



DPDK

DATA PLANE DEVELOPMENT KIT

Release Notes

Release 19.11.10

Sep 06, 2021

CONTENTS

1 DPDK Release 19.11	1
2 DPDK Release 19.08	97
3 DPDK Release 19.05	108
4 DPDK Release 19.02	120
5 DPDK Release 18.11	131
6 DPDK Release 18.08	145
7 DPDK Release 18.05	154
8 DPDK Release 18.02	170
9 DPDK Release 17.11	180
10 DPDK Release 17.08	193
11 DPDK Release 17.05	203
12 DPDK Release 17.02	216
13 DPDK Release 16.11	227
14 DPDK Release 16.07	236
15 DPDK Release 16.04	245
16 DPDK Release 2.2	256
17 DPDK Release 2.1	266
18 DPDK Release 2.0	280
19 DPDK Release 1.8	282
20 Known Issues and Limitations in Legacy Releases	283
21 ABI and API Deprecation	295

1.1 New Features

- **Added support for `-base-virtaddr` EAL option to FreeBSD.**

The FreeBSD version of DPDK now also supports setting base virtual address for mapping pages and resources into its address space.

- **Added Lock-free Stack for aarch64.**

Enabled the lock-free stack implementation for aarch64 platforms.

- **Extended `rte_pktmbuf_mempool_private` structure.**

`rte_pktmbuf_pool_private` structure was extended to include `flags` field for future compatibility. As per 19.11 release this field is reserved and should be set to 0 by the user.

+* **Changed mempool allocation behavior.**

Changed the mempool allocation behaviour so that objects no longer cross pages by default. Note, this may consume more memory when using small memory pages.

- **Added support for dynamic fields and flags in `mbuf`.**

This new feature adds the ability to dynamically register some room for a field or a flag in the `mbuf` structure. This is typically used for specific offload features, where adding a static field or flag in the `mbuf` is not justified.

- **Added support for hairpin queues.**

On supported NICs, we can now setup hairpin queues which will offload packets from the wire, back to the wire.

- **Added flow tag in `rte_flow`.**

The `SET_TAG` action and `TAG` item have been added to support transient flow tag.

- **Extended metadata support in `rte_flow`.**

Flow metadata has been extended to both Rx and Tx.

- Tx metadata can also be set by `SET_META` action of `rte_flow`.
- Rx metadata is delivered to the host via a dynamic field of `rte_mbuf` with `PKT_RX_DYNF_METADATA`.

- **Added `ethdev` API to set supported packet types.**

- Added new API `rte_eth_dev_set_ptypes` which allows an application to inform a PMD about a reduced range of packet types to handle.

- This scheme will allow PMDs to avoid lookup of internal ptype table on Rx and thereby improve Rx performance if the application wishes to do so.
- **Added Rx offload flag to enable or disable RSS update.**
 - Added new Rx offload flag `DEV_RX_OFFLOAD_RSS_HASH` which can be used to enable/disable PMDs write to `rte_mbuf::hash::rss`.
 - PMDs notify the validity of `rte_mbuf::hash::rss` to the application by enabling `PKT_RX_RSS_HASH` flag in `rte_mbuf::ol_flags`.
- **Added Rx/Tx packet burst mode “get” API.**

Added two new functions `rte_eth_rx_burst_mode_get` and `rte_eth_tx_burst_mode_get` that allow an application to retrieve the mode information about Rx/Tx packet burst such as Scalar or Vector, and Vector technology like AVX2.
- **Added Hisilicon hns3 PMD.**

Added the new `hns3` net driver for the inbuilt Hisilicon Network Subsystem 3 (HNS3) network engine found in the Hisilicon Kunpeng 920 SoC. See the `./nics/hns3` guide for more details on this new driver.
- **Added NXP PFE PMD.**

Added the new PFE driver for the NXP LS1012A platform. See the `./nics/pfe` NIC driver guide for more details on this new driver.
- **Updated Broadcom bnxt driver.**

Updated Broadcom `bnxt` driver with new features and improvements, including:

 - Added support for hot firmware upgrade.
 - Added support for error recovery.
 - Added support for querying and using COS classification in hardware.
 - Added LRO support Thor devices.
 - Update HWRM API to version 1.10.1.6
- **Updated the enic driver.**
 - Added support for Geneve with options offload.
 - Added flow API implementation based on VIC Flow Manager API.
- **Updated iavf PMD.**

Enable AVX2 data path for `iavf` PMD.
- **Updated the Intel e1000 driver.**

Added support for the `RTE_ETH_DEV_CLOSE_REMOVE` flag.
- **Updated the Intel ixgbe driver.**

Added support for the `RTE_ETH_DEV_CLOSE_REMOVE` flag.
- **Updated the Intel i40e driver.**

Added support for the `RTE_ETH_DEV_CLOSE_REMOVE` flag.

