DPDK

DATA PLANE DEVELOPMENT KIT

Crypto Device Drivers

Release 17.05.2

September 07, 2017

CRYPTO DEVICE SUPPORTED FUNCTIONALITY MATRICES

1.1 Supported Feature Flags

Table 1.1: Features availability in crypto drivers

Feature	aesni _gcm	aesni _mb	armv8	dpaa2_sec	kasumi	null	openssl	qat	snow3g	zuc
Symmetric crypto	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Asymmetric crypto										
Sym operation chaining	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
HW Accelerated				Y				Y		
CPU SSE		Y								
CPU AVX		Y								
CPU AVX2		Y								
CPU AVX512		Y								
CPU AESNI	Y	Y								
CPU NEON			Y							
CPU ARM CE			Y							

1.2 Supported Cipher Algorithms

Table 1.2: Cipher algorithms in crypto drivers

Cipher algorithm	aesni _gcm	aesni_mb	armv8	dpaa2_sec	kasumi	null	openssl	qat	snow3g	zuc
NULL						Y		Y		
AES CBC (128)		Y	Y	Y			Y	Y		
AES CBC (192)		Y		Y			Y	Y		
AES CBC (256)		Y		Y			Y	Y		
AES CTR (128)		Y					Y	Y		
AES CTR (192)		Y					Y	Y		
AES CTR (256)		Y					Y	Y		
AES DOCSIS BPI		Y						Y		
3DES CBC				Y			Y	Y		
3DES CTR							Y	Y		
DES CBC								Y		
DES DOCSIS BPI							Y	Y		
SNOW3G UEA2								Y	Y	
KASUMI F8					Y			Y		
ZUC EEA3								Y		Y

1.3 Supported Authentication Algorithms

Table 1.3: Authentication algorithms in crypto drivers

Authentication algorithm	aesni _gcm	aesni _mb	armv 8	dpaa 2_sec	kas umi	null	ope nssl	qat	sno w3g	zuc
NULL						Y		Y		
MD5							Y			
MD5 HMAC		Y		Y			Y	Y		
SHA1							Y			
SHA1 HMAC		Y	Y	Y			Y	Y		
SHA224							Y			
SHA224 HMAC		Y		Y			Y	Y		
SHA256							Y			
SHA256 HMAC		Y	Y	Y			Y	Y		
SHA384							Y			
SHA384 HMAC		Y		Y			Y	Y		
SHA512							Y			
SHA512 HMAC		Y		Y			Y	Y		
AES XCBC MAC		Y						Y		
AES GMAC	Y						Y	Y		
SNOW3G UIA2								Y	Y	
KASUMI F9					Y			Y		
ZUC EIA3								Y		Y

1.4 Supported AEAD Algorithms

Table 1.4: AEAD algorithms in crypto drivers

AEAD algorithm	aesni _gcm	aesni _mb	armv 8	dpaa2 _sec	kas umi	null	ope nssl	qat	sno w3g	zuc
AES GCM (128)	Y						Y	Y		
AES GCM (192)							Y	Y		
AES GCM (256)	Y						Y	Y		

AESN-NI MULTI BUFFER CRYPTO POLL MODE DRIVER

The AESNI MB PMD (**librte_pmd_aesni_mb**) provides poll mode crypto driver support for utilizing Intel multi buffer library, see the white paper Fast Multi-buffer IPsec Implementations on Intel® Architecture Processors.

The AES-NI MB PMD has current only been tested on Fedora 21 64-bit with gcc.

2.1 Features

AESNI MB PMD has support for:

Cipher algorithms:

- RTE_CRYPTO_CIPHER_AES128_CBC
- RTE_CRYPTO_CIPHER_AES192_CBC
- RTE_CRYPTO_CIPHER_AES256_CBC
- RTE_CRYPTO_CIPHER_AES128_CTR
- RTE_CRYPTO_CIPHER_AES192_CTR
- RTE_CRYPTO_CIPHER_AES256_CTR
- RTE_CRYPTO_CIPHER_AES_DOCSISBPI

Hash algorithms:

- RTE_CRYPTO_HASH_MD5_HMAC
- RTE_CRYPTO_HASH_SHA1_HMAC
- RTE_CRYPTO_HASH_SHA224_HMAC
- RTE_CRYPTO_HASH_SHA256_HMAC
- RTE_CRYPTO_HASH_SHA384_HMAC
- RTE_CRYPTO_HASH_SHA512_HMAC
- RTE_CRYPTO_HASH_AES_XCBC_HMAC

2.2 Limitations

- Chained mbufs are not supported.

• Only in-place is currently supported (destination address is the same as source address).

2.3 Installation

To build DPDK with the AESNI_MB_PMD the user is required to download the multi-buffer library from here and compile it on their user system before building DPDK. The latest version of the library supported by this PMD is v0.45, which can be downloaded in https://github.com/01org/intel-ipsec-mb/archive/v0.45.zip.

```
make
```

As a reference, the following table shows a mapping between the past DPDK versions and the Multi-Buffer library version supported by them:

Table 2.1: DPDK and Multi-Buffer library version compatibility

DPDK version	Multi-buffer library version
2.2 - 16.11	0.43 - 0.44
17.02	0.44
17.05	0.45

2.4 Initialization

In order to enable this virtual crypto PMD, user must:

• Export the environmental variable AESNI_MULTI_BUFFER_LIB_PATH with the path where the library was extracted.

• Build the multi buffer library (explained in Installation section).

• Set CONFIG_RTE_LIBRTE_PMD_AESNI_MB=y in config/common_base.

To use the PMD in an application, user must:

• Call rte_vdev_init("crypto_aesni_mb") within the application.

• Use –vdev="crypto_aesni_mb" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

• socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).

• max_nb_queue_pairs: Specify the maximum number of queue pairs in the device (8 by default).

• max_nb_sessions: Specify the maximum number of sessions that can be created (2048 by default).

Example:

```
./l2fwd-crypto -l 6 -n 4 --vdev="crypto_aesni_mb,socket_id=1,max_nb_sessions=128"
```

AES-NI GCM CRYPTO POLL MODE DRIVER

The AES-NI GCM PMD (**librte_pmd_aesni_gcm**) provides poll mode crypto driver support for utilizing Intel ISA-L crypto library, which provides operation acceleration through the AES-NI instruction sets for AES-GCM authenticated cipher algorithm.

