

BESS WG  
Internet-Draft  
Updates: [6625](#) (if approved)  
Intended status: Standards Track  
Expires: July 20, 2018

A. Dolganow  
J. Kotalwar  
Nokia  
E. Rosen, Ed.  
Z. Zhang  
Juniper Networks, Inc.  
January 16, 2018

**Explicit Tracking with Wild Card Routes in Multicast VPN**  
**draft-ietf-bess-mvpn-expl-track-04**

Abstract

The MVPN specifications provide procedures to allow a multicast ingress node to invoke "explicit tracking" for a multicast flow or set of flows, thus learning the egress nodes for that flow or set of flows. However, the specifications are not completely clear about how the explicit tracking procedures work in certain scenarios. This document provides the necessary clarifications. It also specifies a new, optimized explicit tracking procedure. This new procedure allows an ingress node, by sending a single message, to request explicit tracking of each of a set of flows, where the set of flows is specified using a wildcard mechanism. This document updates [RFC6625](#).

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 July 20, 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	The Explicit Tracking Flags . . . . .	<a href="#">5</a>
<a href="#">3.</a>	Match for Tracking vs. Match for Reception . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Ingress Node Initiation of Tracking . . . . .	<a href="#">7</a>
<a href="#">5.</a>	Egress Node Response to the Match for Tracking . . . . .	<a href="#">9</a>
<a href="#">5.1.</a>	General Egress Node Procedures . . . . .	<a href="#">9</a>
<a href="#">5.2.</a>	Responding to the LIR-pF Flag . . . . .	<a href="#">10</a>
<a href="#">5.3.</a>	When the Egress Node is an ABR or ASBR . . . . .	<a href="#">12</a>
<a href="#">6.</a>	Acknowledgments . . . . .	<a href="#">13</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">14</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">14</a>
<a href="#">9.</a>	References . . . . .	<a href="#">14</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">14</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">15</a>
	Authors' Addresses . . . . .	<a href="#">15</a>

## [1.](#) Introduction

[RFC6513] and [[RFC6514](#)] define the "Selective Provider Multicast Service Interface Auto-Discovery route" (S-PMSI A-D route).

Per those RFCs, the S-PMSI A-D route contains a Network Layer Reachability Information (NLRI) field that identifies a particular multicast flow. In the terminology of those RFCs, each flow is denoted by (C-S,C-G), where C-S is an IP source address and C-G is an IP multicast address, both in the address space of a VPN customer. The (C-S,C-G) of the multicast flow is encoded into the NLRI field.

An S-PMSI A-D route also carries a PMSI Tunnel attribute (PTA). Typically, the PTA is used to identify a tunnel through the provider backbone network (a "P-tunnel").

By originating an S-PMSI A-D route identifying a particular multicast flow and a particular P-tunnel, a node is advertising the following:



if the node has to transmit packets of the identified flow over the backbone, it will transmit them through the identified tunnel.

[RFC6513] and [[RFC6514](#)] also define a procedure that allows an ingress node of particular multicast flow to determine the set of egress nodes that have requested to receive that flow from that ingress node. The ability of an ingress node to identify the egress nodes for a particular flow is known as "explicit tracking". An ingress node requests explicit tracking by setting a flag (the "Leaf Information Required" flag, or LIR) in the PTA. When an egress node receives an S-PMSI A-D route with LIR set, the egress node originates a Leaf A-D route whose NLRI field contains the NLRI from the corresponding S-PMSI A-D route. In this way, the egress node advertises that it has requested to receive the particular flow identified in the NLRI of that S-PMSI A-D route.

[RFC6513] and [[RFC6514](#)] also allow an ingress node to originate an S-PMSI A-D route whose PTA has LIR set, but which does not identify any P-tunnel. This mechanism can be used when it is desired to do explicit tracking of a flow without at the same time binding that flow to a particular P-tunnel.

