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

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 5, 2012.

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:

- 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.

- 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.

- 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).

- 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.

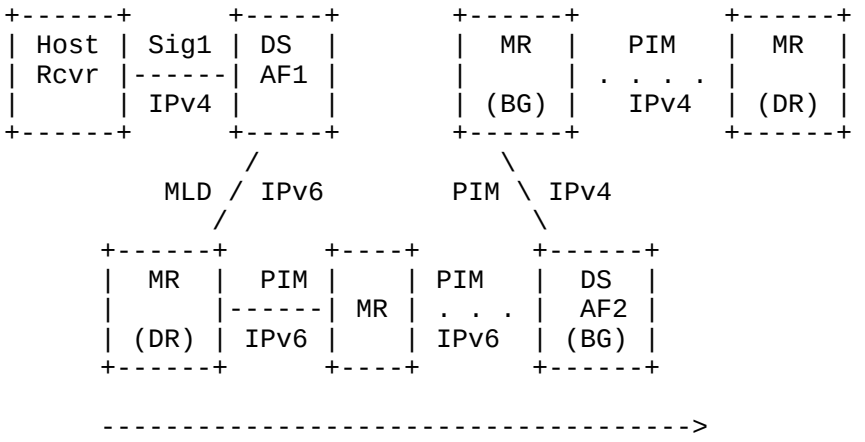

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
BG : Border Gateway

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.

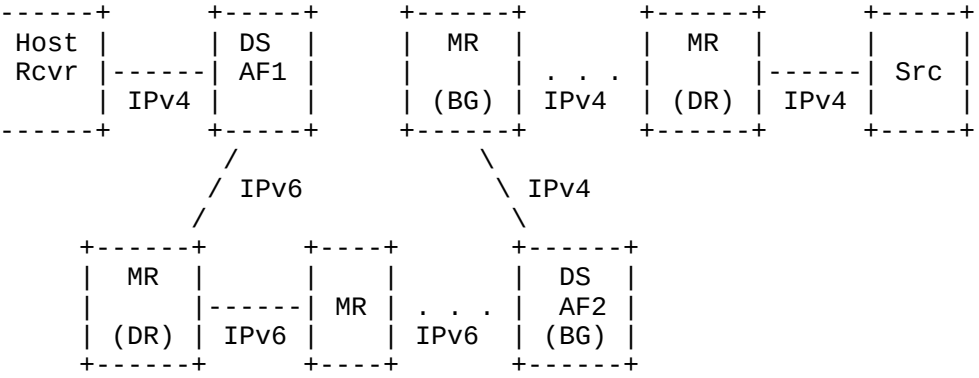

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function
MR : Multicast router
DR : Designated Router
BG : Border Gateway
Src : Multicast source

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.

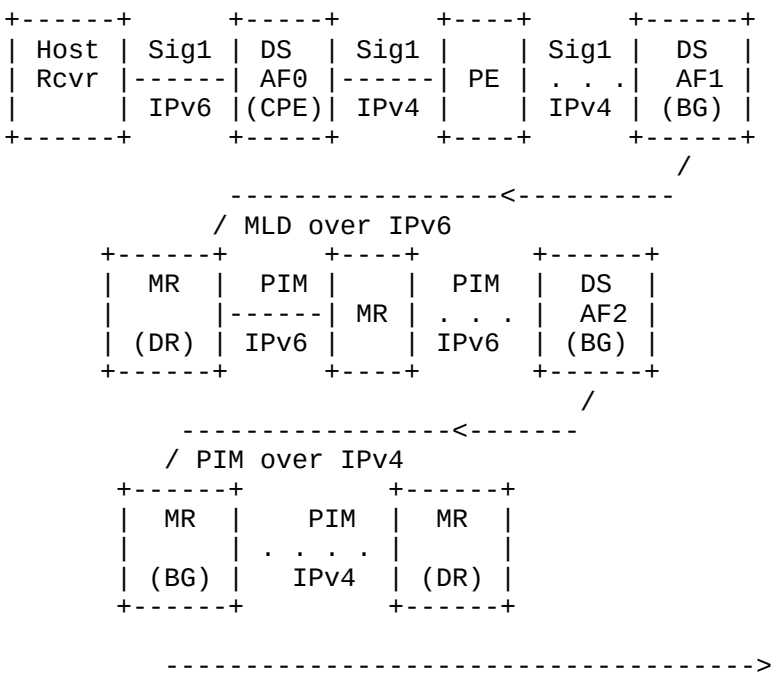

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
CPE : Customer Premises Equipement (Dual Stack router)
PE : Provider Edge router
BG : Border Gateway

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.

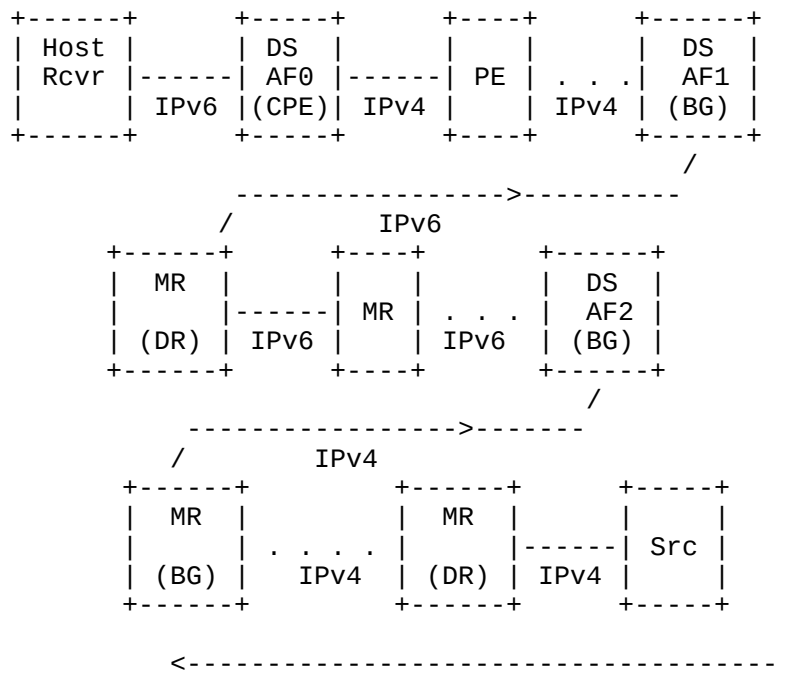

Rcvr: Multicast receiver
Src : Multicast source
DS : Dual Stack
AF : Adaptation function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
CPE : Customer Premises Equipment (Dual Stack router)
PE : Provider Edge router
BG : Border Gateway

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.