3.1 Features

AESNI GCM PMD has support for:

Cipher algorithms:

- RTE_CRYPTO_CIPHER_AES_GCM

Authentication algorithms:

- RTE_CRYPTO_AUTH_AES_GCM
- RTE_CRYPTO_AUTH_AES_GMAC

3.2 Installation

To build DPDK with the AESNI_GCM_PMD the user is required to install the `libisal_crypto` library in the build environment. For download and more details please visit https://github.com/01org/isa-l_crypto.

3.3 Initialization

In order to enable this virtual crypto PMD, user must:

- Install the ISA-L crypto library (explained in Installation section).
- Set CONFIG_RTE_LIBRTE_PMD_AESNI_GCM=y in config/common_base.

To use the PMD in an application, user must:

- Call rte_vdev_init("crypto_aesni_gcm") within the application.
- Use –vdev="crypto_aesni_gcm" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

- socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).

- max_nb_queue_pairs: Specify the maximum number of queue pairs in the device (8 by default).

- max_nb_sessions: Specify the maximum number of sessions that can be created (2048 by default).

Example:

```
./l2fwd-crypto -l 6 -n 4 --vdev="crypto_aesni_gcm,socket_id=1,max_nb_sessions=128"
```

3.4 Limitations

- Chained mbufs are supported but only out-of-place (destination mbuf must be contiguous).

- Hash only is not supported.

- Cipher only is not supported.

ARMV8 CRYPTO POLL MODE DRIVER

This code provides the initial implementation of the ARMv8 crypto PMD. The driver uses ARMv8 cryptographic extensions to process chained crypto operations in an optimized way. The core functionality is provided by a low-level library, written in the assembly code.

4.1 Features

ARMv8 Crypto PMD has support for the following algorithm pairs:

Supported cipher algorithms:

- `RTE_CRYPTO_CIPHER_AES_CBC`

Supported authentication algorithms:

- `RTE_CRYPTO_AUTH_SHA1_HMAC`

- `RTE_CRYPTO_AUTH_SHA256_HMAC`

4.2 Installation

In order to enable this virtual crypto PMD, user must:

- Download ARMv8 crypto library source code from here

- Export the environmental variable ARMV8_CRYPTO_LIB_PATH with the path where the `armv8_crypto` library was downloaded or cloned.

- Build the library by invoking:

```
make -C $ARMV8_CRYPTO_LIB_PATH/
```

- Set CONFIG_RTE_LIBRTE_PMD_ARMV8_CRYPTO=y in config/defconfig_arm64-armv8a-linuxapp-gcc

The corresponding device can be created only if the following features are supported by the CPU:

- `RTE_CPUFLAG_AES`

- `RTE_CPUFLAG_SHA1`

- `RTE_CPUFLAG_SHA2`

- `RTE_CPUFLAG_NEON`

4.3 Initialization

User can use app/test application to check how to use this PMD and to verify crypto processing.

Test name is cryptodev_sw_armv8_autotest. For performance test cryptodev_sw_armv8_perftest can be used.

4.4 Limitations

- Maximum number of sessions is 2048.
- Only chained operations are supported.
- AES-128-CBC is the only supported cipher variant.
- Cipher input data has to be a multiple of 16 bytes.
- Digest input data has to be a multiple of 8 bytes.

NXP DPAA2 CAAM (DPAA2_SEC)

The DPAA2_SEC PMD provides poll mode crypto driver support for NXP DPAA2 CAAM hardware accelerator.

5.1 Architecture

SEC is the SOC's security engine, which serves as NXP's latest cryptographic acceleration and offloading hardware. It combines functions previously implemented in separate modules to create a modular and scalable acceleration and assurance engine. It also implements block encryption algorithms, stream cipher algorithms, hashing algorithms, public key algorithms, run-time integrity checking, and a hardware random number generator. SEC performs higher-level cryptographic operations than previous NXP cryptographic accelerators. This provides significant improvement to system level performance.

DPAA2_SEC is one of the hardware resource in DPAA2 Architecture. More information on DPAA2 Architecture is described in dpaa2_overview.

DPAA2_SEC PMD is one of DPAA2 drivers which interacts with Management Complex (MC) portal to access the hardware object - DPSECI. The MC provides access to create, discover, connect, configure and destroy dpseci objects in DPAA2_SEC PMD.

DPAA2_SEC PMD also uses some of the other hardware resources like buffer pools, queues, queue portals to store and to enqueue/dequeue data to the hardware SEC.

DPSECI objects are detected by PMD using a resource container called DPRC (like in dpaa2_overview).

