






DPDK SUMMIT CHINA 2017



主办方：

参与方： 腾讯云  **ZTE**  美团云


 **Panabit**

 **太一星辰**
Balance Your Networks

 **UNITEDSTACK** 联合云

 **云杉网络**
Yunshan Networks

协办方： **SDNLAB**
专注网络创新技术


视频支持方：

Accelerating the FD.IO/VPP Crypto Workload with the DPDK Cryptodev Framework

FAN ZHANG, PH.D

• NETWORK PLATFORM GROUP, DATA CENTER GROUP
ROY.FAN.ZHANG@INTEL.COM



主办方： 

参与方： 

ZTE

 美团云


 Panabit

 太一星晨
Balance Your Networks



 云杉网络
Yunshan Networks

协办方：  SDNLAB
垂直网络创新技术

视频支持方： 



LEGAL DISCLAIMER

- No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.
- Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.
- This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.
- The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.
- Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm>
- Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.
- *Other names and brands may be claimed as the property of others.
- Copyright © 2017, Intel Corporation. All rights reserved.
- Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804
- Mileage may vary Disclaimer: Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks Test and System Configurations: Estimates are based on internal Intel analysis using at least Data Plane Development Kit IPsec sample application on Intel(R) Xeon(R) CPU E5-2695 v4@ 2.10GHz with at least using Intel(R) Communications Chipset(s) 8955 with Intel(R) QuickAssist Technology.



Agenda

- ▶ Problem Statement
- ▶ DPDK Cryptodev Framework Introduction
- ▶ Enable DPDK Cryptodev Framework in VPP
- ▶ Performance
- ▶ Future work
- ▶ Conclusion





Agenda

- ▶ Problem Statement
- ▶ DPDK Cryptodev Framework Introduction
- ▶ Enable DPDK Cryptodev Framework in VPP
- ▶ Performance
- ▶ Future work
- ▶ Conclusion





Think about security at every step of the process: architecture, implementation, testing, documentation, distribution and deployment

- Dr. Nicko van Someren, CTO, Linux Foundation

With VPP, a single core can do 40G, 100G or even higher throughput L2 forwarding


But what is the throughput after adding security protection?





Let's take IPSec as an example

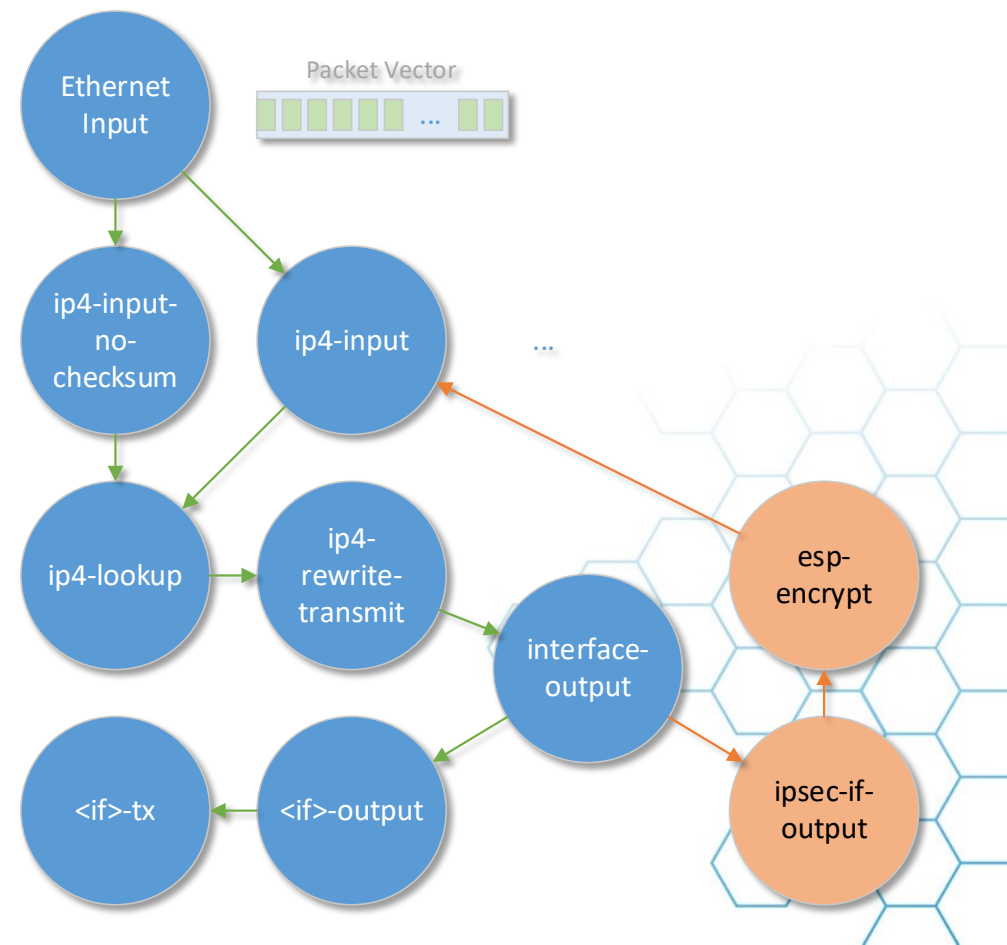
- ▶ > 20 years old but is still extremely popular
- ▶ Playing the role of security guardian in many network applications
- ▶ Requires lots of computations including crypto
- ▶ When traffic rate is high, efficient crypto implementation becomes necessary






