Inter-AS MVPN: Multihoming Considerations

Zhaohui Zhang, IETF104 Prague

draft-zzhang-bess-mvpn-evpn-cmcast-enhancements-01

Summary

- draft-zzhang-bess-mvpn-evpn-cmcast-enhancements-00 covers quite a few clarifications/enhancements/fixes:
  - MVPN C-Bidir Support with VPN Backbone being RPL
  - Inter-AS Propagation of MVPN C-Multicast Routes
  - Provider Tunnel Segmentation with Explicit-Tracking C-Multicast Routes
- -01 adds a new scenario:
  - MVPN multi-homing with Inter-AS

Scenario

- Inter-AS segmentation & per-AS aggregation:
  - Inter-AS I-PMSI A-D (type-2) routes advertised by ASBRs
  - A source is multi-homed to PE1/PE2 in the same AS100
  - Egress PE3/PE4 in remote AS300/400
- PE3 chooses PE1 as upstream PE for the source
  - C-multicast route for (src,grp):
    - RD of ASBR1's Inter-AS I-PMSI A-D route
    - RT that identifies PE1's VRF
- PE4 chooses PE2 as upstream PE for the source
  - C-multicast route for (src,grp):
    - RD of ASBR1's Inter-AS I-PMSI A-D route
    - RT that identifies PE2's VRF
- ASBR22 only re-advertises either PE3 or PE4's C-multicast route towards source AS
  - Because of the same NLRI
  - Targeted to either PE1 or PE2 but not both
    - As a result, only one will transmit packets

