# IETF104-Prague, BMWG

draft-mkonstan-nf-service-density-00

Authors: Maciek Konstantynowicz, Peter Mikus

Presented by: Maciek Konstantynowicz

# Problem

- Challenges
  - Measuring performance of Network Function (NF) running alone on a server is fairly straightforward
  - It gets more complex with many NFs sharing compute server resources in multi-service-instance, multi-tenant environment, a.k.a. "noisy neighbour" problem
  - Using virtual topologies within the server to facilitate "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

# NFV Service Design – with VNFs and CNFs

## Single Compute Node

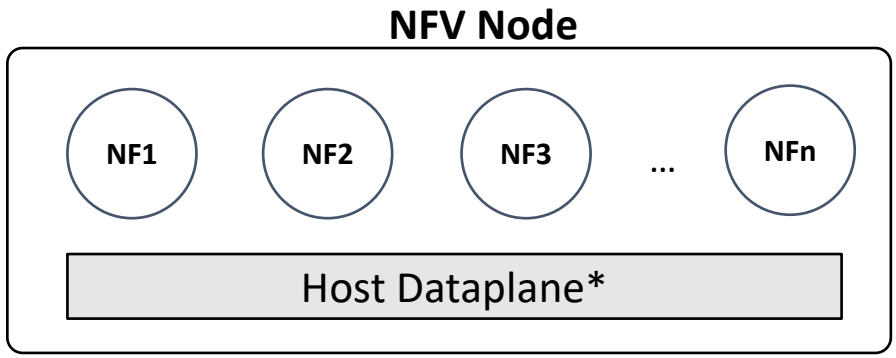


**VNF: NF in VM**
**CNF: NF in Container**

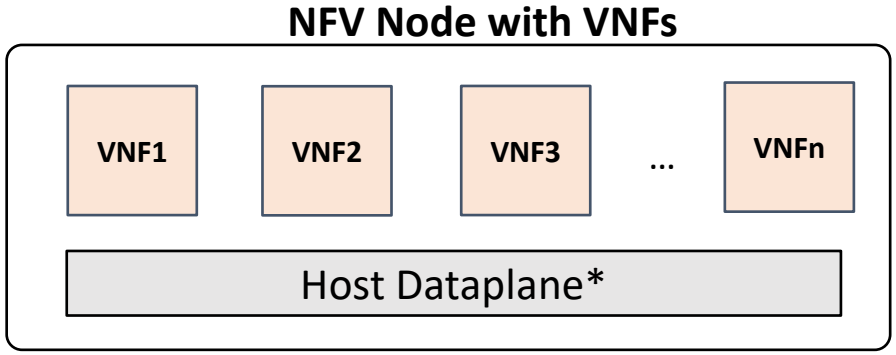


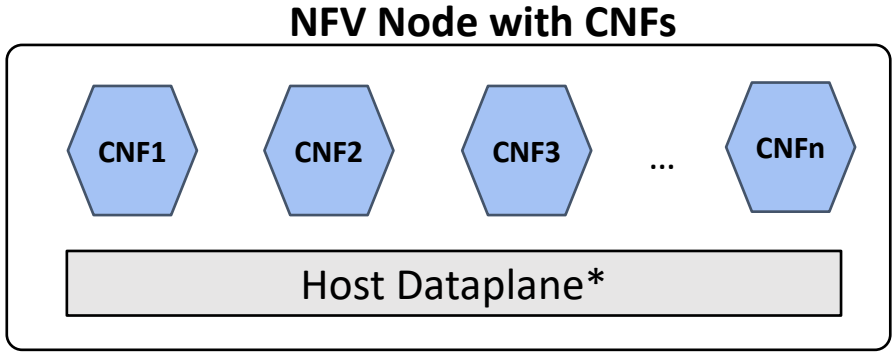


* Linux User-Mode SW Switch, Linux Kernel-Mode Networking, Direct Virtual Function, ..

