



**DPDK**

DATA PLANE DEVELOPMENT KIT

## **Event Device Drivers**

*Release 17.05.2*

September 07, 2017

## CONTENTS

<b>1</b>	<b>Software Eventdev Poll Mode Driver</b>	<b>2</b>
1.1	Features . . . . .	2
1.2	Configuration and Options . . . . .	2
1.3	Limitations . . . . .	3
<b>2</b>	<b>OCTEONTX SSOVF Eventdev Driver</b>	<b>5</b>
2.1	Features . . . . .	5
2.2	Supported OCTEONTX SoCs . . . . .	5
2.3	Prerequisites . . . . .	5
2.4	Pre-Installation Configuration . . . . .	6
2.5	Initialization . . . . .	6
2.6	Limitations . . . . .	7

The following are a list of event device PMDs, which can be used from an application through the eventdev API.

## SOFTWARE EVENTDEV POLL MODE DRIVER

The software eventdev is an implementation of the eventdev API, that provides a wide range of the eventdev features. The eventdev relies on a CPU core to perform event scheduling.

### 1.1 Features

The software eventdev implements many features in the eventdev API;

#### Queues

- Atomic
- Ordered
- Parallel
- Single-Link

#### Ports

- Load balanced (for Atomic, Ordered, Parallel queues)
- Single Link (for single-link queues)

#### Event Priorities

- Each event has a priority, which can be used to provide basic QoS

### 1.2 Configuration and Options

The software eventdev is a vdev device, and as such can be created from the application code, or from the EAL command line:

- Call `rte_vdev_init("event_sw0")` from the application
- Use `--vdev="event_sw0"` in the EAL options, which will call `rte_vdev_init()` internally

Example:

```
./your_eventdev_application --vdev="event_sw0"
```

### 1.2.1 Scheduling Quanta

The scheduling quanta sets the number of events that the device attempts to schedule before returning to the application from the `rte_event_schedule()` function. Note that is a *hint* only, and that fewer or more events may be scheduled in a given iteration.

The scheduling quanta can be set using a string argument to the vdev create call:

```
--vdev="event_sw0,sched_quanta=64"
```

### 1.2.2 Credit Quanta

The credit quanta is the number of credits that a port will fetch at a time from the instance's credit pool. Higher numbers will cause less overhead in the atomic credit fetch code, however it also reduces the overall number of credits in the system faster. A balanced number (eg 32) ensures that only small numbers of credits are pre-allocated at a time, while also mitigating performance impact of the atomics.

Experimentation with higher values may provide minor performance improvements, at the cost of the whole system having less credits. On the other hand, reducing the quanta may cause measurable performance impact but provide the system with a higher number of credits at all times.

A value of 32 seems a good balance however your specific application may benefit from a higher or reduced quanta size, experimentation is required to verify possible gains.

```
--vdev="event_sw0,credit_quanta=64"
```

## 1.3 Limitations

The software eventdev implementation has a few limitations. The reason for these limitations is usually that the performance impact of supporting the feature would be significant.

### 1.3.1 "All Types" Queues

The software eventdev does not support creating queues that handle all types of traffic. An eventdev with this capability allows enqueueing Atomic, Ordered and Parallel traffic to the same queue, but scheduling each of them appropriately.

The reason to not allow Atomic, Ordered and Parallel event types in the same queue is that it causes excessive branching in the code to enqueue packets to the queue, causing a significant performance impact.

The `RTE_EVENT_DEV_CAP_QUEUE_ALL_TYPES` flag is not set in the `event_dev_cap` field of the `rte_event_dev_info` struct for the software eventdev.

### 1.3.2 Distributed Scheduler

The software eventdev is a centralized scheduler, requiring the `rte_event_schedule()` function to be called by a CPU core to perform the required event distribution. This is not really a limitation but rather a design decision.

The `RTE_EVENT_DEV_CAP_DISTRIBUTED_SCHED` flag is not set in the `event_dev_cap` field of the `rte_event_dev_info` struct for the software eventdev.

### 1.3.3 Dequeue Timeout

The eventdev API supports a timeout when dequeuing packets using the `rte_event_dequeue_burst` function. This allows a core to wait for an event to arrive, or until `timeout` number of ticks have passed. Timeout ticks is not supported by the software eventdev for performance reasons.

## OCTEONTX SSOVF EVENTDEV DRIVER

The OCTEONTX SSOVF PMD (`librte_pmd_octeontx_ssovf`) provides poll mode eventdev driver support for the inbuilt event device found in the **Cavium OCTEONTX** SoC family as well as their virtual functions (VF) in SR-IOV context.

More information can be found at [Cavium Networks Official Website](#).

### 2.1 Features

Features of the OCTEONTX SSOVF PMD are:

- 64 Event queues
- 32 Event ports
- HW event scheduler
- Supports 1M flows per event queue
- Flow based event pipelining
- Flow pinning support in flow based event pipelining
- Queue based event pipelining
- Supports ATOMIC, ORDERED, PARALLEL schedule types per flow
- Event scheduling QoS based on event queue priority
- Open system with configurable amount of outstanding events
- HW accelerated dequeue timeout support to enable power management
- SR-IOV VF

### 2.2 Supported OCTEONTX SoCs

- CN83xx

### 2.3 Prerequisites

There are three main pre-perquisites for executing SSOVF PMD on a OCTEONTX compatible board:

## 1. OCTEONTX Linux kernel PF driver for Network acceleration HW blocks

The OCTEONTX Linux kernel drivers (including the required PF driver for the SSOVF) are available on Github at [octeontx-kmod](#) along with build, install and dpdk usage instructions.

## 2. ARM64 Tool Chain

For example, the *aarch64* Linaro Toolchain, which can be obtained from [here](#).

## 3. Rootfile system

Any *aarch64* supporting filesystem can be used. For example, Ubuntu 15.10 (Wily) or 16.04 LTS (Xenial) userland which can be obtained from <http://cdimage.ubuntu.com/ubuntu-base/releases/16.04/release/ubuntu-base-16.04.1-base-arm64.tar.gz>.

As an alternative method, SSOVF PMD can also be executed using images provided as part of SDK from Cavium. The SDK includes all the above prerequisites necessary to bring up a OCTEONTX board.

SDK and related information can be obtained from: [Cavium support site](#).

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

## 2.4 Pre-Installation Configuration

### 2.4.1 Config File Options

The following options can be modified in the `config` file. Please note that enabling debugging options may affect system performance.

- `CONFIG_RTE_LIBRTE_PMD_OCTEONTX_SSOVF` (default `y`)  
Toggle compilation of the `librte_pmd_octeontx_ssovf` driver.
- `CONFIG_RTE_LIBRTE_PMD_OCTEONTX_SSOVF_DEBUG` (default `n`)  
Toggle display of generic debugging messages

### 2.4.2 Driver Compilation

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

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

## 2.5 Initialization

The `octeontx eventdev` is exposed as a `vdev` device which consists of a set of SSO group and work-slot PCIe VF devices. On EAL initialization, SSO PCIe VF devices will be probed and then the `vdev` device can be created from the application code, or from the EAL command line based on the number of probed/bound SSO PCIe VF device to DPDK by

- Invoking `rte_vdev_init("event_octeontx")` from the application
- Using `--vdev="event_octeontx"` in the EAL options, which will call `rte_vdev_init()` internally

Example:

```
./your_eventdev_application --vdev="event_octeontx"
```

## 2.6 Limitations

### 2.6.1 Burst mode support

Burst mode is not supported. Dequeue and Enqueue functions accepts only single event at a time.