Workgroup: Network Working Group
Internet-Draft:
draft-duan-bess-mvpn-ipv6-infras-01
Updates: 6514 (if approved)
Published: 20 May 2022
Intended Status: Standards Track
Expires: 21 November 2022
Authors: F. Duan
Huawei Technologies
J. Xie
Huawei Technologies

BGP MVPN in IPv6 Infrastructure Networks: Problems and Solution Approaches

Abstract

MVPN deployment faces some problems while used in provider's IPv6 infrastructure networks. This document describes these problems, and the solutions to solve these problems.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 21 November 2022.

Copyright Notice

Table of Contents

1. Introduction

BGP MVPN procedure is defined in [RFC6514]. As a mature MVPN technology, it has been accepted by most operators and vendors. In [RFC6515], BGP MVPN is updated for IPv6 infrastructure networks. However, the deployment of BGP MVPN in IPv6 network still faces some problems. This document describes these problems, and the solutions to solve these problems.

2. Terminology

Readers of this document are assumed to be familiar with the terminology and concepts of the documents listed as Normative References.

3. Problems and Solutions

3.1. Problems

In [RFC6514], the following processes are not suitable if IPv6 infrastructure addresses are used in provider's networks:

1. In paragraph 7 of section 11.1.3, it described that "To support non-segmented inter-AS tunnels, the Source AS field in the C-multicast route is set to value of the Originating Router's IP Address field of the found Intra-AS I-PMSI A-D route". In NLRI

of C-multicast route, Source AS field is 4 bytes length, while the Originating Router's IP Address field of Intra-AS I-PMSI A-D route is 16 bytes length in provider's IPv6 networks.

2. In paragraph 2 of section 11.2, it described that "To support non-segmented inter-AS tunnels, instead of matching the RD and Source AS carried in the C-multicast route against the RD and Source AS of an Inter-AS I-PMSI A-D route, the ASBR should match it against the RD and the Originating Router's IP Address of the Intra-AS I-PMSI A-D routes". Source AS field in NLRI of C-multicast route cannot be translated to the Originating Router's IP Address of the Intra-AS I-PMSI A-D routes in provider's IPv6 networks, because of the mismatch of their field length.

In [RFC7716], zero RD is introduced in BGP MVPN NLRIs to enable Global Table Multicast service in provider's networks. In IPv6 infrastructure networks, Leaf PEs cannot send two distinct C-multicast route to two individual upstream root PEs for selctive forwarding, because the RD of the two roots is the same.

In the process of evolution to IPv6, IPv4 and IPv6 infrastructure addresses will coexist in the provider's network. The following figure is an example of BGP MVPN evolution to IPv6.

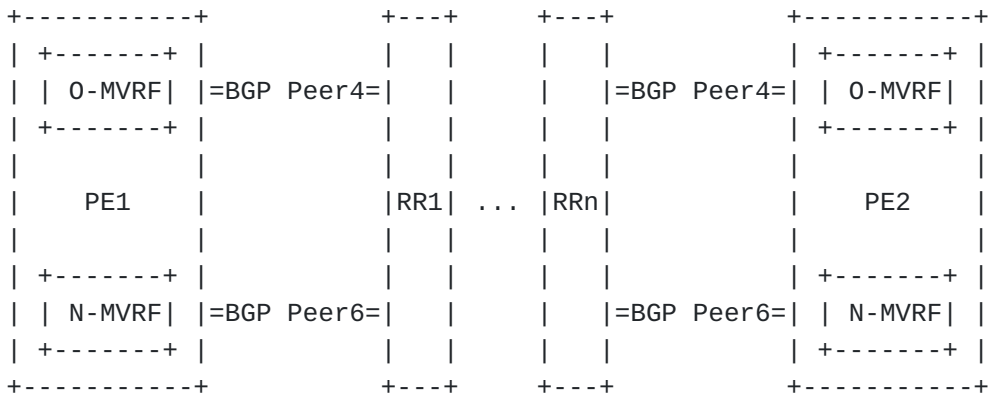

Figure 1: BGP MVPN Evolution to IPv6 Infrastructure

During the evolution process, IPv4 and IPv6 parallel BGP sessions are established between Provider Edge routers and Reflector routers, if the BGP MVPN route send to all IPv4 and IPv6 BGP peer without any control, the number of the PATHs of these routes will be doubled with each reflection while BGP ADD-PATH [RFC7911] procedure is enabled on Reflector routers.

