

L3VPN Working Group
Internet Draft
Intended Status: Proposed Standard
Expires: August 1, 2011

Yiqun Cai
Eric C. Rosen (Editor)
IJsbrand Wijnands
Cisco Systems, Inc.

Arjen Boers
ZTE Corporation

February 1, 2011

MVPN: Using Bidirectional P-Tunnels

[draft-rosen-l3vpn-mvpn-bidir-03.txt](#)

Abstract

The documents specifying multicast support for BGP/MPLS IP VPNs allow customer multicast data to be transported through a service provider's network through a set of multicast tunnels. Such tunnels are advertised by BGP in a BGP attribute known as the "Provider Multicast Service Interface (PMSI) Tunnel Attribute". The base specifications allow the PMSI Tunnel Attribute to advertise bidirectional multicast distribution trees as "PMSI Tunnels"; however, those documents do not provide all the necessary details for using those tunnels. These details are provided in this document. This document also specifies the procedures for assigning customer multicast flows to specific bidirectional PMSI tunnels.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

Copyright and License Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

| | | |
|-----------------------|---|--------------------|
| 1 | Specification of requirements | 3 |
| 2 | Introduction | 4 |
| 3 | Advertising Bidirectional P-tunnels | 5 |
| 3.1 | BIDIR-PIM P-Tunnels | 5 |
| 3.2 | MP2MP LSPs | 6 |
| 3.2.1 | Partial Mesh of MP2MP LSPs | 6 |
| 3.2.2 | Single MP2MP LSP without PE Distinguisher Labels | 7 |
| 3.2.3 | Single MP2MP LSP with PE Distinguisher Labels | 8 |
| 3.2.4 | Identifying a MP2MP LSP | 8 |
| 4 | Associating Received Packets with VRFs | 9 |
| 5 | Data Transmission and Reception | 9 |
| 5.1 | Partial Mesh of MP2MP LSPs | 9 |
| 5.1.1 | Binding (C-S,C-G) to Bidirectional P-tunnels | 9 |
| 5.1.2 | Binding (C-*,C-G) Flows from Unidirectional C-trees ... | 10 |
| 5.1.3 | Binding (C-*,C-G) Flows from Bidirectional C-trees | 10 |
| 5.1.4 | Binding (C-*,C-*) | 11 |
| 5.2 | Single MP2MP LSP With PE Distinguisher Labels | 13 |
| 5.3 | Single MP2MP LSP Without PE Distinguisher Labels | 13 |
| 5.4 | BIDIR-PIM P-Tunnel | 13 |
| 5.4.1 | Single BIDIR-PIM P-Tunnel | 14 |
| 5.4.2 | Partial Mesh of BIDIR-PIM P-Tunnels | 14 |
| 6 | IANA Considerations | 14 |
| 7 | Security Considerations | 14 |
| 8 | Acknowledgments | 14 |
| 9 | Authors' Addresses | 15 |
| 10 | Normative References | 15 |
| 11 | Informative References | 16 |

[1](#). Specification of requirements

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

2. Introduction

The base documents for MVPN, [[MVPN](#)] and [[MVPN-BGP](#)], define a "PMSI Tunnel Attribute" (PTA) that may be carried in the BGP "I-PMSI A-D routes" and BGP "S-PMSI A-D routes" that are defined therein. The base documents define the way that bidirectional P-tunnels are identified in the PTA, and the way in which the identifier of a bidirectional P-tunnel is encoded in the PTA.

However, those documents do not contain the full set of specifications governing the use of the PTA to advertise bidirectional P-Tunnels. These specifications are provided in this document.

This document also specifies the procedures for assigning customer multicast flows to specific bidirectional PMSI tunnels.

Two kinds of bidirectional P-tunnel are discussed in this document:

- Multicast distribution trees that are created through the use of BIDIR-PIM [[BIDIR-PIM](#)].
- Multipoint-to-multipoint Label Switched Paths (MP2MP LSPs), created by Label Distribution Protocol (LDP) Multipoint-to-Multipoint extensions [[mLDP](#)].