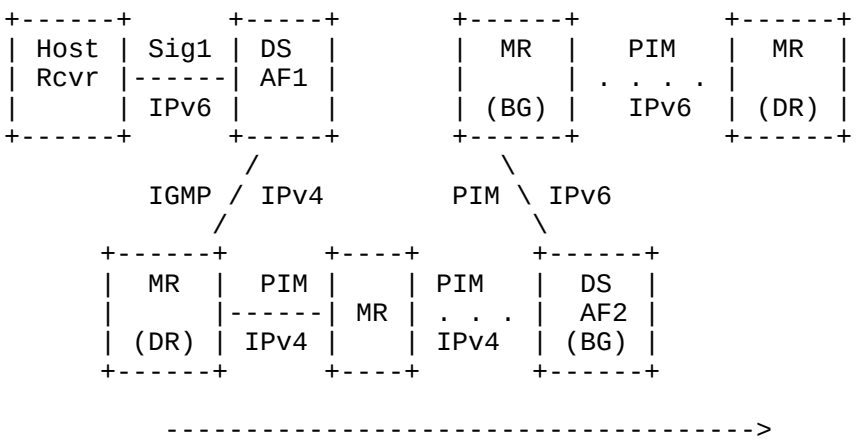

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function (ALG or IWF)
MR : Multicast Router
DR : Designated Router
BG : Border Gateway

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

The multimedia content distribution path is shown in Figure 6.

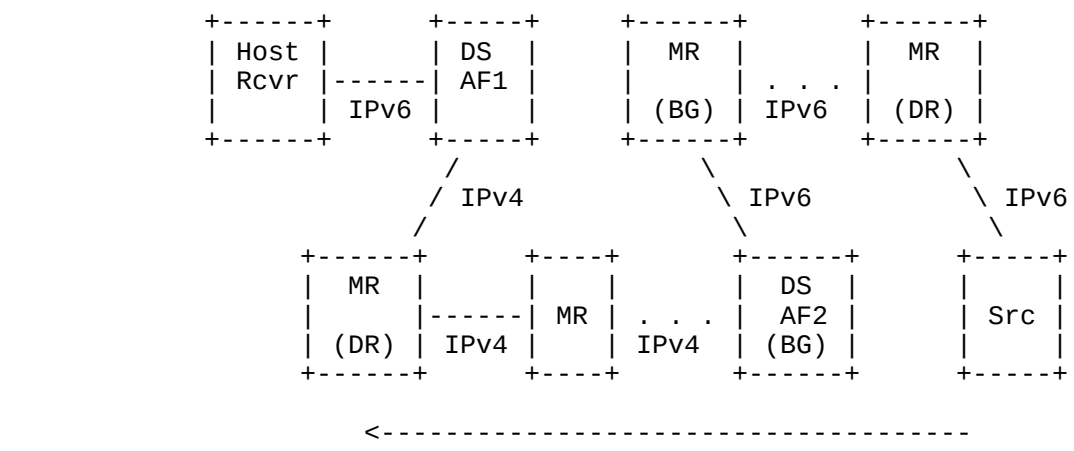

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function
MR : Multicast Router
DR : Designated Router
BG : Border Gateway
Src : Multicast source

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.

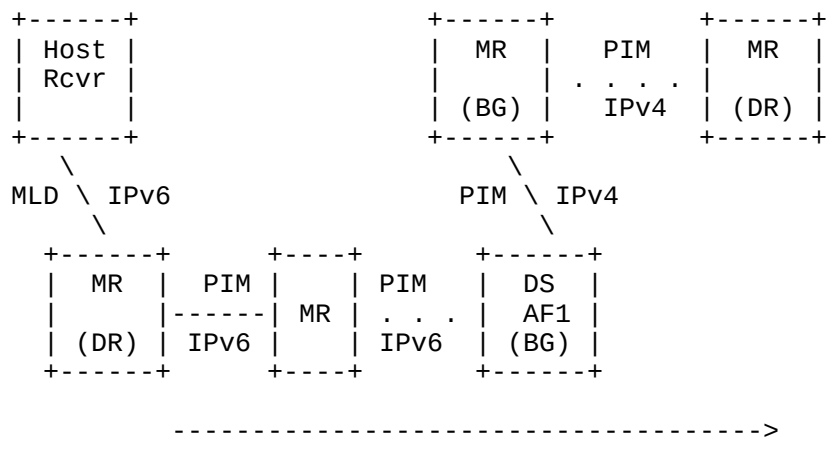

Figure 7: Signalling Path For the 6-4 Scenario.

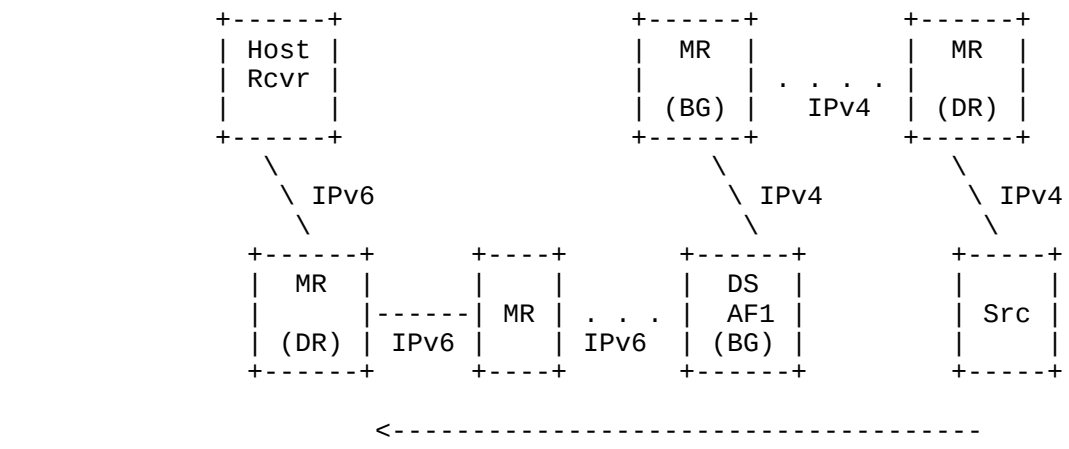

Figure 8: Multicast Content Distribution Path For the 6-4 Scenario.

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.

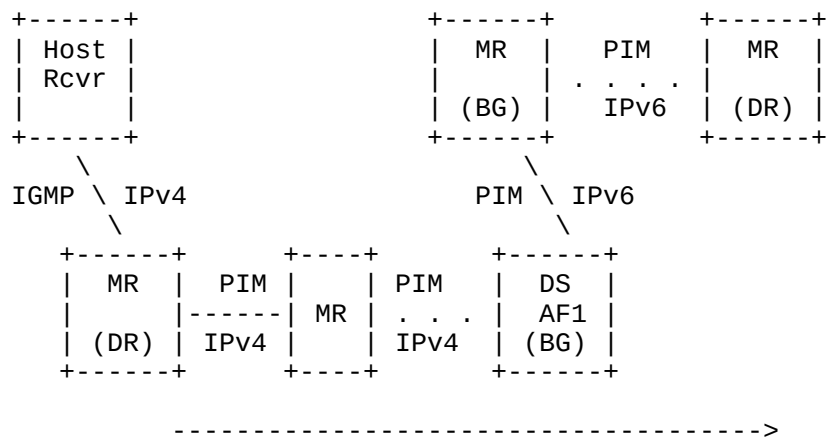

Figure 9: Signalling Path For the 4-6 Scenario.

