

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 1, 2021

H. Chen  
M. McBride  
Futurewei  
Y. Fan  
Casa Systems  
M. Toy  
Verizon  
A. Wang  
China Telecom  
L. Liu  
Fujitsu  
X. Liu  
Volta Networks  
September 28, 2020

SRv6 Point-to-Multipoint Path  
draft-chen-pim-srv6-p2mp-path-01

Abstract

This document describes a solution for a SRv6 Point-to-Multipoint (P2MP) Path/Tree to deliver the traffic from the ingress of the path to the multiple egresses/leaves of the path in a SR domain. There is no state stored in the core of the network for a SR P2MP path like a SR Point-to-Point (P2P) path in this solution.

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 <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 April 1, 2021.

#### Copyright Notice

Copyright (c) 2020 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 . . . . .	2
2. Overview of P2MP Multicast Tree . . . . .	3
3. Encoding P2MP Multicast Tree . . . . .	5
4. Procedures/Behaviors . . . . .	7
4.1. Procedure/Behavior on Ingress Node . . . . .	7
4.2. Procedure/Behavior on Transit Node . . . . .	8
4.3. Procedure/Behavior on Egress Node . . . . .	10
5. Protection . . . . .	10
5.1. Global Protection . . . . .	10
5.2. Local Protection . . . . .	10
6. IANA Considerations . . . . .	11
7. Security Considerations . . . . .	11
8. Acknowledgements . . . . .	11
9. References . . . . .	11
9.1. Normative References . . . . .	11
9.2. Informative References . . . . .	12
Authors' Addresses . . . . .	12

#### 1. Introduction

The Segment Routing (SR) for unicast or Point-to-Point (P2P) path is described in [RFC8402]. For SR multicast or Point-to-Multipoint (P2MP) path/tree, it may be implemented through using multiple SR P2P paths. The function of a SR P2MP path/tree from an ingress node to multiple (say n) egress/leaf nodes is implemented by n SR P2P paths. These n P2P paths are from the ingress to those n egress/leaf nodes of the P2MP path/tree. This solution may waste some network resources such as link bandwidth.

An alternative solution proposed in [I-D.shen-spring-p2mp-transport-chain] uses a number of P2MP chain tunnels to implement a P2MP path/tree from an ingress to n egress/leaf nodes. Each P2MP chain tunnel is a tunnel from the ingress to a leaf node as its tail end and may have some leaf nodes as its bud nodes along the tunnel. This alternative solution improves the usage of network resources over the solution above using pure P2P paths. However, these two solutions are based on SR P2P paths.

A solution for a SR P2MP path/tree using a P2MP multicast tree is proposed in [I-D.voyer-pim-sr-p2mp-policy]. For a SR P2MP path/tree from an ingress/root to multiple egress/leaf nodes, a multicast P2MP tree is created to deliver the traffic from the ingress/root to the egress/leaf nodes. The state of the tree is instantiated in the forwarding plane by a controller such as PCE at Root node, intermediate Replication nodes and Leaf nodes of the tree. This is not consistent with the SR principles in which no state is stored at the core of the network.

This document describes a new solution for a SRv6 Point-to-Multipoint (P2MP) Path/Tree to deliver the traffic from the ingress of the path to the multiple egresses/leaves of the path in a SR domain. This solution uses a P2MP multicast tree without storing its state in the core of the network for a SR P2MP path/tree like a SR P2P path.

## 2. Overview of P2MP Multicast Tree

For a SR P2P path from its ingress to its egress, a segment list for the path is provided to the ingress. The ingress pushes the list into a packet, and the packet is delivered to the egress according to the segment list without any state in the core of the network.

For a SR P2MP path from its ingress to multiple egress/leaf nodes, a segment list for the P2MP path is provided to the ingress. The ingress pushes the list into a packet, and the packet is delivered to the multiple egress/leaf nodes according to the segment list without any state in the core of the network.

Figure 1 shows a SR P2MP path from ingress/root R to four egress/leaf nodes L1, L2, L3 and L4. Nodes P1, P2, P3 and P4 are the transit nodes of the P2MP path.

Suppose that X-m is the segment identifier (SID) of node X. X-m is an adjacent SID or node SID. For simplicity, we assume X-m is a node SID in the illustrations below. R-m, P1-m, P2-m, P3-m, P4-m, L1-m, L2-m, L3-m and L4-m are the SIDs of the nodes on the SR P2MP path. They are multicast SIDs or replication SIDs in general.

A multicast SID is a SID from a multicast SID block. In a SR domain supporting SR multicast, each node has a multicast node SID, which is globally significant; each adjacency of a node has a multicast adjacency SID, which is locally significant. A multicast SID of a node on a SR P2MP path is associated with the SIDs of the next hop (or say downstream) nodes. When the node receives a packet with its multicast SID, it duplicates and sends the packet to each of the next hop nodes according to their SIDs.

If node P on a SR P2MP path has B ( $B > 1$ ) next hop nodes along the path, the SID of node P, P-m, MUST be a multicast SID when it is in the segment list for the P2MP path. The SIDs of the B next hop nodes just follow P-m in the segment list. When node P receives the packet with P-m, it duplicates and sends the packet to each of the B next hop nodes along the P2MP path.




Figure 1: SR P2MP Path from R to L1, L2, L3 and L4

<P1-m, P2-m, P3-m, L1-m, L2-m, P4-m, L3-m, L4-m> is a segment list for the SR P2MP path in Figure 1 to be pushed into a packet at ingress/root R. Node P1 has 2 next hop nodes P2 and P3 along the P2MP path. The next hop nodes' SIDs P2-m and P3-m follow P1-m, which is P1's multicast SID. When P1 receives a packet transported by the P2MP path, it duplicates and sends the packet to the next hop nodes P2 and P3 according to P1-m, P2-m and P3-m.

The number of branches or next hops from node P1 is a value of one argument in P1-m, called N-Branched. The value of N-Branched in P1-m is 2. With this information, node P1 duplicates and sends the packet to 2 next hop nodes P2 and P3, which are indicated by the 2 SIDs P2-m and P3-m following P1-m.

The number of SIDs of the nodes under node P1 is a value of another argument in P1-m, called N-SIDs. The value of N-SIDs in P1-m is 7, indicating that there are 7 SIDs following P1-m in the segment list.

There are 2 branches or next hops (i.e., L1 and L2) from node P2 and 2 SIDs (i.e., L1-m and L2-m) of the nodes under node P2. The values of N-Branched and N-SIDs in P2-m are 2 and 2. with this information, before sending the packet to node P2, node P1 pushes the SIDs under node P2 into the packet (i.e., the packet has a new segment list with the SIDs under node P2. The new segment list replaces the old one in the packet).

There are 1 branch or next hop (i.e., P4) from node P3 and 3 SIDs (i.e., P4-m, L3-m and L4-m) of the nodes under node P3. The values of N-Branched and N-SIDs in P3-m are 1 and 3 respectively. with this information, before sending the packet to node P3, node P1 pushes the SIDs under node P3 into the packet.

Each node on the SR P2MP path sends the packet to its next hop nodes according to the segment list and no state is stored in any transit node (i.e., the core of the network). The packet is delivered to the egress/leaf nodes from the ingress.

### 3. Encoding P2MP Multicast Tree

For each sub-tree  $ST_i$  of a SR P2MP path from the ingress node of the P2MP path, suppose that

- o the multicast SID of the next hop node  $NH_i$  is  $mSID_i$ ;
- o there are  $B_i$  branches (i.e., outgoing interfaces) to the next hop node  $BNH_j$  ( $j = 1, \dots, B_i$ ) from node  $NH_i$  along the sub-tree, the multicast SID of  $BNH_j$  is  $mSID_{ij}$ ;
- o the number of branches (i.e., outgoing interfaces) under the node  $BNH_j$  ( $j = 1, \dots, B_i$ ) is  $BB_j$ ; and the number of SIDs of the nodes under each of the  $B_i$  branches from node  $BNH_j$  is  $NS_j$  ( $j = 1, \dots, B_i$ ).

Sub-tree  $ST_i$  is encoded as segment list

$\langle$	$mSID_i,$	$mSID_{i1}, \dots, mSID_{iB_i},$	$SegSeq_1, \dots,$	$SegSeq_{B_i}$	$\rangle,$
	$\backslash$	$\backslash$	$\backslash$	$\backslash$	
SIDs of	$NH_i$	$B_i$ branches/next-hops $BNH_j$ of node $NH_i$	sub-tree from $BNH_1$	sub-tree from $BNH_{B_i}$	

where  $mSID_i$  contains the number of branches,  $B_i$ , in its N-Branched field, and the number of SIDs under  $mSID_i$  in its N-SIDs field;  $mSID_{ij}$

( $j = 1, \dots, B_i$ ) contains the number of branches,  $B_{Bj}$ , in its N-Branches field and the number of SIDs,  $NS_j$ , in its N-SIDs field;  $SegSeq_j$  ( $j = 1, \dots, B_i$ ) is the SID sequence in the segment list encoding the sub-trees from node  $BNH_j$ .

For the P2MP path in Figure 1 from ingress node R to egress nodes L1, L2, L3 and L4, there is one sub-tree from R.

For this sub-tree,

- o the next hop node is P1 and the multicast SID of P1 is P1-m;
- o there are 2 branches to the next hop nodes P2 and P3 from node P1 along the sub-tree; the number of SIDs of the nodes under P1 is 7; the multicast SIDs of P2 and P3 are P2-m and P3-m respectively;
- o the numbers of SIDs of the nodes under these two branches are 2 and 3 respectively. The SIDs of the nodes under P2 are L1-m and L2-m. The SIDs of the nodes under P3 are P4-m, L3-m and L4-m.

The sub-tree is encoded as segment list

	< P1-m,	P2-m, P3-m,	L1-m, L2-m,	P4-m, L3-m, L4-m > ,
	\___/	\_____/	\_____/	\_____/
SIDs of	P1	2 branches/next-hops P2 and P3 of node P1	sub-tree from P2	sub-tree from P3

where

P1-m's N-Branches field is set to 2 and its N-SIDs field to 7;  
P2-m's N-Branches field is set to 2 and its N-SIDs field to 2;  
P3-m's N-Branches field is set to 1 and its N-SIDs field to 3;

L1-m and L2-m are the SID sequence in the segment list encoding the sub-trees from P2;

P4-m, L3-m and L4-m are the SID sequence in the segment list encoding the sub-trees from P3; and

P4-m's N-Branches field is set to 2 and its N-SIDs field to 2.

Figure 2 shows in details the segment list, which is an encoding of the P2MP multicast tree for the SR P2MP path from R to L1, L2, L3 and L4.

	N-Branches	N-SIDs		
P1's Multicast SID Locator	2	7	Arguments	P1-m
P2's Multicast SID Locator	2	2	Arguments	P2-m
P3's Multicast SID Locator	1	3	Arguments	P3-m
L1's Multicast SID Locator	0	0	Arguments	L1-m
L2's Multicast SID Locator	0	0	Arguments	L2-m
P4's Multicast SID Locator	2	2	Arguments	P4-m
L3's Multicast SID Locator	0	0	Arguments	L3-m
L4's Multicast SID Locator	0	0	Arguments	L4-m

Figure 2: Encoding of P2MP Multicast Tree from R to L1 - L4

SID P1-m indicates that there are 2 branches and 7 SIDs under P1.  
 SID P2-m indicates that there are 2 branches and 2 SIDs under P2.  
 SID P3-m indicates that there are 1 branch and 3 SIDs under P3. SIDs  
 L1-m and L2-m indicate that there is no branch under them. SID P4-m  
 indicates that there are 2 branches and 2 SIDs under P4. L3-m and  
 L4-m indicate that there is no branch under them.

#### 4. Procedures/Behaviors

This section describes the procedures or behaviors on the ingress, transit and egress/leaf node of a SR P2MP path to deliver a packet received from the path to its destinations.

##### 4.1. Procedure/Behavior on Ingress Node

For a packet to be transported by a SR P2MP Path, the ingress of the P2MP path duplicates the packet for each sub-tree of the SR P2MP path branching from the ingress, pushes the segment list encoding the sub-tree into the packet by executing H.Encaps [I-D.ietf-spring-srv6-network-programming] and sends the packet to the next hop node along the sub-tree.

For example, there is one sub-tree from the ingress R of the SR P2MP path in Figure 1 via next hop node P1 towards egress/leaf nodes L1, L2, L3 and L4.

For this sub-tree, the ingress R duplicates the packet, set the destination address (DA) to P1-m (i.e., multicast SID of node P1), pushes the segment list without P1-m (i.e., <P2-m, P3-m, L1-m, L2-m, P4-m, L3-m, L4-m>) encoding the sub-tree in a Segment Routing Header (SRH) of the packet by executing H.Encaps and sends the packet to DA (i.e., node P1). The contents of the multicast SIDs P1-m, P2-m, P3-m, L1-m, L2-m, P4-m, L3-m, L4-m are shown in Figure 2.

Suppose that the duplicated packet is Pkt0 for the sub-tree. The execution of H.Encaps pushes an IPv6 header (i.e., SRH) to Pkt0 and sets some fields in the header to produce an encapsulated packet Pkt'. Pkt' is represented in the following:

$$\text{Pkt}' = (\text{SA}=\text{R}, \text{DA}=\text{P1-m}) ( \underbrace{\text{L4-m, L3-m, \dots, P3-m, P2-m}}_{\text{corresponds to: } \langle \text{P2-m, P3-m, \dots, L3-m, L4-m} \rangle} ; \text{SL}=7 ) \text{Pkt0}$$

where DA=P1-m means that the destination address (DA) is set to P1-m; SA=R means that the source address (SA) is set to R; SL=7 means that the number of Segments Left (SL) is 7.

#### 4.2. Procedure/Behavior on Transit Node

When a transit node of a SR P2MP path receives a packet transported by the P2MP path, the DA of the packet is a multicast SID of the node and the packet contains a segment list for the sub-trees under the transit node. The DA and the segment list comprise the information for encoding the sub-trees.

For example, when node P1 receives a packet transported by the SR P2MP path in Figure 1, the packet's DA is P1-m (which is a multicast SID of node P1) and the segment list in the packet is <P2-m, P3-m, L1-m, L2-m, P4-m, L3-m, L4-m>.

The N-Branched field (which has value of n) of the DA indicates that there are n branches (or say sub-trees) under the transit node. The N-SIDs field of the DA indicates the number of SIDs for these n sub-trees under the transit node. The multicast SIDs of the next hop nodes of these n sub-trees are the first n multicast SIDs of the segment list in the packet.

For example, the N-Branched field (which has value of 2) of DA = P1-m indicates that there are 2 branches (or say sub-trees) under node P1. The N-SIDs field (which has value of 7) of the DA = P1-m indicates that there are 7 SIDs for these 2 sub-trees under node P1.

The first multicast SID (P2-m) of the segment list is the SID of the next hop node (P2) of the first sub-tree; The second multicast SID



(P3-m) of the segment list is the SID of the next hop node (P3) of the second sub-tree.

After the multicast SIDs of the next hop nodes, there are n blocks of SIDs for those n sub-trees. The N-SIDs field (which has value of B1) of the first multicast SID of the next hop nodes indicates that there are B1 SIDs in the first block for the first sub-tree; the N-SIDs field (which has value of B2) of the second multicast SID of the next hop nodes indicates that there are B2 SIDs in the second block for the second sub-tree after the first block; and so on.

For example, there are 2 blocks of SIDs for the 2 sub-trees under node P1 after the multicast SIDs P2-m and P3-m of the next hop nodes P2 and P3. The N-SIDs field of P2-m (the first multicast SID of the next hop nodes) has value of 2, indicating that there are 2 SIDs in the first block for the first sub-tree, which are L1-m and L2-m.

The N-SIDs field of P3-m (the second multicast SID of the next hop nodes) has value of 3, indicating that there are 3 SIDs in the second block for the second sub-tree after the first block, which are P4-m, L3-m and L4-m.

The transit node duplicates the packet without top header for each sub-tree under it and adds a new header with a new segment list built from the SID block for the sub-tree to the duplicated packet by executing H.Encaps. It sets the DA of the packet to the multicast SID of the next hop node along the sub-tree and sends the packet to the DA.

For example, node P1 duplicates the packet for the first sub-tree towards L1 and L2 and adds a new header with a new segment list <L1-m, L2-m>. It sets DA = P2-m (multicast SID of next hop P2), and sends the packet to the DA (i.e., P2).

Suppose that the duplicated packet is Pkt0 for the sub-tree. The execution of H.Encaps pushes a new IPv6 header (i.e., SRH) to Pkt0 and sets some fields in the header to produce an encapsulated packet Pkt'. Pkt' is represented in the following:

$$\text{Pkt}' = (\text{SA}=\text{P1}, \text{DA}=\text{P2-m}) (\text{L2-m}, \text{L1-m}; \text{SL}=2) \text{Pkt0}.$$

corresponds to:  $\underbrace{\hspace{2cm}}_{\langle \text{L1-m}, \text{L2-m} \rangle}$

where DA=P2-m means that the destination address (DA) is set to P2-m; SA=P1 means that the source address (SA) is set to P1; SL=2 means that the number of Segments Left (SL) is 2.

Node P1 duplicates the packet for the second sub-tree via P3 towards L3 and L4 and adds a new header with a new segment list <P4-m, L3-m, L4-m>. It sets DA = P3-m (multicast SID of next hop P3), and sends the packet to the DA (i.e., P3).

#### 4.3. Procedure/Behavior on Egress Node

When an egress node of a SR P2MP path receives a packet transported by the P2MP path, the DA of the packet is a SID of the egress node. The egress node sends the packet to its destination accordingly. If the SID is a multicast SID of the egress, the N-Branches field and N-SIDs field are all zeros.

### 5. Protection

Protections for a SR P2MP path can be classified into two types: global protection and local protection.

#### 5.1. Global Protection

For a primary SR P2MP path from an ingress node R1 to multiple egress nodes  $L_i$  ( $i = 1, \dots, n$ ), a backup SR P2MP path from an ingress node  $R1'$  to multiple egress nodes  $L_i'$  ( $i = 1, \dots, n$ ) is set up to provide global protection for the primary SR P2MP path. If  $R1'$  is the same as  $R1$ , the failure of the ingress node  $R1$  of the primary SR P2MP path is not protected; otherwise (i.e.,  $R1'$  and  $R1$  are different and connected to the same traffic source), the failure of the ingress node  $R1$  is protected. If  $L_i'$  is the same as  $L_i$  ( $i = 1, \dots, n$ ), the failure of the egress nodes  $L_i$  ( $i = 1, \dots, n$ ) of the primary SR P2MP path is not protected; otherwise (i.e.,  $L_i'$  and  $L_i$  are different and connected to the same destination), the failure of the egress nodes  $L_i$  is protected.

When a failure happens on the primary SR P2MP path and is detected by the source of the traffic or other entity, the traffic to be transported by the primary SR P2MP path is switched to the backup SR P2MP path, which sends the traffic from its ingress node  $R1'$  to its egress nodes  $L_i'$  ( $i = 1, \dots, n$ ).

#### 5.2. Local Protection

Local protection or say Fast Reroute (FRR) of a node and adjacency segment on a SR P2P path is proposed in [I-D.ietf-rtgwg-segment-routing-ti-lfa] and [I-D.ietf-rtgwg-srv6-egress-protection]. It can be applied to FRR of a node and adjacency segment on a SR P2MP path in a similar way. But FRR for SR P2MP path is more complicated.

More details will be added later.

## 6. IANA Considerations

TBD

## 7. Security Considerations

TBD

## 8. Acknowledgements

The authors would like to thank Acee Lindem and Daniel Voyer for their valuable comments and suggestions on this draft.

## 9. References

### 9.1. Normative References

- [I-D.ietf-6man-segment-routing-header]  
Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-26 (work in progress), October 2019.
- [I-D.ietf-rtgwg-segment-routing-ti-lfa]  
Litkowski, S., Bashandy, A., Filsfils, C., Decraene, B., Francois, P., Voyer, D., Clad, F., and P. Camarillo, "Topology Independent Fast Reroute using Segment Routing", draft-ietf-rtgwg-segment-routing-ti-lfa-04 (work in progress), August 2020.
- [I-D.ietf-rtgwg-srv6-egress-protection]  
Hu, Z., Chen, H., Chen, H., Wu, P., Toy, M., Cao, C., Liu, L., and X. Liu, "SRv6 Path Egress Protection", draft-ietf-rtgwg-srv6-egress-protection-01 (work in progress), July 2020.
- [I-D.ietf-spring-srv6-network-programming]  
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-ietf-spring-srv6-network-programming-20 (work in progress), September 2020.
- [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>>.

- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

## 9.2. Informative References

- [I-D.shen-spring-p2mp-transport-chain]  
Shen, Y., Zhang, Z., Parekh, R., Bidgoli, H., and Y. Kamite, "Point-to-Multipoint Transport Using Chain Replication in Segment Routing", draft-shen-spring-p2mp-transport-chain-02 (work in progress), April 2020.
- [I-D.voyer-pim-sr-p2mp-policy]  
Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Point-to-Multipoint Policy", draft-voyer-pim-sr-p2mp-policy-02 (work in progress), July 2020.
- [I-D.voyer-spring-sr-replication-segment]  
Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "SR Replication Segment for Multi-point Service Delivery", draft-voyer-spring-sr-replication-segment-04 (work in progress), July 2020.

## Authors' Addresses

Huaimo Chen  
Futurewei  
Boston, MA  
USA

Email: [Huaimo.chen@futurewei.com](mailto:Huaimo.chen@futurewei.com)

Mike McBride  
Futurewei

Email: michael.mcbride@futurewei.com

Yanhe Fan  
Casa Systems  
USA

Email: yfan@casa-systems.com

Mehmet Toy  
Verizon  
USA

Email: mehmet.toy@verizon.com

Aijun Wang  
China Telecom  
Beiqijia Town, Changping District  
Beijing, 102209  
China

Email: wangaj3@chinatelecom.cn

Lei Liu  
Fujitsu

USA

Email: liulei.kddi@gmail.com

Xufeng Liu  
Volta Networks

McLean, VA  
USA

Email: xufeng.liu.ietf@gmail.com

Network Working Group  
Internet-Draft  
Updates: 3376, 3810 (if approved)  
Intended status: Standards Track  
Expires: May 6, 2021

M. Sivakumar  
Juniper Networks  
S. Venaas  
Cisco Systems, Inc.  
Z. Zhang  
ZTE Corporation  
H. Asaeda  
NICT  
November 2, 2020

IGMPv3/MLDv2 Message Extension  
draft-ietf-pim-igmp-mld-extension-02

Abstract

IGMP and MLD protocols are extensible, but no extensions have been defined so far. This document provides a well-defined way of extending IGMP and MLD, using a list of TLVs (Type, Length and Value).

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 6, 2021.

Copyright Notice

Copyright (c) 2020 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 . . . . .	2
2. Conventions used in this document . . . . .	2
3. Extension Format . . . . .	3
3.1. Multicast Listener Query Extension . . . . .	4
3.2. Version 2 Multicast Listener Report Extension . . . . .	5
3.3. IGMP Membership Query Extension . . . . .	6
3.4. IGMP Version 3 Membership Report Extension . . . . .	7
4. Applicability and backwards compatibility . . . . .	8
5. Security Considerations . . . . .	9
6. IANA Considerations . . . . .	9
7. References . . . . .	10
7.1. Normative References . . . . .	10
7.2. Informative References . . . . .	10
Authors' Addresses . . . . .	11

## 1. Introduction

In this document, we describe a generic method to extend IGMPv3 [RFC3376] and MLDv2 [RFC3810] messages to accommodate information other than what is contained in the current message formats. This is done by allowing a list of TLVs (Type, Length and Value) to be used in the Additional Data part of IGMPv3 and MLDv2 messages. This document defines a registry for such TLVs, while other documents will define the specific types and their values, and their semantics. The extension would only be used when at least one TLV is to be added to the message. This extension also applies to the lightweight versions of IGMPv3 and MLDv2 as defined in [RFC5790].

The extension will be part of additional data as mentioned in [RFC3810] Section 5.1.12 (resp. [RFC3376] Section 4.1.10) for query messages and [RFC3810] Section 5.2.12 (resp. [RFC3376] Section 4.2.11) for report messages.

One such TLV is being defined in [I-D.ietf-bier-mlv]

## 2. Conventions used in this document

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. Extension Format

A previously reserved bit in the IGMPv2 and MLDv2 headers is used to indicate whether this extension is used. It is set to 1 if it is used, otherwise 0. When this extension is used, the Additional Data of IGMPv3 and MLDv2 messages would be formatted as follows. Note that this format contains a variable number of TLVs. It MUST contain at least one TLV.

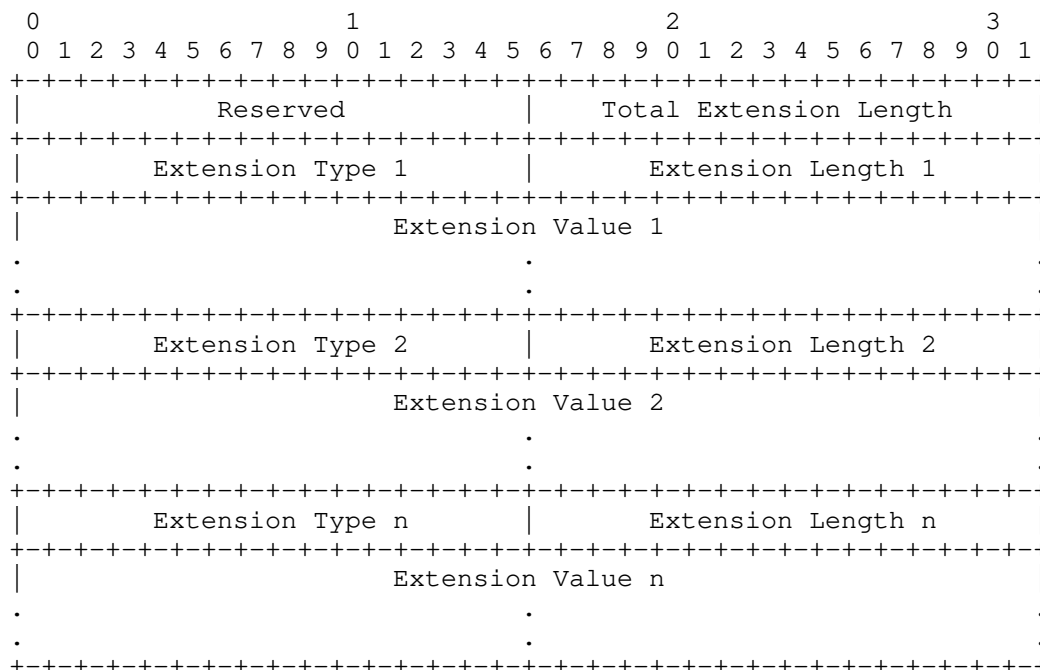


Figure 1: Extension Format

Reserved: 2 octets. Reserved. MUST be set to 0. MUST be ignored when received.

Total Extension Length: 2 octets. The remaining length of the extension. This value MUST be equal to  $((2 + 2) * n) + \text{Extension Length 1} + \text{Extension Length 2} + \dots + \text{Extension Length n}$ . That is, it is the sum of the lengths of all the TLVs, including the type field (2 octets), and the length field (2 octets) of each TLV. The total number of octets used by the extension is the value of



this field plus 4 (including the 2 octets Reserved field and the 2 octets of this field).

**Extension Type:** 2 octets. This identifies a particular Extension Type as defined in the IGMP/MLD Extension Type Registry.

**Extension Length:** 2 octets. This specifies the length in octets of the following Extension Value field.

**Extension Value:** This field contains the value. The length and the contents of this field is according to the specification of the Extension Type as defined in the IGMP/MLD Extension Type Registry. The length MUST be as specified in the Extension Length field.

Note that there may be additional data following this extension. The Total Extension Length field would indicate where this extension ends, and the additional data starts. Also, there is a possibility that an implementation uses the Additional Data part of IGMP/MLD messages, but not according to this extension scheme. When a message is received, it MUST be verified that the Total Extension Length is equal to  $((2 + 2) * n) + \text{Extension Length 1} + \text{Extension Length 2} + \dots + \text{Extension Length n}$ , where n is the number of TLVs. Note that the value of n is not known ahead of time. An implementation would walk through the TLVs and add the 4 octet overhead and the length of each TLV, until the sum is larger or equal to the Total Extension Length, or until the end of the IGMP/MLD message, whichever happens first. Any additional data after this MUST be ignored, except the data MUST be included in checksum computations. If the sum is not equal to the Total Extension Length, then it is assumed that this extension is not being used, and this specification does not apply.

### 3.1. Multicast Listener Query Extension

The MLD query format with extension is shown below. The E-bit is set to 1 to indicate that the extension is present.






Figure 2: MLD Query Extension

3.2. Version 2 Multicast Listener Report Extension

The MLD report format with extension is shown below. The E-bit is set to 1 to indicate that the extension is present.

