



# DPDK

DATA PLANE DEVELOPMENT KIT

## Release Notes

*Release 17.11.10*

Feb 27, 2020

## CONTENTS

<b>1</b>	<b>Description of Release</b>	<b>1</b>
<b>2</b>	<b>DPDK Release 17.11</b>	<b>2</b>
<b>3</b>	<b>DPDK Release 17.08</b>	<b>56</b>
<b>4</b>	<b>DPDK Release 17.05</b>	<b>66</b>
<b>5</b>	<b>DPDK Release 17.02</b>	<b>79</b>
<b>6</b>	<b>DPDK Release 16.11</b>	<b>91</b>
<b>7</b>	<b>DPDK Release 16.07</b>	<b>101</b>
<b>8</b>	<b>DPDK Release 16.04</b>	<b>110</b>
<b>9</b>	<b>DPDK Release 2.2</b>	<b>121</b>
<b>10</b>	<b>DPDK Release 2.1</b>	<b>131</b>
<b>11</b>	<b>DPDK Release 2.0</b>	<b>145</b>
<b>12</b>	<b>DPDK Release 1.8</b>	<b>147</b>
<b>13</b>	<b>Known Issues and Limitations in Legacy Releases</b>	<b>148</b>
<b>14</b>	<b>ABI and API Deprecation</b>	<b>160</b>

## **DESCRIPTION OF RELEASE**

This document contains the release notes for Data Plane Development Kit (DPDK) release version 17.11.10 and previous releases.

It lists new features, fixed bugs, API and ABI changes and known issues.

For instructions on compiling and running the release, see the DPDK Getting Started Guide.

## DPDK RELEASE 17.11

### 2.1 New Features

- **Extended `port_id` range from `uint8_t` to `uint16_t`.**

Increased the `port_id` range from 8 bits to 16 bits in order to support more than 256 ports in DPDK. All `ethdev` APIs which have `port_id` as parameter have been changed.

- **Modified the return type of `rte_eth_stats_reset`.**

Changed return type of `rte_eth_stats_reset` from `void` to `int` so that the caller can determine whether a device supports the operation or not and if the operation was carried out.

- **Added a new driver for Marvell Armada 7k/8k devices.**

Added the new `mrvl` net driver for Marvell Armada 7k/8k devices. See the `./nics/mrvl` NIC guide for more details on this new driver.

- **Updated `mlx4` driver.**

Updated the `mlx4` driver including the following changes:

- Isolated mode (`rte_flow`) can now be enabled anytime, not only during initial device configuration.
- Flow rules now support up to 4096 priority levels usable at will by applications.
- Enhanced error message to help debugging invalid/unsupported flow rules.
- Flow rules matching all multicast and promiscuous traffic are now allowed.
- No more software restrictions on flow rules with the RSS action, their configuration is much more flexible.
- Significantly reduced memory footprint for Rx and Tx queue objects.
- While supported, UDP RSS is temporarily disabled due to a remaining issue with its support in the Linux kernel.
- The new RSS implementation does not automatically spread traffic according to the inner packet of VXLAN frames anymore, only the outer one (like other PMDs).
- Partial (Tx only) support for secondary processes was broken and had to be removed.

- Refactored driver to get rid of dependency on the components provided by Mellanox OFED and instead rely on the current and public rdma-core package and Linux version from now on.
- Removed compile-time limitation on number of device instances the PMD can support.

- **Updated mlx5 driver.**

Updated the mlx5 driver including the following changes:

- Enabled the PMD to run on top of upstream Linux kernel and rdma-core libs, removing the dependency on specific Mellanox OFED libraries.
- Improved PMD latency performance.
- Improved PMD memory footprint.
- Added support for vectorized Rx/Tx burst for ARMv8.
- Added support for secondary process.
- Added support for flow counters.
- Added support for Rx hardware timestamp offload.
- Added support for device removal event.

- **Added SoftNIC PMD.**

Added a new SoftNIC PMD. This virtual device provides applications with software fallback support for traffic management.

- **Added support for NXP DPAA Devices.**

Added support for NXP's DPAA devices - LS104x series. This includes:

- DPAA Bus driver
- DPAA Mempool driver for supporting offloaded packet memory pool
- DPAA PMD for DPAA devices

See the `../nics/dpaa` document for more details of this new driver.

- **Updated support for Cavium OCTEONTX Device.**

Updated support for Cavium's OCTEONTX device (CN83xx). This includes:

- OCTEONTX Mempool driver for supporting offloaded packet memory pool
- OCTEONTX Ethdev PMD
- OCTEONTX Eventdev-Ethdev Rx adapter

See the `../nics/octeontx` document for more details of this new driver.

- **Added PF support to the Netronome NFP PMD.**

Added PF support to the Netronome NFP PMD. Previously the NFP PMD only supported VFs. PF support is just as a basic DPDK port and has no VF management yet.

PF support comes with firmware upload support which allows the PMD to independently work from kernel netdev NFP drivers.

NFP 4000 devices are also now supported along with previous 6000 devices.

- **Updated bnxt PMD.**

Major enhancements include:

- Support for Flow API
- Support for Tx and Rx descriptor status functions

- **Added bus agnostic functions to cryptodev for PMD initialization**

Added new PMD assist, bus independent, functions `rte_cryptodev_pmd_parse_input_args()`, `rte_cryptodev_pmd_create()` and `rte_cryptodev_pmd_destroy()` for drivers to manage creation and destruction of new device instances.

- **Updated QAT crypto PMD.**

Added several performance enhancements:

- Removed atomics from the internal queue pair structure.
- Added coalesce writes to HEAD CSR on response processing.
- Added coalesce writes to TAIL CSR on request processing.

In addition support was added for the AES CCM algorithm.

- **Updated the AESNI MB PMD.**

The AESNI MB PMD has been updated with additional support for:

- The DES CBC algorithm.
- The DES DOCSIS BPI algorithm.

This change requires version 0.47 of the IPsec Multi-buffer library. For more details see the `../cryptodevs/aesni_mb` documentation.

- **Updated the OpenSSL PMD.**

The OpenSSL PMD has been updated with additional support for:

- The DES CBC algorithm.
- The AES CCM algorithm.

- **Added NXP DPAA SEC crypto PMD.**

A new `dpaa_sec` hardware based crypto PMD for NXP DPAA devices has been added. See the `../cryptodevs/dpaa_sec` document for more details.

- **Added MRVL crypto PMD.**

A new crypto PMD has been added, which provides several ciphering and hashing algorithms. All cryptography operations use the MUSDK library crypto API. See the `../cryptodevs/mrvl` document for more details.

- **Add new benchmarking mode to dpdk-test-crypto-perf application.**

Added a new “PMD cyclecount” benchmark mode to the `dpdk-test-crypto-perf` application to display a detailed breakdown of CPU cycles used by hardware acceleration.

- **Added the Security Offload Library.**

Added an experimental library - `rte_security`. This provide security APIs for protocols like IPsec using inline ipsec offload to ethernet devices or full protocol offload with lookaside crypto devices.

See the `../prog_guide/rte_security` section of the DPDK Programmers Guide document for more information.

- **Updated the DPAA2\_SEC crypto driver to support `rte_security`.**

Updated the `dpaa2_sec` crypto PMD to support `rte_security` lookaside protocol offload for IPsec.

- **Updated the IXGBE ethernet driver to support `rte_security`.**

Updated `ixgbe` ethernet PMD to support `rte_security` inline IPsec offload.

- **Updated ipsec-secgw application to support `rte_security`.**

Updated the `ipsec-secgw` sample application to support `rte_security` actions for ipsec inline and full protocol offload using lookaside crypto offload.

- **Added IOMMU support to libvhost-user**

Implemented device IOTLB in the Vhost-user backend, and enabled Virtio's IOMMU feature. The feature is disabled by default, and can be enabled by setting `RTE_VHOST_USER_IOMMU_SUPPORT` flag at vhost device registration time.

- **Added the Event Ethernet Adapter Library.**

Added the Event Ethernet Adapter library. This library provides APIs for eventdev applications to configure the ethdev for eventdev packet flow.

- **Updated DPAA2 Event PMD for the Event Ethernet Adapter.**

Added support for the eventdev ethernet adapter for DPAA2.

- **Added Membership library (`rte_member`).**

Added a new data structure library called the Membership Library.

The Membership Library is an extension and generalization of a traditional filter (for example Bloom Filter) structure that has multiple usages in a wide variety of workloads and applications. In general, the Membership Library is a data structure that provides a “set-summary” and responds to set-membership queries whether a certain member belongs to a set(s).

The library provides APIs for DPDK applications to insert a new member, delete an existing member, and query the existence of a member in a given set, or a group of sets. For the case of a group of sets the library will return not only whether the element has been inserted in one of the sets but also which set it belongs to.

See the `../prog_guide/member_lib` documentation in the Programmers Guide, for more information.

- **Added the Generic Segmentation Offload Library.**

Added the Generic Segmentation Offload (GSO) library to enable applications to split large packets (e.g. MTU is 64KB) into small ones (e.g. MTU is 1500B). Supported packet types are:

- TCP/IPv4 packets.

- VxLAN packets, which must have an outer IPv4 header, and contain an inner TCP/IPv4 packet.
- GRE packets, which must contain an outer IPv4 header, and inner TCP/IPv4 headers.

The GSO library doesn't check if the input packets have correct checksums, and doesn't update checksums for output packets. Additionally, the GSO library doesn't process IP fragmented packets.

- **Added the Flow Classification Library.**

Added an experimental Flow Classification library to provide APIs for DPDK applications to classify an input packet by matching it against a set of flow rules. It uses the `librte_table` API to manage the flow rules.

## 2.2 Resolved Issues

- **Service core fails to call service callback due to atomic lock**

In a specific configuration of multi-thread unsafe services and service cores, a service core previously did not correctly release the atomic lock on the service. This would result in the cores polling the service, but it looked like another thread was executing the service callback. The logic for atomic locking of the services has been fixed and refactored for readability.

## 2.3 API Changes

- **Ethdev device name length increased.**

The size of internal device name has been increased to 64 characters to allow for storing longer bus specific names.

- **Removed the Ethdev RTE\_ETH\_DEV\_DETACHABLE flag.**

Removed the Ethdev `RTE_ETH_DEV_DETACHABLE` flag. This flag is not required anymore, with the new hotplug implementation. It has been removed from the ether library. Its semantics are now expressed at the bus and PMD level.

- **Service cores API updated for usability**

The service cores API has been changed, removing pointers from the API where possible, and instead using integer IDs to identify each service. This simplifies application code, aids debugging, and provides better encapsulation. A summary of the main changes made is as follows:

- Services identified by ID not by `rte_service_spec` pointer
- Reduced API surface by using `set` functions instead of enable/disable
- Reworked `rte_service_register` to provide the service ID to registrar
- Reworked start and stop APIs into `rte_service_runstate_set`
- Added API to set runstate of service implementation to indicate readiness

- **The following changes have been made in the mempool library**



- Moved `flags` datatype from `int` to `unsigned int` for `rte_mempool`.
- Removed `__rte_unused int flag param` from `rte_mempool_generic_put` and `rte_mempool_generic_get` API.
- Added `flags param` in `rte_mempool_xmem_size` and `rte_mempool_xmem_usage`.
- `rte_mem_phy2mch` was used in Xen dom0 to obtain the physical address; remove this API as Xen dom0 support was removed.

- **Added IOVA aliases related to physical address handling.**

Some data types, structure members and functions related to physical address handling are deprecated and have new aliases with IOVA wording. For example:

- `phys_addr_t` can be often replaced by `rte_iova_t` of same size.
- `RTE_BAD_PHYS_ADDR` is often replaced by `RTE_BAD_IOVA` of same value.
- `rte_memseg.phys_addr` is aliased with `rte_memseg.iova_addr`.
- `rte_mem_virt2phy()` can often be replaced by `rte_mem_virt2iova`.
- `rte_malloc_virt2phy` is aliased with `rte_malloc_virt2iova`.
- `rte_memzone.phys_addr` is aliased with `rte_memzone.iova`.
- `rte_mempool_objhdr.physaddr` is aliased with `rte_mempool_objhdr.iova`.
- `rte_mempool_memhdr.phys_addr` is aliased with `rte_mempool_memhdr.iova`.
- `rte_mempool_virt2phy()` can be replaced by `rte_mempool_virt2iova()`.
- `rte_mempool_populate_phys*()` are aliased with `rte_mempool_populate_iova*()`.
- `rte_mbuf.buf_physaddr` is aliased with `rte_mbuf.buf_iova`.
- `rte_mbuf_data_dma_addr*()` are aliased with `rte_mbuf_data_iova*()`.
- `rte_pktmbuf_mtophys*` are aliased with `rte_pktmbuf_iova*()`.

- **PCI bus API moved outside of the EAL**

The PCI bus previously implemented within the EAL has been moved. A first part has been added as an RTE library providing PCI helpers to parse device locations or other such utilities. A second part consisting of the actual bus driver has been moved to its proper subdirectory, without changing its functionalities.

As such, several PCI-related functions are not exposed by the EAL anymore:

- `rte_pci_detach`
- `rte_pci_dump`
- `rte_pci_ioport_map`
- `rte_pci_ioport_read`
- `rte_pci_ioport_unmap`
- `rte_pci_ioport_write`

- `rte_pci_map_device`
- `rte_pci_probe`
- `rte_pci_probe_one`
- `rte_pci_read_config`
- `rte_pci_register`
- `rte_pci_scan`
- `rte_pci_unmap_device`
- `rte_pci_unregister`
- `rte_pci_write_config`

These functions are made available either as part of `librte_pci` or `librte_bus_pci`.

- **Moved vdev bus APIs outside of the EAL**

Moved the following APIs from `librte_eal` to `librte_bus_vdev`:

- `rte_vdev_init`
- `rte_vdev_register`
- `rte_vdev_uninit`
- `rte_vdev_unregister`

- **Add return value to stats\_get dev op API**

The `stats_get dev op` API return value has been changed to be `int`. In this way PMDs can return an error value in case of failure at stats getting process time.

- **Modified the `rte_cryptodev_allocate_driver` function.**

Modified the `rte_cryptodev_allocate_driver()` function in the `cryptodev` library. An extra parameter `struct cryptodev_driver *crypto_drv` has been added.

- **Removed virtual device bus specific functions from `librte_cryptodev`.**

The functions `rte_cryptodev_vdev_parse_init_params()` and `rte_cryptodev_vdev_pmd_init()` have been removed from `librte_cryptodev` and have been replaced by non bus specific functions `rte_cryptodev_pmd_parse_input_args()` and `rte_cryptodev_pmd_create()`.

The `rte_cryptodev_create_vdev()` function was removed to avoid the dependency on `vdev` in `librte_cryptodev`; instead, users can call `rte_vdev_init()` directly.

- **Removed PCI device bus specific functions from `librte_cryptodev`.**

The functions `rte_cryptodev_pci_generic_probe()` and `rte_cryptodev_pci_generic_remove()` have been removed from `librte_cryptodev` and have been replaced by non bus specific functions `rte_cryptodev_pmd_create()` and `rte_cryptodev_pmd_destroy()`.

- **Removed deprecated functions to manage log level or type.**

The functions `rte_set_log_level()`, `rte_get_log_level()`, `rte_set_log_type()` and `rte_get_log_type()` have been removed.

They are respectively replaced by `rte_log_set_global_level()`, `rte_log_get_global_level()`, `rte_log_set_level()` and `rte_log_get_level()`.

- **Removed mbuf flags `PKT_RX_VLAN_PKT` and `PKT_RX_QINQ_PKT`.**

The mbuf flags `PKT_RX_VLAN_PKT` and `PKT_RX_QINQ_PKT` have been removed since their behavior was not properly described.

- **Added mbuf flags `PKT_RX_VLAN` and `PKT_RX_QINQ`.**

Two mbuf flags have been added to indicate that the VLAN identifier has been saved in the mbuf structure. For instance:

- If VLAN is not stripped and TCI is saved: `PKT_RX_VLAN`
- If VLAN is stripped and TCI is saved: `PKT_RX_VLAN | PKT_RX_VLAN_STRIPPED`

- **Modified the `vlan_offload_set_t` function prototype in the `ethdev` library.**

Modified the `vlan_offload_set_t` function prototype in the `ethdev` library. The return value has been changed from `void` to `int` so the caller can determine whether the backing device supports the operation or if the operation was successfully performed.

## 2.4 ABI Changes

- **Extended `port_id` range.**

The size of the field `port_id` in the `rte_eth_dev_data` structure has changed, as described in the *New Features* section above.

- **New parameter added to `rte_eth_dev`.**

A new parameter `security_ctx` has been added to `rte_eth_dev` to support security operations like IPsec inline.

- **New parameter added to `rte_cryptodev`.**

A new parameter `security_ctx` has been added to `rte_cryptodev` to support security operations like lookaside crypto.

## 2.5 Removed Items

- Xen dom0 in EAL has been removed, as well as the `xenvirt` PMD and `vhost_xen`.
- The crypto performance unit tests have been removed, replaced by the `dpdk-test-crypto-perf` application.

## 2.6 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```
librte_acl.so.2
+ librte_bitratestats.so.2
+ librte_bus_dpaa.so.1
+ librte_bus_fslmc.so.1
```

