

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 25, 2011

S. Venaas
UNINETT
H. Asaeda
Keio University
S. SUZUKI
ALAXALA Networks Corporation
T. Fujisaki
NTT PF Lab
December 22, 2010

An IPv4 - IPv6 multicast translator
draft-venaas-behave-mcast46-02.txt

Abstract

This document describes an IPv4 - IPv6 translator device that embeds all IPv4 multicast group addresses into IPv6, and allows IPv6 hosts to receive from and send to any IPv4 multicast group. This mechanism can be also used to allow IPv4 hosts to receive from and send to a subset of the IPv6 multicast groups.

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 June 25, 2011.

Copyright Notice

Copyright (c) 2010 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. Terminology	3
3. Abbreviations	4
4. Architecture	4
5. Address Translation	5
5.1. Embedding IPv4 multicast addresses into IPv6	5
5.2. Translating IPv6 multicast addresses into IPv4	6
5.3. Embedding IPv4 source addresses into IPv6	7
5.4. Translating IPv6 source addresses into IPv4	7
6. Examples	8
6.1. IPv6 host joining a group inside the /96 prefix	8
6.2. IPv6 host sending to group inside the /96 prefix	8
6.3. IPv4 host joining an IPv4 group	9
6.4. IPv4 host sending to a group a.b.c.d	9
7. Acknowledgments	9
8. Security Considerations	10
9. References	10
9.1. Normative References	10
9.2. Informative References	10
Authors' Addresses	11

1. Introduction

IPv4 and IPv6 will co-exist for many years, possibly decades. There are several solutions for how IPv4 and IPv6 hosts and networks can inter-operate. This is usually easy if a host is dual-stack. If however an IPv6-only host needs to communicate with an IPv4-only host, then somewhere along the data path there must be some form of translation. There are several ways of doing this for unicast, but not much work has been done on multicast.

Here we describe a multicast translator solution. This translator could be placed at the border between IPv6-only and IPv4-only networks to allow multicast access between them, or it may also be placed in a dual-stack network, where it can support hosts or other networks that are IPv6-only or IPv4-only. The goal is to give an IPv6 host full access to send to and receive from any IPv4 multicast group by using the usual IPv6 multicast protocols and applications which will then operate on the respective IPv6 groups. It should also allow this for multiple hosts. Multiple IPv4 hosts should be able to use a single IPv4 group, multiple IPv6 hosts a corresponding IPv6 group, and all hosts should be able to send to and receive from all the others. Similar to hosts using the same group from the same address family. The translator solution should work with no changes to other infrastructure.

We will define a one-to-one mapping of multicast IPv4 addresses onto a subset of the IPv6 multicast addresses. An IPv6 host will then be able to receive data from any IPv4 multicast group by joining the corresponding IPv6 group. An IPv6 host can also send, without necessarily joining, to any IPv4 multicast group by sending to the corresponding IPv6 group. Some way of translating unicast addresses is also needed to translate addresses of multicast sources.

The one-to-one mapping also allows an IPv4 host access to send to and receive from the mapped IPv6 multicast groups.

2. Terminology

"IPv6_TRASM_ADDRESS" is an IPv6 ASM multicast address translated by IPv4-IPv6 multicast translator. IPv6_TRASM_ADDRESS is one of the addresss in one specific /96 IPv6 ASM (non-SSM) prefix ("IPv6_TRASM_ADDRESS prefix"). Since this document assumes the translator acts as an RP in IPv6 ASM, this prefix may be used with embedded-RP and be included in FF70::/12 [1].

"IPv6_TRSSM_ADDRESS" is an IPv6 SSM multicast address translated by IPv4-IPv6 multicast translator. IPv6_TRSSM_ADDRESS is one of the

addressss in one specific /96 IPv6 SSM prefix ("IPV6_TRSSM_ADDRESS prefix"). This prefix will be included in FF3x:0::/32 [2].

"IPV4_SSM_ADDRESS" is an IPv4 SSM multicast address. IPV4_SSM_ADDRESS is one of the addresses from the IPv4 SSM address 232/8 [2].

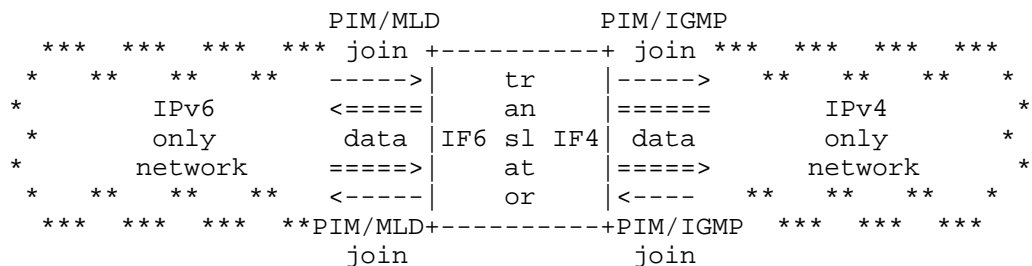
"IPV4_ASM_ADDRESS" is an IPv4 ASM multicast address translated by IPv4-IPv6 multicast translator. It can in principle be any IPv4 multicast address outside the SSM range 232/8.

