Softwire WG
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2012

Q. Wang
China Telecom
J. Qin
ZTE
M. Boucadair
C. Jacquenet
France Telecom
Y. Lee
Comcast

October 31, 2011

# Multicast Extensions to DS-Lite Technique in Broadband Deployments
## draft-ietf-softwire-dslite-multicast-01

Abstract

This document specifies a solution for the delivery of multicast service offerings to DS-Lite serviced customers. The proposed solution relies upon a stateless IPv4-in-IPv6 encapsulation scheme and uses the IPv6 multicast distribution tree to deliver IPv4 multicast traffic over an IPv6 multicast-enabled network.

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 May 3, 2012.

Copyright Notice

Copyright (c) 2011 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
1.1. Requirements Language . . . . . . . . . . . . . . . . . . . 3
2. Terminology . . . . . . . . . . . . . . . . . . . . . . . . . 3
3. Scope . . . . . . . . . . . . . . . . . . . . . . . . . . . . 4
4. Solution Overview . . . . . . . . . . . . . . . . . . . . . . 5
4.1. IPv4-embedded IPv6 Prefixes . . . . . . . . . . . . . . . . 6
4.2. Multicast Distribution Tree Computation . . . . . . . . . . 7
4.3. Multicast Data Forwarding . . . . . . . . . . . . . . . . . 8
5. Address Mapping . . . . . . . . . . . . . . . . . . . . . . . 8
5.1. Prefix Assignment . . . . . . . . . . . . . . . . . . . . . 8
5.2. Examples . . . . . . . . . . . . . . . . . . . . . . . . . 9
6. Multicast B4 (mB4) . . . . . . . . . . . . . . . . . . . . . . 9
6.1. IGMP-MLD Interworking Function . . . . . . . . . . . . . . 9
6.2. Multicast Data Forwarding . . . . . . . . . . . . . . . . . 10
6.3. Fragmentation . . . . . . . . . . . . . . . . . . . . . . . 10
6.4. Host built-in mB4 Function . . . . . . . . . . . . . . . . 10
7. Multicast AFTR (mAFTR) . . . . . . . . . . . . . . . . . . . . 11
7.1. Routing Considerations . . . . . . . . . . . . . . . . . . 11
7.2. Processing PIM Message . . . . . . . . . . . . . . . . . . 11
7.3. Switching from Shared Tree to Shortest Path Tree . . . . . 12
7.4. Multicast Data Forwarding . . . . . . . . . . . . . . . . . 12
7.5. TTL/Scope . . . . . . . . . . . . . . . . . . . . . . . . . 13
8. Security Considerations . . . . . . . . . . . . . . . . . . . 13
8.1. Firewall Configuration . . . . . . . . . . . . . . . . . . 13
9. Acknowledgements . . . . . . . . . . . . . . . . . . . . . . . 13
10. IANA Considerations . . . . . . . . . . . . . . . . . . . . . 13
11. References . . . . . . . . . . . . . . . . . . . . . . . . . 14
11.1. Normative References . . . . . . . . . . . . . . . . . . . 14
11.2. Informative References . . . . . . . . . . . . . . . . . . 15
Appendix A. Use Case: IPTV . . . . . . . . . . . . . . . . . . . 15
Appendix B. Deployment Considerations . . . . . . . . . . . . . . 16
B.1. Load-Balancing . . . . . . . . . . . . . . . . . . . . . . 16
B.2. RP for IPv4-Embedded IPv6 Multicast Groups . . . . . . . . 16
B.3. mAFTR Policy Configuration . . . . . . . . . . . . . . . . 16
B.4. Static vs. Dynamic PIM Triggering . . . . . . . . . . . . . 17
Authors' Addresses . . . . . . . . . . . . . . . . . . . . . . . . 17

# 1. Introduction

DS-Lite [RFC6333] is a technique that rationalizes the usage of the remaining global IPv4 addresses during the transition period by sharing a single IPv4 address with multiple users. A typical DS-Lite scenario is the delivery of an IPv4 service to an IPv4 user over an IPv6 network (denoted as a 4-6-4 scenario). [RFC6333] covers unicast services exclusively.