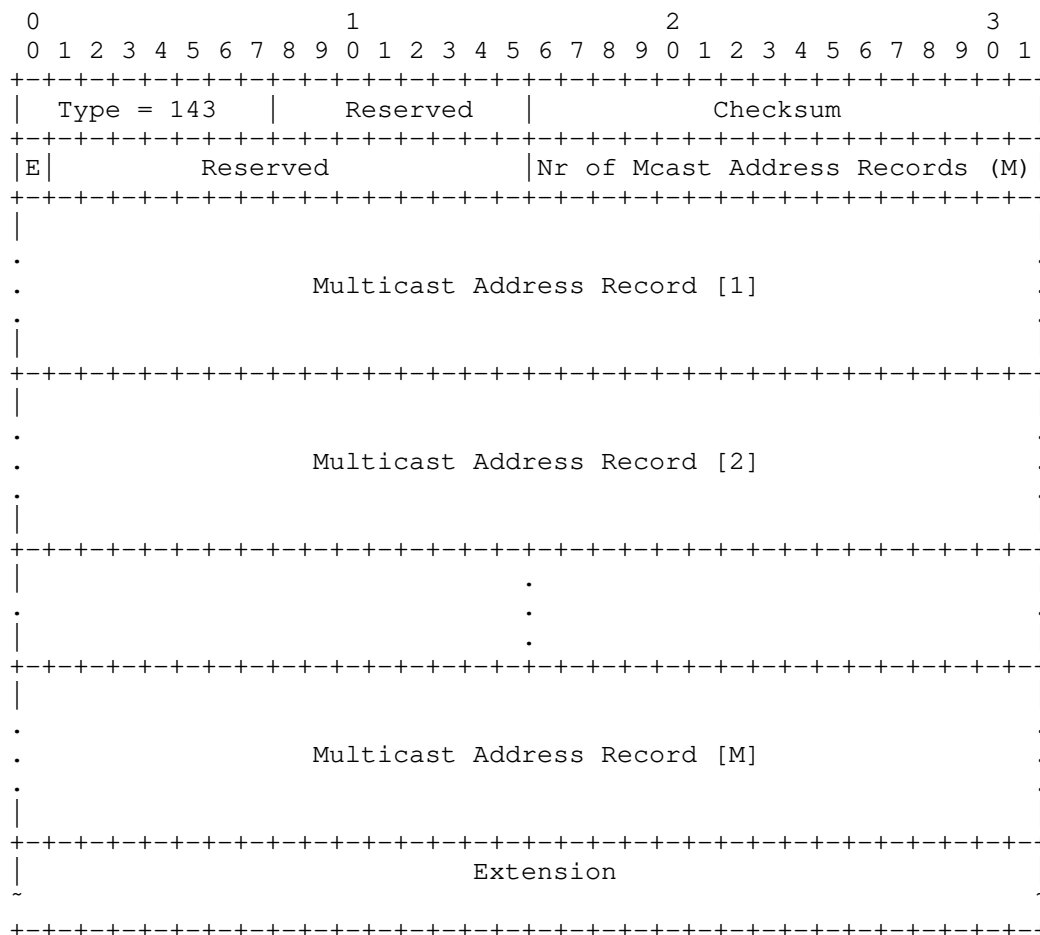


Figure 3: MLD Report Extension

### 3.3. IGMP Membership Query Extension

The IGMP query format with the extension is shown below. The E-bit is set to 1 to indicate that the extension is present.

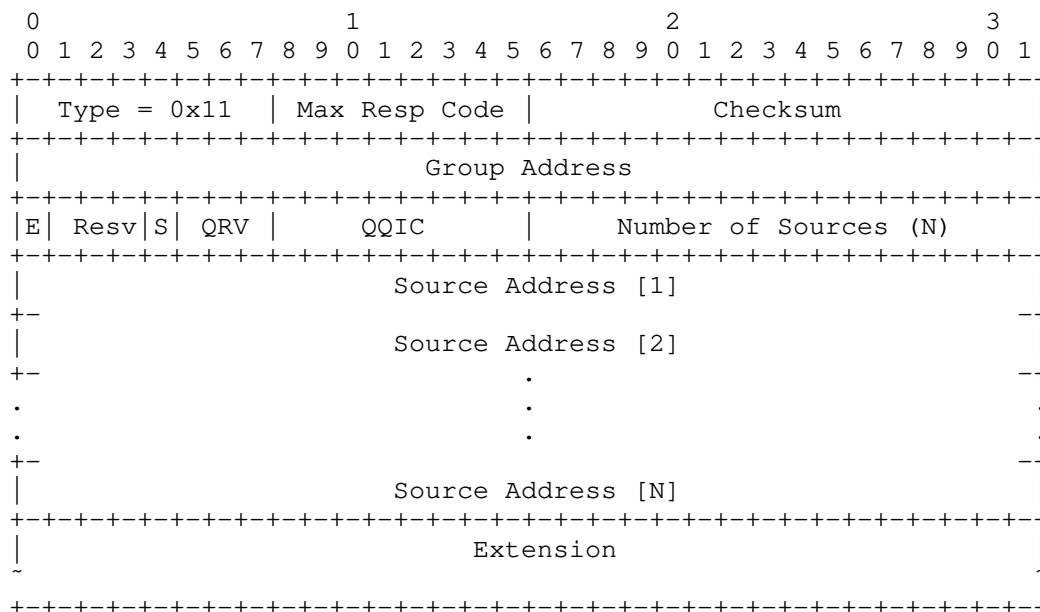


Figure 4: IGMP Query Extension

### 3.4. IGMP Version 3 Membership Report Extension

The IGMP report format with the extension is shown below. The E-bit is set to 1 to indicate that the extension is present.




Figure 5: IGMP Report Extension

4. Applicability and backwards compatibility

IGMP and MLD implementations, host implementations in particular, rarely change, and it is expected to take a long time for them to support this extension mechanism. Also as new extensions are defined, it may take a long time before they are supported. Implementations that do not support this extension mechanism will simply ignore the extension, provided they are compliant with IGMPv3 and MLDv2 RFCs, and behave as if the extension is not present. Implementations that support this extension MUST behave as if it is not present if they support none of the extension types in an IGMP/MLD message. If they support at least one of the types, they will

process the supported types according to the respective type specifications, and ignore any unsupported types.

It is possible that a new extension type only applies to queries, or only reports, or there may be other specific conditions for when it is to be used. A document defining a new type MUST specify clearly under what conditions the new type should be used, including for which message types. It MUST also be considered what the behavior should be if a message is not used in the defined manner, e.g., if it is present in a query message, when it was only expected to be used in reports.

When defining new types, care must be taken to ensure that nodes that support the type can co-exist with nodes that don't, on the same subnet. There could be multiple routers where only some support the extension, or multiple hosts where only some support the extension. Or a router may support it and none of the hosts, or all hosts may support it, but none of the routers. With multiple types being used, it must also be considered that some hosts or routers may only support some of the types, and potentially one node might support only one type, and another node only another type.

Documents defining new types MUST have security considerations relevant to the new types. They MUST also in addition to defining the behavior of host and routers supporting the new types, consider compatibility with hosts and routers on the same subnet that do not support the new types. Further, they MUST consider whether there are any dependencies or restrictions on combinations between the new types and any pre-existing types.

The extension mechanism do not support IGMPv1, IGMPv2 and MLDv1. As nodes may send older version message, they would also not be able to send messages using this extension.

## 5. Security Considerations

This document extends MLD (resp. IGMP) message formats. As such, there is no impact on security or changes to the considerations in [RFC3810] and [RFC3376]. The respective types defined using this extension may impact security and MUST be considered as part of the respective specifications.

## 6. IANA Considerations

A new registry called "IGMP/MLD Extension Types" should be created with registration procedure "IETF Review" as defined in [RFC8126] with this document as a reference. The registry should be common for IGMP and MLD and can perhaps be added to the "Internet Group

Management Protocol (IGMP) Type Numbers" section. The initial content of the registry should be as below.

Type	Length	Name	Reference
------	--------	------	-----------

---

## 7. References

### 7.1. 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>>.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.

### 7.2. Informative References

- [I-D.ietf-bier-mlD] Pfister, P., Wijnands, I., Venaas, S., Wang, C., Zhang, Z., and M. Stenberg, "BIER Ingress Multicast Flow Overlay using Multicast Listener Discovery Protocols", draft-ietf-bier-mlD-04 (work in progress), March 2020.
- [RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, DOI 10.17487/RFC5790, February 2010, <<https://www.rfc-editor.org/info/rfc5790>>.

Authors' Addresses

Mahesh Sivakumar  
Juniper Networks  
64 Butler St  
Milpitas CA 95035  
USA

Email: [sivakumar.mahesh@gmail.com](mailto:sivakumar.mahesh@gmail.com)

Stig Venaas  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: [stig@cisco.com](mailto:stig@cisco.com)

Zheng (Sandy) Zhang  
ZTE Corporation  
No. 50 Software Ave, Yuhuatai District  
Nanjing 210000  
China

Email: [zhang.zheng@zte.com.cn](mailto:zhang.zheng@zte.com.cn)

Hitoshi Asaeda  
National Institute of Information and  
Communications Technology  
4-2-1 Nukui-Kitamachi  
Koganei, Tokyo 184-8795  
Japan

Email: [asaeda@nict.go.jp](mailto:asaeda@nict.go.jp)



PIM Working Group  
Internet Draft  
Intended status: Standards Track  
Expires: April 16, 2021

H. Zhao  
Ericsson  
X. Liu  
Volta  
Y. Liu  
China Mobile  
M. Panchanathan  
Cisco  
M. Sivakumar  
Juniper

October 17, 2020

A Yang Data Model for IGMP/MLD Proxy  
draft-ietf-pim-igmp-mld-proxy-yang-03.txt

Abstract

This document defines a YANG data model that can be used to configure and manage Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) proxy devices. The YANG module in this document conforms to Network Management Datastore Architecture (NMDA).

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on April 16, 2021.

#### Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

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

#### Table of Contents

1. Introduction.....	3
1.1. Terminology.....	3
1.2. Tree Diagrams.....	3
2. Design of Data Model.....	3
2.1. Overview.....	4
2.2. Augment /rt:routing/rt:control-plane-protocols/rt:control-plane-protocol.....	4
3. IGMP/MLD Proxy YANG Module.....	5
4. Security Considerations.....	13
5. IANA Considerations.....	13
6. Normative References.....	14
Authors' Addresses.....	16

## 1. Introduction

This document defines a YANG [RFC6020] data model for the management of Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) proxy devices.

The YANG module in this document conforms to the Network Management Datastore Architecture defined in [RFC8342]. The "Network Management Datastore Architecture" (NMDA) adds the ability to inspect the current operational values for configuration, allowing clients to use identical paths for retrieving the configured values and the operational values.

### 1.1. Terminology

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].

The terminology for describing YANG data models is found in [RFC6020].

### 1.2. Tree Diagrams

A simplified graphical representation of the data model is used in this document. The meaning of the symbols in these diagrams is as follows:

- o Brackets "[" and "]" enclose list keys.
- o Abbreviations before data node names: "rw" means configuration (read-write), and "ro" means state data (read-only).
- o Symbols after data node names: "?" means an optional node, "!" means a presence container, and "\*" denotes a list and leaf-list.
- o Parentheses enclose choice and case nodes, and case nodes are also marked with a colon (":").
- o Ellipsis ("...") stands for contents of subtrees that are not shown.

## 2. Design of Data Model

The model covers Considerations for Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD) - Based Multicast Forwarding ("IGMP/MLD Proxying") [RFC4605].

The goal of this document is to define a data model that provides a common user interface to IGMP/MLD proxy. This document provides freedom for vendors to adapt this data model to their product implementations.

## 2.1. Overview

The IGMP/MLD proxy YANG module defined in this document has all the common building blocks for the IGMP/MLD proxy protocol.

The YANG module augments `/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol` to enable IGMP/MLD proxy and configure other related parameters.

This YANG module follows the Guidelines for YANG Module Authors (NMDA) [draft-dsdt-nmda-guidelines-01]. This NMDA ("Network Management Datastore Architecture") architecture provides an architectural framework for datastores as they are used by network management protocols such as NETCONF [RFC6241], RESTCONF [RFC8040] and the YANG [RFC7950] data modeling language.

## 2.2. Augment `/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol`

The YANG module augments `/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol` to configure IGMP/MLD proxy. The interface list under `igmp-proxy` or `mld-proxy` contains upstream interfaces for IGMP/MLD proxy. There is also a constraint to make sure the upstream interface for IGMP/MLD proxy should not be configured PIM.

To configure a downstream interface for IGMP/MLD proxy, enable IGMP/MLD on that interface. This is defined in the YANG Data Model for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD). In IGMP/MLD proxy model `downstream-interface` is read-only.

```

module: ietf-igmp-mld-proxy
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol:
      +--rw igmp-proxy {igmp-proxy}?
        +--rw interfaces
          +--rw interface* [interface-name]
            +--rw interface-name          if:interface-ref
            +--rw igmp-version?           uint8
            +--rw enable?                  boolean
            +--rw sender-source-address?  inet:ipv4-address
            +--ro group* [group-address]
              +--ro group-address         inet:ipv4-address
              +--ro up-time?              uint32

```

```

        +--ro filter-mode?      enumeration
        +--ro source* [source-address]
            +--ro source-address      inet:ipv4-address
            +--ro up-time?           uint32
            +--ro downstream-interface* [interface-name]
                +--ro interface-name  if:interface-ref
augment /rt:routing/rt:control-plane-protocols
  /rt:control-plane-protocol:
    +--rw mld-proxy {mld-proxy}?
      +--rw interfaces
        +--rw interface* [interface-name]
          +--rw interface-name      if:interface-ref
          +--rw mld-version?        uint8
          +--rw enable?             boolean
          +--rw sender-source-address? inet:ipv6-address
          +--ro group* [group-address]
            +--ro group-address      inet:ipv6-address
            +--ro up-time?          uint32
            +--ro filter-mode?      enumeration
            +--ro source* [source-address]
              +--ro source-address      inet:ipv6-address
              +--ro up-time?          uint32
              +--ro downstream-interface* [interface-name]
                  +--ro interface-name  if:interface-ref

```

### 3. IGMP/MLD Proxy YANG Module

```

<CODE BEGINS> file ietf-igmp-mld-proxy@2020-09-19.yang
module ietf-igmp-mld-proxy {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-igmp-mld-proxy";
  // replace with IANA namespace when assigned
  prefix imp;

  import ietf-inet-types {
    prefix inet;
  }

  import ietf-interfaces {
    prefix if;
  }

  import ietf-routing {
    prefix rt;
  }

  import ietf-pim-base {
    prefix pim-base;
  }

  organization