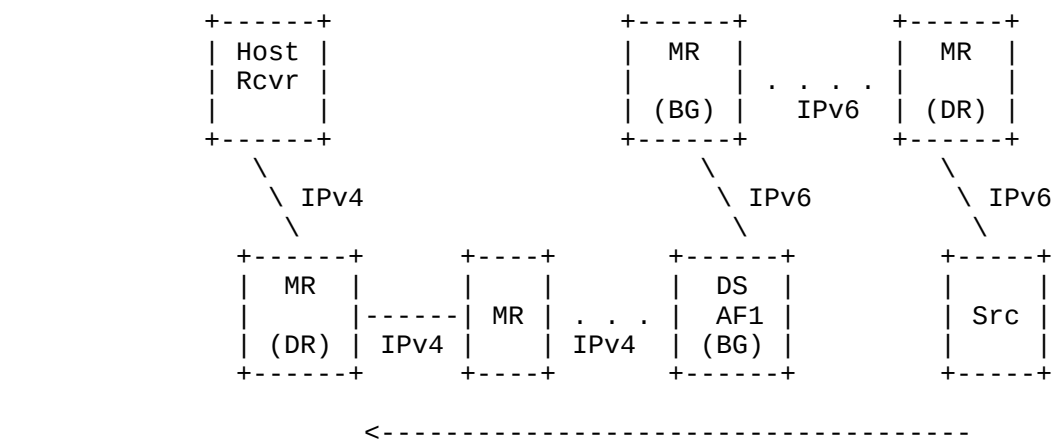

Rcvr: Multicast receiver
DS : Dual Stack
AF : Adaptation Function
MR : Multicast Router
DR : Designated Router
BG : Border Gateway
Src : Multicast source

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

- o 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

v6ops Work Group
Internet Draft
Intended status: Informational
Expires: April 23, 2012

S. Jiang
D. Gu
Y.Fu
Huawei Technologies Co., Ltd
October 24, 2011

Multicast Proxy in IPv6/IPv4 Transition
draft-jiang-v6ops-v4v6mc-proxy-00.txt

Status of this Memo

Copyright Notice

Abstract

During the long co-existing period of IPv6 and IPv4, the interoperation between IPv6 network and IPv4 network is essential. Multicast services across IPv6 and IPv4 networks are also needed. Besides the packet-based multicast translation mechanism, this document describes a multicast proxy solution. The solution is a multicast deployment for transition scenario. It does not propose any new protocol for multicast. The multicast proxy is deployed at the border of IPv6/IPv4 networks. It is mainly based on content cache concept. Without packet-based translation, it retires the content data from IPvX network, caches the data, and multicasts the data in IPvY network. It acts as a multicast leaf in the IPvX network where the data source locates. It also acts as a multicast source in IPvY network where the multicast client locates.

Table of Contents

1. Introduction

The deployment of IPv6 is now in progress, and users with no IPv4 access are likely to appear in increasing numbers in the coming years. However, it is also widely agreed that IPv4 will be still in use for a long period. During the long co-existing period of IPv6 and IPv4, the interoperation between IPv6 network and IPv4 network is essential.

Now, multimedia has been deployed widely, such as IPTV and video conference etc. They also face the IPv6 and IPv4 intercommunication issues. The multicast applications are complicated and face more difficulties than unicast applications deployment.

[I-D.draft-venaas-behave-v4v6mc-framework] proposes a packet-based translation framework between IPv4/IPv6 multicast services. It describes the packet-based translation operations and intercommunication in network layer to support a single source send to multiple receivers in different IP networks.

Besides the packet-based multicast translation mechanism, this document describes a multicast proxy solution, which is mainly based on content cache. It is a multicast deployment for IPv6 transition scenario. It doesn’t propose any new protocol for multicast.

A multicast proxy can be deployed at the border between IPv4/IPv6 networks. Without packet-based translation, it retires the content data from IPvX network, caches the data, and multicasts the data in IPvY network. It acts as a multicast leaf in the IPvX network where the data source locates. It also acts as a multicast source in IPvY network where the multicast client locates.

2. Terminology

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 RFC2119 [RFC2119].

3. Multicast Proxy without packet-based IPv6/IPv4 Translation

3.1. Overview

Within this document, we describe the network where the data source locates as IPvX network and the network where the multicast client locates as IPvY network. When IPvX is IPv6, IPvY must be IPv4, vice versa.

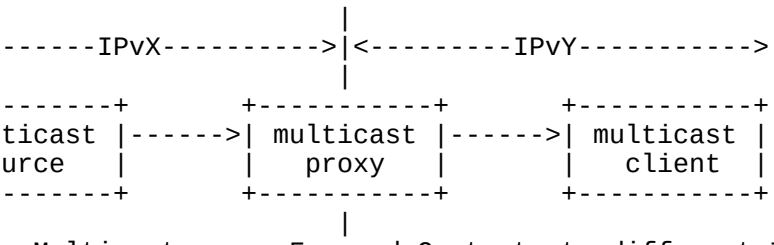

Figure 1: Multicast proxy Forward Contents to different IP networks

As showed in Figure 1, the proposed multicast proxy is deployed at the border of IPv4/IPv6 networks. It MUST support both IPv4 and IPv6. It MUST support both IGMP (Internet Group Management Protocol, [RFC3376]), which is used for IPv4 multicast management functions, and MLD (Multicast Listener Discovery, [RFC3810]), which is used in a similar way in IPv6 Environment. In the IPvX network, the multicast proxy joins the multicast distribution tree as a leaf. In the IPvY network, the multicast proxy broadcasts contents as a multicast source. The establishment of multicast distribution trees obeys the current multicast specifications for each IP family, such as Protocol Independent Multicast (PIM [RFC4601]).

Notice that there are two different multicast distribution trees in two sides of the multicast proxy. They are operational independent from each other in the network layer.

Logically, they are relevant to each other and there are interoperation behaviors between them. The contents published through the multicast distribution tree in IPvY network inherits from the IPvX network. They are received by the multicast proxy, which is a multicast leaf in the multicast distribution tree in IPvX network. Within the multicast proxy, contents are mapped between receiver function and publisher function. The operations of the multicast distribution tree in IPvY network MAY trigger some operations of the multicast distribution tree in IPvX network. For example, a multicast client joins a multicast group in IPvY network, and requests multicast contents may cause the multicast proxy joins a multicast group in IPvX network.

However, as mentioned earlier, in network or IP layer, the two multicast distribution trees are independent from each other. Their operations are separated in two sides of the multicast proxy. Conceptually, the multicast proxy can be presented virtually in functional modules like below Figure 2.

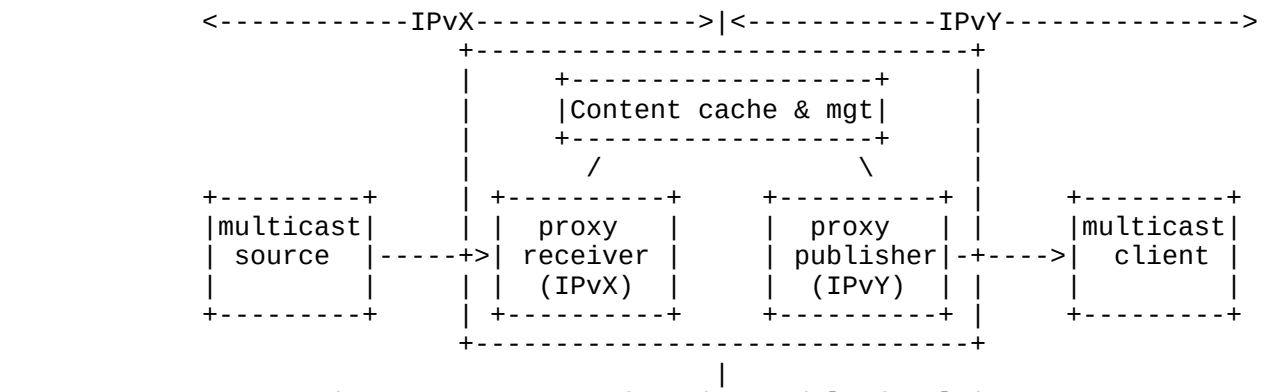

Figure 2: Separate function model of multicast proxy

As shown in Figure 2, the proxy receiver module in IPvX network joins IPvX multicast groups as a receiver client. Thereby it receives packets bound for the IPvX multicast groups, and then hands the content to the content cache and management module. The content cache and management module then forward the on-demand content to the multicast proxy function module in IPvY network, which acts as a publisher and multicast source in IPvY network.