[RFC6625] (and other RFCs that update it) extends the specification of S-PMSI A-D routes, and allows an S-PMSI A-D route to encode a wildcard in its NLRI. Either the C-S or the C-G or both can be replaced by wildcards. These routes are known as (C-\*,C-S) S-PMSI A-D routes, or as (C-S,C-\*) S-PMSI A-D routes, or as (C-\*,C-\*) S-PMSI A-D routes, depending on whether the C-S or C-G or both have been replaced by wildcards. These routes are known jointly as "wildcard S-PMSI A-D routes".

One purpose of this document is to clarify the way that the explicit tracking procedures of [[RFC6513](#)] and [[RFC6514](#)] are applied when wildcard S-PMSI A-D routes are used.

In addition, this document addresses the following scenario, which is not addressed in [[RFC6513](#)], [[RFC6514](#)], or [[RFC6625](#)]. Suppose an ingress node originates an S-PMSI A-D route whose NLRI specifies, for example, (C-\*,C-\*) (i.e., both C-S and C-G are replaced by wildcards), and whose PTA identifies a particular P-tunnel. Now suppose that the ingress node wants explicit tracking for each individual flow that it transmits (following the procedures of [[RFC6625](#)]) on that P-tunnel.

In this example, if the ingress node sets LIR in the PTA of the wildcard S-PMSI A-D route, each egress node that needs to receive a flow from the ingress node will respond with a Leaf A-D route whose NLRI specifies contains the (C-\*,C-\*) wildcard. This allows the



ingress node to determine the set of egress nodes that are interested in receiving flows from the ingress node. However, it does not allow the ingress node to determine exactly which flows are of interest to which egress nodes.

If the ingress node needs to determine which egress nodes are interested in receiving which flows, it needs to originate an S-PMSI A-D route for each individual (C-S,C-G) flow that it is transmitting, and it needs to set LIR in the PTA of each such route. However, since all the flows are being sent through the tunnel identified in the (C-\*,C-\*) S-PMSI A-D route, there is no need to identify a tunnel in the PTA of each (C-S,C-G) S-PMSI A-D route. Per [\[RFC6514\]](#), the PTA of the (C-S,C-G) S-PMSI A-D routes can specify "no tunnel information". This procedure allows explicit tracking of individual flows, even though all those flows are assigned to tunnels in wildcard S-PMSI A-D routes.

However, this procedure requires several clarifications:

- o The procedures of [\[RFC6625\]](#) do not clearly state how to handle an S-PMSI A-D route if its NLRI contains wild cards, but its PTA specifies "no tunnel info".
- o If it is desired to send a set of flows through the same tunnel (where that tunnel is advertised in a wildcard S-PMSI A-D route), but it is also desired to explicitly track each individual flow transmitted over that tunnel, one has to send an S-PMSI A-D route (with LIR set in the PTA) for each individual flow. It would be more optimal if the ingress node could just send a single wildcard S-PMSI A-D route binding the set of flows to a particular tunnel, and have the egress nodes respond with Leaf A-D routes for each individual flow.
- o [\[RFC6513\]](#) and [\[RFC6514\]](#) support the notion of "segmented P-tunnels", where "segmentation" occurs at Autonomous System Border Routers (ASBRs); [\[RFC7524\]](#) extends the notion of segmented P-tunnels so that segmentation can occur at Area Border Routers (ABRs). One can think of a segmented P-tunnel as passing through a number of "segmentation domains". In each segmentation domain, a given P-tunnel has an ingress node and a set of egress nodes. The explicit tracking procedures allow an ingress node of a particular segmentation domain to determine, for a particular flow or set of flows, the egress nodes of that segmentation domain. This has given rise to two further problems:
  - \* The explicit tracking procedures do not allow an ingress node to "see" past the boundaries of the segmentation domain.



- \* The prior specifications do not make it very clear whether a segmented tunnel egress node, upon receiving an S-PMSI A-D route whose PTA specifies "no tunnel information", is expected to forward the S-PMSI A-D route, with the same PTA, to the next segmentation domain.

