DPDK
DATA PLANE DEVELOPMENT KIT

DPDK Summit

DPDK Cryptodev

Deepak Kumar Jain
Technical Project Manager, Intel

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.
-
- 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 E5-2658 v4@ 2.30GHz with atleast using Intel(R) Communications Chipset(s) 8955 with Intel(R) QuickAssist Technology.

Agenda

- DPDK CRYPTODEV INTRODUCTION
- FEATURES
  - SUPPORTED ALGORITHMS
  - CRYPTODEV PACKET PROCESSING FLOW
  - EFFORTLESS MIGRATION
  - SCHEDULER PMD
- VPP + DPDK CRYPTODEV FRAMEWORK
- PERFORMANCE
- FUTURE WORKS
- SUMMARY

Agenda

- DPDK CRYPTODEV INTRODUCTION
- FEATURES
  - SUPPORTED ALGORITHMS
  - CRYPTODEV PACKET PROCESSING FLOW
  - EFFORTLESS MIGRATION
  - SCHEDULER PMD
- VPP + DPDK CRYPTODEV FRAMEWORK
- PERFORMANCE
- FUTURE WORKS
- SUMMARY

DPDK CRYPTODEV INTRODUCTION

- Crypto framework for processing symmetric crypto workloads in DPDK.
- DPDK Cryptodev consists of:
  - Crypto Poll Mode Drivers for hardware accelerated lookaside (Intel® QuickAssist Technology) and software based crypto primitives
  - A standard API supports all PMDs
- Allowing effortless migration of work between hardware and software, even between physical to virtual environments

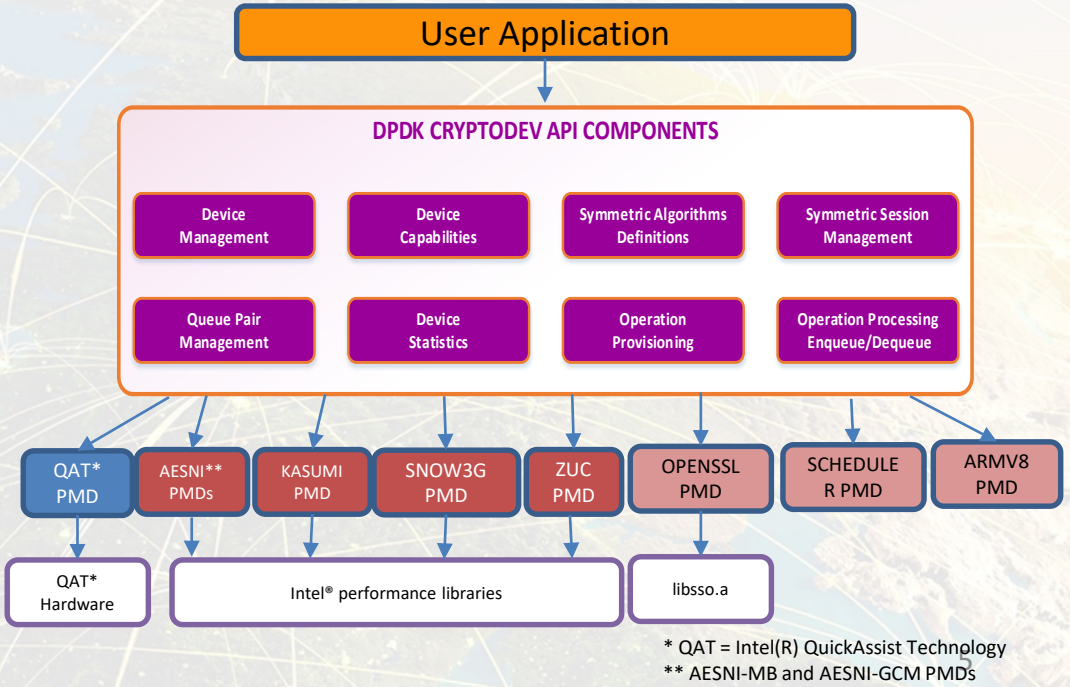

* QAT = Intel(R) QuickAssist Technology
** AESNI-MB and AESNI-GCM PMDs

Cryptodev PMDs

QAT	AESNI MB AESNI GCM ARMv8	SNOW 3G KASUMI ZUC	OpenSSL	Scheduler	NULL
PMD for hardware acceleration	PMDs for optimized software acceleration libraries	PMDs for optimized software acceleration libraries for wireless algorithms	PMD for non-optimized software implementation	PMD to distribute packets across multiple accelerators	PMD for test purposes

Future work includes:

- Extending the API to support asymmetric crypto.
- More advanced Scheduler capabilities.

