

Softwire WG
Internet-Draft
Intended status: Standards Track
Expires: May 19, 2018

M. Xu
Y. Cui
J. Wu
S. Yang
Tsinghua University
C. Metz
G. Shepherd
Cisco Systems
November 15, 2017

IPv4 Multicast over an IPv6 Multicast in Softwire Mesh Network
draft-ietf-softwire-mesh-multicast-19

Abstract

During IPv6 transition, there will be scenarios where a backbone network running one IP address family internally (referred to as internal IP or I-IP), while the attached client networks running another IP address family (referred to as external IP or E-IP). The I-IP backbone should offer both unicast and multicast transit services to the client E-IP networks.

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

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 <https://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 19, 2018.

Copyright Notice

Copyright (c) 2017 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
1.1.	Requirements Language	4
2.	Terminology	4
3.	Scope	6
4.	Mesh Multicast Mechanism	7
4.1.	Mechanism Overview	8
4.2.	Group Address Mapping	8
4.3.	Source Address Mapping	9
4.4.	Routing Mechanism	9
5.	Control Plane Functions of AFBR	10
5.1.	E-IP (*,G) and (S,G) State Maintenance	10
5.2.	I-IP (S',G') State Maintenance	10
5.3.	E-IP (S,G,rpt) State Maintenance	11
5.4.	Inter-AFBR Signaling	11
5.5.	SPT Switchover	13
5.6.	Other PIM Message Types	13
5.7.	Other PIM States Maintenance	13
6.	Data Plane Functions of the AFBR	13
6.1.	Process and Forward Multicast Data	14
6.2.	TTL	14
6.3.	Fragmentation	14
7.	Packet Format and Translation	14
8.	Softwire Mesh Multicast Encapsulation	15
9.	Security Considerations	16
10.	IANA Considerations	16
11.	References	16
11.1.	Normative References	16
11.2.	Informative References	17
Appendix A.	Acknowledgements	17
	Authors' Addresses	17

1. Introduction

During IPv6 transition, there will be scenarios where a backbone network running one IP address family internally (referred to as internal IP or I-IP) will provide transit services to attached client networks running another IP address family (referred to as external IP or E-IP).

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, which 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 softwires 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 client E-IP source-rooted or shared tree to traverse the I-IP backbone network.

One method of accomplishing this is to re-use the multicast VPN approach outlined in [[RFC6513](#)]. MVPN-like schemes can support the softwire 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 [[RFC7899](#)]. Aggregation at the edge contains the (S,G) states for customer's 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 DS-lite [[RFC6333](#)] scenario. This document describes a detailed solution for IPv4-over-IPv6 softwire mesh scenario, where client networks run IPv4 but the backbone network runs IPv6.

Internet-style multicast is somewhat different 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 into one or more backbone multicast trees) does not exist and to date has not been identified. Thus the need for a basic or closer alignment with E-IP and I-IP multicast procedures emerges.

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

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\]](#).

2. Terminology

Figure 1 shows an example of how a softwire 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.