These problems are resolved in [Section 5.3](#).

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL", when and only when appearing in all capital letters, are to be interpreted as described in [[RFC2119](#)].

## **2. The Explicit Tracking Flags**

[RFC6514] defines one flag in the PTA, the "Leaf Info Required" (LIR) flag, that is used for explicit tracking.

This document defines a new flag in the flags field of the PMSI Tunnel attribute. This new flag is known as the "Leaf Info Required per Flow" bit (LIR-pF). This flag MAY be set in the PTA of a (C-\*,C-\*), (C-\*,C-G), or (C-S,C-\*) S-PMSI A-D route. (Use of this flag in a PTA carried by other routes is outside the scope of this document.) Support for this flag is OPTIONAL.

The action taken by an egress node when the LIR-pF bit is set is detailed in [Section 5](#).

If the LIR-pF flag is set in a given PTA, the LIR flag of that PTA SHOULD also be set. (By setting LIR as well as LIR-pF, one forces a response to be sent by an egress node that does not support LIR-pF, and it is possible to tell from that response that the egress node does not support LIR-pF.)

## **3. Match for Tracking vs. Match for Reception**

[Section 3.2 of \[RFC6625\]](#) specifies a set of rules for finding the S-PMSI A-D route that is the "match for data reception" (or more simply, the "match for reception") of a given (C-S,C-G) or (C-\*,C-G) state. These rules do not take into account the fact that some S-PMSI A-D routes may not be carrying PTAs at all, or may be carrying PTAs that do not identify any P-tunnel. (A PTA that does not identify any P-tunnel is one whose "tunnel type" field has been set to "no tunnel information", as specified in [Section 5 of \[RFC6514\]](#).)

The rules for finding a "match for reception" in [[RFC6625](#)] are hereby modified as follows:





When applying the rules of [Section 3.2.1](#) or 3.2.2 of [\[RFC6625\]](#), it is REQUIRED to ignore any S-PMSI A-D route that has no PTA, or whose PTA specifies "no tunnel information".

There are other RFCs that update [\[RFC6625\]](#) and that modify the rules for finding a "match for reception". See, e.g., [\[RFC7582\]](#) and [\[RFC7900\]](#). When applying those modified rules, it is REQUIRED to ignore any S-PMSI A-D route that has no PTA, or whose PTA specifies "no tunnel information".

We also introduce a new notion: the "match for tracking". This differs from the "match for reception" as follows:

For a given C-flow ((C-S,C-G) or (C-\*,C-G)) the "match for tracking" is chosen as follows. Ignore any S-PMSI A-D route that has no PTA. Also ignore any S-PMSI A-D route whose PTA specifies "no tunnel information", but does not have either LIR or LIR-pF set. (That is, DO NOT ignore an S-PMSI A-D route that has a PTA specifying "no tunnel information" unless its LIR and LIR-pF bits are both clear). Then apply the rules (from [\[RFC6625\]](#) and other documents that that update it) for finding the "match for reception". The result (if any) is the match for tracking".

We will clarify this with a few examples. In these examples, we assume that there is only one segmentation domain. In this case, the ingress and egress nodes are Provider Edge (PE) routers.

Suppose a given PE router, PE1, has chosen PE2 as the "upstream PE" ([\[RFC6513\]](#)) for a given flow (C-S1,C-G1). And suppose PE1 has installed the following two routes that were originated by PE2:

- o Route1: A (C-\*,C-\*) S-PMSI A-D route, whose PTA specifies a tunnel.
- o Route2: A (C-S1,C-G1) S-PMSI A-D route, whose PTA specifies "no tunnel info" and has LIR set.

Route1 is (C-S1,C-G1)'s match for reception, and Route2 is (C-S1,C-G1)'s match for tracking.

Continuing this example, suppose:

- o PE1 has chosen PE2 as the upstream PE for a different flow, (C-S2,C-G2).
- o PE2 has not originated an S-PMSI A-D route for (C-S2,C-G2).



