

Internet Engineering Task Force (IETF)
Request for Comments: 6515
Updates: 6514
Category: Standard Track
ISSN: 2070-1721

R. Aggarwal
Juniper Networks, Inc.
E. Rosen
Cisco Systems, Inc.
February 2012

IPv4 and IPv6 Infrastructure Addresses in
BGP Updates for Multicast VPN

Abstract

To provide Multicast VPN (MVPN) service, Provider Edge routers originate BGP Update messages that carry Multicast-VPN ("MCAST-VPN") BGP routes; they also originate unicast VPN routes that carry MVPN-specific attributes. These routes encode addresses from the customer's address space, as well as addresses from the provider's address space. These two address spaces are independent, and the address family (IPv4 or IPv6) of the two spaces may or may not be the same. These routes always contain an "address family" field that specifies whether the customer addresses are IPv4 addresses or whether they are IPv6 addresses. However, there is no field that explicitly specifies the address family of the provider addresses. To ensure interoperability, this document specifies that provider IPv4 addresses are always encoded in these update messages as 4-octet addresses, and that the distinction between IPv4 and IPv6 is signaled solely by the length of the address field. Specific cases are explained in detail. This document updates RFC 6514.

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/rfc6515>.

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. Introduction2
1.1. IPv4 and IPv6 Addresses in MCAST-VPN Routes2
1.2. Specification of Requirements4
1.3. Acronyms Used in This Document4
2. PE Addresses in MCAST-VPN Routes4
3. VRF Route Import Extended Community5
4. PMSI Tunnel Attributes in I-PMSI A-D Routes6
4.1. Relationship to AFI Value6
4.2. Relationship to Next Hop Address Family6
5. IANA Considerations7
6. Security Considerations7
7. Acknowledgments7
8. Normative References7
9. Informative References7

1. Introduction

1.1. IPv4 and IPv6 Addresses in MCAST-VPN Routes

[MVPN-BGP] defines a new set of BGP route types that are used by service providers (SPs) to provide Multicast Virtual Private Network service to their customers. These routes have a newly defined BGP NLRI, the "MCAST-VPN" NLRI. The MCAST-VPN NLRI is carried in the NLRI field of the MP_REACH_NLRI/MP_UNREACH_NLRI attributes defined in [BGP-MP]. The SAFI field of the MP_REACH_NLRI/MP_UNREACH_NLRI attribute is used to identify the NLRI as being an MCAST-VPN NLRI.

When the SAFI field of an MP_REACH_NLRI/MP_UNREACH_NLRI attribute has the "MCAST-VPN" value, the AFI field has two defined values: 1 and 2. AFI 1 indicates that any customer multicast addresses occurring in the MP_REACH_NLRI/MP_UNREACH_NLRI attribute are IPv4 addresses; AFI 2 indicates that such addresses are IPv6 addresses.

However, some of the MCAST-VPN routes also contain addresses of Provider Edge (PE) routers in the SP network. An SP with an IPv4 network may provide MVPN service for customers using IPv6, and an SP with an IPv6 network may provide MVPN service for customers that use IPv4. Therefore, the address family of the PE addresses MUST NOT be inferred from the AFI field of the associated MP_REACH_NLRI/MP_UNREACH_NLRI attribute.

The purpose of this document is to make clear that whenever a PE address occurs in an MCAST-VPN route (whether in the NLRI or in an attribute), the IP address family of that address is determined by the length of the address (a length of 4 octets for IPv4 addresses, a length of 16 octets for IPv6 addresses), NOT by the AFI field of the route.

In particular, if a SP with an IPv4 core network is providing MVPN/IPv6 service to a customer, the PE addresses in the MCAST-VPN routes will be 4-octet IPv4 routes, even though the AFI of those routes will have the value 2.

