

Samsung Achieves 305 Gbps on 5G UPF Core Utilizing Intel® Architecture

Authors

Intel Corporation— Data Center Group

Andriy Glustsov

Khaled Qubaiah

Jianwei Ma

Terence Nally

Huisuk Hong

Henry Jeong

Chetan Hiremath

Samsung Electronics - Networks Business

Ilgee Kang

Kwangseop Hwang

Sungyoon Ryu

Giljung Kim

Namgyun Kim

Yuntae Kim

Gyuil Choi

Yitae Cho

Wonsuk Song

Samsung Research

Jihun Ha

Jihwan Seo

Beomseok Oh

Sewon Oh

Seonjun Park

Vladimir Kuramshin

Dmitry Kandybka

Executive Summary

5G is creating a ripple effect of innovations and developments that will enrich our daily lives. 5G commercial services have begun in 2019 and operators have been continuing to enhance their services with the latest network technologies.

As part of an ongoing partnership, Samsung and Intel have collaborated on performance and latency optimizations of Samsung's 5G Cloud Native UPF, which resulted in a significant performance breakthrough of data throughput measuring 305 Gbps. This was achieved by utilizing software optimizations on Intel® Xeon® Platinum 8280 processor and Intel® Ethernet Network Adapter E810-CQDA2 platforms. This performance and efficiency enables operators to deliver higher quality end user experiences at a lower total cost of ownership.

1. Introduction

We have already had a taste of what 5G will deliver over the next decade from the upgraded 5G New Radio (NR) devices and networks being deployed around the world, providing access to new spectrum and bandwidth. However, the majority of deployments are 5G Non Stand-Alone (NSA) network configurations, which retains much of the 4G core network to support the 5G NR base stations.

To fully unleash the potential of 5G, and to support the new, exciting use cases like Ultra Reliable Low Latency Communication (URLLC), Industrial Control and Fixed Wireless Access (FWA), the entire network will need to be upgraded to 5G Stand Alone (SA). This involves all the 4G network infrastructure being upgraded to a brand-new Service Based Architecture (SBA).

We cannot underestimate the impact 5G will bring to our industry and to our lives, but there is also an evolution happening at the platform level which is having an equally important effect. The shift to 5G SA requires platforms that are providing increases in performance and throughput which are critical to improving the efficiency, scalability and flexibility of the 5G network.

5G User Plane Function (UPF), which is responsible for traffic forwarding and policy enforcement, plays a critical role in 5G network performance. UPF performance can be increased by optimizing the user plane pipelines, packet processing pipelines and software architectures, making it a key enabler for 5G use cases and services.

2. Samsung 5G Core Solution Overview

2.1 Introduction to Samsung 5G Cloud Native Core

Samsung's 5G Cloud Native core allows network operators to launch new services quickly and upgrade frequently according to their business needs while reducing OPEX by providing higher operational efficiency. Samsung 5G core solutions using micro-services architecture, E2E dynamic orchestration and automation, CI/CD, open source platform services, telco-grade performance support, and telco-oriented open sources will deliver an E2E solution that drives success for network operators.

Samsung 4G/5G Common core combines 3GPP network functions from EPC and 5GC architectures into a common Cloud-native Software Platform that enables operators to easily and flexibly deploy 4G, 5G NSA or 5G SA on their network



Figure 1. Top-Down Performance Optimization Methodology

according to network/business requirements.

Samsung core NFs can be provided either as VNFs (Virtualized Network Functions) or CNFs (Containerized Network Functions). When provided as CNFs, they can be deployed on any Container Platform that is aligned with Cloud-native Computing Foundation (CNCf) principles and operated as a Cloud-native implementation.

2.2 Introduction to Samsung User Plane Network Function

Tremendous growth in mobile subscriber traffic and peak data rates have been observed over past years, and this is expected to continue as 5G Mobile deployments roll out. This requires User Plane performance enhancements and optimization with scalable and efficient data traffic handling for 5G use cases such as eMBB, URLLC and mMTC.

