5G Core Network Load Test System with DPDK

Xiaohua Wang @ DPDK China 2019

Content

- Background
  - 5GC and CUPS
  - Extreme high throughput, Ultra low latency
- Challenges
  - Measure and verify 5GC performance
- Load Testing System for 5GC
- Payload Engine with DPDK
- Experiences of developing with DPDK
  - Tips on performance
  - DPDK in docker with SR-IOV
  - DPDK 18.11 on Intel and Mellanox NIC

5GC and CUPS

- CP
  - SBA
  - SBI
    - HTTP2
    - Restful
- UP
  - GTP-U

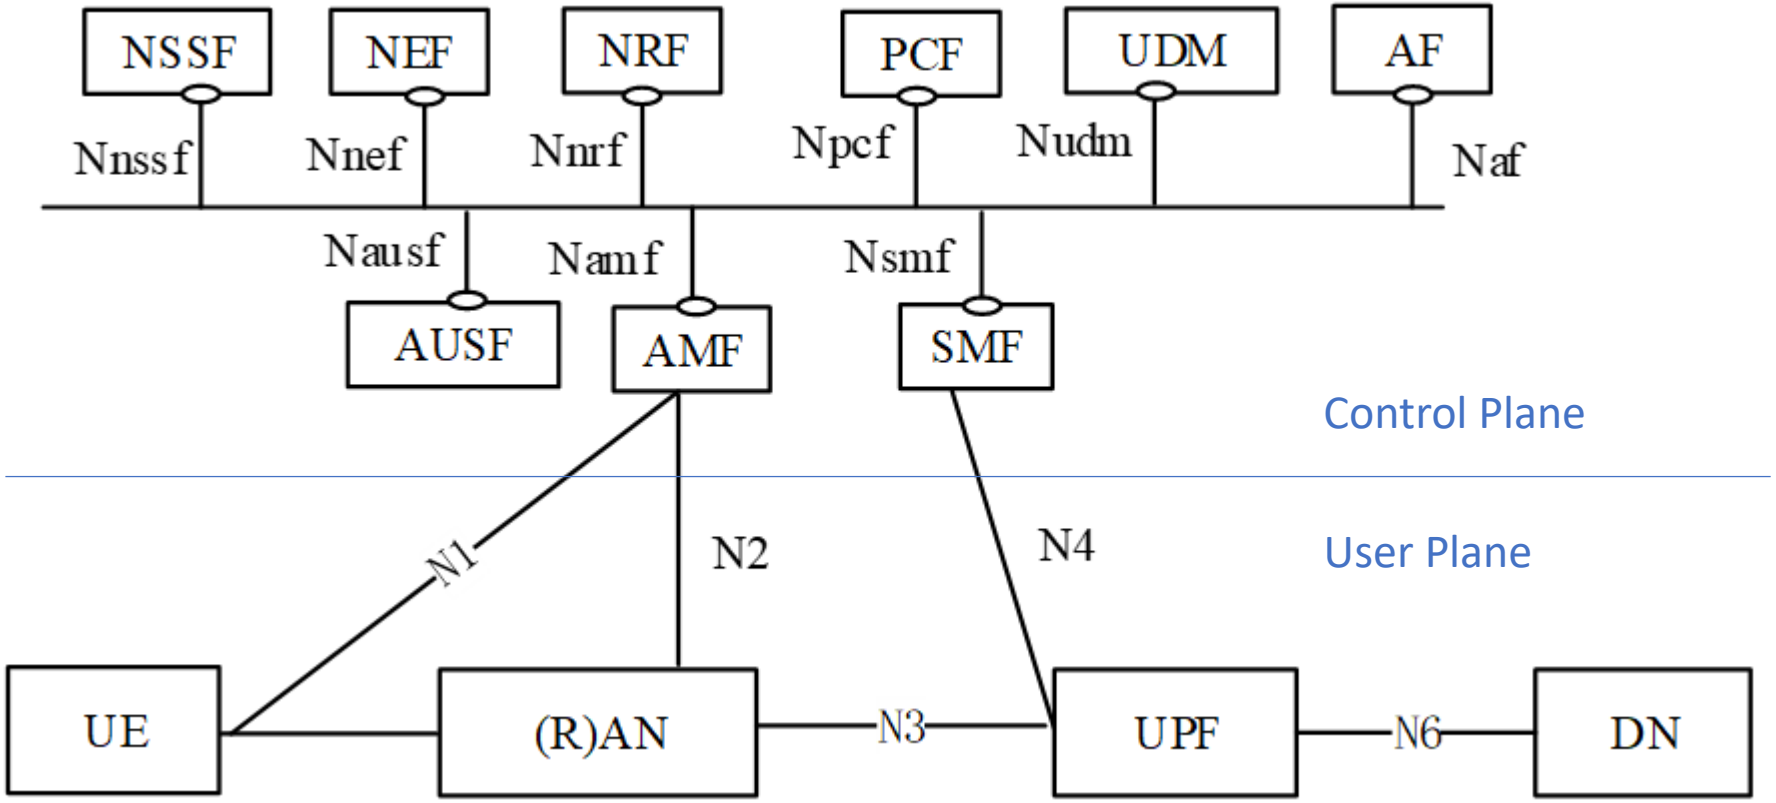

Extreme high throughput, Ultra low latency

- UP Peak Data Rate:
  - Downlink: 20 Gbit/s
  - Uplink: 10 Gbit/s
- UP Latency:
  - 4 ms for eMBB,
  - 1 ms for URLLC

