



Accelerating the Virtualized User Plane for 5G Core Network Readiness

SK Telecom and Intel worked together to demonstrate a virtualized user plane function for the Mobile Packet Core that offers more than 200 Gbps throughput on a single two-socket server. Discover how packets were allocated to worker cores to optimize performance and reduce latency

This solution brief describes how to solve business challenges through investment in innovative technologies.

If you are responsible for...

- **Business strategy:**

You will better understand how the virtualized user plane will enable you to successfully meet your business outcomes.

- **Technology decisions:**

You will learn how the virtualized user plane works to deliver IT and business value.

Executive Summary

With the huge demands that 5G will put on the network, operators will be challenged to satisfy unprecedented performance requirements, without sacrificing agility or flexibility. Network functions can be run in software on standard hardware, but optimizing them for performance may involve introducing hardware dependencies that undermine the flexibility and agility goals of Network Functions Virtualization (NFV).

Intel and SK Telecom propose a highly performant solution for the 5G Core Network user plane function (UPF), based on standard hardware with hardware-assisted network interface cards. It is able to deliver throughput measured more than 200 Gbps¹ on a single two-socket server, powered by the Intel® Xeon® Platinum 8180 processor.

This paper outlines the approach used for routing packets, which ensures they are processed in order and on the same worker core, thereby maximizing the use of the cache and cutting latency.



DongJin Lee, JongHan Park
SK Telecom

**Chetan Hiremath, John Mangan,
Michael Lynch**
Intel Corporation

Solution Benefits

- **Agility and flexibility.** Basing the user plane on commercial off-the-shelf hardware without hardware dependencies means that its functions can be hosted at the edge or at the core, or anywhere in a cloud-like infrastructure. This enables a wider range of applications to be supported, such as augmented reality and virtual reality, which is likely to require user plane functions at the edge. Resources can be scaled up or down as required to meet performance needs
- **Intelligence.** Adding intelligence to the network interface card enables next generation networks to be more easily supported, including support for virtual switching, overlay termination for tenant separation, and more granular load balancing.





Figure 1. Photos of the COTS Rackmount Server

Achieving the Performance and Agility 5G Needs

As operators prepare for 5G, two requirements are of particular importance: the ability of the network to satisfy the demands of new applications; and the flexibility and scalability of the infrastructure.

Mobile subscribers, looking to use the latest and greatest applications and enjoy the highest quality voice and video transmission, will place unprecedented demand on the network. 5G will enable a wide range of new use cases, including ultra-reliable and low-latency communication (URLLC), enhanced Mobile Broadband (eMBB), and massive machine type communications (mMTC). Some of these use cases will drive growth in demand for bandwidth (such as mobile broadband), while others will be more sensitive to latency (such as gaming).

At the same time, upgrading dedicated hardware at each point in the network to handle these new requirements is unsustainable, both technically and economically. Operators are instead adopting a strategy based on NFV, which replaces dedicated hardware with software functions running on commercial off-the-shelf (COTS) servers. Using standard hardware enables resources to be scaled as required, independently for each function, and enables functions to be hosted anywhere in the cloud-like infrastructure. To achieve its full potential, NFV requires there to be no hardware dependencies between the NFV infrastructure and the Virtualized Network Functions (VNFs). That enables live migration and optimal scalability.

SK Telecom had virtualized its entire mobile packet core, including Mobility Management Entity (MME), Serving Gateway (SGW), and Packet Data Network Gateway (PGW) by 2014, which enabled it to reduce its Total Cost of Ownership (TCO) and increase the flexibility of its deployments. Since then, much of the old and new equipment in the core network has been migrated to VNFs including the IP Multimedia Subsystem (IMS), Value-Added Systems (VAS) and management systems that are virtualized and integrated into the Management and Orchestration (MANO) system.

In order to achieve the necessary performance, its implementation was based on Single Root Input/Output Virtualization (SR-IOV) and PCI passthrough (PCI-PT). This enables a VNF to directly access a virtual network device, linked to the physical network device, bypassing the hypervisor. While this is faster than using para-virtualized drivers, it does introduce hardware and software dependencies that break the cloud-like model of NFV. It can only be used in appliance-like deployments.