This document also specifies three methods of using MP2MP LSPs as P-tunnels:

- Partial mesh of MP2MP LSPs. In this method, when a set of PEs have multicast data to send and/or receive to/from each other, each PE becomes the root of a MP2MP LSP. This method is presented in [[MVPN](#)], section 11.2.3. The detailed specification is provided in this document.
- Single MP2MP LSP with PE Distinguisher Labels. This method is presented in [[MVPN](#)], section 11.2.2. The detailed specification is provided in this document.
- Single MP2MP LSP without PE Distinguisher Labels.

In the following, we will sometimes speak of an S-PMSI A-D route being "ignored". When we say the route is "ignored", we do not mean that it's normal BGP processing is not done, but that the route is not considered when determining which P-tunnel to use when sending multicast data, and that the MPLS label values it conveys are not used. We will generally use "ignore" in quotes to indicate this meaning.

3. Advertising Bidirectional P-tunnels

In this specification, we consider the use of bidirectional P-tunnels as advertised in the PTA of a BGP S-PMSI A-D route.

3.1. BIDIR-PIM P-Tunnels

A BIDIR-PIM P-tunnel may be advertised in the PTA of an Intra-AS I-PMSI A-D route or in the PTA of an S-PMSI A-D route.

As is the case with other PIM-created P-tunnels, to transmit packets on a BIDIR-PIM P-tunnel, one uses the GRE encapsulation as specified in [\[MVPN\] section 12](#).

Each BIDIR-PIM P-Tunnel is identified by a unique P-group address [\[MVPN, section 3.1\]](#). (The P-group address is called a "P-Multicast Group" in [\[MVPN-BGP\]](#)). A BIDIR-PIM P-group address is always associated with a unique "Rendezvous Point Address" (RPA).

Every PE that needs to join a particular BIDIR-PIM P-tunnel must be able to determine the RPA that corresponds to the P-tunnel's P-group address. PIM Join/Prune messages are sent along the path from the PE to the RPA. Any P routers along that path must also be able to determine the RPA, so that they too can send PIM Join/Prune messages towards the RPA. The method of mapping a P-group address to an RPA may be static configuration, or some automated means of RPA discovery that is outside the scope of this specification.

If a BIDIR-PIM P-tunnel is to be used, it is RECOMMENDED that the path from each PE in the tunnel to the RPA consist entirely of point-to-point links. On a point-to-point link, there is no ambiguity in determining which router is upstream towards a particular RPA, so the BIDIR-PIM "Designated Forwarder Election" is very quick and simple. Use of a BIDIR-PIM P-tunnel containing multiaccess links is possible, but considerably more complex.

When a BGP A-D route's PTA specifies a BIDIR-PIM P-tunnel, the PE Distinguisher Labels attribute SHOULD NOT be included; if it is included, it MUST be ignored. The PE Distinguisher Labels attribute is not needed because the "distinguished PE" for a particular packet can be identified by placing its IP address in the IP source address field of the GRE encapsulation used to send that packet on the P-tunnel.

There are two different methods of using BIDIR-PIM P-Tunnels, the "Single BIDIR-PIM P-Tunnel method", and the "Partial Mesh of

BIDIR-PIM P-Tunnels method". The choice of method is determined by provisioning.

If the "Single BIDIR-PIM P-Tunnel" method is being used, all PEs in a given MVPN MUST identify the same P-tunnel. The identity of this P-tunnel is known by provisioning. For example, by using this method, and identifying the tunnel in the PTA of the Intra-AS I-PMSI A-D routes, one may use a BIDIR-PIM P-tunnel to instantiate an MI-PMSI.

If the Partial Mesh of BIDIR-PIM P-Tunnels method is being used, the PEs MUST identify different P-Tunnels (by advertising different P-group addresses in their PTAs). If a particular P-group address is advertised by a particular PE, then one of that PE's addresses MUST be the RPA corresponding to that P-group.