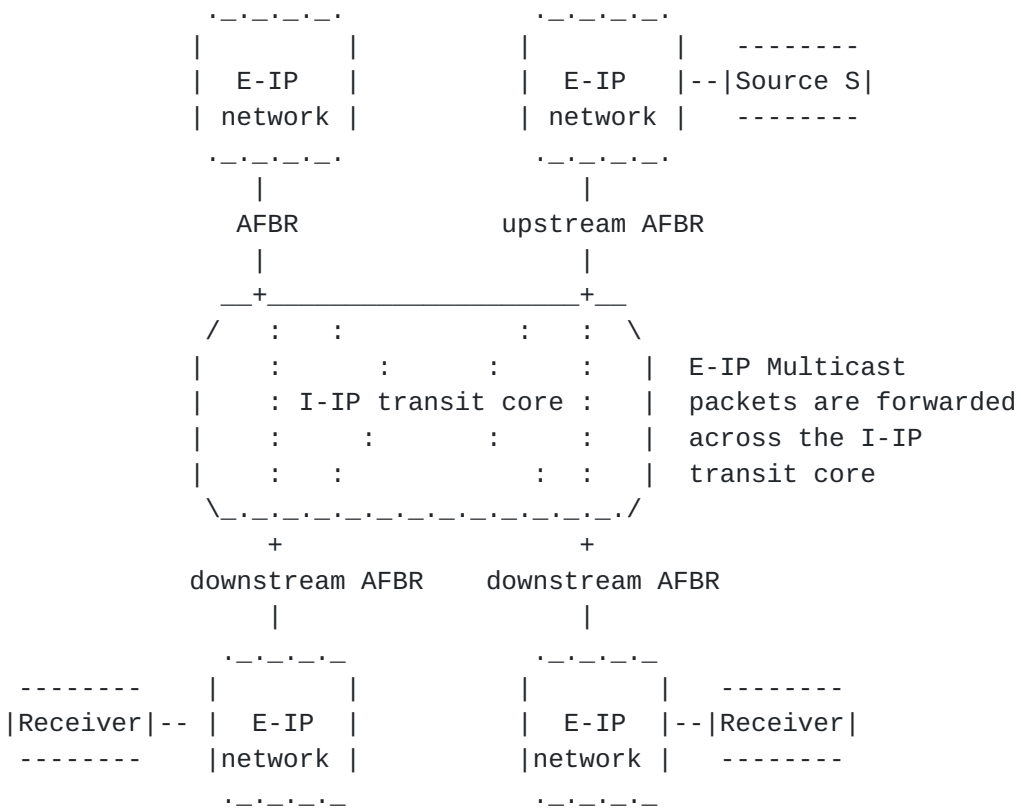


Figure 1: Softwire Mesh Multicast Framework

Terminology used in this document:

- o Address Family Border Router (AFBR) - A router interconnecting two or more networks using different IP address families. It MUST support functions specified in [\[RFC5565\]](#). Besides, 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 located on the upper reaches of a multicast data flow.

- o Downstream AFBR: An AFBR that is located on the lower reaches of a multicast data flow.

- o I-IP (Internal IP): This refers to IP address family (i.e., either IPv4 or IPv6) that is supported by the core network.

- o E-IP (External IP): This refers to the IP address family (i.e. either IPv4 or IPv6) that is supported by the client network(s) attached to the I-IP transit core.
- 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 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 software.
- 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 uPrefix46: The /96 unicast IPv6 prefix for constructing an IPv4-embedded IPv6 unicast address [[RFC6052](#)].
- o mPrefix46: The /96 multicast IPv6 prefix for constructing an IPv4-embedded IPv6 multicast address.
- o PIMv4, PIMv6: refer to [[RFC8114](#)].
- o Inter-AFBR signaling: A mechanism used by downstream AFBRs to send PIMv6 messages to the upstream AFBR.

3. Scope

This document focuses on IPv4-over-IPv6 scenario, the following diagram shows the scenario.

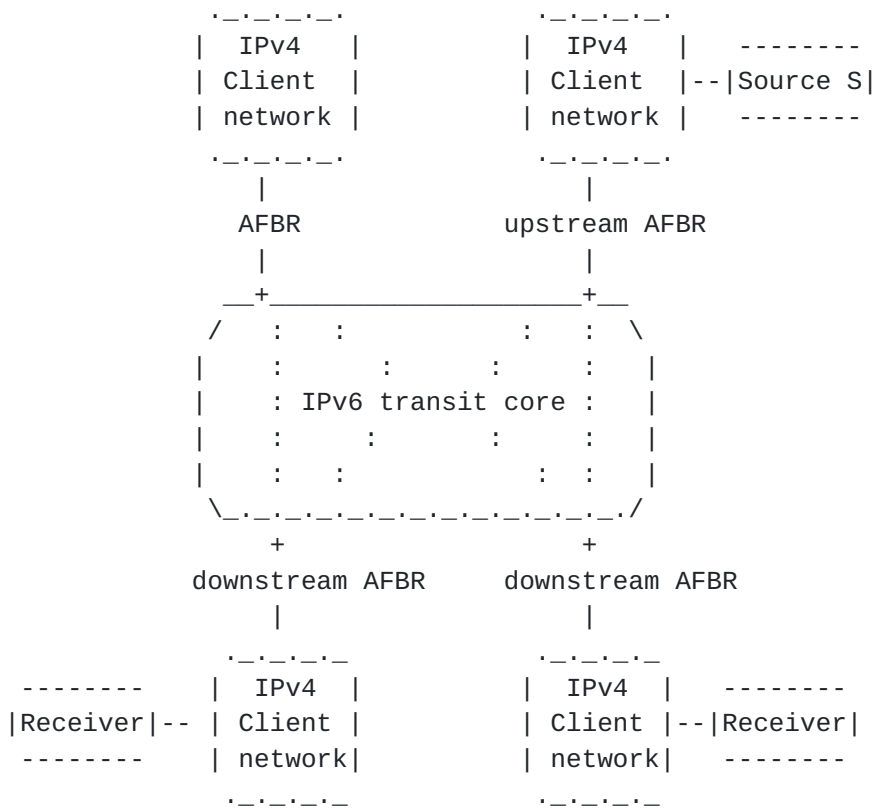


Figure 2: IPv4-over-IPv6 Scenario

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-IPv4 tree can be mapped to a specific I-IPv6 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.

