

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2019

J. Xie
Huawei Technologies
L. Geng
L. Wang
China Mobile
M. McBride
G. Yan
Huawei
October 22, 2018

Segmented MVPN Using IP Lookup for BIER
draft-xie-bier-mvpn-segmented-05

Abstract

This document specifies an alternative of the control plane and data plane procedures that allow segmented MVPN using the more efficient LIR-pF explicit-tracking when BIER is used as the upstream or downstream or both segments. It requires a segmentation point BFR doing an IP header lookup, which is common for the forwarding procedure on BFER, or the forwarding procedure on ABR with local VPN CEs connected. This document updates [[I-D.ietf-bier-mvpn](#)].

Requirements Language

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

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 25, 2019.

Internet-Draft Segmented MVPN Using IP Lookup for BIER October 2018

Copyright Notice

Copyright (c) 2018 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.	Terminology	3
3.	Problem Statement and Considerations	5
3.1.	Problem Statement	5
3.2.	Considerations	5
4.	Upstream BIER and Downstream BIER (BIER-BIER)	6
4.1.	Explicit Tracking using LIR-pF	6
4.2.	Forwarding using IP lookup on Segmentation Point	8
5.	Upstream P2MP and Downstream BIER (P2MP-BIER)	9
5.1.	Explicit Tracking using LIR-pF	9
5.2.	Forwarding using IP lookup on Segmentation Point	10
6.	Upstream BIER and Downstream P2MP (BIER-P2MP)	11
6.1.	Explicit Tracking using LIR-pF	11
6.2.	Forwarding on Segmentation Point	12
7.	Summary and Recommendations	12
8.	Appendix A - Comparison of Solutions	13
9.	Security Considerations	15
10.	IANA Considerations	15
11.	Acknowledgements	15
12.	References	16
12.1.	Normative References	16
12.2.	Informative References	16
	Authors' Addresses	16

[1.](#) Introduction

When using BIER to transport an MVPN data packet through a BIER domain, an ingress PE functions as a BFIR (see [[RFC8279](#)]). The BFIR must determine the set of BFERs to which the packet needs to be delivered. This can be done through an explicit-tracking function using a LIR and/or LIR-pF flag in BGP MVPN routes, per the

Internet-Draft Segmented MVPN Using IP Lookup for BIER October 2018

[[RFC6513](#)], [[RFC6514](#)], [[RFC6625](#)], [[I-D.ietf-bess-mvpn-expl-track](#)], and [[I-D.ietf-bier-mvpn](#)].

Using a LIR-pF Flag will bring some extra benefits, as [[I-D.ietf-bier-mvpn](#)] and [[I-D.ietf-bess-mvpn-expl-track](#)] have stated. But unfortunately, the LIR-pF explicit tracking for a segmented MVPN deployment is not allowed in the current draft [[I-D.ietf-bier-mvpn](#)], because the draft requires a per-flow upstream-assigned label to do the data-plane per-flow lookup on the segmentation point BFR.

This document specifies an alternative of the control plane and data plane procedures that allow segmented MVPN using the more efficient LIR-pF explicit-tracking when BIER is used as the upstream or downstream or both segments. It requires a segmentation point BFR doing an IP header lookup, which is common for the forwarding procedure on BFER, or the forwarding procedure on ABR with local VPN CEs connected. This will bring some significant benefits to the segmented MVPN deployment, including:

- o Getting a much better multicast join latency by eliminating the round trip interaction of S-PMSI AD routes and Leaf AD routes. Especially, the S-PMSI A-D routes may need a data-driven procedure to trigger, and make the multicast join latency even worse.
- o Greatly reducing the number of S-PMSI A-D routes that BFIR and BFERs need to save.
- o Consolidated forwarding procedure of IP lookup for every BIER Overlay functioning routers, such as BFIR, BFER, segmentation point BFR, and segmentation point BFR with BFER function.

[2.](#) Terminology

Readers of this document are assumed to be familiar with the terminology and concepts of the documents listed as Normative References.

P2MP: This document uses the term "P2MP" for "mLDP P2MP, RSVP-TE P2MP, or IR P2MP".

