

Support C-Bidir with Ingress Replication

draft-zzhang-l3vpn-mvpn-bidir-ingress-replication

**Jeffrey Zhang
Yakov Rekhter
Andrew Dolganow**

87th IETF, Berlin



MVPN support for C-bidir

C-bidir: PE-CE multicast protocol being PIM-Bidir

PIM-Bidir: Designated Forwarder election required on LAN

- MVPN backbone is a simulated LAN

Ways to avoid DF election in the MVPN backbone

- Backbone becomes the RPL: Section 11.1, RFC 6513
- Partitioned set of PEs: Section 11.2, RFC 6513

Partitioned set of PEs

PEs are grouped into partitions wrt a particular C-RPA

- A partition includes all PEs selecting the same UMH wrt the C-RPA

A PE only accepts traffic from PEs in the same partition

- Traffic carry a label corresponding to the UMH: Section 11.2.2
 - Advertised as PE Distinguisher (PED) label
 - Upstream allocated by tunnel root
 - Cumbersome for Ingress Replication (IR) P-tunnel
- Each partition uses its own Bidirectional P-tunnel: Section 11.2.3
 - Partial Mesh of MP2MP P-Tunnels

What if one wants to use IR?

Simulating “Partial Mesh of MP2MP P-tunnels” with IR: The “normal” way

An MP2MP tunnel could be simulated by a set of IR tunnels

- One IR tunnel rooted at each PE on the MP2MP tunnel
 - Consisting of a set of P2P LSPs
 - One P2P LSP to each other PE on the tunnel
 - Each PE originates a Leaf A-D route for each IR tunnel
 - N-square

Simulating “Partial Mesh of MP2MP P-tunnels” with IR: Proposed Optimization

One S-PMSI A-D route from the UMH wrt a C-RPA

- Identifying the MP2MP tunnel
- PTA specifies IR and includes a label that the UMH would not allocate for any other PE
 - For other PEs to send traffic to the UMH
 - Typically, different labels are allocated for different PEs
 - So that traffic can be associated with the sending PEs
 - In this case, we want to associate the traffic with the partition (represented by the UMH)

One Leaf A-D route responded from each PE in the same partition

- When it has relevant local states – details later
- Imported by all PEs: Not just by the S-PMSI originator
- PTA includes a label corresponding to the UMH
 - For other PEs to send traffic to the Leaf A-D route’s originator
 - Associate the traffic with the partition (represented by the UMH)

S-PMSI A-D Routes

Originated only by PEs that have local routes (through a VRF interface) to one or more C-RPAs

A (C-*,C-BIDIR) S-PMSI A-D route is always originated

- From each PE that has a local route to any C-RPA
- A single (C-*,C-BIDIR) S-PMSI A-D route is originated even if a PE has local routes to multiple C-RPAs

One or more (C-*,C-G-BIDIR) S-PMSI A-D routes can be originated by a PE with local routes to the C-G-BIDIR's C-RPA

- By typical triggers for S-PMSI

An S-PMSI A-D route identifies an MP2MP tunnel

- With leaves including the originators of the matching S-PMSI and Leaf A-D routes

When to respond with Leaf A-D routes

A PE responds to a (C-*,C-G-BIDIR) S-PMSI A-D route

- If it has local C-G-BIDIR join states learned from its CEs

A PE responds to a (C-*,C-BIDIR) S-PMSI A-D route

- If it has ANY local C-G-BIDIR join states learned from its CEs

Where S-PMSI A-D route's originator is the UMH wrt the C-G-BIDIR's C-RPA

- Optionally, a PE may respond even when the UMH is not the S-PMSI A-D route's originator
 - In this case, traffic will arrive from PEs outside its own petition, with a label corresponding to a PE different from the UMH, and shall be discarded
 - May be useful for live-live protection

PIM-Bidir Forwarding States

PIM-Bidir has (*,G) and (*,G-prefix) forwarding states

- Implementation dependent – but assumed in this proposal
 - for exemplary purpose
- (*,G) states for groups with explicit joins
- (*,G-prefix) states for “sender-only-branches” (no joins)
 - A G-prefix is a group range, where all groups in the range have the same RPA

PIM-Bidir Forwarding States in VRFs

Denoted as (C-*,C-G-Bidir) or (C-*,C-G-Bidir-prefix)

- OIF List = local OIFs + P-Tunnel branches

For a (C-*,C-G-Bidir) S-PMSI A-D route from the UMH

- Install (C-*,C-G-Bidir) forwarding state with P-tunnel branches determined by the S-PMSI A-D route and matching Leaf-AD routes

For a (C-*,C-BIDIR) S-PMSI A-D route

- Install (C-*,C-G-Bidir-prefix) routes, with P-tunnel branches determined by the S-PMSI A-D route and matching Leaf A-D routes
 - If the S-PMSI A-D route's originator is the UMH wrt C-G-Bidir-prefix's C-RPA

For a (C-*,C-G-bidir) local join state w/o (C-*,C-G-BIDIR) S-PMSI A-D route:

- Install (C-*,C-G-Bidir) forwarding state, with P-tunnel branches determined by the (C-*,C-BIDIR) S-PMSI A-D route and matching Leaf A-D routes

Plan

Seek review and comments

- Revision to be posted soon to address comments from Eric Rosen

Seek WG adoption



everywhere