For example, by using this method, and identifying the tunnel in the PTA of an S-PMSI A-D route, one may implement the "Partitioned Sets of PEs" method of supporting C-BIDIR, as discussed in section 11.2 of [\[MVPN\]](#) and [section 3.6](#) of [\[CONSID\]](#).

[3.2. MP2MP LSPs](#)

An MP2MP LSP is identified by an "MP2MP FEC element" [\[mLDP\]](#). The FEC element contains the IP address of the "root", followed by an "opaque value" that identifies the MP2MP LSP uniquely in the context of the root's IP address. This opaque value may be configured or autogenerated, and within an MVPN, there is no need for different roots to use the same opaque value.

The method of using MP2MP LSPs (partial mesh, single with PE Distinguisher Labels, single without PE Distinguisher Labels) is determined by provisioning. It SHOULD be possible to configure this on a per-MVPN basis.

[3.2.1. Partial Mesh of MP2MP LSPs](#)

A partial mesh of MP2MP LSPs is not useful for instantiating an I-PMSI. The LSPs of the partial mesh are therefore only advertised in the PTAs of S-PMSI A-D routes.

A partial mesh of MP2MP LSPs is useful for implementing the "Partitioned Sets of PEs" method of supporting C-BIDIR, as discussed in section 11.2 of [\[MVPN\]](#) and [section 3.6](#) of [\[CONSID\]](#).

[Section 5.1](#) of this document specifies the procedures for

transmitting all kinds of customer data flows, whether unidirectional or bidirectional, on a partial mesh of MP2MP LSPs. The sending and receiving of PE-PE PIM control packets on a partial mesh of MP2MP LSPs is outside the scope of this specification.

When this method is being used:

- Each PE that is participating in the mesh MUST advertise, in the PTA of an S-PMSI A-D route, a MP2MP LSP of which it is the root. (The LSP root address is part of the P-tunnel identifier field carried in the PTA.) A PE MUST "participate in the mesh" if any of the following conditions holds:
 - * The "Partitioned Sets of PEs" method of supporting C-BIDIR traffic is being used, and the PE's route to the Rendezvous Point Address (RPA) for one or more C-BIDIR groups is via a VRF interface.
 - * The "Partitioned Sets of PEs" method is being used, it is desired to transmit some or all of the customer unidirectional multicast traffic (for the given MVPN) on the same LSPs used for carrying C-BIDIR traffic, and the PE has customer multicast traffic to transmit to other PEs.

There may be other conditions under which a PE needs to participate in a partial mesh of MP2MP LSPs; these are outside the scope of the current specification.

- The PE Distinguisher Labels Attribute [[MVPN-BGP](#)] MUST NOT be included; if included, it MUST be ignored.

3.2.2. Single MP2MP LSP without PE Distinguisher Labels

When this method is being used:

- the MP2MP LSP can be advertised in the PTA of an Intra-AS I-PMSI A-D route, or in the PTA of an S-PMSI A-D route. When advertised in the PTA of an Intra-AS I-PMSI A-D route, the MP2MP LSP can be used to instantiate an MI-PMSI.
- The LSP does not have to be advertised by its root. In fact, the root of the LSP does not even need to be a PE router.
- The PE Distinguisher Labels Attribute MUST NOT be included, but if included, it MUST be ignored.

- If two or more PEs of the same MVPN advertise a MP2MP LSP in their Intra-AS I-PMSI A-D routes, they SHOULD advertise the same MP2MP LSP. Any scenario in which they do not advertise the same MP2MP LSP in their Intra-0I A-D routes is outside the scope of this document.

This method cannot be used to support the "Partitioned Set of PEs" method discussed in [[MVPN](#)] [section 11.2](#) and [[CONSID](#)] [section 3.6](#). Also, this method is not compatible with the procedures of [[MVPN](#)] [section 9.1.1](#).

[3.2.3](#). Single MP2MP LSP with PE Distinguisher Labels

In this method, the MP2MP LSP MUST be advertised in the PTA of an Intra-AS I-PMSI A-D route or an S-PMSI A-D route originated by the root of the LSP. That route MUST include a "PE Distinguisher Labels" Attribute. Violation of these conditions MUST cause the route to be ignored.