```

```
"IETF PIM Working Group";
```

```
contact
```

```
"WG Web: <http://tools.ietf.org/wg/pim/>  
WG List: <mailto:pim@ietf.org>
```

```
Editors: Hongji Zhao  
<mailto:hongji.zhao@ericsson.com>
```

```
Xufeng Liu  
<mailto:xufeng.liu.ietf@gmail.com>
```

```
Yisong Liu  
<mailto:liuyisong@chinamobile.com>
```

```
Mani Panchanathan  
<mailto:mapancha@cisco.com>
```

```
Mahesh Sivakumar  
<mailto:sivakumar.mahesh@gmail.com>
```

```
";
```

```
description
```

```
"The module defines a collection of YANG definitions common for  
all Internet Group Management Protocol (IGMP) and Multicast  
Listener Discovery (MLD) Proxy devices.
```

```
Copyright (c) 2020 IETF Trust and the persons identified as  
authors of the code. All rights reserved.
```

```
Redistribution and use in source and binary forms, with or  
without modification, is permitted pursuant to, and subject to  
the license terms contained in, the Simplified BSD License set  
forth in Section 4.c of the IETF Trust's Legal Provisions  
Relating to IETF Documents  
(http://trustee.ietf.org/license-info).
```

```
This version of this YANG module is part of RFC XXXX; see the  
RFC itself for full legal notices.";
```

```
revision 2020-09-19 {
```

```
  description
```

```
    "Initial revision.";
```

```
  reference
```

```
    "RFC XXXX: A YANG Data Model for IGMP and MLD Proxy";
```

```
}
```

```
/*
```

```
 * Features
```

```
*/

feature igmp-proxy {
  description
    "Support IGMP Proxy protocol.";
  reference
    "RFC 4605";
}

feature mld-proxy {
  description
    "Support MLD Proxy protocol.";
  reference
    "RFC 4605";
}

/*
 * Identities
 */

identity igmp-proxy {
  base rt:control-plane-protocol;
  description
    "IGMP Proxy protocol";
}

identity mld-proxy {
  base rt:control-plane-protocol;
  description
    "MLD Proxy protocol";
}

/*
 * Typedefs
 */

/*
 * Groupings
 */

grouping per-interface-config-attributes {

  description "Config attributes under interface view";

  leaf enable {
    type boolean;
    default false;
    description
      "Set the value to true to enable IGMP/MLD proxy";
```

```
    }
  } // per-interface-config-attributes

  grouping state-group-attributes {
    description
      "State group attributes";

    leaf up-time {
      type uint32;
      units seconds;
      description
        "The elapsed time for (S,G) or (*,G).";
    }

    leaf filter-mode {
      type enumeration {
        enum "include" {
          description
            "In include mode, reception of packets sent
             to the specified multicast address is requested
             only from those IP source addresses listed in the
             source-list parameter";
        }
        enum "exclude" {
          description
            "In exclude mode, reception of packets sent
             to the given multicast address is requested
             from all IP source addresses except those
             listed in the source-list parameter.";
        }
      }
      description
        "Filter mode for a multicast group,
         may be either include or exclude.";
    }
  } // state-group-attributes

/* augments */

augment "/rt:routing/rt:control-plane-protocols"+
  "/rt:control-plane-protocol" {

  description
    "IGMP Proxy augmentation to routing control plane protocol
     configuration and state.";

  container igmp-proxy {
    when 'derived-from-or-self(..../rt:type, "imp:igmp-proxy")' {
      description
        "This container is only valid for IGMP Proxy protocol.";
    }
  }
}
```



```
}
if-feature igmp-proxy;
description "IGMP proxy";
container interfaces {
  description
    "Containing a list of upstream interfaces.";

  list interface {
    key "interface-name";
    description
      "List of upstream interfaces.";

    leaf interface-name {
      type if:interface-ref;
      must "not( current() = /rt:routing"+
"/rt:control-plane-protocols/pim-base:pim"+
"/pim-base:interfaces/pim-base:interface"+
"/pim-base:name )" {

        description
          "The upstream interface for IGMP proxy
          should not be configured PIM.";
      }
      description "The upstream interface name.";
    }

    leaf igmp-version {
      type uint8 {
        range "1..3";
      }
      default 2;
      description "IGMP version.";
    }
  }

  uses per-interface-config-attributes;

  leaf sender-source-address {
    type inet:ipv4-address;
    description
      "The sender source address of
      IGMP membership report or leave.";
  }

  list group {
    key "group-address";
    config false;
    description
      "Multicast group membership information
      that joined on the interface.";
  }
}
```

```
leaf group-address {
    type inet:ipv4-address;
    description
        "Multicast group address.";
}

uses state-group-attributes;

list source {
    key "source-address";
    description
        "List of multicast source information
        of the multicast group.";
    leaf source-address {
        type inet:ipv4-address;
        description
            "Multicast source address";
    }

    leaf up-time {
        type uint32;
        units seconds;
        description
            "The elapsed time for (S,G) or (*,G).";
    }

    list downstream-interface {
        key "interface-name";
        description "The downstream interfaces list.";
        leaf interface-name {
            type if:interface-ref;
            description
                "Downstream interfaces for each upstream-interface";
        }
    }
} // list source
} // list group
} // interface
} // interfaces
}

augment "/rt:routing/rt:control-plane-protocols"+
    "/rt:control-plane-protocol" {

    description
        "MLD Proxy augmentation to routing control plane protocol
        configuration and state.";

    container mld-proxy {
        when 'derived-from-or-self(..//rt:type, "imp:mld-proxy")' {
```

```
description
  "This container is only valid for MLD Proxy protocol.";
}
if-feature mld-proxy;
description "MLD proxy";
container interfaces {
  description
    "Containing a list of upstream interfaces.";

  list interface {
    key "interface-name";
    description
      "List of upstream interfaces.";

    leaf interface-name {
      type if:interface-ref;
      must "not( current() = /rt:routing"+
"/rt:control-plane-protocols/pim-base:pim"+
"/pim-base:interfaces/pim-base:interface"+
"/pim-base:name )" {

        description
          "The upstream interface for MLD proxy
          should not be configured PIM.";
      }
      description "The upstream interface name.";
    }

    leaf mld-version {
      type uint8 {
        range "1..2";
      }
      default 2;
      description "MLD version.";
    }
  }

  uses per-interface-config-attributes;

  leaf sender-source-address {
    type inet:ipv6-address;
    description
      "The sender source address of
      MLD membership report or leave.";
  }

  list group {
    key "group-address";
    config false;
    description
      "Multicast group membership information
      that joined on the interface.";
  }
}
```

```
leaf group-address {
    type inet:ipv6-address;
    description
        "Multicast group address.";
}

uses state-group-attributes;

list source {
    key "source-address";
    description
        "List of multicast source information
        of the multicast group.";
    leaf source-address {
        type inet:ipv6-address;
        description
            "Multicast source address";
    }

    leaf up-time {
        type uint32;
        units seconds;
        description
            "The elapsed time for (S,G) or (*,G).";
    }

    list downstream-interface {
        key "interface-name";
        description "The downstream interfaces list.";
        leaf interface-name {
            type if:interface-ref;
            description
                "Downstream interfaces for each upstream-interface";
        }
    }
} // list source
} // list group
} // interface
} // interfaces
}

/* RPCs */

}
<CODE ENDS>
```

#### 4. Security Considerations

The YANG module specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC5246].

The NETCONF access control model [RFC6536] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., config true, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., edit-config) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes and their sensitivity/vulnerability:

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol
```

Unauthorized access to any data node of these subtrees can adversely affect the IGMP/MLD proxy subsystem of both the local device and the network. This may lead to network malfunctions, delivery of packets to inappropriate destinations, and other problems.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. These are the subtrees and data nodes and their sensitivity/vulnerability:

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol
```

Unauthorized access to any data node of these subtrees can disclose the operational state information of IGMP/MLD proxy on this device.

#### 5. IANA Considerations

RFC Ed.: In this section, replace all occurrences of 'XXXX' with the actual RFC number (and remove this note).

This document registers the following namespace URIs in the IETF XML registry [RFC3688]:

---

URI: urn:ietf:params:xml:ns:yang:ietf-igmp-ml-d-proxy

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

---

This document registers the following YANG modules in the YANG Module Names registry [RFC7950]:

---

name: ietf-igmp-ml-d-proxy

namespace: urn:ietf:params:xml:ns:yang:ietf-igmp-ml-d-proxy

prefix: imp

reference: RFC XXXX

---

## 6. Normative References

- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, August 2006.

- [RFC4605] B. Fenner, H. He, B. Haberman and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD) - Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.
- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, July 2013.
- [RFC8342] M. Bjorklund and J. Schoenwaelder, "Network Management Datastore Architecture (NMDA)", RFC 8342, March 2018.
- [RFC8343] M. Bjorklund, "A YANG Data Model for Interface Management", RFC 8343, March 2018.
- [draft-ietf-pim-igmp-mld-yang-06] X. Liu, F. Guo, M. Sivakumar, P. McAllister, A. Peter, "A YANG data model for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD)", draft-ietf-pim-igmp-mld-yang-06, Oct 20, 2017.
- [draft-dsdt-nmda-guidelines-01] M. Bjorklund, J. Schoenwaelder, P. Shafer, K. Watsen, R. Wilton, "Guidelines for YANG Module Authors (NMDA)", draft-dsdt-nmda-guidelines-01, May 2017
- [draft-ietf-netmod-revised-datastores-03] M. Bjorklund, J. Schoenwaelder, P. Shafer, K. Watsen, R. Wilton, "Network Management Datastore Architecture", draft-ietf-netmod-revised-datastores-03, July 3, 2017

Authors' Addresses

Hongji Zhao  
Ericsson (China) Communications Company Ltd.  
Ericsson Tower, No. 5 Lize East Street,  
Chaoyang District Beijing 100102, P.R. China  
Email: hongji.zhao@ericsson.com

Xufeng Liu  
Volta Networks  
USA  
Email: Xufeng.liu.ietf@gmail.com

Yisong Liu  
China Mobile  
China  
Email: liuyisong@chinamobile.com

Mani Panchanathan  
Cisco  
India  
Email: mapancha@cisco.com

Mahesh Sivakumar  
Juniper Networks  
1133 Innovation Way  
Sunnyvale, California  
USA  
Email: sivakumar.mahesh@gmail.com



Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 4, 2021

V. Kamath  
VMware  
R. Chokkanathapuram Sundaram  
Cisco Systems, Inc.  
R. Banthia  
Apstra  
A. Gopal  
Cisco Systems, Inc.  
October 31, 2020

PIM Null-Register packing  
draft-ietf-pim-null-register-packing-06

Abstract

In PIM-SM networks PIM Register messages are sent by the Designated Router (DR) to the Rendezvous Point (RP) to signal the presence of Multicast sources in the network. There are periodic PIM Null-Registers sent by a DR to the RP to keep the state alive at the RP as long as the source is active. The PIM Null-Register message carries information about a single Multicast source and group.

This document defines a standard to send information about multiple multicast sources and multicast groups in a single PIM Null-Register message, in a packed format. This document also discusses the interoperability between PIM routers that support the new packed message format and PIM routers that do not support the new format.

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 4, 2021.

## Copyright Notice

Copyright (c) 2020 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 . . . . .	2
1.1. Conventions used in this document . . . . .	3
1.2. Terminology . . . . .	3
2. Packed Register Capability . . . . .	3
3. PIM Packed Null-Register message . . . . .	4
4. PIM Packed Register-Stop message format . . . . .	5
5. Protocol operation . . . . .	6
6. PIM Anycast RP considerations . . . . .	6
7. PIM RP router version downgrade . . . . .	7
8. Fragmentation consideration . . . . .	7
9. Security Considerations . . . . .	7
10. IANA Considerations . . . . .	8
11. Acknowledgments . . . . .	8
12. Normative References . . . . .	8
Authors' Addresses . . . . .	8

## 1. Introduction

PIM Null-Registers are sent by a DR periodically for Multicast streams to keep the states active on the RP, as long as the source is active. As the number of multicast sources increases, the number of PIM Null-Register messages that are sent also increases. This results in more PIM packet processing at the RP and at the DRs.

The control plane policing (COPP), monitors the packets that are processed by the control plane. Due to the high rate at which Null-Registers are received at the RP, this can lead to COPP drops of Multicast PIM Null-Register messages. This document defines a method to efficiently pack multiple PIM Null-Registers [[RFC7761] (Section 4.4)] and Register-Stops [[RFC7761] (Section 3.2)] into a

single message as these packets anyway do not contain encapsulated data.

The document also discusses interoperability with PIM routers that do not understand the new packet format.

1.1. Conventions used in this document

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] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

RP: Rendezvous Point

DR: Designated Router

2. Packed Register Capability

To ensure compatibility with routers that do not support processing of the packed format, A router (DR) can decide to pack multiple Null-Register messages based on the capability received from the RP as part of Register-Stop. Thus a DR will switch to the packed format only when it learns RP is capable of handling the packed Null-Register messages.

Conversely, a DR that does not support the new format can continue generating the PIM Null-Register using the current format. To exchange the capability information in the Register-Stop message, the "reserved" field is used to indicate this capability in those Register-Stop messages. One bit of the reserved field is used to indicate the "packing" capability (P bit). The rest of the bits in the "Reserved" field will be retained for future use.

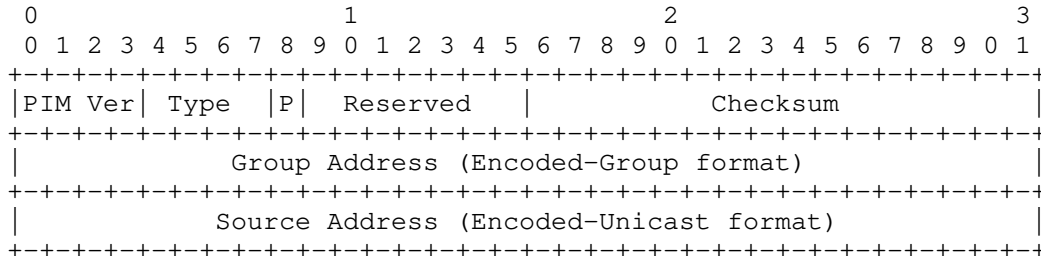


Figure 1: PIM Register-Stop message with capability option

PIM Version, Type, Checksum, Group Address, Source Address:

Same as [RFC7761] (Section 4.9.4)

P:

Capability bit (flag bit 7) used to indicate support for the Packed-Register Capability

3. PIM Packed Null-Register message

PIM Packed Null-Register message format includes a count to indicate the number of Null-Register records in the message.

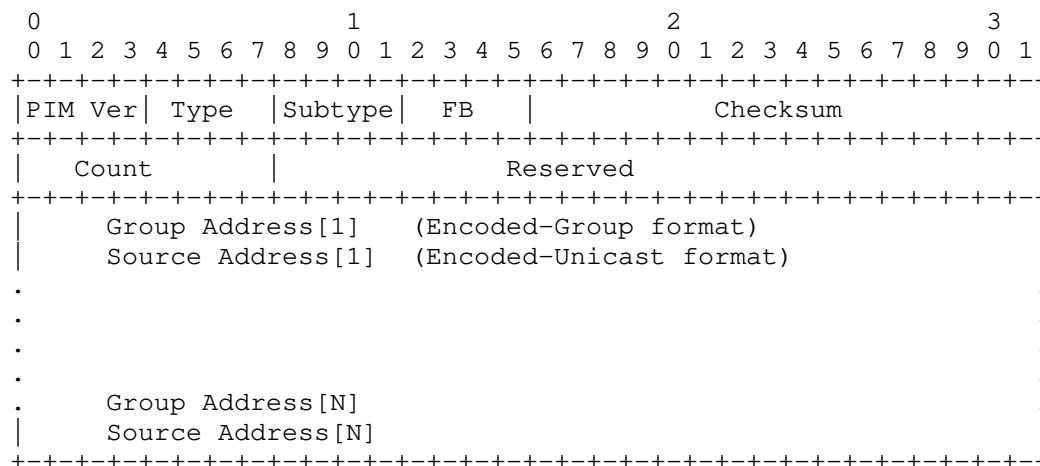


Figure 2: PIM Packed Null-Register message format

PIM Version, Reserved, Checksum:

Same as [RFC7761] (Section 4.9.3)

Type, SubType:

The new packed Null-Register Type and SubType values TBD. [RFC8736]

Count:

The number of packed Null-Register records. A record consists of a Group Address and Source Address pair.

Group Address, Source Address:

Same as [RFC7761] (Section 4.9.4)

4. PIM Packed Register-Stop message format

The PIM Packed Register-Stop message includes a count to indicate the number of records that are present in the message.

