

Low Latency 5G UPF Using Priority Based 5G Packet Classification

Communications Service Providers 5G UPF Packet Core



About the Authors

This whitepaper is jointly developed by SK Telecom and Intel Corporation.

SK Telecom —
5GX Labs, R&D Center
DongJin Lee

5G Core Architect
dongjin@sk.com

JoongGunn Park
5G Core Architect
clark.park@sk.com

Intel Corporation —
Data Platforms Group

Andriy Glustsov
Software Architect
andriy.glustsov@intel.com

Khaled Qubaiah
Software Architect
khaled.qubaiah@intel.com

Peter McCarthy
Software Architect
peter.mccarthy@intel.com

Rory Browne
Platform Solutions Architect
rory.browne@intel.com

Abstract

This paper summarizes a collaboration between Intel and SK Telecom to demonstrate a 5G Standalone (SA) User Plane Function (UPF) capability based on 2nd generation Intel® Xeon® Scalable Processors and Intel® Ethernet 800 Series Network Adapters. The results demonstrate that low latency and jitter can be achieved for a range of applications when using standard server infrastructure for the 5G SA UPF.

Executive Summary

The ability to efficiently deliver on the 5G requirements (eMBB, mMTC, and URLLC) and implement a common NFV infrastructure with high levels of throughput, utilization, and determinism is a key enabler to deploy a virtualized UPF system for 5G and beyond.

The 5G SA UPF, developed by Intel and SK Telecom, shows improved performance in latency and jitter for high priority traffic while still running best effort for lower priority traffic at high throughput rates and high infrastructure utilization. This is accomplished by taking advantage of intelligent packet classification and steering, coupled with enhancements to the 5G User Plane software stack to selectively process high priority traffic. The solution utilizes the Intel® Xeon® Gold CPU SKUs and Intel® Ethernet Network Adapters with Dynamic

Device Personalization (DDP). No other investment in additional acceleration technologies were required.

In summary, we have demonstrated the following results:

With the 5G UPF system loaded up to 87% CPU utilization, we observed that:

- Normal packets resulted in 0.34ms (packet size: 175B) and 0.3ms (packet size: 550B) Round Trip Time (RTT) latency.
- Low Latency packets resulted in 0.07ms (packet size: 175B) and 0.09ms (packet size: 550B) RTT latency; up to **78%** reduction in latency.
- Normal packets resulted in ± 0.1 ms (packet size: 175B) and ± 0.087 ms (packet size: 550B) jitter.
- Low Latency packets resulted in consistent ± 0.014 ms jitter for both 175B and 550B; up to **88%** reduction in jitter.

This showcases how a standards-based 5G SA UPF enables a highly scalable, flexible, and reliable system for deployments in the core, at the central office or on premise for both B2C and B2B services.

Table of Contents

Executive Summary 1

1. Introduction 2

2. Architecture 4

 2.1 Typical Network Model 4

 2.2 Low Latency Model 4

 2.3 Packet Classification and Steering 5

 2.4 Utilizing DDP 6

3. Prototype Implementation 7

 3.1 Hardware Based Priority Aware Packet Steering 7

 3.2 Software Implementation .. 8

4. Test Setup 10

5. Measurements Results 12

 5.1 Results For Test Profile 1 .. 13

 5.2 Results For Test Profile 2 .. 14

 5.3 Results For Test Profile 3 .. 15

 5.4 Results For Test Profile 4 .. 17

Conclusion 18

Glossary 19

1. Introduction

The requirements for the 5G Standalone (SA) wireless network have been standardized by 3GPP. These standards define the Cloud Network Functions / Virtualized Network Functions (CNF / VNF) as microservices which can communicate with each other via APIs in a Service Based Architecture (SBA). A modular 5G core network can be achieved by designing with NFV and SBA concepts and allows for the easy addition and removal of the Network Functions (NFs). Utilizing the SBA, the Control Plane and User Plane functions are separated. The Control Plane could be effectively managed in the Core network or in the Central Office and the User Plane Function (UPF) can be distributed to be located geographically closer to end users to achieve low latency requirements.

Standards development organizations (SDOs) from GSMA, NGMN, and ITU have latency requirements as low as 1ms (ITU M.2083) in order to provide service such as Ultra-Reliable Low Latency Communications (URLLC). Various applications and use cases require URLLC service, such as augmented and virtual reality, telemedicine, intelligent transportation, and industrial automation. While the UPF is just one portion of overall end-to-end network latency and jitter, it is critical that the UPF behaves in a deterministic fashion under heavy load to meet performance requirements.

5G Systems consist of various network elements that contribute to overall end-to-end latency (e.g. RAN, Transport, Core, etc.). The System Architecture caters to a wide range of workloads that individually have varying end-to-end latency needs. Table 1 below illustrates packet delay budget (PDB) requirements in a 5G System. This defines an upper boundary that a packet may be delayed between the UE and the UPF that terminates the N6 interface towards the data network (ref: 3GPP TS 23.501). PDB is used across network elements to support the configuration of scheduling and link layer functions such as priority weights and HARQ target operating points in the gNB. For example, for Guaranteed Flow Bit Rate (GFBR) flows using the delay-critical resource type, packets that are delayed beyond the PDB bounds are considered as lost if the data burst is not exceeding Maximum Data Burst Volume (MDBV) within the period of PDB and the QoS Flow is not exceeding GFBR configuration.

Generally, PDB applicable to the radio interface is determined by subtracting the static value of the Core Network Packet Delay Budget (CN PDB).

Example Services	Overall Packet Delay Budget	Core Network Packet Delay Budget
Low Latency eMBB applications Augmented Reality	10 ms	2 ms
Discrete Automation	10 ms	1 ms
Discrete Automation – V2X messages (UE - RSU Platooning, Advanced Driving: Cooperative Lane Change with low LoA.)	10 ms	1 ms
Intelligent transport systems	30 ms	5 ms
Electricity Distribution - High Voltage, V2X messages (Remote Driving)	5 ms	2 ms
V2X messages (Advanced Driving: Collision Avoidance, Platooning with high LoA)	5 ms	2ms

Table 1.

One of the most stringent communication requirements is for timeliness and availability of a service where transmission occurs over a well-defined time interval with two main attributes of periodicity and determinism. Periodicity means the transmission is repeated over a given time interval, whereas determinism refers to the delay between the transmission of a message and its reception at the target. Typically, message transfers are deemed deterministic when they are bound by a certain threshold for latency, as well as bounds on its variation. Table 2 below enumerates Periodic Deterministic Communication characteristics for select services (re: 3GPP TS 22.104).

SK Telecom's 5G Core network infrastructure for UPF provides such URLLC services depicted in the table and is a key research focus for this paper. The ability for UPF to process packets efficiently when under maximum system utilization can allow operators the flexibility to operate different network topologies including near edge and far edge deployments.

Depending on B2C, B2B, and private network use cases, a well-designed and optimized UPF system could be extended to integrate various functions of the NSA and SA core for optimization. These functions could include:

- NSA Core SGW-U, and PGW-U
- DPI, NAT, and Firewall
- NSA/SA RAN CU-UP
- Edge computing enabling servers and systems

While this paper focuses on the low-latency aspects of the packet processing for the 5G SA core, the same solution provides the agility to process the above functions.

Typically, virtualized systems are overprovisioned and engineered so that the system utilization load does not exceed a certain threshold value to make sure that there are no packet drops, throughput degradation or downtime during service migrations.

To support such URLLC services, research was needed to increase the system utilization while maintaining a

threshold for latency and jitter, for various traffic profiles. The solution shows that it is possible to achieve improved latency and jitter performance in the UPF system under various traffic load scenarios. This was achieved by using run-to-completion (RTC) model on the user plane pipeline, a multi-queue software architecture and Intel's Dynamic Device Personalization (DDP) within the Intel® Ethernet Series 800 Network Adapter.

Example Services	End-To-End Latency: Maximum	Service Bit Rate: User Experienced Data Rate	Message Size [Byte]	Transfer Interval: Target Value
Motion control	< transfer interval value	–	50	500 µs
Motion control	< transfer interval value	–	40	1 ms
Wired-2-wireless 1 Gbit/s link replacement	< transfer interval value	250 Mbit/s		≤ 1 ms
Mobile robots	< transfer interval value	–	40 to 250	1 ms to 50 ms
Mobile control panels – remote control of e.g. assembly robots, milling machines	< transfer interval value	–	40 to 250	4 ms to 8 ms
Mobile Operation Panel: Motion control	< 1 ms	12-16 Mbit/s	10-100	1 ms
Process automation – closed loop control	< transfer interval value	–	20	≥ 10 ms
Robotic Aided Surgery	< 2 ms	2-16 Mbit/s	250 to 2000	1 ms
Robotic Aided Surgery	< 20 ms	2-16 Mbit/s	250 to 2000	1 ms
Robotic Aided Diagnosis	< 20 ms	2-16 Mbit/s	80	1 ms
Cooperative carrying – fragile work pieces; (ProSe communication)	< 0,5 * transfer interval	2,5 Mbit/s	250 500 with localization information	> 5 ms > 2,5 ms > 1,7 ms

Table 2.

