

Internet Engineering Task Force (IETF)
Request for Comments: 8556
Category: Standards Track
ISSN: 2070-1721

E. Rosen, Ed.
M. Sivakumar
T. Przygienda
Juniper Networks, Inc.
S. Aldrin
Google, Inc.
A. Dolganow
Nokia
April 2018

Multicast VPN Using Bit Index Explicit Replication (BIER)

Abstract

The Multicast Virtual Private Network (MVPN) specifications require the use of multicast tunnels ("P-tunnels") that traverse a service provider's backbone network. The P-tunnels are used for carrying multicast traffic across the backbone. A variety of P-tunnel types are supported. Bit Index Explicit Replication (BIER) is a new architecture that provides optimal multicast forwarding through a "multicast domain", without requiring intermediate routers to maintain any per-flow state or to engage in an explicit tree-building protocol. This document specifies the protocol and procedures that allow MVPN to use BIER as the method of carrying multicast traffic over a service provider's backbone network.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8556>.

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 3
- 2. Use of the PMSI Tunnel Attribute in x-PMSI A-D Routes 5
 - 2.1. MPLS Label 7
 - 2.2. Explicit Tracking 9
 - 2.2.1. Using the LIR Flag 10
 - 2.2.2. Using the LIR-pF Flag 10
- 3. Use of the PMSI Tunnel Attribute in Leaf A-D Routes 11
- 4. Data Plane 12
 - 4.1. Encapsulation and Transmission 12
 - 4.2. Disposition 14
 - 4.2.1. At a BFER That Is an Egress PE 14
 - 4.2.2. At a BFER That Is a P-tunnel Segmentation Boundary . 14
- 5. IANA Considerations 15
- 6. Security Considerations 15
- 7. References 15
 - 7.1. Normative References 15
 - 7.2. Informative References 16
- Acknowledgments 16
- Contributors 16
- Authors' Addresses 17

1. Introduction

[RFC6513] and [RFC6514] specify the protocols and procedures that a Service Provider (SP) can use to provide Multicast Virtual Private Network (MVPN) service to its customers. Multicast tunnels are created through an SP's backbone network; these are known as "P-tunnels". The P-tunnels are used for carrying multicast traffic across the backbone. The MVPN specifications allow the use of several different kinds of P-tunnel technology.

Bit Index Explicit Replication (BIER) ([RFC8279]) is an architecture that provides optimal multicast forwarding through a "multicast domain", without requiring intermediate routers to maintain any per-flow state or to engage in an explicit tree-building protocol. The purpose of the current document is to specify the protocols and procedures needed in order to provide MVPN service using BIER to transport the multicast traffic over the backbone.

Although BIER does not explicitly build and maintain multicast tunnels, one can think of BIER as using a number of implicitly created tunnels through a "BIER domain". In particular, one can think of there as being one Point-to-Multipoint (P2MP) tunnel from each Bit Forwarding Ingress Router (BFIR) to all the Bit Forwarding Egress Routers (BFERs) in the BIER domain, where a BIER domain is generally co-extensive with an IGP network. These "tunnels" are not specific to any particular VPN. However, the MVPN architecture provides protocols and procedures that allow the traffic of multiple MVPNs to be aggregated on a single P-tunnel. In this document, we specify how to use these multi-VPN aggregation procedures to enable BIER to transport traffic from multiple MVPNs.

MVPN traffic must sometimes traverse more than one IGP domain, whereas BIER only carries multicast traffic within a single IGP domain. However, the MVPN specifications allow P-tunnels to be segmented (the concept of MVPN segmentation is defined in [RFC6513] and [RFC6514]), where the segmentation points may either be Autonomous System Border Routers (ASBRs) as described in [RFC6514], or Area Border Routers (ABRs) as described in [RFC7524]. As long as the segmentation points are capable of acting as BFIRs and BFERs, BIER can be used to provide some or all of the segments of a P-tunnel.

