

Global Table Multicast with BGP-MVPN Protocol

draft-zhang-mboned-mvpn-global-table-mcast-00

Jeffrey Zhang, Lenny Giuliano – Juniper

Dante Pacella – Verizon

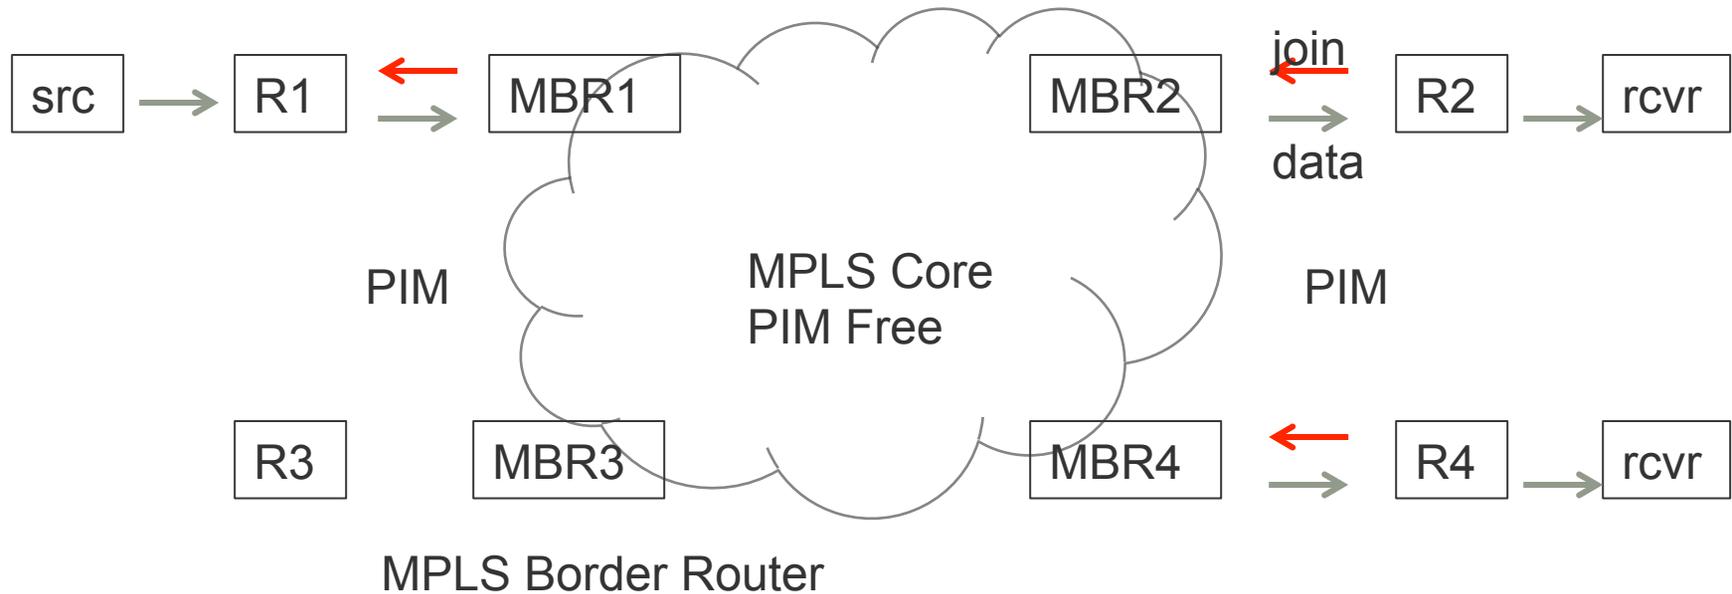
Jason Schiller - Google



Problem Statement

Global Table Multicast over an MPLS Core

- mLDP in-band Signaling
- GTM procedures as specified in draft-ietf-mpls-seamless-mcast
- **BGP-MVPN**



Benefits of GTM with BGP-MVPN

Same protocol & procedures

- Some clarifications on use of RD/RT

Same vendor implementation & operator experiences

Most features & characteristics of BGP-MVPN apply

- Scaling, Aggregation
- Flexible choice of provider tunnels
- Support for PIM-ASM/SSM/Bidir outside the core
- Support for unsolicited flooded data
 - E.g. BSR for Group-to-RP mapping protocol
- Extranet: between VRFs and Global Table

Can co-exist with the other GTM procedure specified in draft-ietf-
mpls-seamless-mcast

Operation

MBRs follow BGP-MVPN protocol & procedure

- Like PEs in MVPN case

Treat global table as a VRF as far as signaling is concerned

- As if an all-zero RD (0:0) is associated with the global table VRF

When an MBR advertises UMH routes to other MBRs, it attaches VRF Route Import and Source AS ECs

- Local Administrator field of the VRF Route Import EC is set to 0
 - Or any value that uniquely maps to the global table on the MBR
- If IBGP session runs between MBR and other routers on the same side of the core,
 - Either MBR needs to reflect UMH routes to the core side, with policy to attach VRF Route Import and Source AS ECs, or
 - RFC 6368 model need to be followed
 - Advertise IBGP learned routes to other IBGP peers

Route Distinguisher

UMH routes do not have RDs

- Single Forwarder Selection procedure can not be used as result

For A-D routes

- Use 0:0 by default, or some other values as appropriate

For C-Multicast routes:

- Use 0:0 if the local and upstream PE are in the same AS
 - RFC 6513 uses VPN-IP UMH route's RD value
- Otherwise use the RD value from a matching Inter-AS I-PMSI A-D route – as in RFC 6513
 - Inter-AS case needs further thoughts

Route Targets For BGP-MVPN Routes

Purpose is to confine importation to Global Table only

Use RT 0:0 for Intra-AS I/S-PMSI and Source Active A-D Routes

- I-PMSI A-D - RFC 6514 allows:
 - having a set of Route Targets used for the Intra-AS I-PMSI A-D routes being distinct from the ones used for the VPN-IP unicast routes
- S-PMSI A-D - RFC 6514 allows:
 - the set of Route Targets carried by the route to be specified by configuration
- Source Active A-D – same as I-PMSI A-D case

C-Multicast – VRF Route Import EC from UMH route

- Same as in RFC 6514

Leaf A-D – derived from matching PMSI A-D route's nexthop

- same as in RFC 6514

Plan

- Seeking review & comments from WG
- Addressing comments from Eric; new revision will be posted soon

Seeking WG adoption

- L3VPN seems to be the right home

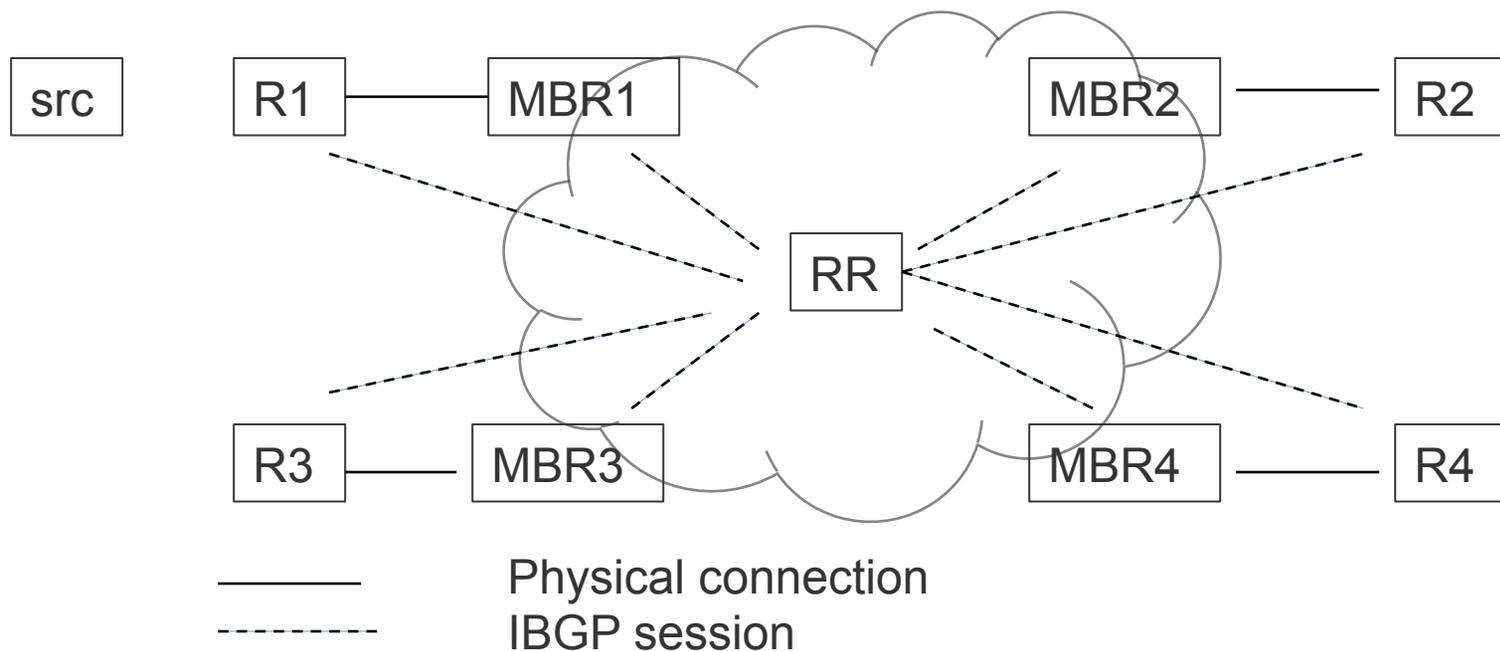
Informational or Standard Track?

- Aimed at Informational
 - Only one small procedure needs clarifying
 - Special deployment scenario may be taken out of scope
 - Is it common and worth standardization effort?
- Open to discussion

Backup Slides: UMH routes learned from non-MBRs over the core

IBGP sessions among all Rx and MBRx

- Full mesh or through RR
- R1 advertise UMH routes to ALL others directly
 - But R1 does not run BGP-MVPN
 - MBR1 runs BGP-MVPN but does not get to attach required ECs



Backup Slides: Solution for the special situation

MBR4 learns UMH route with BGP next hop as R1

- No VRF Route Import and Source AS ECs attached

MBR4 looks up route to R1, which may have been advertised by MBR1 with BGP next hop MBR1

- The route may be a BGP route with the required ECs attached
 - All set – use the attached ECs; otherwise continue

MBR4 looks up route to MBR1

- The route may be via a RSVP tunnel, with endpoint MBR1, and MBR1 has originated a BGP-MVPN A-D route
 - All set – construct a Route Import RT with the Global Admin field being MBR1 and Local Admin field being 0; otherwise continue

Recurse until resolved or reaches end

- Could even work if UMH routes are distributed via IGP throughout