For example:

```
DPRC.1 (bus)
  |
  +--+--------+-------+-------+-------+---------+
     |        |       |       |       |         |
  DPMCP.1  DPIO.1  DPBP.1  DPNI.1  DPMAC.1  DPSECI.1
  DPMCP.2  DPIO.2          DPNI.2  DPMAC.2  DPSECI.2
  DPMCP.3
```
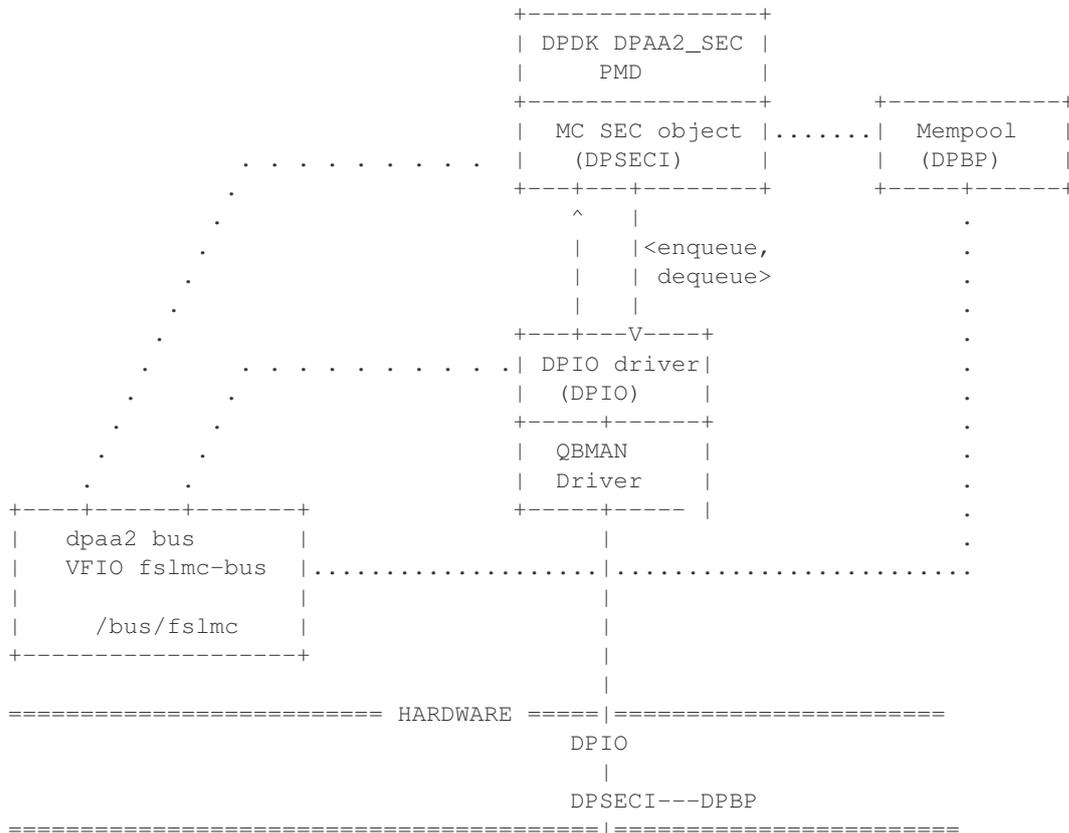
5.2 Implementation

SEC provides platform assurance by working with SecMon, which is a companion logic block that tracks the security state of the SOC. SEC is programmed by means of descriptors (not to be confused with frame descriptors (FDs)) that indicate the operations to be performed and

link to the message and associated data. SEC incorporates two DMA engines to fetch the descriptors, read the message data, and write the results of the operations. The DMA engine provides a scatter/gather capability so that SEC can read and write data scattered in memory. SEC may be configured by means of software for dynamic changes in byte ordering. The default configuration for this version of SEC is little-endian mode.

A block diagram similar to dpaa2 NIC is shown below to show where DPAA2_SEC fits in the DPAA2 Bus model

```
                                +---------------+
                                | DPDK DPAA2_SEC |
                                |     PMD        |
                                +---------------+        +-----------+
                                | MC SEC object |.......| Mempool   |
                . . . . . . . . |   (DPSECI)    |       |  (DPBP)   |
                .               +---+---+-------+        +-----+-----+
                .                   ^   |                      .
                .                   |   |<enqueue,             .
                .                   |   | dequeue>             .
                .                   |   |                      .
                .               +---+---V----+                 .
                . . . . . . . . .| DPIO driver|                .
                .   .           |  (DPIO)    |                 .
                .   .           +-----+------+                 .
                .   .           | QBMAN      |                 .
                .   .           | Driver     |                 .
   +----+------+-------+        +-----+----- |                 .
   |  dpaa2 bus        |             |                         .
   |  VFIO fslmc-bus   |.....................|.......................
   |                   |             |                  |
   |    /bus/fslmc     |             |
   +------------------+             |
                                    |
========================= HARDWARE =====|=======================
                                   DPIO
                                    |
                               DPSECI---DPBP
========================================|=======================
```

5.3 Features

The DPAA2 PMD has support for:

Cipher algorithms:

- RTE_CRYPTO_CIPHER_3DES_CBC

- RTE_CRYPTO_CIPHER_AES128_CBC

- RTE_CRYPTO_CIPHER_AES192_CBC

- RTE_CRYPTO_CIPHER_AES256_CBC

Hash algorithms:

- RTE_CRYPTO_AUTH_SHA1_HMAC

- RTE_CRYPTO_AUTH_SHA224_HMAC

- RTE_CRYPTO_AUTH_SHA256_HMAC

- `RTE_CRYPTO_AUTH_SHA384_HMAC`

- `RTE_CRYPTO_AUTH_SHA512_HMAC`

- `RTE_CRYPTO_AUTH_MD5_HMAC`

5.4 Supported DPAA2 SoCs

- LS2080A/LS2040A

- LS2084A/LS2044A

- LS2088A/LS2048A

- LS1088A/LS1048A

5.5 Limitations

- Chained mbufs are not supported.

- Hash followed by Cipher mode is not supported

- Only supports the session-oriented API implementation (session-less APIs are not supported).

5.6 Prerequisites

DPAA2_SEC driver has similar pre-requisites as described in dpaa2_overview. The following dependencies are not part of DPDK and must be installed separately:

- **NXP Linux SDK**

  NXP Linux software development kit (SDK) includes support for the family of QorIQ® ARM-Architecture-based system on chip (SoC) processors and corresponding boards.

  It includes the Linux board support packages (BSPs) for NXP SoCs, a fully operational tool chain, kernel and board specific modules.

  SDK and related information can be obtained from: NXP QorIQ SDK.

- **DPDK Helper Scripts**

  DPAA2 based resources can be configured easily with the help of ready scripts as provided in the DPDK helper repository.

  DPDK Helper Scripts.

Currently supported by DPDK:

- NXP SDK **2.0+**.

- MC Firmware version **10.0.0** and higher.

- Supported architectures: **arm64 LE**.

- Follow the DPDK Getting Started Guide for Linux to setup the basic DPDK environment.

5.7 Pre-Installation Configuration

5.7.1 Config File Options

Basic DPAA2 config file options are described in dpaa2_overview. In addition to those, the following options can be modified in the `config` file to enable DPAA2_SEC PMD.

Please note that enabling debugging options may affect system performance.

- `CONFIG_RTE_LIBRTE_PMD_DPAA2_SEC` (default `n`) By default it is only enabled in defconfig_arm64-dpaa2-* config. Toggle compilation of the `librte_pmd_dpaa2_sec` driver.