FD.io / VPP IPsec

- Supports IPv4/IPv6 IPsec ESP, tunnel/transport mode, and SA management
- DPDK EthDev integrated
- For crypto it uses OpenSSL by default
- Performance?





FD.io/VPP IPsec with OpenSSL as Crypto Performance^s



- ▶ Does Securing the Network Application have to degrade performance?
- ▶ **Not Really**



Agenda

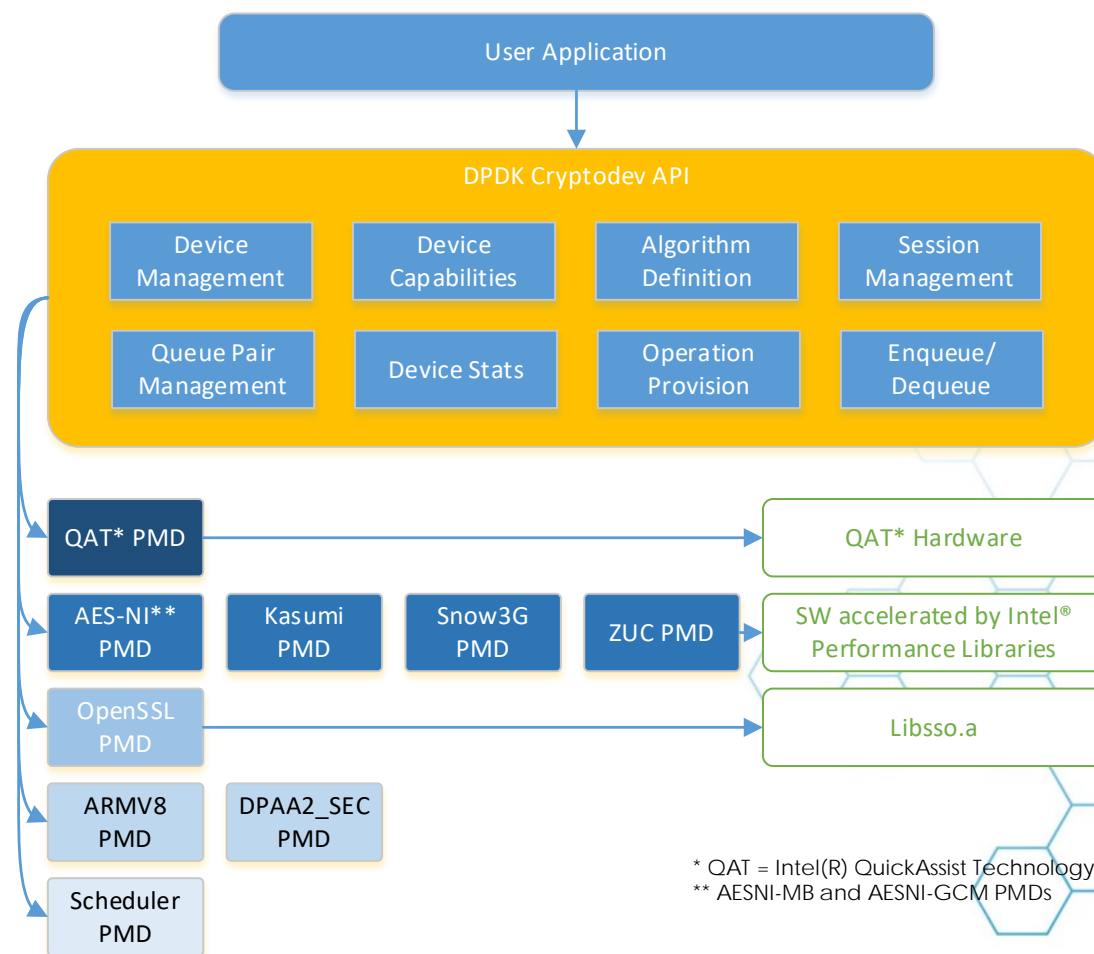
- ▶ Problem Statement
- ▶ **DPDK Cryptodev Framework Introduction**
- ▶ Enable DPDK Cryptodev Framework in VPP
- ▶ Performance
- ▶ Future work
- ▶ Conclusion





