BEHAVE Working Group
Internet-Draft
Intended status: Informational
Expires: May 5, 2012

C. Jacquenet
M. Boucadair
France Telecom
Y. Lee
Comcast
J. Qin
ZTE
T. Tsou
Huawei Technologies (USA)

November 02, 2011

IPv4-IPv6 Multicast: Problem Statement and Use Cases
draft-jaclee-behave-v4v6-mcast-ps-03

Abstract

This document discusses issues and requirements raised by IPv4-IPv6 multicast interconnection and co-existence scenarios. It also discusses various multicast use cases which may occur during IPv6 transitioning.

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].

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

In current deployments, the IP multicast forwarding scheme is used by many service providers to deliver some services, such as live TV broadcasting services. Multiple players intervene in the delivery of these services, including content and service providers. Service providers are responsible for carrying multicast flows from head-ends to receivers. The content can be supplied by a service provider or by other providers (e.g., case of externally paid channels).

Transition to IPv6 raises issues and corresponding requirements. In particular, IPv4 service continuity is an essential requirement from a business perspective. This specifically includes continued receiver access to IPv4-formatted contents even when the assignment of a dedicated global IPv4 address to the receiver is no longer possible and even after the receivers have migrated to IPv6. Likewise, the delivery of IPv6-formatted contents to IPv4 receivers must also be possible.

Finally, in cases where the underlying transport network is of a different address family from that of the source and/or receivers, the delivery of multicast data must still be guaranteed. For example, in DS-Lite environments, the (access) network is IPv6-enabled, but both multicast sources and receivers are likely to remain IPv4-only.

This document does not make any assumption on the techniques used for the delivery of multicast traffic (e.g., native IP multicast with or without traffic isolation features, etc.).

This document further elaborates on the context and discusses multicast-related issues and requirements.

1.1. Goals

The objective of this document is to clarify the problem space. In particular, this document further elaborates on the following issues:

- What are the hurdles encountered for the delivery of multicast-based service offerings when both IPv4 and IPv6 co-exist?

- What standardization effort is needed: are there any missing function and protocol extensions?

- Does the work on multicast transition have to cover both encapsulation and translation schemes, considering the requirement of multicast network performance among others?

1.2. Terminology

This document uses the following terms:

- Multicast Source: Source, in short
- Multicast Receiver: Receiver, in short, e.g., STB (Set Top Box)
- Multicast Delivery Network: Network in short, covers the realm from Designated Routers that are directly connected to sources to IGMP/MLD (Internet Group Management Protocol/Multicast Listener Discovery) Querier devices that process IGMP/MLD signalling traffic exchanged with receivers.

1.3. Organization of the Document

This document is organized as follows:

- Section 2 details basic requirements that should be addressed by providers involved in the delivery of multicast-based services during the transition period,
- Section 3 discusses several use cases that reflect issues raised by the forthcoming transition period,
- Section 4 details design considerations,
- Section 5 summarizes the standardization effort that should be tackled by the IETF.

2. Discussion and Service Requirements

2.1. Scope

Intra-domain only: The delivery of multicast services such as live TV broadcasting often relies upon walled garden designs that restrict the scope to the domain where such services can be subscribed. As a consequence, considerations about inter-domain multicast are out of the scope of this document.

Multicast-enabled networks only: This document assumes that the network is IP multicast-enabled. That is, whatever the IP address family of the content, the latter will be multicast along distribution trees that should be terminated as close to the receivers as possible for the sake of bandwidth optimization. In other words, considerations about forwarding multicast traffic over unicast-only (access) networks is out of the scope of this

document.

Multicast to the receivers, not from the receivers: This document only covers the case where multicast traffic is forwarded by the service provider network to the receivers. This document does not cover the case where the receivers send multicast traffic to the network.

2.2. Issues Raised by the Transition Period

Global IPv4 address depletion inevitably challenges service providers who must guarantee IPv4 service continuity during the forthcoming transition period. In particular, access to IPv4 contents that are multicast to IPv4 receivers becomes an issue when the forwarding of multicast data assumes the use of global IPv4 addresses.

