

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 11, 2013

B. Sarikaya
Huawei USA
T. Kiviniemi
CSC
January 7, 2013

Multicast Support for NAT64
draft-sarikaya-behave-mcast4nat64-07.txt

Abstract

This memo specifies modifications required to NAT64 so that IPv6 only hosts can receive multicast data from IPv4 only servers. The protocol is based on translating IPv4 multicast data before delivering it to the host in IPv6. The protocol also allows IPv6 only host to join IPv4 any source/ source specific multicast group in IPv6 using Multicast Listener Discovery protocol.

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 July 11, 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

1.	Introduction	3
2.	Terminology	3
3.	Requirements	3
4.	Architecture	3
5.	NAT64 Multicast Operation	5
5.1.	Address Translation	5
5.1.1.	Learning Multicast Prefixes for IPv4-embedded IPv6 Multicast Addresses	6
5.2.	Protocol Translation	7
5.2.1.	Differences with IP/ICMP Translation	8
5.2.2.	PIM versus IGMP	8
6.	Security Considerations	9
7.	IANA Considerations	9
8.	Acknowledgements	9
9.	References	9
9.1.	Normative References	9
9.2.	Informative references	10
	Authors' Addresses	11

1. Introduction

With IPv4 address depletion on the horizon, many techniques are being standardized for IPv6 migration including NAT64 [[RFC6146](#)]. NAT64 together with DNS64 [[RFC6147](#)] and the translation algorithm [[RFC6145](#)] enables IPv6-only hosts to communicate with IPv4-only servers.

NAT64 currently supports only unicast communication [[RFC6146](#)], [[RFC6145](#)], [[RFC6052](#)]. With the advent of IPTV and Mobile IPTV, there is a need to provide support for multicast communication as well.

The document continues in [Section 3](#) with a set of requirements on a solution for NAT64 multicast support. In [Section 4](#) the architecture is presented. Multicast translation protocol is explained in [Section 5](#).

2. Terminology

This document uses the terminology defined in [[RFC6146](#)], [[RFC6145](#)], [[RFC6052](#)], [[I-D.ietf-mboned-64-multicast-address-format](#)], [[RFC3810](#)] and [[RFC3376](#)].

3. Requirements

This section states requirements on NAT64 translation protocol.

The protocol MUST support IPv4-embedded IPv6 multicast addresses as defined in [[I-D.ietf-mboned-64-multicast-address-format](#)]. The translation protocol MUST enable an IPv6 only host to join IPv4 multicast groups where IPv6 only host identifies IPv4 groups using IPv4-embedded IPv6 multicast addresses.

Both any source multicast (ASM) and source specific multicast (SSM) MUST be supported.

In IPv4 network, Protocol Independent Multicast routing MAY be supported. In IPv4 network, Internet Group Management Protocol routing MAY be supported.

User Datagram Protocol (UDP) MUST be supported. Transmission Control Protocol (TCP) MAY be supported.

4. Architecture

We consider an IPv6 only host (Host 1, H1) that wishes to receive

multicast data sent to IPv4 multicast groups, sent by an IPv4 only host (Host 2, H2). Multicast data sent to an IPv4 multicast group such as 232.1.2.3 must be translated into an IPv6 multicast group data such as FF3E::232.1.2.3. So a translator element is needed in the architecture. The translator has to be connected simultaneously into IPv4 network and IPv6 network.

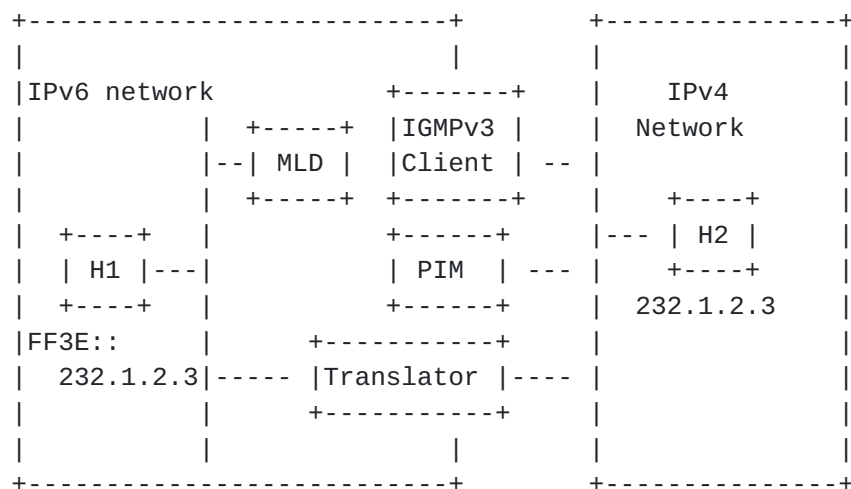


Figure 1: Key elements of NAT64 Multicast Translator

In order to receive multicast data, the host H1 must first subscribe to the multicast group of interest. In IPv6 this is done using MLD protocol [[RFC3810](#)] by sending MLD Membership Report message indicating the group address which should have an IPv4 multicast group address embedded such as FF3E::232.1.2.3. MLD entity has to communicate the group membership information to an entity that supports wide-area multicast routing protocol such as PIM [[RFC3973](#)], [[RFC4601](#)], [[RFC5015](#)]. PIM supports both IPv4 and IPv6.