BIER tunnel: An unidirectional tunnel indicated by an upstream-assigned VpnLabel and the according context such as the RootIP where the VpnLabel is generated. In control plane the BIER tunnel is the PTA of type<BIER> including the MPLS Label of the PTA. One BIER tunnel is bound to an BGP-MVPN SPMSI route. There are no two BGP-MVPN SPMSI routes of different VPN binding the same BIER tunnel, but there are possibly two BGP-MVPN SPMSI routes of the same VPN binding the same BIER tunnel.

P2MP tunnel: An unidirectional tunnel indicated by an downstream label. In control plane the P2MP tunnel is the PTA of type<mLDP> or type <RSVP-TE> or type <IR>. The MPLS Label of the PTA is zero in real implementation. One P2MP tunnel is bound to an BGP-MVPN SPMSI route. There are no two BGP-MVPN SPMSI routes in different VPNs binding the same BIER tunnel, but there are possibly two BGP-MVPN SPMSI routes in the same VPN binding the same P2MP tunnel.

BGP-MVPN FEC: Either a (OrigIP, RD) tuple indicating by an SPMS(OrigIP, RD, *, *), or a (OrigIP, RD, S, G) tuple indicating by an SPMSI(OrigIP, RD, S, G). The BGP-MVPN FEC is used for control plane to build some form of forwarding state on an segment point ABR node. One example of such forwarding state is to use the upstream P2MP tunnel bound to an SPMSI route, and the downstream P2MP tunnel(s) bound to an SPMSI route, with the same BGP-MVPN FEC. Another example of such forwarding state is to use the upstream BIER tunnel bound to an SPMSI route, and the downstream P2MP tunnel(s) and/or BIER tunnel(s).

VRF: Virtual Routing Forwarding. It is usually a number to indicate a VPN/MVPN instance locally when L3VPN is configured. It is also called local VRF.

Pseudo VRF: a Pseudo VRF indicates a VPN/MVPN instance locally on the segment point ABR when downstream BIER forwarding using IP lookup is required. The segment point ABR can be PE of a VPN/MVPN instance at the same time, and then there would have VRF and Pseudo VRF(s) simultaneously. Note that the Pseudo VRF(s) is per the tunnel (BIER tunnel or P2MP tunnel). Numbers of RootIP per a VPN/MVPN, and

numbers of tunnels per a VPN/MVPN and a RootIP, in the BGP-MVPN FEC(s), will cause numbers of the Pseudo VRF allocated.

Local VpnLabel: a Local VpnLabel indicates a upstream-assigned MPLS Label locally on the segment point ABR when downstream BIER forwarding using IP lookup is required. Local VpnLabel is 1:1 to Pseudo VRF, and it is assumed that a Pseudo VRF always indicates a Local VpnLabel within it.

- o Getting a much better multicast join latency by eliminating the round trip interaction of S-PMSI AD routes and Leaf AD routes. Especially, the S-PMSI A-D routes may need a data-driven procedure to trigger, and make the multicast join latency even worse.
- o Greatly reducing the number of S-PMSI A-D routes that BFIR and BFERs need to save.

- o Consolidated forwarding procedure of IP lookup for every BIER Overlay functioning routers, such as BFIR, BFER, segmentation point BFR, and segmentation point BFR with BFER function.

[3.](#) Problem Statement and Considerations

[3.1.](#) Problem Statement

BIER is a stateless multicast forwarding by introducing a multicast-specific BIER header in the data plane. The maximal number of BFERs a packet can reach is limited by the bit string length of a BIER header. For a network with many routers in multiple IGP areas (typically an Inter-Area network), it may be expected to use a segmented MVPN when deploying BIER because operators don't want to replicate packets to many BIER Sets. It is also possible for an incremental deployment of BIER on one area (e.g. one Metro network area) while keeping the P2MP in other areas (e.g. the Core network area, and the other Metro network areas).