The rarefaction of global IPv4 addresses may indeed affect the multicast delivery of IPv4-formatted contents to IPv4 receivers. For example, the observed evolution of ADSL broadband access infrastructures from a service-specific, multi-PVC (Permanent Virtual Circuit) scheme towards a "service-agnostic", single PVC scheme, assumes the allocation of a globally unique IPv4 address on the WAN (Wide Area Network) interface of the CPE (Customer Premises Equipment), or to a mobile terminal), whatever the number and the nature of the services the customer has subscribed to.

Likewise, the global IPv4 address depletion encourages the development of IPv6 receivers while contents may very well remain IPv4-formatted. There is therefore a need to make sure such IPv6 receivers can access IPv4-formatted contents during the transition period.

During the transition period, the usage of the remaining global IPv4 address blocks will have to be rationalized for the sake of IPv4 service continuity. The current state-of-the-art suggests the introduction of NAT (Network Address Translation) capabilities (generally denoted as CGN, for Carrier-Grade NAT) in providers' networks, so that a global IPv4 address will be shared between several customers. As a consequence, CPE or mobile UE (User Equipment) devices will no longer be assigned a dedicated global IPv4 address anymore, and IPv4 traffic will be privately-addressed until it reaches one of the CGN capabilities deployed in the network.

From a multicast delivery standpoint, this situation suggests the following considerations:

- o The current design of some multicast-based services like TV broadcasting often relies upon the use of a private IPv4 addressing scheme because of a walled garden approach. Privately-addressed IGMP [RFC2236] [RFC3376] traffic sent by IPv4 receivers is generally forwarded over a specific (e.g. "IPTV") PVC towards an IGMP Querier located in the access infrastructure, e.g.- in some deployments it is hosted by a BRAS (BRoadband Access Server) device that is the PPP (Point-to-Point Protocol) session endpoint and which may also act as a PIM DR (Protocol Independent Multicast Designated Router)[RFC4601] router. This design does not suffer from global IPv4 address depletion by definition (since multicast traffic relies upon the use of a private IPv4 addressing scheme), but it is inconsistent with migrating the access infrastructure towards a publicly-addressed single PVC design scheme.

- o Likewise, other deployments (e.g., cable operators' environments) rely upon the public CPE's address for multicast delivery and will therefore suffer from IPv4 address depletion.

- o The progressive introduction of IPv6 as the only perennial solution to global IPv4 address depletion does not necessarily assume that multicast-based IPv4 services will be migrated accordingly. Access to IPv4 multicast contents when no global IPv4 address can be assigned to a customer anymore raises several issues: (1) The completion of the IGMP-based multicast group subscription procedure, (2) The computation of the IPv4 multicast distribution tree, and (3) The IPv4-inferred addressing scheme to be used in a networking environment which will progressively become IPv6-enabled.

2.3. Service Requirements

Given the issues highlighted in Section 2.2, the delivery of multicast contents during the forthcoming transition period needs to address the following requirements. Note that some of these requirements are not necessarily specific to the IPv4-to-IPv6 transition context, but rather apply to a wide range of multicast-based services whatever the environmental constraints.

But the forthcoming transition period further stresses these requirements (see Section 4.4.1 for more details).

- o Service_REQ-1: Optimize bandwidth. Contents SHOULD NOT be multicast twice (using both versions of IP) for the sake of bandwidth optimization. Injecting multicast content using both IPv4 and IPv6 raises some dimensioning issues that should be carefully evaluated by service providers during network planning operations. For instance, if only few IPv6- enabled receivers are

in use, it can be more convenient to convey multicast traffic over IPv4 rather than doubling the consumed resources for few users. IPv4/IPv6 co-existence solutions SHOULD be designed to optimize network resource utilization.

- Service_REQ-2: Zap rapidly. The time it takes to switch from one content to another MUST be as short as possible. For example, zapping times between two TV channels should be in the magnitude of a few seconds at most, whatever the conditions to access the multicast network. A procedure called "IGMP fast-leave" is sometimes used to minimize this issue so that the corresponding multicast stream is stopped as soon as the IGMP Leave message is received by the Querier. In current deployments, IGMP fast-leave often assumes the activation of the IGMP Proxy function in DSLAMs. The complexity of such design is aggravated within a context where IPv4 multicast control messages are encapsulated in IPv6.

