CASE STUDY

Implementation of ZTE's High-Performance 5G Core Network UPF

Testing using 2nd generation Intel® Xeon® Scalable processors with Intel® Speed Select Technology and Intel® Ethernet Network Adapter XXV710 with Dynamic Device Personalization (DDP) technology showed that in a telecom carrier's standard traffic test model with a packet length of 690 bytes, the highest overall performance reached 173 Gbps and 252 Gbps without any additional hardware acceleration devices.

Authors
Zhang, Jerry (INTEL)
Lou, Fangliang (ZTE)
Ma, Jianwei (INTEL)
Zhang, Yong (ZTE)
Ma, Jiangying (ZTE)
Xu, Yang (ZTE)
Zou, Jinchun (ZTE)
Cheng, Peng (ZTE)
Xu, Zhou (ZTE)
Yao, Chenghu (ZTE)
Zhang, Haoming (ZTE)

Contents

Foreword

Currently, the NFV-based core network has been used commercially at a large scale worldwide. The most important feature of NFV is that it utilizes universal hardware and NFV infrastructure (NFVI) to provide unified operating resources for upper-layer virtualized network functions (VNFs). This technology provides a foundation for flexible deployment, rapid launch, intelligent & automated operation and maintenance as well as network transformation. As one of the key elements of the 5G core network, the user plane function (UPF) enables 5G users' network data forwarding and its deployment on Intel® architecture-based high-throughput servers has won industry-wide recognition in the current virtualization trend. It has become a pressing issue for telecom carriers and OEMs to improve the performance of virtualized UPF based on Intel architecture.

This paper describes the excellent performance achieved by ZTE's 5G core network UPF products based on 2nd generation Intel® Xeon® Scalable processors—Intel Xeon Gold 6230N processors and Intel Xeon Platinum 8280 processors—and Intel® Ethernet Network Adapter XXV710 with Dynamic Device Personalization (DDP) technology. **The test results show that in a telecom carrier's standard traffic test model with a packet length of 690 bytes, the highest overall performance achieved was 173 Gbps on Intel Xeon Gold 6230N CPUs and 252 Gbps on Intel Xeon Platinum 8280 CPUs without any additional hardware acceleration devices.**

The testing for these findings was conducted by ZTE on 20 December 2019. See System Test Environment section for configurations. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. Refer to http://software.intel.com/en-us/articles/optimization-notice for more information regarding performance and optimization choices in Intel software products.

1. ZTE's 5G Core Network Solution

1.1. Introduction to the Overall Solution

1.1.1 Introduction to ZTE's TECS Cloud Management Platform

TECS (Tulip Elastic Cloud System) is cloud platform management software developed by ZTE, based on the OpenStack open source standard interface. It mainly comprises two types of components—OpenStack and Extension—and is used for managing, monitoring, and fault reporting of infrastructure-layer hardware resources and virtualization resources. It also provides virtualized resource pools for upper-layer business applications. TECS supports global resource scheduling, flexible capacity scaling, and flexible network adjustment, and fully supports the flexible deployment of upper-layer IT and CT applications. It also provides IaaS business operation capabilities—including componentized micro-service platforms, rapidly iterative business development models, and intelligent analysis based on big data—and enables rapid service launch and innovation. In addition, TECS has enhanced performance, reliability, and architecture advancement to help telecom carriers cope with the challenges of the 5G/IoT era.

1.1.2 Introduction to ZTE's 5G Common Core

5G commercial deployments have been launched globally. As the "brain" of the 5G network, the core network is undoubtedly one of the critical points of the network construction. Hence, ZTE has taken the lead in launching the 5G Common Core solution based on its deep understanding of network design, network construction, and network operation. The 5G Common Core incorporates cloud native, network slicing, MEC, AI and other technologies to create a "stronger brain", establishing a minimalist, intelligent, integrated, and efficient network for telecom carriers.

ZTE's 5G Common Core supports full access to 2G, 3G, 4G, 5G and fixed networks and achieves four integrations (data integration, strategy integration, control integration, and forwarding integration) for the 5G network function and service. Based on SBA+, it enables functional components to be shared across networks and one network for multiple purposes. While simplifying the network topology, it helps improve service KPIs and delivers an excellent user experience. At the same time, it supports five deployment modes of 3GPP R15 SA and NSA (Options 2/3/4/5/7). Telecom carriers can deploy flexibly SA, NSA or SA & NSA dual modes on the same network according to their network requirements, complete the architecture in one step, re-utilize resources, and avoid multiple network upgrades and transformations, which reduces the costs of network construction and meets telecom carriers' 5G deployment goals.

