5G Transformation with Open Source

5G Cloud Native from RAN to Core

Christian Maciocco, Intel
Shilpa Talwar, Intel
Saikrishna Edupuganti, Intel
Muhammad (Asim) Jamshed, Intel

2020

Agenda

- Cloud Native Disaggregated Network Infrastructure
- Transition to 5G
- Near Real-Time RAN Information Controller & Services
- Demo of Dual Mode 5G UPF

Open 5G Network Infrastructure to Accelerate Edge Deployment

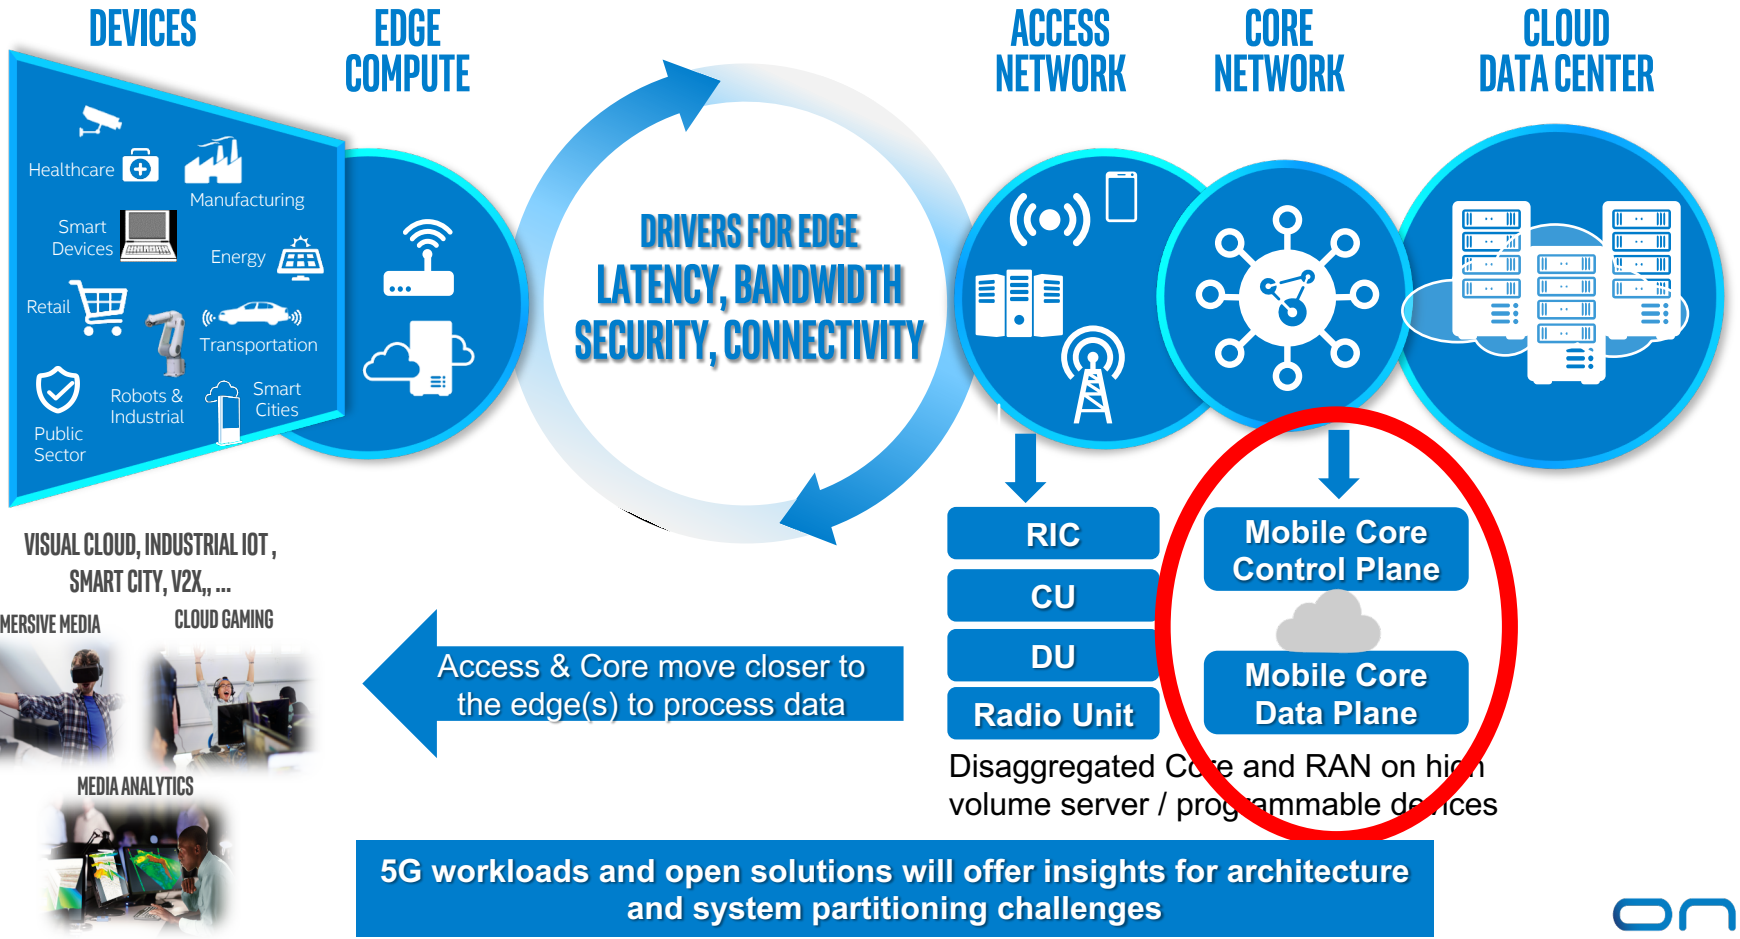

Building on ONF Success Disaggregating a 4G/LTE Core

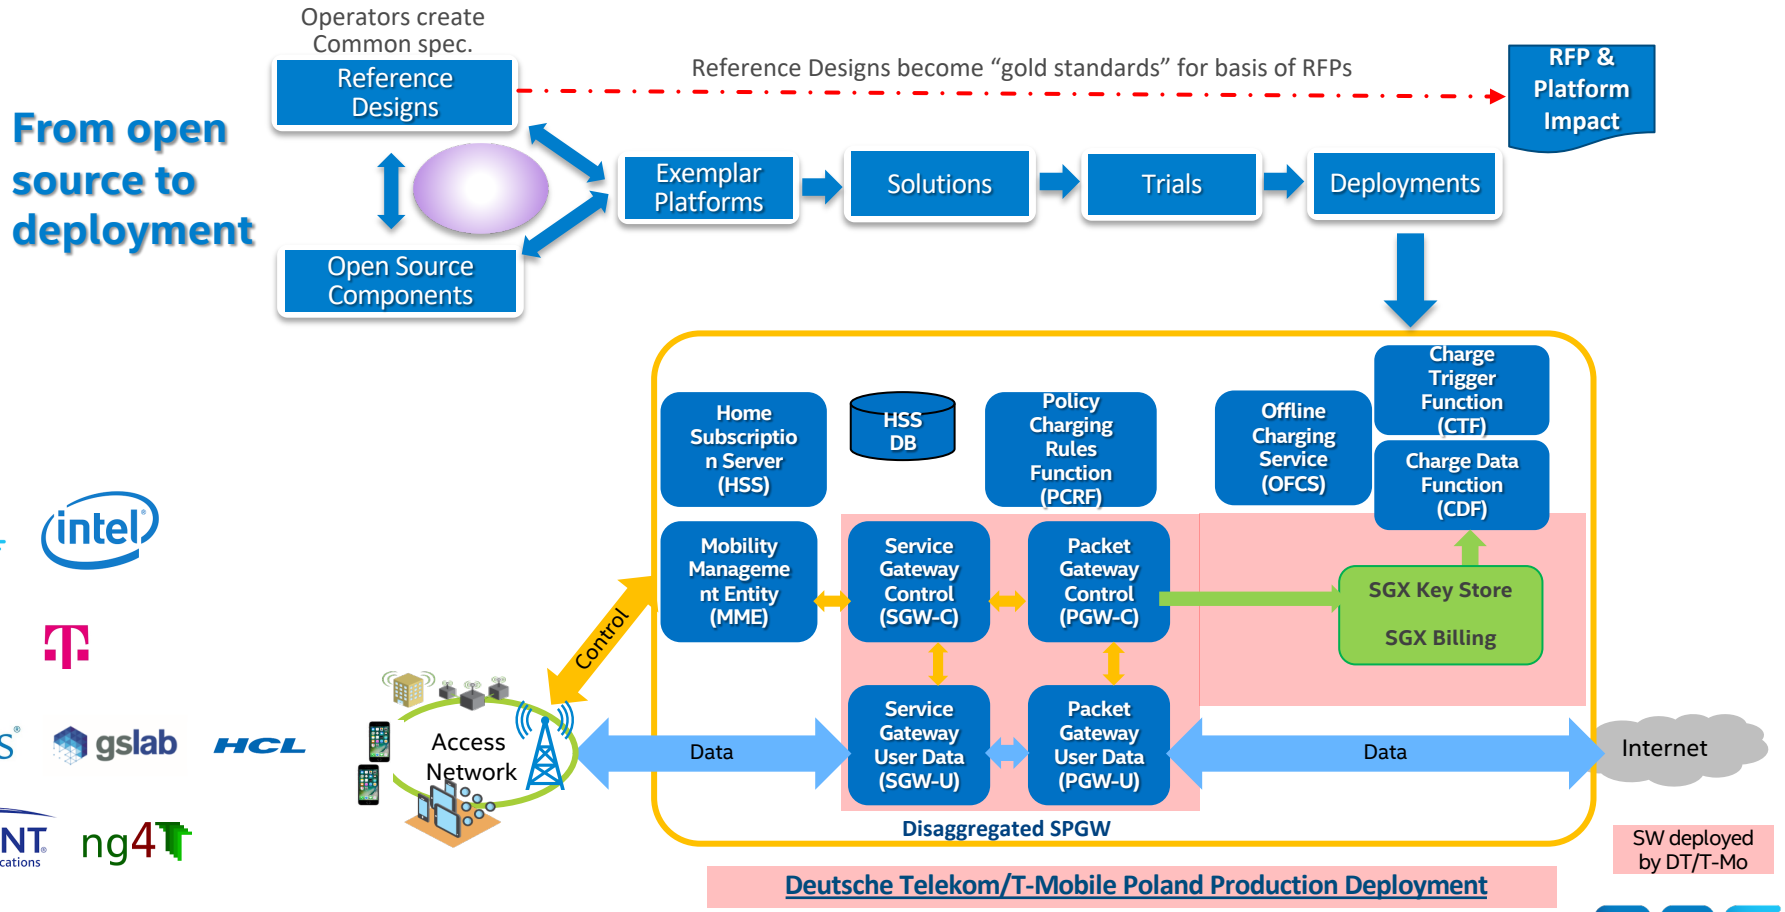

Towards 5G SA - A Dual Mode 5G/LTE UPF

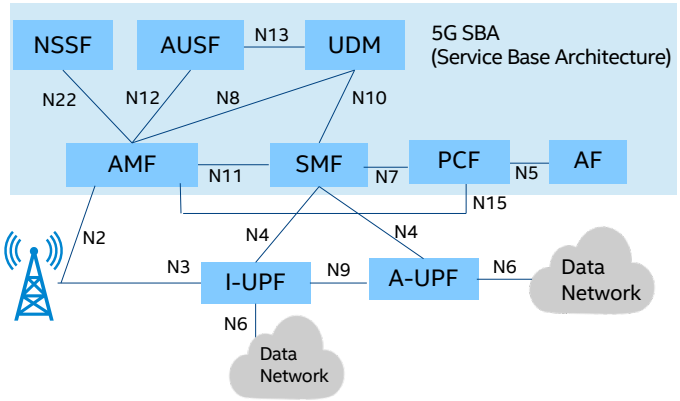

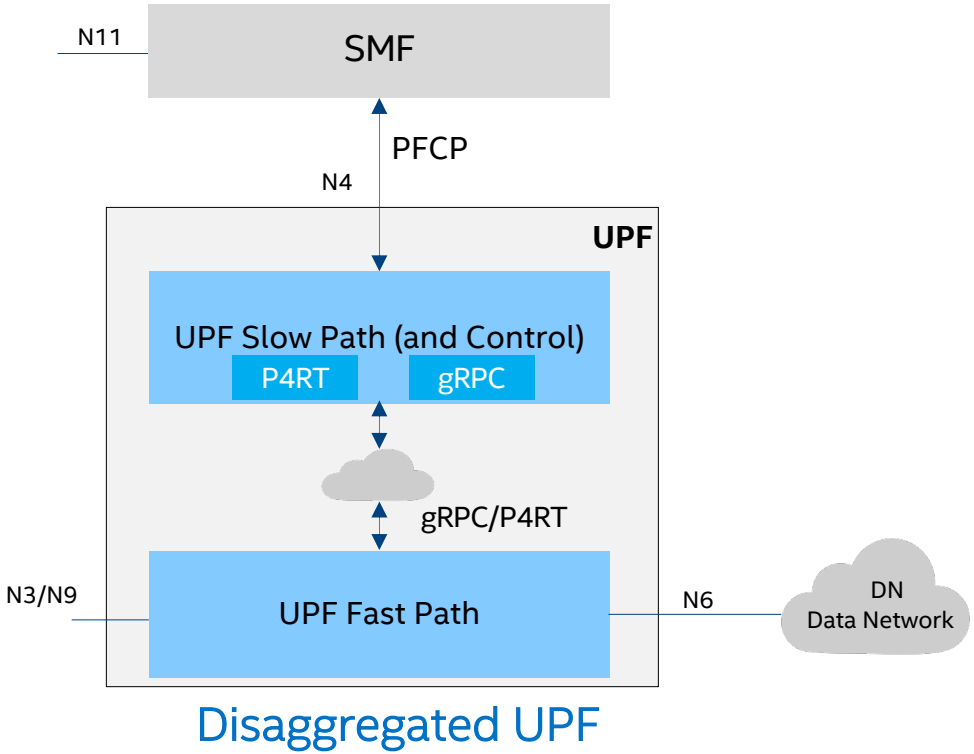

Disaggregated UPF

UPF with One Slow Path, Fast Path Options

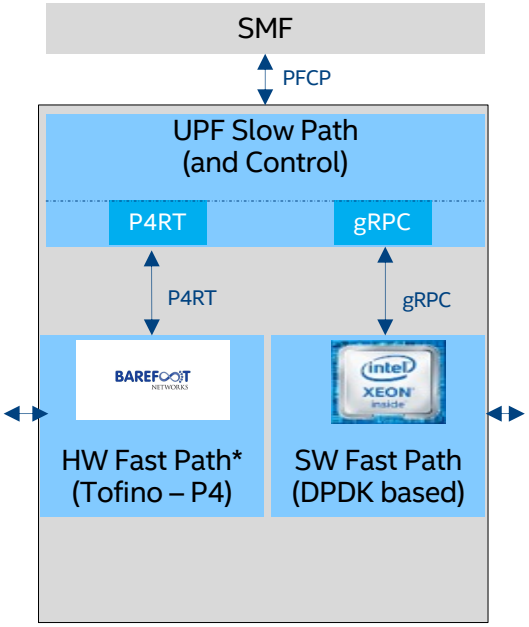

* P4 Pipeline developed at ONF

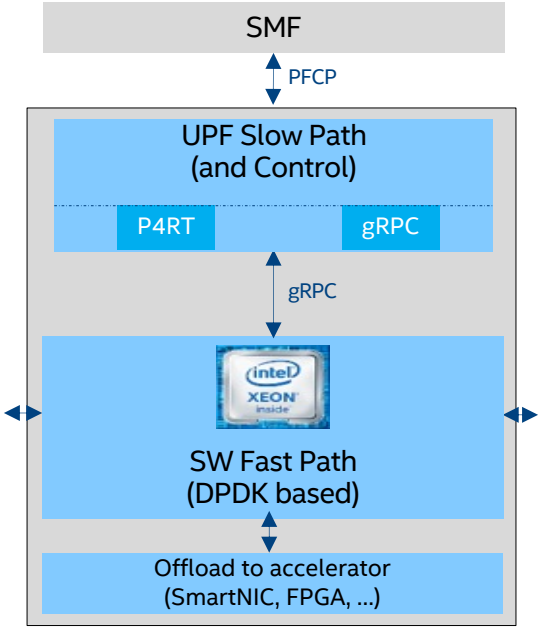

SW Fast Path Pros:
- Flexibility & support all features including hQoS, DDN, DPI, FW
- Support very large users' table
- Use of platform features : DDP, DLB, SGX

Limitations vs. HW Fast Path:
- Aggregate throughput
