

MVPN IPv6-Infras

draft-duan-bess-mvpn-ipv6-infras-00

Fanghong Duan @Huawei

Jingrong Xie @Huawei

IETF 113

Mar. 2022


BGP MVPN gaps for provider's IPv6 network

- Non-segmented inter-AS tunnel establishment procedure which defined in [RFC6514] faces following problems.
 - 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.
 - 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.

Proposed solutions

- Modification of C-Multicast route NLRI

```
+-----+
|      RD      (8 octets) |
+-----+
| Source AS (4 octets) |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (variable) |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable) |
+-----+
```



```
+-----+
|      RD      (8 octets) |
+-----+
| Root Distinguisher (4 octets) |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (variable) |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable) |
+-----+
```

In IPv6 underlay networks, “Root Distinguisher” field is an root index locally assigned by the originator of the C-Multicast route

Proposed solutions (continue)

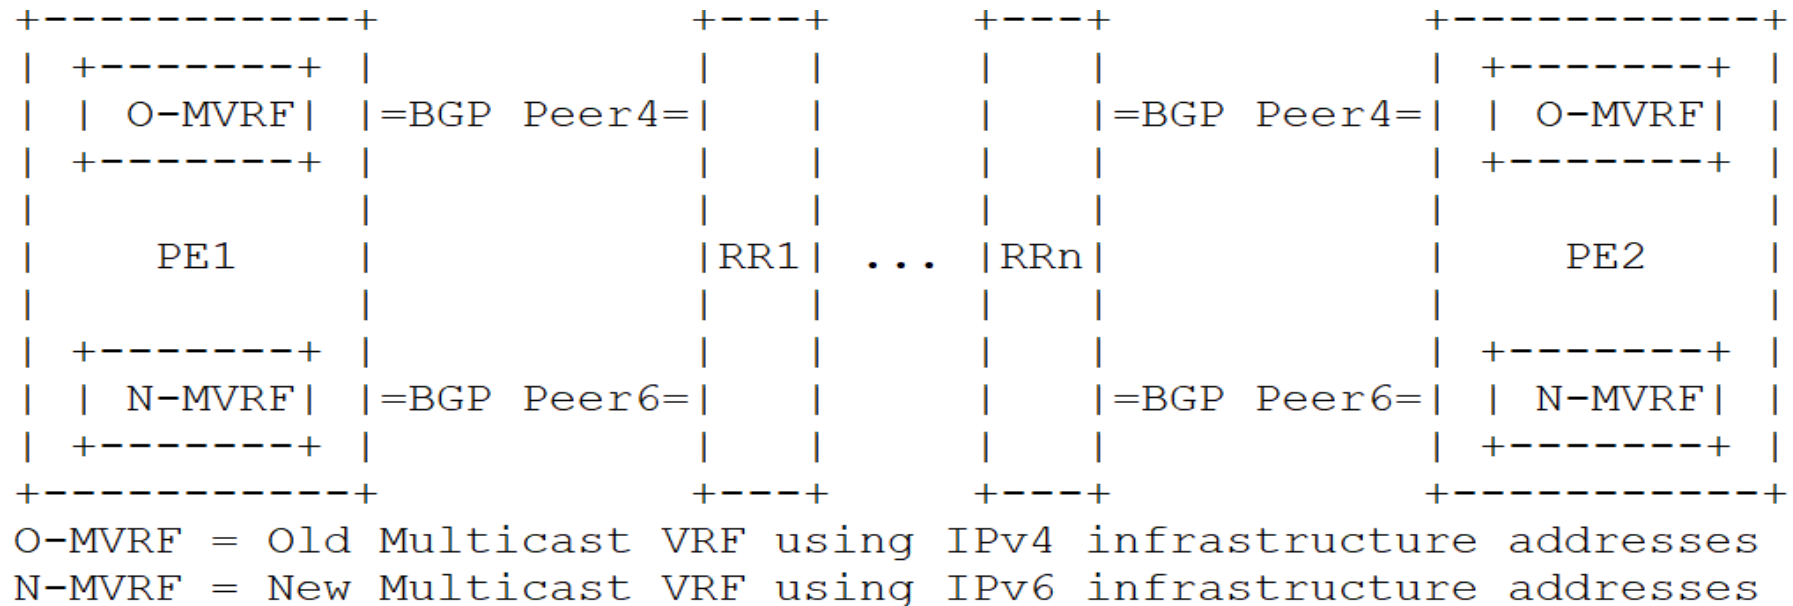
- Modification of the procedure for the originators of C-Multicast route .
 - If the Originating Router's IP Address field of the found Intra-AS I-PMSI A-D route is an IPv4 address, the Root Distinguisher field MUST be treated as Source AS field and section 11.1.3 of [RFC6514] MUST be followed.
 - If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

Proposed solutions (continue)

- Modification of the procedure for ASBR. When receiving a C-Multicast route from E-BGP peers, the ASBR checks whether an IPv6 VRF Route Import Extended Community is include in this route and takes following actions.
 - 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].
 - 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.

Problems for evolutions to IPv6 underlay

- 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.



- IPv4 and IPv6 parallel BGP sessions are established, the number of the PATHs of these routes will be doubled with each reflection.

Proposed solutions

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

Next Steps

- Comments and discussion.
- Call for workgroup adoption.

Thanks