- Service_REQ-3: Preserve the integrity of contents. Some contract agreements may prevent a service provider from altering the content owned by the content provider, because of copyright, confidentiality and SLA assurance reasons. Multicast streams SHOULD be delivered without altering their content.

- Service_REQ-4: Preserve service quality. Crossing a CGN or performing multicast packet encapsulation may lead to fragmentation or extra delays and may therefore impact the perceived quality of service. Such degradation MUST be avoided.

- Service_REQ-5: Optimize IPv4-IPv6 inter-working design. In some operational networks, a source-based stateful NAT function is sometimes used for load balancing purposes, for example. Because of the operational issues raised by such a stateful design, the deployment of stateless IPv4-IPv6 interworking functions SHOULD be privileged.

3. Use Cases

During the forthcoming IPv4-to-IPv6 transition period, there might be a mix of multicast receivers, sources, and networks running in different address families. However, service providers must guarantee the delivery of multicast services to IPv4 receivers and, presumably, IPv6 receivers. Because of the inevitable combination of different IP version-related environments (sources, receivers and networks), service providers should carefully plan and choose the right transition technique that will optimize the network resources to deliver multicast-based services.

Concretely, several use cases can be considered during the IPv4/ IPv6 co-existence period. Some of them are depicted in the following sub-sections.

3.1. IPv4 Receiver and Source Connected to an IPv6-Only Network

We refer to this scenario as 4-6-4. An example of such use case is a DS-Lite environment, where customers will share global IPv4 addresses. IPv4 receivers are connected to CPE devices that are provisioned with an IPv6 prefix only. Delivering multicast content sent by an IPv4 source to such receivers requires the activation of some adaptation functions (AFs). These may operate at the network layer (interworking functions (IWF) or at the application layer (application level gateways (ALGs)).

The signalling flow for the 4-6-4 use case is shown in Figure 1.

```
      +------+      +-----+       +------+         +------+
      | Host | Sig1 | DS  |       |  MR  |   PIM   |  MR  |
      | Rcvr |------| AF1 |       |      | . . . . |      |
      |      | IPv4 |     |       | (BG) |   IPv4  | (DR) |
      +------+      +-----+       +------+         +------+
                     /               \
                MLD / IPv6        PIM \ IPv4
                   /                   \
            +------+      +----+       +------+
            |  MR  |  PIM |    | PIM   |  DS  |
            |      |------| MR | . . . |  AF2 |
            | (DR) | IPv6 |    | IPv6  | (BG) |
            +------+      +----+       +------+

            -------------------------------------->

      Rcvr: Multicast receiver
      DS  : Dual Stack
      AF  : Adaptation Function (ALG or IWF)
      MR  : Multicast Router
      DR  : Designated Router
      BG  : Border Gateway
```

Figure 1: Signalling Path for the 4-6-4 Scenario.

Sig1 denotes the signalling protocol used by the host. If the adaptation function AF1 to which it sends this signalling is an ALG, Sig1 will be an application-layer protocol such as HTTP or SIP. If the adaptation function is an interworking function, Sig1 will be IGMP. If the adaptation function is collocated with the multicast router at the edge of the IPv6 network, the intermediate MLD step can

be avoided.

Another dual stack adaptation function AF2 is needed at the border between the IPv6 multicast domain and the IPv4 multicast domain where the source resides. This device acts as a multicast router either terminating or interworking PIM signalling in the IPv6 network directed toward the source, depending on whether it is acting as an ALG or as an IWF.

On the IPv4 side, AF2 also acts as a multicast router, and uses PIMv4 signalling to join the IPv4 multicast group. If AF2 is acting as an ALG, the PIMv4 signalling is triggered by application-level signalling or management action. If AF2 is acting as an interworking function, the PIMv4 signalling is triggered by the arrival of PIMv6 signalling directed toward the source.

The return path taken by the multicast content is shown in Figure 2.

```
   +------+      +-----+     +------+       +------+      +-----+
   | Host |      | DS  |     |  MR  |       |  MR  |      |     |
   | Rcvr |------| AF1 |     |      | . . . |      |------| Src |
   |      | IPv4 |     |     | (BG) | IPv4  | (DR) | IPv4 |     |
   +------+      +-----+     +------+       +------+      +-----+
                  /               \
                 / IPv6            \ IPv4
                /                   \
         +------+      +----+       +------+
         |  MR  |      |    |       |  DS  |
         |      |------| MR | . . . |  AF2 |
         | (DR) | IPv6 |    | IPv6  | (BG) |
         +------+      +----+       +------+

             <-------------------------------------

   Rcvr: Multicast receiver
   DS  : Dual Stack
   AF  : Adaptation Function
   MR  : Multicast router
   DR  : Designated Router
   BG  : Border Gateway
   Src : Multicast source
```

Figure 2: Multicast Content Distribution Path for the 4-6-4 Scenario.