Samsung provides a common UP with UPF, PGW-U, SGW-U and several value-added services such as Content Filtering, DPI, NAT, Firewall etc.

With combined 4G/5G session management and user traffic handling, HW resources (e.g. CPU, Memory) can be reduced with optimized resource pooling between 4G and 5G functions.

Samsung User Plane NF is designed and implemented to be highly scalable and flexible, and can be deployed in an optimized dimensioning with variety of physical scenarios.

Samsung utilizes DPDK (Data Plane Development Kit)/VPP (Vector Packet Processing) technologies, NUMA binding, and huge pages for higher performance processing and packet acceleration technologies such as SR-IOV for telco-grade I/O performance. It also improves packet transfer performance by leveraging parallel packet processing that processes both QoS control and packet transfer stipulations simultaneously.

3. Key Intel Technologies for UPF Optimization

Tasked with optimizing Samsung's 5GC for Intel platforms, a joint team was formed. We targeted COTS servers powered by second generation Intel® Xeon® Scalable processors to run Samsung 5GC UPF pipeline to take advantage of Intel's technologies to address extremely high bandwidth for I/O, and to address the tight and strict time budget and latency defined by 3GPP standards to process user plane packets.

We used classical top-down methodologies to improve the performance including lower latency, higher throughput, and improved CPU utilization as shown in Figure 1.

Broadly speaking, performance optimizations are carried out

in three main areas: IA platform, Packet IO and distribution, and Packet processing.

3.1 IA Platform Optimization

The Intel® Xeon® Scalable processor family provides a unique core micro-architecture that offers a unified LLC (L3 Cache) for all cores in a CPU, this design will guarantee a deterministic and consistent memory latency access from all cores in a CPU, which eventually will allow scaling UPF pipeline to all available cores/threads in a very linear fashion.

Since 5GC UPF is a very intensive I/O bandwidth application, the memory access and cache system need to be carefully optimized. Using Intel DDIO technology to enable direct communication between the ethernet controller and CPU last level cache eliminates many memory accesses. This technology helps to reduce latency and reduce memory bandwidth and power consumption, which are all required resources that must be minimized so they can be used to process more packets and push UPF performance to the limits.

Intel® Advanced Vector Extensions 512 (Intel® AVX-512) is a set of new instructions that can accelerate performance for 5GC UPF workloads by delivering improvements to the VPP infrastructure library. This enables execution of one instruction on multiple data sets simultaneously which is beneficial in TX and RX traffic operations.

Another important Intel technology being used is the cache pre-fetching via hardware prefetcher and specifically for the cache lines from main system level DRAM to CPU internal caches before it is needed, Intel offers multiple HW prefetcher instructions for different use cases, this technology was used heavily in performance optimization for the used 5GC UPF. The usage of cache prefetchers in the right places enhance the memory access time thus the instructions are executed in less clock cycles.

2nd Gen Intel® Xeon® Scalable processors offer six DDR4-2933 memory channels per CPU, the 5GC UPF was configured to utilize all the available memory channels with the maximum memory speed available, enhancing the memory access and reducing latencies further which helped to push the UPF performance to new limits.

The CPU enables a new level of consistent, pervasive, foundational enhancements including higher per-core performance; higher Memory Bandwidth/Capacity; Intel Advanced Vector Extensions 512: Intel Speed Select technology etc.




Figure 2. Advanced Packet Distribution Capability for UPF

Intel's rich ecosystem and development tools like Intel® VTune™ Profiler speed up the optimization cycle. By immediately working on the bottlenecks in the used workload, they speed up the development cycle and shorten the time to market. The specialized monitoring counters can be collected to expose information on the hardware resource consumption to identify the critical processes, threads, modules, functions, and lines of code.

