



WHITE PAPER

Optimize 5G Network Access with DPDK and XDP on uCPE Platform

www.advantech.com/nc

Authors:

Intel Corporation

Yang Zhang

Jiale Lu

Frank Liu

Xuekun Hu

Advantech

Lin Liu

Bo Liu

Fibocom

Zhentaο (James) Liang

Xi (Simon) Tao

Introduction

5G wireless networks introduced a lot of new features like high bandwidth, low latency, and network slicing to name a few. These enable solution providers to build innovative applications from the cloud to the edge and drive the evolution of connectivity technologies used by enterprise networks and WANs (Wide Area Networks). Modern network applications built on top of DPDK (Data Plane Development Kit) still rely on a Linux kernel based “slow path” to access to a 5G network, with the cost of sacrificing CPU cycles and performance. This white paper introduces a turnkey solution for uCPE Edge hardware platforms, 5G modules, and device drivers for DPDK development environments, to provide a “fast track” between DPDK applications and 5G networks.

Table of Contents

1. Advantech uCPE Solution	3
1.1. uCPE Introduction	3
1.2. uCPE White-Box.....	3
2. Fibocom 5G Client Solutions	6
2.1. Introduction.....	6
2.2. 5G uCPE use scenarios	7
2.2.1 SD-WAN and SASE	7
2.2.2 IOT monitor application for Government.	8
2.2.3 Financial business experience and security application	9
2.3. Device Driver Enhancement.....	9
3. Software stack for performance improvement	10
4. Validation environment	13
4.1. Network Topology	13
4.2. Packets transmission path optimization.....	13
4.3. Packets exchange optimization between Driver and Stack.....	13
4.4. Validation Platform Configuration.....	14
5. Conclusion.....	14

1. Advantech uCPE Solution

1.1. uCPE Introduction

Network and computing are fundamental capabilities for Edge Computing. uCPE (Universal CPE) devices are an on-prem edge platform combined with advanced networking technologies and computing power, benefitting from highly integrated Intel network platforms. uCPE devices can host Virtualized Network Functions (VNFs) such as SD-WAN, Firewall, DPI (Deep Packet Inspection) and enterprise applications in a single platform, which simplifies end to end connectivity between enterprise networks and the cloud. With a unified CPU architecture, compatible software environment, and multiple platform options, application developers can enjoy the experience of one-time development, run anywhere.

For over 30 years, the world's leading brands have chosen to embed Advantech computing platforms and IoT intelligent systems into their products, empowering Industry 4.0 and smart city applications and transforming network infrastructures. As an early innovator in new technologies such as AI, IoT and NFV, Advantech helps co-create new business ecosystems that enable an intelligent planet.

Advantech's uCPE designs provide a range of innovative platforms needed by service providers to transform the network using new disaggregated models. uCPE and NFV extend the cloud to the enterprise edge where technologies such as SD-WAN, IoT, and virtual RAN (vRAN)) enable a converged edge architecture securely connecting people and things. Over 100 dedicated engineers design our networking products to address new market needs following strict quality design rules and test criteria. Certifications and regional homologation services ensure products can be safely deployed globally. All this is backed by a solid financial base and an extensive networks of more than 8,000 employees globally. That is why we are a trusted uCPE and SD-WAN hardware partner that service providers and system integrators rely on across all major continents.

1.2. uCPE White-Box

uCPE platforms are “white-box” network devices, with standardized hardware, firmware and software interface. This dramatically saves on product development cycles and supply chain management.

Advantech offers a full range of high-quality uCPE products. Advantech's white box uCPE range is built on standard Intel® processors in feature-flexible appliances, that offer multiple configurations and price points scaling from 2 to 28 cores and maximum achievable throughputs up to 220 Gbps. Optional networking modules offer a highly flexible WAN connectivity choice of hybrid 5G, 4G LTE, WiFi6, WiFi5, xDSL & SFP+ configurations.