Procedures to support MVPN customers who are using Bidirectional PIM (BIDIR-PIM) are outside the scope of this document.

This document uses the following terminology from [RFC8279]:

- o BFR: Bit-Forwarding Router.
- o BFIR: Bit-Forwarding Ingress Router.
- o BFER: Bit-Forwarding Egress Router.

This document uses the following terminology from [RFC6513]:

- o MVPN: Multicast Virtual Private Network -- a VPN [RFC4364] in which multicast service is offered.
- o P-tunnel: A multicast tunnel through the network of one or more SPs. P-tunnels are used to transport MVPN multicast data
- o PMSI: Provider Multicast Service Interface. PMSI is an abstraction that represents a multicast service for carrying packets. A PMSI is instantiated via one or more P-tunnels.
- o C-S: A multicast source address, identifying a multicast source located at a VPN customer site.
- o C-G: A multicast group address used by a VPN customer.
- o C-flow: A customer multicast flow. Each C-flow is identified by the ordered pair (source address, group address), where each address is in the customer's address space. The identifier of a particular C-flow is usually written as (C-S,C-G).

Sets of C-flows can be identified by the use of the "C-*" wildcard (see [RFC6625]), e.g., (C-*,C-G).

- o I-PMSI A-D Route: Inclusive PMSI Auto-Discovery route. Carried in BGP Update messages, these routes are used to advertise the default P-tunnel for a particular MVPN.
- o S-PMSI A-D route: Selective PMSI Auto-Discovery route. Carried in BGP Update messages, these routes are used to advertise the fact that particular C-flows are bound to (i.e., are traveling through) particular P-tunnels.
- o x-PMSI A-D route: A route that is either an I-PMSI A-D route or an S-PMSI A-D route.
- o Leaf A-D route: A route that a multicast egress node sends in order to join a particular P-tunnel.

- o PMSI Tunnel attribute (PTA): In an x-PMSI A-D route, the Network Layer Reachability Information (NLRI) of the route identifies a PMSI. The BGP attribute known as the PMSI Tunnel attribute is attached to such a route in order to identify a particular P-tunnel that is associated with the PMSI. When C-flows of multiple VPNs are carried in a single P-tunnel, this attribute also carries the information needed to multiplex and demultiplex the C-flows. A PTA can also be carried by a Leaf A-D root. In this case, it contains information that is needed in order for the originator of the route to join the specified P-tunnel.

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.

2. Use of the PMSI Tunnel Attribute in x-PMSI A-D Routes

As defined in [RFC6514], the PMSI Tunnel attribute (PTA) carried by an x-PMSI A-D route identifies the P-tunnel that is used to instantiate a particular PMSI. If a PMSI is to be instantiated by BIER, the PTA is constructed by a BFIR.

If segmented P-tunnels are not being used, the PTA attached to a given x-PMSI A-D route is constructed by the router that originated the route (typically by the ingress Provider Edge (PE) router), and the PTA is not changed as the route is propagated.

If segmented P-tunnels are being used, the PTA attached to a given x-PMSI A-D route by the route's originator may be replaced at a segmentation point (a BFER) by a PTA identifying the next segment of the P-tunnel. If the next segment of the P-tunnel is instantiated by BIER, the segmentation point serves as the BFIR for that next segment.

In either case, a PTA is constructed by a BFIR as follows (see Figure 1):

The PTA contains the following fields:

- o Tunnel Type: IANA has assigned 0x0B as the tunnel type codepoint for "BIER" in the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry. This codepoint is used to indicate that the PMSI is instantiated by BIER.

Although BIER does not actually create tunnels, the MVPN procedures treat BIER as if it were a type of tunnel.