Again, adaptation functions are needed whenever the IP protocol version changes. The adaptation function instance AF2 at the boundary between the source network and the IPv6 network may either encapsulate or translate the headers of the IPv4 packets to allow the

content to cross the IPv6 network - note that encapsulation requires knowledge that the receiver is IPv4. The adaptation function instance at the boundary between the IPv6 network and the receiver network performs the reverse operation to deliver IPv4 packets.

Given the current state-of-the-art where multicast content is likely to remain IPv4-formatted while receiver devices such as Set Top Boxes will also remain IPv4-only for quite some time, this scenario is prioritized by some service providers, including those that are deploying or will deploy DS-Lite CGN capabilities for the sake of IPv4 service continuity.

3.2. IPv6 Receiver Connected to an IPv4 Source Through an IPv4 Multicast-Disabled Access Network and an IPv6 Multicast Network

One major provider faces a complex transitional situation where the receiver is IPv6, the CPE router is dual stack but is provided by the customer, and the IPv4 access network is not multicast capable. The IPv4 access network connects to an IPv6 network that is multicast capable, and which in turn connects to IPv4 sources.

This scenario is denoted as the 6-4-6-4 scenario.

Because the provider does not manage the CPE router, encapsulation of IPv6 packets across the IPv4 network is unlikely. Figure 3 shows the signalling path for this scenario.

```
      +------+      +-----+      +----+      +------+
      | Host | Sig1 | DS  | Sig1 |    | Sig1 |  DS  |
      | Rcvr |------| AF0 |------| PE | . . .|  AF1 |
      |      | IPv6 |(CPE)| IPv4 |    | IPv4 | (BG) |
      +------+      +-----+      +----+      +------+
                                                /
                    -----------------<----------
                   / MLD over IPv6
            +------+      +----+       +------+
            |  MR  |  PIM |    |  PIM  |  DS  |
            |      |------| MR | . . . |  AF2 |
            | (DR) | IPv6 |    | IPv6  | (BG) |
            +------+      +----+       +------+
                                          /
                 -----------------<-------
                / PIM over IPv4
             +------+         +------+
             |  MR  |    PIM  |  MR  |
             |      | . . . . |      |
             | (BG) |   IPv4  | (DR) |
             +------+         +------+

                -------------------------------------->
```

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
CPE : Customer Premises Equipement (Dual Stack router)
PE : Provider Edge router
BG : Border Gateway

Figure 3: Signalling Path For the 6-4-6-4 Scenario.

The major challenge of this scenario is how to ensure that signalling packets from the CPE (AF0) reach the adaptation function instance at the boundary between the IPv4 multicast-disabled access network and the IPv6 multicast network (AF1). If the signalling Sig1 from the receiver is MLD, the CPE router has to translate the MLD destination address ff02::16 into the address of AF1.

This requires some sort of configuration by the provider. Alternatively, Sig1 could be an application-layer protocol. in that case, the CPE router can use DNS to get the address of AF1. The adaptation function AF2 between the IPv6 multicast network and the IPv4 network where the multicast source is connected is similar to AF2 in the 4-6-4 scenario.

