

draft-mkonstan-nf-service-density-01 (expired)

IETF-109 Online BMWG Meeting

Authors: Maciek Konstantynowicz, Peter Mikus

Presented by: Maciek Konstantynowicz

Draft Status

- No changes since -02 (expired January 9, 2020).
- Discussion at IETF-107 about progressing this work as part of a wider NFV service chains benchmarking standardization effort.
 - But no conclusions / concrete steps. (Apologies in case we missed them.)
 - BMWG position about this problem space?
- Views and comments welcome.

No Changes After IETF-105 and -01, Why?

- Target problem space:
 - Benchmarking network function service topologies (a.k.a. chains) realized in a virtualized compute infrastructure including compute “clouds”.
- The target space is dynamic and continuously evolving:
 - New virtualization and optimization approaches proliferating.
 - Variations of virtual network infrastructure involved in connecting the service topologies: vSwitch, SRIOV / SIOV, SmartNICs, virtual (memory) interfaces.
 - Hard to define a systematic and universal approach to benchmarking clusters of services often mapped to shared physical resources.
- Number of parallel industry standard efforts going after the same space:
 - ETSI, LFN OPNFV, CNTT, CNCF TUG.



Benchmarking Many NF Services on a COTS Server

- Challenges

- Measuring performance of a Network Function applications (NF) running solo on a server is fairly straightforward
- Benchmarking gets much more complex with many NFs sharing server resources in multi-instance, multi-tenant scenarios a.k.a. the "noisy neighbour" problem
- Using virtual topologies within the server to facilitate NF "service chains" further exacerbates the problem

- Proposal


- Define NFV service benchmarking methodology that yields repeatable and portable benchmarking results and use it to aid in deriving deterministic operating range of NFV designs with many NFs/NFV services running on a server



Proposed solution

- Separate out three aspects of NFV design
 - NFV service packet processing
 - Shared virtualization infrastructure
 - Shared host networking infrastructure

- Resource allocation scheme to address “noisy neighbour” aspects
 - Use common practice to start with
 - Processor core pinning
 - Observe NUMA affinity
 - Evolve from there e.g. manage LLC resources



But, is this the right approach ?

THANK YOU !

draft-mkonstan-nf-service-density-01 (expired)

IETF-109 Online BMWG Meeting

Authors: Maciek Konstantynowicz, Peter Mikus

Presented by: Maciek Konstantynowicz