Encryption acceleration is supported using Intel® QuickAssist on Intel® Atom™ and Intel® Xeon® based platforms, with DPDK providing the technology needed to accelerate packet handling by up to 10x. As a result, secure branch connectivity including end-to-end encryption can be provided without compromising VNF performance or increasing cost. The 1U higher end platforms have been designed for high-availability networks with integrated failsafe redundancy and advanced remote security and management features that minimize system down time.

Advantech’s networking platforms have been specifically designed to run high-availability telecommunication services and minimize costly downtime. Advantech’s Advanced Platform Management provides IPMI v2.0 Baseboard Management Controller (BMC) functionality and also additional features that allow local and remote users to early detect system degradation, avoid system interruption, and shorten mean time to repair.

 <p>FWA-1212VC</p>	 <p>FWA-2112</p>	 <p>FWA-3050</p>
<p>Intel® Atom™ Processor C3000 Versatile platform for cloud-based /small branch deployment</p>	<p>Intel® Atom™ Processor C3000 Flexible platform for medium branch deployment</p>	<p>Intel® Xeon® Processor D-2100 Fully configurable to fit any enterprise deployment</p>

Advantech’s FWA-1212VC (as the test platform for DPDK accelerating 5G on uCPE) fanless edge network appliance offers innovative key features to meet the requirements of mainstream entry and mid-range software-defined WAN (SD-WAN) and universal Customer Premises Equipment (uCPE) solutions. FWA-1212VC, with multiple software defined LED indicators plus a software defined reset button, is an optimized performance/price uCPE ideal for enterprise communications service providers. Empowered with 5G and Wi-Fi6, it accelerates time-to-cloud access and helps enterprises to expedite the integration of cloud and network deployment. In addition, FWA-1212VC provides high-speed, reliable, and flexible wireless network connectivity for 5G SD-WAN solutions, offering a one-stop total solution for the digitalization of enterprises in the 5G era.



Figure 1. FWA-1212VC Rear View

Specifications

System PN		FWA-1212VCL-2CA1S	FWA-1212VC-2CA1S	FWA-1212VC-4CA1S	FWA-1212VCL-4CA1S
Form factor		Tabletop			
Processor System	Processor	Intel® Atom® C3336	Intel® Atom® C3336	Intel® Atom® C3558	Intel® Atom® C3436L
	Core Number	2C	2C	4C	4C
	Frequency	1.5GHz	1.5GHz	2.2GHz	1.3GHz
	L2 Cache	4MB	4MB	8MB	8MB
	BIOS	AMI EFI 128 Mbit			
Memory	Technology	DDR4 1866/2133/2400MHz			
	Max. Capacity	16GB			
	Socket	1 x 260pin SODIMM			
	ECC Option	Support both ECC or Non-ECC SODIMM Memory			
Networking	Controller	4 x Intel i211 for 1GbE RJ45 LAN ports	1 x Marvell 88E1543 for 2 combo ports		
			4 x Intel i211 for 1GbE RJ45 LAN ports		
	1GbE	4 x 1GbE RJ45 ports via Intel i211	2x 1GbE SFP/RJ45 combo ports (auto-negotiation) via Marvell 88E1543 PHY		
			2x 1GbE SFP/RJ45 combo ports (auto-negotiation) via Marvell 88E1543 PHY		
Expansion	Mini-PCle	1 x Full-size Mini-PCle with SIM slot for Mini-PCIE LTE, Wi-Fi module or encryption card;			
	m.2	1 x m.2 2230 E-key for Wi-Fi module			
		1 x m.2 3052 B-key for 5G module (compatible with m.2 3042 B-key for LTE/4G module)			
Storage	eMMC	eMMC 16GB default onboard			
	mSATA	1 x mSATA slot(share the same connector with Mini-PCle)			
	2.5" HDD bay	1 x 2.5" SSD bay(optional)			
I/O	USB	2 x USB2.0 Type A on rear			
	Serial	1 x Console port (RJ45 connector)			
	LED Indicator	Rear IO: Power, System, Cloud			
		Front IO: HDD, Power, System, Cloud, Wi-Fi, LTE,LTE/5G and LTE RF Signal LED			
	Others	1 x DC Jack with lock on rear			
		1 x Power Button with LED indicator			
		1 x Software definable button			
TPM	TPM	LPC connector reserved for TPM module option			