In this case, PE1 would consider Route1 to be (C-S2,C-G2)'s match for tracking as well as its match for reception.

Also note that if a match for tracking does not have the LIR flag or the LIR-pF flag set, no explicit tracking information will be generated. See [Section 5](#).

As another example, suppose PE1 has installed the following two routes that were originated by PE2:

- o Route1: A (C-\*,C-\*) S-PMSI A-D route (irrespective of whether the PTA specifies a tunnel)
- o Route2: A (C-S1,C-G1) S-PMSI A-D route whose PTA specifies a tunnel.

Then Route2 is both the "match for reception" and the "match for tracking" for (C-S1,C-G1).

Note that for a particular C-flow, PE1's match for reception might be the same route as its match for tracking, or its match for reception might be a "less specific" route than its match for tracking. But its match for reception can never be a "more specific" route than its match for tracking.

#### **4. Ingress Node Initiation of Tracking**

An ingress node that needs to initiate explicit tracking for a particular flow or set of flows can do so by performing one of the following procedures:

1. An ingress node can initiate explicit tracking for (C-S1,C-G1) by originating an S-PMSI A-D route that identifies (C-S1,C-G1) in its NLRI, including a PTA in that route, and setting the LIR flag in that PTA. The PTA may specify a particular tunnel, or may specify "no tunnel info".

However, the PTA of the (C-S1,C-G1) S-PMSI A-D route SHOULD NOT specify "no tunnel info" unless the ingress node also originates an A-D route carrying a PTA that specifies the tunnel to be used for carrying (C-S1,C-G1) traffic. Such a route could be an I-PMSI A-D route, a (C-\*,C-G1) S-PMSI A-D route, a (C-S1,C-\*) S-PMSI A-D route, or a (C-\*,C-\*) S-PMSI A-D route. (There is no point in requesting explicit tracking for a given flow if there is no tunnel on which the flow is being carried.)

Note that if the ingress node originates a wildcard S-PMSI A-D route carrying a PTA specifying the tunnel to be used for



carrying (C-S1,C-G1) traffic, and if that PTA has the LIR-pF bit set, then explicit tracking for (C-S1,C-G1) is requested by that S-PMSI A-D route. In that case, the ingress node SHOULD NOT originate a (C-S1,C-G1) S-PMSI A-D route whose PTA specifies "no tunnel info"; such a route would not provide any additional functionality.

To terminate explicit tracking that has been initiated by an S-PMSI A-D route whose PTA specifies "no tunnel info", the ingress node withdraws the route.

To terminate explicit tracking that has been initiated by an S-PMSI A-D route whose PTA specifies a tunnel, the ingress node re-originates the route without the LIR flag set.

2. The following procedure can be used if and only if it is known that the egress nodes support the optional LIR-pF flag. If the ingress node originates a wildcard S-PMSI A-D route, it can initiate explicit tracking for the individual flows that match the wildcard route by setting the LIR-pF flag in the PTA of the wildcard route. If an egress node needs to receive one or more flows for which that wildcard route is a match for tracking, the egress node will originate a Leaf A-D route for each such flow, as specified in [Section 5.2](#)).

When following this procedure, the PTA of the S-PMSI A-D route may specify a tunnel, or may specify "no tunnel info". The choice between these two options is determined by considerations that are outside the scope of this document.

To terminate explicit tracking that has been initiated by an S-PMSI A-D route whose PTA specifies "no tunnel info", the ingress node withdraws the route.

To terminate explicit tracking that has been initiated by an S-PMSI A-D route whose PTA specifies a tunnel, the ingress node re-originates the route without either the LIR or LIR-pF flags set.

Note that this procedure (procedure 2 of [Section 4](#)) may not yield the expected results if there are egress nodes that do not support the LIR-pF flag, and hence SHOULD NOT be used in that case.



