White Paper

Network Operators
Cloud Native 5G

ZTE uses Intel Container Bare Metal Reference Architecture (BMRA) to Improve Performance in Cloud Native 5G Core UPF Solution

ZTE is at the forefront of the transition to cloud native architecture in the 5G mobile core using Intel® technologies.

ZTE and Intel repeated performance tests less than a year apart using Intel® Xeon® Gold 6330N processor, the Intel® Xeon® Platinum 8380 processor, and Intel® Ethernet 800 network adapters. For the 2022 test, the companies incorporated the ZTE TECS Cloud Foundation* (TCF) platform, 5G core network UPF with CNFs, and Intel's Container Bare Metal Reference Architecture (BMRA) to accelerate the cloudification of the 5G mobile core and leverage the scale of cloud native.

ZTE's cloud native 5G core UPF achieved 389 Gbps with the Intel Xeon Gold 6330N processor and 656 Gbps with the Intel Xeon Platinum 8380 processor, representing year-over-year improvement of 34.5% and 42% respectively.

The test highlights the performance and scale benefits of cloud native architecture and applications.

Preface

After network operators successfully virtualized network architectures with networks function virtualization (NFV) deployments, many believe cloud-native SA core architecture is the next step towards monetizing 5G services, easing deployments, and achieving the operational efficiency, service agility, and scale needed for TCO savings. In fact, 90 percent of operators claim Cloud Native 5G SA Core readiness by 2023[1].

The industry will continue to evolve from virtual network functions (VNFs) to cloud native functions (CNFs) because of these advantages in performance, resource savings, and elasticity. To advance optimized cloud native solutions readiness with the latest Intel® platforms and open-source software innovation, Intel delivers the Network and Cloud Edge Reference System Architectures portfolio. These Reference System Architectures address 5G Core deployment scenarios and workloads by integrating and validating Intel's latest platform releases and open-source software innovation with Intel's best-known configurations and practices. As a result, forward-looking cloud native applications and implementations can be achieved quicker and with confidence.

This paper describes the ZTE TECS Cloud Foundation (TCF) bare metal reference architecture (BMRA) platform and 5G core network UPF products based on 3rd Gen Intel® Xeon® Scalable processors, Intel® Ethernet 800 Series Network Adapters with Dynamic Device Personalization (DDP), and the Intel Container BMRA. Test results show that the maximum overall performance can reach 389 Gbps on the Intel® Xeon® Gold 6330N processor and 656 Gbps on the Intel® Xeon® Platinum 8380 processor under an operator's standard traffic test model with 690 bytes packet length[2].

