

Internet Engineering Task Force (IETF)
Request for Comments: 6625
Updates: 6514
Category: Standards Track
ISSN: 2070-1721

E. Rosen, Ed.
Cisco Systems, Inc.
Y. Rekhter, Ed.
Juniper Networks, Inc.
W. Henderickx
Alcatel-Lucent
R. Qiu
Huawei
May 2012

Wildcards in Multicast VPN Auto-Discovery Routes

Abstract

In Multicast Virtual Private Networks (MVPNs), customer multicast flows are carried in "tunnels" through a service provider's network. The base specifications for MVPN define BGP multicast VPN "auto-discovery routes" and specify how to use an auto-discovery route to advertise the fact that an individual customer multicast flow is being carried in a particular tunnel. However, those specifications do not provide a way to specify, in a single such route, that multiple customer flows are being carried in a single tunnel. Those specifications also do not provide a way to advertise that a particular tunnel is to be used by default to carry all customer flows, except in the case where that tunnel is joined by all the provider edge routers of the MVPN. This document eliminates these restrictions by specifying the use of "wildcard" elements in the customer flow identifiers. With wildcard elements, a single auto-discovery route can refer to multiple customer flows or even to all customer flows.

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 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6625>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Wildcards in S-PMSI A-D Routes	5
1.3. Use Cases	5
2. Encoding of Wildcards	7
3. Finding the Matching S-PMSI A-D Route	8
3.1. Finding the Match for Data Transmission	8
3.2. Finding the Match for Data Reception	9
3.2.1. Finding the Match for (C-S,C-G)	9
3.2.2. Finding the Wildcard Match for (C-*,C-G)	9
4. Procedures for S-PMSI A-D Routes with Wildcards	10
4.1. Procedures for All Kinds of Wildcards	10
4.2. Procedures for (C-*,C-G) S-PMSI A-D Routes	11
4.3. Procedures for (C-S,C-*) S-PMSI A-D Routes	12
4.4. Procedures for (C-*,C-*) S-PMSI A-D Routes	13
5. Security Considerations	15
6. Acknowledgments	15
7. Normative References	15

1. Introduction

In Multicast Virtual Private Networks (MVPNs), customer multicast flows are carried in tunnels through a service provider's network. The base specifications for MVPN define BGP multicast VPN "auto-discovery routes" and specify how to use an auto-discovery route to advertise the fact that an individual customer multicast flow is being carried in a particular tunnel. However, those specifications do not provide a way to specify, in a single such route, that multiple customer flows are being carried in a single tunnel. Those specifications also do not provide a way to advertise that a particular tunnel is to be used by default to carry all customer flows, except in the case where that tunnel is joined by all the provider edge routers of the MVPN. This document eliminates these restrictions by specifying the use of "wildcard" elements in the customer flow identifiers. With wildcard elements, a single auto-discovery route can refer to multiple customer flows or even to all customer flows.

1.1. Terminology

This document uses terminology from [MVPN] and, in particular, uses the prefixes "C-" and "P-", as specified in Section 3.1 of [MVPN], to distinguish addresses in the "customer address space" from addresses in the "provider address space". The following terminology and acronyms are particularly important in this document:

- MVPN

Multicast Virtual Private Network -- a VPN [L3VPN] in which multicast service is offered.

- VRF

VPN Routing and Forwarding table [L3VPN].

- SP

Service Provider.

- P-tunnel

A tunnel through the network of one or more SPs.

- C-S

Multicast Source. A multicast source address, in the address space of a customer network.

- C-G

Multicast Group. A multicast group address (destination address) in the address space of a customer network.

- C-multicast flow or 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).

- RP

A "Rendezvous Point", as defined in [PIM].

- C-RP

A Rendezvous Point whose address is in the customer's address space.

- Selective P-tunnel

A P-tunnel that is joined only by Provider Edge (PE) routers that need to receive one or more of the C-flows that are traveling through that P-tunnel.

- Inclusive P-tunnel

A P-tunnel that is joined by all PE routers that attach to sites of a given MVPN.

- S-PMSI A-D route

Selective Provider Multicast Service Interface 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.

Familiarity with multicast concepts and terminology [PIM] is also presupposed.

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].

1.2. Wildcards in S-PMSI A-D Routes

As specified in [MVPN] and [MVPN-BGP], an S-PMSI A-D route advertises that a particular C-flow is bound to a particular selective P-tunnel.