- `CONFIG_RTE_LIBRTE_DPAA2_SEC_DEBUG_INIT` (default `n`) Toggle display of initialization related driver messages

- `CONFIG_RTE_LIBRTE_DPAA2_SEC_DEBUG_DRIVER` (default `n`) Toggle display of driver runtime messages

- `CONFIG_RTE_LIBRTE_DPAA2_SEC_DEBUG_RX` (default `n`) Toggle display of receive fast path run-time message

- `CONFIG_RTE_DPAA2_SEC_PMD_MAX_NB_SESSIONS` By default it is set as 2048 in defconfig_arm64-dpaa2-* config. It indicates Number of sessions to create in the session memory pool on a single DPAA2 SEC device.

5.8 Installations

To compile the DPAA2_SEC PMD for Linux arm64 gcc target, run the following `make` command:

```
cd <DPDK-source-directory>
make config T=arm64-dpaa2-linuxapp-gcc install
```

KASUMI CRYPTO POLL MODE DRIVER

The KASUMI PMD (**librte_pmd_kasumi**) provides poll mode crypto driver support for utilizing Intel Libsso library, which implements F8 and F9 functions for KASUMI UEA1 cipher and UIA1 hash algorithms.

6.1 Features

KASUMI PMD has support for:

Cipher algorithm:

  • RTE_CRYPTO_CIPHER_KASUMI_F8

Authentication algorithm:

  • RTE_CRYPTO_AUTH_KASUMI_F9

6.2 Limitations

  • Chained mbufs are not supported.

  • KASUMI(F9) supported only if hash offset field is byte-aligned.

  • In-place bit-level operations for KASUMI(F8) are not supported (if length and/or offset of data to be ciphered is not byte-aligned).

6.3 Installation

To build DPDK with the KASUMI_PMD the user is required to download the export controlled `libsso_kasumi` library, by requesting it from https://networkbuilders.intel.com/network-technologies/dpdk.   Once approval has been granted, the user needs to log in https://networkbuilders.intel.com/dpdklogin and click on "Kasumi Bit Stream crypto library" link, to download the library. After downloading the library, the user needs to unpack and compile it on their system before building DPDK:

```
make
```

Note: To build the PMD as a shared library, the libsso_kasumi library must be built as follows:

```
make KASUMI_CFLAGS=-DKASUMI_C
```

6.4 Initialization

In order to enable this virtual crypto PMD, user must:

- Export the environmental variable LIBSSO_KASUMI_PATH with the path where the library was extracted (kasumi folder).

- Build the LIBSSO library (explained in Installation section).

- Set CONFIG_RTE_LIBRTE_PMD_KASUMI=y in config/common_base.

To use the PMD in an application, user must:

- Call rte_vdev_init("crypto_kasumi") within the application.

- Use –vdev="crypto_kasumi" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

- socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).

- max_nb_queue_pairs: Specify the maximum number of queue pairs in the device (8 by default).

- max_nb_sessions: Specify the maximum number of sessions that can be created (2048 by default).

Example:

```
./l2fwd-crypto -l 6 -n 4 --vdev="crypto_kasumi,socket_id=1,max_nb_sessions=128"
```

OPENSSL CRYPTO POLL MODE DRIVER

This code provides the initial implementation of the openssl poll mode driver. All cryptography operations are using Openssl library crypto API. Each algorithm uses EVP interface from openssl API - which is recommended by Openssl maintainers.

For more details about openssl library please visit openssl webpage: https://www.openssl.org/

7.1 Features

OpenSSL PMD has support for:

Supported cipher algorithms: * RTE_CRYPTO_CIPHER_3DES_CBC * RTE_CRYPTO_CIPHER_AES_CBC * RTE_CRYPTO_CIPHER_AES_CTR * RTE_CRYPTO_CIPHER_3DES_CTR * RTE_CRYPTO_CIPHER_AES_GCM * RTE_CRYPTO_CIPHER_DES_DOCSISBPI

Supported authentication algorithms: * RTE_CRYPTO_AUTH_AES_GMAC * RTE_CRYPTO_AUTH_MD5 * RTE_CRYPTO_AUTH_SHA1 * RTE_CRYPTO_AUTH_SHA224 * RTE_CRYPTO_AUTH_SHA256 * RTE_CRYPTO_AUTH_SHA384 * RTE_CRYPTO_AUTH_SHA512 * RTE_CRYPTO_AUTH_MD5_HMAC * RTE_CRYPTO_AUTH_SHA1_HMAC * RTE_CRYPTO_AUTH_SHA224_HMAC * RTE_CRYPTO_AUTH_SHA256_HMAC * RTE_CRYPTO_AUTH_SHA384_HMAC * RTE_CRYPTO_AUTH_SHA512_HMAC

7.2 Installation

To compile openssl PMD, it has to be enabled in the config/common_base file and appropriate openssl packages have to be installed in the build environment.

The newest openssl library version is supported: * 1.0.2h-fips 3 May 2016. Older versions that were also verified: * 1.0.1f 6 Jan 2014 * 1.0.1 14 Mar 2012

For Ubuntu 14.04 LTS these packages have to be installed in the build system: sudo apt-get install openssl sudo apt-get install libc6-dev-i386 (for i686-native-linuxapp-gcc target)

This code was also verified on Fedora 24. This code was NOT yet verified on FreeBSD.

7.3 Initialization

User can use app/test application to check how to use this pmd and to verify crypto processing.

Test name is cryptodev_openssl_autotest. For performance test cryptodev_openssl_perftest can be used.

To verify real traffic l2fwd-crypto example can be used with this command:

```
sudo ./build/l2fwd-crypto -l 0-1 -n 4 --vdev "crypto_openssl"
--vdev "crypto_openssl"-- -p 0x3 --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
--iv 00:01:02:03:04:05:06:07:08:09:0a:0b:0c:0d:0e:ff
--auth_op GENERATE --auth_algo SHA1_HMAC
--auth_key 11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11
:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11
:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11:11
```

7.4 Limitations

- Maximum number of sessions is 2048.
- Chained mbufs are supported only for source mbuf (destination must be contiguous).
- Hash only is not supported for GCM and GMAC.
- Cipher only is not supported for GCM and GMAC.

