Next Gen Infrastructure Core (NGIC) & M-CORD

Christian Maciocco

Principal Engineer, Director of Telecom Systems Research – Intel Labs

Ashok Sunder Rajan, Kannan Babu Ramia, Sai Edupuganti, Jacob Cooper, Karla Saur – Intel Labs

CORD Build Event - November 7-9th, 2017 – San Jose, CA

Intel Labs

Legal Disclaimer

- This presentation contains the general insights and opinions of Intel Corporation (“Intel”). The information in this presentation is provided for information only and is not to be relied upon for any other purpose than educational. Use at your own risk! Intel makes no representations or warranties regarding the accuracy or completeness of the information in this presentation. Intel accepts no duty to update this presentation based on more current information. Intel is not liable for any damages, direct or indirect, consequential or otherwise, that may arise, directly or indirectly, from the use or misuse of the information in this presentation.
- Intel technologies’ features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.
- No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.
- Intel, the Intel logo and Xeon are trademarks of Intel Corporation in the United States and other countries.
- *Other names and brands may be claimed as the property of others.
-

Agenda

- vEPC Background
- Next Generation Infrastructure Core (NGIC)
- CORD High Level View
- NGIC and M-CORD
- Summary / Next Steps

vEPC Background (1/2)

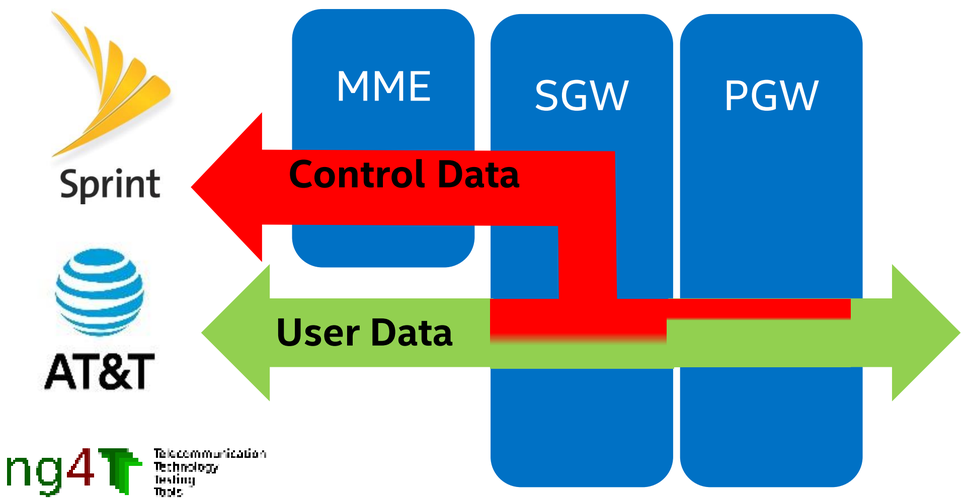

Dimension	Parameter	Description	Value
Control Plane	λ_C	NAS event arrival rate	6400 events/sec
	T_C	Total time event spends in the system (Service + Wait)	156.3 µsec
	S_{TWC}	Average event wait time	42.9 µsec
User Plane	λ_U	Packet arrival rate	5.533 MPPS
	T_U	Total time packet spends in the system (Service + Wait)	181 ηsec
	S_{TWU}	Average packet wait time	29 ηsec

Control & User Plane Simulation Parameters

- Operators' real traffic (San Jose, Houston, Chicago, …)
- Identified system's bottleneck
  - "Understanding Bottlenecks in Virtualizing Cellular Core Network functions", IEEE LANMAN '15
- No independent control or data plane scaling

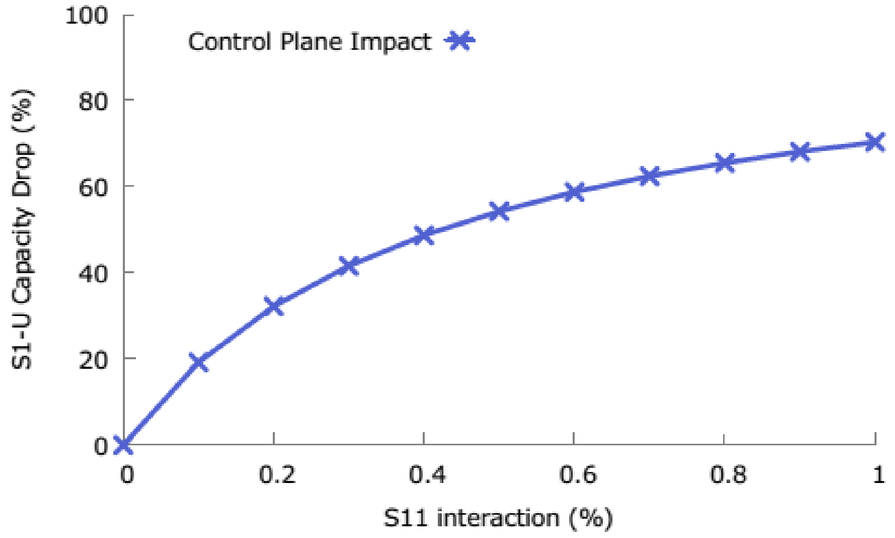

User plane capacity reduction as control plane interference increases

vEPC Background (2/2)

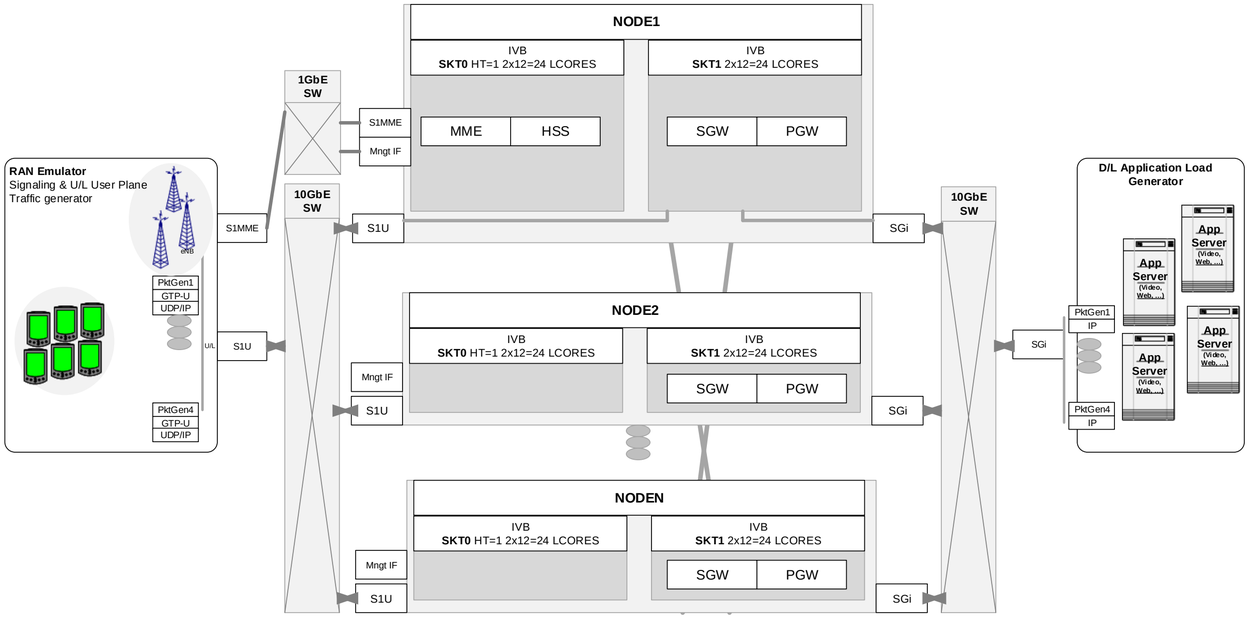

EPC Setup

Direction	Traffic type	Traffic mix %	Pkt size (bytes)	Pkt rate (pkt/sec)
Downlink	VoLTE	5.1	72	2.1M
	Web	52.2	1200	1.3M
	Video	29.5	1440	602K
	Apps	6.7	675	290K
	Email/other	6.7	1440	136K
Uplink	VoLTE	32.3	72	2.1M
	Web	43.1	690	289K
	Video	12.1	240	232K
	Apps	6.5	400	76K
	Email/other	6.2	1000	29K

User Plane Traffic Model

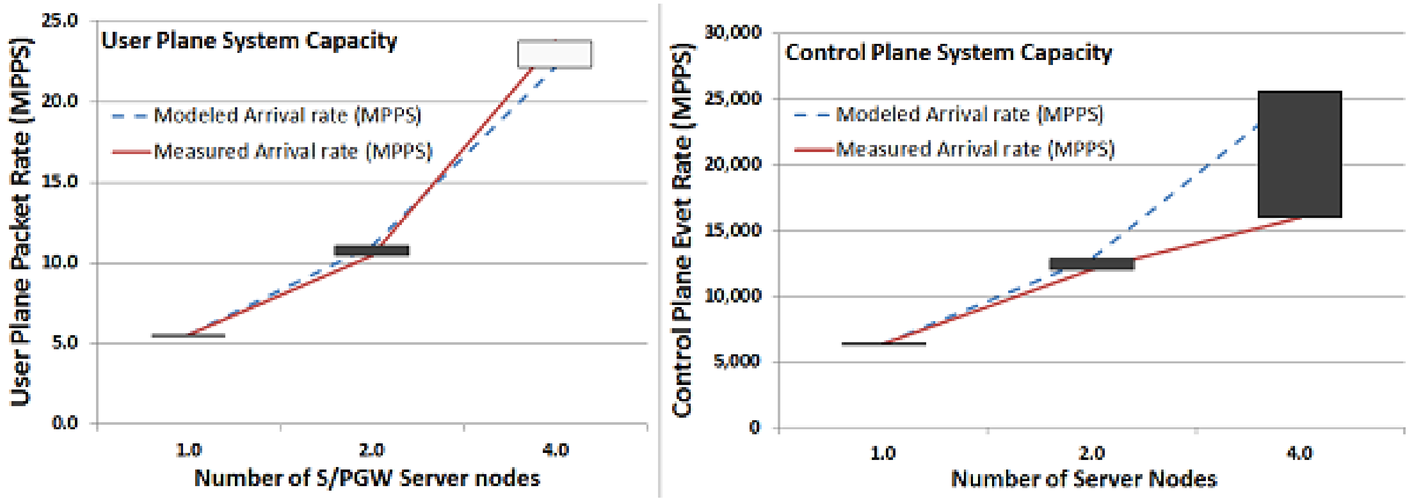

User plane and control plane queue flooding

Agenda

- vEPC Background
- Next Generation Infrastructure Core (NGIC)
- CORD High Level View
- NGIC and M-CORD
- Summary / Next Steps

Next Generation Infrastructure Core (NGIC)

What is it ? Why ?

- Research and collaborations to understand telecom workloads behavior on Intel Architecture
  - Research to address software/hardware platform improvements
  - Released learnings and optimized reference software into open source community (DPDK, CORD)
    - S/P-GW, Cuckoo Hash for optimized lookup from collaborations with industry or academic partners
    - MME, HSS, PCRF will be released later

  - Investigate functionality for new usage models, e.g. Connectionless IOT, Multi-Radio Access Technologies, etc

- NGIC is not a product and will not be a product from Intel

https://gerrit.opencord.org/#/admin/projects/ngic

Next Generation Infrastructure Core (NGIC)

