

Softwires
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2012

C. Bao
X. Li
Y. Zhai
CERNET Center/Tsinghua
University
T. Murakami, Ed.
IP Infusion
W. Dec, Ed.
Cisco Systems
March 9, 2012

MAP Translation (MAP-T) - specification
draft-mdt-softwire-map-translation-01

Abstract

This document specifies the "Mapping of Address and Port" (MAP) double stateless translation based solution (MAP-T) for providing IPv4 hosts connectivity to and across an IPv6 domain.

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 10, 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

Internet-Draft

Map Translation

March 2012

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.	Extended Contributors List	3
2.	Introduction	4
3.	Requirements Language	5
4.	Terminology	5
5.	MAP-T Translation Framework	7
6.	MAP-T Node Behavior	8
6.1.	Provisioning of MAP-T CE	8
6.2.	Packet Forwarding Behavior of MAP-T CE	9
6.2.1.	IPv4 to IPv6	9
6.2.2.	IPv6 to IPv4	9
6.3.	Provisioning of MAP-T BR	9
6.4.	Packet Forwarding Behavior on MAP-T BR	10
6.4.1.	IPv6 to IPv4	10
6.4.2.	IPv4 to IPv6	10
7.	MAP-T IPv4/IPv6 Translation Specifications	10
7.1.	Address Formats	11
7.2.	Translating IPv4 Address and Port Number into IPv6 Address and Port Number at the BR	11
7.3.	Translating IPv6 Address and Port Number into IPv4 Address and Port Number at the BR	12
7.4.	Translating IPv4 Address and Port Number into IPv6 Address and Port Number at the CE	12
7.5.	Translating IPv6 Address and Port Number into IPv4 Address and Port Number at the CE	13
7.6.	Translating ICMP/ICMPv6 Headers	13
7.7.	Path MTU Discovery and Fragmentation	13
8.	MAP-T Packet Forwarding considerations	14
8.1.	Mesh Model	14
8.2.	Hub & Spoke model	15
8.3.	Communication with IPv6 servers in the MAP-T domain	15
9.	NAT44 considerations	15
10.	Security Considerations	15
11.	IANA Consideration	16
12.	Acknowledgements	16
13.	References	16

13.1 . Normative References	16
13.2 . Informative References	17
Appendix A . Example of MAP-T translation	19
Authors' Addresses	21

[1](#). Extended Contributors List

This document is the result of the IETF Softwire MAP design team effort and numerous previous individual contributions in this area initiated by dIVI [[I-D.xli-behave-divi](#)] along with a similar idea proposed by [[I-D.murakami-softwire-4v6-translation](#)]. The following are the authors who contributed in a major way to this document:

Chongfeng Xie (China Telecom)

Room 708, No.118, Xizhimennei Street Beijing 100035 CN

Phone: +86-10-58552116

Email: xiechf@ctbri.com.cn

Chongfeng Xie (China Telecom)

Room 708, No.118, Xizhimennei Street Beijing 100035 CN

Phone: +86-10-58552116

Email: xiechf@ctbri.com.cn

Qiong Sun (China Telecom)

Room 708, No.118, Xizhimennei Street Beijing 100035 CN

Phone: +86-10-58552936

Email: sunqiong@ctbri.com.cn

Satoru Matsushima (Softbank Telecom)

1-9-1 Higashi-Shinbashi, Munato-ku, Tokyo, Japan

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

Gang Chen (China Mobile)

53A,Xibianmennei Ave. Beijing 100053 P.R.China

Email: chengang@chinamobile.com

Wentao Shang (CERNET Center/Tsinghua University)

Room 225, Main Building, Tsinghua University Beijing 100084 CN

Internet-Draft

Map Translation

March 2012

Guoliang Han (CERNET Center/Tsinghua University)
Room 225, Main Building, Tsinghua University Beijing 100084 CN
Email: bupthgl@gmail.com

Rajiv Asati (Cisco Systems)
7025-6 Kit Creek Road Research Triangle Park NC 27709 USA
Email: rajiva@cisco.com

[2.](#) Introduction

Experiences from several years of IPv6 deployment [[RFC6219](#)] indicates that transitioning a service providers' domain fully to IPv6-only requires not only the continued support of legacy IPv4 communication across that domain, but also the need for an ultimate IPv4 exit strategy allowing communication between IPv4 and IPv6 address families in that domain. The use of an IPv4/IPv6 translation based solution is an optimal way to address these requirements, particularly in combination with stateless translation techniques that seek to minimize complexities as described in [[I-D.operators-softwire-stateless-4v6-motivation](#)]. The double Pv4/IPv6 translation based solution, MAP-T, is such a solution, and one that builds on existing stateless IPv4/IPv6 address translation techniques specified in [[RFC6052](#)], [[RFC6144](#)], and [[RFC6145](#)], by:

- o Extending stateless IPv4/IPv6 translation with algorithmic address and port mapping rules as defined in MAP MAP [[I-D.mdt-softwire-mapping-address-and-port](#)].
- o Introducing the notion of stateless double IPv4/IPv6 translation that can restore the original IPv4 address.
- o Allowing IPv4-translatable addresses to be either fully or partially encoded in IPv6 prefixes (or addresses) assigned to customers.

