Considerations about mVPN standardization

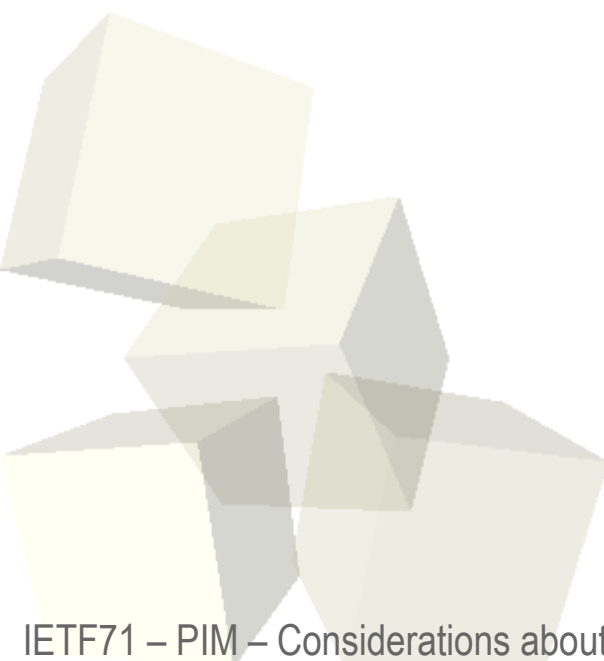draft-morin-l3vpn-mvpn-considerations-02

Thomas Morin - FT Orange

Benjamin Niven-Jenkins - BT

Nicolaï Leymann - DT

Yuji Kamite - NTT

Raymond Zang - BT

- The l3vpn working group is working to produce a *"specification of IPv4 multicast over BGP/MPLS VPNs for publication"* (charter)
- draft-ietf-l3vpn-2547bis-mcast intends to be that specification, but describes multiple approaches for certains building blocks
  - in its current state, it is more a "framework" document
  - not enough to privode solution interoperability
- we believe there is a need for progressing this work toward a good standard candidate, by identifying a set of mandatory procedures
- we would consider the recourse to a "multiple solutions" approach as a last resort

- **Goals**
  - ◆ discuss the different options proposed in draft-ietf-l3vpn-2547bis-mcast, in the lights of requirements formulated in RFC4834
  - ◆ identify the better candidates for a set of mandatory mVPN procedures
- **Non goals**
  - ◆ it is not the aim of this document to remove options
  - ◆ non-mandatory approaches would still be described in the final solution spec as OPTIONNAL procedures
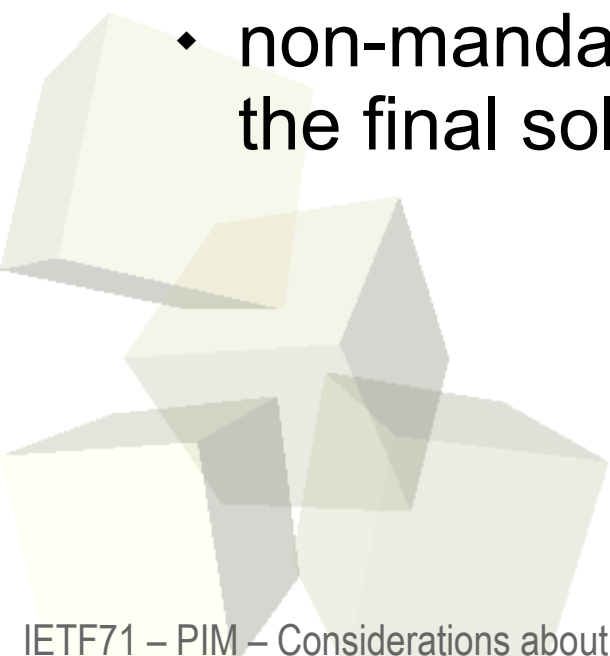
- **Draft -00 submitted one year ago for Prague**
  - co-authored by 4 operators
- **Draft -01 submitted in October**
  - Vancouver showed good support from the working group to adopt as a WG item
  - extensive discussion in December on the mailing list
- **Draft -02 submitted for this meeting**
  - updated to address the issue raised
  - presented during l3vpn WG on Monday
  - still some friction...
    - some debate around the arguments and conclusions
    - it appears that some disagree with the goal of this draft