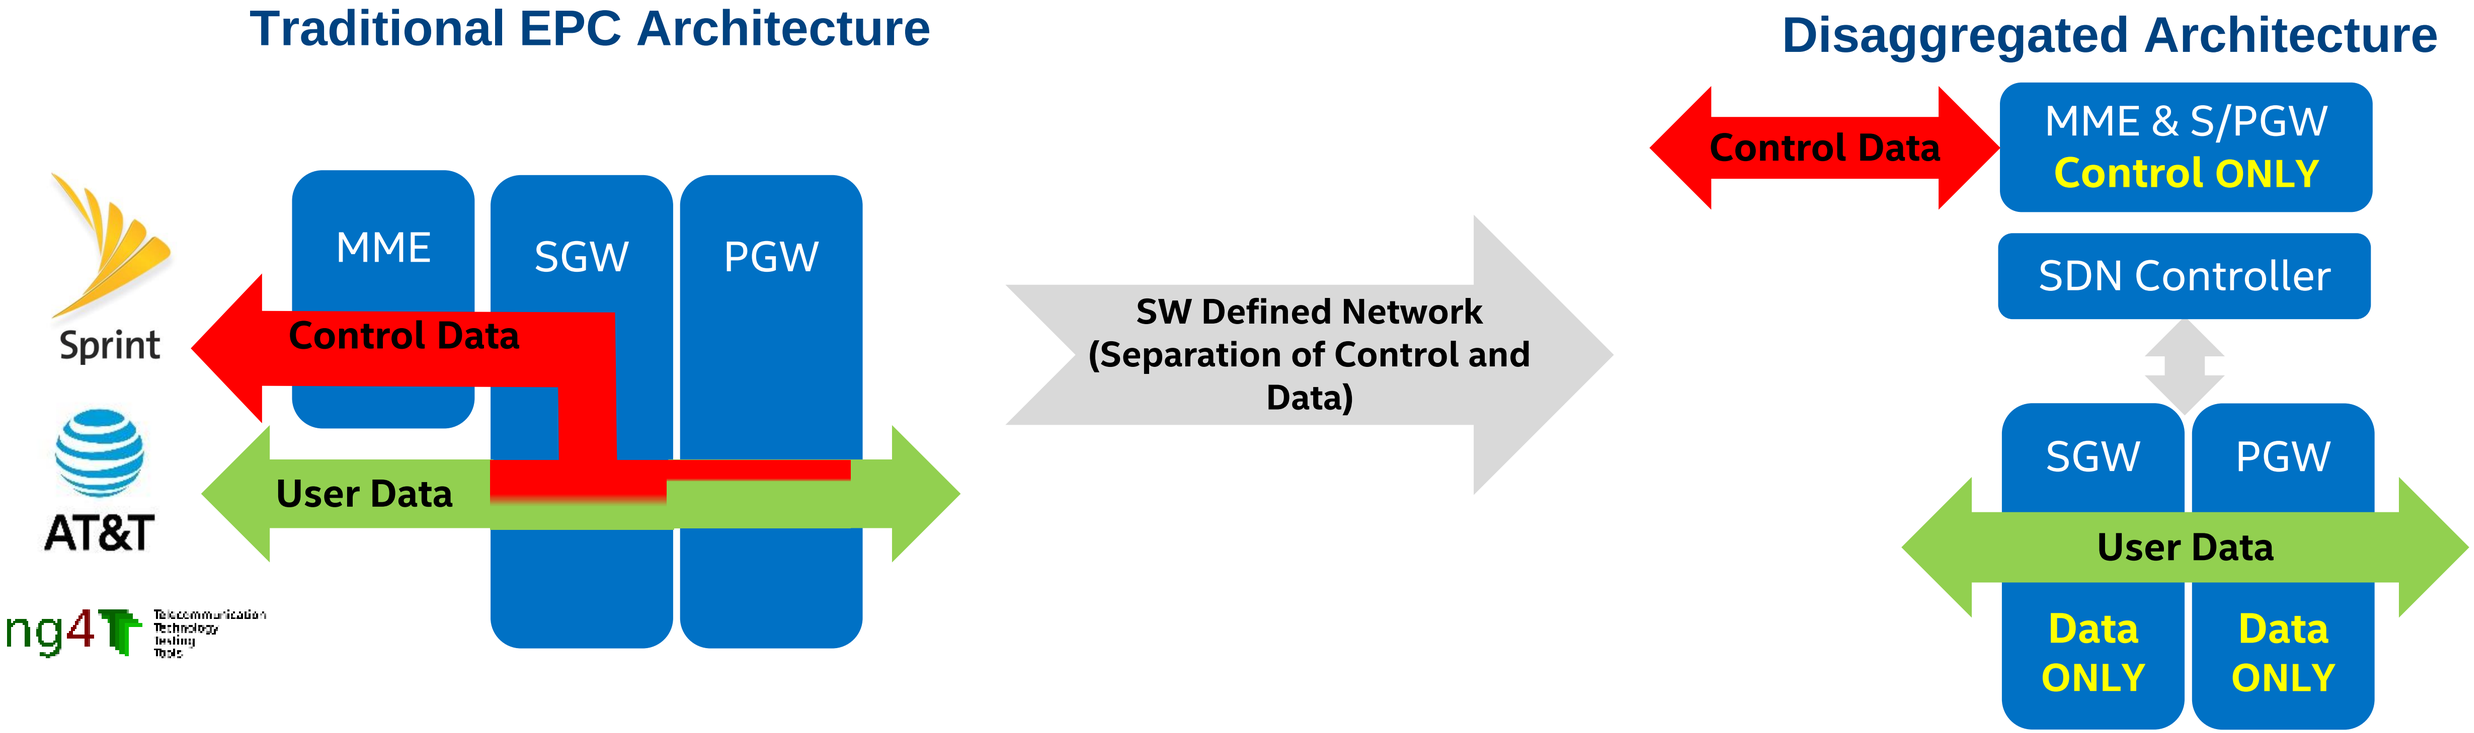

- Operators' real traffic (San Jose, Houston, Chicago, …)
- Identified system's bottleneck
  - "Understanding Bottlenecks in Virtualizing Cellular Core Network functions", IEEE LANMAN '15
- No independent control or data scaling

- SDN based architecture
- Match/Action semantic data plane
- Independent & scalable control & data
- Functional EPC per operator's requirements

