Draft-rosen vs L3VPN WG documents

- Since 2000 we have 9 iterations of draft-rosen
  - From draft-rosen-00 to draft-rosen-09
- Mechanisms described in draft-rosen-06 form a subset of the mechanisms specified in L3VPN WG documents
- Additional mechanisms introduced in draft-rosen-07, draft-rosen-08, and draft-rosen-09 to support I-PMSI auto-discovery (using MDT SAFI) and inter-AS operations option B (using the Connector attribute) are **not** part of the mechanisms specified in L3VPN WG documents, and are **not interoperable** with the corresponding mechanisms specified in L3VPN WG documents

Multi-vendor interoperable implementations

- Multi-vendor interoperable implementations support:
  - I-PMSI using PIM-SM only
    - No support for PIM-SSM, or PIM-Bidir
      - Support for I-PMSI using PIM-SSM requires BGP MDT SAFI
  - S-PMSI using PIM-SSM
  - Inter-AS operations options (A) and (C) only
    - No support for Inter-AS operations option (B)
  - Based on draft-rosen-06 + subset of draft-rosen-07 (S-PMSI with PIM-SSM)
- No multi-vendor interoperable implementations based on draft-rosen-07, or draft-rosen-08, or draft-rosen-09

Differences between unicast and multicast with draft-rosen: architecture

BGP/MPLS VPN – unicast:

- Based on the Aggregated Routing architecture[1]
- Inter-AS/inter-provider scenario allows to constrain exchange of routing information to only ASBRs
- Control plane is decoupled from its data plane

draft-rosen – multicast:

- Based on the Virtual Router architecture[1]
- Inter-AS/inter-provider scenario requires PEs in different ASes/providers to have (direct) routing peering
  - As long as these PEs have at least one VPN in common
- Control plane is coupled with the data plane
  - The same inter-PE tunnels are used to exchange both control and data

[1]See RFC4110 for more on the comparison between the Aggregated Routing and the Virtual Routers architectures

Differences between unicast and multicast with draft-rosen: mechanisms

BGP/MPLS VPN – unicast:

- Uses BGP to exchange VPN unicast routing information among PEs
- Supports both MPLS and GRE for inter-PE tunnels
- Uses LDP or RSVP-TE for setting up inter-PE (MPLS) tunnels

draft-rosen – multicast:

- Uses PIM to exchange VPN multicast routing information among PEs
- Supports only GRE and IP-in-IP for inter-PE tunnels
  - Only GRE is implemented by multiple vendors
- Uses PIM for setting up inter-PE (GRE) tunnels

Control Plane scalability considerations (1)

- Given VRF on a PE has to maintain PIM adjacency with every other VRF of that MVPN
  - Granularity of such PIM adjacency is **per MVPN per PE**, not just per PE
  - Direct consequence of the Virtual Router model
- VRF also has to maintain PIM adjacencies with all the locally connected CEs of the VRF's MVPN
- Usually for a given MVPN the number of sites per PE is (much) less than the number of PEs that have sites of that MVPN => the overhead of maintaining PIM adjacencies with other PEs dominates the overhead of maintaining PIM adjacencies with the locally connected CEs

Control Plane scalability considerations (2)

Multicast with draft-rosen:

PE-PE PIM adjacencies

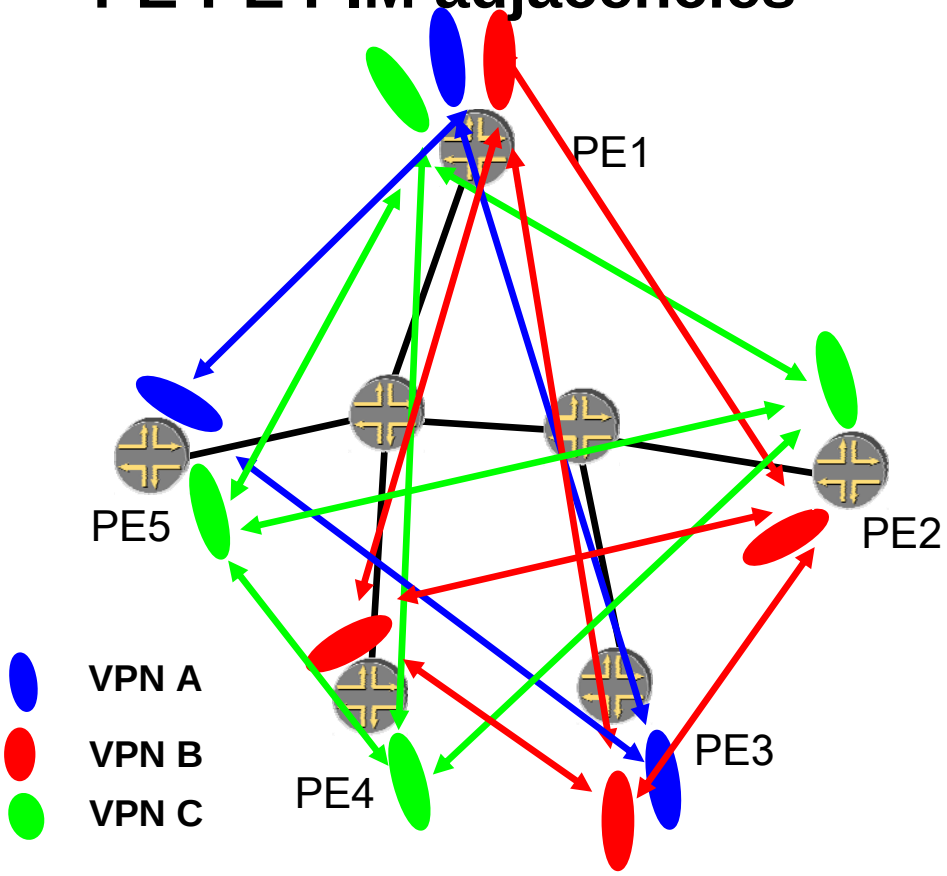

Unicast:

BGP adjacencies

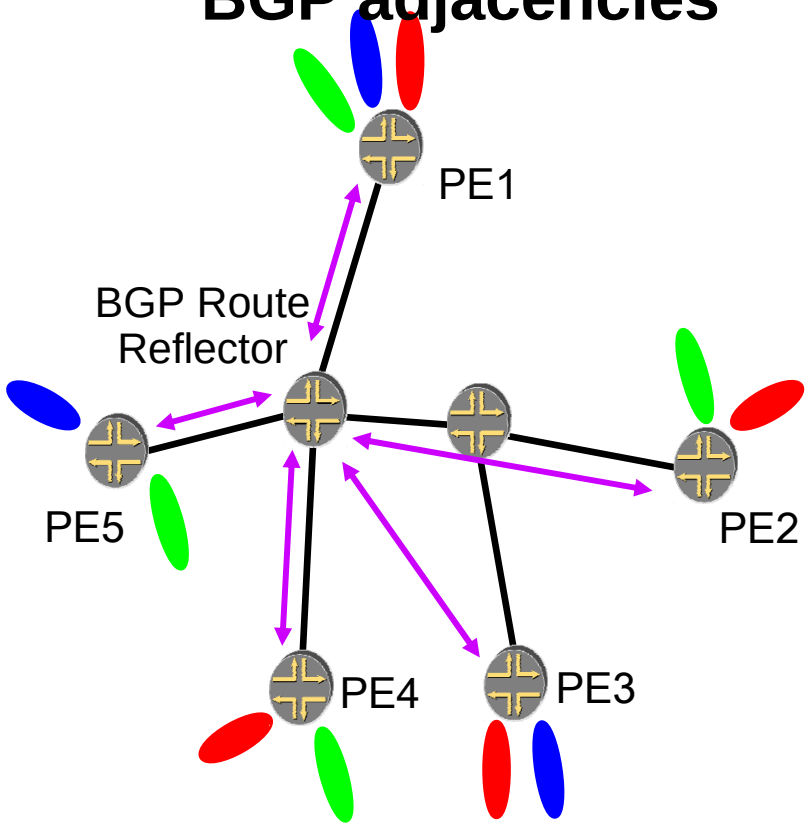

(Diagram shows routing adjacencies on all PEs)