1.2. Introduction to ZTE's NFV UPF Network Element

As a key network element of the 5G core network, UPF undertakes the important functions of user data processing and forwarding. With the commercial advancement of 5G, there is an urgent demand for eMBB/URLLC applications such as 4K/8K high-definition video, cloud gaming, remote driving, and industrial control. These applications require a network with ultra-low latency and ultra-high bandwidth. These performance requirements are closely related to the UPF, which is needed to provide efficient data processing and forwarding.

The pure software UPF launched by ZTE is based on a server using Intel Xeon processors and incorporates performance optimization technologies that are used widely across the industry, such as DPDK (Data Plane Development Kit), NUMA binding, OS optimization, and huge pages. In addition, ZTE has been optimizing the software architecture and the logic of UPF. Based on the VPP (Vector Packet Processing) principle, ZTE adopts message multi-queue lock-free processing, service first packet DPI, and flow table to achieve batch processing and forwarding of hotspot messages, which effectively reduces the CPU consumption by business logic. The performance of the UPF has reached an outstanding level under the same conditions.

2. Key Features of Intel® Platforms and Technologies

2.1. Introduction to Intel Platform

With excellent performance and rich features, a server featuring an Intel® platform delivers a complete ecosystem and a robust roadmap and is the mainstream NFV platform in the industry. To meet the flexibility, security and accessibility requirements of application services, Intel has been working with ecosystem partners to develop solutions based on NFV and SDN to provide an open network infrastructure platform. Together, Intel® processors, Intel® network adapters, DPDK, and other relevant technologies form a platform that reduces device complexity and improves network utilization to meet telecom carriers' requirements for the flexibility, scalability, reliability, and cost efficiency of telecommunication networks, accelerating technological innovation and network transformation.

2.2. Intel® Speed Select Technology – Core Power (Intel® SST-CP)

The 2nd generation Intel Xeon Scalable processor is equipped with the enhanced Intel® Speed Select Technology (Intel® SST) to allow the system to provide more control for CPU performance, thus improving performance and optimizing configuration costs. With Intel Speed Select Technology, one server can do more.

The 2nd generation Intel Xeon Scalable processor "N" SKUs, which are specialized for Networking/NFV, support Intel® Speed Select Technology - Core Power (Intel® SST-CP). This function can be used to configure and adjust the priority of the CPU cores flexibly, so high-priority CPU cores run at higher frequencies. This enables the system's critical missions to run on high-priority CPU cores to ensure performance and maximize the use of CPU rated power.

2.3. Dynamic Device Personalization (DDP)

Intel's Dynamic Device Personalization (DDP) is an advanced feature of Intel® Ethernet 700 Series and Intel Ethernet 800 Series network adapters. This feature allows dynamic reconfiguration of the data packet processing pipelines by loading firmware profiles to meet specific scenario needs. In other words, these network adapters are partially programmable. By downloading specific firmware profiles, they provide users with parsing support for specific communication network protocols. Coupled with their flow-director FDIR (stream guidance) and RSS (hashing technique) features, they enable hardware offloading of network message parsing and distribution, improving network performance. Intel® Ethernet 700 Series provides multiple industrial-grade profiles, which currently cover multiple types of protocols, including PPPOE, GTP-meC, GTP-U, and L2TP. These profiles can be loaded easily through generic ethtool or DPDK libraries.

DDP allows telecommunications protocols to be steered into specific queues directly from the NIC. This means that the application does not need receiving (RX) or transmitting (TX) cores in order to steer subscriber traffic, so these cores may be freed up for application workloads. As this steering is done in hardware, a more deterministic platform is seen from an NFVI point of view: round-trip delay and jitter are reduced, and with the liberation of load-balancing RX/TX cores, the overall throughput of the system also increases.

3. ZTE NFV UPF Solution with Intel Platform and Technologies

3.1. Application of Intel SST-CP Acceleration Technology

To achieve excellent performance in practical applications, Intel® Turbo Boost Technology is often enabled, so the CPU cores can run at a higher Turbo frequency. Turbo Boost technology is a function of thermal headroom. Intel SST-CP acceleration technology can be used to lock cores for critical missions at high frequencies and cores for other missions at low frequencies, achieving excellent performance while keeping the CPU frequency stable.