The MAP-T solution presents an operator with the prospect of a full transition of a domain to IPv6-only, in a manner that:

Internet-Draft

Map Translation

March 2012

- o Retains the ability for IPv4 end hosts to communicate across the IPv6 domain with other IPv4 hosts.
- o Permits both individual IPv4 address assignment as well as IPv4 address sharing, with predefined port ranges, to be enacted using IPv6.
- o Allows communication between IPv4-only, as well as any IPv6 enabled end hosts, to native IPv6-only servers in the domain that are using IPv4-mapped IPv6 address.
- o Does not require the operation of an IPv4 overlay network, nor the introduction of non native-IPv6 network device or server functionality.
- o Allows the use of IPv6 native network operations, including the ability to classify IP traffic, as well as to perform IP traffic routing optimization policies, e.g. routing optimization based on peering policies for Internet IPv4 destinations outside of the domain.

[3.](#) Requirements Language

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](#)].

[4.](#) Terminology

MAP-T: Mapping of Address and Port - Translation mode. MAP-T utilizes IPv4/IPv6 translation as per [[RFC6145](#)] along with the MAP extensions for mapping between IPv4 and IPv6 defined in MAP [[I-D.mdt-software-mapping-address-and-port](#)] and this draft.

MAP-T domain (Domain): A set of MAP-T CEs and BRs,. A service provider may deploy MAP-T with a single MAP-T domain, or may utilize multiple MAP-T domains. Each domain requires a separate MAP-T rule set.

MAP-T Border Relay (BR): A MAP-T enabled router/translator at the edge of a MAP-T domain, providing connectivity to the MAP-T domain. A Border Relay router has at least an IPv6-enabled interface and an IPv4 interface connected to the native IPv4 network,

and it can serve multiple MAP-T domains.

MAP-T Customer Edge (CE): A router/translator node functioning as a Customer Edge Router/translator in a MAP-T domain. This type of device is sometimes referred to as a "Residential Gateway" (RG) or "Customer Premises Equipment" (CPE). A typical MAP-T CE adopting MAP rules will serve a residential site with one WAN side interface, one or more LAN side interfaces. A MAP-T CE may also be referred to simply as a "CE" within the context of MAP-T.

Shared IPv4 address: An IPv4 address that is shared among multiple MAP CE nodes. Each node has a separate part of the transport layer port space.

MAP-T Rule: A MAP rule defining the mapping relationship for a given MAP-T domain between IPv4 and IPv6, defined in MAP [\[I-D.mdt-software-mapping-address-and-port\]](#). Three such rules are the BMR, DMR, and FMR.

Basic Mapping Rule (BMR): A mandatory rule governing the relationship between the IPv4 prefix, address or port set IPv6 address and MAP domain configuration information. The BMR is used for configuring the MAP CE. The BMR is effectively a type of FMR.

Default Mapping Rule (DMR): A mandatory rule used for mapping of IPv4 information into IPv6 for destinations outside a MAP domain. Can be thought of as representing an IPv4 0.0.0.0/0 default route.

Forward Mapping Rule (FMR): An optional rule for mapping between

specific IPv4 and IPv6 destinations within a MAP domain. Can be thought of as representing a more specific IPv4 route in the MAP domain. Finds application primarily on CEs where forwarding using more specific routes is desired. To a BR, the BMR and FMR are effectively the same.

5. MAP-T Translation Framework

Figure 1 depicts the overall MAP-T architecture with IPv4 users (N and M) networks connected to a routed IPv6 network.