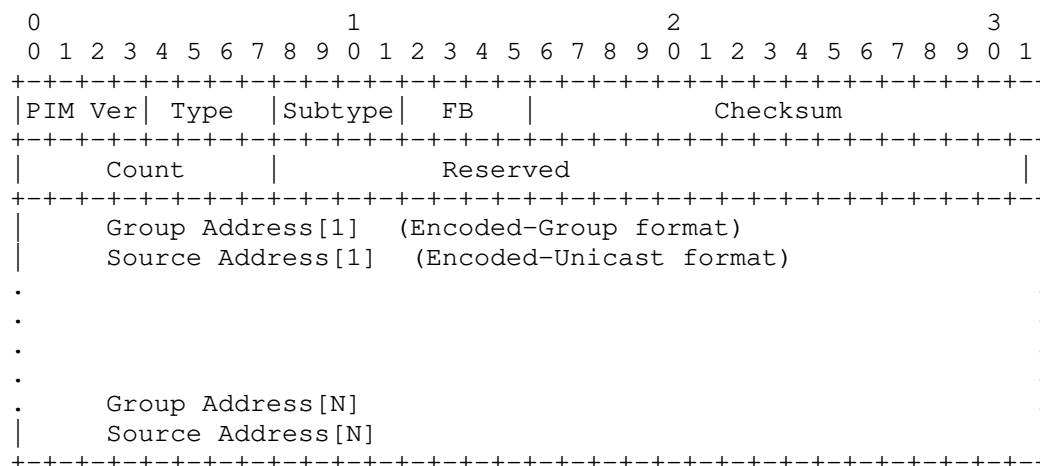


Figure 3: PIM Packed Register-Stop message format

PIM Version, Reserved, Checksum:

Same as [RFC7761] (Section 4.9.4)

Type:

The new Register Stop Type and SubType values TBD

Count:

The number of packed Register-Stop records. A record consists of a Group Address and Source Address pair.

Group Address, Source Address:

Same as [RFC7761] (Section 4.9.4)

## 5. Protocol operation

The following combinations are possible:

1. DR and RP both support the PIM Packed Register format
  - \* As specified in [[RFC7761]], the DR sends PIM Register messages towards the RP when a new source is detected.
  - \* An RP supporting this specification SHOULD set the P-bit in the corresponding Register-Stop messages.
  - \* When a Register-Stop message with the P-bit set is received, the DR SHOULD send Packed Null-Register messages (Section 3) to the RP instead of multiple Register messages with the N-bit set ([[RFC7761]]).
  - \* The RP, after receiving a Packed Null-Register message SHOULD start sending Packed Register-Stop messages (Section 4) to the corresponding DR instead of individual Register-Stop messages.
2. DR supports but RP does not support PIM Packed Register format
  - \* As specified in [[RFC7761]], DR sends PIM Register towards the RP.
  - \* RP sends a Register-Stop in the [[RFC7761]] without any capability information.
  - \* DR then sends Null-Registers in the [[RFC7761]] format.
3. RP supports but DR doesn't support the PIM Packed Register format
  - \* As specified in [[RFC7761]], DR sends the PIM Register towards the RP.
  - \* RP sends a PIM Packed Register-Stop towards the DR that includes capability information.
  - \* Since the the DR doesn't support the PIM Packet Register format, it will always send registers according to [[RFC7761]].

## 6. PIM Anycast RP considerations

The PIM Packed Register format should be enabled only if it is supported by all PIM Anycast RP [[RFC4610]] members in the RP set for An RP address.

## 7. PIM RP router version downgrade

Consider a PIM RP router that supports PIM Register Packing and then downgrades to a software version which does not support PIM Register Packing. The DR that sends the PIM Packed Register message will not get a PIM Register-Stop message back. In such scenarios the DR MUST send an unpacked PIM Register and check the PIM Register-Stop to see if the capability bit (P-bit) for PIM Packed Register is set or not. If it is not set then the DR will continue sending unpacked PIM Register messages.

## 8. Fragmentation consideration

When building a PIM Packed Register message using the packed-register format, a router should include as many records as possible based on the path MTU towards RP, if path MTU discovery is done. Otherwise, the number of records should be limited to the MTU of the outgoing interface.

## 9. Security Considerations

General Register messages security considerations from [[RFC7761]] apply. As mentioned in [[RFC7761]], Register messages and Register-Stop messages are forwarded by intermediate routers to their destination using normal IP forwarding. Without data origin authentication, an attacker who is located anywhere in the network may be able to forge a Register or Register-Stop message. We next consider the effect of a forgery of each of these messages. By forging a Register message, an attacker can cause the RP to inject forged traffic onto the shared multicast tree.

By forging a Register-Stop message, an attacker can prevent a legitimate DR from registering packets to the RP. This can prevent local hosts on that LAN from sending multicast packets. The above two PIM messages are not changed by intermediate routers and need only be examined by the intended receiver. Thus, these messages can be authenticated end-to-end. Attacks on Register and Register-Stop messages do not apply to a PIM-SSM-only implementation, as these messages are not required for PIM-SSM.

There is another case where a spoof Register can be sent to make it appear that it is from the RP, and that the RP supports this new packed capability when it does not. This can cause Null-Registers to not be received by the RP. But standard methods to prevent spoofing should take care of this case. Spoofing methods like uRPF, or on domain boundary, filter out packets coming from the outside from addresses that belong to routers inside.

All these considerations apply equally to the [[RFC7761]] messages.

#### 10. IANA Considerations

This document requires the assignment of "Capability bit" (P-bit), flag bit 7 in the PIM Null-Register message.

This document requires the assignment of 2 new PIM message types for the "PIM Packed Register" and "PIM Register Stop" in the PIM Message Types registry.

#### 11. Acknowledgments

The authors would like to thank Stig Venaas, Anish Peter, Zheng Zhang and Umesh Dudani for their helpful comments on the draft.

#### 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>>.
- [RFC4610] Farinacci, D. and Y. Cai, "Anycast-RP Using Protocol Independent Multicast (PIM)", RFC 4610, DOI 10.17487/RFC4610, August 2006, <<https://www.rfc-editor.org/info/rfc4610>>.
- [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>>.
- [RFC8736] Venaas, S. and A. Retana, "PIM Message Type Space Extension and Reserved Bits", RFC 8736, DOI 10.17487/RFC8736, February 2020, <<https://www.rfc-editor.org/info/rfc8736>>.

#### Authors' Addresses

Vikas Ramesh Kamath  
VMware  
3401 Hillview Ave  
Palo Alto CA 94304  
USA

Email: [vkamath@vmware.com](mailto:vkamath@vmware.com)



Ramakrishnan Chokkanathapuram Sundaram  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: ramaksun@cisco.com

Raunak Banthia  
Apstra  
333 Middlefield Rd STE 200  
Menlo Park CA 94025  
USA

Email: rbanthia@apstra.com

Ananya Gopal  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: ananygop@cisco.com

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 3, 2021

D. Voyer, Ed.  
Bell Canada  
C. Filsfils  
R. Parekh  
Cisco Systems, Inc.  
H. Bidgoli  
Nokia  
Z. Zhang  
Juniper Networks  
October 30, 2020

Segment Routing Point-to-Multipoint Policy  
draft-ietf-pim-sr-p2mp-policy-01

Abstract

This document describes an architecture to construct a Point-to-Multipoint (P2MP) tree to deliver Multi-point services in a Segment Routing domain. A SR P2MP tree is constructed by stitching a set of Replication segments together. A SR Point-to-Multipoint (SR P2MP) Policy is used to define and instantiate a P2MP tree which is computed by a PCE.

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 <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 3, 2021.

## Copyright Notice

Copyright (c) 2020 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. P2MP Tree . . . . .	3
2.1. Sharing Replication segments across P2MP trees . . . . .	4
3. SR P2MP Policy . . . . .	5
4. Using Controller to build a P2MP Tree . . . . .	6
4.1. Provisioning SR P2MP Policy Creation . . . . .	6
4.1.1. API . . . . .	6
4.1.2. Invoking API . . . . .	7
4.2. P2MP Tree Computation . . . . .	7
4.2.1. Topology Discovery . . . . .	8
4.2.2. Capability and Attribute Discovery . . . . .	8
4.3. Instantiating P2MP tree on nodes . . . . .	8
4.3.1. PCEP . . . . .	8
4.3.2. BGP . . . . .	9
4.3.3. NetConf . . . . .	9
4.4. Protection . . . . .	9
4.4.1. Local Protection . . . . .	9
4.4.2. Path Protection . . . . .	9
5. IANA Considerations . . . . .	9
6. Security Considerations . . . . .	9
7. Acknowledgements . . . . .	9
8. Contributors . . . . .	10
9. References . . . . .	11
9.1. Normative References . . . . .	11
9.2. Informative References . . . . .	11
Appendix A. Illustration of SR P2MP Policy and P2MP Tree . . . . .	12
A.1. P2MP Tree with non-adjacent Replication Segments . . . . .	12
A.1.1. SR-MPLS . . . . .	13
A.2. P2MP Tree with adjacent Replication Segments . . . . .	14
A.2.1. SR-MPLS . . . . .	14
Authors' Addresses . . . . .	16

## 1. Introduction

A Multi-point service delivery could be realized via P2MP trees in a Segment Routing domain [RFC8402]. A P2MP tree spans from a Root node to a set of Leaf nodes via intermediate Replication nodes. It consists of a Replication segment [I-D.ietf-spring-sr-replication-segment] at the root node, one or more Replication segments at Leaf nodes and intermediate Replication nodes. The Replication segments are stitched together.

A Segment Routing P2MP policy, a variant of the SR Policy [I-D.ietf-spring-segment-routing-policy], is used to define a P2MP tree. A PCE is used to compute the tree from the Root node to the set of Leaf nodes via a set of replication nodes. The PCE then instantiates the P2MP tree in the SR domain by signaling Replication segments to Root, replication and Leaf nodes using various protocols (PCEP, BGP, NetConf etc.).

## 2. P2MP Tree

A P2MP tree in a SR domain connects a Root to a set of Leaf nodes via a set of intermediate Replication nodes. It consists of a Replication segment at the root stitched to Replication segments at intermediate Replication nodes eventually reaching the Leaf nodes.

The Replication SID of the Replication Segment at Root node is called Tree-SID. The Tree-SID SHOULD also be used as Replication SID of Replication segments at Replication and Leaf nodes. The Replication segments at Replication and Leaf nodes MAY use Replication SIDs that are not same as the Tree-SID.

The Replication segment at Root of a P2MP tree MUST be associated with that P2MP tree (i.e. <Root, Tree-ID> identifier in SR P2MP policy section below) to map a Multi-point service to the tree. A Replication segment that terminates a P2MP tree at a Leaf node MUST be associated with the P2MP tree to determine the context for a Multi-point service. The information that can be used to derive this association is specific to encoding of the protocol (PCEP, BGP, NetConf etc.) used to instantiate the Replication segment for a P2MP tree. Replication segments at intermediate Replication nodes of a tree are also associated with that tree.

A PCE MAY decide not instantiate Replication segments at Leaf nodes of a P2MP tree if it is known a priori that Multi-point services mapped to the P2MP tree can be identified using a context that is globally unique in SR domain. A Multi-point service context assigned from "Domain-wide Common Block" (DCB) [I-D.ietf-bess-mvpn-evpn-aggregation-label] is an example of globally

unique context. SR-MPLS Segment Routing Global Block (SRGB) [RFC8402] MAY be used to allocate globally unique Multi-point service contexts, but it is NOT RECOMMENDED to do so as the service contexts only need to be unique at service edge nodes. In this case, Replication nodes connecting to Leaf nodes SHOULD use Penultimate-Hop Pop (PHP) behavior to pop Tree-SID from a packet.

A packet steered into a P2MP tree is replicated by the Replication segment at Root node to each downstream node in the Replication segment, with the Replication SID of the Replication Segment at the downstream node. A downstream node could be a Leaf node or an intermediate Replication node. In the latter case, replication continues with the Replication segments until all Leaf nodes are reached. A packet is steered into a P2MP tree in two ways:

- o Based on a local policy-based routing at the Root node.
- o Based on steering via the Tree-SID at the Root node.

#### 2.1. Sharing Replication segments across P2MP trees

Two or more P2MP trees MAY share a Replication segment at Root or Replication nodes if at minimum the first condition below is satisfied. A tree always has its own Replication segment at its root even if shares another Replication segment. A tree that shares another Replication segment may or may not have its own Replication segment on its Leaf nodes. If not, the second and third conditions apply to such situations.

1. The Leaf nodes reached via a shared Replication segment must be subset of Leaf or Replication nodes of the P2MP trees that share this segment. Note if a Replication segment is shared, all its downstream Replication segments are also shared.
2. Some Multi-point services realized by the P2MP trees may need service context (e.g. packets are for certain VPNs, and/or from certain nodes). If the trees do not have their own Replication segments at their Leaf nodes then the packets transported on the P2MP trees MUST carry a service context that does not rely on the tree or root identification, e.g. a service label assigned from Domain-wide Common Block or common SRGB for SR-MPLS.
3. For some Multi-point services using P2MP trees that share Replication segments, packets transported on these trees MAY require a Tree context (e.g. MVPN Extranet [RFC7900] to avoid certain ambiguities - see Section 2.3.1 of RFC 7900). In this case, the trees MUST have their own Replication segments on the

Leaf nodes. For SR-MPLS, this is similar to "tunnel stacking" concept.

Sharing of a Replication segment for P2MP trees is OPTIONAL. Exact procedures to ensure validity of above conditions across PM2P services on nodes of a Segment Routing domain are outside the scope of this document.

### 3. SR P2MP Policy

The SR P2MP policy is a variant of an SR policy [I-D.ietf-spring-segment-routing-policy] and is used to instantiate SR P2MP trees.

A SR P2MP Policy is identified by the tuple <Root, Tree-ID>, where:

- o Root: The address of Root node of P2MP tree instantiated by the SR P2MP Policy
- o Tree-ID: A identifier that is unique in context of the Root. This is an unsigned 32-bit number.

A SR P2MP Policy is defined by following elements:

- o Leaf nodes: A set of nodes that terminate the P2MP trees.
- o Candidate Paths: See below.

A SR P2MP policy is provisioned on a PCE to instantiate the P2MP tree. The Tree-SID SHOULD be used as Binding SID of the P2MP policy. A PCE computes the P2MP tree and instantiates Replication segments at Root, Replication and Leaf nodes. When Replication segments are not shared across P2MP trees, the Root and Tree-ID of the SR P2MP policy are mapped to Replication-ID element of the Replication segment identifier i.e the SR Replication segment identifier is <Root, Tree-ID, Node-ID>. A shared Replication segment MAY be identified with zero Root-ID address (0.0.0.0 for IPv4 and :: for IPv6) and a Replication-ID that is unique in context of Node address where the Replication segment is instantiated when it is not associated a particular tree.

A SR P2MP Policy has one or more Candidate paths. The active Candidate path is selected based on the tie breaking rules amongst the candidate-paths as specified in [I-D.ietf-spring-segment-routing-policy]. Each candidate path has a set of topological/resource constraints and/or optimization objectives which determine the P2MP tree for that Candidate path. Tree-SID is an identifier of the P2MP tree of the candidate path in

the forwarding plane. It is instantiated in the forwarding plane at Root node, intermediate Replication nodes and Leaf nodes. The Tree-SID MAY be different at Replication and Leaf nodes.

4. Using Controller to build a P2MP Tree

A P2MP tree can be built using a Path Computation Element (PCE). This section outlines a high-level architecture for such an approach.