Table of Contents

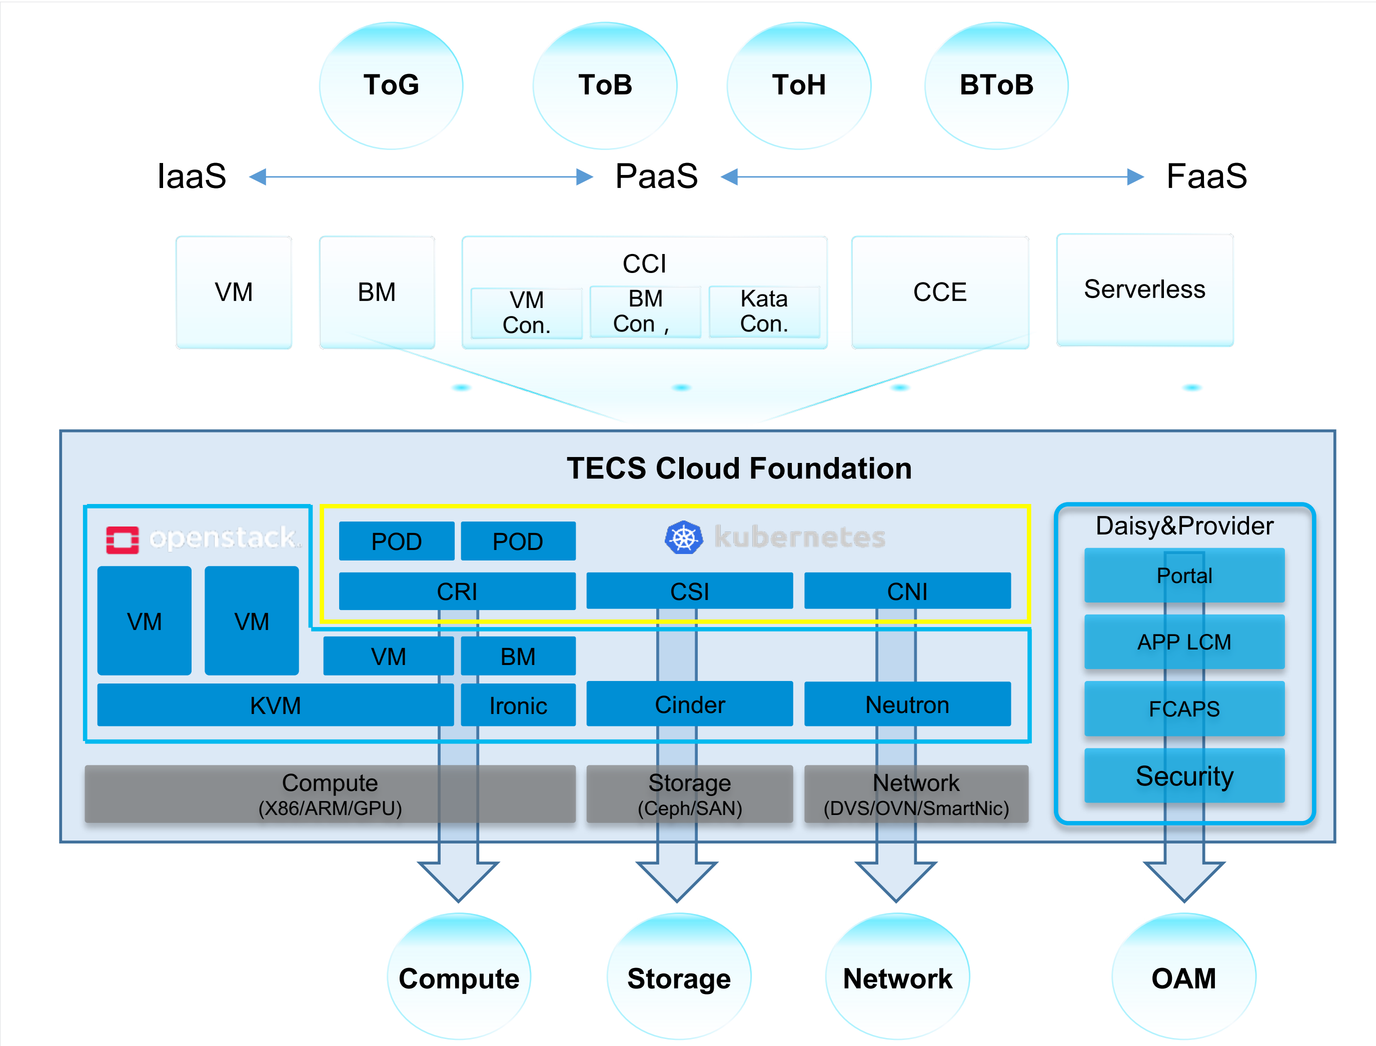

Figure 1. ZTE TCF Platform Architecture

1. ZTE Cloud Native UPF solution

1.1. ZTE TECS Cloud Foundation Introduction

ZTE TCF is an ICT-oriented, integrated cloud platform based on deep integration of the industry's two major open-source platforms, OpenStack* and Kubernetes*. The platform is widely compatible with various standardized, enhanced, and customized hardware devices, provides unified computing power (including virtual machines, containers, and bare metal), security and operation and maintenance management services, and a general cloud infrastructure for applications. The TECS Cloud Foundation* provides a microservice platform and general technical middleware services to enable agile development and rapid deployment of applications that address unique requirements of every industry.

TECS Cloud Foundation supports different deployment and operation forms according to needs and delivery scenarios. By adapting resource configuration, installation components, and related features, the platform interface provides basic cloud services, which can be flexibly expanded or contracted, and supports horizontal expansion and flexible switching of scenes.

TECS Cloud Foundation can provide a unified and consistent infrastructure, services, and operation, and maintenance whether it is deployed as a telecom cloud, an industry cloud, or an IT cloud. At the same time, TECS Cloud Foundation also supports the provisioning of technical services and cloud native technology stacks that run on third-party, cloud facilities.

1.2. ZTE 5G Cloud Native UPF Introduction

ZTE's 5G converged user-plane function is based on a fully virtualized cloud-native architecture and supports cross-data center deployment based on slicing requirements.

The cloud native, microservice architecture is highly automated from blueprint design and provides resource scheduling and lifecycle management, application status monitoring, control policy updates, and more. The effective connection between each link provides a closed-loop feedback mechanism that enables one-click deployment and installation, full autonomy, and efficient management of services.

Compared with virtualization, containers are more lightweight, spin-up faster, and require fewer IT resources to deploy and manage. Cloud-native applications and underlying virtualization technologies are decoupled and can be deployed in container technology to achieve improved resource utilization, rapid delivery, and agile maintenance of services. ZTE's 5G converged core network supports both virtual machine-based and container-based deployments with the following benefits:

- Optimization based on virtual machines and containers meets the requirements of 5G core network deployment resources.
- Performance optimization includes open-source virtual switches with open architecture and interfaces.
- Optimization improves user plane forwarding performance to support carrier-class network performance requirements.

2. Intel Container Bare Metal Reference Architecture

2.1. Intel Container BMRA Introduction

The Intel Container BMRA is an open, Kubernetes cluster architecture designed to support the convergence of key applications and services, control plane, and high-performance, packet processing functions using Intel technologies (starting with Early Access hardware platforms) and open-source platform software capabilities.

This solution template enables deployment of Kubernetes clusters based on multiple worker nodes, which can be managed by one or more Kubernetes control nodes. The user may configure and tune the features of Intel technologies and scale the number of nodes per cluster based on their defined variables. All servers relate to one or two switches that provide connectivity within the cluster and to the cloud. The Ansible playbook allows auto-provisioning and auto-configuration of the BMRA and can be installed on any connected host server.

The main elements forming the reference architecture are:

- **Hardware Components:** Provides multiple platform hardware options, including a variety of 2nd Generation Intel® Xeon® Scalable Processor SKUs, 3rd Generation Intel® Xeon® Scalable Processor SKUs, and Intel® Ethernet Network Adapters. Select and deploy the most optimal BIOS values before the cluster provisioning.
- **Software Capabilities:** Container environment uses a Docker containers runtime. The software capabilities are based on open-source software enabled in communities such as DPDK, FD.io. OVS, Kubernetes, and through Intel GitHub. Three Linux based operating system options are available: CentOS, Red Hat Enterprise Linux (RHEL), and Ubuntu.