Power	Power Type	Adapter
	Watts	36W
	Input	100~240V AC
	Connector	DC-Jack with lock
	Power Adapter	12V 3A
Environment	Operating Temperature	0 ~ 40° C (32 ~ 104° F)
	Operating Temperature	-20 ~ 60 °C (-40 ~ 140 °F) and 40° C @5~95% RH Non-Condensing
Mechanical	Construction	Iron
	Dimensions (W x H x D)	270 x 44 x 168 mm , 8.1" x 1.32 " x 5.04"
	Weight	about 3kg
OS Support		Linux (CentOS7.6, Ubuntu 18.4, etc.)
Advantech Software Packages	QuickStart Linux Image (CentOS based reference BSP)	Afru, Sensor-reader, Wi-Fi driver, 5G driver, Intel DPDK, Intel QAT,
Certification/Safety		CCC/CE/FCC/UL/CB/VCCI/RCM (RF certification is not included)

2. Fibocom 5G Client Solutions

2.1. Introduction

This 5G uCPE solution integrated Fibocom's first global 5G module FM150-AE. It supports both 5G SA and NSA network architectures with faster transmission speed, better carrying capacity, and lower network latency.

FM150-AE supports the 5G NR Sub6 band, and is compatible with LTE and WCDMA standards, eliminating customers' investment concerns in the initial stage of 5G construction. These modules provide a perfect high-speed experience for customers using a uCPE solution.

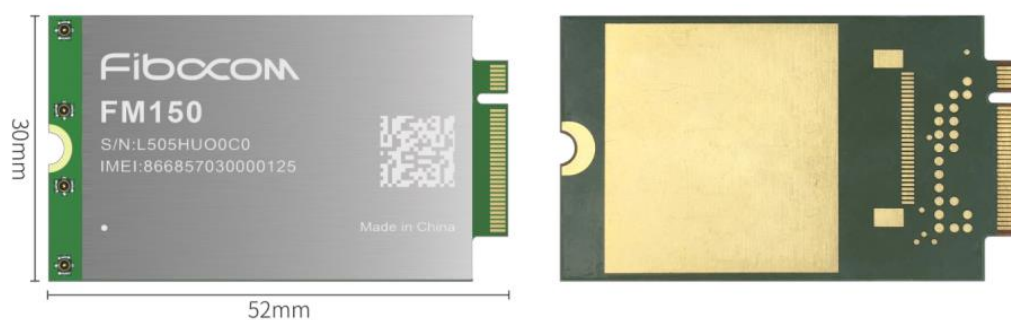


Figure 2. FM150-AE

Basic Parameters

Factor	M.2; 30 x 52 x 2.3mm
Weight	~ 8.2g
Operating Voltage	3.135V~4.4V, Typical 3.8V
Operating Temperature	-30 ~ +75°C
Extend Temperature	-40 ~ +85°C
DL	SA Peak 2.1Gbps
UL	Peak 900Mbps
AT Command Set	3GPP TS 27.007 and 27.005, proprietary FIBOCOM AT commands
Antenna	4 Antennas
Support FOAT/DFOTA/VoLTE/Audio/eSIM (optional)	

FM150-AE provides USIM*2, USB3.1/3.0*1 and PCIE3.0*1 interface in standard M.2 factor for the network communication, which means it's easy to be integrated into uCPE solution and be able to achieve high speed performance with accelerate technology.

Besides, FM150 uses 4*Antennas to adapt the global 5G network. This streamlined antenna design can extremely help optimize customer's whole machine design complexity and product size. Upstream is enhanced to 2x2MIMO Under SA condition, and 5G NR/LTE downstream is designed with 4x4MIMO.