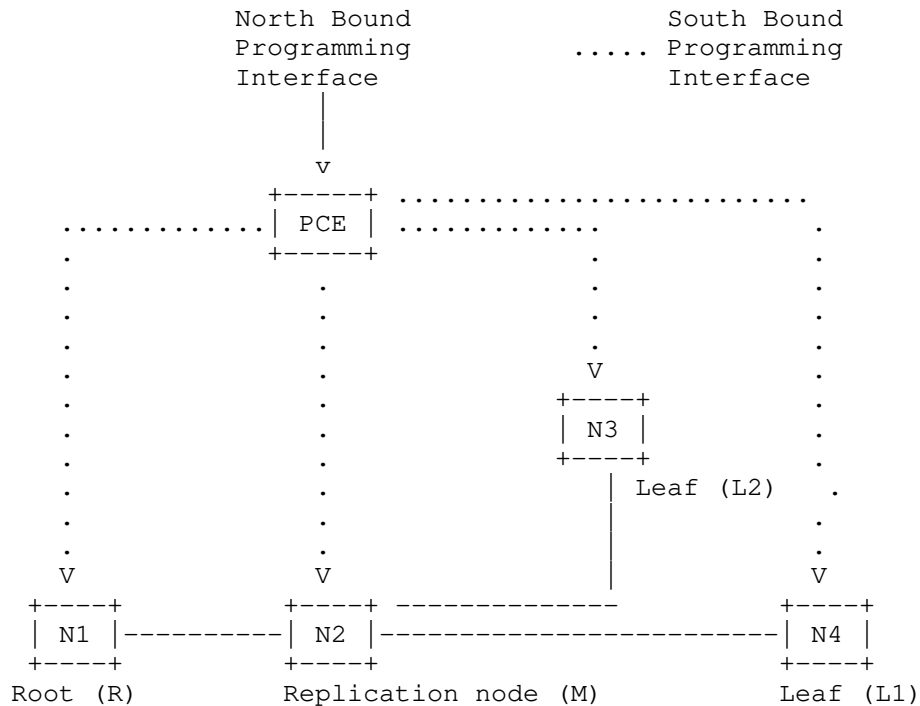


Figure 1: Centralized Control Plane Model

4.1. Provisioning SR P2MP Policy Creation

A SR P2MP policy can be instantiated and maintained in a centralized fashion using a Path Computation Element (PCE).

4.1.1. API

North-bound APIs on a PCE can be used to:

1. Create SR P2MP policy: CreateSRP2MPPolicy<Root, Tree-ID>

2. Delete SR P2MP policy: DeleteSRP2MPPolicy<Root, Tree-ID>
3. Modify SR P2MP policy Leaf Set: SRP2MPPolicyLeafSetModify<Root, Tree-ID, {Leaf Set}>
4. Create a Candidate Path for SR P2MP policy:  
CreateSRP2MPCandidatePath<Root, Tree-ID, <CP-ID>>
5. Delete a Candidate Path for SR P2MP policy:  
DeleteSRP2MPCandidatePath<Root, Tree-ID, <CP-ID>>
6. Update a Candidate Path for SR P2MP policy:  
UpdateSRP2MPCandidatePath<Root, Tree-ID, <CP-ID>, Preference, [Constraints], [Optimization], ...>

CP-ID is identifier of a Candidate Path within a SR P2MP policy. One possible identifier is the tuple <Protocol-Origin, originator, discriminator> as specified in [I-D.ietf-spring-segment-routing-policy].

Note these are conceptual APIs. Actual implementations may offer different APIs as long as they provide same functionality. For example, API might allow symbolic name to be assigned for a P2MP policy or APIs might allow individual Leaf nodes to be added or deleted from a policy instead of an update operation.

#### 4.1.2. Invoking API

Interaction with a PCE can be via PCEP, REST, Netconf, gRPC, CLI. Yang model shall be developed for this purpose as well.

#### 4.2. P2MP Tree Computation

An entity (an operator, a network node or a machine) provisions a SR P2MP policy by specifying the addresses of the root (R) and set of leaves {L} as well as Traffic Engineering (TE) attributes of Candidate paths via a suitable North-Bound API. The PCE computes the tree of Active candidate path. The PCE MAY compute P2MP trees for all Candidate paths., If tree computation is successful, PCE instantiates the P2MP tree(s) using Replication segments on Root, Replication, and Leaf nodes.

Candidate path constraints shall include link color affinity, bandwidth, disjointness (link, node, SRLG), delay bound, link loss, etc. Candidate path shall be optimized based on IGP or TE metric or link latency.



The Tree SID of Candidate path of a SR P2MP policy can be either dynamically allocated by the PCE or statically assigned by entity provisioning the SR P2MP policy. Ideally, same Tree-SID SHOULD be used for Replication segments at Root, Replication, and Leaf nodes. Different Tree-SIDs MAY be used at replication node(s) if it is not feasible to use same Tree SID.

A PCE can modify a P2MP tree following network element failure or in case a better path can be found based on the new network state. In this case, the PCE may want to setup the new instance of the tree and remove the old instance of the tree from the network in order to minimize traffic loss. In this case, the instances of trees for all the Candidate paths of a P2MP policy can be identified by an Instance-ID which is unique in context of the P2MP policy. As such, the identifier of non-shared Replication segments used to instantiate these trees becomes <Root-ID, Tree-ID, Node-ID, Instance-ID>.

A PCE shall be capable of computing paths across multiple IGP areas or levels as well as Autonomous Systems (ASs).

#### 4.2.1. Topology Discovery

A PCE shall learn network topology, TE attributes of link/node as well as SIDs via dynamic routing protocols (IGP and/or BGP-LS). It may be possible for entities to pass topology information to PCE via north-bound API.

#### 4.2.2. Capability and Attribute Discovery

It shall be possible for a node to advertise SR P2MP tree capability via IGP and/or BGP-LS. Similarly, a PCE can also advertise its P2MP tree computation capability via IGP and/or BGP-LS. Capability advertisement allows a network node to dynamically choose one or more PCE(s) to obtain services pertaining to SR P2MP policies, as well as PCE to dynamically identify SR P2MP tree capable nodes.

#### 4.3. Instantiating P2MP tree on nodes

Once a PCE computes a P2MP tree for Candidate path of SR P2MP policy, it needs to instantiate the tree on the relevant network nodes via Replication segments. The PCE can use various protocols to program the Replication segments as described below.

##### 4.3.1. PCEP

PCE Protocol (PCEP) has been traditionally used:

1. For a head-end to obtain paths from a PCE.

## 2. A PCE to instantiate SR policies.

PCEP protocol can be stateful in that a PCE can have a stateful control of an SR policy on a head-end which has delegated the control of the SR policy to the PCE. PCEP shall be extended to provision and maintain SR P2MP trees in a stateful fashion.

### 4.3.2. BGP

BGP has been extended to instantiate and report SR policies. It shall be extended to instantiate and maintain P2MP trees for SR P2MP policies.

### 4.3.3. NetConf

TBD

## 4.4. Protection

### 4.4.1. Local Protection

A network link, node or path on the tree of a P2MP tree can be protected using SR policies computed by PCE. The backup SR policies shall be programmed in forwarding plane in order to minimize traffic loss when the protected link/node fails. It is also possible to use node local Fast Re-Route protection mechanisms (LFA) to protect link/nodes of P2MP tree.

### 4.4.2. Path Protection

It is possible for PCE create a disjoint backup tree for providing end-to-end path protection.

## 5. IANA Considerations

This document makes no request of IANA.

## 6. Security Considerations

There are no additional security risks introduced by this design.

## 7. Acknowledgements

The authors would like to acknowledge Siva Sivabalan, Mike Koldychev and Vishnu Pavan Beeram for their valuable inputs..

8. Contributors

Clayton Hassen  
Bell Canada  
Vancouver  
Canada

Email: clayton.hassen@bell.ca

Kurtis Gillis  
Bell Canada  
Halifax  
Canada

Email: kurtis.gillis@bell.ca

Arvind Venkateswaran  
Cisco Systems, Inc.  
San Jose  
US

Email: arvvenka@cisco.com

Zafar Ali  
Cisco Systems, Inc.  
US

Email: zali@cisco.com

Swadesh Agrawal  
Cisco Systems, Inc.  
San Jose  
US

Email: swaagraw@cisco.com

Jayant Kotalwar  
Nokia  
Mountain View  
US

Email: jayant.kotalwar@nokia.com

Tanmoy Kundu  
Nokia  
Mountain View  
US

Email: tanmoy.kundu@nokia.com

Andrew Stone  
Nokia  
Ottawa  
Canada

Email: andrew.stone@nokia.com

Tarek Saad  
Juniper Networks  
Canada

Email:tsaad@juniper.net

## 9. References

### 9.1. Normative References

[I-D.ietf-spring-segment-routing-policy]  
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-08 (work in progress), July 2020.

[I-D.ietf-spring-sr-replication-segment]  
Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "SR Replication Segment for Multi-point Service Delivery", draft-ietf-spring-sr-replication-segment-00 (work in progress), July 2020.

[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>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

### 9.2. Informative References

[I-D.ietf-bess-mvpn-evpn-aggregation-label]  
Zhang, Z., Rosen, E., Lin, W., Li, Z., and I. Wijnands, "MVPN/EVPN Tunnel Aggregation with Common Labels", draft-ietf-bess-mvpn-evpn-aggregation-label-03 (work in progress), October 2019.

[RFC7900] Rekhter, Y., Ed., Rosen, E., Ed., Aggarwal, R., Cai, Y., and T. Morin, "Extranet Multicast in BGP/IP MPLS VPNs", RFC 7900, DOI 10.17487/RFC7900, June 2016, <<https://www.rfc-editor.org/info/rfc7900>>.

#### Appendix A. Illustration of SR P2MP Policy and P2MP Tree

Consider the following topology:




Figure 1

In these examples, the Node-SID of a node Rn is N-SIDn and Adjacency-SID from node Rm to node Rn is A-SIDmn. Interface between Rm and Rn is Lmn.

Assume PCE is provisioned following SR P2MP policy at Root R1 with Tree-ID T-ID:

```

SR P2MP Policy <R1,T-ID>:
  Leaf Nodes: {R2, R6, R7}
  Candidate-path 1:
    Optimize: IGP metric
    Tree-SID: T-SID1
  
```

The PCE is responsible for P2MP tree computation. Assume PCE instantiates P2MP trees by signalling non-shared Replication segments i.e. Replication-ID of these Replication Segments is <Root, Tree-ID>. If a Candidate-path can have multiple instances of P2MP trees, the Replication-ID is <Root, Tree-ID, Instance-ID>. In this example, we assume one instance of P2MP tree for a candidate-path. All Replication Segments use the Tree-SID T-SID1 as Replication-SID.

##### A.1. P2MP Tree with non-adjacent Replication Segments

Assume PCE computes a P2MP tree with Root node R1, Intermediate and Leaf node R2, and Leaf nodes R6 and R7. The PCE instantiates the P2MP tree by stitching Replication Segments at R1, R2, R6 and R7. Replication Segment at R1 replicates to R2. Replication Segment at R2 replicates to R6 and R7. Note nodes R3, R4 and R5 do not have any Replication Segment state for the tree.

## A.1.1. SR-MPLS

The Replication Segment state at nodes R1, R2, R6 and R7 is shown below.

Replication Segment at R1:

Replication Segment <R1,T-ID,R1>:

Replication SID: T-SID1

Replication State:

R2: <T-SID1->L12>

Replication to R2 steers packet directly to the node on interface L12.

Replication Segment at R2:

Replication Segment <R1,T-ID,R2>:

Replication SID: T-SID1

Replication State:

R2: <Leaf>

R6: <N-SID6, T-SID1>

R7: <N-SID7, T-SID1>

R2 is a Bud-Node. It performs role of Leaf as well as a transit node replicating to R6 and R7. Replication to R6, using N-SID6, steers packet via IGP shortest path to that node. Replication to R7, using N-SID7, steers packet via IGP shortest path to R7 via either R5 or R4 based on ECMP hashing.

Replication Segment at R6:

Replication Segment <R1,T-ID,R6>:

Replication SID: T-SID1

Replication State:

R6: <Leaf>

Replication Segment at R7:

Replication Segment <R1,T-ID,R7>:

Replication SID: T-SID1

Replication State:

R7: <Leaf>

When a packet is steered into the SR P2MP Policy at R1:

- o Since R1 is directly connected to R2, R1 performs PUSH operation with just <T-SID1> label for the replicated copy and sends it to R2 on interface L12.
- o R2, as Leaf, performs NEXT operation, pops T-SID1 label and delivers the payload. For replication to R6, R2 performs a PUSH operation of N-SID6, to send <N-SID6,T-SID1> label stack to R3. R3 is the penultimate hop for N-SID6; it performs penultimate hop popping, which corresponds to the NEXT operation and the packet is then sent to R6 with <T-SID1> in the label stack. For replication to R7, R2 performs a PUSH operation of N-SID7, to send <N-SID7,T-SID1> label stack to R4, one of IGP ECMP nexthops towards R7. R4 is the penultimate hop for N-SID6; it performs penultimate hop popping, which corresponds to the NEXT operation and the packet is then sent to R7 with <T-SID1> in the label stack.
- o R6, as Leaf, performs NEXT operation, pops T-SID1 label and delivers the payload.
- o R7, as Leaf, performs NEXT operation, pops R-SID7 label and delivers the payload.

#### A.2. P2MP Tree with adjacent Replication Segments

Assume PCE computes a P2MP tree with Root node R1, Intermediate and Leaf node R2, Intermediate nodes R3 and R5, and Leaf nodes R6 and R7. The PCE instantiates the P2MP tree by stitching Replication Segments at R1, R2, R3, R5, R6 and R7. Replication Segment at R1 replicates to R2. Replication Segment at R2 replicates to R3 and R5. Replication segment at R3 replicates to R6. Replication segment at R5 replicates to R7. Note node R4 does not have any Replication Segment state for the tree.

##### A.2.1. SR-MPLS

The Replication Segment state at nodes R1, R2, R3, R5, R6 and R7 is shown below.

Replication Segment at R1:

```
Replication Segment <R1,T-ID,R1>:
  Replication SID: T-SID1
  Replication State:
    R2: <T-SID1->L12>
```

Replication to R2 steers packet directly to tje node on interface L12.

Replication Segment at R2:

Replication Segment <R1,T-ID,R2>:

Replication SID: T-SID1

Replication State:

R2: <Leaf>

R3: <T-SID1->L23>

R5: <T-SID1->L25>

R2 is a Bud-Node. It performs role of Leaf as well as a transit node replicating to R3 and R5. Replication to R3, steers packet directly to the node on L23. Replication to R5, steers packet directly to the node on L25.

Replication Segment at R3:

Replication Segment <R1,T-ID,R3>:

Replication SID: T-SID1

Replication State:

R6: <T-SID1->L36>

Replication to R6, steers packet directly to the node on L36.

Replication Segment at R5:

Replication Segment <R1,T-ID,R5>:

Replication SID: T-SID1

Replication State:

R7: <T-SID1->L57>

Replication to R7, steers packet directly to the node on L57.