DPDK Cryptodev Framework

- ▶ Crypto framework for processing symmetric crypto workloads in DPDK.
- ▶ DPDK Cryptodev consists of:
 - ▶ SW and HW Crypto PMDs
 - ▶ A standard API supports all PMDs
 - ▶ Multi-queues for multi-thread sharing
- ▶ Effortless migration (SW-HW, PHY-VIRT)
- ▶ Asynchronous enqueue/dequeue



* QAT = Intel(R) QuickAssist Technology

** AESNI-MB and AESNI-GCM PMDs



Supported Algorithms In Cryptodev

Cipher Algorithms

- **AES CBC/CTR**
128/192/256 bit
- **Snow3G (UEA2)**
- **KASUMI F8,**
- **ZUC EEA3**
- **AES_CFB**
- **NULL**

Hash Algorithms

- **MD5_HMAC***
- **SHA1/224*/256/384*/512,**
- **AES XCBC,**
- **Snow3G UIA2,**
- **KASUMI F9,**
- **ZUC EIA3,**
- **NULL**

AEAD Algorithms

- **AES GCM 128/192**/256 bit**

* QAT = Intel(R) QuickAssist Technology

** AESNI-MB and AESNI-GCMPMDs



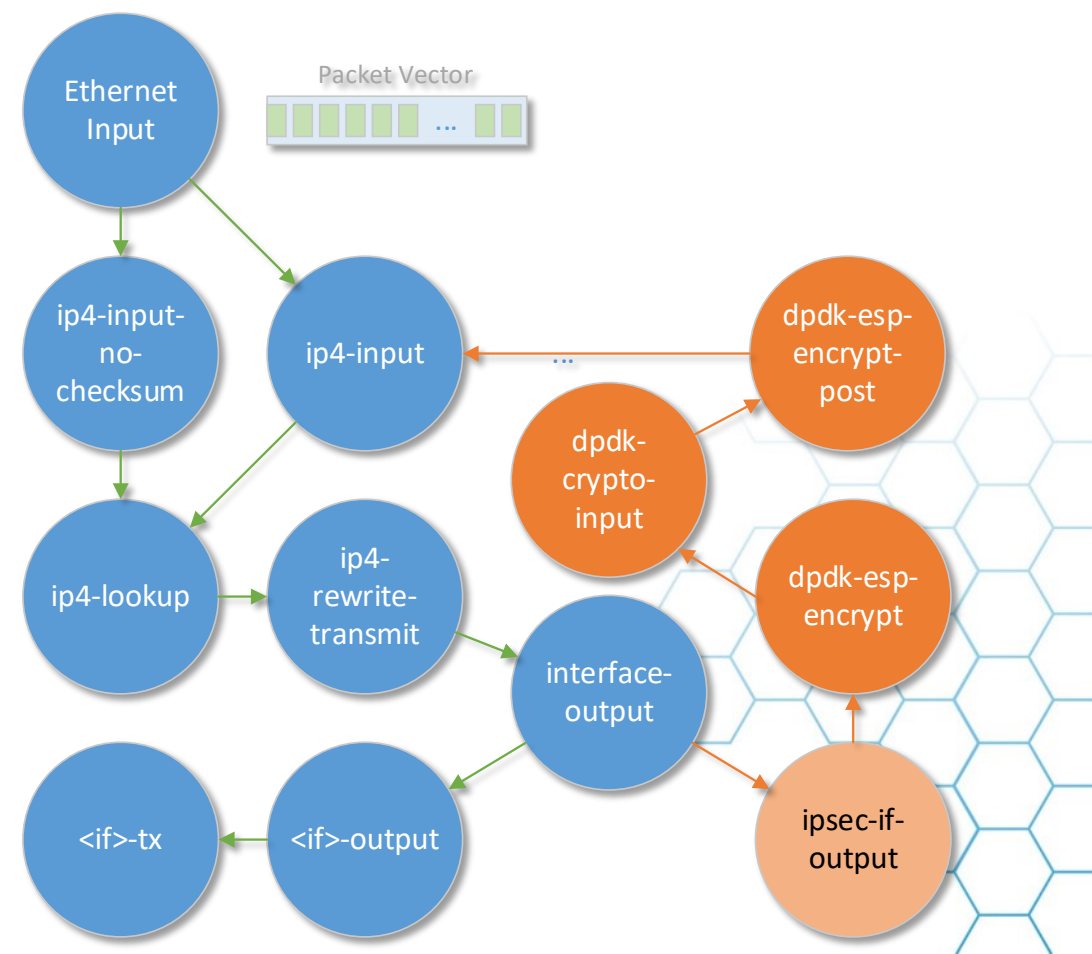
Agenda

- ▶ Problem Statement
- ▶ DPDK Cryptodev Framework Introduction
- ▶ **Enable DPDK Cryptodev Framework in VPP**
- ▶ Performance
- ▶ Future work
- ▶ Conclusion



Enable DPDK Cryptodev in VPP IPsec

- ▶ Replaced 2 nodes:
 - ▶ esp-encrypt → dpdk-esp-encrypt
 - ▶ esp-decrypt → dpdk-esp-decrypt
- ▶ Added 3 nodes:
 - ▶ dpdk-crypto-input
 - ▶ dpdk-esp-encrypt-post
 - ▶ dpdk-esp-decrypt-post





VPP Configuration for DPDK Cryptodev

- ▶ Environmental option:
 - ▶ For software PMD:
vpp_use_dpdk_cryptodev_sw=yes
- ▶ User only needs to provide Cryptodevs in startup.conf file
- ▶ Allocate crypto resources on best effort approach
- ▶ No special IPsec configuration is required
- ▶ More information can be found [here](#)

Sample Configuration:

```
dpdk {  
    ...  
    #HW PMDs  
    enable-cryptodev  
    dev 0000:85:01.0  
    dev 0000:85:01.1  
    #SW PMDs  
    vdev cryptodev_aesni_mb_pmd0,socket_id=1  
    vdev cryptodev_aesni_mb_pmd1,socket_id=1  
}
```