A multicast translator attaches both IPv6-only and IPv4-only networks with two different interfaces. In this document, each interface is named "IF6" and "IF4".

3. Abbreviations

ASM	Any Source Multicast
DR	Designated Router
IGMP	Internet Group Management Protocol
MLD	Multicast Listener Discovery
MSDP	Multicast Source Discovery Protocol
PIM-SM	Protocol Independent Multicast - Sparse Mode
RP	Rendezvous Point
RPF	Reverse Path Forwarding
SSM	Source-Specific Multicast

4. Architecture



We propose that the translator makes use of PIM-SM (Sparse Mode) [3] for IPv6. For ASM it should then be the RP for the /96 IPv6 prefix used for ASM, "IPV6_TRASM_ADDRESS prefix". This allows the

translator to know which IPv4 groups the IPv6 hosts join, and also to learn of IPv6 sources for those groups. It is sufficient to support MLD if there are no IPv6 PIM neighbors (e.g. a single link or MLD proxies [4]).

When the translator receives a PIM or MLD join for an IPv6 group in "IPV6_TRASM_ADDRESS prefix" on IF6, it will need to join the corresponding IPv4 multicast group on IF4. It may behave as an IPv4 host and send an IGMP join for the correspondig IPv4 group, or it might be an IPv4 PIM router and send an IPv4 PIM join.

If the translator learns of an IPv6 source for an "IPV6_TRASM_ADDRESS" it needs to receive the data on IF6, and send the translated data out on IF4. If the translator is an IPv6 RP, it may receive IPv6 PIM. As a regular IPv6 RP, it may then join towards the source to receive packets natively on IF6. It may also be a directly connected source or a source behind an MLD proxy [4], in that case packets are also received on IF6. If the translator behaves as an IPv4 host, it sends any such IPv6 packets out on IF4.

One can improve on this by making the translator behave as an IPv4 RP, or be an IPv4 PIM router running MSDP [5] to exchange information about active IPv4 sources. The translator can then use MSDP to signal its active IPv4 sources (that may be translated IPv6 sources) so that it will receive PIM joins if there are IPv4 receivers for the groups. It can also use MSDP to see if there are IPv4 sources for IPv4 groups that IPv6 hosts have joined.

Note that for SSM this is much simpler with no RP nor MSDP involved. It may still be an advantage to act as an IPv4 PIM router, in order to only do translation from IPv6 to IPv4 when there are IPv4 listeners.

5. Address Translation

When IPv4 packets are resent as IPv6 we will need to replace the source and destination addresses with suitable IPv6 addresses. And similar replacement going from IPv6 to IPv4. The source addresses are always unicast addresses, and the destination addresses are always multicast addresses.

5.1. Embedding IPv4 multicast addresses into IPv6

We need a way of referring to an IPv4 multicast group using an IPv6 address. IPv4 multicast addresses are embedded into IPv6 by simply prepending them with a specific /96 IPv6 prefix such that for each IPv4 multicast address we have a respective IPv6 multicast address.

However, both IPv4 and IPv6 have special ranges for SSM usage, and one might want to take scoping into account.

This document proposes the use of one specific /96 IPv6 SSM prefix (IPV6_TRSSM_ADDRESS prefix) for all IPv4 SSM addresses, and one specific /96 IPv6 ASM (non-SSM) prefix (IPV6_TRASM_ADDRESS prefix) for all IPv4 ASM (non-SSM) addresses. Hence IPv4 multicast addresses are embedded into IPv6 by appending them with a /96 IPV6_TRSSM_ADDRESS prefix if the IPv4 multicast addresses are in the IPV4_SSM_ADDRESS range, or with a /96 IPV6_TRASM_ADDRESS prefix if they are in the IPV4_ASM_ADDRESS range.

An administrator may choose the exact prefixes used, and depending on the prefix, also which IPv6 scope. The prefix must be in accordance with the IPv6 multicast address format defined in section 2.7 of [6]. The addresses used will then be of the form FFxx:<blah>:<IPv4> where flags, scope and the value of "blah" are chosen by the administrator. "IPv4" is the last 32 bits specifying the IPv4 address of the IPv4 multicast group. For ASM it may be useful to use an Embedded-RP [1] prefix based on an IPv6 unicast address of the translator.