The identifier of the specified C-flow, e.g., (C-S,C-G), is encoded into the Network Layer Reachability Information (NLRI) of the S-PMSI A-D route. The identifier of the specified P-tunnel is encoded into an attribute (the "PMSI Tunnel Attribute") of the S-PMSI A-D route. Each S-PMSI A-D route thus specifies a single C-flow. To bind multiple C-flows to a single P-tunnel, it is necessary to advertise one S-PMSI A-D route for each C-flow, specifying the same P-tunnel in each such route.

This document defines OPTIONAL extensions to the procedures and encodings specified in [MVPN] and [MVPN-BGP]. These extensions enable a single S-PMSI A-D route to advertise that multiple C-multicast flows are bound to a single P-tunnel.

The extensions specified in this document are based on the notion of allowing the NLRI of an S-PMSI A-D route to contain a "wildcard". In the NLRI encoding, a wildcard can replace the C-S, the C-G, or both. We use the notation "C-*" to denote a wildcard. The extensions allow the NLRI to encode three kinds of wildcards: (C-*,C-*), (C-S,C-*), and (C-*,C-G).

By using wildcards, a PE may be able to reduce the number of S-PMSI A-D routes it originates, thereby improving the scalability of the control plane. There is, however, no impact on data plane scalability, as the number of P-tunnels is not reduced.

Encoding and detailed procedures are specified in subsequent sections of this document.

1.3. Use Cases

There are a number of situations in which it can be useful to use wildcards in the NLRI of an S-PMSI A-D route.

- Using a selective P-tunnel as the default tunnel.

There are procedures in [MVPN] and [MVPN-BGP] that allow a PE to advertise that it is going to use an inclusive P-tunnel as the P-tunnel on which it will transmit all C-flows by "default". However, those documents do not provide any way for a PE to advertise that it is going to use a selective P-tunnel as the P-tunnel on which it will transmit all C-flows by "default". Using the extensions defined in this document, a PE can

advertise that it is going to use a selective P-tunnel as its default P-tunnel. It does so by advertising an S-PMSI A-D route whose NLRI contains (C-*,C-*).

- Binding multiple C-flows traveling along a customer's Protocol Independent Multicast - Sparse Mode (PIM-SM) shared tree to a single P-tunnel.

A PE router may be connected to an MVPN site that contains a customer RP (C-RP). The C-RP may be the root of one or more shared trees. In multicast terminology, these are known as (*,G) trees. By advertising a single S-PMSI A-D route whose NLRI contains the (C-*,C-G) wildcard, the PE can bind all the C-flows traveling along a customer's (*,G) tree to a single P-tunnel. This use case applies only when C-G is a non-bidirectional ASM (Any Source Multicast) group.

- Binding multiple C-flows with the same C-group address to a single P-tunnel, even if each such C-flow is traveling along a customer's PIM source tree.

A PE router may be connected to an MVPN site containing several multicast sources that are all sending to a common multicast group, along a customer's PIM source trees. Alternatively, the PE may be connected to several sites, each containing at least one source sending to the common multicast group. By advertising a single S-PMSI A-D route whose NLRI contains (C-*,C-G), the PE can bind these C-flows to a single P-tunnel.

This use case applies only when the C-group is a non-bidirectional ASM group.

- Binding multiple C-flows with the same C-group address to a single P-tunnel, when those C-flows are traveling along a customer's BIDIR-PIM shared tree.

This use case applies only when the C-group is a BIDIR-PIM group.

- Binding multiple C-flows from a given C-source to a given P-tunnel, irrespective of whether those C-flows all have the same C-group address.

This can be useful when the C-group addresses are SSM (Single Source Multicast) addresses. Suppose, for example, that a given source transmits multiple "channels" of information, each with

its own C-group address. It may be desirable to bind all these channels to a single P-tunnel, without having to advertise an S-PMSI A-D route for each one.

Of course, a specific C-flow, (C-S,C-G), can always be assigned individually to a particular P-tunnel by advertising an S-PMSI A-D route whose NLRI contains (C-S,C-G).

In Section 4, 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 its normal BGP processing is not done, but that the route is not considered when determining which P-tunnel to use when receiving multicast data, and that the MPLS label values it conveys are not used. We will use "ignore" in quotes to indicate this meaning.