2. Architecture

2.1 Typical Network Model

Figure 1 below illustrates a typical packet scheduling mechanism in many of today's networks. The network adapter determines the target worker core and ensures UE to core pinning. It then posts packets into a queue for a worker core to pull packets from and process them in sequential order while ensuring packet ordering for every packet of a given UE. This typical model has an unintended consequence in which higher priority packets (e.g. VoLTE, URLLC) might get backed up behind low priority packets in the queue and is known as Head Of-Line-Blocking (HOLB). This can result in potentially higher latency/jitter for such packet flows, including drops and packet loss under high CPU load and traffic conditions. Furthermore, low priority traffic processing might also need higher compute complexity capabilities like Deep Packet Inspection (DPI), or Application Detection that can add to the overall latency for packet processing of high priority traffic.

2.2 Low Latency Model

To mitigate the unintended consequences that might occur in the typical network processing model, a solution could be to separate out packets based on QoS flows (e.g. Guaranteed Bit Rate, Non-Guaranteed Bit Rate, Delay Critical Guaranteed Bit Rate), and then pin the processing of the packets belonging to such flows to dedicated worker cores. The implication of this approach could result in the network infrastructure not being optimally utilized and being over provisioned. Hence, the goal of this work is to improve packet latency/jitter characteristics of designated low latency traffic types under all CPU load levels and traffic conditions in the 5G system. This is implemented by:

1. Enhancements to packet parsing classification capabilities of the Intel® Ethernet 800 Series Network Adapter to be able to identify QFI values in GTP headers' PDU Session Container as well as DSCP codes in IP headers. These classifications are done by software configurable header field values via Dynamic Device Personalization (DDP) capabilities.

2. Steering of packets into one or more queue groups, with multiple queues in each queue group. An example would be steering high priority packets into one queue group, medium priority packets into a second queue group, and low priority packets into a third queue group.
3. Receive Side Scaling (RSS) based load distribution of packets within queue groups, such that queues within queue groups can be assigned to worker cores.
4. Enhancements to packet processing scheduling logic in software that runs on each of the worker cores, that polls packets from one or more queues based on its priority/weighted-priority and processes packets.

Each worker core keeps track of its utilization (i.e. system loading at low granularity) and implements scheduling to pull packets from queues based on loading. Under high loading, more weight is given to higher priority flows versus low priority flows; there by ensuring latency/jitter characteristic requirements are met.

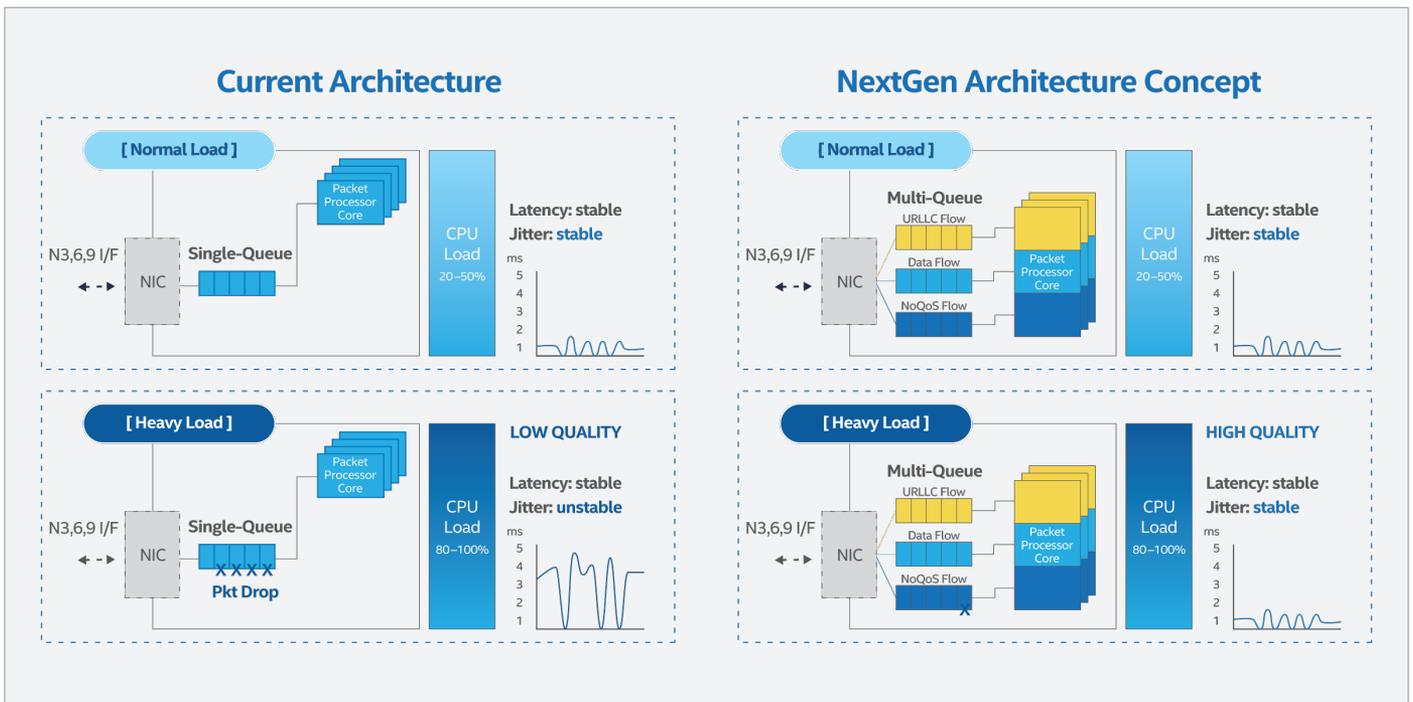


Figure 1. Typical and Desired UPF Latency Behavior

This implementation supports QoS characteristics for all low latency traffic independent of CPU load. The system maps Differentiated Services Code Points (DSCPs) for downlink traffic and QoS Flow Identifiers (QFIs) for uplink to ensure packet steering into the correct priority queue group. After that, the scheduling logic (software algorithm) ensures that the traffic processing is priority-aware when dealing with multiple traffic types in the same system.

2.3 Packet Classification and Steering

Intel® Ethernet 700 Series and Intel® Ethernet 800 Series, are relatively low cost, low power network interface cards that enable a pure software VNF/CNF model. Foundational NICs do not typically offer offloads for functions such as vSwitch acceleration, VXLAN TEP or inline IPsec. They do, however, offer advanced features in order to scale the VNF performance by optimizing packet steering in the server.

Dynamic Device Personalization (DDP) is a capability that was introduced with Intel® Ethernet 700 Series Network Adapters to load an additional package to enable classification and steering of additional specified packet types and performance of additional inline actions. DDP can be used to optimize packet processing performance for different network functions, native or running in a virtual environment. By applying a DDP profile to the network controller the following use cases can be addressed.

Extended support for protocols:

- 5G GTP support for 5G user plane.
- 5G SDAP/PDCP support for 5G NR user plane.
- 5G/4G PFCP (CP-UP separation) support.
- IP protocols as new flow types, for example L2TPv3, ESP/AH for IPsec.
- Legacy protocols: PPPoE, PPPoL2TPv2.

- New protocols/standards: eCPRI/ORAN, Radio over Ethernet (RoE).
- Extensibility for custom protocol parsing/classification.

In Figure 2 below, we can see how this type of sophisticated traffic control may be used to realize 5G core functions. DDP can effectively classify and steer traffic within the server based on control plane (N1/2, N4), user plane (N3, N6) or between UPF handover (N9) interfaces. For example, the foundational NIC can steer control plane protocols such as PFCP into the SMF or control plane part of UPF and can steer UE session either based on PDU session, flow, QoS class etc. on N3 and N6. Furthermore, DDP may be used to support extended header (EH) for 5G user plane traffic.