If customers have to access IPv4 multicast-based services through DS-Lite environment, Address Family Transition Router (AFTR) devices will have to process all the IGMP Report messages [RFC2236] [RFC3376] that have been forwarded by the CPE into the IPv4-in-IPv6 tunnels. From that standpoint, AFTR devices are likely to behave as a replication point for downstream multicast traffic. And the multicast packets will be replicated for each tunnel endpoint where IPv4 receivers are connected to.

This kind of DS-Lite environment raises two major issues:

1. The IPv6 network loses the benefits of the multicast traffic forwarding efficiency because it is unable to deterministically replicate the data as close to the receivers as possible. As a consequence, the downstream bandwidth in the IPv6 network will be vastly consumed by sending multicast data over a unicast infrastructure.

2. The AFTR is responsible for replicating multicast traffic and forwarding it into each tunnel endpoint connecting IPv4 receivers that have explicitly asked for the corresponding contents. This process may greatly consume AFTR's resources and overload the AFTR.

This document specifies an extension to the DS-Lite model to deliver IPv4 multicast services to IPv4 clients over an IPv6 multicast-enabled network.

## 1.1. 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].

# 2. Terminology

This document makes use of the following terms:

- IPv4-embedded IPv6 address: is an IPv6 address which embeds a 32-bit-encoded IPv4 address. An IPv4-embedded IPv6 address can be unicast or multicast.

- mPrefix64: is a dedicated multicast IPv6 prefix for constructing IPv4-embedded IPv6 multicast addresses [I-D.boucadair-behave-64-multicast-address-format]. mPrefix64 can be of two types: ASM_mPrefix64 used in Any Source Multicast (ASM) mode or SSM_mPrefix64 used in Source Specific Multicast (SSM) mode [RFC4607].

- uPrefix64: is a dedicated IPv6 unicast prefix for constructing IPv4- embedded IPv6 unicast addresses [RFC6052].

- Multicast AFTR (mAFTR): is a functional entity which supports IPv4- IPv6 multicast interworking function (refer to Figure 3). It receives and encapsulates the IPv4 multicast packets into IPv4-in-IPv6 packets and behaves as the corresponding IPv6 multicast source for the encapsulated IPv4-in-IPv6 packets.

- Multicast B4 (mB4): is a functional entity embedded in a CPE, which supports an IGMP-MLD interworking function (refer to Section 6.1) that relays information conveyed in IGMP messages by forwarding the corresponding MLD messages towards the MLD Querier in the IPv6 network. In addition, the mB4 decapsulates IPv4-in-IPv6 multicast packets.

- PIMv4: refers to PIM when deployed in an IPv4 infrastructure (i.e., IPv4 transfer capabilities are used to exchange PIM messages).

- PIMv6: refers to PIM when deployed in an IPv6 infrastructure (i.e., IPv6 transfer capabilities are used to exchange PIM messages).

## 3. Scope

This document focuses only on issues raised by a DS-Lite environment: subscription to an IPv4 multicast group and the delivery of IPv4-formatted content to IPv4 receivers over an IPv6-only network. In particular, only the following case is covered:

> An IPv4 receiver accesses IPv4 multicast contents over an IPv6-only multicast-enabled network.

This document does not cover the source/receiver heuristics, where as IPv4 receiver can also behave as an IPv4 multicast source. This document assumes that hosts behind the mB4 are IPv4 multicast receivers only.

## 4. Solution Overview

In the original DS-Lite specification [RFC6333], an IPv4-in-IPv6 tunnel is used to carry bidirectional IPv4 unicast traffic between a B4 and an AFTR. An extension to DS-Lite is proposed in this document which specifies an IPv4-in-IPv6 encapsulation scheme to deliver unidirectional IPv4 multicast traffic from a mAFTR to a mB4.