- Higher latency & jitter

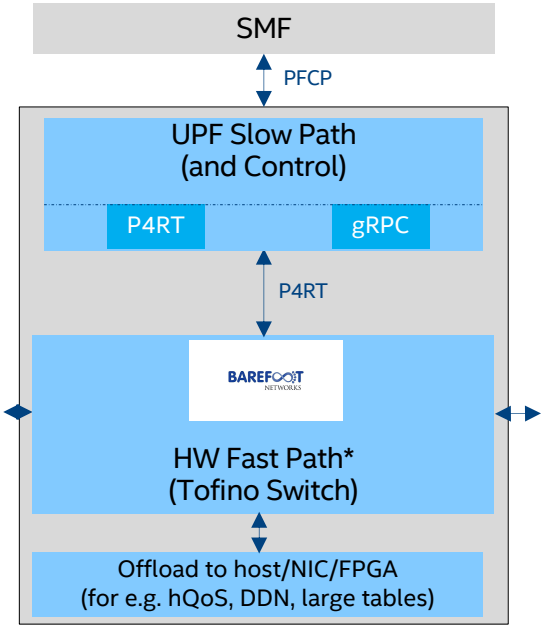

HW Fast Path Pros:
- Aggregate throughput
- Latency & jitter

Limitations vs. SW Fast path:
- Need to offload to CPU/FPGA/SmartNIC to support hQoS, DDN, DPI, FW
- Support for large number of users (flows in/out of TCAM create exception)

DDN: Downlink Data Notification
hQoS: Hierarchical QOS
DPI: Deep Packet Inspection
FW: Firewall
DDP: Dynamic Device Personalization
DLB: Dynamic Load Balancing
SGX: Secure Enclave

A flexible 5G UPF architecture optimized for specific deployment, e.g. edge or Central Office

UPF Processing Pipeline

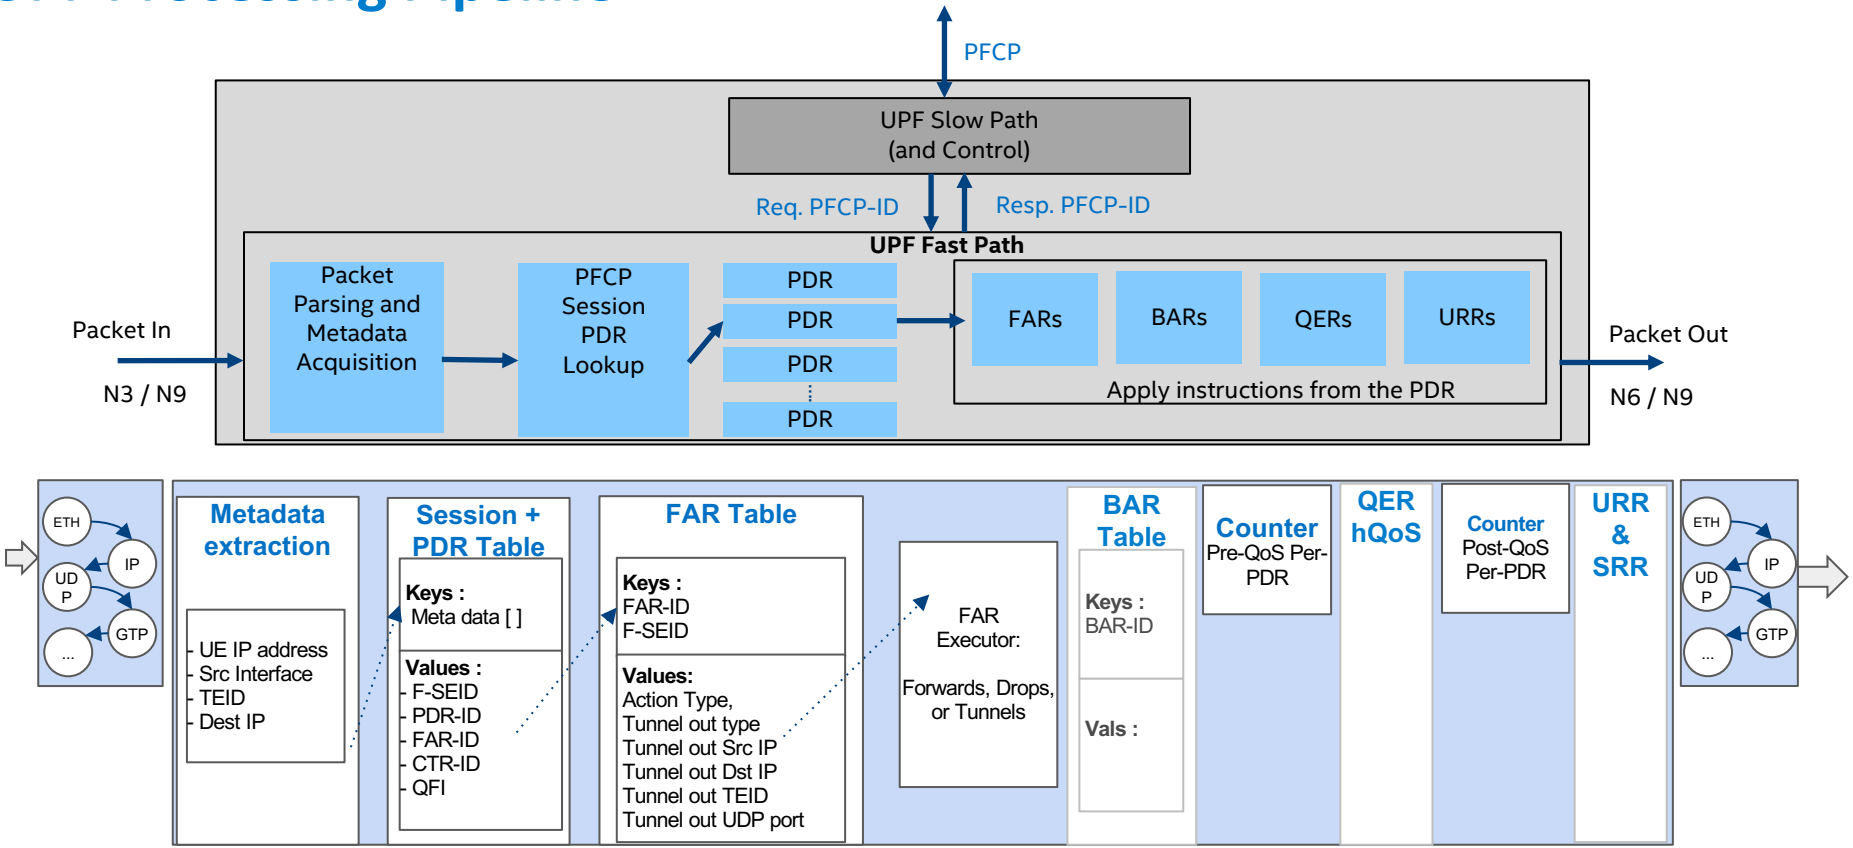

PFCP Session: PDR [ ], FAR [ ], BAR [ ], URR [ ], QER [ ], SRR [ ], ...
PDR : Packet Detection Rule [ ]
FAR : Forwarding Action Rule [ ] e.g. drop, forward, buffer, notify CP, duplicate, ...
BAR : Buffering Action Rule, e.g. how much data to buffer and how to notify the CP
QER : QoS Enforcement Rule [ ] -- Flow and service level marking
URR: Usage Reporting Rule [ ] -- Generate reports to enable charging functionality

