

**Multicast Support for MAP-E**  
**draft-sarikaya-software-4rdmulticast-06.txt**

**Abstract**

This memo specifies MAP-E (together with MAP-T and 4rd)'s multicast component so that IPv4 hosts can receive multicast data from IPv4 servers over an IPv6 network. In the encapsulation solution for encapsulation variant of Mapping of Address and Port (MAP), MAP-E, IGMP Proxy at the MAP-E Customer Edge router uses IPv4-in-IPv6 tunnel established by MAP-E to exchange IGMP messages to establish multicast state at MAP-E Border Relay so that MAP-E Border Relay can tunnel IPv4 multicast data to IPv4 hosts connected to MAP-E Customer Edge device. In the Translation Multicast solution for the translation variant of MAP, MAP-T and 4rd, IGMP messages are translated into MLD messages at the CE router which is IGMP/MLD Proxy and sent to the network in IPv6. MAP-T/4rd Border Relay does the reverse translation and joins IPv4 multicast group for MAP-T/4rd hosts. Border Relay as multicast router receives IPv4 multicast data and translates the packet into IPv6 multicast data and sends downstream on the multicast tree. Member CEs receive multicast data, translate it back to IPv4 and transmit to the hosts.

**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 23, 2013.

**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](#) 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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Requirements</a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Architecture</a>	<a href="#">5</a>
<a href="#">4.1.</a>	<a href="#">Encapsulation Multicast Architecture</a>	<a href="#">5</a>
<a href="#">4.2.</a>	<a href="#">MAP-T and 4rd Translation Architecture</a>	<a href="#">6</a>
<a href="#">5.</a>	<a href="#">Encapsulation Multicast Operation</a>	<a href="#">7</a>
<a href="#">5.1.</a>	<a href="#">Encapsulation Interface Considerations</a>	<a href="#">9</a>
<a href="#">5.2.</a>	<a href="#">Avalanche Problem Considerations</a>	<a href="#">10</a>
<a href="#">6.</a>	<a href="#">MAP-T and 4rd Translation Multicast Operation</a>	<a href="#">10</a>
<a href="#">6.1.</a>	<a href="#">Address Translation</a>	<a href="#">11</a>
<a href="#">6.2.</a>	<a href="#">Protocol Translation</a>	<a href="#">12</a>
6.3.	<a href="#">Supporting IPv6 Multicast in MAP-T and 4rd Translation Multicast</a>	<a href="#">13</a>
6.4.	<a href="#">Learning Multicast Prefixes for IPv4-embedded IPv6 Multicast Addresses</a>	<a href="#">14</a>
<a href="#">7.</a>	<a href="#">Security Considerations</a>	<a href="#">15</a>
<a href="#">8.</a>	<a href="#">IANA Considerations</a>	<a href="#">15</a>
<a href="#">9.</a>	<a href="#">Acknowledgements</a>	<a href="#">15</a>
<a href="#">10.</a>	<a href="#">References</a>	<a href="#">15</a>
<a href="#">10.1.</a>	<a href="#">Normative References</a>	<a href="#">15</a>
<a href="#">10.2.</a>	<a href="#">Informative references</a>	<a href="#">17</a>
<a href="#">Appendix A.</a>	<a href="#">Group Membership Message Translation Details</a>	<a href="#">18</a>
	<a href="#">Author's Address</a>	<a href="#">20</a>



## 1. Introduction

With IPv4 address depletion on the horizon, many techniques are being standardized for IPv6 migration including Mapping of Address and Port (MAP) - Encapsulation, - Translation and 4rd [[I-D.ietf-softwire-map](#)], [[I-D.ietf-softwire-map-t](#)], [[I-D.ietf-softwire-4rd](#)]. MAP/4rd enables IPv4 hosts to communicate with external hosts using IPv6 only ISP network. MAP/4rd Customer Edge (CE) device's LAN side is dual stack and WAN side is IPv6 only. CE tunnels/translates IPv4 packets received from the LAN side to 4rd Border Relays (BR). BRs have anycast IPv6 addresses and receive encapsulated/translated packets from CEs over a virtual interface. MAP/4rd operation is stateless. Packets are received/ sent independent of each other and no state needs to be maintained except for NAT44 operation on IPv4 packets received from the user.

It should be noted that there is no depletion problem for IPv4 address space allocated for any source multicast and source specific multicast [[RFC3171](#)]. This document is not motivated by the depletion of IPv4 multicast addresses.

MAP-E, MAP-T and 4rd are unicast only. They do not support multicast. In this document we specify how multicast from home IPv4 users can be supported in MAP-E (as well as MAP-T and 4rd).