3.2. Operation procedure

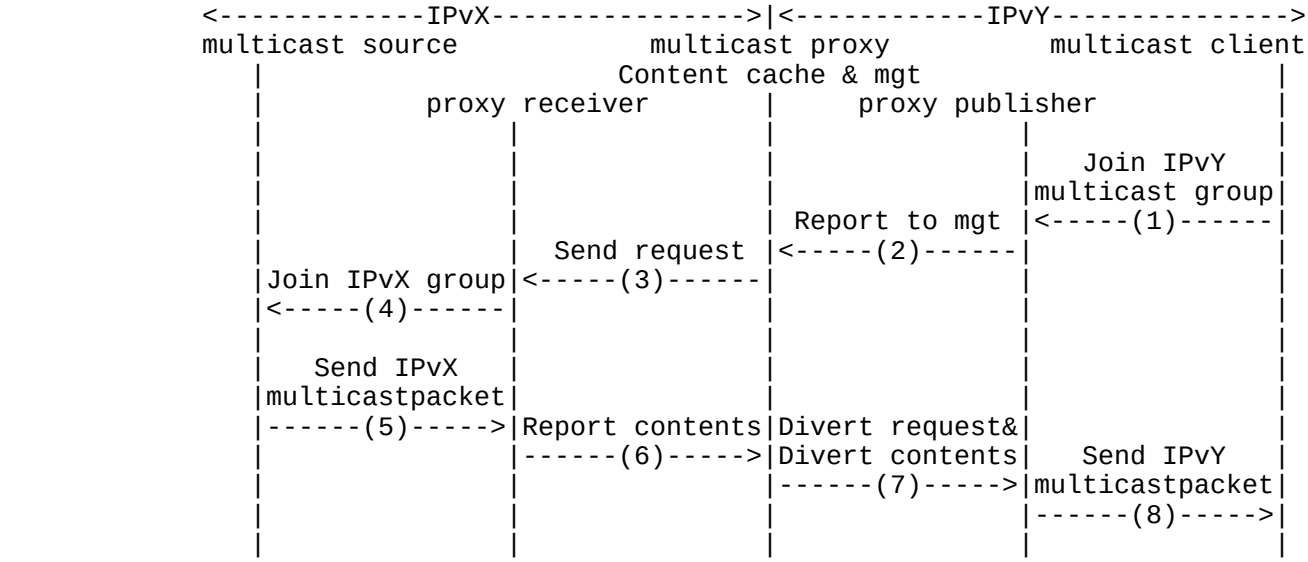

Figure. 3 The interaction communicating from IPvX to IPvY

As shown in Figure 3, a client, which locates in IPvY network, connects to the multicast proxy, requesting a multicast service whose source locates in the IPvX network.

First, the client sends a join IPvY multicast group report (1) to the multicast proxy. The proxy publisher module, which also locates in the IPvY network, receives this report, then forwards the content request to the content cache & management module (2). The content cache & mgt module maintains a content & multicast service table, including all available multicast services from IPvX network. The content cache & mgt module searches the client request in its dynamically updated table.

If the requested content is already multicasted in the IPvY network, the content cache & mgt module diverts the user report back to the proxy publisher module (7). The proxy publisher module adds the new client into its existing multicast tree. Then the requested content can be sent to the client (8).

If the requested content is available but not multicasted in IPvY network yet, the content cache & mgt module sends a request to the proxy receiver module, which locates in the IPvX network (3). It initiates the proxy receiver module to send a join IPvX multicast group report (4) to the multicast source. The multicast source adds the multicast proxy into its multicast tree and sends IPvX multicast packets (5). When receiving the multicast packets, the proxy receiver module drops all network layer information, such as IP headers, etc., and only reports contents to the content cache & mgt module (6). The content cache & mgt module then diverts contents to the proxy publisher module (7). The proxy publisher module builds up a new multicast tree in the IPvY network, and sends multicast packets to clients (8).

If all the clients, requesting a certain multicast service in the IPvY network, leave the IPvY multicast group, the multicast proxy MAY leave the IPvX multicast group in IPvX network.

Multicast proxies MAY also perform load-balancing, user authentication and other additional functions.

4. Security Considerations

The multicast proxy solution actually separate the IPv4 and IPv6 multicast services effectively. It prevents the attacks at only one side of it.

However, multicast proxy itself is as vulnerable as normal multicast sources and multicast leafs in each IPv4 or IPv6 environment. The security mechanisms for IGMP/MLD can be used to enhance the security of multicast proxy.

5. IANA Considerations

This draft does not request any IANA action.

6. Acknowledgments

The authors would like to thank Stig Venaas, Cisco for his valuable comments.

7. References

7.1. Normative References

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

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

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

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

7.2. Informative References

[I-D.draft-venaas-behave-v4v6mc-framework] S. Venaas, X. Li and C. Bao, "Framework for IPv4/IPv6 Multicast Translation ", draft-venaas-behave-v4v6mc-framework, working in progress, June 2011.

Author's Addresses

Sheng Jiang
Huawei Technologies Co., Ltd
156 Bei-Qing Road, Hai-Dian District, Beijing 100095
P.R. China
Email: jiangsheng@huawei.com

Dujuan Gu
Huawei Technologies Co., Ltd
156 Bei-Qing Road, Hai-Dian District, Beijing 100095
P.R. China
Email: gudujuan@huawei.com

Yu Fu
Huawei Technologies Co., Ltd
156 Bei-Qing Road, Hai-Dian District, Beijing 100095
P.R. China
Email: eleven.fuyu@huawei.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 23, 2011

S. Venaas
cisco Systems
X. Li
C. Bao
CERNET Center/Tsinghua University

June 21, 2011

Framework for IPv4/IPv6 Multicast Translation
draft-venaas-behave-v4v6mc-framework-03.txt

Abstract

This draft describes how IPv4/IPv6 multicast translation may be used in various scenarios and attempts to be a framework for possible solutions. This can be seen as a companion document to the document "Framework for IPv4/IPv6 Translation" by Baker et al. When considering scenarios and solutions for unicast translation, one should also see how they may be extended to provide multicast translation.

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 December 23, 2011.

Copyright Notice

Table of Contents

1. Introduction

There will be a long period of time where IPv4 and IPv6 systems and networks need to coexist. There are various solutions for how this can be done for unicast, some of which are based on translation. The document [RFC6144] discusses the needs and provides a framework for unicast translation for various scenarios. Here we discuss the need for multicast translation for those scenarios.