4. Mesh Multicast Mechanism

4.1. Mechanism Overview

Routers in the client E-IPv4 networks have routes to all other client E-IPv4 networks. Through PIMv4 messages, E-IPv4 hosts and routers have discovered or learnt of (S,G) or (*,G) IPv4 addresses. Any I-IPv6 multicast state instantiated in the core is referred to as (S',G') or (*,G') and is certainly separated from E-IPv4 multicast state.

Suppose a downstream AFBR receives an E-IPv4 PIM Join/Prune message from the E-IPv4 network for either an (S,G) tree or a (*,G) tree. The AFBR translates the PIMv4 message into an 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 an PIMv4 message. The result of these actions is the construction of E-IPv4 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.

4.2. Group Address Mapping

For the IPv4-over-IPv6 scenario, a simple algorithmic mapping between IPv4 multicast group addresses and IPv6 group addresses is performed. Figure 4 shows the remainder of the format:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| 0-----32--40--48--56--64--72--80--88--96-----127|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               mPrefix46                      |group address |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 4: IPv4-Embedded IPv6 Multicast Address Format

An IPv6 multicast prefix (mPrefix46) 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 mPrefix46 for 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 into an IPv6 multicast address (with the assigned prefix), and each IPv6 multicast address with the assigned prefix can be mapped into an IPv4 multicast address. The group address translation algorithm can be referred in [Section 5.2 of \[RFC8114\]](#).

4.3. Source Address Mapping

There are two kinds of multicast modes: 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 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 only differ in that the latter has both the WC and RPT bits of the Encoded-Source-Address set, while the former is all cleared (See [Section 4.9.5.1 of \[RFC7761\]](#)). So we can translate source list entries in (*,G) messages into source list entries in (S'G') messages by applying the format specified in Figure 5 and clearing both the WildCard (WC) and RPT bits at downstream AFBRs, and vice-versa for the reverse translation at upstream AFBRs.

4.4. Routing Mechanism

In the mesh multicast scenario, extra multicast routing information is REQUIRED to be distributed among AFBRs to make sure that the PIMv6 messages that a downstream AFBR propagates reach the right upstream AFBR.

Every AFBR MUST know the /32 prefix in "IPv4-Embedded IPv6 Virtual Source Address Format". To achieve this, every AFBR should announce one of its E-IPv4 interfaces in the "v4" field, and the corresponding uPrefix46. The announcement SHOULD be sent to the other AFBRs

through MBGP [[RFC4760](#)]. Since every IP address of upstream AFBR's E-IPv4 interface is different from each other, every uPrefix46 that AFBR announces MUST be different, and uniquely identifies each AFBR. "uPrefix46" is an IPv6 prefix, and the distribution mechanism is the same as the traditional mesh unicast scenario. But "v4" field is an E-IPv4 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 I-IPv6, whose NLRI and NH are of E-IPv4 address family.

In this way, when a downstream AFBR receives an E-IPv4 PIM (S,G) message, it can translate this message into (S',G') by looking up the IP address of the corresponding AFBR's E-IPv4 interface. Since the uPrefix46 of S' is unique, and is known to every router in the I-IPv6 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-IPv4 PIM (*,G) message, S' can be generated according to the format specified in Figure 4, with "source address" field set to *(the IPv4 address of RP). The translated message will be forwarded to the corresponding upstream AFBR. Since every PIM router within a PIM domain MUST be able to map a particular multicast group address to the same RP (see [Section 4.7 of \[RFC7761\]](#)), when the upstream AFBR checks the "source address" field of the message, it finds the IPv4 address of the RP, and ascertains that this is originally a (*,G) message. This is then translated back to the (*,G) message and processed.

5. Control Plane Functions of AFBR

AFBRs are responsible for the following functions:

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

E-IP (*,G) and (S,G) state maintenance on AFBR is the same as E-IP (*,G) and (S,G) state maintenance on mAFTR described in [Section 7.2 of \[RFC8114\]](#)

5.2. 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 SHOULD NOT be used by other 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 MUST translate this message back to PIMv4 message and process it.