In case of MAP-E we integrate the multicast solution into the MAP-E tunnel resulting in a multicast tunneling protocol. Multicast tunneling protocol has the advantage of not requiring multicast enabled IPv6 network between CE routers and MAP-E BRs.

When MAP-E CE router receives an IGMP join message to an Any-Source Multicast (ASM) [[RFC1112](#)] or Source-Specific Multicast (SSM) group [[RFC4607](#)], it sends an aggregated IGMP membership report message in the IPv4-in-IPv6 tunnel to the border relay. MAP-E BR joins the source in the multicast infrastructure and sends multicast data downstream to all member CEs in the IPv4-in-IPv6 tunnel. When a CE has no membership state, i.e. after all member hosts leave the group(s), its state with the BR expires and the CE can send the next join message in anycast. IPv4 multicast data received at the BR is tunneled to the member CE in IPv6 and CE decapsulates the packet and sends IPv4 multicast data packet to the member hosts.

In case IPv6 network is multicast enabled, MAP-T/4rd can provide multicast service to the hosts using MAP-T/4rd Multicast Translation based solution. A Multicast Translator can be used that receives IPv4 multicast group management messages in IGMP and generates corresponding IPv6 group management messages in MLD and sends them to IPv6 network towards MAP-T/4rd Border Relay. We use



[I-D.ietf-softwire-map-t] or [I-D.ietf-softwire-4rd] for sending IPv4 multicast data in IPv6 to the CE routers. At MAP-T/4rd CE router another translator is needed to translate IPv6 multicast data into IPv4 multicast data.

It should be noted that if IPv6 network is multicast enabled the translation multicast solution presented in [Section 6](#) can also be used for MAP-E.

In this document we address MAP-E (and MAP-T/4rd) multicast problem and propose two architectures: Multicast Tunneling and Multicast Translation based solutions. [Section 2](#) is on terminology, [Section 3](#) is on requirements, [Section 4](#) is on architecture, [Section 5](#) is on multicast tunneling protocol [Section 6](#) is on multicast translation protocol, and [Section 7](#) states security considerations.

## 2. Terminology

This document uses the terminology defined in [I-D.ietf-softwire-map], [I-D.ietf-softwire-map-t], [I-D.ietf-softwire-4rd], [RFC3810] and [RFC3376].

## 3. Requirements

This section states requirements on MAP-E, MAP-T and 4rd multicast support protocol.

IPv4 hosts connected to MAP-E, MAP-T and 4rd CE router MUST be able to join multicast groups in IPv4 and receive multicast data.

Any source multicast (ASM) SHOULD be supported and source specific multicast (SSM) MUST be supported.

In case of encapsulation solution, MAP-E CE MUST support IGMP Proxy as defined in [RFC4605]. MAP-E BR MUST support IGMP querier downstream and MAP-E BR may support PIM protocol or IGMP router upstream.

In case of translation solution, MAP-T and 4rd CE MUST support IGMP to MLD translation. MAP-T and 4rd CE MUST be MLD Proxy as defined in [RFC4605]. MAP-T and 4rd BR MUST support MLD Querier. MAP-T and 4rd BR MUST support join/leave operations in IPv4 multicast upstream.

## **4. Architecture**

In MAP-E, MAP-T and 4rd, there are hosts (possibly IPv4/ IPv6 dual stack) served by MAP-E, MAP-T and 4rd Customer Edge device. CE is dual stack facing the hosts and IPv6 only facing the network or WAN side. MAP-E, MAP-T and 4rd CE may be local IPv4 Network Address and Port Translation (NAPT) box [[RFC3022](#)] by assigning private IPv4 addresses to the hosts. MAP-E, MAP-T and 4rd CEs in the same domain may use shared public IPv4 addresses on their WAN side and if they do they should avoid ports outside of the allocated port set for NAPT operation. At the boundary of the network there is MAP-E, MAP-T and 4rd Border Relay. BR receives IPv4 packets tunneled in IPv6 from CE and decapsulates them and sends them out to IPv4 network.

Unicast MAP-E, MAP-T and 4rd are stateless except for the local NAPT at the CE. Each IPv4 packet sent by CE treated separately and different packets from the same CE may go to different BRs or CEs. CE encapsulates IPv4 packet in IPv6 with destination address set to BR address (usually anycast IPv6 address). BR receives the encapsulated packet and decapsulates and send it to IPv4 network. CEs are configured with Rule IPv4 Prefixes, Rule IPv6 Prefixes and with an BR IPv6 anycast address. BR receives IPv4 packets addressed to this ISP and from the destination address it extracts the destination host's IPv4 address and uses this address as destination address and encapsulates the IPv4 packet in IPv6 and sends it to IPv6-only network.

