

Upcall rate limiting in OVS-DPDK

Anju Thomas

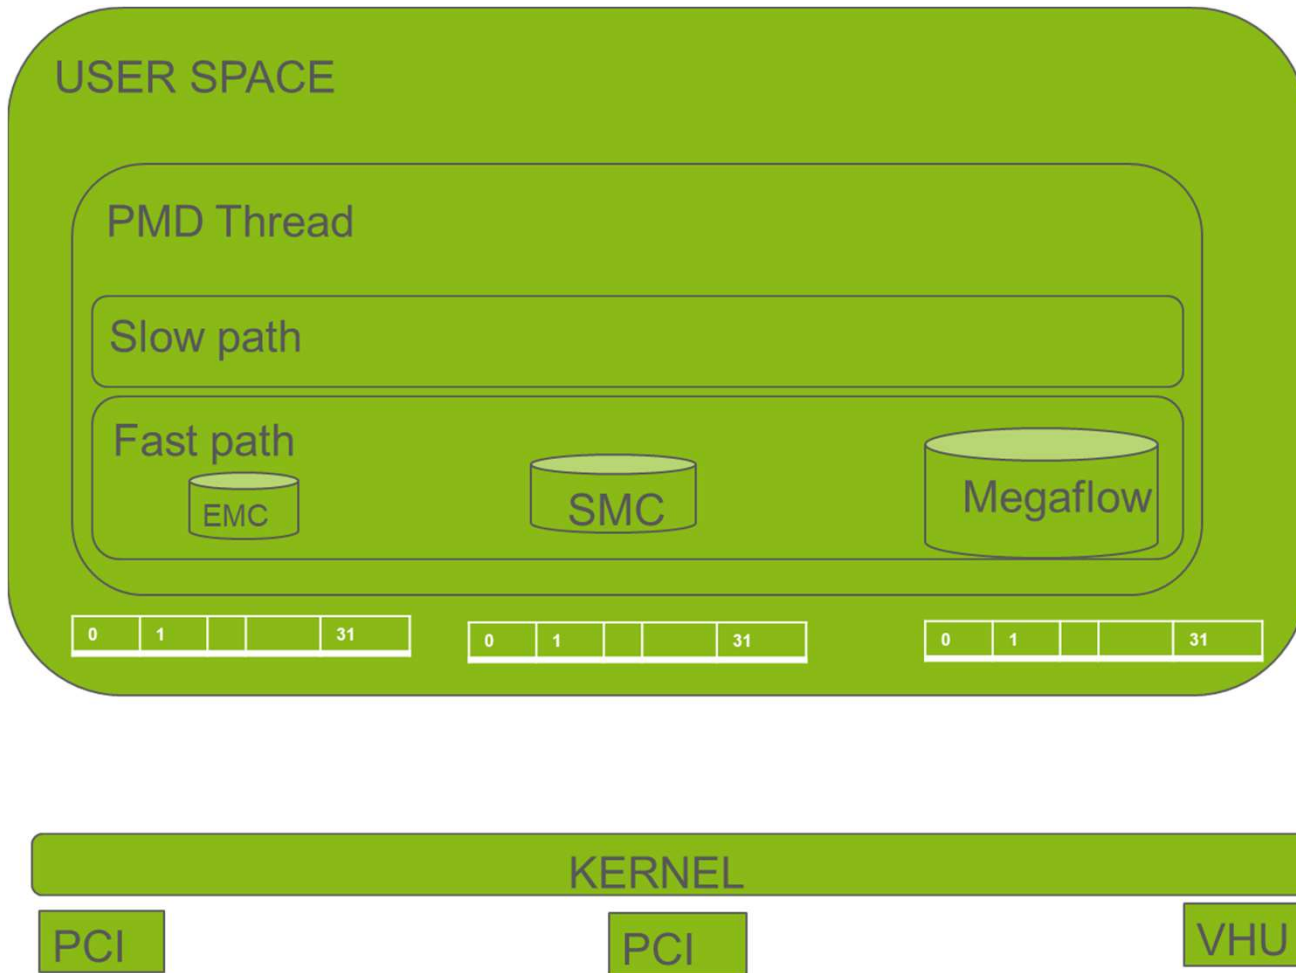
Nitin Katiyar

Vishal Deep Ajmera

Venkatesan Pradeep

AGENDA

- Slow path and fast path in OVS+DPDK
- Why rate limit
- Solution



WHY RATE LIMIT?

- In OVS-DPDK both slow path and fast path execute in the context of the PMD thread.
- Lookup cost of EMC:Megaflow:Slowpath is 1:10:100
- This means that new flows leading to upcalls can increase the latency of other packets with learnt flows which are serviced by the same PMD.
- Sudden burst of such flows can also lead intermittent traffic drop for already learnt flows .
- This also make OVS-DPDK vulnerable for DoS attacks

PROPOSED SOLUTION



- Limit rate of packets going into slowpath
- Simple token bucket policer per PMD thread
- Upcall policer can be enabled/disabled using the below knob:
- *ovs-vsctl set Open_vSwitch . other_config:upcall-rl=true*
- How to configure upcall rate and burst
 - *ovs-vsctl set Open_vSwitch . other_config:upcall-rate=xxx*
 - *ovs-vsctl set Open_vSwitch . other_config:upcall-burst=xxx*
- A packet entering slow-path has to take a token to get into slow-path and if no tokens are available, the packet is dropped and accounted per PMD

PROPOSED SOLUTION

- Configure upcall rate limiting per port
 - *ovs-vsctl set Interface vm1p1 other_config:upcall_rate=xxx*
 - *ovs-vsctl set Interface vm1p1 other_config:upcall_burst=xxx*
- Or per queue
 - *ovs-vsctl set Interface vm1p1 other_config:queue_upcall_rate="0:xxx,1:xxx,2:xxx"*
 - *ovs-vsctl set Interface vm1p1 other_config:queue_burst_rate="0:xxx,1:xxx,2:xxx"*

SUMMARY & STATUS

Summary

- Importance of rate limiting in OVS+DPDK
- DOS attacks
- Setting upcall rate limits at PMD level
- Setting upcall rate limits at Interface/queue levels

Status

- Patch in the community . Need to work on the port/interface solution
 - <https://patchwork.ozlabs.org/patch/909676/>

THANK YOU