

BMWG -
Containerized
Infrastructure
Benchmarking

IETF 106

Nov. 16-17, 2019

Singapore



KJ SUN (Soongsil Univ.)
gomjae@dcn.ssu.ac.kr

Hackathon Plan

- The main goal is to figure out container networking performance impacts by various resource options.

- Related Draft:

Considerations for Benchmarking Network Performance in Containerized Infrastructures

<https://tools.ietf.org/html/draft-dcn-bmwg-containerized-infra>

- Verifying words in draft:

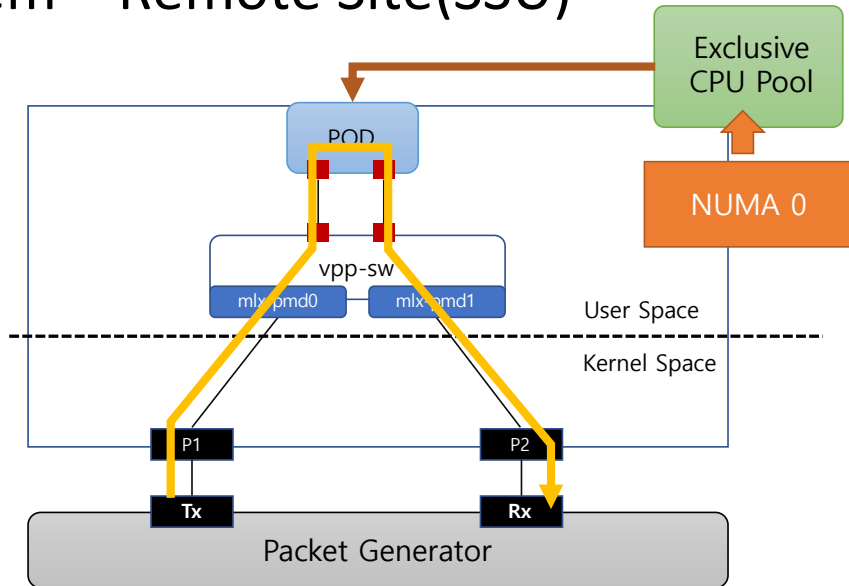
“The instantiation of C-VNFs is somewhat non-deterministic and apparently NUMA-Node agnostic, which is one way of saying that performance will likely vary..”

Hackathon Plan

- Verifying CPU allocation using current container orchestration engine
 - Kubernetes CPU Scheduler (v1.6.1)
 - Compare with CPU pinning technology
 - CMK(CPU Manager for K8S)
- Allocating same number of CPUs to POD
 - measuring throughput

Hackathon Plan

- System – Remote Site(SSU)



[POD]

Suricata

Simple rule to pass incoming traffic to output port

[CPU Pinning]

CMK

Assign dedicated CPU Core in specific NUMA Zone

Kubernetes native

[Container Networking]

DPDK / Contiv-VPP

[Traffic Generator] – Trex

On Bare-metal

Send IMIX traffic (255 clients—255 servers)

What got done

- Unfortunately, we couldn't get expected result..
 - So many issues
 - Install/config – HW/OS/NIC dependencies
 - Base packet throughput/drop were bad
 - Tx: 1/5/10 Gbps -> Rx: ~ 1 Mbps
 - UDP/Ethernet Checksum Error
 - VPP Switch configuration(?)
- Trying to solve problem ASAP – during this week
 - And sharing results to the BMWG

What we learned

- Differences between vSwitch and VNF benchmarking
 - We thought that just forwarding from switch to POD
 - But it's not..!
 - Difficult to find the problem point
- But in other hand, we learned more than simple/quick success
 - Understanding for technologies and configurations
 - K8S, VPP, TRex, ...
- TRex Power consumption
 - Our server was down when running TRex with 100% CPU usage

Wrap Up

Team members:

- KJ Sun, Youngki Park, Hojoon Won, Younghan Kim (SSU)
- Wangbong Lee (ETRI)
- Derrick Lim (Rakuten)
- Somers-Harris David (Rakuten)
- Hogeun Lim (SSU) – Remote

First timers @ IETF/Hackathon:

