

MAC binding aging in OVN

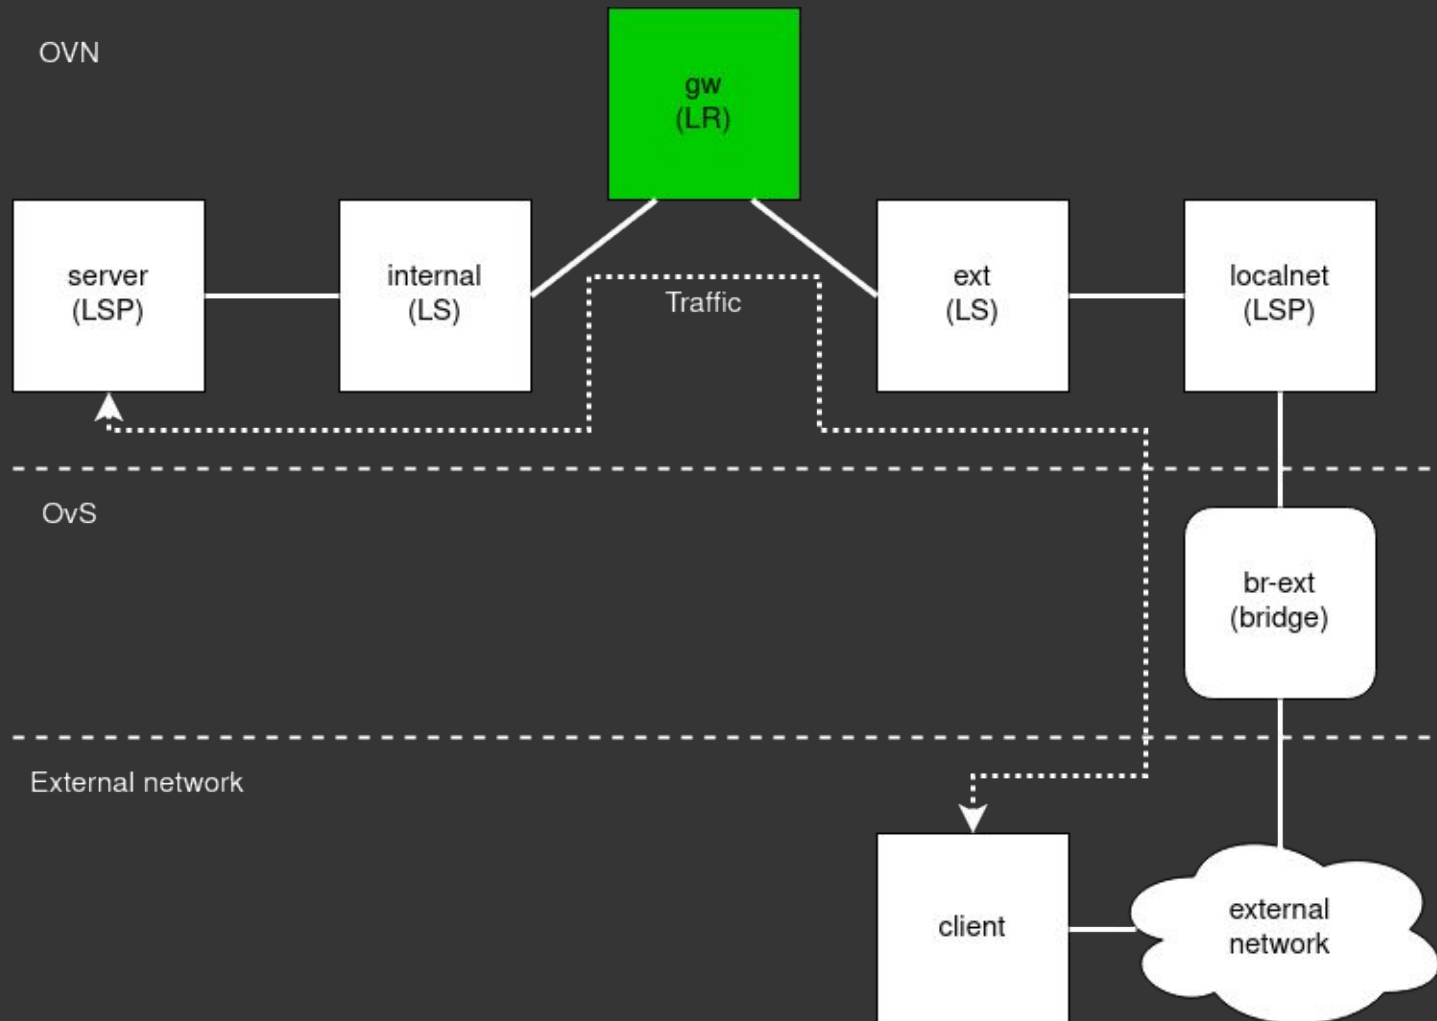
Ales Musil

Senior Software Engineer | OVN Core

Agenda

- ▶ MAC binding
- ▶ Motivation
- ▶ Considered solutions
- ▶ Measurements
- ▶ Current solution
- ▶ Possible improvements

Topology



MAC binding row

- ▶ `_uuid` : ad00de8f-99bf-41a1-8058-4af13f6056b4
- ▶ `datapath` : 5f99830a-0c21-40b4-8f7d-83956447abc7
- ▶ `ip` : "192.168.10.15"
- ▶ `logical_port` : gw-public
- ▶ `mac` : "00:00:00:00:11:00"
- ▶ `timestamp` : 1666246695935

Motivation

- ▶ When?
 - Large scale environments
 - Environments with a lot of changes in topology
- ▶ Why?
 - Nothing stops the MAC binding table to grow infinitely
 - Size of SB has direct impact on OVN performance

Idle age column

- ▶ Column in MAC binding
- ▶ Pros
 - Traffic is not disturbed
 - Easy to implement
- ▶ Cons
 - Possible source of race conditions
 - High contention over SB DB