For unicast the problem is basically how two hosts can communicate when they are not able to use the same IP protocol. For multicast we can restrict ourselves to looking at how a single source can efficiently send to multiple receivers. When using a single IP protocol one builds a multicast distribution tree from the source to the receivers, and independent of the number of receivers, one sends each data packet only once on each link. We wish to maintain the same characteristics when there are different IP protocols used. That is, when the nodes of the tree (source, receivers and routers) cannot all use the same IP protocol. In general there may be multiple sources sending to a multicast group, but that can be thought of as separate trees, one per source. We will focus on the case where the source and the receivers cannot all use the same IP protocol. If the issue is the network in between, encapsulation may be a better alternative. Note that if the source supports both IPv4 and IPv6, then one alternative could be for the source to send two streams. This need not be the same host. There could be two different hosts, and in different locations/networks, sending the same content.

2. Translation scenarios

We will consider six different translation scenarios. For each of the scenarios we will look at how host in one network can receive multicast from a source in another network. For unicast one might consider the following six scenarios as described in [RFC6144]:

Scenario 1: An IPv6 network to the IPv4 Internet

Scenario 2: The IPv4 Internet to an IPv6 network

Scenario 3: The IPv6 Internet to an IPv4 network

Scenario 4: An IPv4 network to the IPv6 Internet

Scneario 5: An IPv6 network to an IPv4 network

Scenario 6: An IPv4 network to an IPv6 network

We have intentionally left out how one might connect the entire IPv4 Internet with the entire IPv6 Internet. In these scenarios one would look at how a host in one network initiates a uni- or bi-directional flow to another network. The initiator needs to somehow know which address to send the initial packet to, and the initial packet gets translated before reaching its destination.

For unicast translation it is quite natural to talk about networks and the Internet. For multicast this is not so clear, since there is limited use of multicast on the Internet. Certain parts of the Internet, e.g. academic and research networks and the links connecting them do carry multicast though. Also, the challenges and ideas described here regarding the Internet, also applies in other cases where there are multiple connected networks exchanging multicast.

For multicast one generally need a receiver to signal the group (and sometimes also the source) it wants to receive from. The signalling generally goes hop-by-hop towards the source to build multicast forwarding state that later is used to forward multicast in the reverse direction. This means that for the receiving host to receive multicast, it must first somehow know which group (and possibly source) it should signal that it wants to receive. These signals would then probably go hop-by-hop to a translator, and then the translated signalling would go hop-by-hop from the translator to the source. Note that this description is correct for SSM (source-specific multicast), but is in reality more complex for ASM (any-source multicast). An anology to unicast might perhaps be TCP streaming where a SYN is sent from the host that wants to receive the

stream to the source of the stream. Then the application data flows in the reverse direction of the initial signal. Hence we argue that the above unicast scenarios correspond to the following multicast scenarios, respectively:

Scenario 1: An IPv6 network receiving multicast from the IPv4 Internet

Scenario 2: The IPv4 Internet receiving multicast from an IPv6 network

Scneario 3: The IPv6 Internet receiving multicast from an IPv4 network

Scenario 4: An IPv4 network receiving multicast from the IPv6 Internet

Scenario 5: An IPv6 network receiving multicast from an IPv4 network

Scenario 6: An IPv4 network receiving multicast from an IPv6 network

2.1. Scenario 1: An IPv6 network receiving multicast from the IPv4 Internet

Here we have a network, say ISP or enterprise, that for some reason is IPv6-only, but the hosts in the IPv6-only network should be able to receive multicast from sources in the IPv4 internet. The unicast equivalent is "IPv6 network to the IPv4 Internet".

This is simple because the global IPv4 address space can be embedded into IPv6 [RFC6052]. Unicast addresses according to the unicast translation in use. For multicast one may embed all IPv4 multicast addresses inside a single IPv6 multicast prefix. Or one may have multiple embeddings to allow for appropriate mapping of scopes and ASM versus SSM. Using this embedding, the IPv6 host (or an application running on the host) can send IPv6 MLD reports for IPv6 groups (and if SSM, also sources) that specify which IPv4 source and groups that it wants to receive. The usual IPv6 state (including MLD and possibly PIM) needs to be created. If PIM is involved we need to use RPF to set up the tree and accept data, so the source addresses must be routed towards some translation device. This is likely to be the same device that would do the unicast translation. The translation device can in this case be completely stateless. There is some multicast state, but that is similar to the state in a multicast router when translation is not performed. Basically if the translator receives MLD or PIM messages asking for a specific group (or source and group), it uses these mappings to find out which IPv4 group (or source and group) it needs to send IGMP or PIM messages

for. This is no different than multicast in general, except for the translation. Whenever the translator receives data from the IPv4 source, it checks if it has anyone interested in the respective IPv6 group (or source and group), and if so, translates and forwards the data packets.

IPv6 applications need to somehow learn which IPv6 group (or source and group) to join. This is further discussed in Section 3.4.

2.2. Scenario 2: The IPv4 Internet receiving multicast from an IPv6 network

Here we will consider an IPv6 network connected to the IPv4 internet, and how any IPv4 host may receive multicast from a source in the IPv6 network. The unicast equivalent is "the IPv4 Internet to an IPv6 network".

This is difficult since the IPv6 multicast address space cannot be embedded into IPv4. Indeed this case has many similarities with how IPv4 networks can receive from the IPv6 Internet. See scenario (4), Section 2.4. However, in this case, all IPv4 hosts on the Internet should use the same mapping, and it might make sense to have additional requirements on the IPv6 network, rather than to add requirements for the IPv4 Internet.

One solution here might be for the IPv6 source application to somehow register with the translator to set up a mapping and receive an IPv4 address. The application could then possibly send SDP that includes both its IPv6 source and group, and the IPv4 source and group it got from the translator. Of course the signalling could also be done by manually adding a static mapping to the translator and specifying that address to the application. If instead we were to do signalling on the IPv4 side, then an IPv4 receiver would probably need a mechanism for finding an IPv4 address of the translator for a given IPv6 group. The IPv4 address could perhaps be embedded in the IPv6 group address? Or with say SDP there could be a way of specifying the IPv4 translator address. The IPv4 host could then communicate with the translator to establish a mapping (unless one exists) and learn which IPv4 group to join.

The best alternative might be to restrict the IPv6 multicast groups that should be accessible on the IPv4 internet to a certain IPv6 prefix. This may allow stateless translation. This could also be used in the reverse direction, for an IPv6 host to receive from an IPv4 source. Or in other words, the same mapping can be used in both directions. This has similarities with IVI [I-D.xli-behave-ivi], [RFC6145], [RFC6052] and also [I-D.venaas-behave-mcast46]. By using IVI source addresses (IPv4-translatable addresses) and a similar

technique for multicast addresses, the correct IPv4 source and group addresses can be derived from those. This method has many benefits, the main issue is that it cannot work for arbitrary IPv6 multicast addresses.

2.3. Scenario 3: The IPv6 Internet receiving multicast from an IPv4 network

We here consider the case where the Internet is IPv6, but there is some network of perhaps legacy IPv4 hosts that is IPv4-only. We want any IPv6 host on the Internet to be able to receive multicast from an IPv4 source. The unicast equivalent is "the IPv6 Internet to an IPv4 network".