- o Tunnel Identifier: When the tunnel type is BIER, this field contains three subfields:
 1. The first subfield is a single octet, containing a BIER sub-domain-id (see [RFC8279]). This indicates that packets sent on the PMSI will be sent on the specified BIER sub-domain. How that sub-domain is chosen is outside the scope of this document.
 2. The second subfield is a two-octet field containing the BFR-id in the sub-domain identified in the first subfield of the router that is constructing the PTA.
 3. The third subfield is the BFR-prefix (see [RFC8279]) of the router (a BFIR) that is constructing the PTA. The BFR-prefix will either be a /32 IPv4 address or a /128 IPv6 address. Whether the address is IPv4 or IPv6 can be inferred from the total length of the PTA.

The BFR-prefix need not be the same IP address that is carried in any other field of the x-PMSI A-D route, even if the BFIR is the originating router of the x-PMSI A-D route.

Failure to properly set the Tunnel Identifier field cannot be detected by the protocol and will result in improper delivery of the data packets sent on the PMSI.

- o MPLS Label: This field MUST contain an upstream-assigned non-zero MPLS label. It is assigned by the router (a BFIR) that constructs the PTA. Constraints on the way in which a BFIR selects this label are discussed in Section 2.1.

Failure to follow the constraints on label assignment cannot be detected by the protocol and may result in improper handling of data packets by the egress PE routers.

- o Flags: When the tunnel type is BIER, two of the flags in the PTA Flags field are meaningful. Details about the use of these flags can be found in Section 2.2.
- * Leaf Information Required per Flow (LIR-pF): This flag is introduced in [RFC8534]. A BFIR SHOULD NOT set this flag UNLESS it knows that all the BFERs in the BIER domain (or at least all the BFERs to which it needs to transmit) support this flag. (How this is known is outside the scope of this document.) Procedures for the use of this flag are given in Section 2.2.2. Support for this flag is OPTIONAL.

* Leaf Information Required (LIR): see Section 2.2.1.

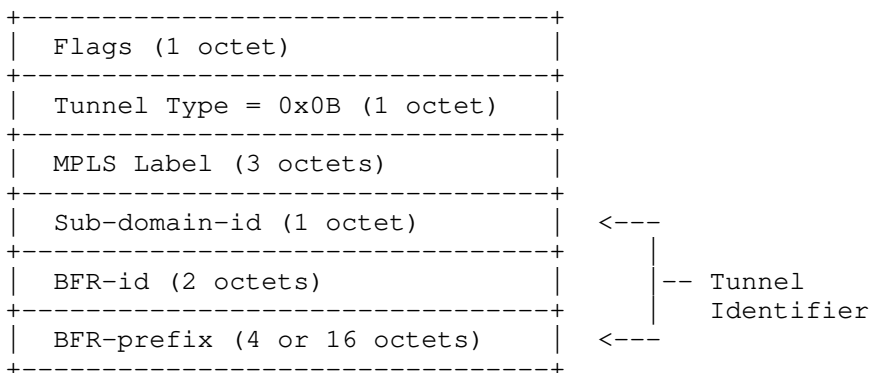


Figure 1: PMSI Tunnel Attribute for BIER

If a PTA specifying tunnel type BIER is attached to an x-PMSI A-D route, the route MUST NOT be distributed beyond the boundaries of a BIER domain. That is, any routers that receive the route must be in the same BIER domain as the originator of the route. If the originator is in more than one BIER domain, the route must be distributed only within the BIER domain in which the BFR-prefix in the PTA uniquely identifies the originator. As with all MVPN routes, distribution of these routes is controlled by the provisioning of Route Targets (RTs). Thus, the requirement expressed in this paragraph is really a requirement on the way the Route Targets are provisioned.

2.1. MPLS Label

The MPLS Label carried in the PTA is an upstream-assigned label.

If two PTAs contain the same BFR-prefix in their respective Tunnel Identifier fields, then the labels carried in those PTAs MUST come from the same label space (see Section 7 of [RFC5331]). An implementation may choose to use this fact when setting up the tables it uses to interpret the upstream-assigned labels.

Suppose that a BFIR attaches a PTA to each of two x-PMSI A-D routes, and both PTAs specify a tunnel type of BIER.