The PE at the root of the LSP SHOULD use the PE Distinguisher Labels Attribute to bind an upstream-assigned MPLS label to the IP address of some or all of the other PEs that attach to the same MVPN (as determined by the RTs of the A-D route). That is, the PE at the root of the P-tunnel assigns a distinct label to some or all of the other PEs attaching to the same MVPN.

An MP2MP LSP with PE Distinguisher Labels can be used to instantiate an MI-PMSI. In this case, the PE at the root MUST use the PE Distinguisher Labels Attribute to bind an upstream-assigned MPLS label to the IP address of each other PE that attaches to the same MVPN. This set of PEs is learned via the reception of Intra-AS I-PMSI A-D routes.

An MP2MP LSP with PE Distinguisher Labels can also be used to support the "Partitioned Set of PEs" method discussed in [[MVPN](#)] [section 11.2](#) and [[CONSID](#)] [section 3.6](#). In this case, it is not necessary to bind an upstream-assigned label to the IP address of a particular PE in the MVPN unless that PE has advertised a unicast VPN-IP route to one of the C-RPAs of that MVPN.

[3.2.4](#). Identifying a MP2MP LSP

To identify a MP2MP LSP, the PTA of a BGP A-D route contains an MP2MP FEC Element [[mLDP](#)] in its "Tunnel Identifier" field. This contains the IP address of the root of the LSP, as well as an "Opaque Value" which is unique at the root. The mLDP specification supports the use

of several different ways of constructing the tunnel identifiers. This specification does not place any restrictions on the types of tunnel identifier that might be used. However, a given implementation might not support every possible type of tunnel identifier. Future revisions of this specification will establish one or two types of tunnel identifier as being "mandatory to support".

4. Associating Received Packets with VRFs

When a packet is received from a bidirectional P-tunnel, the packet is first associated one or more VRFs, and then processed in the context of that VRF or VRFs. If the bidirectional P-tunnel was advertised in the PTA of an A-D route that did not specify an MPLS label, then all packets received from the P-tunnel are associated with the same set of VRFs. If the bidirectional P-tunnel was advertised in the PTA of an A-D route, and the PTA does specify an MPLS label, then received packets will carry a label that must be processed in order to determine the context. If the P-tunnel is a MP2MP LSP, this label appears below the label that identifies the LSP itself.

5. Data Transmission and Reception

5.1. Partial Mesh of MP2MP LSPs

5.1.1. Binding (C-S,C-G) to Bidirectional P-tunnels

When PE1 advertises an S-PMSI A-D route that binds a (C-S,C-G) flow to a bidirectional P-tunnel, or when PE1 sends an S-PMSI Join message that binds a (C-S,C-G) flow to a bidirectional P-tunnel, the semantics are as follows. PE1 is stating that any (C-S,C-G) traffic that it needs to transmit to other PEs will be transmitted on the specified P-tunnel. Any other PE that needs to receive such traffic from PE1 (i.e., any other PE that needs to receive (C-S,C-G) traffic and which has selected PE1 as the upstream PE for C-S) MUST join that P-tunnel.

If a PE has joined the P-tunnel, but does not need to receive the (C-S,C-G) traffic, or if it needs to receive (C-S,C-G) traffic but has not selected PE1 as the upstream PE for C-S, then the PE MUST discard any such received traffic.

5.1.2. Binding (C-*,C-G) Flows from Unidirectional C-trees

When PE1 advertises an S-PMSI A-D route or sends an S-PMSI Join message that binds (C-*,C-G) [[MVPN-WILD](#)] to a bidirectional P-tunnel, where C-G is not a "Source-Specific Multicast" (SSM) group, and the (C-*,C-G) traffic is traveling on a unidirectional shared C-tree, the semantics are as follows. PE1 is stating that any traffic to C-G that is traveling the shared C-tree and which PE1 needs to transmit to other PEs will be transmitted on the specified P-tunnel.