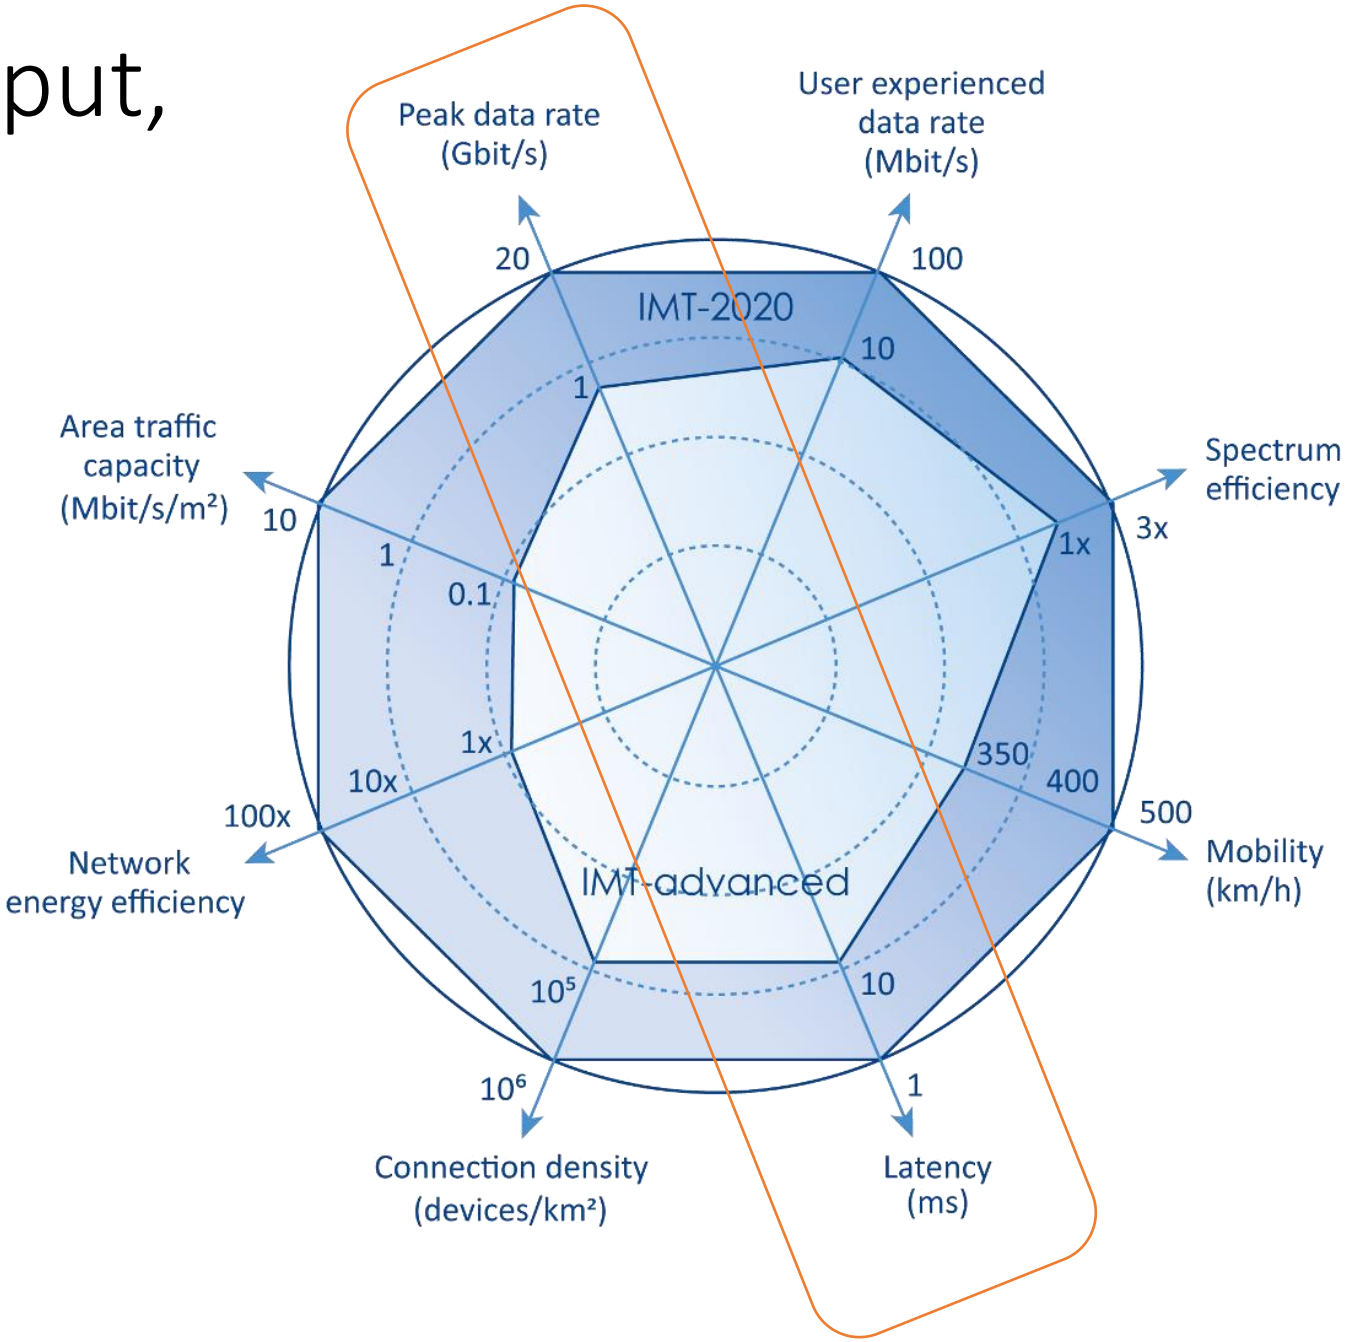

Challenges

- How to measure and verify UP?
  - Extreme high throughput
  - Ultra low latency
- Complex Service and Traffic Model
  - Interaction with control plane

Performance requirements for high data rate and traffic density scenarios.

	Scenario	Experienced data rate (DL)	Experienced data rate (UL)	Area traffic capacity (DL)	Area traffic capacity (UL)	Overall user density	Activity factor	UE speed	Coverage
1	Urban macro	50 Mbps	25 Mbps	100 Gbps/km^2 (note 4)	50 Gbps/km^2 (note 4)	10 000/km^2	20%	Pedestrians and users in vehicles (up to 120 km/h	Full network (note 1)
2	Rural macro	50 Mbps	25 Mbps	1 Gbps/km^2 (note 4)	500 Mbps/km^2 (note 4)	100/km^2	20%	Pedestrians and users in vehicles (up to 120 km/h	Full network (note 1)
3	Indoor hotspot	1 Gbps	500 Mbps	15 Tbps/km^2	2 Tbps/km^2	250 000/km^2	note 2	Pedestrians	Office and residential (note 2) (note 3)
4	Broadband access in a crowd	25 Mbps	50 Mbps	[3,75] T bps/km^2	[7,5] T bps/km^2	[500 000]/km^2	30%	Pedestrians	Confined area
5	Dense urban	300 Mbps	50 Mbps	750 Gbps/km^2 (note 4)	125 Gbps/km^2 (note 4)	25 000/km^2	10%	Pedestrians and users in vehicles (up to 60 km/h)	Downtown (note 1)
6	Broadcast-like services	Maximum 200 Mbps (per TV channel)	N/A or modest (e.g., 500 kbps per user)	N/A	N/A	[15] TV channels of [20 Mbps] on one carrier	N/A	Stationary users, pedestrians and users in vehicles (up to 500 km/h)	Full network (note 1)
7	High-speed train	50 Mbps	25 Mbps	15 Gbps/train	7,5 Gbps/train	1 000/train	30%	Users in trains (up to 500 km/h)	Along railways (note 1)
8	High-speed vehicle	50 Mbps	25 Mbps	[100] Gbps/km^2	[50] Gbps/km^2	4 000/km^2	50%	Users in vehicles (up to 250 km/h)	Along roads (note 1)
9	Airplanes connectivity	15 Mbps	7,5 Mbps	1,2 Gbps/plane	600 Mbps/plane	400/plane	20%	Users in airplanes (up to 1 000 km/h)	(note 1)