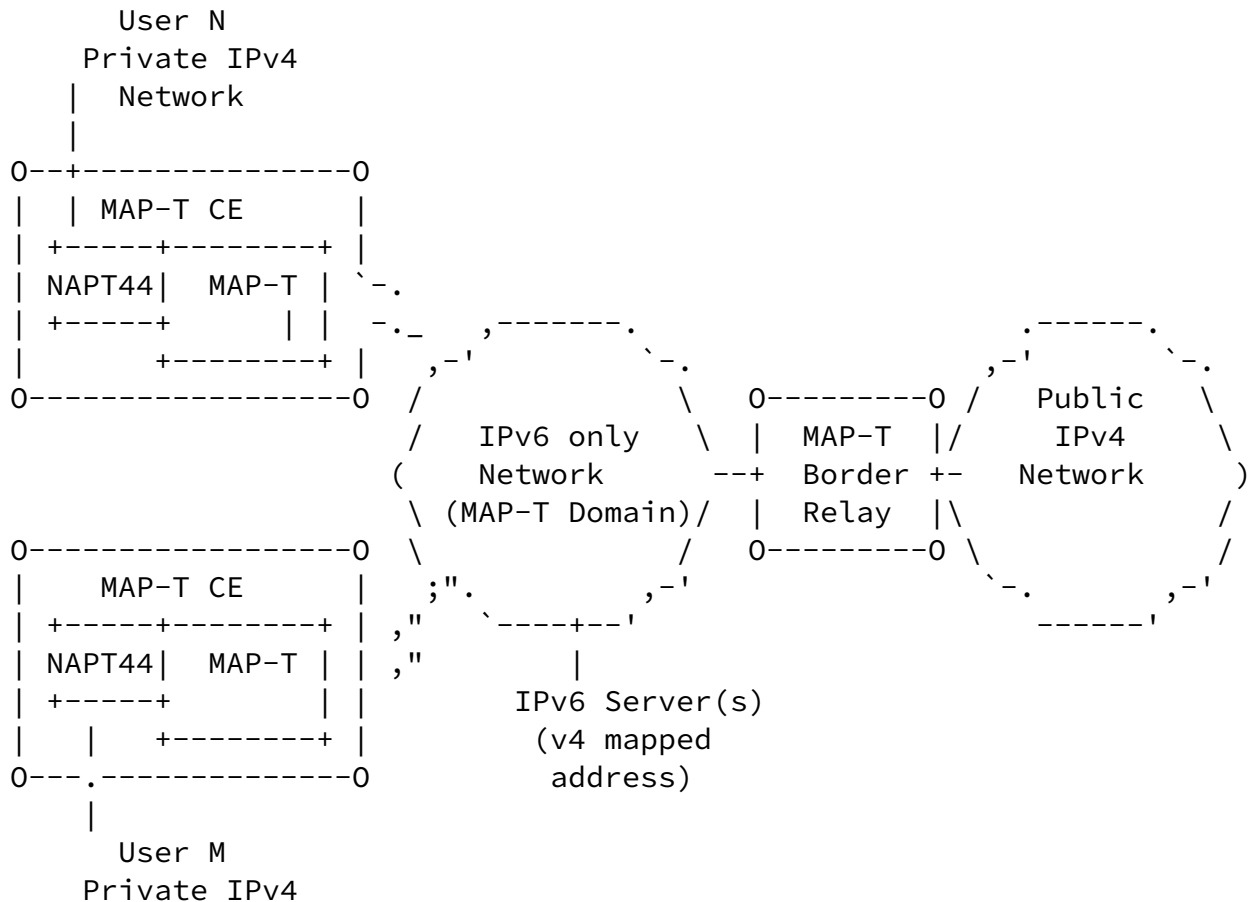


Figure 1: Network Topology

Figure 1: Network Topology

The MAP-T solution relies on IPv4/IPv6 translating components, the MAP-T CE and MAP-T BR, connected to a MAP-T domain. The MAP-T CE is responsible for connecting a users' private IPv4, along with any native IPv6 network to the IPv6-only MAP-T domain. To multiplex multiple IPv4 user hosts, the CE relies on regular NAT44 functionality, which is however configured based on MAP-T settings. The CE's stateless IPv4/IPv6 translation function [[RFC6145](#)], again configured to operate based on MAP-T settings, completes the model of the CE defined in Figure 1. The CE's MAP-T domain facing interface is configured with a regular operator assigned IPv6 prefix that can be the same as that used to address any native IPv6 (non MAP-T) user network devices i.e. MAP-T does not require more than one IPv6 prefix per user network, and supports regular IPv6 prefix or address

assignment mechanism including SLAAC and DHCPv6 stateful.

The MAP-T BR is responsible for connecting external IPv4 networks to all devices in one or more MAP-T domains, using stateless IPv4/IPv6 translation [[RFC6145](#)] extended by the MAP-T rules as per this document. Besides the CE and BR, the MAP-T domain can contain any regular IPv6-only hosts/servers that have an IPv4 mapped IPv6 address (IPv4-translatable address per [[RFC6052](#)]) using a prefix assigned to the MAP-T domain. Communication with such devices is naturally possible using native IPv6 means from inside or outside the domain as well as from any IPv4-only hosts inside or outside of the MAP-T domain.

The IPv4 in IPv6 address mapping scheme employed by the MAP-T solution, along with the avoidance of using any additional encapsulating headers allows the MAP-T domain to be operated using regular native IPv6 functionality. This includes also the ability to classify traffic based on specific source and destination addresses (including any IPv4 in IPv6 mapped source and destinations), and higher layer packet payload. Similarly, the address mapping characteristic allows IPv6 traffic forwarding in the MAP-T domain to be optimized in line with an operators' policies, e.g. native IPv6

routing selection of MAP-T domain egress points based on peering policies bound to IPv4 destination. IP Traffic between CEs in any MAP-T can flow either in hub & spoke modes, with a BR acting as the spoke, or in mesh mode directly between the CEs.