Agenda

- DPDK CRYPTODEV INTRODUCTION
- FEATURES
  - SUPPORTED ALGORITHMS
  - CRYPTODEV PACKET PROCESSING FLOW
  - EFFORTLESS MIGRATION
  - SCHEDULER PMD
- VPP + DPDK CRYPTODEV FRAMEWORK
- PERFORMANCE
- FUTURE WORKS
- SUMMARY

SUPPORTED ALGORITHMS IN CRYPTODEV

Cipher Algorithms	Hash Algorithms	AEAD Algorithms
• *AES CBC/CTR 128/192/256 bit*	• *MD5_HMAC**	• *AES GCM 128/192**/256 bit*
• *Snow3G (UEA2)*	• *SHA1/224*/256/384*/512,*	
• *KASUMI F8,*	• *AES XCBC,*	
• *ZUC EEA3*	• *Snow3G UIA2,*	
• AES_CFB	• *KASUMI F9,*	
	• *ZUC EIA3,*	
	• *NULL*	

* QAT = Intel(R) QuickAssist Technology
** AESNI-MB and AESNI-GCM PMDs

CRYPTODEV PACKET PROCESSING FLOW

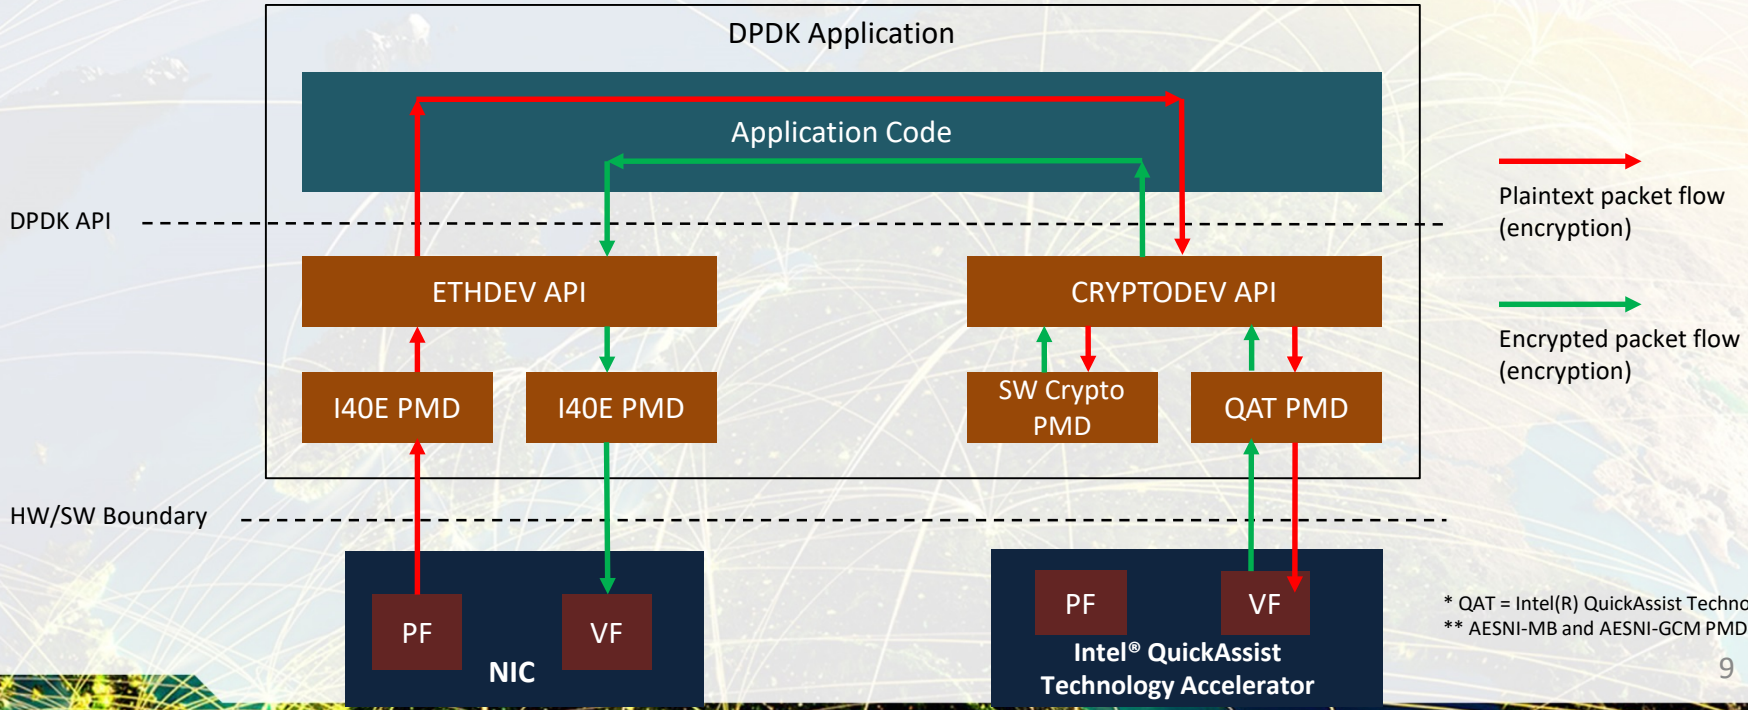

* QAT = Intel(R) QuickAssist Technology
** AESNI-MB and AESNI-GCM PMDs

Effortless Migration (HW & SW)

- ./l2fwd-crypto -l 11 -n 4 --socket-mem 0,2048 -w 88:01.0 -w 88:01.1 -w 84:00.0 -- -p 0x1 --chain CIPHER_HASH --cipher_op ENCRYPT --cipher_algo AES_CBC --cipher_key 00:01:02:03:04:05:06:07:08:09:0a:0b:0c:0d:0e:0f --auth_op GENERATE --auth_algo SHA1_HMAC --auth_key 10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23
- ./l2fwd-crypto -l 11 -n 4 --socket-mem 0,2048 -w 88:01.0 -w 88:01.1 --vdev "crypto_aesni_mb" -- -p 0x1 --chain CIPHER_HASH --cipher_op ENCRYPT --cipher_algo AES_CBC --cipher_key 00:01:02:03:04:05:06:07:08:09:0a:0b:0c:0d:0e:0f --auth_op GENERATE --auth_algo SHA1_HMAC --auth_key 10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23

Same application can be used on both SW PMD and QAT PMD, simply address the device in the EAL commandline option

Scheduler PMD

- Distributing crypto ops to multiple crypto PMDs (slaves)
- Supports multiple distribution modes:
  - Round-robin mode to balance workload across multiple slaves. (DPDK 17.02)
  - Packet Size based mode (DPDK 17.05 RC1)
  - More modes are planned for future releases
- Provides API to manage slaves, set modes, and enable/disable ordering
- Provided API for user to use his own crafted mode