```
+ librte_bus_pci.so.1
+ librte_bus_vdev.so.1
  librte_cfgfile.so.2
  librte_cmdline.so.2
+ librte_cryptodev.so.4
  librte_distributor.so.1
+ librte_eal.so.6
+ librte_ethdev.so.8
+ librte_eventdev.so.3
+ librte_flow_classify.so.1
  librte_gro.so.1
+ librte_gso.so.1
  librte_hash.so.2
  librte_ip_frag.so.1
  librte_jobstats.so.1
  librte_kni.so.2
  librte_kvargs.so.1
  librte_latencystats.so.1
  librte_lpm.so.2
  librte_mbuf.so.3
+ librte_mempool.so.3
  librte_meter.so.1
  librte_metrics.so.1
  librte_net.so.1
+ librte_pci.so.1
+ librte_pdump.so.2
  librte_pipeline.so.3
+ librte_pmd_bnxt.so.2
+ librte_pmd_bond.so.2
+ librte_pmd_i40e.so.2
+ librte_pmd_ixgbe.so.2
  librte_pmd_ring.so.2
+ librte_pmd_softnic.so.1
+ librte_pmd_vhost.so.2
  librte_port.so.3
  librte_power.so.1
  librte_reorder.so.1
  librte_ring.so.1
  librte_sched.so.1
+ librte_security.so.1
+ librte_table.so.3
  librte_timer.so.1
  librte_vhost.so.3
```

## 2.7 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
  - CPU
    - \* Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
    - \* Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU D-1541 @ 2.10GHz
    - \* Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz

- \* Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz
- \* Intel(R) Xeon(R) CPU E5-2658 v2 @ 2.40GHz
- \* Intel(R) Xeon(R) CPU E5-2658 v3 @ 2.20GHz
- OS:
  - \* CentOS 7.2
  - \* Fedora 25
  - \* Fedora 26
  - \* FreeBSD 11
  - \* Red Hat Enterprise Linux Server release 7.3
  - \* SUSE Enterprise Linux 12
  - \* Wind River Linux 8
  - \* Ubuntu 16.04
  - \* Ubuntu 16.10
- NICs:
  - \* Intel(R) 82599ES 10 Gigabit Ethernet Controller
    - Firmware version: 0x61bf0001
    - Device id (pf/vf): 8086:10fb / 8086:10ed
    - Driver version: 5.2.3 (ixgbe)
  - \* Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
    - Firmware version: 0x800003e7
    - Device id (pf/vf): 8086:15ad / 8086:15a8
    - Driver version: 4.4.6 (ixgbe)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
    - Firmware version: 6.01 0x80003205
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 2.1.26 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
    - Firmware version: 6.01 0x80003204
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 2.1.26 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
    - Firmware version: 6.01 0x80003221
    - Device id (pf/vf): 8086:158b
    - Driver version: 2.1.26 (i40e)

- \* Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
  - Firmware version: 6.01 0x8000321c
  - Device id (pf/vf): 8086:1583 / 8086:154c
  - Driver version: 2.1.26 (i40e)
- \* Intel(R) Corporation I350 Gigabit Network Connection
  - Firmware version: 1.63, 0x80000dda
  - Device id (pf/vf): 8086:1521 / 8086:1520
  - Driver version: 5.3.0-k (igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
  - Platform details:
    - \* Intel(R) Xeon(R) CPU E5-2697A v4 @ 2.60GHz
    - \* Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz
    - \* Intel(R) Xeon(R) CPU E5-2640 @ 2.50GHz
    - \* Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz
  - OS:
    - \* Red Hat Enterprise Linux Server release 7.3 (Maipo)
    - \* Red Hat Enterprise Linux Server release 7.2 (Maipo)
    - \* Ubuntu 16.10
    - \* Ubuntu 16.04
    - \* Ubuntu 14.04
  - MLNX\_OFED: 4.2-1.0.0.0
  - NICs:
    - \* Mellanox(R) ConnectX(R)-3 Pro 40G MCX354A-FCC\_Ax (2x40G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1007
      - Firmware version: 2.42.5000
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.21.1000
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
      - Host interface: PCI Express 3.0 x8

- Device ID: 15b3:1013
- Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000

- \* Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.21.1000
- \* Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.21.1000
- \* Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.21.1000
- \* Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1017
  - Firmware version: 16.21.1000
- \* Mellanox(R) ConnectX-5 Ex EN 100G MCX516A-CDAT (2x100G)
  - Host interface: PCI Express 4.0 x16
  - Device ID: 15b3:1019
  - Firmware version: 16.21.1000
- ARM platforms with Mellanox(R) NICs combinations
  - Platform details:
    - \* Qualcomm ARM 1.1 2500MHz
  - OS:
    - \* Ubuntu 16.04
  - MLNX\_OFED: 4.2-1.0.0.0
  - NICs:
    - \* Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1015
      - Firmware version: 14.21.1000
    - \* Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
      - Host interface: PCI Express 3.0 x16
      - Device ID: 15b3:1017



· Firmware version: 16.21.1000

## 2.8 Fixes in 17.11 LTS Release

### 2.8.1 17.11.1

- app/procinfo: add compilation option in config
- app/testpmd: fix crash of txonly with multiple segments
- app/testpmd: fix flow director filter
- app/testpmd: fix flowgen forwarding offload flags
- app/testpmd: fix invalid Rx queue number setting
- app/testpmd: fix invalid Tx queue number setting
- app/testpmd: fix port configuration print
- app/testpmd: fix port id allocation
- app/testpmd: fix port index in RSS forward config
- app/testpmd: fix port topology in RSS forward config
- app/testpmd: fix port validation
- app/testpmd: remove xenvirt again
- bus/dpaa: fix ARM big endian build
- bus/dpaa: fix build when assert enabled
- bus/dpaa: fix default IOVA mode
- bus/fslmc: fix build with latest glibc
- bus/fslmc: fix the cplusplus macro closure
- bus/pci: fix interrupt handler type
- bus/pci: forbid IOVA mode if IOMMU address width too small
- bus/vdev: continue probing after a device failure
- cmdline: avoid garbage in unused fields of parsed result
- cmdline: fix dynamic tokens parsing
- cryptodev: add missing CPU flag string
- cryptodev: fix function prototype
- cryptodev: fix session pointer cast
- crypto/dpaa2\_sec: fix enum conversion for GCM
- crypto: fix pedantic compilation
- crypto/qat: fix allocation check and leak
- crypto/qat: fix null auth algo overwrite

- crypto/qat: fix out-of-bounds access
- crypto/qat: fix parameter type
- crypto/scheduler: fix strncpy
- doc: fix format in OpenSSL installation guide
- doc: fix lists of supported crypto algorithms
- drivers: change the deprecated memseg physaddr to IOVA
- eal/arm64: remove the braces in memory barrier macros
- eal/ppc64: revert arch-specific TSC freq query
- eal/ppc: remove the braces in memory barrier macros
- ethdev: fix link autonegotiation value
- ethdev: fix missing imissed counter in xstats
- ethdev: fix port data reset timing
- ethdev: fix port id allocation
- eventdev: fix doxygen comments
- eventdev: set error code in port link/unlink functions
- event/octeontx: fix Rx adapter port id mapping
- event/sw: fix debug logging config option
- event/sw: fix queue memory leak and multi-link bug
- examples/bond: check mbuf allocation
- examples/bond: fix vdev name
- examples/ip\_pipeline: fix timer period unit
- examples/ipsec-secgw: fix corner case for SPI value
- examples/ipsec-secgw: fix missing ingress flow attribute
- examples/ipsec-secgw: fix SPI byte order in flow item
- examples/ipsec-secgw: fix usage of incorrect port
- examples/l3fwd-power: fix frequency detection
- examples/l3fwd-power: fix Rx without interrupt
- examples/vhost: fix sending ARP packet to self
- examples/vhost: fix startup check
- flow\_classify: fix ISO C in exported header
- igb\_uio: allow multi-process access
- keepalive: fix state alignment
- kni: fix build dependency
- kni: fix build with kernel 4.15

- lib: fix missing includes in exported headers
- log: fix memory leak in regexp level set
- lpm: fix ARM big endian build
- malloc: fix end for bounded elements
- malloc: protect stats with lock
- mbuf: fix NULL freeing when debug enabled
- mbuf: fix performance of freeing with non atomic refcnt
- member: fix ISO C in exported header
- member: fix memory leak on error
- mempool: fix first memory area notification
- mempool: fix physical contiguous check
- mempool/octeontx: fix improper memory barrier
- mempool/octeontx: fix memory area registration
- mempool/octeontx: fix natural alignment being optimized out
- memzone: fix leak on allocation error
- mk: fix external build
- mk: remove TILE-Gx machine type
- mk: support renamed Makefile in external project
- net/bnxt: fix check for ether type
- net/bnxt: fix double increment of idx during Tx ring alloc
- net/bnxt: fix duplicate filter pattern creation error
- net/bnxt: fix duplicate pattern for 5tuple filter
- net/bnxt: fix group info usage
- net/bnxt: fix link speed setting with autoneg off
- net/bnxt: fix number of pools for RSS
- net/bnxt: fix return code in MAC address set
- net/bnxt: fix Rx checksum flags
- net/bnxt: fix size of Tx ring in HW
- net/bnxt: free the aggregation ring
- net/bnxt: parse checksum offload flags
- net/bonding: check error of MAC address setting
- net/bonding: fix activated slave in 8023ad mode
- net/bonding: fix bonding in 8023ad mode
- net/bonding: fix setting slave MAC addresses

- net/dpaa: fix FW version code
- net/dpaa: fix potential memory leak
- net/dpaa: fix the mbuf packet type if zero
- net/dpaa: fix uninitialized and unused variables
- net/e1000: fix null pointer check
- net/e1000: fix VF Rx interrupt enabling
- net/ena: do not set Tx L4 offloads in Rx path
- net/enic: fix crash due to static max number of queues
- net/enic: fix L4 Rx ptype comparison
- net/failsafe: fix invalid free
- net/failsafe: fix Rx safe check compiler hint
- net: fix ESP header byte ordering definition
- net/fm10k: fix logical port delete
- net/i40e: add debug logs when writing global registers
- net/i40e: add FDIR NVGRE parameter check
- net/i40e: check multi-driver option parsing
- net/i40e: exclude LLDP packet count
- net/i40e: fix ARM big endian build
- net/i40e: fix FDIR input set conflict
- net/i40e: fix FDIR rule confliction issue
- net/i40e: fix flag for MAC address write
- net/i40e: fix flow director Rx resource defect
- net/i40e: fix interrupt conflict with multi-driver
- net/i40e: fix ISO C in exported header
- net/i40e: fix memory leak
- net/i40e: fix multiple DDP packages conflict
- net/i40e: fix multiple driver support
- net/i40e: fix packet type for X722
- net/i40e: fix port segmentation fault when restart
- net/i40e: fix Rx interrupt
- net/i40e: fix setting MAC address of VF
- net/i40e: fix setting of MAC address on i40evf
- net/i40e: fix VF reset stats crash
- net/i40e: fix VF Rx interrupt enabling

- net/i40e: fix VLAN offload setting
- net/i40e: fix VLAN offload setting issue
- net/i40e: fix VSI MAC filter on primary address change
- net/i40e: warn when writing global registers
- net/igb: fix Tx queue number assignment
- net/ixgbe: fix ARM big endian build
- net/ixgbe: fix max queue number for VF
- net/ixgbe: fix parsing FDIR NVGRE issue
- net/ixgbe: fix reset error handling
- net/ixgbe: fix the failure of number of Tx queue check
- net/ixgbe: fix tunnel filter fail problem
- net/ixgbe: fix VF Rx interrupt enabling
- net/ixgbe: fix wrong PBA setting
- net/mlx4: fix drop flow resources leak
- net/mlx4: fix Rx offload non-fragmented indication
- net/mlx4: fix Tx packet drop application report
- net/mlx4: fix unnecessary include
- net/mlx4: revert workaround for broken Verbs
- net/mlx5: cleanup allocation of ethtool stats
- net/mlx5: fix calculation of flow ID flag
- net/mlx5: fix deadlock of link status alarm
- net/mlx5: fix flow item validation
- net/mlx5: fix flow priority on queue action
- net/mlx5: fix flow RSS configuration
- net/mlx5: fix handling link status event
- net/mlx5: fix HW checksum offload for outer IP
- net/mlx5: fix link state on device start
- net/mlx5: fix memory region boundary checks
- net/mlx5: fix memory region cache last index
- net/mlx5: fix memory region cache lookup
- net/mlx5: fix memory region lookup
- net/mlx5: fix Memory Region registration
- net/mlx5: fix missing attribute size for drop action
- net/mlx5: fix missing RSS capability

- net/mlx5: fix overflow of Memory Region cache
- net/mlx5: fix overwriting bit-fields in SW Rx queue
- net/mlx5: fix port stop by verify flows are still present
- net/mlx5: fix return value of start operation
- net/mlx5: fix RSS key configuration
- net/mlx5: fix secondary process verification
- net/mlx5: fix Tx checksum offloads
- net/mlx5: fix UAR remapping on non configured queues
- net/mlx5: fix un-supported RSS hash fields use
- net/mlx5: fix VLAN configuration after port stop
- net/mlx5: remove parser/flow drop queue
- net/mlx5: use PCI address as port name
- net/mrvl: fix HIF objects allocation
- net/mrvl: fix multiple probe
- net/mrvl: fix oversize bpool handling
- net/mrvl: fix shadow queue tail and size calculations
- net/mrvl: keep shadow Txqs inside PMD Txq
- net/nfp: fix CRC strip check behaviour
- net/nfp: fix jumbo settings
- net/nfp: fix MTU settings
- net/octeontx: add channel to port id mapping
- net/pcap: fix the NUMA id display in logs
- net/qede/base: fix VF LRO tunnel configuration
- net/qede: check tunnel L3 header
- net/qede: fix clearing of queue stats
- net/qede: fix few log messages
- net/qede: fix MTU set and max Rx length
- net/qede: fix to enable LRO over tunnels
- net/qede: fix to reject config with no Rx queue
- net/qede: fix tunnel header size in Tx BD configuration
- net/qede: replace config option with run-time arg
- net/sfc: do not hold management event queue lock while MCDI
- net/sfc: fix DMA memory leak after kvarg processing failure
- net/sfc: fix flow RSS check in error handling

- net/sfc: fix incorrect bitwise ORing of L3/L4 packet types
- net/sfc: fix initialization of flow structure
- net/sfc: fix label name to be consistent
- net/sfc: fix main MAC address handling
- net/sfc: fix multicast address list copy memory leak
- net/sfc: stop periodic DMA if MAC stats upload fails
- net/szedata2: fix check of mmap return value
- net/tap: fix cleanup on allocation failure
- net/tap: remove unused kernel version definitions
- net/thunderx: fix multi segment Tx function return
- net/virtio: fix incorrect cast
- net/virtio: fix memory leak when reinitializing device
- net/virtio: fix queue flushing with vector Rx enabled
- net/virtio: fix Rx and Tx handler selection for ARM32
- net/virtio: fix typo in LRO support
- net/virtio: fix vector Rx flushing
- net/virtio-user: fix crash as features change
- pdump: fix error check when creating/canceling thread
- pmdinfo: fix cross compilation for ARM big endian
- security: fix device operation type
- security: fix enum start value
- security: fix pedantic compilation
- service: fix lcore role after delete
- service: fix number mapped cores count
- service: fix possible mem leak on initialize
- service: fix service core launch
- test/bitmap: fix memory leak
- test/crypto: fix missing include
- test/eventdev: use CPU event type
- test/memzone: fix freeing test
- test/memzone: fix NULL freeing
- test/memzone: fix wrong test
- test: register test as failed if setup failed
- test/reorder: fix memory leak

- test/ring: fix memory leak
- test/ring\_perf: fix memory leak
- test/table: fix memory leak
- test/table: fix uninitialized parameter
- test/timer\_perf: fix memory leak
- timer: fix reset on service cores
- usertools/devbind: fix kernel module reporting
- vfio: fix enabled check on error
- vhost: fix crash
- vhost: fix dequeue zero copy with virtio1
- vhost: fix error code check when creating thread
- vhost: fix IOTLB pool out-of-memory handling
- vhost: fix mbuf free
- vhost: protect active rings from async ring changes
- vhost: remove pending IOTLB entry if miss request failed

### **2.8.2 17.11.2**

- examples/vhost: move to safe GPA translation API
- examples/vhost\_scsi: move to safe GPA translation API
- vhost: add support for non-contiguous indirect desc tables (fixes CVE-2018-1059)
- vhost: check all range is mapped when translating GPAs (fixes CVE-2018-1059)
- vhost: deprecate unsafe GPA translation API (fixes CVE-2018-1059)
- vhost: ensure all range is mapped when translating QVAs (fixes CVE-2018-1059)
- vhost: fix indirect descriptors table translation size (fixes CVE-2018-1059)
- vhost: handle virtually non-contiguous buffers in Rx (fixes CVE-2018-1059)
- vhost: handle virtually non-contiguous buffers in Rx-mrg (fixes CVE-2018-1059)
- vhost: handle virtually non-contiguous buffers in Tx (fixes CVE-2018-1059)
- vhost: introduce safe API for GPA translation (fixes CVE-2018-1059)

### **2.8.3 17.11.3**

- app/crypto-perf: check minimum lcore number
- app/crypto-perf: fix excess crypto device error
- app/crypto-perf: fix IOVA translation
- app/crypto-perf: fix parameters copy



- app/crypto-perf: use strcpy for allocated string
- app/procinfo: fix strncpy usage in args parsing
- app/testpmd: fix burst stats reporting
- app/testpmd: fix command token
- app/testpmd: fix empty list of RSS queues for flow
- app/testpmd: fix forward ports Rx flush
- app/testpmd: fix forward ports update
- app/testpmd: fix removed device link status asking
- app/testpmd: fix slave port detection
- app/testpmd: fix synchronic port hotplug
- app/testpmd: fix valid ports prints
- bus/dpaa: fix resource leak
- bus/fslmc: fix find device start condition
- bus/pci: fix find device implementation
- bus/vdev: fix finding device by name
- cryptodev: fix supported size check
- crypto/dpaa2\_sec: fix HMAC supported digest sizes
- crypto/scheduler: fix 64-bit mask of workers cores
- crypto/scheduler: fix memory leak
- crypto/scheduler: fix multicore rings re-use
- crypto/scheduler: fix possible duplicated ring names
- crypto/scheduler: set null pointer after freeing
- crypto/zuc: batch ops with same transform
- crypto/zuc: do not set default op status
- doc: add timestamp offload to mlx5 features
- doc: fix NFP NIC guide grammar
- drivers/net: fix link autoneg value for virtual PMDs
- eal/ppc: remove braces in SMP memory barrier macro
- ethdev: fix port accessing after release
- ethdev: fix queue start
- event/dpaa2: remove link from info structure
- examples/exception\_path: limit core count to 64
- examples/l2fwd-crypto: fix the default aead assignments
- examples/performance-thread: fix return type of threads

- examples/quota\_watermark: fix return type of threads
- hash: fix missing spinlock unlock in add key
- ip\_frag: fix double free of chained mbufs
- kni: fix build on CentOS 7.4
- kni: fix build on RHEL 7.5
- mbuf: fix Tx checksum offload API doc
- mbuf: improve tunnel Tx offloads API doc
- mem: do not use physical addresses in IOVA as VA mode
- mempool: fix leak when no objects are populated
- mempool: fix virtual address population
- mk: fix make defconfig on FreeBSD
- net: add IPv6 header fields macros
- net/bnx2x: do not cast function pointers as a policy
- net/bnx2x: fix for PCI FLR after ungraceful exit
- net/bnx2x: fix KR2 device check
- net/bnx2x: fix memzone name overrun
- net/bnxt: avoid invalid vnic id in set L2 Rx mask
- net/bnxt: fix endianness of flag
- net/bnxt: fix license header
- net/bnxt: fix LRO disable
- net/bnxt: fix Rx checksum flags
- net/bnxt: fix Rx checksum flags for tunnel frames
- net/bnxt: fix Rx drop setting
- net/bnxt: fix Rx mbuf and agg ring leak in dev stop
- net/bnxt: fix usage of vnic id
- net/bnxt: free memory allocated for VF filters
- net/bnxt: set padding flags in Rx descriptor
- net/bonding: clear started state if start fails
- net/bonding: export mode 4 slave info routine
- net/bonding: fix primary slave port id storage type
- net/bonding: fix setting VLAN ID on slave ports
- net/bonding: fix slave activation simultaneously
- net/bonding: free mempool used in mode 6
- net/dpaa2: fix xstats

- net/dpaa: fix oob access
- net/enic: allocate stats DMA buffer upfront during probe
- net/enic: fix crash on MTU update with non-setup queues
- net/failsafe: fix duplicate event registration
- net/failsafe: fix probe cleanup
- net/failsafe: fix removed sub-device cleanup
- net/i40e: fix DDP profile DEL operation
- net/i40e: fix failing to disable FDIR Tx queue
- net/i40e: fix intr callback unregister by adding retry
- net/i40e: fix link status update
- net/i40e: fix link update no wait
- net/i40e: fix shifts of signed values
- net/ixgbe: enable vector PMD for icc 32 bits
- net/ixgbe: fix busy wait during checking link status
- net/ixgbe: fix DCB configuration
- net/ixgbe: fix intr callback unregister by adding retry
- net/ixgbe: fix too many interrupts
- net/liquidio: fix link state fetching during start
- net/mlx4: avoid constant recreations in function
- net/mlx4: fix a typo in header file
- net/mlx4: fix broadcast Rx
- net/mlx4: fix removal detection of stopped port
- net/mlx4: fix RSS resource leak in case of error
- net/mlx4: fix Rx resource leak in case of error
- net/mlx4: fix single port configuration
- net/mlx4: fix UDP flow rule limitation enforcement
- net/mlx4: store RSS hash result in mbufs
- net/mlx5: add data-plane debug message macro
- net/mlx5: add missing function documentation
- net/mlx5: add packet type index for TCP ack
- net/mlx5: change device reference for secondary process
- net/mlx5: change non failing function return values
- net/mlx5: change pkt burst select function prototype
- net/mlx5: change tunnel flow priority

- net/mlx5: enforce RSS key length limitation
- net/mlx5: fix allocation when no memory on device NUMA node
- net/mlx5: fix build with clang on ARM
- net/mlx5: fix calculation of Tx TSO inline room size
- net/mlx5: fix close after start failure
- net/mlx5: fix count in xstats
- net/mlx5: fix CRC strip capability query
- net/mlx5: fix disabling Tx packet inlining
- net/mlx5: fix double free on error handling
- net/mlx5: fix ethtool link setting call order
- net/mlx5: fix existing file removal
- net/mlx5: fix flow creation with a single target queue
- net/mlx5: fix flow director conversion
- net/mlx5: fix flow director drop rule deletion crash
- net/mlx5: fix flow director mask
- net/mlx5: fix flow director rule deletion crash
- net/mlx5: fix flow validation
- net/mlx5: fix icc build
- net/mlx5: fix invalid flow item check
- net/mlx5: fix IPv6 header fields
- net/mlx5: fix link status behavior
- net/mlx5: fix link status initialization
- net/mlx5: fix link status to use wait to complete
- net/mlx5: fix probe return value polarity
- net/mlx5: fix reception of multiple MAC addresses
- net/mlx5: fix resource leak in case of error
- net/mlx5: fix RSS flow action bounds check
- net/mlx5: fix RSS key length query
- net/mlx5: fix secondary process mempool registration
- net/mlx5: fix socket connection return value
- net/mlx5: fix sriov flag
- net/mlx5: fix synchronization on polling Rx completions
- net/mlx5: improve flow error explanation
- net/mlx5: map UAR address around huge pages

- net/mlx5: mark parameters with unused attribute
- net/mlx5: name parameters in function prototypes
- net/mlx5: normalize function prototypes
- net/mlx5: prefix all functions with mlx5
- net/mlx5: refuse empty VLAN flow specification
- net/mlx5: remove 32-bit support
- net/mlx5: remove assert un-accessible from secondary process
- net/mlx5: remove control path locks
- net/mlx5: remove excessive data prefetch
- net/mlx5: remove get priv internal function
- net/mlx5: remove kernel version check
- net/mlx5: remove useless empty lines
- net/mlx5: setup RSS regardless of queue count
- net/mlx5: split L3/L4 in flow director
- net/mlx5: standardize on negative errno values
- net/mlx5: use dynamic logging
- net/mlx5: use port id in PMD log
- net/mlx5: warn for unsuccessful memory registration
- net/mlx: control netdevices through ioctl only
- net/mrvl: fix crash when port is closed without starting
- net/mrvl: fix Rx descriptors number
- net/nfp: fix assigning port id in mbuf
- net/nfp: fix barrier location
- net/nfp: fix link speed capabilities
- net/nfp: fix mbufs releasing when stop or close
- net/octeontx: fix null pointer dereference
- net/octeontx: fix uninitialized speed variable
- net/octeontx: fix uninitialized variable in port open
- net/qede/base: fix to support OVLAN mode
- net/qede: fix alloc from socket 0
- net/qede: fix device stop to remove primary MAC
- net/qede: fix L2-handles used for RSS hash update
- net/qede: fix memory alloc for multiple port reconfig
- net/qede: fix missing loop index in Tx SG mode

- net/qede: fix multicast filtering
- net/qede: fix to prevent overwriting packet type
- net/qede: fix unicast filter routine return code
- net/qede: fix VF port creation sequence
- net/sfc: add missing defines for SAL annotation
- net/sfc: add missing Rx fini on RSS setup fail path
- net/sfc/base: fix comparison always true warning
- net/sfc: fix mbuf data alignment calculation
- net/sfc: fix type of opaque pointer in perf profile handler
- net/sfc: ignore spec bits not covered by mask
- net/sfc: process RSS settings on Rx configure step
- net/szedata2: fix format string for PCI address
- net/szedata2: fix total stats
- net/tap: fix icc build
- net/vhost: fix crash when creating vdev dynamically
- net/vhost: fix invalid state
- net/vhost: initialise device as inactive
- net/vmxnet3: set the queue shared buffer at start
- nfp: allow for non-root user
- nfp: restore the unlink operation
- nfp: unlink the appropriate lock file
- pci: remove duplicated symbol from map file
- test/distributor: fix return type of thread function
- test: fix memory flags test for low NUMA nodes number
- test/mempool: fix autotest retry
- test/pipeline: fix return type of stub miss
- test/pipeline: fix type of table entry parameter
- test/reorder: fix freeing mbuf twice
- vfio: do not needlessly check for IOVA mode
- vhost: check cmsg not null
- vhost: fix compilation issue when vhost debug enabled
- vhost: fix dead lock on closing in server mode
- vhost: fix device cleanup at stop
- vhost: fix message payload union in setting ring address

- vhost: fix offset while mmaping log base address
- vhost: fix realloc failure
- vhost: fix ring index returned to master on stop

#### 2.8.4 17.11.4

- app/crypto-perf: fix auth IV offset
- app/testpmd: fix buffer leak in TM command
- app/testpmd: fix DCB config
- app/testpmd: fix VLAN TCI mask set error for FDIR
- bitrate: add sanity check on parameters
- bus/dpaa: fix buffer offset setting in FMAN
- bus/dpaa: fix build
- bus/dpaa: fix phandle support for Linux 4.16
- bus/pci: use IOVAs check when setting IOVA mode
- crypto/qat: fix checks for 3GPP algo bit params
- doc: fix bonding command in testpmd
- doc: update qede management firmware guide
- eal: fix bitmap documentation
- eal: fix return codes on thread naming failure
- eal/linux: fix invalid syntax in interrupts
- eal/linux: fix uninitialized value
- ethdev: fix a doxygen comment for port allocation
- ethdev: fix queue statistics mapping documentation
- eventdev: add event buffer flush in Rx adapter
- eventdev: fix internal port logic in Rx adapter
- eventdev: fix missing update to Rx adapter WRR position
- eventdev: fix port in Rx adapter internal function
- eventdev: fix Rx SW adapter stop
- event: fix ring init failure handling
- event/octeonx: remove unnecessary port start and stop
- examples/exception\_path: fix out-of-bounds read
- examples: fix strncpy error for GCC8
- examples/flow\_filtering: add flow director config for i40e
- examples/ipsec-secgw: fix bypass rule processing

- examples/ipsec-secgw: fix IPv4 checksum at Tx
- examples/l2fwd-crypto: check return value on IV size check
- examples/l2fwd-crypto: fix digest with AEAD algo
- examples/l2fwd-crypto: skip device not supporting operation
- examples/l3fwd: remove useless include
- hash: fix a multi-writer race condition
- hash: fix doxygen of return values
- hash: fix key slot size accuracy
- hash: fix multiwriter lock memory allocation
- kni: fix build on RHEL 7.5
- kni: fix build with gcc 8.1
- kni: fix crash with null name
- maintainers: claim maintainership for ARM v7 and v8
- maintainers: update for Mellanox PMDs
- mem: add function for checking memsegs IOVAs addresses
- mem: fix max DMA maskbit size
- mem: use address hint for mapping hugepages
- metrics: add check for invalid key
- metrics: disallow null as metric name
- metrics: do not fail silently when uninitialised
- mk: fix cross build
- mk: fix permissions when using make install
- mk: remove unnecessary test rules
- mk: update targets for classified tests
- net/bnx2x: fix FW command timeout during stop
- net/bnx2x: fix poll link status
- net/bnx2x: fix to set device link status
- net/bnxt: add missing ids in xstats
- net/bnxt: check access denied for HWRM commands
- net/bnxt: check for invalid vNIC id
- net/bnxt: fix filter freeing
- net/bnxt: fix HW Tx checksum offload check
- net/bnxt: fix lock release on NVM write failure
- net/bnxt: fix memory leaks in NVM commands



- net/bnxt: fix RETA size
- net/bnxt: fix Rx ring count limitation
- net/bnxt: fix set MTU
- net/bnxt: fix to move a flow to a different queue
- net/bnxt: use correct flags during VLAN configuration
- net/bonding: always update bonding link status
- net/bonding: do not clear active slave count
- net/bonding: fix MAC address reset
- net/bonding: fix race condition
- net/cxgbe: fix init failure due to new flash parts
- net/cxgbe: fix Rx channel map and queue type
- net/dpaa2: remove loop for unused pool entries
- net/ena: change memory type
- net/ena: check pointer before memset
- net/ena: fix GENMASK\_ULL macro
- net/ena: fix SIGFPE with 0 Rx queue
- net/ena: set link speed as none
- net/enic: add devarg to specify ingress VLAN rewrite mode
- net/enic: do not overwrite admin Tx queue limit
- net/i40e: fix check of flow director programming status
- net/i40e: fix link speed
- net/i40e: fix packet type parsing with DDP
- net/i40e: fix setting TPID with AQ command
- net/i40e: fix shifts of 32-bit value
- net/i40e: revert fix of flow director check
- net/i40e: workaround performance degradation
- net/ixgbe: add support for VLAN in IP mode FDIR
- net/ixgbe: fix mask bits register set error for FDIR
- net/ixgbe: fix tunnel id format error for FDIR
- net/ixgbe: fix tunnel type set error for FDIR
- net/mlx4: check RSS queues number limitation
- net/mlx4: fix minor resource leak during init
- net/mlx5: add missing sanity checks for Tx completion queue
- net/mlx5: fix assert for Tx completion queue count

- net/mlx5: fix build with old kernels
- net/mlx5: fix compilation for rdma-core v19
- net/mlx5: fix crash in device probe
- net/mlx5: fix error number handling
- net/mlx5: fix flow search on FDIR deletion
- net/mlx5: fix queue rollback when starting device
- net/mlx5: fix return value when deleting fdir filter
- net/mlx5: fix Rx buffer replenishment threshold
- net/mlx5: fix secondary process resource leakage
- net/mlx5: fix TCI mask filter
- net/mlx5: preserve allmulticast flag for flow isolation mode
- net/mlx5: preserve promiscuous flag for flow isolation mode
- net/mvpp2: check pointer before using it
- net/nfp: check hugepages IOVAs based on DMA mask
- net/nfp: fix field initialization in Tx descriptor
- net/nfp: support IOVA VA mode
- net/octeontx: fix stop clearing Rx/Tx functions
- net/pcap: fix multiple queues
- net/qede/base: fix GRC attention callback
- net/qede/base: fix to clear HW indication
- net/qede: fix default extended VLAN offload config
- net/qede: fix for devargs
- net/qede: fix incorrect link status update
- net/qede: fix interrupt handler unregister
- net/qede: fix legacy interrupt mode
- net/qede: fix link change event notification
- net/qede: fix MAC address removal failure message
- net/qede: fix ntuple filter configuration
- net/qede: fix unicast MAC address handling in VF
- net/qede: fix VF MTU update
- net/qede: remove primary MAC removal
- net/sfc: cut non VLAN ID bits from TCI
- net/sfc: fix assert in set multicast address list
- net/sfc: handle unknown L3 packet class in EF10 event parser

- net/tap: fix zeroed flow mask configurations
- net/thunderx: avoid sq door bell write on zero packet
- net/thunderx: fix build with gcc optimization on
- ring: fix sign conversion warning
- security: fix crash on destroy null session
- test/crypto: fix device id when stopping port
- test: fix code on report
- test: fix EAL flags autotest on FreeBSD
- test: fix result printing
- test: fix uninitialized port configuration
- test/flow\_classify: fix return types
- test/hash: fix multiwriter with non consecutive cores
- test/hash: fix potential memory leak
- test: improve filtering
- test: make autotest runner python 2/3 compliant
- test: print autotest categories
- vfiio: fix PCI address comparison
- vhost: fix missing increment of log cache count
- vhost: flush IOTLB cache on new mem table handling
- vhost: improve dirty pages logging performance
- vhost: release locks on RARP packet failure
- vhost: retranslate vring addr when memory table changes

### **2.8.5 17.11.5**

- acl: forbid rule with priority zero
- app/pdump: fix port id storage size
- app/procinfo: fix sprintf overrun
- app/test-crypto-perf: fix check for auth key
- app/test-crypto-perf: fix check for cipher IV
- app/test-crypto-perf: fix double allocation of memory
- app/testpmd: check Rx VLAN offload flag to print VLAN TCI
- app/testpmd: fix csum parse-tunnel command invocation
- app/testpmd: fix duplicate exit
- app/testpmd: fix L4 length for UDP checksum

- app/testpmd: fix memory leak for TM object
- app/testpmd: fix metering and policing commands
- app/testpmd: fix physical port socket initialization
- app/testpmd: fix printf format in event callback
- app/testpmd: fix RED byte stats
- app/testpmd: fix shaper profile parameters
- app/testpmd: fix vdev socket initialization
- app/testpmd: optimize mbuf pool allocation
- app/testpmd: reserve NUMA node per port and per ring
- build: enable ARM NEON flag when `__aarch64__` defined
- bus/dpaa: fix build with gcc 9.0
- bus/dpaa: fix inconsistent struct alignment
- bus/pci: compare kernel driver instead of interrupt handler
- bus/pci: fix allocation of device path
- bus/pci: fix config r/w access
- bus/pci: replace `strncpy` by `strncpy`
- crypto/aesni\_mb: fix possible array overrun
- crypto/mvsam: update hash digest sizes
- crypto/scheduler: fix build with gcc 8.2
- devtools: provide more generic grep in git check
- doc: add cross-compilation in sample apps guide
- doc: add VFIO in ENA guide
- doc: clarify L3 Tx checksum prerequisite
- doc: clarify L4 Tx checksum prerequisite
- doc: clarify TSO Tx offload prerequisite
- doc: fix formatting in IP reassembly app guide
- doc: fix missing CCM to QAT feature list
- doc: fix NUMA library name in Linux guide
- doc: fix spelling in PMD guides
- doc: fix style and syntax in flow API guide
- doc: fix typo in testpmd guide
- doc: fix typos in the flow API guide
- doc: fix wrong usage of bind command
- eal/arm64: fix intrinsic for GCC < 4.9

- eal: declare trace buffer at top of own block
- eal: explicit cast in constant byte swap
- eal: explicit cast in rwlock functions
- eal: explicit cast of builtin for bsf32
- eal: explicit cast of core id when getting index
- eal: explicit cast of strcpy return
- eal: fix build
- eal: fix build with gcc 9.0
- eal: fix build with -O1
- eal: fix casts in random functions
- eal: introduce rte version of fls
- eal/linux: fix memory leak of logid
- eal/linux: handle UIO read failure in interrupt handler
- eal: support strcpy function
- eal: use correct data type for bitmap slab operations
- eal/x86: fix type of variable in memcpy function
- eal/x86: remove unused memcpy file
- efd: fix write unlock during ring creation
- ethdev: explicit cast of buffered Tx number
- ethdev: explicit cast of queue count return
- ethdev: fix doxygen comment to be with structure
- ethdev: fix queue start and stop
- ethdev: fix type and scope of variables in Rx burst
- eventdev: fix eth Rx adapter hotplug incompatibility
- eventdev: fix unlock in Rx adapter
- event/sw: fix cq index check for unlink usecases
- examples/flow\_filtering: remove VLAN item
- examples/ipv4\_multicast: fix leak of cloned packets
- examples/vhost: remove unnecessary constant
- fix dpdk.org URLs
- gro: fix overflow of TCP payload calculation
- hash: explicit casts for truncation in CRC32c
- hash: move stack declaration at top of CRC32c function
- igb\_uio: fix refcount if open returns error

- ip\_frag: fix overflow in key comparison
- ip\_frag: use key length for key comparison
- kni: fix build on Linux < 3.14
- kni: fix build on Linux 4.19
- kni: fix kernel FIFO synchronization
- kni: fix possible uninitialized variable
- kvars: fix processing a null list
- latency: fix timestamp marking and latency calculation
- mbuf: avoid implicit demotion in 64-bit arithmetic
- mbuf: avoid integer promotion in prepend/adj/chain
- mbuf: explicit cast of headroom on reset
- mbuf: explicit cast of size on detach
- mbuf: explicit casts of reference counter
- mbuf: fix reference counter integer promotion
- mbuf: fix Tx offload mask
- mbuf: fix type of private size in detach
- mbuf: fix type of variables in linearize function
- mem: fix memory initialization time
- mem: fix undefined behavior in NUMA-aware mapping
- mk: disable gcc AVX512F support
- net/bnx2x: fix call to link handling periodic function
- net/bnx2x: fix logging to include device name
- net/bnx2x: fix to add PHY lock
- net/bnx2x: fix to disable further interrupts
- net/bnx2x: fix VF link state update
- net/bnxt: fix registration of VF async event completion ring
- net/bnxt: fix uninitialized pointer access in Tx
- net/bnxt: set MAC filtering as outer for non tunnel frames
- net/bnxt: set VLAN strip mode before default VNIC cfg
- net/bonding: do not ignore RSS key on device config
- net/bonding: fix crash when stopping mode 4 port
- net/bonding: fix possible silent failure in config
- net/bonding: fix Rx slave fairness
- net/bonding: stop and deactivate slaves on stop

- net/bonding: support matching QinQ ethertype
- net/bonding: use evenly distributed default RSS RETA
- net/e1000/base: fix uninitialized variable
- net/e1000: do not error out if Rx drop enable is set
- net/ena: fix passing RSS hash to mbuf
- net/enic: do not use non-standard integer types
- net/enic: fix flow API memory leak
- net/enic: set Rx VLAN offload flag for non-stripped packets
- net: explicit cast in L4 checksum
- net: explicit cast of IP checksum to 16-bit
- net: explicit cast of multicast bit clearing
- net: explicit cast of protocol in IPv6 checksum
- net/failsafe: add checks for deferred queue setup
- net/failsafe: fix crash on slave queue release
- net/failsafe: remove not supported multicast MAC filter
- net: fix build with pedantic
- net: fix Intel prepare function for IP checksum offload
- net/i40e/base: correct global reset timeout calculation
- net/i40e/base: fix comment referencing internal data
- net/i40e/base: gracefully clean the resources
- net/i40e/base: properly clean resources
- net/i40e: enable loopback function for X722 MAC
- net/i40e: fix link status update
- net/i40e: fix offload not supported mask
- net/i40e: fix send admin queue command before init
- net/i40e: fix X710 Rx after reading some registers
- net/i40e: keep promiscuous on if allmulticast is enabled
- net/i40e: update Tx offload mask
- net/igb: update Tx offload mask
- net/ixgbe: fix busy polling while fiber link update
- net/ixgbe: fix maximum wait time in comment
- net/ixgbe: stop link setup alarm handler before device start
- net/ixgbe: update Tx offload mask
- net/ixgbev: fix link state

- net/ixgbe: wait longer for link after fiber MAC setup
- net/mlx4: fix possible uninitialized variable
- net/mlx5: add Bluefield device id
- net/mlx5: disable ConnectX-4 Lx Multi Packet Send by default
- net/mlx5: fix initialization of struct members
- net/mlx5: fix interrupt completion queue index wrapping
- net/mlx5: fix multi-chunk mempool support
- net/mlx5: make vectorized Tx threshold configurable
- net/mlx5: optimize Rx buffer replenishment threshold
- net: move stack variable at top of VLAN strip function
- net/mvpp2: fix array initialization
- net/nfp: fix mbuf flags with checksum good
- net/nfp: fix memcpy out of source range
- net/nfp: fix misuse of strncpy
- net/nfp: fix RSS
- net/nfp: replace strncpy by strncpy
- net/octeontx: fix failures when available ports > queues
- net/octeontx: fix mbuf corruption with large private sizes
- net/octeontx: fix packet corruption on Tx
- net/qede/base: fix MFW FLR flow
- net/qede/base: fix to handle stag update event
- net/qede: fix crash when configure fails
- net/qede: fix ethernet type in HW registers
- net/qede: fix flow director for IPv6 filter
- net/qede: fix Rx buffer size calculation
- net/qede: fix strncpy
- net/qede: fix Tx offload mask
- net/qede: fix Tx tunnel offload support mask
- net/qede: replace strncpy by strncpy
- net/sfc: allow to query RSS key and HF in isolated mode
- net/sfc: allow to query RSS key and HF when RSS is disabled
- net/sfc/base: avoid usage of too big arrays on stack
- net/sfc/base: check size of memory to read sensors data to
- net/sfc/base: fix a typo in unicast filter insertion comment



- net/sfc/base: fix build because of no declaration
- net/sfc/base: fix ID retrieval in v3 licensing
- net/sfc/base: fix MAC Tx stats for less or equal to 64 bytes
- net/sfc/base: fix out of bounds read when dereferencing sdup
- net/sfc/base: fix PreFAST warnings because of unused return
- net/sfc/base: fix SAL annotation for input buffers
- net/sfc/base: make last byte of module information available
- net/sfc/base: prevent access to the NIC config before probe
- net/sfc/base: properly align on line continuation
- net/sfc/base: remove Falcon-specific concurrency check
- net/sfc: do not skip RSS configuration step on reconfigure
- net/sfc: fix an Rx queue double release possibility
- net/sfc: fix a Tx queue double release possibility
- net/sfc: make sure that stats name is nul-terminated
- net/sfc: receive prepared packets even in Rx exception case
- net/softnic: fix undefined dev info fields
- net/tap: fix file descriptor check
- net/thunderx: fix Tx desc corruption in scatter-gather mode
- net/vhost: fix parameters string
- net/virtio: fix PCI config error handling
- net/virtio: fix unchecked return value
- net/virtio: register/unregister intr handler on start/stop
- net/virtio-user: check negotiated features before set
- net/virtio-user: do not reset owner when driver resets
- net/virtio-user: fix typo in error message
- pci: fix parsing of address without function number
- ring: remove signed type flip-flopping
- ring: remove useless variables
- spinlock/x86: move stack declaration before code
- table: fix casting cuckoo hash function
- test/crypto: fix number of queue pairs
- test/event: check burst mode capability
- test/hash: fix bucket size in perf test
- test/hash: fix build

- test: release ring resources after PMD perf test
- test/reorder: fix out of bound access
- usertools: check for lspci dependency
- vfio: do not needlessly setup device in secondary process
- vhost: fix corner case for enqueue operation
- vhost: remove unneeded null pointer check

### 2.8.6 17.11.6

- app/testpmd: expand RED queue thresholds to 64 bits
- app/testpmd: fix quit to stop all ports before close
- crypto/qat: fix block size error handling
- devtools: fix wrong headline lowercase for arm
- doc: add dependency for PDF in contributing guide
- doc: add missing loopback option in testpmd guide
- doc: fix a parameter name in testpmd guide
- doc: fix garbage text in generated HTML guides
- doc: fix references in power management guide
- drivers/crypto: fix PMDs memory leak
- drivers: fix sprintf with snprintf
- drivers/net: fix several Tx prepare functions
- eal: check string parameter lengths
- eal: fix core number validation
- eal: fix out of bound access when no CPU available
- efd: fix tail queue leak
- ethdev: clear ethdev data upon detach
- ethdev: fix errno to have positive value
- ethdev: fix typo in queue setup error log
- eventdev: fix xstats documentation typo
- examples/bond: fix crash when there is no active slave
- examples/flow\_filtering: fix example documentation
- examples/ipsec-secgw: fix outbound codepath for single SA
- examples/ipsec-secgw: make local variables static
- examples/tep\_term: remove unused constant
- examples/vhost: fix path allocation failure handling

- gro: check invalid TCP header length
- gso: fix VxLAN/GRE tunnel checks
- kni: fix build for dev\_open in Linux 5.0
- kni: fix build for igb\_ndo\_bridge\_setlink in Linux 5.0
- kni: fix build on RHEL 8
- kni: fix build on RHEL8 for arm and Power9
- mk: fix scope of disabling AVX512F support
- net/af\_packet: fix setting MTU decrements sockaddr twice
- net/bnx2x: cleanup info logs
- net/bonding: fix possible null pointer reference
- net/dpaa2: fix device init for secondary process
- net/dpaa: fix secondary process
- net/ena: add supported RSS offloads types
- net/ena: fix dev init with multi-process
- net/ena: fix errno to positive value
- net/ena: update completion queue after cleanup
- net/enic: remove useless include
- net: fix underflow for checksum of invalid IPv4 packets
- net/fm10k: fix internal switch initial status
- net/i40e: fix get RSS conf
- net/i40e: fix getting RSS configuration
- net/i40e: fix queue region DCB configure
- net/i40e: fix statistics inconsistency
- net/i40e: fix using recovery mode firmware
- net/i40e: fix VF overwrite PF RSS LUT for X722
- net/i40e: remove redundant reset of queue number
- net/i40e: revert fix offload not supported mask
- net/igb: fix LSC interrupt when using MSI-X
- net/ixgbe: fix over using multicast table for VF
- net/ixgbe: fix overwriting RSS RETA
- net/mlx5: fix validation of Rx queue number
- net/qede: fix performance bottleneck in Rx path
- net/qede: remove prefetch in Tx path
- net/sfc/base: fix Tx descriptor max number check

- net/sfc: discard last seen VLAN TCI if Tx packet is dropped
- net/sfc: fix datapath name references in logs
- net/sfc: fix port ID log
- net/sfc: fix Rx packets counter
- net/sfc: fix typo in preprocessor check
- net/sfc: fix VF error/missed stats mapping
- net/sfc: pass HW Tx queue index on creation
- net/tap: add buffer overflow checks before checksum
- net/tap: fix possible uninitialized variable access
- net/virtio: add barrier before reading the flags
- net/virtio: fix resuming port with Rx vector path
- net/virtio-user: fix used ring in cvq handling
- test/crypto: fix misleading trace message
- test/memzone: fix typo
- test/memzone: handle previously allocated memzones
- timer: fix race condition
- vfiio: fix error message
- vhost: fix access for indirect descriptors
- vhost: fix crash after mmap failure
- vhost: fix error handling when mem table gets updated
- vhost: fix memory leak on realloc failure
- vhost: fix payload size of reply
- vhost: fix possible out of bound access in vector filling
- vhost: fix race condition when adding fd in the fdset

## 2.9 17.11.7 Release Notes

### 2.9.1 17.11.7 Fixes

- app/crypto-perf: check range of socket id
- app/test: fix build with musl libc
- app/test: fix sprintf with strlcat
- app/testpmd: fix a typo in log message
- app/testpmd: fix stdout flush after printing stats
- app/testpmd: fix typo in comment

- app/testpmd: remove unused field from port struct
- app/testpmd: remove useless casts on statistics
- app/testpmd: revert fixed flag for exact link speed
- app/testpmd: set fixed flag for exact link speed
- bitrate: fix unchecked return value
- build: fix crash by disabling AVX512 with binutils 2.31
- bus/dpaa: fix Rx discard register mask
- bus/fslmc: fix build with musl libc
- bus/fslmc: fix warning with GCC 9
- bus/fslmc: remove unused include of error.h
- bus/vdev: fix debug message on probing
- cfgfile: replace strcat with strlcat
- cryptodev: fix driver name comparison
- crypto/dpaa2\_sec: fix session clearing
- crypto/qat: fix not included algs for zero counter
- doc: fix a minor typo in testpmd guide
- doc: fix broken link in LPM guide
- doc: fix examples in bonding guide
- doc: fix interactive commands in testpmd guide
- doc: fix link in Linux getting started guide
- doc: fix links to doxygen and sphinx sites
- doc: fix typo in mlx5 guide
- doc: remove reference to rte.doc.mk in programmers guide
- drivers/net: do not use private ethdev data
- drivers/net: fix possible overflow using strlcat
- drivers/net: fix shifting 32-bit signed variable 31 times
- eal: fix check when retrieving current CPU affinity
- eal: fix typo in comment of vector function
- eal: improve musl compatibility of string functions
- eal/linux: fix log levels for pagemap reading failure
- eal/ppc: fix global memory barrier
- eal: remove dead code in core list parsing
- eal: support strlcat function
- eal: tighten permissions on shared memory files

- ethdev: fix a typo
- event/sw: fix enqueue checks in self-test
- examples/ethtool: fix two typos
- examples/ipsec-secgw: fix AES-CTR block size
- examples/ipsec-secgw: fix build error log
- examples/ipsec-secgw: fix debug logs
- examples/l2fwd-cat: fix build on FreeBSD
- examples/vhost\_scsi: fix null-check for parameter
- examples/vm\_power\_manager: fix PMD specific code
- hash: fix doc about thread/process safety
- igb\_uio: fix build on Linux 5.3 for fall through
- kni: fix build on RHEL8
- kni: fix build with Linux 5.1
- maintainers: update for IBM POWER
- malloc: fix documentation of realloc function
- mbuf: fix a typo
- mk: disable warnings for packed mem config data structure
- mk: fix AVX512 disabled warning on non x86
- mk: fix build of shared library with libbsd
- net/bnx2x: fix DMAE timeout
- net/bnx2x: fix memory leak
- net/bnx2x: fix optic module verification
- net/bnx2x: fix race for periodic flags
- net/bnx2x: fix ramrod timeout
- net/bnx2x: fix segfaults due to stale interrupt status
- net/bnxt: fix Rx VLAN offload flags
- net/bnxt: support IOVA VA mode
- net/bonding: avoid warning for invalid port
- net/bonding: fix buffer length when printing strings
- net/bonding: fix LACP negotiation
- net/bonding: fix port id types
- net/bonding: fix queue index types
- net/bonding: fix reset active slave
- net/cxgbe: fix missing checksum flags and packet type

- net/cxgbe: update Chelsio T5/T6 NIC device ids
- net/enic: check for unsupported flow item types
- net/enic: fix flow director SCTP matching
- net/enic: fix SCTP match for flow API
- net: fix Tx VLAN flag for offload emulation
- net/fm10k: fix VLAN strip offload flag
- net/i40e: fix dereference before null check in mbuf release
- net/i40e: fix link speed for X722
- net/i40e: fix time sync for 25G
- net/ixgbe: fix crash on remove
- net/ixgbe: fix warning with GCC 9
- net/kni: fix return value check
- net/mlx5: check Tx queue size overflow
- net/mlx5: fix comments mixing Rx and Tx
- net/mlx5: fix hex dump of error completion
- net/mlx5: fix instruction hotspot on replenishing Rx buffer
- net/mlx5: fix max number of queues for NEON Tx
- net/mlx5: fix packet inline on Tx queue wraparound
- net/mlx5: fix release of Rx queue object
- net/mlx5: fix sync when handling Tx completions
- net/nfp: fix RSS query
- net/null: add MAC address setting fake operation
- net/octeontx: fix vdev name
- net/qede: fix Tx packet prepare for tunnel packets
- net/qede: support IOVA VA mode
- net/ring: avoid hard-coded length
- net/ring: check length of ring name
- net/ring: fix return value check
- net/ring: use calloc style where appropriate
- net/sfc: log port ID as 16-bit unsigned integer on panic
- net/sfc: remove control path logging from Rx queue count
- net/virtio: add barrier in interrupt enable
- net/virtio: fix buffer leak on VLAN insert
- net/virtio: fix dangling pointer on failure

- net/virtio: fix duplicate naming of include guard
- net/virtio: remove forward declaration
- net/virtio: remove useless condition
- power: fix frequency list buffer validation
- ring: fix an error message
- ring: fix namesize macro documentation block
- spinlock: reimplement with atomic one-way barrier
- test/bonding: assign non-zero MAC to null devices
- test/distributor: replace sprintf with strlcpy
- test/hash: replace sprintf with snprintf
- test/rwlock: add missing inttypes include
- test/rwlock: benchmark on all available cores
- test/spinlock: amortize the cost of getting time
- test/spinlock: remove delay for correct benchmarking
- test/virtual\_pmd: add MAC address setting fake op
- version: 17.11.7-rc1
- vhost: fix device leak on connection add failure
- vhost: fix potential use-after-free for memory region
- vhost: fix potential use-after-free for zero copy mbuf
- vhost: fix silent queue enabling with legacy guests
- vhost: restore mbuf first when freeing zmbuf

### 2.9.2 17.11.7 Validation

- Mellanox(R) Testing
  - Basic functionality
    - \* Send and receive multiple types of traffic
    - \* testpmd xstats counter test
    - \* testpmd timestamp test
    - \* Changing/checking link status through testpmd
    - \* RTE flow and flow\_director tests
    - \* Some RSS tests
    - \* VLAN stripping and insertion tests
    - \* Checksum and TSO tests
    - \* ptype tests



- \* Port interrupt testing
- \* Multi-process testing
- OFED versions tested
  - \* rdma-core v25.0
  - \* MLNX\_OFED\_LINUX-4.6-1.0.1.1
- Tested NICs
  - \* ConnectX-4 Lx (fw 14.25.6406).
  - \* ConnectX-5 (fw 16.25.6406).
- OS tested
  - \* RHEL7.4 (kernel 5.3.0-rc4).
  - \* RHEL7.4 (kernel 3.10.0-693.el7.x86\_64).
- Red Hat(R) Testing
  - Functionality
    - \* PF
    - \* VF
    - \* vhost single/multi queues and cross-NUMA
    - \* vhostclient reconnect
    - \* vhost live migration with single/multi queues and cross-NUMA
    - \* OVS PVP
- Intel(R) Testing
  - Basic Intel(R) NIC(ixgbe and i40e) testing
    - \* vlan
    - \* vxlan
    - \* Jumbo frames
    - \* Generic filter
    - \* Flow director
    - \* PF and VF
    - \* Intel NIC single core/NIC performance
  - Basic cryptodev and virtio testing
    - \* cryptodev
    - \* vhost/virtio basic loopback, PVP and performance test

### 2.9.3 17.11.7 Known Issues

- DPDK 17.11.7 contains fixes up to DPDK v19.08. Issues identified/fixed in DPDK master branch after DPDK v19.08 may be present in DPDK 17.11.7

### 2.9.4 Fixes skipped and status unresolved

## 2.10 17.11.8 Release Notes

### 2.10.1 17.11.8 Fixes

- vhost: add number of fds to vhost-user messages
- vhost: fix possible denial of service by leaking FDs - CVE-2019-14818
- vhost: fix possible denial of service on SET\_VRING\_NUM - CVE-2019-14818
- vhost: validate virtqueue size

### 2.10.2 17.11.8 Validation

- Tested with two testpmd instances, one with Vhost PMD, the other with Virtio-user PMD. Initialization goes well, and packets flow.

## 2.11 17.11.9 Release Notes

### 2.11.1 17.11.9 Fixes

- vhost: fix vring requests validation broken if no FD

### 2.11.2 17.11.9 Validation

- virtio/vhost regression tests by Intel: \* [http://doc.dpdk.org/dts/test\\_plans/virtio\\_pvp\\_regression\\_test\\_plan.html](http://doc.dpdk.org/dts/test_plans/virtio_pvp_regression_test_plan.html) \* [http://doc.dpdk.org/dts/test\\_plans/vhost\\_dequeue\\_zero\\_copy\\_test\\_plan.html](http://doc.dpdk.org/dts/test_plans/vhost_dequeue_zero_copy_test_plan.html) \* [http://doc.dpdk.org/dts/test\\_plans/vm2vm\\_virtio\\_pmd\\_test\\_plan.html](http://doc.dpdk.org/dts/test_plans/vm2vm_virtio_pmd_test_plan.html)

## 2.12 17.11.10 Release Notes

### 2.12.1 17.11.10 Fixes

- app/crypto-perf: fix input of AEAD decrypt
- app/proc\_info: fix string copying to use strlcpy
- app/testpmd: block xstats for hidden ports
- app/testpmd: fix crash on port reset

- app/testpmd: fix help for loop topology option
- app/testpmd: fix Tx checksum when TSO enabled
- app/testpmd: use better randomness for Tx split
- bus/pci: align next mapping address on page boundary
- bus/pci: fix Intel IOMMU sysfs access check
- bus/pci: remove useless link dependency on ethdev
- cryptodev: fix checks related to device id
- cryptodev: fix initialization on multi-process
- cryptodev: fix missing device id range checking
- crypto/dpaa2\_sec: fix length retrieved from hardware
- crypto/openssl: use local copy for session contexts
- crypto/qat: fix digest length in XCBC capability
- crypto/qat: fix null auth issues when using vfio\_pci
- doc: fix a common typo in NIC guides
- doc: fix AESNI-GCM limitations in crypto guide
- doc: fix description of versioning macros
- doc: fix format in virtio guide
- doc: fix tap guide
- doc: fix typo in l2fwd-crypto guide
- doc/guides: clean repeated words
- doc: robustify PDF build
- drivers/crypto: remove some invalid comments
- eal: remove dead code on NUMA node detection
- ethdev: fix endian annotation for SPI item
- ethdev: fix include of ethernet header file
- ethdev: fix typos for ENOTSUP
- ethdev: limit maximum number of queues
- ethdev: remove redundant device info cleanup before get
- event/dpaa2: fix default queue configuration
- event/octeontx: fix partial Rx packet handling
- event/sw: fix xstats reset value
- examples/ipsec-secgw: fix default configuration
- examples/ipsec-secgw: fix GCM IV length
- examples/ipsec-secgw: fix SHA256-HMAC digest length

- examples/l3fwd-power: fix Rx interrupt disabling
- examples/multi\_process: fix client crash with sparse ports
- examples/vm\_power: fix build without i40e
- examples/vm\_power: fix type of cmdline token in cli
- kni: fix build with kernel 5.3
- lib/distributor: fix deadlock on aarch64
- lib: fix doxygen typos
- lib: fix log typos
- maintainers: update for stable branches
- malloc: fix realloc copy size
- malloc: fix realloc padded element size
- malloc: set pad to 0 on free
- mempool/dpaa2: report error on endless loop in mbuf release
- mk: remove library search path from binary
- net/af\_packet: fix stale sockets
- net/af\_packet: improve Tx statistics accuracy
- net/bnxt: cleanup comments
- net/bnxt: enforce IO barrier for doorbell command
- net/bnxt: expose some missing counters in port stats
- net/bnxt: fix async link handling and update
- net/bnxt: fix coding style
- net/bnxt: fix crash in secondary process
- net/bnxt: fix crash in xstats get
- net/bnxt: fix dereference before null check
- net/bnxt: fix log message level
- net/bnxt: fix mbuf free when clearing Tx queue
- net/bnxt: fix memory leak
- net/bnxt: fix multicast filter programming
- net/bnxt: fix Rx queue count
- net/bnxt: fix setting default MAC address
- net/bnxt: get default HWRM command timeout from FW
- net/bnxt: move macro definitions to header file
- net/bnxt: remove commented out code
- net/bnxt: remove duplicate barrier

- net/bnxt: remove redundant header file inclusion
- net/bnxt: remove unnecessary variable assignment
- net/bnxt: return error if setting link up fails
- net/bonding: fix LACP fast queue Rx handler
- net/bonding: fix link speed update in broadcast mode
- net/bonding: fix OOB access in other aggregator modes
- net/bonding: fix port ID check
- net/bonding: fix selection logic
- net/bonding: fix slave id types
- net/bonding: fix unicast packets filtering
- net/bonding: use non deprecated PCI API
- net/cxgbe: fix prefetch for non-coalesced Tx packets
- net/dpaa2: add retry and timeout in packet enqueue API
- net/dpaa2: fix possible use of uninitialized vars
- net/dpaa2: set port in mbuf
- net/e1000: fix link status
- net/fm10k: fix mbuf free in vector Rx
- net/fm10k: fix stats crash in multi-process
- net/i40e: downgrade error log
- net/i40e: fix address of first segment
- net/i40e: remove compiler barrier from NEON Rx
- net/i40e: remove memory barrier from NEON Rx
- net/igb: fix global variable multiple definitions
- net/igb: fix PHY status if PHY reset is not blocked
- net/ixgbe: enable new PF host mbox version
- net/ixgbe: fix address of first segment
- net/ixgbe: fix link status
- net/ixgbe: fix link status
- net/ixgbe: fix queue interrupt for X552/557
- net/ixgbe: fix VF RSS offloads configuration
- net/ixgbe: fix X553 speed capability
- net/ixgbe: remove memory barrier from NEON Rx
- net/ixgbe: remove redundant assignment
- net/ixgbe: support packet type with NEON

- net/mlx4: fix build on ppc64
- net/mlx5: fix Rx CQ doorbell synchronization on aarch64
- net/mlx: fix build with make and recent gcc
- net/mlx: fix debug build with icc
- net/qede: fix setting MTU
- net/qede: fix setting VLAN strip mode
- net/qede: limit Rx ring index read for debug
- net/sfc: fix adapter lock usage on rule creation
- net/tap: fix blocked Rx packets
- net/vhost: fix redundant queue state event
- net/virtio: fix descriptor addressed in Tx
- net/virtio: fix mbuf data and packet length mismatch
- net/virtio: fix Tx checksum offloads
- net/virtio: get all pending Rx packets in vectorized paths
- net/virtio: init MTU in case no control channel
- net/virtio: reject deferred Rx start
- net/virtio: reject deferred Tx start
- net/virtio-user: fix setting filters
- power: fix socket indicator value
- power: handle frequency increase with turbo disabled
- ring: enforce reading tail before slots
- security: fix doxygen fields
- service: use log for error messages
- test/bonding: fix LSC related cases
- test: fix global variable multiple definitions
- test/interrupt: account for race with callback
- test/lpm: fix measured cycles for delete
- test/mbuf: fix forged mbuf in clone test
- test: optimise fd closing in forks
- test/service: fix wait for service core
- test/test: fix test app defining unused variable
- usertools: fix pmdinfo with python 3 and pyelftools>=0.24
- version: 17.11.10-rc1
- vfiio: fix truncated BAR offset for 32-bit

- vhost: convert buffer addresses to GPA for logging
- vhost: fix IPv4 checksum
- vhost: fix slave request fd leak
- vhost: fix virtqueue not accessible
- vhost: fix vring address handling during live migration
- vhost: fix vring memory partially mapped
- vhost: forbid reallocation when running
- vhost: prevent zero copy mode if IOMMU is on
- vhost: protect vring access done by application
- vhost: translate incoming log address to GPA
- vhost: un-inline dirty pages logging functions

### 2.12.2 17.11.10 Validation

- Red Hat(R) Testing
  - RHEL 7.8
  - Functionality
    - \* PF assignment
    - \* VF assignment
    - \* vhost single/multi queues and cross-NUMA
    - \* vhostclient reconnect
    - \* vhost live migration with single/multi queues and cross-NUMA
    - \* OVS PVP
- Intel(R) Testing
  - Basic Intel(R) NIC(ixgbe and i40e) testing
    - \* PF (i40e)
    - \* PF (ixgbe)
    - \* VF
    - \* Compile Testing
    - \* Intel NIC single core/NIC performance
  - Basic cryptodev and virtio testing
    - \* cryptodev
    - \* vhost/virtio basic loopback, PVP and performance test
- Mellanox(R) Testing
  - Basic functionality with testpmd

- \* Tx/Rx
- \* xstats
- \* Timestamps
- \* Link status
- \* RTE flow and flow\_director
- \* RSS
- \* VLAN stripping and insertion
- \* Checksum/TSO
- \* ptype
- \* Multi-process
- ConnectX-5
  - \* RHEL 7.4
  - \* Kernel 3.10.0-693.el7.x86\_64
  - \* Driver MLNX\_OFED\_LINUX-4.7-3.2.9.0
  - \* fw 16.26.4012
- ConnectX-4 Lx
  - \* RHEL 7.4
  - \* Kernel 3.10.0-693.el7.x86\_64
  - \* Driver MLNX\_OFED\_LINUX-4.7-3.2.9.0
  - \* fw 14.26.4012
- Intel(R) Testing with Open vSwitch
  - OVS testing with OVS branches 2.10 and 2.19 with VSPERF
  - Tested NICs
    - \* i40e (X710)
    - \* ixgbe (82599ES)
  - Functionality
    - \* P2P
    - \* PVP
    - \* Hotplug
    - \* Multiqueue
    - \* Vhostuserclient reconnect
    - \* Vhost cross-NUMA awareness
    - \* Jumbo frames
    - \* Rate limiting



- \* QoS policer

### 2.12.3 17.11.10 Known Issues

- MLX4/5 and rte\_flow: count actions and rules with vlan items are not supported, and MLNX\_OFED\_LINUX 4.6-1.0.1.1 is not compatible
- DPDK 17.11.10 contains fixes up to DPDK v19.11. Issues identified/fixed in DPDK master branch after DPDK v19.11 may be present in DPDK 17.11.10

## 3.1 New Features

- **Increase minimum x86 ISA version to SSE4.2.**

Starting with version 17.08, DPDK requires SSE4.2 to run on x86. Previous versions required SSE3.

- **Added Service Core functionality.**

The service core functionality added to EAL allows DPDK to run services such as software PMDs on lcores without the application manually running them. The service core infrastructure allows flexibility of running multiple services on the same service lcore, and provides the application with powerful APIs to configure the mapping from service lcores to services.

- **Added Generic Receive Offload API.**

Added Generic Receive Offload (GRO) API support to reassemble TCP/IPv4 packets. The GRO API assumes all input packets have the correct checksums. GRO API doesn't update checksums for merged packets. If input packets are IP fragmented, the GRO API assumes they are complete packets (i.e. with L4 headers).

- **Added Fail-Safe PMD**

Added the new Fail-Safe PMD. This virtual device allows applications to support seamless hotplug of devices. See the `/nics/fail_safe` guide for more details about this driver.

- **Added support for generic flow API (`rte_flow`) on igb NICs.**

This API provides a generic means of configuring hardware to match specific ingress or egress traffic, altering its behavior and querying related counters according to any number of user-defined rules.

Added generic flow API support for Ethernet, IPv4, UDP, TCP and RAW pattern items with QUEUE actions. There are four types of filter support for this feature on igb.

- **Added support for generic flow API (`rte_flow`) on enic.**

Added flow API support for outer Ethernet, VLAN, IPv4, IPv6, UDP, TCP, SCTP, VxLAN and inner Ethernet, VLAN, IPv4, IPv6, UDP and TCP pattern items with QUEUE, MARK, FLAG and VOID actions for ingress traffic.

- **Added support for Chelsio T6 family of adapters**

The CXGBE PMD was updated to run Chelsio T6 family of adapters.

- **Added latency and performance improvements for cxgbe**

the Tx and Rx path in cxgbe were reworked to improve performance. In addition the latency was reduced for slow traffic.

- **Updated the bnxt PMD.**

Updated the bnxt PMD. The major enhancements include:

- Support MTU modification.
- Add support for LRO.
- Add support for VLAN filter and strip functionality.
- Additional enhancements to add support for more dev\_ops.
- Added PMD specific APIs mainly to control VF from PF.
- Update HWRM version to 1.7.7

- **Added support for Rx interrupts on mlx4 driver.**

Rx queues can be now be armed with an interrupt which will trigger on the next packet arrival.

- **Updated mlx5 driver.**

Updated the mlx5 driver including the following changes:

- Added vectorized Rx/Tx burst for x86.
- Added support for isolated mode from flow API.
- Reworked the flow drop action to implement in hardware classifier.
- Improved Rx interrupts management.

- **Updated szedata2 PMD.**

Added support for firmware with multiple Ethernet ports per physical port.

- **Updated dpaa2 PMD.**

Updated dpaa2 PMD. Major enhancements include:

- Added support for MAC Filter configuration.
- Added support for Segmented Buffers.
- Added support for VLAN filter and strip functionality.
- Additional enhancements to add support for more dev\_ops.
- Optimized the packet receive path

- **Reorganized the symmetric crypto operation structure.**

The crypto operation (`rte_crypto_sym_op`) has been reorganized as follows:

- Removed the `rte_crypto_sym_op_sess_type` field.
- Replaced the pointer and physical address of IV with offset from the start of the crypto operation.
- Moved length and offset of cipher IV to `rte_crypto_cipher_xform`.

- Removed “Additional Authentication Data” (AAD) length.
  - Removed digest length.
  - Removed AAD pointer and physical address from `auth` structure.
  - Added `aead` structure, containing parameters for AEAD algorithms.
- **Reorganized the crypto operation structure.**

The crypto operation (`rte_crypto_op`) has been reorganized as follows:

    - Added the `rte_crypto_op_sess_type` field.
    - The enumerations `rte_crypto_op_status` and `rte_crypto_op_type` have been modified to be `uint8_t` values.
    - Removed the field `opaque_data`.
    - Pointer to `rte_crypto_sym_op` has been replaced with a zero length array.
  - **Reorganized the crypto symmetric session structure.**

The crypto symmetric session structure (`rte_cryptodev_sym_session`) has been reorganized as follows:

    - The `dev_id` field has been removed.
    - The `driver_id` field has been removed.
    - The mempool pointer `mp` has been removed.
    - Replaced `private` marker with array of pointers to private data sessions `sess_private_data`.
  - **Updated cryptodev library.**
    - Added AEAD algorithm specific functions and structures, so it is not necessary to use a combination of cipher and authentication structures anymore.
    - Added helper functions for crypto device driver identification.
    - Added support for multi-device sessions, so a single session can be used in multiple drivers.
    - Added functions to initialize and free individual driver private data with the same session.
  - **Updated dpaa2\_sec crypto PMD.**

Added support for AES-GCM and AES-CTR.
  - **Updated the AESNI MB PMD.**

The AESNI MB PMD has been updated with additional support for:

    - 12-byte IV on AES Counter Mode, apart from the previous 16-byte IV.
  - **Updated the AES-NI GCM PMD.**

The AES-NI GCM PMD was migrated from the ISA-L library to the Multi Buffer library, as the latter library has Scatter Gather List support now. The migration entailed adding additional support for 192-bit keys.

- **Updated the Cryptodev Scheduler PMD.**

Added a multicore based distribution mode, which distributes the enqueued crypto operations among several slaves, running on different logical cores.

- **Added NXP DPAA2 Eventdev PMD.**

Added the new dpaa2 eventdev driver for NXP DPAA2 devices. See the “Event Device Drivers” document for more details on this new driver.

- **Added dpdk-test-eventdev test application.**

The dpdk-test-eventdev tool is a Data Plane Development Kit (DPDK) application that allows exercising various eventdev use cases. This application has a generic framework to add new eventdev based test cases to verify functionality and measure the performance parameters of DPDK eventdev devices.

## 3.2 Known Issues

- **Starting with version 17.08, libnuma is required to build DPDK.**

## 3.3 API Changes

- **Modified the `_rte_eth_dev_callback_process` function in the ethdev library.**

The function `_rte_eth_dev_callback_process()` has been modified. The return value has been changed from void to int and an extra parameter `void *ret_param` has been added.

- **Moved bypass functions from the `rte_ethdev` library to `ixgbe` PMD**

- The following `rte_ethdev` library functions were removed:

- \* `rte_eth_dev_bypass_event_show()`
- \* `rte_eth_dev_bypass_event_store()`
- \* `rte_eth_dev_bypass_init()`
- \* `rte_eth_dev_bypass_state_set()`
- \* `rte_eth_dev_bypass_state_show()`
- \* `rte_eth_dev_bypass_ver_show()`
- \* `rte_eth_dev_bypass_wd_reset()`
- \* `rte_eth_dev_bypass_wd_timeout_show()`
- \* `rte_eth_dev_wd_timeout_store()`

- The following `ixgbe` PMD functions were added:

- \* `rte_pmd_ixgbe_bypass_event_show()`
- \* `rte_pmd_ixgbe_bypass_event_store()`
- \* `rte_pmd_ixgbe_bypass_init()`
- \* `rte_pmd_ixgbe_bypass_state_set()`