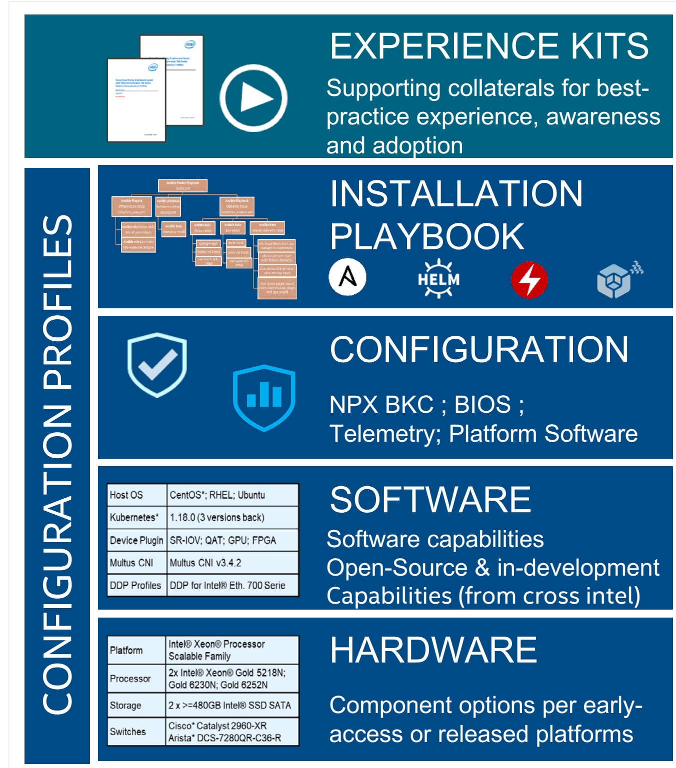

Figure 2. Intel Bare Metal Reference Architecture Overview

- **Configuration:** Specific hardware and software configurations based on Intel assessment and verification support two modes of operation: "base" and "plus" (for high performance).
- **Installation Playbook:** Details best-practice setup of hardware and software for each BMRA configuration, but Ansible playbooks are available to automate and shorten the setup time from day to a few hours.

2.2. Intel BMRA Accelerates ZTE TCF

Container BMRA is an extensible template solution that packages multiple Intel hardware solutions, optimizations, and cloud-native capabilities proven to accelerate and simplify the development and deployment of next-generation cloud-native services.

Containers are increasingly important in cloud computing and fundamental to cloud native adoption. Containers are lightweight, agile, and portable. They are easy to create, update, and remove. Kubernetes (K8s) is the leading open-source system for automating deployment, scaling, and management of containerized applications. ZTE integrates Intel BMRA key device plugins/CNIs into TCF platform to enhance Kubernetes for network functions virtualization (NFV) and networking usage for high performance cloud native 5GC UPF implementation.

Intel and ZTE advanced the discovery, scheduling, and isolation of server hardware features and developed capabilities and methodologies that expose Intel platform features for increased and deterministic application and network performance. Some of these platform features include:

- **Dynamic Device Personalization (DDP)** is one of the key technologies of the Intel® Ethernet 700 and 800 Series because it supports a broad range of traffic types by enabling workload-specific optimizations using the programmable packet processing pipeline.
- **Intel® Speed Select Technology – Core Power (Intel® SST-CP)** enables greater system performance because it allows the OS/VMM to control which cores can utilize any available power headroom across the system.
- **Multus** supports multiple network interfaces per pod, which enables the separation of control, management, and data planes – a key VNF requirement. Multiple network interfaces support different protocols, software stacks, as well as tuning and configuration requirements to further expand the networking capability of Kubernetes.
- **Topology manager** is a kubelet component that coordinates the set of components that are responsible for making topology- aligned resource allocations.
- **Node Feature Discovery (NFD)** enables generic hardware capability discovery in Kubernetes, including Intel® Advanced Vector Extensions 512 (Intel® AVX-512), Intel® Software Guard Extension (Intel® SGX), Intel® Dynamic Load Balancer (Intel® DLB), and others.
- **Device plugins**, such as Intel® QuickAssist Technology (Intel® QAT) and SR-IOV device plugins, boost performance and platform efficiency.
- **Native CPU Manager and Intel CPU Manager for Kubernetes** provide mechanisms for CPU core pinning and isolation of containerized workloads.
- **Telemetry Aware Scheduling (TAS)** makes telemetry data available to scheduling and de-scheduling decisions in Kubernetes.
- **Ethernet operator** orchestrates and manages the configuration of the capabilities exposed by the Intel® Ethernet E810 Series NICs. The operator is a state machine that configures certain functions, monitors the status, and acts autonomously based on user interaction.
- **Intel® QuickAssist Technology (Intel® QAT)** provides support for secure applications by directing the requested computation of cryptographic operations to the available hardware acceleration or instruction acceleration present on the platform.

BMRA offers comprehensive features and a repeatable deployment vehicle to accelerate cloud native 5G mobile core deployments.

3. System Test Environment

3.1. A Telecom Carrier's Traffic Test Model

Number of users[3]	Access User: 600000 Data User: 6000
Content billing configuration rules (DPI)	L7 Networking: 40000 L3 Networking: 10000
PCC strategy	Static: 45; Dynamic: 5
Volume ratio	HTTP: 85% UDP: 15%
Average packet length[4]	690 byte

3.2. Test topology

The Ixia IxNetworks-XGS2 was used in the test, conducted by ZTE on 2022-03-25, to simulate the control and data planes when 5G mobile users access them through the telecom carrier's base stations