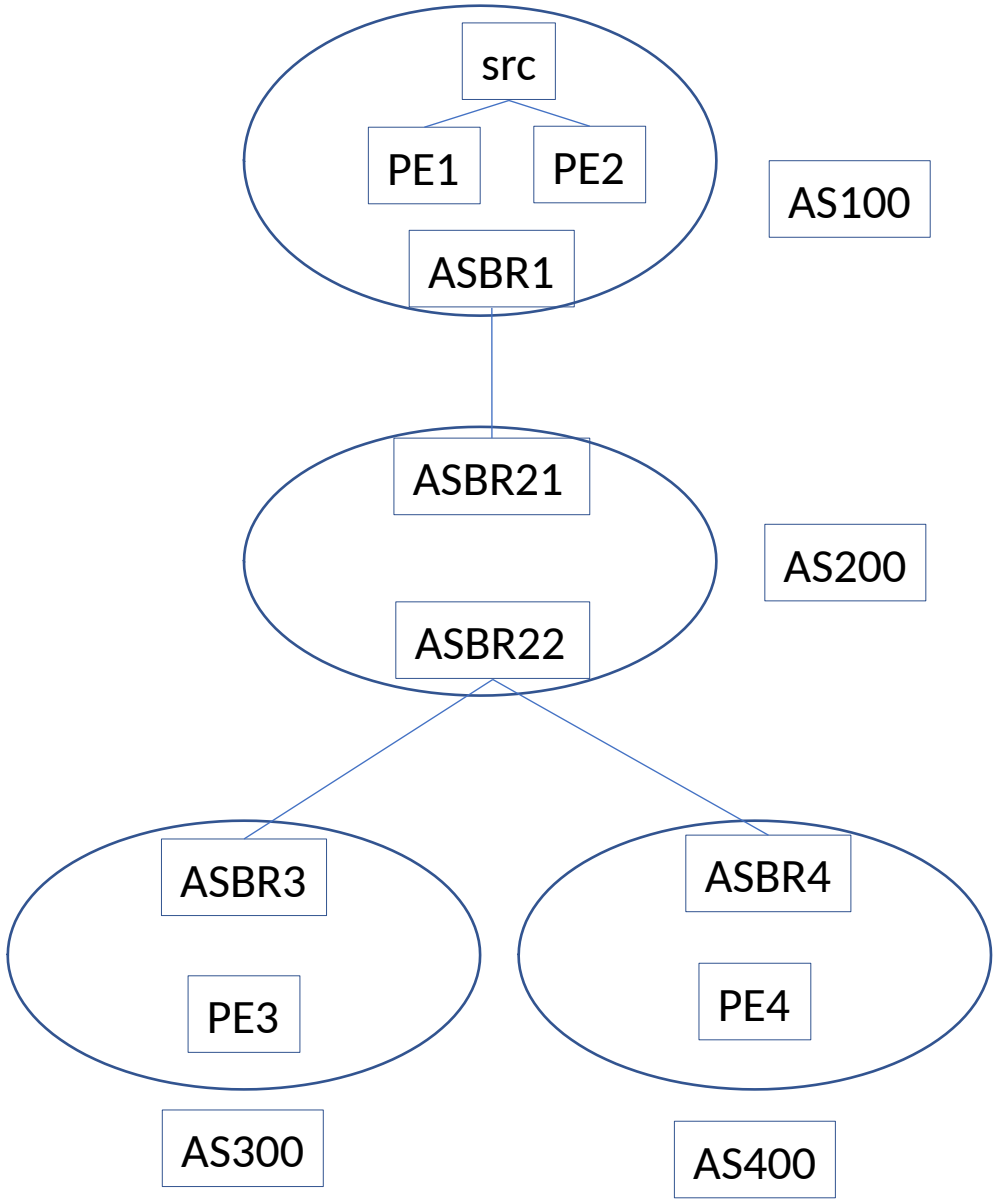

Problem 1 & Solution

- If selective tunnels are used, PE3/PE4 only joins the tunnel rooted at its chose upstream PE
  - One of them will not receive traffic
- Solution
  - With previously updated C-multicast Inter-AS propagation procedure (in -00 revision), C-multicast route's construction can be done as in intra-AS case:
    - In particular, using RD from the UMH route advertised by the chosen upstream PE
  - Now PE3 and PE4 constructs C-multicast routes with different RDs
    - Specifically, PE1's and PE2's RD respectively
    - Both routes will be propagated, and both PE1 and PE2 will transmit packets
      - PE3/PE4 receives and accepts packets from their chose upstream PE respectively

Problem 2 & Solution

- Two problems after problem 1 is resolved:
  - Problem 2a
    - Two copies from AS100 to ASBR22 - Inefficient use of inter-as resources
    - Solution: ***Single Forwarder Selection*** for sources in remote ASes
      - So that PE3 and PE4 will select the same upstream PE
      - Upstream PE selection based on installed unicast route can still be used for sources in the local AS if configured so
        - This is so that egress PEs can receive traffic from closest upstream PE in the same AS
  - Problem 2b
    - If inclusive inter-as tunnels are used and PE1/PE2 both transmit, PE3/PE4 will receive duplicate traffic
      - PE3/PE4 have no way to tell which copy is from its chosen upstream PE

Problem 2b Solutions

- Option 1: Ingress ASBR attaches a PE Distinguisher label when it sends traffic from local PEs into its inter-AS tunnel
  - PE Distinguisher labels advertised in the Inter-AS I-PMSI A-D route
  - From the PED label in packets an egress PEs knows the source PE of the traffic
- Option 2: Ingress ASBR does IP forwarding and only accepts/forwards traffic from the upstream PE of its own choice
  - Ingress ASBR needs to receive C-multicast routes and treat as PIM/IGMP joins from a local PE-CE interface
    - To do that, egress PEs need to attach a RT that identifies the VRF on the ingress ASBR
    - The RT's value comes from a VRF Route Import EC attached to the Inter-AS I-PMSI route
      - Just like that PEs attach VRF Route Import EC to UMH routes

Next Steps

- Seeking comments
  - For existing and new aspects of this document
- Seeking WG adoption
  - The document covers quite a few clarifications/enhancements/fixes for MVPN/EVPN
  - The authors believe the document is ready for adoption