However, it is not allowed in the [[I-D.ietf-bier-mvpn](#)] to use a LIR-pF explicit-tracking when deploying a segmented MVPN. This will lead to a low efficiency of explicit-tracking, and cause a worse multicast

join latency. Here we first take a scenario of inter-area segmented MVPN with both segments using BIER in [section 4](#) for a generic description. The second scenario is inter-area segmented MVPN with upstream segment using P2MP and downstream segment using BIER in [section 5](#). The third scenario is inter-area segmented MVPN with upstream segment using BIER and downstream segment using P2MP in [section 6](#).

[3.2.](#) Considerations

A BFIR always need to know the BFERs interested in a specific flow. This is a function of a BIER overlay defined in [\[RFC8279\]](#). A segmentation point BFR in a segmented MVPN deployment, saying ABR, will play similar roles of both BFIR and BFER. It needs to do a disposition of a BIER Header, and then do an imposition of a new BIER Header. It requires the ABR router to maintain per-flow states, and especially, such per-flow states always include a set of BFERs who are intrested in a specific flow by using an explicit-tracking procedure.

This behavior is completely different from a traditional segmented MVPN deployment, e.g, with both of the two segments using P2MP label switch. In a traditional segmented MVPN with both segments using P2MP label switch, it is expected to receive a MPLS packet and replicate to downstream routers after swap the MPLS Label. A lookup of IP packet is not expected. Also, in a traditional segmented MVPN

deployment, an MPLS label represents a P-tunnel, which may carry one, many or even all multicast flow(s) of a VPN, so it is not always a per-flow state on the segmentation point router.

In conclusion, the pattern of forwarding packets on segmentation points only by lookup of MPLS label mapped from multicast flow(s) is significantly unnecessary when BIER is introduced. Instead, doing a per-flow lookup of IP header on segmentation points is more efficient and consolidated.

[4.](#) Upstream BIER and Downstream BIER (BIER-BIER)

[4.1.](#) Explicit Tracking using LIR-pF

In a scenario of Inter-area Segmented MVPN with both segments using

BIER, the determination of the set of BFERs that need to receive a specific multicast flow of (C-S1,C-G1) in each segment, can be obtained by using a LIR-pF flag.

Suppose a topology of this:

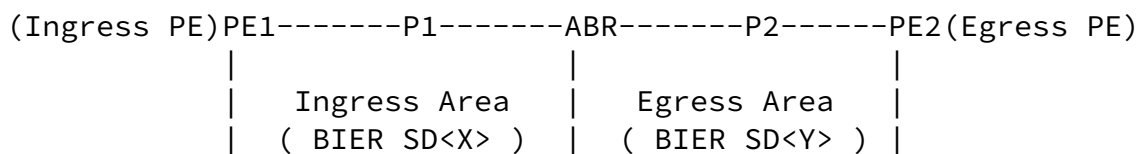


Figure 1: Example topology

PE1 is Ingress PE, and the area of { PE1 -- P1 -- ABR } is called an Ingress Area.

PE2 is Egress PE, and the area of { ABR -- P2 -- PE2 } is called an Egress Area.

The Ingress PE is configured to use a BIER tunnel type for an MVPN instance for the Ingress Area, and the ABR is configured to use a BIER tunnel type for the MVPN instance for the Egress Area.

The Ingress PE originates a wildcard S-PMSI A-D route (C-*,C-*) and the PTA of that route has the following settings:

- o The LIR-pF and LIR flags be set.
- o The tunnel type be set to "BIER".
- o A non-zero MPLS label be specified.

ABR receives the S-PMSI A-D route from the Ingress PE, and re-advertises the route to the Egress PE, with a PTA type "BIER", and PTA flags of LIR and LIR-pF, and a new non-zero upstream-assigned MPLS label allocated by ABR per-VPN.

Egress PE receives the S-PMSI A-D route from the ABR, and checks if it need to response with a Leaf A-D route to this S-PMSI A-D route using the process of the "match for reception" and "match for

tracking" as defined in [I-D.bess-mvpn-expl-track]. In this example, for a C-flow of (C-S1, C-G1), the checking result of "matched for tracking" is the S-PMSI(C-*, C-*), and the checking result of "matched for reception" is also the S-PMSI(C-*, C-*). Egress PE will then send a Leaf A-D route (RD, C-S1, C-G1, Root=PE1, Leaf=PE2) to the ABR with a PTA flag LIR-pF, and a Leaf A-D route (RD, C-*, C-*, Root=PE1, Leaf=PE2) without a PTA flag LIR-pF.