Any other PE that needs to receive such traffic from PE1 (i.e., any other PE that needs to receive (C-*,C-G) traffic and which has selected PE1 as the upstream PE for the C-RP corresponding to the C-G group) MUST join that P-tunnel.

If a PE has joined the P-tunnel, but does not need to receive the (C-*,C-G) traffic, or if it needs to receive (C-*,C-G) traffic but has not selected PE1 as the upstream PE for the C-RP that corresponds to C-G, then the PE MUST discard any such received traffic.

5.1.3. Binding (C-*,C-G) Flows from Bidirectional C-trees

When PE1 advertises an S-PMSI A-D route or sends an S-PMSI Join message that binds (C-*,C-G) to a bidirectional P-tunnel, where C-G is not an SSM group, and the (C-*,C-G) traffic is traveling on a bidirectional shared C-tree, the semantics are as follows:

- PE1 is stating that any traffic to C-G that it (PE1) needs to send downstream will be sent on the specified P-tunnel
- Any other PE that is interested in receiving (C-*,C-G) traffic MUST join the specified P-tunnel
- Any other PE, say PE2, that (a) has traffic to C-G to send upstream and (b) has selected PE1 as its upstream PE for the C-RPA corresponding to C-G, MUST join the specified P-tunnel, and MUST send such traffic on the specified P-tunnel.
- If a PE, say PE3, has joined the specified P-tunnel, but does not need to receive the (C-*,C-G) traffic, or has not selected PE1 as the upstream PE for the C-RPA corresponding to C-G, then PE3 MUST NOT send any (C-*,C-G) traffic on that P-tunnel, and MUST discard any (C-*,C-G) traffic it received on that P-tunnel.

These procedures implement the "Partitioned Set of PEs" scheme described in section 11.2 of [[MVPN](#)].

The specification given so far requires an S-PMSI A-D route or an S-PMSI Join message to be sent for each (C-*,C-G) that is using a bidirectional C-tree. A more efficient method is given in the next section.

5.1.4. Binding (C-*,C-*)

With regard to the procedures of this section, the "all BIDIR-PIM" wildcard [[MVPN-WILD](#)] is treated identically to the (C-*,C-*) wildcard, except that it applies only to customer BIDIR-PIM flows.

When PE1 advertises an S-PMSI A-D route or sends an S-PMSI Join message that binds (C-*,C-*) to a specified bidirectional P-tunnel of which PE1 is the root, the semantics are as that the bidirectional P-tunnel is to be used to carry C-multicast traffic in the following sets of cases:

1. If PE1 has (C-S,C-G) traffic that is traveling on a source-specific C-tree, and PE1 needs to transmit that data to one or more other PEs, and PE1 has not bound (C-S,C-G) or (C-*,C-G) to a different P-tunnel, then the (C-S,C-G) traffic is sent by PE1 on the specified bidirectional P-tunnel.
2. If PE1 has (C-*,C-G) traffic that is traveling on a unidirectional shared C-tree, and PE1 needs to transmit that data to one or more other PEs, and PE1 has not bound (C-*,C-G) to a different P-tunnel, then the (C-*,C-G) traffic is sent by PE1 on the specified bidirectional P-tunnel.
3. If PE1 has (C-*,C-G) traffic that is traveling on a bidirectional shared C-tree, and PE1 needs to transmit that data to one or more other PEs, and PE1 has not bound (C-*,C-G) to a different P-tunnel, then the (C-*,C-G) traffic is sent by PE1 on the specified bidirectional P-tunnel.
4. Consider some other PE, PE2, that has received the S-PMSI A-D route or S-PMSI Join message from PE1. If PE2 has (C-*,C-G) traffic that is traveling on a bidirectional shared C-tree, and PE2 needs to transmit that traffic UPSTREAM, and PE2 has selected PE1 as the upstream PE for the C-RPA corresponding to C-G, and PE1 has not bound (C-*,C-G) to any other P-tunnel, then the (C-*,C-G) traffic is sent by PE2 on the specified bidirectional P-tunnel.

