address: 2001:db8:0012:3400:00c0:0002:1200:3400

5.4. Default mapping rule (DMR)

The Default Mapping rule is used to reach IPv4 destinations outside of the MAP domain. Traffic using this rule will be sent from a CE to a BR.

The Rule IPv4 prefix in the DMR is: 0.0.0.0/0. The Rule IPv6 prefix is the IPv6 address or prefix of the BR. Which is used, is dependent on the mode used. For example translation requires that the IPv4 destination address is encoded in the BR IPv6 address, so only a

prefix is used in the DMR to allow for a generated interface identifier. For the encapsulation mode the Rule IPv6 prefix can be the full IPv6 address of the BR.

There MUST be only one Default Mapping Rule within a MAP domain.

Default Mapping Rule:

```
{2001:db8:0001:0000:<interface-id>:/128 (Rule IPv6 prefix),  
  0.0.0.0/0 (Rule IPv4 prefix),  
  192.0.2.1 (BR IPv4 address)}
```

Example 3: Default Mapping Rule

In most implementations of a routing table, the next-hop address must be of the same address family as the prefix. To satisfy this requirement a BR IPv4 address is included in the rule. Giving a default route in the IPv4 routing table:

```
0.0.0.0 -> 192.0.2.1, MAP-Interface0
```

6. The IPv6 Interface Identifier

The Interface identifier format is based on the format specified in [section 2.2 of \[RFC6052\]](#), with the added PSID format field.

In an encapsulation solution, an IPv4 address and port is mapped to an IPv6 address. This is the address of the tunnel end point of the receiving MAP CE. For traffic outside the MAP domain, the IPv6 tunnel end point address is the IPv6 address of the BR. The interface-id used for all MAP nodes in the domain MUST be deterministic.

When translating, the destination IPv4 address is translated into a corresponding IPv6 address. In the case of traffic outside of the MAP domain, it is translated to the BR's IPv6 prefix. For the BR to be able to reverse the translation, the full destination IPv4 address must be encoded in the IPv6 address. The same thing applies if an IPv4 prefix is encoded in the IPv6 address, then the reverse translator needs to know the full destination IPv4 address, which has to be encoded in the interface-id.

The encoding of the full IPv4 address into the interface identifier, both for the source and destination IPv6 addresses have been shown to be useful for troubleshooting.


```

+---+---+---+---+---+---+---+---+
|PL|  8  16  24  32  40  48  56  |
+---+---+---+---+---+---+---+---+
|64| u | IPv4 address  |  PSID | 0 |
+---+---+---+---+---+---+---+---+

```

Figure 7

In the case of an IPv4 prefix, the IPv4 address field is right-padded with zeroes up to 32 bits. The PSID field is left-padded to create a 16 bit field. For an IPv4 prefix or a complete IPv4 address, the PSID field is zero.

If the End-user IPv6 prefix length is larger than 64, the most significant parts of the interface identifier is overwritten by the prefix. For translation mode the End-user IPv6 prefix MUST be 64 or shorter.

7. IANA Considerations

This specification does not require any IANA actions.

8. Security Considerations

Specific security considerations with the MAP mechanism are detailed in the encapsulation and translation documents [I-D.mdt-map-t/I-D.mdt-map-e].

[RFC6269] outlines general issues with IPv4 address sharing.

9. Contributors

Mohamed Boucadair, Gang Chen, Maoke Chen, Wojciech Dec, Xiaohong Deng, Jouni Korhonen, Tomasz Mrugalski, Jacni Qin, Chunfa Sun, Qiong Sun, Leaf Yeh.

10. Acknowledgements

This document is based on the ideas of many. In particular Remi Despres, who has tirelessly worked on generalized mechanisms for stateless address mapping.

The authors would like to thank Guillaume Gottard, Dan Wing, Jan

Zorz, Necj Scoberne, Tina Tsou for their thorough review and comments.

11. References

11.1. Normative References

- [I-D.mdt-software-map-dhcp-option]
Mrugalski, T., Boucadair, M., Deng, X., Troan, O., and C. Bao, "DHCPv6 Options for Mapping of Address and Port", [draft-mdt-software-map-dhcp-option-02](#) (work in progress), January 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC6346] Bush, R., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", [RFC 6346](#), August 2011.

11.2. Informative References

- [I-D.ietf-software-stateless-4v6-motivation]
Boucadair, M., Matsushima, S., Lee, Y., Bonness, O., Borges, I., and G. Chen, "Motivations for Stateless IPv4 over IPv6 Migration Solutions", [draft-ietf-software-stateless-4v6-motivation-00](#) (work in progress), September 2011.
- [I-D.ietf-tsvwg-iana-ports]
Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", [draft-ietf-tsvwg-iana-ports-10](#) (work in progress), February 2011.
- [RFC1933] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", [RFC 1933](#), April 1996.
- [RFC2529] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", [RFC 2529](#), March 1999.
- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.
- [RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address

Translation - Protocol Translation (NAT-PT)", [RFC 2766](#), February 2000.

- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", [RFC 4380](#), February 2006.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", [RFC 5214](#), March 2008.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", [RFC 5969](#), August 2010.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), October 2010.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), April 2011.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [RFC 6147](#), April 2011.
- [RFC6250] Thaler, D., "Evolution of the IP Model", [RFC 6250](#), May 2011.
- [RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), June 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [RFC 6333](#), August 2011.

Authors' Addresses

Ole Troan
cisco
Oslo
Norway

Email: ot@cisco.com

Satoru Matsushima
SoftBank Telecom
1-9-1 Higashi-Shinbashi, Munato-ku
Tokyo
Japan

Email: satoru.matsushima@tm.softbank.co.jp

Tetsuya Murakami
IP Infusion
1188 East Arques Avenue
Sunnyvale
USA

Email: tetsuya@ipinfusion.com

Xing Li
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Email: xing@cernet.edu.cn

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

Email: congxiao@cernet.edu.cn