# Factors impacting NF performance (given a fixed set of resources)

- Main direct factors
  - NF code and packet processing configurations such as routing, rule set of a firewall or IPS - size and complexity of configuration,
  - Virtual topologies of multiple interconnected NFs, and their quantity (density)
  - Packet traffic profile including distributions of flow arrivals, flow sizes, packet sizes and packet interarrival times
- Plus two indirect factors
  - Virtualization technology - workload scheduling and virtual packet interfaces efficiency
  - Host networking - access to shared physical NIC resources

# 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

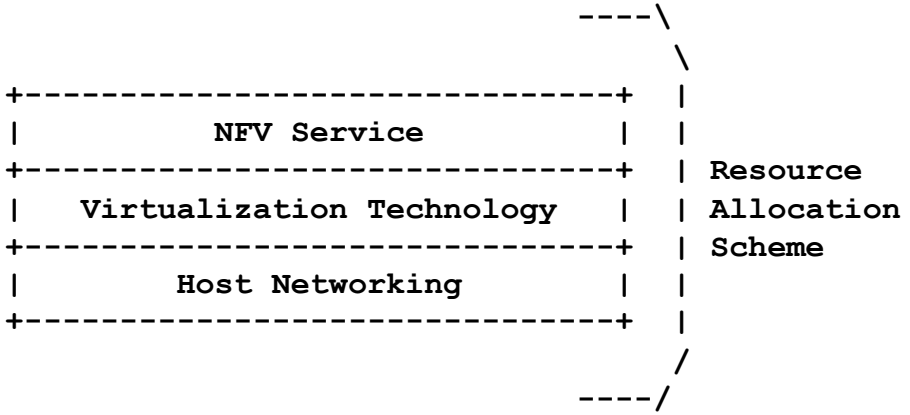

# NFV Service Abstraction

**Packet Path(s)**

How are the network functions and network devices forwarding packets

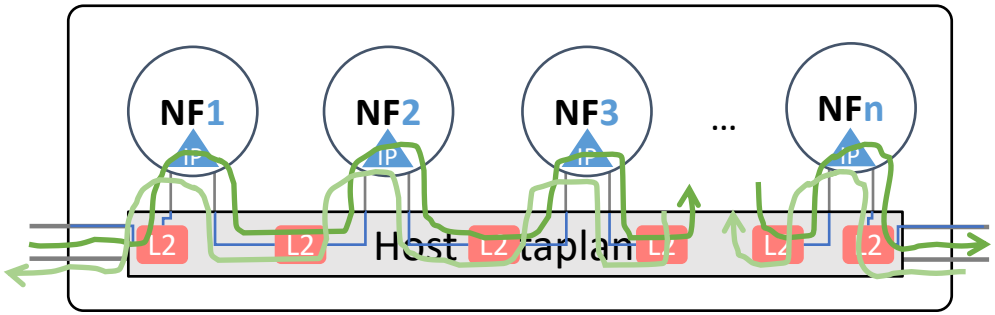


**Configuration**

How are the network functions and network devices configured

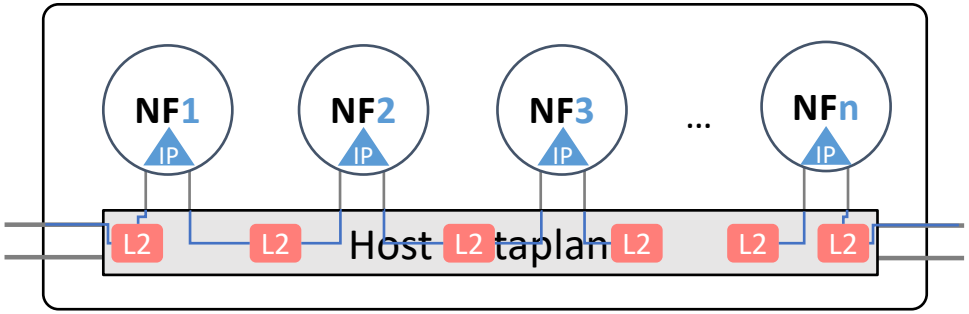