- **Updated the Intel fm10k driver.**

Added support for the `RTE_ETH_DEV_CLOSE_REMOVE` flag.

- **Updated the Intel ice driver.**

Updated the Intel ice driver with new features and improvements, including:

- Added support for device-specific DDP package loading.
- Added support for handling Receive Flex Descriptor.
- Added support for protocol extraction on per Rx queue.
- Added support for Flow Director filter based on generic filter framework.
- Added support for the `RTE_ETH_DEV_CLOSE_REMOVE` flag.
- Generic filter enhancement - Supported pipeline mode. - Supported new packet type like PPPoE for switch filter.
- Supported input set change and symmetric hash by `rte_flow` RSS action.
- Added support for GTP Tx checksum offload.
- Added new device IDs to support E810_XXV devices.

- **Updated the Huawei hinic driver.**

Updated the Huawei hinic driver with new features and improvements, including:

- Enabled SR-IOV - Partially supported at this point, VFIO only.
- Supported VLAN filter and VLAN offload.
- Supported Unicast MAC filter and Multicast MAC filter.
- Supported Flow API for LACP, VRRP, BGP and so on.
- Supported FW version get.

- **Updated Mellanox mlx5 driver.**

Updated Mellanox mlx5 driver with new features and improvements, including:

- Added support for VLAN pop flow offload command.
- Added support for VLAN push flow offload command.
- Added support for VLAN set PCP offload command.
- Added support for VLAN set VID offload command.
- Added support for matching on packets with Geneve tunnel header.
- Added hairpin support.
- Added ConnectX-6 Dx support.
- Flow engine selected based on RDMA Core library version. DV flow engine selected if version is `rdma-core-24.0` or higher. Verbs flow engine selected otherwise.

- **Updated the AF_XDP PMD.**

Updated the AF_XDP PMD. The new features include:

- Enabled zero copy between application mempools and UMEM by enabling the `XDP_UMEM_UNALIGNED_CHUNKS` UMEM flag.
- **Added cryptodev asymmetric session-less operation.**

Added a session-less option to the cryptodev asymmetric structure. It works the same way as symmetric crypto, and the corresponding transform is used directly by the crypto operation.
- **Added Marvell NITROX symmetric crypto PMD.**

Added a symmetric crypto PMD for Marvell NITROX V security processor. See the `./cryptodevs/nitrox` guide for more details on this new PMD.
- **Added asymmetric support to Marvell OCTEON TX crypto PMD.**

Added support for asymmetric operations to Marvell OCTEON TX crypto PMD. Supports RSA and modexp operations.
- **Added Marvell OCTEON TX2 crypto PMD.**

Added a new PMD driver for hardware crypto offload block on OCTEON TX2 SoC.
See `./cryptodevs/octeontx2` for more details
- **Updated NXP crypto PMDs for PDCP support.**

Added PDCP support to the DPAA_SEC and DPAA2_SEC PMDs using `rte_security` APIs. Support has been added for all sequence number sizes for control and user plane. Test and test-crypto-perf applications have been updated for unit testing.
- **Updated the AESNI-MB PMD.**
 - Added support for intel-ipsec-mb version 0.53.
- **Updated the AESNI-GCM PMD.**
 - Added support for intel-ipsec-mb version 0.53.
 - Added support for in-place chained mbufs with AES-GCM algorithm.
- **Enabled Single Pass GCM acceleration on QAT GEN3.**

Added support for Single Pass GCM, available on QAT GEN3 only (Intel QuickAssist Technology C4xxx). It is automatically chosen instead of the classic 2-pass mode when running on QAT GEN3, significantly improving the performance of AES GCM operations.
- **Updated the Intel QuickAssist Technology (QAT) asymmetric crypto PMD.**
 - Added support for asymmetric session-less operations.
 - Added support for RSA algorithm with pair (n, d) private key representation.
 - Added support for RSA algorithm with quintuple private key representation.
- **Updated the Intel QuickAssist Technology (QAT) compression PMD.**

Added stateful decompression support in the Intel QuickAssist Technology PMD. Please note that stateful compression is not supported.
- **Added external buffers support for dpdk-test-compress-perf tool.**

Added a command line option to the `dpdk-test-compress-perf` tool to allocate and use memory zones as external buffers instead of keeping the data directly in mbuf areas.
- **Updated the IPSec library.**

- Added Security Associations (SA) Database API to `librte_ipsec`. A new `test-sad` application has also been introduced to evaluate and perform custom functional and performance tests for an IPsec SAD implementation.
- Support fragmented packets in inline crypto processing mode with fallback `lookaside-none` session. Corresponding changes are also added in the IPsec Security Gateway application.
- **Introduced FIFO for NTB PMD.**