UPF supports dual-mode 5G and LTE Core

5G : UPF Interoperating with Spirent 5G Emulator and other emulators
4G : Deployed on Aether's edges

Cloud Native SW w/ Enhance Platform Awareness (EPA) (1/3)

- CPU Core isolation & pinning
- Huge Pages
- Containers with multi-network interfaces & SR-IOV support in K8s

- Core pinning/affinity and isolation
  - CPU Manager for K8s
  - Automated CPU core mask for DPDK apps

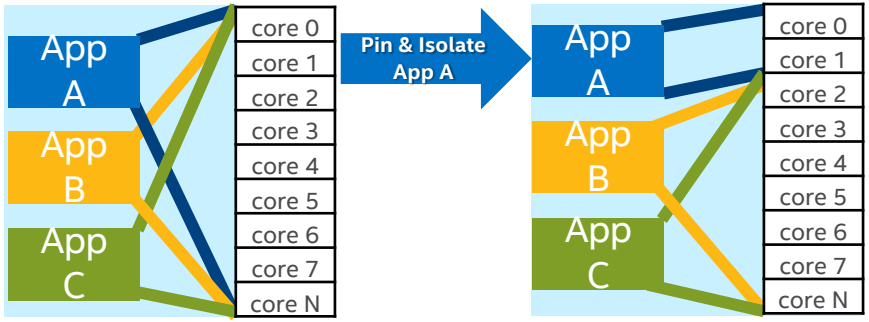

- Huge pages

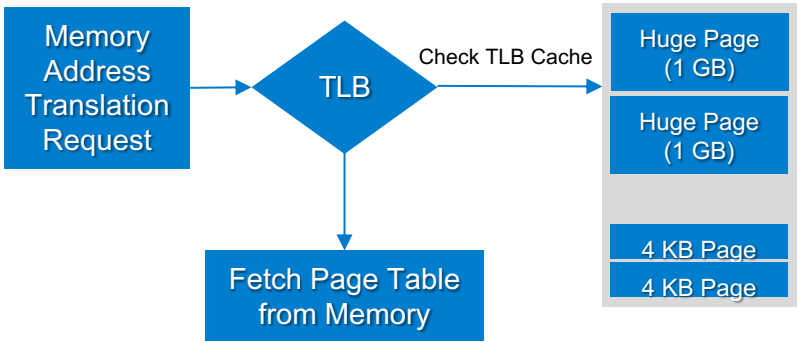

https://networkbuilders.intel.com/docs/kilo-a-path-to-line-rate-capable-nfv-deployments-with-intel-architecture-and-the-openstack-kilo-release.pdf

Cloud Native SW w/ Enhance Platform Awareness (EPA) (2/3)

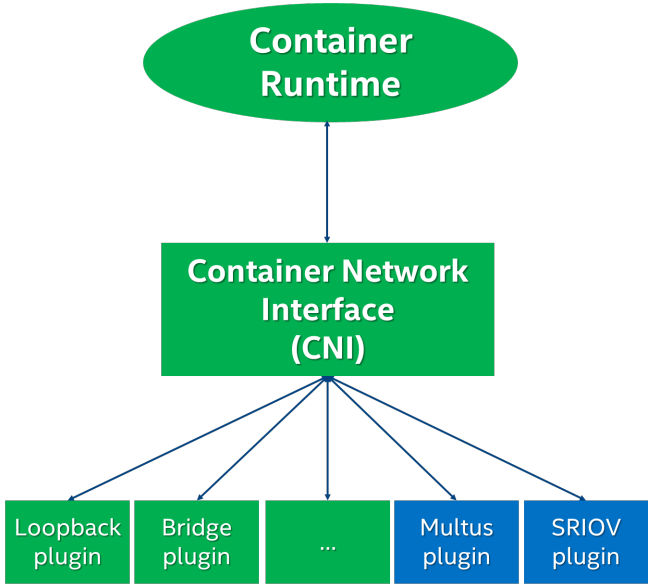

Logical

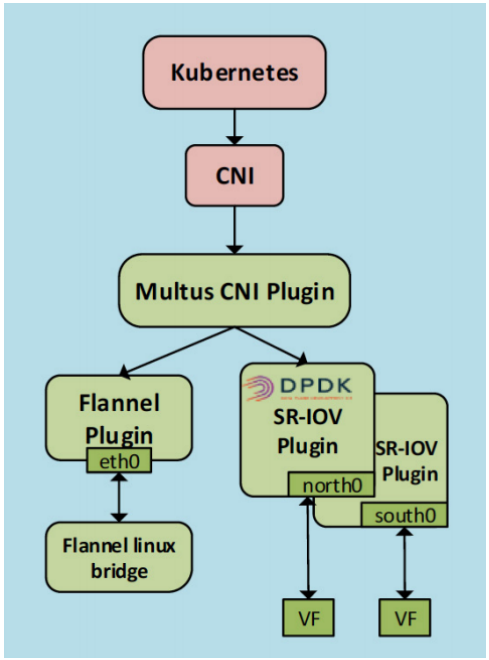

Physical Manifestation

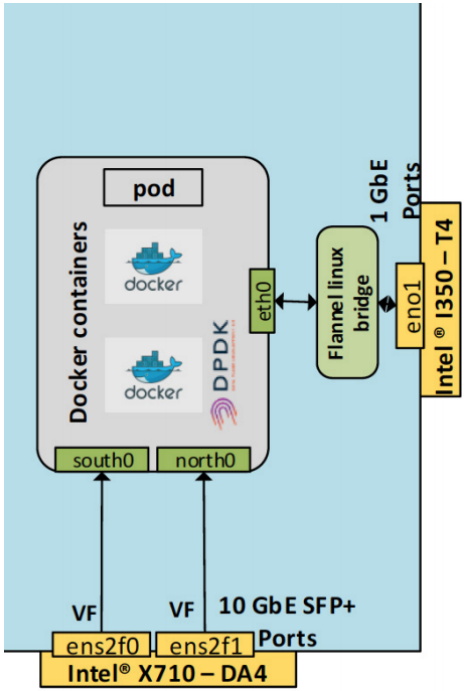

- Multiple networks and high throughput I/O for DP
- Multus CNI plugin and SR-IOV CNI plugin (enables VFs + DPDK user space drivers)

https://builders.intel.com/docs/networkbuilders/enabling_new_features_in_kubernetes_for_NFV.pdf

Cloud Native SW w/ Enhance Platform Awareness (EPA) (3/3)

- Native – Bare metal processes, no containers, no orchestration
- K8s – Docker containers orchestrated by K8s with EPA knobs ON / OFF

Test	User Space Driver	CPU Pinning	Huge Page	Pkts/sec*	(w/ noise)
Native	Yes	Yes	Yes	**1,550K**	(1,100K)
K8s	Yes	Yes	Yes	**1450K**	(1.150K)
K8s	**No**	Yes	Yes	**750K**	(650K)
K8s	Yes	**No**	Yes	**1450K**	**(400K)**
K8s	Yes	Yes	**No**	**1200K**	(1100K)

*50K Granularity, 1 CPU Core

- Cloud Native SW w/ EPA achieves performance similar to bare-metal processes
- Supporting additional features like AF-XDP, DDP (Device Data Personalization)

Deployment in Aether : Enterprise Edge-as-a-Service

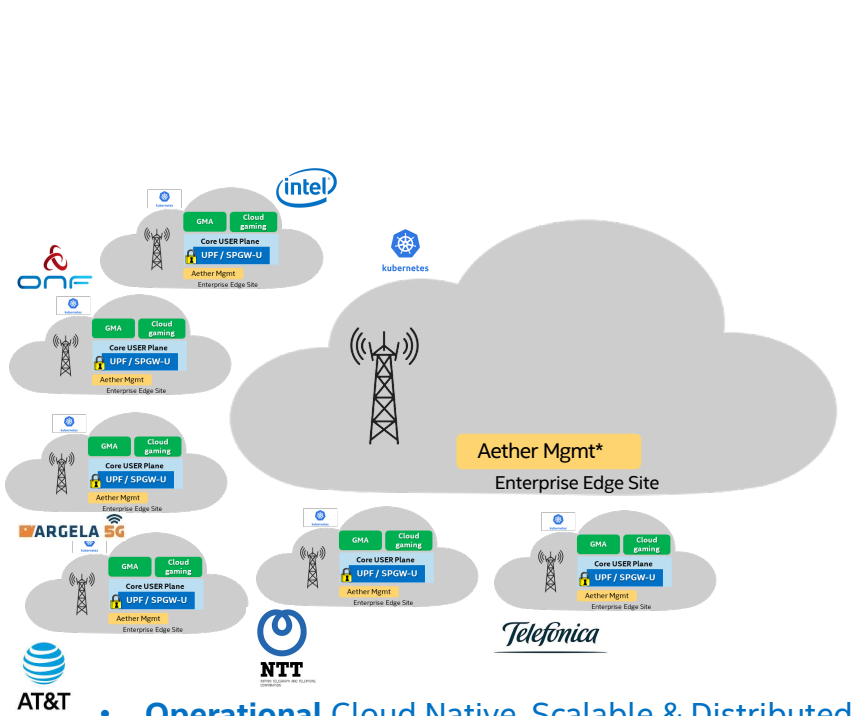

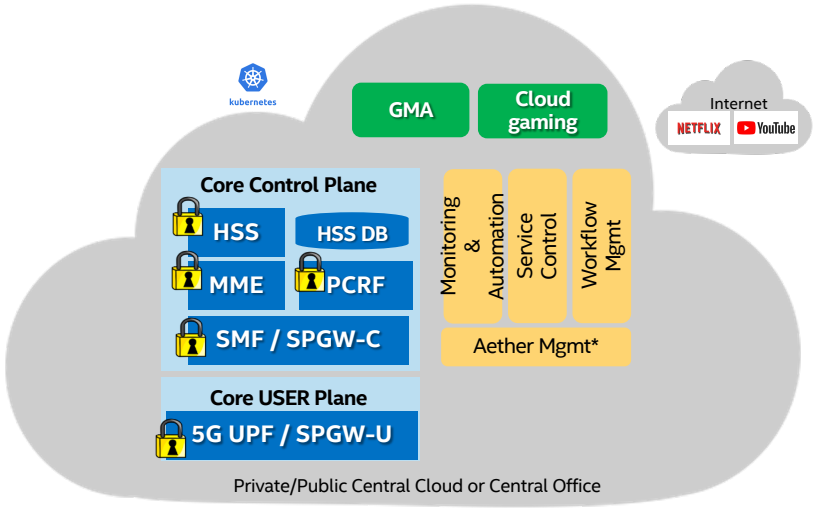

Intel SGX based secure container

- **Operational** Cloud Native, Scalable & Distributed Gateways, with central control in private/public clouds (Google / Azure)
  - Multiple Aether edges deployed e.g. AT&T, NTT, Telefonica, Argela, Ciena, Intel, ONF – More to come
  - To be deployed as part of DARPA "Verifiable Closed Loop Control Network" with Stanford, Princeton and Cornell
  - ONF acts as "operator" for Day-0, Day-1, Day-2 (Deployment, reliability, support) - Benefit platform maturity
  - Deploy and evaluate benefits of edge applications or capabilities, e.g. Cloud Gaming, GMA, etc