### **4.1. Encapsulation Multicast Architecture**

Encapsulation variant of MAP called MAP-E network lends itself easily to the Multicast Tunneling architecture. Dual stack hosts are connected to the Customer Edge router and it is multicast enabled. It is assumed that IPv6 only network is the unicast only network and that IPv6 multicast is not enabled or IPv6 multicast is partially enabled. At the boundary of the network MAP-E Border Relay is connected to the native multicast backbone infrastructure.

We place IGMP Proxy at the CE router. CE router serves all the connected hosts. For multicast traffic, CE Router uses MAP-E tunneling interface with MAP-E BR to send/receive IGMP messages using IPv4-in-IPv6 tunnel [[RFC2473](#)].

MAP-E BR is IGMP Router towards the CEs and it could be IGMP Router or PIM router upstream. A given relay and all CEs connected to it can be considered to be on a separate logical link. On this link, gateways and relay communicate using IPv4-in-IPv6 tunneling to transmit and receive multicast control messages for membership management and multicast data from the relay to the gateways.





All the elements of MAP-E multicast support system with tunneling are shown in Figure 1.

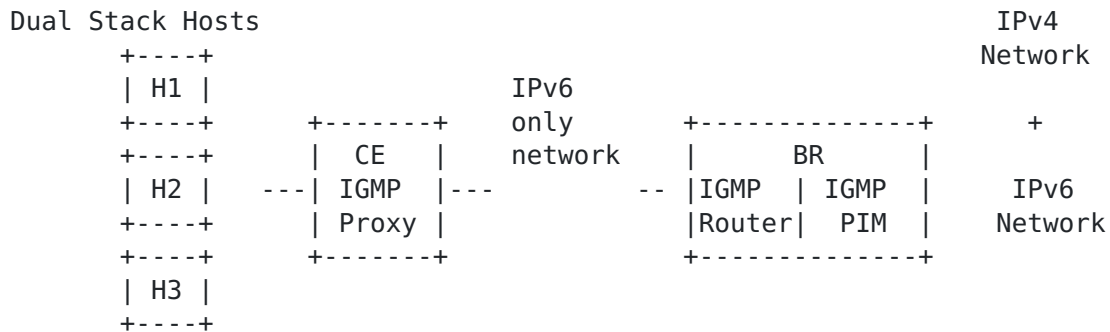


Figure 1: Architecture of MAP-E Multicast Tunneling

#### 4.2. MAP-T and 4rd Translation Architecture

In case IPv6 only network is multicast enabled, translation multicast architecture can be used. CE implements IGMP Proxy function [RFC4605] towards the LAN side and MLD Proxy on its WAN interface. IPv4 hosts send their join requests (IGMP Membership Report messages) to CE. CE as a MLD proxy sends aggregated MLD Report messages upstream towards BR. CE replies MLD membership query messages with MLD membership report messages based on IGMP membership state in the IGMP/MLD proxy.

BR is MLD querier on its WAN side. On its interface to IPv4 network BR may either have IGMP client or PIM. PIM being able to support both IPv4 and IPv6 multicast should be preferred. BR receives MLD join requests, extracts IPv4 multicast group address and then joins the group upstream, possibly by issuing a PIM join message.

IPv4 multicast data received by the BR as a leaf node in IPv4 multicast distribution tree is translated into IPv6 multicast data by the translator using [I-D.ietf-softwire-map-t], [I-D.ietf-softwire-4rd] and then sent downstream to the IPv6 part of the multicast tree to all downstream routers that are members. IPv6 data packet eventually gets to the CE. At the CE, a reverse [I-D.ietf-softwire-map-t], [I-D.ietf-softwire-4rd] operation takes place by the translator and then IPv4 multicast data packet is sent to the member hosts on the LAN interface. [I-D.ietf-softwire-map-t], [I-D.ietf-softwire-4rd] are modified to handle multicast addresses.

In order to support SSM, IGMPv3 MUST be supported by the host, CE and



BR. For ASM, BR MUST be the Rendezvous Point (RP).

MAP-T and 4rd Translation Multicast solution uses the multicast 46 translator in not one but two places in the architecture: at the CE router and at the Border Relay. IPv4 multicast data received at 4rd BR goes through a [[I-D.ietf-softwire-4rd](#)] header-mapping into IPv6 multicast data at the BR and another [[I-D.ietf-softwire-4rd](#)] header-mapping back to IPv4 multicast data at the CE router. Encapsulation variant of [[I-D.ietf-softwire-4rd](#)] is not used. In case of MAP-T, IPv4 data packet is translated using v4 to v6 header translation using multicast addresses instead of the mapping algorithm used in [[I-D.ietf-softwire-map-t](#)].