Functional EPC per Operator's Requirements - NGIC

https://gerrit.opencord.org/#/admin/projects/ngic

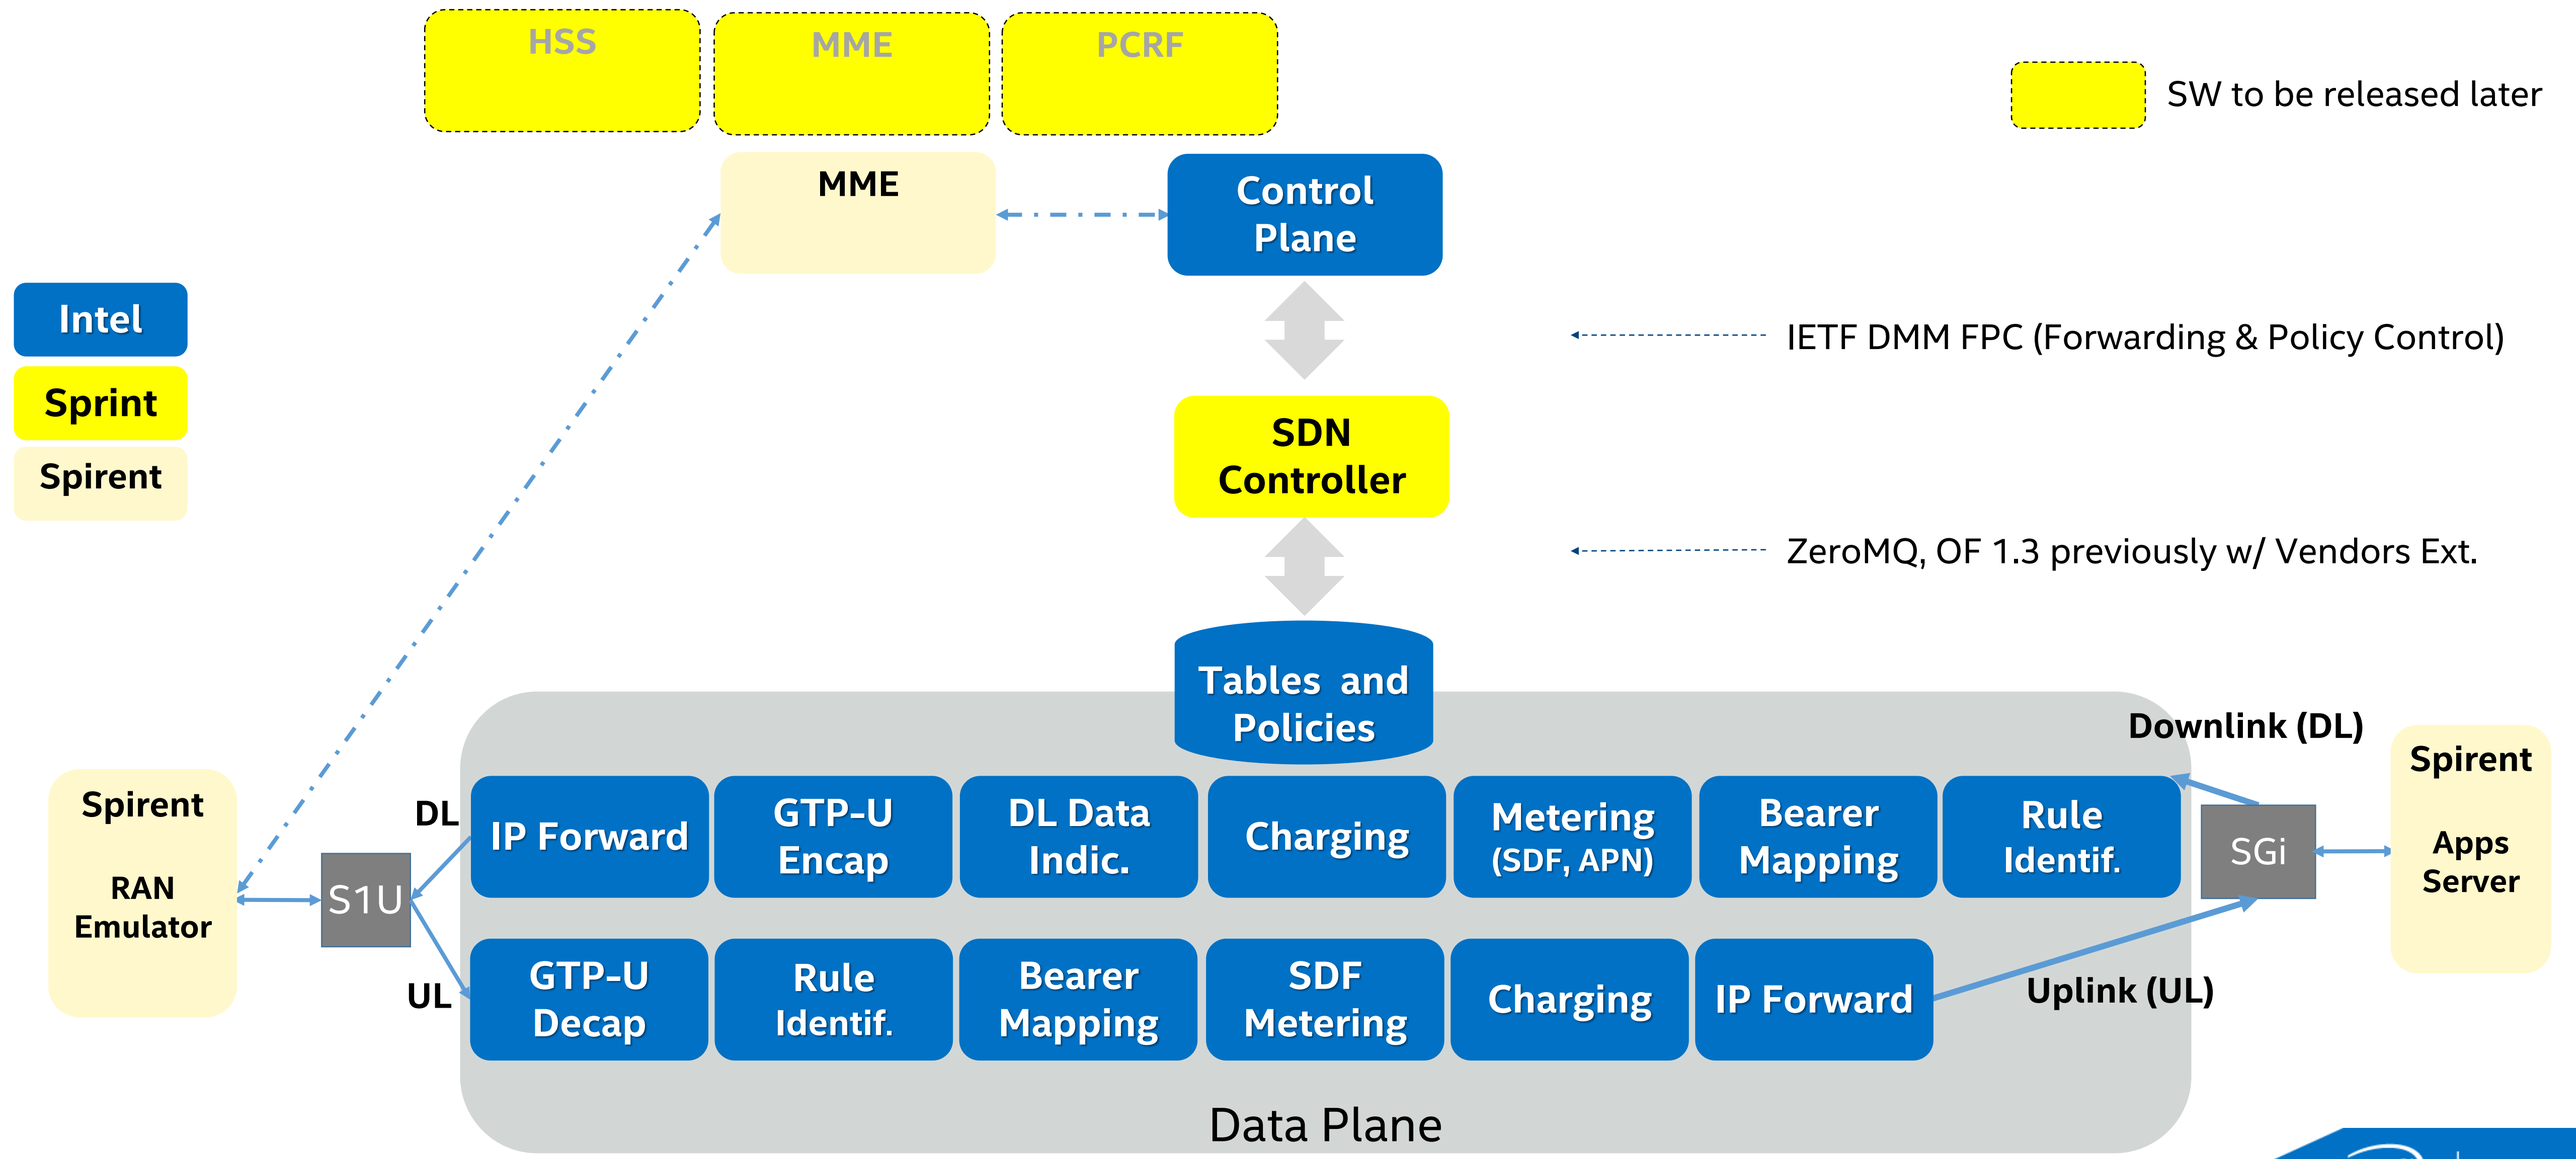

Linear Scaling Characteristics

- 1 packet processing core
  - 2.125 Million pps (50k subs)
  - 1.625 Million pps (500k subs)
- 4 packet processing cores
  - 8.2 Million pps (50k subs)
  - 6.5 Million pps (500k subs)

All packet processing cores implement all network functions (PGW, SGW, DPI, Child Protection, CG-NAT, static FW & SFC)

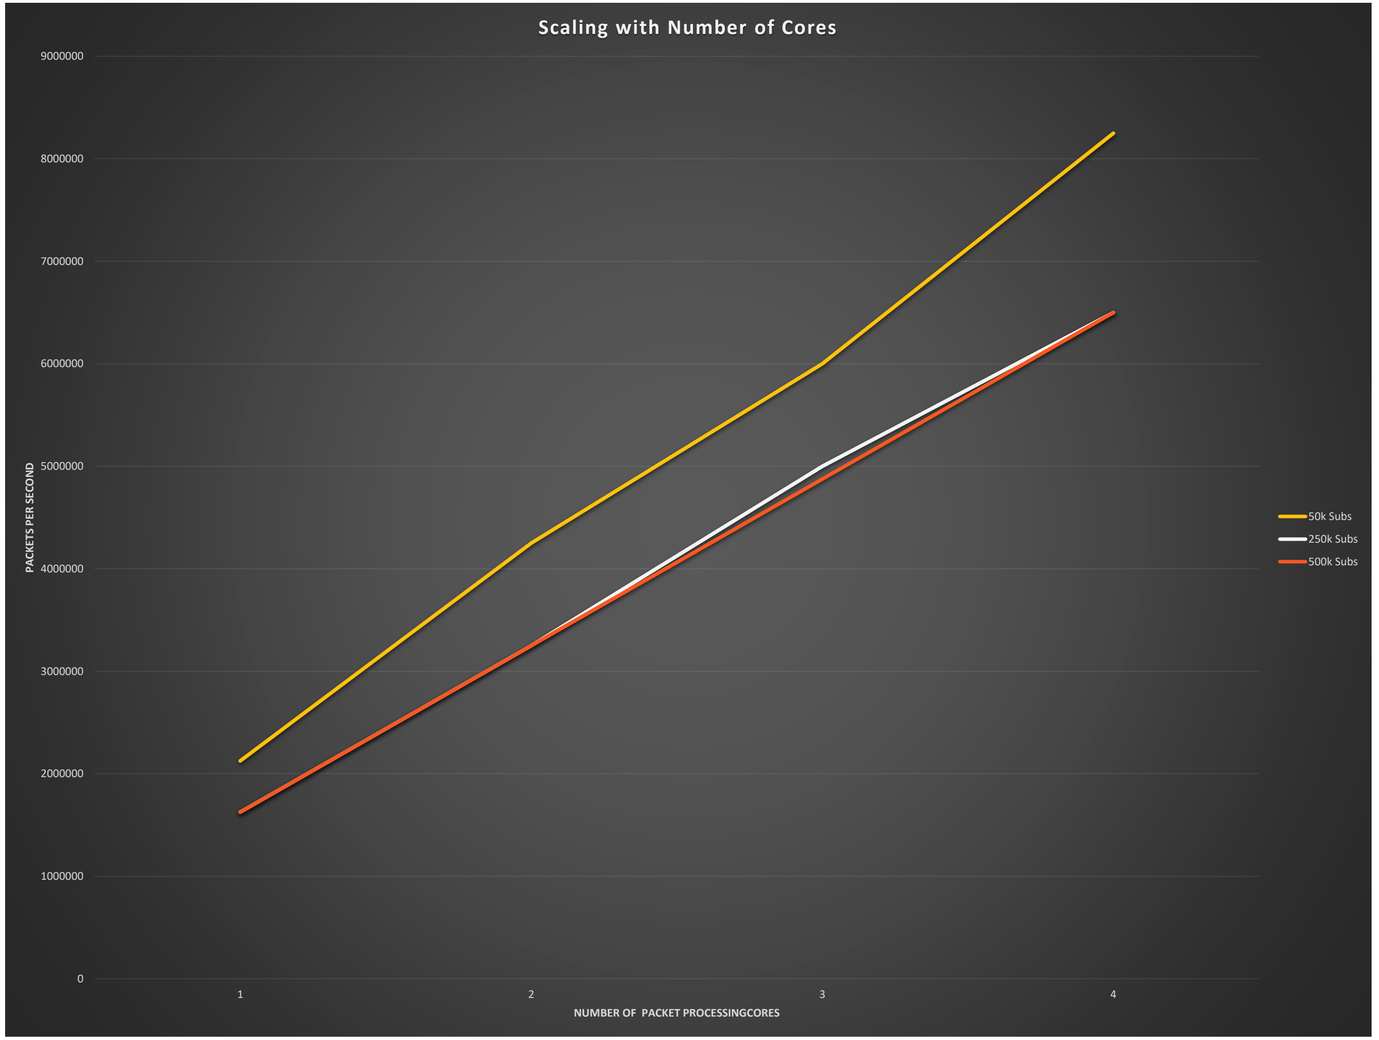

- Additional (overhead) cores for C3PO EPC Data Plane Node
  - Rx Core
  - Tx Core
  - Load-balancing Core
  - Master Core
  - Interface Core
- Packet processing cores are in addition to the overhead cores

Agenda

- vEPC Background
- Next Generation Infrastructure Core (NGIC)
- CORD High Level View
- NGIC and M-CORD
- Summary / Next Steps