* : Aether SW developed by ONF

Edge Service - Cloud Gaming

Today: Interactive Frame Streaming Model

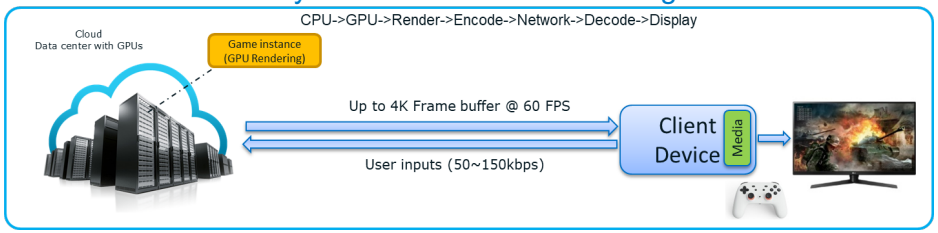

E2E latency about ~166ms @1080p

Edge based deployment models

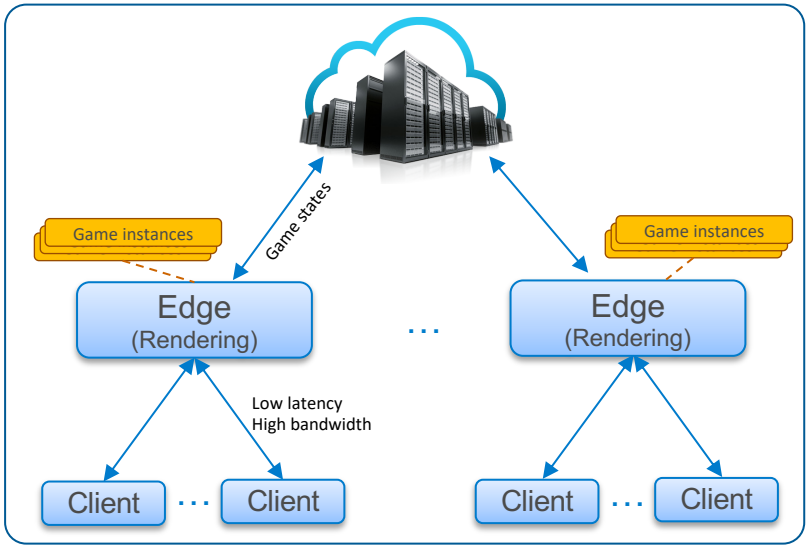

Key Challenges:

- Reduce E2E Latency *(as close to client gaming <70ms)*

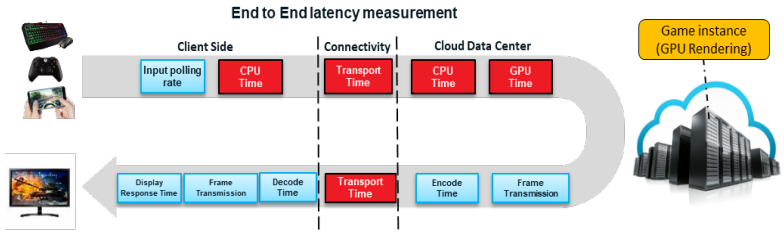

- Improve Quality *(4K or higher without aliasing / encoding artifacts)*
- Provide constant throughput *(60 or higher FPS)*

Source: Selvakumar Panneer, Intel Labs

Edge Service : Generic Multi-Access (GMA) Ref. Design

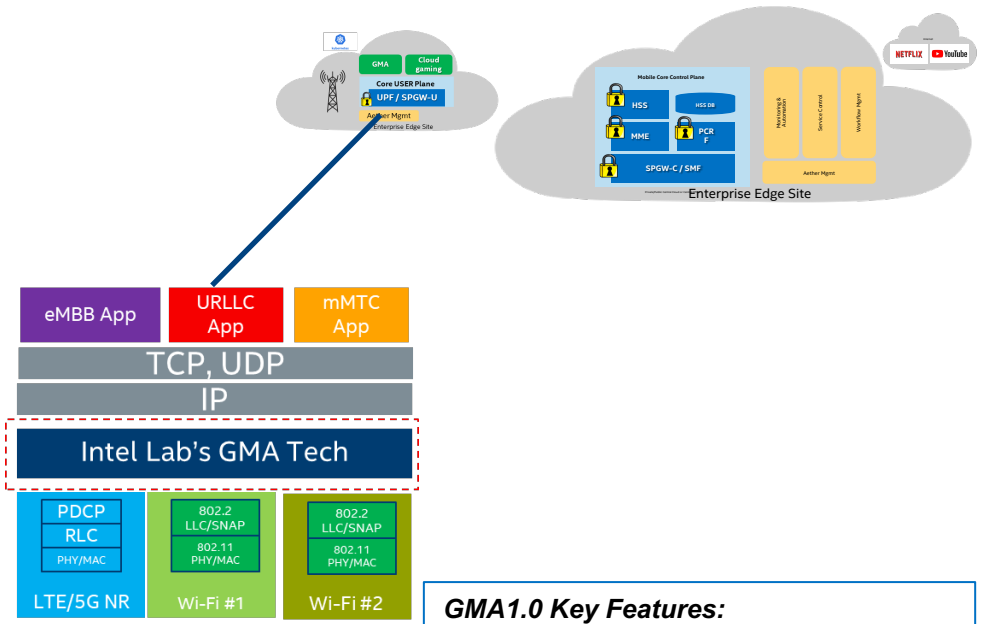

GMA Data-Plane
Packet splitting, reordering, retransmission, fragmentation, concatenation, coding

GMA Control Plane
- **Management:** e2e signaling/protocols
- **Measurements:** signal strength, traffic load, mobility, QoS, packet loss, latency

GMA1.0 Key Features:
- **seamless handover**: moving traffic seamlessly from Wi-Fi to Cellular when detecting weak Wi-Fi signal
- **downlink boost**: using both Wi-Fi and Cellular to increase the download speed when detecting congestion over Wi-Fi link
- **uplink redundancy**: sending uplink traffic over both Wi-Fi and Cellular to increase reliability and reduce latency

(ready for trial and ecosystem engagement)

Source: Jing Zhu, Intel Labs

Application: File Download (Iperf)