All the elements of MAP-T and 4rd translation-based multicast support system are shown in Figure 2.




Figure 2: Architecture of MAP-T and 4rd Translation Multicast

## 5. Encapsulation Multicast Operation

When a host (H1, H2 or H3 in Figure 1) wants to join an IPv4 multicast group G or (S,G), it sends an IGMP report (IGMPv3 report for a source-specific group) to CE router.

CE encapsulates IGMP report messages in IPv6 and sends it over the tunnel to BR in anycast. CE router uses BR's anycast address this CE router is configured with. After CE receives unicast address of BR, it sends all subsequent IGMP messages for G or (S,G) in unicast.

BR (topologically closest to this CE router) receives the message, decapsulates it and then lets IGMP router to process it. On the upstream, an IGMP Join message is sent to subscribe group G or (S,G) or a PIMv4 Join message is sent if PIM is supported. BR establishes



membership state for group G or (S,G). BR sends all related IGMP messages to this CE in unicast using IPv4-in-IPv6 tunneling.

CE now has BR's unicast address which it uses to send all IGMP packets for group G for any source multicast or (S,G) for source specific multicast. If CE receives multiple join messages for the same group G, CE sends an aggregated join message to BR.

If CE receives another join message for a different group G', (S',G') CE encapsulates it and sends it in anycast to the BR. This enables the use of multiple BRs that may be deployed as anchor points and makes downstream multicast data delivery more efficient.

A CE is required to assist in IGMP signaling and data forwarding between the hosts that it serves and the corresponding BRs that are handling the multicast group G or (S,G). CE must have IGMP Proxy for each upstream tunnel interface that has been established with the BR. The CE decides on the mapping of downstream links to a proxy instance connected to an upstream link to a BR based on the unicast source IPv6 address in the packets received from BR. Because of this BRs MUST use the unicast source IPv6 address in packets sent to CEs. Encapsulation at the CE is according to [[RFC2473](#)] with an IPv4 payload carrying IGMP messages.

On the reception of IGMP reports from the hosts, the CE must identify the corresponding proxy instance from the incoming interface and perform regular IGMP proxy operations of inserting, updating or removing multicast forwarding state on the incoming interface and will merge state updates into the IGMP proxy membership database. It will then send an aggregated Report via the upstream tunnel to the BR when the membership database changes.

On the reception of IGMP queries, the CE proxy instance will answer the Queries on behalf of all active downstream receivers maintained in its membership database. Queries sent by the BR do not force the CE to trigger corresponding messages immediately towards hosts.

BR acts as the default multicast querier for the corresponding CE. It implements the function of the designated multicast router or a further IGMP proxy. After BR receives IGMP Join message it adds the tunnel to the CE in its outgoing interface list for the group (G) or the source, group (S,G) that the host wants to join. BR establishes group-/source-specific multicast forwarding states at its corresponding downstream tunnel interfaces. Afterwards, BR maintains/removes these group-/source-specific multicast forwarding states. BR treats its tunnel interfaces as multicast-enabled downstream links, serving zero to many listening nodes. BR will send a join message upstream towards the source of the multicast group to



build a multicast tree in the native multicast infrastructure and becomes a leaf node in the multicast tree.

BR will send any group management messages (IGMP Report or Query messages) downstream to specific CEs on the tunnel interface by encapsulating these IGMP messages in IPv6 using [\[RFC2473\]](#).

As for multicast data, the data packets from the source received at the BR will be replicated to all interfaces in its outgoing interface list as well as the tunnel outgoing interface for all member CEs. BR sends multicast data in IPv4-in-IPv6 tunnel to the CE with the data packet encapsulated. Encapsulation is according to [\[RFC2473\]](#) with an IPv4 payload.

CE receives Multicast Data message over the tunnel interface associated with the tunnel to BR. After decapsulation, multicast traffic arriving at the CE on an upstream interface will be forwarded according to the group-specific or source-specific forwarding states as acquired for each downstream interface within the IGMP proxy instance.

### **[5.1](#). Encapsulation Interface Considerations**

Legacy IPv4 in IPv6 tunneling is performed as in [\[RFC2473\]](#). Packets upstream from CE carry only IGMP signaling messages and they are not expected to be subject to fragmentation. However packets downstream, i.e. multicast data to CE may be subject to fragmentation.

Source and destination addresses of IGMP messages in IPv4-in-IPv6 software from CE is as follows:

Source address of IPv6 header is CE IPv6 address, e.g. 2001:db8:0:1::1, destination address is BR anycast address, possibly shared of the MAP domain.

Source address of IGMP messages is CE's IPv4 interface address, e.g. 192.0.0.2, destination address is the all-systems multicast address of 224.0.0.1 for IGMP Query, all IGMPv3-capable multicast routers of 224.0.0.22 for IGMPv3 Report, the multicast group specified in the Group Address field of the Report for IGMPv1 or IGMPv2 Report.