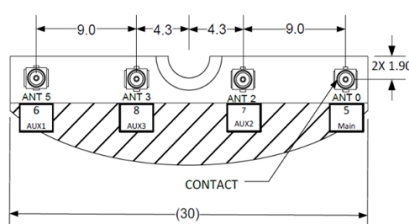


Figure 3. Antenna design

2.2. 5G uCPE Use Scenarios

2.2.1 SD-WAN and SASE

uCPE are general purpose “white box” networking edge devices with advanced network capabilities and application services. One typical workload deployed on a uCPE device is SDWAN

(Software Defined WAN), which facilitates the connectivity between cloud and enterprise branches. SD-WAN is widely used to improve Enterprise Network operation efficiency and Cloud Service experience.

SASE (Security Access Service Edge) is the new trend in the network security market. SASE is highly integrated with SDWAN to provide cloud and edge coordinated security services to enterprises.

Integrating 5G modules into uCPE empowers SD-WAN to bring secure, reliable, and wireless broadband to the enterprise network on demand. It also enables SD-WAN to establish measurable data transmission paths between edge and cloud, as well as improve application QoS (Quality of Experience) by utilize new 5G network features such as low latency and network slicing.

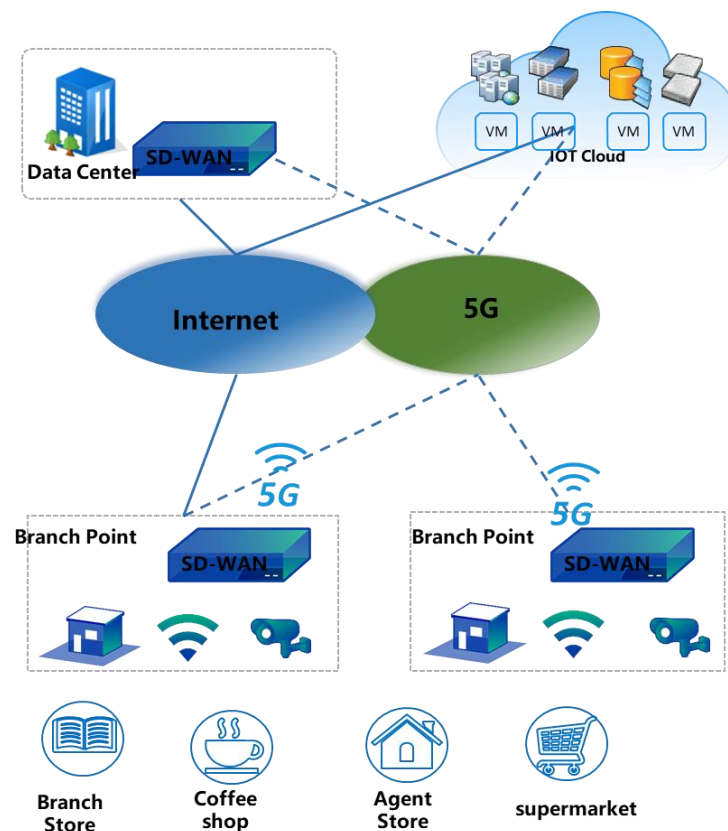


Figure 4. SD-WAN Structure

2.2.2 IOT Monitor Application for Government.

5G on uCPE devices can be applied to IOT monitoring projects. This can help headquarters remotely perform global operations and management.

