

# TC update

- Or Gerlitz, Mellanox
- Netdev 2.2, Nov 2017, Seoul, Korea

# Outline:

- Flows insertion rate optimizations
  - work by Chris and Jiri – merged to 4.14
- Flows insertion rate further optimizations
  - design by Rony and Co. presented @ netdev, to be continued
- vxlan tunnels sriov offload - proper handling of neigh events/changes
  - work by Hadar, 4.12
- Hairpin

# NIC Hairpin

- the term is borrowed from switches...
- ability of NIC HW to fwd packet from port1 to port2 (or port to itself)
  - implemented by pairing a receive queue to transmit queue
- Flow based.. packets that match certain rules
- Use ingress tc/flower rules with the mirrored redirect action
  - these are NIC and not e-switch rules
- other actions can apply to packets before the redirect, e.g header re-write

# Hairpin – continued

- allows to fully offload the data-path of various SW/DDoS gateways / load-balancers to HW
- packets go through all the required processing in HW (encap/decap, push/pop vlan, header re-write) and then forwarded
- CPU stays at practically zero usage
- HW Flow counters are used by the control plane for monitoring, billing, etc
- coming up soon 😊