User Plane Virtualization: Results

Throughput:	More than 200 Gbps
Latency:	70 microseconds
CPU usage:	60 percent




Figure 2. Typical load distribution in a virtual EPC system, where packets for a particular user device may be processed by different virtual machines, or different servers

The challenge is to achieve cloud-like flexibility in the virtualized user plane, while achieving a level of throughput that would help to satisfy the requirements of 5G. Such a system can help to deliver efficiencies today, and help to prepare the network for the demands of tomorrow.

Solution Value: Delivering a High-Performance Virtualized User Plane

Intel and SK Telecom worked together to investigate how the virtualized Evolved Packet Core (vEPC) can be implemented on standard hardware to handle the performance and scalability needs of 5G. Hardware-assisted Network Interface Cards (NICs) were used to handle user plane bottlenecks and enable better network utilization, including support for overlay networks. The solution is based on the Intel® Xeon® Platinum 8180 processor, which has 28 cores and 56 hyperthreads, enabling the required level of performance and virtual machine density.

The implementation described in this paper has achieved more than 200 Gbps throughput on a single dual socket rackmount server, across four different user profiles¹. The study measured latency of approximately 70 microseconds¹, and ran at approximately 60 percent CPU usage¹. Figure 1 shows the photos of the rackmount server.

This prototype implementation demonstrates that it is possible to achieve a level of performance that can satisfy the demands of 5G, and still enjoy the flexibility and agility of NFV using a standard server, free from hardware and software dependencies. The servers used in the test are designed to be highly scalable and flexible, and can be deployed in the core or at the edge of the infrastructure.

Solution Architecture: Virtualized User Plane

The User Plane Function (UPF) handles the processing of data packets between the Radio Access Network (RAN) and the data network (for example, Internet). The UPF includes a

pipeline of functionalities, including for access control, tunnel encapsulation and decapsulation, routing, quality of service, and charging. Additionally, operators may need to customize or add their own functionalities to meet different use cases or regulatory requirements.

There are several different options for implementing a UPF pipeline. Intel and SK Telecom used a software-based approach where the entire pipeline for one packet is executed on the same worker core (or thread) of the processor, and then the next packet is processed in the same way. This run-to-completion (RTC) approach improves latency and throughput, because frequently used data (such as the subscriber context) can be stored in the cache associated with the core. Using a software-based approach provides the highest degree of flexibility for modifying or adding functionalities to the pipeline. A software approach also enables the UPF function to be deployed using standard hardware, at the edge or in the core of the network. Because the software is independent of the underlying platform hardware, hardware can be scaled out to increase performance.

UPF servers typically run one or more Virtual Machines (VMs), or containers. In this way, the traffic can be evenly distributed across servers. In a sub-optimal deployment, this can result in situations where traffic for a particular mobile device might end up on different VMs or containers on a server, or worse, on different servers, as illustrated in Figure 2. The result is that additional context information must be brought in from the control plane to process the packet successfully, which increases latency. It also incurs delay when checking data accuracy, because it requires data to be managed between multiple VMs. This approach also means that the order of packets is not maintained, which can result in retransmissions from application layers.

Intel and SK Telecom propose a solution that enables all of the packets for a particular user device to be directed to the same worker core, first time. Hardware-assisted NICs




Figure 3. The User Plane Processing Pipeline

use hashing to direct packets to the correct worker core. To do this, they pick the correct fields from the inner IP packet, primarily the Src-IP address that identifies the user equipment. This method for steering packets to cores is implemented for ingress traffic on both the S1-U (mobile network side) and SGi (Internet side) interfaces. This enables packet order to be maintained and helps to optimize throughput. The Data Plane Development Kit (DPDK) is used to accelerate the movement of packets from the NIC to the worker cores.