```
* rte_pmd_ixgbe_bypass_state_show()
* rte_pmd_ixgbe_bypass_ver_show()
* rte_pmd_ixgbe_bypass_wd_reset()
* rte_pmd_ixgbe_bypass_wd_timeout_show()
* rte_pmd_ixgbe_bypass_wd_timeout_store()
```

- **Reworked `rte_cryptodev` library.**

The `rte_cryptodev` library has been reworked and updated. The following changes have been made to it:

- The crypto device type enumeration has been removed from `cryptodev` library.
- The function `rte_crypto_count_devtype()` has been removed, and replaced by the new function `rte_crypto_count_by_driver()`.
- Moved crypto device driver names definitions to the particular PMDs. These names are not public anymore.
- The `rte_cryptodev_configure()` function does not create the session mempool for the device anymore.
- The `rte_cryptodev_queue_pair_attach_sym_session()` and `rte_cryptodev_queue_pair_detach_sym_session()` functions require the new parameter `device id`.
- Parameters of `rte_cryptodev_sym_session_create()` were modified to accept `mempool`, instead of `device id` and `rte_crypto_sym_xform`.
- Removed `device id` parameter from `rte_cryptodev_sym_session_free()`.
- Added a new field `session_pool` to `rte_cryptodev_queue_pair_setup()`.
- Removed `aad_size` parameter from `rte_cryptodev_sym_capability_check_auth()`.
- Added `iv_size` parameter to `rte_cryptodev_sym_capability_check_auth()`.
- Removed `RTE_CRYPTO_OP_STATUS_ENQUEUED` from enum `rte_crypto_op_status`.

### 3.4 ABI Changes

- Changed type of `domain` field in `rte_pci_addr` to `uint32_t` to follow the PCI standard.
- Added new `rte_bus` experimental APIs available as operators within the `rte_bus` structure.
- Made `rte_devargs` structure internal device representation generic to prepare for a bus-agnostic EAL.
- **Reorganized the crypto operation structures.**

Some fields have been modified in the `rte_crypto_op` and `rte_crypto_sym_op` structures, as described in the [New Features](#) section.

- **Reorganized the crypto symmetric session structure.**

Some fields have been modified in the `rte_cryptodev_sym_session` structure, as described in the [New Features](#) section.

- **Reorganized the `rte_crypto_sym_cipher_xform` structure.**

- Added cipher IV length and offset parameters.
- Changed field size of key length from `size_t` to `uint16_t`.

- **Reorganized the `rte_crypto_sym_auth_xform` structure.**

- Added authentication IV length and offset parameters.
- Changed field size of AAD length from `uint32_t` to `uint16_t`.
- Changed field size of digest length from `uint32_t` to `uint16_t`.
- Removed AAD length.
- Changed field size of key length from `size_t` to `uint16_t`.

- Replaced `dev_type` enumeration with `uint8_t driver_id` in `rte_cryptodev_info` and `rte_cryptodev` structures.

- Removed `session_mp` from `rte_cryptodev_config`.

### 3.5 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```

librte_acl.so.2
librte_bitratestats.so.1
librte_cfgfile.so.2
librte_cmdline.so.2
+ librte_cryptodev.so.3
librte_distributor.so.1
+ librte_eal.so.5
+ librte_ethdev.so.7
+ librte_eventdev.so.2
+ librte_gro.so.1
librte_hash.so.2
librte_ip_frag.so.1
librte_jobstats.so.1
librte_kni.so.2
librte_kvargs.so.1
librte_latencystats.so.1
librte_lpm.so.2
librte_mbuf.so.3
librte_mempool.so.2
librte_meter.so.1
librte_metrics.so.1
librte_net.so.1
librte_pdump.so.1
librte_pipeline.so.3
librte_pmd_bond.so.1
librte_pmd_ring.so.2
librte_port.so.3
librte_power.so.1
librte_reorder.so.1
librte_ring.so.1