Source and destination addresses of IGMP messages in IPv4-in-IPv6 software from BR is as follows:

Source address of IPv6 header is BR's unicast IPv6 address, e.g. 2001:db8:0:2::1, destination address is CE IPv6 address, e.g. 2001:db8:0:1::1.





Source address of IGMP messages is CE's IPv4 interface address, e.g. 192.0.2.1, destination address is the all-systems multicast address of 224.0.0.1 for IGMP Query, all IGMPv3-capable multicast routers of 224.0.0.22 for IGMPv3 Report, the multicast group specified in the Group Address field of the Report for IGMPv1 or IGMPv2 Report.

Source and destination addresses of multicast data messages in IPv4-in-IPv6 software is as follows:

Source address of IPv6 header is BR IPv6 unicast address, e.g. 2001:db8:0:2::1, destination address is CE IPv6 address, e.g. 2001:db8:0:1::1.

Source address of IPv4 multicast data is unicast IPv4 address of the multicast source, e.g. the content provider, destination address is IPv4 multicast group address.

BR decapsulates datagrams carrying IGMP messages from CE's and then IGMP/PIM router processing takes over. Network Address Translation (NAT) is not applied on IGMP messages.

## **5.2. Avalanche Problem Considerations**

In [Section 5](#) BR replicates the data packets from the source received to all outgoing interfaces for all member CEs. This replication (often called avalanche problem) can be very costly if there are very large number of downstream member CEs such as in IPTV application. Note that the avalanche problem is faced by all multicast solutions that use tunneling to bypass non-multicast enabled access network.

In multicast MAP-E, one approach that can be used is to deploy MAP-E BRs close to the user. BRs colocated at the access network gateway such as at the Border Network Gateway (BNG) could reduce the packet duplication bottleneck considerably.

In multicast MAP-E, another approach is to exploit multiple BRs that can be deployed in the network. MAP-E CE can use BR anycast address when sending an encapsulated upstream IGMP join request and then use the unicast source address of this BR in subsequent IGMP messages.

## **6. MAP-T and 4rd Translation Multicast Operation**

In this section we specify how the host can subscribe and receive IPv4 multicast data from IPv4 content providers based on the architecture defined in Figure 2 in two parts: address translation and protocol translation. Translation details are given in [Appendix A](#).



## 6.1. Address Translation

IPv4-only host, H1 will join IPv4 multicast group by sending IGMP Membership Report message upstream towards the IGMP Proxy in Figure 2. MLD Proxy first creates a synthesized IPv6 address of IPv4 multicast group address using IPv4-embedded IPv6 multicast address format [[I-D.ietf-mboned-64-multicast-address-format](#)]. ASM\_MPREFIX64 for any source multicast groups and SSM\_MPREFIX64 for source specific multicast groups are used. Both are /96 prefixes.

SSM\_MPREFIX64 is set to ff3x:0:8000::/96, with 'x' set to any valid scope. ASM\_MPREFIX64 values are formed as shown in Figure 3. M bit MUST BE set to 1. "flgs" and "scop" fields are defined in [[RFC3956](#)]. The usage of the "rsv" bits is the same as defined in [[RFC3306](#)]. "sub-group-id" field MUST follow the recommendations specified in [[RFC3306](#)] if unicast-based prefix is used or the recommendations specified in [[RFC3956](#)] if embedded-RP is used. The default value is all zeros.



Figure 3: ASM\_MPREFIX64 Formation

Each translator in the upstream BR is assigned a unique ASM\_MPREFIX64 prefix. CE (MLD Proxy in CE) can learn this value by means out of scope with this document. With this, CE can easily create an IPv6 multicast address from the IPv4 group address a.b.c.d that the host wants to join.

Source-Specific Multicast (SSM) can also be supported similar to the Any Source Multicast (ASM) described above. In case of SSM, IPv4 multicast addresses use 232.0.0.0/8 prefix. IPv6 SSM\_MPREFIX64 is set to FF3x:0:8000::/96.

Since SSM translation requires a unique address for each IPv4 multicast source, an IPv6 unicast prefix must be configured to the translator in the upstream BR to represent IPv4 sources. This prefix is prepended to IPv4 source addresses in translated packets.

The join message from the host for the group ASM\_MPREFIX64:a.b.c.d or SSM\_MPREFIX64:a.b.c.d or an aggregate join message will be received by MLD querier at the BR. BR as multicast anchor checks the group address and recognizes ASM\_MPREFIX64 or SSM\_MPREFIX64 prefix. It next checks the last 32 bits is an IPv4 multicast address in range 224/8 - 239/8. If all checks succeed, IGMPv4 Client joins a.b.c.d



using IGMP on its IPv4 interface.