## **5. Egress Node Response to the Match for Tracking**

### **5.1. General Egress Node Procedures**

There are four cases to consider:

1. With regard to a particular (C-S,C-G) or (C-\*,C-G) multicast state, the egress node's match for tracking is same as its match for reception, and neither LIR nor LIR-pF flags are on.

In this case, the egress node does not originate a Leaf A-D route in response to the match for reception/tracking, and there is no explicit tracking of the flow. This document specifies no new procedures for this case.

2. With regard to a particular (C-S,C-G) or (C-\*,C-G) multicast state, the egress node's match for tracking is the same as its match for reception, LIR is set, but LIR-pF is not set.

In this case, a Leaf A-D route is originated by the egress node, corresponding to the S-PMSI A-D route that is the match for reception/tracking. Construction of the Leaf A-D route is as specified in [[RFC6514](#)]; this document specifies no new procedures for this case.

3. With regard to a particular (C-S,C-G) or (C-\*,C-G) multicast state, the egress node's match for tracking is the same as its match for reception, and LIR-pF is set. The egress node MUST follow whatever procedures are required by other specifications, based on the match for reception. If the egress node supports the LIR-pF flag, it MUST also follow the procedures of [Section 5.2](#).

4. With regard to a particular (C-S,C-G) or (C-\*,C-G) multicast state, the egress node's match for tracking is not the same as its match for reception. This can only happen if the match for tracking has a PTA specifying "no tunnel info", with either LIR or LIR-pF set. In this case, the egress node must respond, separately, BOTH to the match for tracking and to the match for reception.

When responding to the match for reception, the egress node MUST ignore the LIR-pF flag. However, the LIR flag is processed normally per the procedures for the match for reception.

If the match for tracking has LIR set and if either (a) the egress node does not support LIR-pF, or (b) LIR-pF is not set, then the egress node must respond to the match for tracking,





following procedures specified in other documents for the case where LIR is set.

If the match for tracking has LIR-pF set, and the egress node supports the LIR-pF flag, the egress node must originate one or more Leaf A-D routes, as specified in [Section 5.2](#).

Note that if LIR is set in the PTA of the match for reception, the egress node may need to originate one or more Leaf A-D routes corresponding to the match for tracking, as well as originating a Leaf A-D route corresponding to the match for reception.

## **[5.2](#). Responding to the LIR-pF Flag**

To respond to a match for tracking that has LIR-pF set, an egress node originates one or more Leaf A-D routes.

Suppose the egress node has multicast state for a (C-S,C-G) or a (C-\*,C-G) flow, and has determined a particular S-PMSI A-D route, which has the LIR-pF flag set, to be the match for tracking for that flow. Then if the egress node supports the LIR-pF flag, it **MUST** originate a Leaf A-D route whose NLRI identifies that particular flow. Note that if a single S-PMSI A-D route (with wild cards) is the match for tracking for multiple flows, the egress node may need to originate multiple Leaf A-D routes, one for each such flow. We say that, from the perspective of a given egress node, a given S-PMSI A-D route tracks the set of flows for which it is the match for tracking. Each of the Leaf A-D routes originated in response to that S-PMSI A-D route tracks a single such flow.

The NLRI of each the Leaf A-D route that tracks a particular flow is constructed as follows. The "route key" field of the NLRI will have the following format:



```

+-----+
|      RD      (8 octets)      |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (Variable)      |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (Variable)      |
+-----+
| Ingress PE's IP address          |
+-----+

```

Figure 1: NLRI of S-PMSI A-D Route

- o The "ingress PE" address is taken from the "originating router" field of the NLRI of the S-PMSI A-D route that is the match for tracking.
- o The multicast source and group fields specify the S and G of one of the flow being tracked by this Leaf A-D route. If a (C-\*,C-G) is being tracked by this Leaf A-D route, the source field is omitted, and its length is set to 0.
- o The Route Distinguisher (RD) field is constructed as follows:
  - \* Take the RD value from the NLRI of the S-PMSI A-D route.
  - \* Per [\[RFC4364\]](#), every RD begins with a two-octet type field that is either 0, 1, or 2. Change the value of this two-octet type field to TBD1, TBD2, or TBD3, respectively. (See [Section 7](#), IANA Considerations.)