M-CORD

Enabling the Programmable 5G Edge Cloud

Source: Open Networking Foundation

M-CORD: Enabling the Programmable 5G Edge Cloud

Source: Open Networking Foundation

Edge Cloud

Wireless Access **Edge Cloud** **Core**

PROGRAMMATICALLY instantiate slices for different use cases

Virtualized RAN Control

Virtualized EPC

DISAGGREGATE RAN

VIRTUALIZE RAN components in the Cloud

SDN-ize it for **PROGRAMMABILITY**

DISAGGREGATE Core

VIRTUALIZE all components in the Cloud

SDN-ize it for **PROGRAMMABILITY**

After with SDN & NFV

Before with dedicated hardware

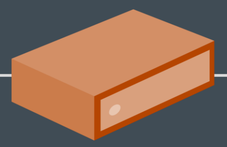

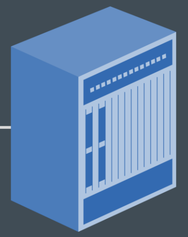

Radio

eNodeB

EPC

Source: Open Networking Foundation

M-CORD as the Edge Cloud

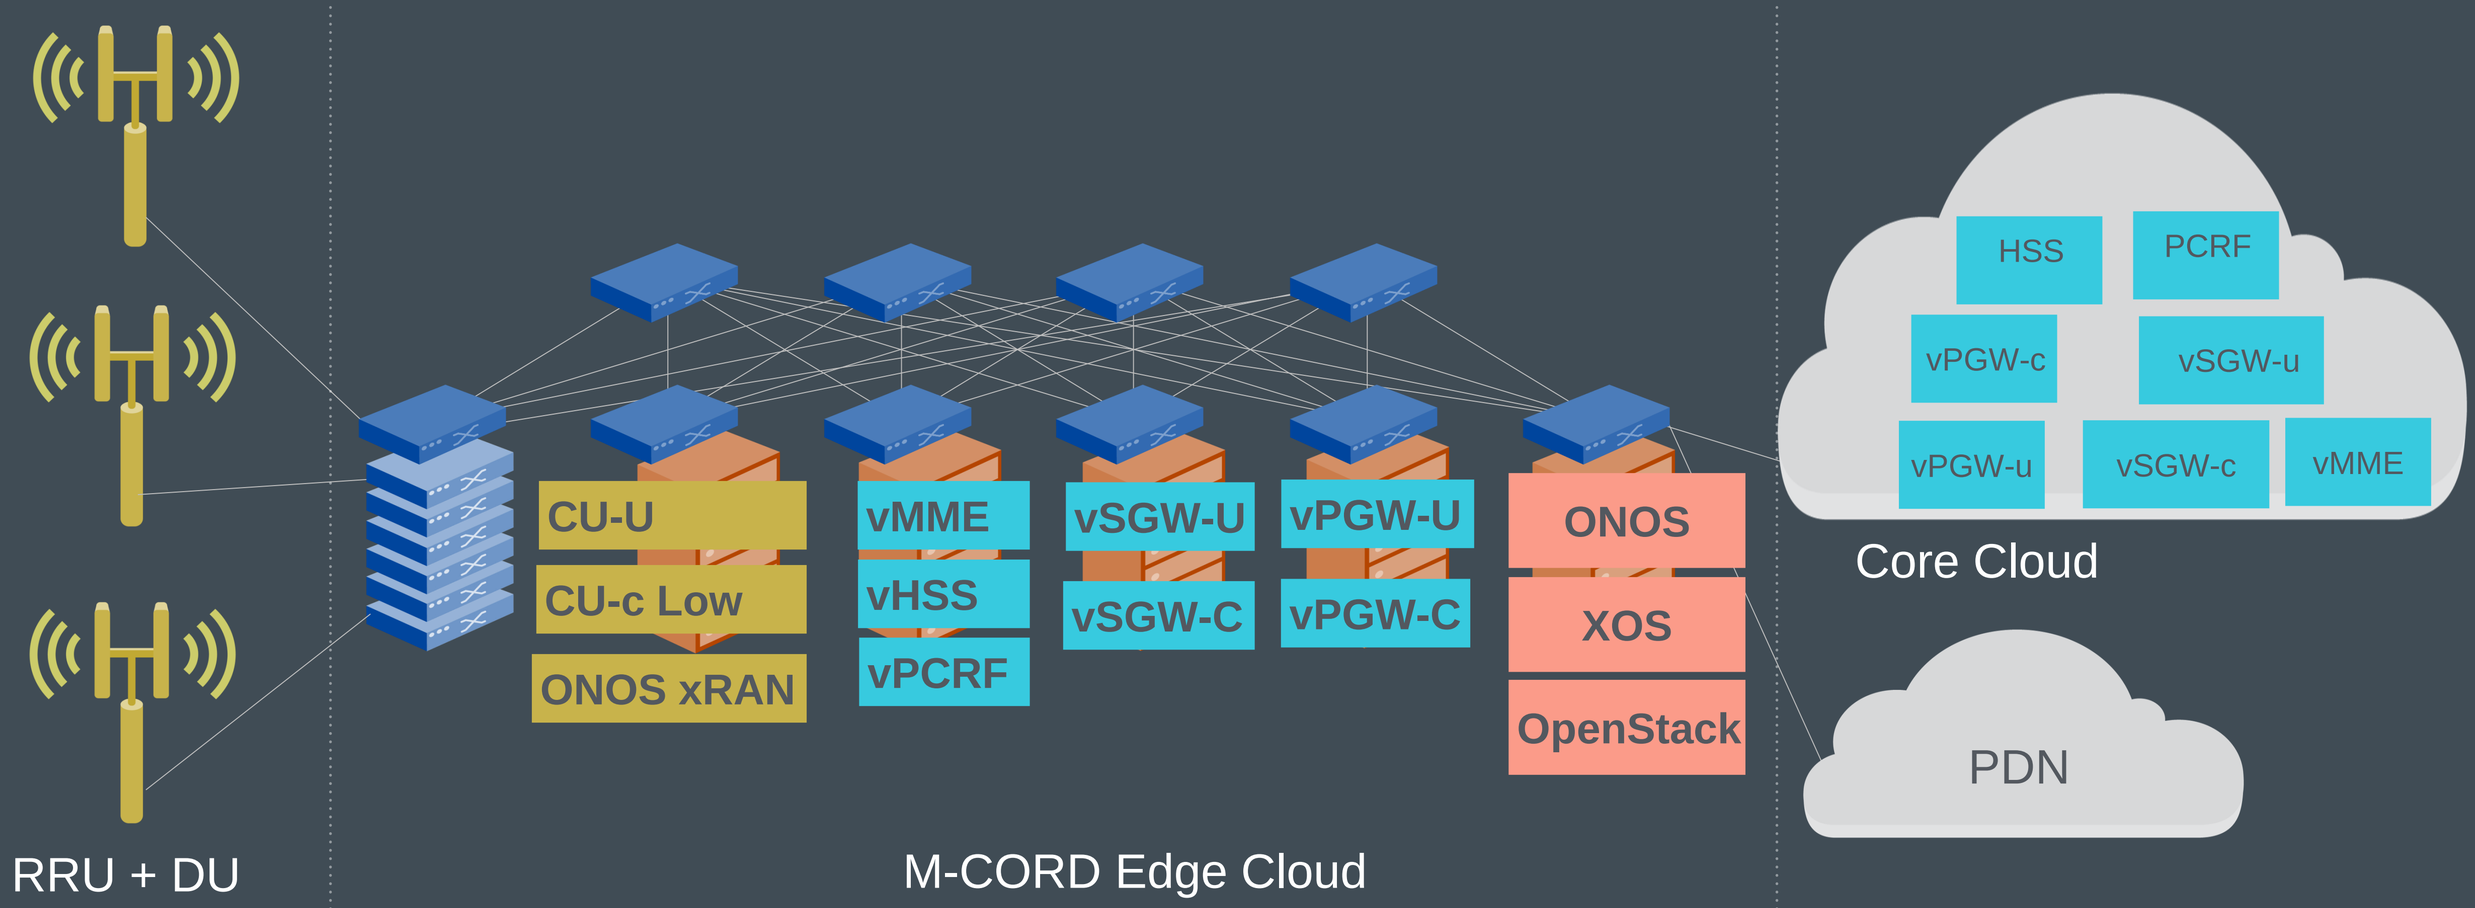

Source: Open Networking Foundation

M-CORD Usage Models

RAN and Access

SD-RAN:
xRAN + ONOS

Multi-Radio (MRAT)
WiFi / LTE Steering

Commercial Small Cells

CBRS Spectrum

WiFi RAN Slicing

End-to-End Solution PoCs & Capabilities

Multi-Access CORD

End-to-End Network Slice Stitching

Diversification of Hardware Choices

Virtualized RAN Control

Virtualized EPC

Mobile Core

Open Source EPC

Infrastructure Acceleration

Source: Open Networking Foundation

Agenda

- vEPC Background
- Next Generation Infrastructure Core (NGIC)
- CORD High Level View
- NGIC and M-CORD
- Summary / Next Steps

NGIC (vEPC) & M-CORD

https://gerrit.opencord.org/#/admin/projects/ngic

NGIC & M-CORD

- Millions of Any-G Connections
- Full Mobility with Traffic
- Call Modeling
- Topology Emulation
- Full virtual
- ETSI Compliant