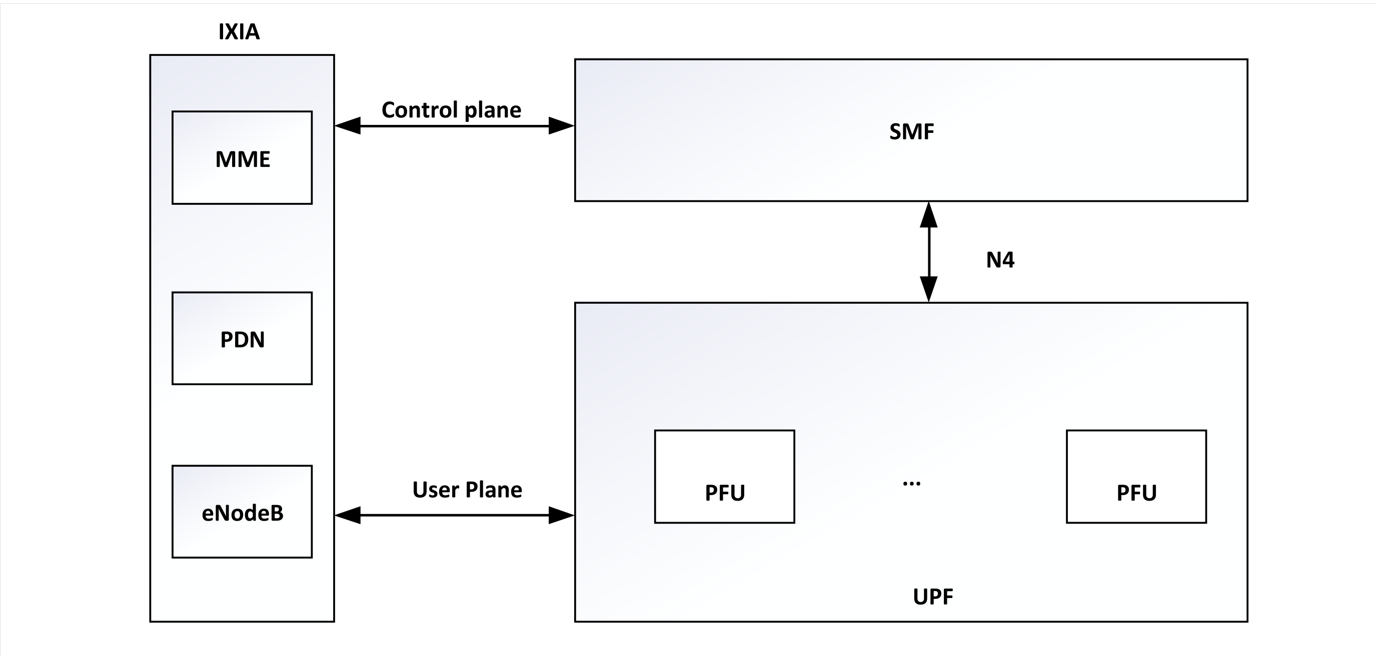

Figure 3. ZTE UPF test topology

3.3. Hardware Configuration

The servers used were ZTE's self-developed 5300-G4X, powered by 2-socket 3rd Gen Intel Xeon Scalable processors with each socket being connected to two or three Intel® Ethernet Network Adapters E810-CQDA2.

SERVER	ZTE 5300G4X
CPU	Intel® Xeon® Gold 6330N Intel® Xeon® Platinum 8380
uCode	0d000280
Number of CPU	2
MEMORY	512G DDR4@2933MHZ
NIC	Intel® E810-2CQDA2*4
TOR	ZTE 5960-4M

3.4. Software Configuration

OS	ZTE CGSL 4.18.0-193.14.2.el8.x86_64
TCF Platform	TECS OpenPalette V7.21.20.06P05
UPF CNF	ZXUN-xGW(GUL)V7.22.10
ICE Driver	Driver Version: 1.3.2 Firmware-version: 2.30 0x80005d1b 1.2877.0 DDP Profile:ice_comms-1.3.24
DPDK	20.11

3.5. Network Topology

The cloud management platform used in the test was ZTE TCF platform. Two servers with Intel processors 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 6330N processor-based server and the Intel® Xeon® Platinum 8380 processor-based server are shown in Figure 4 and Figure 5. Four PFs belonging to one socket were bonded together, and each PFU have one VF from each of PFs.