Joining IPv4 groups can also be done using PIM since PIM supports both IPv4 and IPv6. The advantage of using PIM is that there is no need to enable IGMP support in neighboring IPv4 routers. The advantage of using IGMP is that IGMP is a simpler protocol and it is supported by a wider range of routers. The use of PIM or IGMP is left as an implementation choice.

## 6.2. Protocol Translation

The hosts will send their subscription requests for IPv4 multicast groups upstream to the default router, i.e. Customer Edge device. After subscribing the group, the host can receive multicast data from the CE. The host implements IGMP protocol's host part.

Customer Edge device is IGMP Proxy facing the LAN interface. After receiving the first IGMP Report message requesting subscription to an IPv4 multicast group, a.b.c.d, MLD Proxy in the CE's WAN interface synthesizes an IPv6 multicast group address corresponding to a.b.c.d and sends an MLD Report message upstream to join the group.

When CE is a NAT or NAPT box assigning private IPv4 addresses to the hosts, IP Multicast requirements for a Network Address Translator (NAT) and a Network Address Port Translator (NAPT) stated in [\[RFC5135\]](#) apply to IGMP messages and IPv4 multicast data packets.

When MAP-T or 4rd BR receives IPv4 multicast data for an IPv4 group a.b.c.d it [\[I-D.ietf-softwire-4rd\]](#) translates/encapsulates IPv4 packet into IPv6 multicast packet and sends it to IPv6 synthesized address corresponding to a.b.c.d using ASM\_MPREFIX64 or SSM\_MPREFIX64. The header mapping described in [\[I-D.ietf-softwire-4rd\]](#) [Section 4.2](#) (using Table 1) is used except for mapping the source and destination addresses. In this document we use the multicast address translation described in [Section 6.1](#) and propose it as a complementary enhancement to the translation algorithm in [\[I-D.ietf-softwire-4rd\]](#).

The IP/ICMP translation translates IPv4 packets into IPv6 using minimum MTU size of 1280 bytes (Section 4.3 in [\[I-D.ietf-softwire-4rd\]](#)) but this can be changed for multicast. Path MTU discovery for multicast is possible in IPv6 so 4rd BR can perform path MTU discovery for each ASM group and use these values instead of 1280. For SSM, a different MTU value MUST be kept for each SSM channel. Because of this 8 bytes added by IPv6 fragment header in each data packet can be tolerated.

Since multicast address translation does not preserve checksum



neutrality, [[I-D.ietf-softwire-4rd](#)] translator/encapsulator at 4rd BR must however modify the UDP checksum to replace the IPv4 addresses with the IPv6 source and destination addresses in the pseudo-header which consists of source address, destination address, protocol and UDP length fields before calculating the new checksum.

IPv6 multicast data must be translated back to IPv4 at the 4rd CE (e.g. using Table 2 in Section 4.3 of [[I-D.ietf-softwire-4rd](#)]). Such a task is much simpler than the translation at 4rd BR because IPv6 header is much simpler than IPv4 header and IPv4 link on the LAN side of 4rd CE is a local link. The packet is sent on the local link to IPv4 group address a.b.c.d for IPv6 group address of ASM\_MPREFIX64:a.b.c.d or SSM\_MPREFIX64:a.b.c.d.

In case an IPv4 multicast source sends multicast data with the don't fragment (DF) flag set to 1, [[I-D.ietf-softwire-4rd](#)] header mapping sets the D bit in IPv6 fragment header before sending the packet downstream as in Fig. 3 in Section 4.3 of [[I-D.ietf-softwire-4rd](#)]. This feature of [[I-D.ietf-softwire-4rd](#)] preserves the semantics of DF flag at the BR and CE.

Because MAP-T/4rd is stateless, Multicast MAP-T/4rd should stay faithful to this as much as possible. Border Relay acts as the default multicast querier for all CEs that have established multicast communication with it. In order to keep a consistent multicast state between a CE and BR, CE MUST use the same IPv6 multicast prefixes (ASM\_MPREFIX64/SSM\_REFIX64) until the state becomes empty. After that point, the CE may obtain different values for these prefixes, effectively changing to a different 4rd BR.

### **6.3. Supporting IPv6 Multicast in MAP-T and 4rd Translation Multicast**

IPv6 multicast can be supported natively since IPv6-only network is assumed to be multicast enabled. MAP-T or 4rd Customer Edge device has MLD Proxy function. Proxy operation for MLD [[RFC3810](#)] is described in [[RFC4605](#)].