ZTE's 5G NFV UPF solution uses the Intel SST-CP technology, which sets the CPU cores processing the UPF forwarding threads to run at a high priority and a higher frequency. As shown in the figure below, on an Intel Xeon Gold 6230N processor-based platform, Cores 0-7 of each CPU run at 1.5 GHz, and Cores 8-19 run at 2.8 GHz.

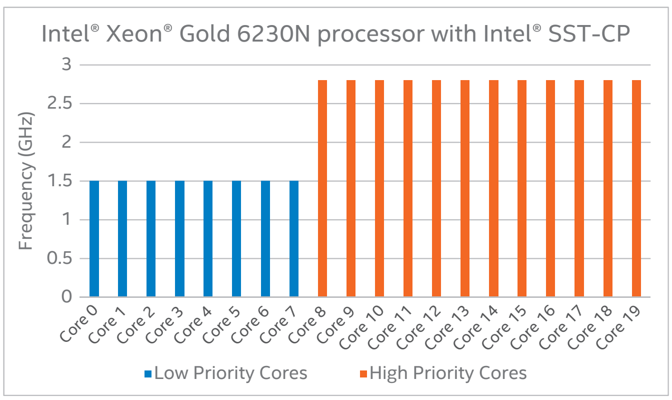

Figure 1. Intel® Xeon® Gold 6230N processor with Intel SST-CP core priority

3.2. Application of DDP Acceleration Technology

Without DDP, the inner header of a GTP (GPRS tunneling protocol) message cannot be parsed by the network adapter, as the network adapter only assumes that it is an ordinary Layer 4 message by default. In this way, for the UPF of the 5G core network, the main types of messages processed are GTP (uplink) messages and ordinary TCP/IP (downlink) messages.

GTP messages cannot be hashed to different queues and bound to different processor cores for parallel processing by using the FDIR or RSS distribution function of the network adapter. To solve this challenge, traditional solutions use software for processor cores to parse the GTP inner headers and distribute the messages to different processing cores. In this scenario, the performance of the software distribution cores often causes a bottleneck for the UPF processing, while software distribution increases the system's overall processing latency overhead.

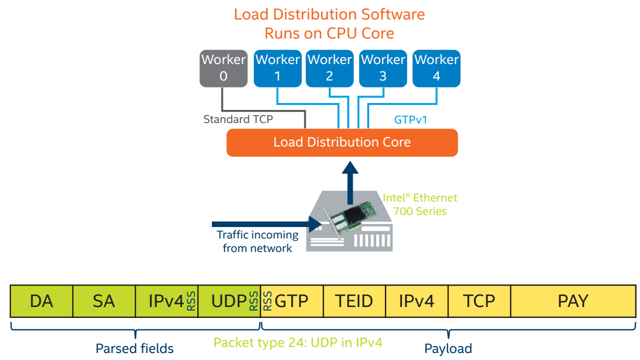

Figure 2. Software architecture without DDP technology

With DDP, the network adapter can extend the message identification depth to the transmission layer of the GTP inner messages, so the FDIR and RSS distribution functions can be applied directly to the inner headers of the GTP messages. With the network adapter, the GTP messages are hashed directly to different network adapter queues and bound to different cores for parallel processing. In this way, hardware offload of the distribution function can be achieved, performance improved, and the overall system processing latency reduced.

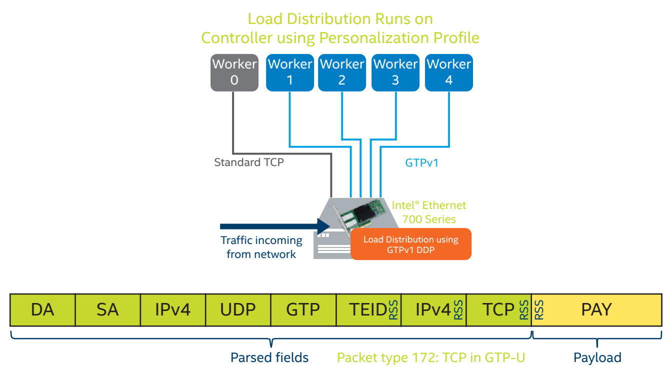

Figure 3. Software architecture with DDP technology

ZTE's 5G NFV UPF solution demonstrates the in-depth redevelopment and optimization of the DDP function carried out by ZTE and Intel. The two companies have utilized DDP features to distribute the identification of GTP messages on Intel Ethernet Network Adapter XXV710 SR-IOV virtual functions.

