

2022-03-23 ViNePerf Agenda and Minutes

Attendees

Sridhar Rao

Luc Provoost

Daniele Zulberti

Al Morton

Federica Paganelli

Sl. No	Topic	Presenters	Notes
1	Development Update		<p>Trying to successfully deploy L3 network without manual configuration...</p> <p>Limited to 2 CNIs: SRIOV and Userspace</p> <ol style="list-style-type: none">1. Userspace OVS - Only Bridge mode is supported (L2).2. Userspace VPP - CNI is not adding IP/routes to the Interfaces in VPP. It assigns the IP on the container-side, but not on the host-side. Expect some forwarding rules to be added on the host-side.
2.	Discussion of the results		<p>Summary :</p> <ol style="list-style-type: none">1. VPP performs better than OVS. OVS performance is also inconsistent. OVS: Latency too is high ???2. Prox performs better than T-Rex3. Bi-Directional Prox is more inconsistent (throughput) than unidirectional4. Trex results are not consistent.5. Increasing hops, even with single node, affects the performance. Better summary for varying topologies6. Prox has some limitations - mainly w.r.t core assignments - which affects the performance. <p>To Try:</p> <p>Topology-1</p> <p>VPP + Prox</p> <p>Bidirectional, 64Bytes, ...</p> <p>Variables: Cores (6), Rx_queues, Queue_size, Tx_Queues - double these values.</p>
3.	Prox-L3		<p>Prox-Tgen and Swap respond to arp - Include this in core configurations.</p> <p>sub mode=l3</p> <p>If You want an interface with only latency configured - but still need to respond to ARPs, configure a generator but set bps=0.</p> <p>Need an interface that is configured for both receiving and sending</p> <p>sub mode=l2 is also possible (that's a lower-case L)</p> <p>bps=0 on a core configured as a generator means that the interface will still respond to ARP packets</p> <p><code>\${local_hex_mac1} 00 00 00 00 00 00 08 00 45 00 00 2e 00 01 00 00 40 11 f7 7d \${local_hex_ip2}</code> <code>\${local_hex_ip1} 0b b8 0b b9 00 1a 55 7b</code></p> <p>Source-MAC, Dest-Mac,</p> <p>Copy/paste the source dest MAC and higher addresses into https://hpd.gasmi.net/ for analysis</p> <p>(a rough rule of thumb is the a core using DPDK requires 100 cycles to deal with a packet. so take the CPU clock and divide by 100 to get the Maximum packet header processing rate. the denominator could be higher in practice, Luc indicated it could be 200 or so)</p>
4.	There Is a Korean team working on Benchmarking CN scenarios		<p>https://datatracker.ietf.org/meeting/113/materials/slides-113-bmwg-considerations-for-benchmarking-network-performance-in-containerized-infrastructures-slides-113-bmwg-considerations-for-benchmar-00</p>