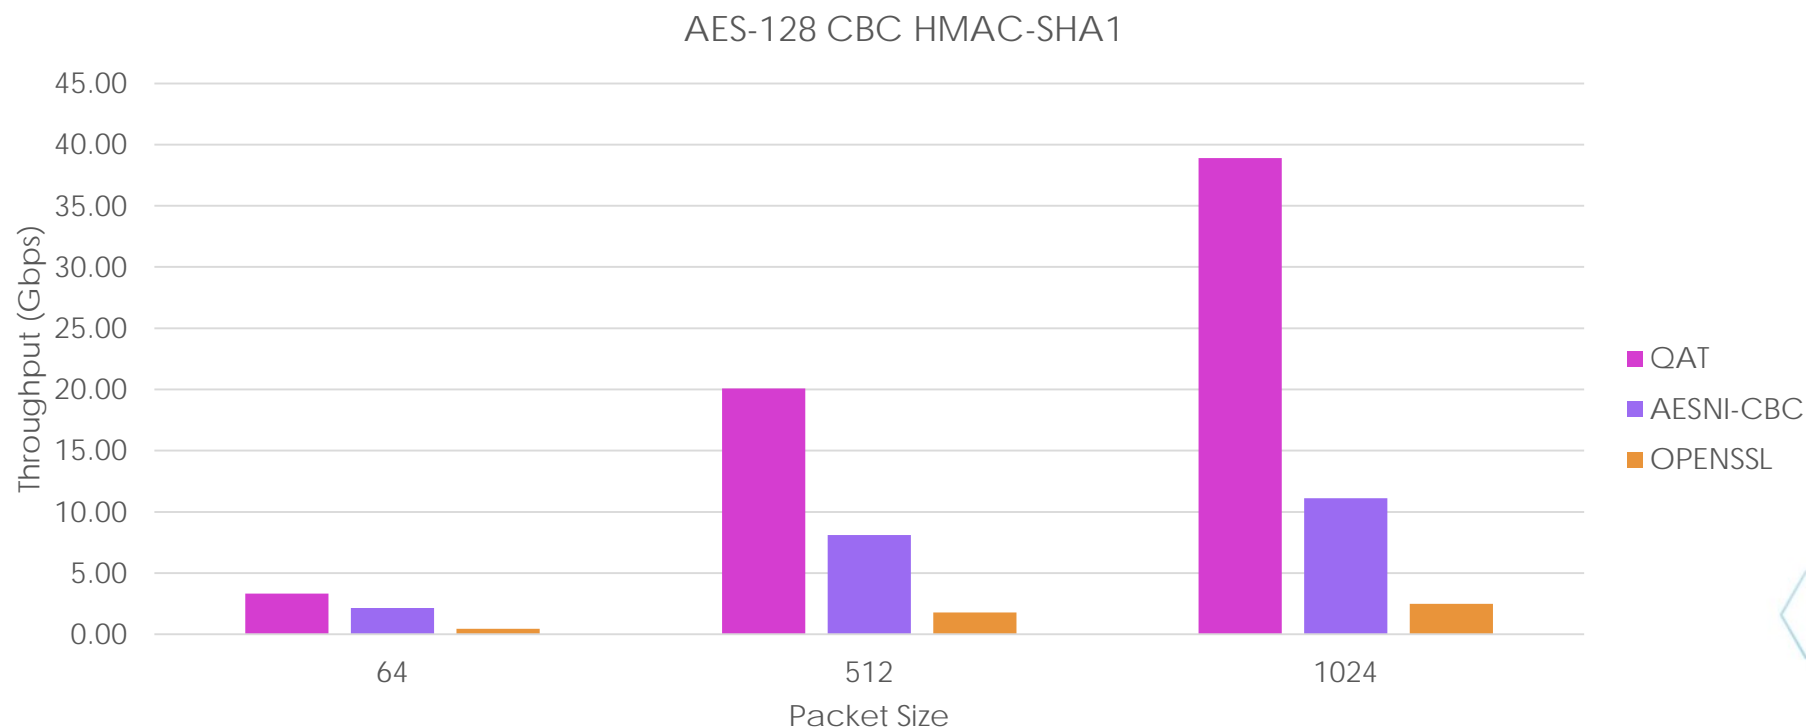
Agenda

- ▶ Problem Statement
- ▶ DPDK Cryptodev Framework Introduction
- ▶ Enable DPDK Cryptodev Framework in VPP
- ▶ **Performance**
- ▶ Future work
- ▶ Conclusion





Performance§ from VPP IPsec



*QAT = Intel(R) QuickAssist Technology

§ Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks

Test and System Configurations: Estimates are based on internal Intel analysis using at least Data Plane Development Kit IPsec sample application on Intel(R) Xeon(R) CPU E5-2695 v4@ 2.10GHz with atleast using Intel(R) Communications Chipset(s) 8955 with Intel(R) QuickAssist Technology.





Agenda

- ▶ Problem Statement
- ▶ DPDK Cryptodev Framework Introduction
- ▶ Enable DPDK Cryptodev Framework in VPP
- ▶ Performance
- ▶ **Future work**
- ▶ Conclusion





Future Work

- ▶ DPDK Cryptodev Optimization
- ▶ Enable DPDK Cryptodev Framework in VPP IKEv2.
- ▶ VPP IPsec Performance Tuning
- ▶ Enable DPDK Cryptodev Scheduler PMD to increase crypto workload processing capability per-worker thread