```

```
librte_sched.so.1  
librte_table.so.2  
librte_timer.so.1  
librte_vhost.so.3
```

## 3.6 Tested Platforms

- Intel(R) platforms with Mellanox(R) NICs combinations
  - Platform details:
    - \* Intel(R) Xeon(R) CPU E5-2697A v4 @ 2.60GHz
    - \* Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2640 @ 2.50GHz
  - OS:
    - \* Red Hat Enterprise Linux Server release 7.3 (Maipo)
    - \* Red Hat Enterprise Linux Server release 7.2 (Maipo)
    - \* Ubuntu 16.10
    - \* Ubuntu 16.04
    - \* Ubuntu 14.04
  - MLNX\_OFED: 4.1-1.0.2.0
  - NICs:
    - \* Mellanox(R) ConnectX(R)-3 Pro 40G MCX354A-FCC\_Ax (2x40G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1007
      - Firmware version: 2.40.5030
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013



- Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)

- Host interface: PCI Express 3.0 x8
- Device ID: 15b3:1015
- Firmware version: 14.18.2000
- \* Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.18.2000
- \* Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1017
  - Firmware version: 16.19.1200
- \* Mellanox(R) ConnectX-5 Ex EN 100G MCX516A-CDAT (2x100G)
  - Host interface: PCI Express 4.0 x16
  - Device ID: 15b3:1019
  - Firmware version: 16.19.1200
- Intel(R) platforms with Intel(R) NICs combinations
  - CPU
    - \* Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
    - \* Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU D-1541 @ 2.10GHz
    - \* Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
    - \* Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz
    - \* Intel(R) Xeon(R) CPU E5-2658 v2 @ 2.40GHz
    - \* Intel(R) Xeon(R) CPU E5-2658 v3 @ 2.20GHz
  - OS:
    - \* CentOS 7.2
    - \* Fedora 25
    - \* FreeBSD 11
    - \* Red Hat Enterprise Linux Server release 7.3
    - \* SUSE Enterprise Linux 12
    - \* Wind River Linux 8
    - \* Ubuntu 16.04

- \* Ubuntu 16.10
- NICs:
  - \* Intel(R) 82599ES 10 Gigabit Ethernet Controller
    - Firmware version: 0x61bf0001
    - Device id (pf/vf): 8086:10fb / 8086:10ed
    - Driver version: 4.0.1-k (ixgbe)
  - \* Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
    - Firmware version: 0x800001cf
    - Device id (pf/vf): 8086:15ad / 8086:15a8
    - Driver version: 4.2.5 (ixgbe)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
    - Firmware version: 6.01 0x80003205
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 2.0.19 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
    - Firmware version: 6.01 0x80003204
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 2.0.19 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
    - Firmware version: 6.01 0x80003221
    - Device id (pf/vf): 8086:158b
    - Driver version: 2.0.19 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
    - Firmware version: 6.01 0x8000321c
    - Device id (pf/vf): 8086:1583 / 8086:154c
    - Driver version: 2.0.19 (i40e)
  - \* Intel(R) Corporation I350 Gigabit Network Connection
    - Firmware version: 1.48, 0x800006e7
    - Device id (pf/vf): 8086:1521 / 8086:1520
    - Driver version: 5.2.13-k (igb)

## 4.1 New Features

- **Reorganized mbuf structure.**

The mbuf structure has been reorganized as follows:

- Align fields to facilitate the writing of `data_off`, `refcnt`, and `nb_segs` in one operation.
- Use 2 bytes for port and number of segments.
- Move the sequence number to the second cache line.
- Add a timestamp field.
- Set default value for `refcnt`, `next` and `nb_segs` at mbuf free.

- **Added mbuf raw free API.**

Moved `rte_mbuf_raw_free()` and `rte_pktmbuf_prefree_seg()` functions to the public API.

- **Added free Tx mbuf on demand API.**

Added a new function `rte_eth_tx_done_cleanup()` which allows an application to request the driver to release mbufs that are no longer in use from a Tx ring, independent of whether or not the `tx_rs_thresh` has been crossed.

- **Added device removal interrupt.**

Added a new ethdev event `RTE_ETH_DEV_INTR_RMV` to signify the sudden removal of a device. This event can be advertised by PCI drivers and enabled accordingly.

- **Added EAL dynamic log framework.**

Added new APIs to dynamically register named log types, and control the level of each type independently.

- **Added descriptor status ethdev API.**

Added a new API to get the status of a descriptor.

For Rx, it is almost similar to the `rx_descriptor_done` API, except it differentiates descriptors which are held by the driver and not returned to the hardware. For Tx, it is a new API.

- **Increased number of next hops for LPM IPv6 to 2<sup>21</sup>.**

The next\_hop field has been extended from 8 bits to 21 bits for IPv6.
- **Added VFIO hotplug support.**

Added hotplug support for VFIO in addition to the existing UIO support.
- **Added PowerPC support to pci probing for vfio-pci devices.**

Enabled sPAPR IOMMU based pci probing for vfio-pci devices.
- **Kept consistent PMD batching behavior.**

Removed the limit of fm10k/i40e/ixgbe Tx burst size and vhost Rx/Tx burst size in order to support the same policy of “make an best effort to Rx/Tx pkts” for PMDs.
- **Updated the ixgbe base driver.**

Updated the ixgbe base driver, including the following changes:

  - Add link block check for KR.
  - Complete HW initialization even if SFP is not present.
  - Add VF xcast promiscuous mode.
- **Added PowerPC support for i40e and its vector PMD.**

Enabled i40e PMD and its vector PMD by default in PowerPC.
- **Added VF max bandwidth setting in i40e.**

Enabled capability to set the max bandwidth for a VF in i40e.
- **Added VF TC min and max bandwidth setting in i40e.**

Enabled capability to set the min and max allocated bandwidth for a TC on a VF in i40e.
- **Added TC strict priority mode setting on i40e.**

There are 2 Tx scheduling modes supported for TCs by i40e HW: round robin mode and strict priority mode. By default the round robin mode is used. It is now possible to change the Tx scheduling mode for a TC. This is a global setting on a physical port.
- **Added i40e dynamic device personalization support.**
  - Added dynamic device personalization processing to i40e firmware.
- **Added Cloud Filter for QinQ steering to i40e.**
  - Added a QinQ cloud filter on the i40e PMD, for steering traffic to a VM using both VLAN tags. Note, this feature is not supported in Vector Mode.
- **Updated mlx5 PMD.**

Updated the mlx5 driver, including the following changes:

  - Added Generic flow API support for classification according to ether type.
  - Extended Generic flow API support for classification of IPv6 flow according to Vtc flow, Protocol and Hop limit.
  - Added Generic flow API support for FLAG action.
  - Added Generic flow API support for RSS action.

- Added support for TSO for non-tunneled and VXLAN packets.
- Added support for hardware Tx checksum offloads for VXLAN packets.
- Added support for user space Rx interrupt mode.
- Improved ConnectX-5 single core and maximum performance.

- **Updated mlx4 PMD.**

Updated the mlx4 driver, including the following changes:

- Added support for Generic flow API basic flow items and actions.
- Added support for device removal event.

- **Updated the sfc\_efx driver.**

- Added Generic Flow API support for Ethernet, VLAN, IPv4, IPv6, UDP and TCP pattern items with QUEUE action for ingress traffic.
- Added support for virtual functions (VFs).

- **Added LiquidIO network PMD.**

Added poll mode driver support for Cavium LiquidIO II server adapter VFs.

- **Added Atomic Rules Arkville PMD.**

Added a new poll mode driver for the Arkville family of devices from Atomic Rules. The net/ark PMD supports line-rate agnostic, multi-queue data movement on Arkville core FPGA instances.

- **Added support for NXP DPAA2 - FSLMC bus.**

Added the new bus “fslmc” driver for NXP DPAA2 devices. See the “Network Interface Controller Drivers” document for more details of this new driver.

- **Added support for NXP DPAA2 Network PMD.**

Added the new “dpaa2” net driver for NXP DPAA2 devices. See the “Network Interface Controller Drivers” document for more details of this new driver.

- **Added support for the Wind River Systems AVP PMD.**

Added a new networking driver for the AVP device type. These devices are specific to the Wind River Systems virtualization platforms.

- **Added vmxnet3 version 3 support.**

Added support for vmxnet3 version 3 which includes several performance enhancements such as configurable Tx data ring, Receive Data Ring, and the ability to register memory regions.

- **Updated the TAP driver.**

Updated the TAP PMD to:

- Support MTU modification.
- Support packet type for Rx.
- Support segmented packets on Rx and Tx.
- Speed up Rx on TAP when no packets are available.

- Support capturing traffic from another netdevice.
  - Dynamically change link status when the underlying interface state changes.
  - Added Generic Flow API support for Ethernet, VLAN, IPv4, IPv6, UDP and TCP pattern items with DROP, QUEUE and PASSTHRU actions for ingress traffic.
- **Added MTU feature support to Virtio and Vhost.**  
Implemented new Virtio MTU feature in Vhost and Virtio:
    - Add `rte_vhost_mtu_get()` API to Vhost library.
    - Enable Vhost PMD's MTU get feature.
    - Get max MTU value from host in Virtio PMD
  - **Added interrupt mode support for virtio-user.**  
Implemented Rxq interrupt mode and LSC support for virtio-user as a virtual device. Supported cases:
    - Rxq interrupt for virtio-user + vhost-user as the backend.
    - Rxq interrupt for virtio-user + vhost-kernel as the backend.
    - LSC interrupt for virtio-user + vhost-user as the backend.
  - **Added event driven programming model library (`rte_eventdev`).**  
This API introduces an event driven programming model.  
  
In a polling model, lcores poll ethdev ports and associated Rx queues directly to look for a packet. By contrast in an event driven model, lcores call the scheduler that selects packets for them based on programmer-specified criteria. The Eventdev library adds support for an event driven programming model, which offers applications automatic multicore scaling, dynamic load balancing, pipelining, packet ingress order maintenance and synchronization services to simplify application packet processing.  
  
By introducing an event driven programming model, DPDK can support both polling and event driven programming models for packet processing, and applications are free to choose whatever model (or combination of the two) best suits their needs.
  - **Added Software Eventdev PMD.**  
Added support for the software eventdev PMD. The software eventdev is a software based scheduler device that implements the eventdev API. This PMD allows an application to configure a pipeline using the eventdev library, and run the scheduling workload on a CPU core.
  - **Added Cavium OCTEONTX Eventdev PMD.**  
Added the new octeontx ssovf eventdev driver for OCTEONTX devices. See the “Event Device Drivers” document for more details on this new driver.
  - **Added information metrics library.**  
Added a library that allows information metrics to be added and updated by producers, typically other libraries, for later retrieval by consumers such as applications. It is intended to provide a reporting mechanism that is independent of other libraries such as ethdev.
  - **Added bit-rate calculation library.**

Added a library that can be used to calculate device bit-rates. Calculated bitrates are reported using the metrics library.

- **Added latency stats library.**

Added a library that measures packet latency. The collected statistics are jitter and latency. For latency the minimum, average, and maximum is measured.

- **Added NXP DPAA2 SEC crypto PMD.**

A new “dpaa2\_sec” hardware based crypto PMD for NXP DPAA2 devices has been added. See the “Crypto Device Drivers” document for more details on this driver.

- **Updated the Cryptodev Scheduler PMD.**

- Added a packet-size based distribution mode, which distributes the enqueued crypto operations among two slaves, based on their data lengths.
- Added fail-over scheduling mode, which enqueues crypto operations to a primary slave first. Then, any operation that cannot be enqueued is enqueued to a secondary slave.
- Added mode specific option support, so each scheduling mode can now be configured individually by the new API.

- **Updated the QAT PMD.**

The QAT PMD has been updated with additional support for:

- AES DOCSIS BPI algorithm.
- DES DOCSIS BPI algorithm.
- ZUC EEA3/EIA3 algorithms.

- **Updated the AESNI MB PMD.**

The AESNI MB PMD has been updated with additional support for:

- AES DOCSIS BPI algorithm.

- **Updated the OpenSSL PMD.**

The OpenSSL PMD has been updated with additional support for:

- DES DOCSIS BPI algorithm.

## 4.2 Resolved Issues

- **l2fwd-keepalive: Fixed unclean shutdowns.**

Added clean shutdown to l2fwd-keepalive so that it can free up stale resources used for inter-process communication.

## 4.3 Known Issues

- **LSC interrupt doesn't work for virtio-user + vhost-kernel.**



LSC interrupt cannot be detected when setting the backend, tap device, up/down as we fail to find a way to monitor such event.

## 4.4 API Changes

- The LPM `next_hop` field is extended from 8 bits to 21 bits for IPv6 while keeping ABI compatibility.

- **Reworked `rte_ring` library.**

The `rte_ring` library has been reworked and updated. The following changes have been made to it:

- Removed the build-time setting `CONFIG_RTE_RING_SPLIT_PROD_CONS`.
- Removed the build-time setting `CONFIG_RTE_LIBRTE_RING_DEBUG`.
- Removed the build-time setting `CONFIG_RTE_RING_PAUSE_REP_COUNT`.
- Removed the function `rte_ring_set_water_mark` as part of a general removal of watermarks support in the library.
- Added an extra parameter to the burst/bulk enqueue functions to return the number of free spaces in the ring after enqueue. This can be used by an application to implement its own watermark functionality.
- Added an extra parameter to the burst/bulk dequeue functions to return the number elements remaining in the ring after dequeue.
- Changed the return value of the enqueue and dequeue bulk functions to match that of the burst equivalents. In all cases, ring functions which operate on multiple packets now return the number of elements enqueued or dequeued, as appropriate. The updated functions are:

```
* rte_ring_mp_enqueue_bulk
* rte_ring_sp_enqueue_bulk
* rte_ring_enqueue_bulk
* rte_ring_mc_dequeue_bulk
* rte_ring_sc_dequeue_bulk
* rte_ring_dequeue_bulk
```

NOTE: the above functions all have different parameters as well as different return values, due to the other listed changes above. This means that all instances of the functions in existing code will be flagged by the compiler. The return value usage should be checked while fixing the compiler error due to the extra parameter.

- **Reworked `rte_vhost` library.**

The `rte_vhost` library has been reworked to make it generic enough so that the user could build other vhost-user drivers on top of it. To achieve this the following changes have been made:

- The following vhost-pmd APIs are removed:
  - \* `rte_eth_vhost_feature_disable`

- \* `rte_eth_vhost_feature_enable`
- \* `rte_eth_vhost_feature_get`
- The vhost API `rte_vhost_driver_callback_register(ops)` is reworked to be per vhost-user socket file. Thus, it takes one more argument: `rte_vhost_driver_callback_register(path,ops)`.
- The vhost API `rte_vhost_get_queue_num` is deprecated, instead, `rte_vhost_get_vring_num` should be used.
- The following macros are removed in `rte_virtio_net.h`
  - \* `VIRTIO_RXQ`
  - \* `VIRTIO_TXQ`
  - \* `VIRTIO_QNUM`
- The following net specific header files are removed in `rte_virtio_net.h`
  - \* `linux/virtio_net.h`
  - \* `sys/socket.h`
  - \* `linux/if.h`
  - \* `rte_ether.h`
- The vhost struct `virtio_net_device_ops` is renamed to `vhost_device_ops`
- The vhost API `rte_vhost_driver_session_start` is removed. Instead, `rte_vhost_driver_start` should be used, and there is no need to create a thread to call it.
- The vhost public header file `rte_virtio_net.h` is renamed to `rte_vhost.h`

## 4.5 ABI Changes

- **Reorganized the mbuf structure.**

The order and size of the fields in the `mbuf` structure changed, as described in the [New Features](#) section.

- The `rte_cryptodev_info.sym` structure has a new field `max_nb_sessions_per_qp` to support drivers which may support a limited number of sessions per `queue_pair`.

## 4.6 Removed Items

- KNI vhost support has been removed.
- The `dpdk_qat` sample application has been removed.

## 4.7 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```

librte_acl.so.2
+ librte_bitratestats.so.1
librte_cfgfile.so.2
librte_cmdline.so.2
librte_cryptodev.so.2
librte_distributor.so.1
+ librte_eal.so.4
librte_ethdev.so.6
+ librte_eventdev.so.1
librte_hash.so.2
librte_ip_frag.so.1
librte_jobstats.so.1
librte_kni.so.2
librte_kvargs.so.1
+ librte_latencystats.so.1
librte_lpm.so.2
+ librte_mbuf.so.3
librte_mempool.so.2
librte_meter.so.1
+ librte_metrics.so.1
librte_net.so.1
librte_pdump.so.1
librte_pipeline.so.3
librte_pmd_bond.so.1
librte_pmd_ring.so.2
librte_port.so.3
librte_power.so.1
librte_reorder.so.1
librte_ring.so.1
librte_sched.so.1
librte_table.so.2
librte_timer.so.1
librte_vhost.so.3

```

## 4.8 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
  - CPU
    - \* Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
    - \* Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
    - \* Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz
    - \* Intel(R) Xeon(R) CPU E5-2658 v2 @ 2.40GHz
    - \* Intel(R) Xeon(R) CPU E5-2658 v3 @ 2.20GHz
  - OS:

- \* CentOS 7.2
  - \* Fedora 25
  - \* FreeBSD 11
  - \* Red Hat Enterprise Linux Server release 7.3
  - \* SUSE Enterprise Linux 12
  - \* Wind River Linux 8
  - \* Ubuntu 16.04
  - \* Ubuntu 16.10
- NICs:
- \* Intel(R) 82599ES 10 Gigabit Ethernet Controller
    - Firmware version: 0x61bf0001
    - Device id (pf/vf): 8086:10fb / 8086:10ed
    - Driver version: 4.0.1-k (ixgbe)
  - \* Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
    - Firmware version: 0x800001cf
    - Device id (pf/vf): 8086:15ad / 8086:15a8
    - Driver version: 4.2.5 (ixgbe)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XL710-QDA1 (1x40G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1584 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1583 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Corporation I350 Gigabit Network Connection

- Firmware version: 1.48, 0x800006e7
- Device id (pf/vf): 8086:1521 / 8086:1520
- Driver version: 5.2.13-k (igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
  - Platform details:
    - \* Intel(R) Xeon(R) CPU E5-2697A v4 @ 2.60GHz
    - \* Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2640 @ 2.50GHz
  - OS:
    - \* Red Hat Enterprise Linux Server release 7.3 (Maipo)
    - \* Red Hat Enterprise Linux Server release 7.2 (Maipo)
    - \* Ubuntu 16.10
    - \* Ubuntu 16.04
    - \* Ubuntu 14.04
  - MLNX\_OFED: 4.0-2.0.0.0
  - NICs:
    - \* Mellanox(R) ConnectX(R)-3 Pro 40G MCX354A-FCC\_Ax (2x40G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1007
      - Firmware version: 2.40.5030
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)

- Host interface: PCI Express 3.0 x8
- Device ID: 15b3:1013
- Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015

- Firmware version: 14.18.2000
- \* Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.18.2000
- \* Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1017
  - Firmware version: 16.19.1200
- \* Mellanox(R) ConnectX-5 Ex EN 100G MCX516A-CDAT (2x100G)
  - Host interface: PCI Express 4.0 x16
  - Device ID: 15b3:1019
  - Firmware version: 16.19.1200
- IBM(R) Power8(R) with Mellanox(R) NICs combinations
  - Platform details:
    - \* Processor: POWER8E (raw), AltiVec supported
    - \* type-model: 8247-22L
    - \* Firmware FW810.21 (SV810\_108)
  - OS: Ubuntu 16.04 LTS PPC le
  - MLNX\_OFED: 4.0-2.0.0.0
  - NICs:
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
      - Host interface: PCI Express 3.0 x8
      - Device ID: 15b3:1013
      - Firmware version: 12.18.2000
    - \* Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)

- Host interface: PCI Express 3.0 x8
- Device ID: 15b3:1013
- Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000
- \* Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.2000



## DPDK RELEASE 17.02

## 5.1 New Features

- **Added support for representing buses in EAL**

The `rte_bus` structure was introduced into the EAL. This allows for devices to be represented by buses they are connected to. A new bus can be added to DPDK by extending the `rte_bus` structure and implementing the scan and probe functions. Once a new bus is registered using the provided APIs, new devices can be detected and initialized using bus scan and probe callbacks.

With this change, devices other than PCI or VDEV type can be represented in the DPDK framework.

- **Added generic EAL API for I/O device memory read/write operations.**

This API introduces 8 bit, 16 bit, 32 bit and 64 bit I/O device memory read/write operations along with “relaxed” versions.

Weakly-ordered architectures like ARM need an additional I/O barrier for device memory read/write access over PCI bus. By introducing the EAL abstraction for I/O device memory read/write access, the drivers can access I/O device memory in an architecture-agnostic manner. The relaxed version does not have an additional I/O memory barrier, which is useful in accessing the device registers of integrated controllers which is implicitly strongly ordered with respect to memory access.

- **Added generic flow API (`rte_flow`).**

This API provides a generic means to configure hardware to match specific ingress or egress traffic, alter its behavior and query related counters according to any number of user-defined rules.

In order to expose a single interface with an unambiguous behavior that is common to all poll-mode drivers (PMDs) the `rte_flow` API is slightly higher-level than the legacy filtering framework, which it encompasses and supersedes (including all functions and filter types) .

See the Generic flow API documentation for more information.

- **Added firmware version get API.**

Added a new function `rte_eth_dev_fw_version_get()` to fetch the firmware version for a given device.

- **Added APIs for MACsec offload support to the ixgbe PMD.**

Six new APIs have been added to the ixgbe PMD for MACsec offload support. The declarations for the APIs can be found in `rte_pmd_ixgbe.h`.

- **Added I219 NICs support.**

Added support for I219 Intel 1GbE NICs.

- **Added VF Daemon (VFD) for i40e. - EXPERIMENTAL**

This is an EXPERIMENTAL feature to enhance the capability of the DPDK PF as many VF management features are not currently supported by the kernel PF driver. Some new private APIs are implemented directly in the PMD without an abstraction layer. They can be used directly by some users who have the need.

The new APIs to control VFs directly from PF include:

- Set VF MAC anti-spoofing.
- Set VF VLAN anti-spoofing.
- Set TX loopback.
- Set VF unicast promiscuous mode.
- Set VF multicast promiscuous mode.
- Set VF MTU.
- Get/reset VF stats.
- Set VF MAC address.
- Set VF VLAN stripping.
- Vf VLAN insertion.
- Set VF broadcast mode.
- Set VF VLAN tag.
- Set VF VLAN filter.

VFD also includes VF to PF mailbox message management from an application. When the PF receives mailbox messages from the VF the PF should call the callback provided by the application to know if they're permitted to be processed.

As an EXPERIMENTAL feature, please be aware it can be changed or even removed without prior notice.

- **Updated the i40e base driver.**

Updated the i40e base driver, including the following changes:

- Replace existing legacy `memcpy()` calls with `i40e_memcpy()` calls.
- Use `BIT()` macro instead of bit fields.
- Add clear all WoL filters implementation.
- Add broadcast promiscuous control per VLAN.
- Remove unused `X722_SUPPORT` and `I40E_NDIS_SUPPORT` macros.

- **Updated the enic driver.**

- Set new Rx checksum flags in mbufs to indicate unknown, good or bad checksums.

- Fix set/remove of MAC addresses. Allow up to 64 addresses per device.
- Enable TSO on outer headers.

- **Added Solarflare libefx-based network PMD.**

Added a new network PMD which supports Solarflare SFN7xxx and SFN8xxx family of 10/40 Gbps adapters.

- **Updated the mlx4 driver.**

- Addressed a few bugs.

- **Added support for Mellanox ConnectX-5 adapters (mlx5).**

Added support for Mellanox ConnectX-5 family of 10/25/40/50/100 Gbps adapters to the existing mlx5 PMD.

- **Updated the mlx5 driver.**

- Improve Tx performance by using vector logic.
- Improve RSS balancing when number of queues is not a power of two.
- Generic flow API support for Ethernet, IPv4, IPv6, UDP, TCP, VLAN and VXLAN pattern items with DROP and QUEUE actions.
- Support for extended statistics.
- Addressed several data path bugs.
- As of MLNX\_OFED 4.0-1.0.1.0, the Toeplitz RSS hash function is not symmetric anymore for consistency with other PMDs.

- **virtio-user with vhost-kernel as another exceptional path.**

Previously, we upstreamed a virtual device, virtio-user with vhost-user as the backend as a way of enabling IPC (Inter-Process Communication) and user space container networking.

Virtio-user with vhost-kernel as the backend is a solution for the exception path, such as KNI, which exchanges packets with the kernel networking stack. This solution is very promising in:

- Maintenance: vhost and vhost-net (kernel) is an upstreamed and extensively used kernel module.
- Features: vhost-net is designed to be a networking solution, which has lots of networking related features, like multi-queue, TSO, multi-seg mbuf, etc.
- Performance: similar to KNI, this solution would use one or more kthreads to send/receive packets from user space DPDK applications, which has little impact on user space polling thread (except that it might enter into kernel space to wake up those kthreads if necessary).

- **Added virtio Rx interrupt support.**

Added a feature to enable Rx interrupt mode for virtio pci net devices as bound to VFIO (noiommu mode) and driven by virtio PMD.

With this feature, the virtio PMD can switch between polling mode and interrupt mode, to achieve best performance, and at the same time save power. It can work on both legacy

and modern virtio devices. In this mode, each `rxq` is mapped with an excluded MSIx interrupt.

See the Virtio Interrupt Mode documentation for more information.

- **Added ARMv8 crypto PMD.**

A new crypto PMD has been added, which provides combined mode cryptographic operations optimized for ARMv8 processors. The driver can be used to enhance performance in processing chained operations such as cipher + HMAC.

- **Updated the QAT PMD.**

The QAT PMD has been updated with additional support for:

- DES algorithm.
- Scatter-gather list (SGL) support.

- **Updated the AESNI MB PMD.**

- The Intel(R) Multi Buffer Crypto for IPsec library used in AESNI MB PMD has been moved to a new repository, in GitHub.
- Support has been added for single operations (cipher only and authentication only).

- **Updated the AES-NI GCM PMD.**

The AES-NI GCM PMD was migrated from the Multi Buffer library to the ISA-L library. The migration entailed adding additional support for:

- GMAC algorithm.
- 256-bit cipher key.
- Session-less mode.
- Out-of place processing
- Scatter-gather support for chained mbufs (only out-of place and destination mbuf must be contiguous)

- **Added crypto performance test application.**

Added a new performance test application for measuring performance parameters of PMDs available in the crypto tree.

- **Added Elastic Flow Distributor library (`rte_efd`).**

Added a new library which uses perfect hashing to determine a target/value for a given incoming flow key.

The library does not store the key itself for lookup operations, and therefore, lookup performance is not dependent on the key size. Also, the target/value can be any arbitrary value (8 bits by default). Finally, the storage requirement is much smaller than a hash-based flow table and therefore, it can better fit in CPU cache and scale to millions of flow keys.

See the Elastic Flow Distributor Library documentation in the Programmers Guide document, for more information.

## 5.2 Resolved Issues

### 5.2.1 Drivers

- **net/virtio: Fixed multiple process support.**

Fixed a few regressions introduced in recent releases that break the virtio multiple process support.

### 5.2.2 Examples

- **examples/ethtool: Fixed crash with non-PCI devices.**

Fixed issue where querying a non-PCI device was dereferencing non-existent PCI data resulting in a segmentation fault.

## 5.3 API Changes

- **Moved five APIs for VF management from the ethdev to the ixgbe PMD.**

The following five APIs for VF management from the PF have been removed from the ethdev, renamed, and added to the ixgbe PMD:

```
rte_eth_dev_set_vf_rate_limit()
rte_eth_dev_set_vf_rx()
rte_eth_dev_set_vf_rxmode()
rte_eth_dev_set_vf_tx()
rte_eth_dev_set_vf_vlan_filter()
```

The API's have been renamed to the following:

```
rte_pmd_ixgbe_set_vf_rate_limit()
rte_pmd_ixgbe_set_vf_rx()
rte_pmd_ixgbe_set_vf_rxmode()
rte_pmd_ixgbe_set_vf_tx()
rte_pmd_ixgbe_set_vf_vlan_filter()
```

The declarations for the API's can be found in `rte_pmd_ixgbe.h`.

## 5.4 ABI Changes

## 5.5 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```
librte_acl.so.2
librte_cfgfile.so.2
librte_cmdline.so.2
librte_cryptodev.so.2
librte_distributor.so.1
librte_eal.so.3
+ librte_ethdev.so.6
librte_hash.so.2
librte_ip_frag.so.1
```

```
librte_jobstats.so.1
librte_kni.so.2
librte_kvargs.so.1
librte_lpm.so.2
librte_mbuf.so.2
librte_mempool.so.2
librte_meter.so.1
librte_net.so.1
librte_pdump.so.1
librte_pipeline.so.3
librte_pmd_bond.so.1
librte_pmd_ring.so.2
librte_port.so.3
librte_power.so.1
librte_reorder.so.1
librte_ring.so.1
librte_sched.so.1
librte_table.so.2
librte_timer.so.1
librte_vhost.so.3
```

## 5.6 Tested Platforms

This release has been tested with the below list of CPU/device/firmware/OS. Each section describes a different set of combinations.