This scenario can be solved using the same techniques as in Scenario 1, Section 2.1. There may however be differences regarding exactly which mappings are used and how applications may become aware of them. To obtain full benefit of multicast, all IPv6 hosts need to use the same mappings.

2.4. Scenario 4: An IPv4 network receiving multicast from the IPv6 Internet

Here we consider how an IPv4-only host in an IPv4 network may receive from an IPv6 multicast sender on the Internet. The unicast equivalent is "an IPv4 network to the IPv6 Internet".

For dual-stack hosts in an IPv4 network one should consider tunneling. This is difficult since we cannot embed the entire IPv6 space into IPv4. One might consider some of the techniques from scenario (2), Section 2.2. That scenario is however much easier since one may restrict which IPv6 groups are used and there is a limited number of sources.

For unicast one might use a DNS-ALG for this, where the ALG would instantiate translator mappings as needed. This is the technique used in NAT-PT [RFC2766], which was deprecated by [RFC4966].

However, for multicast one generally does not use DNS. One could consider doing the same with an ALG for some other protocol. E.g. translate addresses in SDP files when they pass the translator, or in any other protocol that might transfer multicast addresses. This would be very complicated and not recommended.

Rather than using an ALG that translates addresses in application protocol payload, one could consider new signalling mechanisms for more explicit signalling. The additional signalling could be either on the IPv6 or the IPv4 side. It may however not be a good idea to

require additional behavior by host and applications on the IPv6 Internet to accomodate legacy IPv4 networks. Also, since one may not be able to provide unique IPv4 multicast addresses for all the IPv6 multicast groups that are in use, it makes more sense that the mappings are done locally in each of the IPv4 networks, where IPv4 multicast addresses might be assigned on-demand. An IPv4 receiver might somehow request an IPv4 mapping for an IPv6 group (and possibly source). This creates a mapping in the translator so that when the IPv4 receiver joins the IPv4 group, the translator knows which IPv6 group (and possibly source) to translate it into. Of course the signalling could also be done manually by adding a static mapping to the translator and somehow specifying the right IPv4 address to the application.

2.5. Scenario 5: An IPv6 network receiving multicast from an IPv4 network

In this scenario we consider IPv4 and IPv6 networks belonging to the same organization. The unicast equivalent is "an IPv6 network to an IPv4 network".

We would like any IPv6 host to receive from any IPv4 sources. Here one can use the same techniques as for an IPv6 network receiving from the IPv4 internet. It is really a special case of scenario (1), Section 2.1.

The fact that the number of hosts are limited and that there is common management might simplify things. Due to the limited scale, one could perhaps just manually configure all the static mappings needed in the translator and manually create the necessary announcements or in some cases have the applications create the necessary announcements. But it might be better to use a stateless approach where IPv4 unicast and multicast addresses are embedded into IPv6. Like IVI [I-D.xli-behave-ivi], or [I-D.venaas-behave-mcast46].

2.6. Scenario 6: An IPv4 network receiving multicast from an IPv6 network

In this scenario we consider IPv4 and IPv6 networks belonging to the same organization. The unicast equialent is "an IPv4 network to an IPv6 network".

We would like any IPv4 host to receive from any IPv6 source. This can be seen as special cases of either scenario (2), Section 2.2 or scenario (4), Section 2.4, where any of those techniques might apply. However, as discussed in scenario (5) Section 2.5 where we looked at how to do multicast in the reverse direction; the limited number of hosts and common managment might allow us to just use static mappings

or a stateless approach by restricting which IPv6 addresses are used. By using these techniques one may be able to create mappings that can be used for multicast in both directions, combining this scenario with scenario (5).

3. Framework

Having considered some possible scenarios for where and how we may use multicast translation, we will now consider some general issues and the different components of such solutions.

3.1. Addressing

When doing stateless translation, one need to somehow encode IPv4 addresses inside IPv6 addresses so that there is a well defined way for the translator to transform an IPv6 address into IPv4. This can be done with techniques like IVI [I-D.xli-behave-ivi] and [I-D.venaas-behave-mcast46].

There are two types of addressing schemes related to the IPv4/IPv6 multicast translation. The source addressing and the group addressing.

3.1.1. Source addressing

Source addressing issues is the same as in the unicast IPv4/IPv6 translation defined in [RFC6052]. The IPv4-mapped address is used for representing IPv4 in IPv6 and the IPv4-translatable address is used for representing IPv6 in IPv4 when the stateless translator is used. The multicast RPF relies on the source address to build the distribution tree. Therefore, depending on the operation mode of the IPv4/IPv6 translator and receiving directions, the IPv4-mapped or the IPv4-translatable addresses will be used.

3.1.2. Group addressing

Group addressing issue is unique to the IPv4/IPv6 multicast translation. The entire IPv4 group addresses can be uniquely represented by the IPv6 group addresses, while the entire IPv6 group addresses cannot be uniquely represented by the IPv6 group addresses. Therefore, special group address mapping rule between IPv4 group addresses and IPv6 group addresses should be defined for the IPv4/ IPv6 multicast translation.

3.2. Routing

The actual translation of multicast packets may not be very complicated, in particular if it can be stateless. For the multicast to actually go through the translator we need to have routes for the multicast source addresses involved, so that multicast packets both on their way to and from the translator satisfy RPF checks. These routes are also needed for protocols like PIM-SM to establish a multicast tree, since RPF is used to determine where to send join

messages. To go into more detail we need to look at different scenarios like SSM (Source-Specific Multicast) and ASM (Any-Source Multicast), and PIM versus IGMP/MLD.

3.2.1. Translation with PIM and SSM

When doing SSM, a receiver specifies both source and group addresses. If the receiver is to receive translated packets, it must do an IGMP/MLD join for the source and group address that the data packets will have after translation. We will later look at how it may learn those addresses. For the source address it joins, the unicast routing (or it may be an alternate topology specific to multicast), must point towards the translator. With this in place, PIM should build a tree hop-by-hop from the last-hop router to the translator. The translator then maps the source and group addresses in the PIM join to the source and group the data packets have before translation. The translator then does a PIM join for that source and group. Provided the routing is correct, this will then build a tree all the way to the source. Finally when these joins reach the source, any data sent by the source will follow this path to the translator, get translated, and then continue to the receiver.

3.2.2. Translation with PIM and ASM

Let us first consider PIM Sparse Mode. In this case a receiver just joins a group. If this group is to be received via the translator we need to send joins towards the translator, but initially PIM will send joins towards the RP (Rendezvous-Point) for the group. The most efficient solution is probably to make sure that the translator is configured as an RP for all groups that one may receive through it. That is, for the groups it translates to. E.g. if IPv4 groups are embedded into an IPv6 multicast prefix, then the translator could be an RP for that specific prefix. The translator may then translate the group and join towards the group address that is used before translation. Note that if the translator also is an RP for the addresses used before translation, it should know which sources exist and join each of these. If it is not an RP, it needs to join towards the RP. If the translator did not know the sources, it may join each of the sources as soon as it receives from them (that is, switching to Shortest Path Trees). When the translator receives data, it translates it and then sends the translated data. This then follows the joins for the translated groups to the receivers. When the last-hop routers start receiving, they will probably (this is usually the default behavior) switch to SPTs (Shortest Path Trees). These trees also need to go to the translator and would probably follow the same path as the previously built shared tree. One might argue here that switching to the SPT has no benefit if it is the same path anyway. Also with shared trees, RPF is not an issue, so the translated source

addresses don't need to be routed towards the translator.

