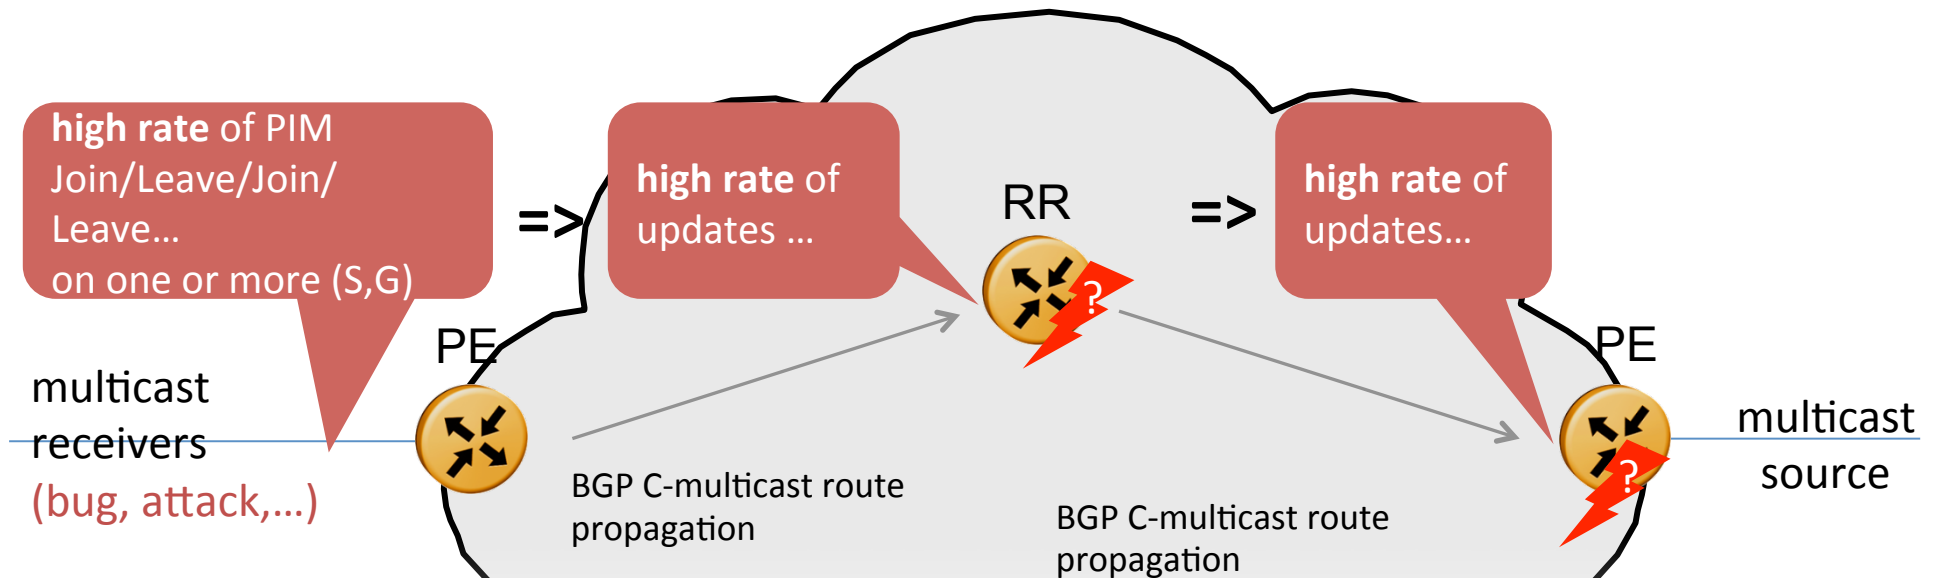


Multicast state damping

[draft-morin-multicast-damping-00](#)

Thomas Morin, Stéphane Litkowski, Keyur Patel,
Jeffrey Zhang, Robert Kebler

Problem statement



- **High dynamicity in membership updates from CEs can result in high control plane load in the infrastructure (PEs, RRs)**

- issue exist for PIM-based and BGP-based C-multicast routing

- This is not only in theory: we can significantly load the CPU of upstream routers in the lab

- Rate limiting Join/Prune messages received at the edge ?

- too crude: unpredictable impact on legitimate uses of the service (or limited efficiency)

- Applying BGP route damping ?

- as is, it has too much impact on the service delivered

- Cases in which we need to address this problem:

- multicast in VPNs **and** multicast in the global routing table (similar)

Solution proposed

- Principle: delay the propagation of prunes
 - if too much Join/Prune activity on (C-S,C-G), stop propagating Prune(C-S,C-G) toward the upstream router, for some time
 - for BGP C-mcast routing, it means: delay before withdrawing the route
- Benefit:
 - if the number of (C-S,C-G) is limited, this result in an upper bound of the average rate of Join/Prunes sent to the upstream
 - ⇒ **protects the upstream router from excessive Join/prune activity**
 - all Join/Prunes take effect locally as they did before
 - ⇒ **no impact on the service delivered**
- Side effect:
 - average increase of bandwidth in the core
 - traffic present on a P-tunnel for a longer time
 - minor increase => acceptable trade-off

Proposed procedures [1/2]

- We could apply dampening on VRF PIM states
 - we are proposing generic PIM damping in mboned
 - it does not allow to protect against dynamicity coming from inter-AS C-multicast route redistribution
 - it does not provide the option of protecting upstream PEs at the RRs
- We recommend using BGP route damping, with a few twists:
 - [keep the principle of exponential decay, increments, high/low threshold]
 - **when a BGP C-multicast route is damped, keep advertising it** (instead of withdrawing it)
 - use specific damping parameters and default values for C-multicast routes
 - and require times to be configurable in seconds

Proposed procedures [2/2]

- Selective provider tunnels bound to a specific S-PMSI also follow group membership dynamicity
 - (C-S,C-G) S-PMSI
 - but also true for wildcard S-PMSI

⇒ **the state of the provider tunnels need also be damped**

- There are different ways to do it
 - build damping in the P-tunnel protocols (mLDP, PIM)
 - damp Leaf A-D route (applies to P2MP RSVP-TE only)
 - join/leave P-tunnel based on BGP C-multicast routes, not based on VRF C-PIM states

Conclusions, next steps

- To do:
 - ASM states
 - default and max values
- Feedback welcome on the principle and proposed procedures
- We would like this draft to find a home
 - problem and proposed solution are similar for VPN and non-VPN cases
 - mboned looks like a better home than PIM or L3VPN (even if these WGs would have to be involved)
 - the alternative is to progress the two separately