An overview of the solution is provided in this section which is intended as an introduction to how it works, but is NOT normative. For the normative specifications of the two new functional elements: mB4 and mAFTR (Figure 1), refer to Section 6 and Section 7.

```
                             ------------
                           /              \
                          |  IPv4 network  |
                           \              /
                             ------------
               IPv4 multicast  :   |   ^  PIMv4 Join
                               v   |   :
                            +-------------+
                            |    mAFTR    |
                            +-------------+
              IPv6 multicast  |:|  |   ^  PIMv6 Join (PIMv6
              (IPv4 embedded) |:|  |   :   routers in between)
                             ------------
                           /              \
                          |  IPv6 network  |
                           \              /
                             ------------
                              |:|  |   :  MLD Report
                              |v|  |   :
                             +-----------+
                             |    mB4    |
                             +-----------+
               IPv4 multicast  :   |   ^  IGMP Report
                               v   |   :
                             +-----------+
                             |   IPv4    |
                             | receiver  |
                             +-----------+
```

Figure 1: Functional Architecture

### 4.1. IPv4-embedded IPv6 Prefixes

In order to map the addresses of IPv4 multicast traffic with IPv6 multicast addresses, an IPv6 multicast prefix (mPrefix64) and an IPv6 unicast prefix (uPrefix64) are provided to mAFTR and mB4 elements, both of which contribute to the computation and the maintenance of the IPv6 multicast distribution tree that extends the IPv4 multicast distribution tree into the IPv6 multicast network.

The mAFTR and mB4 use mPrefix64 to convert an IPv4 multicast address (G4) to an IPv4-embedded IPv6 multicast address (G6). The mAFTR and mB4 use uPrefix64 to convert an IPv4 multicast source address (S4) to an IPv4-embedded IPv6 address (S6). The mAFTR and mB4 MUST use the same mPrefix64 and uPrefix64, as well as run the same algorithm for building IPv4-embedded IPv6 addresses. Refer to Section 5 for more details about the address mapping.

## 4.2. Multicast Distribution Tree Computation

When an IPv4 receiver connected to the CPE that embeds mB4 wants to subscribe to an IPv4 multicast group, it sends an IGMP Report message to the mB4. The mB4 creates the IPv6 multicast group (G6) address using mPrefix64 and the original IPv4 multicast gorup address. If the receiver sends a source-specific IGMPv3 Report message, the mB4 will create the IPv6 source address (S6) using uPrefix64 and the original IPv4 source address.

The mB4 uses the G6 (and both S6 and G6 in SSM) to create the corresponding MLD Report message. The mB4 sends the Report message to the MLD Querier in the IPv6 network. The MLD Querier (typically acts as the PIMv6 Designated Router) receives the MLD Report message and sends the PIMv6 Join to join the IPv6 multicast distribution tree. The MLD Querier can send either PIMv6 Join (*,G6) in ASM or PIMv6 Join (S6,G6) in SSM to the mAFTR.

The mAFTR acts as the DR to which the uPrefix64-derived S6 is connected. The mAFTR will receive the source-specific PIMv6 Join message (S6,G6) from the IPv6 multicast network. If the mAFTR is the Rendezvous Point (RP) of G6, it will receive the any-source PIMv6 Join message (*,G6) from the IPv6 multicast network. If the mAFTR is not the RP of G6, it will send the PIM Register message to the RP of G6 located in the IPv6 multicast network.

When the mAFTR receives the PIMv6 Join message (*,G6), it will extract the IPv4 multicast group address (G4). If the mAFTR is the RP of G4 in the IPv4 multicast network, it will create a (*,G4) entry (if there is not yet an existing one) in its own IPv4 multicast routing table. If the mAFTR is not the RP of G4, it will send the corresponding PIMv4 Join message (*,G4) towards the RP of G4 in the IPv4 multicast network.

When the mAFTR receives the PIMv6 Join message (S6,G6), it will extract the IPv4 multicast group address (G4) and IPv4 source address (S4) and send the corresponding (S4,G4) PIMv4 Join message directly to the IPv4 source.

A branch of the multicast distribution tree is then grafted, comprising both an IPv4 part (from the mAFTR upstream) and an IPv6 part (from mAFTR downstream to the mB4).

The mAFTR MUST advertise the route of uPrefix64 with an IPv6 IGP, so as to represent the IPv4-embedded IPv6 source in the IPv6 multicast network.