- Intel(R) platforms with Mellanox(R) NICs combinations
  - Platform details
    - \* Intel(R) Xeon(R) CPU E5-2697 v2 @ 2.70GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz
  - OS:
    - \* CentOS 7.0
    - \* Fedora 23
    - \* Fedora 24
    - \* FreeBSD 10.3
    - \* Red Hat Enterprise Linux 7.2
    - \* SUSE Enterprise Linux 12
    - \* Ubuntu 14.04 LTS
    - \* Ubuntu 15.10
    - \* Ubuntu 16.04 LTS
    - \* Wind River Linux 8
  - MLNX\_OFED: 4.0-1.0.1.0
  - NICs:
    - \* Mellanox(R) ConnectX(R)-3 Pro 40G MCX354A-FCC\_Ax (2x40G)

- Host interface: PCI Express 3.0 x8
- Device ID: 15b3:1007
- Firmware version: 2.40.5030
- \* Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013

- Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.18.1000
- \* Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.18.1000
- \* Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1017
  - Firmware version: 16.18.1000
- \* Mellanox(R) ConnectX-5 Ex EN 100G MCX516A-CDAT (2x100G)
  - Host interface: PCI Express 4.0 x16
  - Device ID: 15b3:1019
  - Firmware version: 16.18.1000
- IBM(R) Power8(R) with Mellanox(R) NICs combinations
  - Machine:
    - \* Processor: POWER8E (raw), AltiVec supported
      - type-model: 8247-22L
      - Firmware FW810.21 (SV810\_108)



- OS: Ubuntu 16.04 LTS PPC le
- MLNX\_OFED: 4.0-1.0.1.0
- NICs:
  - \* Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
    - Host interface: PCI Express 3.0 x16
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - Firmware version: 12.18.1000
  - \* Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013

- Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - Firmware version: 12.18.1000
- \* Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.18.1000
- \* Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - Firmware version: 14.18.1000
- \* Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1017
  - Firmware version: 16.18.1000
- Intel(R) platforms with Intel(R) NICs combinations
  - Platform details
    - \* Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
    - \* Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
    - \* Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
    - \* Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
    - \* Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz
    - \* Intel(R) Xeon(R) CPU E5-2658 v2 @ 2.40GHz

- OS:
  - \* CentOS 7.2
  - \* Fedora 25
  - \* FreeBSD 11
  - \* Red Hat Enterprise Linux Server release 7.3
  - \* SUSE Enterprise Linux 12
  - \* Wind River Linux 8
  - \* Ubuntu 16.04
  - \* Ubuntu 16.10
- NICs:
  - \* Intel(R) 82599ES 10 Gigabit Ethernet Controller
    - Firmware version: 0x61bf0001
    - Device id (pf/vf): 8086:10fb / 8086:10ed
    - Driver version: 4.0.1-k (ixgbe)
  - \* Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
    - Firmware version: 0x800001cf
    - Device id (pf/vf): 8086:15ad / 8086:15a8
    - Driver version: 4.2.5 (ixgbe)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XL710-QDA1 (1x40G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1584 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  - \* Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1583 / 8086:154c
    - Driver version: 1.5.23 (i40e)

- \* Intel(R) Corporation I350 Gigabit Network Connection
  - Firmware version: 1.48, 0x800006e7
  - Device id (pf/vf): 8086:1521 / 8086:1520
  - Driver version: 5.2.13-k (igb)

## DPDK RELEASE 16.11

## 6.1 New Features

- **Added software parser for packet type.**

- Added a new function `rte_pktmbuf_read()` to read the packet data from an mbuf chain, linearizing if required.
- Added a new function `rte_net_get_ptype()` to parse an Ethernet packet in an mbuf chain and retrieve its packet type from software.
- Added new functions `rte_get_ptype_*` to dump a packet type as a string.

- **Improved offloads support in mbuf.**

- Added a new function `rte_raw_cksum_mbuf()` to process the checksum of data embedded in an mbuf chain.
- Added new Rx checksum flags in mbufs to describe more states: unknown, good, bad, or not present (useful for virtual drivers). This modification was done for IP and L4.
- Added a new Rx LRO mbuf flag, used when packets are coalesced. This flag indicates that the segment size of original packets is known.

- **Added vhost-user dequeue zero copy support.**

The copy in the dequeue path is avoided in order to improve the performance. In the VM2VM case, the boost is quite impressive. The bigger the packet size, the bigger performance boost you may get. However, for the VM2NIC case, there are some limitations, so the boost is not as impressive as the VM2VM case. It may even drop quite a bit for small packets.

For that reason, this feature is disabled by default. It can be enabled when the `RTE_VHOST_USER_DEQUEUE_ZERO_COPY` flag is set. Check the VHost section of the Programming Guide for more information.

- **Added vhost-user indirect descriptors support.**

If the indirect descriptor feature is enabled, each packet sent by the guest will take exactly one slot in the enqueue virtqueue. Without this feature, as in the current version, even 64 bytes packets take two slots with Virtio PMD on guest side.

The main impact is better performance for 0% packet loss use-cases, as it behaves as if the virtqueue size was enlarged, so more packets can be buffered in the case of system

perturbations. On the downside, small performance degradations were measured when running micro-benchmarks.

- **Added vhost PMD xstats.**

Added extended statistics to vhost PMD from a per port perspective.

- **Supported offloads with virtio.**

Added support for the following offloads in virtio:

- Rx/Tx checksums.
- LRO.
- TSO.

- **Added virtio NEON support for ARM.**

Added NEON support for ARM based virtio.

- **Updated the ixgbe base driver.**

Updated the ixgbe base driver, including the following changes:

- Added X550em\_a 10G PHY support.
- Added support for flow control auto negotiation for X550em\_a 1G PHY.
- Added X550em\_a FW ALEF support.
- Increased mailbox version to `ixgbe_mbox_api_13`.
- Added two MAC operations for Hyper-V support.

- **Added APIs for VF management to the ixgbe PMD.**

Eight new APIs have been added to the ixgbe PMD for VF management from the PF. The declarations for the API's can be found in `rte_pmd_ixgbe.h`.

- **Updated the enic driver.**

- Added update to use interrupt for link status checking instead of polling.
- Added more flow director modes on UCS Blade with firmware version  $\geq 2.0(13e)$ .
- Added full support for MTU update.
- Added support for the `rte_eth_rx_queue_count` function.

- **Updated the mlx5 driver.**

- Added support for RSS hash results.
- Added several performance improvements.
- Added several bug fixes.

- **Updated the QAT PMD.**

The QAT PMD was updated with additional support for:

- MD5\_HMAC algorithm.
- SHA224-HMAC algorithm.
- SHA384-HMAC algorithm.

- GMAC algorithm.
- KASUMI (F8 and F9) algorithm.
- 3DES algorithm.
- NULL algorithm.
- C3XXX device.
- C62XX device.
- **Added openssl PMD.**

A new crypto PMD has been added, which provides several ciphering and hashing algorithms. All cryptography operations use the Openssl library crypto API.
- **Updated the IPsec example.**

Updated the IPsec example with the following support:

  - Configuration file support.
  - AES CBC IV generation with cipher forward function.
  - AES GCM/CTR mode.
- **Added support for new gcc -march option.**

The GCC 4.9 `-march` option supports the Intel processor code names. The config option `RTE_MACHINE` can be used to pass code names to the compiler via the `-march` flag.

## 6.2 Resolved Issues

### 6.2.1 Drivers

- **enic: Fixed several flow director issues.**
- **enic: Fixed inadvertent setting of L4 checksum ptype on ICMP packets.**
- **enic: Fixed high driver overhead when servicing Rx queues beyond the first.**

## 6.3 Known Issues

- **L3fwd-power app does not work properly when Rx vector is enabled.**

The L3fwd-power app doesn't work properly with some drivers in vector mode since the queue monitoring works differently between scalar and vector modes leading to incorrect frequency scaling. In addition, L3fwd-power application requires the mbuf to have correct packet type set but in some drivers the vector mode must be disabled for this.

Therefore, in order to use L3fwd-power, vector mode should be disabled via the config file.
- **Digest address must be supplied for crypto auth operation on QAT PMD.**

The cryptodev API specifies that if the `rte_crypto_sym_op.digest.data` field, and by inference the `digest.phys_addr` field which points to the same location, is not set for an auth

operation the driver is to understand that the digest result is located immediately following the region over which the digest is computed. The QAT PMD doesn't correctly handle this case and reads and writes to an incorrect location.

Callers can workaround this by always supplying the digest virtual and physical address fields in the `rte_crypto_sym_op` for an auth operation.

## 6.4 API Changes

- The driver naming convention has been changed to make them more consistent. It especially impacts `--vdev` arguments. For example `eth_pcap` becomes `net_pcap` and `cryptodev_aesni_mb_pmd` becomes `crypto_aesni_mb`.

For backward compatibility an alias feature has been enabled to support the original names.

- The log history has been removed.
- The `rte_ivshmem` feature (including library and EAL code) has been removed in 16.11 because it had some design issues which were not planned to be fixed.
- The `file_name` data type of `struct rte_port_source_params` and `struct rte_port_sink_params` is changed from `char *` to `const char *`.
- **Improved device/driver hierarchy and generalized hotplugging.**

The device and driver relationship has been restructured by introducing generic classes. This paves the way for having PCI, VDEV and other device types as instantiated objects rather than classes in themselves. Hotplugging has also been generalized into EAL so that Ethernet or crypto devices can use the common infrastructure.

- Removed `pmd_type` as a way of segregation of devices.
- Moved `numa_node` and `devargs` into `rte_driver` from `rte_pci_driver`. These can now be used by any instantiated object of `rte_driver`.
- Added `rte_device` class and all PCI and VDEV devices inherit from it
- Renamed `devinit/devuninit` handlers to `probe/remove` to make it more semantically correct with respect to the device  $\Leftrightarrow$  driver relationship.
- Moved hotplugging support to EAL. Hereafter, PCI and vdev can use the APIs `rte_eal_dev_attach` and `rte_eal_dev_detach`.
- Renamed helpers and support macros to make them more synonymous with their device types (e.g. `PMD_REGISTER_DRIVER` => `RTE_PMD_REGISTER_PCI`).
- Device naming functions have been generalized from `ethdev` and `cryptodev` to EAL. `rte_eal_pci_device_name` has been introduced for obtaining unique device name from PCI Domain-BDF description.
- Virtual device registration APIs have been added: `rte_eal_vdrv_register` and `rte_eal_vdrv_unregister`.



## 6.5 ABI Changes

## 6.6 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```
librte_acl.so.2
librte_cfgfile.so.2
librte_cmdline.so.2
+ librte_cryptodev.so.2
librte_distributor.so.1
+ librte_eal.so.3
+ librte_ethdev.so.5
librte_hash.so.2
librte_ip_frag.so.1
librte_jobstats.so.1
librte_kni.so.2
librte_kvargs.so.1
librte_lpm.so.2
librte_mbuf.so.2
librte_mempool.so.2
librte_meter.so.1
librte_net.so.1
librte_pdump.so.1
librte_pipeline.so.3
librte_pmd_bond.so.1
librte_pmd_ring.so.2
librte_port.so.3
librte_power.so.1
librte_reorder.so.1
librte_ring.so.1
librte_sched.so.1
librte_table.so.2
librte_timer.so.1
librte_vhost.so.3
```

## 6.7 Tested Platforms

1. SuperMicro 1U
  - BIOS: 1.0c
  - Processor: Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
2. SuperMicro 1U
  - BIOS: 1.0a
  - Processor: Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
  - Onboard NIC: Intel(R) X552/X557-AT (2x10G)
    - Firmware-version: 0x800001cf
    - Device ID (PF/VF): 8086:15ad /8086:15a8
  - kernel driver version: 4.2.5 (ixgbe)
3. SuperMicro 2U

- BIOS: 1.0a
  - Processor: Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
4. Intel(R) Server board S2600GZ
    - BIOS: SE5C600.86B.02.02.0002.122320131210
    - Processor: Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
  5. Intel(R) Server board W2600CR
    - BIOS: SE5C600.86B.02.01.0002.082220131453
    - Processor: Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
  6. Intel(R) Server board S2600CWT
    - BIOS: SE5C610.86B.01.01.0009.060120151350
    - Processor: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
  7. Intel(R) Server board S2600WTT
    - BIOS: SE5C610.86B.01.01.0005.101720141054
    - Processor: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
  8. Intel(R) Server board S2600WTT
    - BIOS: SE5C610.86B.11.01.0044.090120151156
    - Processor: Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz
  9. Intel(R) Server board S2600WTT
    - Processor: Intel(R) Xeon(R) CPU E5-2697 v2 @ 2.70GHz
  10. Intel(R) Server
    - Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz
  11. IBM(R) Power8(R)
    - Machine type-model: 8247-22L
    - Firmware FW810.21 (SV810\_108)
    - Processor: POWER8E (raw), AltiVec supported

## 6.8 Tested NICs

1. Intel(R) Ethernet Controller X540-AT2
  - Firmware version: 0x80000389
  - Device id (pf): 8086:1528
  - Driver version: 3.23.2 (ixgbe)
2. Intel(R) 82599ES 10 Gigabit Ethernet Controller
  - Firmware version: 0x61bf0001
  - Device id (pf/vf): 8086:10fb / 8086:10ed

- Driver version: 4.0.1-k (ixgbe)
3. Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
    - Firmware version: 0x800001cf
    - Device id (pf/vf): 8086:15ad / 8086:15a8
    - Driver version: 4.2.5 (ixgbe)
  4. Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  5. Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  6. Intel(R) Ethernet Converged Network Adapter XL710-QDA1 (1x40G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1584 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  7. Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2x40G)
    - Firmware version: 5.05
    - Device id (pf/vf): 8086:1583 / 8086:154c
    - Driver version: 1.5.23 (i40e)
  8. Intel(R) Corporation I350 Gigabit Network Connection
    - Firmware version: 1.48, 0x800006e7
    - Device id (pf/vf): 8086:1521 / 8086:1520
    - Driver version: 5.2.13-k (igb)
  9. Intel(R) Ethernet Multi-host Controller FM10000
    - Firmware version: N/A
    - Device id (pf/vf): 8086:15d0
    - Driver version: 0.17.0.9 (fm10k)
  10. Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
    - Host interface: PCI Express 3.0 x8
    - Device ID: 15b3:1013
    - MLNX\_OFED: 3.4-1.0.0.0
    - Firmware version: 12.17.1010

11. Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
12. Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
13. Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
14. Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
15. Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
  - Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
16. Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
17. Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
  - Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0

- Firmware version: 12.17.1010
18. Mellanox(R) ConnectX(R)-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
- Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
19. Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
- Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
20. Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
- Host interface: PCI Express 3.0 x16
  - Device ID: 15b3:1013
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 12.17.1010
21. Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
- Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 14.17.1010
22. Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
- Host interface: PCI Express 3.0 x8
  - Device ID: 15b3:1015
  - MLNX\_OFED: 3.4-1.0.0.0
  - Firmware version: 14.17.1010

## 6.9 Tested OSes

- CentOS 7.2
- Fedora 23
- Fedora 24
- FreeBSD 10.3
- FreeBSD 11

- Red Hat Enterprise Linux Server release 6.7 (Santiago)
- Red Hat Enterprise Linux Server release 7.0 (Maipo)
- Red Hat Enterprise Linux Server release 7.2 (Maipo)
- SUSE Enterprise Linux 12
- Wind River Linux 6.0.0.26
- Wind River Linux 8
- Ubuntu 14.04
- Ubuntu 15.04
- Ubuntu 16.04

## 7.1 New Features

- **Removed the mempool cache memory if caching is not being used.**

The size of the mempool structure is reduced if the per-lcore cache is disabled.

- **Added mempool external cache for non-EAL thread.**

Added new functions to create, free or flush a user-owned mempool cache for non-EAL threads. Previously the caching was always disabled on these threads.

- **Changed the memory allocation scheme in the mempool library.**

- Added the ability to allocate a large mempool in fragmented virtual memory.
- Added new APIs to populate a mempool with memory.
- Added an API to free a mempool.
- Modified the API of the `rte_mempool_obj_iter()` function.
- Dropped the specific Xen Dom0 code.
- Dropped the specific anonymous mempool code in `testpmd`.

- **Added a new driver for Broadcom NetXtreme-C devices.**

Added the new `bnxt` driver for Broadcom NetXtreme-C devices. See the “Network Interface Controller Drivers” document for more details on this new driver.

- **Added a new driver for ThunderX nicvf devices.**

Added the new `thunderx` net driver for ThunderX nicvf devices. See the “Network Interface Controller Drivers” document for more details on this new driver.

- **Added mailbox interrupt support for ixgbe and igb VFs.**

When the physical NIC link comes up or down, the PF driver will send a mailbox message to notify each VF. To handle this link up/down event, support have been added for a mailbox interrupt to receive the message and allow the application to register a callback for it.

- **Updated the ixgbe base driver.**

The ixgbe base driver was updated with changes including the following:

- Added `sgmii` link for X550.
- Added MAC link setup for X550a SFP and SFP+.

- Added KR support for X550em\_a.
  - Added new PHY definitions for M88E1500.
  - Added support for the VLVF to be bypassed when adding/removing a VFTA entry.
  - Added X550a flow control auto negotiation support.
- **Updated the i40e base driver.**  
Updated the i40e base driver including support for new devices IDs.
  - **Updated the enic driver.**  
The enic driver was updated with changes including the following:
    - Optimized the Tx function.
    - Added Scattered Rx capability.
    - Improved packet type identification.
    - Added MTU update in non Scattered Rx mode and enabled MTU of up to 9208 with UCS Software release 2.2 on 1300 series VICs.
  - **Updated the mlx5 driver.**  
The mlx5 driver was updated with changes including the following:
    - Data path was refactored to bypass Verbs to improve RX and TX performance.
    - Removed compilation parameters for inline send, `MLX5_MAX_INLINE`, and added command line parameter instead, `txq_inline`.
    - Improved TX scatter gather support: Removed compilation parameter `MLX5_PMD_SGE_WR_N`. Scatter-gather elements is set to the maximum value the NIC supports. Removed linearization logic, this decreases the memory consumption of the PMD.
    - Improved jumbo frames support, by dynamically setting RX scatter gather elements according to the MTU and mbuf size, no need for compilation parameter `MLX5_PMD_SGE_WR_N`
  - **Added support for virtio on IBM POWER8.**  
The ioports are mapped in memory when using Linux UIO.
  - **Added support for Virtio in containers.**  
Add a new virtual device, named `virtio_user`, to support virtio for containers.  
Known limitations:
    - Control queue and multi-queue are not supported yet.
    - Doesn't work with `--huge-unlink`.
    - Doesn't work with `--no-huge`.
    - Doesn't work when there are more than `VHOST_MEMORY_MAX_NREGIONS (8)` hugepages.
    - Root privilege is required for sorting hugepages by physical address.
    - Can only be used with the vhost user backend.



- **Added vhost-user client mode.**

DPDK vhost-user now supports client mode as well as server mode. Client mode is enabled when the `RTE_VHOST_USER_CLIENT` flag is set while calling `rte_vhost_driver_register`.

When DPDK vhost-user restarts from a normal or abnormal exit (such as a crash), the client mode allows DPDK to establish the connection again. Note that QEMU version v2.7 or above is required for this feature.

DPDK vhost-user will also try to reconnect by default when:

- The first connect fails (for example when QEMU is not started yet).
- The connection is broken (for example when QEMU restarts).

It can be turned off by setting the `RTE_VHOST_USER_NO_RECONNECT` flag.

- **Added NSH packet recognition in i40e.**

- **Added AES-CTR support to AESNI MB PMD.**

Now AESNI MB PMD supports 128/192/256-bit counter mode AES encryption and decryption.

- **Added AES counter mode support for Intel QuickAssist devices.**

Enabled support for the AES CTR algorithm for Intel QuickAssist devices. Provided support for algorithm-chaining operations.

- **Added KASUMI SW PMD.**

A new Crypto PMD has been added, which provides KASUMI F8 (UEA1) ciphering and KASUMI F9 (UIA1) hashing.

- **Added multi-writer support for RTE Hash with Intel TSX.**

The following features/modifications have been added to `rte_hash` library:

- Enabled application developers to use an extra flag for `rte_hash` creation to specify default behavior (multi-thread safe/unsafe) with the `rte_hash_add_key` function.
- Changed the Cuckoo Hash Search algorithm to breadth first search for multi-writer routines and split Cuckoo Hash Search and Move operations in order to reduce transactional code region and improve TSX performance.
- Added a hash multi-writer test case to the test app.

- **Improved IP Pipeline Application.**

The following features have been added to the `ip_pipeline` application:

- Configure the MAC address in the routing pipeline and automatic route updates with change in link state.
- Enable RSS per network interface through the configuration file.
- Streamline the CLI code.

- **Added keepalive enhancements.**

Added support for reporting of core states other than “dead” to monitoring applications, enabling the support of broader liveness reporting to external processes.

- **Added packet capture framework.**

- A new library `librte_pdump` is added to provide a packet capture API.
- A new `app/pdump` tool is added to demonstrate capture packets in DPDK.

- **Added floating VEB support for i40e PF driver.**

A “floating VEB” is a special Virtual Ethernet Bridge (VEB) which does not have an upload port, but instead is used for switching traffic between virtual functions (VFs) on a port.

For information on this feature, please see the “I40E Poll Mode Driver” section of the “Network Interface Controller Drivers” document.

- **Added support for live migration of a VM with SRIOV VF.**

Live migration of a VM with Virtio and VF PMD’s using the bonding PMD.

## 7.2 Resolved Issues

### 7.2.1 EAL

- **igb\_uio: Fixed possible mmap failure for Linux >= 4.5.**

The mmaping of the iomem range of the PCI device fails for kernels that enabled the `CONFIG_IO_STRICT_DEVMEM` option. The error seen by the user is as similar to the following:

```
EAL: pci_map_resource():
      cannot mmap(39, 0x7f1c51800000, 0x100000, 0x0):
      Invalid argument (0xffffffffffffffff)
```

The `CONFIG_IO_STRICT_DEVMEM` kernel option was introduced in Linux v4.5.

The issues was resolve by updating `igb_uio` to stop reserving PCI memory resources. From the kernel point of view the iomem region looks like idle and mmap works again. This matches the `uio_pci_generic` usage.

### 7.2.2 Drivers

- **i40e: Fixed vlan stripping from inner header.**

Previously, for tunnel packets, such as VXLAN/NVGRE, the vlan tags of the inner header will be stripped without putting vlan info to descriptor. Now this issue is fixed by disabling vlan stripping from inner header.

- **i40e: Fixed the type issue of a single VLAN type.**

Currently, if a single VLAN header is added in a packet, it’s treated as inner VLAN. But generally, a single VLAN header is treated as the outer VLAN header. This issue is fixed by changing corresponding register for single VLAN.

- **enic: Fixed several issues when stopping then restarting ports and queues.**

Fixed several crashes related to stopping then restarting ports and queues. Fixed possible crash when re-configuring the number of Rx queue descriptors.

- **enic: Fixed Rx data mis-alignment if mbuf data offset modified.**

Fixed possible Rx corruption when mbufs were returned to a pool with data offset other than `RTE_PKTMBUF_HEADROOM`.

- **enic: Fixed Tx IP/UDP/TCP checksum offload and VLAN insertion.**
- **enic: Fixed Rx error and missed counters.**

### 7.2.3 Libraries

- **mbuf: Fixed refcnt update when detaching.**

Fix the `rte_pktmbuf_detach()` function to decrement the direct mbuf's reference counter. The previous behavior was not to affect the reference counter. This lead to a memory leak of the direct mbuf.

### 7.2.4 Examples

### 7.2.5 Other

## 7.3 Known Issues

## 7.4 API Changes

- The following counters are removed from the `rte_eth_stats` structure:
  - `ibadcrc`
  - `ibadlen`
  - `imcasts`
  - `fdirmatch`
  - `fdirmiss`
  - `tx_pause_xon`
  - `rx_pause_xon`
  - `tx_pause_xoff`
  - `rx_pause_xoff`
- The extended statistics are fetched by ids with `rte_eth_xstats_get` after a lookup by name `rte_eth_xstats_get_names`.
- The function `rte_eth_dev_info_get` fill the new fields `nb_rx_queues` and `nb_tx_queues` in the structure `rte_eth_dev_info`.
- The vhost function `rte_vring_available_entries` is renamed to `rte_vhost_avail_entries`.
- All existing vhost APIs and callbacks with `virtio_net` struct pointer as the parameter have been changed due to the ABI refactoring described below. It is replaced by `int vid`.

- The function `rte_vhost_enqueue_burst` no longer supports concurrent enqueueing packets to the same queue.
- The function `rte_eth_dev_set_mtu` adds a new return value `-EBUSY`, which indicates the operation is forbidden because the port is running.
- The script `dpdk_nic_bind.py` is renamed to `dpdk-devbind.py`. And the script `setup.sh` is renamed to `dpdk-setup.sh`.

## 7.5 ABI Changes

- The `rte_port_source_params` structure has new fields to support PCAP files. It was already in release 16.04 with `RTE_NEXT_ABI` flag.
- The `rte_eth_dev_info` structure has new fields `nb_rx_queues` and `nb_tx_queues` to support the number of queues configured by software.
- A Vhost ABI refactoring has been made: the `virtio_net` structure is no longer exported directly to the application. Instead, a handle, `vid`, has been used to represent this structure internally.

## 7.6 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```
+ libethdev.so.4
  librte_acl.so.2
  librte_cfgfile.so.2
  librte_cmdline.so.2
  librte_cryptodev.so.1
  librte_distributor.so.1
  librte_eal.so.2
  librte_hash.so.2
  librte_ip_frag.so.1
  librte_ivshmem.so.1
  librte_jobstats.so.1
  librte_kni.so.2
  librte_kvargs.so.1
  librte_lpm.so.2
  librte_mbuf.so.2
+ librte_mempool.so.2
  librte_meter.so.1
  librte_pdump.so.1
  librte_pipeline.so.3
  librte_pmd_bond.so.1
  librte_pmd_ring.so.2
+ librte_port.so.3
  librte_power.so.1
  librte_reorder.so.1
  librte_ring.so.1
  librte_sched.so.1
  librte_table.so.2
  librte_timer.so.1
+ librte_vhost.so.3
```

## 7.7 Tested Platforms

1. SuperMicro 1U
  - BIOS: 1.0c
  - Processor: Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
2. SuperMicro 1U
  - BIOS: 1.0a
  - Processor: Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
  - Onboard NIC: Intel(R) X552/X557-AT (2x10G)
    - Firmware-version: 0x800001cf
    - Device ID (PF/VF): 8086:15ad /8086:15a8
  - kernel driver version: 4.2.5 (ixgbe)
3. SuperMicro 2U
  - BIOS: 1.0a
  - Processor: Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
4. Intel(R) Server board S2600GZ
  - BIOS: SE5C600.86B.02.02.0002.122320131210
  - Processor: Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
5. Intel(R) Server board W2600CR
  - BIOS: SE5C600.86B.02.01.0002.082220131453
  - Processor: Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
6. Intel(R) Server board S2600CWT
  - BIOS: SE5C610.86B.01.01.0009.060120151350
  - Processor: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
7. Intel(R) Server board S2600WTT
  - BIOS: SE5C610.86B.01.01.0005.101720141054
  - Processor: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
8. Intel(R) Server board S2600WTT
  - BIOS: SE5C610.86B.11.01.0044.090120151156
  - Processor: Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz

## 7.8 Tested NICs

1. Intel(R) Ethernet Controller X540-AT2
  - Firmware version: 0x80000389

- Device id (pf): 8086:1528
  - Driver version: 3.23.2 (ixgbe)
2. Intel(R) 82599ES 10 Gigabit Ethernet Controller
    - Firmware version: 0x61bf0001
    - Device id (pf/vf): 8086:10fb / 8086:10ed
    - Driver version: 4.0.1-k (ixgbe)
  3. Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
    - Firmware version: 0x800001cf
    - Device id (pf/vf): 8086:15ad / 8086:15a8
    - Driver version: 4.2.5 (ixgbe)
  4. Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
    - Firmware version: 5.04
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.4.26 (i40e)
  5. Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
    - Firmware version: 5.04
    - Device id (pf/vf): 8086:1572 / 8086:154c
    - Driver version: 1.4.25 (i40e)
  6. Intel(R) Ethernet Converged Network Adapter XL710-QDA1 (1x40G)
    - Firmware version: 5.04
    - Device id (pf/vf): 8086:1584 / 8086:154c
    - Driver version: 1.4.25 (i40e)
  7. Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2x40G)
    - Firmware version: 5.04
    - Device id (pf/vf): 8086:1583 / 8086:154c
    - Driver version: 1.4.25 (i40e)
  8. Intel(R) Corporation I350 Gigabit Network Connection
    - Firmware version: 1.48, 0x800006e7
    - Device id (pf/vf): 8086:1521 / 8086:1520
    - Driver version: 5.2.13-k (igb)
  9. Intel(R) Ethernet Multi-host Controller FM10000
    - Firmware version: N/A
    - Device id (pf/vf): 8086:15d0
    - Driver version: 0.17.0.9 (fm10k)

## 7.9 Tested OSes

- CentOS 7.0
- Fedora 23
- Fedora 24
- FreeBSD 10.3
- Red Hat Enterprise Linux 7.2
- SUSE Enterprise Linux 12
- Ubuntu 15.10
- Ubuntu 16.04 LTS
- Wind River Linux 8

## 8.1 New Features

- **Added function to check primary process state.**

A new function `rte_eal_primary_proc_alive()` has been added to allow the user to detect if a primary process is running. Use cases for this feature include fault detection, and monitoring using secondary processes.

- **Enabled bulk allocation of mbufs.**

A new function `rte_pktmbuf_alloc_bulk()` has been added to allow the user to bulk allocate mbufs.