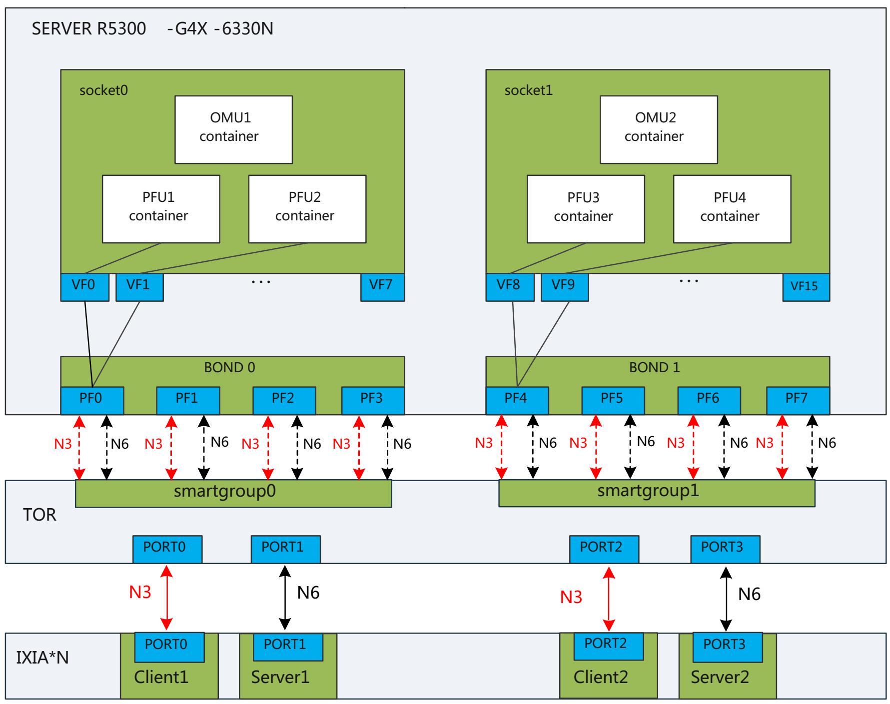

Figure 4. Test topology for Intel® Xeon® Gold 6330N processor server

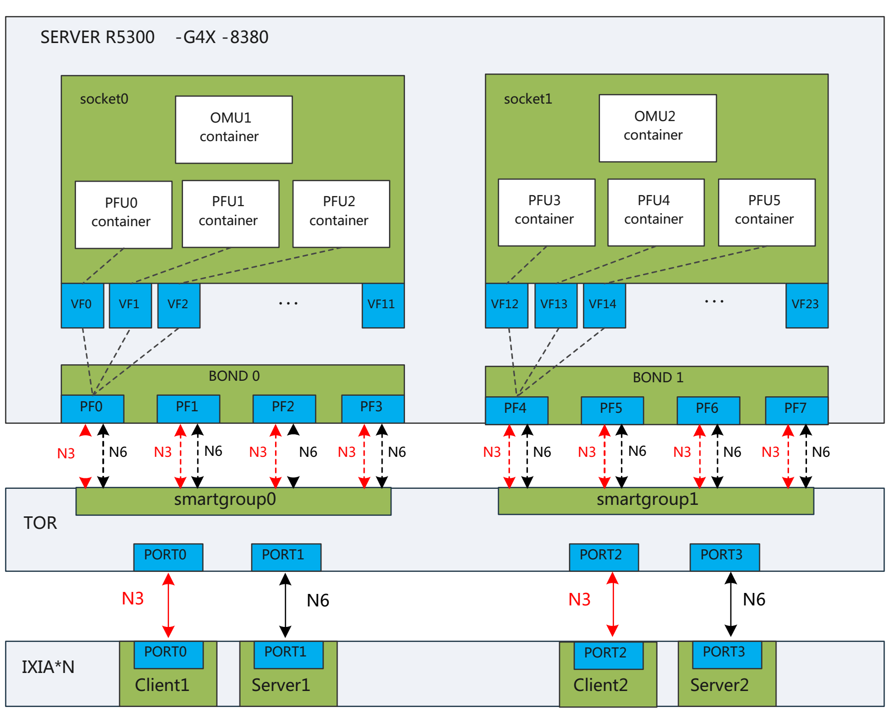

Figure 5. Test topology for Intel® Xeon® Platinum 8380 processor server

3.6. BIOS Configuration for a Compute Node

The BIOS configuration 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 TECH	ENABLED
		INTEL® VIRTUALIZATION TECHNOLOGY	ENABLED
POWER CONFIGURATION	ADVANCED -> POWER&PERFORMANCE	CPU POWER & PERFORMANCE POLICY	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

3.7. UPF Computing Resource Configuration

1. UPF cores mapping based on Intel Xeon Gold 6330N

One set of OMU containers and two PFU containers are deployed on each socket of the Intel Xeon Gold 6330N processor server. The OMU container occupies 4 vCPUs, and each PFU occupies 24 vCPUs (18 vCPUs run worker threads, 4 vCPUs run DPI threads, 1 vCPU run IPSEC threads, and 1 vCPU run control threads).

	OS		OMU		PFU1												PFU2											
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
6330N	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55

Figure 6. ZTE container UPF deployment in Intel® Xeon® Gold 6330N

2. UPF cores mapping based on Intel Xeon Platinum 8380

One set of OMU containers and three PFU containers are deployed on each socket of the Intel Xeon Platinum 8380 processor server. OMU occupies 4 vCPUs, and each PFU occupies 24 vCPUs (18 vCPUs run worker threads, 4 vCPUs run DPI threads, 1 vCPU run IPSEC threads, and 1 vCPU run control threads).

	OS		OMU		PFU1												PFU2												PFU3											
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39
8380	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79

Figure 7. ZTE container UPF deployment in Intel® Xeon® Platinum 8380

4. Performance Test Results

4.1. Overall Forwarding Performance

Basic forwarding performance of the UPF solution based on dual Intel Xeon Gold 6330N processors:

- When offline billing and DPI were disabled, the forwarding performance reached 389 Gbps. Worker forwarding threads' average core utilization was 84%.
- When offline billing and DPI were enabled, the forwarding performance reached 246 Gbps Worker forwarding threads' average core utilization was 85%.

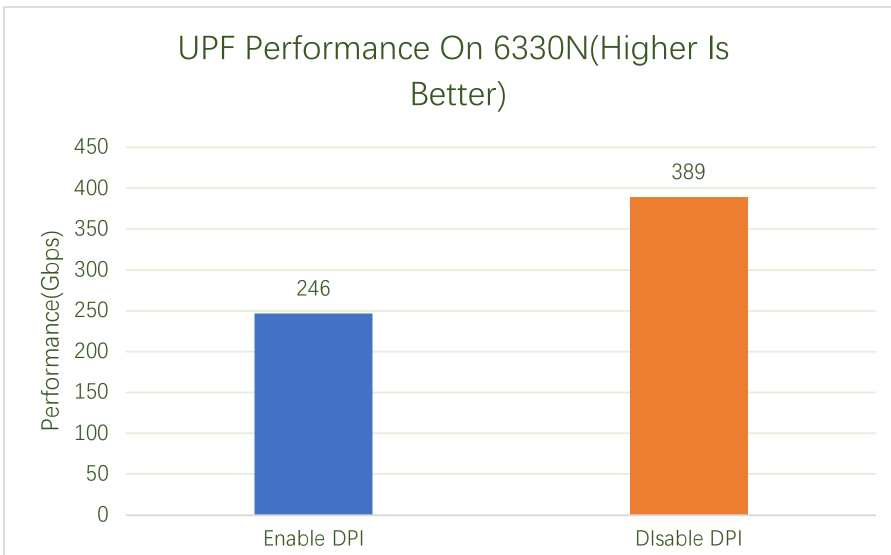

Figure 8. ZTE UPF performance result on Intel® Xeon® Gold 6330N processor

Basic overall forwarding performance of the UPF solution based on dual Intel® Xeon® Platinum 8380 processors:

- When offline billing and DPI were disabled, the forwarding performance reached 656 Gbps. Worker forwarding threads' average core utilization was 83%.
- When offline billing and DPI were enabled, the forwarding performance reached 395 Gbps. Worker forwarding threads' average core utilization was 84%.