ABR then has an explicit-tracking result of a new per-flow information of (RD, C-S1, C-G1, Root=PE1) with Egress PE as its leaf or BFER. ABR's "matched for tracking" result to this flow(RD, C-S1, C-G1, PE1) will then be updated with a new record, and ABR then sends a Leaf A-D route (RD, C-S1, C-G1, Root=PE1, Leaf=ABR) to Ingress PE.

Ingress PE then has an explicit-tracking result of a new per-flow information of (RD, C-S1, C-G1, Root=PE1) with ABR as its leaf or BFER.

From this procedure description one can see that:

1. The S-PMSI A-D(C-*, C-*) route is functioning as a per-VPN anchor of the upstream and the downstream(s), which can be called an BGP-MVPN FEC in this document, saying BGP-MVPN FEC(OrigIP, RD).
2. The Leaf A-D(S,G) routes are functioning as a per-flow anchor of the downstream(s) and the upstream, which are also BGP-MVPN FECs accordingly, saying BGP-MVPN FEC(OrigIP, RD, S,G).
3. The Tuple of (Root=PE1, RD) in S-PMSI (C-*, C-*) or Leaf AD(C-*, C-*) or Leaf AD(C-S, C-G) represents an VRF on the ABR implicitly.

ABR knows the per-vpn information of a (Root=PE1, RD) tuple when receiving and re-advertising the S-PMSI A-D(*,*) route bound with a PTA, where:

- o Inbound SD (InSD): in PTA of the received S-PMSI(*,*) route.
- o Inbound VpnLabel (InVpnLabel): in PTA of the received S-PMSI(*,*) route.

- o Inbound BfirId (InBfirId): in PTA of the received S-PMSI(*,*)

route.

- o Outbound SD(OutSD): in PTA of the re-advertised S-PMSI(*,*) route.
- o Outbound VpnLabel (OutVpnLabel): in PTA of the re-advertised S-PMSI(*,*) route.
- o Outbound BfirId (OutBfirId): in PTA of the re-advertised S-PMSI(*,*) route.

ABR establishes a per-flow control-plane state accordingly like this:

- o Per-flow upstream state, according to the Leaf A-D (C-S, C-G) route send to the Ingress PE: (PE1, RD, C-S1, C-G1, InSD, InBfirId, InVpnLabel).
- o Per-flow downstream state(s), according to the Leaf A-D(C-S, C-G) route(s) received by the ABR from Egress PE(s): (PE1, RD, C-S1, C-G1, Leaf, OutSD, OutBfirId, OutVpnLabel).

ABR knows the BIER Label(s) it allocated for InSD and OutSD, saying InBierLabel for InSD<X> and OutBierLabel for OutSD<Y>, and thus it can establish the per-flow forwarding state:

- o Per-flow upstream forwarding state: (InBierLabel, InBfirId, InVpnLabel, C-S1, C-G1).
- o Per-flow downstream(s) forwarding state: (InBierLabel, InBfirId, InVpnLabel, C-S1, C-G1, Leaf, OutBfirId, OutVpnLabel, OutBitString)

[4.2.](#) Forwarding using IP lookup on Segmentation Point

The Forwarding procedure of a segmentation point, ABR, have the following steps:

1. Do a disposition of BIER tunnel to get the Pseudo-VRF, and a following IP lookup in the Pseudo-VRF to get the appropriate BIER encapsulation for Area<Y>, including the Local-VpnLabel indicated by the Pseudo-VRF.
2. Do a disposition of BIER tunnel to get the VRF, and a following IP lookup to get the appropriate local VRF interfaces if ABR has local VPN CEs.
3. The said "disposition of BIER tunnel" requires a lookup of upstream-assigned VpnLabel in special context.