The CPU cores in the UPF have two roles: receivers (RX) and workers. Receivers are responsible for obtaining data packets from the network adapter's RX ring queue, classifying and distributing them to specific workers as well as balancing workload among workers. The workers implement the user plane stack function and handle uplink (UL, from UE/eNB to PDN) and downlink (DL, from PDN to eNB/UE) traffic. They process data packets in the run-to-completion (RTC) mode.

The UPF classifies each received data packet and distributes it to a worker for processing. To obtain better cache utilization and improve performance, the UPF binds all data traffic from the same UE IP to the same fixed worker on uplink and downlink.

To fix the UE IP to the same worker, the UPF uses the UE IP address as the key for worker identification. The uplink traffic comprises GTP-U encapsulated IP messages, so the source address is extracted from the encapsulated IP messages as the UE IP address. The downlink traffic comprises normal IP messages, so the UE IP address is the destination IP address of the message. By modifying the matching keywords of the GTP-U protocol and common IP message types using the DDP technology, the UPF binds all data traffic from the same UE IP to a fixed worker to enable symmetric hash of asymmetric messages and affinity of bidirectional flows, minimizing receivers and improving system performance.

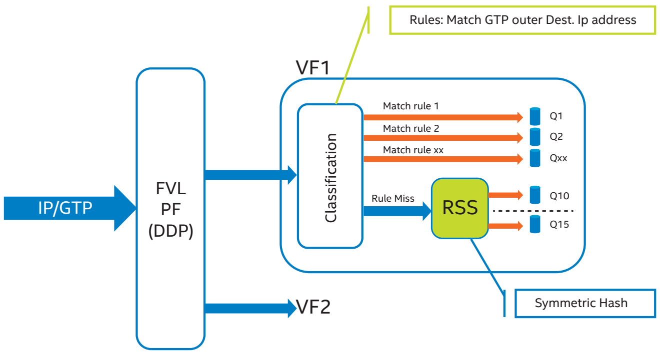

Figure 4. DDP practice in ZTE UPF solution

4. System Test Environment

4.1. A Telecom Carrier's Traffic Test Model

The main configuration of a telecom carrier's standard traffic test model is as follows:

Number of users[1]	Access users: 600,000 Data users: 6,000
Content billing configuration rules (DPI)	L7 Networking: 40,000 L3 Networking: 10,000
PCC strategy	Static: 45, dynamic: 5
Volume ratio	HTTP: 85% UDP: 15%
Average packet length[2]	690 bytes

4.2. Test Instruments (Ixia)

In the test, which was conducted 20 December 2019 by ZTE, the Ixia IxNetworks-XGS2 was used to simulate the control and data planes when 5G mobile users access them through the telecom carrier's base stations.

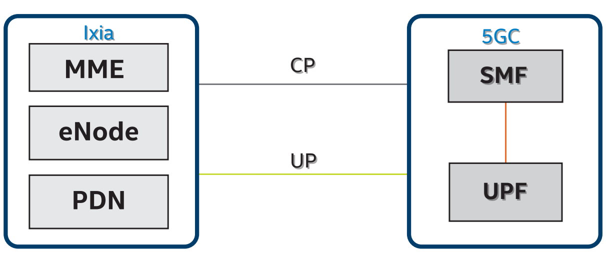

Figure 5. ZTE UPF test instruments

4.3. Hardware Configuration

The servers used were ZTE's self-developed R5300-G4, which are powered by 2-socket Intel® CPUs with each socket being connected to two or three Intel® Ethernet Network Adapters XXV710-DA2.