- **Added device link speed capabilities.**

The structure `rte_eth_dev_info` now has a `speed_capa` bitmap, which allows the application to determine the supported speeds of each device.

- **Added bitmap of link speeds to advertise.**

Added a feature to allow the definition of a set of advertised speeds for auto-negotiation, explicitly disabling link auto-negotiation (single speed) and full auto-negotiation.

- **Added new poll-mode driver for Amazon Elastic Network Adapters (ENA).**

The driver operates for a variety of ENA adapters through feature negotiation with the adapter and upgradable commands set. The ENA driver handles PCI Physical and Virtual ENA functions.

- **Restored vmxnet3 TX data ring.**

TX data ring has been shown to improve small packet forwarding performance on the vSphere environment.

- **Added vmxnet3 TX L4 checksum offload.**

Added support for TCP/UDP checksum offload to vmxnet3.

- **Added vmxnet3 TSO support.**

Added support for TSO to vmxnet3.

- **Added vmxnet3 support for jumbo frames.**

Added support for linking multi-segment buffers together to handle Jumbo packets.



- **Enabled Virtio 1.0 support.**

Enabled Virtio 1.0 support for Virtio pmd driver.

- **Supported Virtio for ARM.**

Enabled Virtio support for ARMv7/v8. Tested for ARM64. Virtio for ARM supports VFIO-noiommu mode only. Virtio can work with other non-x86 architectures as well, like PowerPC.

- **Supported Virtio offload in vhost-user.**

Added the offload and negotiation of checksum and TSO between vhost-user and vanilla Linux Virtio guest.

- **Added vhost-user live migration support.**

- **Added vhost driver.**

Added a virtual PMD that wraps `librte_vhost`.

- **Added multicast promiscuous mode support on VF for ixgbe.**

Added multicast promiscuous mode support for the ixgbe VF driver so all VFs can receive the multicast packets.

Please note if you want to use this promiscuous mode, you need both PF and VF driver to support it. The reason is that this VF feature is configured in the PF. If you use kernel PF driver and the dpdk VF driver, make sure the kernel PF driver supports VF multicast promiscuous mode. If you use dpdk PF and dpdk VF ensure the PF driver is the same version as the VF.

- **Added support for E-tag on X550.**

E-tag is defined in [802.1BR - Bridge Port Extension](#).

This feature is for the VF, but the settings are on the PF. It means the CLIs should be used on the PF, but some of their effects will be shown on the VF. The forwarding of E-tag packets based on GRP and E-CID\_base will have an effect on the PF. Theoretically, the E-tag packets can be forwarded to any pool/queue but normally we'd like to forward the packets to the pools/queues belonging to the VFs. And E-tag insertion and stripping will have an effect on VFs. When a VF receives E-tag packets it should strip the E-tag. When the VF transmits packets, it should insert the E-tag. Both actions can be offloaded.

When we want to use this E-tag support feature, the forwarding should be enabled to forward the packets received by the PF to the indicated VFs. And insertion and stripping should be enabled for VFs to offload the effort to hardware.

Features added:

- Support E-tag offloading of insertion and stripping.
- Support Forwarding E-tag packets to pools based on GRP and E-CID\_base.

- **Added support for VxLAN and NVGRE checksum off-load on X550.**

- Added support for VxLAN and NVGRE RX/TX checksum off-load on X550. RX/TX checksum off-load is provided on both inner and outer IP header and TCP header.
- Added functions to support VxLAN port configuration. The default VxLAN port number is 4789 but this can be updated programmatically.

- **Added support for new X550EM\_a devices.**

Added support for new X550EM\_a devices and their MAC types, X550EM\_a and X550EM\_a\_vf. Updated the relevant PMD to use the new devices and MAC types.

- **Added x550em\_x V2 device support.**

Added support for x550em\_x V2 device. Only x550em\_x V1 was supported before. A mask for V1 and V2 is defined and used to support both.

- **Supported link speed auto-negotiation on X550EM\_X**

Normally the auto-negotiation is supported by firmware and software doesn't care about it. But on x550em\_x, firmware doesn't support auto-negotiation. As the ports of x550em\_x are 10GbE, if we connect the port with a peer which is 1GbE, the link will always be down. We added the support for auto-negotiation by software to avoid this link down issue.

- **Added software-firmware sync on X550EM\_a.**

Added support for software-firmware sync for resource sharing. Use the PHY token, shared between software-firmware for PHY access on X550EM\_a.

- **Updated the i40e base driver.**

The i40e base driver was updated with changes including the following:

- Use RX control AQ commands to read/write RX control registers.
- Add new X722 device IDs, and removed X710 one was never used.
- Expose registers for HASH/FD input set configuring.

- **Enabled PCI extended tag for i40e.**

Enabled extended tag for i40e by checking and writing corresponding PCI config space bytes, to boost the performance. The legacy method of reading/writing sysfile supported by kernel module igb\_uio is now deprecated.

- **Added i40e support for setting mac addresses.**

- **Added dump of i40e registers and EEPROM.**

- **Supported ether type setting of single and double VLAN for i40e**

- **Added VMDQ DCB mode in i40e.**

Added support for DCB in VMDQ mode to i40e driver.

- **Added i40e VEB switching support.**

- **Added Flow director enhancements in i40e.**

- **Added PF reset event reporting in i40e VF driver.**

- **Added fm10k RX interrupt support.**

- **Optimized fm10k TX.**

Optimized fm10k TX by freeing multiple mbufs at a time.

- **Handled error flags in fm10k vector RX.**

Parse error flags in RX descriptor and set error bits in mbuf with vector instructions.

- **Added fm10k FTAG based forwarding support.**
- **Added mlx5 flow director support.**

Added flow director support (`RTE_FDIR_MODE_PERFECT` and `RTE_FDIR_MODE_PERFECT_MAC_VLAN`).

Only available with Mellanox OFED  $\geq 3.2$ .
- **Added mlx5 RX VLAN stripping support.**

Added support for RX VLAN stripping.

Only available with Mellanox OFED  $\geq 3.2$ .
- **Added mlx5 link up/down callbacks.**

Implemented callbacks to bring link up and down.
- **Added mlx5 support for operation in secondary processes.**

Implemented TX support in secondary processes (like mlx4).
- **Added mlx5 RX CRC stripping configuration.**

Until now, CRC was always stripped. It can now be configured.

Only available with Mellanox OFED  $\geq 3.2$ .
- **Added mlx5 optional packet padding by HW.**

Added an option to make PCI bus transactions rounded to a multiple of a cache line size for better alignment.

Only available with Mellanox OFED  $\geq 3.2$ .
- **Added mlx5 TX VLAN insertion support.**

Added support for TX VLAN insertion.

Only available with Mellanox OFED  $\geq 3.2$ .
- **Changed szedata2 driver type from vdev to pdev.**

Previously szedata2 device had to be added by `--vdev` option. Now szedata2 PMD recognizes the device automatically during EAL initialization.
- **Added szedata2 functions for setting link up/down.**
- **Added szedata2 promiscuous and allmulticast modes.**
- **Added af\_packet dynamic removal function.**

An `af_packet` device can now be detached using the API, like other PMD devices.
- **Increased number of next hops for LPM IPv4 to  $2^{24}$ .**

The `next_hop` field has been extended from 8 bits to 24 bits for IPv4.
- **Added support of SNOW 3G (UEA2 and UIA2) for Intel Quick Assist devices.**

Enabled support for the SNOW 3G wireless algorithm for Intel Quick Assist devices. Support for cipher-only and hash-only is also provided along with algorithm-chaining operations.

- **Added SNOW3G SW PMD.**

A new Crypto PMD has been added, which provides SNOW 3G UEA2 ciphering and SNOW3G UIA2 hashing.

- **Added AES GCM PMD.**

Added new Crypto PMD to support AES-GCM authenticated encryption and authenticated decryption in software.

- **Added NULL Crypto PMD**

Added new Crypto PMD to support null crypto operations in software.

- **Improved IP Pipeline Application.**

The following features have been added to ip\_pipeline application;

- Added CPU utilization measurement and idle cycle rate computation.
- Added link identification support through existing port-mask option or by specifying PCI device in every LINK section in the configuration file.
- Added load balancing support in passthrough pipeline.

- **Added IPsec security gateway example.**

Added a new application implementing an IPsec Security Gateway.

## 8.2 Resolved Issues

### 8.2.1 Drivers

- **ethdev: Fixed overflow for 100Gbps.**

100Gbps in Mbps (100000) was exceeding the 16-bit max value of `link_speed` in `rte_eth_link`.

- **ethdev: Fixed byte order consistency between fdir flow and mask.**

Fixed issue in ethdev library where the structure for setting fdir's mask and flow entry was not consistent in byte ordering.

- **cxgbe: Fixed crash due to incorrect size allocated for RSS table.**

Fixed a segfault that occurs when accessing part of port 0's RSS table that gets overwritten by subsequent port 1's part of the RSS table due to incorrect size allocated for each entry in the table.

- **cxgbe: Fixed setting wrong device MTU.**

Fixed an incorrect device MTU being set due to the Ethernet header and CRC lengths being added twice.

- **ixgbe: Fixed zeroed VF mac address.**

Resolved an issue where the VF MAC address is zeroed out in cases where the VF driver is loaded while the PF interface is down. The solution is to only set it when we get an ACK from the PF.

- **ixgbe: Fixed setting flow director flag twice.**

Resolved an issue where packets were being dropped when switching to perfect filters mode.

- **ixgbe: Set MDIO speed after MAC reset.**

The MDIO clock speed must be reconfigured after the MAC reset. The MDIO clock speed becomes invalid, therefore the driver reads invalid PHY register values. The driver now set the MDIO clock speed prior to initializing PHY ops and again after the MAC reset.

- **ixgbe: Fixed maximum number of available TX queues.**

In IXGBE, the maximum number of TX queues varies depending on the NIC operating mode. This was not being updated in the device information, providing an incorrect number in some cases.

- **i40e: Generated MAC address for each VFs.**

It generates a MAC address for each VFs during PF host initialization, and keeps the VF MAC address the same among different VF launch.

- **i40e: Fixed failure of reading/writing RX control registers.**

Fixed i40e issue of failing to read/write rx control registers when under stress with traffic, which might result in application launch failure.

- **i40e: Enabled vector driver by default.**

Previously, vector driver was disabled by default as it couldn't fill packet type info for l3fwd to work well. Now there is an option for l3fwd to analyze the packet type so the vector driver is enabled by default.

- **i40e: Fixed link info of VF.**

Previously, the VF's link speed stayed at 10GbE and status always was up. It did not change even when the physical link's status changed. Now this issue is fixed to make VF's link info consistent with physical link.

- **mlx5: Fixed possible crash during initialization.**

A crash could occur when failing to allocate private device context.

- **mlx5: Added port type check.**

Added port type check to prevent port initialization on non-Ethernet link layers and to report an error.

- **mlx5: Applied VLAN filtering to broadcast and IPv6 multicast flows.**

Prevented reception of multicast frames outside of configured VLANs.

- **mlx5: Fixed RX checksum offload in non L3/L4 packets.**

Fixed report of bad checksum for packets of unknown type.

- **aesni\_mb: Fixed wrong return value when creating a device.**

The `cryptodev_aesni_mb_init()` function was returning the device id of the device created, instead of 0 (on success) that `rte_eal_vdev_init()` expects. This made it impossible to create more than one aesni\_mb device from the command line.

- **qat: Fixed AES GCM decryption.**

Allowed AES GCM on the cryptodev API, but in some cases gave invalid results due to incorrect IV setting.

## 8.2.2 Libraries

- **hash: Fixed CRC32c hash computation for non multiple of 4 bytes sizes.**

Fix crc32c hash functions to return a valid crc32c value for data lengths not a multiple of 4 bytes.

- **hash: Fixed hash library to support multi-process mode.**

Fix hash library to support multi-process mode, using a jump table, instead of storing a function pointer to the key compare function. Multi-process mode only works with the built-in compare functions, however a custom compare function (not in the jump table) can only be used in single-process mode.

- **hash: Fixed return value when allocating an existing hash table.**

Changed the `rte_hash*_create()` functions to return `NULL` and set `rte_errno` to `EEXIST` when the object name already exists. This is the behavior described in the API documentation in the header file. The previous behavior was to return a pointer to the existing object in that case, preventing the caller from knowing if the object had to be freed or not.

- **lpm: Fixed return value when allocating an existing object.**

Changed the `rte_lpm*_create()` functions to return `NULL` and set `rte_errno` to `EEXIST` when the object name already exists. This is the behavior described in the API documentation in the header file. The previous behavior was to return a pointer to the existing object in that case, preventing the caller from knowing if the object had to be freed or not.

- **librte\_port: Fixed segmentation fault for ring and ethdev writer nodrop.**

Fixed core dump issue on txq and swq when droptless is set to yes.

## 8.2.3 Examples

- **l3fwd-power: Fixed memory leak for non-IP packet.**

Fixed issue in l3fwd-power where, on receiving packets of types other than IPv4 or IPv6, the mbuf was not released, and caused a memory leak.

- **l3fwd: Fixed using packet type blindly.**

l3fwd makes use of packet type information without querying if devices or PMDs really set it. For those devices that don't set ptypes, add an option to parse it.

- **examples/vhost: Fixed frequent mbuf allocation failure.**

The vhost-switch often fails to allocate mbuf when dequeue from vring because it wrongly calculates the number of mbufs needed.

## 8.3 API Changes

- The ethdev statistics counter `imissed` is considered to be independent of `ierrors`. All drivers are now counting the missed packets only once, i.e. drivers will not increment `ierrors` anymore for missed packets.
- The ethdev structure `rte_eth_dev_info` was changed to support device speed capabilities.
- The ethdev structures `rte_eth_link` and `rte_eth_conf` were changed to support the new link API.
- The functions `rte_eth_dev_udp_tunnel_add` and `rte_eth_dev_udp_tunnel_delete` have been re-named into `rte_eth_dev_udp_tunnel_port_add` and `rte_eth_dev_udp_tunnel_port_delete`.
- The `outer_mac` and `inner_mac` fields in structure `rte_eth_tunnel_filter_conf` are changed from pointer to struct in order to keep code's readability.
- The fields in ethdev structure `rte_eth_fdir_masks` were changed to be in big endian.
- A parameter `vlan_type` has been added to the function `rte_eth_dev_set_vlan_ether_type`.
- The `af_packet` device init function is no longer public. The device should be attached via the API.
- The LPM `next_hop` field is extended from 8 bits to 24 bits for IPv4 while keeping ABI compatibility.
- A new `rte_lpm_config` structure is used so the LPM library will allocate exactly the amount of memory which is necessary to hold application's rules. The previous ABI is kept for compatibility.
- The prototype for the pipeline input port, output port and table action handlers are updated: the pipeline parameter is added, the packets mask parameter has been either removed or made input-only.

## 8.4 ABI Changes

- The RETA entry size in `rte_eth_rss_reta_entry64` has been increased from 8-bit to 16-bit.
- The ethdev flow director structure `rte_eth_fdir_flow` structure was changed. New fields were added to extend flow director's input set.
- The cmdline buffer size has been increase from 256 to 512.

## 8.5 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```
+ libethdev.so.3
  librte_acl.so.2
  librte_cfgfile.so.2
+ librte_cmdline.so.2
  librte_distributor.so.1
  librte_eal.so.2
  librte_hash.so.2
  librte_ip_frag.so.1
  librte_ivshmem.so.1
  librte_jobstats.so.1
  librte_kni.so.2
  librte_kvargs.so.1
  librte_lpm.so.2
  librte_mbuf.so.2
  librte_mempool.so.1
  librte_meter.so.1
+ librte_pipeline.so.3
  librte_pmd_bond.so.1
  librte_pmd_ring.so.2
  librte_port.so.2
  librte_power.so.1
  librte_reorder.so.1
  librte_ring.so.1
  librte_sched.so.1
  librte_table.so.2
  librte_timer.so.1
  librte_vhost.so.2
```

## 8.6 Tested Platforms

1. SuperMicro 1U
  - BIOS: 1.0c
  - Processor: Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
2. SuperMicro 1U
  - BIOS: 1.0a
  - Processor: Intel(R) Xeon(R) CPU D-1540 @ 2.00GHz
  - Onboard NIC: Intel(R) X552/X557-AT (2x10G)
    - Firmware-version: 0x800001cf
    - Device ID (PF/VF): 8086:15ad /8086:15a8
  - kernel driver version: 4.2.5 (ixgbe)
3. SuperMicro 1U
  - BIOS: 1.0a
  - Processor: Intel(R) Xeon(R) CPU E5-4667 v3 @ 2.00GHz
4. Intel(R) Server board S2600GZ
  - BIOS: SE5C600.86B.02.02.0002.122320131210
  - Processor: Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
5. Intel(R) Server board W2600CR



- BIOS: SE5C600.86B.02.01.0002.082220131453
  - Processor: Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz
6. Intel(R) Server board S2600CWT
    - BIOS: SE5C610.86B.01.01.0009.060120151350
    - Processor: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
  7. Intel(R) Server board S2600WTT
    - BIOS: SE5C610.86B.01.01.0005.101720141054
    - Processor: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
  8. Intel(R) Server board S2600WTT
    - BIOS: SE5C610.86B.11.01.0044.090120151156
    - Processor: Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz

## 8.7 Tested NICs

1. Intel(R) Ethernet Controller X540-AT2
  - Firmware version: 0x80000389
  - Device id (pf): 8086:1528
  - Driver version: 3.23.2 (ixgbe)
2. Intel(R) 82599ES 10 Gigabit Ethernet Controller
  - Firmware version: 0x61bf0001
  - Device id (pf/vf): 8086:10fb / 8086:10ed
  - Driver version: 4.0.1-k (ixgbe)
3. Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
  - Firmware version: 0x800001cf
  - Device id (pf/vf): 8086:15ad / 8086:15a8
  - Driver version: 4.2.5 (ixgbe)
4. Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
  - Firmware version: 5.02 0x80002284
  - Device id (pf/vf): 8086:1572 / 8086:154c
  - Driver version: 1.4.26 (i40e)
5. Intel(R) Ethernet Converged Network Adapter X710-DA2 (2x10G)
  - Firmware version: 5.02 0x80002282
  - Device id (pf/vf): 8086:1572 / 8086:154c
  - Driver version: 1.4.25 (i40e)
6. Intel(R) Ethernet Converged Network Adapter XL710-QDA1 (1x40G)

- Firmware version: 5.02 0x80002281
  - Device id (pf/vf): 8086:1584 / 8086:154c
  - Driver version: 1.4.25 (i40e)
7. Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
- Firmware version: 5.02 0x80002285
  - Device id (pf/vf): 8086:1583 / 8086:154c
  - Driver version: 1.4.25 (i40e)
8. Intel(R) 82576EB Gigabit Ethernet Controller
- Firmware version: 1.2.1
  - Device id (pf): 8086:1526
  - Driver version: 5.2.13-k (igb)
9. Intel(R) Ethernet Controller I210
- Firmware version: 3.16, 0x80000500, 1.304.0
  - Device id (pf): 8086:1533
  - Driver version: 5.2.13-k (igb)
10. Intel(R) Corporation I350 Gigabit Network Connection
- Firmware version: 1.48, 0x800006e7
  - Device id (pf/vf): 8086:1521 / 8086:1520
  - Driver version: 5.2.13-k (igb)
11. Intel(R) Ethernet Multi-host Controller FM10000
- Firmware version: N/A
  - Device id (pf/vf): 8086:15d0
  - Driver version: 0.17.0.9 (fm10k)

## 9.1 New Features

- **Introduce ARMv7 and ARMv8 architectures.**

- It is now possible to build DPDK for the ARMv7 and ARMv8 platforms.
- ARMv7 can be tested with virtual PMD drivers.
- ARMv8 can be tested with virtual and physical PMD drivers.

- **Enabled freeing of ring.**

A new function `rte_ring_free()` has been added to allow the user to free a ring if it was created with `rte_ring_create()`.

- **Added keepalive support to EAL and example application.**

- **Added experimental cryptodev API**

The cryptographic processing of packets is provided as a preview with two drivers for:

- Intel QuickAssist devices
- Intel AES-NI multi-buffer library

Due to its experimental state, the API may change without prior notice.

- **Added ethdev APIs for additional IEEE1588 support.**

Added functions to read, write and adjust system time in the NIC. Added client slave sample application to demonstrate the IEEE1588 functionality.

- **Extended Statistics.**

Defined an extended statistics naming scheme to store metadata in the name string of each statistic. Refer to the Extended Statistics section of the Programmers Guide for more details.

Implemented the extended statistics API for the following PMDs:

- `igb`
- `igbvf`
- `i40e`
- `i40evf`
- `fm10k`

- virtio

- **Added API in ethdev to retrieve RX/TX queue information.**

- Added the ability for the upper layer to query RX/TX queue information.
- Added new fields in `rte_eth_dev_info` to represent information about RX/TX descriptors min/max/align numbers, per queue, for the device.

- **Added RSS dynamic configuration to bonding.**

- **Updated the e1000 base driver.**

The e1000 base driver was updated with several features including the following:

- Added new i218 devices
- Allowed both ULP and EEE in Sx state
- Initialized 88E1543 (Marvell 1543) PHY
- Added flags to set EEE advertisement modes
- Supported inverted format ETrackId
- Added bit to disable packetbuffer read
- Added defaults for i210 RX/TX PBSIZE
- Check more errors for ESB2 init and reset
- Check more NVM read errors
- Return code after setting receive address register
- Removed all NAHUM6LP\_HW tags

- **Added e1000 RX interrupt support.**

- **Added igb TSO support for both PF and VF.**

- **Added RSS enhancements to Intel x550 NIC.**

- Added support for 512 entry RSS redirection table.
- Added support for per VF RSS redirection table.

- **Added Flow director enhancements on Intel x550 NIC.**

- Added 2 new flow director modes on x550. One is MAC VLAN mode, the other is tunnel mode.

- **Updated the i40e base driver.**

The i40e base driver was updated with several changes including the following:

- Added promiscuous on VLAN support
- Added a workaround to drop all flow control frames
- Added VF capabilities to virtual channel interface
- Added TX Scheduling related AQ commands
- Added additional PCTYPES supported for FortPark RSS
- Added parsing for CEE DCBX TLVs

- Added FortPark specific registers
- Added AQ functions to handle RSS Key and LUT programming
- Increased PF reset max loop limit
- **Added i40e vector RX/TX.**
- **Added i40e RX interrupt support.**
- **Added i40e flow control support.**
- **Added DCB support to i40e PF driver.**
- **Added RSS/FD input set granularity on Intel X710/XL710.**
- **Added different GRE key length for input set on Intel X710/XL710.**
- **Added flow director support in i40e VF.**
- **Added i40e support of early X722 series.**

Added early X722 support, for evaluation only, as the hardware is alpha.
- **Added fm10k vector RX/TX.**
- **Added fm10k TSO support for both PF and VF.**
- **Added fm10k VMDQ support.**
- **New NIC Boulder Rapid support.**

Added support for the Boulder Rapid variant of Intel's fm10k NIC family.
- **Enhanced support for the Chelsio CXGBE driver.**
  - Added support for Jumbo Frames.
  - Optimized forwarding performance for Chelsio T5 40GbE cards.
- **Improved enic TX packet rate.**

Reduced frequency of TX tail pointer updates to the NIC.
- **Added support for link status interrupts in mlx4.**
- **Added partial support (TX only) for secondary processes in mlx4.**
- **Added support for Mellanox ConnectX-4 adapters (mlx5).**

The mlx5 poll-mode driver implements support for Mellanox ConnectX-4 EN and Mellanox ConnectX-4 Lx EN families of 10/25/40/50/100 Gb/s adapters.

Like mlx4, this PMD is only available for Linux and is disabled by default due to external dependencies (libibverbs and libmlx5).
- **Added driver for Netronome nfp-6xxx card.**

Support for using Netronome nfp-6xxx with PCI VFs.
- **Added virtual szedata2 driver for COMBO cards.**

Added virtual PMD for COMBO-100G and COMBO-80G cards. PMD is disabled in default configuration.
- **Enhanced support for virtio driver.**

- Virtio ring layout optimization (fixed avail ring)
- Vector RX
- Simple TX
- **Added vhost-user multiple queue support.**
- **Added port hotplug support to vmxnet3.**
- **Added port hotplug support to xenvirt.**
- **Added ethtool shim and sample application.**
- **Added experimental performance thread example application.**

The new sample application demonstrates L3 forwarding with different threading models: pthreads, cgroups, or lightweight threads. The example includes a simple cooperative scheduler.

Due to its experimental state this application may change without notice. The application is supported only for Linux x86\_64.

- **Enhancements to the IP pipeline application.**

The following features have been added to the `ip_pipeline` application;

- Added Multiple Producers/Multiple Consumers (MPSC) and fragmentation/reassembly support to software rings.
- Added a dynamic pipeline reconfiguration feature that allows binding a pipeline to other threads at runtime using CLI commands.
- Added enable/disable of `promisc` mode from `ip_pipeline` configuration file.
- Added check on RX queues and TX queues of each link whether they are used correctly in the `ip_pipeline` configuration file.
- Added flow id parameters to the flow-classification table entries.
- Added more functions to the routing pipeline: ARP table enable/disable, Q-in-Q and MPLS encapsulation, add color (traffic-class for QoS) to the MPLS tag.
- Added flow-actions pipeline for traffic metering/marketing (for e.g. Two Rate Three Color Marker (trTCM)), policer etc.
- Modified the pass-through pipeline's actions-handler to implement a generic approach to extract fields from the packet's header and copy them to packet metadata.

## 9.2 Resolved Issues

### 9.2.1 EAL

- **eal/linux: Fixed epoll timeout.**

Fixed issue where the `rte_epoll_wait()` function didn't return when the underlying call to `epoll_wait()` timed out.

## 9.2.2 Drivers

- **e1000/base: Synchronize PHY interface on non-ME systems.**

On power up, the MAC - PHY interface needs to be set to PCIe, even if the cable is disconnected. In ME systems, the ME handles this on exit from the Sx (Sticky mode) state. In non-ME, the driver handles it. Added a check for non-ME system to the driver code that handles it.

- **e1000/base: Increased timeout of reset check.**

Previously, in `check_reset_block` RSPCIPHY was polled for 100 ms before determining that the ME veto was set. This was not enough and it was increased to 300 ms.

- **e1000/base: Disabled IPv6 extension header parsing on 82575.**

Disabled IPv6 options as per hardware limitation.

- **e1000/base: Prevent ULP flow if cable connected.**

Enabling ULP on link down when the cable is connected caused an infinite loop of link up/down indications in the NDIS driver. The driver now enables ULP only when the cable is disconnected.

- **e1000/base: Support different EEARBC for i210.**

EEARBC has changed on i210. It means EEARBC has a different address on i210 than on other NICs. So, add a new entity named `EEARBC_I210` to the register list and make sure the right one is being used on i210.

- **e1000/base: Fix K1 configuration.**

Added fix for the following updates to the K1 configurations: TX idle period for entering K1 should be 128 ns. Minimum TX idle period in K1 should be 256 ns.

- **e1000/base: Fix link detect flow.**

Fix link detect flow in case where auto-negotiate is not enabled, by calling `e1000_setup_copper_link_generic` instead of `e1000_phy_setup_autoneg`.

- **e1000/base: Fix link check for i354 M88E1112 PHY.**

The `e1000_check_for_link_media_swap()` function is supposed to check PHY page 0 for copper and PHY page 1 for "other" (fiber) links. The driver switched back from page 1 to page 0 too soon, before `e1000_check_for_link_82575()` is executed and was never finding the link on the fiber (other).

If the link is copper, as the M88E1112 page address is set to 1, it should be set back to 0 before checking this link.

- **e1000/base: Fix beacon duration for i217.**

Fix for I217 Packet Loss issue - The Management Engine sets the FEXTNVM4 Beacon Duration incorrectly. This fix ensures that the correct value will always be set. Correct value for this field is 8 usec.

- **e1000/base: Fix TIPG for non 10 half duplex mode.**

TIPG value is increased when setting speed to 10 half duplex to prevent packet loss. However, it was never decreased again when speed changed. This caused performance issues in the NDIS driver. Fix this to restore TIPG to default value on non 10 half duplex.

- **e1000/base: Fix reset of DH89XXCC SGMII.**

For DH89XXCC\_SGMII, a write flush leaves registers of this device trashed (0xFFFFFFFF). Add check for this device.

Also, after both Port SW Reset and Device Reset case, the platform should wait at least 3ms before reading any registers. Remove this condition since waiting is conditionally executed only for Device Reset.

- **e1000/base: Fix redundant PHY power down for i210.**

Bit 11 of PHYREG 0 is used to power down PHY. The use of PHYREG 16 is no longer necessary.

- **e1000/base: fix jumbo frame CRC failures.**

Change the value of register 776.20[11:2] for jumbo mode from 0x1A to 0x1F. This is to enlarge the gap between read and write pointers in the TX FIFO.

- **e1000/base: Fix link flap on 82579.**

Several customers have reported a link flap issue on 82579. The symptoms are random and intermittent link losses when 82579 is connected to specific switches. The issue was root caused as an inter-operability problem between the NIC and at least some Broadcom PHYs in the Energy Efficient Ethernet wake mechanism.

To fix the issue, we are disabling the Phase Locked Loop shutdown in 100M Low Power Idle. This solution will cause an increase of power in 100M EEE link. It may cost an additional 28mW in this specific mode.

- **igb: Fixed IEEE1588 frame identification in I210.**

Fixed issue where the flag `PKT_RX_IEEE1588_PTP` was not being set in the Intel I210 NIC, as the EtherType in RX descriptor is in bits 8:10 of Packet Type and not in the default bits 0:2.

- **igb: Fixed VF start with PF stopped.**

VF needs the PF interrupt support initialized even if not started.

- **igb: Fixed VF MAC address when using with DPDK PF.**

Assign a random MAC address in VF when not assigned by PF.

- **igb: Removed CRC bytes from byte counter statistics.**

- **ixgbe: Fixed issue with X550 DCB.**

Fixed a DCB issue with x550 where for 8 TCs (Traffic Classes), if a packet with user priority 6 or 7 was injected to the NIC, then the NIC would only put 3 packets into the queue. There was also a similar issue for 4 TCs.

- **ixgbe: Removed burst size restriction of vector RX.**

Fixed issue where a burst size less than 32 didn't receive anything.

- **ixgbe: Fixed VF start with PF stopped.**

VF needs the PF interrupt support initialized even if not started.

- **ixgbe: Fixed TX hang when RS distance exceeds HW limit.**



Fixed an issue where the TX queue can hang when a lot of highly fragmented packets have to be sent. As part of that fix, `tx_rs_thresh` for ixgbe PMD is not allowed to be greater than 32 to comply with HW restrictions.

- **ixgbe: Fixed rx error statistic counter.**

Fixed an issue that the rx error counter of ixgbe was not accurate. The mac short packet discard count (`mispdc`) was added to the counter. Mac local faults and mac remote faults are removed as they do not count packets but errors, and jabber errors were removed as they are already accounted for by the CRC error counter. Finally the XEC (I3 / I4 checksum error) counter was removed due to errata, see commit 256ff05a9cae for details.

- **ixgbe: Removed CRC bytes from byte counter statistics.**

- **i40e: Fixed base driver allocation when not using first numa node.**

Fixed i40e issue that occurred when a DPDK application didn't initialize ports if memory wasn't available on socket 0.

- **i40e: Fixed maximum of 64 queues per port.**

Fixed an issue in i40e where it would not support more than 64 queues per port, even though the hardware actually supports it. The real number of queues may vary, as long as the total number of queues used in PF, VFs, VMDq and FD does not exceeds the hardware maximum.

- **i40e: Fixed statistics of packets.**

Added discarding packets on VSI to the stats and rectify the old statistics.

- **i40e: Fixed issue of not freeing memzone.**

Fixed an issue of not freeing a memzone in the call to free the memory for adminq DMA.

- **i40e: Removed CRC bytes from byte counter statistics.**

- **mlx: Fixed driver loading.**

The mlx drivers were unable to load when built as a shared library, due to a missing symbol in the mempool library.

- **mlx4: Performance improvements.**

Fixed bugs in TX and RX flows that improves mlx4 performance.

- **mlx4: Fixed TX loss after initialization.**

- **mlx4: Fixed scattered TX with too many segments.**

- **mlx4: Fixed memory registration for indirect mbuf data.**

- **vhost: Fixed Qemu shutdown.**

Fixed issue with libvirt `virsh destroy` not killing the VM.

- **virtio: Fixed crash after changing link state.**

Fixed IO permission in the interrupt handler.

- **virtio: Fixed crash when releasing queue.**

Fixed issue when releasing null control queue.

### 9.2.3 Libraries

- **hash: Fixed memory allocation of Cuckoo Hash key table.**

Fixed issue where an incorrect Cuckoo Hash key table size could be calculated limiting the size to 4GB.

- **hash: Fixed incorrect lookup if key is all zero.**

Fixed issue in hash library that occurred if an all zero key was not added to the table and the key was looked up, resulting in an incorrect hit.

- **hash: Fixed thread scaling by reducing contention.**

Fixed issue in the hash library where, using multiple cores with hardware transactional memory support, thread scaling did not work, due to the global ring that is shared by all cores.

### 9.2.4 Examples

- **I3fwd: Fixed crash with IPv6.**
- **vhost\_xen: Fixed compile error.**

### 9.2.5 Other

- This release drops compatibility with Linux kernel 2.6.33. The minimum kernel requirement is now 2.6.34.

## 9.3 Known Issues

- Some drivers do not fill in the packet type when receiving. As the I3fwd example application requires this info, the i40e vector driver must be disabled to benefit of the packet type with i40e.
- Some (possibly all) VF drivers (e.g. i40evf) do not handle any PF reset events/requests in the VF driver. This means that the VF driver may not work after a PF reset in the host side. The workaround is to avoid triggering any PF reset events/requests on the host side.
- 100G link report support is missing.
- **Mellanox PMDs (mlx4 & mlx5):**
  - PMDs do not support CONFIG\_RTE\_BUILD\_COMBINE\_LIBS and CONFIG\_RTE\_BUILD\_SHARED\_LIB simultaneously.
  - There is performance degradation for small packets when the PMD is compiled with SGE\_WR\_N = 4 compared to the performance when SGE\_WR\_N = 1. If scattered packets are not used it is recommended to compile the PMD with SGE\_WR\_N = 1.
  - When a Multicast or Broadcast packet is sent to the SR-IOV mlx4 VF, it is returned back to the port.
  - PMDs report “bad” L4 checksum when IP packet is received.

- mlx5 PMD reports “bad” checksum although the packet has “good” checksum. Will be fixed in upcoming MLNX\_OFED release.

## 9.4 API Changes

- The deprecated flow director API is removed. It was replaced by `rte_eth_dev_filter_ctrl()`.
- The `dcb_queue` is renamed to `dcb_tc` in following dcb configuration structures: `rte_eth_dcb_rx_conf`, `rte_eth_dcb_tx_conf`, `rte_eth_vmdq_dcb_conf`, `rte_eth_vmdq_dcb_tx_conf`.
- The `rte_eth_rx_queue_count()` function now returns “int” instead of “uint32\_t” to allow the use of negative values as error codes on return.
- The function `rte_eal_pci_close_one()` is removed. It was replaced by `rte_eal_pci_detach()`.
- The deprecated ACL API `ipv4vlan` is removed.
- The deprecated hash function `rte_jhash2()` is removed. It was replaced by `rte_jhash_32b()`.
- The deprecated KNI functions are removed: `rte_kni_create()`, `rte_kni_get_port_id()` and `rte_kni_info_get()`.
- The deprecated ring PMD functions are removed: `rte_eth_ring_pair_create()` and `rte_eth_ring_pair_attach()`.
- The devargs union field `virtual` is renamed to `virt` for C++ compatibility.

## 9.5 ABI Changes

- The EAL and ethdev structures `rte_intr_handle` and `rte_eth_conf` were changed to support RX interrupt. This was already included in 2.1 under the `CONFIG_RTE_NEXT_ABI` #define.
- The ethdev flow director entries for SCTP were changed. This was already included in 2.1 under the `CONFIG_RTE_NEXT_ABI` #define.
- The ethdev flow director structure `rte_eth_fdir_flow_ext` structure was changed. New fields were added to support flow director filtering in VF.
- The size of the ethdev structure `rte_eth_hash_filter_info` is changed by adding a new element `rte_eth_input_set_conf` in a union.
- New fields `rx_desc_lim` and `tx_desc_lim` are added into `rte_eth_dev_info` structure.
- For debug builds, the functions `rte_eth_rx_burst()`, `rte_eth_tx_burst()`, `rte_eth_rx_descriptor_done()` and `rte_eth_rx_queue_count()` will no longer be separate functions in the DPDK libraries. Instead, they will only be present in the `rte_ethdev.h` header file.
- The maximum number of queues per port `CONFIG_RTE_MAX_QUEUES_PER_PORT` is increased to 1024.

- The mbuf structure was changed to support the unified packet type. This was already included in 2.1 under the `CONFIG_RTE_NEXT_ABI` #define.
- The dummy malloc library is removed. The content was moved into EAL in 2.1.
- The LPM structure is changed. The deprecated field `mem_location` is removed.
- `librte_table` LPM: A new parameter to hold the table name will be added to the LPM table parameter structure.
- `librte_table` hash: The key mask parameter is added to the hash table parameter structure for 8-byte key and 16-byte key extendable bucket and LRU tables.
- `librte_port`: Macros to access the packet meta-data stored within the packet buffer has been adjusted to cover the packet mbuf structure.
- `librte_cfgfile`: Allow longer names and values by increasing the constants `CFG_NAME_LEN` and `CFG_VALUE_LEN` to 64 and 256 respectively.
- `vhost`: a new field `enabled` is added to the `vhost_virtqueue` structure.
- `vhost`: a new field `virt_qp_nb` is added to `virtio_net` structure, and the `virtqueue` field is moved to the end of `virtio_net` structure.
- `vhost`: a new operation `vring_state_changed` is added to `virtio_net_device_ops` structure.
- `vhost`: a few spaces are reserved both at `vhost_virtqueue` and `virtio_net` structure for future extension.

## 9.6 Shared Library Versions

The libraries prepended with a plus sign were incremented in this version.