Some previous specifications (e.g., [RFC4659] and [RFC4798]) have taken a different approach, requiring that in any routes containing IPv6 or VPN-IPv6 customer addresses, the IPv4 PE addresses be represented as IPv6-mapped IPv4 addresses [RFC4291]. This document does not use that approach. Rather, this specification uses the approach adopted in [RFC4684] and [RFC5549]. The MCAST-VPN routes contain enough information to enable the IP address family of the PE addresses to be inferred from the address lengths.

[MVPN-BGP] also defines an attribute, the "VRF Route Import Extended Community", that is attached to unicast VPN-IPv4 or VPN-IPv6 routes. This extended community contains a PE address, and this document specifies how to encode an IPv6 address in this attribute, independent of whether the attribute is attached to a VPN-IPv4 route or a VPN-IPv6 route.

This document also clarifies an issue with respect to the significance of the Address Family field of an Intra-AS I-PMSI A-D route that carries a PMSI Tunnel Attribute.

1.2. 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].

1.3. Acronyms Used in This Document

This document uses a number of acronyms, mostly taken directly from the BGP and VPN specifications.

- A-D Route: Auto-Discovery Route [MVPN]
- AFI: Address Family Identifier [BGP-MP]
- AS: Autonomous System [BGP]
- I-PMSI: Inclusive PMSI [RFC4364]
- MVPN: Multicast Virtual Private Network [MVPN]
- MCAST-VPN routes: BGP routes of "MCAST-VPN" Subsequent Address Family, as defined in [MVPN-BGP]. The NLRI of such routes may be referred to as MCAST-VPN NLRI.
- MP_REACH_NLRI: Multiprotocol Reachable NLRI [BGP-MP]
- MP_UNREACH_NLRI: Multiprotocol Unreachable NLRI [BGP-MP]
- PMSI: Provider Multicast Service Interface [MVPN]
- NLRI: Network Layer Reachability Information [BGP]
- PE: Provider Edge [RFC4364]
- S-PMSI: Selective PMSI [RFC4364]
- SAFI: Subsequent Address Field Identifier [BGP-MP]
- SP: Service Provider

2. PE Addresses in MCAST-VPN Routes

PE addresses occur in MCAST-VPN routes in the following places:

1. "Network Address of Next Hop" field in the MP_REACH_NLRI attribute, as defined in Section 3 of [BGP-MP]. This field is preceded by a "length of next hop address" field. Hence, it is

always clear whether the address is an IPv4 address (length is 4) or an IPv6 address (length is 16). If the length of the next hop address is neither 4 nor 16, the MP_REACH_NLRI attribute MUST be considered to be "incorrect", and MUST be handled as specified in Section 7 of [BGP-MP].

2. "Intra-AS I-PMSI A-D route", defined in Section 4.1 of [MVPN-BGP]. All MCAST-VPN routes begin with a 1-octet route type field, followed by a 1-octet "NLRI length" field. In the Intra-AS I-PMSI A-D route, the length is followed by an 8-octet Route Distinguisher (RD), which is then followed by the "Originating Router's IP Address" field. The length of this field (4 octets for IPv4 or 16 octets for IPv6) can thus be inferred from the NLRI length field (which will be either 12 or 24, respectively). If the inferred length of the "Originating Router's IP Address" field is neither 4 nor 16, the MP_REACH_NLRI attribute MUST be considered to be "incorrect", and MUST be handled as specified in Section 7 of [BGP-MP].
3. "S-PMSI A-D Route", defined in Section 4.3 of [MVPN-BGP]. In this route, the "NLRI length" field is followed by an 8-octet RD, a variable-length "multicast source" field, a variable-length "multicast group" field, and an "Originating Router's IP Address" field. The two variable-length fields have their own length fields. From these two length fields and the NLRI length field, one can compute the length of the "Originating Router's IP Address" field, which again is either 4 for IPv4 or 16 for IPv6. If the computed length of the "Originating Router's IP Address" field is neither 4 nor 16, the MP_REACH_NLRI attribute MUST be considered to be "incorrect", and MUST be handled as specified in Section 7 of [BGP-MP].
4. "Leaf A-D Route", defined in Section 4.4 of [MVPN-BGP]. In this route, the "NLRI length" field is followed by a variable-length "route key", which is followed by the "Originating Router's IP Address" field. The Route Key has its own length field. From the NLRI length and the route key length, one can compute the length of the "Originating Router's IP Address" field. If the computed length of the "Originating Router's IP Address" field is neither 4 nor 16, the MP_REACH_NLRI attribute MUST be considered to be "incorrect", and MUST be handled as specified in Section 7 of [BGP-MP].