### 4.3. Multicast Data Forwarding

When the mAFTR receives an IPv4 multicast packet, it will encapsulate the packet into an IPv6 multicast packet using the IPv4-embedded IPv6 multicast address as the destination address and an IPv4-embedded IPv6 unicast address as the source address. The encapsulated IPv6 multicast packet will be forwarded down the IPv6 multicast distribution tree and the mB4 will eventually receive the packet.

The IPv6 multicast network treats the IPv4-in-IPv6 encapsulated multicast packets as native. The IPv6 multicast routers use the outer IPv6 header to make forwarding decisions.

When the mB4 receive the IPv6 multicast packet (to G6) derived by mPrefix64, it MUST decapsulate it and forward the original IPv4 multicast packet to the receivers subscribing to G4.

## 5. Address Mapping

### 5.1. Prefix Assignment

A dedicated IPv6 multicast prefix (mPrefix64) is provisioned to the mAFTR and the mB4. The mAFTR and the mB4 use the mPrefix64 to form an IPv6 multicast group address from an IPv4 multicast group address. The mPrefix64 can be of two types: ASM_mPrefix64 (a mPrefix64 used in ASM mode) or SSM_mPrefix64 (a mPrefix64 used in SSM mode). The mPrefix64 Must be derived from the corresponding IPv6 multicast address space (e.g., the SSM_mPrefix64 MUST be in the range of multicast address space specified in [RFC4607]).

The IPv6 part of the multicast distribution tree can be seen as an extension of the IPv4 part of the multicast distribution tree. The IPv4 multicast source address MUST be mapped to an IPv6 multicast source address. An IPv6 unicast prefix (uPrefix64) is provisioned to the mAFTR and the mB4. The mAFTR and the mB4 use the uPrefix64 to form an IPv6 multicast source address from an IPv4 multicast source address. The uPrefix-formed IPv6 multicast source address will represent the original IPv4 multicast source in the IPv6 multicast network. The uPrefix64 MUST be derived from the IPv6 unicast address space.

The address format to be used is left to the responsibility of the network provider. The address synthesizing MUST follow [I-D.boucadair-behave-64-multicast-address-format] and [RFC6052].

The mPrefix64 and uPrefix64 can be configured in the mB4 using a variety of methods, including an out-of-band mechanism, manual configuration, or a dedicated provisioning protocol (e.g., using DHCPv6 [I-D.qin-softwire-multicast-prefix-option]).

### 5.2. Examples

Group address mapping example when a /96 is used:

| mPrefix64 | IPv4 address | IPv4-Embedded IPv6 address |
| --- | --- | --- |
| ffxx:abc::/96 | 230.1.2.3 | ffxx:abc::230.1.2.3 |

Source address mapping example when a /96 is used:

| uPrefix64 | IPv4 address | IPv4-Embedded IPv6 address |
| --- | --- | --- |
| 2001:db8::/96 | 192.1.2.3 | 2001:db8::192.1.2.3 |

## 6. Multicast B4 (mB4)

### 6.1. IGMP-MLD Interworking Function

The IGMP-MLD Interworking Function combines the IGMP/MLD Proxying function and the IGMP/MLD translation function. The IGMP/MLD Proxying function is specified in [RFC4605]. The IGMP/MLD translation function translates the contents of IGMP messages into MLD messages by using a stateless algorithm. The address synthesizing MUST comply with the rules documented in Section 5. MLD messages will be forwarded natively towards the MLD Querier located upstream in the IPv6 network. The mB4 performs the IGMP-MLD Interworking Function to relay between the IGMP messages and the MLD messages.

```
+----------+   IGMP  +-------+   MLD   +---------+
|   IPv4   |---------|  mB4  |---------|   MLD   |
| Receiver |         |       |         | Querier |
+----------+         +-------+         +---------+
```

Figure 2: IGMP-MLD Interworking

When the mB4 receives an IGMP Report message from a receiver to subscribe to multicast group (and optionally associated to a source in SSM mode), it MUST translate the IGMP Report message into a MLD Report message and send to the MLD Querier. The mB4 MUST construct the IPv6 multicast group address using the mPrefix64.

When the mB4 receives an MLD Listener Query message from the MLD Querier, it MUST convert the MLD listener Query message to the IGMP Query message and send it to the IPv4 receiver(s). The mB4 MUST retrieve the IPv4 multicast group address using the mPrefix64.