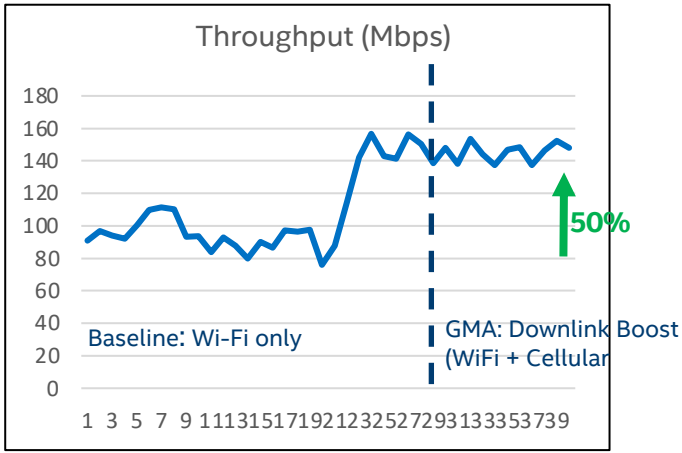

OWD: One-Way-Delay

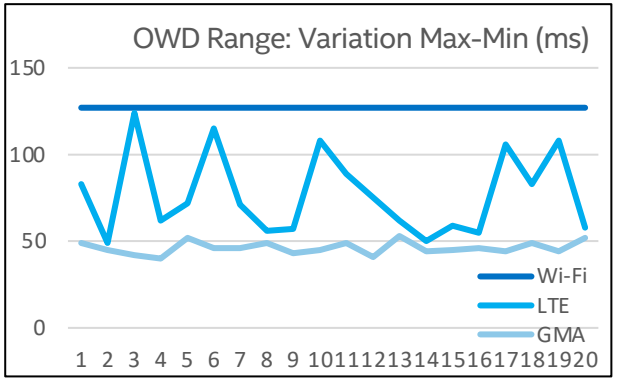

App: Google Stadia over Intel Hotspot Wi-Fi + AT&T LTE Cellular

Open 5G Network Infrastructure to Accelerate Edge Deployment

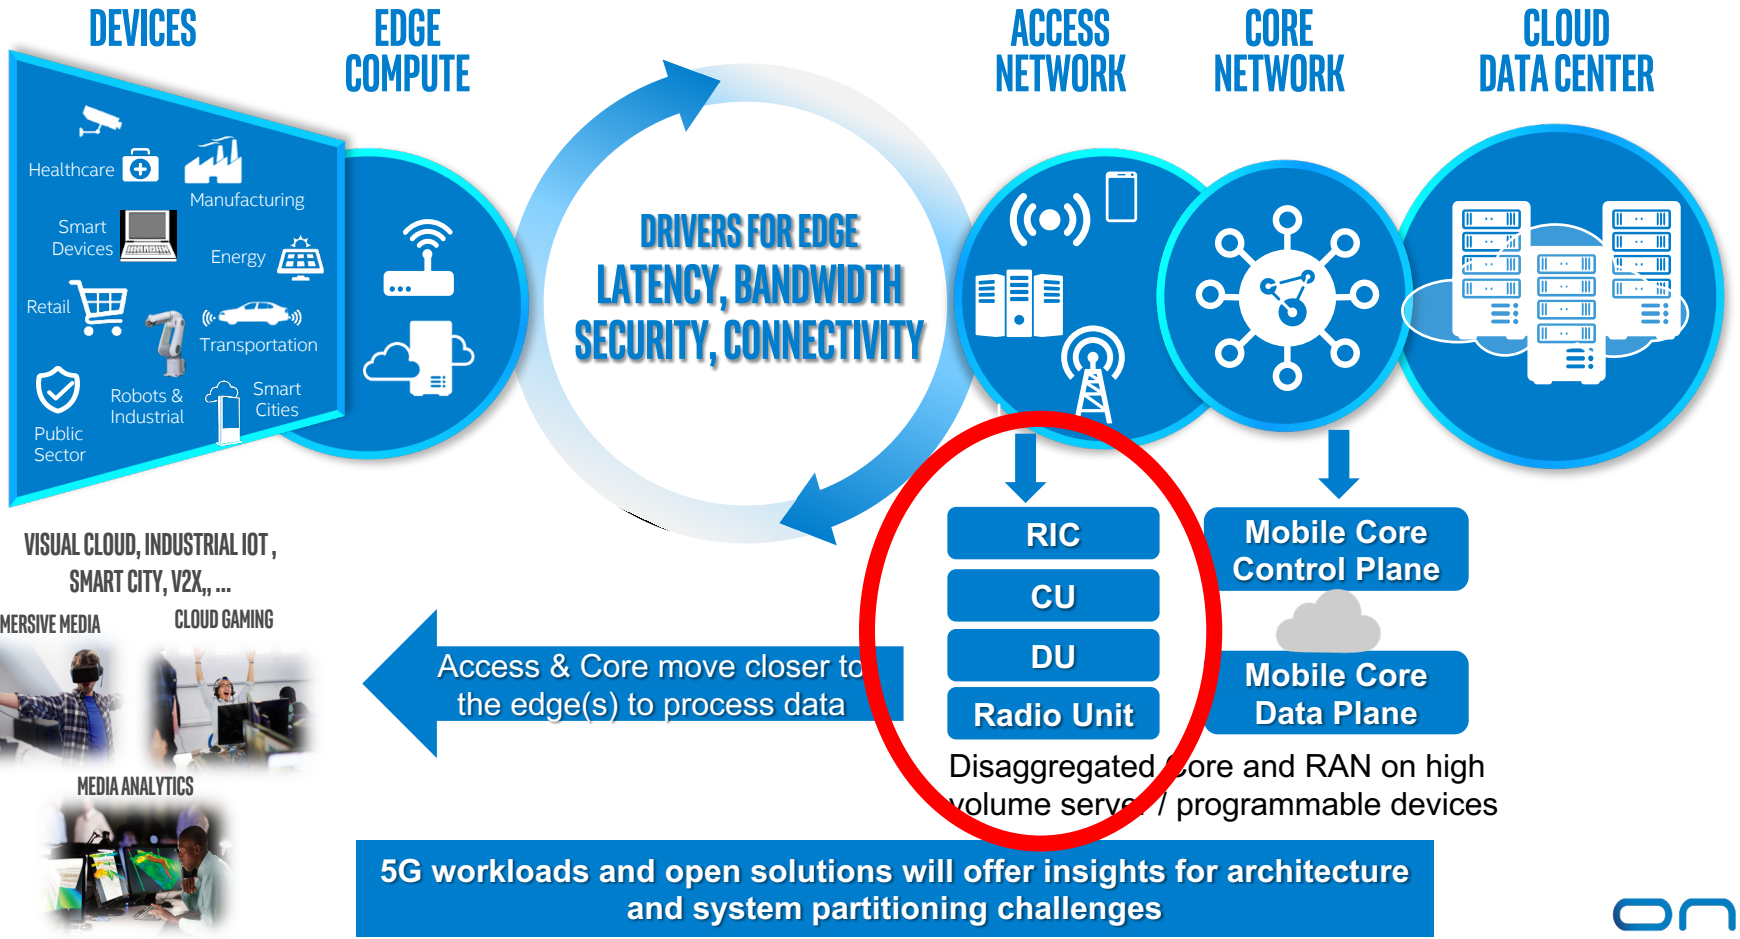

Ran Intelligence Services for Near Real-Time RIC

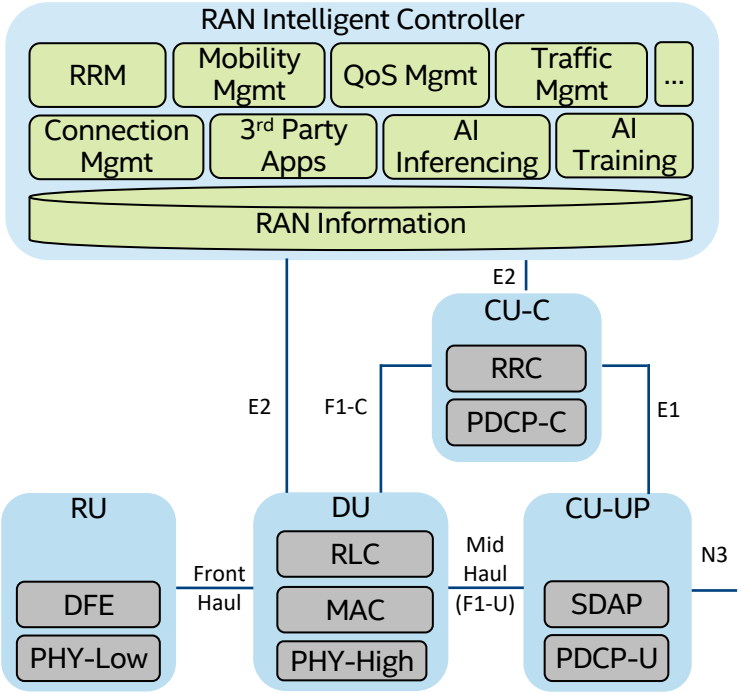

Intel Labs working with ONF on value-add services for near-RT RIC

- **2 Initial services planned:** Connection management & multi-access traffic management
- **Integrated using open interfaces,** but not open sourced
- **Extensions of E2 & A1 interfaces** to enable above services
- **Extensions to AI/ML framework**

Source: Shilpa Talwar, Jing Zhu, Hosein Nikopour, Shu-ping Yeh, Meryem Simsek, Mahima Mehta – Intel Labs

Near RT-RIC Service #1 : Connection Management in RAN

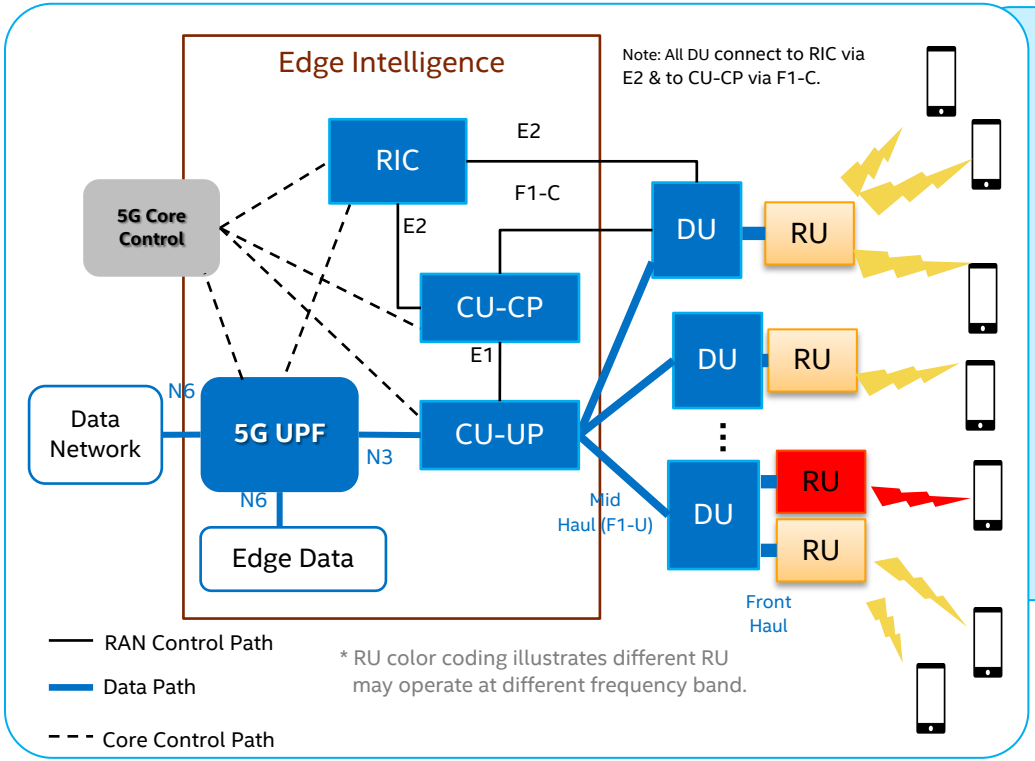

Semi-static connection management for band/cell selection

- **Problem**: Select best band/cell(s) for each UE based on radio conditions, traffic load and QoS requirements
- Update time scale >50ms
- **Actions**: UE cell association via UE initial access or handover
- **Target**: load balancing and QoS management

Source: Shilpa Talwar, Jing Zhu, Hosein Nikopour, Shu-ping Yeh, Meryem Simsek, Mahima Mehta – Intel Labs

Near RT-RIC Service #2 : Multi-Access Traffic Management

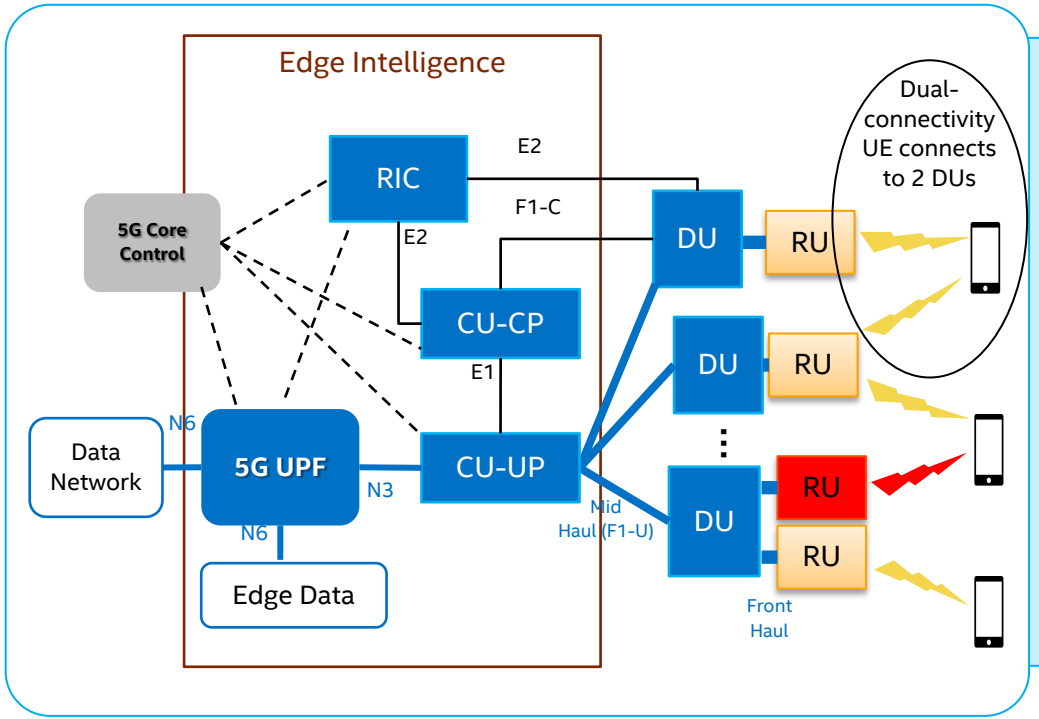

Dynamic traffic management and packet routing for multi-connectivity or multi-RAT

- **Problem**: Determine the best packet routing strategy from multiple diverse connections (DU or WiFi AP) to UE based on radio quality and QoS targets
- Update time scale: 10-50ms
- **Actions**: Add/Remove 2nd connection, change of packet routing rules.
- **Target**: load balancing and QoS management

Note: Possible DC configurations: LTE+LTE, LTE+NR, or NR+NR. Solutions also applicable to cellular unlicensed convergence.

Source: Shilpa Talwar, Jing Zhu, Hosein Nikopour, Shu-ping Yeh, Meryem Simsek, Mahima Mehta – Intel Labs

Ran Intelligence Services for Near Real-Time RIC

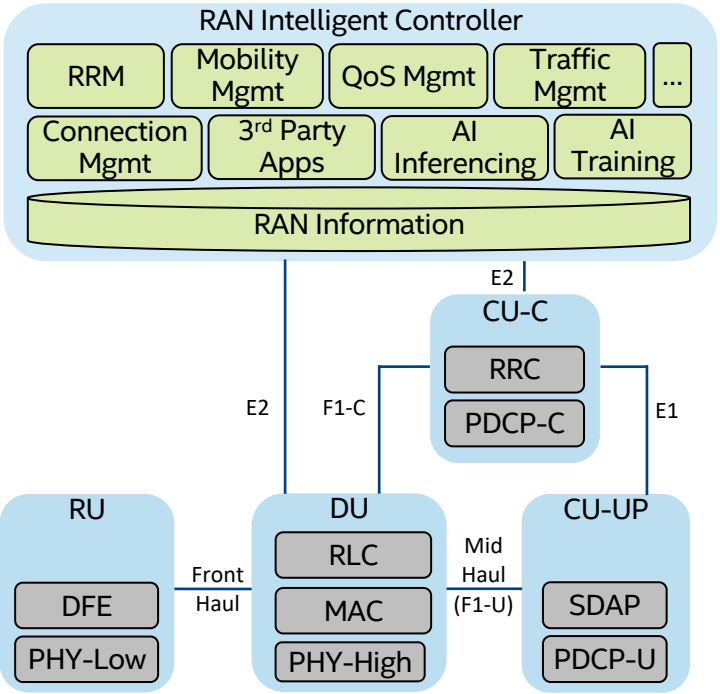

Contributions to ORAN WG-3

RAN Control / Configuration

- **Dual-connectivity Control:** Change of bearer termination point, bearer types & control of bearer split ratio
- **Reliability enhancement Configuration:** Packet duplication, rate selection with lower target BLER