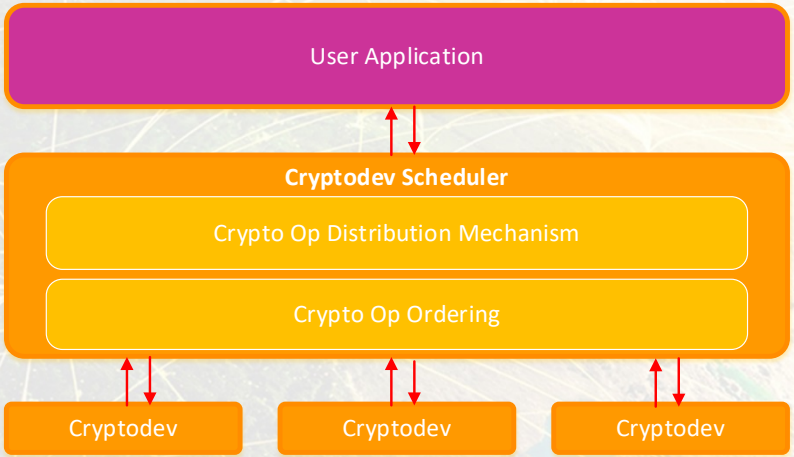

Agenda

- DPDK CRYPTODEV INTRODUCTION
- FEATURES
  - SUPPORTED ALGORITHMS
  - CRYPTODEV PACKET PROCESSING FLOW
  - EFFORTLESS MIGRATION
  - SCHEDULER PMD
- **VPP + DPDK CRYPTODEV FRAMEWORK**
- PERFORMANCE
- FUTURE WORK
- SUMMARY

FD.io/VPP

- Open-source Linux Foundation Project.
- Highly performant data plane platform.
- VPP is a packet processing engine using DPDK as the network I/O.
- Run-to-completion mode of VPP

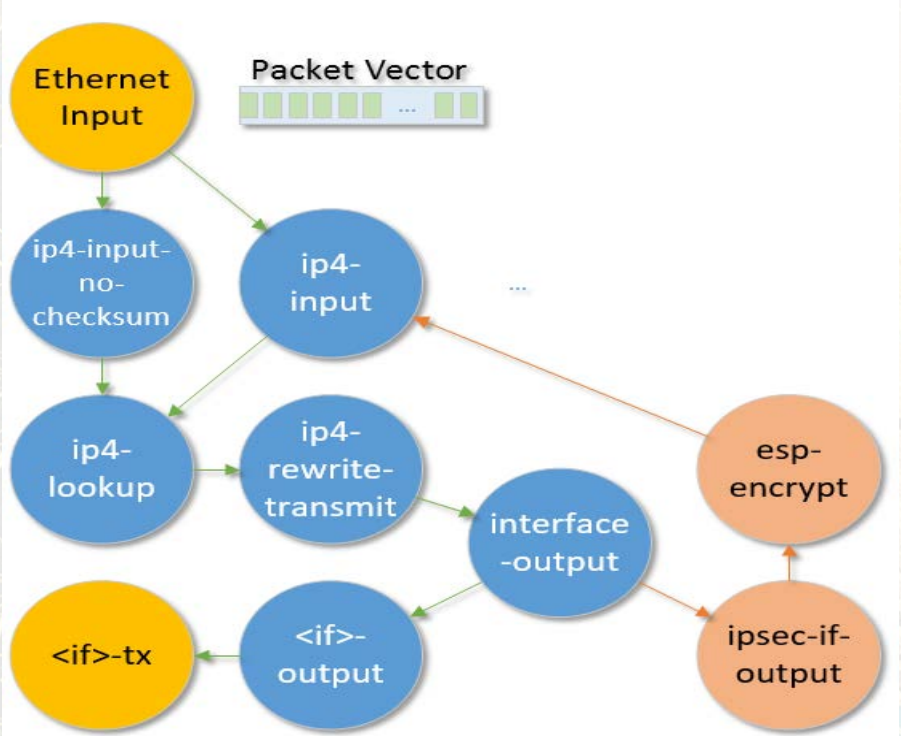

FD.io/VPP + DPDK CRYPTODEV FRAMEWORK

- FD.io/VPP supports IPv4/IPv6 IPsec ESP, tunnel/transport mode, and SA management.
- It has DPDK EthDev integrated, but didn't have DPDK Cryptodev enabled.
- We integrated DPDK Cryptodev Framework into VPP, to accelerate VPP IPsec with Intel® Performance Library and/or Intel® QuickAssist Technology.