Load Testing System for 5GC

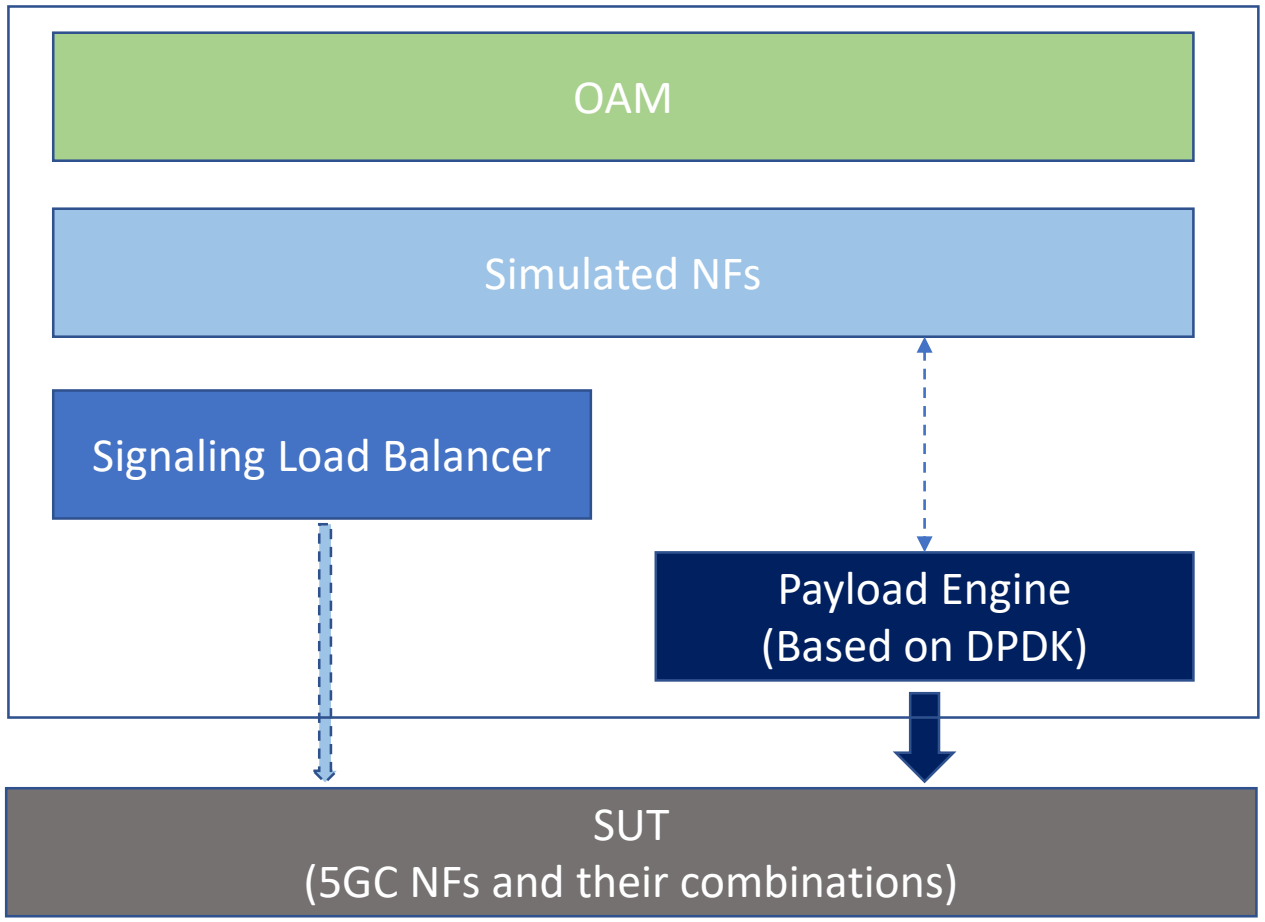

- Support to test towards all 5GC core NFs and their combinations
- Support to work on bare metal, VM and docker
- NF: network function