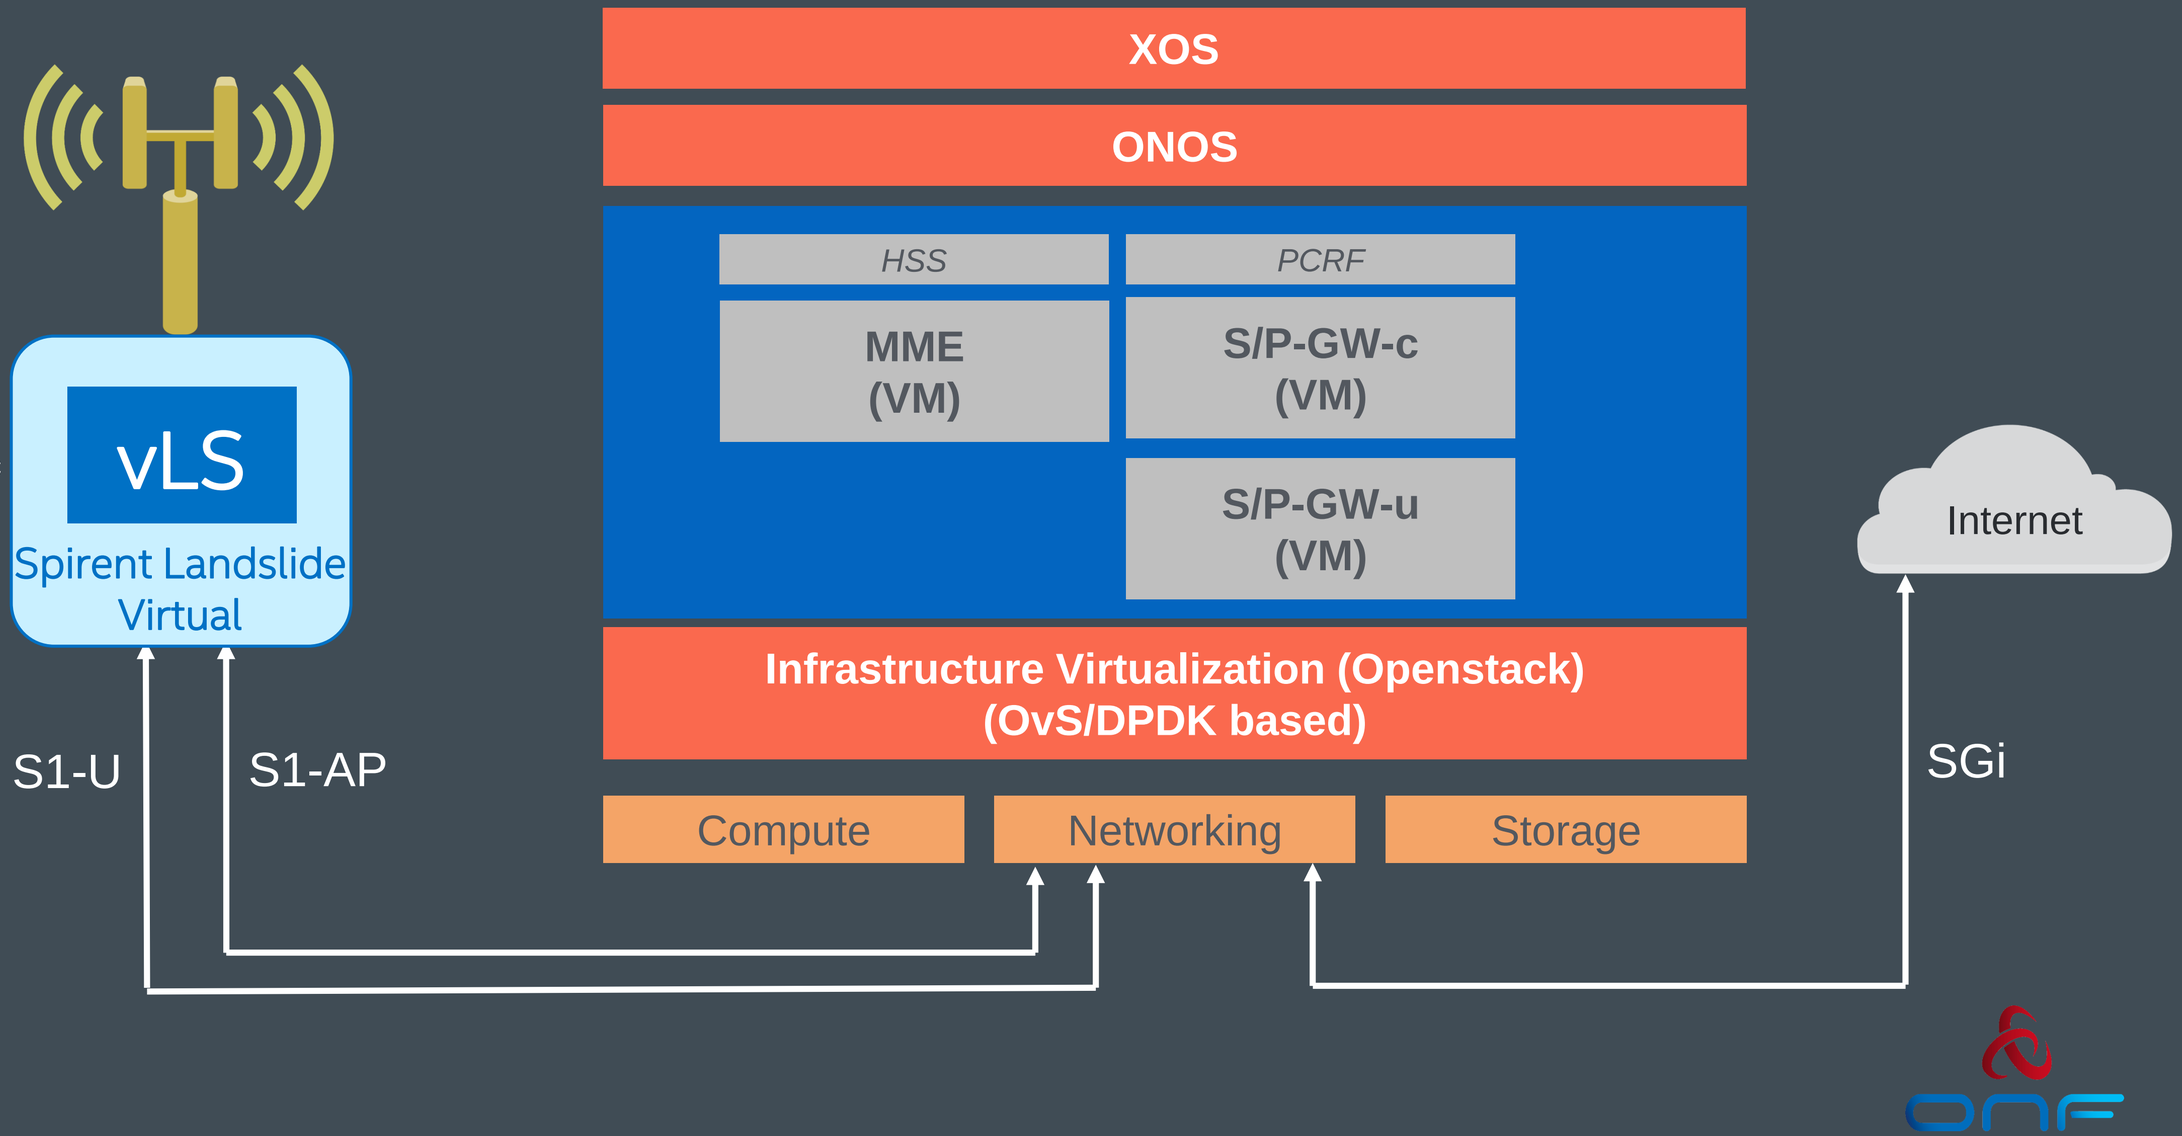

Source: Open Networking Foundation

LOGICAL NETWORK CONNECTIVITY IN CORD

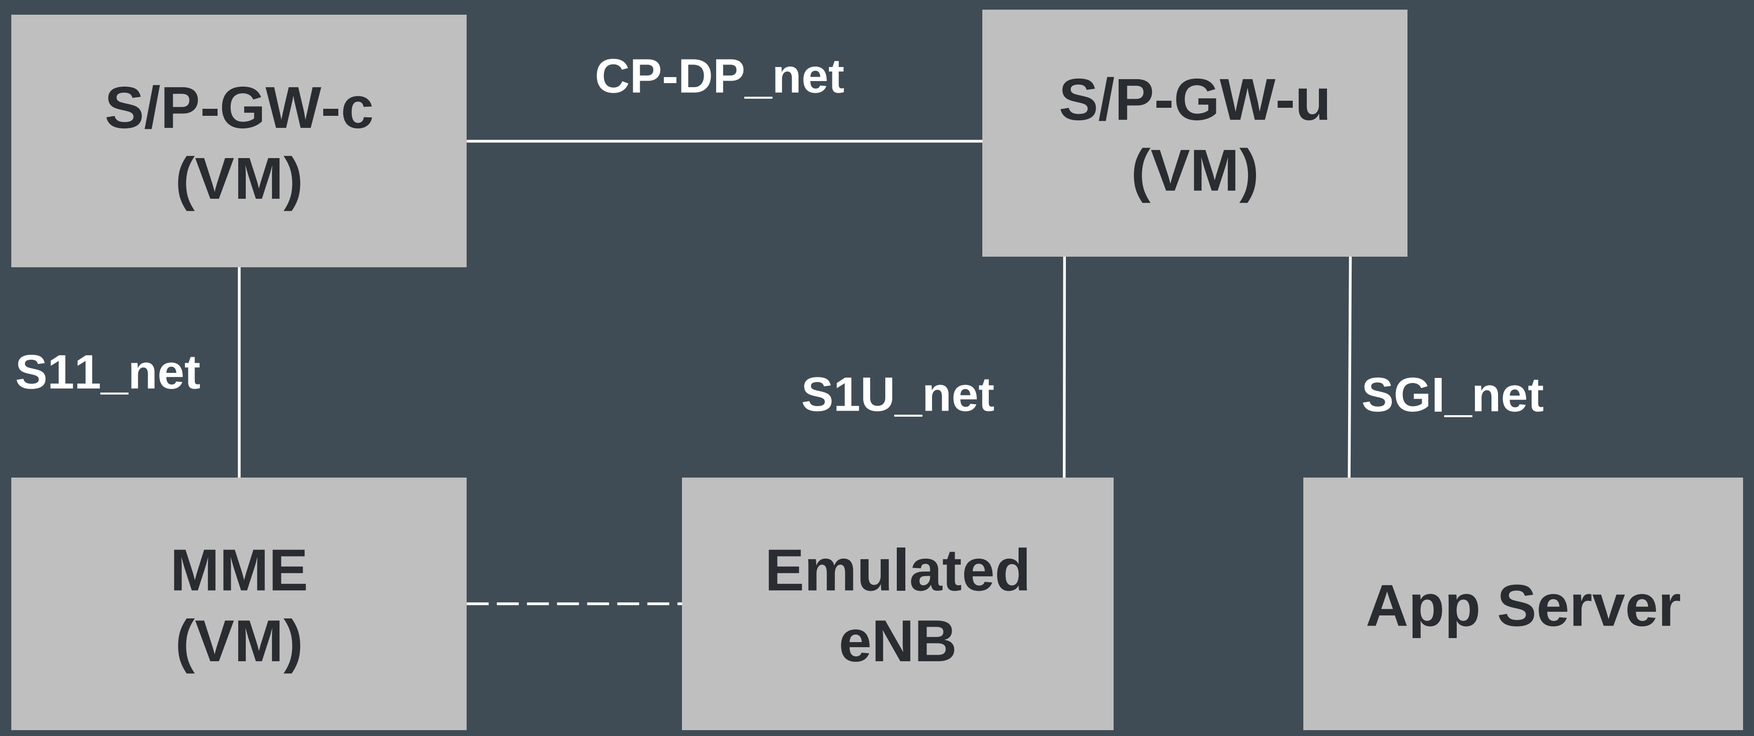

PHYSICAL NETWORK CONNECTIVITY

Legend:

A,D	management	MANAGEMENT
B,E	CP-DP_net	PRIVATE
C	S11_net	ACCESS_AGENT
F	S1U_net	ACCESS_AGENT
G	SGI_net	ACCESS_AGENT

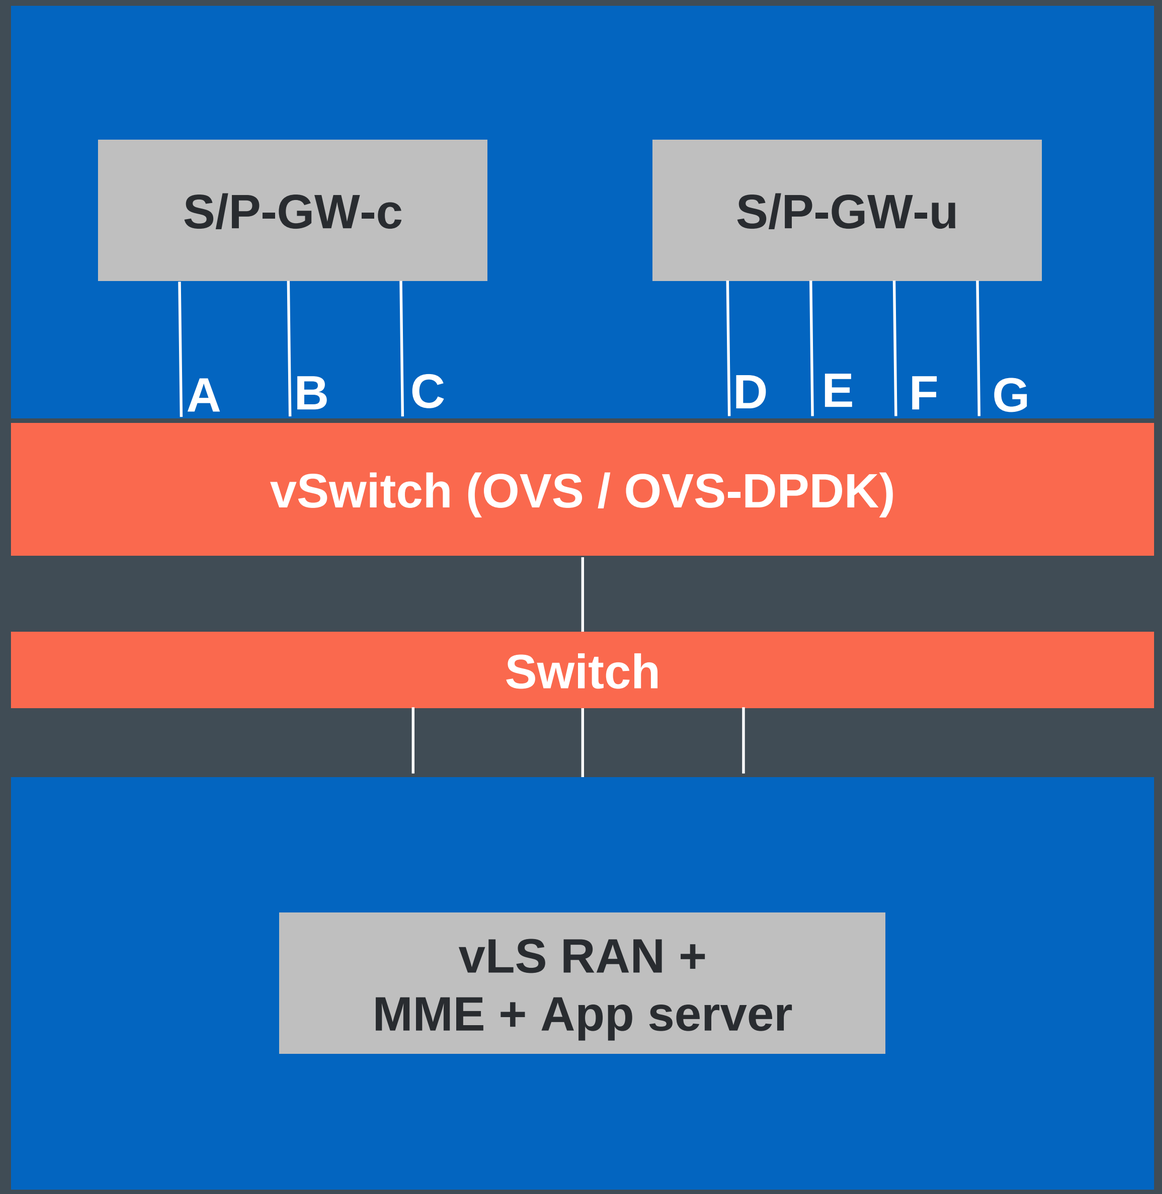

M-CORD using OvS/DPDK and Enhanced Platform Awareness

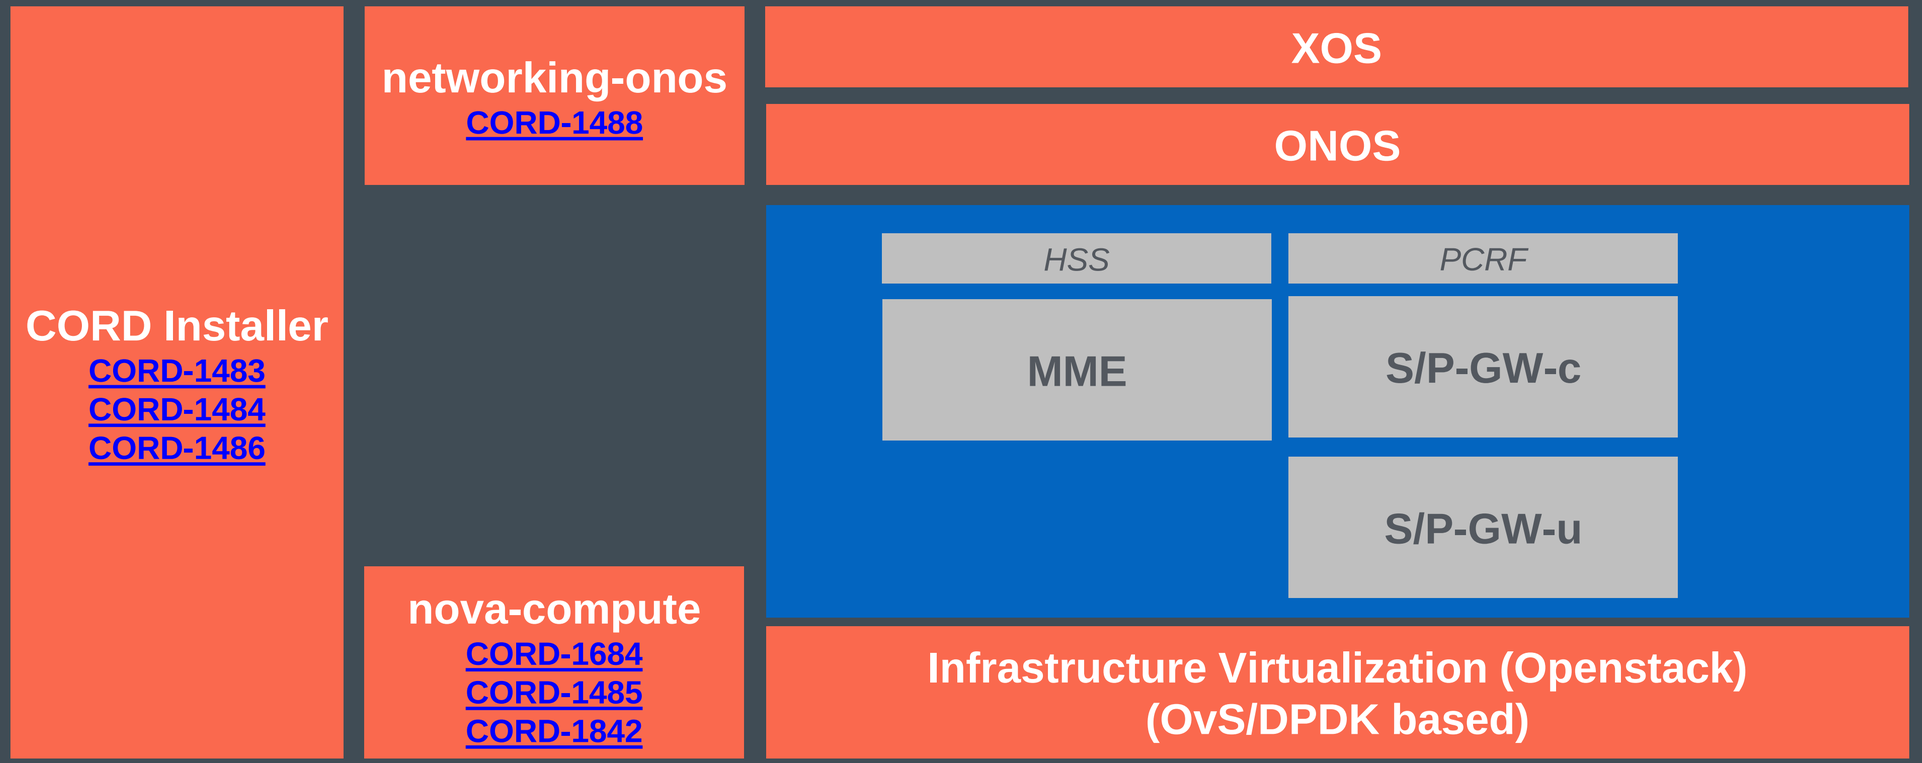

CORD-1483: Configure Nova with DPDK-enabled flavors
CORD-1484: Add nova config options to juju_config.yml
CORD-1485: Install OVS-DPDK through nova-compute charm
CORD-1486: Set kernel commandline parameters for compute nodes
CORD-1488: Add VIF details to handle VHOST_USER port type.
CORD-1684: Backport libvirt driver changes, to enable memAccess=shared tag required for correct functioning of VHOST_USER
CORD-1842: Replace default QEMU with patched QEMU 2.7.0 for correct functioning of VHOST_USER