Compared with Intel Xeon Gold 6230N and Intel Platinum 8280 processor, the forwarding performance of the UPF solutions are largely improved in the test.

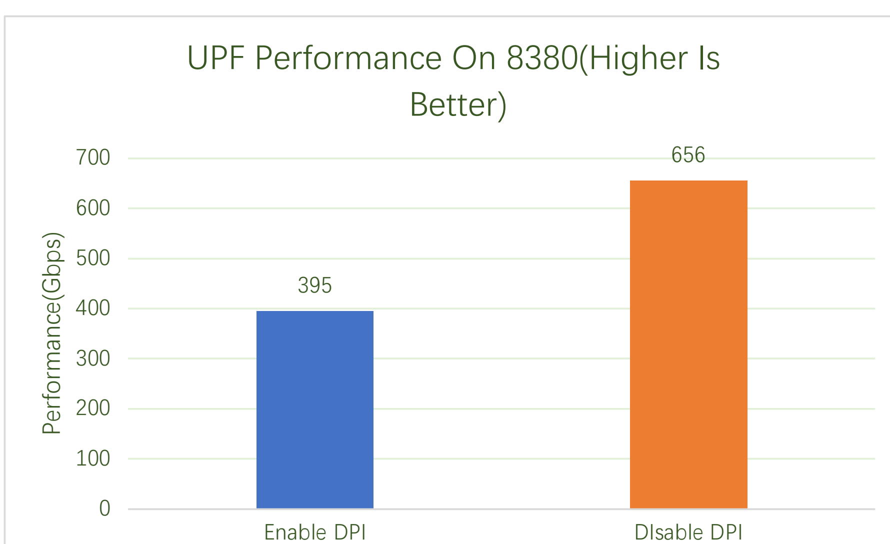

Figure 9. ZTE UPF performance result on Intel® Xeon® Platinum 8380 processor

4.2. VNF/CNF UPF Performance Comparison

Compared with the VNF UPF on the TECS OpenStack platform (Referring to white paper), the forwarding performance of the CNF UPF based on the TCF Cloud Native platform was greatly improved.