If SSM is deployed, the mB4 MUST construct the IPv6 source address (or retrieve the IPv4 source address) using the uPrefix64. The mB4 may create a membership database which associates the IPv4-IPv6 multicast groups with the interfaces (e.g., Wi-Fi and Wired Ethernet) facing IPv4 multicast receivers.

### 6.2. Multicast Data Forwarding

When the mB4 receives an IPv6 multicast packet, it MUST check the group address and the source address. If the IPv6 multicast group prefix is mPrefix64 and the IPv6 source prefix is uPrefix64, the mB4 MUST de-capsulate the IPv6 header and forward the IPv4 multicast packet through each relevant interface. Otherwise, the mB4 MUST drop the packet silently.

As an illustration, if a packet is received from source 2001:db8::192.1.2.3 and to be forwarded to group ffxx:abc::230.1.2.3, the mB4 will de-capsulate it into an IPv4 multicast packet using 192.1.2.3 as the IPv4 multicast source address and using 230.1.2.3 as the IPv4 destination address.

### 6.3. Fragmentation

Encapsulating IPv4 multicast packets into IPv6 multicast packets that will be forwarded by the mAFTR to the mB4 along the IPv6 multicast distribution tree reduces the effective MTU size by the size of an IPv6 header. In this specification, the data flow is unidirection from mAFTR to mB4, the mAFTR must fragment the oversized IPv6 packet after the encapsulation into two IPv6 packets. The mB4 MUST reassemable the IPv6 packets, decapsulate the IPv6 packet, and forward the IPv4 packet to the hosts subscribing the multicast group. Further considerations about fragmentation issues are documented in [RFC6333].

### 6.4. Host built-in mB4 Function

If the mB4 function is implemented in the host which is directly connected to an IPv6-only network. If an IPv4 application running in the host requests to subscribe an IPv4 multicast stream, the host MUST implement Section 6.1, Section 6.2, and Section 6.3. The host MAY optimize the implemntation to provide an Application Programming

Interface (API) or kernel module to skip the IGMP-MLD Interworking Function. The optimization is out of scope of the specification.

## 7. Multicast AFTR (mAFTR)

### 7.1. Routing Considerations

The mAFTR is responsible for interconnecting the IPv4 multicast distribution tree with the corresponding IPv6 multicast distribution tree. The mAFTR MUST use the uPrefix64 to build the IPv6 source addresses of the multicast group address derived from mPrefix64. In other words, the mAFTR MUST be the multicast source derived from mPrefix64.

The mAFTR MUST advertise the route of uPrefix64 to the IPv6 IGP. This is needed for the IPv6 multicast router to have routing information to discover the source. In order to pass the Reverse Path Forwarding (RPF) check, the IPv6 routers MUST enable PIM on the interfaces which has the shortest path to the uPrefix64.

### 7.2. Processing PIM Message

The mAFTR MUST interwork PIM Join/Prune messages for (*, G6) and (S6, G6) on their corresponding (*, G4) and (S4, G4). The following text specifies the expected behavior of mAFTR for PIM Join message.

```
                  +---------+
         ---------|  mAFTR  |---------
           PIMv6  |uPrefix64|  PIMv4
                  |mPreifx64|
                  +---------+
```

Figure 3: PIMv6-PIMv4 Interworking Function

The mAFTR contains two separate multicast routing table (mRIB): IPv4 multicast routing table (mRIB4) and IPv6 multicast routing table (mRIB6), which are bridged by one IPv4-in-IPv6 virtual interface. It should be noted that the implementations may vary (e.g., using one integrated mRIB without any virtual interface), while they should follow the specification herein for the consistency of overall functionality.

When a mAFTR receives a PIMv6 Join message (*,G6) with an IPv6 multicast group address (G6) that is derived from the mPrefix64, it MUST check its IPv6 multicast routing table (mRIB6). If there is an entry for this G6, it MUST check whether the interface through which the PIMv6 Join message has been received is on the outgoing interface list. If not, the mAFTR MUST add the interface to the outgoing interface list. If there is no entry in the mRIB6, the mAFTR MUST create a new entry (*,G6) for the multicast group. While, whether or not to set the IPv4-in-IPv6 virtual interface as the incoming interface of the newly created entry is up to the implementation but should comply with the mAFTR's behavior of multicast data forwarding, see Section 7.4.