NULL CRYPTO POLL MODE DRIVER

The Null Crypto PMD (**librte_pmd_null_crypto**) provides a crypto poll mode driver which provides a minimal implementation for a software crypto device. As a null device it does not modify the data in the mbuf on which the crypto operation is to operate and it only has support for a single cipher and authentication algorithm.

When a burst of mbufs is submitted to a Null Crypto PMD for processing then each mbuf in the burst will be enqueued in an internal buffer for collection on a dequeue call as long as the mbuf has a valid rte_mbuf_offload operation with a valid rte_cryptodev_session or rte_crypto_xform chain of operations.

8.1 Features

Modes:

- RTE_CRYPTO_XFORM_CIPHER ONLY
- RTE_CRYPTO_XFORM_AUTH ONLY
- RTE_CRYPTO_XFORM_CIPHER THEN RTE_CRYPTO_XFORM_AUTH
- RTE_CRYPTO_XFORM_AUTH THEN RTE_CRYPTO_XFORM_CIPHER

Cipher algorithms:

- RTE_CRYPTO_CIPHER_NULL

Authentication algorithms:

- RTE_CRYPTO_AUTH_NULL

8.2 Limitations

- Only in-place is currently supported (destination address is the same as source address).

8.3 Installation

The Null Crypto PMD is enabled and built by default in both the Linux and FreeBSD builds.

8.4 Initialization

To use the PMD in an application, user must:

- Call rte_vdev_init("crypto_null") within the application.
- Use –vdev="crypto_null" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

- socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).
- max_nb_queue_pairs: Specify the maximum number of queue pairs in the device (8 by default).
- max_nb_sessions: Specify the maximum number of sessions that can be created (2048 by default).

Example:

```
./l2fwd-crypto -l 6 -n 4 --vdev="crypto_null,socket_id=1,max_nb_sessions=128"
```

CRYPTODEV SCHEDULER POLL MODE DRIVER LIBRARY

Scheduler PMD is a software crypto PMD, which has the capabilities of attaching hardware and/or software cryptodevs, and distributes ingress crypto ops among them in a certain manner.
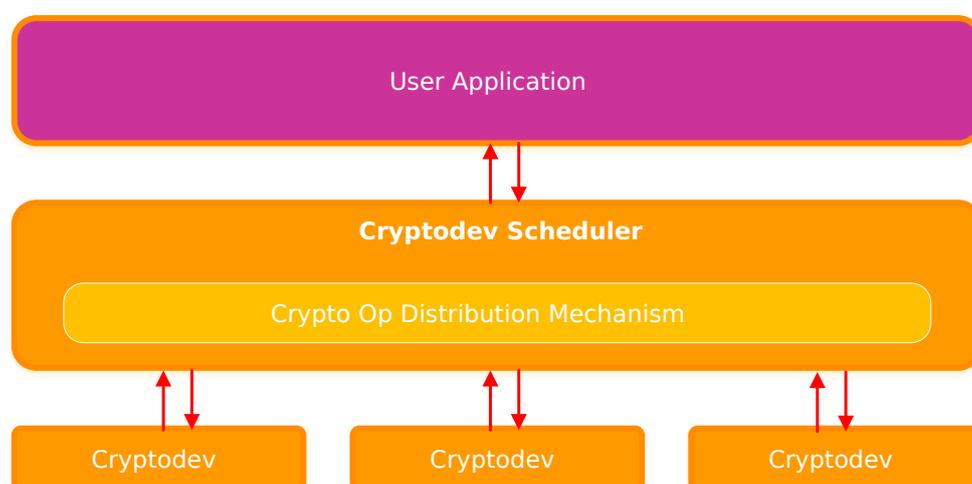

Fig. 9.1: Cryptodev Scheduler Overview

The Cryptodev Scheduler PMD library (**librte_pmd_crypto_scheduler**) acts as a software crypto PMD and shares the same API provided by librte_cryptodev. The PMD supports attaching multiple crypto PMDs, software or hardware, as slaves, and distributes the crypto workload to them with certain behavior. The behaviors are categorizes as different "modes". Basically, a scheduling mode defines certain actions for scheduling crypto ops to its slaves.

The librte_pmd_crypto_scheduler library exports a C API which provides an API for attaching/detaching slaves, set/get scheduling modes, and enable/disable crypto ops reordering.

9.1 Limitations

- Sessionless crypto operation is not supported
- OOP crypto operation is not supported when the crypto op reordering feature is enabled.

9.2 Installation

To build DPDK with CRYTPO_SCHEDULER_PMD the user is required to set CON-FIG_RTE_LIBRTE_PMD_CRYPTO_SCHEDULER=y in config/common_base, and recompile DPDK

9.3 Initialization

To use the PMD in an application, user must:

- Call rte_vdev_init("crpyto_scheduler") within the application.
- Use –vdev="crpyto_scheduler" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

- socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).
- max_nb_sessions: Specify the maximum number of sessions that can be created. This value may be overwritten internally if there are too many devices are attached.
- slave: If a cryptodev has been initialized with specific name, it can be attached to the scheduler using this parameter, simply filling the name here. Multiple cryptodevs can be attached initially by presenting this parameter multiple times.
- mode: Specify the scheduling mode of the PMD. The supported scheduling mode parameter values are specified in the "Cryptodev Scheduler Modes Overview" section.
- ordering: Specify the status of the crypto operations ordering feature. The value of this parameter can be "enable" or "disable". This feature is disabled by default.

Example:

```
... --vdev "crypto_aesni_mb_pmd,name=aesni_mb_1" --vdev "crypto_aesni_mb_pmd,name=aesni_mb_2" -
```

Note:

- The scheduler cryptodev cannot be started unless the scheduling mode is set and at least one slave is attached. Also, to configure the scheduler in the run-time, like attach/detach slave(s), change scheduling mode, or enable/disable crypto op ordering, one should stop the scheduler first, otherwise an error will be returned.

- The crypto op reordering feature requires using the userdata field of every mbuf to be processed to store temporary data. By the end of processing, the field is set to pointing to NULL, any previously stored value of this field will be lost.