Payload Engine (PE), Fast Path for user plane

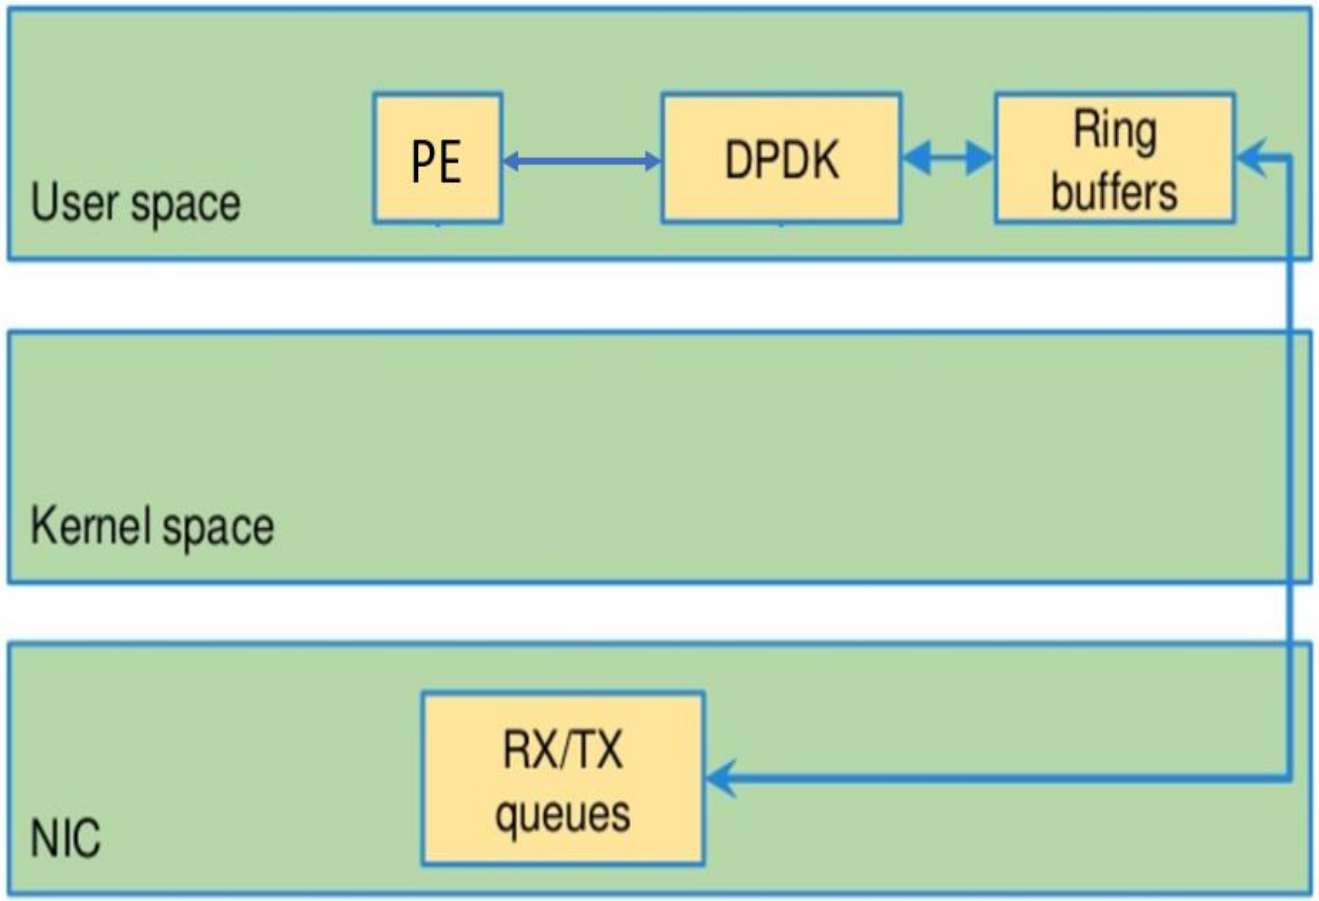

- Developed based on DPDK
- Support to test UP in 5GC
- Support to test throughput and latency from 5GC e2e