SERVER	ZTE R5300-G4
CPU	Intel® Xeon® Gold 6230N processor Intel® Xeon® Platinum 8280 processor
Number of CPUs	2
MEMORY	384 G DDR4 @ 2666 MHz
NIC	Intel® Ethernet Network Adapters XXV710-DA2 * 4 on Intel® Xeon® Gold 6230N CPU-based server Intel® Ethernet Network Adapters XXV710-DA2 * 6 on Intel® Xeon® Platinum 8280 CPU-based server
TOR	ZTE 5960-4M

4.4. Software Configuration

OS	ZTE CGSL 3.10.0-693.21.1.el7.x86_64
OpenStack Platform	ZTE TECS 3.0 (OpenStack 3.8.1)
QEMU	QEMU 2.5.0
UPF VNF	ZXUN-xGW (GUL) V7.19.13
i40e Driver	Driver Version: 2.7.29 Firmware Version: 6.02 0x80003620 1.1747.0
DPDK	18.11

4.5. Network Topology

In the test, ZTE's TECS was used as the cloud management platform and two Intel processor-based servers were used as the control node and the compute node respectively. The UPF was deployed on the compute node while the UPF network element was connected to the Ixia tester with a 5960-4M switch.

The test topologies for the Intel Xeon Gold 6230N processor-based server and the Intel Xeon Platinum 8280 processor-based server are shown in Figure 6 and Figure 7.

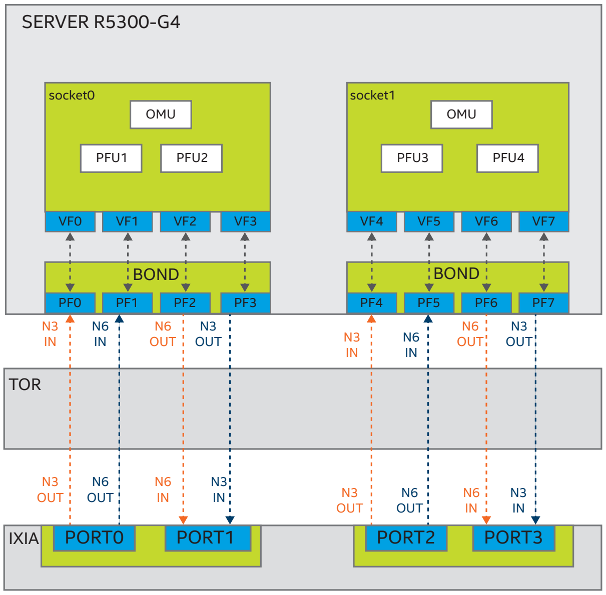

Figure 6. Test topology for the Intel® Xeon® Gold 6230N processor-based server

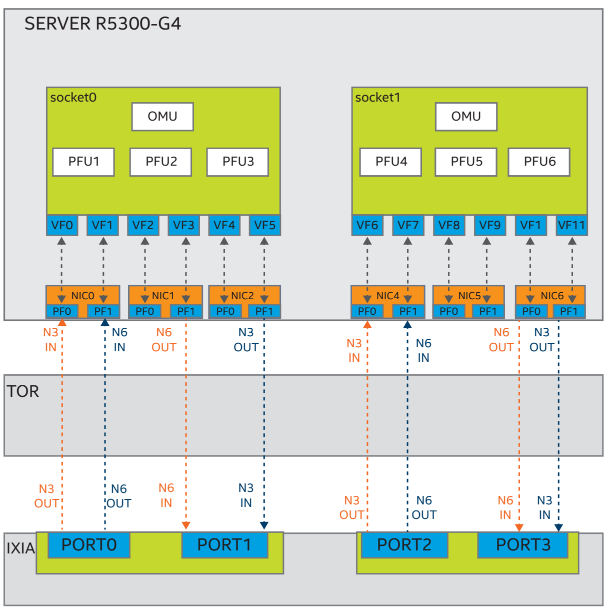

Figure 7. Test topology for the Intel® Xeon® Platinum 8280 processor-based server

4.6. BIOS Configuration for a Compute Node

The BIOS configuration that was used for the server as a compute node is shown in the following table:

Menu	Path to BIOS Setting	BIOS Setting	Required Settings
CPU CONFIGURATION	ADVANCED -> PROCESSOR CONFIGURATION	INTEL® HYPER THREADING TECHNOLOGY	ENABLED
		INTEL® VIRTUALIZATION TECHNOLOGY	ENABLED
POWER CONFIGRATION	ADVANCED -> POWER & PERFORMANCE	CPU POWER & PERFORMANCE POLCY	PERFORMANCE
	ADVANCED -> POWER & PERFORMANCE -> CPU P STATE CONTROL	ENHANCED INTEL SPEEDSTEP TECH	ENABLED
		INTEL® TURBO BOOST TECHNOLOGY	ENABLED
	ADVANCED -> POWER & PERFORMANCE -> HARDWARE P STATES	HARDWARE P-STATES	DISABLED
	ADVANCED -> POWER & PERFORMANCE -> CPU C STATE CONTROL	PACKAGE C-STATE	C0/C1 STATE
		CIE	DISABLED
		PROCESSOR C6	DISABLED
IO CONFIGURATION	ADVANCED -> INTEGRATED IO CONFIGURATION	INTEL VT FOR DIRECTED I/O	ENABLED

4.7. Virtual Machine Configuration for UPF

One OMU VM and two PFU VMs were deployed on each socket of the Intel Xeon Gold 6230N processor-based server. The OMU occupied 4 vCPUs and each PFU occupied 16 vCPUs (12 vCPUs ran worker forwarding threads). Intel SST-CP technology was used to set the core frequency of worker threads to 2.8 GHz and the remaining core frequencies to 1.5 GHz.

One OMU VM and three PFU VMs were deployed on each socket of the Intel Xeon Platinum 8280 processor-based server. The OMU occupied 4 vCPUs and each PFU occupied 16 vCPUs (14 vCPUs ran worker forwarding threads).

5. Performance Test Results

5.1. Overall Forwarding Performance

The performance tests verified respectively the basic forwarding capabilities of the UPF solution and the forwarding capabilities of the UPF solution with business processing capabilities including billing and DPI. The bladed (dual processor) test results are as follows:

1) The basic forwarding performance of the UPF solution based on dual Intel Xeon Gold 6230N processors was as follows:

   a) Each Intel Xeon Gold 6230 CPU totally used 12 cores (24 HTs) used for UPF worker forwarding threads.

   b) When offline billing and DPI was disabled, the forwarding performance reached 173 Gbps (30.6 MPPS). Worker forwarding threads' average core utilization was 83%.

   c) When offline billing and DPI was enabled, the forwarding performance reached 132 Gbps (23.4 MPPS). Worker forwarding threads' average cores utilization was 85%.

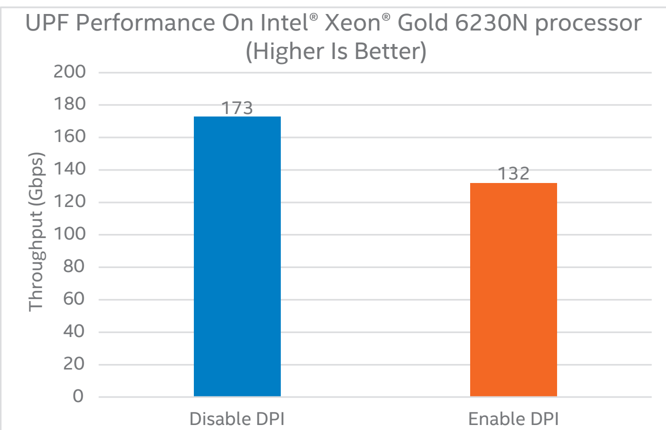

Figure 8. ZTE UPF performance result on Intel® Xeon® Gold 6230N processor (For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. Refer to http://software.intel.com/en-us/articles/optimization-notice for more information regarding performance and optimization choices in Intel software products)

2) The basic overall forwarding performance of the UPF solution based on dual Intel Xeon Platinum 8280 processors was as follows:

   a) Each Intel Xeon Platinum 8280 CPU totally used 21 cores (42 HTs) for UPF worker forwarding threads.

   b) When offline billing and DPI was disabled, the forwarding performance reached 252 Gbps (44.5 MPPS). Worker forwarding threads' average core utilization was 53%.

   c) When offline billing and DPI was enabled, the forwarding performance reached 201 Gbps (35.3 MPPS). Worker forwarding threads' average core utilization was 85%.