IPv4 group address in MLD membership Report message should be communicated to an entity that supports IGMP protocol [[RFC3376](#)]. So an IGMP Client is needed to handle joining and leaving IPv4 multicast groups by sending IGMPv3 Report messages to IPv4 network. Once IGMP Client subscribes to an IPv4 multicast group, all IPv4 multicast packets can be received from the interface connected to IPv4 network. The translator translates such packets into IPv6 multicast data packets and forwards them to IPv6 network which delivers it to IPv6 hosts that have joined the corresponding IPv6 multicast group.

All the elements of NAT64 multicast translation system are shown in Figure 1. Not shown in the figure are MLD Proxies which are located in Host 1's first hop router. MLD Proxy optimizes MLD operation by providing only aggregate multicast group membership information to the upstream MLD router and duplicating multicast data at a place

close to the hosts.

Note that the architecture in Figure 1 is generic and does not prescribe any solution as to where in a real network the different components can be hosted or whether the architecture can be duplicated in the same network. While in unicast communication multiple NAT64 boxes can be supported in an operator's network using multiple Pref64s, in multicast NAT64 the same does not hold because IPv6 only hosts do not send multicast data. The elements in the architecture in Figure 1 are best placed in where the designated MLD router/ querier is hosted. In broadband networks Broadband Network Gateway (BNG), in 3GPP networks Packet Data Network Gateway (P-GW) are the candidates for such a placement. This implies that NAT64 multicast translator may be hosted in a different network element than NAT64 unicast translator [[RFC6146](#)].

5. NAT64 Multicast Operation

In this section we specify how the host can receive IPv4 multicast data from IPv4-only content provider based on the architecture defined in [Section 4](#). The reverse translation of IPv6 multicast data for IPv4-only receivers is out of scope. Multicast translation involves address translation defined in [Section 5.1](#) and protocol (IPv4 to IPv6) translation defined in [Section 5.2](#).

5.1. Address Translation

IPv6-only H1 will join IPv4 multicast group by sending MLD Membership Report message upstream towards the MLD entity in Figure 1. H1 MUST use 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.

In both ASM_MPREFIX64 and SSM_MPREFIX64, M bit MUST be set to 1 to indicate that an IPv4 address is embedded in the last 32 bits of the multicast IPv6 address. ASM_MPREFIX64 values are formed by setting flgs bits to make it an embedded RP prefix by setting R bit to 1 and P and T bits to 1 as shown in Figure 2 [[RFC4291](#)], [[RFC3306](#)], [[RFC3956](#)].

8	4	4	4	76	32
+-----+-----+-----+-----+-----+-----+					
11111111	0111	scop	1000	sub-group-id	v4 address
+-----+-----+-----+-----+-----+-----+					

Figure 2: ASM_MPREFIX64 Formation

Each translator is assigned a unique ASM_MPREFIX64 prefix. The hosts can learn this value by means out of scope with this document. With this, the host can easily create an IPv6 multicast address from the IPv4 group address a.b.c.d that it 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 values are formed by setting R bit to zero, P and T bits to 1. This gives FF3x00008x::/96 as the SSM prefix. This prefix is referred to as SSM_PREFIX64 Figure 3.

8	4	4	16	4	60	32
+-----+-----+-----+-----+-----+-----+-----+						
11111111	0011	scop	00.....00	1000	sub-group-id	v4 address
+-----+-----+-----+-----+-----+-----+-----+						

Figure 3: SSM_MPREFIX64 Formation

Since SSM translation requires a unique address for each IPv4 multicast source, an IPv6 unicast prefix must be configured to the translator 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 entity in the translator. The translator 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.

5.1.1. Learning Multicast Prefixes for IPv4-embedded IPv6 Multicast Addresses

The hosts 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.

[[I-D.sarikaya-softwire-6man-raoptions](#)] describes a Router

Advertisement based method which is suitable for mobile hosts accessing 3GPP networks.