Join/Prune Latency

- Direct consequence of unreliable transport used by PIM
- Loss of the first PIM Join for a given (C-S, /C-RP, C-G) that a PE sends to another PE results in Join latency of up to 30 secs
- Loss of the last PIM Prune for a given (C-S/C-RP, C-G) that a PE sends to another PE results in (a) the receiver PE receiving unwanted traffic, and (b) the upstream PE maintaining unnecessary state
  - For roughly 3 minutes (see [PORT])
- Service provider(s) infrastructure has different performance characteristics than a single LAN => do not extrapolate experience with PIM over a single LAN to PIM over service provider(s) infrastructure that emulates a single LAN

Packet losses when switching to S-PMSI

- Direct consequence of using unreliable transport for signaling switching to S-PMSI
- Loss of the first S-PMSI advertisement results in losses of multicast data for up to 57 secs (MDT-INTERVAL – MDT-DATA-DELAY)

Anycast Customer RP (C-RP)

- If an MVPN customer uses anycast RP, where several C-RPs use the same (anycast) address, but are reachable via different PEs, then for a given C-G at any given point in time only one of these C-RPs can be used to deliver (multicast) traffic to other sites of that MVPN
  - Direct consequence of treating service provider(s) infrastructure as an (emulated) LAN

Multi-homed Multicast Sources

- If an MVPN site contains a (multicast) source for a given C-G, and the site is multi-homed to several PEs, then at any given point in time only one of these PEs can be used to deliver (multicast) C-G traffic from the source to other sites of that MVPN
  - This is in contrast to unicast, where unicast traffic could be forwarded to/from the source via all of these PEs
  - Direct consequence of treating service provider(s) infrastructure as an (emulated) LAN

Mandatory I-PMSI

- Each MVPN requires its own I-PMSI
  - Even if most/all of the multicast data is sent using S-PMSIs
- Results in extra overhead of maintaining I-PMSI
  - Both in the control and in the data plane
- The overhead is especially significant in the inter-AS scenario when I-PMSI is realized using PIM-SSM
  - As the number of point-to-multipoint tunnels required by a single I-PMSI is equal to the number of PEs that span that I-PMSI

QoS support

- DiffServ QoS mechanism for multicast with draft-rosen is different from DiffServ QoS mechanism for unicast
  - Multicast uses IP-based DiffServ, unicast uses MPLS-based DiffServ
- MPLS traffic engineering, which is available for unicast, is not available for multicast with draft-rosen
- MPLS DiffServ Traffic Engineering, which is available for unicast, is not available for multicast with draft-rosen

Protection/restoration

- MPLS Fast re-route mechanisms that are available for unicast VPN traffic, are **not** available for multicast with draft-rosen

Security Considerations (1)

- Security considerations in RFC4797 (unicast BGP/MPLS IP VPNs with PE-PE tunnels realized via GRE) are also applicable in the context of draft-rosen
- draft-rosen presents additional security considerations:
  - Inability to restrict joining I-PMSI of a given MVPN to only the PEs that have VRFs of that MVPN may result in:
    - Leaking multicast traffic originated within that MVPN to the receivers outside of that MVPN
    - Various forms of packet spoofing

Security Considerations (2)

- Implications of packet spoofing in the context of draft-rosen are more significant than in the context of RFC4797
  - While in the context of RFC4797 spoofing can impact only the data traffic, in the context of draft-rosen spoofing can impact both the data and the control traffic associated with the exchange of MVPN routing information among PEs
    - This is because in the context of draft-rosen the same GRE tunnels are used to exchange both control and data traffic
- Protection against packet spoofing by securing control traffic associated with the exchange of MVPN routing information among PEs by applying security mechanisms specified in RFC4601 is problematic
  - See Section 10 of draft-rekhter-mboned-mvpn-deploy-00.txt for more details

Next Step

- Merge this document with draft-ycai-mboned-mvpn-pim-deploy-02.txt, and present the resulting document to MBONED for considerations as an MBONED WG document