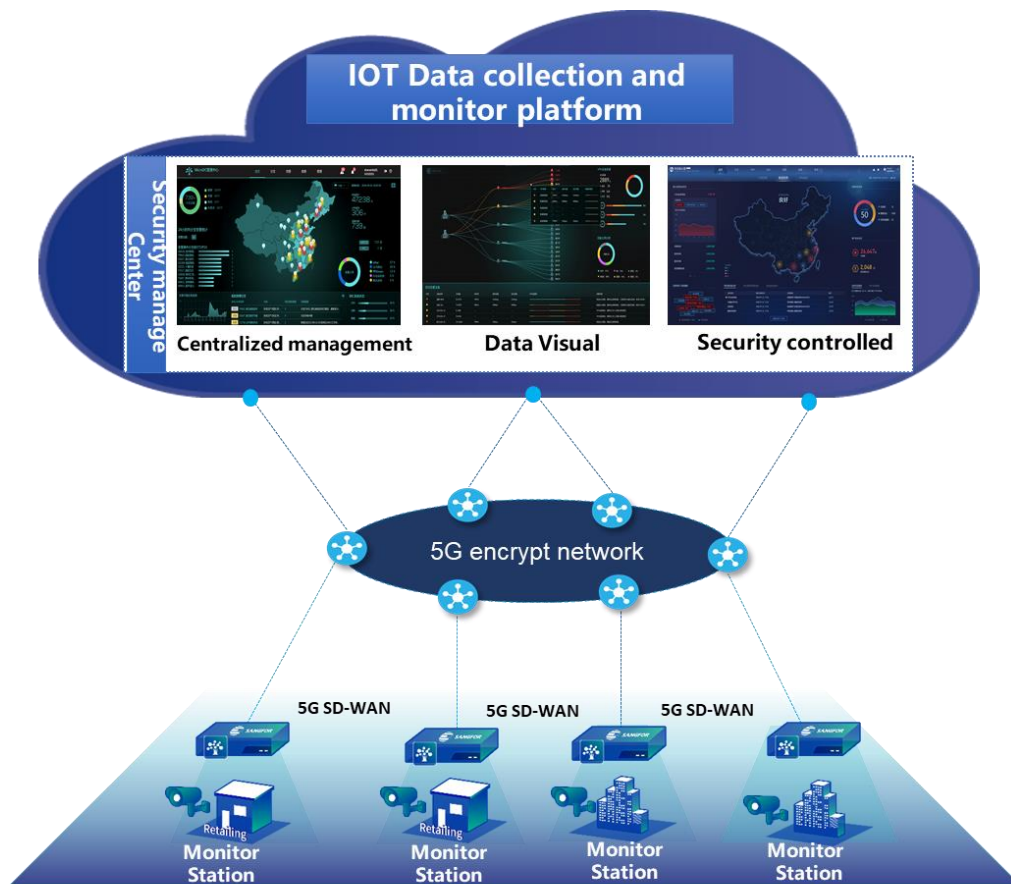


Figure 5. IoT Data Collection & Monitor Platform

2.2.3 Financial Business Experience and Security Application

Based on the ability of SD-WAN application identification, the transmission of services can be identified and classified, and then the intelligent path selection is carried out based on traffic categories.

For example, Internet traffic takes the “network slice for internet”, audio and video traffic goes to the “slice for low latency bandwidth” to ensure real-time data transmission, and the “production business” goes to the slice for security to ensure that the traffic is protected.

2.3. Device Driver Enhancement

DPDK packages have implemented the AF_XDP PMD driver to directly send packets to user-space applications. Base on this interface, 5G modem drivers support both generic XDP and native XDP modes. In generic XDP mode, packet will pass through kernel stack, softirq, or a poll mechanism to the application. In native XDP mode, the packet doesn’t go into the kernel stack, instead, it is handled by the XDP socket, and a copy (or zero copy) goes to the receive ring of AF_XDP and is detected by application.

For each packet, the decision of whether use the generic XDP or native XDP mode is made by the eBPF action resulting in: A XDP_PASS or XDP_REDIRECT return value. The first will choose the

generic XDP mode, otherwise, native XDP mode is selected. However, native XDP mode is preferred with its better performance.

3. Software Stack for Performance Improvement

Data Plane Development Kit (DPDK) is a widely used open source software library in modern network applications, mainly for network devices. It is developed by intel and dramatically increases the packet throughput by optimizing the software framework using many CPU features, such as huge-page, SIMD instruction, software prefetch, DDIO, etc.

Vector Packet Processing (VPP) is a user-space network stack application developed by Cisco. It is built on a packet processing graph concept, which lowers the instruction cache miss and increases the instruction-level parallelism. Integrated with DPDK, VPP also benefits from the optimization features of DPDK, which increases the IO rates, and thus provides high performance packet processing capability.

