

Considerations for Benchmarking Network Performance in Containerized Infrastructure

draft-dcn-bmwg-containerized-infra-03

Kyoungjae Sun, Hyunsik Yang, Youngki Park, Younghan Kim
IISTR, Soong-Sil University

Wangbong Lee

ETRI

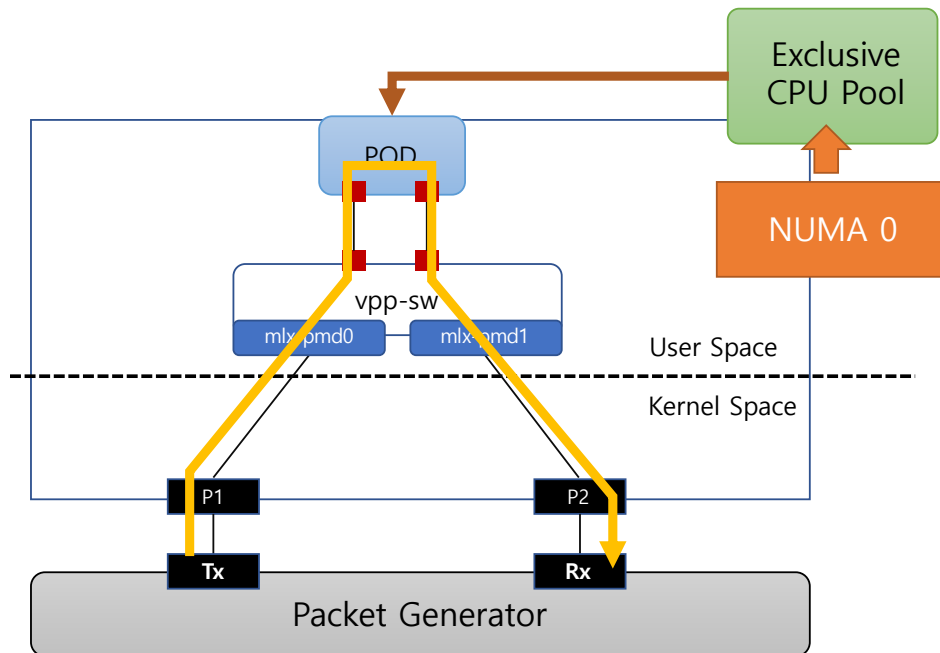
Draft - minor updates

- 3.3. Resource Considerations (more explanation)
 - Huge page
 - In the containerized infrastructure, the container is isolated in the application level and administrators can set huge pages more granular level. - Kubernetes allows to use of 512M bytes huge pages for the container as default values.
 - NUMA (Non-Uniform Memory Access)
 - In containerized infrastructure, it is difficult to expose the NUMA topology to the container and currently, it is hard to guarantee the locality of memory when the container is deployed to host that has multiple NUMA nodes.

from Hackathon

- NUMA / Huge page – Network Performance Impacts
 - Verifying CPU allocation using current container orchestration engine
 - Compare with CPU pinning technology
 - Kubernetes CPU Scheduler (v1.6.1)
 - CMK(CPU Manager for K8S)
 - measuring throughputs
 - k8s sched vs CMK (with CPU variation)
 - Verifying Huge page values
 - measuring throughputs

from Hackathon - Test bed



[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)

Next Step

- Keep trying to update new technologies, resource considerations
- Any comments or feedbacks are welcome
- IETF BMWG Hackathon
 - Proof our draft scenarios and features
 - Sharing results to the BMWG