1. "Adding DC related DRB control for QoS UCR," O-RAN WG3 Web Conf. #62
2. "Include reliability enhancement control for QoS UCR," O-RAN WG3 Web Conf. #64

RAN Measurements

- PRB usage at DU, buffer status, data volume, location/velocity of UE, delay, packet loss

3. "Additional E2 Requirements for Traffic Steering," O-RAN WG3 Web Conf. #61
4. "UE Location and Velocity information for Traffic Steering use case," O-RAN WG3 Web Conf. #63

Source: Shilpa Talwar, Jing Zhu, Hosein Nikopour, Shu-ping Yeh, Meryem Simsek, Mahima Mehta – Intel Labs

Demo of host based Dual-Mode 5G/LTE UPF

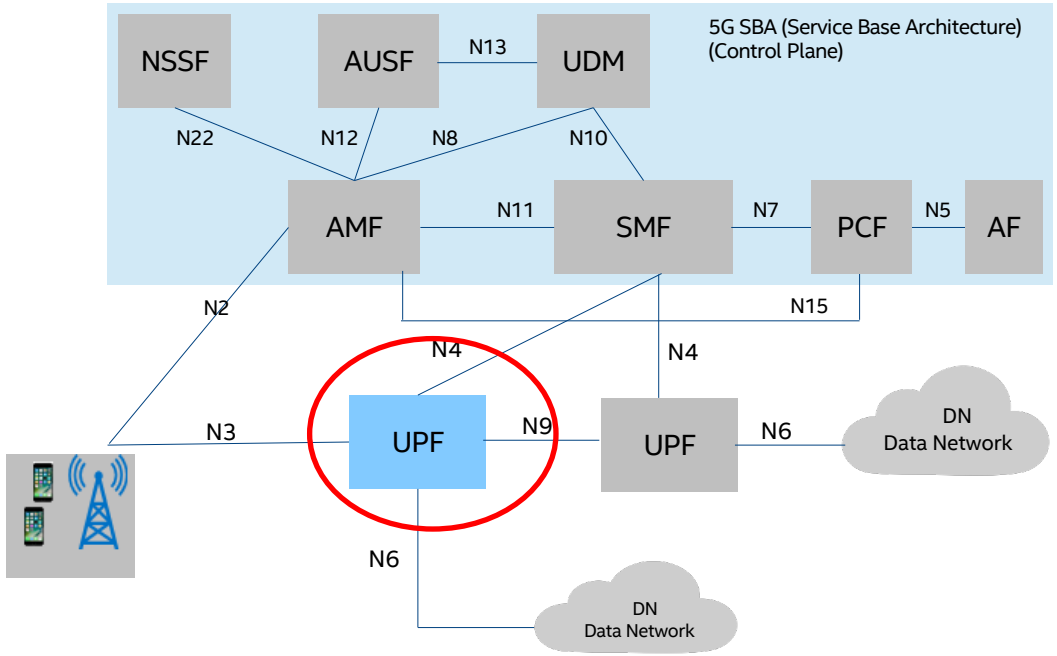

Spirent Emulator

SPIRENT Communications Inspired Innovation

UPF SW Architecture

Youtube Video

Spirent Details

Topologies for User Plane Testing

(vNF/cNF)-UPF Isolation

spirent

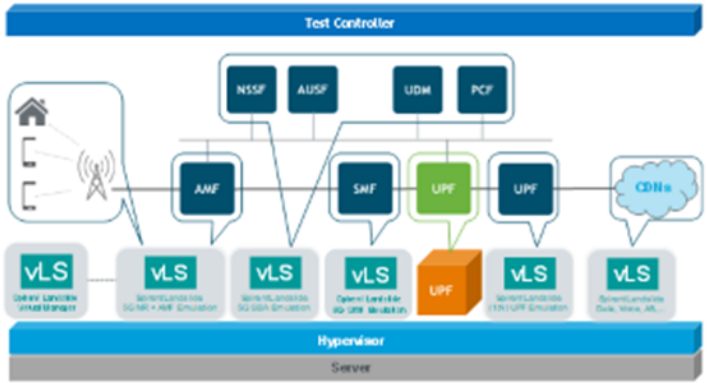

Features:

- ✓ Multi-node 5G SA core radio access(gNB) and control plane emulation (AMF+SMF)
- ✓ Integrated 5G Service-nodes (PCF, UDM, AuSF, ...)
- ✓ Multi-peer N3,N4 and N6 interface support
- ✓ High session/QoS Flows scale
- ✓ High throughput
- ✓ UPF emulation for N9 testing (optional)
- ✓ Content network emulation (optional)

Landslide Test Cases:

- UPF Nodal
- UPF, NRF Node Emulators (optional)

Opportunities to Contribute to development & deployment

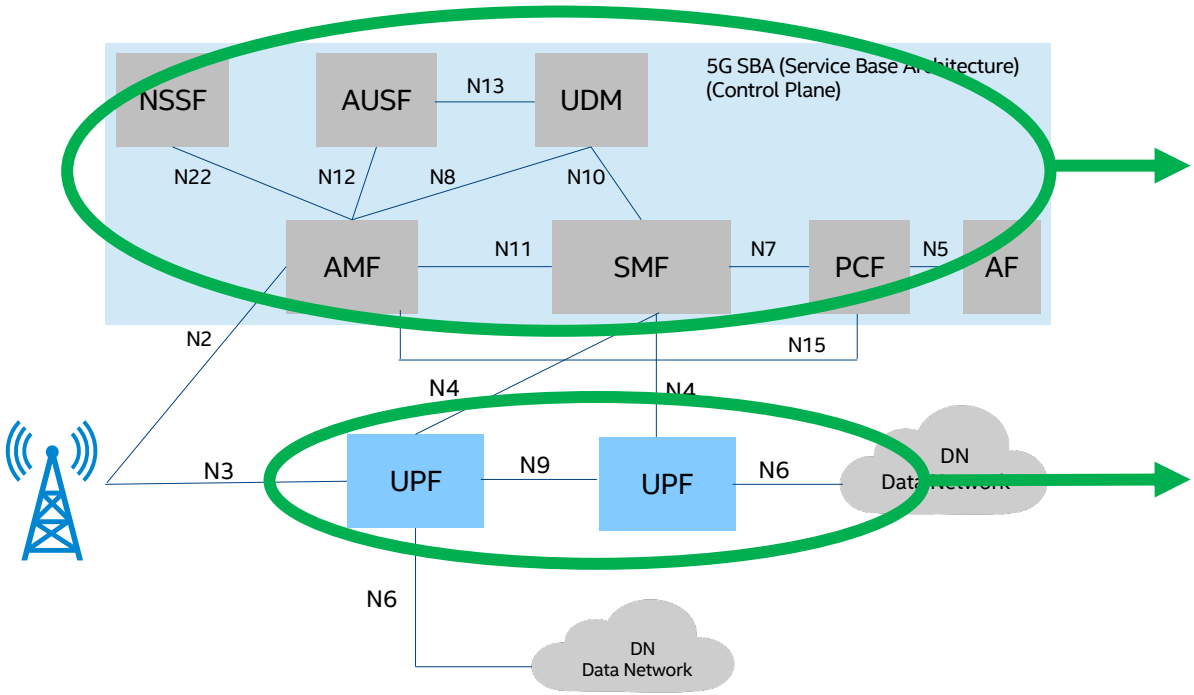

- Significant opportunities to collaborate & contribute with System Integrators, Operators and other partners
- Opportunities to add functionality enabling specific usage model(s), e.g. TSN (Time Sensitive Network), …

Summary

- An open solution from RAN to Core will create a vibrant and healthy eco-system
- Upcoming 5G workloads and open solutions will offer a unique insights for architecture and system partitioning challenges
- You have opportunities to join, contribute and make it a successful architecture & technology evolution

5G Transformation with Open Source

Thank You

Christian.maciocco@intel.com

https://www.opennetworking.org/omec/
https://www.opennetworking.org/aether/

2020

Dual-mode 5G/LTE UPF SW Architecture

Berkeley Extensible Soft Switch: Revisiting the data plane

- Monolithic framework
  - Static + Dynamic lib linkages
- Compile-time config

- Modularize the framework
  - Graph-based modular architecture
- Run-time config
- Debugging ability

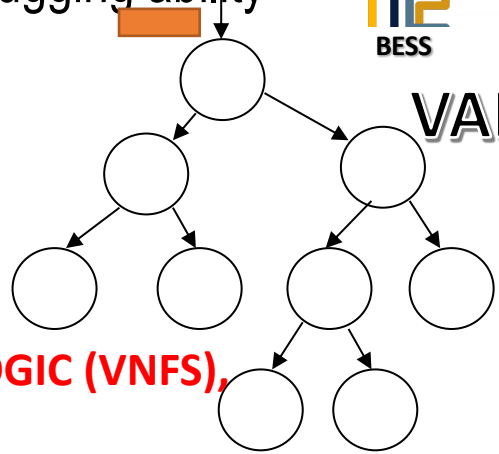

DEVELOPER ONLY FOCUSES ON THE CORE BUSINESS LOGIC (VNFS), & NOT THE SOFTWARE INFRASTRUCTURE

BESS in Industry & Academia

- Red Hat unsolicited Data Planes Review:
  - CTO Control and Data Plane Full Investigation Doc
  - Data Plane Findings - Slide Presentation
  - Data Plane Performance Test Plan

- ACM CoNEXT '19: "Comparing the Performance of State-of-the-Art Software Switches for NFV," Institut PolyTech de Paris, Nokia Bell Labs

"BESS achieves both high throughput and low latency in phy-to-phy, phy-2-virtual, and 1-VNF loopback scenarios."

- Arista vEOS Dataplane router in DPDK mode
  - https://www.arista.com/en/cg-veos-router/veos-router-dpdk-mode
  - https://www.arista.com/en/cg-veos-router/veos-router-general-troubleshooting

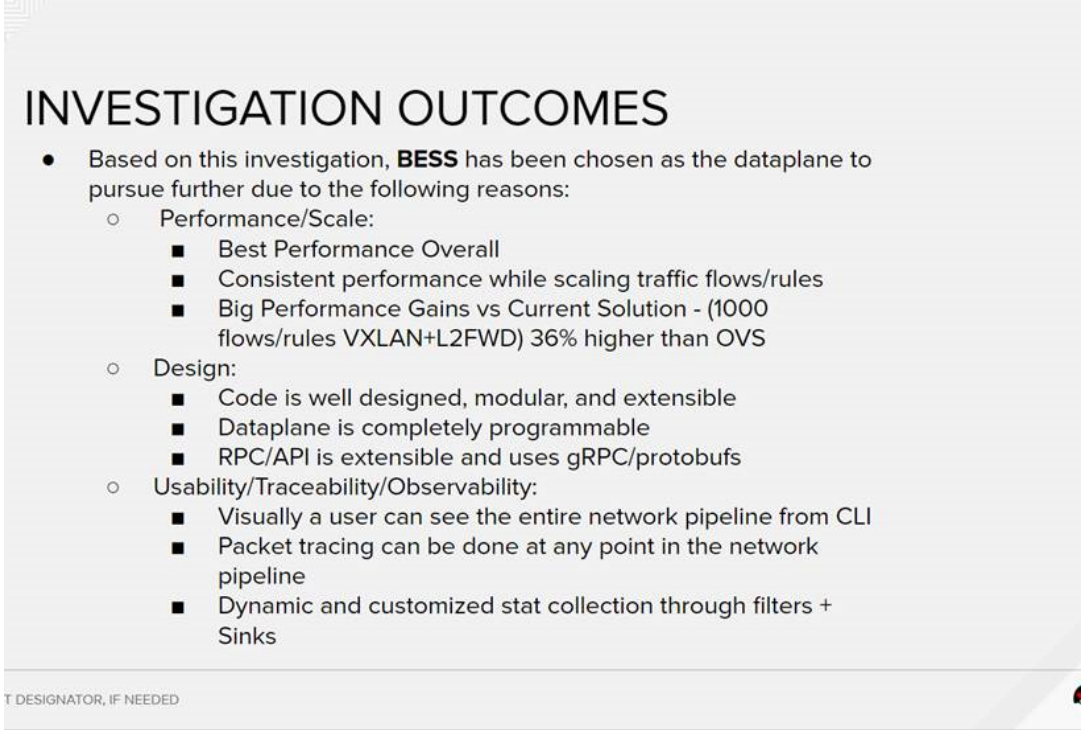

BESS Motivation: desired feature set

- Graph-based framework
  - Modularity
    - Addition of modules within the NF pipeline
    - Composability of functionality specific to the use case without invasive code changes
  - Abstract infrastructure complexities from the NFs
    - Model: run-to-completion←→ Pipelining (inter-changeable)
    - Dual interface (S1U/N3, SGi/N6) to single interface
    - CPU, mem allocation
  - Debug capabilities