3.2 DDP for Packet IO and Distribution

Intel® Ethernet 800 Series is the next generation of Intel® Ethernet Controllers and Network Adapters. The Intel® Ethernet Network Adapter E810 is designed with an enhanced programmable pipeline, allowing deeper and more diverse protocol header processing. This on-chip capability is called Dynamic Device Personalization (DDP). This allows the consistent parsing and steering of traffic from a given UE or UE flow to a worker core.

A network controller packet pipeline is responsible for packet identification and reporting protocol information on the packet's receive (Rx) descriptor. This information is used by the filters and queue management on the controller and by upper layers of software.

To optimize UPF performance, it is important for the NIC to understand the protocols and tunnels that are received to allow for filtering and various stateless offloads to assist in the packet processing.

3.2.1 Intel® Ethernet Network Adapter E810 Enhanced DDP Package for Telecommunications

The Enhanced DDP Package for Telecommunications used for UPF supports GPRS Tunneling Protocol (GTP). Metadata can be extracted from the GTP headers, then used in the subsequent steps of the packet processing engine of the network adapter including the switch, Receive Side Scaling (RSS) and Intel® Ethernet Flow Director (Intel® Ethernet FD).

3.2.2 DDP Based Packet I/O Distribution

One of the most critical requirements in the telecom environment is determinism. One way to ensure this is by utilizing a Run to Completion (RTC) pipeline implementation,

where packets can be received from the NIC queues and fully processed by the same CPU core. Packet parsing and classification capabilities of the NIC allow Receive Side Scaling (RSS) based load distribution of packets between receive queues, such that queues can be assigned to worker cores.

With DDP profiles the packet classification capabilities of the Intel® Ethernet Network Adapter E810 are extended, the GTP flow types are defined and encapsulated frame fields (including GTP TEID, QFI, fields of encapsulated IP header) can be used for Flow Director, queue group selection and RSS-based load distribution of packets within queue groups.

As seen in Figure 2, the network adapter has full visibility of header fields and can perform packet distribution to receive queues based on this improved classification capability carrying out load distribution inline, removing the need for a load distribution cores to perform the same function in software.

3.2.3 Generic Flow Rule

The DPDK Generic flow API (`rte_flow`) is used by the network adapter to match specific ingress traffic and forward it to specified queues. The specific ingress traffic is identified by a matching pattern that is composed of one or more Pattern items. Once a match has been determined, one or more associated Actions will be performed.

Several flow rules can be combined such that one rule directs traffic to a queue group based on QFI/DSCP, a second rule distributes matching packets within that queue group using RSS.

3.3 Packet Processing

The packet processing pipeline gets performance improvement from the network adapter-based distribution of received packets made possible by DDP support for 5G UPF: all packets of the same IP flow (both plain IP packets received on N6 interface and GTP-U encapsulated IP packets received on N3/N9 interface) are delivered to the application stack over the same RX queue and processed by the same application thread bound to the CPU core ensuring packet buffers and flow-related control structures are not shared between cores.

