

Multicast VPN Upstream Designated Forwarder Selection

draft-wang-bess-mvpn-upstream-df-selection-05

Fanghong Duan@Huawei

Siyu Chen@Huawei

IETF 116

Mar. 2023

Background

IETF113 to 115, to combine advantages of both hot and warm root standby, we proposed:

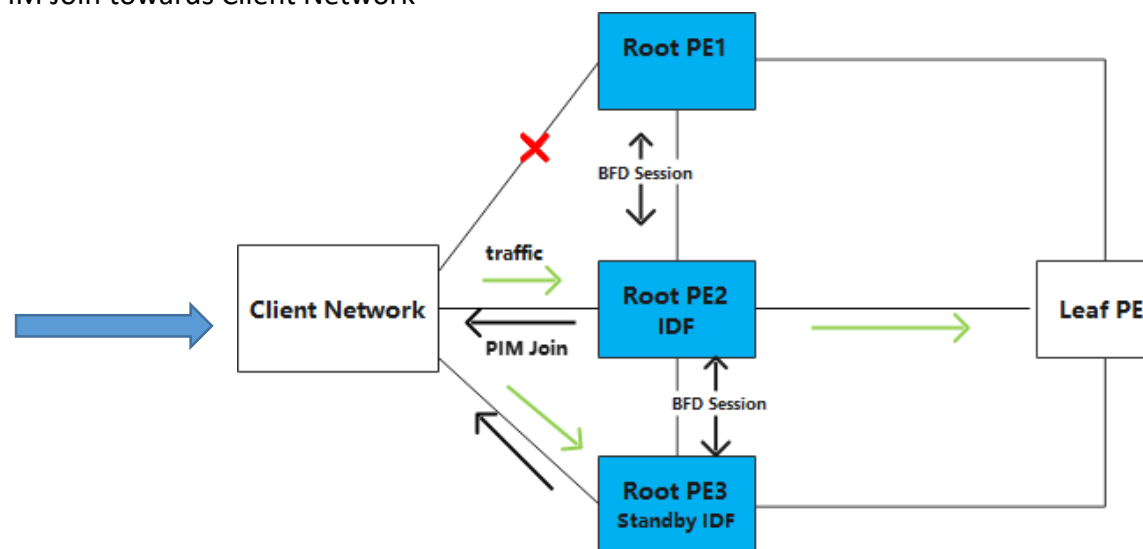
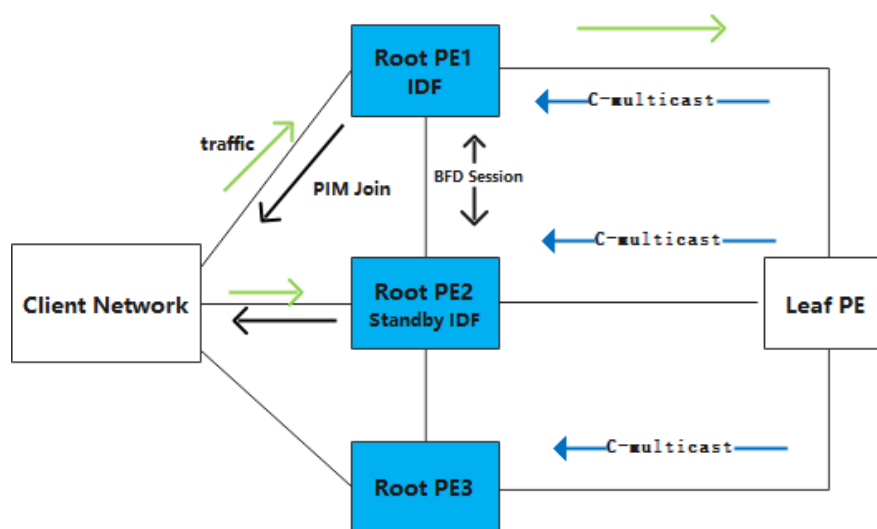
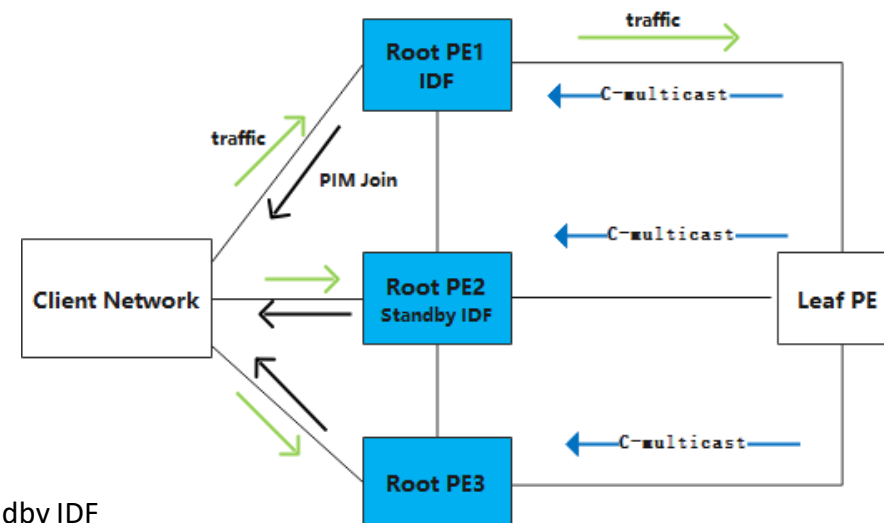
- MVPN Extensions: IDF negotiation Community and BFD Discriminator Attribute
- IDF Election Procedure:
 - Primary and Standby DF elected in ingress site
 - RPF Checklists instead of RPF interface in egress site for fast failover
 - Failure detection with BFD session between primary and standby PEs