5. If a PE receives traffic from a particular bidirectional P-tunnel, and the traffic is traveling a unidirectional (C-*,C-G) or (C-S,C-G) tree, and the root of the bidirectional P-tunnel is not the PE's selected upstream PE for the (C-*,C-G) or (C-S,C-G), the PE MUST discard the traffic.
6. If a PE receives traffic from a particular bidirectional P-tunnel, and the traffic is traveling a bidirectional (C-*,C-G) tree, and the PE's selected upstream PE for the C-RPA corresponding to C-G is not the root of the bidirectional P-tunnel, then the PE MUST discard the traffic.

With respect to traffic traveling a bidirectional C-tree, these procedures implement, for S-PMSIs, the "partitioning" scheme described in section 11.2 of [\[MVPN\]](#), without the need to send an S-PMSI A-D route for each (C-*,C-G) that is using a bidirectional C-tree. Each PE becomes the root of a bidirectional P-tunnel, and binds the double wildcard selector to it. The bidirectional P-tunnels serve as the "partitions". The bidirectional P-tunnel rooted at PE1 becomes the default P-tunnel for all traffic that PE1 needs to send downstream to other PEs. It also becomes the default P-tunnel for all traffic that others PEs need to send upstream, as long as those other PEs have selected PE1 as the upstream PE for the C-RPA corresponding to that traffic.

Note that other PEs SHOULD NOT join the specified bidirectional P-tunnel unless they have a need to send or receive data over it. A PE knows when it needs to receive data by virtue of having certain multicast state in its C-PIM instance. With regard to multicast data traveling on a bidirectional (C-*,C-G) tree, a PE may not know whether it has to send data until such data actually arrives over a VRF interface; the PE may be on a "sender-only" branch. However, the PE in this case would have to know, through provisioning or some automatic procedure such as "Bootstrap Routing Protocol for PIM" (BSR) [\[BSR\]](#), the set of C-RPAs that are being used to support (C-*,C-G) traffic. For each C-RPA, the PE could join the bidirectional P-tunnel advertised by its selected upstream PE for that C-RPA. Alternatively the PE could defer joining the P-tunnel until it actually has data to send.

5.2. Single MP2MP LSP With PE Distinguisher Labels

The procedures for transmitting data on a single MP2MP LSP with PE Distinguisher Labels differ from the procedures for transmitting data on a partial mesh of MP2MP LSPs only in the following way. Let PE1 be the root of the P-tunnel. When a packet that is traveling on a unidirectional C-tree is transmitted on the P-tunnel by a particular PE, say PE2, PE2 must push on the packet's label stack the label that PE1 assigned to PE2 via the procedure above. When a packet that is traveling on a bidirectional C-tree is transmitted on the P-tunnel by PE2, PE2 must push on the packet's label stack the label that PE1 assigned to PE3, where PE3 is the upstream PE that PE2 has selected for the C-RPA corresponding to C-G.

For unidirectional flows, this allows the transmitter to be identified, and for bidirectional flows, this allows the partition to be identified. Packets received from the wrong upstream PE or from the wrong partition MUST be discarded. (In effect, this is a case of tunnel hierarchy, where the PE Distinguisher Labels represent a set of MP2MP LSPs that are being tunneled through a single bidirectional P-tunnel.)

If the PTA identifying the bidirectional P-tunnel contains an MPLS label, then that label shall appear in the label stack immediately preceding the label specified in the PE Distinguisher Labels attribute.

5.3. Single MP2MP LSP Without PE Distinguisher Labels

No special rules are needed for this case; the general procedures specified in [[MVPN](#)] and [[MVPN-BGP](#)] are used.

5.4. BIDIR-PIM P-Tunnel

Packets are transmitted using GRE encapsulation as described in sections [12.1.1](#) and [12.2.1](#) of [[MVPN](#)].