3.2. Modification of C-Multicast route NLRI

To support non-segmented inter-AS tunnels in IPv6 infrastructure network, the C-Multicast route NLRI is redefined as following:

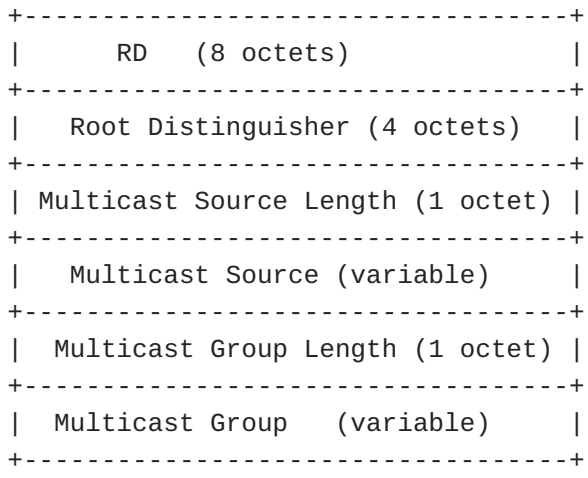

In the above figure, the Root Distinguisher field replaces the Source As field defined in [RFC6514]. When constructing a C-Multicast route, leaf PE follows the following specification:

1. If the Originating Router's IP Address field of the found Intra-AS I-PMSI A-D route is an IPv4 address, the Root Distinguisher field MUST be treated as Source AS field and section 11.1.3 of [RFC6514] MUST be followed.

2. If the Originating Router's IP Address field of the found Intra-AS I-PMSI A-D route is an IPv6 address and the root PE and leaf PE are in the same AS, the Root Distinguisher field MUST be treated as Source AS field and section 11.1.3 of [RFC6514] MUST be followed.

3. If the Originating Router's IP Address field of the found Intra-AS I-PMSI A-D route is an IPv6 address and the root PE and leaf PE are in the different ASs, a four bytes distinct value MUST be assigned by leaf PE for each root PE, the Root Distinguisher field in C-Multicast NLRI is filled with this value and a distinct C-multicast route will be send to individual upstream root PE.

When receiving a C-Multicast route from E-BGP neighbors, the ASBR checks whether an IPv6 VRF Route Import Extended Community is included in this route and takes following actions:

1. If the IPv6 VRF Route Import Extended Community does not exist in the C-Multicast route, the ASBR treats the Root Distinguisher field as Source AS field and follows the description in section 11.2 of [RFC6514].

2. If the IPv6 VRF Route Import Extended Community does exist in the C-Multicast route, the ASBR match the IPv6 address carried in this extended community and the RD in C-Multicast route NLRI against the Originating Router's IP Address and the RD of the Intra-AS I-PMSI A-D routes. If the corresponding Intra-AS I-PMSI A-D route exists, the ASBR propagates the C-Multicast route in its local AS.

3.3. Route reflection control

To reduce BGP MVPN routes in Parallel IPv4 and IPv6 BGP sessions scenario, the following actions should be taken by sender PEs:

1. For Intra-AS I-PMSI A-D Route, S-PMSI A-D Route and Leaf A-D Route, if the Originating Router's IP Address field in the route is filled with an IPv6 address, it is sent to the IPv6 BGP neighbors; otherwise, it is sent to the IPv4 BGP neighbors.

2. For Inter-AS I-PMSI A-D Route and Source Active A-D Route, it is sent to both IPv6 BGP neighbors and IPv4 BGP neighbors.

3. For C-Multicast Route, if the IPv6 VRF Route Import Extended Community exists in the route, it is sent to the IPv6 BGP neighbors; otherwise, it is sent to the IPv4 BGP neighbors.

In the reflector routers, the part of routes which are received from IPv6 BGP neighbors will be reflected to other IPv6 BGP neighbors and the other part of routes which are received from IPv4 BGP neighbors will be reflected to other IPv4 BGP neighbors.

4. Security Considerations

This document introduces no new security considerations beyond those already specified in [RFC6514] and [RFC6515].

5. IANA Considerations

This document contains no actions for IANA.

6. Acknowledgements

Your name here

7. 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>.

[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>.

[RFC6515] Aggarwal, R. and E. Rosen, "IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPN", RFC 6515, DOI 10.17487/RFC6515, February 2012, <https://www.rfc-editor.org/info/rfc6515>.

[RFC7716] Zhang, J., Giuliano, L., Rosen, E., Ed., Subramanian, K., and D. Pacella, "Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures", RFC 7716, DOI 10.17487/RFC7716, December 2015, <https://www.rfc-editor.org/info/rfc7716>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <https://www.rfc-editor.org/info/rfc7911>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Authors' Addresses

Fanghong Duan
Huawei Technologies

Email: duanfanghong@huawei.com

Jingrong Xie
Huawei Technologies

Email: xiejingrong@huawei.com