This time:

- Updates on position where PIM Join sent towards
- New format of UMH route and “Root Distinguisher Extended Community”
- Explicit tracking in inter-As segmented scenario
- Backward compatibility

Updates on Ingress PE

- **Previously,**
 - C-multicast routes sent towards all root PEs
 - Traffic received on all root PEs
- **This time:**
 - Egress PE still sends C-multicast route towards all ingress PE
 - Send PIM Join towards Client Network only if current PE is primary or standby IDF
 - When primary IDF is down, new standby IDF need to send PIM Join towards Client Network



New UMH Route Format

- **UMH Route:**

```
+-----+
| Key: RD (8 octets) |
+-----+
| IP Prefix Length (1 octet, 0 to 32 / 128) |
+-----+
| Key:IP Prefix (4 / 16 octets ) |
+-----+
```

- **RD may not carried or not distinct:**

- In [RFC 7716], RD is not carried in GTM scenario
- In certain scenario, same RD is configured by different PEs

- **Same RD and same Multicast Source -----> Same UMH Route**

- **To distinguish different UMH routes, new UMH route format proposed :**

```
+-----+
| Key: RD (8 octets) |
+-----+
| IP Prefix Length (1 octet, 0 to 32 / 128) |
+-----+
| Key: IP Prefix (4 / 16 octets ) |
+-----+
| Key: Originating Router's IP Addr (4/16 octets) |
+-----+
```

- Multi-homed PEs exchange new UMH routes with each other
- High priority than existing UMH route

Root Distinguisher Extended Community

- **Non-segmented Inter-AS P-Tunnel over IPv6 infrastructure scenarios**

As for C-multicast Route:

4-octet Source AS cannot hold an IPv6 address ---> Hard to distinguish two C-Multicast routes

- **Solution: Root Distinguisher Extended Community**

Global Administrator Field	Local Administrator Field
4-octet Global Unique Value	0