**Topology**

How are the network functions and network devices interconnected

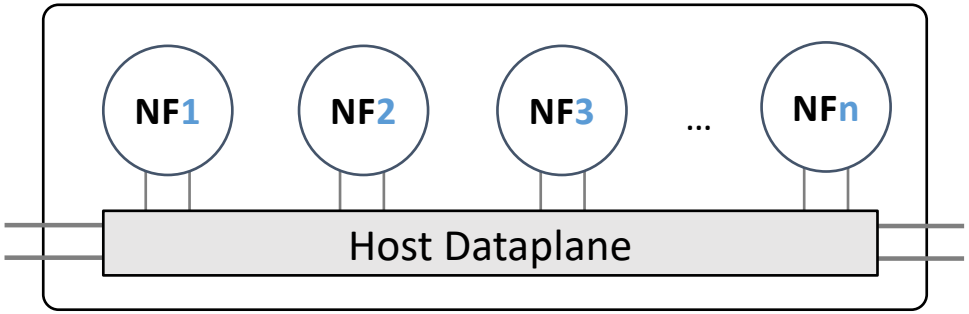

# Test Scenarios: Defined Topologies, Configuration and Packet Paths - So Far..

**Topology & Configuration** — **V**NF **S**ervice **C**hain (**VSC**)

**Packet Path** — Snake Forwarding

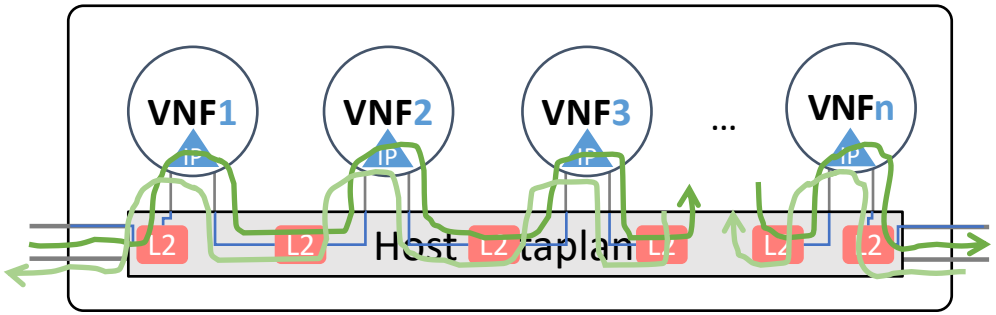


**Topology & Configuration** — **C**NF **S**ervice **C**hain (**CSC**)

**Packet Path** — Snake Forwarding

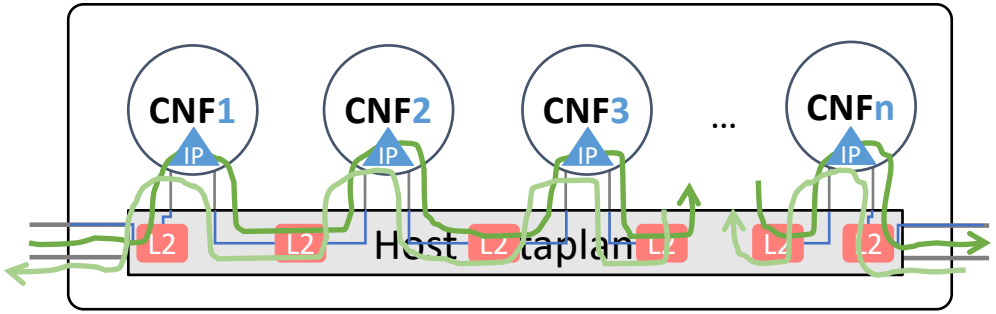


**Topology & Configuration** — **C**NF **S**ervice **P**ipeline (**CSP**)

**Packet Path** — Pipeline Forwarding

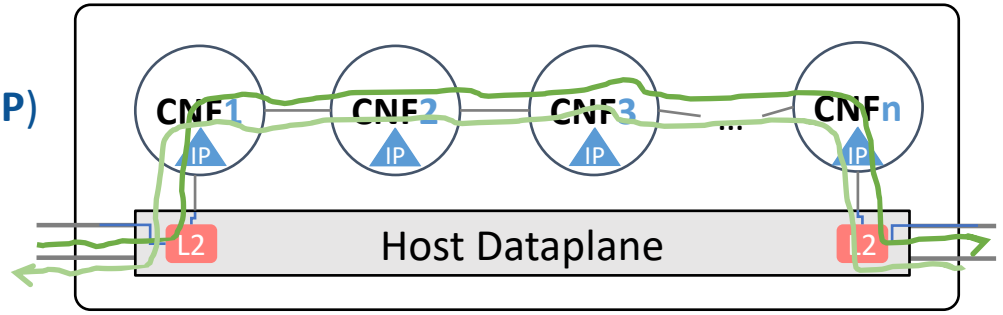

# Composing Service Density Matrix: NF Count View

| SVC | 001 | 002 | 004 | 006 | 008 | N |
|---|---|---|---|---|---|---|
| 001 | 1 | 2 | 4 | 6 | 8 | 1∗N |
| 002 | 2 | 4 | 8 | 12 | 16 | 2∗N |
| 004 | 4 | 8 | 16 | 24 | 32 | 4∗N |
| 006 | 6 | 12 | 24 | 36 | 48 | 6∗N |
| 008 | 8 | 16 | 32 | 48 | 64 | 8∗N |
| M | M∗1 | M∗2 | M∗4 | M∗6 | M∗8 | M∗N |