Figure 4 shows the path taken by multicast content flowing from the source to the receiver. Again, AF2 can either encapsulate or translate the headers of the incoming packets. AF1 performs the reverse action, and forwards unencapsulated IPv4 packets towards AF0. AF0 then performs header translation to convert the incoming packets into IPv6 multicast packets before sending them on to the receiver.

```
+------+      +-----+      +----+      +------+
| Host |      | DS  |      |    |      |  DS  |
| Rcvr |------| AF0 |------| PE | . . .|  AF1 |
|      | IPv6 |(CPE)| IPv4 |    | IPv4 | (BG) |
+------+      +-----+      +----+      +------+
                                          /
              ----------------->----------
             /         IPv6
      +------+      +----+       +------+
      |  MR  |      |    |       |  DS  |
      |      |------| MR | . . . |  AF2 |
      | (DR) | IPv6 |    | IPv6  | (BG) |
      +------+      +----+       +------+
                                    /
           ----------------->-------
          /        IPv4
       +------+         +------+      +-----+
       |  MR  |         |  MR  |      |     |
       |      | . . . . |      |------| Src |
       | (BG) |   IPv4  | (DR) | IPv4 |     |
       +------+         +------+      +-----+

          <-------------------------------------
```

Rcvr: Multicast receiver
Src : Multicast source
DS : Dual Stack
AF : Adaptation function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
CPE : Customer Premises Equipment (Dual Stack router)
PE : Provider Edge router
BG : Border Gateway

Figure 4: Multicast Content Distribution Path For the 6-4-6-4 Scenario.

3.3. IPv6 Receiver and Source Connected to an IPv4-Only Network

We refer to this scenario as 6-4-6. According to a BEHAVE WG consensus when elaborating the transition unicast scenarios, servers are likely to remain IPv4-enabled in a first stage. This is also

true for multicast. Additionally, content providers who own the content may not be ready for IPv6 migration for some reason. Therefore, the content is likely to remain IPv4-formatted.

As a consequence, this 6-4-6 scenario is of lower priority than the 4-6-4 scenario.

The signalling path for the 6-4-6 scenario is illustrated in Figure 5.

```
   +------+      +-----+       +------+         +------+
   | Host | Sig1 | DS  |       |  MR  |   PIM   |  MR  |
   | Rcvr |------| AF1 |       |      | . . . . |      |
   |      | IPv6 |     |       | (BG) |   IPv6  | (DR) |
   +------+      +-----+       +------+         +------+
                  /               \
         IGMP / IPv4       PIM \ IPv6
                /                   \
         +------+      +----+       +------+
         |  MR  |  PIM |    | PIM   |  DS  |
         |      |------| MR | . . . |  AF2 |
         | (DR) | IPv4 |    | IPv4  | (BG) |
         +------+      +----+       +------+

             ------------------------------------->
```

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
BG : Border Gateway

Figure 5: Signalling Path For the 6-4-6 Scenario.

The multimedia content distribution path is shown in Figure 6.

```
   +------+      +-----+     +------+       +------+
   | Host |      | DS  |     |  MR  |       |  MR  |
   | Rcvr |------| AF1 |     |      | . . . |      |
   |      | IPv6 |     |     | (BG) | IPv6  | (DR) |
   +------+      +-----+     +------+       +------+
                  /               \               \
                 / IPv4           \ IPv6          \ IPv6
                /                   \               \
         +------+      +----+       +------+      +-----+
         |  MR  |      |    |       |  DS  |      |     |
         |      |------| MR | . . . |  AF2 |      | Src |
         | (DR) | IPv4 |    | IPv4  | (BG) |      |     |
         +------+      +----+       +------+      +-----+

             <-------------------------------------

   Rcvr: Multicast receiver
   DS  : Dual Stack
   AF  : Adaptation Function
   MR  : Multicast Router
   DR  : Designated Router
   BG  : Border Gateway
   Src : Multicast source
```

Figure 6: Multicast Content Distribution Path For the 6-4-6 Scenario.

3.4. IPv6 Receiver and IPv4 Source

We refer to this scenario as 6-4. An example of such use case is the context of some mobile networks, where terminal devices are only provisioned with an IPv6 prefix. Accessing IPv4-formatted multicast content from an IPv6-only receiver requires additional functions to be enabled.