For Step 1 and step 2, One can think them as one-to-many swapping of a big 'BIER header' instead of a small 'Label' in P2MP case:

- o swapping the InBierLabel with an OutBierLabel.
- o swapping the InBfirId with an OutBfirId.
- o swapping the InVpnLabel with an OutVpnLabel.
- o swapping the InBitString with an OutBitString.

The key of a per-flow lookup on ABR is a tuple of (InBierLabel, InBfirId, InVpnLabel) and a tuple of (C-S1, C-G1), representing a BIER tunnel and a flow respectively. All the elements are from a BIER packet, and such an IP lookup is very similar to an MFIB lookup, if the (InBierLabel, InBfirId, InVpnLabel) tuple is mapped to a Pseudo-VRF locally on the ABR per the upstream tunnel (a BIER tunnel in this example). The ABR can be configured as PE of an MVPN instance as well, and then there would be an MFIB lookup of the real VRF, and an MFIB lookup of the Pseudo-VRF, independently, without any impact to each other when configuration is changed. The MFIB lookup of the Pseudo-VRF is not required to do RPF check.

5. Upstream P2MP and Downstream BIER (P2MP-BIER)

5.1. Explicit Tracking using LIR-pF

The procedures described in chapter 4.1 is also suitable for this case, provided the LIR-pF explicit-tracking is used appropriately.

Suppose a topology of this:

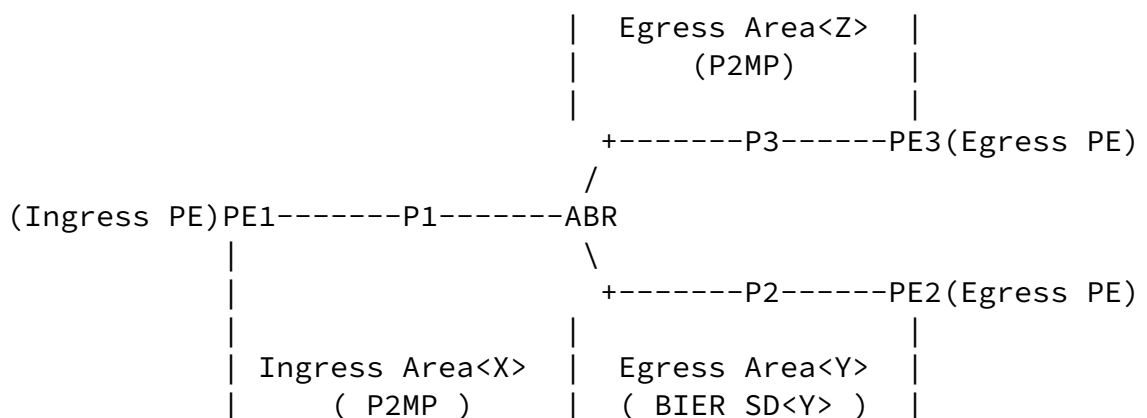


Figure 2: Example topology

The Ingress PE is configured to use an P2MP tunnel type for a MVPN instance for the Ingress Area<X>, and the ABR is configured to use a

BIER tunnel type for the MVPN instance for the Egress Area <Y>, and ABR may be configured to use a P2MP tunnel type for another Egress Area<Z>.

Example 1: Use Inclusive P-tunnel for traditional areas.

PE1 may configure to use one SPMSI (*,*,PTA<mLDP, Flag=LIR+LIRpF>) route , for one unidirectional Inclusive mLDP P2MP tunnel.

ABR may configure to reflect only the SPMSI(*,*) route with the PTA type changed to BIER for Area<Y>, and reflect the SPMSI(*,*) route with the PTA type changed to mLDP, RSVP-TE or IR for Area<Z>.

Example 2: Use Selective P-tunnel for traditional areas.

PE1 may configure to use one SPMSI (*,*,PTA<mLDP, Flag=LIR+LIRpF>) route and a number of SPMSI (S,G,PTA<mLDP, Flag=LIR+LIFpF>) routes, for one unidirectional Inclusive mLDP P2MP tunnel and a number of Selective mLDP P2MP tunnels respectively. The different P2MP tunnels bound in the SPMSI routes enable the ABR to initialize the different downstream P2MP tunnels for Area<Z> or BIER-P2MP tunnels for area <Y>, per the upstream P2MP tunnels.