6. MAP-T Node Behavior

6.1. Provisioning of MAP-T CE

A MAP-T CE requires the following parameters for provisioning:

- o The MAP Domain IPv4 and IPv6 prefix, and their lengths (Basic Mapping Rule)
- o The MAP EA-bits (CE index), including IPv4 suffix, length and any port-range (including any excluded ports and the port number continuity parameter)
- o The MAP domain BR IPv6 prefix and its length (Default Mapping Rule)

A MAP-T CE that receives a MAP DHCP option [[I-D.mdt-software-map-dhcp-option](#)] and performs the following (MAP initialization) functions:

- o Configures the IPv4 address along with any applicable NAT44 port-

range function parameters (BMR)

- o Configures additional IPv4/IPv6 stateless translation parameters - optional FMRs.

6.2. Packet Forwarding Behavior of MAP-T CE

6.2.1. IPv4 to IPv6

A MAP-T CE receiving IPv4 packets SHOULD perform NAT44 function first and create appropriate NAT44 stateful bindings. The resulting IPv4 packets MUST contain the source IPv4 address and source transport number defined by MAP-T. The resulting IPv4 packet is forwarded to the CE's MAP-T function that performs IPv4 to IPv6 stateless translation. The IPv6 source and destination addresses MUST then be

derived as per [Section 6](#) of this draft, and the IPv4 header MUST be replaced with an IPv6 header following [\[RFC6145\]](#).

[6.2.2.](#) IPv6 to IPv4

A MAP-T CE receiving an IPv6 packet performs its regular IPv6 operations, whereby only packets that are addressed to the MAP-T CE's MAP derived BMR address are forwarded to the CE's MAP-T function. All other IPv6 traffic is forwarded as per the CE's IPv6 routing rules. The CE SHOULD check that MAP-T received packets' transport-layer destination port number is in the range configured by MAP for the CE and the CE SHOULD drop any non conforming packet and respond with an ICMPv6 "Address Unreachable" (Type 1, Code 3). In other cases, the MAP-T function MUST derive the IPv4 source and destination addresses as per [Section 6](#) of this draft and MUST replace the IPv6 header with an IPv4 header in accordance with [\[RFC6145\]](#). The resulting IPv4 packet is then forwarded to the CE's NAT44 function where the destination port number MUST be checked against the stateful port mapping session table and the destination port number MUST be mapped to its original value.

[6.3.](#) Provisioning of MAP-T BR

The MAP-T BR needs to be provisioned with information for the MAP-T domain or domains it is expected to handle, along with any necessary routing processes. For each MAP-T domain, the BR will have the following parameters:

- o The MAP Domain IPv4 and IPv6 prefix and their lengths (Basic Mapping Rule).
- o The BR prefix and its length (Default Mapping Rule)

- o Optionally, any specific Forward Mapping Rules applicable to the domain.

[6.4.](#) Packet Forwarding Behavior on MAP-T BR

[6.4.1.](#) IPv6 to IPv4

A MAP-T BR receiving IPv6 packets selects a best matching MAP-T

domain rule based on a longest address match of the packets' source address against the BR's configured MAP-T BMR prefix(es), as well as a match of the packet destination address against the configured BR prefixes or FMR prefix(es). The selected MAP rule allows the BR to determine the CE-index from the source IPv6 address. The BR MUST perform a validation of the consistency of the source IPv6 address and source port number for the packet using BMR. If the packets source port number is found to be outside the range allowed for this CE-index and the BMR, the BR MUST drop the packet and respond with an ICMPv6 "Destination Unreachable, Source address failed ingress/egress policy" (Type 1, Code 5).

For packets that are to be forwarded outside of a MAP-T domain, the BR MUST derive the source and destination IPv4 addresses as per [Section 7](#) of this draft and translate the IPv6 to IPv4 headers following [\[RFC6145\]](#). The resulting IPv4 packets are then passed to regular IPv4 forwarding.

[6.4.2](#). IPv4 to IPv6

A MAP-T BR receiving IPv4 packets uses a longest match IPv4 lookup to select the target MAP-T domain and rule. The BR MUST then derive the IPv6 source and destination addresses from the IPv4 source and destination address and port as per [Section 7](#) of this draft. Following this, the BR MUST translate the IPv4 to IPv6 headers following [\[RFC6145\]](#). The resulting IPv6 packets are then passed to regular IPv6 forwarding.

Note that the operation of a BR when forwarding to MAP-T domains that do not utilize IPv4 address sharing, is the same as stateless IPv4/IPv6 translation.