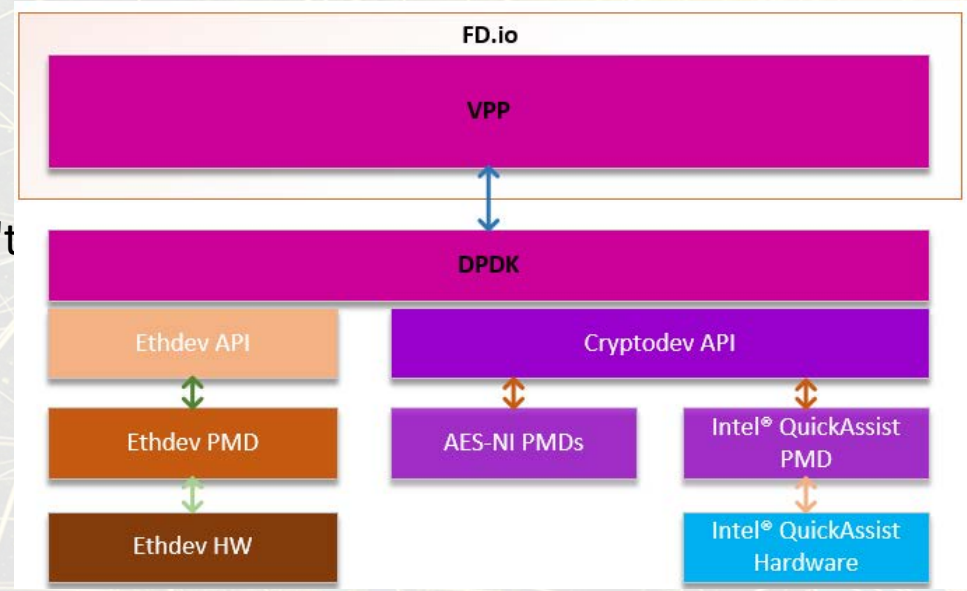

ENABLE DPDK CRYPTODEV IN VPP IPSEC

- Replace VPP esp-encrypt and esp-decrypt nodes with dpdk-esp-encrypt and dpdk-esp-decrypt
- Added 3 nodes:
  - dpdk-crypto-input: polling input node, dequeuing from crypto PMDs
  - dpdk-esp-encrypt-post and dpdk-esp-decrypt-post: encapsulate to valid packet vectors from dequeued packets

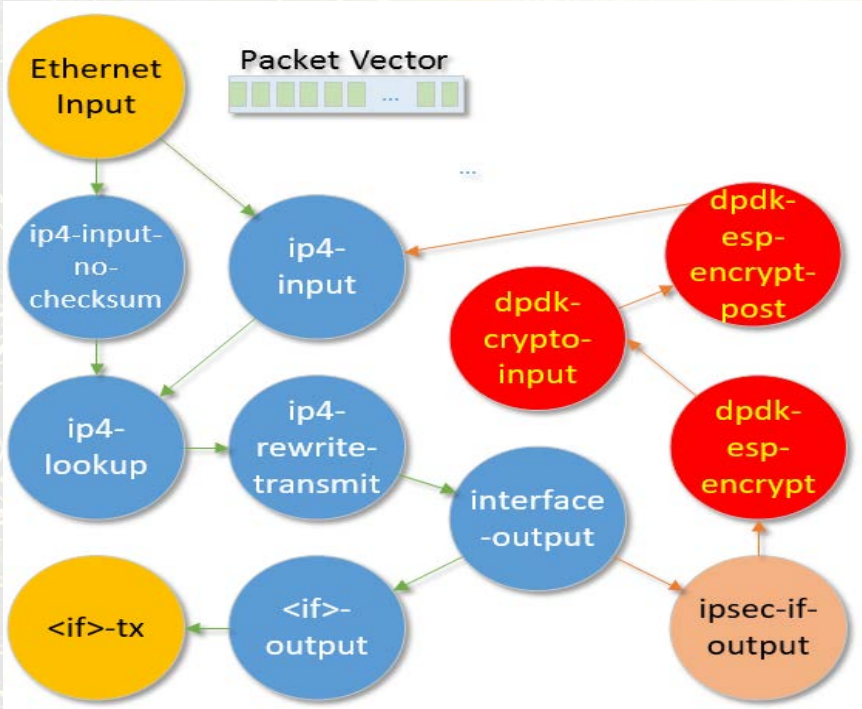

VPP CONFIGURATION FOR DPDK CRYPTODEV

- Environmental option: **vpp_use_dpdk_cryptodev=yes**
- No special IPsec configuration is required
- Allocate crypto resources on best effort approach: hardware first, then software. If there is not enough crypto resource for every worker, drop all packets.
- User only needs to provide Cryptodevs in startup.conf file.