Hardware-assisted NICs (for example, based on FPGAs) can offer a number of advantages, compared to using NICs without embedded intelligence:

- **Hashing:** To achieve greater granularity in the hashing and ensure packets are processed in order, it is necessary to look inside the packet to identify the user device, and not just the originating base station as indicated by the packet's header. Most NICs, including OpenFlow based SmartNICs, do not have the ability to parse within the GTP tunnel. A hardware-assisted NIC can be configured to parse on both the outer header and inner GTP tunnel, eliminating any requirement for a software based load balancer core and enabling worker cores to scale linearly as performance requirements increase.
- **Network overlays:** Network overlays provide a way to create many virtual networks that can coexist on the same physical network. Two of the common network overlay tunneling protocols are VxLAN and NVGRE. The tunnel endpoint capability with complete encapsulation/decapsulation of the VxLAN overlay can be implemented

in a hardware-assisted NIC. For the test setup, each VM is considered part of its own virtual network. Other mapping schemes could be considered based on deployment and VNF design.

- **Virtual switching:** This is used for intra-VM interactions when co-located on the physical server (for example, control plane and user plane VMs). It also assists in deploying, orchestrating, and managing VMs on NFV platforms. In many cases, virtual switching is performed solely in software. Accelerating some or all of the virtual switching in the NIC can reduce software overhead, particularly in the domain of lookups and performing actions such as encapsulation and address translation.
- **IPSec termination:** Transport Security in the form of IPSec or equivalent is now commonplace in networking to ensure isolation between virtual networks, particularly in a multi-tenant environment. Implementing IPSec in the hardware-assisted NIC provides two benefits. First, it's a relatively cycle-intensive task on a CPU, and a hardware implementation can reduce the overall number of CPU cores required. Second, an IPSec implementation in the NIC allows inline packet processing of a decrypted packet and enables hardware-based load balancing/distribution capability on the decrypted packets.

In the prototype implementation, the CPU subsystem comprises a Linux*-based host operating system running a KVM hypervisor on which user plane VMs are deployed. Each identical VM runs multiple user plane threads. Each user plane thread executes a complete instance of the pipeline, as shown in Figure 3.




Figure 4. Mapping virtual machines to the physical and logical cores of the 2-socket server

Figure 4 shows how VMs and housekeeping activities were allocated across the 56 logical cores of the server. Each VM is configured with a 25 GbE interface, so there is a total of ten 25 GbE ports on the system. The LTE EPC stack used for the prototype implementation is based on a commercial implementation from ASTRI. This software stack includes MME, SGW and PGW functionalities.

Figure 5 shows the hardware platform is a dual-socket server based on the Intel Xeon Platinum 8180 processor. Ten virtual machines were hosted across the two sockets, and virtual CPU threads were pinned to specific logical cores.

VM housekeeping activities were shared on a minimal number of logical and physical cores to reduce packet loss and control plane/user plane interactions. A number of platform optimizations were carried out, including powering down unused cores and starting VMs with minimal services and kernel configurations as required. Each VM's memory was configured from pre-allocated 1G huge pages to guarantee physically contiguous memory. Resources for the VMs were localized to the socket to take advantage of CPU cache and local memory channel bandwidth.

Processor	Product	Intel® Xeon® Platinum 8180 processor
	Frequency	2.5 GHz
	Cores per processor	28 cores/56 hyperthreads
Memory	DIMM slots per processor	6 channels per processor
	Capacity	192 GB
	Memory speed	2667 MHz, DDR4
Network	Number of ports	8 x 25 GbE using prototype FPGA-based SmartNIC from Intel 2 x 25 GbE using on board LOM
1U Server	Vendor	HPE* ProLiant DL360 Gen 10
Host OS	Vendor/version	Wind River* OVP6
KVM	Vendor/version	CentOS*-6.8 generic cloud image

Figure 5. The platform specification for the device under test

DUT User Plane Server Node:




Figure 6. The benchmark setup

Figure 6 shows the overall structure of the setup used to exercise various traffic profiles to measure the performance and latency. To measure performance and latency in the system, Spirent Landslide* systems were used to emulate eNodeB and mobile subscribers. The device under test (DUT) system running user plane VMs is interconnected with traffic generators (source, sink) and control plane functions (SGW-C/PGW-C, MME) via an Ethernet switch.