This scenario is privileged by mobile operators who deploy NAT64 capabilities in their network. It is illustrated in Figures 7 (signalling path) and 8 (distribution of multicast contents). Only one adaptation function instance is needed, at the IPv4/IPv6 boundary.

```
      +------+                  +------+         +------+
      | Host |                  |  MR  |   PIM   |  MR  |
      | Rcvr |                  |      | . . . . |      |
      |      |                  | (BG) |   IPv4  | (DR) |
      +------+                  +------+         +------+
         \                           \
      MLD \ IPv6                  PIM \ IPv4
           \                           \
        +------+      +----+       +------+
        |  MR  |  PIM |    | PIM   |  DS  |
        |      |------| MR | . . . |  AF1 |
        | (DR) | IPv6 |    | IPv6  | (BG) |
        +------+      +----+       +------+

                -------------------------------------->

      Rcvr: Multicast receiver
      DS  : Dual Stack
      AF  : Adaptation Function (ALG or IWF)
      MR  : Multicast Router
      DR  : Designated Router
      BG  : Border Gateway
```

Figure 7: Signalling Path For the 6-4 Scenario.

```
    +------+                  +------+         +------+
    | Host |                  |  MR  |         |  MR  |
    | Rcvr |                  |      | . . . . |      |
    |      |                  | (BG) |   IPv4  | (DR) |
    +------+                  +------+         +------+
       \                           \              \
        \ IPv6                      \ IPv4         \ IPv4
         \                           \              \
      +------+      +----+       +------+         +-----+
      |  MR  |      |    |       |  DS  |         |     |
      |      |------| MR | . . . |  AF1 |         | Src |
      | (DR) | IPv6 |    | IPv6  | (BG) |         |     |
      +------+      +----+       +------+         +-----+

              <-------------------------------------

    Rcvr: Multicast receiver
    DS  : Dual Stack
    AF  : Adaptation Function
    MR  : Multicast Router
    DR  : Designated Router
    BG  : Border Gateway
    Src : Multicast source
```

Figure 8: Multicast Content Distribution Path For the 6-4 Scenario.

3.5. IPv4 Receiver and IPv6 Source

We refer to this scenario as 4-6. According to a BEHAVE WG consensus when elaborating the transition unicast scenarios, multicast sources are likely to remain IPv4-enabled in a first stage; therefore, the content is likely to remain IPv4-formatted.

As a consequence, this scenario is unlikely to occur during the first years of the transition period, and has been assigned a lower priority compared to the use cases depicted in Sections 3.1, 3.2 and 3.4.

The signalling path for this scenario is shown in Figure 9. The multicast content distribution path is shown in Figure 10. There are similarities with the 6-4 scenario but address mapping across IP version boundaries is more challenging.

```
      +------+                  +------+         +------+
      | Host |                  |  MR  |   PIM   |  MR  |
      | Rcvr |                  |      | . . . . |      |
      |      |                  | (BG) |   IPv6  | (DR) |
      +------+                  +------+         +------+
          \                           \
      IGMP \ IPv4                  PIM \ IPv6
            \                            \
         +------+      +----+       +------+
         |  MR  |  PIM |    | PIM   |  DS  |
         |      |------| MR | . . . |  AF1 |
         | (DR) | IPv4 |    | IPv4  | (BG) |
         +------+      +----+       +------+

                ------------------------------------->

      Rcvr: Multicast receiver
      DS  : Dual Stack
      AF  : Adaptation Function (ALG or IWF)
      MR  : Multicast Router
      DR  : Designated Router
      BG  : Border Gateway
```

Figure 9: Signalling Path For the 4-6 Scenario.

```
   +------+                 +------+         +------+
   | Host |                 |  MR  |         |  MR  |
   | Rcvr |                 |      | . . . . |      |
   |      |                 | (BG) |   IPv6  | (DR) |
   +------+                 +------+         +------+
      \                          \              \
       \ IPv4                     \ IPv6         \ IPv6
        \                          \              \
     +------+      +----+       +------+         +-----+
     |  MR  |      |    |       |  DS  |         |     |
     |      |------| MR | . . . |  AF1 |         | Src |
     | (DR) | IPv4 |    | IPv4  | (BG) |         |     |
     +------+      +----+       +------+         +-----+

             <-------------------------------------
```

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function
MR : Multicast Router
DR : Designated Router
BG : Border Gateway
Src : Multicast source

Figure 10: Multicast Content Distribution Path For the 4-6 Scenario.

3.6. Summary

To summarize, the use cases of highest priority are those involving IPv4 sources, i.e., the 4-6-4, 6-4-6-4, and 6-4 scenarios.

4. Design Considerations

4.1. Group and Source Discovery Considerations

Multicast applications that embed address information in the payload may require Application Level Gateway (ALG) during the transition period. An ALG is application-specific by definition, and may therefore be unnecessary depending on the nature of the multicast service.