Note that as specified in [7] the IPv4 address will become the Group ID, and since all IPv4 multicast addresses have the leading bit set, the IPv6 multicast addresses will become "server allocated" addresses. We can regard the translator as the "allocation server".

The unicast addresses of multicast sources also need to be translated. We recommend embedding all IPv4 unicast addresses into a /96 IPv6 prefix. This allows different IPv4 unicast addresses to be mapped to different IPv6 unicast addresses, and for IPv6 SSM joins to address specific IPv4 SSM sources. Note that for ASM use, it may be sufficient to map all IPv4 sources to one single IPv6 address. For translating IPv6 sources into IPv4 sources, one may use a single address, or a pool of IPv4 addresses. The same IPv4 address may need to be re-used for different IPv6 sources. If the translator also translates unicast packets, then it should use the same unicast translation mechanism for source addresses in multicast packets. Due to multicast RPF checks, the IPv4 and IPv6 unicast addresses used need to be routed towards the translator.

5.2. Translating IPv6 multicast addresses into IPv4

For the multicast address translation from IPv6 to IPv4, we simply extract the last 32 bits. However, if the IPv6 multicast address is not in the range of either IPV6_TRSSM_ADDRESS or IPV6_TRASM_ADDRESS range, IPv4 hosts cannot join the multicast whose destination address is not in these address ranges. Therefore the translator does not translate and does not forward such data from the interface IF4.

The translator can be implemented on a host (bumped into a stack), or a layer 3 device. In the latter case, link-local multicast addresses MUST NOT be translated, since a CPE has to treat IPv4 link-local scope and IPv6 link-local scope as different ones. (Please note that this is specific to an IGMP/MLD-based translator, since PIM does not generate a join for link-local multicast addresses.)

5.3. Embedding IPv4 source addresses into IPv6

Unicast addresses of multicast sources also need to be translated. An IPv4 unicast address of a multicast source is embedded into a /96 IPv6 unicast prefix. The /96 IPv6 unicast prefix will be prepared as the address pool of the translator. It will be same or part of the unicast prefix assigned at the translator's IF6. This allows PIM join messages to be forwarded to the translator, and also enables IPv6 SSM joins to be translated to IPv4 SSM joins.

One could consider using just a single IPv6 unicast address for all IPv4 multicast translated into IPv6. For ASM use, it may be sufficient to map all IPv4 sources to one single IPv6 address, and this single IPv6 address can be the translator's IPv6 global unicast address assigned on IF6, because this document assumes that the translator acts as an RP for IPv6 PIM. On the other hand, for SSM, each IPv6 SSM join should be translated to uniquely specify a corresponding IPv4 SSM join. In order to do this, the simplest way is that an IPv4 unicast address of a multicast source is embedded into a /96 IPv6 unicast prefix the translator prepared.

An alternative to using a /96 IPv6 unicast prefix could also be to dynamically allocate IPv6 unicast addresses from a pool, see [8].

5.4. Translating IPv6 source addresses into IPv4

For translating IPv6 sources into IPv4 sources, one may use a single address, or a pool of IPv4 addresses. The same IPv4 address may need to be re-used for different IPv6 sources. If the translator also translates unicast packets, then it should use the same unicast translation mechanism for source addresses in multicast packets.

If unicast traffic is translated, then similar translation should be used for the multicast source addresses. Note that for RTP the application can know the real source and tell streams apart, even if they are translated into the same multicast source address. The translation mechanism for IGMP/MLD/PIM MUST be same as the one for the multicast data.

6. Examples

To illustrate how the translator works, we will look at some examples. In all the examples we assume that there is no previous state in the translator.

6.1. IPv6 host joining a group inside the /96 prefix

An IPv6 host joins the group FFxx:<blah>:a.b.c.d. If the translator is the DR for the host, it will receive an MLD membership report. If not, it will receive a PIM join since it is the RP for the group. The translator checks whether the address is inside the /96 prefix, and whether the last 32 bits (a.b.c.d) is an IPv4 multicast address. If it is, it joins a.b.c.d using IGMP or PIM, and stays joined as long as there are IPv6 receivers.

For SSM the translator would in addition check if the source in the join is inside the /96 unicast prefix used. If this is the case, it then uses the last 32 bits as the IPv4 source. It can then do a source-specific IPv4 join.