9.4 Cryptodev Scheduler Modes Overview

Currently the Crypto Scheduler PMD library supports following modes of operation:

- **CDEV_SCHED_MODE_ROUNDROBIN:**

Initialization mode parameter: **round-robin**

Round-robin mode, which distributes the enqueued burst of crypto ops among its slaves in a round-robin manner. This mode may help to fill the throughput gap between the physical core and the existing cryptodevs to increase the overall performance.

- **CDEV_SCHED_MODE_PKT_SIZE_DISTR:**

Initialization mode parameter: **packet-size-distr**

Packet-size based distribution mode, which works with 2 slaves, the primary slave and the secondary slave, and distributes the enqueued crypto operations to them based on their data lengths. A crypto operation will be distributed to the primary slave if its data length is equal to or bigger than the designated threshold, otherwise it will be handled by the secondary slave.

A typical usecase in this mode is with the QAT cryptodev as the primary and a software cryptodev as the secondary slave. This may help applications to process additional crypto workload than what the QAT cryptodev can handle on its own, by making use of the available CPU cycles to deal with smaller crypto workloads.

The threshold is set to 128 bytes by default. It can be updated by calling function **rte_cryptodev_scheduler_option_set**. The parameter of **option_type** must be **CDEV_SCHED_OPTION_THRESHOLD** and **option** should point to a rte_cryptodev_scheduler_threshold_option structure filled with appropriate threshold value. Please NOTE this threshold has be a power-of-2 unsigned integer.

- **CDEV_SCHED_MODE_FAILOVER:**

Initialization mode parameter: **fail-over**

Fail-over mode, which works with 2 slaves, the primary slave and the secondary slave. In this mode, the scheduler will enqueue the incoming crypto operation burst to the primary slave. When one or more crypto operations fail to be enqueued, then they will be enqueued to the secondary slave.

SNOW 3G CRYPTO POLL MODE DRIVER

The SNOW 3G PMD (**librte_pmd_snow3g**) provides poll mode crypto driver support for utilizing Intel Libsso library, which implements F8 and F9 functions for SNOW 3G UEA2 cipher and UIA2 hash algorithms.

10.1 Features

SNOW 3G PMD has support for:

Cipher algorithm:

- RTE_CRYPTO_CIPHER_SNOW3G_UEA2

Authentication algorithm:

- RTE_CRYPTO_AUTH_SNOW3G_UIA2

10.2 Limitations

- Chained mbufs are not supported.

- SNOW 3G (UIA2) supported only if hash offset field is byte-aligned.

- In-place bit-level operations for SNOW 3G (UEA2) are not supported (if length and/or offset of data to be ciphered is not byte-aligned).

10.3 Installation

To build DPDK with the SNOW3G_PMD the user is required to download the export controlled `libsso_snow3g` library, by requesting it from https://networkbuilders.intel.com/network-technologies/dpdk. Once approval has been granted, the user needs to log in https://networkbuilders.intel.com/dpdklogin and click on "Snow3G Bit Stream crypto library" link, to download the library. After downloading the library, the user needs to unpack and compile it on their system before building DPDK:

```
make snow3G
```

10.4 Initialization

In order to enable this virtual crypto PMD, user must:

- Export the environmental variable LIBSSO_SNOW3G_PATH with the path where the library was extracted (snow3g folder).

- Build the LIBSSO_SNOW3G library (explained in Installation section).

- Set CONFIG_RTE_LIBRTE_PMD_SNOW3G=y in config/common_base.

To use the PMD in an application, user must:

- Call rte_vdev_init("crypto_snow3g") within the application.

- Use –vdev="crypto_snow3g" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

- socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).

- max_nb_queue_pairs: Specify the maximum number of queue pairs in the device (8 by default).

- max_nb_sessions: Specify the maximum number of sessions that can be created (2048 by default).

Example:

```
./l2fwd-crypto -l 6 -n 4 --vdev="crypto_snow3g,socket_id=1,max_nb_sessions=128"
```

INTEL(R) QUICKASSIST (QAT) CRYPTO POLL MODE DRIVER

The QAT PMD provides poll mode crypto driver support for the following hardware accelerator devices:

- Intel QuickAssist Technology DH895xCC

- Intel QuickAssist Technology C62x

- Intel QuickAssist Technology C3xxx

- Intel QuickAssist Technology D15xx

11.1 Features

The QAT PMD has support for:

Cipher algorithms:

- RTE_CRYPTO_CIPHER_3DES_CBC

- RTE_CRYPTO_CIPHER_3DES_CTR

- RTE_CRYPTO_CIPHER_AES128_CBC

- RTE_CRYPTO_CIPHER_AES192_CBC

- RTE_CRYPTO_CIPHER_AES256_CBC

- RTE_CRYPTO_CIPHER_AES128_CTR

- RTE_CRYPTO_CIPHER_AES192_CTR

- RTE_CRYPTO_CIPHER_AES256_CTR

- RTE_CRYPTO_CIPHER_SNOW3G_UEA2

- RTE_CRYPTO_CIPHER_AES_GCM

- RTE_CRYPTO_CIPHER_NULL

- RTE_CRYPTO_CIPHER_KASUMI_F8

- RTE_CRYPTO_CIPHER_DES_CBC

- RTE_CRYPTO_CIPHER_AES_DOCSISBPI

- RTE_CRYPTO_CIPHER_DES_DOCSISBPI

- RTE_CRYPTO_CIPHER_ZUC_EEA3

Hash algorithms:

- `RTE_CRYPTO_AUTH_SHA1_HMAC`

- `RTE_CRYPTO_AUTH_SHA224_HMAC`

- `RTE_CRYPTO_AUTH_SHA256_HMAC`

- `RTE_CRYPTO_AUTH_SHA384_HMAC`

- `RTE_CRYPTO_AUTH_SHA512_HMAC`

- `RTE_CRYPTO_AUTH_AES_XCBC_MAC`

- `RTE_CRYPTO_AUTH_SNOW3G_UIA2`

- `RTE_CRYPTO_AUTH_MD5_HMAC`

- `RTE_CRYPTO_AUTH_NULL`

- `RTE_CRYPTO_AUTH_KASUMI_F9`

- `RTE_CRYPTO_AUTH_AES_GMAC`

- `RTE_CRYPTO_AUTH_ZUC_EIA3`

11.2 Limitations

- Hash only is not supported except SNOW 3G UIA2 and KASUMI F9.

