

BIER Penultimate Hop Popping

draft-zhang-bier-php-00

102th IETF, Montreal

Zhaohui (Jeffrey) Zhang

Use Case

- An MVPN/EVPN deployment with enough P/PE routers capable of BIER to benefit from using BIER as provider tunnel
- We know how to deal with BIER incapable P routers
- What if some PE routers are not BIER capable?

BIER Incapable Egress PE

- A BIER capable Ingress PE would have to send traffic via BIER to BIER PEs and via traditional tunnels to incapable PEs
 - Complicated & inefficient
- What if an incapable egress PE pretends it supports BIER, but requests the upstream BFR to pop the BIER header?
 - Transparent to other PEs

BIER PHP

- A BIER incapable router signals BIER information but requests other BFRs to pop the BIER header and send traffic “natively”
 - Those BFRs do not have to be directly connected
 - There could be incapable P routers in between – traffic could be tunneled
- PHP requested via
 - PHP sub-sub-TLV in BIER sub-TLV
 - MPLS and non-MPLS encapsulation
 - Implicit Null Label as label range base
 - In MPLS Encapsulation sub-sub-TLV

Conditions for PHP

- Egress PE must be able to demultiplex the payload w/o the BIER header
 - IP payload: must be in the address space for the BIER routing underlay
 - MPLS payload: top of the label stack must be “downstream assigned” by the egress PE
 - DCB labels considered as “downstream assigned”
 - Draft-ietf-bess-mvpn-evpn-aggregation-label
 - VXLAN/NVGRE: should be fine
 - Others: For further study

BIER Incapable Ingress PE

- Have BIER capable PEs relay traffic from incapable Ingress PEs
 - MVPN: Virtual Hub & Spoke (RFC 7024)
 - EVPN-mpls: Virtual Hub & Spoke
 - draft-keyupate-bess-evpn-virtual-hub
 - EVPN-overlay: Assisted Replication
 - draft-ietf-bess-evpn-optimized-ir

Next Steps

- Seek WG comments
- Seek WG adoption