A new DHCPv6 option, `OPTION_AFT_PREFIX_DHCP`, can be defined for this purpose. The option contains IPv6 ASM and SSM prefixes. The host can request these prefixes by sending this option in its request to the DHCP server and the server replies with the option containing the prefixes.

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.

5.2. Protocol Translation

Translator, after processing the addresses will then translate IPv4 multicast data packet into an IPv6 multicast data packet. The destination address is IPv6 group address `ASM_MPREFIX64::a.b.c.d` and source address is the translator's IPv6 interface address. The value in Type of Service (TOS) field of IPv4 packet is copied into IPv6 Traffic Class field. IPv4 Protocol and TTL fields are copied into IPv6 Next Header and Hop Limit fields respectively. IPv4 payload is copied into IPv6 payload. UDP checksum is updated which completes the packet translation process [[Thesis](#)]. The packet is sent towards the host and on its way it may be duplicated for each member of this group and then sent to the individual host separately.

Any IPv4 fragments sent by the routers must be translated into IPv6 packets with IPv6 Fragment Header. Fragmentation Offset field is copied into the corresponding field in the Fragment Header. 16-bit Identification field is copied into the low-order 16 bits of IPv6 Fragment Header Identification field. The high-order bits of the 32-bit IPv6 Fragment Header Identification field are set to zero. More Fragments (MF) flag is copied to the corresponding field in IPv6 Fragment Header [[Thesis](#)].

Multicast translation described in this section is not specific to the hosts. Translator gets the join message from the host and then updates the membership database. Translator and any MLD Proxies downstream have to know all members of each IPv4 group so that they can correctly duplicate the data packets and deliver to the individual hosts.

Also this prefix must be routed towards the translator on the IPv6 network, to enable reverse path forwarding for multicast, and to inform other PIM routers about the correct destination for PIM (S,G) Join messages [[Thesis](#)].

5.2.1. Differences with IP/ICMP Translation

The Stateless IP/ICMP translation [[RFC6145](#)] is designed for unicast communication. IP/ICMP Translation uses a different address translation than multicast address translation described [Section 5.1](#). However some parts of IP/ICMP Translation can be used in multicast translation. In this section we describe the differences of IP/ICMP Translation with NAT64 translation.

IP/ICMP Translation translates IPv4 packets into IPv6 using minimum MTU size of 1280 bytes. However in DVB-IPTV data streams, 1364 byte IPv6 packets need to be supported. NAT64 translator must perform IPv6 path MTU discovery and set the MTU size accordingly, not necessarily always to 1280 byte MTU size. Note that IPv4 routers must not send ICMPv4 error message in response to a multicast packet [[RFC1812](#)] while in IPv6 this is not the case which enables path MTU discovery for multicast. Path MTU values are kept in the translator for each multicast group. However, for SSM, a different MTU value MUST be kept for each SSM channel.

IP/ICMP Translation does not require transport layer checksum modifications because the prefixes used in IP/ICMP Translation are checksum neutral. NAT64 translator 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.

5.2.2. PIM versus IGMP

To handle joining and leaving IPv4 multicast groups, PIM client instead of IGMPv3 client as mentioned in [Section 4](#) can be used. With IGMPv3, SSM support requires that all neighboring routers support IGMPv3. If they support IGMPv1 or IGMPv2 the translator has to also support IGMPv1 or IGMPv2 for compatibility and as a result, SSM could not be supported.

SSM can be supported with a PIM Client at the translator. There are other advantages of having PIM client at the translator. PIM can communicate with neighboring IPv4 routers over flexible connections. The need to enable IGMP in the neighboring routers is removed if PIM is used.

On the other hand, IGMP, being a simpler protocol than PIM, is more widely supported. IGMP does not need to know the location of rendezvous points (RP) which makes it easier to configure the translator.

6. Security Considerations

Security considerations for IPv4 interface of the translator is similar to [[RFC6146](#)] and the considerations stated there apply.

7. IANA Considerations

TBD.

8. Acknowledgements

TBD.

9. References

9.1. 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.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", [RFC 1812](#), June 1995.
- [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.
- [RFC3973] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", [RFC 3973](#), January 2005.
- [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.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", [RFC 4601](#), August 2006.

[RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", [RFC 5015](#), October 2007.

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

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), April 2011.

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

[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.

9.2. Informative references

[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.

[Thesis] Teemu Kiviniemi, Helsinki University of Technology, Master's Thesis, "Implementation of an IPv4 to IPv6 Multicast Translator", October 2009.

Authors' Addresses

Behcet Sarikaya
Huawei USA
5340 Legacy Drive, Suite 175
Plano, TX 75074

Phone: +1 469-277-5700
Email: sarikaya@ieee.org

Teemu Kiviniemi
CSC

Phone:
Email: teemu.kiviniemi@csc.fi

