

Internet Engineering Task Force (IETF)
Request for Comments: 8638
Category: Standards Track
ISSN: 2070-1721

M. Xu
Y. Cui
J. Wu
Tsinghua University
S. Yang
Shenzhen University
C. Metz
Cisco Systems
September 2019

IPv4 Multicast over an IPv6 Multicast in Software Mesh Networks

Abstract

During the transition to IPv6, there are scenarios where a backbone network internally running one IP address family (referred to as the internal IP or I-IP family) connects client networks running another IP address family (referred to as the external IP or E-IP family). In such cases, the I-IP backbone needs to offer both unicast and multicast transit services to the client E-IP networks.

This document describes a mechanism for supporting multicast across backbone networks where the I-IP and E-IP protocol families differ. The document focuses on the IPv4-over-IPv6 scenario, due to lack of real-world use cases for the IPv6-over-IPv4 scenario.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8638>.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 3
- 2. Requirements Language 5
- 3. Terminology 5
- 4. Scope 6
- 5. Mesh Multicast Mechanism 7
 - 5.1. Mechanism Overview 7
 - 5.2. Group Address Mapping 7
 - 5.3. Source Address Mapping 8
 - 5.4. Routing Mechanism 9
- 6. Control-Plane Functions of AFBR 10
 - 6.1. E-IP (*,G) and (S,G) State Maintenance 10
 - 6.2. I-IP (S',G') State Maintenance 10
 - 6.3. E-IP (S,G,rpt) State Maintenance 10
 - 6.4. Inter-AFBR Signaling 10
 - 6.5. SPT Switchover 13
 - 6.6. Other PIM Message Types 13
 - 6.7. Maintenance of Other PIM States 13
- 7. Data-Plane Functions of the AFBR 13
 - 7.1. Process and Forward Multicast Data 13
 - 7.2. TTL or Hop Count 14
 - 7.3. Fragmentation 14
- 8. Packet Format and Translation 14
- 9. Software Mesh Multicast Encapsulation 16
- 10. Security Considerations 16
- 11. IANA Considerations 16
- 12. Normative References 16
- Acknowledgements 18
- Authors' Addresses 18

1. Introduction

During the transition to IPv6, there are scenarios where a backbone network internally running one IP address family (referred to as the internal IP or I-IP family) connects client networks running another IP address family (referred to as the external IP or E-IP family).

One solution is to leverage the multicast functions inherent in the I-IP backbone to efficiently forward client E-IP multicast packets inside an I-IP core tree. The I-IP tree is rooted at one or more ingress Address Family Border Routers (AFBRs) [RFC5565] and branches out to one or more egress AFBRs.

[RFC4925] outlines the requirements for the software mesh scenario and includes support for multicast traffic. It is likely that client E-IP multicast sources and receivers will reside in different client E-IP networks connected to an I-IP backbone network. This requires the source-rooted or shared tree of the client E-IP to traverse the I-IP backbone network.

This could be accomplished by reusing the multicast VPN (MVPN) approach outlined in [RFC6513]. MVPN-like schemes can support the software mesh scenario and achieve a "many-to-one" mapping between the E-IP client multicast trees and the transit-core multicast trees. The advantage of this approach is that the number of trees in the I-IP backbone network scales less than linearly with the number of E-IP client trees. Corporate enterprise networks, and by extension multicast VPNs, have been known to run applications that create too many (S,G) states, which are source-specific states related to a specified multicast group [RFC7761] [RFC7899]. Aggregation at the edge contains the (S,G) states for customers' VPNs and these need to be maintained by the network operator. The disadvantage of this approach is the possibility of inefficient bandwidth and resource utilization when multicast packets are delivered to a receiving AFBR with no attached E-IP receivers.

[RFC8114] provides a solution for delivering IPv4 multicast services over an IPv6 network, but it mainly focuses on the DS-Lite scenario [RFC6333], where IPv4 addresses assigned by a broadband service provider are shared among customers. This document describes a detailed solution for the IPv4-over-IPv6 software mesh scenario, where client networks run IPv4 and the backbone network runs IPv6.

Internet-style multicast is somewhat different from the scenario in [RFC8114] in that the trees are source-rooted and relatively sparse. The need for multicast aggregation at the edge (where many customer

multicast trees are mapped to one or more backbone multicast trees) does not exist and to date has not been identified. Thus, the need for alignment between the E-IP and I-IP multicast mechanisms emerges.

