

# Multicast VPN using BIER

E. Rosen, Ed.

M. Sivakumar

IJ. Wijnands

S. Aldrin

A. Dolganow

T. Przygienda

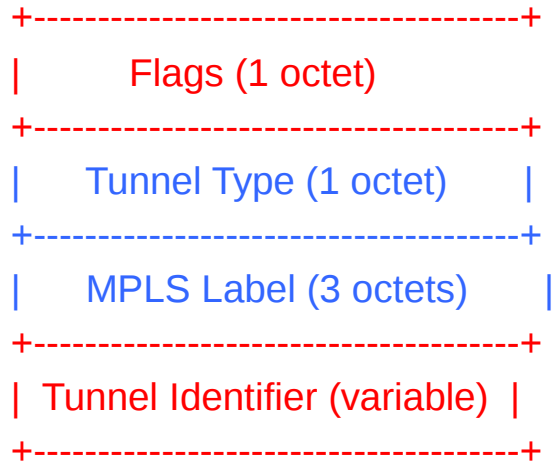
(<https://tools.ietf.org/html/draft-ietf-bier-mvpn-01>)

# BIER P-Tunnel

- No explicit multicast tunnel building with BIER
- Can however be modeled as an implicit P2MP tunnel through a **BIER domain** (from a **BFIR** to all the **BFERs**)
  - BIER domain is coextensive with an IGP network
  - Tunnel not specific to any VPN
  - Traffic from multiple MVPNs aggregated on a single BIER P-tunnel
- BIER carries traffic within one IGP domain
  - MVPN allows P-tunnels to be “**segmented**”
  - Segmentation points are **Autonomous System Border Routers (ASBR)** or **Area Border Routers (ABR)**
  - As long as segmentation end-points are capable of acting as BFIR and BFERs, BIER to be used on one or more segments of a P-tunnel

# PMSI Tunnel Attribute (PTA)

- Identifies the P-tunnel to which one or more flows are bound



- **Tunnel Type**
  - New tunnel type for “BIER”
- **Tunnel Identifier**
  - **First sub-field, a single octet, contains the sub-domain-id of the sub-domain that the BFIR will assign the packets that it transmits on the PMSI**
  - Second sub-field has the BFR-Prefix of the originator of the route carrying the PTA
- **Leaf-Info-Required bit**
  - Set in S-PMSI A-D route and clear in I-PMSI A-D route

# Flags

- When tunnel type is BIER, two bits in PTA flags are significant
  - Leaf Info Required per Flow (LIR-pF)
    - Introduced in [EXPLICIT\_TRACKING]
    - **SHOULD** be **SET** only if BFIR knows all BFERs (at-least ones to which it needs to transmit)
    - **MAY** be **SET** in (C-\*, C-\*) S-PMSI AD route but **MUST NOT** be **SET** in I-PMSI A-D routes
  - Leaf Info Required (LIR)
    - **SHOULD** be **CLEAR** in I-PMSI or (C-\*, C-\*) S-PMSI AD routes unless LIR-pF is set. **SHOULD** be **CLEAR** in (C-S, C-\*) and (C-\*, C-G) S-PMSI AD routes
    - **SHOULD** be **SET** in other S-PMSI AD routes

# Upstream Assigned Label

- MPLS Label
  - Upstream assigned by the router originating the PMSI route to which PTA is attached
  - MPLS label in two x-PMSIs ( $x \rightarrow I$  or  $S$ ) originated by an ingress PE for BIER MUST be different if
    - they carry a different set of Route Targets (RTs)
    - Ingress PE supports extranet and the routes are from two different VRFs
    - Ingress PE supports “Extranet separation” and only one of the routes carry the “Extranet Separation” EC
    - If segmented p-tunnels are used, the MPLS label can be used by the BFER to identify the C-flow to which a data packet belongs.

# Explicit Tracking

- BFIR can determine the set of BFERs to which packet needs to be sent in two ways
  - by using explicit tracking mechanism based on the LIR flag, as specified in RFC 6513 and 6514
  - by using the explicit tracking mechanism based on the LIR-pF flag as specified in [EXPLICIT TRACKING]
    - may be used if and only if segmented P-tunnels are not being used

# Using the LIR-pF flag

- BFIR originates (C-\*, C-\*) S-PMSI AD route with LIR-pF flag set in the PTA
- Each BFER responds with one or more Leaf AD routes identifying the flows interested in receiving traffic from the BFIR
  - Significant reduction in S-PMSI AD routes needed
  - One disadvantage is that it does not provide a way for BFIR to assign unique label to each c-flow

# Encapsulation & Transmission

- BFIR needs to find the S-PMSI or I-PMSI AD route that is a “match for transmission”
- To create BIER header, BFIR needs to know all egress PEs
  - If matching S-PMSI AD route has LIR flag set
    - Leaf AD routes have “route key” identical to the NLRI of the S-PMSI AD route
  - If matching S-PMSI AD route has LIR-pF flag
    - Leaf AD routes have “route key” derived from NLRI of the S-PMSI AD route according to section 5.2 of [EXPLICIT TRACKING]



# Decapsulation (egress PE)

- If BFER is an egress PE
  - determines, using upstream-assigned MPLS label that it is an egress PE
  - delivers the following to the multicast flow layer
    - BFR-prefix (corresponding to the bier-id and sub-domain-id in the header) and
    - payload (MPLS packet with top label as upstream assigned MPLS label)

# Decapsulation (segment-boundary)

- Packet needs to be forwarded on to its next tunnel segment
  - choice of next tunnel depends on the c-flow to which packet belongs
  - accomplished by assigning distinct upstream-assigned MPLS label to each c-flow

# References

- RFC 6513
- RFC 6514
- [EXPLICIT TRACKING]
  - Dolganow, A., Kotalwar, J., Rosen E., and Z. Zhang, “Explicit Tracking with Wild Card Routes in Multicast VPN” internet-draft draft-ietf-bess-mvpn-expl-track-00, March 2015