[7](#). MAP-T IPv4/IPv6 Translation Specifications

This section specifies the MAP-T IPv6 address format and IPv4-IPv6 address mapping behaviour, based on the MAP MAP [\[I-D.mdt-software-mapping-address-and-port\]](#) specification. Numeric examples of the MAP-T address translation in action are given in [Appendix A](#).

[7.1](#). Address Formats

The MAP-T address format of the (mapped) CE address adopts the format defined in MAP [[I-D.mdt-softwire-mapping-address-and-port](#)]. It is used in mapping rules operations to construct the source and destination IPv6 addresses. An example, is shown in Figure 2 for the specific case of n+m bits less or equal to 64 bit length, where the (optional) well known m subnet-Id bits are used to auto-complete a prefix up to the 64th bit. In cases where the End-user IPv6 prefix n+o bits exceed a length of 64, the excess bits are "bit spread" across the fixed u-octet boundary as needed, however for practical purposes operators may find it easier to work at octet aligned boundaries. In any case, the maximum length of the End-user IPv6 prefix is 96 minus the length of PSID, to allow for the encoding of the IPv4 address and PSID. The EA bits are composed of the IPv4 suffix and PSID as per MAP

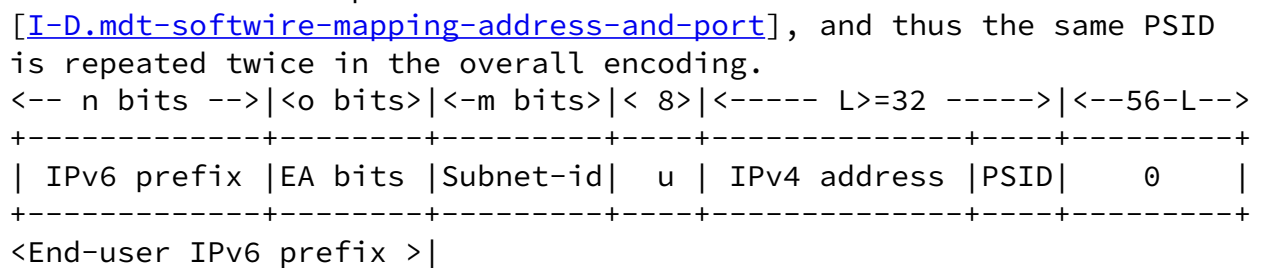


Figure 2: IPv4-translatable address for BMR and FMR

The address format used by the MAP-T Default Mapping Rule (DMR, IPv4 converted address used to represent IPv4 destinations outside of the MAP-T domain) is specific to MAP-T. An example is as shown in Figure 3. Note that the BR-prefix length is variable and can be both shorter or longer than 64 bits, up to 96 bits. In the respective cases the IPv4 address and the BR prefix are shifted and "bit spread" across the fixed u-octet boundary as per [[RFC6052](#)]. All trailing bits after the IPv4 address are set to 0x0.




Figure 3: Example of IPv4-converted address for DMR

In all cases the "u-octet" is taken to be 0x00.

7.2. Translating IPv4 Address and Port Number into IPv6 Address and Port Number at the BR

IPv6 Source Address and Source Port Number:

At the BR, the IPv6 source address (IPv4-converted address) MUST be derived from the IPv4 source Address as per DMR. The source Layer 4 TCP/UDP port number MUST be unchanged.

IPv6 Destination Address and Destination Port Number:

At the BR, the IPv6 destination address (IPv4-translatable address) MUST be derived from the IPv4 destination address and the destination port number as per FMR. The destination port number MUST be unchanged.

[7.3.](#) Translating IPv6 Address and Port Number into IPv4 Address and Port Number at the BR

IPv4 Source Address and Source Port Number:

At the BR, the IPv4 source address MUST be derived from the IPv6 source address (IPv4-translatable address) as per FMR. The source port number MUST be unchanged.

IPv4 Destination Address and Destination Port Number:

At the BR, the IPv4 destination address MUST be derived from the IPv6 destination address (IPv4-converted address) as per DMR. The destination port number MUST be unchanged.

[7.4.](#) Translating IPv4 Address and Port Number into IPv6 Address and Port Number at the CE

IPv6 Source Address and Source Port Number:

At the CE, the IPv6 source address (IPv4-translatable address) MUST be derived from the IPv4 source address as per BMR. The source port number MUST be unchanged.

IPv6 Destination Address and Destination Port Number:

At the CE, if Forwarding Mapping Rules (FMRs) are enabled, the IPv4 packet MUST be checked to see if the IPv4 destination address matches the FMR. If matching, the IPv6 destination address (IPv4-translatable address) MUST be derived from the IPv4 destination address and the destination port number per FMR. Otherwise, the IPv6 destination address (IPv4-translatable address) MUST be derived from the received IPv4 destination address per DMR. The destination port number MUST be unchanged.