This document provides procedures only for the case where the P-tunnels are "unidirectional", i.e., point-to-multipoint. The use of "bidirectional" (multipoint-to-multipoint) P-tunnels is outside the scope of this document.

2. Encoding of Wildcards

Per [MVPN-BGP] Section 4.3, the MCAST-VPN NLRI in an S-PMSI A-D route is encoded as follows:

```

+-----+
|      RD      (8 octets)      |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (variable)      |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable)       |
+-----+
| Originating Router's IP Addr    |
+-----+

```

where the "source length" and "group length" fields always have a non-zero value. This document specifies that a "zero-length" source or group represents the corresponding wildcard. Specifically,

- A source wildcard is encoded as a zero-length source field. That is, the "multicast source length" field contains the value 0x00, and the "multicast source" field is omitted.

- A group wildcard is encoded as a zero-length group field. That is, the "multicast group length" field contains the value 0x00, and the "multicast group" field is omitted.

3. Finding the Matching S-PMSI A-D Route

This section gives the precise rules for determining the S-PMSI A-D route that is "matched" by a given (C-S,C-G) or (C-*,C-G). The procedures in Section 4 will make use of the matching rules defined in this section.

All matching rules assume the context of a given VRF at a given PE.

The rules that a PE applies to find the S-PMSI A-D route that matches a (C-S,C-G) C-flow that it needs to transmit are slightly different than the rules it applies to find the S-PMSI A-D route that matches a (C-S,C-G) C-flow that it needs to receive. These rules are specified in Sections 3.1 and 3.2, respectively.

The S-PMSI A-D route that is matched by a given (C-S,C-G) may change over time, as the result of S-PMSI A-D routes being withdrawn or as a result of new S-PMSI A-D routes being originated and/or advertised. In particular, if (C-S,C-G) matches an S-PMSI A-D route whose NLRI contains (C-*,C-*), the origination or reception of an S-PMSI A-D route whose NLRI contains (C-S,C-G) may cause (C-S,C-G) to match the latter route instead. Note also that the S-PMSI A-D route that matches a given (C-S,C-G) is independent of the order in which the routes were originated or received.

3.1. Finding the Match for Data Transmission

Consider a given PE; call it PE1. At any given time, for a given VRF at PE1, there is a (possibly empty) set of S-PMSI A-D routes that PE1 has originated and advertised, but not withdrawn. We will refer to these routes as "currently originated" by PE1. Suppose that PE1 needs to transmit a particular C-flow (C-S,C-G) to one or more other PEs. We use the following algorithm to find the S-PMSI A-D route that the C-flow "matches":

- If there is an S-PMSI A-D route currently originated by PE1, whose NLRI contains (C-S,C-G), the (C-S,C-G) C-flow matches that route.
- Otherwise, if there is an S-PMSI A-D route currently originated by PE1, whose NLRI contains (C-S,C-*), AND if C-G is an SSM group address, the (C-S,C-G) C-flow matches that route.

- Otherwise, if there is an S-PMSI A-D route currently originated by PE1, whose NLRI contains (C-*,C-G), AND if C-G is an ASM group address, the (C-S,C-G) C-flow matches that route.
- Otherwise, if there is an S-PMSI A-D route currently originated by PE1, whose NLRI contains (C-*,C-*), the (C-S,C-G) C-flow matches that route.

3.2. Finding the Match for Data Reception

We refer to an S-PMSI A-D route as being "installed" (in a given VRF) if it has been selected by the BGP decision process as the preferred route for its NLRI.

An S-PMSI A-D route is considered to be "originated by a given PE" if that PE's IP address is contained in the "Originating Router's IP Address" field in the MCAST-VPN NLRI of the route.

3.2.1. Finding the Match for (C-S,C-G)

Suppose that a PE router (call it PE1) needs to receive (C-S,C-G), and that PE1 has chosen another PE router (call it PE2) as the "upstream PE" [MVPN] for that flow.

- If there is an installed S-PMSI A-D route originated by PE2, whose NLRI contains (C-S,C-G), then (C-S,C-G) matches that route.
- Otherwise, if there is an installed S-PMSI A-D route originated by PE2, whose NLRI contains (C-S,C-*), AND if C-G is an SSM multicast group address, then (C-S,C-G) matches that route.
- Otherwise, if there is an installed S-PMSI A-D route originated by PE2, whose NLRI contains (C-*,C-G), AND if C-G is an ASM multicast group address, then (C-S,C-G) matches that route.
- Otherwise, if there is an installed S-PMSI A-D route originated by PE2, whose NLRI contains (C-*,C-*), then (C-S,C-G) matches that route.