RowIndex: Number of NFV Service Instances, 1..M.
ColumnIndex: Number of NFs per NFV Service Instance, 1..N.
Value: Total number of NFs running the system.

# Resource Allocation – Processor Cores

## Methodology

| application | thread type | app:core ratio | threads/pcores (SMT disabled) | threads/lcores map (SMT enabled) |
|---|---|---|---|---|
| vSwitch-1c | data | 1:1 | 1DT/1PC | 2DT/2LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| vSwitch-2c | data | 1:2 | 2DT/2PC | 4DT/4LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| vSwitch-4c | data | 1:4 | 4DT/4PC | 8DT/8LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| NF-0.5c | data | 1:S2 | 1DT/S2PC | 1DT/1LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| NF-1c | data | 1:1 | 1DT/1PC | 2DT/2LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| NF-2c | data | 1:2 | 2DT/2PC | 4DT/4LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |

**Legend to table:**

- **Sn** - shared core, sharing ratio of (n).
- **DT** - data-plane thread.
- **MT** - main-control thread.
- **PC** - physical core, with SMT/HT enabled has many (mostly 2 today) logical cores associated with it.
- **LC** - logical core, if more than one lc get allocated in sets of two sibling logical cores running on the same physical core.
- **SnPC** - shared physical core, sharing ratio of (n).

# NFV Density
# Sample Measurements

# Resource Allocation – Processor Cores

## Methodology

**Tested Core Allocation: vSwitch-{1c,2c}, NF-1c**

| application | thread type | app:core ratio | threads/pcores (SMT disabled) | threads/lcores map (SMT enabled) |
|---|---|---|---|---|
| **vSwitch-1c** | **data** | **1:1** | 1DT/1PC | **2DT/2LC** |
| | **main** | **1:S2** | 1MT/S2PC | **1MT/1LC** |
| **vSwitch-2c** | **data** | **1:2** | 2DT/2PC | **4DT/4LC** |
| | **main** | **1:S2** | 1MT/S2PC | **1MT/1LC** |
| vSwitch-4c | data | 1:4 | 4DT/4PC | 8DT/8LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| NF-0.5c | data | 1:S2 | 1DT/S2PC | 1DT/1LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |
| **NF-1c** | **data** | **1:1** | 1DT/1PC | **2DT/2LC** |
| | **main** | **1:S2** | 1MT/S2PC | **1MT/1LC** |
| NF-2c | data | 1:2 | 2DT/2PC | 4DT/4LC |
| | main | 1:S2 | 1MT/S2PC | 1MT/1LC |

**Legend to table:**

- **Sn** - shared core, sharing ratio of (n).
- **DT** - data-plane thread.
- **MT** - main-control thread.
- **PC** - physical core, with SMT/HT enabled has many (mostly 2 today) logical cores associated with it.
- **LC** - logical core, if more than one lc get allocated in sets of two sibling logical cores running on the same physical core.
- **SnPC** - shared physical core, sharing ratio of (n).

# Service Density Matrix: Sample Core Usage View

**Tested Core Allocation: vSwitch-{1c,2c}, NF-1c**

| SVC | 001 | 002 | 004 | 006 | 008 | 010 |
|---|---|---|---|---|---|---|
| 001 | 2 | 3 | 6 | 9 | 12 | 15 |
| 002 | 3 | 6 | 12 | 18 | 24 | 30 |
| 004 | 6 | 12 | 24 | 36 | 48 | 60 |
| 006 | 9 | 18 | 36 | 54 | 72 | 90 |
| 008 | 12 | 24 | 48 | 72 | 96 | 120 |
| 010 | 15 | 30 | 60 | 90 | 120 | 150 |

Not tested due to restricted core count per NUMA

(SUT = SuperMicro with 2x Intel Xeon 8180 28c 2.5GHz)

RowIndex: Number of NFV Service Instances, 1..10.
ColumnIndex: Number of NFs per NFV Service Instance, 1..10.
Value: Total number of physical processor cores used for NFs.

