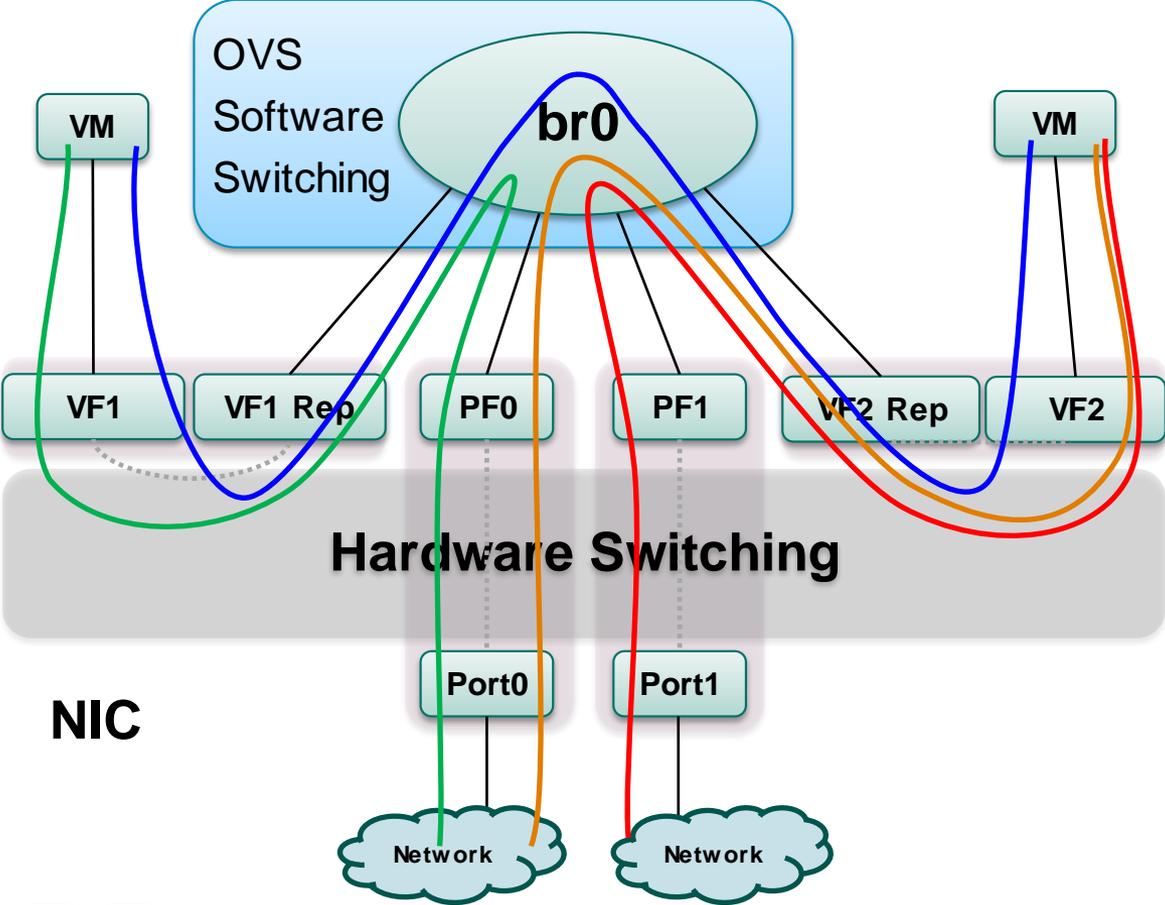


bnxt_en TC Flower offload

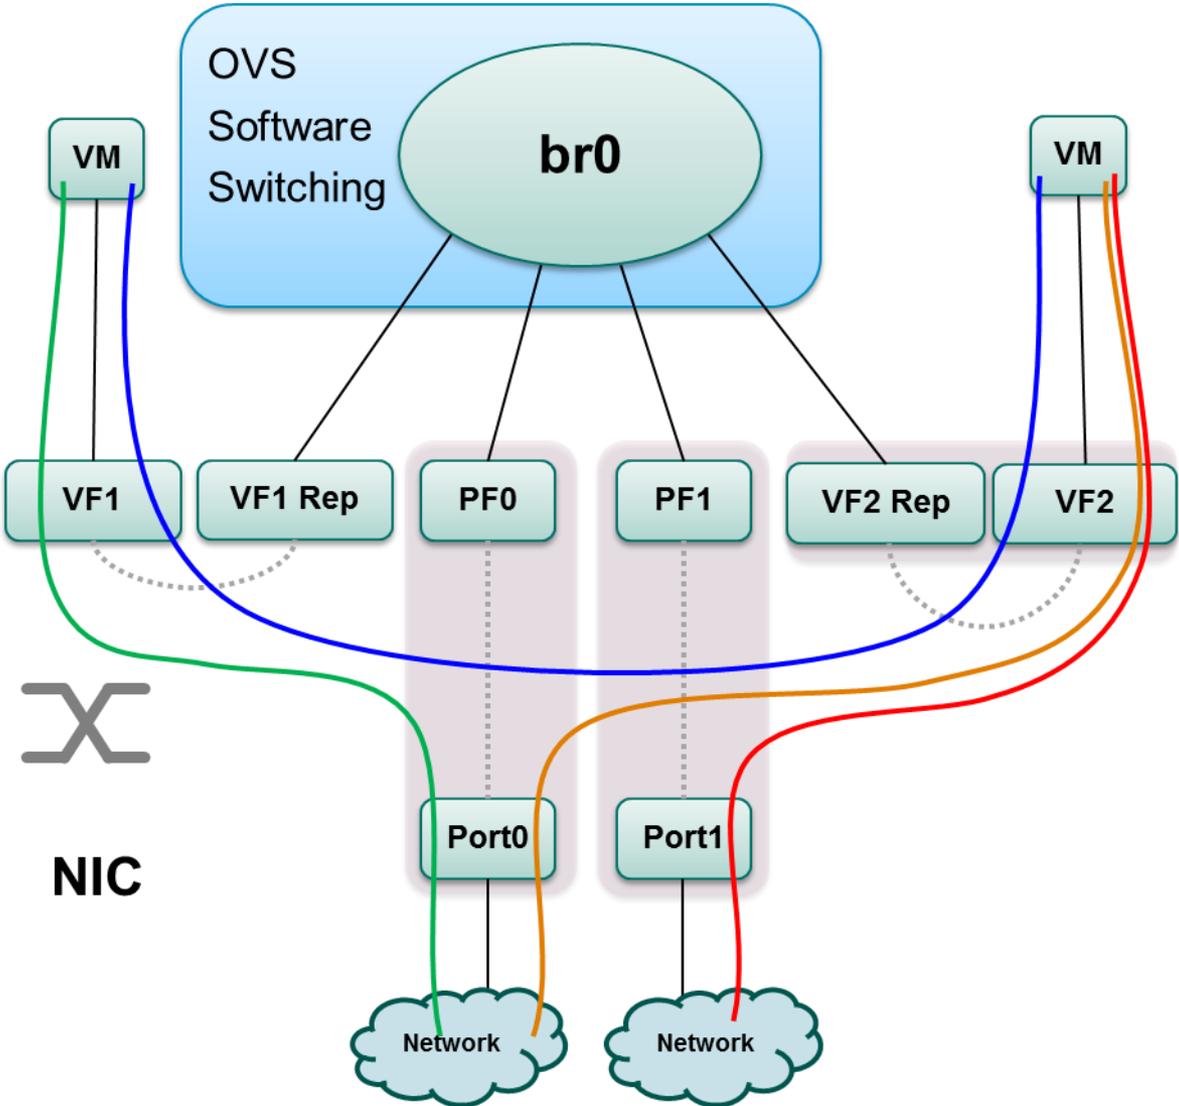
Michael Chan



bnxt_en switchdev mode (before offload)



bnxt_en switchdev mode (after offload)



bnxt_en TC Flower offload overview

- **Implementation was done by Sathya Perla**
- **Code went upstream in July 2017**
- **TC Flower Classifier (ingress):**
 - L2, VLANs, ipv4, ipv6, TCP/UDP, ICMP, VXLAN (ipv4)
- **TC Actions:**
 - Drop, redirect, push/pop VLAN
- **Flow statistics**

Comparing TC Flower with NTUPLE

- **NTUPLE**

- Directing RX flows to rings
- Dissect RX packet to get flow key
- Create exact match HW filter from flow key
- Action: redirect to a ring

- Use simple hash table to store the filter using RSS hash
- No hardware statistics

- **TC Flower**

- Directing flows in SRIOV eswitch
- TC subsystem provides flow key
- Create HW filter from flow key and flow mask
- Action: drop, redirect, push/pop VLAN, tunnel encap/decap
- Use rhashtable to hash the filter multiple ways
- Hardware flow statistics

Future plans

- **Switching across the ASIC**
- **Uplink representation**
- **More actions, such as pedit**
- **Improve on the statistics gathering**
- **Improve on Flow insertion rate**