At the end of the previous paragraph we pointed out that there is no benefit in switching to shortest path trees if they have to go via the translator anyway. A possibility here could be to use Bidirectional PIM where there is no source specific state and data always go through the RP. It is possible to use Bidir just for those groups that are translated, and then make the translator the RP.

3.2.3. Translation with IGMP/MLD

For translation taking place close to the edge, e.g. a home gateway, one may consider just using IGMP and MLD, and no PIM. In that case the translator should for any received MLD reports for IPv6 groups that correspond to translated IPv4 groups, map those into IGMP reports that it sends out on the IPv4 side. And vice versa for data in the other direction. Note that a translator implementation could also choose to do this in just one direction. For SSM it would also need to translate the source addresses.

3.3. Translation in operation

Currently, the proposed solutions for IPv6/IPv4 translation are classified into stateless translation and stateful translation.

3.3.1. Stateless Translation

For stateless translation, the translation information is carried in the address itself, permitting both IPv4->IPv6 and IPv6-<IPv4 sessions establishment. The stateless translation supports end-to-end address transparency and has better scalability compared with the stateful translation. See [RFC6145] and [I-D.xli-behave-ivi].

Stateless translation can be used for Scenarios 1, 2, 5 and 6, i.e. it supports "An IPv6 network receiving multicast from the IPv4 Internet", "the IPv4 Internet receiving multicast from an IPv6 network", "An IPv6 network receiving multicast from an IPv4 network" and "An IPv4 network receiving multicast from an IPv6 network".

In the stateless translation, an IPv6 network uses the IPv4-translatable addresses, while the IPv4 Internet or an IPv4 network can be represented by IPv4-mapped addresses.

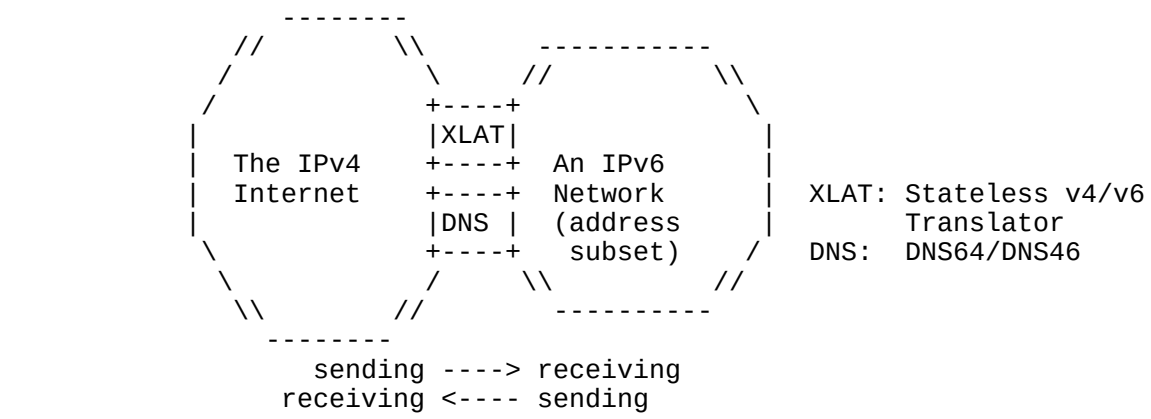

Figure 1: Stateless translation for Scenarios 1 and 2

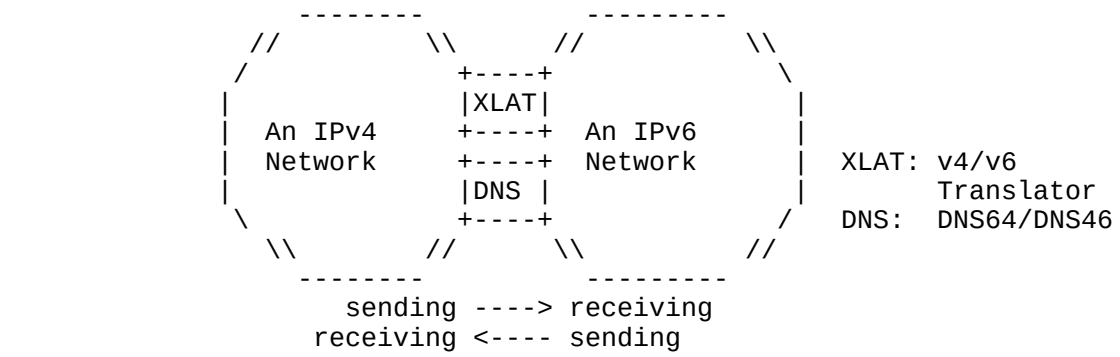

Figure 2: Stateless translator for Scenarios 5 and 6

3.3.2. Stateful Translation

For stateful translation, the translation state is maintained between IPv4 address/port pairs and IPv6 address/port pairs, enabling IPv6 systems to open sessions with IPv4 systems. See [RFC6145] and [RFC6146].

Stateful translator can be used for Scenarios 1, 3 and 5, i.e. it supports "An IPv6 network receiving multicast from the IPv4 Internet", "The IPv6 Internet receiving multicast from an IPv4 network" and "An IPv6 network receiving multicast from an IPv4 network".

In the stateful translation, an IPv6 network or the IPv6 Internet use any IPv6 addresses, while the IPv4 Internet or an IPv4 network can be represented by IPv4-mapped addresses.

