Multicast state damping

draft-morin-multicast-damping-00

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

Problem statement

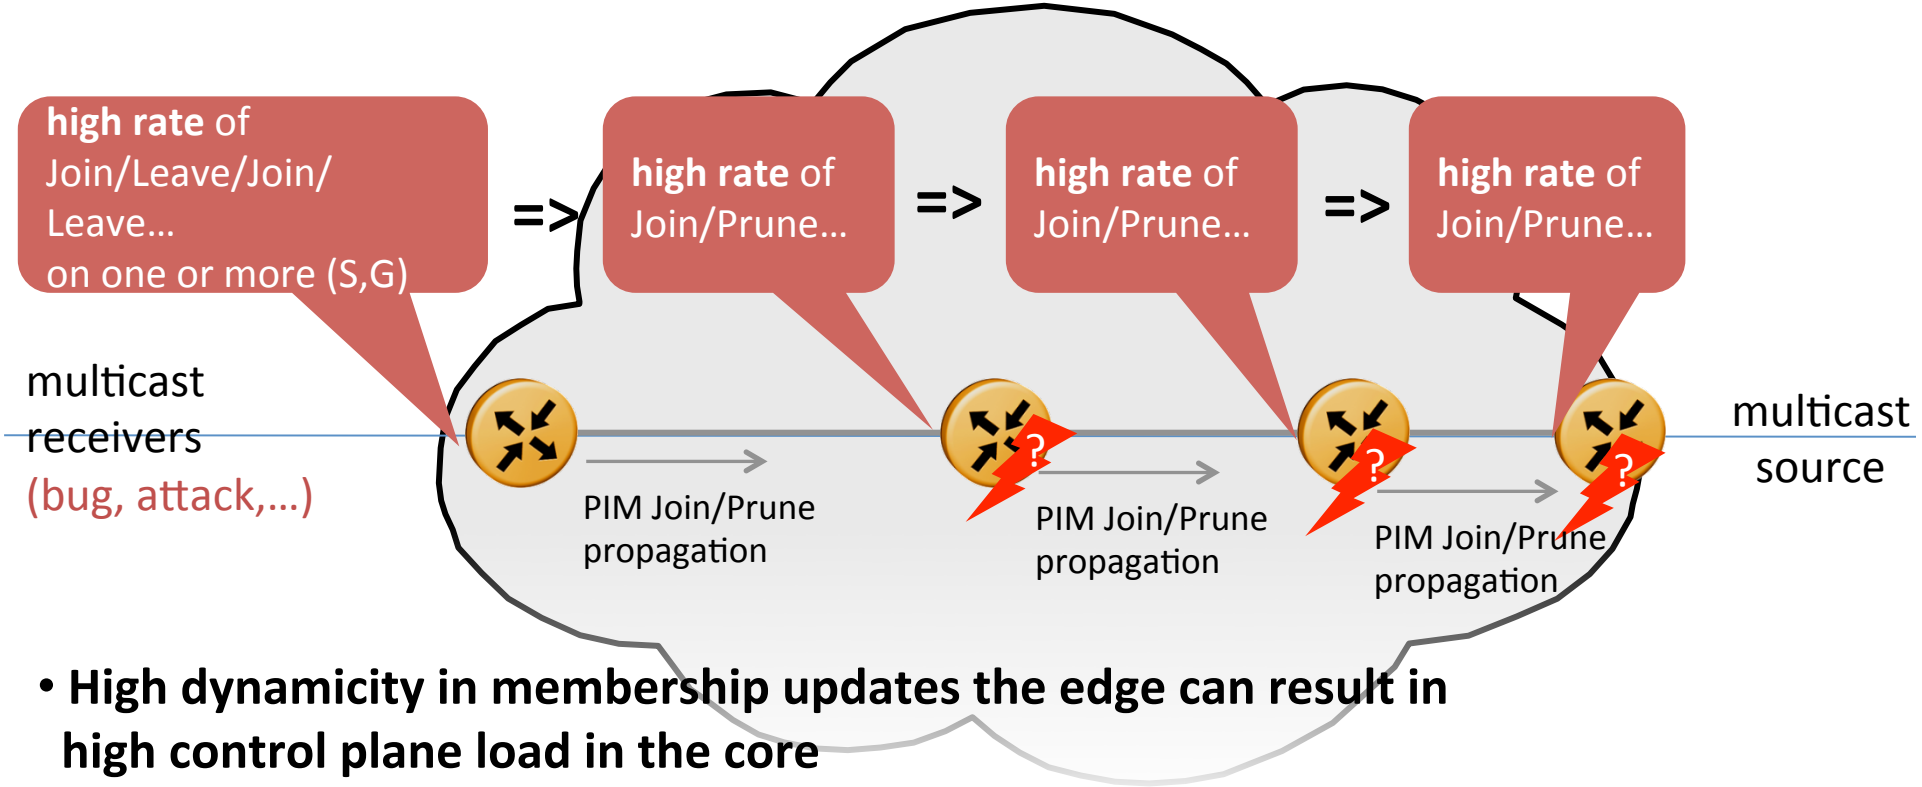

- **High dynamicity in membership updates the edge can result in high control plane load in the core**
- This is not only in theory: we can significantly load the CPU of upstream routers in the lab
- Rate limiting Join/Prune messages received at the edge ? (as proposed in RFC4609)
  - too crude: unpredictable side-effects impacting legitimate uses of the service (or limited efficiency)
- Cases in which we need to address this problem:
  - multicast in the global routing table and multicast in VPNs

Solution proposed

- Principle: delay the propagation of prunes
  - if too much Join/Prune activity on (S,G), stop propagating Prune(S,G) toward the upstream router, for some time
- Benefit:
  - if the number of (S,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 on the upstream link
  - minor increase => acceptable trade-off

Proposed procedures

for multicast routing in the global routing table

- Reuse the concept (not the specs) of BGP dampening…
  - an arbitrary number is associated to each route
  - it increases at each routing event on the route
  - it decreases with time, with an exponential decay
  - crossing a high threshold triggers damping
    - router stops advertising the route
  - (crossing a low threshold => damping stops)
  - thresholds and how fast the decays happens is configurable
- …with a few twists:
  - what BGP damping does: stop advertising a damped route
  - to delay the propagation of Prune, we want:
    - omit Pruning an (S,G) with too much activity
    - multicast-specific parameters and values
- [detailed procedures are in the draft]

Proposed procedures
for multicast routing in VPNs

- The same problem needs to be solved for multicast in VPNs (RFC6513)
- Different options exists for multicast VPN routing
  - PIM-based
    - ⇒ PIM procedures for the global routing table could be applied in VRFs
  - BGP-based
    - ⇒ doing damping in the VRF is not enough: does not cover inter-AS
    - ⇒ **need to adapt BGP dampening so that it does what we want for BGP routes carrying C-multicast routing information:**
      - dampen route withdrawals, not route announcements
      - multicast-specific parameters and default values
- Damping provider tunnels membership is also needed
  - provider tunnels carry multicast VPN traffic
    - can be based e.g. on PIM, mLDP, or P2MP RSVP-TE
    - can be dynamic (S-PMSI, aka "Data-MDTs")

=> the state of the provider tunnels need also be damped

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)