3. VRF Route Import Extended Community

The "VRF Route Import Extended Community", specified in [MVPN-BGP], is an attribute carried by unicast VPN-IPv4 or VPN-IPv6 routes. It is an "IPv4 Address Specific Extended Community" of type "VRF Route

Import"; hence, it can only carry an IPv4 address. To carry an IPv6 address, an "IPv6 Address Specific Extended Community" [RFC5701], of type "VRF Route Import", must be used. A code point for this type of extended community has been allocated by IANA.

4. PMSI Tunnel Attributes in I-PMSI A-D Routes

When a PMSI Tunnel Attribute occurs in an I-PMSI A-D route originated by a particular PE or Autonomous System Border Router (ASBR), it identifies a tunnel that the PE/ASBR uses by default for carrying the multicast traffic of a particular customer MVPN. The proper encoding and interpretation of the PMSI Tunnel attribute is affected by both the AFI and "Network Address of Next Hop" fields.

4.1. Relationship to AFI Value

When the PMSI Tunnel Attribute occurs in a BGP Update message with a MP_REACH_NLRI attribute whose AFI is 1, the meaning is that the identified tunnel is used by default to carry IPv4 MVPN traffic for a particular customer MVPN. When the PMSI Tunnel Attribute occurs in a BGP Update message with a MP_REACH_NLRI attribute whose AFI is 2, the meaning is that the identified tunnel is used by default to carry IPv6 MVPN traffic for a particular customer MVPN. To assign both IPv4 and IPv6 MVPN traffic to an I-PMSI tunnel, two I-PMSI A-D routes MUST be used -- one whose MP_REACH_NLRI has an AFI of 1 and one whose MP_REACH_NLRI has an AFI of 2. To use the same tunnel for both IPv4 and IPv6 traffic, the same value of the PMSI Tunnel attribute can be used in each route.

4.2. Relationship to Next Hop Address Family

If the "Network Address of Next Hop" field in the MP_REACH_NLRI attribute contains an IPv4 address, then any IP addresses appearing in the "Tunnel Identifier" field of the PMSI Tunnel Attribute MUST be IPv4 addresses.

If the "Network Address of Next Hop" field in the MP_REACH_NLRI attribute contains an IPv6 address, then any IP addresses appearing in the "Tunnel Identifier" field of the PMSI Tunnel Attribute MUST be IPv6 addresses.

If these conditions are not met, the PMSI Tunnel Attribute MUST be handled as a "malformed" PMSI Tunnel Attribute, as specified in Section 5 of [MVPN-BGP].

5. IANA Considerations

IANA has assigned the code point 0x000b for "VRF Route Import" in the "IPv6 Address Specific Extended Community" registry in the "transitive communities" portion of the namespace. The references are to this document and to [MVPN-BGP].

6. Security Considerations

This document does not raise any security considerations beyond those raised by [MVPN-BGP].

7. Acknowledgments

The authors wish to thank Dongling Duan, Keyur Patel, Yakov Rekhter, and Karthik Subramanian.

8. Normative References

- [BGP] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [BGP-MP] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [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.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9. Informative References

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, September 2006.

- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, February 2007.
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, November 2006.
- [RFC5549] Le Faucheur, F. and E. Rosen, "Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop", RFC 5549, May 2009.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, November 2009.

Authors' Addresses

Rahul Aggarwal
Juniper Networks
1194 North Mathilda Avenue
Sunnyvale, CA 94089
EMail: raggarwa_1@yahoo.com

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