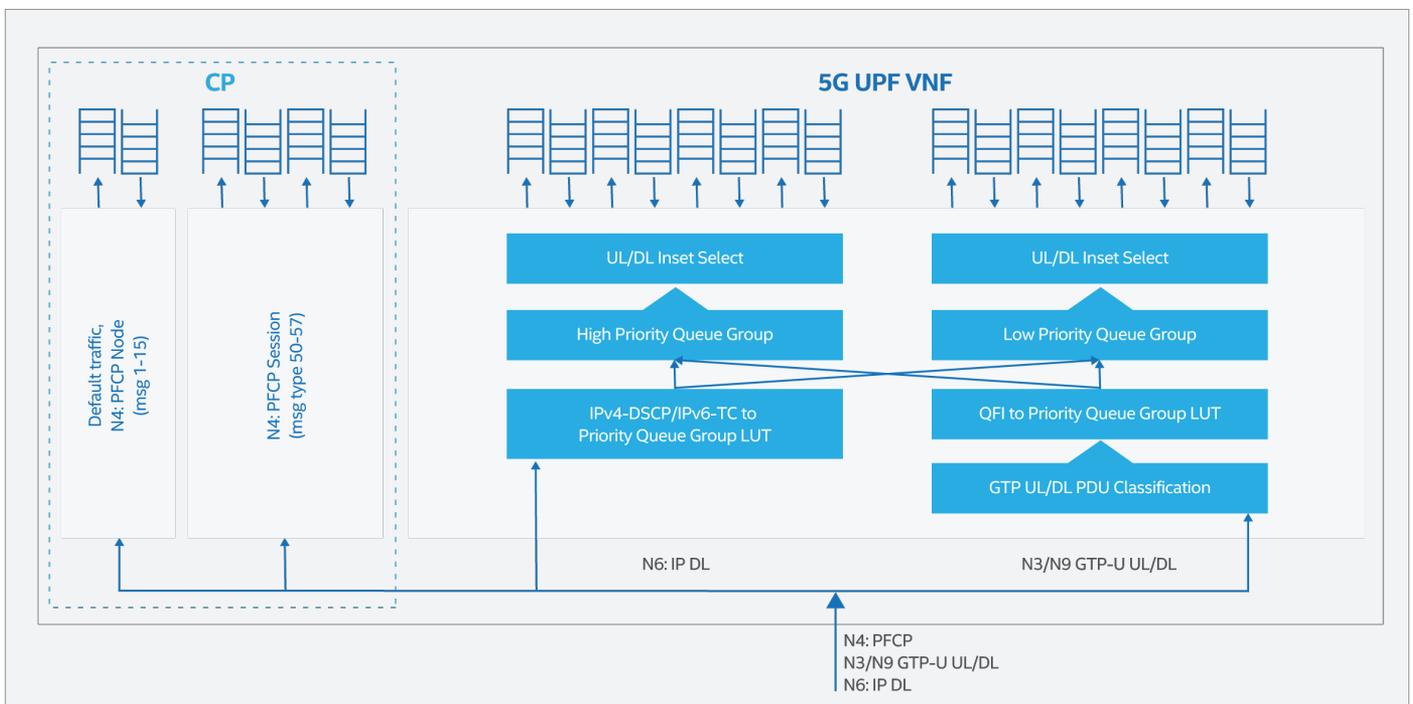


Figure 2. DDP Classification for 5G UPF

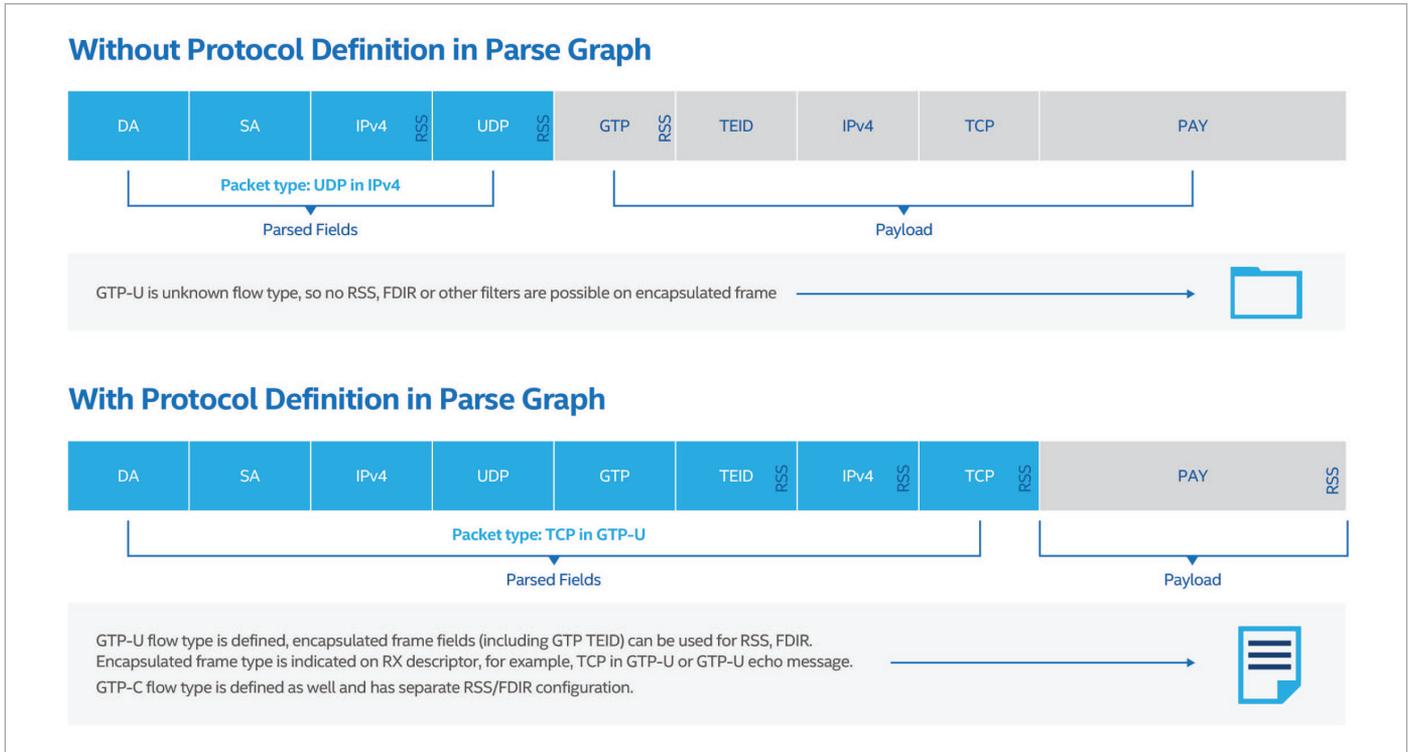


Figure 3. DDP Classification Example

2.4 Utilizing DDP

With no DDP profile, the NIC has a coarse traffic steering capability that steers packets to the appropriate load distribution function running in RX-Cores based on outer five tuples (as shown on upper part of Figure 3). Then, the load balancer core distributes UE packets to worker cores that executes 5GC UPF pipeline functions. The load distribution function on RX-Cores adds extra latency and jitter due to the increased number of RX-Queue levels, and software function usage. This creates a bottleneck for performance and introduces even more latency and jitter at high traffic loads.

The lower part of the Figure 3 explains how DDP may be used to parse deeper into the packet and steer packets directly to worker cores based on inner header fields such as GTP-TEID and/or source/destination IP address. This enables the NIC to distribute UE's

packets in a very deterministic way to the appropriate function for 5GC UPF pipeline processing without using specific cores for load distribution. This frees up valuable resources (cores) and decreases RX queue levels, which will make the system more deterministic with low latency and low jitter at high traffic and CPU load.

The Intel® Ethernet 800 Series Intel® Network Adapter carries over the Intel® Ethernet 700 Series Network Adapter features described in Figure 3 and adds more options. The Intel® Ethernet 800 Series Network Adapter focuses on meeting customer requirements and targets for connectivity and latency along with providing a 100Gbps connection. This results in reducing the variability in application response time, improving predictability, and increasing throughput.

Intel is accomplishing this through two technologies - Application Device

Queues (ADQ) and Dynamic Device Personalization (DDP). Intel continues to develop more DDP profiles based on customers demand and these are available for download here: <https://www.intel.com/content/www/us/en/search.html?ws=text#q=DDP&t=Downloads&layout=table>.

DDP may also be used to extend functionality into other areas based on steering and applying preferential QoS to other IP protocols. Typically, Operators want to ensure that control plane, management plane and OAM protocols get prioritized in the infrastructure. Examples may include, but are not limited to:

- BGP, OSPF, or ISIS for routing control planes and fast convergence.
- OAM protocols like Bidirectional Forwarding Detection (BFD), Multi-Protocol Label Switching Operations and Maintenance (MPLS-OAM), Two-Way Active Measurement Protocol

(TWAMP), Virtual Router Redundancy Protocol (VRRP).

- Service Based Architecture (SBA) interfaces for consistently high control plane processing in 5G Core, or related critical OSS/BSS application traffic.
- IEEE1588 PTP (Precision Time Protocol) for synchronous operations.

algorithms for packet placement into receive queues at ingress.

The application layer may be bare metal, VM-based or containers and is decoupled from the multi-queue architecture. Figure 5 illustrates just two priorities, but the architecture is extendable to support a higher number of priority queue groups.

Figure 5 below describes logical processing of GTP-U and IP based traffic in the system. The NIC reads the DSCP value of the IP packet or QFI value in the PDU session container extension header of GTP-U to identify the priority queue group and steers the traffic to a specific queue within that priority group based on the inner source IP address for the N3 and N9 uplink or inner destination IP address for N9 or N6 downlink. This is effectively the UE IP address. By consistently steering the UE traffic to a fixed queue the system ensures that all user plane processing for that UE can be handled by the same CPU core using run-to-completion (RTC) methods.

3. Prototype Implementation

3.1 Hardware-Based Priority Aware Packet Steering

Figure 5 shows the multi-queue architecture implemented on foundational NICs from Intel. Traffic is prioritized and steered via configurable policies into one of two queue groups which are served by RSS

PDU Type for uplink and downlink selection and QFI values are parsed from the PDU session information in "PDU session container" extension headers of GTP-U packets as per Figure 4 below. The system maps the 64 possible QFI values into two or more queue groups.