# Service Density Matrix: MRR Throughput Results

vSwitch-1c

**VNF Service Chains (VSC)**

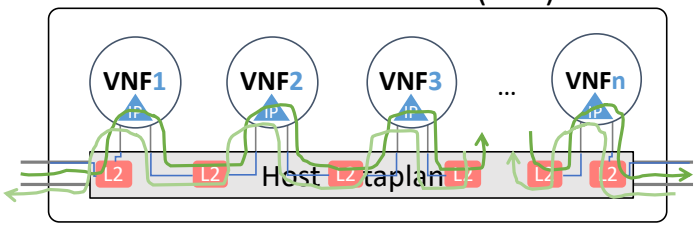


| | |
|---|---|
| NF-1c | DPDK-L3FWD v18.11 |
| vSwitch-1c | VPP v19.01-release |

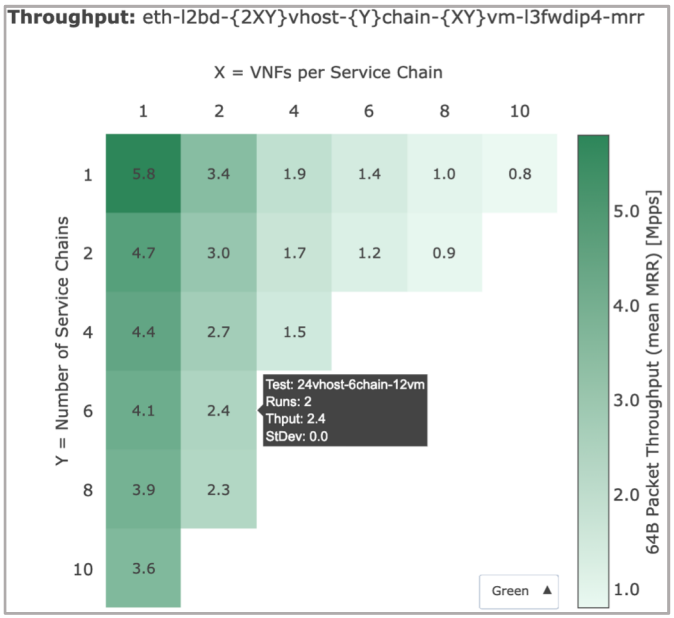


**CNF Service Chains (CSC)**

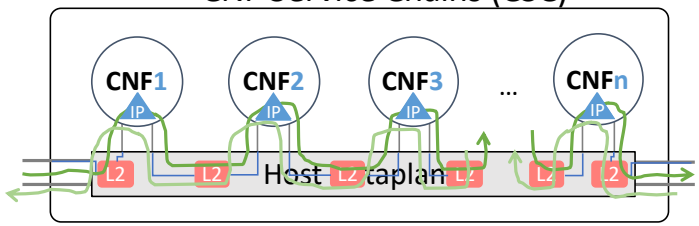


| | |
|---|---|
| NF-1c | VPP v19.01-release |
| vSwitch-1c | VPP v19.01-release |