→ **Ability to configure dataplane at run-time**

BESS Intro

Programmable platform for data plane development

- Clean-slate internal architecture with NFV in mind
  - Highly flexible & customizable
- Creating BESS applications
  - Modular pipeline represented as a directed acyclic graph
  - Each module can run arbitrary code
  - Independently extensible & optimizable
- Configure & control BESS
  - Via NF controller

BESS Architecture Overview

NET_CONTROLLER
Policy updates
via CP

DAG of interconnecting modules

HOST_CONTROLLER
Neighbor updates
via OS

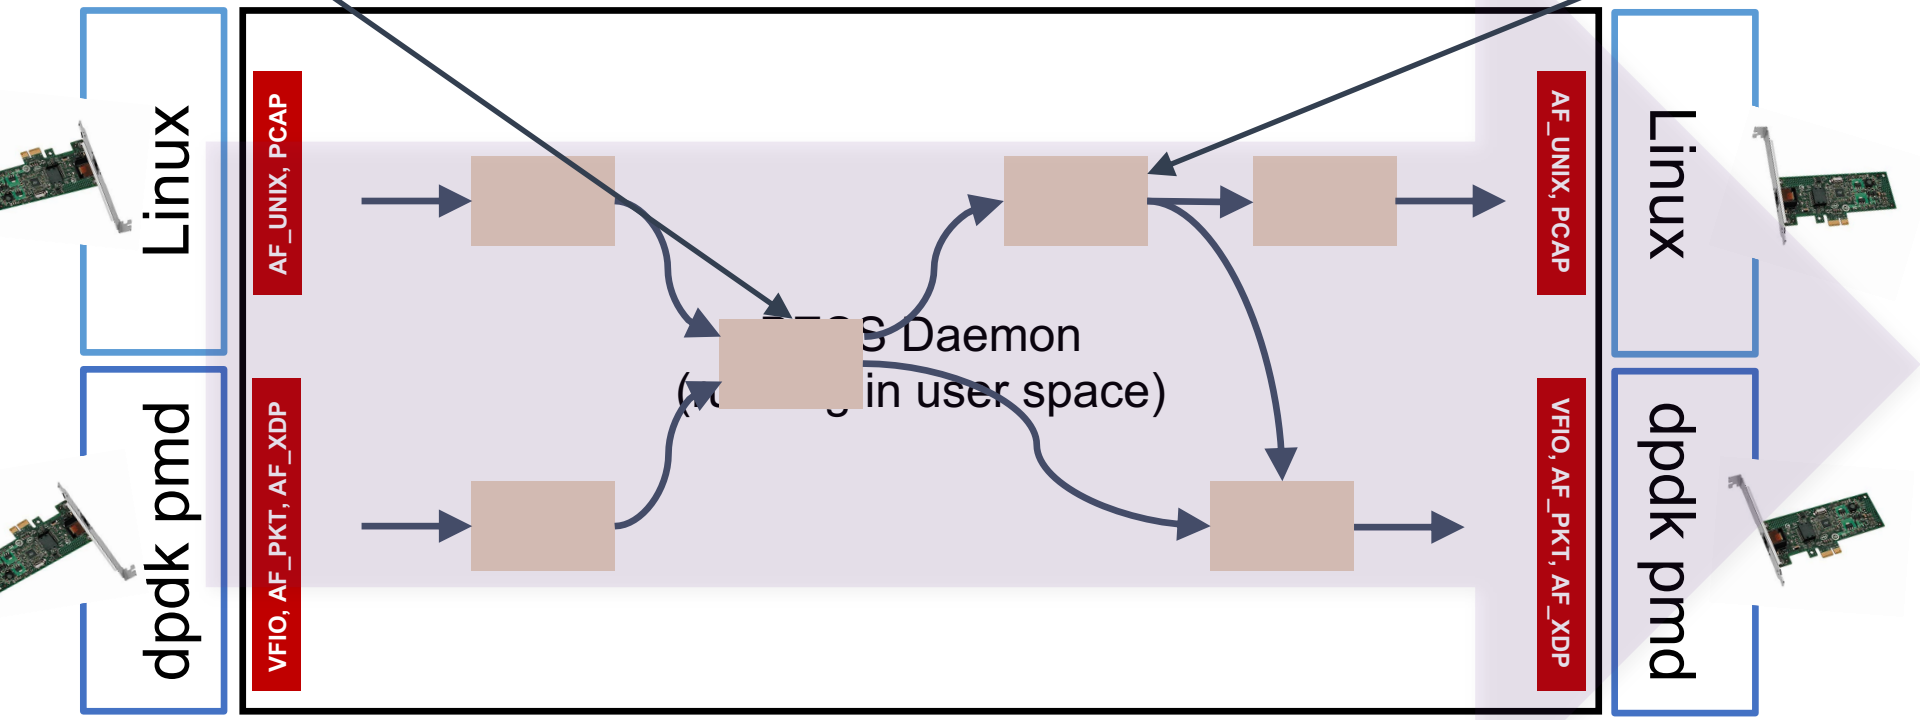

UPF-EPC over BESS: Resource Aware CPU Scheduling

Allows flexible scheduling policies for the data path

- In terms of CPU utilization & bandwidth

UPF-EPC over BESS: Resource Aware CPU Scheduling

Allows flexible scheduling policies for the data path

- In terms of CPU utilization & bandwidth

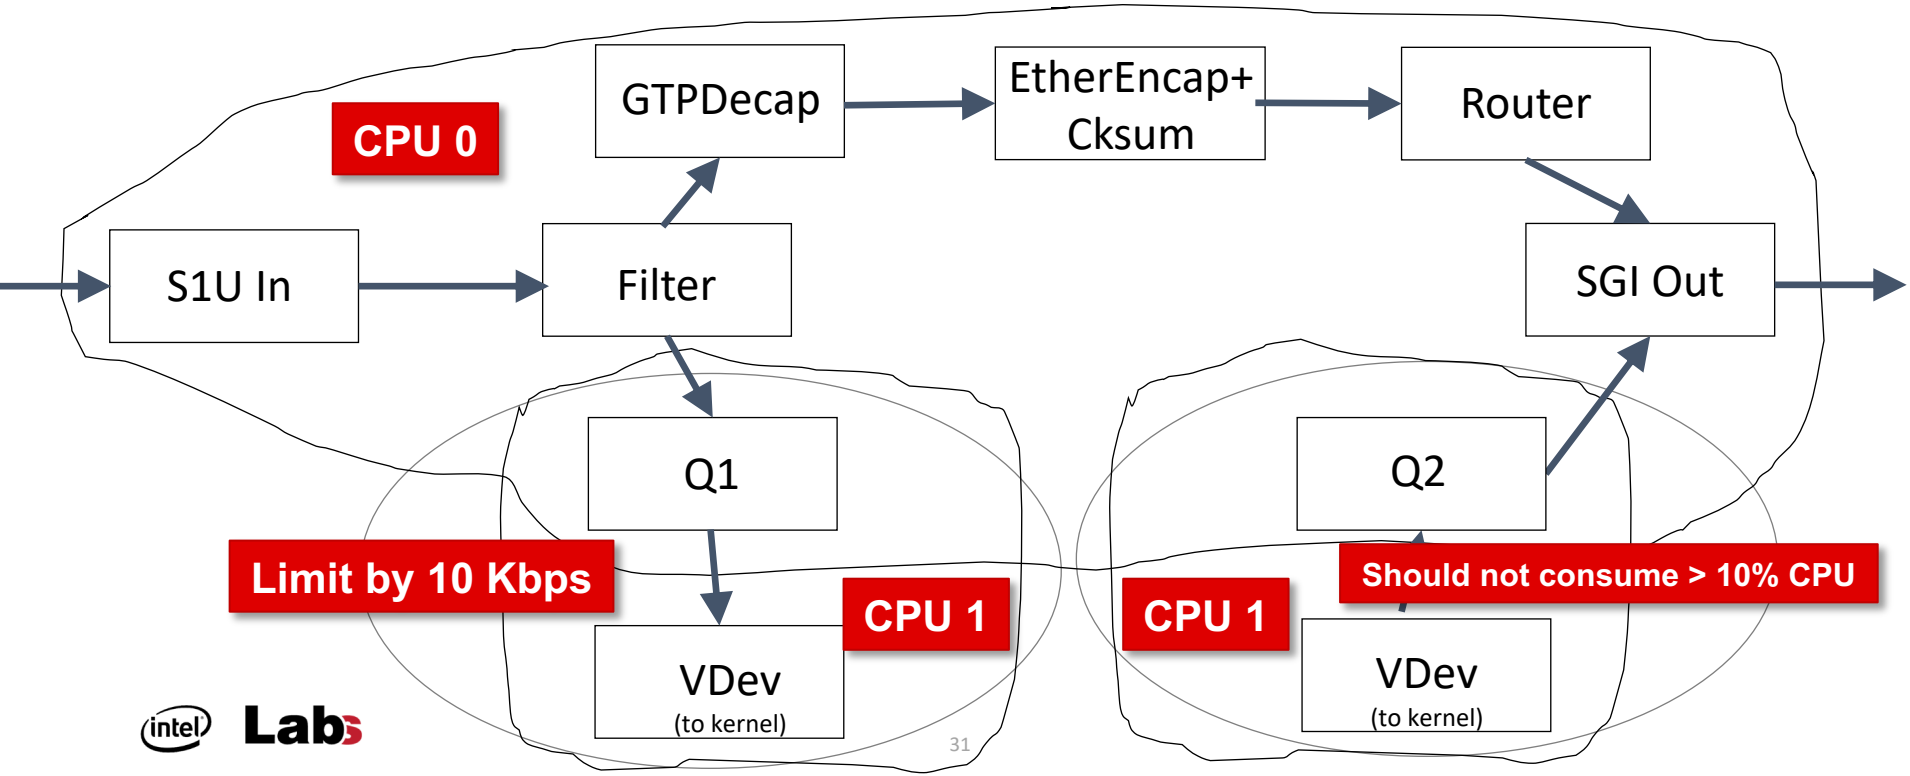

UPF-EPC over BESS (1/3)

Key benefits of architecting user-plane with BESS

- Modular data plane
  - Developers concentrate only on core business logic (i.e. VNF development) and not the software infrastructure development
    - Mostly rely on built-in BESS modules resulting in a thin stack
    - Controllers can be created in any gRPC-supported language
      - (Route+L2 neighbor) python controller based on pyroute2: SLOC ~= 350
  - Ease of customizing pipeline at runtime
    - *e.g.* CPU scheduling, adding/removing specific modules
- Configuration ease
  - Multi-workers enable/disable at ease
    - Economical usage of CPU usage
    - Can run individual modules on different CPUs
      - Run-to-completion vs pipeline vs hybrid become run-time choices (& not compile-time)
  - No need to restart the daemon process for configuration updates
- Monitoring ease at runtime
  - `tcpdump`
    - Monitor traffic over any module
  - Visualization tool
    - Web interface