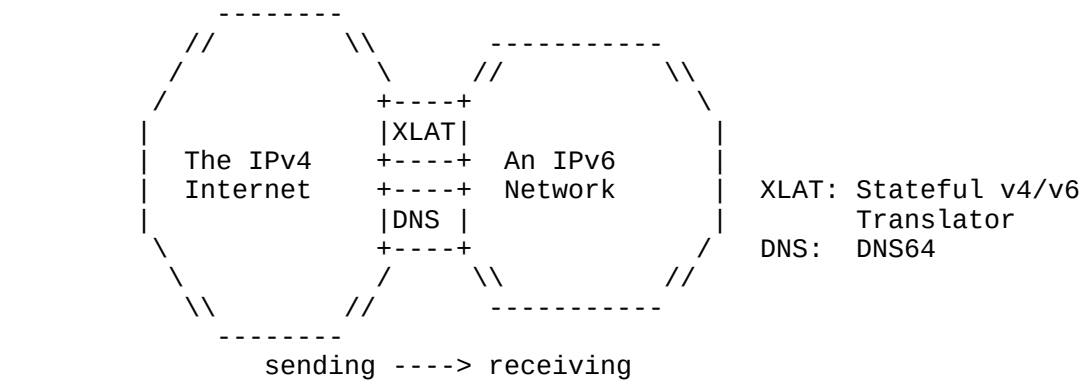

Figure 3: Stateful translator for Scenario 1

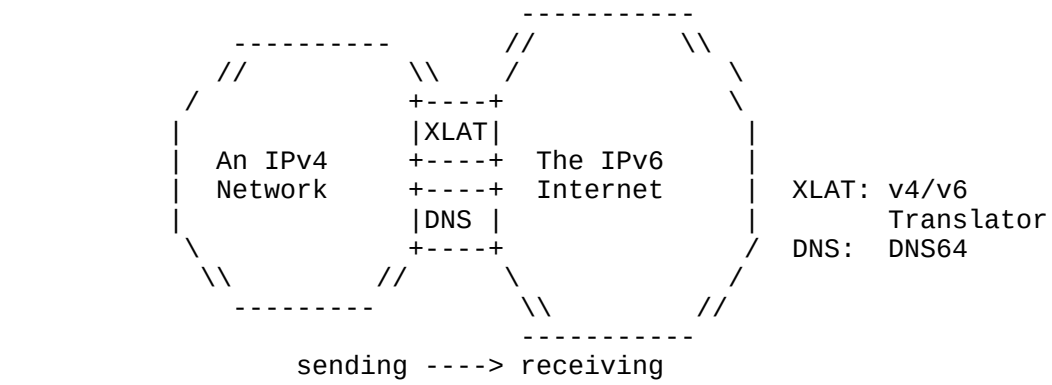

Figure 4: Stateful translator for Scenario 3

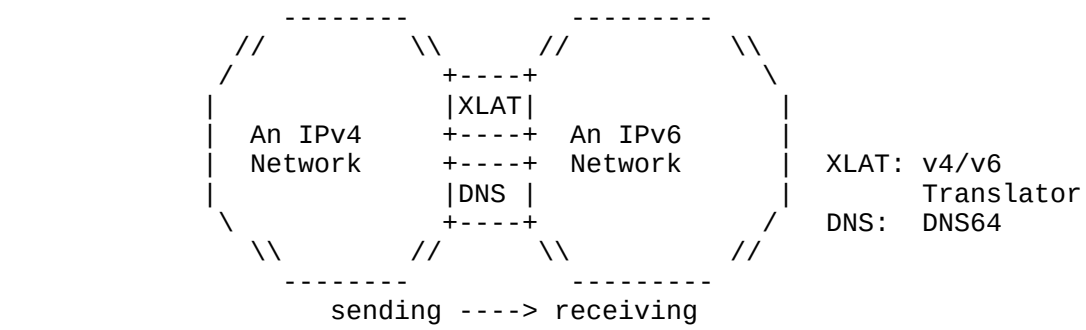

Figure 5: Stateful translator for Scenario 5

3.4. Application layer issues

The main application layer issue is perhaps how the applications learn what groups (or sources and groups) to join. For unicast, applications may often obtain addresses via DNS and a DNS-ALG. For multicast, DNS is usually not used, and there are a wide range of different ways applications learn addresses. It can be through configuration or user input, it can be URLs on a web page, it can be SDP files (via SAP or from web page or mail etc), or also via protocols like RTSP/SIP. It is no easy task to handle all of these possible methods using ALGs.

SDP is maybe the most common way for applications to learn which multicast addresses (and other parameters) to use in order to receive a multicast session. Inside the SDP files it is common to use literal IP addresses, but it is also possible to specify domain names. Applications would then query the DNS for the addresses, and a DNS-ALG could perform the necessary translation. There is however a problem with this.

Here is a typical SDP taken from RFC 2327:

```
v=0
o=mhandley 2890844526 2890842807 IN IP4 126.16.64.4
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.cs.ucl.ac.uk/staff/M.Handley/sdp.03.ps
e=mjh@isi.edu (Mark Handley)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
m=audio 49170 RTP/AVP 0
m=video 51372 RTP/AVP 31
m=application 32416 udp wb
a=orient:portrait
```

The line of interest here is "c=IN IP4 224.2.17.12/127". It is legal to use a domain name, this line would then become e.g. "c=IN IP4 mcast.example.com/127". The problem here is that the application is told to use IPv4. It will expect the name to resolve to an IPv4 address, and may ignore any IPv6 replies. One could argue that it would be incorrect to use IPv6, since IPv4 is specified. For DNS to solve our problem, we would need a new IP neutral SDP syntax, and applications would need to be updated to support it.

An alternative to rewriting addresses in the network is to make the applications aware of the translation and mappings in use. One approach could be for the source to create say SDP that includes both the original and the translated addresses. This may require use of techniques like CCAP [I-D.boucadair-mmusic-ccap] for specifying both IPv4 and IPv6 multicast addresses, allowing the receiver to choose which one to use. The other alternative would be for the receiving application to be aware of the translation and the mapping in use. This means that the receiving application can receive the original SDP, but then apply the mapping to those addresses.

As we just discussed, it may be useful for applications to perform the mappings. The next question is how they may learn those mappings. The easiest would be if there was a standard way used for all mappings, e.g. a well-known IPv6 prefix for embedding IPv4 addresses. But that does not work in all scenarios. There could be a way for applications to learn which prefix to use, see [I-D.wing-behave-learn-prefix]. But note that there may be different multicast prefixes depending on whether we are doing SSM or ASM and scope. In addition we need the unicast prefix for the multicast source addresses. Alternatively one could imagine applications requesting mappings for specific addresses on demand from the translator. The translator could have static mappings, or install

mappings as requested by applications.

An alternative to making applications aware of the translation and rewriting addresses as needed, could be to do translation in the API or stack, so that e.g. an application joins an IPv4 group, the API or stack rewrites that into IPv6 and sends the necessary MLD reports. When IPv6 packets arrive, the API/stack can rewrite those packets back to IPv4. This could allow legacy IPv4 applications to run on a dual-stack node (or IPv6-only with translation in the API) to receive IPv4 packets through an IPv6-only network. But in this case it might be better to just use tunneling.

3.5. Further Work

There are some special cases and scenarios that should be added to this document. One is addressing. Are there certain types of IPv6 multicast addresses that could make translation easier? What happens if there are multiple translators? And also more details on translation in the host, e.g. bump-in-the-stack or bump-in-the-API.

The document layout of the IPv4/IPv6 multicast translation should be presented in this document.

4. IANA Considerations

This document requires no IANA assignments.

5. Security Considerations

This requires more thought, but the author is not aware of any obvious security issues specific to multicast translation.

6. Acknowledgements

Dan Wing provided early feedback that helped shape this document. Dave Thaler also provided good feedback that unfortunately still has not been addressed in this document. See Section 3.5.

7. References

7.1. Normative References

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

[RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.

[RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.

[RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

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

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.

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

[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.

7.2. Informative References

[I-D.xli-behave-ivi]
Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition", draft-xli-behave-ivi-07 (work in progress), January 2010.

[I-D.venaas-behave-mcast46]
Venaas, S., Asaeda, H., SUZUKI, S., and T. Fujisaki, "An IPv4 - IPv6 multicast translator", draft-venaas-behave-mcast46-02 (work in progress), December 2010.

[I-D.wing-behave-learn-prefix]
Wing, D., "Learning the IPv6 Prefix of a Network's IPv6/ IPv4 Translator", draft-wing-behave-learn-prefix-04 (work in progress), October 2009.

[I-D.boucadair-mmusic-ccap]
Boucadair, M. and H. Kaplan, "Session Description Protocol (SDP) Connectivity Capability (CCAP) Attribute", draft-boucadair-mmusic-ccap-00 (work in progress), July 2009.

Authors' Addresses

Stig Venaas
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: stig@cisco.com

Xing Li
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 10-62785983
Email: xing@cernet.edu.cn

Congxiao Bao
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 10-62785983
Email: congxiao@cernet.edu.cn