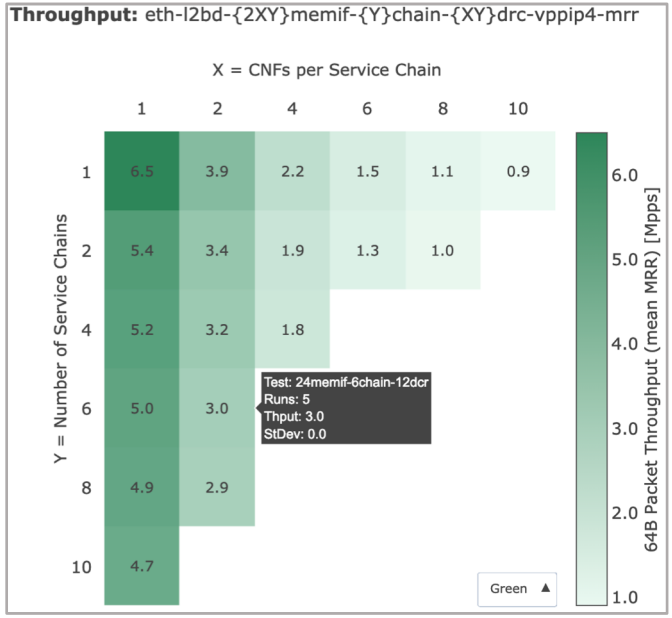


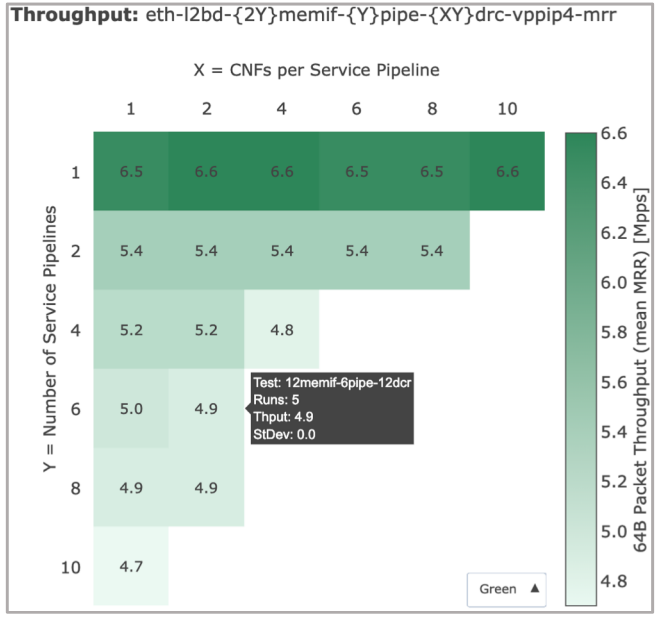


Source: https://docs.fd.io/csit/rls1901/report/vpp_performance_tests/nf_service_density/index.html

# Service Density Matrix: MRR Throughput Results

vSwitch-2c

**VNF Service Chains (VSC)**

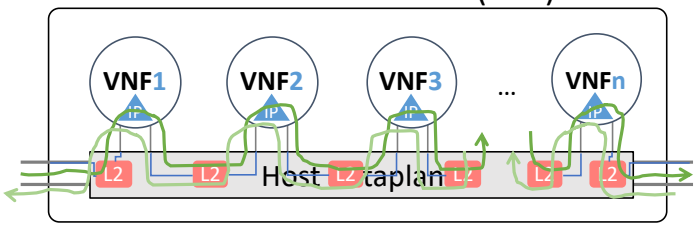


| | |
|---|---|
| NF-1c | DPDK-L3FWD v18.11 |
| vSwitch-2c | VPP v19.01-release |

**CNF Service Chains (CSC)**

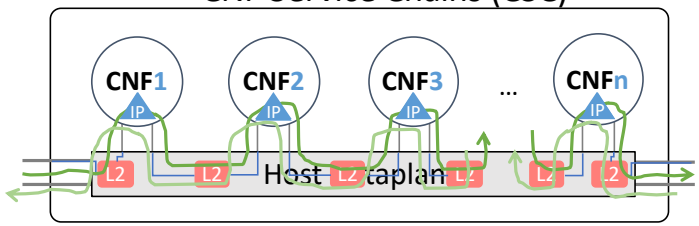


| | |
|---|---|
| NF-1c | VPP v19.01-release |
| vSwitch-2c | VPP v19.01-release |