- For mVPN auto-discovery
  - proposed approaches
    - BGP-based auto-discovery
    - Discovery using PIM Hellos through an MI-PMSI implemented with an any-to-any tunnel (ASM/PIM-SM, or Bidir-PIM or MP2MP LDP)
  - the BGP-based auto-discovery is seen as the approach to make mandatory
    - because it is a prerequisite to support certain type of tunnels (PIM-SSM, P2MP mLDP, P2MP RSVP-TE...)
    - because it is a prerequisite for the BGP-based approach for C-multicast routing to be usable
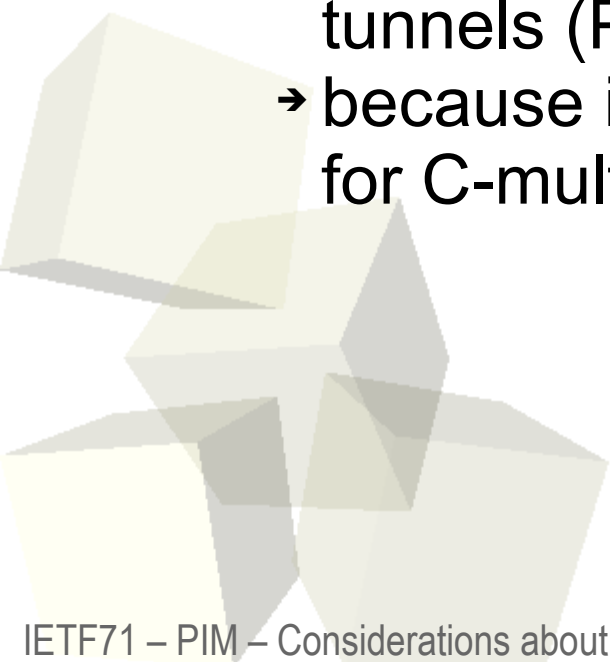
- For the signalling of S-PMSI
    - candidates
        - UDP-based signalling over MI-PMSI
        - BGP-based signalling
    - we see the BGP-based signalling as the right candidate for the set of mandatory procedures
    - key reasons :
        - no need for another protocol
            - BGP S-PMSI signalling is very close to the BGP auto-discovery procedures
        - the UDP-based signalling does not do more, but less
            - restricted to deployments using an MI-PMSI which uses more state in the P routers when one-to-many tunnels are used to built an MI-PMSI
            - limitations wrt. Inter-AS deployments
    - other pros/cons still being debated, but currently seen as not changing the conclusion

- **For C-multicast routing**
  - two approaches are focused on:
    - PIM LAN procedures on an MI-PMSI
    - BGP-based procedures
  - current state of our document is that:
    - the BGP-approach is seen as having a number of significant advantages, for scaling, architecture consistency, and reduced deployments contraints
    - in the mean time...
      - not all of the above is "free" : e.g. BGP mVPN may push for dedicated RRs, and/or multi-session BGP
      - unlike with the PIM-based approach which is known since deployments of draft-rosen, there are "unknowns" related to the BGP approach due to lack of experience : it is debated on how BGP handles the dynamic nature of multicast routing, and its impact on join latency
    - the above is detailed in the draft
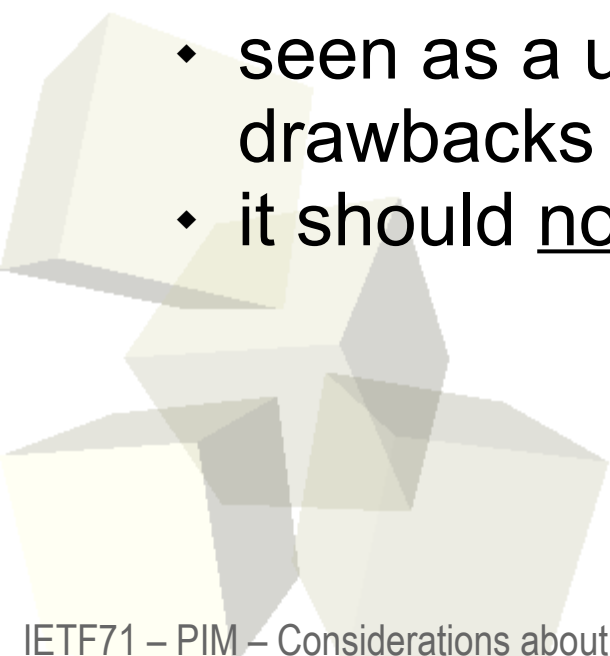 - **Conclusion : for now, recommend to implement both, delay mandating one until further experience is gained**

- **For inter-AS**
  - two approaches are proposed
    - segmented inter-AS trees
    - non-segmented inter-AS trees
  - we see the segmented approach as the right candidate for the mandatory set of procedures
    - because it offers the largest degree of deployment flexibility to operators
    - because of the scalability improvements for P-routers state and C-multicast routing
    - for VPN inter-AS option B deployments, it is the approach that has the best fit, and which isolate the most two ASes
  - it is also identified that the non-segmented approach, similar to the one used in the deployed draft-rosen, can be helpful for some scenarios and migrations
    - still recommend the implementation of the non-segmented approach

- Also discussed...
- Type of tunnels
  - just like for unicast VPNs, given the diversity of backbone engineering choices and constraints, there doesn't appear to be one and only "best" tunnel type
  - conclusion : suggest prioritizing mLDP, P2MP RSVP-TE and GRE/IP-Multicast
- RP "collocated" in the PE
  - seen as a useful optional feature, but also has drawbacks
  - it should <u>not</u> be required

- Feedback from the multicast-knowledgeable community is welcome:
    - to review the solutions considered
        - see Eric presentation about draft-ietf-l3vpn-2547bis-mcast
    - to help us pursue the comparisons between the different considered approaches
        - please review draft-morin-l3vpn-mvpn-considerations
- mboned BCP on mVPN
    - can be useful to help distinguish between approaches proposed, based on experience gained on implementations !

Thank you !
Comments ?