Agenda

- DPDK CRYPTODEV INTRODUCTION
- FEATURES
  - SUPPORTED ALGORITHMS
  - DPDK CRYPTODEV WORKFLOW
  - EFFORTLESS MIGRATION
  - SCHEDULER PMD
- VPP + DPDK CRYPTODEV FRAMEWORK
- **PERFORMANCE**
- FUTURE WORK
- SUMMARY

Performance§ from DPDK IPSec sample application

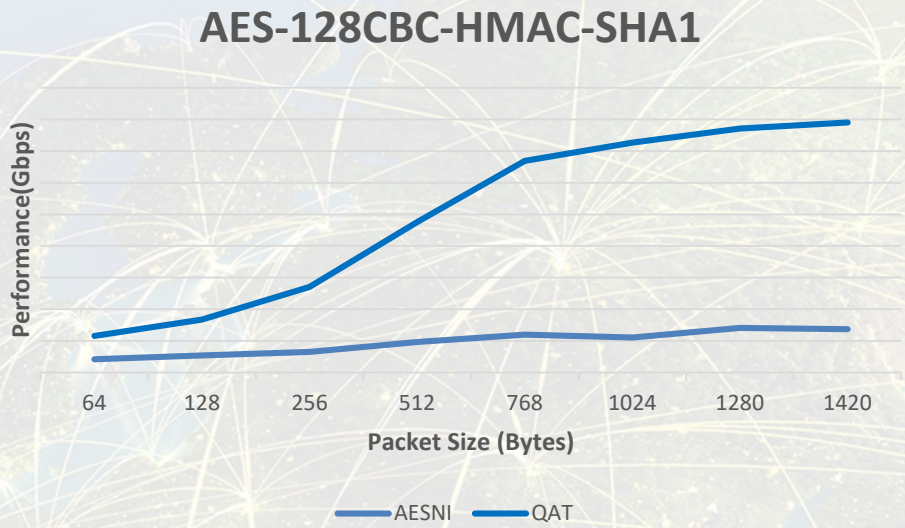

*QAT = Intel(R) QuickAssist Technology

§ 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 E5-2658 v4@ 2.30GHz with atleast using Intel(R) Communications Chipset(s) 8955 with Intel(R) QuickAssist Technology.

Agenda

- DPDK CRYPTODEV INTRODUCTION
- FEATURES
  - SUPPORTED ALGORITHMS
  - DPDK CRYPTODEV WORKFLOW
  - EFFORTLESS MIGRATION
  - SCHEDULER PMD
- PERFORMANCE
- **FUTURE WORK**
- SUMMARY

PLANNED FEATURES IN FUTURE RELEASES

Performance

- QAT* PMD optimizations
- SW PMD optimizations
- Refactoring & Clean up

Algorithm support

- QAT* PMD
  - DES
  - ZUC
  - AES-CFB64/ECB
- **SW PMD**
  - DES
  - AES-CFB64/ECB

Scheduler

- Packet-size based scheduling mode
- Multi-core scheduling mode

VPP IPsec integration

- Performance Optimization
- Enable DPDK cryptodev in VPP IKEv2
- Enable DPDK Cryptodev scheduler PMD

*QAT = Intel(R) QuickAssist Technology

Summary

- Cryptodev currently provides support for symmetric algorithms.
- Provides both Software(SW) and Hardware (Intel® QuickAssist Technology) implementations.
- Healthy pipeline of features planned for future release.
- HW provides provides major boost in performance over SW implementation.

QUESTIONS?

Deepak Kumar Jain

deepak.k.jain@intel.com

BACKUP

VPP IPsec with DPDK Cryptodev: How it works

- Assign available Cryptodev resources to each worker

VPP IPsec with DPDK Cryptodev: How it works

- Assign available Cryptodev resources to each worker
- When adding an SA, create sessions for each worker/crypto resource based on the specified algorithms and store them along with queue pair info with same index of SA.