[RFC5565] describes the "Software Mesh Framework". This document provides a more detailed description of how one-to-one mapping schemes ([RFC5565], Section 11.1) for IPv4-over-IPv6 multicast can be achieved.

Figure 1 shows an example of how a software mesh network can support multicast traffic. A multicast source S is located in one E-IP client network, while candidate E-IP group receivers are located in the same or different E-IP client networks that all share a common I-IP transit network. When E-IP sources and receivers are not local to each other, they can only communicate with each other through the I-IP core. There may be several E-IP sources for a single multicast group residing in different client E-IP networks. In the case of shared trees, the E-IP sources, receivers, and rendezvous points (RPs) might be located in different client E-IP networks. In the simplest case, a single operator manages the resources of the I-IP core, although the inter-operator case is also possible and so not precluded.

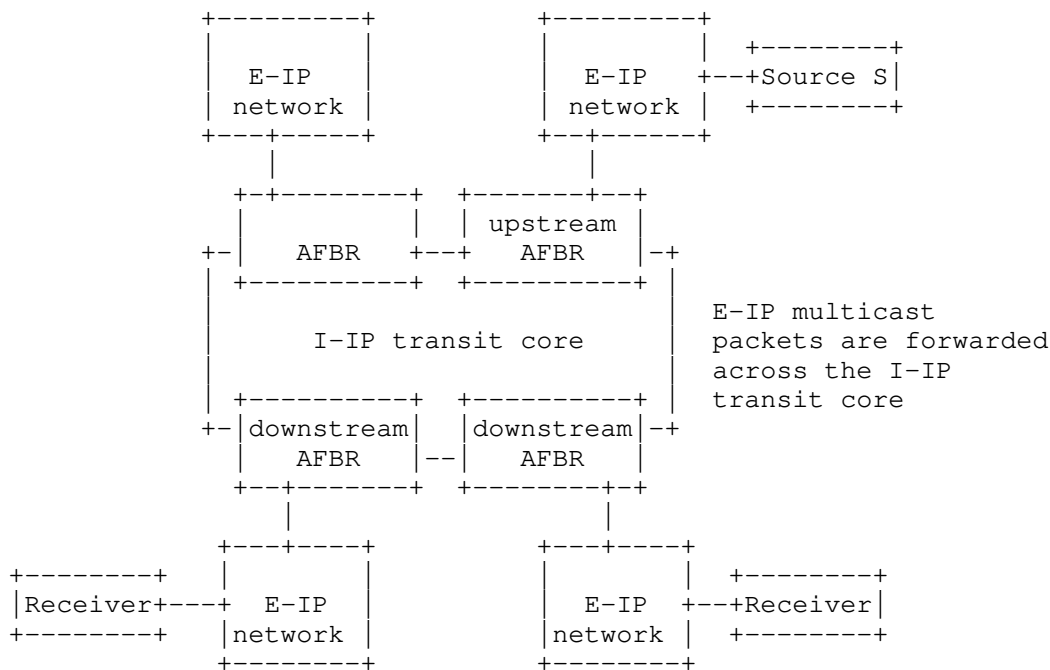


Figure 1: Software Mesh Multicast Framework

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology

The following terminology is used in this document.

- o Address Family Border Router (AFBR) - A router interconnecting two or more networks using different IP address families. Additionally, in the context of softwire mesh multicast, the AFBR runs E-IP and I-IP control planes to maintain E-IP and I-IP multicast states respectively and performs the appropriate encapsulation/decapsulation of client E-IP multicast packets for transport across the I-IP core. An AFBR will act as a source and/or receiver in an I-IP multicast tree.
- o Upstream AFBR: An AFBR that is closer to the source of a multicast data flow.
- o Downstream AFBR: An AFBR that is closer to a receiver of a multicast data flow.
- o I-IP (Internal IP): This refers to the IP address family that is supported by the core network. In this document, the I-IP is IPv6.
- o E-IP (External IP): This refers to the IP address family that is supported by the client network(s) attached to the I-IP transit core. In this document, the E-IP is IPv4.
- o I-IP core tree: A distribution tree rooted at one or more AFBR source nodes and branched out to one or more AFBR leaf nodes. An I-IP core tree is built using standard IP or MPLS multicast signaling protocols (in this document, we focus on IP multicast) operating exclusively inside the I-IP core network. An I-IP core tree is used to forward E-IP multicast packets belonging to E-IP trees across the I-IP core. Another name for an I-IP core tree is multicast or multipoint softwire.
- o E-IP client tree: A distribution tree rooted at one or more hosts or routers located inside a client E-IP network and branched out to one or more leaf nodes located in the same or different client E-IP networks.