ABR may configure to reflect only the SPMSI(*,*) route with the PTA type changed to BIER for Area<Y>, and reflect the SPMSI(*,*) route and SPMSI(S,G) routes with the PTA type changed to mLDP, RSVP-TE or IR for Area<Z>.

[5.2.](#) Forwarding using IP lookup on Segmentation Point

The Forwarding procedure of a segmentation point, ABR, have 3 conditions:

1. Do a disposition of P2MP tunnel to get the Pseudo-VRF, and a following IP lookup in Pseudo-VRF to get the appropriate BIER encapsulation for Area<Y>, including the Local-VpnLabel indicated by the Pseudo-VRF.
2. Do a lookup of P2MP tunnel to get the appropriate P2MP

downstreams for Area<Z>.

3. Do a disposition of P2MP tunnel to get the VRF, and a following IP lookup to get the appropriate local VRF interfaces if ABR has local VPN CEs.
4. The said "lookup/disposition of P2MP tunnel" requires a lookup of downstream-assigned P2MP Label lookup.

6. Upstream BIER and Downstream P2MP (BIER-P2MP)

6.1. Explicit Tracking using LIR-pF

The procedures described in chapter 4.1 is also suitable for this case, provided the LIR-pF explicit-tracking is used appropriately.

Suppose a topology of this:

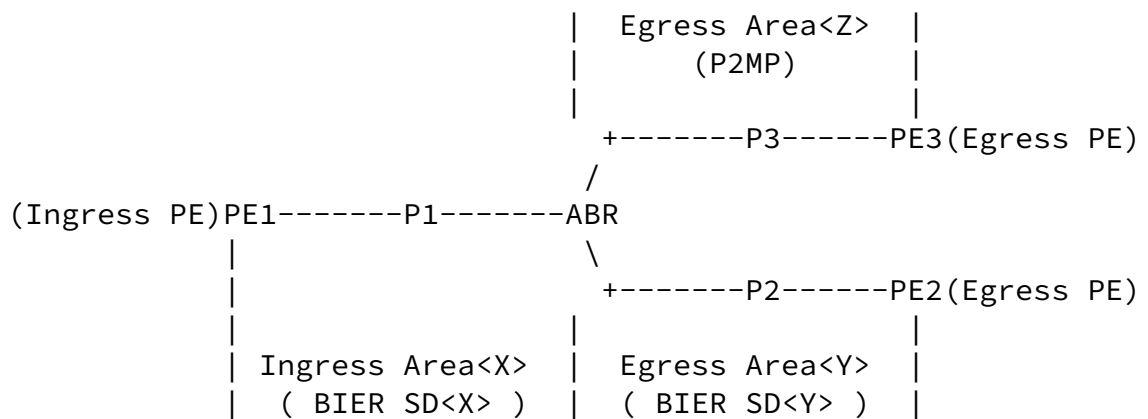


Figure 3: Example topology

The Ingress PE is configured to use a BIER tunnel type for an MVPN instance for the Ingress Area<X>, and the ABR is configured to use a BIER tunnel type for the MVPN instance for the Egress Area <Y>, and ABR may be configured to use a P2MP tunnel type for another Egress Area<Z>.

Example 1: Use Inclusive P-tunnel for traditional areas.

PE1 may configure to use one SPMSI (*,*,PTA<BIER, Flag=LIR+LIRpF>)

route , for one BIER tunnel.

ABR may configure to reflect only the SPMSI(*,*) route with the PTA type unchanged for Area<Y>, and reflect the SPMSI(*,*) route with the PTA type changed to mLDP, RSVP-TE or IR for Area<Z>.

Example 2: Use Selective P-tunnel for traditional areas.

PE1 may configure to use one SPMSI (*,*,PTA<BIER, Flag=LIR+LIRpF, VpnLabel1>) route and a number of SPMSI (S,G,PTA<BIER, Flag=LIR+LIFpF>, VpnLabelX) routes. The different VpnLabels represent different BIER tunnels, and thus enable the ABR to initialize the different downstream P2MP tunnels for Area<Z> or BIER-P2MP tunnels for area <Y>, per the upstream BIER tunnels.