[7.5.](#) Translating IPv6 Address and Port Number into IPv4 Address and Port Number at the CE

IPv4 Source Address and Source Port Number:

At the CE, the IPv4 source address MUST be derived from the IPv6 source address (IPv6-converted address) as per the DMR, or as per the FMR. The source port number MUST be unchanged.

IPv4 Destination Address and Destination Port Number: At the CE, the IPv4 destination address MUST be derived from the IPv6 destination address (IPv6-translatable address) as per BMR. The destination port number MUST be unchanged.

[7.6.](#) Translating ICMP/ICMPv6 Headers

MAP-T CEs and BRs MUST follow ICMP/ICMPv6 translation as per [\[RFC6145\]](#), with the following extension to cover the address sharing/port-range feature.

Unlike TCP and UDP, which each provide two port fields to represent both source and destination, the ICMP/ICMPv6 Query message header has only one ID field [\[RFC0792\]](#), [\[RFC4443\]](#). Thus, if the ICMP Query message is originated from an IPv4 host behind a MAP-T CE, the ICMP ID field SHOULD be used to exclusively identify that IPv4 host. This means that the MAP-T CE SHOULD rewrite the ID field to a port-set value obtained via the BMR during the IPv4 to IPv6 ICMPv6 translation operation. A BR can translate the resulting ICMPv6 packets back to ICMP preserving the ID field on its way to an IPv4 destination. In the return path, when MAP-T BR receives an ICMP packet containing an ID field which is bound for a shared address in the MAP-T domain, the MAP-T BR SHOULD use the ID value as a substitute for the destination port in determining the IPv6 destination address according to [Section 5.1](#). In all other cases, the MAP-T BR MUST derive the destination IPv6 address by simply mapping the destination IPv4 address without additional port info.

[7.7.](#) Path MTU Discovery and Fragmentation

Due to the different sizes of the IPv4 and IPv6 header, which are 20+ octets and 40 octets respectively, handling the maximum packet size is relevant for the operation of any system connecting the two address families. There are three mechanisms to handle this issue: path MTU discovery (PMTUD), fragmentation, and transport-layer negotiation such as the TCP Maximum Segment Size (MSS) option [[RFC0897](#)]. MAP-T uses all three mechanisms to deal with different cases.

Following [[RFC6145](#)], when an IPv4 node performs path MTU discovery (by setting the Don't Fragment (DF) bit in the header), path MTU discovery can operate end-to-end across the MAP-T BR and CE translators. In this case, either IPv4 or IPv6 routers (including the translators) can send back ICMP Packet Too Big messages to the sender. When IPv6 routers send these as ICMPv6 errors, these packets will pass through the translator and will result in an appropriate ICMP error sent to the IPv4 sender. When the IPv4 sender does not set the DF bit, the translator MUST ensure that the packet does not exceed the path MTU on the IPv6 side. This is done, if necessary, by fragmenting the IPv4 packet and including with Fragment Headers to fit in the minimum MTU 1280-byte IPv6 packets. When the IPv4 sender does not set the DF bit, the translator SHOULD include an IPv6 Fragment Header to indicate that the sender allows fragmentation. The rules defined in [[RFC6145](#)] ensure that when packets are fragmented, either by the sender or by IPv4 routers, the low-order 16 bits of the fragment identification are carried end-to-end, ensuring that packets are correctly reassembled. The above mechanism ensures that the Don't Fragment (DF) bit in the IPv4 header can be carried end-to-end via double stateless translation in most of the cases. For example, the IPv4 packets with DF=1 will be translated to IPv6 packets without fragmentation header and will be translated back to IPv4 packets with DF=1. The IPv4 packets with DF=0 will be translated to IPv6 packets with fragmentation header (keeping the ID value) and will be translated back to IPv4 packets with DF=0. An open corner case left up for specific handling by implementations [[RFC6145](#)] is for when IPv4 packets with DF=1 and MF=1 are received by a translator. MAP-T devices SHOULD translate such IPv4 packets into IPv6 with a fragmentation header present. Experimental evidence [[operational-exp](#)] and [IMC-07}, indicate that only 27,474 packets observed with DF=1/MF=1 among 10 billion samples. This indicates that IPv4 packets with DF=1 and MF=1 are rare in production networks

(10e-5) and that their handling by this rule causes no negative effects in practice.

8. MAP-T Packet Forwarding considerations

8.1. Mesh Model

MAP-T allows the use of the mesh model in order for all CEs to communicate with each other directly (i.e. bypassing the BR). When a CE receives an IPv4 packet from its LAN side, the CE looks up a mapping rule corresponding to an IPv4 destination address in the received IPv4 packet. If the corresponding mapping rule is found, CE can communicate to another CE directly based on the mapping rule defined as Forwarding mapping rule (FMR) in MAP. If the corresponding mapping rule is not found, CE must forward the packet

Bao, et al.

Expires September 10, 2012

[Page 14]

Internet-Draft

Map Translation

March 2012

to a given BR.