Slow Path using TAP device for control plane

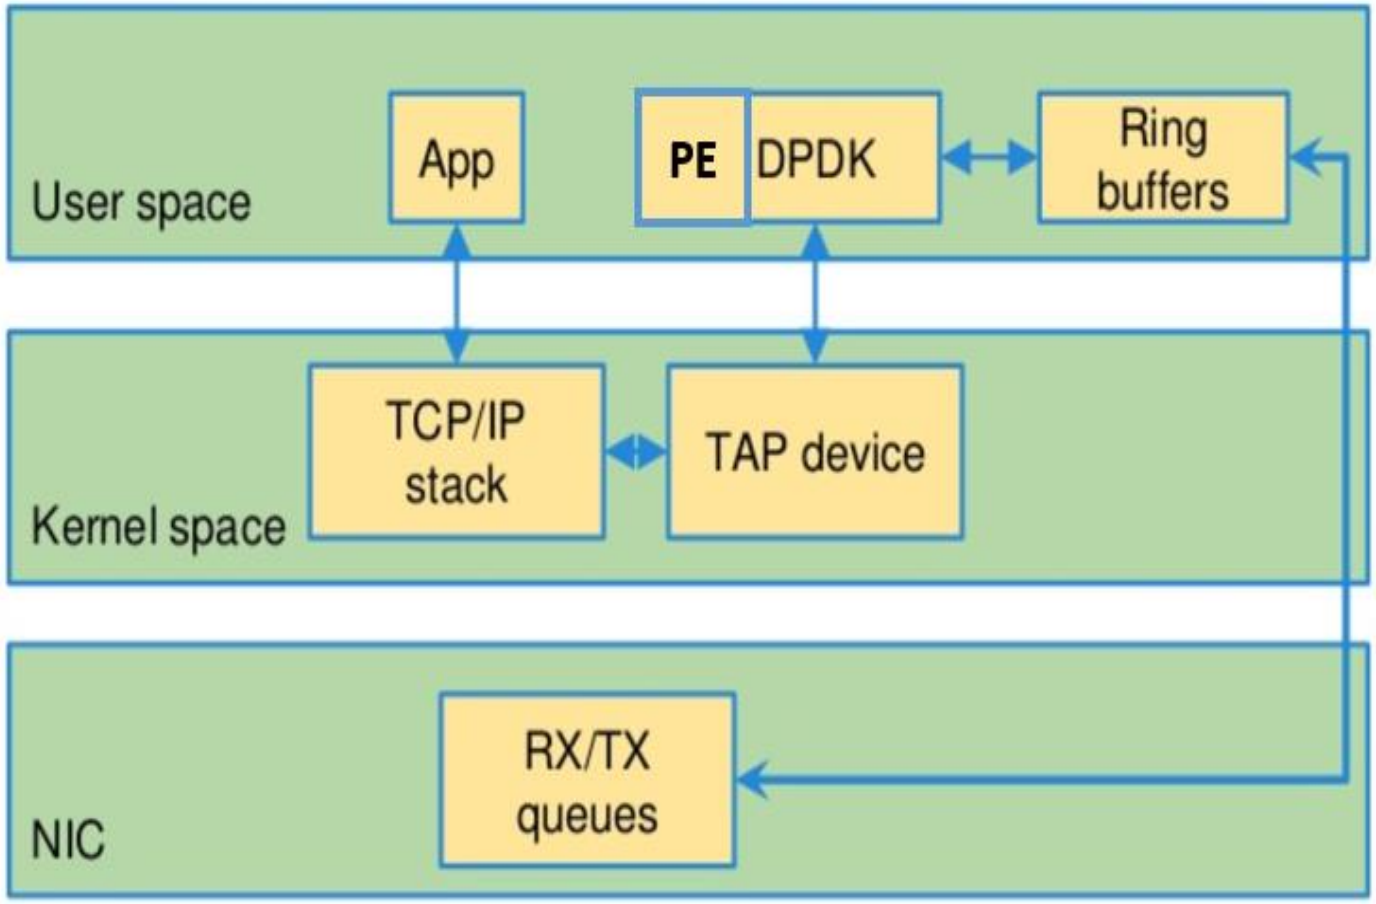

- Support to test control plane and signaling (GTP-C/TCP/SCTP based protocol) collocated with payload engine developed with DPDK

Experiences of developing with DPDK

- **Tips on performance**
- DPDK in docker with SR-IOV
- DPDK 18.11 on Intel and Mellanox NIC

Tips on performance

NUMA	Global Lock	Prefetch

• Memory pool
• Hash table
• Rte_ring
• Rte_malloc
	• Design for localization
	  • Allocation on stack
	  • Per-core data structure
	• Minimize lock scope
	• Sperate reading and writing
		• Used on array usually
		• Consider the data size

Experiences of developing with DPDK

- Tips on performance
- **DPDK in docker with SR-IOV**
- DPDK 18.11 on Intel and Mellanox NIC