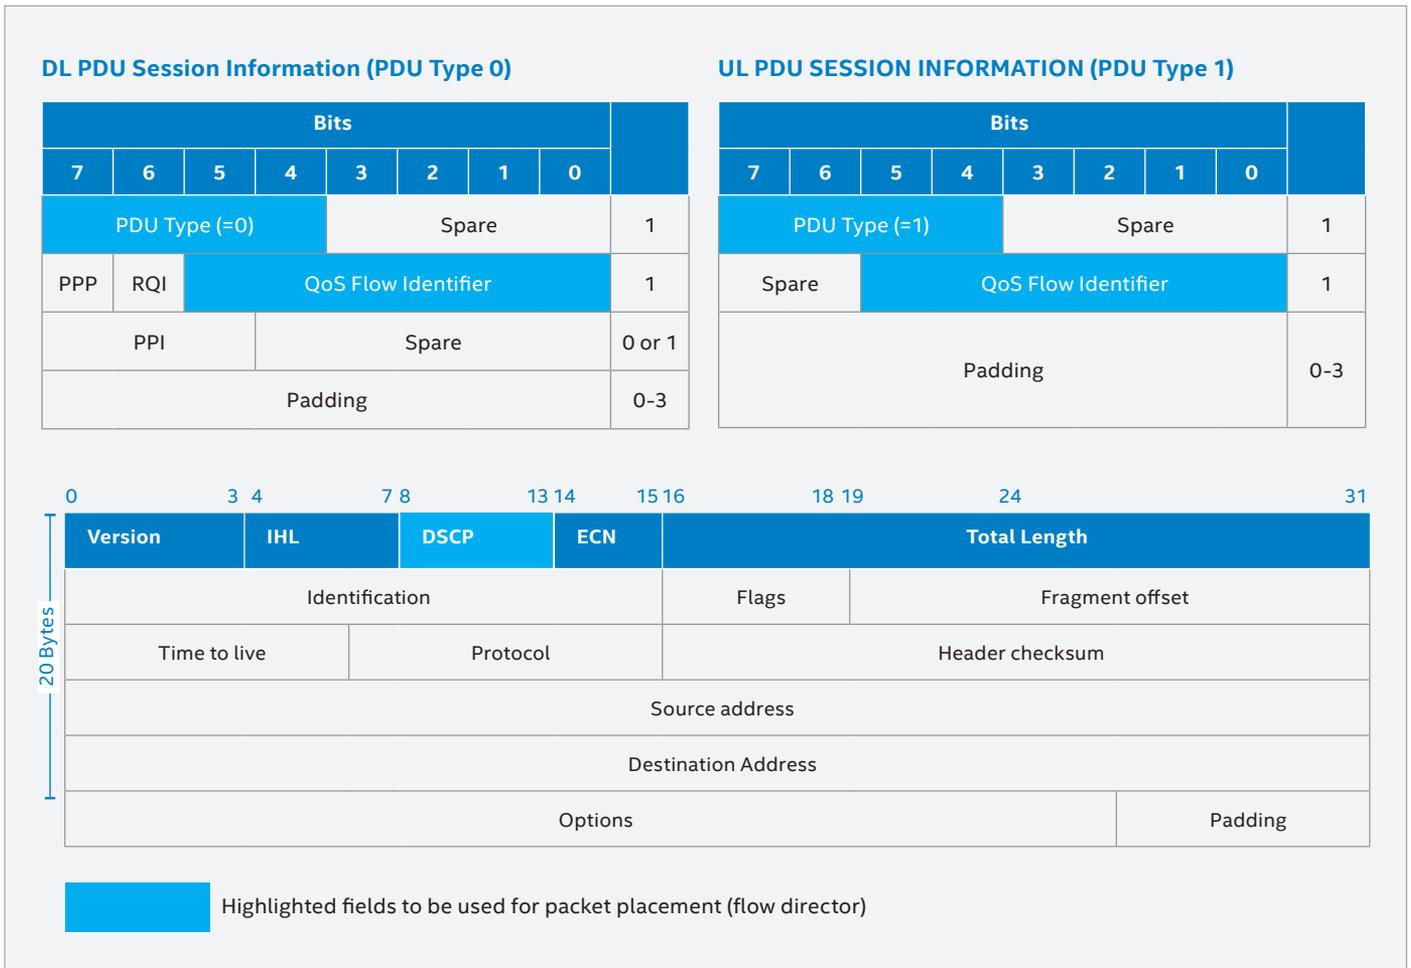


Figure 4. DSCP (RFC 2474) and PDU Session Parsing (TS 38.415)

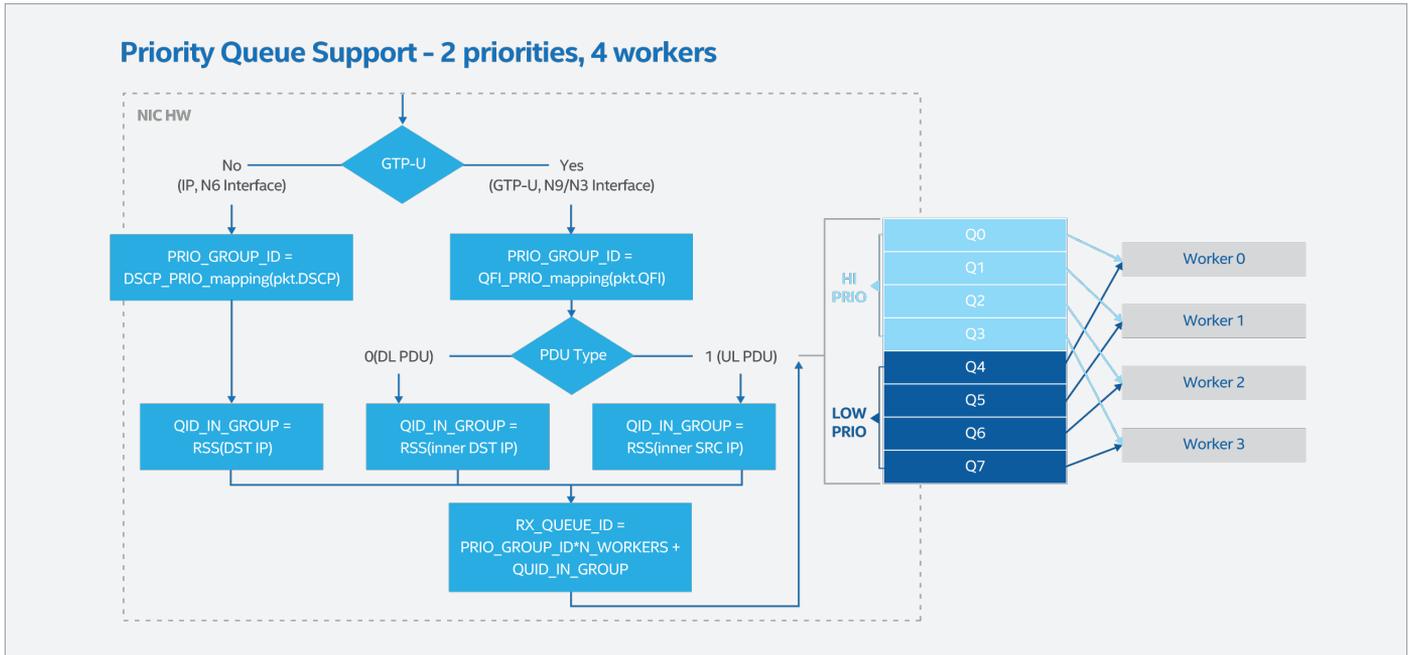


Figure 5. N3, N9 & N6 Queue Selection example

This capability ensures that:

1. UEs stay on the same cores irrespective of mobility events or other session events.
2. UE context is in the local cache and the full pipeline executes on that core.
3. The system becomes very linear and deterministic in terms of performance.
4. The system becomes very easy to dimension.

Similarly, on the N6, downlink traffic is steered and queued based on a combination of destination IP and DSCP into either high or low priority RSS queues serving consistent CPU resources. If desired, the system can implement even further granularity by binding a UE flow type or application to a core.

3.2 Software Implementation

The Intel® Ethernet 800 Series provides the capability to steer packets of different priority into specific queue groups as described in Figure 5. The software receiving and processing

packets must be aware of the receive queue priority mechanism to enable the efficient handling of high priority packets.

The User Plane Function (UPF) application used in the context of this work is based on the FD.io Vector Packet Processor (VPP) framework. This framework utilizes Data Plane Development Kit (DPDK) functionality to fetch received packets from the NIC queues and deliver them for further processing. The DPDK plugin is a part of the VPP project that exposes packet receive functionality over the dpdk-input node. Default implementation of the dpdk-input node enables handling of multiple RX queues in the context of a single worker thread. However, the DPDK plugin has no notion of RX queue priority and handles packets from all RX queues with the same priority.

The packet processing logic of the VPP framework can be simplified as a continuous cycle where:

- Input node generates a vector of packets (e.g. dpdk-input node fetches packets from NIC receive queues).

- VPP framework dispatches nodes connected in the form of directed graph to process vector where the vector itself can be subdivided during processing.

Packets to the graph nodes of VPP framework are delivered in the form of frames where the maximum number of packets in a single frame is limited by VLIB_FRAME_SIZE compile time parameter. The average time spent on one packet processing in VPP graph node decreases when the size of the frame (number of packets in the frame) increases. Thus, the VPP framework implements self-adjusting mechanisms for packet handling. For low packet rate and low CPU utilization level, the average frame size becomes lower. High packet rates result in larger frames and more efficient CPU usage yielding improved performance.

In this model, the packet processing time also depends on the size of the vector generated by the input node. It takes more time to process a frame of N packets in the VPP stack than to process a frame with a single packet. If a single high priority

packet is bundled with many low priority eMBB packets in one frame, the average packet processing time for the high priority packet will be the same as for the eMBB packet.

Additional functionality is introduced in the DPDK plugin of the VPP framework in order to enable priority-aware receive queue handling and latency control. This functionality is enabled and configured over two basic parameters:

- Receive queue priority (RQP) parameter is a numeric value assigned to a receive queue where the value 0 identifies the highest priority. Each receive queue has one priority value assigned and multiple queues with the same priority value are allowed.
- Maximum read burst size (MAX_RBS) parameter defines maximum number of packets that could be fetched from receive queues in one call to the dpdk-input node function. Unlike VLIB_FRAME_SIZE, this parameter could be changed at runtime. The MAX_RBS parameter value cannot exceed VLIB_FRAME_SIZE.