- o If the two routes do not carry the same set of RTs, then their respective PTAs MUST contain different MPLS label values.

- o If the two routes do not have the same Address Family Identifier (AFI) value, then their respective PTAs MUST contain different MPLS label values. This ensures that when an egress PE receives a data packet with the given label, the egress PE can infer from the label whether the payload is an IPv4 packet or an IPv6 packet.
- o If the BFIR is an ingress PE supporting MVPN extranet ([RFC7900]) functionality, and if the two routes originate from different VPN Routing and Forwarding tables (VRFs) on this ingress PE, then the respective PTAs of the two routes MUST contain different MPLS label values.
- o If the BFIR is an ingress PE supporting the "Extranet Separation" feature of MVPN extranet (see Section 7.3 of [RFC7900]), and if one of the routes carries the "Extranet Separation" extended community but the other does not, then the respective PTAs of the two routes MUST contain different MPLS label values.
- o If segmented P-tunnels are being used, then the respective PTAs of the two routes MUST contain different MPLS label values whenever the respective NLRIs of the two routes are not identical. The MPLS label can then be used at the next segmentation point to switch packets from one P-tunnel segment directly to the next, without requiring the segmentation points to contain any other multicast forwarding state. This is explained further below; see also Section 4.

When segmented P-tunnels are being used, a segmentation point, call it "B1", may receive an x-PMSI A-D route whose PTA specifies BIER from within a given BIER domain. This means that BIER is being used for the previous segment of a segmented P-tunnel. If the next segment is also of type BIER, B1 will be the BFIR for the next segment. That is, B1 is a BFIR of one BIER domain (corresponding to the previous segment) and a BFIR of another BIER domain (corresponding to the next segment). B1 needs to replace the PTA of the x-PMSI A-D route with a new PTA, specifying its own BFR-prefix and specifying an upstream-assigned label assigned by B1 itself.

Suppose that B1 has received two x-PMSI A-D routes, R1 and R2, where:

- o R1 and R2 each have a PTA specifying BIER.
- o R1's PTA specifies BFR-prefix B2 and Label L2.
- o R2's PTA specifies BFR-prefix B3 and Label L3.

Suppose B1 decides to propagate both R1 and R2, replacing each PTA with a new PTA specifying BIER. Suppose these new PTAs specify labels L4 and L5, respectively. Then L4 and L5 MUST be different (upstream-assigned) label values, UNLESS both of the following conditions hold:

- o R1 and R2 have the same value in the Originating Router field of their respective NLRIs, and
- o B2 is equal to B3, and
- o L2 is equal to L3.

The segmentation point (B1, in this example) MUST also program its data plane appropriately. For example, when:

- o B1 receives a BIER packet for which it is a BFER, and
- o the BIER header specifies the BFIR-id that corresponds to B2, and
- o the BIER payload is an MPLS packet with upstream-assigned label, and
- o the top label value is L2,

then the data plane must be programmed to replace L2 with L4 and to re-encapsulate the packet in a BIER header with B1's BFR-id in the BFIR-id field. The BitString of the new BIER header is determined by applying the procedures for MVPN explicit tracking (see Section 2.2) in the BIER domain of the next segment, i.e., in the BIER domain for which B1 is the BFIR).

2.2. Explicit Tracking

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 in either of two ways:

1. Using the explicit tracking mechanism based on the "Leaf Information Required" flag specified in [RFC6513] and [RFC6514]. This method is further described in Section 2.2.1.
2. Using the OPTIONAL explicit tracking mechanism based on the LIR-pF flag specified in [RFC8534]. This method, further described in Section 2.2.2, may be used if (and only if) segmented P-tunnels are not being used.