However, many traditional applications are based on the socket framework. The socket receives/transmits packets from/to the kernel space. For applications running on the user space, this path contains at least one memory copy operation, one interrupt handle routine, and a pass-through kernel network stack, causing negative impact on both throughput and latency for user-space applications. To overcome these problems, and be compatible with traditional applications, many software frameworks have been developed, including eBPF, XDP and AF_XDP.

eBPF is the expansion of the Berkeley Packet Filter (BPF), which acts as a virtual machine on the abstract layer within the Linux kernel, allowing developers to trace and even modify each packet. eXpress Data Path (XDP) is a framework which inserts the eBPF application at an early point, to process received packets as soon as possible before moving the kernel forward to the network stack, resulting in an improvement of both latency and throughput. The insert point of the XDP program can be divided into:

- Generic: after driver receives a packet and before entering the network kernel stack.
- Native: in the kernel driver.
- Offload: in the NIC hardware.

XDP performs four operations after analyzing a received packet, including drop packet, direct forward to transmit port, redirect to kernel network stack, and bypass kernel stack. The main procedure of XDP framework is as follows:

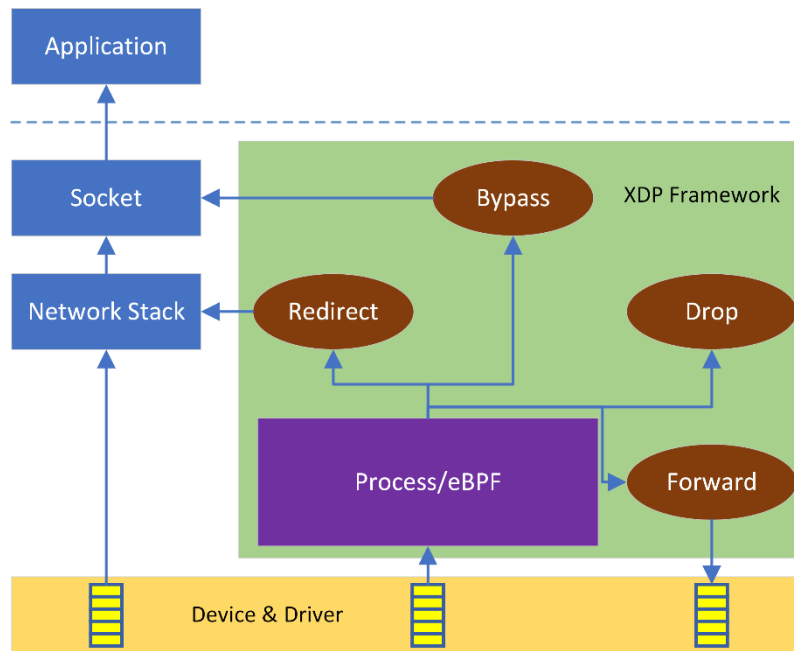


Figure 6. XDP Framework Procedure

AF_XDP is a new address family in the Linux kernel, the position of AF_XDP is shown as below:

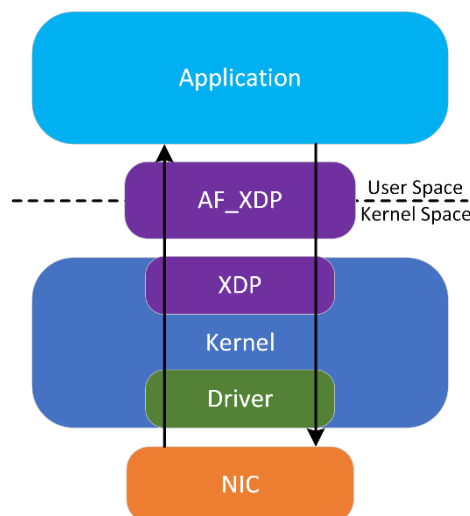


Figure 7. AF_XDP Framework Procedure

Similar to the bypass path in one of the XDP operations, in AF_XDP, the packets can bypass the network stack. One of the implementations is that, the packet buffer is allocated in user-space and mapped to the kernel-space, thus enabling zero-copy movement of packet data. Moreover, a poll operation instead of interrupt handling can be used for high-throughput scenarios. With all these optimizations, AF_XDP can easily achieve nearly closed performance as in DPDK.