Such ALG (Application Level Gateway) may also be required to help an IPv6 receiver select the appropriate multicast group address when only the IPv4 address is advertised (e.g., when the SDP (Session Description Protocol) protocol is used to advertise some contents); otherwise, access to IPv4 multicast content from an IPv6 receiver may be compromised.

ALGs may be located upstream in the network. As a consequence, these ALGs do not know in advance whether the receiver is dual-stack or IPv6-only. In order to avoid the use of an ALG in the path, an IPv4-only source can advertise both an IPv4 multicast group address and the corresponding IPv4-embedded IPv6 multicast group address [I-D.boucadair-behave-64-multicast-address-format].

However, a dual-stack receiver may prefer to use the IPv6 address to receive the multicast content. The selection of the IPv6 multicast address would then require multicast flows to cross an IPv4-IPv6 interworking function.

The receiver should therefore be able to unambiguously distinguish an IPv4-embedded IPv6 multicast address from a native IPv6 multicast address.

4.2. Subscription

Multicast distribution trees are receiver-initiated. IPv4 receivers that wish to subscribe to an IPv4 multicast group will send the corresponding IGMP Report message towards the relevant IGMP Querier. In case the underlying access network is IPv6, the information conveyed in IGMP messages should be relayed by corresponding MLD messages.

4.3. Multicast Tree Computation

Grafting to an IPv4 multicast distribution tree through an IPv6 multicast domain suggests that IPv4 multicast traffic will have to be conveyed along an "IPv6-equivalent" multicast distribution tree. That is, part of the multicast distribution tree along which IPv4 multicast traffic will be forwarded SHOULD be computed and maintained by means of the PIMv6 machinery, so that the distribution tree can be terminated as close to the IPv4 receivers as possible for the sake of the multicast forwarding efficiency. This assumes a close interaction between the PIM designs enforced in both IPv4 and IPv6 multicast domains, by means of specific Inter-Working Functions that are further discussed in Section 4.4.

Such interaction may be complicated by different combinations: the IPv4 multicast domain is SSM-enabled (with no RP (Rendezvous Point) routers), while the IPv6 multicast domain may support both ASM (Any Source Multicast) and SSM (Source Specific Multicast) [RFC3569] modes.

4.4. Multicast Interworking Functions (IWF)

IPv4-IPv6 multicast interworking functions are required for both translation (one address family to another) and traversal (one address family over another) contexts.

Given the multiple versions of Group Membership management protocols, issues may be raised when, for example, IGMPv2 is running in the IPv4 multicast domain that is connected to the IPv6 multicast domain by means of an IWF, while MLDv2 is running in the IPv6 multicast domain. To solve these problems, the design of the IWF function SHOULD adhere to the IP version-independent, protocol interaction approach documented in Section 8 of [RFC3810] and Section 7 of [RFC3376].

Note that, for traversal cases, to improve the efficiency of the multicast service delivery, traffic will be multicast along distribution trees that should be terminated as close to the receivers as possible for the sake of bandwidth optimization. As a reminder, the traversal of unicast-only (access) networks is not considered in this draft.

4.4.1. IWF For Control Flows

The IWF to process multicast signalling flows (such as IGMP or MLD Report messages) should be independent of the IP version and consist mainly of an IPv4-IPv6 adaptation element and an IP address translation element. The message format adaptation must follow what is specified in [RFC3810] or [RFC4601], and the device that embeds the IWF device must be multicast-enabled, i.e., support IGMP, MLD and/or PIM, depending on the context (address family-wise) and the design (e.g., this device could be a PIM DR in addition to a MLD Querier).

The IWF can then be operated in the following modes: IGMP-MLD, PIMv4-PIMv6, MLD-PIMv4 and IGMP-PIMv6. In particular, Source-Specific Multicast (SSM) must be supported (i.e., IGMPv3/MLDv2 signalling traffic as well as the ability to directly send PIM (S, G) Join messages towards the source).

The following sub-sections describe some interworking functions which may be solicited, depending on the environment.

4.4.1.1. IGMP-MLD Interworking

The IGMP-MLD Interworking Function combines the IGMP/MLD Proxying function specified in [RFC4605] and the IGMP/MLD adaptation function which is meant to reflect the contents of IGMP messages into MLD messages, and vice versa.

