

# Update on draft-ietf-bess-mvpn-expl-track

A. Dolganow  
J. Kotalwar  
*E. Rosen, Ed.*  
Z. Zhang

# Two Major Topics

- First topic
  - Use of *PMSI Tunnel Attribute (PTA)* as follows:
    - Tunnel type of *no tunnel information*,
    - *Leaf Info Required (LIR)* flag set
  - This provide mechanism for ingress node to learn:
    - which egress nodes are interested in receiving which multicast flows from the ingress node
    - without impacting the binding of flows to tunnels
  - Useful information for monitoring
- Second topic
  - New *PTA* flag: *Leaf Info Required per Flow (LIR-pF)*

# *Leaf Info Required per Flow (LIR-pF)*

- Ingress-PE can send *wildcard S-PMSI* route with *LIR-pF* set
- Egress-PEs respond with a *Leaf* route for each flow
  - *route key* field identifies individual flow
- Without *LIR-pF*, egress-PE responds with only one *Leaf* route
  - *route key* field contains the wildcard (the *S-PMSI*'s *NLRI*)
  - cannot identify individual flows
- Useful with BIER or Ingress Replication, since ingress-PE:
  - benefits from per-flow response,
  - can save the overhead of sending per-flow *S-PMSI* route.
- **N.B.:** this changes the ingress PE processing of received *Leaf* routes, because now a valid *Leaf* route may have a *route key* that is not identical to an *S-PMSI* route's *NLRI*

# Recent Changes

- Many clarifications and textual improvements
  - Thanks to Stephane Litkowski for thorough shepherd's review
- Changes in the way *Leaf* routes are constructed when sent in response to *LIR-pF*
  - New and improved way of identifying that a *Leaf* route is a response to *LIR-pF*
  - Much more detailed spec of how the *PTA* of such a *Leaf* route is constructed
    - depending upon the *PTA* of the *S-PMSI* route to which it is a reply
- New section detailing how an ingress node processes a received *Leaf* route whose route key is not identical to the NLRI of a sent x-PMSI route

# Marking a *Leaf* Route to Indicate that it is a Response to *LIR-pF*

- New technique:
  - If and only if *Leaf* route is response to *LIR-pF*, it carries a *PTA* with *LIR-pF* set
    - (so *PTA* is now mandatory in such *Leaf* routes)
- Previous drafts instead required modification of the *RD* in the *route key* field:
  - that was a more complex solution with more unintended side-effects

# More Detail about Constructing *PTA* of *Leaf* Route in Response to *LIR-pF*

- New details about how to set *Tunnel type* field:
  - when *Leaf's PTA* can have same *tunnel type* field as corresponding *S-PMSI* route
  - vs. when it can specify *no tunnel type*
- New details about how to set *MPLS Label* field:
  - For *Ingress Replication*, specifies procedure that allows:
    - Ingress to send one wildcard *S-PMSI* route
    - Egress to not only reply with one *Leaf* route per flow, but also to **optionally** specify an MPLS label for each flow
- Informational reference added to bier-mvpn document, which covers these details for the case where the ingress specifies a tunnel type of BIER

# New Section On Ingress Node Processing of Received *Leaf* Routes

- If *Leaf* route *PTA* has *LIR-pF* set, route key field doesn't have to match *NLRI* of any *S-PMSI* route originated by the ingress
  - but the *Leaf* route does have to be a valid response to some *S-PMSI* route, or it won't impact the multicast processing on the ingress node
  - ingress node has to match the *Leaf* route to the right *S-PMSI* route (or at least keep an eye on the *Leaf* route until there is a right *S-PMSI* route)
- Details provided for handling received *Leaf* routes, with *LIR-pF* set, that specify *IR* tunnels and per-flow MPLS labels

# What Documents Does this Update?

- Updates RFC 6625 (*MVPN Wild Cards*),
  - RFC 6625 neglects to handle combination of wild cards with *PTA* that specifies *no tunnel type*
- Modifies ingress node processing of received *Leaf* routes, therefore updates:
  - RFC 6514 (*Base spec for BGP-MVPN*), which says that *Leaf* route *root key* must be identical to *NLRI* of corresponding *S-PMSI* route
  - RFC 7524 (*Inter-Area Segmentation*) which also has *Leaf* routes whose *route keys* don't correspond exactly to the *NLRI* of an *S-PMSI* route.



# Next Steps

- Document is ready for WG LC
- It has been claimed that this document is a normative reference of the BIER-MVPN spec, which is ready to go to the RFC Editor
  - We hope we can move this document along quickly, so as not to block the BIER documents.