The mAFTR MUST extract the IPv4 multicast group address (G4) from the IPv4-embedded IPv6 multicast address (G6) contained in the PIMv6 Join message. The mAFTR MUST check its IPv4 multicast routing table (mRIB4). If there is an entry for G4, it MUST check whether the IPv4-in-IPv6 virtual interface is on the outgoing interface list. If not, the mAFTR MUST add the interface to the outgoing interface list. If there is no entry for G4, the mAFTR MUST create a new (*,G4) entry in its mRIB4 and initiate the procedure for building the shared tree in the IPv4 multicast network without any additional requirement.

If mAFTR receives a source-specific Join message, the (S6, G6) will be processed rather than (*,G6). The procedures of processing (S6,G6) and (*,G6) are almost the same. Differences have been detailed in [RFC4601].

### 7.3. Switching from Shared Tree to Shortest Path Tree

When the mAFTR receives the first IPv4 multicast packet, it may extract the multicast source address (S4) from the packet and send an Explicit PIMv4 (S4,G4) Join message directly to S4. The mAFTR will switch from the shared Rendezvous Point Tree (RPT) to the Shortest Path Tree (SPT) for G4.

For IPv6 multicast routers to switch to the SPT, there is no new requirement. IPv6 multicast routers may send an Explicit PIMv6 Join to mAFTR once the first (S6,G6) multicast packet arrives from upstream multicast routers.

### 7.4. Multicast Data Forwarding

When the mAFTR receives an IPv4 multicast packet, it will look up the mRIB4 to find a matching entry and then forward the packet to the interface(s) on the outgoing interface list. If the IPv4-in-IPv6 virtual interface also belongs to this list, the packet will be encapsulated with the mPrefix64-derived and uPrefix64-derived IPv4-embedded IPv6 addresses to form an IPv6 multicast packet. Then another lookup is executed to find a matching entry in the mRIB6, while whether or not to perform RPF check for the second lookup is up to the implementation and is out of the scope of this document. The

IPv6 multicast packet is forwarded along the IPv6 multicast distribution tree, based upon the outgoing interface list of the matching entry in the mRIB6.

As an illustration, if a packet is received from source 192.1.2.3 and to be forwarded to group 230.1.2.3, the mAFTR encapsulates it into an IPv6 multicast packet using ffxx:abc::230.1.2.3 as the IPv6 destination address and using 2001:db8::192.1.2.3 as the IPv6 multicast source address.

### 7.5. TTL/Scope

The Scope field of IPv4-in-IPv6 multicast addresses can be valued to "E" (Global scope) or to "8" (Organization-local scope). This specification does not discuss the scope value that should be used.

## 8. Security Considerations

This document does not introduce any new security concern in addition to what is discussed in Section 5 of [RFC6052], Section 10 of [RFC3810] and Section 6 of [RFC4601].

### 8.1. Firewall Configuration

The CPE should be configured to accept incoming MLD messages and traffic forwarded to multicast groups subscribed by receivers located in the customer premises.

## 9. Acknowledgements

The authors would like to thank Dan Wing for his guidance in the early discussions which initiated this work. We also thank Peng Sun, Jie Hu, Qiong Sun, Lizhong Jin, Alain Durand, Dean Cheng, Behcet Sarikaya, Tina Tsou, Rajiv Asati, and Xiaohong Deng for their valuable comments.

## 10. IANA Considerations

This document includes no request to IANA.

## 11. References

### 11.1. Normative References

[I-D.boucadair-behave-64-multicast-address-format]
Boucadair, M., Qin, J., Lee, Y., Venaas, S., Li, X., and M. Xu, "IPv4-Embedded IPv6 Multicast Address Format", draft-boucadair-behave-64-multicast-address-format-03 (work in progress), October 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

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

[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.

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

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

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

### 11.2. Informative References

[I-D.ietf-mboned-multiaaa-framework]
Satou, H., Ohta, H., Hayashi, T., Jacquenet, C., and H. He, "AAA and Admission Control Framework for Multicasting", draft-ietf-mboned-multiaaa-framework-12 (work in progress), August 2010.