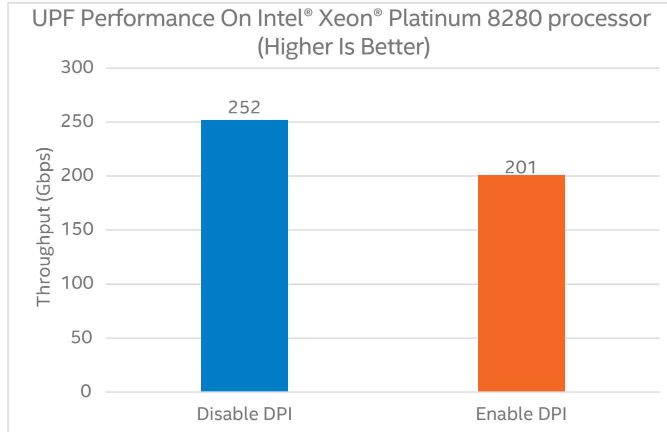

Figure 9. ZTE UPF performance result on Intel® Xeon® Platinum 8280 processor

3) The tests demonstrated great improvement in the forwarding performance of the UPF solutions including the Intel SST and DDP acceleration technologies compared with a UPF solution using Intel Xeon Gold 6138 processors,[3] which does not include these technologies.

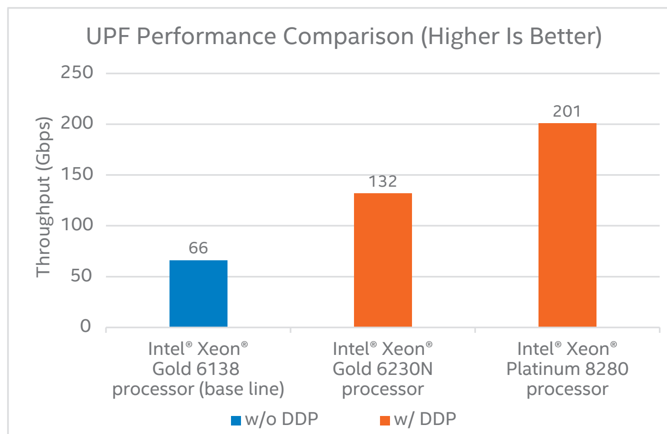

Figure 10. ZTE UPF performance comparison

5.2. Average Forwarding Latency

With the DDP acceleration technology, it is unnecessary to use software to distribute messages among cores, reducing the forwarding latency greatly. The test results show that the one-way average forwarding latency of UPF messages was reduced from 150 µs to 74 µs.

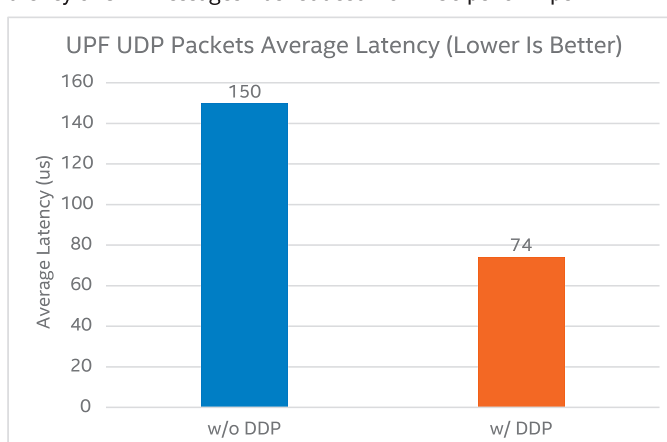

Figure 11. ZTE UPF UDP packets average latency

5.3. Analysis of the System's Total Cost of Ownership (TCO)

What Intel SST and DDP technologies bring to UPF is more than just performance improvement and latency reduction. The chart below estimates the device costs and ten-year electricity expenditures for three different solutions with different performance requirements. It shows that with the same requirements for processing capabilities, the UPF solutions based on Intel Xeon Gold 6230N processors and Intel Xeon Platinum 8280 processors, both of which are equipped with Intel SST and DDP technologies, are superior to the UPF solution based on Intel Xeon Gold 6138 processor without Intel SST and DDP technologies in terms of overall system cost. Over time, the TCO advantage of UPF solutions based on Intel SST and DDP technologies becomes increasingly obvious.