- o uPrefix64: The /96 unicast IPv6 prefix for constructing an IPv4-embedded IPv6 unicast address [RFC8114].
- o mPrefix64: The /96 multicast IPv6 prefix for constructing an IPv4-embedded IPv6 multicast address [RFC8114].
- o PIMv4, PIMv6: Refer to [RFC8114].
- o Inter-AFBR signaling: A mechanism used by downstream AFBRs to send PIMv6 messages to the upstream AFBR.

4. Scope

This document focuses on the IPv4-over-IPv6 scenario, as shown in the following diagram.




Figure 2: IPv4-over-IPv6 Scenario

In Figure 2, the E-IP client networks run IPv4, and the I-IP core runs IPv6.

Because of the much larger IPv6 group address space, the client E-IP tree can be mapped to a specific I-IP core tree. This simplifies operations on the AFBR because it becomes possible to algorithmically

map an IPv4 group/source address to an IPv6 group/source address and vice versa.

The IPv4-over-IPv6 scenario is an emerging requirement as network operators build out native IPv6 backbone networks. These networks support native IPv6 services and applications, but, in many cases, support for legacy IPv4 unicast and multicast services will also need to be accommodated.

5. Mesh Multicast Mechanism

5.1. Mechanism Overview

Routers in the client E-IP networks have routes to all other client E-IP networks. Through PIMv4 messages, E-IP hosts and routers have discovered or learnt of IPv4 addresses that are in (S,G) or (*,G) state [RFC7761]. Any I-IP multicast state instantiated in the core is referred to as (S',G') or (*,G') and is separated from E-IP multicast state.

Suppose a downstream AFBR receives an E-IP PIM Join/Prune message from the E-IP network for either an (S,G) tree or a (*,G) tree. The AFBR translates the PIMv4 message into a PIMv6 message with the latter being directed towards the I-IP IPv6 address of the upstream AFBR. When the PIMv6 message arrives at the upstream AFBR, it is translated back into a PIMv4 message. The result of these actions is the construction of E-IP trees and a corresponding I-IP tree in the I-IP network. An example of the packet format and translation is provided in Section 8.

In this case, it is incumbent upon the AFBRs to perform PIM message conversions in the control plane and IP group address conversions or mappings in the data plane. The AFBRs perform an algorithmic, one-to-one mapping of IPv4 to IPv6.

5.2. Group Address Mapping

A simple algorithmic mapping between IPv4 multicast group addresses and IPv6 group addresses is performed. Figure 3 is provided as a reminder of the format:




Figure 3: IPv4-Embedded IPv6 Multicast Address Format

An IPv6 multicast prefix (mPrefix64) is provisioned on each AFBR. AFBRs will prepend the prefix to an IPv4 multicast group address when translating it to an IPv6 multicast group address.

The construction of the mPrefix64 for Source-Specific Multicast (SSM) is the same as the construction of the mPrefix64 described in Section 5 of [RFC8114].

With this scheme, each IPv4 multicast address can be mapped to an IPv6 multicast address (with the assigned prefix), and each IPv6 multicast address with the assigned prefix can be mapped to an IPv4 multicast address. The group address translation algorithm is specified in Section 5.2 of [RFC8114].

5.3. Source Address Mapping

There are two kinds of multicast: Any-Source Multicast (ASM) and SSM. Considering that the I-IP network and E-IP network may support different kinds of multicast, the source address translation rules needed to support all possible scenarios may become very complex. But since SSM can be implemented with a strict subset of the PIM-SM protocol mechanisms [RFC7761], we can treat the I-IP core as SSM-only to make it as simple as possible. There then remain only two scenarios to be discussed in detail:

- o E-IP network supports SSM

One possible way to make sure that the translated PIMv6 message reaches the upstream AFBR is to set S' to a virtual IPv6 address that leads to the upstream AFBR. The unicast address translation should be achieved according to [RFC6052].