The presence of one of these values will indicate that the Leaf A-D route was constructed in response to a less specific S-PMSI A-D route that had the LIR-pF bit set. (That is, it distinguishes the routes from "ordinary" MVPN Leaf A-D routes.)

The encoding of these Leaf A-D routes is similar to the encoding of the Leaf A-D routes described in [section 6.2.2 of \[RFC7524\]](#), which were designed for the support of "global table multicast". However, that document sets the RD to either 0 or -1; following the procedures of this document, the RD will never be 0 or -1. Therefore Leaf A-D routes constructed according to the procedures of this section can always be distinguished from the Leaf A-D routes constructed according to the procedures of [section 6.2.2 of \[RFC7524\]](#). Also,



Leaf A-D routes constructed according to the procedures of this section are VPN-specific routes, and will always carry an IP-address-specific Route Target, as specified in [\[RFC6514\]](#).

If a Leaf A-D route is originated as a response to a match for tracking whose PTA specifies "no tunnel info", a PTA SHOULD NOT be attached to the Leaf A-D route; if a PTA is attached, it MUST specify "no tunnel info".

In the case where the match for tracking and the match for reception are the same, the PTA of the match may have both the LIR and the LIR-pF flags set. This may cause the egress node to originate one Leaf A-D route in response to the LIR bit, and one or more Leaf A-D routes in response to the LIR-pF bit. A PTA SHOULD NOT be attached to the Leaf A-D routes that are originated in response to the LIR-pF bit.

When a Leaf A-D route constructed according to the procedures of this section is received, it MUST be processed by the node identified in its IP-address-specific Route Target, even though its "route key" field does not correspond to the NLRI of any S-PMSI A-D route.

Of course, an egress node that originates such Leaf A-D routes needs to remember which S-PMSI A-D route caused these Leaf A-D routes to be originated; if that S-PMSI A-D route is withdrawn, those Leaf A-D routes MUST be withdrawn.

Similarly, a Leaf A-D route needs to be withdrawn (either implicitly or explicitly) if the egress node changes its Upstream Multicast Hop (UMH) ([\[RFC6513\]](#)) for the flow that is identified in the Leaf A-D route's NLRI, or if the egress node that originated the route no longer needs to receive the flow identified in the NLRI of the route.

Note that an egress node may acquire (C-S,C-G) state or (C-\*,C-G) state after it has already received the S-PMSI A-D that is the match for tracking for that state. In this case, a Leaf A-D route needs to be originated at that time, and the egress node must remember that the new Leaf A-D route corresponds to that match for tracking.

Note that if a particular S-PMSI A-D route is a match for tracking but not a match for reception, the LIR bit in its PTA is ignored if the LIR-pF bit is set.

### **[5.3.](#) When the Egress Node is an ABR or ASBR**

When segmented P-tunnels are used, the ingress and egress nodes may be ABRs or ASBRs. An egress ABR/ASBR that receives and installs an S-PMSI A-D route also forwards that route. If the PTA of an installed S-PMSI A-D route specifies a tunnel, the egress ABR/ASBR



MAY change the PTA to specify a different tunnel type (as discussed in [[RFC6514](#)] and/or [[RFC7524](#)]).

However, if the PTA of the installed S-PMSI A-D route specifies "no tunnel info", the egress ABR/ASBR MUST pass the PTA along unchanged when it forwards the S-PMSI A-D route. (That is, a PTA specifying "no tunnel info" MUST NOT be changed into a PTA specifying a tunnel.) Furthermore, if the PTA specifies "no tunnel info", the LIR and LIR-pF flags in the PTA MUST be passed along unchanged.