8.2. Hub & Spoke model

In order to allow the mesh topology so that all CEs can communicate each others directly, all CE should know all mapping rules applied to a given MAP-T domain or MAP-T domains. However, if a CE knows only a subset of the mapping rules applied to a given MAP-T domain, a CE can not communicate directly with some of the CEs in that domain due to the lack of mapping rules. In this case, an IPv4 packet toward to these CEs must be forwarded to a given BR. In order to achieve the hub & spoke mode fully, the Forwarding mapping rule (FMR) defined in MAP need to be disabled (not defined).

8.3. Communication with IPv6 servers in the MAP-T domain

MAP-T allows communication between both IPv4-only and any IPv6 enabled end hosts, with native IPv6-only servers which are using IPv4-mapped IPv6 address based on DMR in the MAP-T domain. In this mode, the IPv6-only servers SHOULD have both A and AAAA records in the authorities DNS server [[RFC6219](#)]. DNS64 [[RFC6147](#)] become required only when IPv6 servers in the MAP-T domain are expected themselves to initiate communication to external IPv4-only hosts.

9. NAT44 considerations

The NAT44 implemented in MAP-T CE SHOULD conform with the behavior and best current practice documented in [[RFC4787](#)], [[RFC5508](#)] and [[RFC5382](#)]. In MAP-T address sharing mode (determined by the MAP-T configuration parameters) the operation of the NAT44 must be restricted to the available port numbers derived via BMR.

10. Security Considerations

Spoofing attacks: With consistency checks between IPv4 and IPv6 sources that are performed on IPv4/IPv6 packets received by BR's and CE's ([Section 6](#)), MAP-T does not introduce any opportunity for spoofing attack that would not pre-exist in IPv6.

Denial-of-service attacks: In MAP-T domains where IPv4 addresses are shared, the fact that IPv4 datagram reassembly may be necessary introduces an opportunity for DOS attacks. This is inherent to address sharing, and is common with other address sharing approaches such as DS-Lite and NAT64/ DNS64. The best protection against such attacks is to accelerate IPv6 enablement in both clients and servers so that, where MAP-T is supported, it is less

and less used.

Routing-loop attacks: This attack may exist in some automatic-tunneling scenarios are documented in [[I-D.ietf-v6ops-tunnel-loops](#)]. They cannot exist with MAP-T because each BRs checks that the IPv6 source address of a received IPv6 packet is a CE address based on Forwarding Mapping Rule defined in MAP [[I-D.mdt-software-mapping-address-and-port](#)].

Attacks facilitated by restricted port set: From hosts that are not subject to ingress filtering of [[RFC2827](#)], some attacks are possible by intervening with faked packets during ongoing transport connections ([[RFC4953](#)], [[RFC5961](#)], [[RFC6056](#)]). The attacks depend on guessing which ports are currently used by target hosts, and using an unrestricted port set is preferable, i.e. using native IPv6 connections that are not subject to MAP port range restrictions. To minimize this type of attacks when using a restricted port set, the MAP CE's NAT44 filtering behavior

SHOULD be "Address-Dependent Filtering". Furthermore, the MAP CEs SHOULD use a DNS transport proxy function to handle DNS traffic, and source such traffic from IPv6 interfaces not assigned to MAP-T. Practicalities of these methods are discussed in [Section 5.9](#) of [[I-D.dec-stateless-4v6](#)].

[11.](#) IANA Consideration

This document has no IANA actions.

[12.](#) Acknowledgements

The authors would like to thank Maoke Chen, Leaf Yeh and Senthil Sivakumar for their review and comments.

[13.](#) References

[13.1.](#) Normative References

- [I-D.mdt-softwire-mapping-address-and-port]
Troan, O., "Mapping of Address and Port (MAP)",
[draft-mdt-softwire-mapping-address-and-port-02](#) (work in progress), November 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

Bao, et al. Expires September 10, 2012 [Page 16]

Internet-Draft Map Translation March 2012

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

[13.2.](#) Informative References

- [I-D.dec-stateless-4v6]
Dec, W., Asati, R., and H. Deng, "Stateless 4Via6 Address Sharing", [draft-dec-stateless-4v6-04](#) (work in progress), October 2011.
- [I-D.ietf-v6ops-tunnel-loops]

Nakibly, G. and F. Templin, "Routing Loop Attack using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", [draft-ietf-v6ops-tunnel-loops-07](#) (work in progress), May 2011.

[I-D.mdt-software-map-dhcp-option]

Mrugalski, T., Boucadair, M., and O. Troan, "DHCPv6 Options for Mapping of Address and Port", [draft-mdt-software-map-dhcp-option-00](#) (work in progress), October 2011.

[I-D.murakami-software-4v6-translation]

Murakami, T., Chen, G., Deng, H., Dec, W., and S. Matsushima, "4via6 Stateless Translation", [draft-murakami-software-4v6-translation-00](#) (work in progress), July 2011.