- o E-IP network supports ASM

The (S,G) source list entry and the (*,G) source list entry differ only in that the latter has both the WildCard (WC) and RPT bits of the Encoded-Source-Address set, while with the former, the bits are cleared. (See Section 4.9.5.1 of [RFC7761].) As a result, the source list entries in (*,G) messages can be translated into source list entries in (S',G') messages by clearing both the WC and RPT bits at downstream AFBRs, and vice versa for the reverse translation at upstream AFBRs.

5.4. Routing Mechanism

With mesh multicast, PIMv6 messages originating from a downstream AFBR need to be propagated to the correct upstream AFBR, and every AFBR needs the /96 prefix in the IPv4-embedded IPv6 source address format [RFC6052].

To achieve this, every AFBR MUST announce the address of one of its E-IPv4 interfaces in the "v4" field [RFC6052] alongside the corresponding uPrefix46. The announcement MUST be sent to the other AFBRs through Multiprotocol BGP (MBGP) [RFC4760]. Every uPrefix64 that an AFBR announces MUST be unique. "uPrefix64" is an IPv6 prefix, and the distribution mechanism is the same as the traditional mesh unicast scenario.

As the "v4" field is an E-IP address, and BGP messages are not tunneled through softwires or any other mechanism specified in [RFC5565], AFBRs MUST be able to transport and encode/decode BGP messages that are carried over the I-IP, and whose Network Layer Reachability Information (NLRI) and next hop (NH) are of the E-IP address family.

In this way, when a downstream AFBR receives an E-IP PIM (S,G) message, it can translate this message into (S',G') by looking up the IP address of the corresponding AFBR's E-IP interface. Since the uPrefix64 of S' is unique and is known to every router in the I-IP network, the translated message will be forwarded to the corresponding upstream AFBR, and the upstream AFBR can translate the message back to (S,G).

When a downstream AFBR receives an E-IP PIM (*,G) message, S' can be generated with the "source address" field set to * (the wildcard value). The translated message will be forwarded to the corresponding upstream AFBR. Every PIM router within a PIM domain MUST be able to map a particular multicast group address to the same RP when the source address is set to the wildcard value. (See Section 4.7 of [RFC7761].) So, when the upstream AFBR checks the "source address" field of the message, it finds the IPv4 address of the RP and ascertains that this was originally a (*,G) message. This is then translated back to the (*,G) message and processed.

6. Control-Plane Functions of AFBR

AFBRs are responsible for the functions detailed in the subsections that follow.

6.1. E-IP (*,G) and (S,G) State Maintenance

E-IP (*,G) and (S,G) state maintenance for an AFBR is the same as E-IP (*,G) and (S,G) state maintenance for a multicast AFTR (mAFTR) described in Section 7.2 of [RFC8114].

6.2. I-IP (S',G') State Maintenance

It is possible that the I-IP transit core runs another, non-transit, I-IP PIM-SSM instance. Since the translated source address starts with the unique "Well-Known" prefix or the ISP-defined prefix that MUST NOT be used by another service provider, mesh multicast will not influence non-transit PIM-SSM multicast at all. When an AFBR receives an I-IP (S',G') message, it MUST check S'. If S' starts with the unique prefix, then the message is actually a translated E-IP (S,G) or (*,G) message, and the AFBR translates this message back to a PIMv4 message and processes it.

6.3. E-IP (S,G,rpt) State Maintenance

When an AFBR wishes to propagate a Join/Prune(S,G,rpt) message [RFC7761] to an I-IP upstream router, the AFBR MUST operate as specified in Sections 6.5 and 6.6.

6.4. Inter-AFBR Signaling

Assume that one downstream AFBR has joined an RPT of (*,G) and a shortest path tree (SPT) of (S,G) and decided to perform an SPT switchover. (See Section 4.2.1 of [RFC7761].) According to [RFC7761], it should propagate a Prune(S,G,rpt) message along with the periodic Join(*,G) message upstream towards the RP. However, routers in the I-IP transit core do not process (S,G,rpt) messages since the I-IP transit core is treated as SSM only. As a result, the downstream AFBR is unable to prune S from this RPT, so it will receive two copies of the same data for (S,G). In order to solve this problem, we introduce a new mechanism for downstream AFBRs to inform upstream AFBRs of pruning any given S from an RPT.

When a downstream AFBR wishes to propagate an (S,G,rpt) message upstream, it SHOULD encapsulate the (S,G,rpt) message, then send the encapsulated unicast message to the corresponding upstream AFBR, which we call "RP'".