M-CORD using OvS/DPDK and Enhanced Platform Awareness

L2FWD Throughput in CORD w/ OvS/DPDK vs. CORD w/ OvS
(Higher is better)

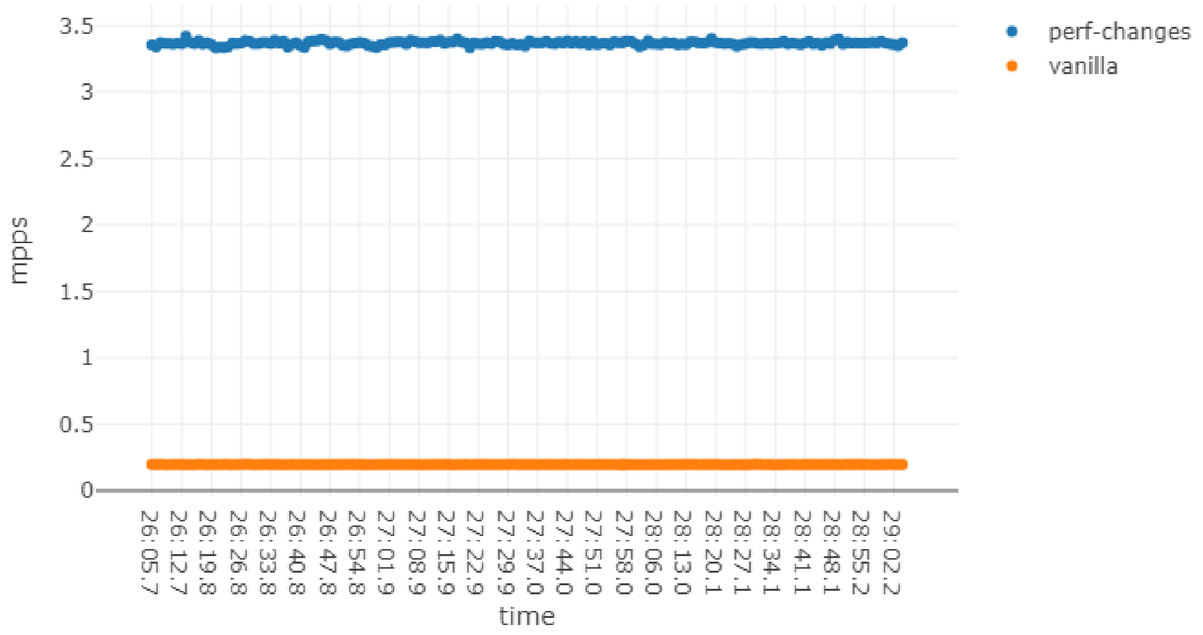

L2 FWD Latency in CORD w/ OvS/DPDK vs. CORD w/ OvS
(Lower is better)

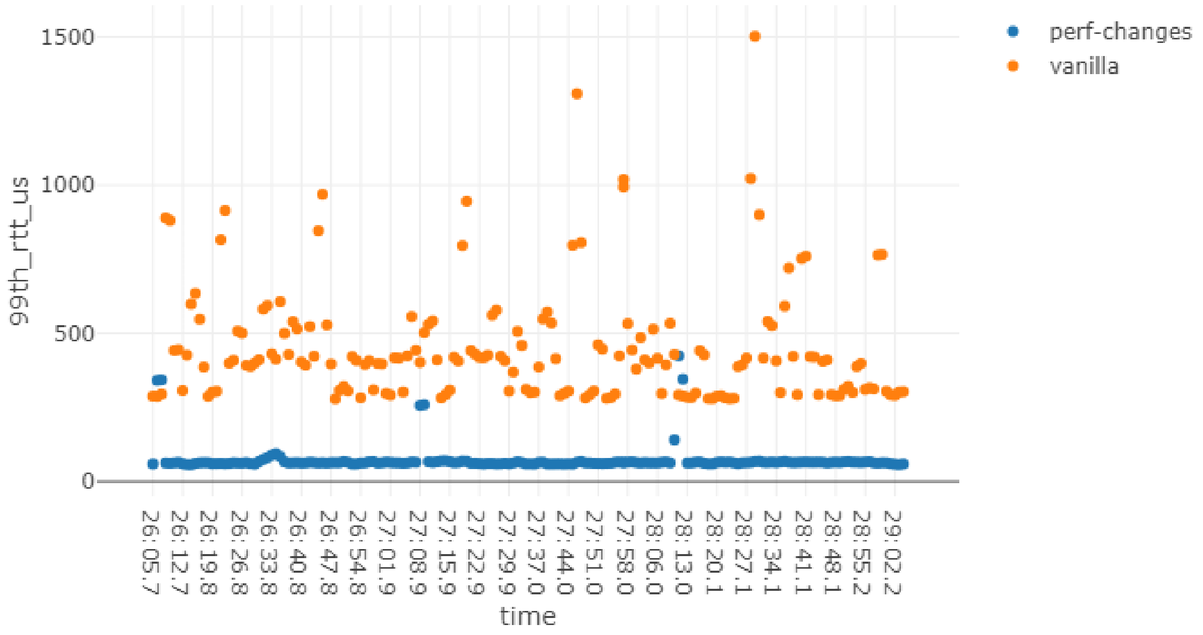

- **Features used to achieved high performance on CORD**
- OvS/DPDK
- Enhanced Platform Awareness:
- VM memory backed by Huge Pages
- CPU isolation
- Core pinning
- vHost user data path

Modules: vEPC & M-RAT (1/2)