[I-D.operators-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-operators-software-stateless-4v6-motivation-02](#) (work in progress), June 2011.

[I-D.xli-behave-divi]

Shang, W., Li, X., Zhai, Y., and C. Bao, "dIVI: Dual-Stateless IPv4/IPv6 Translation", [draft-xli-behave-divi-04](#) (work in progress), October 2011.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), September 1981.

[RFC0897] Postel, J., "Domain name system implementation schedule", [RFC 897](#), February 1984.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source

Address Spoofing", [BCP 38](#), [RFC 2827](#), May 2000.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol

- Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", [BCP 127](#), [RFC 4787](#), January 2007.
- [RFC4953] Touch, J., "Defending TCP Against Spoofing Attacks", [RFC 4953](#), July 2007.
- [RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", [BCP 142](#), [RFC 5382](#), October 2008.
- [RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", [BCP 148](#), [RFC 5508](#), April 2009.
- [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks", [RFC 5961](#), August 2010.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), October 2010.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", [BCP 156](#), [RFC 6056](#), January 2011.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", [RFC 6144](#), 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.
- [RFC6219] Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The China Education and Research Network (CERNET) IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition", [RFC 6219](#), May 2011.
- [operational-exp] John, Wolfgang; Tafvelin, Sven: Analysis of Internet Backbone Traffic and Header Anomalies Observed. IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, pp. 111-116. ISBN/ISSN: 978-1-59593-908-1
<http://conferences.sigcomm.org/imc/2007/papers/imc91.pdf>

[Appendix A](#). Example of MAP-T translation

The following is a MAP-T example derived from the general MAP example in MAP [[I-D.mdt-software-mapping-address-and-port](#)].

Example 1.

Given the MAP domain information and an IPv6 address of an endpoint:

IPv6 prefix assigned to the end user: 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: $256 (16 - (32 - 24) = 8 \cdot 2^8 = 256)$

PSID offset: 4

A MAP node (CE or BR) can via the BMR 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: $o - p = 16 (56 - 40) - 8 = 8$

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

The BMR information allows a MAP-T CE also to determine (complete) its IPv6 address within the indicated IPv6 prefix.

Bao, et al.

Expires September 10, 2012

[Page 19]

Internet-Draft

Map Translation

March 2012

IPv6 address of MAP-T CE: 2001:db8:0012:3400:00c0:0002:1200:3400

Example 2.

Another example can be made of a hypothetical MAP-T BR, configured with the following FMR when receiving a packet with the following characteristics:

IPv4 source address: 1.2.3.4 (0x01020304)

IPv4 source port: 80

IPv4 destination address: 192.0.2.18 (0xc0000212)

IPv4 destination port: 9030

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

MAP-T BR Prefix 2001:db8:ffff::/64

The above information allows the BR to derive as follows the mapped destination IPv6 address for the corresponding MAP-T CE, and also the mapped source IPv6 address for the IPv4 source.

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

PSID length: 8

PSID: 0x34 (9030 (0x2346))

The resulting IPv6 packet will have the following key fields

IPv6 source address 2001:db8:ffff:0:0001:0203:0400::

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

IPv6 source Port: 80

IPv6 destination Port: 9030

Example 3:

An IPv4 host behind the MAP-T CE (addressed as per the previous examples) corresponding with IPv4 host 1.2.3.4 will have its packets converted into IPv6 using the DMR configured on the MAP-T CE as follows:

Bao, et al.

Expires September 10, 2012

[Page 20]

Internet-Draft

Map Translation

March 2012

Default Mapping Rule used by MAP-T CE: {2001:db8:ffff::/64 (Rule IPv6 prefix), 0.0.0.0/0 (Rule IPv4 prefix), null (BR IPv4 address)}

IPv4 source address (post NAT44 if present) 192.0.2.18

IPv4 destination address: 1.2.3.4

IPv4 source port (post NAT44 if present): 9030

IPv4 destination port: 80

IPv6 source address of MAP-T CE: 2001:db8:0012:3400:00c0:0002:1200:3400

IPv6 destination address: 2001:db8:ffff:0:0001:0203:0400::

Authors' Addresses

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

Email: congxiao@cernet.edu.cn

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

Email: xing@cernet.edu.cn

Yu Zhai
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Email: jacky.zhai@gmail.com

Bao, et al.

Expires September 10, 2012

[Page 21]

Internet-Draft

Map Translation

March 2012

Tetsuya Murakami (editor)
IP Infusion
1188 East Arques Avenue
Sunnyvale
USA

Email: tetsuya@ipinfusion.com

Wojciech Dec (editor)
Cisco Systems
Haarlerbergpark Haarlerbergweg 13-19
Amsterdam, NOORD-HOLLAND 1101 CH
Netherlands

Phone:

Email: wdec@cisco.com