Introduced FIFO for NTB (Non-transparent Bridge) PMD to support packet based processing.
- **Added eBPF JIT support for arm64.**

Added eBPF JIT support for arm64 architecture to improve the eBPF program performance.
- **Added RIB and FIB (Routing/Forwarding Information Base) libraries.**

Added Routing and Forwarding Information Base (RIB/FIB) libraries. RIB and FIB can replace the LPM (Longest Prefix Match) library with better control plane (RIB) performance. The data plane (FIB) can be extended with new algorithms.
- **Updated testpmd with a command for ptypes.**
 - Added a console command to testpmd app, `show port (port_id) ptypes` which gives ability to print port supported ptypes in different protocol layers.
 - Packet type detection disabled by default for the supported PMDs.
- **Added new l2fwd-event sample application.**

Added an example application `l2fwd-event` that adds event device support to the traditional `l2fwd` example. It demonstrates usage of poll and event mode IO mechanism under a single application.
- **Added build support for Link Time Optimization.**

LTO is an optimization technique used by the compiler to perform whole program analysis and optimization at link time. In order to do that compilers store their internal representation of the source code that the linker uses at the final stage of the compilation process.

See `../prog_guide/lto` for more information:
- **Added IOVA as VA support for KNI.**
 - Added IOVA = VA support for KNI. KNI can operate in IOVA = VA mode when `iova-mode=va` EAL option is passed to the application or when bus IOVA scheme is selected as `RTE_IOVA_VA`. This mode only works on Linux Kernel versions 4.10.0 and above.
 - Due to IOVA to KVA address translations, based on the KNI use case there can be a performance impact. For mitigation, forcing IOVA to PA via EAL `--iova-mode=pa` option can be used, `IOVA_DC` bus iommu scheme can also result in IOVA as PA.

1.2 Removed Items

- Removed library-level ABI versions. These have been replaced with a single project-level ABI version for non-experimental libraries and an ABI version of 0 for experimental libraries. Review the `../contributing/abi_policy` and `../contributing/abi_versioning` guides for more information.

- Removed duplicated set of commands for Rx offload configuration from testpmd:

```
port config all crc-strip|scatter|rx-cksum|rx-timestamp|
              hw-vlan|hw-vlan-filter|hw-vlan-strip|hw-vlan-extend on|off
```

The testpmd command set that can be used instead in order to enable or disable Rx offloading on all Rx queues of a port is:

```
port config <port_id> rx_offload crc_strip|scatter|
                              ipv4_cksum|udp_cksum|tcp_cksum|timestamp|
                              vlan_strip|vlan_filter|vlan_extend on|off
```

- Removed AF_XDP pmd_zero copy vdev argument. Support is now auto-detected.
- The following sample applications have been removed in this release:
 - Exception Path
 - L3 Forwarding in a Virtualization Environment
 - Load Balancer
 - Netmap Compatibility
 - Quota and Watermark
 - vhost-scsi
- Removed arm64-dpaa2-* build config. arm64-dpaa-* can now build for both dpaa and dpaa2 platforms.

1.3 API Changes

- eal: made the lcore_config struct and global symbol private.
- eal: removed the rte_cpu_check_supported function, replaced by rte_cpu_is_supported since dpdk v17.08.
- eal: removed the rte_malloc_virt2phy function, replaced by rte_malloc_virt2iova since v17.11.
- eal: made the rte_config struct and rte_eal_get_configuration function private.
- mem: hid the internal malloc_heap structure and the rte_malloc_heap.h header.
- vfio: removed rte_vfio_dma_map and rte_vfio_dma_unmap that have been marked as deprecated in release 19.05. rte_vfio_container_dma_map and rte_vfio_container_dma_unmap can be used as substitutes.
- pci: removed the following functions deprecated since dpdk v17.11:
 - 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
- The network structure esp_tail has been prefixed by rte_.
- The network definitions of PPPoE ethertypes have been prefixed by RTE_.
- The network structure for MPLS has been prefixed by rte_.