It is possible to implement the "Partitioned Sets of PEs" scheme ([[MVPN](#)] [section 11.2](#) and [[CONSID](#)] [section 3.6](#)) using either a single BIDIR-PIM P-Tunnel or using a partial mesh of BIDIR-PIM P-Tunnels.

[5.4.1.](#) Single BIDIR-PIM P-Tunnel

In this method, the rules for transmitting packets of a given C-flow on a BIDIR-PIM P-Tunnel are essentially the same as the rules given in [section 5.2](#), except that a particular "distinguished PE" is identified not by the use of a PE Distinguisher Label, but by the use of the IP Source Address field of the GRE header. For unidirectional C-flows, the IP source address field of the GRE header identifies the PE that transmitted the packet onto the P-tunnel. For bidirectional C-flows, suppose that PE1 is transmitting a packet over the P-tunnel, that the packet's C-group address is C-G, and that PE1 has selected PE2 as the upstream PE corresponding to the C-RPA that corresponds to C-G. Then when PE1 transmits the packet over the P-tunnel, the IP source address field of the GRE header will contain the IP address of PE2.

[5.4.2.](#) Partial Mesh of BIDIR-PIM P-Tunnels

In this method, the rules for for transmitting packets of a given C-flow on a BIDIR-PIM P-Tunnel are essentially the same as the rules given in [section 5.1](#).

[6.](#) IANA Considerations

Both [[MVPN](#)] and [[MVPN-BGP](#)] discuss the use of the "PE Distinguisher Labels" Attribute, but neither document has asked IANA to define a codepoint for it. We now ask IANA to assign a codepoint for this attribute, as an optional transitive attribute, referencing [[MVPN](#)], [[MVPN-BGP](#)], and this document.

[7.](#) Security Considerations

There are no additional security considerations beyond those of [[MVPN](#)] and [[MVPN-BGP](#)].

[8.](#) Acknowledgments

The authors wish to thank Karthik Subramanian, Rajesh Sharma, and Apoorva Karan for their input. We also thank Yakov Rekhter for his valuable critique.

9. Authors' Addresses

Arjen Boers
ZTE Corporation
114 Rue Gallieni
92100 Boulogne Billancourt
France

E-mail: arjen.boers@zte.com.cn

Yiqun Cai
Cisco Systems, Inc.
170 Tasman Drive
San Jose, CA, 95134
E-mail: ycai@cisco.com

Eric C. Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA, 01719
E-mail: erosen@cisco.com

IJsbrand Wijnands
Cisco Systems, Inc.
De kleetlaan 6a Diegem 1831
Belgium
E-mail: ice@cisco.com

10. Normative References

[BIDIR-PIM] "Bidirectional Protocol Independent Multicast", Handley, Kouvelas, Speakman, Vicisano, [RFC 5015](#), October 2007

[mLDP] "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", Minei, Kompella, Wijnands, Thomas, [draft-ietf-mpls-ldp-p2mp-11.txt](#), October 2010

[MVPN] "Multicast in MPLS/BGP IP VPNs", Rosen, Aggarwal, et. al., [draft-ietf-l3vpn-2547bis-mcast-10.txt](#), January 2010

[MVPN-BGP] "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", Aggarwal, Rosen, Morin, Rekhter, [draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt](#), September 2009

[MVPN-WILD] "MVPN: S-PMSI Wild Card Selectors", Cai, Rosen, Wijnands, [draft-rosen-l3vpn-mvpn-wildcards-03.txt](#), February 2011

[RFC2119] "Key words for use in RFCs to Indicate Requirement Levels.", Bradner, March 1997

11. Informative References

[BSR] "Bootstrap Router (BSR) Mechanism for PIM", N. Bhaskar, et.al., [RFC 5059](#), January 2008

[CONSID] "Mandatory Features in a Layer 3 Multicast BGP/MPLS VPN Solution", Morin, Niven-Jenkins, Kamite, Zhang, Leymann, Bitar, [draft-ietf-l3vpn-mvpn-considerations-06.txt](#), February 2010