2.2.1. Using the LIR Flag

To determine the set of BFERs to which the packets of a given C-flow must be sent, a BFIR MUST originate a (C-S,C-G) S-PMSI A-D route for the given C-flow. It MUST attach a PTA to that route and MUST set the Leaf Information Required (LIR) flag in the PTA. Per [RFC6514], the BFERs that need to receive that C-flow will respond with (C-S,C-G) Leaf A-D routes. By matching the received Leaf A-D routes to the originated S-PMSI A-D routes, the originator of the S-PMSI A-D route determines the set of BFERs that need to receive the multicast data flow that is identified in the NLRI of S-PMSI A-D route.

Suppose that an ingress PE has originated an I-PMSI A-D route or a wildcard S-PMSI A-D route [RFC6625] with a PTA specifying a tunnel type of BIER. Now suppose that the ingress PE originates an S-PMSI A-D route specifying (C-S,C-G), where (C-S,C-G) "matches" (according to the rules of [RFC6625]) the wildcard S-PMSI A-D route or the I-PMSI A-D route. Instead of attaching a PTA specifying BIER to the (C-S,C-G) route, the ingress PE MAY attach a PTA specifying a tunnel type of "no tunnel information". This is equivalent to attaching the same PTA attached to the matching "less specific" route.

2.2.2. Using the LIR-pF Flag

If segmented P-tunnels are not being used, the BFIR can determine the set of BFERs that need to receive the packets of a given (C-S,C-G) C-flow as follows. The BFIR MUST originate a wildcard S-PMSI A-D route (either (C-*,C-*), (C-*,C-G), or (C-S,C-G)), and the PTA of that route MUST use the following settings:

- o The LIR-pF flag MUST be set.
- o The tunnel type MUST be set to BIER.
- o A non-zero MPLS label MUST be specified.

Per [RFC8534], a BFER that needs to receive (C-S,C-G) traffic from the BFIR will respond with a Leaf A-D route.

A BFIR MUST NOT use this method of finding the set of BFERs needing to receive a given C-flow unless it knows that all those BFERs support the LIR-pF flag. How this is known is outside the scope of this document.

This method greatly reduces the number of S-PMSI A-D routes that a BFIR needs to originate; it can now originate as few as one such route (a (C-*,C-*) S-PMSI A-D route), rather than one for each C-flow. However, the method does not provide a way for the BFIR to assign a distinct label to each C-flow. Therefore, it cannot be used when segmented P-tunnels are in use (see Section 4 for an explanation).

Note: If a BFIR originates a (C-*,C-*) S-PMSI A-D route with the LIR-pF flag set but also originates a more specific wildcard route that matches a particular (C-S,C-G), the BFERs will not originate Leaf A-D routes for that (C-S,C-G) unless the LIR-pF flag is also set in the more specific wildcard route. If the BFIR also originates a (C-S,C-G) S-PMSI A-D route without the LIR flag set, the BFERs will not originate Leaf A-D routes for that (C-S,C-G) unless the LIR flag is also set in that route.

3. Use of the PMSI Tunnel Attribute in Leaf A-D Routes

Before an egress PE can receive a (C-S,C-G) flow from a given ingress PE via BIER, the egress PE must have received one of the following x-PMSI A-D routes from the ingress PE:

- o A (C-S,C-G) S-PMSI A-D route (i.e., an S-PMSI A-D route whose NLRI encodes (C-S,C-G)) and whose PTA specifies a tunnel type of BIER. If such a route is found, we refer to it as the "matching x-PMSI A-D route."
- o A "less specific" x-PMSI A-D route (one specifying (C-*,C-*), (C-*,C-G), or (C-S,C-G)) whose PTA specifies a tunnel type of BIER, and that is the egress PE's "match for reception" of (C-S,C-G).

The rules for determining which x-PMSI A-D route is the match for reception are given in [RFC6625]. However, these rules are modified here to exclude any x-PMSI A-D route that does not have a PTA or whose PTA specifies "no tunnel type".

If such a route is found, we refer to it as the "matching x-PMSI A-D route."

If no matching x-PMSI A-D route for (C-S,C-G) is found, the egress PE cannot receive the (C-S,C-G) flow from the ingress PE via BIER until such time as a matching route is received.