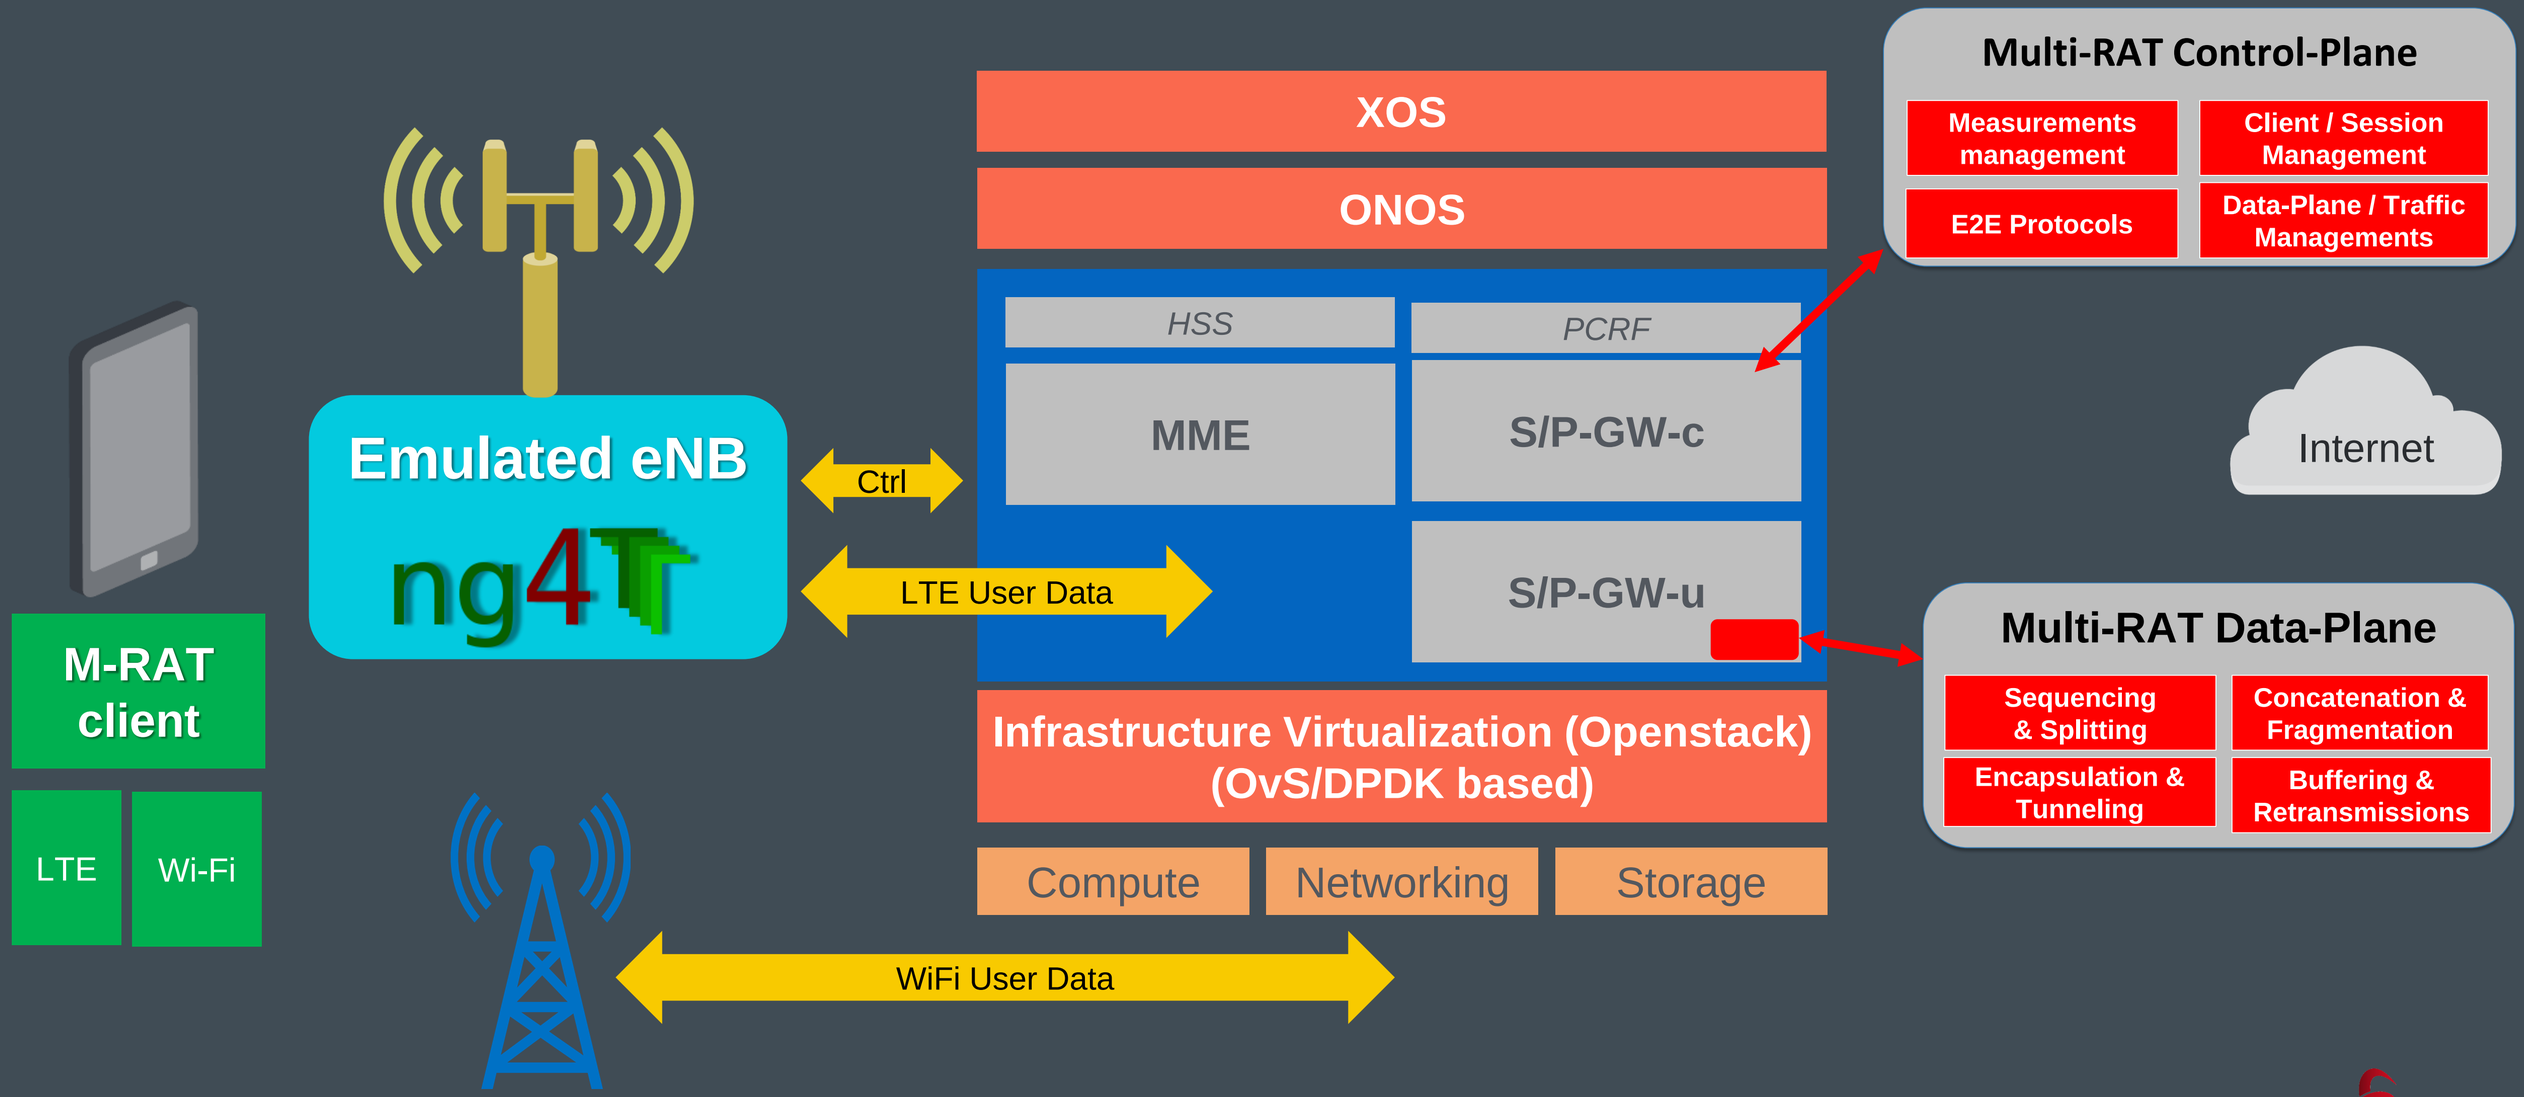

Modules: vEPC & M-RAT (2/2)

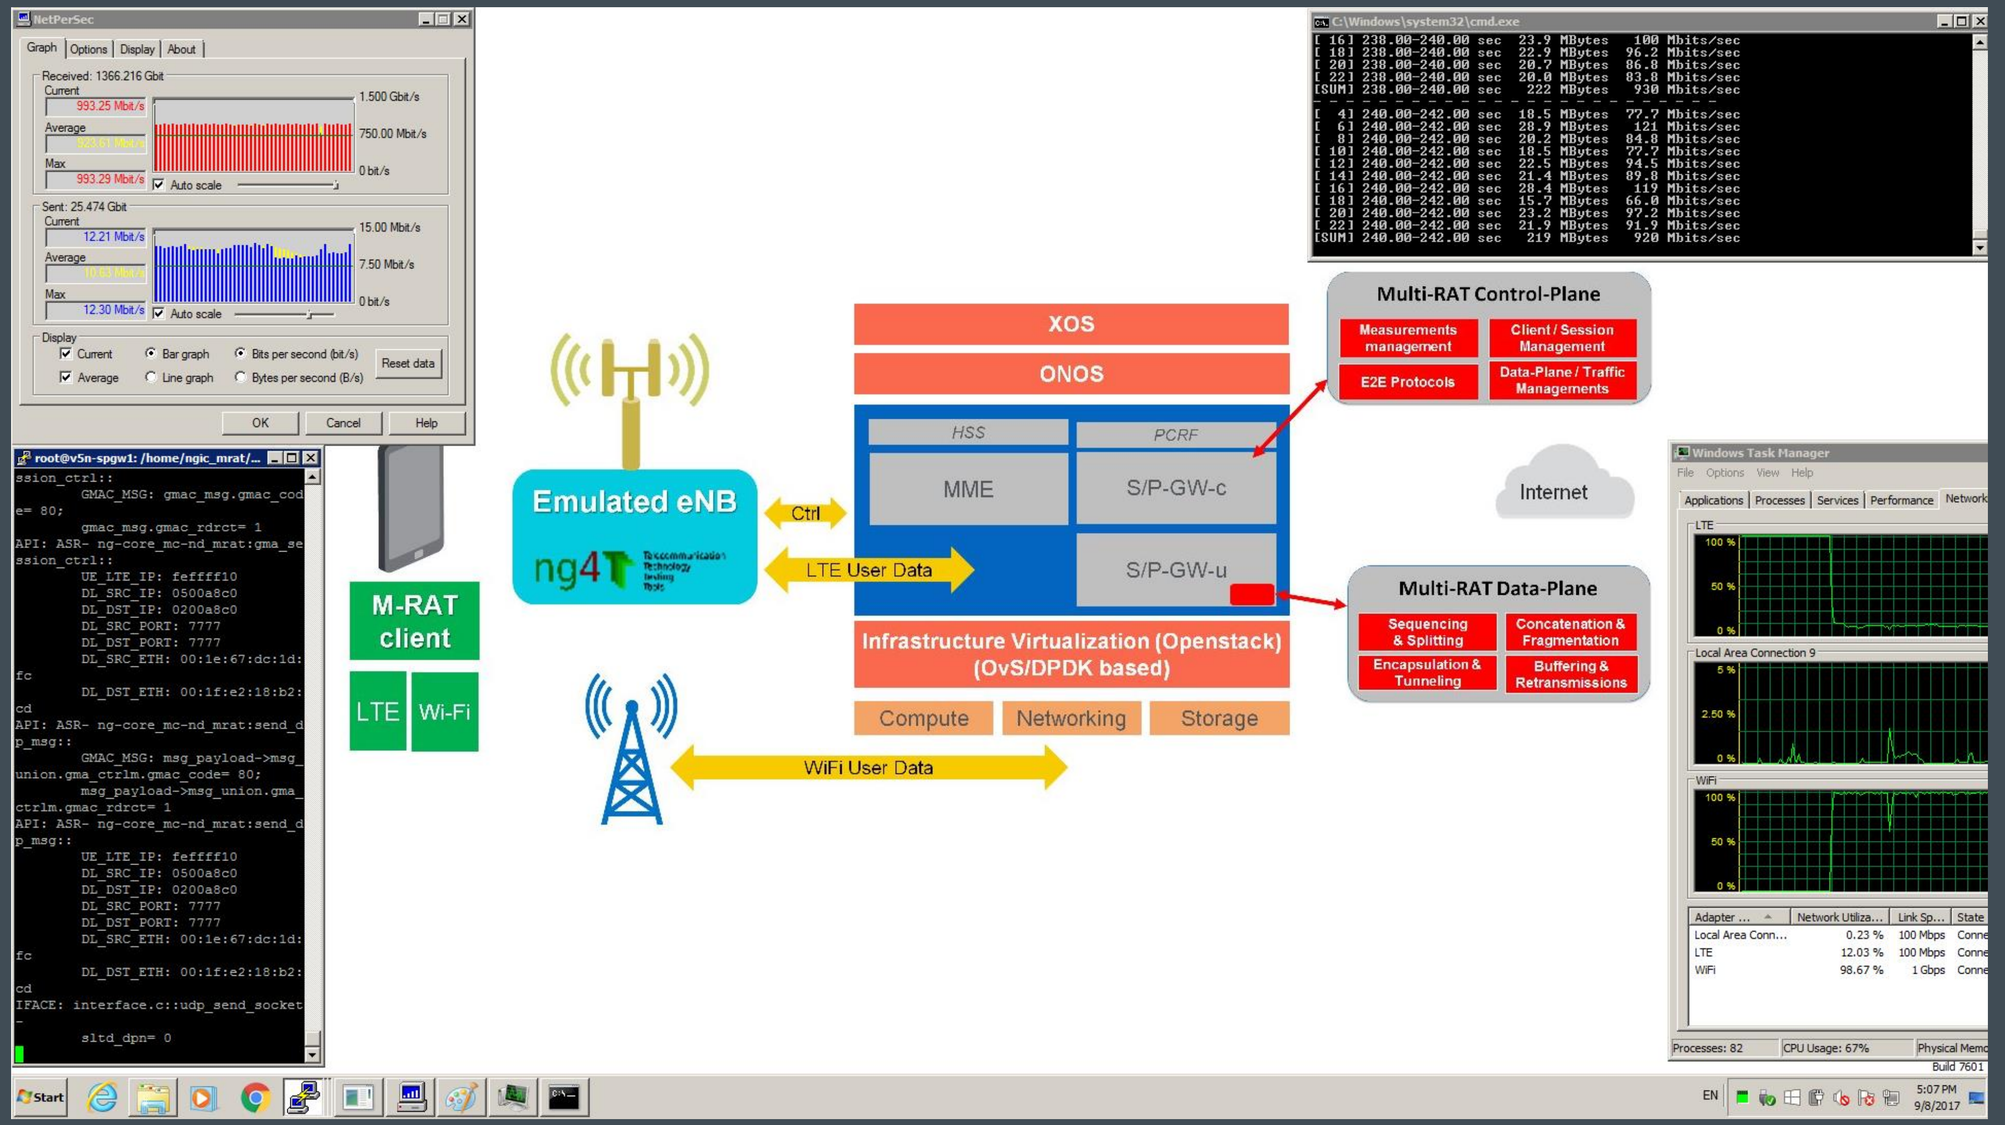

Agenda

- vEPC Background
- Next Generation Infrastructure Core (NGIC)
- CORD High Level View
- NGIC and M-CORD
- Summary / Next Steps

Summary / Next Steps

- NGIC is a functional, DPDK based, EPC operating in bare-metal, VMs, or containers and orchestrated in M-CORD

- Additional EPC components, MME, HSS, PCRF, will be made available to CORD over time

NGIC Hands-on Tutorial Next

Prerequisites for Hands-on Tutorial

- Install Docker (1.13 or higher) and Docker Compose

- Docker images for NGIC control, data plane and traffic generation

```
docker pull ngiccorddemo/ngic-cp
docker pull ngiccorddemo/ngic-dp
docker pull ngiccorddemo/ngic-traffic
```

- Demo folder

```
git clone https://github.com/ngiccorddemo/cordbuild2017.git
```

Thank You

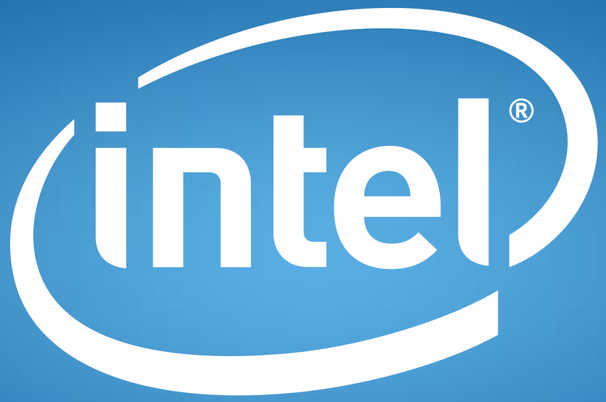

experience
what's inside™

Backup

Efficient MATCH/Action Semantic Data Plane (1/2)

Match/Action : Optimized Table Lookup with Cuckoo Hashing[Pagh 01]

- "Scalable, High Performance Ethernet Forwarding with CuckooSwitch", Dong Zhu, Bin Fan, Dave Anderson (CMU), Michael Kaminsky (Intel)