Networking

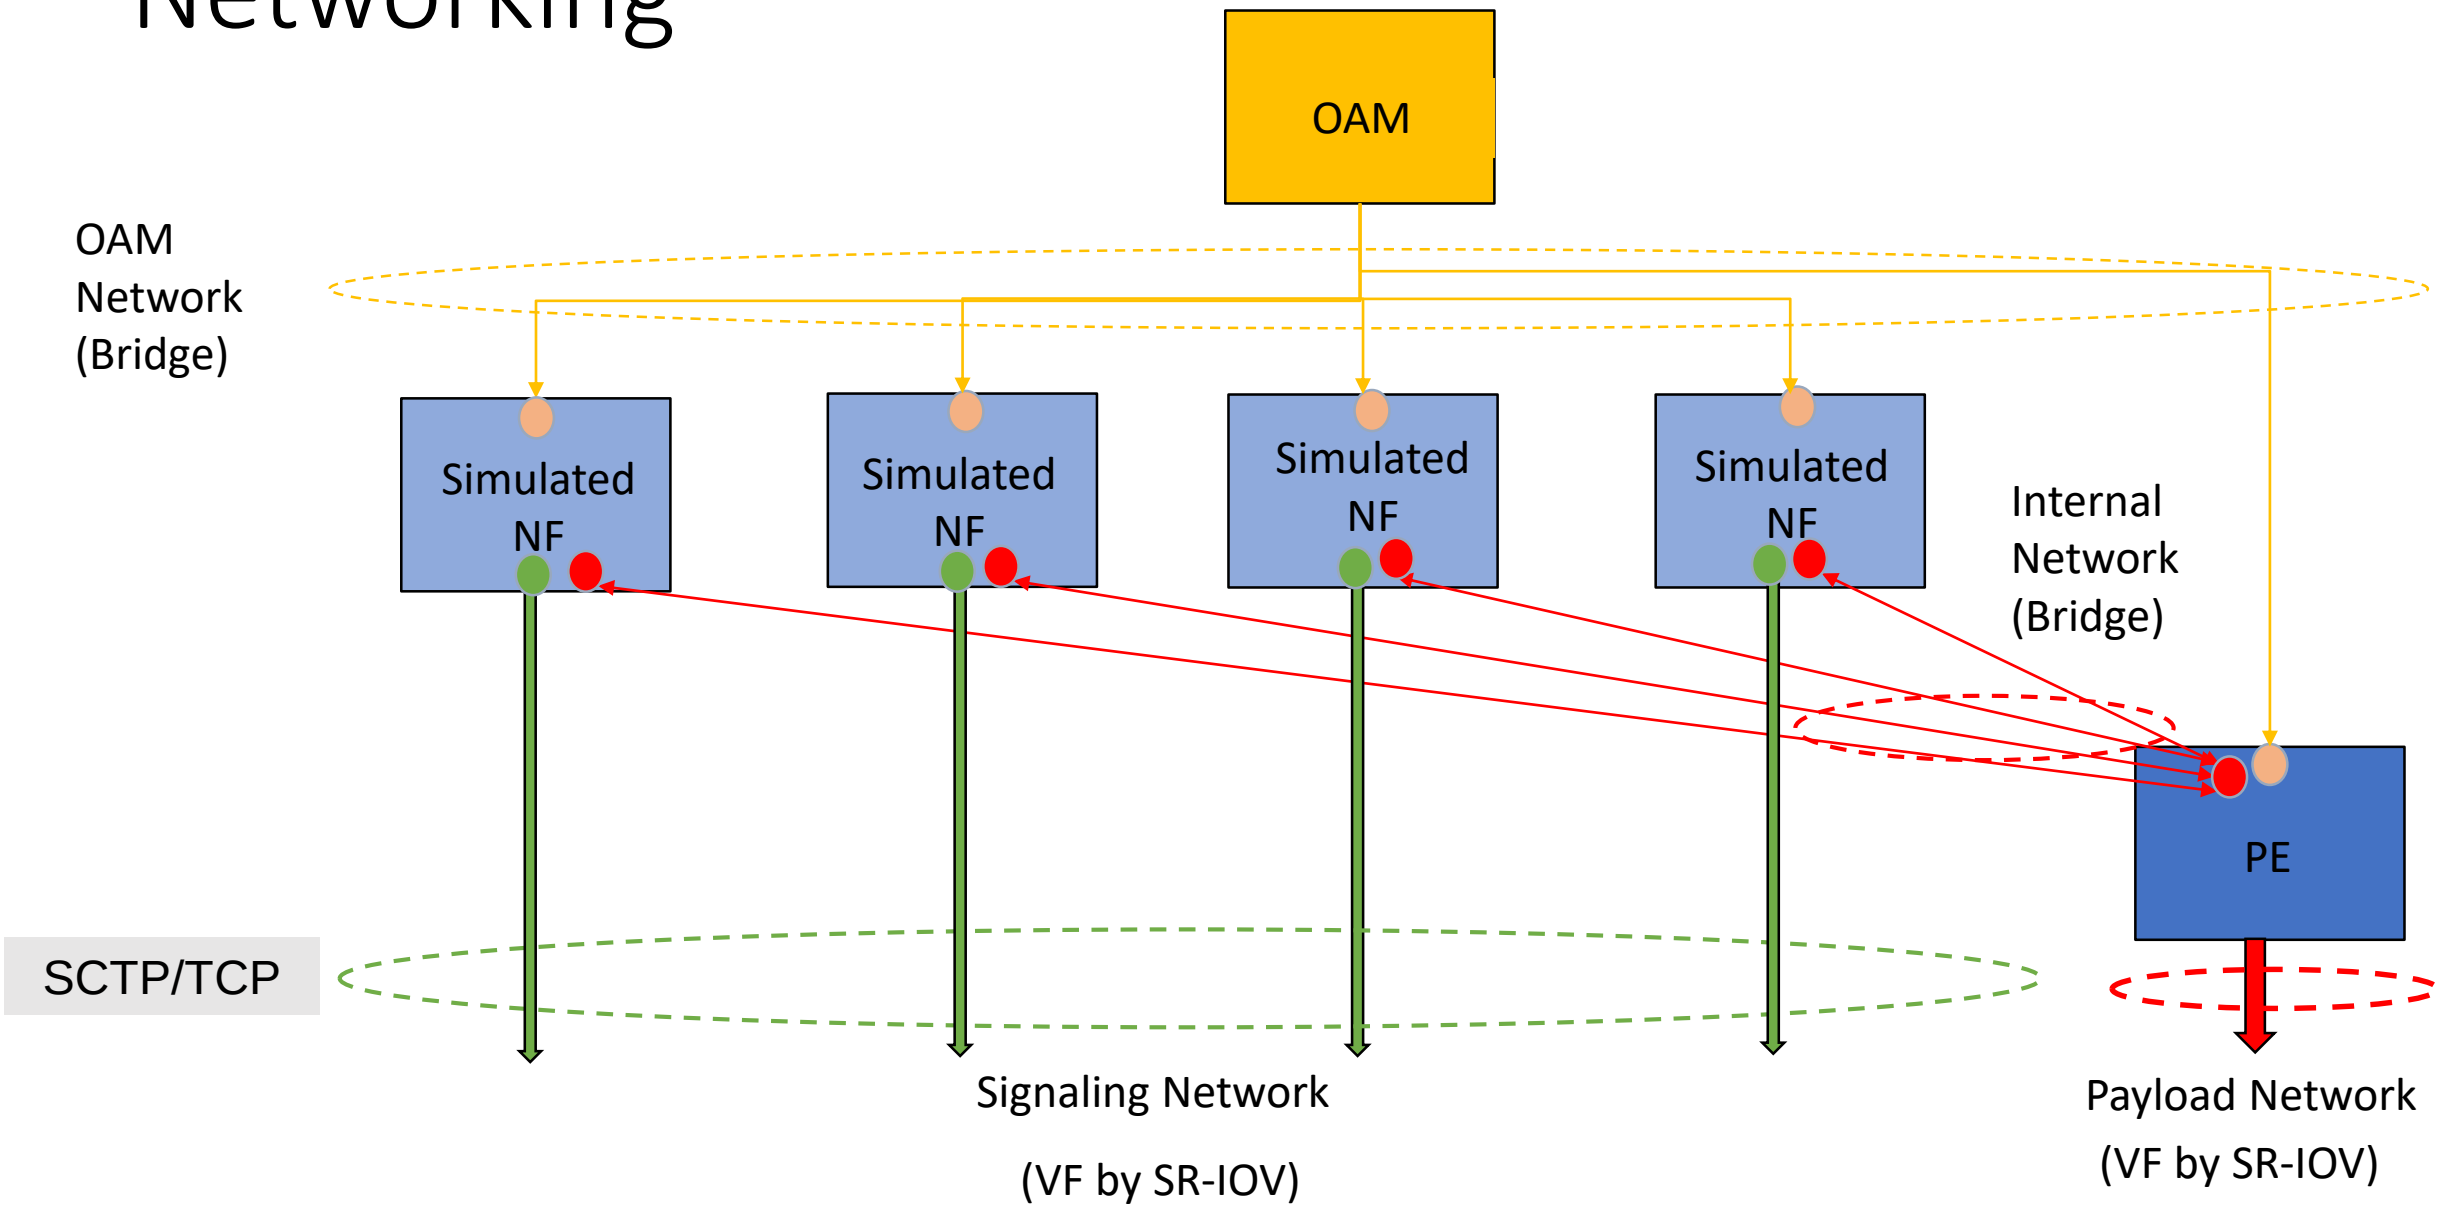
OAM
OAM Network (Bridge)
Simulated NF
Simulated NF
Simulated NF
Simulated NF
Internal Network (Bridge)
PE
SCTP/TCP
Signaling Network
(VF by SR-IOV)
Payload Network
(VF by SR-IOV)