When RP' receives this encapsulated message, it MUST decapsulate the message as in the unicast scenario and retrieve the original (S,G,rpt) message. The incoming interface of this message may be different from the outgoing interface that propagates multicast data to the corresponding downstream AFBR, and there may be other downstream AFBRs that need to receive multicast data for (S,G) from this incoming interface, so RP' should not simply process this message as specified in [RFC7761] on the incoming interface.

To solve this problem, we introduce an "interface agent" to process all the encapsulated (S,G,rpt) messages the upstream AFBR receives. The interface agent's RP' should prune S from the RPT of group G when no downstream AFBR is subscribed to receive multicast data for (S,G) along the RPT.

In this way, we ensure that downstream AFBRs will not miss any multicast data that they need. The cost of this is that multicast data for (S,G) will be duplicated along the RPT received by AFBRs affected by the SPT switchover, if at least one downstream AFBR exists that has not yet sent Prune(S,G,rpt) messages to the upstream AFBR.

In certain deployment scenarios (e.g., if there is only a single downstream router), the interface agent function is not required.

The mechanism used to achieve this is left to the implementation. The following diagram provides one possible solution for an "interface agent" implementation:

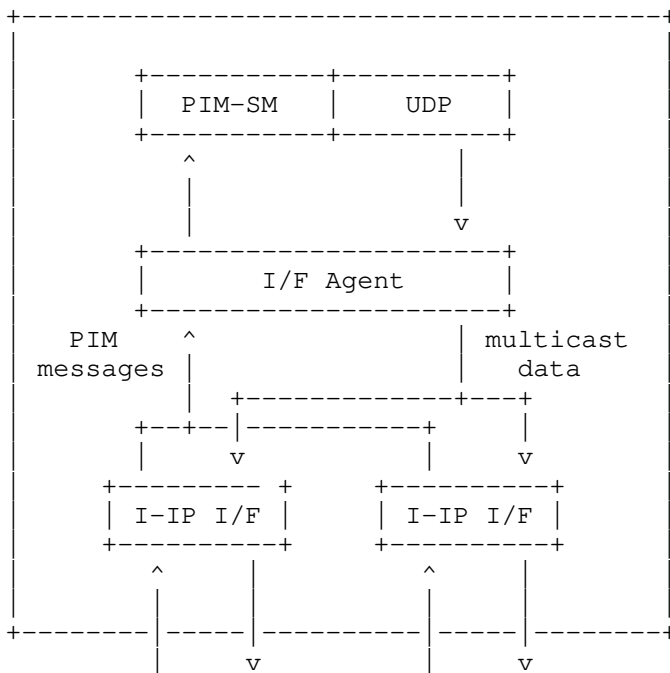


Figure 4: Interface Agent Implementation Example

Figure 4 shows an example of an interface agent implementation using UDP encapsulation. The interface agent has two responsibilities: In the control plane, it should work as a real interface that has joined (*,G), representing all the I-IP interfaces that are outgoing interfaces of the (*,G) state machine, and it should process the (S,G,rpt) messages received from all the I-IP interfaces.

The interface agent maintains downstream (S,G,rpt) state machines for every downstream AFBR, and it submits Prune(S,G,rpt) messages to the PIM-SM module only when every (S,G,rpt) state machine is in the Prune(P) or PruneTmp(P') state, which means that no downstream AFBR is subscribed to receive multicast data for (S,G) along the RPT of G. Once a (S,G,rpt) state machine changes to NoInfo (NI) state, which means that the corresponding downstream AFBR has switched to receive multicast data for (S,G) along the RPT again, the interface agent MUST send a Join(S,G,rpt) to the PIM-SM module immediately.

In the data plane, upon receiving a multicast data packet, the interface agent MUST encapsulate it at first, then propagate the encapsulated packet from every I-IP interface.

NOTICE: It is possible that an E-IP neighbor of RP' has joined the RPT of G, so the per-interface state machine for receiving E-IP Join/Prune(S,G,rpt) messages should be preserved.

6.5. SPT Switchover

After a new AFBR requests the receipt of traffic destined for a multicast group, it will receive all the data from the RPT at first. At this time, every downstream AFBR will receive multicast data from any source from this RPT, in spite of whether they have switched over to an SPT or not.