- Only supports the session-oriented API implementation (session-less APIs are not supported).

- SNOW 3G (UEA2), KASUMI (F8) and ZUC (EEA3) supported only if cipher length and offset fields are byte-multiple.

- SNOW 3G (UIA2), KASUMI (F9) and ZUC (EIA3) supported only if hash length and offset fields are byte-multiple.

- No BSD support as BSD QAT kernel driver not available.

- ZUC EEA3/EIA3 is not supported by dh895xcc devices

- Maximum additional authenticated data (AAD) for GCM is 240 bytes long.

11.3 Installation

To enable QAT in DPDK, follow the instructions for modifying the compile-time configuration file as described here.

Quick instructions are as follows:

```
cd to the top-level DPDK directory
make config T=x86_64-native-linuxapp-gcc
sed -i 's,\(CONFIG_RTE_LIBRTE_PMD_QAT\)=n,\1=y,' build/.config
make
```

To use the DPDK QAT PMD an SRIOV-enabled QAT kernel driver is required. The VF devices exposed by this driver will be used by the QAT PMD. The devices and available kernel drivers and device ids are :

Table 11.1: QAT devices and drivers

Device	Driver	Kernel Module	Pci Driver	PF Did	Num PFs	Vf Did	VFs per PF
DH895xCC	01.org	icp_qa_al	n/a	435	1	443	32
DH895xCC	4.4+	qat_dh895xcc	dh895xcc	435	1	443	32
C62x	4.5+	qat_c62x	c6xx	37c8	3	37c9	16
C3xxx	4.5+	qat_c3xxx	c3xxx	19e2	1	19e3	16
D15xx	p	qat_d15xx	d15xx	6f54	1	6f55	16

The `Driver` column indicates either the Linux kernel version in which support for this device was introduced or a driver available on Intel's 01.org website. There are both linux and 01.org kernel drivers available for some devices. p = release pending.

If you are running on a kernel which includes a driver for your device, see *Installation using kernel.org driver* below. Otherwise see *Installation using 01.org QAT driver*.

11.4 Installation using kernel.org driver

The examples below are based on the C62x device, if you have a different device use the corresponding values in the above table.

In BIOS ensure that SRIOV is enabled and either:

- Disable VT-d or

- Enable VT-d and set `"intel_iommu=on iommu=pt"` in the grub file.

Check that the QAT driver is loaded on your system, by executing:

```
lsmod | grep qa
```

You should see the kernel module for your device listed, e.g.:

```
qat_c62x                5626  0
intel_qat              82336  1 qat_c62x
```

Next, you need to expose the Virtual Functions (VFs) using the sysfs file system.

First find the BDFs (Bus-Device-Function) of the physical functions (PFs) of your device, e.g.:

```
lspci -d : 37c8
```

You should see output similar to:

```
1a:00.0 Co-processor: Intel Corporation Device 37c8
3d:00.0 Co-processor: Intel Corporation Device 37c8
3f:00.0 Co-processor: Intel Corporation Device 37c8
```

Enable the VFs for each PF by echoing the number of VFs per PF to the pci driver:

```
echo 16 > /sys/bus/pci/drivers/c6xx/0000:1a:00.0/sriov_numvfs
echo 16 > /sys/bus/pci/drivers/c6xx/0000:3d:00.0/sriov_numvfs
echo 16 > /sys/bus/pci/drivers/c6xx/0000:3f:00.0/sriov_numvfs
```

Check that the VFs are available for use. For example `lspci -d:37c9` should list 48 VF devices available for a `C62x` device.

To complete the installation follow the instructions in *Binding the available VFs to the DPDK UIO driver*.

Note: If the QAT kernel modules are not loaded and you see an error like `Failed to load MMP firmware qat_895xcc_mmp.bin` in kernel logs, this may be as a result of not using a distribution, but just updating the kernel directly.

Download firmware from the kernel firmware repo.

Copy qat binaries to `/lib/firmware`:

```
cp qat_895xcc.bin /lib/firmware
cp qat_895xcc_mmp.bin /lib/firmware
```

Change to your linux source root directory and start the qat kernel modules:

```
insmod ./drivers/crypto/qat/qat_common/intel_qat.ko
insmod ./drivers/crypto/qat/qat_dh895xcc/qat_dh895xcc.ko
```

Note: If you see the following warning in `/var/log/messages` it can be ignored: `IOMMU should be enabled for SR-IOV to work correctly.`

11.5 Installation using 01.org QAT driver

Download the latest QuickAssist Technology Driver from 01.org. Consult the *Getting Started Guide* at the same URL for further information.

The steps below assume you are:

- Building on a platform with one `DH895xCC` device.

- Using package `qatmux.l.2.3.0-34.tgz`.

- On Fedora21 kernel `3.17.4-301.fc21.x86_64`.

In the BIOS ensure that SRIOV is enabled and VT-d is disabled.

Uninstall any existing QAT driver, for example by running:

- `./installer.sh uninstall` in the directory where originally installed.

- or `rmmod qat_dh895xcc; rmmod intel_qat`.

Build and install the SRIOV-enabled QAT driver:

```
mkdir /QAT
cd /QAT

# Copy qatmux.l.2.3.0-34.tgz to this location
tar zxof qatmux.l.2.3.0-34.tgz

export ICP_WITHOUT_IOMMU=1
./installer.sh install QAT1.6 host
```

You can use `cat /proc/icp_dh895xcc_dev0/version` to confirm the driver is correctly installed. You can use `lspci -d:443` to confirm the of the 32 VF devices available per `DH895xCC` device.

To complete the installation - follow instructions in *Binding the available VFs to the DPDK UIO driver*.