When an egress PE determines that it needs to receive a (C-S,C-G) flow from a particular ingress PE via BIER, it originates a Leaf A-D route. Construction of the Leaf A-D route generally follows the procedures specified in [RFC6514] or, optionally, the procedures specified in [RFC8534]. However, when BIER is being used, the Leaf A-D route MUST carry a PTA that is constructed as follows:

1. The tunnel type MUST be set to BIER.
2. The MPLS Label field SHOULD be set to zero.
3. The sub-domain-id subfield of the Tunnel Identifier field (as defined in Section 2) MUST be set to the corresponding value from the PTA of the matching x-PMSI A-D route.
4. The BFR-id subfield of the Tunnel Identifier field MUST be set to the BFR-id in the sub-domain identified by the sub-domain-id subfield of the egress PE (BFER).
5. The BFR-prefix field of the Tunnel Identifier field (as defined in Section 2) MUST be set to the egress PE's (BFER's) BFR-prefix.

The BFR-prefix need not be the same IP address that is carried in any other field of the Leaf A-D route.

When an ingress PE receives such a Leaf A-D route, it learns the BFR-prefix of the egress PE from the PTA. The ingress PE does not make any use the value of the PTA's MPLS label field.

Failure to properly construct the PTA cannot always be detected by the protocol and will cause improper delivery of the data packets.

4. Data Plane

The MVPN application plays the role of the "multicast flow overlay" as described in [RFC8279].

4.1. Encapsulation and Transmission

To transmit an MVPN data packet, an ingress PE follows the rules of [RFC6625] to find the x-PMSI A-D route that is a "match for transmission" for that packet. (In applying the rules of [RFC6625], any S-PMSI A-D route with a PTA specifying "no tunnel information" is ignored.) If the matching route has a PTA specifying BIER, the (upstream-assigned) MPLS label from that PTA is pushed on the packet's label stack. Then the packet is encapsulated in a BIER

header. That is, the ingress PE functions as a BFIR. The BIER sub-domain used for transmitting the packet is specified in the PTA of the above-mentioned x-PMSI A-D route.

In order to create the proper BIER header for a given packet, the BFIR must know all the BFERs that need to receive that packet. It determines this by finding all the Leaf A-D routes that correspond to the S-PMSI A-D route that is the packet's match for transmission. There are two different cases to consider:

1. The S-PMSI A-D route that is the match for transmission carries a PTA that has the LIR flag set but does not have the LIR-pF flag set.

In this case, the corresponding Leaf A-D routes are those whose "route key" field is identical to the NLRI of the S-PMSI A-D route.

2. The S-PMSI A-D route that is the match for transmission carries a PTA that has the LIR-pF flag.

In this case, the corresponding Leaf A-D routes are those whose "route key" field is derived from the NLRI of the S-PMSI A-D route according to the procedures described in Section 5.2 of [RFC8534].

The Leaf A-D route from a given BFER will contain a PTA that specifies the BFER's BFR-prefix. With this information, the BFIR can construct the BIER BitString.

However, if the PTA of the Leaf A-D route from a given BFER specifies a sub-domain other than the one being used for transmitting the packet, the bit for that BFER cannot be determined and that BFER will not receive the packet.

The BIER-encapsulated packet is then forwarded, according to the procedures described in [RFC8279] and [RFC8296]. (See especially Section 3, "Imposing and Processing the BIER Encapsulation" in [RFC8296].)

4.2. Disposition

When a BFER receives an MVPN multicast data packet that has been BIER-encapsulated, the BIER layer passes the following information to the multicast flow overlay:

- o The sub-domain-id and the BFIR-id from the BIER header. (As the sub-domain-id is inferred from the BIFT-id field of the BIER header, an implementation might choose to pass the BIFT-id rather than the sub-domain-id; this is an implementation matter.)
- o The "payload", which is an MPLS packet whose top label is an upstream-assigned label. In the data plane, the BFIR-id and the sub-domain-id provide the context in which the upstream-assigned label is interpreted.