3.2.2. Finding the Wildcard Match for (C-*,C-G)

Suppose that a PE router (call it PE1) needs to receive (C-*,C-G) traffic. Note that even if (C-*,C-G) matches a non-wildcard S-PMSI A-D route (as detailed in Section 12.3 of [MVPN-BGP]), it may also match one or more wildcard S-PMSI A-D routes, as specified below.

If on PE1 there is an installed S-PMSI A-D route originated by PE2, whose NLRI contains (C-*,C-G), then (C-*,C-G) matches this route if one of the following conditions holds:

- PE1 determines that PE2 is the "upstream" PE [MVPN] for the C-RP of C-G, or
- PE1 has installed one or more Source Active A-D routes for C-G originated by PE2, and for at least one of these routes, PE1 does not have a corresponding (C-S,C-G) state, or
- C-G is a BIDIR-PIM group, or
- Source Active A-D routes are not being used.

If (C-*,C-G) does not match a (C-*,C-G) S-PMSI A-D route from PE2, but PE1 has an installed (C-*,C-*) S-PMSI A-D route from PE2, then (C-*,C-G) matches the (C-*,C-*) route if one of the following conditions holds:

- PE1 determines that PE2 is the "upstream" PE [MVPN] for the C-RP of C-G, or
- PE1 has installed one or more Source Active A-D routes for C-G originated by PE2, and for at least one of these routes, PE1 does not have a corresponding (C-S,C-G) state, or
- C-G is a BIDIR-PIM group, or
- Source Active A-D routes are not being used.

4. Procedures for S-PMSI A-D Routes with Wildcards

4.1. Procedures for All Kinds of Wildcards

This document defines procedures for the following uses of the wildcard in the NLRI of an S-PMSI A-D route:

- (C-*,C-G): Source wildcard, group specified.
- (C-S,C-*): Source specified, group wildcard.
- (C-*,C-*): Source wildcard, group wildcard.

All other wildcard functionality is outside the scope of this document.

The ability to originate S-PMSI A-D routes with a particular kind of wildcard is OPTIONAL. However, if a PE has the ability to originate S-PMSI A-D routes with a particular kind of wildcard, it MUST have the ability to interpret and correctly process S-PMSI A-D routes with that kind of wildcard, and it SHOULD have the ability to interpret and correctly process all three kinds of wildcards.

For a given MVPN, A PE MUST NOT originate S-PMSI A-D routes with a particular kind of wildcard unless it is known a priori that all PEs attached to that MVPN have the ability to interpret and correctly process that kind of wildcard.

The criteria for originating and withdrawing S-PMSI A-D routes with wildcards are local to the originating PE.

As specified in [MVPN-BGP], an S-PMSI A-D route is carried in the NLRI field of an MP_REACH_NLRI attribute (see [BGP-MP]). Every S-PMSI A-D route has a particular address family (IPv4 or IPv6), as specified in the Address Family Identifier (AFI) field of the MP_REACH_NLRI attribute. A wildcard in a particular S-PMSI A-D route always refers only to multicast flows of that same address family.

The procedures specified in this document apply only when the PMSI Tunnel Attribute of an S-PMSI A-D route specifies a "unidirectional" P-tunnel. The use of "bidirectional" P-tunnels (e.g., Multipoint-to-Multipoint Label Switched Paths, BIDIR-PIM trees) is outside the scope of this document.

In the following sections, an S-PMSI A-D route whose NLRI contains (C-*,C-G), (C-S,C-*), or (C-*,C-*) will be referred to as a "(C-*,C-G) route", a "(C-S,C-*) route", or a "(C-*,C-*)" route, respectively.

4.2. Procedures for (C-*,C-G) S-PMSI A-D Routes

This document specifies the use of (C-*,C-G) S-PMSI A-D routes only in the case where C-G is an ASM group address. Use of (C-*,C-G) S-PMSI A-D routes where C-G is an SSM group address is outside the scope of this document. If a PE receives a (C-*,C-G) S-PMSI A-D route, and the PE can determine that C-G is an SSM group address, the PE SHOULD "ignore" this S-PMSI A-D route.