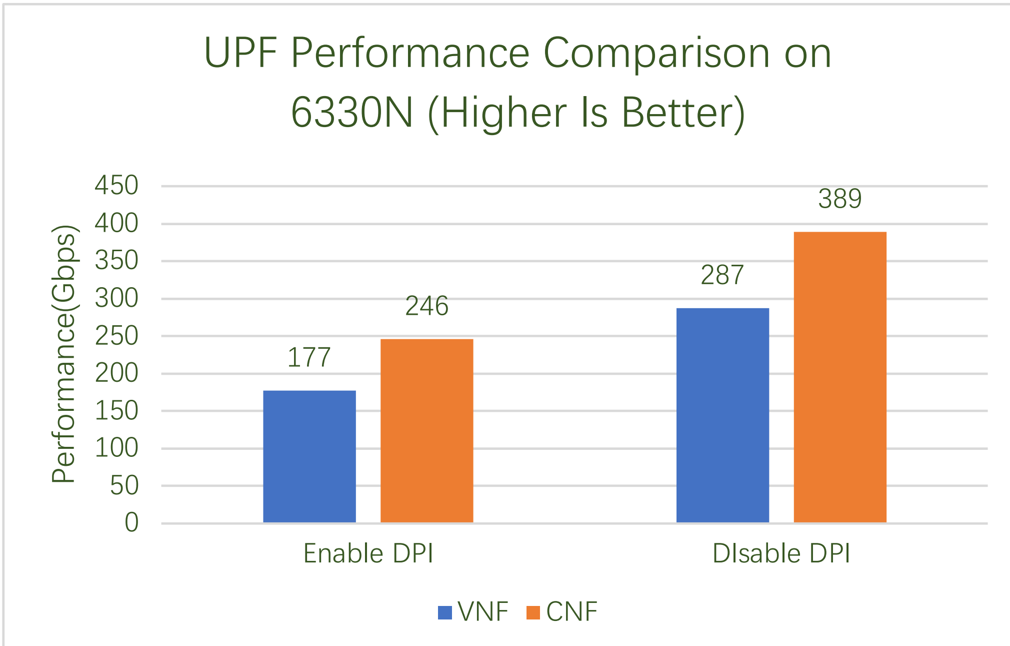

Figure 10. ZTE VNF/CNF UPF performance comparison on Intel Xeon Gold 6330N

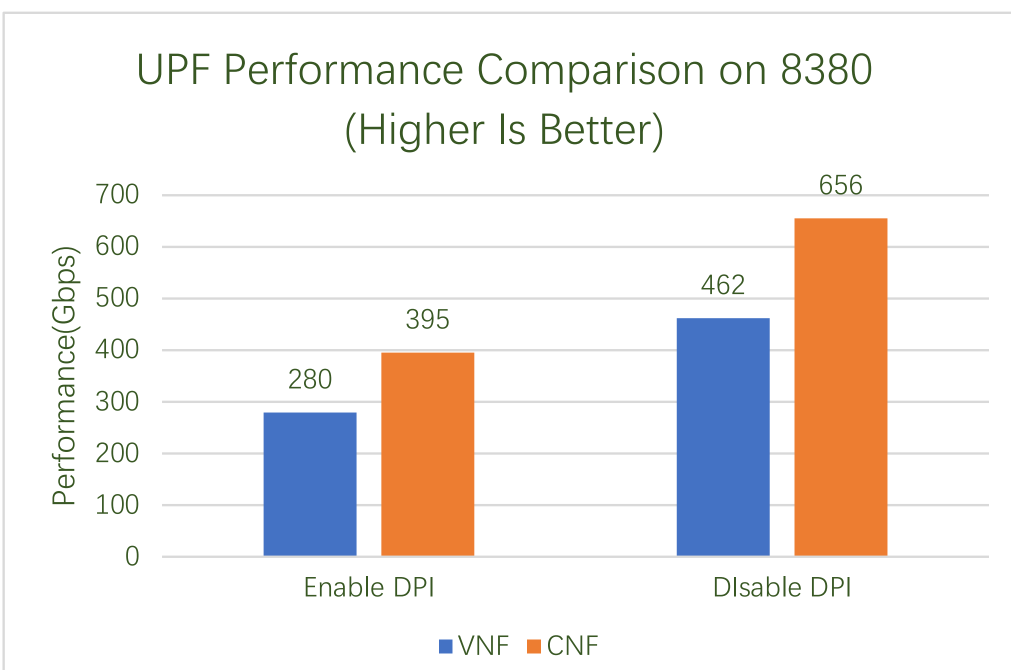

Figure 11. ZTE VNF/CNF UPF performance comparison on Intel Xeon Platinum 8380

4.3. CNF UPF Performance lift Analysis

Overall forwarding performance of ZTE's cloud-native 5GC solution CNF UPF improved by about 40% compared to the previous VNF solution. The newly designed CNF UPF software architecture makes full use of software offloading, cache prefetching, and other mechanisms that reduces unnecessary packet processing and CPU idle waiting, and performance continues to improve.

Below is an analysis of the factors that influenced performance:

- **TCF BMRA Enhanced Platform: 5% CNF UPF Improvement**

  Using the TCF BMRA enhanced platform as the underlying platform reduces the use of a virtual machine as the middle layer, which is case with the traditional OpenStack virtualization platform. UPF business units are carried in the form of micro-services and lightweight containers, which significantly reduces the management overhead of the system, fully releases the overall performance of the server, and brings about a 5% performance improvement for CNF UPF.

- **More stages of Flow table cache prefetch: 15% worker thread performance improvement**

  Cache prefetching is widely used in UPF code design to reduce cache-miss. Conventional cache prefetching applications mainly target cyclic operations of the same type of data, and prefetching the next batch of data. The CPU load of the database table lookup interface in the offloading process is too high based on our analysis of UPF CNF performance. Combined with the low hash conflict rate in the table lookup algorithm, a fast table lookup interface is added to prefetch the first record in the hash bucket. If the verification fails, the original table lookup interface is used for ordinary table lookup. After prefetch optimization, worker thread performance improved by nearly 15%.

- **More Offload Enhancement Design: Improve worker thread performance ~7%**

  Software offloading reduces unnecessary processing flow to realize fast forwarding of packets. In terms of software design, UPF continues to enrich the types of packets that can be offloaded and improve the processing efficiency of the offloading operation. The packets of the raw channel, as the inter-transfer packets inside the UPF, have a different processing flow from the outgoing packets, so a separate offload processing is designed for them. By designing the offload process for more types of packets, the performance of worker threads is improved by nearly 7%.

- **More Forwarding Cores: Another 12% performance improvement for CNF UPF**

  As the performance bottleneck of the system, the computing power of the worker thread has a decisive impact on the forwarding performance of the whole machine. The software architecture optimization enhances the performance of UPF from the perspective of single-thread performance. Increasing the number of worker threads can achieve horizontal expansion of the forwarding performance of the entire UPF. CNF UPF adjusts the ratio of different types of threads, making full use of the CPU cores released after the DDP enhancement. CNF UPF increases the number of worker threads from 16 to 18, an increase of 12.5%, which brings about a 12% performance improvement for CNF UPF.

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

Intel BMRA brings more than just performance improvement and latency reduction to UPF. Figure 12 estimates the device costs and ten-year electricity expenditures for three different solutions with different performance capabilities. Given the same processing requirements, the UPF solutions based on Intel Xeon Gold 6330N processors and Intel Xeon Platinum 8380 processors, both of which are equipped with Intel SST and DDP technologies, deliver substantially lower total cost of ownership. The UPF solution based on Intel Xeon Gold 6230N processor is not equipped with Intel SST and DDP technologies, resulting in a higher overall total cost of ownership. Over time, ZTE found that the TCO of Intel-based UPF solutions becomes even more advantageous.

Another TCO consideration for a telecom carrier's UPF element is the number of servers required for their performance requirements. The number of servers (excluding N+1 redundant servers and control plane servers) can be reduced if the total system performance is improved, as seen in UPF solutions based on Intel BMRA.