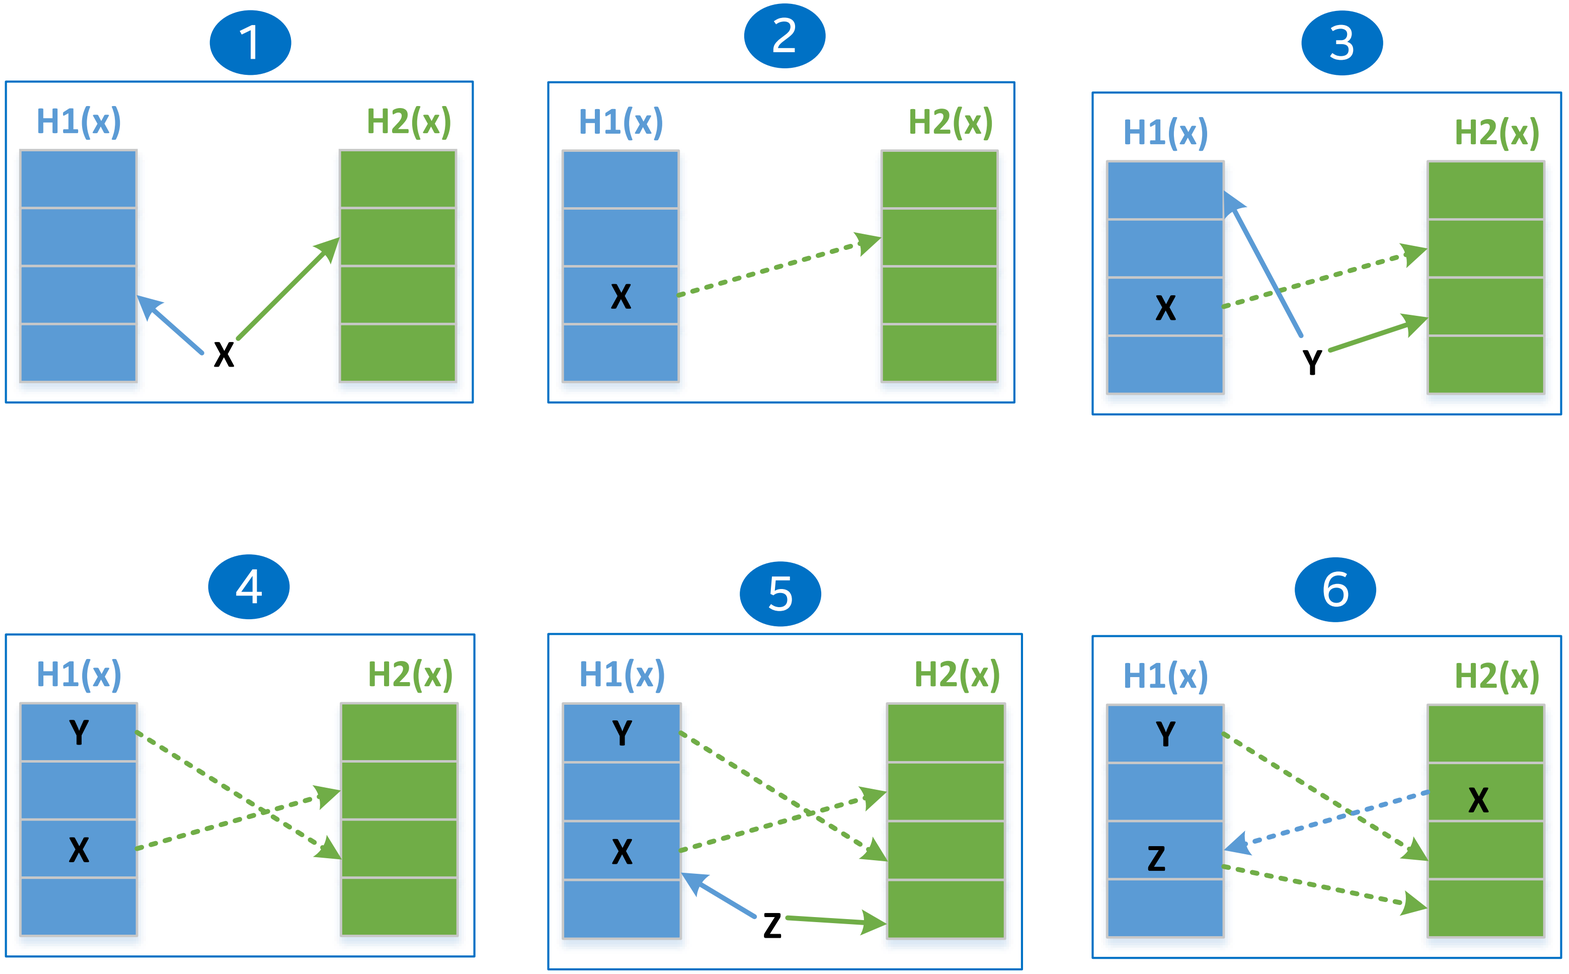

One `Insert` may move a lot of items especially at high table occupancy. Optimal multi-writer insertion using Intel® TSX

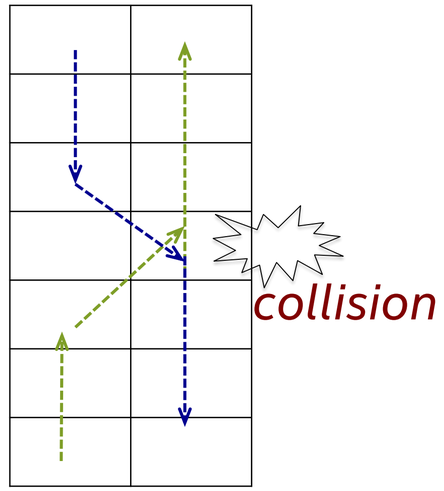

Efficient MATCH/Action Semantic Data Plane (2/2)

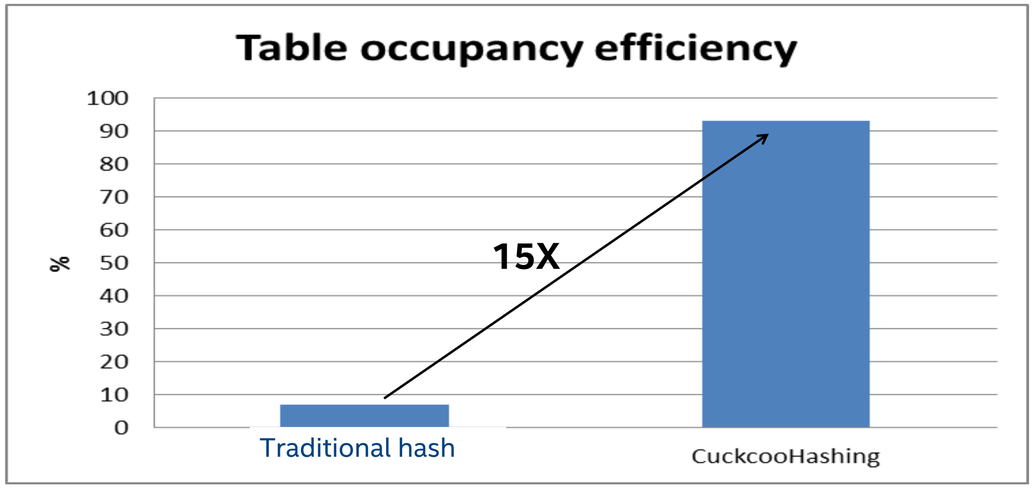

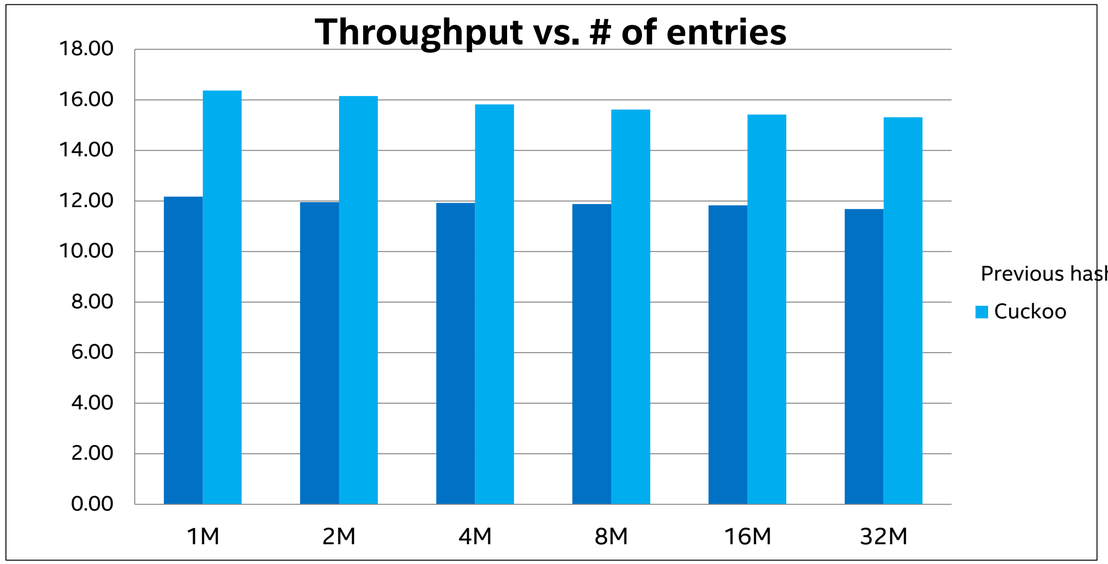

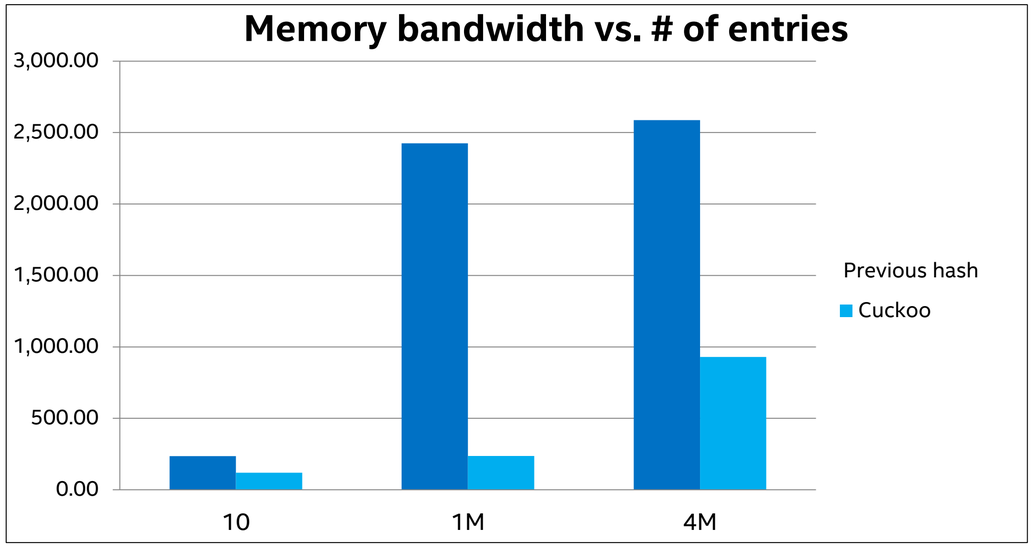

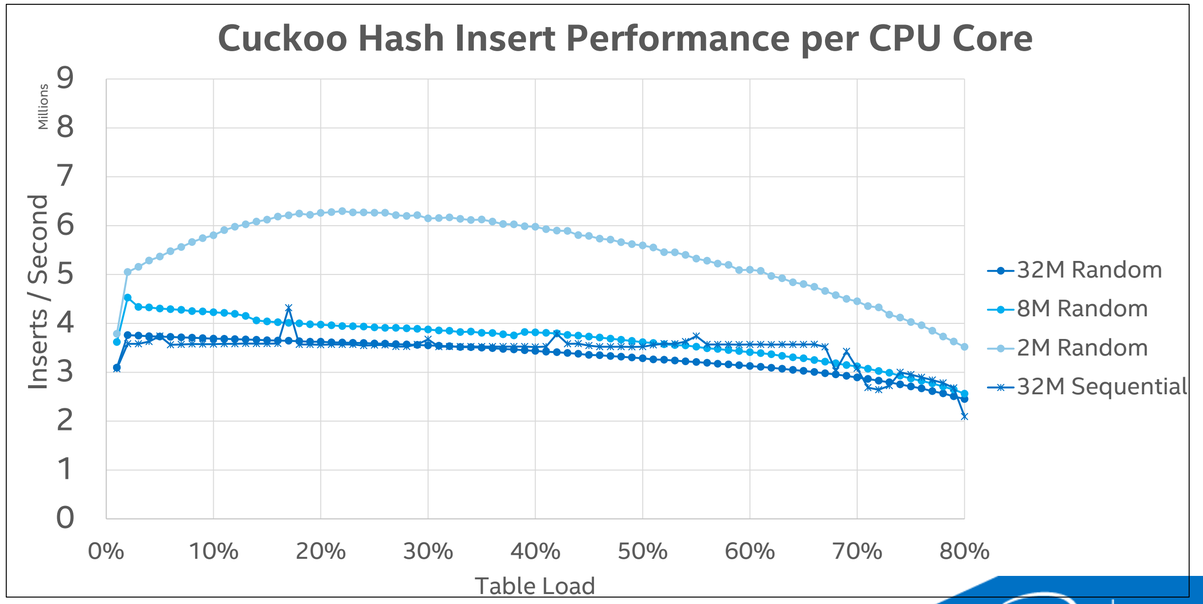

Modules: vEPC & M-RAT (2/2)