By default, the set of Route Targets carried by a (C-*,C-G) S-PMSI A-D route originated by a given VRF is the same as the set of Route Targets carried in the (unicast) VPN-IP routes that originated from that VRF. An implementation MUST allow the set of Route Targets

carried by the (C-*,C-G) S-PMSI A-D route to be specified by configuration. In the absence of a configured set of Route Targets, the route MUST carry the default set of Route Targets.

If a PE needs to transmit packets of a (C-S,C-G) C-flow, and if (C-S,C-G) matches a (C-*,C-G) S-PMSI A-D route according to the rules of Section 3.1, then the PE MUST use the P-tunnel advertised in this route for transmitting that C-flow. (Note that it is impossible for a given (C-S,C-G) to match both a (C-*,C-G) wildcard and a (C-S,C-*) wildcard.)

If PIM is being used as the PE-PE control protocol, then if the PE has (C-*,C-G) and/or (C-S,C-G) state that matches (according to the procedures of Section 3.2) an S-PMSI A-D route, the PE MUST join the P-tunnel specified in the PMSI Tunnel Attribute of that route.

If BGP is being used as the PE-PE control protocol, then

- If a given PE has currently originated a C-multicast Shared Tree Join for (C-*,C-G), and if (C-*,C-G) matches a (C-*,C-G) S-PMSI A-D route, then the PE applies the procedures of Section 12.3 ("Receiving S-PMSI A-D Routes by PEs") of [MVPN-BGP] to that S-PMSI A-D route.
- Otherwise (the given PE does not have a currently originated C-multicast Shared Tree Join for (C-*,C-G)), if there are one or more values of C-S for which the PE has a currently originated Source Tree Join C-multicast route for (C-S,C-G), the PE MUST join the tunnels advertised by the S-PMSI A-D routes that match (according to Section 3.2) each such (C-S,C-G).
- Otherwise, the PE "ignores" the route.

4.3. Procedures for (C-S,C-*) S-PMSI A-D Routes

This document covers the use of (C-S,C-*) S-PMSI A-D routes for only the C-multicast flows where C-G is an SSM group address. Use of (C-S,C-*) S-PMSI A-D routes for other C-multicast flows is outside the scope of this document. Specifically, if a PE receives a (C-S,C-*) S-PMSI A-D route, and the PE can determine that C-G is not an SSM group address, the PE SHOULD "ignore" this S-PMSI A-D route.

By default, the set of Route Targets carried by a (C-S,C-*) S-PMSI A-D route originated by a given VRF is an intersection between the set of Route Targets carried in the Intra-AS I-PMSI A-D route that originated from that VRF, and the set of Route Targets carried by the unicast VPN-IP route to C-S originated from that VRF. An implementation MUST allow the set of Route Targets carried by the

(C-S,C-*) S-PMSI A-D route to be specified by configuration. In the absence of a configured set of Route Targets, the route MUST carry the default set of Route Targets.

If a PE needs to transmit packets of a (C-S,C-G) C-flow, and if (C-S,C-G) matches a (C-S,C-*) S-PMSI A-D route according to the rules of Section 3.1, then the PE MUST use the P-tunnel advertised in this route for transmitting that C-flow. (Note that it is impossible for a given (C-S,C-G) to match both a (C-*,C-G) wildcard and a (C-S,C-*) wildcard.)

If PIM is being used as the PE-PE control protocol for distributing C-multicast routing, and if a given PE needs to receive a (C-S,C-G) flow, and if (C-S,C-G) matches the (C-S,C-*) S-PMSI A-D route (according to the procedures of Section 3.2), then the PE MUST join the P-tunnel specified in the PMSI Tunnel Attribute of that route.

If BGP is being used as the PE-PE control protocol for distributing C-multicast routing, and if there is some (C-S,C-G) such that (a) the PE has a currently originated (C-S,C-G) Source Tree Join C-multicast route, AND (b) the given (C-S,C-G) matches (according to the procedures of Section 3.2) a (C-S,C-*) S-PMSI A-D route, then PE1 applies the procedures of Section 12.3 ("Receiving S-PMSI A-D Routes by PEs") of [MVPN-BGP] to the matching S-PMSI A-D route.