```

+ libethdev.so.2
+ librte_acl.so.2
+ librte_cfgfile.so.2
  librte_cmdline.so.1
  librte_distributor.so.1
+ librte_eal.so.2
+ librte_hash.so.2
  librte_ip_frag.so.1
  librte_ivshmem.so.1
  librte_jobstats.so.1
+ librte_kni.so.2
  librte_kvargs.so.1
+ librte_lpm.so.2
+ librte_mbuf.so.2
  librte_mempool.so.1
  librte_meter.so.1
+ librte_pipeline.so.2
  librte_pmd_bond.so.1
+ librte_pmd_ring.so.2
+ librte_port.so.2
  librte_power.so.1
  librte_reorder.so.1
  librte_ring.so.1
  librte_sched.so.1
+ librte_table.so.2
  librte_timer.so.1
+ librte_vhost.so.2

```

## DPDK RELEASE 2.1

## 10.1 New Features

- **Enabled cloning of indirect mbufs.**

This feature removes a limitation of `rte_pktmbuf_attach()` which generated the warning: “mbuf we’re attaching to must be direct”.

Now, when attaching to an indirect mbuf it is possible to:

- Copy all relevant fields (address, length, offload, ...) as before.
- Get the pointer to the mbuf that embeds the data buffer (direct mbuf), and increase the reference counter.

When detaching the mbuf, we can now retrieve this direct mbuf as the pointer is determined from the buffer address.

- **Extended packet type support.**

In previous releases mbuf packet types were indicated by 6 bits in the `ol_flags`. This was not enough for some supported NICs. For example i40e hardware can recognize more than 150 packet types. Not being able to identify these additional packet types limits access to hardware offload capabilities

So an extended “unified” packet type was added to support all possible PMDs. The 16 bit `packet_type` in the mbuf structure was changed to 32 bits and used for this purpose.

To avoid breaking ABI compatibility, the code changes for this feature are enclosed in a `RTE_NEXT_ABI` ifdef. This is enabled by default but can be turned off for ABI compatibility with DPDK R2.0.

- **Reworked memzone to be allocated by malloc and also support freeing.**

In the memory hierarchy, memsegs are groups of physically contiguous hugepages, memzones are slices of memsegs, and malloc slices memzones into smaller memory chunks.

This feature modifies `malloc()` so it partitions memsegs instead of memzones. Now memzones allocate their memory from the malloc heap.

Backward compatibility with API and ABI are maintained.

This allow memzones, and any other structure based on memzones, for example mem-pools, to be freed. Currently only the API from freeing memzones is supported.

- **Interrupt mode PMD.**

This feature introduces a low-latency one-shot RX interrupt into DPDK. It also adds a polling and interrupt mode switch control example.

DPDK userspace interrupt notification and handling mechanism is based on UIO/VFIO with the following limitations:

- Per queue RX interrupt events are only allowed in VFIO which supports multiple MSI-X vectors.
- In UIO, the RX interrupt shares the same vector with other interrupts. When the RX interrupt and LSC interrupt are both enabled, only the former is available.
- RX interrupt is only implemented for the linuxapp target.
- The feature is only currently enabled for tow PMDs: ixgbe and igb.

- **Packet Framework enhancements.**

Several enhancements were made to the Packet Framework:

- A new configuration file syntax has been introduced for IP pipeline applications. Parsing of the configuration file is changed.
- Implementation of the IP pipeline application is modified to make it more structured and user friendly.
- Implementation of the command line interface (CLI) for each pipeline type has been moved to the separate compilation unit. Syntax of pipeline CLI commands has been changed.
- Initialization of IP pipeline is modified to match the new parameters structure.
- New implementation of pass-through pipeline, firewall pipeline, routing pipeline, and flow classification has been added.
- Master pipeline with CLI interface has been added.
- Added extended documentation of the IP Pipeline.

- **Added API for IEEE1588 timestamping.**

This feature adds an ethdev API to enable, disable and read IEEE1588/802.1AS PTP timestamps from devices that support it. The following functions were added:

- `rte_eth_timesync_enable()`
- `rte_eth_timesync_disable()`
- `rte_eth_timesync_read_rx_timestamp()`
- `rte_eth_timesync_read_tx_timestamp()`

The “ieee1588” forwarding mode in testpmd was also refactored to demonstrate the new API.

- **Added multicast address filtering.**

Added multicast address filtering via a new ethdev function `set_mc_addr_list()`.

This overcomes a limitation in previous releases where the receipt of multicast packets on a given port could only be enabled by invoking the `rte_eth_allmulticast_enable()` function. This method did not work for VFs

in SR-IOV architectures when the host PF driver does not allow these operation on VFs. In such cases, joined multicast addresses had to be added individually to the set of multicast addresses that are filtered by the [VF] port.

- **Added Flow Director extensions.**

Several Flow Director extensions were added such as:

- Support for RSS and Flow Director hashes in vector RX.
- Added Flow Director for L2 payload.

- **Added RSS hash key size query per port.**

This feature supports querying the RSS hash key size of each port. A new field `hash_key_size` has been added in the `rte_eth_dev_info` struct for storing hash key size in bytes.

- **Added userspace ethtool support.**

Added userspace ethtool support to provide a familiar interface for applications that manage devices via kernel-space `ethtool_op` and `net_device_op`.

The initial implementation focuses on operations that can be implemented through existing `netdev` APIs. More operations will be supported in later releases.

- **Updated the ixgbe base driver.**

The ixgbe base driver was updated with several changes including the following:

- Added a new 82599 device id.
- Added new X550 PHY ids.
- Added SFP+ dual-speed support.
- Added wait helper for X550 IOSF accesses.
- Added X550em features.
- Added X557 PHY LEDs support.
- Commands for flow director.
- Issue firmware command when resetting X550em.

See the git log for full details of the ixgbe/base changes.

- **Added additional hotplug support.**

Port hotplug support was added to the following PMDs:

- e1000/igb.
- ixgbe.
- i40e.
- fm10k.
- ring.
- bonding.
- virtio.

Port hotplug support was added to BSD.

- **Added ixgbe LRO support.**

Added LRO support for x540 and 82599 devices.

- **Added extended statistics for ixgbe.**

Implemented `xstats_get()` and `xstats_reset()` in `dev_ops` for ixgbe to expose detailed error statistics to DPDK applications.

These will be implemented for other PMDs in later releases.

- **Added proc\_info application.**

Created a new `proc_info` application, by refactoring the existing `dump_cfg` application, to demonstrate the usage of retrieving statistics, and the new extended statistics (see above), for DPDK interfaces.

- **Updated the i40e base driver.**

The i40e base driver was updated with several changes including the following:

- Support for building both PF and VF driver together.
- Support for CEE DCBX on recent firmware versions.
- Replacement of `i40e_debug_read_register()`.
- Rework of `i40e_hmc_get_object_va`.
- Update of shadow RAM read/write functions.
- Enhancement of polling NVM semaphore.
- Enhancements on adminq init and sending asq command.
- Update of get/set LED functions.
- Addition of AOC phy types to case statement in `get_media_type`.
- Support for iSCSI capability.
- Setting of `FLAG_RD` when sending driver version to FW.

See the git log for full details of the i40e/base changes.

- **Added support for port mirroring in i40e.**

Enabled mirror functionality in the i40e driver.

- **Added support for i40e double VLAN, QinQ, stripping and insertion.**

Added support to the i40e driver for offloading double VLAN (QinQ) tags to the mbuf header, and inserting double vlan tags by hardware to the packets to be transmitted. Added a new field `vlan_tci_outer` in the `rte_mbuf` struct, and new flags in `ol_flags` to support this feature.

- **Added fm10k promiscuous mode support.**

Added support for promiscuous/allmulticast enable and disable in the fm10k PF function. VF is not supported yet.



- **Added fm10k jumbo frame support.**

Added support for jumbo frame less than 15K in both VF and PF functions in the fm10k pmd.

- **Added fm10k mac vlan filtering support.**

Added support for the fm10k MAC filter, only available in PF. Updated the VLAN filter to add/delete one static entry in the MAC table for each combination of VLAN and MAC address.

- **Added support for the Broadcom bnx2x driver.**

Added support for the Broadcom NetXtreme II bnx2x driver. It is supported only on Linux 64-bit and disabled by default.

- **Added support for the Chelsio CXGBE driver.**

Added support for the CXGBE Poll Mode Driver for the Chelsio Terminator 5 series of 10G/40G adapters.

- **Enhanced support for Mellanox ConnectX-3 driver (mlx4).**

- Support Mellanox OFED 3.0.
- Improved performance for both RX and TX operations.
- Better link status information.
- Outer L3/L4 checksum offload support.
- Inner L3/L4 checksum offload support for VXLAN.

- **Enabled VMXNET3 vlan filtering.**

Added support for the VLAN filter functionality of the VMXNET3 interface.

- **Added support for vhost live migration.**

Added support to allow live migration of vhost. Without this feature, qemu will report the following error: “migrate: Migration disabled: vhost lacks VHOST\_F\_LOG\_ALL feature”.

- **Added support for pcap jumbo frames.**

Extended the PCAP PMD to support jumbo frames for RX and TX.

- **Added support for the TILE-Gx architecture.**

Added support for the EZchip TILE-Gx family of SoCs.

- **Added hardware memory transactions/lock elision for x86.**

Added the use of hardware memory transactions (HTM) on fast-path for rwlock and spinlock (a.k.a. lock elision). The methods are implemented for x86 using Restricted Transactional Memory instructions (Intel(r) Transactional Synchronization Extensions). The implementation fall-backs to the normal rwlock if HTM is not available or memory transactions fail. This is not a replacement for all rwlock usages since not all critical sections protected by locks are friendly to HTM. For example, an attempt to perform a HW I/O operation inside a hardware memory transaction always aborts the transaction since the CPU is not able to roll-back should the transaction fail. Therefore, hardware transactional locks are not advised to be used around `rte_eth_rx_burst()` and `rte_eth_tx_burst()` calls.

- **Updated Jenkins Hash function**

Updated the version of the Jenkins Hash (jhash) function used in DPDK from the 1996 version to the 2006 version. This gives up to 35% better performance, compared to the original one.

Note, the hashes generated by the updated version differ from the hashes generated by the previous version.

- **Added software implementation of the Toeplitz RSS hash**

Added a software implementation of the Toeplitz hash function used by RSS. It can be used either for packet distribution on a single queue NIC or for simulating RSS computation on a specific NIC (for example after GRE header de-encapsulation).

- **Replaced the existing hash library with a Cuckoo hash implementation.**

Replaced the existing hash library with another approach, using the Cuckoo Hash method to resolve collisions (open addressing). This method pushes items from a full bucket when a new entry must be added to it, storing the evicted entry in an alternative location, using a secondary hash function.

This gives the user the ability to store more entries when a bucket is full, in comparison with the previous implementation.

The API has not been changed, although new fields have been added in the `rte_hash` structure, which has been changed to internal use only.

The main change when creating a new table is that the number of entries per bucket is now fixed, so its parameter is ignored now (it is still there to maintain the same parameters structure).

Also, the maximum burst size in `lookup_burst` function hash been increased to 64, to improve performance.

- **Optimized KNI RX burst size computation.**

Optimized KNI RX burst size computation by avoiding checking how many entries are in `kni->rx_q` prior to actually pulling them from the fifo.

- **Added KNI multicast.**

Enabled adding multicast addresses to KNI interfaces by adding an empty callback for `set_rx_mode` (typically used for setting up hardware) so that the `ioctl` succeeds. This is the same thing as the Linux tap interface does.

- **Added cmdline polling mode.**

Added the ability to process console input in the same thread as packet processing by using the `poll()` function.

- **Added VXLAN Tunnel End point sample application.**

Added a Tunnel End point (TEP) sample application that simulates a VXLAN Tunnel Endpoint (VTEP) termination in DPDK. It is used to demonstrate the offload and filtering capabilities of Intel XL710 10/40 GbE NICs for VXLAN packets.

- **Enabled combining of the “-m” and “-no-huge” EAL options.**

Added option to allow combining of the `-m` and `--no-huge` EAL command line options.

This allows user application to run as non-root but with higher memory allocations, and removes a constraint on `--no-huge` mode being limited to 64M.

## 10.2 Resolved Issues

- **acl: Fix ambiguity between test rules.**

Some test rules had equal priority for the same category. That could cause an ambiguity in building the trie and test results.

- **acl: Fix invalid rule wildness calculation for bitmask field type.**

- **acl: Fix matching rule.**

- **acl: Fix unneeded trie splitting for subset of rules.**

When rebuilding a trie for limited rule-set, don't try to split the rule-set even further.

- **app/testpmd: Fix crash when port id out of bound.**

Fixed issues in testpmd where using a port greater than 32 would cause a seg fault.

Fixes: edab33b1c01d ("app/testpmd: support port hotplug")

- **app/testpmd: Fix reply to a multicast ICMP request.**

Set the IP source and destination addresses in the IP header of the ICMP reply.

- **app/testpmd: fix MAC address in ARP reply.**

Fixed issue where in the `icmpecho` forwarding mode, ARP replies from testpmd contain invalid zero-filled MAC addresses.

Fixes: 31db4d38de72 ("net: change arp header struct declaration")

- **app/testpmd: fix default flow control values.**

Fixes: 422a20a4e62d ("app/testpmd: fix uninitialized flow control variables")

- **bonding: Fix crash when stopping inactive slave.**

- **bonding: Fix device initialization error handling.**

- **bonding: Fix initial link status of slave.**

On Fortville NIC, link status change interrupt callback was not executed when slave in bonding was (re-)started.

- **bonding: Fix socket id for LACP slave.**

Fixes: 46fb43683679 ("bond: add mode 4")

- **bonding: Fix device initialization error handling.**

- **cmdline: Fix small memory leak.**

A function in `cmdline.c` had a return that did not free the buf properly.

- **config: Enable same drivers options for Linux and BSD.**

Enabled vector `ixgbe` and `i40e` bulk alloc for BSD as it is already done for Linux.

Fixes: 304caba12643 (“config: fix bsd options”) Fixes: 0ff3324da2eb (“ixgbe: rework vector pmd following mbuf changes”)

- **devargs: Fix crash on failure.**

This problem occurred when passing an invalid PCI id to the blacklist API in devargs.

- **e1000/i40e: Fix descriptor done flag with odd address.**

- **e1000/igb: fix ieee1588 timestamping initialization.**

Fixed issue with e1000 ieee1588 timestamp initialization. On initialization the IEEE1588 functions read the system time to set their timestamp. However, on some 1G NICs, for example, i350, system time is disabled by default and the IEEE1588 timestamp was always 0.

- **eal/bsd: Fix inappropriate header guards.**

- **eal/bsd: Fix virtio on FreeBSD.**

Closing the `/dev/io` fd caused a SIGBUS in inb/outb instructions as the process lost the IOPL privileges once the fd is closed.

Fixes: 8a312224bcde (“eal/bsd: fix fd leak”)

- **eal/linux: Fix comments on vfio MSI.**

- **eal/linux: Fix irq handling with igb\_uio.**

Fixed an issue where the the introduction of `uio_pci_generic` broke interrupt handling with `igb_uio`.

Fixes: c112df6875a5 (“eal/linux: toggle interrupt for uio\_pci\_generic”)

- **eal/linux: Fix numa node detection.**

- **eal/linux: Fix socket value for undetermined numa node.**

Sets zero as the default value of pci device `numa_node` if the socket could not be determined. This provides the same default value as FreeBSD which has no NUMA support, and makes the return value of `rte_eth_dev_socket_id()` be consistent with the API description.

- **eal/ppc: Fix cpu cycle count for little endian.**

On IBM POWER8 PPC64 little endian architecture, the definition of `tsc` union will be different. This fix enables the right output from `rte_rdtsc()`.

- **ethdev: Fix check of threshold for TX freeing.**

Fixed issue where the parameter to `tx_free_thresh` was not consistent between the drivers.

- **ethdev: Fix crash if malloc of user callback fails.**

If `rte_zmalloc()` failed in `rte_eth_dev_callback_register` then the NULL pointer would be dereferenced.

- **ethdev: Fix illegal port access.**

To obtain a detachable flag, `pci_drv` is accessed in `rte_eth_dev_is_detachable()`. However `pci_drv` is only valid if port is enabled. Fixed by checking `rte_eth_dev_is_valid_port()` first.

- **ethdev: Make tables const.**
- **ethdev: Rename and extend the mirror type.**
- **examples/distributor: Fix debug macro.**

The macro to turn on additional debug output when the app was compiled with `-DDEBUG` was broken.

Fixes: 07db4a975094 (“examples/distributor: new sample app”)

- **examples/kni: Fix crash on exit.**
- **examples/vhost: Fix build with debug enabled.**

Fixes: 72ec8d77ac68 (“examples/vhost: rework duplicated code”)

- **fm10k: Fix RETA table initialization.**

The fm10k driver has 128 RETA entries in 32 registers, but it only initialized the first 32 when doing multiple RX queue configurations. This fix initializes all 128 entries.

- **fm10k: Fix RX buffer size.**
- **fm10k: Fix TX multi-segment frame.**
- **fm10k: Fix TX queue cleaning after start error.**
- **fm10k: Fix Tx queue cleaning after start error.**
- **fm10k: Fix default mac/vlan in switch.**
- **fm10k: Fix interrupt fault handling.**
- **fm10k: Fix jumbo frame issue.**
- **fm10k: Fix mac/vlan filtering.**
- **fm10k: Fix maximum VF number.**
- **fm10k: Fix maximum queue number for VF.**

Both PF and VF shared code in function `fm10k_stats_get()`. The function worked with PF, but had problems with VF since it has less queues than PF.

Fixes: a6061d9e7075 (“fm10k: register PF driver”)

- **fm10k: Fix queue disabling.**
- **fm10k: Fix switch synchronization.**
- **i40e/base: Fix error handling of NVM state update.**
- **i40e/base: Fix hardware port number for pass-through.**
- **i40e/base: Rework virtual address retrieval for lan queue.**
- **i40e/base: Update LED blinking.**
- **i40e/base: Workaround for PHY type with firmware < 4.4.**
- **i40e: Disable setting of PHY configuration.**
- **i40e: Fix SCTP flow director.**

- **i40e: Fix check of descriptor done flag.**

Fixes: 4861cde46116 (“i40e: new poll mode driver”) Fixes: 05999aab4ca6 (“i40e: add or delete flow director”)

- **i40e: Fix condition to get VMDQ info.**
- **i40e: Fix registers access from big endian CPU.**
- **i40evf: Clear command when error occurs.**
- **i40evf: Fix RSS with less RX queues than TX queues.**
- **i40evf: Fix crash when setup TX queues.**
- **i40evf: Fix jumbo frame support.**
- **i40evf: Fix offload capability flags.**

Added checksum offload capability flags which have already been supported for a long time.

- **ivshmem: Fix crash in corner case.**

Fixed issues where depending on the configured segments it was possible to hit a segmentation fault as a result of decrementing an unsigned index with value 0.

Fixes: 40b966a211ab (“ivshmem: library changes for mmaping using ivshmem”)

- **ixgbe/base: Fix SFP probing.**
  - **ixgbe/base: Fix TX pending clearing.**
  - **ixgbe/base: Fix X550 CS4227 address.**
  - **ixgbe/base: Fix X550 PCIe master disabling.**
  - **ixgbe/base: Fix X550 check.**
  - **ixgbe/base: Fix X550 init early return.**
  - **ixgbe/base: Fix X550 link speed.**
  - **ixgbe/base: Fix X550em CS4227 speed mode.**
  - **ixgbe/base: Fix X550em SFP+ link stability.**
  - **ixgbe/base: Fix X550em UniPHY link configuration.**
  - **ixgbe/base: Fix X550em flow control for KR backplane.**
  - **ixgbe/base: Fix X550em flow control to be KR only.**
  - **ixgbe/base: Fix X550em link setup without SFP.**
  - **ixgbe/base: Fix X550em mux after MAC reset.**
- Fixes: d2e72774e58c (“ixgbe/base: support X550”)
- **ixgbe/base: Fix bus type overwrite.**
  - **ixgbe/base: Fix init handling of X550em link down.**
  - **ixgbe/base: Fix lan id before first i2c access.**
  - **ixgbe/base: Fix mac type checks.**

- **ixgbe/base: Fix tunneled UDP and TCP frames in flow director.**
- **ixgbe: Check mbuf refcnt when clearing a ring.**

The function to clear the TX ring when a port was being closed, e.g. on exit in testpmd, was not checking the mbuf refcnt before freeing it. Since the function in the vector driver to clear the ring after TX does not setting the pointer to NULL post-free, this caused crashes if mbuf debugging was turned on.

- **ixgbe: Fix RX with buffer address not word aligned.**

Niantic HW expects the Header Buffer Address in the RXD must be word aligned.

- **ixgbe: Fix RX with buffer address not word aligned.**

- **ixgbe: Fix Rx queue reset.**

Fix to reset vector related RX queue fields to their initial values.

Fixes: c95584dc2b18 (“ixgbe: new vectorized functions for Rx/Tx”)

- **ixgbe: Fix TSO in IPv6.**

When TSO was used with IPv6, the generated frames were incorrect. The L4 frame was OK, but the length field of IPv6 header was not populated correctly.

- **ixgbe: Fix X550 flow director check.**

- **ixgbe: Fix check for split packets.**

The check for split packets to be reassembled in the vector ixgbe PMD was incorrectly only checking the first 16 elements of the array instead of all 32.

Fixes: cf4b4708a88a (“ixgbe: improve slow-path perf with vector scattered Rx”)

- **ixgbe: Fix data access on big endian cpu.**

- **ixgbe: Fix flow director flexbytes offset.**

Fixes: d54a9888267c (“ixgbe: support flexpayload configuration of flow director”)

- **ixgbe: Fix number of segments with vector scattered Rx.**

Fixes: cf4b4708a88a (ixgbe: improve slow-path perf with vector scattered Rx)

- **ixgbe: Fix offload config option name.**

The RX\_OLFLAGS option was renamed from DISABLE to ENABLE in the driver code and Linux config. It is now renamed also in the BSD config and documentation.

Fixes: 359f106a69a9 (“ixgbe: prefer enabling olflags rather than not disabling”)

- **ixgbe: Fix release queue mbufs.**

The calculations of what mbufs were valid in the RX and TX queues were incorrect when freeing the mbufs for the vector PMD. This led to crashes due to invalid reference counts when mbuf debugging was turned on, and possibly other more subtle problems (such as mbufs being freed when in use) in other cases.

Fixes: c95584dc2b18 (“ixgbe: new vectorized functions for Rx/Tx”)

- **ixgbe: Move PMD specific fields out of base driver.**

Move `rx_bulk_alloc_allowed` and `rx_vec_allowed` from `ixgbe_hw` to `ixgbe_adapter`.

Fixes: 01fa1d6215fa (“ixgbe: unify Rx setup”)

- **ixgbe: Rename TX queue release function.**
- **ixgbev: Fix RX function selection.**

The logic to select ixgbe the VF RX function is different than the PF.

- **ixgbev: Fix link status for PF up/down events.**
- **kni: Fix RX loop limit.**

Loop processing packets dequeued from rx\_q was using the number of packets requested, not how many it actually received.

- **kni: Fix ioctl in containers, like Docker.**
- **kni: Fix multicast ioctl handling.**
- **log: Fix crash after log\_history dump.**
- **lpm: Fix big endian support.**
- **lpm: Fix depth small entry add.**

- **mbuf: Fix cloning with private mbuf data.**

Added a new `priv_size` field in mbuf structure that should be initialized at mbuf pool creation. This field contains the size of the application private data in mbufs.

Introduced new static inline functions `rte_mbuf_from_indirect()` and `rte_mbuf_to_baddr()` to replace the existing macros, which take the private size into account when attaching and detaching mbufs.

- **mbuf: Fix data room size calculation in pool init.**

Deduct the mbuf data room size from `mempool->elt_size` and `priv_size`, instead of using an hardcoded value that is not related to the real buffer size.

To use `rte_pktmbuf_pool_init()`, the user can either:

- Give a NULL parameter to `rte_pktmbuf_pool_init()`: in this case, the private size is assumed to be 0, and the room size is `mp->elt_size - sizeof(struct rte_mbuf)`.
- Give the `rte_pktmbuf_pool_private` filled with appropriate `data_room_size` and `priv_size` values.

- **mbuf: Fix init when private size is not zero.**

Allow the user to use the default `rte_pktmbuf_init()` function even if the mbuf private size is not 0.

- **mempool: Add structure for object headers.**

Each object stored in mempools are prefixed by a header, allowing for instance to retrieve the mempool pointer from the object. When debug is enabled, a cookie is also added in this header that helps to detect corruptions and double-frees.

Introduced a structure that materializes the content of this header, and will simplify future patches adding things in this header.

- **mempool: Fix pages computation to determine number of objects.**



- **mempool: Fix returned value after counting objects.**  
Fixes: 148f963fb532 (“xen: core library changes”)
- **mlx4: Avoid requesting TX completion events to improve performance.**  
Instead of requesting a completion event for each TX burst, request it on a fixed schedule once every MLX4\_PMD\_TX\_PER\_COMP\_REQ (currently 64) packets to improve performance.
- **mlx4: Fix compilation as a shared library and on 32 bit platforms.**
- **mlx4: Fix possible crash on scattered mbuf allocation failure.**  
Fixes issue where failing to allocate a segment, `mlx4_rx_burst_sp()` could call `rte_pktmbuf_free()` on an incomplete scattered mbuf whose next pointer in the last segment is not set.
- **mlx4: Fix support for multiple vlan filters.**  
This fixes the “Multiple RX VLAN filters can be configured, but only the first one works” bug.
- **pcap: Fix storage of name and type in queues.**  
`pcap_rx_queue/pcap_tx_queue` should store it’s own copy of name/type values, not the pointer to temporary allocated space.
- **pci: Fix memory leaks and needless increment of map address.**
- **pci: Fix uio mapping differences between linux and bsd.**
- **port: Fix unaligned access to metadata.**  
Fix `RTE_MBUF_METADATA` macros to allow for unaligned accesses to meta-data fields.
- **ring: Fix return of new port id on creation.**
- **timer: Fix race condition.**  
Eliminate problematic race condition in `rte_timer_manage()` that can lead to corruption of per-icore pending-lists (implemented as skip-lists).
- **vfio: Fix overflow of BAR region offset and size.**  
Fixes: 90a1633b2347 (“eal/Linux: allow to map BARs with MSI-X tables”)
- **vhost: Fix enqueue/dequeue to handle chained vring descriptors.**
- **vhost: Fix race for connection fd.**
- **vhost: Fix virtio freeze due to missed interrupt.**
- **virtio: Fix crash if CQ is not negotiated.**  
Fix NULL dereference if virtio control queue is not negotiated.
- **virtio: Fix ring size negotiation.**  
Negotiate the virtio ring size. The host may allow for very large rings but application may only want a smaller ring. Conversely, if the number of descriptors requested exceeds the virtio host queue size, then just silently use the smaller host size.  
  
This fixes issues with virtio in non-QEMU environments. For example Google Compute Engine allows up to 16K elements in ring.

- **vmxnet3: Fix link state handling.**

## 10.3 Known Issues

- When running the `vmdq` sample or `vhost` sample applications with the Intel(R) XL710 (i40e) NIC, the configuration option `CONFIG_RTE_MAX_QUEUES_PER_PORT` should be increased from 256 to 1024.
- VM power manager may not work on systems with more than 64 cores.

## 10.4 API Changes

- The order that user supplied RX and TX callbacks are called in has been changed to the order that they were added (fifo) in line with end-user expectations. The previous calling order was the reverse of this (lifo) and was counter intuitive for users. The actual API is unchanged.

## 10.5 ABI Changes

- The `rte_hash` structure has been changed to internal use only.

## 11.1 New Features

- Poll-mode driver support for an early release of the PCIE host interface of the Intel(R) Ethernet Switch FM10000.
  - Basic Rx/Tx functions for PF/VF
  - Interrupt handling support for PF/VF
  - Per queue start/stop functions for PF/VF
  - Support Mailbox handling between PF/VF and PF/Switch Manager
  - Receive Side Scaling (RSS) for PF/VF
  - Scatter receive function for PF/VF
  - Reta update/query for PF/VF
  - VLAN filter set for PF
  - Link status query for PF/VF

---

**Note:** The software is intended to run on pre-release hardware and may contain unknown or unresolved defects or issues related to functionality and performance. The poll mode driver is also pre-release and will be updated to a released version post hardware and base driver release. Should the official hardware release be made between DPDK releases an updated poll-mode driver will be made available.

---

- Link Bonding
  - Support for adaptive load balancing (mode 6) to the link bonding library.
  - Support for registration of link status change callbacks with link bonding devices.
  - Support for slaves devices which do not support link status change interrupts in the link bonding library via a link status polling mechanism.
- PCI Hotplug with NULL PMD sample application
- ABI versioning
- x32 ABI
- Non-EAL Thread Support

- Multi-pthread Support
- Re-order Library
- ACL for AVX2
- Architecture Independent CRC Hash
- uio\_pci\_generic Support
- KNI Optimizations
- Vhost-user support
- Virtio (link, vlan, mac, port IO, perf)
- IXGBE-VF RSS
- RX/TX Callbacks
- Unified Flow Types
- Indirect Attached MBUF Flag
- Use default port configuration in TestPMD
- Tunnel offloading in TestPMD
- Poll Mode Driver - 40 GbE Controllers (librte\_pmd\_i40e)
  - Support for Flow Director
  - Support for ethertype filter
  - Support RSS in VF
  - Support configuring redirection table with different size from 1GbE and 10 GbE
  - 128/512 entries of 40GbE PF
  - 64 entries of 40GbE VF
  - Support configuring hash functions
  - Support for VXLAN packet on Intel® 40GbE Controllers
- Poll Mode Driver for Mellanox ConnectX-3 EN adapters (mlx4)

---

**Note:** This PMD is only available for Linux and is disabled by default due to external dependencies (libibverbs and libmlx4). Please refer to the NIC drivers guide for more information.

---

- Packet Distributor Sample Application
- Job Stats library and Sample Application.
- Enhanced Jenkins hash (jhash) library

---

**Note:** The hash values returned by the new jhash library are different from the ones returned by the previous library.

---

## 12.1 New Features

- Link Bonding
  - Support for 802.3ad link aggregation (mode 4) and transmit load balancing (mode 5) to the link bonding library.
  - Support for registration of link status change callbacks with link bonding devices.
  - Support for slaves devices which do not support link status change interrupts in the link bonding library via a link status polling mechanism.
- Poll Mode Driver - 40 GbE Controllers (librte\_pmd\_i40e)
  - Support for Flow Director
  - Support for ethertype filter
  - Support RSS in VF
  - Support configuring redirection table with different size from 1GbE and 10 GbE
  - 128/512 entries of 40GbE PF
  - 64 entries of 40GbE VF
  - Support configuring hash functions
  - Support for VXLAN packet on Intel 40GbE Controllers
- Packet Distributor Sample Application

## KNOWN ISSUES AND LIMITATIONS IN LEGACY RELEASES

This section describes known issues with the DPDK software that aren't covered in the version specific release notes sections.

### 13.1 Unit Test for Link Bonding may fail at `test_tlb_tx_burst()`

**Description:** Unit tests will fail in `test_tlb_tx_burst()` function with error for uneven distribution of packets.

**Implication:** Unit test `link_bonding_autotest` will fail.

**Resolution/Workaround:** There is no workaround available.

**Affected Environment/Platform:** Fedora 20.

**Driver/Module:** Link Bonding.

### 13.2 Pause Frame Forwarding does not work properly on igb

**Description:** For igb devices `rte_eth_flow_ctrl_set` does not work as expected. Pause frames are always forwarded on igb, regardless of the `RFCE`, `MPMCF` and `DPF` registers.

**Implication:** Pause frames will never be rejected by the host on 1G NICs and they will always be forwarded.

**Resolution/Workaround:** There is no workaround available.

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

### 13.3 In packets provided by the PMD, some flags are missing

**Description:** In packets provided by the PMD, some flags are missing. The application does not have access to information provided by the hardware (packet is broadcast, packet is multicast, packet is IPv4 and so on).

**Implication:** The `ol_flags` field in the `rte_mbuf` structure is not correct and should not be used.

**Resolution/Workaround:** The application has to parse the Ethernet header itself to get the information, which is slower.

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

## 13.4 The `rte_malloc` library is not fully implemented

**Description:** The `rte_malloc` library is not fully implemented.

**Implication:** All debugging features of `rte_malloc` library described in architecture documentation are not yet implemented.

**Resolution/Workaround:** No workaround available.

**Affected Environment/Platform:** All.

**Driver/Module:** `rte_malloc`.

## 13.5 HPET reading is slow

**Description:** Reading the HPET chip is slow.

**Implication:** An application that calls `rte_get_hpet_cycles()` or `rte_timer_manage()` runs slower.

**Resolution/Workaround:** The application should not call these functions too often in the main loop. An alternative is to use the TSC register through `rte_rdtsc()` which is faster, but specific to an Icore and is a cycle reference, not a time reference.

**Affected Environment/Platform:** All.

**Driver/Module:** Environment Abstraction Layer (EAL).

## 13.6 HPET timers do not work on the Osage customer reference platform

**Description:** HPET timers do not work on the Osage customer reference platform which includes an Intel® Xeon® processor 5500 series processor) using the released BIOS from Intel.

**Implication:** On Osage boards, the implementation of the `rte_delay_us()` function must be changed to not use the HPET timer.

**Resolution/Workaround:** This can be addressed by building the system with the `CONFIG_RTE_LIBEAL_USE_HPET=n` configuration option or by using the `--no-hpet` EAL option.

**Affected Environment/Platform:** The Osage customer reference platform. Other vendor platforms with Intel® Xeon® processor 5500 series processors should work correctly, provided the BIOS supports HPET.

**Driver/Module:** `lib/librte_eal/common/include/rte_cycles.h`

## 13.7 Not all variants of supported NIC types have been used in testing

**Description:** The supported network interface cards can come in a number of variants with different device ID's. Not all of these variants have been tested with the DPDK.

The NIC device identifiers used during testing:

- Intel® Ethernet Controller XL710 for 40GbE QSFP+ [8086:1584]
- Intel® Ethernet Controller XL710 for 40GbE QSFP+ [8086:1583]
- Intel® Ethernet Controller X710 for 10GbE SFP+ [8086:1572]
- Intel® 82576 Gigabit Ethernet Controller [8086:10c9]
- Intel® 82576 Quad Copper Gigabit Ethernet Controller [8086:10e8]
- Intel® 82580 Dual Copper Gigabit Ethernet Controller [8086:150e]
- Intel® I350 Quad Copper Gigabit Ethernet Controller [8086:1521]
- Intel® 82599 Dual Fibre 10 Gigabit Ethernet Controller [8086:10fb]
- Intel® Ethernet Server Adapter X520-T2 [8086: 151c]
- Intel® Ethernet Controller X540-T2 [8086:1528]
- Intel® 82574L Gigabit Network Connection [8086:10d3]
- Emulated Intel® 82540EM Gigabit Ethernet Controller [8086:100e]
- Emulated Intel® 82545EM Gigabit Ethernet Controller [8086:100f]
- Intel® Ethernet Server Adapter X520-4 [8086:154a]
- Intel® Ethernet Controller I210 [8086:1533]

**Implication:** Risk of issues with untested variants.

**Resolution/Workaround:** Use tested NIC variants. For those supported Ethernet controllers, additional device IDs may be added to the software if required.

**Affected Environment/Platform:** All.

**Driver/Module:** Poll-mode drivers

## 13.8 Multi-process sample app requires exact memory mapping

**Description:** The multi-process example application assumes that it is possible to map the hugepage memory to the same virtual addresses in client and server applications. Occasionally, very rarely with 64-bit, this does not occur and a client application will fail on startup. The Linux “address-space layout randomization” security feature can sometimes cause this to occur.

**Implication:** A multi-process client application fails to initialize.

**Resolution/Workaround:** See the “Multi-process Limitations” section in the DPDK Programmer’s Guide for more information.

**Affected Environment/Platform:** All.



**Driver/Module:** Multi-process example application

### 13.9 Packets are not sent by the 1 GbE/10 GbE SR-IOV driver when the source MAC is not the MAC assigned to the VF NIC

**Description:** The 1 GbE/10 GbE SR-IOV driver can only send packets when the Ethernet header's source MAC address is the same as that of the VF NIC. The reason for this is that the Linux `ixgbe` driver module in the host OS has its anti-spoofing feature enabled.

**Implication:** Packets sent using the 1 GbE/10 GbE SR-IOV driver must have the source MAC address correctly set to that of the VF NIC. Packets with other source address values are dropped by the NIC if the application attempts to transmit them.

**Resolution/Workaround:** Configure the Ethernet source address in each packet to match that of the VF NIC.

**Affected Environment/Platform:** All.

**Driver/Module:** 1 GbE/10 GbE VF Poll Mode Driver (PMD).

### 13.10 SR-IOV drivers do not fully implement the `rte_ethdev` API

**Description:** The SR-IOV drivers only supports the following `rte_ethdev` API functions:

- `rte_eth_dev_configure()`
- `rte_eth_tx_queue_setup()`
- `rte_eth_rx_queue_setup()`
- `rte_eth_dev_info_get()`
- `rte_eth_dev_start()`
- `rte_eth_tx_burst()`
- `rte_eth_rx_burst()`
- `rte_eth_dev_stop()`
- `rte_eth_stats_get()`
- `rte_eth_stats_reset()`
- `rte_eth_link_get()`
- `rte_eth_link_get_no_wait()`

**Implication:** Calling an unsupported function will result in an application error.

**Resolution/Workaround:** Do not use other `rte_ethdev` API functions in applications that use the SR-IOV drivers.

**Affected Environment/Platform:** All.

**Driver/Module:** VF Poll Mode Driver (PMD).

## 13.11 PMD does not work with `--no-huge` EAL command line parameter

**Description:** Currently, the DPDK does not store any information about memory allocated by `malloc()` (for example, NUMA node, physical address), hence PMD drivers do not work when the `--no-huge` command line parameter is supplied to EAL.

**Implication:** Sending and receiving data with PMD will not work.

**Resolution/Workaround:** Use huge page memory or use VFIO to map devices.

**Affected Environment/Platform:** Systems running the DPDK on Linux

**Driver/Module:** Poll Mode Driver (PMD).

## 13.12 Some hardware off-load functions are not supported by the VF Driver

**Description:** Currently, configuration of the following items is not supported by the VF driver:

- IP/UDP/TCP checksum offload
- Jumbo Frame Receipt
- HW Strip CRC

**Implication:** Any configuration for these items in the VF register will be ignored. The behavior is dependent on the current PF setting.

**Resolution/Workaround:** For the PF (Physical Function) status on which the VF driver depends, there is an option item under PMD in the config file. For others, the VF will keep the same behavior as PF setting.

**Affected Environment/Platform:** All.

**Driver/Module:** VF (SR-IOV) Poll Mode Driver (PMD).

## 13.13 Kernel crash on IGB port unbinding

**Description:** Kernel crash may occur when unbinding 1G ports from the `igb_uio` driver, on 2.6.3x kernels such as shipped with Fedora 14.

**Implication:** Kernel crash occurs.

**Resolution/Workaround:** Use newer kernels or do not unbind ports.

**Affected Environment/Platform:** 2.6.3x kernels such as shipped with Fedora 14

**Driver/Module:** IGB Poll Mode Driver (PMD).

### 13.14 Twinpond and Ironpond NICs do not report link status correctly

**Description:** Twin Pond/Iron Pond NICs do not bring the physical link down when shutting down the port.

**Implication:** The link is reported as up even after issuing `shutdown` command unless the cable is physically disconnected.

**Resolution/Workaround:** None.

**Affected Environment/Platform:** Twin Pond and Iron Pond NICs

**Driver/Module:** Poll Mode Driver (PMD).

### 13.15 Discrepancies between statistics reported by different NICs

**Description:** Gigabit Ethernet devices from Intel include CRC bytes when calculating packet reception statistics regardless of hardware CRC stripping state, while 10-Gigabit Ethernet devices from Intel do so only when hardware CRC stripping is disabled.

**Implication:** There may be a discrepancy in how different NICs display packet reception statistics.

**Resolution/Workaround:** None

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

### 13.16 Error reported opening files on DPDK initialization

**Description:** On DPDK application startup, errors may be reported when opening files as part of the initialization process. This occurs if a large number, for example, 500 or more, or if hugepages are used, due to the per-process limit on the number of open files.

**Implication:** The DPDK application may fail to run.

**Resolution/Workaround:** If using 2 MB hugepages, consider switching to a fewer number of 1 GB pages. Alternatively, use the `ulimit` command to increase the number of files which can be opened by a process.

**Affected Environment/Platform:** All.

**Driver/Module:** Environment Abstraction Layer (EAL).

### 13.17 Intel® QuickAssist Technology sample application does not work on a 32-bit OS on Shumway

**Description:** The Intel® Communications Chipset 89xx Series device does not fully support NUMA on a 32-bit OS. Consequently, the sample application cannot work properly on Shumway, since it requires NUMA on both nodes.

**Implication:** The sample application cannot work in 32-bit mode with emulated NUMA, on multi-socket boards.

**Resolution/Workaround:** There is no workaround available.

**Affected Environment/Platform:** Shumway

**Driver/Module:** All.

### 13.18 Differences in how different Intel NICs handle maximum packet length for jumbo frame

**Description:** 10 Gigabit Ethernet devices from Intel do not take VLAN tags into account when calculating packet size while Gigabit Ethernet devices do so for jumbo frames.

**Implication:** When receiving packets with VLAN tags, the actual maximum size of useful payload that Intel Gigabit Ethernet devices are able to receive is 4 bytes (or 8 bytes in the case of packets with extended VLAN tags) less than that of Intel 10 Gigabit Ethernet devices.

**Resolution/Workaround:** Increase the configured maximum packet size when using Intel Gigabit Ethernet devices.

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

### 13.19 Binding PCI devices to `igb_uio` fails on Linux kernel 3.9 when more than one device is used

**Description:** A known bug in the `uio` driver included in Linux kernel version 3.9 prevents more than one PCI device to be bound to the `igb_uio` driver.

**Implication:** The Poll Mode Driver (PMD) will crash on initialization.

**Resolution/Workaround:** Use earlier or later kernel versions, or apply the following [patch](#).

**Affected Environment/Platform:** Linux systems with kernel version 3.9

**Driver/Module:** `igb_uio` module

### 13.20 GCC might generate Intel® AVX instructions for processors without Intel® AVX support

**Description:** When compiling DPDK (and any DPDK app), `gcc` may generate Intel® AVX instructions, even when the processor does not support Intel® AVX.

**Implication:** Any DPDK app might crash while starting up.

**Resolution/Workaround:** Either compile using `icc` or set `EXTRA_CFLAGS='-O3'` prior to compilation.

**Affected Environment/Platform:** Platforms which processor does not support Intel® AVX.

**Driver/Module:** Environment Abstraction Layer (EAL).

## 13.21 Ethertype filter could receive other packets (non-assigned) in Niantic

**Description:** On Intel® Ethernet Controller 82599EB When Ethertype filter (priority enable) was set, unmatched packets also could be received on the assigned queue, such as ARP packets without 802.1q tags or with the user priority not equal to set value. Launch the testpmd by disabling RSS and with multiply queues, then add the ethertype filter like the following and then start forwarding:

```
add_ethertype_filter 0 ethertype 0x0806 priority enable 3 queue 2 index 1
```

When sending ARP packets without 802.1q tag and with user priority as non-3 by tester, all the ARP packets can be received on the assigned queue.

**Implication:** The user priority comparing in Ethertype filter cannot work probably. It is a NIC's issue due to the following: "In fact, ETQF.UP is not functional, and the information will be added in errata of 82599 and X540."

**Resolution/Workaround:** None

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

## 13.22 Cannot set link speed on Intel® 40G Ethernet controller

**Description:** On Intel® 40G Ethernet Controller you cannot set the link to specific speed.

**Implication:** The link speed cannot be changed forcibly, though it can be configured by application.

**Resolution/Workaround:** None

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

## 13.23 Devices bound to igb\_uio with VT-d enabled do not work on Linux kernel 3.15-3.17

**Description:** When VT-d is enabled (`iommu=pt intel_iommu=on`), devices are 1:1 mapped. In the Linux kernel unbinding devices from drivers removes that mapping which result in IOMMU errors. Introduced in Linux kernel 3.15 commit, solved in Linux kernel 3.18 commit.

**Implication:** Devices will not be allowed to access memory, resulting in following kernel errors:

```
dmar: DRHD: handling fault status reg 2
dmar: DMAR:[DMA Read] Request device [02:00.0] fault addr a0c58000
DMAR:[fault reason 02] Present bit in context entry is clear
```

**Resolution/Workaround:** Use earlier or later kernel versions, or avoid driver binding on boot by blacklisting the driver modules. I.e., in the case of `ixgbe`, we can pass the kernel command line option: `modprobe.blacklist=ixgbe`. This way we do not need to unbind the device to bind it to `igb_uio`.

**Affected Environment/Platform:** Linux systems with kernel versions 3.15 to 3.17.

**Driver/Module:** `igb_uio` module.

## 13.24 VM power manager may not work on systems with more than 64 cores

**Description:** When using VM power manager on a system with more than 64 cores, VM(s) should not use cores 64 or higher.

**Implication:** VM power manager should not be used with VM(s) that are using cores 64 or above.

**Resolution/Workaround:** Do not use cores 64 or above.

**Affected Environment/Platform:** Platforms with more than 64 cores.

**Driver/Module:** VM power manager application.

## 13.25 DPDK may not build on some Intel CPUs using clang < 3.7.0

**Description:** When compiling DPDK with an earlier version than 3.7.0 of clang, CPU flags are not detected on some Intel platforms such as Intel Broadwell/Skylake (and possibly future CPUs), and therefore compilation fails due to missing intrinsics.

**Implication:** DPDK will not build when using a clang version < 3.7.0.

**Resolution/Workaround:** Use clang 3.7.0 or higher, or gcc.

**Affected Environment/Platform:** Platforms with Intel Broadwell/Skylake using an old clang version.

**Driver/Module:** Environment Abstraction Layer (EAL).

## 13.26 The last EAL argument is replaced by the program name in argv[]

**Description:** The last EAL argument is replaced by program name in `argv[]` after `eal_parse_args` is called. This is the intended behavior but it causes the pointer to the last EAL argument to be lost.

**Implication:** If the last EAL argument in `argv[]` is generated by a malloc function, changing it will cause memory issues when freeing the argument.

**Resolution/Workaround:** An application should not consider the value in `argv[]` as unchanged.

**Affected Environment/Platform:** ALL.

**Driver/Module:** Environment Abstraction Layer (EAL).

## 13.27 I40e VF may not receive packets in the promiscuous mode

**Description:** Promiscuous mode is not supported by the DPDK i40e VF driver when using the i40e Linux kernel driver as host driver.

**Implication:** The i40e VF does not receive packets when the destination MAC address is unknown.

**Resolution/Workaround:** Use an explicit destination MAC address that matches the VF.

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

## 13.28 uio\_pci\_generic module bind failed in X710/XL710/XXV710

**Description:** The `uio_pci_generic` module is not supported by XL710, since the errata of XL710 states that the Interrupt Status bit is not implemented. The errata is the item #71 from the [xl710 controller spec](#). The hw limitation is the same as other X710/XXV710 NICs.

**Implication:** When use `--bind=uio_pci_generic`, the `uio_pci_generic` module probes device and check the Interrupt Status bit. Since it is not supported by X710/XL710/XXV710, it return a *failed* value. The statement that these products don't support INTx masking, is indicated in the related [linux kernel commit](#).

**Resolution/Workaround:** Do not bind the `uio_pci_generic` module in X710/XL710/XXV710 NICs.

**Affected Environment/Platform:** All.

**Driver/Module:** Poll Mode Driver (PMD).

## 13.29 virtio tx\_burst() function cannot do TSO on shared packets

**Description:** The standard TX function of virtio driver does not manage shared packets properly when doing TSO. These packets should be read-only but the driver modifies them.

When doing TSO, the virtio standard expects that the L4 checksum is set to the pseudo header checksum in the packet data, which is different than the DPDK API. The driver patches the L4 checksum to conform to the virtio standard, but this solution is invalid when dealing with shared packets (clones), because the packet data should not be modified.

**Implication:** In this situation, the shared data will be modified by the driver, potentially causing race conditions with the other users of the mbuf data.

**Resolution/Workaround:** The workaround in the application is to ensure that the network headers in the packet data are not shared.

**Affected Environment/Platform:** Virtual machines running a virtio driver.

**Driver/Module:** Poll Mode Driver (PMD).

### 13.30 igb\_uio legacy mode can not be used in X710/XL710/XXV710

**Description:** X710/XL710/XXV710 NICs lack support for indicating INTx is asserted via the interrupt bit in the PCI status register. Linux deleted them from INTx support table. The related [commit](#).

**Implication:** When `insmod igb_uio` with `intr_mode=legacy` and test link status interrupt. Since INTx interrupt is not supported by X710/XL710/XXV710, it will cause Input/Output error when reading file descriptor.

**Resolution/Workaround:** Do not bind `igb_uio` with legacy mode in X710/XL710/XXV710 NICs, or do not use kernel version >4.7 when you bind `igb_uio` with legacy mode.

**Affected Environment/Platform:** ALL.

**Driver/Module:** Poll Mode Driver (PMD).

### 13.31 igb\_uio can not be used when running I3fwd-power

**Description:** Link Status Change(LSC) interrupt and packet receiving interrupt are all enabled in I3fwd-power APP. Because of UIO only support one interrupt, so these two kinds of interrupt need to share one, and the receiving interrupt have the higher priority, so can't get the right link status.

**Implication:** When `insmod igb_uio` and running I3fwd-power APP, link status getting doesn't work properly.

**Resolution/Workaround:** Use `vfiopci` when LSC and packet receiving interrupt enabled.

**Affected Environment/Platform:** ALL.

**Driver/Module:** `igb_uio` module.

### 13.32 AVX-512 support disabled

**Description:** AVX-512 support has been disabled on some conditions. This shouldn't be confused with `CONFIG_RTE_ENABLE_AVX512` config option which is already disabled by default. This config option defines if AVX-512 specific implementations of some file to be used or not. What has been disabled is compiler feature to produce AVX-512 instructions from any source code.

On DPDK v18.11 AVX-512 is disabled for all GCC builds which reported to cause a performance drop.

On DPDK v19.02 AVX-512 disable scope is reduced to GCC and `binutils` version 2.30 based on information accured from the GCC community defect.

**Reason:** Generated AVX-512 code cause crash: [https://bugs.dpdk.org/show\\_bug.cgi?id=97](https://bugs.dpdk.org/show_bug.cgi?id=97)  
[https://gcc.gnu.org/bugzilla/show\\_bug.cgi?id=88096](https://gcc.gnu.org/bugzilla/show_bug.cgi?id=88096)

**Resolution/Workaround:**



- Update `binutils` to newer version than 2.30.

OR

- Use different compiler, like `clang` for this case.

**Affected Environment/Platform:** GCC and `binutils` version 2.30.

**Driver/Module:** ALL.

## ABI AND API DEPRECATION

See the guidelines document for details of the ABI policy. API and ABI deprecation notices are to be posted here.

### 14.1 Deprecation Notices

- **eal:** several API and ABI changes are planned for `rte_devargs` in v18.02. The format of device command line parameters will change. The bus will need to be explicitly stated in the device declaration. The enum `rte_devtype` was used to identify a bus and will disappear. The structure `rte_devargs` will change. The `rte_devargs_list` will be made private. The following functions are deprecated starting from 17.08 and will either be modified or removed in 18.02:
  - `rte_eal_devargs_add`
  - `rte_eal_devargs_type_count`
  - `rte_eal_parse_devargs_str`, replaced by `rte_eal_devargs_parse`
- **pci:** Several exposed functions are misnamed. The following functions are deprecated starting from v17.11 and are replaced:
  - `eal_parse_pci_BDF` replaced by `rte_pci_addr_parse`
  - `eal_parse_pci_DomBDF` replaced by `rte_pci_addr_parse`
  - `rte_eal_compare_pci_addr` replaced by `rte_pci_addr_cmp`
- **ethdev:** a new Tx and Rx offload API was introduced on 17.11. In the new API, offloads are divided into per-port and per-queue offloads. Offloads are disabled by default and enabled per application request. The old offloads API is target to be deprecated on 18.05. This includes:
  - removal of `ETH_TXQ_FLAGS_NO*` flags.
  - removal of `txq_flags` field from `rte_eth_txconf` struct.
  - removal of the offloads bit-field from `rte_eth_rxmode` struct.
- **ethdev:** the legacy filter API, including `rte_eth_dev_filter_supported()`, `rte_eth_dev_filter_ctrl()` as well as filter types `MACVLAN`, `ETHERTYPE`, `FLEXIBLE`, `SYN`, `NTUPLE`, `TUNNEL`, `FDIR`, `HASH` and `L2_TUNNEL`, is superseded by the generic flow API (`rte_flow`) in PMDs that implement the latter. Target release for removal of the legacy API will be defined once most PMDs have switched to `rte_flow`.

- i40e: The default flexible payload configuration which extracts the first 16 bytes of the payload for RSS will be deprecated starting from 18.02. If required the previous behavior can be configured using existing flow director APIs. There is no ABI/API break. This change will just remove a global configuration setting and require explicit configuration.
- librte\_meter: The API will change to accommodate configuration profiles. Most of the API functions will have an additional opaque parameter.