Performance results

Based on the measured packet distribution and throughput data, and looking forward to 5G requirements, the following four profiles were chosen for test purposes:

- **Profile 1** (4.2 — 5.7): 42.5 percent uplink throughput, 57.5 percent downlink throughput
- **Profile 2** (4.0 — 6.0): 40 percent uplink throughput, 60 percent downlink throughput
- **Profile 3** (2.5 — 7.5): 25 percent uplink throughput, 75 percent downlink throughput
- **Profile 4** (1.0 — 9.0): 10 percent uplink throughput, 90 percent downlink throughput

Figure 7 shows the uplink, downlink packet, and throughput distribution.




Figure 7. Packet rate (left), and bitrate throughput (right) for each of the traffic profiles analyzed




Figure 8. User plane worker core usages along with measured throughput (left), and measured core utilization percentages for each traffic profile (right)

The chart in Figure 8 shows the packet throughput and overall system-level throughput for each of the scenarios considered. The system was implemented by September 2017, and all performance validation related to the proof of concept was conducted and completed during October 2017. As illustrated, more than 200 Gbps throughput was consistently measured for packet rates between 30 and 42 million packets per second¹. The chart also shows the CPU usage of between 52 and 66 percent on the user plane worker cores. The system under test has a total of 56 cores (2 processors), with 40 cores used for user plane worker cores. There is significant headroom (approximately 40 percent) on the user plane processing cores, while minimizing the use of the host OS and VM housekeeping cores.

In addition, because the ten cores not allocated for user plane processing are entirely free, they can be repurposed for additional system functionality, such as for value-added service and all-active redundancy mechanisms.

"Our current telco data center has to be more efficient," said JoongGun Park, Core Network Lab Leader, Network Technology R&D Center, SK Telecom. "We have been endeavoring to accelerate user plane packet processing in the virtualized environment. This study with Intel is important because we were able to thoroughly analyze 3GPP-specified user plane modules and model different packet sizes with different resource topology. The study is a foundation for us to develop a new architecture from the iterative implementations using optimal software and hardware configuration."

He adds: "One of surprising result was that the carefully designed user plane function in a 1U rackmount proof-of-concept server was able to process more than 200 Gb/s. This is mainly the result of rigorous tests and validations from Intel and SK Telecom on the commercial mobile traffic profiles and 5G service patterns."

Conclusion

As the industry is preparing for early 5G services that will generate much more traffic, the mobile packet core system expects to receive exponential amounts of user plane loads, and so it must be able to scale and perform efficiently for 5G use cases such as eMBB, URLLC, mMTC.

This solution brief presents a new architecture for the user plane function (UPF), which shows how a combination of software and COTS hardware, together with hardware-assisted NICs, can be used to deliver the performance required for user plane virtualization. Park said: "With Intel, we were able to validate cutting-edge NIC technology that assisted with packet processing, such as offloading for intelligent load balancing and packet forwarding."

This performance can be achieved without compromising on hardware-independence, essential for achieving the flexibility and agility required for scaling and innovation in the 5G network.

Solutions Proven by Your Peers

Intel Solutions Architects are technology experts who work with the world's largest and most successful companies to design business solutions that solve pressing business challenges. These solutions are based on real-world experience gathered from customers who have successfully tested, piloted, and/or deployed these solutions in specific business use cases. Solutions architects and technology experts for this solution brief are listed on the front cover.

Find the solution that is right for your organization. Contact your Intel representative or visit [intel.com/communications](https://www.intel.com/communications).

Learn More

- [White paper: Towards Achieving High Performance in 5G Mobile Packet Core's User Plane Function](#)
- [Intel® Xeon® Platinum 8180 processor](#)
- [Data Plane Development Kit \(DPDK\)](#)
- [ASTRI](#)

Solution Provided By:



¹ Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or learn more at <https://www.intel.co.uk/content/www/uk/en/products/processors/xeon/scalable/platinumprocessors/platinum-8180.html>.

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 complete information visit www.intel.com/benchmarks

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Copyright © 2018 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon 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.

0518/SH/CAT/PDF

337478-001EN