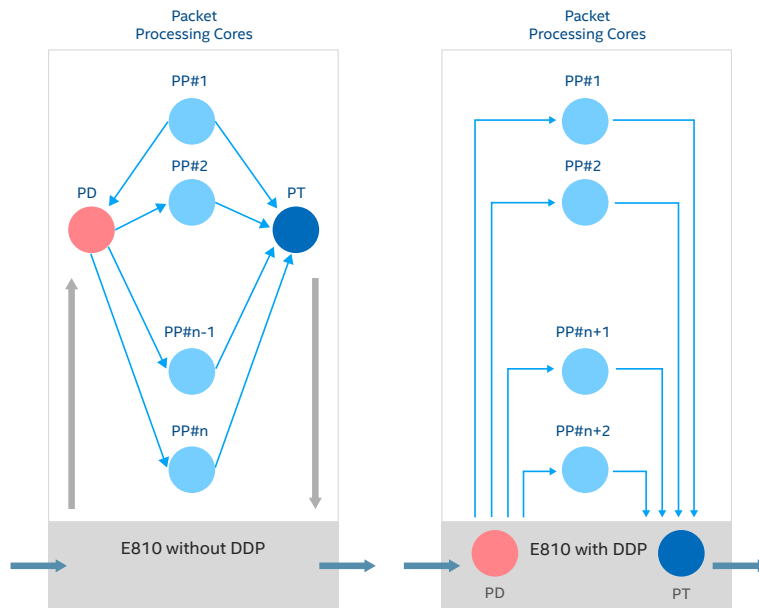


Figure 3. Packet distribution and processing without and with DDP

Further performance improvement is achieved by UPF software stack optimization including software architecture changes and code optimizations.

Architectural changes are done in two main areas:

- Packet processing graph nodes combination & Protocol-based pipeline re-organization
- Repurposing of packet receive and load distribution cores and packet transmit cores to perform packet processing in run to completion mode as shown in Figure 3.

Code optimization is executed in several iterations by collecting profiling data and measuring maximum UPF performance at zero packet loss using selected test traffic profiles at every iteration. The manual code review and optimization process are facilitated by Intel® VTune™ Profiler tool used to collect and analyze key profiling data from the target platform in order to identify the areas for optimization, and to measure and control the impact of the introduced code changes on the UPF performance. The code optimization involved changes in several areas listed below.

- Algorithm optimization and use of optimized implementations from VPP and DPDK for basic functions like hashing, lookups, etc.
- Relocating selected application data structures frequently accessed by the UPF code into the huge page memory region in order to reduce Translation Lookaside Buffer (TLB) misses.
- Data structure layout adjustment and cache alignment of frequently used structure members to reduce memory footprint and improve cache utilization.
- Ensuring cache line alignment for data in the application to minimize cache line sharing and to improve prefetching efficiency.
- Identifying and avoiding false cache line sharing among multiple threads of the application. Restructuring data modified from multiple CPU cores (like statistic counters)

to keep them in memory on per-core basis so that data updated from different cores are not located in the same cache line.

- Prefetching packets headers and context data structures to cache. Efficient memory prefetching is made possible by the nature of vectorized packet processing where code has knowledge not only of the next processing stages for the current packet but also information on the next packets to be processed.

4. Samsung NGCore UPF Solution with Intel Platform and Technologies

4.1 Samsung NFV UPF with Intel® Ethernet Network Adapter E810 with Enhanced DDP Technology

Without DDP acceleration technology, the UPF had to have multiple dedicated cores (Packet Distribution/Packet Transmission) to steer the same flow packets to fixed packet processing cores and transmit them to the network adapter in order. PD (Packet Distribution) cores parsed inner/outer IP headers of all incoming UL(GTP-U)/DL packets to get a unique hash value from 5 tuples of each IP header and deliver them to the fixed cores. But with the increasing need for performance improvement, the distribution/transmission load requires more and more processing and it becomes more of a challenge to avoid forwarding bottlenecks and maintaining target latency at a high load.

Using DDP for UPF, the Intel® Ethernet Network Adapter E810 has the ability to parse more in-depth layer protocol to get inner IP header fields of incoming GTP-U packet for RSS hash calculation. UPF can offload its packet distribution/transmission workloads to the network adaptor and reuse packet distribution cores as PP cores, boosting UPF performance.