Dual Mode 5G/LTE UPF BESS Pipeline - (A subset of the pipeline in the picture)

- I-UPF and A-UPF
- Interoperating with Spirent Emulator
  - PFCP based N4 I/F
  - N3, N6, N9 for data traffic

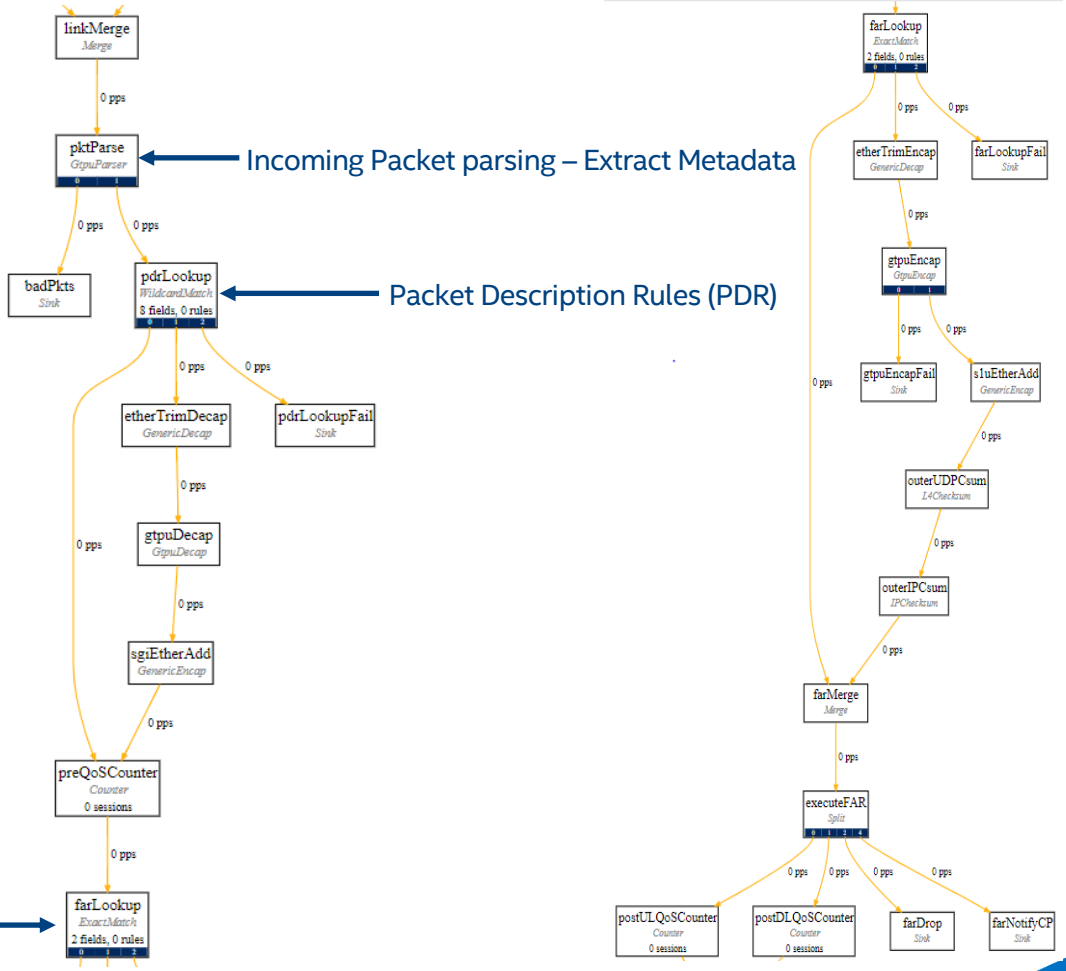

Spirent Landslide

5G User Plane Testing

Technical implications and challenges

100x peak data rates (20Gbps+)

100x traffic capacity

10x experienced throughput

10x decrease in latency

New goals

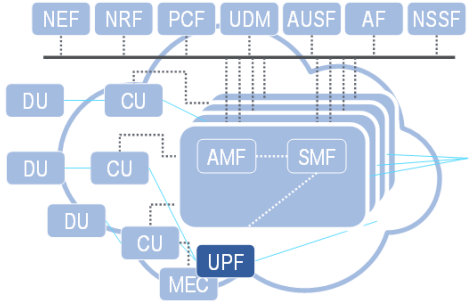

Network deployment

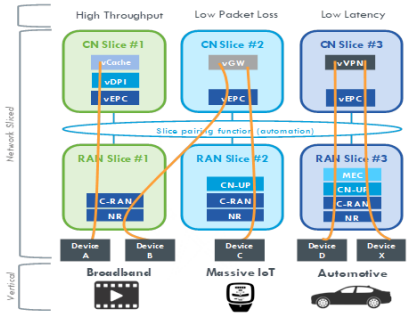

Network slicing

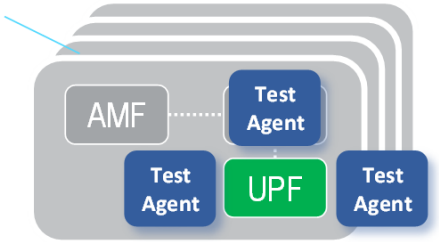

User plane design

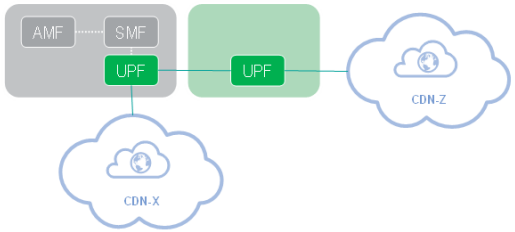

Chaining

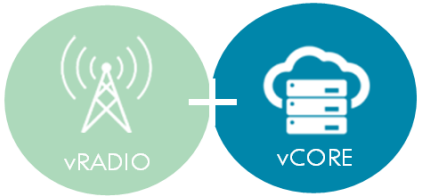

Successful virtualization

What is Spirent Landslide?

High Scale ‘Any-G’ Mobility Layer Tester and Emulator

C100/C50

Millions of Any-G Connections with Full Mobility

Line Rate Traffic Generation

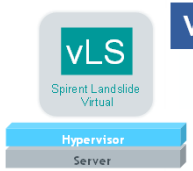

Virtual

Real World User Traffic (Voice, Video, Internet, Apps…)

Carrier and Smartphone Call Modeling

E10 (OTA)

Node and Network Emulation

Topologies for User Plane Testing

(vNF/cNF)-UPF Isolation

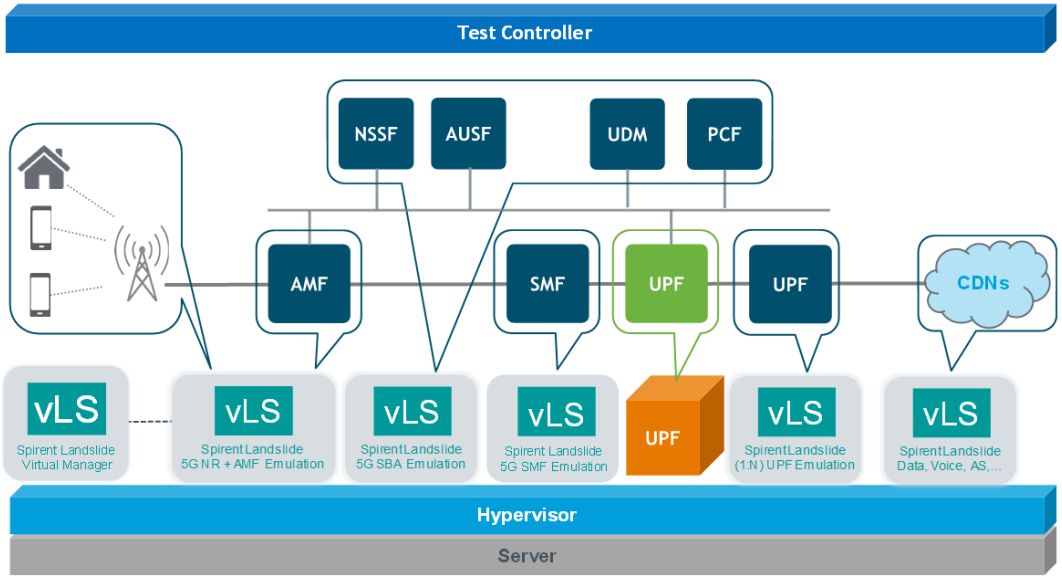

Features:

- ✓ Multi-node 5G SA core radio access(gNB) and control plane emulation (AMF+SMF)
- ✓ Integrated 5G Service-nodes (PCF, UDM, AuSF, ...)
- ✓ Multi-peer N3,N4 and N6 interface support
- ✓ High session/QoS Flows scale
- ✓ High throughput
- ✓ UPF emulation for N9 testing (optional)
- ✓ Content network emulation (optional)

Landslide Test Cases:

- UPF Nodal
- UPF, NRF Node Emulators (optional)

Busy Hour Call Modeling

Control & User Plane call modeling

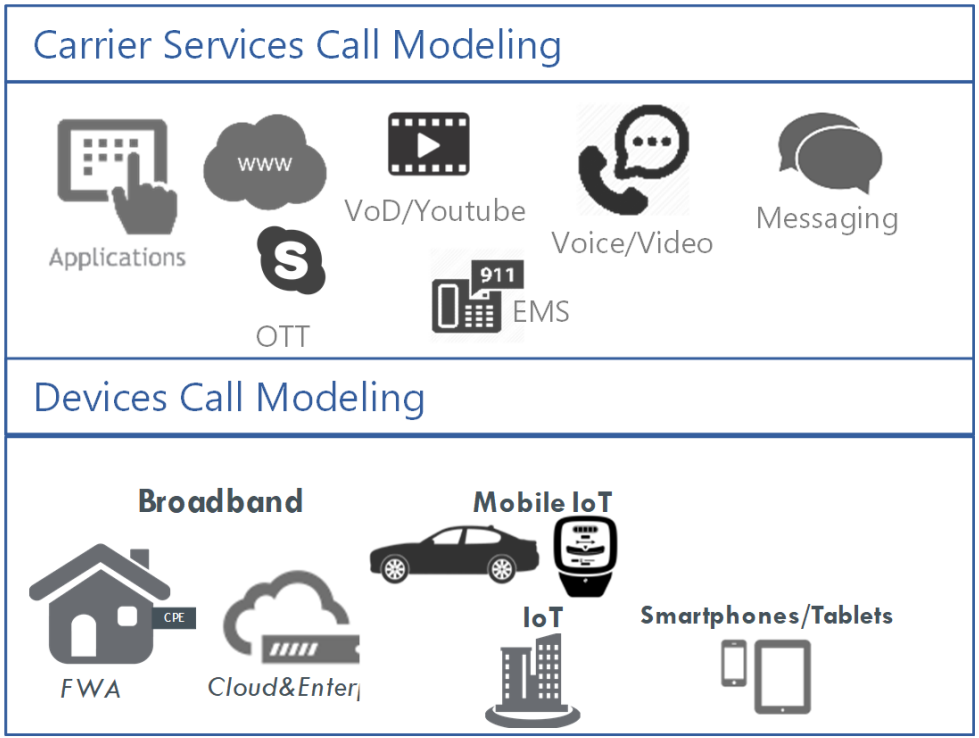

Features:

- ✓ High throughput and PPS on all topologies
- ✓ User Plane call modeling based on subscriber profiles (eg; FWA)
- ✓ User Plane call modeling based on service types (eg; Video on Demand)
- ✓ Control Plane call modeling impacting user plane behavior (eg; IRAT HO sequences)