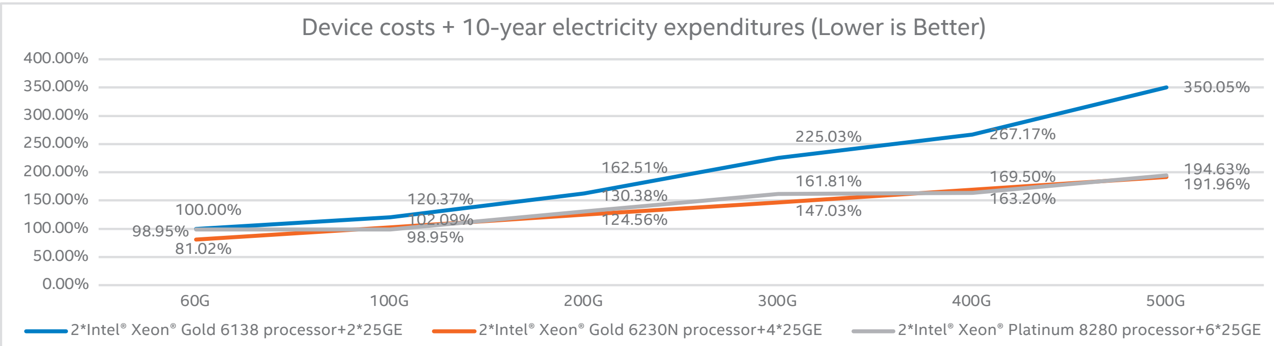

Figure 12. Device costs and ten years' electricity expenditures (For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. Refer to http://software.intel.com/en-us/articles/optimization-notice for more information regarding performance and optimization choices in Intel software products)

In addition, the number of servers needed to meet a telecom carrier's UPF element performance requirements also must be considered. The table below lists the minimal "round up" number of servers for different UPF element performance requirements. Obviously, the number of servers (excluding N+1 redudent servers and control plane servers) can be reduced if the system total performance is improved. The performance improvements seen in UPF solutions based on Intel SST and DDP technologies thus present another TCO advantage.

UPF element reqs (Gbps) / Server Performance (Gbps)	100	200	300
66 (Intel Xeon Gold 6138 CPU)	2 servers	4 servers	5 servers
132 (Intel Xeon Gold 6230N CPU)	1 server	2 servers	3 servers
201 (Intel Xeon Platinum 8280 CPU)	1 server	1 server	2 severs

6. Conclusions

ZTE's NFV UPF solution based on Intel platform and acceleration technologies has achieved over 100 Gbps forwarding capability in the telecom carriers' real-world traffic test models on the 2nd generation Intel Xeon Scalable processors, namely Intel Xeon Gold 6230N processors and Intel Xeon Platinum 8280 processors. The solution delivers the outstanding performance that telecom carriers are seeking for their 5G UPD deployments. Besides performance improvement, the solution greatly reduces the forwarding latency and fulfills the end-to-end low-latency requirements in 5G business scenarios[4]. Moreover, the solution is also highly advantageous in terms of the system's total cost of ownership.

This testing also demonstrated that without any additional hardware acceleration devices, it is possible to achieve extraordinary performance with an Intel® technology-based platform by using the latest general-purpose Intel processors, standard low-cost and low-power Intel Ethernet 700 Series network adapters, and related technologies such as the Data Plane Development Kit (DPDK). With these technologies together, the processing capabilities of VNFs can be improved, which helps the telecom carriers' and equipment manufacturers' in network function virtualization scenarios.

7. Abbreviations

eNB	eNodeB
FDIR	Fault Detection, Isolation, and Recovery
GPRS	General Packet Radio Service
GTP	GPRC Tunneling Protocol
IMIX	Internet Mix
IP	Internet Protocol
L2TP	Layer Two Tunneling Protocol
OEMs	Original Equipment Manufacturers
OMU	Operations Manager (Unix)
PDN	Packet Data Network
PFU	Packet Forwarding Unit
PPPOE	Point-to-Point Protocol Over Ethernet
RSS	Receive Side Scaling
SR-IOV	Single Root I/O Virtualization
TCP	Transmission Control Protocol
UE	User Equipment

[1] The Access Users number is the maximum number of users supported by the UPF system. The Data Users number is the actual number of users for the test.

[2] IMIX traffic.

[3] Testing conducted by ZTE on 2019-12-20. Configurations: Except for the processors used and SST/DDP technologies, the other test setup was the same as Intel Xeon Gold 6230N CPU and Intel Xeon Platinum 8280 CPU (see System Test Environment section for full configurations).

[4] See 3GPP 23.501 chapter 5.7.4