AF_XDP has many implementations in the user-space, this section takes DPDK as a reference to introduce the framework of AF_XDP, also assume the NIC driver supports XDP. AF_XDP socket (XSK) is based on the AF_XDP family. DPDK has added the AF_XDP poll mode driver (PMD), which acts as a

device and interface to DPDK device interfaces. Compared with a normal DPDK application, while a PMD driver in user-space directly receive/transmit packets from/to rings which are mapped by a general kernel driver, such as vfio-pci. The AF_XDP PMD directly communicates with the NIC kernel driver.

After the initialization of XSK, four rings will be allocated: RX, TX, FILL, and COMPLETION rings. Among them, receive ring and fill ring is used for receiving directions, while the other two rings are used for transmitting directions.

The notification between kernel space and user space is done by moving the producer and consumer point of each ring, for example, kernel space can move forward the producer point to indicate there exist available data on the ring, then the user-space app can consume this data and adjust the consumer point to align the points. The FILL and COMPLETION rings are managed by a so-called UMEM, which is a buffer management. The structure of these rings is achieved through XSK in the kernel space by using getsockopt syscall. This information is used to allocate the memory for these structures. By using mmap syscall, the virtual memory of these rings is returned to the user-space. The framework is shown as below:

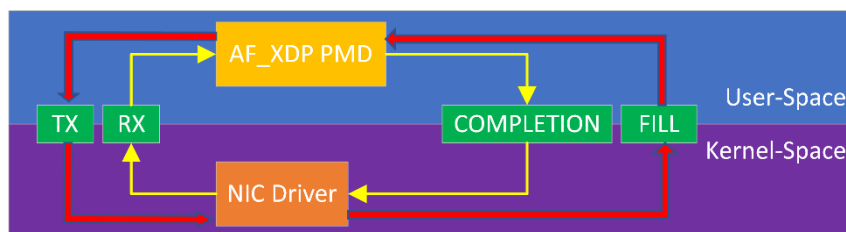


Figure 8. The Notification Framework between Kernel Space and User Space

When the NIC driver receives packets, these packets will pass the pre-loaded eBPF filter, assuming a XDP_REDIRECT action will be returned, resulting in redirection of packets to XSK. The main job of XSK is to fill these packet addresses to receive a ring in DPDK, and move the producer point forward. The DPDK receive functions checks the distance between consumer and producer to figure out if there are packets waiting to be received, then move the consumer point forward and receive these packets. Meanwhile, DPDK will fill the FILL ring with a new allocated buffer address to ensure enough buffer addresses in the kernel driver. The NIC driver can now fill these addresses to its own receive descriptor.

The transmit side is similar to the receive sides, except that the RX/FILL ring is replaced with the TX/COMPLETION rings. DPDK first acquires the un-used address from the COMPLETION ring, moves the consumer point forward. Then fills the TX ring with data addresses, and moves the producer point forward. An NIC driver acts as a consumer of the TX ring, getting the data address then filling them to the transmit descriptor, and moving the consumer point of the TX ring. After these packets are transmitted outside successfully, the producer point of the COMPLETION ring is moved forward.

Note that, AF_XDP implements many features such as zero-copy, polling mode, wakeup mechanism, etc. However, the basic framework of this AF_XDP has not changed. Taking non-zero-copy scenario as an example, there is a little difference. Right after the DPDK receives the packets, it will copy these packets to a new allocated address, the old address is marked in a buffer ring, then, at the right timing, DPDK will fill the FILL ring with the address of this buffer ring.