Tips for SR-IOV

- On Host (NIC: Intel XL710)
  - dpdk-devbind -b i40e "0000:81:00.0“
  - ip link set dev p1p1 up
  - echo 0 > /sys/bus/pci/devices/"0000:81:00.0"/sriov_numvfs
  - echo 3 > /sys/bus/pci/devices/"0000:81:00.0"/sriov_numvfs
  - dpdk-devbind --bind=igb_uio "0000:81:02.0" "0000:81:02.1" "0000:81:02.2“

  - VF mac address
    - Need to set mac address to enable VF to route by mac address
  - VLAN tags
    - Need to take care of VLAN tag in DPDK or on VF setting

- Docker container
  - privileged: true
  - volumes: - /dev/hugepages:/dev/hugepages

Experiences of developing with DPDK

- Tips on performance
- DPDK in docker with SR-IOV
- DPDK 18.11 on Intel and Mellanox NIC

DPDK 18.11 on Intel and Mellanox NIC

- Intel XL710 (40Gbps)
  - For high performance, to use **“Legacy memory mode”** instead of **“Dynamic memory mode”**
  - Enabled by specifying --legacy-mem command-line switch to the EAL
- Mellanox ConnectX-5 (100Gpbs)
  - Can’t work properly with legacy mode, need to use **“Dynamic memory mode”** instead of **“Legacy memory mode”**
  - But introduces some performance degradation
  - Enabled by default on Linux

Q & A