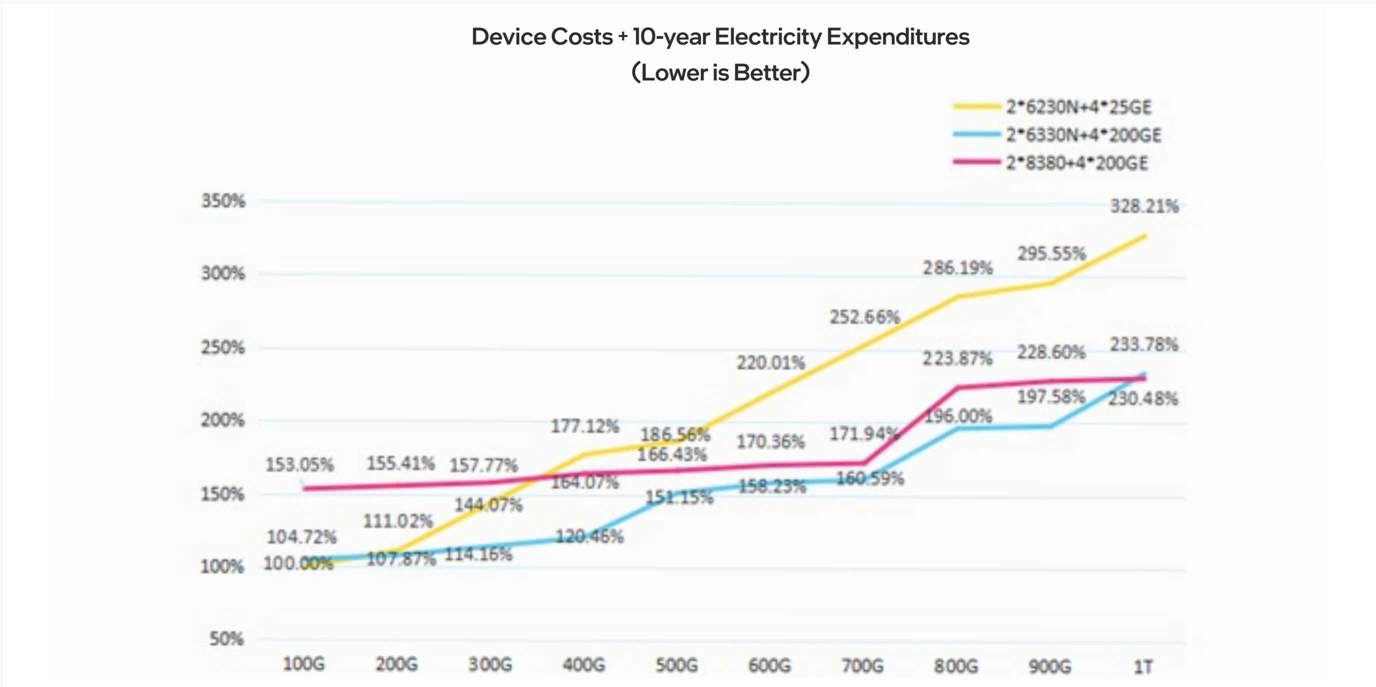

Figure 12. Device costs and ten years' electricity expenditures

The test with ZTE did not address energy consumption, which represents a significant percentage of network OPEX. The server configurations did not include available optimizations and features of Intel Xeon Scalable processors that support power management. In the future we will demonstrate how these capabilities reduce power consumption and costs without performance penalties and align with sustainability and Net Zero initiatives that may generate additional carrier revenue in the form of green credits.

	UPF element reqs (Gbps)		
Server performance (Gbps)	200	300	400
132 (Intel Xeon Gold 6230N CPU)	2 servers	3 servers	4 servers
246 (Intel Xeon Gold 6330N CPU)	1 server	2 servers	2 servers
395 (Intel Xeon Platinum 8380 CPU)	1 server	1 server	2 servers

5. Conclusions

The Intel Container BMRA is a cloud-native, Kubernetes-cluster template solution that enables operators to maximize underlying Intel architecture for network implementations. As a result, the ZTE TECS Cloud Foundation (TCF) platform provides high-performance packet processing for network applications, including a containerized, 5G core UPF solution that achieved over 650 Gbps forwarding capability using the 3rd Generation Intel Xeon Platinum processor.

The platform improves the processing capability of the same CPU and NIC through the implementation of containerized network functions (CNFs) and other hardware and software optimizations included in the Container BMRA. ZTE is delivering a powerful performance advantage to telecommunication operators and equipment manufacturers to meet cloud native 5G core network requirements. ZTE is at the leading edge of the transition to cloud native in partnership with Intel.

6. Abbreviations

BMRA	Container Bare Metal Reference Architecture
CNF	Containerized Network Function
FDIR	Flow Director
GPRS	General Packet Radio Service
GTP	GPRS Tunneling Protocol
IP	Internet Protocol
OEMs	Original Equipment Manufacturers
OMU	Operations Manager (Unit)
PDN	Packet Data Network
PFU	Packet Forwarding Unit
RSS	Receive Side Scaling
SR-IOV	Single Root I/O Virtualization
TECS	Telecom Elastic Cloud System
UE	User Equipment
VNF	Virtual Network Function

Authors

Intel

Yong Li

Jianwei Ma

Jack Ma

Bide Xu

Guoliang Gu

Xiaodong Ying

ZTE

Fangliang Lou

Weimin Li

Tao Wang

Jinchun Zou

Hongbo Cai

Chen Wu

Notices & Disclaimers

[1] "Cloud-Native 5G Core Operator Survey" Heavy Reading, March 2021

[2] This test was completed by ZTE on March 30th, 2022. See the System Test Environment section for specific test configurations

[3] The access user is the maximum number of users supported by the UPF system, and the traffic user is the actual number of users during the test.

[4] mixed average length

Performance varies by use, configuration and other factors. Learn more on the Performance Index site.

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 technologies may require enabled hardware, software or service activation.

© 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.

0922/TM/HO9/PDF ♻ Please Recycle 351850-001US