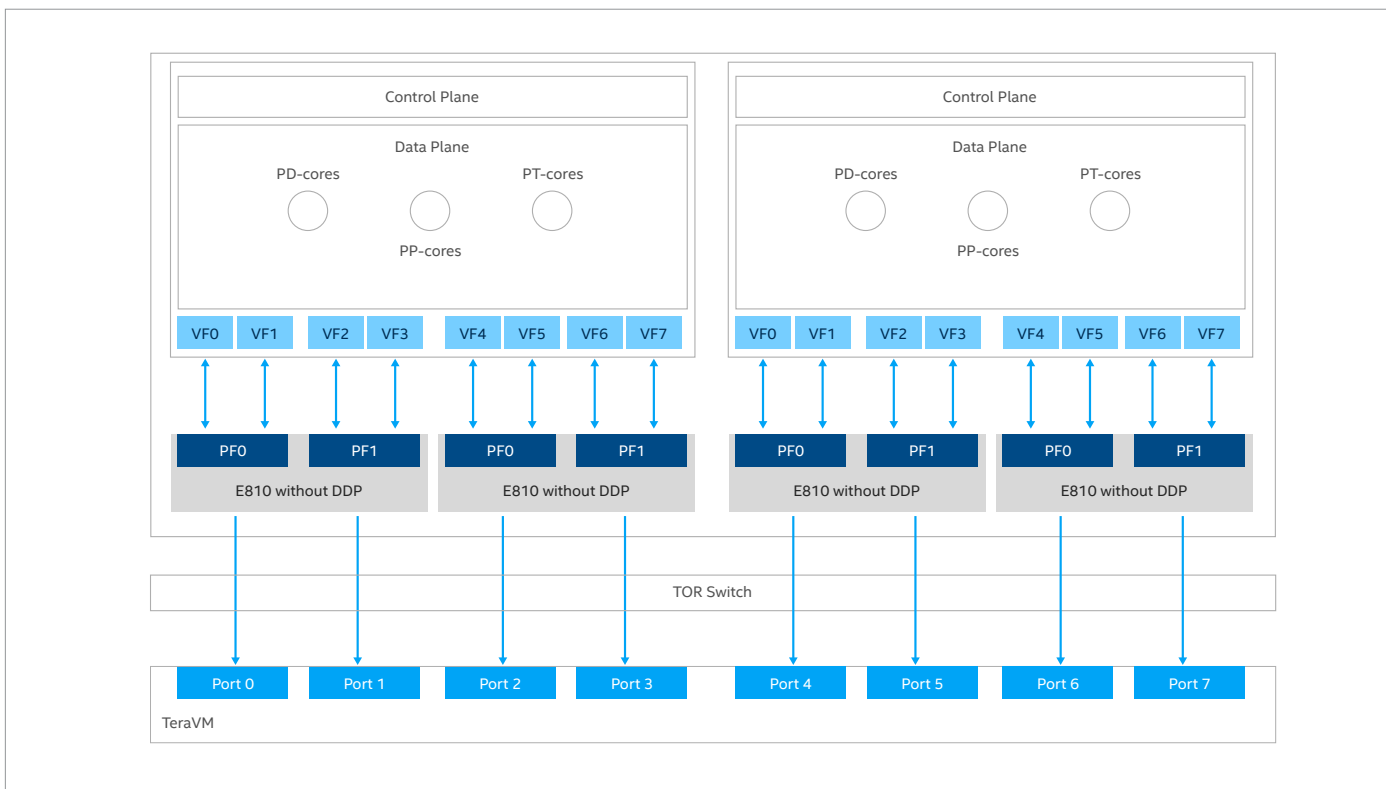


Figure 4. Network Adapter IO without DDP and CPU Core Mapping




Figure 5. Network Adapter IO with DDP and CPU Core Mapping

5. System Test Environment

5.1 Traffic Test Model

The tests were carried out in Samsung's lab with a typical CoSP traffic model configuration as follows.

| | |
|-----------------------------|-----------|
| Total number of subscribers | 600,000 |
| Active Subscribers | 70,400 |
| Volume Ratio | UDP: 100% |
| Average packet length | 690 bytes |

5.2 Test Tools (TeraVM)

VIAVI's TeraVM was used to simulate gNodeB and PDN Network server and inject data traffic into Samsung UPF system. It sets up calls and puts Uplink/Downlink high-volume traffic into Samsung 5G-UPF.