ABR may configure to reflect only the SPMSI(*,*) route with the PTA type changed to BIER for Area<Y>, and reflect the SPMSI(*,*) route and SPMSI(S,G) routes with the PTA type changed to mLDP, RSVP-TE or IR for Area<Z>.

[6.2.](#) Forwarding on Segmentation Point

The Forwarding procedure of a segmentation point, ABR, have the following conditions:

1. Do a disposition of BIER tunnel to get the Pseudo-VRF, and a following IP lookup in the Pseudo-VRF to get the appropriate BIER encapsulation for Area<Y>, including the Local-VpnLabel indicated by the Pseudo-VRF.
2. Do a disposition of BIER tunnel to get the appropriate P2MP downstreams for Area<Z>.
3. Do a disposition of BIER tunnel to get the VRF, and a following IP lookup to get the appropriate local VRF interfaces if ABR has local VPN CEs.
4. The said "disposition of BIER tunnel" requires a lookup of upstream-assigned VpnLabel in special context.

7. Summary and Recommendations

Summary:

1. When BIER is used as the tunnel of the upstream segment, then a upstream-assigned VpnLabel lookup in a special context is required, instead of a downstream-assigned label lookup.
2. When BIER is used as the tunnel of the downstream segment, then a per-flow IP lookup is required to get the BIER encapsulation, following the label lookup to map 1:1 from the upstream tunnel (BIER tunnel or P2MP tunnel) to the downstream BIER tunnel.
3. It is possible to require a per-flow VpnLabel allocation for the whole Segmented MVPN, then a per-flow IP lookup can be the same result as the per-tunnel VpnLabel lookup, and then one can optimize not to use IP lookup. But this may not be allowed when BIER-P2MP or P2MP-BIER deployed and the P2MP part can not support per-flow tunnels. One obvious example is when the P2MP segment uses Inclusive P2MP tunnel for all or for part of the multicast flows.

4. An IP lookup following a VpnLabel lookup in special context is a mandatory capability for a BFER function, and an IP lookup following a P2MP label lookup is a mandatory capability for a P2MP BUD function. Use of IP lookup on ABR for downstream BIER segment can leverage the already required IP lookup capability.
5. Use of LIR-pF explicit-tracking in segmented MVPN deployment can make the same benefits as in non segmented MVPN deployment. Both can make the tunnel allocation (P2MP Label or VpnLabel allocation) and the per-flow states decoupled.

Recommendations are:

1. that implementations support the IP lookup for Segmented MVPN when downstream segment using BIER.
2. that implementations support the LIR-pF explicit tracking for Segmented MVPN when BIER being deployed in any one segment.

3. that implementations support the optimization of not using IP lookup on segmentation point ABR when end to end distinct tunnels (P2MP tunnels or BIER tunnels) for distinct C-flows is not a limit.

8. [Appendix A](#) - Comparison of Solutions

The table below provides a side-by-side comparison of explicit tracking recommended, for non segmented MVPN and segmented MVPN cases.

<Expl-tracking>	draft-ietf-bier-mvpn	this draft
Non segmented MVPN	LIR-pF Per-flow Label NotReq	LIR-pF Per-flow Label NotReq
Segmented MVPN	LIR Per-flow Label Required	LIR-pF Per-flow Label NotReq
(BIER-BIER)	LIR	LIR-pF
(BIER-P2MP)	Not specified	LIR-pF
(BIER-BIER)	Not specified	LIR-pF

Figure 4: Comparison of Segmented and non segmented MVPN

The table below provides a side-by-side comparison of forwarding procedure on BFER/segmentation point ABR about whether to use IP Lookup, for BFER, BIER-BIER, P2MP-BIER and BIER-P2MP cases.