4. Validation environment

4.1. Network Topology

From the view of 5G modem, one server with 10G NIC plug-in acted as the transmitter, who sends packets with DPDK software to the Keysight 5G tester, which forwards the packets based on the 5G protocol and receives data via the 5G modem located in the receiver. The transmit direction is similar, the client send the packets through 5G modem, received by 5G tester, then forward to the server NIC, handled by DPDK. The whole diagram shown as:

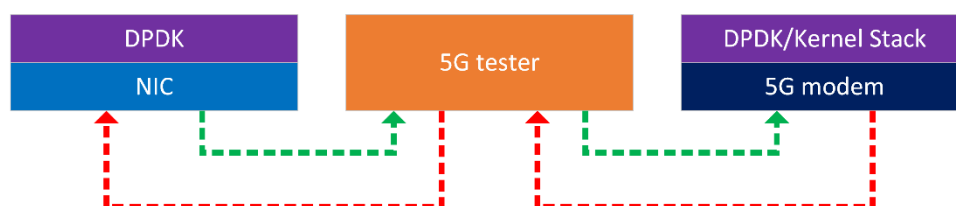


Figure 9. The 5G Tester Network Topology

4.2. Packets Transmission Path Optimization

For DPDK AF_XDP receive side, when the 5G modem receive packets, the eBPF program is executed to determine the action. For the XDP_REDIRECT result, the packets are directly received by the DPDK program through the UMEM framework. For comparison, the receive side is also replaced with the kernel driver. After the packets are received by the 5G modem driver, it directly forwards these packets to the kernel stack received by a normal socket framework. All these paths are reversed for transmit direction.

It will show that even the zero-copy or other optimization featured is not used, the performance will also benefit from the short receive/transmit path when bypass the kernel stack implemented by the DPDK AF_XDP.

4.3. Packets Exchange Optimization between Driver and Stack

DPDK and 5G modem driver can support zero-copy mechanism to transfer packets. When this feature is enabled, the memory occupied by packets in the driver are also seen by the DPDK program, which directly operates these packets on the user-space. However, if zero-copy is disabled, there exists a memory copy operation when the packets transfer between kernel and user space. The simulation result also shows this feature will bring huge performance increase.

4.4. Validation Platform Configuration

Component	Description
uCPE	Advantech's FWA-1212VC with Intel Atom C3558, 4GB DDR4
5G Module	Fibocom FM150-AE
OS	Ubuntu
DPDK	21.11
Kernel	5.7
5G Tester	KeySight

5. Conclusion

The 5G network coverage is growing rapidly and brings a lot of new possibilities and network experience to edge computing and IoT. Simply converting the 5G network to Ethernet Network will not release the power of 5G for edge computing. At the heart of this solution is optimizing the transmission path of packets between the Intel processor and the 5G modem from the hardware platform level using new software technologies. The validated approach can eliminate the performance barriers between 5G carrier network, edge applications and the DPDK ecosystem. The 5G uCPE can bring differentiated 5G experiences and new value to service providers and ecosystem developers. 5G uCPE will facilitate new edge application innovations.

For more information, please visit our website ucpe.tech,
or contact Advantech representatives: cloud.iot@advantech.com

About Advantech

Advantech has the corporate vision to "Enable an Intelligent Planet". The company is a global leader in the fields of IoT intelligent systems and embedded platforms. To support the trends of IoT, big data, and artificial intelligence, Advantech promotes IoT hardware and software solutions with the industrial IoT platform - WISE-PaaS core to assist business partners and clients in connecting their industrial chains. Advantech is also working with business partners to co-create business ecosystems that accelerate the goal of industrial intelligence. www.advantech.com

About Advantech Cloud-IoT Group

Advantech Cloud-IoT Group helps service, solution providers and enterprises extend the cloud to the IoT edge by providing scalable platforms that reliably execute business and mission critical services anywhere in the network. Our innovative servers and appliances enable a seamless transformation toward cloud native operations and are supported by a vibrant software ecosystem and supply chain, offering key foundational building blocks that accelerate digital transformation. From SD-WAN, SASE and Private Wireless to Open RAN, Edge and Hybrid Cloud, Advantech Cloud-IoT Group is enabling the co-creation of products and services that will form the backbone of the new 5G & AIoT economy.

Notices & Disclaimers

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

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 does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© 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.



WHITE PAPER

Optimize 5G Network Access with DPDK and XDP on uCPE Platform

www.advantech.com/nc

intel
partner
Networking
Specialist

intel.
partner
Titanium

ADVANTECH