CE receives MLD join requests from the hosts and then sends aggregated MLD Report messages upstream towards BR. No address or protocol translation is needed at the CE or at the BR. IPv6 Hosts in MAP-T or 4rd domain use any source multicast block FF0X [[RFC4291](#)] or source specific multicast block FF3X::8000:0-FF3X::FFFF:FFFF for dynamic allocation by a host [[RFC4607](#)], [[RFC3307](#)].

MAP-T or 4rd Border Relay is MLD querier. It serves all CEs downstream. After receiving an MLD join message, BR sends PIM join message upstream to join IPv6 multicast group. Multicast membership database is maintained based on the aggregated Reports received from





downstream interfaces in the multicast tree.

MAP-T or 4rd Border Relay is a Rendezvous Point (RP) for ASM groups. For SSM, BR MUST support MLDv2.

IPv6 multicast data received from the Single Source Multicast or Any Source Multicast sources are replicated according to the multicast membership database and the data packets are sent downstream on the multicast tree and eventually received by the CEs that have one of more members of this multicast group.

MLD Proxy in the CE receives multicast data then forwards the packet downstream. Each member host receives IPv6 multicast data packet from its Layer 2 interface.

#### **6.4. Learning Multicast Prefixes for IPv4-embedded IPv6 Multicast Addresses**

CE can be pre-configured with Multicast Prefix64 of ASM\_MPREFIX64 and SSM\_MPREFIX64 that are supported in their network. However automating this process is also desired.

A new router advertisement option, a Multicast ASM Translation Prefix option, can be defined for this purpose. The option contains IPv6 ASM multicast translation prefix, ASM\_MPREFIX64. A new router advertisement option, a Multicast SSM Translation Prefix option, can be defined for this purpose. The option contains IPv6 SSM multicast prefix translation prefix SSM\_MPREFIX64.

After the host gets the multicast prefixes, when an application in the host wishes to join an IPv4 multicast group the host MUST use ASM\_MPREFIX64 or SSM\_MPREFIX64 and then obtain the synthesized IPv6 group address before sending MLD join message.

Source-specific multicast (SSM) group membership message payloads in IGMPv3 and MLDv2 contain address literals and their translation requires another multicast translation prefix option. IPv4 source addresses in IGMPv3 Membership Report message are unicast addresses of IPv4 sources and they have to be translated into unicast IPv6 source addresses in MLDv2 Membership Report message. A new router advertisement option, a Multicast Translation Unicast Prefix option can be defined for this purpose. The option contains IPv6 unicast Network-Specific Prefix U\_PREFIX64. The host can be configured by its default router using router advertisements containing the prefixes [[I-D.sarikaya-softwire-6man-raoptions](#)]. 64:ff9b::/96 is the global value called well-known prefix that is assigned to U\_PREFIX64 [[RFC6052](#)]. Organization specific values called Network-Specific Prefixes can also be used. Since multicast is potentially inter-



domain, the use of well-known prefix for U\_PREFIX64 is recommended.

Note that U\_PREFIX64 is also used in multicast data packet address translation. Source-specific multicast source address in multicast data packets coming from SSM sources MUST be translated using U\_PREFIX64.

## **7. Security Considerations**

4rd Encapsulation Multicast control and data message security can be provided by the security architecture, mechanisms, and services described in [\[RFC4301\]](#), [\[RFC4302\]](#) and [\[RFC4303\]](#). 4rd Translation Multicast control and data message security are as described in [\[RFC4607\]](#) for source specific multicast.

## **8. IANA Considerations**

TBD.

## **9. Acknowledgements**

TBD.

## **10. References**

### **10.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), June 1999.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, [RFC 1112](#), August 1989.
- [RFC2113] Katz, D., "IP Router Alert Option", [RFC 2113](#), February 1997.
- [RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", [RFC 2711](#), October 1999.
- [RFC3171] Albanna, Z., Almeroth, K., Meyer, D., and M. Schipper, "IANA Guidelines for IPv4 Multicast Address Assignments",

[RFC 3171](#), August 2001.

- [RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", [RFC 4605](#), August 2006.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", [RFC 4607](#), August 2006.
- [RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", [RFC 3307](#), August 2002.
- [RFC2491] Armitage, G., Schulter, P., Jork, M., and G. Harter, "IPv6 over Non-Broadcast Multiple Access (NBMA) networks", [RFC 2491](#), January 1999.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", [RFC 2473](#), December 1998.
- [RFC2765] Nordmark, E., "Stateless IP/ICMP Translation Algorithm (SIIT)", [RFC 2765](#), February 2000.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", [RFC 3376](#), October 2002.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC4302] Kent, S., "IP Authentication Header", [RFC 4302](#), December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.
- [RFC5135] Wing, D. and T. Eckert, "IP Multicast Requirements for a Network Address Translator (NAT) and a Network Address Port Translator (NAPT)", [BCP 135](#), [RFC 5135](#), February 2008.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation



Algorithm", [RFC 6145](#), April 2011.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), October 2010.