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

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

5.4. Inter-AFBR Signaling

Assume that one downstream AFBR has joined a RPT of (*,G) and a SPT of (S,G), and decide to perform an SPT switchover. According to [\[RFC7761\]](#), it SHOULD propagate a Prune(S,G,rpt) message along with the periodical Join(*,G) message upstream towards 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 of (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 a (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 SHOULD 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 to the outgoing interface which propagates multicast data to the corresponding downstream AFBR, and there may be other downstream AFBRs that need to receive multicast data of (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 of (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 of (S,G) will be duplicated along the RPT received by SPT-switched-over downstream AFBRs, 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 7: Interface Agent Implementation Example

Figure 7 shows an example of 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 of all the I-IP interfaces which are outgoing interfaces of the $(*,G)$ state machine, and process the (S,G,rpt) messages received from all the I-IP interfaces.

The interface agent maintains downstream (S,G,rpt) state machines of every downstream AFBR, and submits Prune (S,G,rpt) messages to the PIM-SM module only when every (S,G,rpt) state machine is at Prune(P) or PruneTmp(P') state, which means that no downstream AFBR is subscribed to receive multicast data of (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 of (S,G) along the RPT again, the interface agent SHOULD 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 SHOULD 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' that 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.

5.5. 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 of some source(s) 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.

5.6. 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.

5.7. Other PIM States Maintenance

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.

6. Data Plane Functions of the AFBR

6.1. 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 mPrefix46-derived and uPrefix46-derived IPv6 address to form an IPv6 multicast packet.

6.2. TTL

Processing of TTL information in protocol headers depends on the tunneling technology, and it is out of scope of this document.

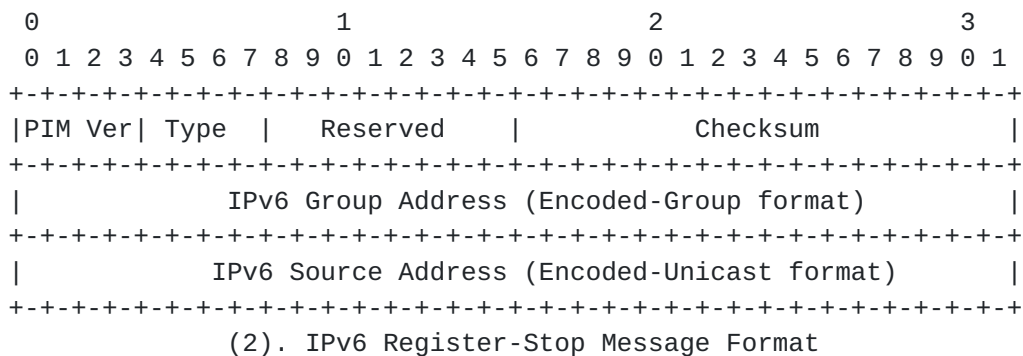
6.3. Fragmentation

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. As it is not always possible for core operators to increase the MTU of every link. Fragmentation after encapsulation and reassembling of encapsulated packets MUST be supported by AFBRs [[RFC5565](#)].

7. 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 AFBR.

For example, Figure 8 shows the register-stop message format in IPv4 and IPv6 address family.



In Figure 8, the semantics of fields "PIM Ver", "Type", "Reserved", and "Checksum" remain the same.

IPv4 Source Address (Encoded-Group format): The encoded-unicast format of the IPv4 source address described in [Section 4.3](#).

IPv6 Source Address (Encoded-Group format): The encoded-unicast format of the IPv6 source address described in [Section 4.3](#).

Software 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.

9. Security Considerations

The security concerns raised in [RFC4925] and [RFC7761] are applicable here. In addition, the additional workload associated with some schemes could be exploited by an attacker to perform a out 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.

10. IANA Considerations

This document includes no request to IANA.

11. References

11.1. 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>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [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., "Softwire 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, "Softwire 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>>.
- [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>>.

11.2. Informative References

- [RFC7371] Boucadair, M. and S. Venaas, "Updates to the IPv6 Multicast Addressing Architecture", [RFC 7371](#), DOI 10.17487/RFC7371, September 2014, <<https://www.rfc-editor.org/info/rfc7371>>.

Appendix A. Acknowledgements

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

Authors' Addresses

Mingwei Xu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

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

Yong Cui
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

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

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

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

Shu Yang
Tsinghua University
Graduate School at Shenzhen
Shenzhen 518055
P.R. China

Phone: +86-10-6278-5822
Email: yangshu@csnet1.cs.tsinghua.edu.cn

Chris Metz
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

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

Greg Shepherd
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Phone: +1-541-912-9758
Email: shep@cisco.com