<Forwarding>	draft-ietf-bier-mvpn	this draft
BFER	VpnLabel Lookup + IP Lookup	VpnLabel Lookup + IP Lookup
ABR(BIER-BIER)	VpnLabel lookup	VpnLabel lookup

		+ IP Lookup
ABR(P2MP-BIER)	P2MP-Label Lookup	P2MP-Label Lookup + IP Lookup
ABR(BIER-P2MP)	VpnLabel lookup	VpnLabel Lookup
ABR(P2MP-P2MP)	P2MP-Label lookup	P2MP-Label Lookup

Figure 5: Comparison of Various Forwarding cases

The table below provides a list of the forwarding on ABR in a 3 segments deployment cases.

Topology Deployed	Forwarding
PE1--(Tnl1)--ABR1--(Tnl2)--ABR2--(Tnl3)--PE2	on ABR1

BIER	BIER	ANY	*(1)
BIER	IR	ANY	*(2)
BIER	MLDP	ANY	*(2)
BIER	RSVP-TE	ANY	*(2)
+-----+			
IR	BIER	ANY	*(3)
IR	IR	ANY	P2MP Swap
IR	MLDP	ANY	P2MP Swap
IR	RSVP-TE	ANY	P2MP Swap
+-----+			
MLDP	BIER	ANY	*(3)
MLDP	IR	ANY	P2MP Swap
MLDP	MLDP	ANY	P2MP Swap
MLDP	RSVP-TE	ANY	P2MP Swap
+-----+			
RSVP-TE	BIER	ANY	*(3)
RSVP-TE	IR	ANY	P2MP Swap
RSVP-TE	MLDP	ANY	P2MP Swap
RSVP-TE	RSVP-TE	ANY	P2MP Swap
+-----+			

*(1) VpnLabel Lookup for PseudoVRF, and IP Lookup for per-flow encapsulation.

*(2) VpnLabel Lookup for per-tunnel P2MP downstreams.

*(3) P2MP-Label Lookup for PseudoVRF, and IP Lookup for per-flow encapsulation.

*(4) P2MP-Label Lookup for per-tunnel P2MP downstreams.

Figure 6: Comparison of Various segmented combinations

9. Security Considerations

The procedures of this document do not, in themselves, provide privacy, integrity, or authentication for the control plane or the data plane.

10. IANA Considerations

No IANA allocation is required.

11. Acknowledgements

TBD.

[12.](#) References

[12.1.](#) Normative References

- [I-D.ietf-bess-mvpn-expl-track]
Dolganow, A., Kotalwar, J., Rosen, E., and Z. Zhang,
"Explicit Tracking with Wild Card Routes in Multicast
VPN", [draft-ietf-bess-mvpn-expl-track-12](#) (work in
progress), October 2018.
- [I-D.ietf-bier-mvpn]
Rosen, E., Sivakumar, M., Aldrin, S., Dolganow, A., and T.
Przygienda, "Multicast VPN Using BIER", [draft-ietf-bier-
mvpn-11](#) (work in progress), March 2018.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/
BGP IP VPNs", [RFC 6513](#), DOI 10.17487/RFC6513, February
2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP
Encodings and Procedures for Multicast in MPLS/BGP IP
VPNs", [RFC 6514](#), DOI 10.17487/RFC6514, February 2012,
<<https://www.rfc-editor.org/info/rfc6514>>.
- [RFC6625] Rosen, E., Ed., Rekhter, Y., Ed., Hendrickx, W., and R.
Qiu, "Wildcards in Multicast VPN Auto-Discovery Routes",
[RFC 6625](#), DOI 10.17487/RFC6625, May 2012,
<<https://www.rfc-editor.org/info/rfc6625>>.
- [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>>.

[12.2.](#) Informative 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>>.

Authors' Addresses

Jingrong Xie
Huawei Technologies

Email: xiejingrong@huawei.com

Xie, et al.

Expires April 25, 2019

[Page 16]

Internet-Draft

Segmented MVPN Using IP Lookup for BIER

October 2018

Liang Geng
China Mobile
Beijing 10053

Email: gengliang@chinamobile.com

Lei Wang
China Mobile
Beijing 10053

Email: wangleiyjy@chinamobile.com

Mike McBride
Huawei

Email: mmcbride7@gmail.com

Gang Yan
Huawei

Email: yangang@huawei.com

Xie, et al.

Expires April 25, 2019

[Page 17]