**CNF Service Pipelines (CSP)**

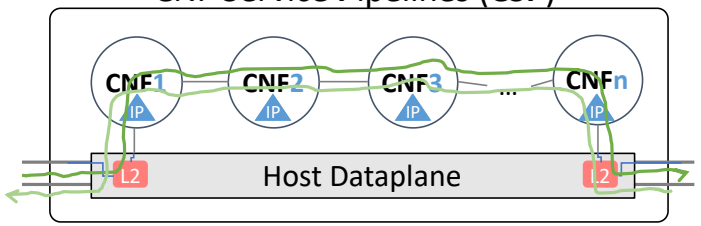


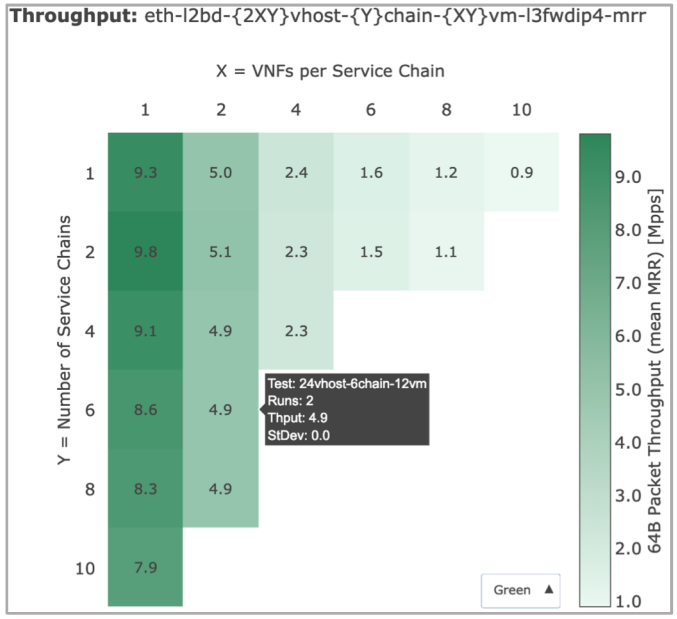


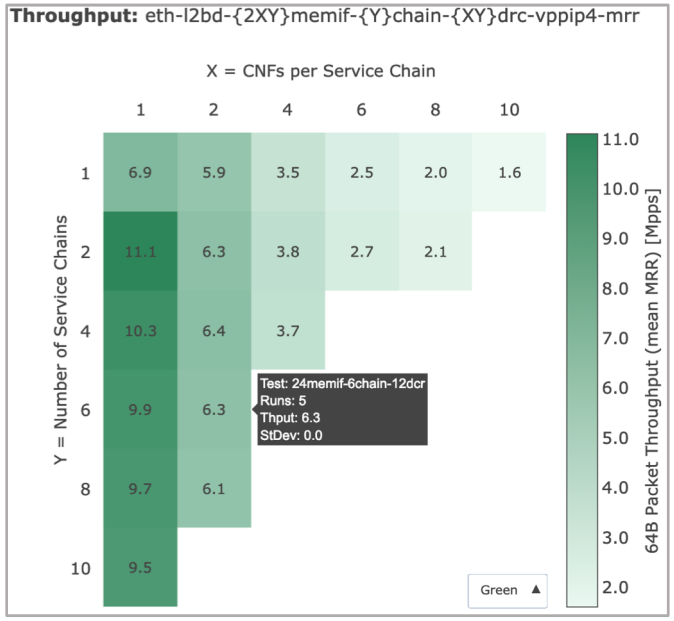


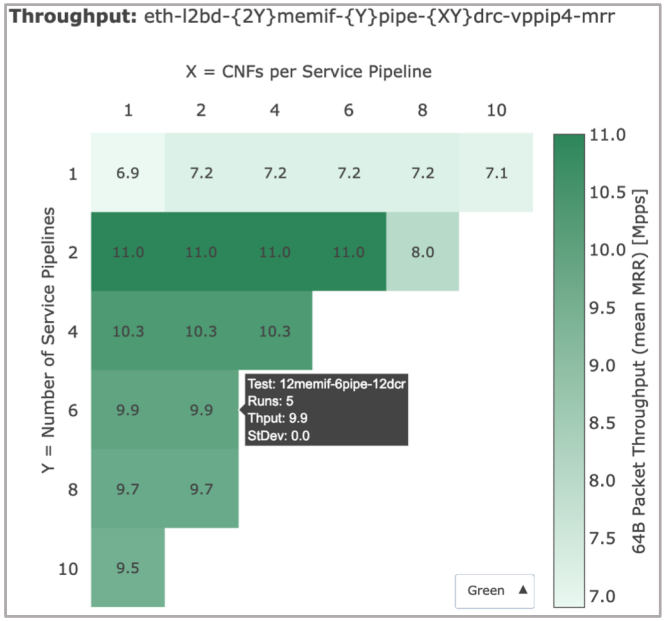


Source: https://docs.fd.io/csit/rls1901/report/vpp_performance_tests/nf_service_density/index.html

# Next Steps

- Review comments
- BMWG workgroup adoption