[I-D.ietf-softwire-map]

Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., and T. Murakami, "Mapping of Address and Port with Encapsulation (MAP)", [draft-ietf-softwire-map-04](#) (work in progress), February 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-01](#) (work in progress), February 2013.

[I-D.ietf-softwire-4rd]

Jiang, S., Despres, R., Penno, R., Lee, Y., Chen, G., and M. Chen, "IPv4 Residual Deployment via IPv6 - a Stateless Solution (4rd)", [draft-ietf-softwire-4rd-04](#) (work in progress), October 2012.

[I-D.ietf-mboned-64-multicast-address-format]

Boucadair, M., Qin, J., Lee, Y., Venaas, S., Li, X., and M. Xu, "IPv6 Multicast Address With Embedded IPv4 Multicast Address", [draft-ietf-mboned-64-multicast-address-format-04](#) (work in progress), August 2012.

## **10.2. Informative references**

[RFC3306] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", [RFC 3306](#), August 2002.

[RFC3956] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", [RFC 3956](#), November 2004.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.

[I-D.sarikaya-softwire-6man-raoptions]

Sarikaya, B., "IPv6 RA Options for Translation Multicast Prefixes", [draft-sarikaya-softwire-6man-raoptions-00](#) (work in progress), August 2012.



[I-D.perreault-mboned-igmp-mld-translation]

Perreault, S. and T. Tsou, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Translation ("IGMP/MLD Translation")", [draft-perreault-mboned-igmp-mld-translation-01](#) (work in progress), April 2012.

## **[Appendix A.](#) Group Membership Message Translation Details**

IGMP Report messages (IGMP type number 0x12 and 0x16, in IGMPv1 and IGMPv2 and 0x22 in IGMPv3) are translated into MLD Report messages (MLDv1 ICMPv6 type number 0x83 and MLDv2 type number 0x8f). IGMP Query message (IGMP type number 0x11) is translated into MLD Query message (ICMPv6 type number 0x82)

[[I-D.perreault-mboned-igmp-mld-translation](#)].

Destination address in ASM, i.e. IGMPv1, IGMPv2 and MLDv1, is the multicast group address so the destination address in IGMP message is translated into the destination address in MLD message using [[I-D.ietf-mboned-64-multicast-address-format](#)].

Destination address in SSM, i.e. IGMPv3 and MLDv2 is translated as follows: it could be all nodes on link, which has the value of 224.0.0.1 (IGMPv3) and ff02::1 (MLDv2), all routers on link, which has the value of 224.0.0.2 (IGMPv3) and ff02::2 (MLDv2), all IGMP/MLD-capable routers on link, which has the value of 224.0.0.22 (IGMPv3) and ff02::16 (MLDv2).

Source address of MLD message that CE sends is set to link-local IPv6 address of CE's WAN side interface. Source address of MLD message that BR sends is set to link-local IPv6 address of BR's downstream interface.

Multicast Address or Group Address field in IGMP message payloads is translated using [[I-D.ietf-mboned-64-multicast-address-format](#)] as described above into the corresponding field in MLD message.

Source Address in IGMPv3 message payloads is translated using U\_PREFIX64, the IPv6 unicast prefix to be used by SSM source. [[RFC6052](#)] defines in [Section 2.3](#) the address translation algorithm of embedding an IPv4 source address and obtaining an IPv6 source address using a network specific prefix like U\_PREFIX64. At the BR on its upstream interface or at the CE on its LAN interface, IPv4 addresses are extracted from the IPv4-embedded IPv6 addresses.

Maximum Response Time (MRT) field in IGMPv2 and IGMPv3 queries are translated into Maximum Response Delay (MRD) in MLDv1 and MLDv2





queries, respectively. In the corresponding MLD message, MRD is set to 100 times the value of MRT. At the BR on its upstream interface or at the CE on its LAN interface, MRT value is obtained by dividing MRD into 100 and rounding it to the nearest integer.

IGMP messages are sent with a Router Alert IPv4 option [[RFC2113](#)]. The translated MLD message are sent with a Router Alert option in a Hop-By-Hop IPv6 extension header [[RFC2711](#)]. In both cases, 2-octet value is set to 0.

Author's Address

Behcet Sarikaya  
Huawei USA  
5340 Legacy Dr. Building 175  
Plano, TX 75024

Phone:

Email: [sarikaya@ieee.org](mailto:sarikaya@ieee.org)