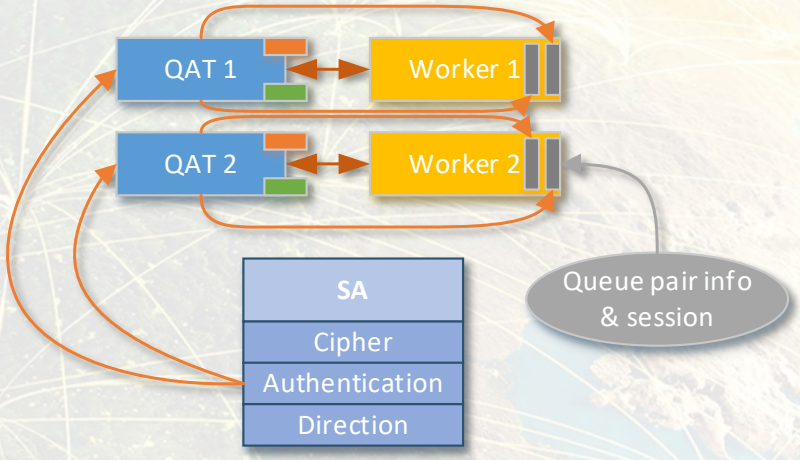

VPP IPsec with DPDK Cryptodev: How it works

- Assign available Cryptodev resources to every worker.
- When adding an SA, create sessions for each worker/crypto resource based on the specified algorithms and store them along with queue pair info with same index of SA.
- For each packet, retrieve SA, and hence get session for this worker.

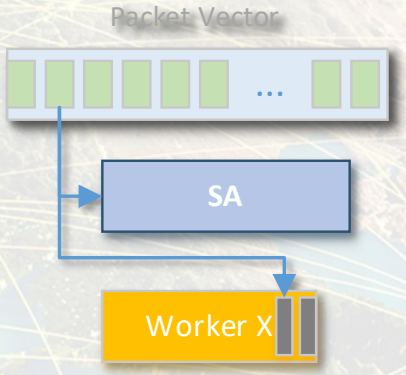

VPP IPsec with DPDK Cryptodev: How it works

- Create crypto op for each packet, attach key, pass data and digest pointer, etc.
- Attach session to each crypto op
- Enqueue the burst of crypto ops to the destination crypto device / queue

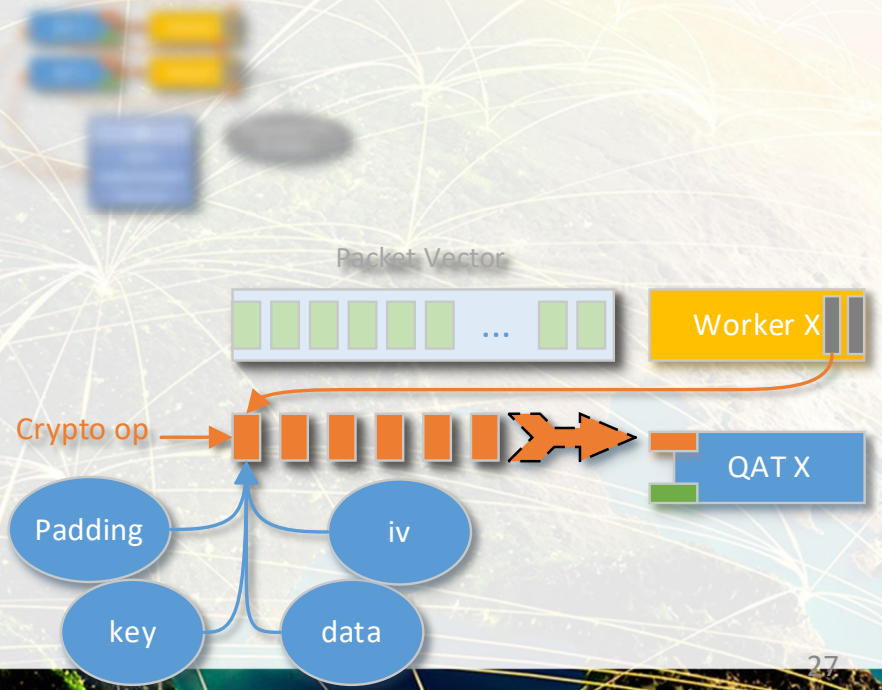

VPP IPsec with DPDK Cryptodev: How it works

- Dequeue the burst of processed crypto ops from the same device/queue in dpdk-crypto-input node