In the case where the egress node of a given P-tunnel is a PE, it will know whether it needs to receive a given flow by virtue of its having received a PIM or IGMP Join for that flow from a Customer Edge (CE) router. In the case where the egress node is not a PE, but rather an ABR or ASBR, it will not know whether it needs to receive a given flow unless it receives a Leaf A-D route whose NLRI specifies that flow, and which carries an IP-address-specific Route Target (RT) Extended Community that specifies the address of the egress ABR/ASBR. Therefore an egress ABR/ASBR MUST NOT originate a Leaf A-D route for a given flow UNLESS it has an installed Leaf A-D route for that flow, received from further downstream.

This will ensure that an egress ABR/ASBR only sends a Leaf A-D route in response to a "match for tracking" if it is on the path to an egress PE for the flow(s) identified in the corresponding S-PMSI A-D route.

We can thus establish the following rule for egress ABRs/ASBRs. Suppose an egress ABR/ASBR receives an S-PMSI A-D route whose NLRI is X, and whose PTA (a) specifies "no tunnel info" and (b) has LIR set. The egress ABR/ASBR should not immediately originate a Leaf A-D route in response. Rather it should wait until it receives a Leaf A-D route whose NLRI contains X in the "route key" field. If it receives such a Leaf A-D route, it redistributes that route, but first it changes that route's RT. The "global administrator" field of the modified RT will be set to the IP address taken either from the S-PMSI A-D route's next hop field, or from its Segmented P2MP Next Hop Extended Community. (This is the same rule that is used for when the PTA does specify a tunnel type.)

## **6. Acknowledgments**

The authors wish to thank Robert Kebler for his ideas and comments.





## 7. IANA Considerations

The LIR-pF flag needs to be added to the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Attribute Flags" in the "Border Gateway Protocol (BGP) Parameters" registry. This registry is defined in [RFC7902]. The requested value is Bit Position 2. This document should be the reference.

IANA is requested to allocate three new types from the Route Distinguisher Type Field registry:

- o TBD1: Administrator field is two-byte Autonomous System Number. To be used only in certain MCAST-VPN Leaf A-D routes.
- o TBD2: Administrator field is four-byte IP Address. To be used only in certain MCAST-VPN Leaf A-D routes.
- o TBD3: Administrator field is four-byte Autonomous System Number. To be used only in certain MCAST-VPN Leaf A-D routes.

The requested values are 16, 17, and 18 respectively.

## 8. Security Considerations

The Security Considerations of [RFC6513] and [RFC6514] apply.

By setting the LIR-pF flag in a single wildcard S-PMSI A-D route, a large number of Leaf A-D routes can be elicited. If this flag is set when not desired (through either error or malfeasance), a significant increase in control plane overhead can result.

## 9. References

### 9.1. 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>>.
- [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>>.
- [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>>.
- [RFC7902] Rosen, E. and T. Morin, "Registry and Extensions for P-Multicast Service Interface Tunnel Attribute Flags", [RFC 7902](#), DOI 10.17487/RFC7902, June 2016, <<https://www.rfc-editor.org/info/rfc7902>>.

## **9.2. Informative References**

- [RFC7582] Rosen, E., Wijnands, IJ., Cai, Y., and A. Boers, "Multicast Virtual Private Network (MVPN): Using Bidirectional P-Tunnels", [RFC 7582](#), DOI 10.17487/RFC7582, July 2015, <<https://www.rfc-editor.org/info/rfc7582>>.
- [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>>.

## **Authors' Addresses**

Andrew Dolganow

Nokia

438B Alexandra Rd #08-07/10

Alexandra Technopark

Singapore 119968

Singapore

Email: [andrew.dolganow@nokia.com](mailto:andrew.dolganow@nokia.com)



Jayant Kotalwar  
Nokia  
701 East Middlefield Rd  
Mountain View, California 94043  
United States of America

Email: jayant.kotalwar@nokia.com

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

Email: erosen@juniper.net

Zhaohui Zhang  
Juniper Networks, Inc.  
10 Technology Park Drive  
Westford, Massachusetts 01886  
United States of America

Email: zzhang@juniper.net