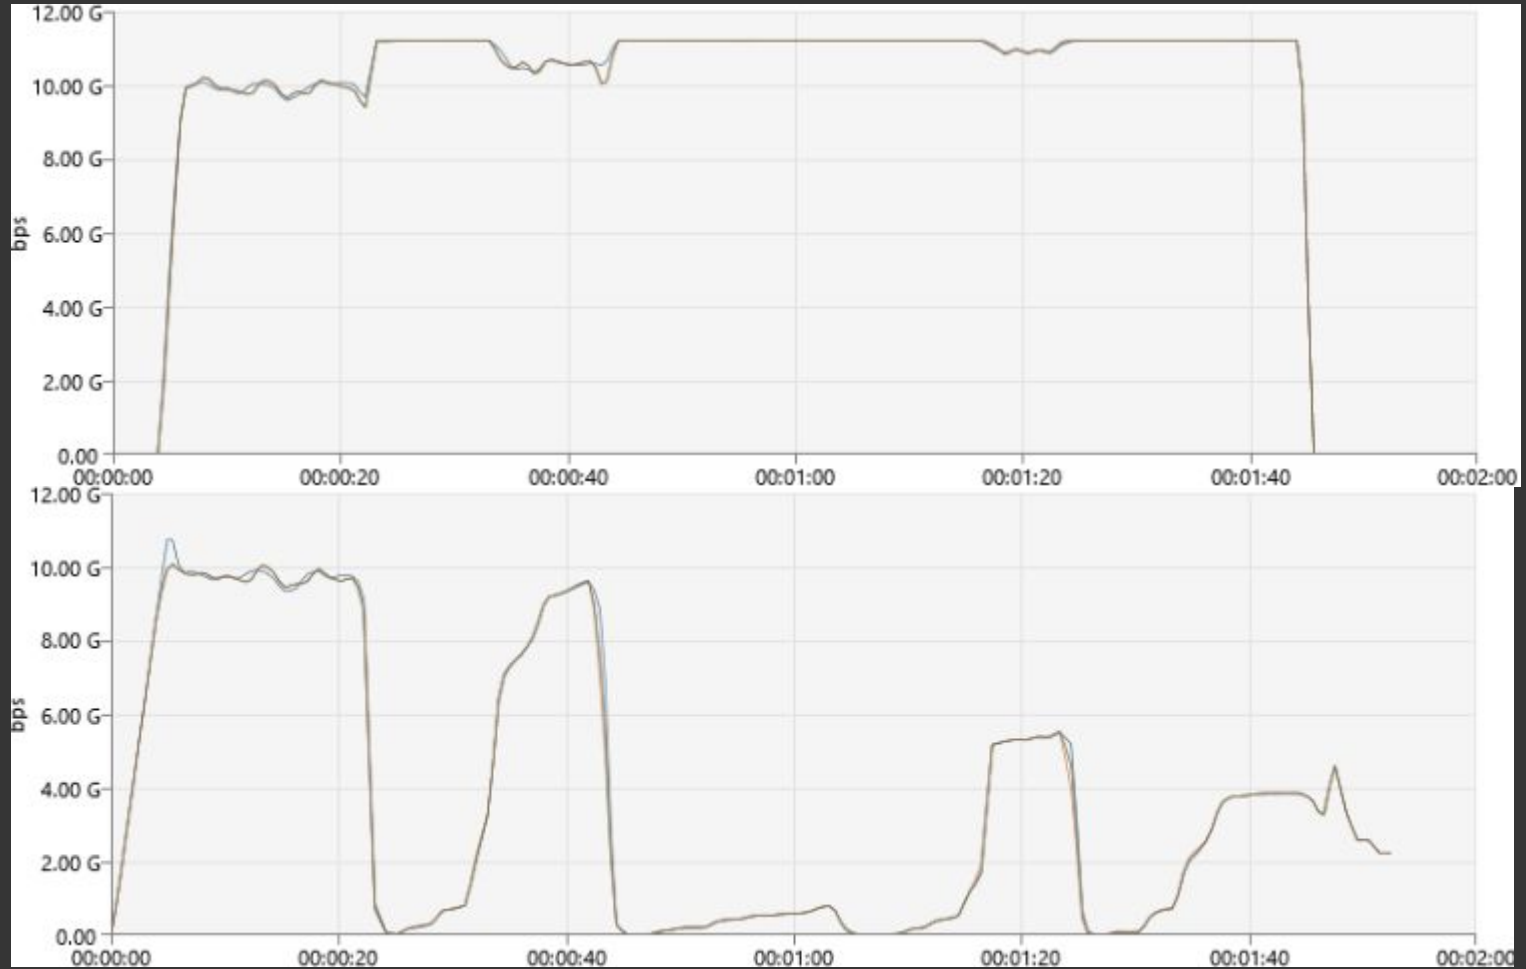
“Validating” flow

- ▶ Flow before the MAC binding check
- ▶ Pros
 - Traffic is not disturbed
 - We don't have to check “idle_age” of corresponding flows
- ▶ Cons
 - Possible source of race conditions
 - High contention over SB DB
 - Possibly a lot of controller actions


Owner

- ▶ ovn-controller would “own” the MAC binding row
- ▶ Pros
 - Easy to implement
 - Single controller is in charge
- ▶ Cons
 - It would cause traffic disruption
 - A lot of IPC with OvS to check the idle_age

TCP traffic



UDP traffic



Northd

- ▶ Northd monitors all MAC binding rows
- ▶ Pros
 - Very easy to implement
 - Only northd is in charge of aging
- ▶ Cons
 - It can cause traffic disruptions
 - It can cause SB DB contention on bulk removal

Configuration

- ▶ Configurable per logical router
 - `ovn-nbctl set logical_router <LR>`
`options:mac_binding_age_threshold=<THRESHOLD>`
- ▶ Bulk removal limit
 - `ovn-nbctl set nb_global .`
`options:mac_binding_removal_limit=<LIMIT>`

Improvements

- ▶ Getting feedback from users
- ▶ Avoid removal of used MAC bindings
 - ovn-controllers updating the timestamp
 - Active approach, sending arp for MAC bindings that are about to expire

Resources

- ▶ [Bugzilla RFE](#)
- ▶ [Measurement results](#)
- ▶ [Patch series](#)
- ▶ [Solution discussion](#)

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



twitter.com/RedHat