4.4. Procedures for (C-*,C-*) S-PMSI A-D Routes

(C-*,C-*) S-PMSI A-D routes are used when, for a given MVPN, a PE has a policy not to use an I-PMSI for carrying multicast data traffic originated in the MVPN's site(s) connected to that PE. When the (C-*,C-*) wildcard is used together with BGP C-multicast routing, this results in the "S-PMSI only" model, where no I-PMSIs are used at all for the given MVPN.

A (C-*,C-*) S-PMSI A-D route is originated for a given MVPN by a given PE only if that PE has been provisioned with the policy to do so.

When so provisioned, the PE MAY originate the (C-*,C-*) S-PMSI A-D route as soon as it is enabled to support the given MVPN. Alternatively, the PE MAY delay originating the route until one of the following conditions holds:

- The PE-PE protocol for distributing C-multicast routing is PIM, and for the given MVPN, the PE has some (C-S,C-G) or (C-*,C-G) state for which the upstream interface is one of the VRF interfaces for the given MVPN.

- The PE-PE protocol for distributing C-multicast routing is BGP, and the given PE has received and installed either of the following:
 - * a Source Tree Join C-multicast route, with the C-S contained in the route's NLRI being reachable via one of the given MVPN's VRF interfaces, or
 - * a Shared Tree Join C-multicast route, with the C-RP carried in that route being reachable via one of the given MVPN's VRF interfaces.

By default, the set of Route Targets carried by a (C-*,C-*) S-PMSI A-D route originated from a given VRF is the same as the set of Route Targets carried in the VPN-IP unicast routes originated from that VRF. An implementation MUST allow the set of Route Targets carried by the (C-*,C-*) S-PMSI A-D route to be specified by configuration. In the absence of a configured set of Route Targets, the route MUST carry the default set of Route Targets, as specified above.

If a PE needs to transmit packets of a (C-S,C-G) C-flow, and if (C-S,C-G) matches a (C-*,C-*) S-PMSI A-D route according to the rules of Section 3.1, then the PE MUST use the P-tunnel advertised in this route for transmitting that C-flow. (Note that it is impossible for a given (C-S,C-G) to match both a (C-*,C-*) wildcard and any other wildcard.)

If PIM is being used as the PE-PE control protocol for distributing C-multicast routing, and if a given PE, say PE1, needs to receive a (C-S,C-G) flow, and if (C-S,C-G) matches the (C-*,C-*) S-PMSI A-D route (according to the procedures of Section 3.2), then PE1 MUST join the P-tunnel specified in the PMSI Tunnel Attribute of that route.

If BGP is being used as the PE-PE control protocol for distributing C-multicast routing, then if (and only if) one of the following conditions holds, the PE applies the procedures of Section 12.3 ("Receiving S-PMSI A-D Routes by PEs") of [MVPN-BGP] to the matching S-PMSI A-D route. The conditions are as follows:

- The PE has a currently originated C-multicast Source Tree Join route for (C-S,C-G) that matches (according to the procedures of Section 3.2) the (C-*,C-*) S-PMSI A-D route, or
- The PE has a currently originated a C-multicast Shared Tree Join route for (C-*,C-G) that matches (according to the procedures of Section 3.2) the (C-*,C-*) S-PMSI A-D route.

5. Security Considerations

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

6. Acknowledgments

The authors wish to thank Arjen Boers, Dongling Duan, Apoorva Karan, Thomas Morin, Keyur Patel, Karthik Subramanian, and Kurt Windisch for many helpful discussions.

7. Normative References

- [BGP-MP] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [L3VPN] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [MVPN] Rosen, E., Ed., and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, February 2012.
- [MVPN-BGP] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, February 2012.
- [PIM] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Authors' Addresses

Rahul Aggarwal
Arktan

E-Mail: raggarwa_1@yahoo.com

Yiqun Cai
Microsoft
1065 La Avenida
Mountain View, CA 94043

E-Mail: yiqunc@microsoft.com

Wim Henderickx
Alcatel-Lucent

E-Mail: wim.henderickx@alcatel-lucent.be

Praveen Muley
Alcatel-Lucent

E-Mail: Praveen.Muley@alcatel-lucent.com

Ray (Lei) Qiu
2330 Central Expressway
Santa Clara, CA 95050
USA

E-Mail: rayq@huawei.com

Yakov Rekhter (editor)
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089

E-Mail: yakov@juniper.net

Eric C. Rosen (editor)
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