In every call to the dpdk-input node function, it executes an algorithm simplified to these steps:

1. Set current priority P=highest priority (e.g. P=0).
2. Try to fetch up to MAX_RBS packets from the queue(s) with the current priority P.
3. If one or more (up to MAX_RBS) packets received from queues with priority P -return control to the VPP framework to allow further processing of received packets.
4. If there are no packet(s) from queues with priority P received and more queues available switch to next priority P=P+1 and goto (2).
5. Return control to the VPP framework with actual number of fetched packets (zero or more).

The algorithm in Figure 6 implements strict priority logic where packets from higher priority queues are always fetched first and next priority queues are not served until all packets from high priority queues are in the processing stage. This ensures that high-priority traffic is served even if low priority traffic has to be dropped. The MAX_RBS parameter limits the maximum number of lower priority packets that are delivered to the VPP framework in a single call to the input node and as a result, it limits maximum interval between fetching high-priority packets.

Latency results in this paper are based on strict priority receive queue logic described above.

Other improvements can be added to the current implementation, such as:

- Enabling mix of packets of different priorities in one frame, e.g. fetching up to MAX_RBS packets from receive queues of multiple priorities in one dpdk-input node call.
- Using flexible MAX_RBS parameter (e.g. which value depends on CPU load or packet rate) or configuring MAX_RBS on per priority queue basis.
- Using weighted receive queue priority mechanism.
- Dedicating cores to serve high-priority queues.
- Ensuring high priority packets also given priority in transmit path.

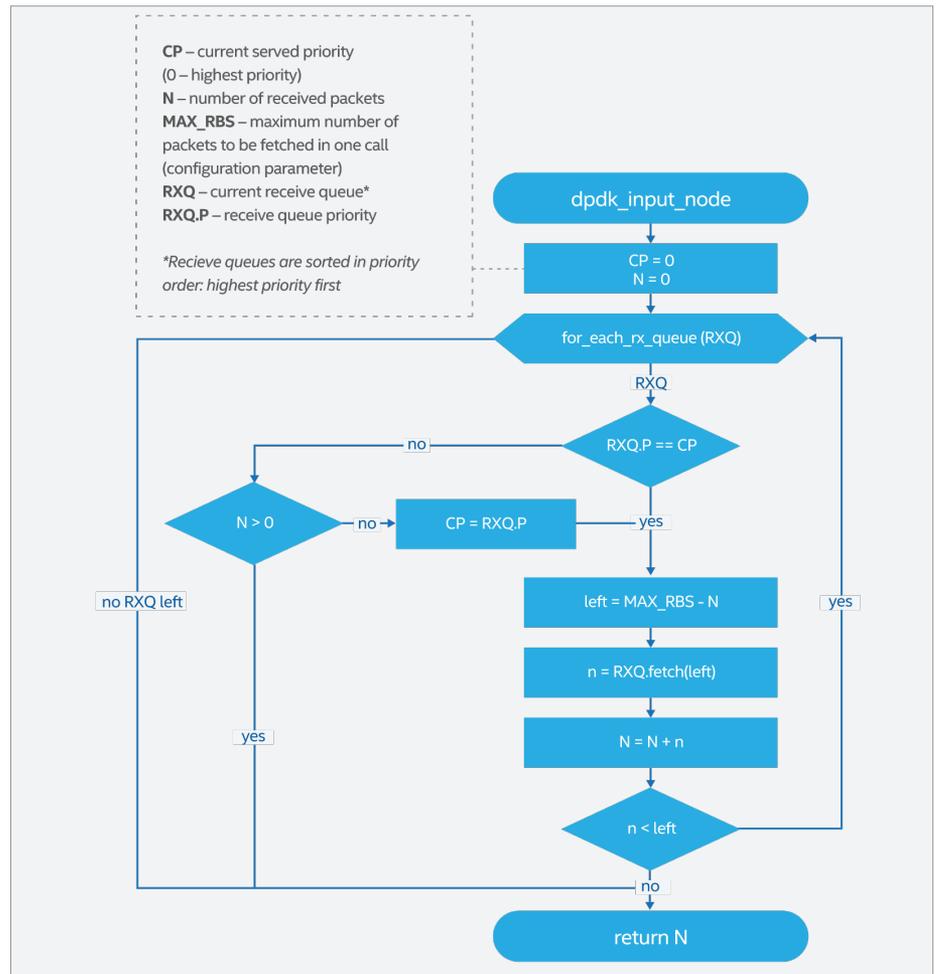


Figure 6. Input Node Packet Processing Flow Chart

4. Test Setup

The development and testing environment for this solution is shown in Figure 7 and consist of the following:

- 5G Core Network User Plane Function (UPF) reference stack from ASTRI (Applied Science and Telecoms Research Institute - www.astri.org). The stack is used to demonstrate a typical pipeline performance (via compiled binaries) which we have additionally patched and configured for our Low-Latency research.
- 5G Core Network Control Plane components include Access and Mobility Management Function (AMF) and Session Management Function (SMF) for testing the user plane as sessions are established, deactivated, and moved due to mobility events.
- A combination of standard test equipment (Spirent Landslide) and in-house tools in order to characterize

the system. The UE and RAN elements are simulated by Landslide test equipment. Please note that a full 5GCN characterization is not reported here as the effort focuses on the user plane performance for specific call models and use cases.

In our testing we stress the performance of the UPF as this is the main forwarding element in the core architecture – and thus the main component in jitter and latency for user plane traffic. The effort concentrates on technological advances that aid the implementation of low-latency on the Intel architecture.

The 5GCN test harness is shown in Figure 7. The device under test (DUT) hosts the UPF top middle server. AMF and SMF are hosted in the top right server.

The Spirent Landslide systems are connected into the ToR switch and generate the UL traffic into the UPF.

On the N6 side, the traffic terminates on the uplink (UL) sink server and downlink (DL) replies are generated with in-house software tools. Downlink traffic is sent through the UPF and terminated in the DL sink server on the bottom right. Using this harness, we can easily change traffic profile configurations and measure throughput latency and jitter in real time.

Details of the 5G UPF configuration are shown in Table 3 below.

We deployed UPF functionality on 16 physical cores from CPU2 on the Intel S2600WFT server shown in Figure 8 and 9 below. This is an internal Intel reference board with 2x16PCIe Gen3 lanes and 96GB of DRAM (6x16GB) per socket. One Intel® Ethernet 800 Series Network Adapter (1x100G) feeds CPU2. For optimal performance the PCIe x16 slot on the same NUMA that is running the UPF is used.