Replication Segment at R6:

Replication Segment <R1,T-ID,R6>:

Replication SID: T-SID1

Replication State:

R6: <Leaf>

Replication Segment at R7:

Replication Segment <R1,T-ID,R7>:

Replication SID: T-SID1

Replication State:

R7: <Leaf>

When a packet is steered into the SR P2MP Policy at R1:



- o Since R1 is directly connected to R2, R1 performs PUSH operation with just <T-SID1> label for the replicated copy and sends it to R2 on interface L12.
- o R2, as Leaf, performs NEXT operation, pops T-SID1 label and delivers the payload. It also performs PUSH operation on T-SID1 for replication to R3 and R5. For replication to R6, R2 sends <T-SID1> label stack to R3 on interface L23. For replication to R5, R2 sends <T-SID1> label stack to R5 on interface L25.
- o R3 performs NEXT operation on T-SID1 and performs a PUSH operation for replication to R6 and sends <T-SID1> label stack to R6 on interface L36.
- o R5 performs NEXT operation on T-SID1 and performs a PUSH operation for replication to R7 and sends <T-SID1> label stack to R7 on interface L57.
- o R6, as Leaf, performs NEXT operation, pops T-SID1 label and delivers the payload.
- o R7, as Leaf, performs NEXT operation, pops R-SID7 label and delivers the payload.

#### Authors' Addresses

Daniel Voyer (editor)  
Bell Canada  
Montreal  
CA

Email: [daniel.voyer@bell.ca](mailto:daniel.voyer@bell.ca)

Clarence Filsfils  
Cisco Systems, Inc.  
Brussels  
BE

Email: [cfilsfil@cisco.com](mailto:cfilsfil@cisco.com)

Rishabh Parekh  
Cisco Systems, Inc.  
San Jose  
US

Email: [riparekh@cisco.com](mailto:riparekh@cisco.com)

Hooman Bidgoli  
Nokia  
Ottawa  
CA

Email: [hooman.bidgoli@nokia.com](mailto:hooman.bidgoli@nokia.com)

Zhaohui Zhang  
Juniper Networks

Email: [zzhang@juniper.net](mailto:zzhang@juniper.net)

PIM Working Group  
Internet-Draft  
Intended status: Informational  
Expires: April 2, 2021

O. Komolafe  
Arista Networks  
September 29, 2020

IGMPv3 and MLDv2 Survey Report  
draft-komolafe-pim-igmp-ml-d-survey-report-00

Abstract

The PIM WG intends to progress IGMPv3 and MLDv2 from Proposed Standards to Internet Standards. The WG decided to conduct a survey of operators, vendors and implementors of these and related protocols to gather information about their implementation and deployment. This document presents the results of the survey and briefly summarizes the key findings. The survey indicates that there is widespread deployment and usage of IGMPv3 and MLDv2, with numerous independent implementations interoperating successfully. No major issues with either protocol were identified and, similarly, no major unused features in the specifications were highlighted. These findings suggest that IGMPv3 and MLDv2 are indeed ready for progression to Internet Standards.

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 April 2, 2021.

Copyright Notice

Copyright (c) 2020 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 . . . . .	2
2. Approach . . . . .	3
2.1. Methodology . . . . .	3
3. Responses for Vendors or Host Implementors . . . . .	4
3.1. Protocols Implemented . . . . .	4
3.2. Features Supported . . . . .	4
3.3. Issues Identified . . . . .	4
3.4. Suggestions . . . . .	4
4. Responses for Network Operators . . . . .	5
4.1. Protocols Deployed . . . . .	5
4.2. Features Enabled . . . . .	5
4.3. Interoperability Issues . . . . .	5
4.4. Fallback Mechanism . . . . .	5
4.5. Strengths and Weaknesses of IGMPv3 and MLDv2 . . . . .	5
5. Conclusions . . . . .	6
6. Acknowledgements . . . . .	6
7. References . . . . .	6
7.1. Normative References . . . . .	6
7.2. Informative References . . . . .	7
Appendix A. Questionnaire . . . . .	7
A.1. Questionnaire for Vendors or Host Implementors . . . . .	7
A.1.1. Implementation Status . . . . .	7
A.1.2. Implementation Specifics . . . . .	7
A.1.3. Implementation Perspectives . . . . .	8
A.2. Questionnaire for Network Operators . . . . .	8
A.2.1. Deployment Status . . . . .	8
A.2.2. Deployment Specifics . . . . .	9
A.2.3. Deployment Perspectives . . . . .	10
Author's Address . . . . .	10

## 1. Introduction

Internet Group Management Protocol Version 3 (IGMPv3) [RFC3376] and Multicast Listener Discovery Version 2 (MLDv2) for IPv6 [RFC3810] are currently Proposed Standards. Given the fact that multiple independent implementations of these protocols exist and they have been successfully and widely used operationally, the PIM WG is keen to progress these protocols to Internet Standards. As such, it is

critical to establish if there are features specified in [RFC3376] and [RFC3810] that have not been widely used and also to determine any interoperability issues that have arisen from using the protocols.

Following approach taken for PIM-SM, documented in [RFC7063], the PIM WG has decided that conducting a comprehensive survey on implementations and deployment of IGMPv3 and MLDv2 will provide valuable information to facilitate their progression to Internet Standard.

This document summarizes the findings of the survey.

## 2. Approach

### 2.1. Methodology

The raw survey questions are shown in Appendix A. In order to make the submission and processing of responses as convenient as possible, Tim Chown kindly formatted and posted the survey online using the JISC online surveys tool. The PIM WG chairs subsequently announced the survey, publicizing the URL at which the survey could be completed. In addition to announcing the survey on the relevant IETF WG mailing lists, effort was made to distribute the survey to other forums such as NANOG.

The survey was targeted at:

- Network operators

- Router vendors

- Switch vendors

- Host implementors

Once the deadline for the survey elapsed, Tim Chown collated the responses, anonymizing the data so the responses from a specific operator, vendor or implementor could not be identified.

The questions targeted at vendors or host implementors were answered by 10 respondents. The network operators questions were answered by 14 respondents. (These numbers are comparable with the number of responses to the PIM-SM survey [RFC7063].)

### 3. Responses for Vendors or Host Implementors

#### 3.1. Protocols Implemented

80% or more of the respondents had implemented each of IGMPv1, IGMPv2, IGMPv3, MLDv1 and MLDv2, with IGMPv3 being the only protocol that had been implemented by all the respondents. In contrast, Lightweight IGMPv3 and Lightweight MLDv2 had been implemented by only 20% of the respondents.

#### 3.2. Features Supported

All the respondents supported source filtering with include list. Snooping querier was also a popular feature, with 80% of respondents supporting it. Source filtering with exclude list, snooping proxy, snooping filtering, L2 report flooding, host proxy were moderately popular, with 40%–70% of respondents supporting each of these features. Unicast queries/reports were supported by only 20% of the respondents.

#### 3.3. Issues Identified

No ambiguities or inconsistencies in [RFC3376] and [RFC3810] that made the implementation challenging were identified by any respondent.

#### 3.4. Suggestions

A number of respondents made suggestions to the PIM WG regarding progressing IGMPv3 and MLDv2 to full standards:

- o Add source discovery mechanism to SSM in addition to existing application-based source discovery
- o Improve scalability of query/response messages
- o Deprecate older versions and streamline IGMPv3
- o Allow reports to be sent without a querier
- o Remove source filtering with exclude list as it is not widely used and makes state machine unnecessarily complicated

Each of these points was raised by a different respondent, apart from the last point which was raised by two separate respondents.

## 4. Responses for Network Operators

### 4.1. Protocols Deployed

IGMPv2 was the most widely deployed protocol, with 86% of respondents indicating it is running in their network. Next was IGMPv3 with 79% of respondents indicating it is deployed. However, between only 20% and 36% of respondents indicated they had deployed IGMPv1, MLDv1 and MLDv2. Lightweight IGMPv3 and Lightweight MLDv2 were undeployed.

### 4.2. Features Enabled

Between 20% and 30% of respondents indicated that had enabled Source filtering with include list, source filtering with exclude list, snooping querier, snooping filtering or unicast queries/reports. Snooping proxy and L2 report flooding were enabled by 7% of respondents. No respondent was using host proxy.

### 4.3. Interoperability Issues

Half the respondents indicated they were using equipment with multi-vendor implementations in their network. No interoperability issues were identified.

### 4.4. Fallback Mechanism

36% of respondents indicated there are dependent on the fallback mechanisms between the different protocol versions. 7% of respondents have experienced issues related to this fallback mechanism.

### 4.5. Strengths and Weaknesses of IGMPv3 and MLDv2

A respondent indicated that a significant strength of IGMPv3 was the simplicity introduced by using SSM, avoiding the complexities associated with ASM. The weaknesses associated with IGMPv3 which were identified were:

- o No CPE implementations
- o Automatic fallback makes deployments challenging
- o ASM provides better source filtering (by potentially restricting the acceptance of register messages at the RP) whereas SSM allows only data plane filtering using multicast boundary

## 5. Conclusions

There were a total of 24 respondents to the survey which asked vendors/implementors and network operators questions about IGMPv1, IGMPv2, IGMPv3, Lightweight IGMPv3, MLDv1, MLDv2 and Lightweight MLDv2. A reasonable number of responses were gathered to the survey, allowing some interesting observations to be made. Firstly, and perhaps unsurprisingly, operators use a lower number of protocols and protocol features than have been implemented. Furthermore, there is a relatively lower deployment of the different MLD versions, suggesting that IPv6 multicast is less widely used than IPv4 multicast. No major flaws, inconsistencies or ambiguity in the IGMPv3 [RFC3376] and MLDv2 [RFC3810] specifications were identified. However, a number of issues were raised about the usage of these protocols, notably concerns about the automatic fallback from IGMPv3 to IGMPv2 being sometimes problematic and the loss of certain useful features offered by the ASM control plane with the transition to SSM.

These findings suggest that IGMPv3 and MLDv2 are indeed ready for progression to Internet Standards.

## 6. Acknowledgements

The authors are grateful to Tim Chown for posting the survey online, and for collating and anonymizing the responses.

## 7. References

### 7.1. Normative References

- [RFC1112] Deering, S., "Host Extensions for IP Multicasting", RFC 1112, August 1989.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.



[RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, February 2010.

## 7.2. Informative References

[RFC7063] Zheng, L., Zhang, Z., and R. Parekh, "Survey Report on Protocol Independent Multicast - Sparse Mode (PIM-SM) Implementations and Deployments", RFC 7063, December 2013.

## Appendix A. Questionnaire

### A.1. Questionnaire for Vendors or Host Implementors

Name:

Affiliation/Organization:

Contact Email:

Do you wish to complete the survey anonymously?: Y/N

#### A.1.1. Implementation Status

Which of the following have you implemented?

1. IGMPv1 [RFC1112]?
2. IGMPv2 [RFC2236]?
3. IGMPv3 [RFC3376]?
4. Lightweight IGMPv3 [RFC5790]?
5. MLDv1 [RFC2710]?
6. MLDv2 [RFC3810]?
7. Lightweight MLDv2 [RFC5790]?

#### A.1.2. Implementation Specifics

1. Which IGMPv3 and MLDv2 features have you implemented?
  - A. Source filtering with include list?
  - B. Source filtering with exclude list?

- C. Snooping proxy?
  - D. Snooping querier?
  - E. Snooping filtering?
  - F. L2 Report flooding?
  - G. Host proxy?
  - H. Unicast queries/reports?
2. Have you carried out IGMPv3 or MLDv2 interoperability tests with other implementations?
- A. What issues, if any, arose during these tests?
  - B. How could [RFC3376] and [RFC3810] have helped minimize these issues?

#### A.1.3. Implementation Perspectives

- 1. Which ambiguities or inconsistencies in RFC 3376 or RFC 3810 made the implementation challenging?
- 2. What suggestions would you make to the PIM WG as it seeks to progress IGMPv3 and MLDv2 to Internet Standard?

#### A.2. Questionnaire for Network Operators

Name:

Affiliation/Organization:

Contact Email:

Do you wish to complete the survey anonymously?: Y/N:

##### A.2.1. Deployment Status

Which of the following have you deployed in your network?

- 1. IGMPv1 [RFC1112]?
- 2. IGMPv2 [RFC2236]?
- 3. IGMPv3 [RFC3376]?

4. Lightweight IGMPv3 [RFC5790]?
5. MLDv1 [RFC2710]?
6. MLDv2 [RFC3810]?
7. Lightweight MLDv2 [RFC5790]?

#### A.2.2. Deployment Specifics

1. Which IGMPv3 and MLDv2 features do you use?
  - A. Source filtering with include list?
  - B. Source filtering with exclude list?
  - C. Snooping proxy?
  - D. Snooping querier?
  - E. Snooping filtering?
  - F. L2 Report flooding?
  - G. Host proxy?
  - H. Unicast queries/reports?
2. Are you using equipment with multi-vendor implementations in your IGMPv3/MLDv2 deployment?
  - A. What inter-operability issues, if any, have you experienced?
  - B. How could [RFC3376] and [RFC3810] have helped minimize these issues?
3. Are you using different IGMP versions or different MLD versions in your network?
  - A. Are you dependent on the fallback mechanism between the different versions?
  - B. Have you experienced any issues related to the fallback mechanism between the different versions?
  - C. How could [RFC3376] and [RFC3810] have helped minimize these issues?

A.2.3. Deployment Perspectives

1. Based on your operational experience, What have you found to be the strengths of IGMPv3 or MLDv2?
2. What have you found to be the weaknesses of IGMPv3 or MLDv2?
3. What suggestions would you make to the PIM WG as it seeks to progress IGMPv3 and MLDv2 to Internet Standard?

Author's Address

Olufemi Komolafe  
Arista Networks  
UK

Email: femi@arista.com

IPPM  
Internet-Draft  
Intended status: Standards Track  
Expires: April 23, 2021

H. Song  
M. McBride  
Futurewei Technologies  
G. Mirsky  
ZTE Corp.  
October 20, 2020

Multicast On-path Telemetry Solutions  
draft-song-multicast-telemetry-06

Abstract

This document discusses the requirement of on-path telemetry for multicast traffic. The existing solutions are examined and their issues are identified. Solution modifications are proposed to allow the original multicast tree to be correctly reconstructed without unnecessary replication of telemetry information.

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.

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 April 23, 2021.

## Copyright Notice

Copyright (c) 2020 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 . . . . .	2
2. Requirements for Multicast Traffic Telemetry . . . . .	3
3. Issues of Existing Techniques . . . . .	4
4. Proposed Modifications to Existing Techniques . . . . .	4
4.1. Per-hop postcard using IOAM DEX . . . . .	5
4.2. Per-section postcard . . . . .	7
5. Considerations for Different Multicast Protocols . . . . .	8
5.1. Application in PIM . . . . .	8
5.2. Application in P2MP . . . . .	9
5.3. Application in BIER . . . . .	9
6. Security Considerations . . . . .	10
7. IANA Considerations . . . . .	10
8. Contributors . . . . .	10
9. Acknowledgments . . . . .	10
10. References . . . . .	10
10.1. Normative References . . . . .	10
10.2. Informative References . . . . .	11
Authors' Addresses . . . . .	12

## 1. Introduction

Multicast traffic is an important traffic type in today's Internet. Multicast provides services that are often real time (e.g., online meeting) or have strict QoS requirements (e.g., IPTV, Market Data). Multicast packet drop and delay can severely affect the application performance and user experience.