Carried by Intra-AS A-D or wildcard S-PMSI A-D routes

Fill Source AS with Root Distinguishing value of root PE

Root PE Explicit Tracking in Segmented Inter-AS Scenario

- **Segmented Inter-AS Scenario:**

- Different Intra-AS A-D routes aggregated and Originator's IP discarded

- Leaf PE in downstream AS cannot distinguish multicast traffic from different root PEs

- **To perform root PE explicit tracking:**

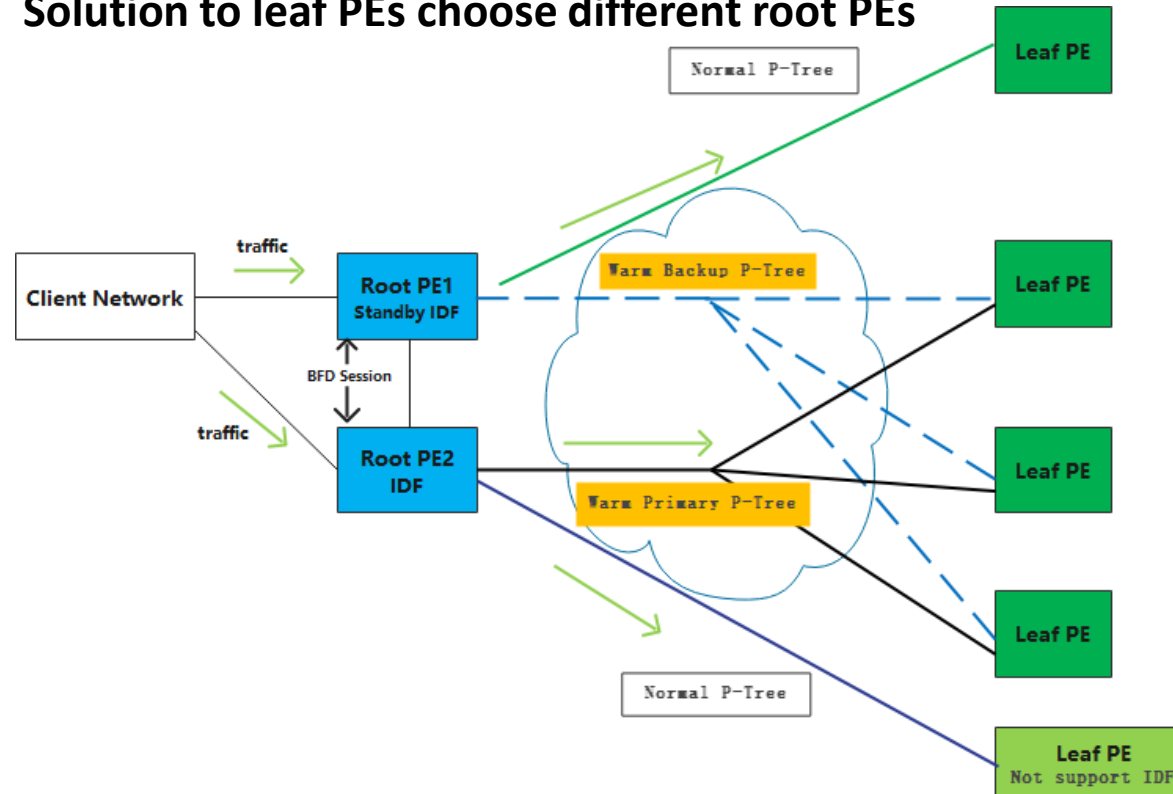
- Solution A : Wildcard S-PMSI AD route in [RFC6625]

- Solution B : PE Distinguisher Labels Attribute in [RFC6514]

- Distributed in Inter-AS A-D route

Backward Compatibility

- **Procedures for**
 - Root Not Support IDF Election
 - Leaf Not Support IDF Election
- **Solution to leaf PEs choose different root PEs**



Next Steps

- BESS WG Reviews and Comments
- Call for Adoption

Thanks