To minimize this redundancy, it is recommended that every AFBR's SwitchToSptDesired(S,G) function employs the "switch on first packet" policy. In this way, the delay in switchover to SPT is kept as small as possible, and after the moment that every AFBR has performed the SPT switchover for every S of group G, no data will be forwarded in the RPT of G, thus no more unnecessary duplication will be produced.

6.6. Other PIM Message Types

In addition to Join or Prune, other message types exist, including Register, Register-Stop, Hello and Assert. Register and Register-Stop messages are sent by unicast, while Hello and Assert messages are only used between directly linked routers to negotiate with each other. It is not necessary to translate these for forwarding, thus the processing of these messages is out of scope for this document.

6.7. Maintenance of Other PIM States

In addition to states mentioned above, other states exist, including (*,*,RP) and I-IP (*,G') state. Since we treat the I-IP core as SSM only, the maintenance of these states is out of scope for this document.

7. Data-Plane Functions of the AFBR

7.1. Process and Forward Multicast Data

Refer to Section 7.4 of [RFC8114]. If there is at least one outgoing interface whose IP address family is different from the incoming interface, the AFBR MUST encapsulate this packet with mPrefix64-derived and uPrefix64-derived IPv6 addresses to form an IPv6 multicast packet.

7.2. TTL or Hop Count

Upon encapsulation, the TTL and hop count in the outer header SHOULD be set by policy. Upon decapsulation, the TTL and hop count in the inner header SHOULD be modified by policy; it MUST NOT be incremented and it MAY be decremented to reflect the cost of tunnel forwarding. Besides, processing of TTL and hop count information in protocol headers depends on the tunneling technology, which is out of scope of this document.

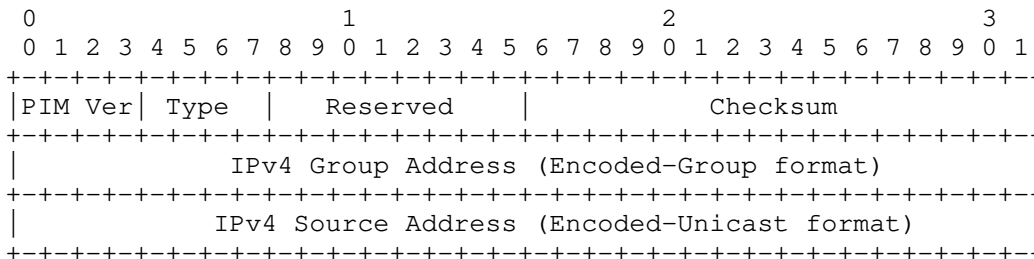
7.3. Fragmentation

The encapsulation performed by an upstream AFBR will increase the size of packets. As a result, the outgoing I-IP link MTU may not accommodate the larger packet size. It is not always possible for core operators to increase the MTU of every link, thus source fragmentation after encapsulation and reassembling of encapsulated packets MUST be supported by AFBRs [RFC5565]. Path MTU Discovery (PMTUD) [RFC8201] SHOULD be enabled, and ICMPv6 packets MUST NOT be filtered in the I-IP network. Fragmentation and tunnel configuration considerations are provided in Section 8 of [RFC5565]. The detailed procedure can be referred in Section 7.2 of [RFC2473].

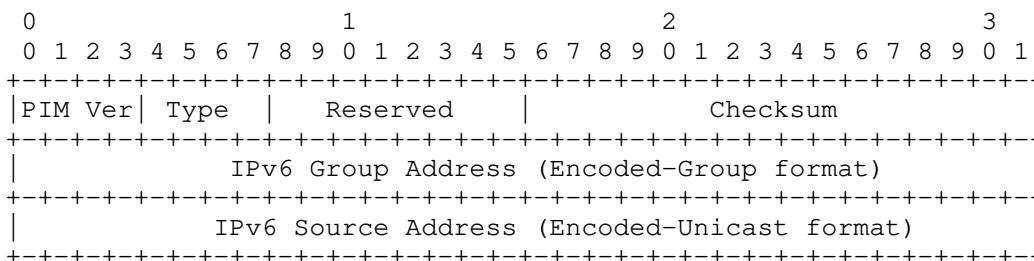
8. Packet Format and Translation

Because the PIM-SM specification is independent of the underlying unicast routing protocol, the packet format in Section 4.9 of [RFC7761] remains the same, except that the group address and source address MUST be translated when traversing an AFBR.

For example, Figure 5 shows the register-stop message format in the IPv4 and IPv6 address families.