It is important to monitor the performance of the multicast traffic. Existing OAM techniques cannot gain direct and accurate information about the multicast traffic. New on-path telemetry techniques such as In-situ OAM [I-D.ietf-ippm-ioam-data], Postcard-based Telemetry

[I-D.song-ippm-postcard-based-telemetry], and Hybrid Two-Step (HTS) [I-D.mirsky-ippm-hybrid-two-step] provide promising means to directly monitor the network experience of multicast traffic. However, multicast traffic has some unique characteristics which pose some challenges on efficiently applying such techniques.

When a network contains multicast (p2mp) trees there will be redundant data as data is replicated at branch points. The IP Multicast S,G data is identical from one branch to another on it's way to multiple receivers. When adding iOAM trace data, to multicast packets, we enlarge data packets thus consuming more network bandwidth. Instead of adding iOAM trace data, it could be more efficient to collect the telemetry information using solutions, such as iOAM postcard or HTS, to cut down on the redundant iOAM data. The problem is that a postcard type solution doesn't have a branch identifier.

This draft proposes a set of solutions to this iOAM data redundancy problem. The requirements for multicast traffic telemetry are discussed along with the issues of the existing on-path telemetry techniques. We propose modifications to make these techniques adapt to multicast in order for the original multicast tree to be correctly reconstructed while eliminating redundant data.

## 2. Requirements for Multicast Traffic Telemetry

Multicast traffic is forwarded through a multicast tree. With PIM and P2MP (MLDP, RSVP-TE) the forwarding tree is established and maintained by the multicast routing protocol. With BIER, no state is created in the network to establish a forwarding tree, instead, a bier header provides the necessary information for each packet to know the egress points. Multicast packets are only replicated at each tree branch node for efficiency.

There are several requirements for multicast traffic telemetry, a few of which are:

- o Reconstruct and visualize the multicast tree through data plane monitoring.
- o Gather the multicast packet delay and jitter performance.
- o Find the multicast packet drop location and reason.
- o Gather the VPN state and tunnel information in case of P2MP multicast.

In order to meet these requirements, we need the ability to directly monitor the multicast traffic and derive data from the multicast packets. The conventional OAM mechanisms, such as multicast ping and trace, may not be sufficient to meet these requirements.

### 3. Issues of Existing Techniques

On-path Telemetry techniques that directly retrieve data from multicast traffic's live network experience are ideal to address the above mentioned requirements. The representative techniques include In-situ OAM (IOAM) Trace option [I-D.ietf-ippm-ioam-data], IOAM Direct Export (DEX) option [I-D.ioamteam-ippm-ioam-direct-export], and Postcard-based Telemetry with Packet Marking (PBT-M) [I-D.song-ippm-postcard-based-telemetry]. However, unlike unicast, multicast poses some unique challenges to applying these techniques.

Multicast packets are replicated at each branch node in the corresponding multicast tree. Therefore, there are multiple copies of packets in the network.

If the IOAM trace option is used for on-path data collection, the partial trace data will also be replicated into multiple copies. The end result is that each copy of the multicast packet has a complete trace. Most of the data, however, is redundant. Data redundancy introduces unnecessary header overhead, wastes network bandwidth, and complicates the data processing. In case the multicast tree is large, and the path is long, the redundancy problem becomes severe.

The PBT solutions, including the IOAM DEX option and PBT-M, can be used to eliminate such data redundancy, because each node on the tree only sends a postcard covering local data. However, they cannot track the tree branches properly so it can bring confusion about the multicast tree topology. For example, Node A has two branches, one to Node B and the other to node D, and Node B leads to Node C and Node D leads to Node E. From the received postcards, one cannot tell whether or not Node C(E) is the next hop of Node B(D).

The fundamental reason for this problem is that there is not an identifier (either implicit or explicit) to correlate the data on each branch.

### 4. Proposed Modifications to Existing Techniques

Two solutions are proposed to address the above issues. One is built on PBT and requires augmentation or modification to the instruction header of the IOAM Direct Export Option; the other combines the IOAM trace option and PBT for an optimized solution.



4.1. Per-hop postcard using IOAM DEX

One way to mitigate PBT's multiple tree tracking weakness is to augment it with a branch identifier field. Note that this works for the IOAM DEX option but not for PBT-M because the IOAM DEX option uses an instruction header. To make the branch identifier globally unique, the branch node ID plus an index is used. For example, if Node A has two branches, one to Node B and one to Node C, Node A will use [A, 0] as the branch identifier for the branch to B, and [A, 1] for the branch to C. The identifier is unchanged for each multicast tree instance and carried with the multicast packet until the next branch node. Each postcard needs to include the branch identifier in the export data. The branch identifier, along with the other fields such as flow ID and sequence number, is sufficient for the data analyzer to reconstruct the topology of the multicast tree.

Figure 1 shows an example of this solution. "P" stands for the postcard packet. The square brackets contains the branch identifier. The curly brace contains the telemetry data about a specific node.




Figure 1: Per-hop Postcard

Each branch fork node need to generate the branch ID for each branch in its multicast tree instance and include it in the IOAM DEX option header so the downstream node can learn it. The branch ID contains two parts: the branch fork node ID and a unique branch index.

Figure 2 shows that the branch ID is carried as an optional field after the flow ID and sequence number optional fields in the IOAM DEX option header. A bit "M" in the Flags field is reserved to indicate

the presence of the branch index field. The "M" flag position will be determined later after the other flags are specified in [I-D.ioamteam-ippm-ioam-direct-export].

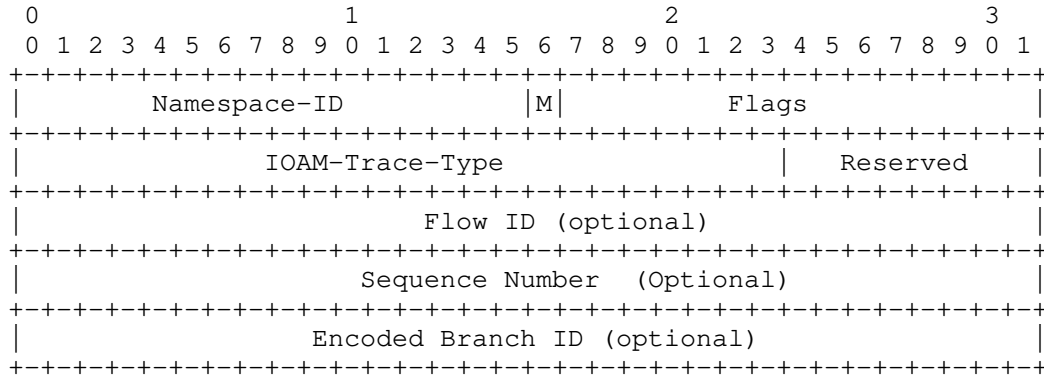


Figure 2: Carry Branch Index in IOAM DEX option header

To avoid introducing a new type of data field to the IOAM DEX option header, we can encode the branch identifier using the existing node ID data field as defined in [I-D.ietf-ippm-ioam-data]. Currently, the node ID field occupies three octets. A simple solution is to shorten the node ID field so a number of bits can be saved to encode the branch index, as shown in Figure 3.




Figure 3: Encode Branch Index with Node ID Method 1

Another encoding method is to use the sum of the node ID and the branch index as the new node ID, as shown in Figure 4. As long as the node IDs are assigned with large enough gap, the telemetry data analyzer can still successfully recover the original node ID and branch index.




Figure 4: Encode Branch Index with Node ID Method 2

Once a node gets the branch ID information from the upstream, it MUST carry this information in its telemetry data export postcards, so the original multicast tree can be correctly reconstructed based on the postcards.

#### 4.2. Per-section postcard

The second solution is a combination of the IOAM trace mode and PBT. To avoid data redundancy at each branch node, the trace data accumulated, to that point, is exported by a postcard before the packet is replicated. In this case, each branch still needs to maintain some identifier to help correlate the postcards for each tree section. The natural way to accomplish this is to simply carry the branch node's data (including its ID) in the trace of each branch. This is also necessary because each replicated multicast packet can have different telemetry data pertaining to this particular copy (e.g., node delay, egress timestamp, and egress interface). As a consequence, the local data exported by each branch node can only contain partial data (e.g., ingress interface and ingress timestamp).

Figure 5 shows an example in a segment of a multicast tree. Node B and D are two branch nodes and they will export a postcard covering the trace data for the previous section. The end node of each path will also need to export the data of the last section as a postcard.




Figure 5: Per-section Postcard

There is no need to modify the IOAM trace mode header format. We just need to configure the branch node to export the postcard and refresh the IOAM header and data.

5. Considerations for Different Multicast Protocols

MTRACEv2 [RFC8487] provides an active probing approach for the tracing of an IP multicast routing path. Mtrace can also provide information such as the packet rates and losses, as well as other diagnostic information. New on-path telemetry techniques will enhance Mtrace, and other existing OAM solutions, with more granular and realtime network status data through direct measurements. There are various multicast protocols that are used to forward the multicast data. Each will require their own unique on-path telemetry solution.

5.1. Application in PIM

PIM-SM [RFC7761] is the most widely used multicast routing protocol deployed today. Of the various PIM modes (PIM-SM, PIM-DM, BIDIR-PIM, PIM-SSM), PIM-SSM is the preferred method due to its simplicity and removal of network source discovery complexity. With all PIM modes, control plane state is established in the network in order to forward multicast UDP data packets. But with PIM-SSM, the discovery of multicast sources is performed outside of the network via HTTP, SDN, etc. IP Multicast packets fall within the range of 224.0.0.0 through

239.255.255.255. The telemetry solution will need to work within this address range and provide telemetry data for this UDP traffic.

The proposed solutions for encapsulating the telemetry instruction header and metadata in IPv4/IPv6 UDP packets are described in [I-D.herbert-ipv4-udpencap-eh] and [I-D.ioametal-ippm-6man-ioam-ipv6-deployment].

## 5.2. Application in P2MP

Multicast Label Distribution Protocol (MLDP) and P2MP RSVP-TE are commonly used within a Multicast VPN (MVPN) environment. MLDP provides extensions to LDP to establish point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) label switched paths (LSPs) in MPLS networks. P2MP RSVP-TE provides extensions to RSVP-TE for establish traffic-engineered P2MP LSPs in MPLS networks. The telemetry solution will need to be able to follow these P2MP paths. The telemetry instruction header and data should be encapsulated into MPLS packets on P2MP paths. A corresponding proposal is described in [I-D.song-mpls-extension-header].

## 5.3. Application in BIER

BIER [RFC8279] adds a new header to multicast packets and allows the multicast packets to be forwarded according to the header only. By eliminating the requirement of maintaining per multicast group state, BIER is more scalable than the traditional multicast solutions.

OAM Requirements for BIER [I-D.ietf-bier-oam-requirements] lists many of the requirements for OAM at the BIER layer which will help in the forming of on-path telemetry requirements as well.

There is also current work to provide solutions for BIER forwarding in ipv6 networks. For instance, a solution, BIER in Non-MPLS IPv6 Networks [I-D.xie-bier-ipv6-encapsulation], proposes a new bier Option Type codepoint from the "Destination Options and Hop-by-Hop Options" IPv6 sub-registry. This is similar to what IOAM proposes for IPv6 transport.

Depending on how the BIER header is encapsulated into packets with different transport protocols, the method to encapsulate the telemetry instruction header and metadata also varies. It is also possible to make the instruction header and metadata a part of the BIER header itself, such as in a TLV.

## 6. Security Considerations

No new security issues are identified other than those discovered by the IOAM and PBT drafts.

## 7. IANA Considerations

The document makes no request of IANA.

## 8. Contributors

TBD

## 9. Acknowledgments

The authors would like to thank Frank Brockners, Tianran Zhou for the comments and advice.

## 10. References

### 10.1. 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>>.
- [RFC4687] Yasukawa, S., Farrel, A., King, D., and T. Nadeau, "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", RFC 4687, DOI 10.17487/RFC4687, September 2006, <<https://www.rfc-editor.org/info/rfc4687>>.
- [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>>.
- [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>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.

[RFC8487] Asaeda, H., Meyer, K., and W. Lee. Ed., "Mtrace Version 2: Traceroute Facility for IP Multicast", RFC 8487, DOI 10.17487/RFC8487, October 2018, <<https://www.rfc-editor.org/info/rfc8487>>.

## 10.2. Informative References

- [I-D.herbert-ipv4-udpencap-eh]  
Herbert, T., "IPv4 Extension Headers and UDP Encapsulated Extension Headers", draft-herbert-ipv4-udpencap-eh-01 (work in progress), March 2019.
- [I-D.ietf-bier-oam-requirements]  
Mirsky, G., Nainar, N., Chen, M., and S. Pallagatti, "Operations, Administration and Maintenance (OAM) Requirements for Bit Index Explicit Replication (BIER) Layer", draft-ietf-bier-oam-requirements-10 (work in progress), May 2020.
- [I-D.ietf-ippm-ioam-data]  
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-10 (work in progress), July 2020.
- [I-D.ioametal-ippm-6man-ioam-ipv6-deployment]  
Bhandari, S., Brockners, F., Mizrahi, T., Kfir, A., Gafni, B., Spiegel, M., Krishnan, S., and M. Smith, "Deployment Considerations for In-situ OAM with IPv6 Options", draft-ioametal-ippm-6man-ioam-ipv6-deployment-03 (work in progress), March 2020.
- [I-D.ioamteam-ippm-ioam-direct-export]  
Song, H., Gafni, B., Zhou, T., Li, Z., Brockners, F., Bhandari, S., Sivakolundu, R., and T. Mizrahi, "In-situ OAM Direct Exporting", draft-ioamteam-ippm-ioam-direct-export-00 (work in progress), October 2019.
- [I-D.mirsky-ippm-hybrid-two-step]  
Mirsky, G., Lingqiang, W., and G. Zhui, "Hybrid Two-Step Performance Measurement Method", draft-mirsky-ippm-hybrid-two-step-05 (work in progress), April 2020.
- [I-D.song-ippm-postcard-based-telemetry]  
Song, H., Zhou, T., Li, Z., Shin, J., and K. Lee, "Postcard-based On-Path Flow Data Telemetry", draft-song-ippm-postcard-based-telemetry-07 (work in progress), April 2020.

[I-D.song-mpls-extension-header]

Song, H., Li, Z., Zhou, T., and L. Andersson, "MPLS Extension Header", draft-song-mpls-extension-header-02 (work in progress), February 2019.

[I-D.xie-bier-ipv6-encapsulation]

Xie, J., Geng, L., McBride, M., Asati, R., Dhanaraj, S., Zhu, Y., Qin, Z., Shin, M., Mishra, G., and X. Geng, "Encapsulation for BIER in Non-MPLS IPv6 Networks", draft-xie-bier-ipv6-encapsulation-08 (work in progress), July 2020.

#### Authors' Addresses

Haoyu Song  
Futurewei Technologies  
2330 Central Expressway  
Santa Clara  
USA

Email: [hsong@futurewei.com](mailto:hsong@futurewei.com)

Mike McBride  
Futurewei Technologies  
2330 Central Expressway  
Santa Clara  
USA

Email: [mmcbride@futurewei.com](mailto:mmcbride@futurewei.com)

Greg Mirsky  
ZTE Corp.

Email: [gregimirsky@gmail.com](mailto:gregimirsky@gmail.com)