- ▶ Virtio-Crypto Enabling





Agenda

- ▶ Problem Statement
- ▶ DPDK Cryptodev Framework Introduction
- ▶ Enable DPDK Cryptodev Framework in VPP
- ▶ Performance
- ▶ Future work
- ▶ **Conclusion**





Summary

- ▶ Achieved VPP IPsec Performance boost by enabling DPDK Cryptodev Framework
- ▶ QAT hardware accelerated VPP IPsec has more performance boost than the software alternative
- ▶ Seamlessly integrated into VPP, easy to enable and configure, no extra IPsec configuration is required
- ▶ Migration between Software and Hardware, Physical and Virtual, is effortless





Acknowledgement

Arkadiuszx kusztal (arkadiuszx.kusztal@intel.com)
Declan Doherty (declan.doherty@intel.com)
Fiona Trahe (fiona.trahe@intel.com)
Jain Deepak (deepak.k.jain@intel.com)
John Griffin (john.griffin@intel.com)
Kirill Rybalchenko (kirill.rybalchenko@intel.com)
Pablo D. L. Guarch (pablo.de.lara.guarch@intel.com)
Radu Nicolau (radu.nicolau@intel.com)
Sergio G. M (sergio.gonzalez.monroy@intel.com)





Q&A

Thanks!!