For example, when an IGMP Report message is received to subscribe to a given multicast group (which may be associated to a source address if SSM mode is used), the IGMP-MLD Interworking Function MUST send an MLD Report message to subscribe to the corresponding IPv6 multicast group.

4.4.1.2. IPv4-IPv6 PIM Interworking

[RFC4601] allows the computation of PIM-based IPv4 or IPv6 distribution trees; PIM is IP version agnostic. There is no specific IPv6 PIM machinery that would work differently than an IPv4 PIM machinery. The new features needed for the IPv4-IPv6 PIM Interworking Function consist in dynamically triggering the PIM message of Address Family 1 upon receipt of the equivalent PIM message of Address Family 2.

The address mapping MUST be performed similarly to that of the IGMP-MLD Interworking Function.

4.4.1.3. MLD-IPv4 PIM Interworking

This IWF function is required when the MLD Querier is connected to an IPv4 PIM domain.

The address mapping MUST be performed similarly to that of the IGMP-MLD Interworking Function.

4.4.1.4. IGMP-IPv6 PIM Interworking

The address mapping MUST be performed similarly to that of the IGMP-MLD Interworking Function.

4.4.2. IWF For Data Flows

The IWF to be used for multicast data flows is operated at the boundary between IPv4 and IPv6 multicast networks. Either encapsulation/de-capsulation or translation modes can be enforced, depending on the design. Note that translation operations must follow the algorithm specified in [RFC6145].

4.4.3. Address Mapping

The address mapping mechanisms to be used in either a stateful or stateless fashion need to be specified for the translation from one address family to the other.

The address formats have been defined in [I-D.boucadair-behave-64-multicast-address-format] and [RFC6052] for IPv4-embedded IPv6

multicast and unicast addresses. Mapping operations are performed in a stateless manner by the algorithms specified in the aforementioned documents.

In this context, the IPv6 prefixes required for embedding IPv4 addresses can be assigned to devices that support IWF features by various means (e.g., static or dynamic configuration, out-of-band mechanisms, etc.).

If stateful approaches are used, it is recommended to carefully investigate the need to synchronize mapping states between multiple boxes, and the coordination of the IWF and source/group discovery elements is also required, at the cost of extra complexity.

4.5. Combination of ASM and SSM Modes

The ASM (Any Source Multicast) mode could be used to optimize the forwarding of IPv4 multicast traffic sent by different sources into the IPv6 multicast domain by selecting RP routers that could be located at the border between the IPv6 and the IPv4 multicast domains. This design may optimize the multicast forwarding efficiency in the IPv6 domain when access to several IPv4 multicast sources needs to be granted.

[To be further elaborated.]

5. What Is Expected From The IETF

This document highlights the following IETF standardization needs:

- Specify the inter-working function as described in Sections 4.4.1 and 4.4.2. In particular:

  - Specify the algorithms used by various inter-working functions, covering both encapsulation and translation approaches

  - Specify the multicast IPv4-embedded address format

  - Document a 6-4 multicast architecture

  - Document a 6-4-6-4 multicast architecture

  - Document a 4-6-4 multicast architecture

- Document a Management Information Base (MIB) to be used for the management of IWF functions

- Encourage the publication of various Applicability Statement documents to reflect IWF operational experience in different contexts

6. IANA Considerations

This document makes no request to IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

Access to contents in a multicast-enabled environment raises different security issues that have been already documented. This draft does not introduce any specific security issue.

8. Acknowledgments

Special thanks to T. Taylor for providing the figures and some of the text that illustrate the use cases depicted in Section 3. Thanks also to N. Leymann and S. Venaas for their comments.

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.

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

[RFC3569] Bhattacharyya, S., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.

[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.

9.2. Informative References

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

Authors' Addresses

Christian Jacquenet
France Telecom

Email: christian.jacquenet@orange.com

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Yiu Lee
Comcast
US

Email: Yiu_Lee@Cable.Comcast.com

Jacni Qin
ZTE
China

Email: jacniq@gmail.com

Tina Tsou
Huawei Technologies (USA)
2330 Central Expressway
Santa Clara, CA 95050
USA

Phone: +1 408 330 4424
Email: tena@huawei.com