(a) IPv4 Register-Stop Message Format



(b) IPv6 Register-Stop Message Format

Figure 5: Register-Stop Message Format

In Figure 5, the semantics of fields "PIM Ver", "Type", "Reserved", and "Checksum" are specified in Section 4.9 of [RFC7761].

IPv4 Group Address (Encoded-Group format): The encoded-group format of the IPv4 group address described in Section 4.9.1 of [RFC7761].

IPv4 Source Address (Encoded-Group format): The encoded-unicast format of the IPv4 source address described in Section 4.9.1 of [RFC7761].

IPv6 Group Address (Encoded-Group format): The encoded-group format of the IPv6 group address described in Section 5.2.

IPv6 Source Address (Encoded-Group format): The encoded-unicast format of the IPv6 source address described in Section 5.3.

9. Softwire Mesh Multicast Encapsulation

Softwire mesh multicast encapsulation does not require the use of any one particular encapsulation mechanism. Rather, it MUST accommodate a variety of different encapsulation mechanisms and allow the use of encapsulation mechanisms mentioned in [RFC4925]. Additionally, all of the AFBRs attached to the I-IP network MUST implement the same encapsulation mechanism and follow the requirements mentioned in Section 8 of [RFC5565].

10. Security Considerations

The security concerns raised in [RFC4925] and [RFC7761] are applicable here.

The additional workload associated with some schemes, such as interface agents, could be exploited by an attacker to perform a DDoS attack.

Compared with [RFC4925], the security concerns should be considered more carefully: An attacker could potentially set up many multicast trees in the edge networks, causing too many multicast states in the core network. To defend against these attacks, BGP policies SHOULD be carefully configured, e.g., AFBRs only accept Well-Known prefix advertisements from trusted peers. Besides, cryptographic methods for authenticating BGP sessions [RFC7454] could be used.

11. IANA Considerations

This document has no IANA actions.

12. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <<https://www.rfc-editor.org/info/rfc2473>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

- [RFC4925] Li, X., Ed., Dawkins, S., Ed., Ward, D., Ed., and A. Durand, Ed., "Software Problem Statement", RFC 4925, DOI 10.17487/RFC4925, July 2007, <<https://www.rfc-editor.org/info/rfc4925>>.
- [RFC5565] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Software Mesh Framework", RFC 5565, DOI 10.17487/RFC5565, June 2009, <<https://www.rfc-editor.org/info/rfc5565>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <<https://www.rfc-editor.org/info/rfc6333>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <<https://www.rfc-editor.org/info/rfc7454>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC7899] Morin, T., Ed., Litkowski, S., Patel, K., Zhang, Z., Kebler, R., and J. Haas, "Multicast VPN State Damping", RFC 7899, DOI 10.17487/RFC7899, June 2016, <<https://www.rfc-editor.org/info/rfc7899>>.
- [RFC8114] Boucadair, M., Qin, C., Jacquenet, C., Lee, Y., and Q. Wang, "Delivery of IPv4 Multicast Services to IPv4 Clients over an IPv6 Multicast Network", RFC 8114, DOI 10.17487/RFC8114, March 2017, <<https://www.rfc-editor.org/info/rfc8114>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed.,
"Path MTU Discovery for IP version 6", STD 87, RFC 8201,
DOI 10.17487/RFC8201, July 2017,
<<https://www.rfc-editor.org/info/rfc8201>>.

Acknowledgements

Wenlong Chen, Xuan Chen, Alain Durand, Yiu Lee, Jacni Qin, and Stig Venaas provided useful input to this document.

Authors' Addresses

Mingwei Xu
Tsinghua University
Department of Computer Science
Beijing 100084
China

Phone: +86-10-6278-5822
Email: xumw@tsinghua.edu.cn

Yong Cui
Tsinghua University
Department of Computer Science
Beijing 100084
China

Phone: +86-10-6278-5822
Email: cuiyong@tsinghua.edu.cn

Jianping Wu
Tsinghua University
Department of Computer Science
Beijing 100084
China

Phone: +86-10-6278-5983
Email: jianping@cernet.edu.cn

Shu Yang
Shenzhen University
South Campus
Shenzhen 518060
China

Phone: +86-755-2653-4078
Email: yang.shu@szu.edu.cn

Chris Metz
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
United States of America

Phone: +1-408-525-3275
Email: chmetz@cisco.com