By looking up the upstream-assigned label in the appropriate context, the multicast flow overlay determines whether the BFER is an egress PE for the packet.

Note that if segmented P-tunnels are in use, a BFER might be a P-tunnel segmentation border router rather than an egress PE, or a BFER might be both an egress PE and a P-tunnel segmentation border router. Depending upon the role of the BFER for the given packet, it may need to follow the procedures of Section 4.2.1, the procedures of Section 4.2.2, or both.

4.2.1. At a BFER That Is an Egress PE

From looking up the packet's upstream-assigned label in the context of the packet's BFIR-prefix, the egress PE determines the egress VRF for the packet. From the IP header of the payload, the multicast states of the VRF, the upstream-assigned label, and the BFR-prefix, the egress PE can determine whether the packet needs to be forwarded out one or more VRF interfaces.

4.2.2. At a BFER That Is a P-tunnel Segmentation Boundary

When segmented P-tunnels are being used, a BFER that receives a BIER-encapsulated MVPN multicast data packet may need to be forwarded on its next P-tunnel segment. The choice of the next P-tunnel segment for the packet depends upon the C-flow to which the packet belongs. As long as the BFIR has assigned the MPLS label according to the constraints specified in Section 2.1, the BFIR will have assigned distinct upstream-assigned MPLS labels to distinct C-flows. The BFER can thus select the proper "next P-tunnel segment" for a given packet simply by looking up the upstream-assigned label that immediately follows the BIER header.

5. IANA Considerations

IANA has assigned the codepoint 0x0B to BIER in the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry.

6. Security Considerations

The procedures of this document do not, in themselves, provide privacy, integrity, or authentication for the control plane or the data plane. For a discussion of the security considerations regarding the use of BIER, please see [RFC8279] and [RFC8296]. Security considerations regarding VPN technology based on [RFC4364], [RFC6513], and [RFC6514] can be found in those RFCs.

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>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, DOI 10.17487/RFC5331, August 2008, <<https://www.rfc-editor.org/info/rfc5331>>.
- [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>>.

- [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>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [RFC8534] Dolganow, A., Kotalwar, J., Rosen, E., Ed., and Z. Zhang, "Explicit Tracking with Wildcard Routes in Multicast VPN", RFC 8534, DOI 10.17487/RFC8534, February 2019, <<https://www.rfc-editor.org/info/rfc8534>>.

7.2. Informative References

- [RFC7524] Rekhter, Y., Rosen, E., Aggarwal, R., Morin, T., Grosclaude, I., Leymann, N., and S. Saad, "Inter-Area Point-to-Multipoint (P2MP) Segmented Label Switched Paths (LSPs)", RFC 7524, DOI 10.17487/RFC7524, May 2015, <<https://www.rfc-editor.org/info/rfc7524>>.
- [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>>.

Acknowledgments

The authors wish to thank Jeffrey Zhang for his ideas and contributions to this work. We also thank Stig Venaas for his review and comments.

Contributors

IJsbrand Wijnands
Cisco Systems, Inc.
De Kleetlaan 6a
Diegem 1831
Belgium
Email: ice@cisco.com

Authors' Addresses

Eric C. Rosen (editor)
Juniper Networks, Inc.
10 Technology Park Drive
Westford, Massachusetts 01886
United States of America

Email: erosen52@gmail.com

Mahesh Sivakumar
Juniper Networks, Inc.
1137 Innovation Way
Sunnyvale, California 94089
United States of America

Email: sivakumar.mahesh@gmail.com

Tony Przygienda
Juniper Networks, Inc.
1137 Innovation Way
Sunnyvale, California 94089
United States of America

Email: prz@juniper.net

Sam K Aldrin
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, California
United States of America

Email: aldrin.ietf@gmail.com

Andrew Dolganow
Nokia
438B Alexandra Rd #08-07/10
Alexandra Technopark
Singapore 119968

Email: andrew.dolganow@nokia.com