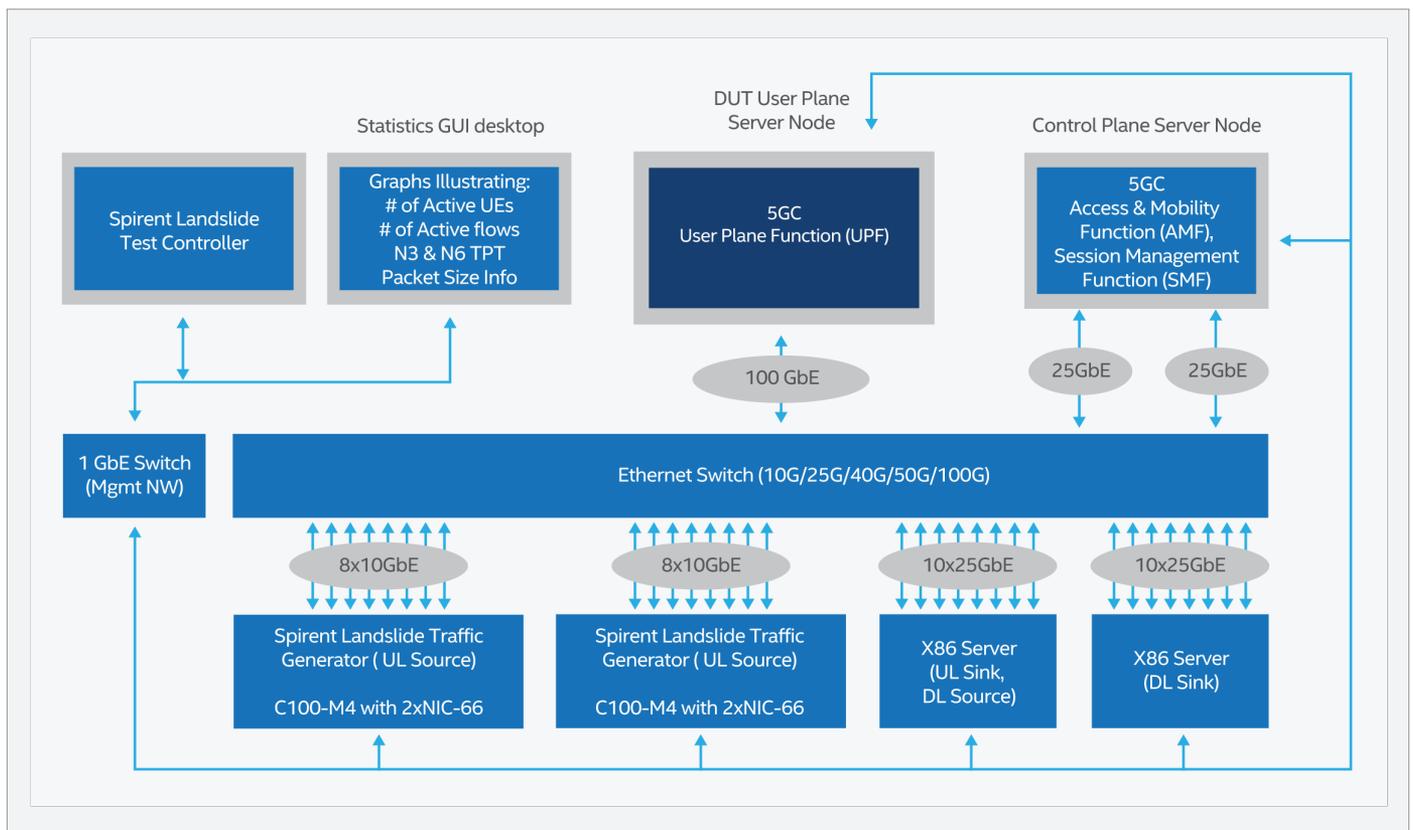


Figure 7. 5G Test Harness

Category	Description	
Processor	Product	Intel® Xeon® Gold 6252N Processor
	Frequency	2.3 to 3.5GHz
	Cores per Processor	24 Cores/48 Hyper threads
Memory	DIMM Slots per Processor	6 Channels per Processor
	Capacity	192GB DRAM, 1.5TB DCPMM (FW version 01.02.00.5346)
	Memory Speed	2666 MHz, DDR4
Network	NIC	Intel Foundational NIC: E810_CQDA2 100GbE QSFP CVL-FW: FW-1.1.16.40 NVM-1.02 0x80002b68 DDP: ice-1.3.10.0.pkg Driver: ice-0.12.34
	Number of Ports	1 port from E810_CQDA2 100GbE QSFP NIC
Server	Vendor	Intel S2600WFT
Host OS	Vendor/Version	Ubuntu 18.04.2 LTS 4.15.0-20-generic x86_64
BIOS	Vendor/Version	Intel Corporation SE5C620.86B.02.01.0010.010620200716 Release Date: 01/06/2020
5G UPF	Vendor/Version	ASTRI rel. 19.07-rc2 with code changes from Intel

Table 3.

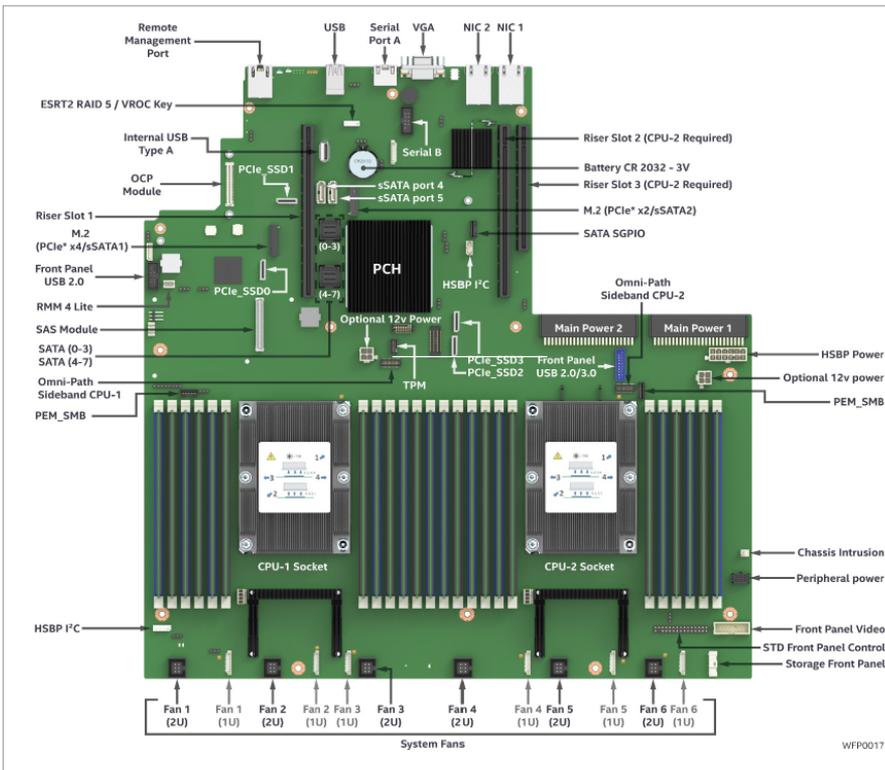


Figure 8. Intel S2600WFT Layout



Figure 9. UPF Server (2x6252N with Intel® Network Adapters)

5. Measurements Results

We tested four scenarios and measured latency and jitter with the Low-Latency implementation. Each scenario has a different traffic profile, and measurements were done under different CPU loads. All measurements were based on the Intel® Xeon® 6252N (24 core, 2.3Ghz) CPU based platform. All latencies reported included the delay in traffic emulator transmit and receive stacks and ethernet switch components. All tests involved emulating 50,000 UEs with up to 100Gbits of aggregated throughput on a single CPU. Each profile tested different packet size values and priority ratios as explained in Table 4 to the right.

Jitter measurements in this document are based on the algorithm described in RFC3550, Appendix A.8. This is implemented on a per-flow basis where an estimate of the statistical variance of the packet interarrival times is made. Mean deviation of the relative transit time of packets is measured in microseconds. Areas for potential future optimization of VPP performance were noted with respect to deterministic performance across different traffic profiles.

Profile	Traffic Mix	Low Latency Traffic Characteristic
1	100Gbps, ~20MPPS total UPF traffic ~1.8 MPPS for low latency traffic, 2.5 Gbps ~18.9 MPPS for low priority traffic, 97.5 Gbps	175 Byte Low-Latency packet size, ~ 1:1 UL/DL Utilizing 2.5% of aggregated TPT
2	100Gbps, ~22.7MPPS total traffic thru the UPF ~2.2 MPPS for low latency traffic, 10 Gbps ~20 MPPS for low priority traffic, 90 Gbps	550 Byte Low-Latency packet size, 1:1 UL/DL, Utilizing 10% of aggregated TPT
3	1:3 UL to DL traffic profile 100Gbps, ~21MPPS total UPF traffic 3.9 MPPS for low latency traffic, 44 Gbps 17 MPPS for low priority traffic, 56 Gbps	1400 Byte Low-Latency packet size, 1:9 UL/DL Utilizing 44% of aggregated TPT
4	1:3 UL to DL traffic profile 100Gbps, ~21MPPS total UPF traffic 705 KPPS for low latency traffic, 4.4 Gbps 20 MPPS for low priority traffic, 95.6 Gbps	780 Byte Low-Latency packet size, 1:9 UL/DL Utilizing 4.4% of aggregated TPT

Table 4.

5.1 Results for Test Profile – 1

This test profile has approximately 2.5% of high priority traffic (175B packets), while the rest of the traffic (645B packets) is regular mobile broadband (i.e., eMBB). The overall packet size, packet rate, and per UE packet rate are described in Table 5 below.

With this traffic profile, the peak performance at zero packet loss was

measured, along with the latency and jitter characteristics. Similar measurements were also taken for additional CPU loading points by modulating the packet rate/packet size mix to check for latency and jitter at each of these sample points. Results are shown in Figure 10 below. Consistent latency measurements of ~31-38 microseconds were noted with ~13 microseconds jitter for high priority.

It was noted that the latency and jitter characteristics for high priority traffic stay consistent across the load line of various measurement points of CPU loading, demonstrating a 78% reduction in latency for high priority traffic over low priority traffic, while the jitter reductions are up to 88%.

UL TPT Mbps (N6)	N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
1,230	175	878,571	17.57
6,712	645	1,300,775	26.02
7,942		2,179,347	43.59
DL TPT Mbps (N6)	N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
1,312	175	937,143	18.74
90,746	645	17,586,434	351.73
92,058		18,523,577	370.47
TOTAL TPT Mbps (N6)	DL to UL TPT Ratio	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
100,000	12	20,702,924	414