Note: If using a later kernel and the build fails with an error relating to `strict_stroul` not being available apply the following patch:

```
/QAT/QAT1.6/quickassist/utilities/downloader/Target_CoreLibs/uclo/include/linux/uclo_platform.h
+ #if LINUX_VERSION_CODE >= KERNEL_VERSION(3,18,5)
+ #define STR_TO_64(str, base, num, endPtr) {endPtr=NULL; if (kstrtoul((str), (base), (num))) p
+ #else
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,38)
#define STR_TO_64(str, base, num, endPtr) {endPtr=NULL; if (strict_strtoull((str), (base), (num
#else
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,25)
#define STR_TO_64(str, base, num, endPtr) {endPtr=NULL; strict_strtoll((str), (base), (num));}
#else
#define STR_TO_64(str, base, num, endPtr)                                      \
      do {                                                                     \
            if (str[0] == '-')                                                 \
            {                                                                  \
                  *(num) = -(simple_strtoull((str+1), &(endPtr), (base)));     \
            }else {                                                            \
                  *(num) = simple_strtoull((str), &(endPtr), (base));          \
            }                                                                  \
      } while(0)
+ #endif
#endif
#endif
```

Note: If the build fails due to missing header files you may need to do following:

```
sudo yum install zlib-devel
sudo yum install openssl-devel
```

Note: If the build or install fails due to mismatching kernel sources you may need to do the following:

```
sudo yum install kernel-headers-`uname -r`
sudo yum install kernel-src-`uname -r`
sudo yum install kernel-devel-`uname -r`
```

11.6 Binding the available VFs to the DPDK UIO driver

Unbind the VFs from the stock driver so they can be bound to the uio driver.

11.6.1 For an Intel(R) QuickAssist Technology DH895xCC device

The unbind command below assumes BDFs of 03:01.00-03:04.07, if your VFs are different adjust the unbind command below:

```
for device in $(seq 1 4); do \
    for fn in $(seq 0 7); do \
        echo -n 0000:03:0${device}.${fn} > \
        /sys/bus/pci/devices/0000\:03\:0${device}.${fn}/driver/unbind; \
    done; \
done
```

11.6.2 For an Intel(R) QuickAssist Technology C62x device

The unbind command below assumes `BDFs` of `1a:01.00-1a:02.07`, `3d:01.00-3d:02.07` and `3f:01.00-3f:02.07`, if your VFs are different adjust the unbind command below:

```
for device in $(seq 1 2); do \
    for fn in $(seq 0 7); do \
        echo -n 0000:1a:0${device}.${fn} > \
        /sys/bus/pci/devices/0000\:1a\:0${device}.${fn}/driver/unbind; \

        echo -n 0000:3d:0${device}.${fn} > \
        /sys/bus/pci/devices/0000\:3d\:0${device}.${fn}/driver/unbind; \

        echo -n 0000:3f:0${device}.${fn} > \
        /sys/bus/pci/devices/0000\:3f\:0${device}.${fn}/driver/unbind; \
    done; \
done
```

11.6.3 For Intel(R) QuickAssist Technology C3xxx or D15xx device

The unbind command below assumes `BDFs` of `01:01.00-01:02.07`, if your VFs are different adjust the unbind command below:

```
for device in $(seq 1 2); do \
    for fn in $(seq 0 7); do \
        echo -n 0000:01:0${device}.${fn} > \
        /sys/bus/pci/devices/0000\:01\:0${device}.${fn}/driver/unbind; \
    done; \
done
```

11.6.4 Bind to the DPDK uio driver

Install the DPDK igb_uio driver, bind the VF PCI Device id to it and use lspci to confirm the VF devices are now in use by igb_uio kernel driver, e.g. for the C62x device:

```
cd to the top-level DPDK directory
modprobe uio
insmod ./build/kmod/igb_uio.ko
echo "8086 37c9" > /sys/bus/pci/drivers/igb_uio/new_id
lspci -vvd:37c9
```

Another way to bind the VFs to the DPDK UIO driver is by using the `dpdk-devbind.py` script:

```
cd to the top-level DPDK directory
./usertools/dpdk-devbind.py -b igb_uio 0000:03:01.1
```

ZUC CRYPTO POLL MODE DRIVER

The ZUC PMD (**librte_pmd_zuc**) provides poll mode crypto driver support for utilizing Intel Libsso library, which implements F8 and F9 functions for ZUC EEA3 cipher and EIA3 hash algorithms.

12.1 Features

ZUC PMD has support for:

Cipher algorithm:

- RTE_CRYPTO_CIPHER_ZUC_EEA3

Authentication algorithm:

- RTE_CRYPTO_AUTH_ZUC_EIA3

12.2 Limitations

- Chained mbufs are not supported.

- ZUC (EIA3) supported only if hash offset field is byte-aligned.

- ZUC (EEA3) supported only if cipher length, cipher offset fields are byte-aligned.

- ZUC PMD cannot be built as a shared library, due to limitations in in the underlying library.

12.3 Installation

To build DPDK with the ZUC_PMD the user is required to download the export controlled `libsso_zuc` library, by requesting it from https://networkbuilders.intel.com/network-technologies/dpdk. Once approval has been granted, the user needs to log in https://networkbuilders.intel.com/dpdklogin and click on "ZUC Library" link, to download the library. After downloading the library, the user needs to unpack and compile it on their system before building DPDK:

```
make
```

12.4 Initialization

In order to enable this virtual crypto PMD, user must:

- Export the environmental variable LIBSSO_ZUC_PATH with the path where the library was extracted (zuc folder).

- Build the LIBSSO_ZUC library (explained in Installation section).

- Build DPDK as follows:

```
make config T=x86_64-native-linuxapp-gcc
sed -i 's,\(CONFIG_RTE_LIBRTE_PMD_ZUC\)=n,\1=y,' build/.config
make
```

To use the PMD in an application, user must:

- Call rte_vdev_init("crypto_zuc") within the application.

- Use –vdev="crypto_zuc" in the EAL options, which will call rte_vdev_init() internally.

The following parameters (all optional) can be provided in the previous two calls:

- socket_id: Specify the socket where the memory for the device is going to be allocated (by default, socket_id will be the socket where the core that is creating the PMD is running on).

- max_nb_queue_pairs: Specify the maximum number of queue pairs in the device (8 by default).

- max_nb_sessions: Specify the maximum number of sessions that can be created (2048 by default).

Example:

```
./l2fwd-crypto -l 6 -n 4 --vdev="crypto_zuc,socket_id=1,max_nb_sessions=128"
```