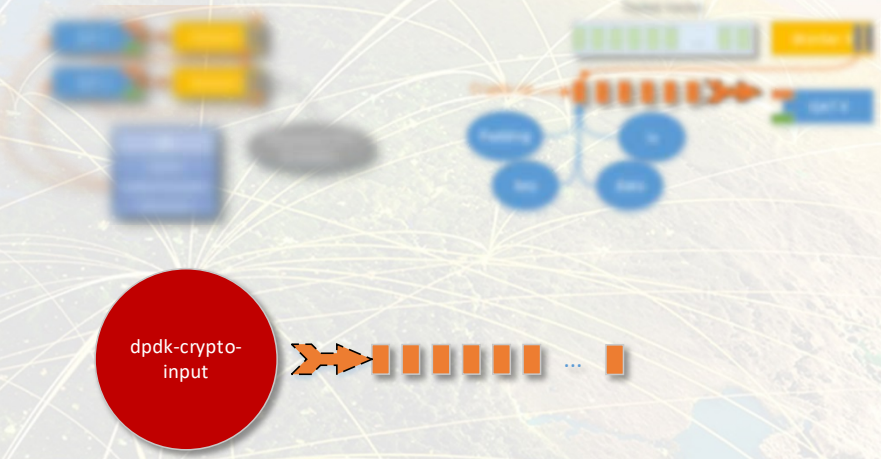

- The asynchronous enqueue/dequeue action helps amortize the cost of crypto operations over multiple packets and also greatly maximize the performance when offloading to hardware lookaside.

VPP IPsec with DPDK Cryptodev: How it works

- Dequeue the burst of processed crypto ops from the same device/queue in dpdk-crypto-input node
- Encapsulate the crypto ops to a valid IPsec packet vector in dpdk-encrypt/decrypt-post node, and pass to next graph node.

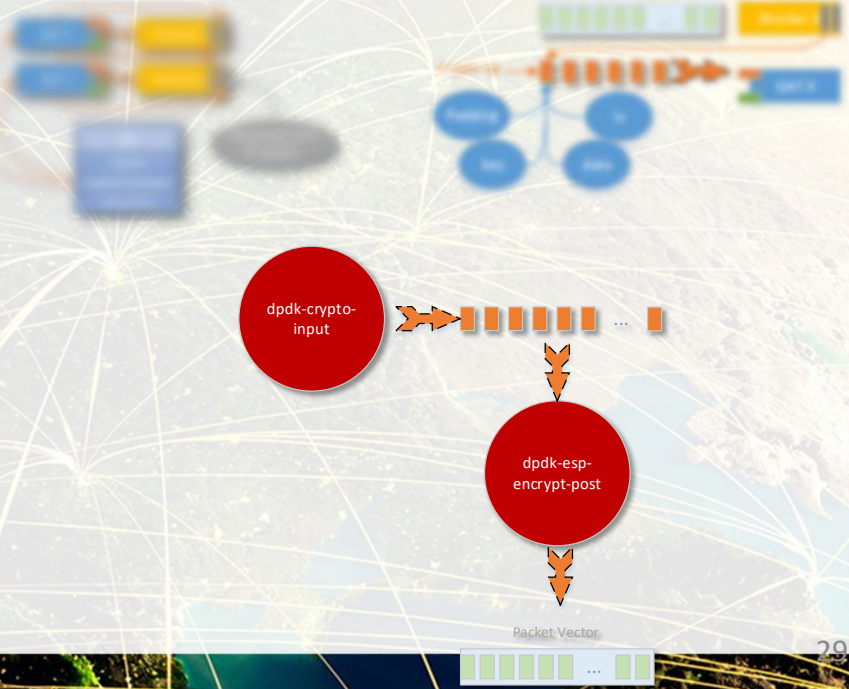

THANK YOU