



Contribution ID: 160

Type: **not specified**

TC SW Datapath: A Performance Analysis

Wednesday, November 14, 2018 4:00 PM (35 minutes)

Currently the Linux kernel implements two distinct datapaths for Open vSwitch: the `ovskdp` and the TC DP. The latter has been added recently mainly to allow HW offload, while the former is usually preferred for SW based forwarding due to functional and performance reasons.

We evaluate both datapaths in a typical forwarding scenario - the PVP test - using the `perf` tool to identify bottlenecks in the TC SW dp. While similar steps usually incur in similar costs, the TC SW DP requires an additional, per packet, `skb_clone`, due to a TC actions constraint.

We propose to extend the existing act infrastructure, leveraging the `ACT_REDIRECT` action and the `bpf redirect` code, to allow clone-free forwarding from the mirrored action and then re-evaluate the datapaths performances: the gap is then almost already closed.

Nevertheless, TC SW performance can be further improved by completing the RCU-ification of the TC actions and expanding the recent listification infrastructure to the TC (ingress) hook. We plan also to compare the TC/SW datapath with a custom eBPF program implementing the equivalent flow set to tag a reference value for the target performances.

Presenters: ABENI, Paolo (Red Hat); CARATTI, Davide (Red Hat); CHAUDRON, Eelco (Red Hat); LEITNER, Marcelo Ricardo (Red Hat)

Session Classification: Networking Track