[I-D.jaclee-behave-v4v6-mcast-ps]
Jacquenet, C., Boucadair, M., Lee, Y., Qin, J., and T. ZOU), "IPv4-IPv6 Multicast: Problem Statement and Use Cases", draft-jaclee-behave-v4v6-mcast-ps-02 (work in progress), June 2011.

[I-D.qin-softwire-multicast-prefix-option]
Qin, J., Boucadair, M., and T. Tsou, "DHCPv6 Options for IPv6 DS-Lite Multicast Prefix", draft-qin-softwire-multicast-prefix-option-01 (work in progress), October 2011.

[RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.

[RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, August 2006.

[RFC4608] Meyer, D., Rockell, R., and G. Shepherd, "Source-Specific Protocol Independent Multicast in 232/8", BCP 120, RFC 4608, August 2006.

## Appendix A. Use Case: IPTV

IPTV generally includes two categories of service offerings:

- o Video on Demand (VoD) that unicast video content to receivers.

- o Multicast live TV broadcast services.

Two players intervene in the delivery of this service:

- o Content Providers, who usually own the contents that is multicast to receivers. Content providers may contractually define an agreement with network providers to deliver contents to receivers.

- o Network Providers, who provide network connectivity services (e.g., network providers are responsible for carrying multicast flows from head-ends to receivers). Refer to [I-D.ietf-mboned-multiaaa-framework].

Note that some contract agreements prevent a network provider from altering the content as sent by the content provider for various reasons. Under the contract, multicast streams should be delivered unaltered to the requesting users.

Many current IPTV contents are likely to remain IPv4-formatted and out of control of the network providers. Additionally, there are numerous legacy receivers (e.g., IPv4-only Set Top Boxes (STB)) that can't be upgraded or be easily replaced to support IPv6. As a consequence, IPv4 service continuity MUST be guaranteed during the transition period, including the delivery of multicast services such as Live TV Broadcasting to users.

## Appendix B. Deployment Considerations

### B.1. Load-Balancing

For robustness and load distribution purposes, several nodes in the network can embed the mAFTR function. In such case, the same IPv6 prefixes (i.e., mPrefix64 and uPrefix64) and algorithm to build IPv4-embedded IPv6 addresses MUST be configured on those nodes.

### B.2. RP for IPv4-Embedded IPv6 Multicast Groups

For the sake of simplicity, it is RECOMMENDED to configure mAFTR as the RP for the IPv4-embedded IPv6 multicast groups it manages. No registration procedure is required under this configuration.

### B.3. mAFTR Policy Configuration

mAFTR may be configured with a list of IPv4 multicast groups and sources. Only multicast flows bound to the configured addresses should be handled by the mAFTR. Otherwise, packets are silently drooped.

### B.4. Static vs. Dynamic PIM Triggering

To optimize the usage of network resources in current deployments, all multicast streams are conveyed in the core network while only popular ones are continuously conveyed in the aggregation/access network (static mode). Non-popular streams are conveyed in the access network upon request (dynamic mode). Depending on the location of the mAFTR in the network, two modes can be envisaged: static and dynamic.

- o Static Mode: the mAFTR is configured to instantiate permanent (S6, G6) and (*, G6) entries in its MRIBv6 using a pre-configured (S4, G4) list.
- o Dynamic Mode: the instantiation and deletion of (S6, g6) or (*, G6) is triggered by the receipt of PIMv6 messages.

Authors' Addresses

Qian Wang
China Telecom
No.118, Xizhimennei
Beijing, 100035
China

Phone: +86 10 5855 2177
Email: wangqian@ctbri.com.cn

Jacni Qin
ZTE
Shanghai,
China

Phone: +86 1391 8619 913
Email: jacni@jacni.com

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Phone:
Email: mohamed.boucadair@orange.com

Christian Jacquenet
France Telecom
Rennes,   35000
France

Phone:
Email: christian.jacquenet@orange.com

Yiu L. Lee
Comcast
U.S.A.

Phone:
Email: yiu_lee@cable.comcast.com
URI:   http://www.comcast.com