Table 5. Profile 1

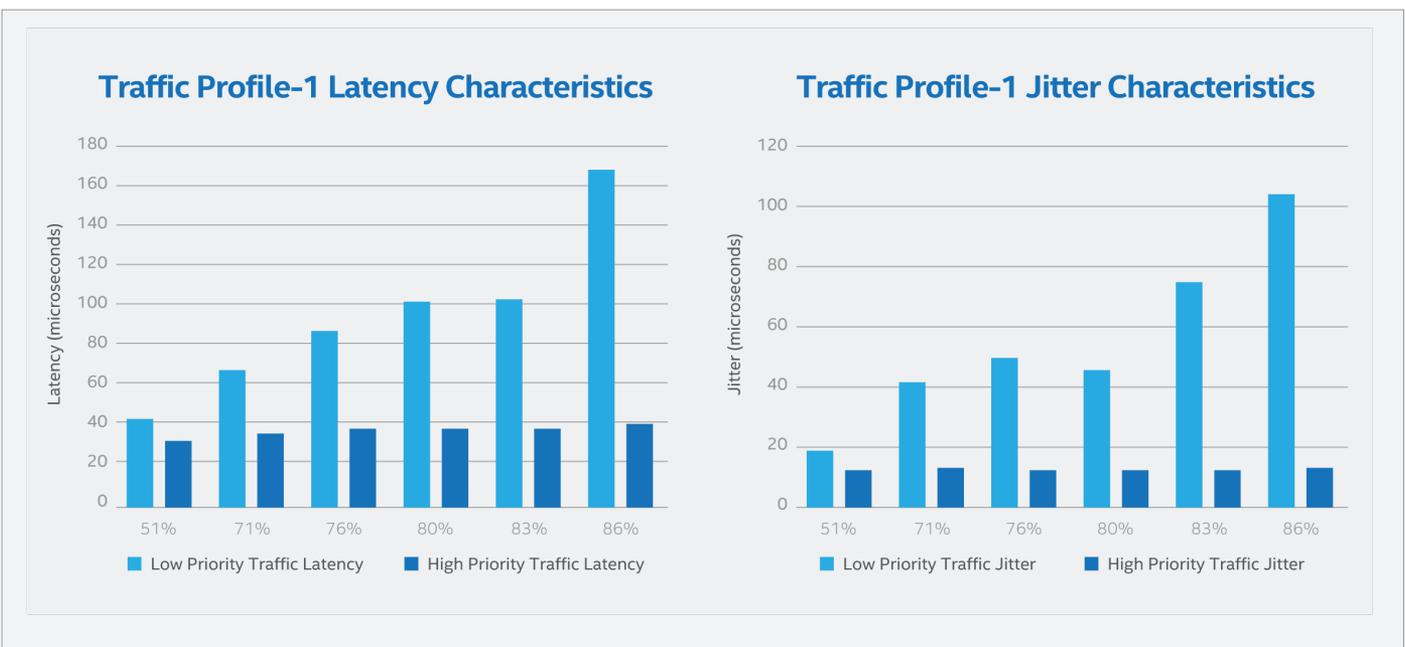


Figure 10. Profile 1

5.2 Results for Test Profile – 2

This test profile has approximately 10% of high priority traffic (550B packets), while the rest of the traffic is regular mobile broadband (i.e., eMBB). The overall packet size, packet rate, and per UE packet rate are described in Table 6 below.

With this traffic profile, the peak performance at zero packet loss was

measured, and along with it, the latency and jitter characteristics were noted. Similar measurements were also taken for additional CPU loading points by modulating the packet rate/packet size mix to check for latency and jitter at each of these sample points with results shown in Figure 11 below. Consistent latency of ~32-45 microseconds latency were measured

with ~12 microseconds jitter for high priority traffic. It is notable that the latency and jitter characteristics for high priority traffic stay consistent across the load line of various measurement points of CPU loading, demonstrating a 69% reduction in latency for high priority traffic over low priority traffic, while the jitter reductions are up to 84%.

UL TPT Mbps (N6)	N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
5,000	550	1,136,364	22.73
5,000	550	1,136,364	22.73
10,000		2,272,727	45.45
DL TPT Mbps (N6)	N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
5,000	550	1,136,364	22.73
85,000	550	19,318,182	386.36
90,000		20,454,545	409.09
TOTAL TPT Mbps (N6)	DL to UL TPT Ratio	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
100,000	9	22,727,273	445

Table 6. Profile 2

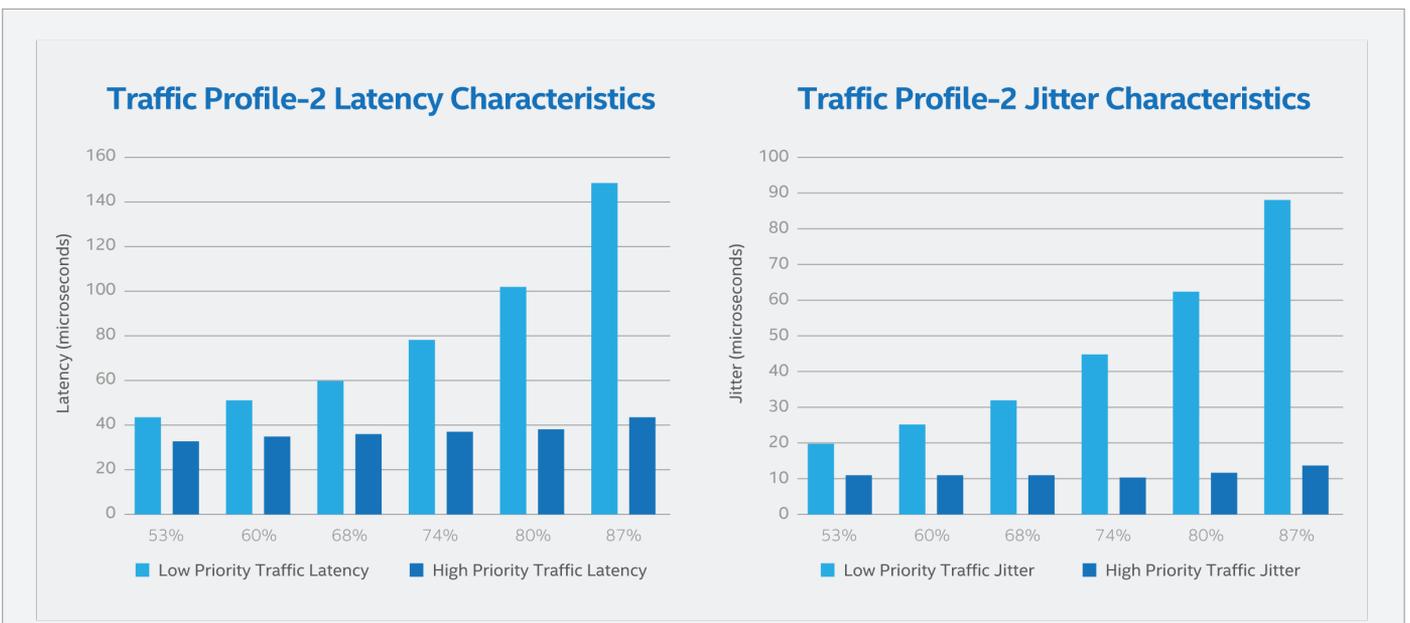


Figure 11. Profile 2

5.3 Results for Test Profile – 3

This test profile has approximately 44% of high priority traffic (1400B packets), while the rest of the traffic is regular mobile broadband (i.e., eMBB). The overall packet size, packet rate, and per UE packet rate are described in Table 7 below.

With this traffic profile below, the peak performance at zero packet loss

was measured, and the latency and jitter characteristics were noted. Similar measurements were also taken for additional CPU loading points by modulating the packet rate/packet size mix to check for latency and jitter at each of these sample points with results shown in Figure 12 on the following page. Consistent latency measurements of ~32-40 microseconds was noticed with ~12 microseconds jitter for high

priority traffic. It was noted that the latency and jitter characteristics for high priority traffic stay consistent across the load line of various measurement points of CPU loading, demonstrating a 61% reduction in latency for high priority traffic over low priority traffic, while the jitter reductions are up to 69%.

UL TPT Mbps (N6)	N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
481	64	940,370	18.81
884	128	863,534	17.27
3,060	256	1,493,927	29.88
3,671	512	896,240	17.92
420	780	67,308	1.35
12,353	1,024	1,507,935	30.16
4,201	1,400	375,087	7.50
25,070		6,144,399	122.89
DL TPT Mbps (N6)	N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
69	64	134,817	2.70
4,424	128	4,319,853	86.40
6,856	256	3,347,418	66.95
2,660	512	649,414	12.99
3,976	780	637,179	12.74
17,486	1,024	2,134,521	42.69
39,760	1,400	3,550,000	71.00
75,230		14,773,202	295.46
TOTAL TPT Mbps (N6)	DL to UL TPT Ratio	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
100,300	3.00	20,917,601	418