When the translator receives a multicast packet for a.b.c.d it prepends the /96 prefix to form the IPv6 address FFxx:<blah>:a.b.c.d. If the translator has outgoing interfaces for this group, it will send an IPv6 packet to the same interfaces to which it would have forwarded an IPv6 packet for the group. The destination address will be FFxx:<blah>:a.b.c.d, and the source address will be computed using the /96 unicast prefix. For SSM, the translator would also check that it got an outgoing interface for the specific source.

6.2. IPv6 host sending to group inside the /96 prefix

An IPv6 host sends to the group FFxx:<blah>:a.b.c.d. If the translator is the DR for the host, it will receive the data natively. If not, it will receive PIM register messages containing the data since it is the RP. For each packet received, either natively or inside register messages, it will first check that the destination address is inside the /96 prefix and that the last 32 bits (a.b.c.d) is an IPv4 multicast address. If this is okay, it will resend the packet to the IPv4 address a.b.c.d. The source address would be chosen from a given pool of IPv4 unicast addresses (this may just be a single fixed address).

If the translator is also an IPv4 PIM router, then we do some further steps. For ASM, if the translator is an RP and uses MSDP, it should announce the translated source in MSDP, and only forward translated packets if it has a join for the group. For SSM, it should only forward translated packets if it has a join for the specific source

and group.

6.3. IPv4 host joining an IPv4 group

In some cases, ISP network supports IPv6-multicast, but a host or an application only supports IPv4-multicast. IGMP/MLD-proxy based translator helps such case by accommodating such hosts. (IGMP/PIM-based translator would also work, but it is not described here since the behavior is almost same as Section 6.1 and Section 6.2.)

We have mapped IPv4 multicast groups to a subset of the IPv6 multicast groups by embedding them in /96 IPv6 prefixes. Typically one prefix for ASM and one for SSM. An IPv4 host joins to the group a.b.c.d, and the translator will receive the IGMP join and resend an MLD join for FFxx:<blah>::a.b.c.d to IF6. When an MLD query arrives from IF6, the translator replies with MLD reports based on the IGMP membership database (a.b.c.d) and the /96 prefix (FFxx:<blah>::).

In case of SSM, the translator in addition synthesizes an IPv6 source address from the IPv4 source address in the IGMP join (see Section 5.3), and resends a source-specific MLD join for the synthesized IPv6 address.

When the translator receives an IPv6 multicast packet, it checks whether the group address is inside the /96 prefix, and whether the last 32bits (a.b.c.d) is an IPv4 multicast address. If both hold true and a.b.c.d exists in the IGMP membership database, the translator will convert the received IPv6 packet into IPv4 and send it for the IF4 interfaces where a.b.c.d is subscribed. The source address of the IPv4 packet is synthesised as in Section 5.4, and its destination address is a.b.c.d.

6.4. IPv4 host sending to a group a.b.c.d

When the translator receives a packet for a.b.c.d and it is also a DR for the host, it will convert the received IPv4 packet into IPv6 and send it to IF6. The source and destination address of the IPv6 packet is synthesized as in Section 5.3 and Section 5.1, respectively.

7. Acknowledgments

The authors thank Michal Przybylski and Pekka Savola for valuable comments, and also people from the M6Bone community for testing a prototype implementation. Also thanks to Teemu Kiviniemi who has implemented and deployed a second translator implementation based on this document.

8. Security Considerations

When using such a translator one needs to take some care of scoping and TTL values. Due to differences in IPv4 and IPv6 scoping, a narrow scope might be translated into a wider one.

One may wish to limit who can access the translator. If for instance one wishes to restrict it to a site, one can use a /96 prefix of site-local scope, and then filter at the site border, just like one would for multicast in general. A translator implementation could also offer a way of restricting which groups and sources should be accepted.

9. References

9.1. Normative References

- [1] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, November 2004.
- [2] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.
- [3] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [4] 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.
- [5] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.
- [6] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [7] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

9.2. Informative References

- [8] Tsuchiya, K., Higuchi, H., Sawada, S., and S. Nozaki, "An IPv6/IPv4 Multicast Translator based on IGMP/MLD Proxying (mtp)", draft-tsuchiya-mtp-01 (work in progress), February 2003.

Authors' Addresses

Stig Venaas
UNINETT
Trondheim NO-7465
Norway

Email: venaas@uninett.no

Hitoshi Asaeda
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-8520
Japan

Email: asaeda@wide.ad.jp

Shinsuke SUZUKI
ALAXALA Networks Corporation
Shinkawasaki Mitsui Bldg.
890 Kashimada
Saiwai-ku, Kawasaki, Kanagawa 212-0058
Japan

Email: suz@alaxala.net

Tomohiro Fujisaki
NTT PF Lab
3-9-11 Midori-Cho
Musashino-shi, Tokyo 180-8585
Japan

Email: fujisaki@nttv6.net