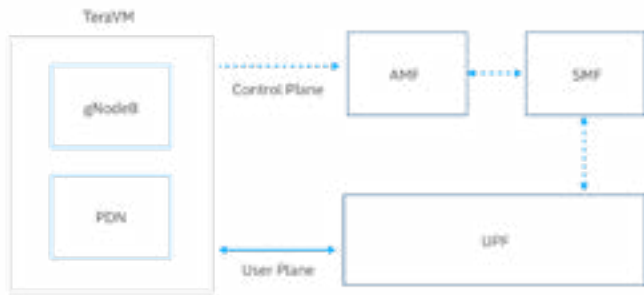


Figure 6. Logical topology of test setup

5.3 Hardware Configuration

The tests used HPE ProLiant DL380 Gen10 Server powered by two Intel® Xeon® CPUs with each socket being connected to two Intel® Ethernet Network Adapters. Tests were conducted by Samsung Electronics at Samsung Electronics Labs on Nov 23rd, 2020.

| | HP Server |
|-----------------|--|
| CPU | Intel® Xeon® Platinum 8280 Processor |
| Number of CPUs | 2 |
| Memory | 384GB DDR4 |
| Network Adapter | 2 x Dual-port 100Gb Intel® Ethernet Network Adapter E810-CQDA2 |
| Microcode | x5002f00 |
| BIOS version | Vendor : HPE, Version : U30 |

5.4 Software Configuration

The following software configuration was used for the device under test.

| | |
|----------------|--|
| OS | RHEL 7.5 |
| UPF(VNF) | SVR (Samsung VNF Release) 20A |
| DPDK | 20.5 |
| Host OS | RHEL 7.6 |
| Host OS kernel | Linux G5-U24-NOVA 3.10.0-862.11.6.el7.x86_64 |

5.5 Network Topology of Performance Test

With DDP for UPF technology, the Data Plane was able to remove the need to allocate vCPUs for PD (Packet Distribution) and PT (Packet Transmission) and allowing them to be used for PP (Packet Processing) which significantly boosted performance of Samsung's 5G-UPF. In addition, removed PDs led to the reduction of number of VFs and made UPF system configuration simpler, as seen in the table below.

6. Performance Results

With DDP for UPF technology, Samsung UPF forwarding performance reached 305 Gbps while (one way) latency also saw an improvement, decreasing to 69 usec with DDP in subsequent tests.

7. Summary

Through this industry collaboration, Samsung and Intel achieved 305 Gbps forwarding capability on Samsung's 5G UPF solution. This demonstrates how the optimization of packet processing pipelines and software architectures, along with a high performance user plane pipeline, can result in meaningful UPF performance increases for 5G use cases and services. This allows CoSPs to serve the increasing traffic demand of 5G more efficiently by getting more performance from their infrastructure which delivers a higher return on their investment.

Abbreviations

| Term | Description |
|-----------|--|
| CI/CD | Continuous Integration and Continuous Delivery |
| DDP | Dynamic Device Personalization |
| GNB | gNodeB |
| LUT | Lookup Table |
| msg 1-15 | Message Type 1-15 (subtype of PFCP) |
| msg 50-57 | Message Type 50-57 (subtype of PFCP) |
| PD | Packet Distribution |
| PFCP | Packet Forwarding Control Protocol |
| PP | Packet Processing |
| Prio | Priority |
| PT | Packet Transmission |
| QFI | Qos Flow Identifier |
| Qgroup | Queue Group |
| SR-IOV | Single Root I/O Virtualization |
| SMF | Session Management Function |
| TLB | Translation Lookaside Buffer |
| UPF | User Plane Function |



Notices & Disclaimers

See backup for workloads and configurations. Results may vary.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

This test result was conducted at optimized lab environment. Actual field result may vary under circumstances. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel and Samsung technologies may require enabled hardware, software or service activation.

Samsung logo and other Samsung marks are trademarks of Samsung Electronics or its subsidiaries.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.