Table 7. Profile 3

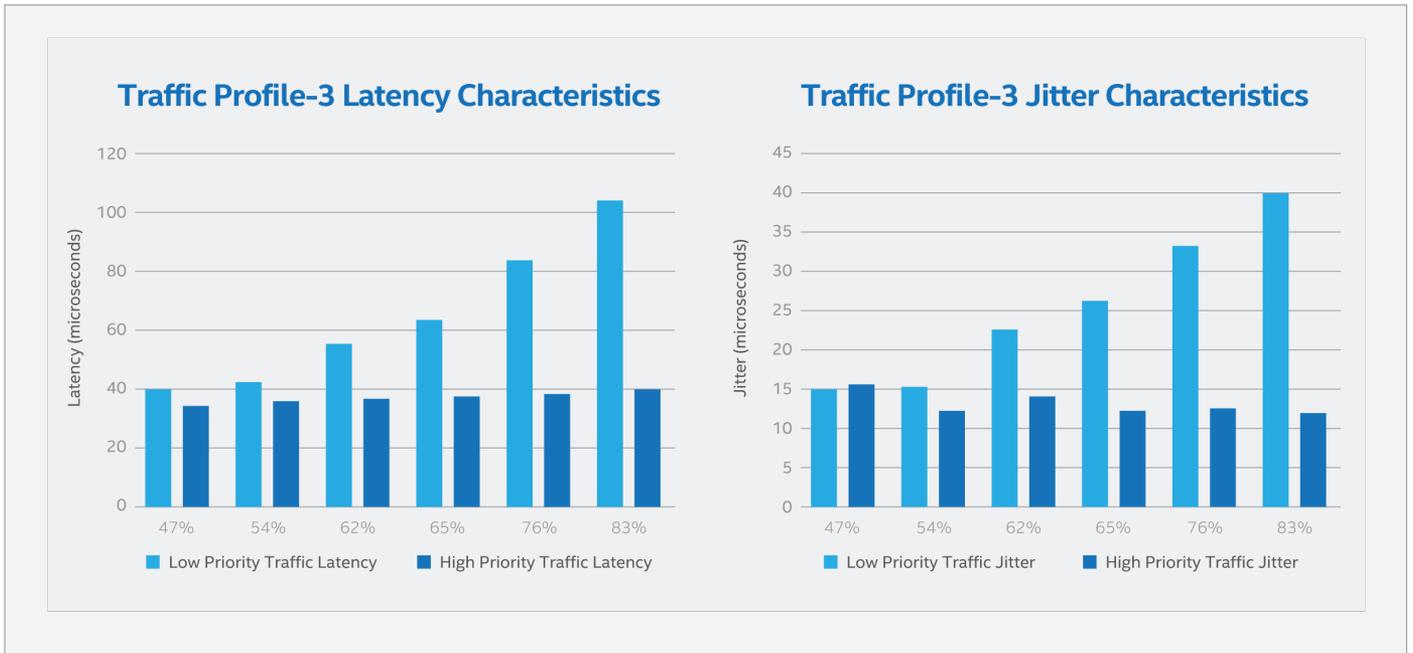


Figure 12. Profile 3

5.4 Results for Test Profile – 4

This test profile has approximately 4.4% of high priority traffic (780B packets), while the rest of the traffic is regular mobile broadband (i.e., eMBB). The overall packet size, packet rate, and per UE packet rate are described in Table 8 below.

With the traffic profiles below, the peak performance at zero packet loss was measured and the latency and jitter characteristics were noted. Similar measurements were also taken for additional CPU loading

points by modulating the packet rate/ packet size mix to check for latency and jitter at each of these sample points. Results are shown In Figure 13 on the following page. Consistent measurements of latency of ~32-38 microseconds was measured with ~14 microseconds jitter for high priority traffic. It was noted that the latency and jitter characteristics for high priority traffic stay consistent across the load line of various measurement points of CPU loading, demonstrating a 59% reduction in latency for high priority traffic over low priority

while the jitter reductions are up to 62%. It was observed that the UL to DL TPT ratio in this test case is significantly higher than first 2 traffic profiles and does result in slightly lower latency and jitter reductions. This has been attributed to traffic generators used in the setup that generate 'n' times the number of packets in the DL for every packet received in the UL direction (i.e. UL to DL packet ratio). In an actual live network, the actual reductions are expected to be similar to the first two traffic profiles.

UL TPT Mbps (N6)		N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
481		64	940,370	18.81
884		128	863,534	17.27
3,060		256	1,493,927	29.88
3,671		512	896,240	17.92
420		780	67,308	1.35
12,353		1,024	1,507,935	30.16
4,201		1,400	375,087	7.50
25,070			6,144,399	122.89
DL TPT Mbps (N6)		N6 PKT SIZE	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
69		64	134,817	2.70
4,424		128	4,319,853	86.40
6,856		256	3,347,418	66.95
2,660		512	649,414	12.99
3,976		780	637,179	12.74
17,486		1,024	2,134,521	42.69
39,760		1,400	3,550,000	71.00
75,230			14,773,202	295.46
TOTAL TPT Mbps (N6)		DL to UL TPT Ratio	TOTAL NO PKTS (N6)	PACKETS/SEC (PER UE)
100,300		3.00	20,917,601	418

Table 8. Profile 4

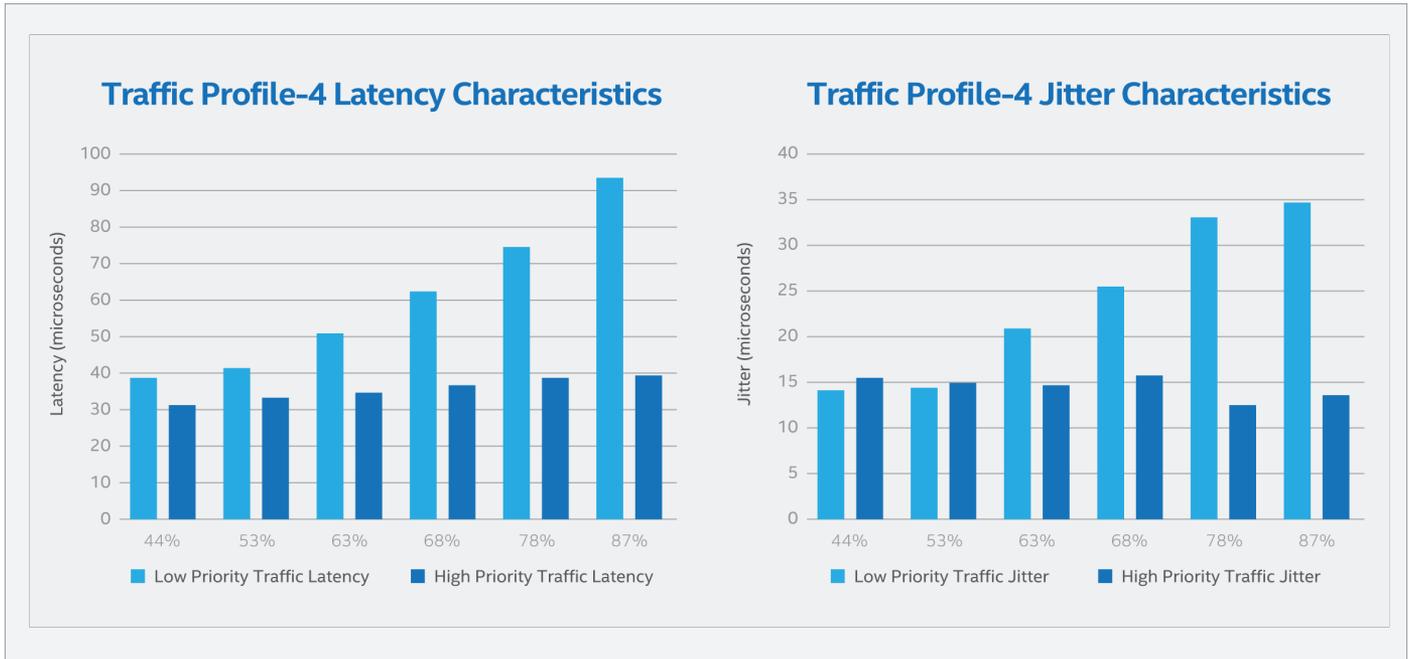


Figure 13. Profile 4

Conclusion

The 5G SA UPF collaboration by Intel and SK Telecom demonstrates how a standards-based solution can be used to achieve low latency and jitter and deliver a scalable and flexible system for network deployments.

The ability to efficiently service eMBB, MTC, and URLLC traffic types on the same NFV infrastructure with high levels of throughput, utilization, and determinism is a key enabler of the deployment for virtualized and containerized packet core systems for 5G and beyond.

Using intelligent classification, steering, and processing of traffic, this solution, based on the Intel® Xeon® processor 6252N processor and Intel® Ethernet 800 Series Network Adapters, demonstrates low latency packet processing in UPF, with significant reduction in latency and jitter of the 5G user plane. This is accomplished while still running lower priority traffic at high rates and infrastructure utilization. Based on the test profiles executed, we demonstrated up to 78% reduction in

latency and 88% reduction in jitter. Hardware-based packet steering and software-based prioritization were used to achieve results.

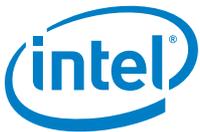
These capabilities show how it is possible for MNOs to utilize Intel architecture and software optimizations to meet customer demand and generate new revenue streams for latency sensitive 5G applications and services. Use cases that can benefit include factory automation, AI-enabled vision processing, and video analytics. This flexible architecture with associated capabilities can also reduce CapEx as different services with different latency or jitter requirements can own different slices of the same architecture without compromising each other.

Looking forward, Intel and SK Telecom will continue to collaborate on the 5G core with new CPU architectures, Intel NIC technologies, and software optimization techniques to demonstrate further performance improvements for 5G SA UPF.

Glossary

Term	Description
AMF	Access and Mobility Management Function
CAGR	Compound Annual Growth Rate
COTS	Commercial Off-The-Shelf
CPU	Central Processing Unit
CoSP	Communication Service Provider
CUPS	Control and User Plane Separation
DDP	Dynamic Device Personalization
DNN	Data Network Name
DPDK	Data Plane Development Kit
DPI	Deep Packet Inspection
EPC	Evolved Packet Core
FWA	Fixed Wireless Access
HQoS	Hierarchical Quality of Service
NEF	Network Exposure Function
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
NIC	Network Interface Controller
RAN	Radio Access Network
RSS	Receive Side Scaling
RTC	Run To Completion

SMF	Session Management Function
SP	Service Providers
SR-IOV	Single Root I/O Virtualization
TEM	Telecom Equipment Manufacturers
URLLC	Ultra Reliable Low Latency Communication
UPF	User Plane Function
VM	Virtual Machine
VNF	Virtual Network Function



Copyright © 2020 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries. Other names and brands may be claimed as the property of others.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks.

Performance results are based on testing as of March 2020 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at www.intel.com.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804.