- ethdev: changed `rte_eth_dev_infos_get` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: changed `rte_eth_promiscuous_enable` and `rte_eth_promiscuous_disable` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: changed `rte_eth_allmulticast_enable` and `rte_eth_allmulticast_disable` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: changed `rte_eth_dev_xstats_reset` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: changed `rte_eth_link_get` and `rte_eth_link_get_nowait` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: changed `rte_eth_macaddr_get` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: changed `rte_eth_dev_owner_delete` return value from `void` to `int` to provide a way to report various error conditions.
- ethdev: The deprecated function `rte_eth_dev_count` was removed. The function `rte_eth_dev_count_avail` is a drop-in replacement. If the intent is to iterate over ports, `RTE_ETH_FOREACH_*` macros are better port iterators.
- ethdev: `RTE_FLOW_ITEM_TYPE_META` data endianness altered to host one. Due to the new dynamic metadata field in mbuf is host-endian either, there is a minor compatibility issue for applications in case of 32-bit values supported.
- ethdev: the `tx_metadata` mbuf field is moved to dynamic one. `PKT_TX_METADATA` flag is replaced with `PKT_TX_DYNF_METADATA`. `DEV_TX_OFFLOAD_MATCH_METADATA` offload flag is removed, now metadata support in PMD is engaged on dynamic field registration.
- event: The function `rte_event_eth_tx_adapter_enqueue` takes an additional input as `flags`. Flag `RTE_EVENT_ETH_TX_ADAPTER_ENQUEUE_SAME_DEST` which has been introduced in this release is used when all the packets enqueued in the Tx adapter are destined for the same Ethernet port and Tx queue.
- sched: The pipe nodes configuration parameters such as number of pipes, pipe queue sizes, pipe profiles, etc., are moved from port level structure to subport level. This allows different subports of the same port to have different configuration for the pipe nodes.

1.4 ABI Changes

- policy: Please note the revisions to the `../contributing/abi_policy` introducing major ABI versions, with DPDK 19.11 becoming the first major version `v20`. ABI changes to add new features continue to be permitted in subsequent releases, with the condition that ABI compatibility with the major ABI version is maintained.
- net: The Ethernet address and other header definitions have changed attributes. They have been modified to be aligned on 2-byte boundaries. These changes should not impact normal usage because drivers naturally align the Ethernet header on receive and all known encapsulations preserve the alignment of the header.

- **security:** The field `replay_win_sz` has been moved from the ipsec library based `rte_ipsec_sa_prm` structure to security library based structure `rte_security_ipsec_xform`, which specify the anti-replay window size to enable sequence replay attack handling.
- **ipsec:** The field `replay_win_sz` has been removed from the structure `rte_ipsec_sa_prm` as it has been added to the security library.
- **ethdev:** Added 32-bit fields for maximum LRO aggregated packet size, in struct `rte_eth_dev_info` for the port capability and in struct `rte_eth_rxmode` for the port configuration. Application should use the new field in struct `rte_eth_rxmode` to configure the requested size. PMD should use the new field in struct `rte_eth_dev_info` to report the supported port capability.

1.5 Shared Library Versions

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

```

librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
librte_bpf.so.1
librte_bus_dpaa.so.2
librte_bus_fslmc.so.2
librte_bus_ifpga.so.2
librte_bus_pci.so.2
librte_bus_vdev.so.2
librte_bus_vmbus.so.2
librte_cfgfile.so.2
librte_cmdline.so.2
librte_compressdev.so.1
librte_cryptodev.so.8
librte_distributor.so.1
+ librte_eal.so.12
librte_efd.so.1
+ librte_ethdev.so.13
+ librte_eventdev.so.8
+ librte_fib.so.1
librte_flow_classify.so.1
librte_gro.so.1
librte_gso.so.1
librte_hash.so.2
librte_ip_frag.so.1
+ librte_ipsec.so.2
librte_jobstats.so.1
librte_kni.so.2
librte_kvargs.so.1
librte_latencystats.so.1
librte_lpm.so.2
librte_mbuf.so.5
librte_member.so.1
librte_mempool.so.5
librte_meter.so.3
librte_metrics.so.1
librte_net.so.1
+ librte_pci.so.2
librte_pdump.so.3
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_dpaa2_qdma.so.1
librte_pmd_ring.so.2
librte_pmd_softnic.so.1
librte_pmd_vhost.so.2
librte_port.so.3
librte_power.so.1
librte_rawdev.so.1
+ librte_rib.so.1
librte_rcu.so.1
librte_reorder.so.1
librte_ring.so.2
+ librte_sched.so.4
+ librte_security.so.3
librte_stack.so.1
librte_table.so.3
librte_timer.so.1
librte_vhost.so.4
```

1.6 Known Issues

1.7 Tested Platforms

- Intel® platforms with Intel® NICs combinations
 - CPU
 - * Intel® Atom™ CPU C3758 @ 2.20GHz
 - * Intel® Atom™ CPU C3858 @ 2.00GHz
 - * Intel® Atom™ CPU C3958 @ 2.00GHz
 - * Intel® Xeon® CPU D-1541 @ 2.10GHz
 - * Intel® Xeon® CPU D-1553N @ 2.30GHz
 - * Intel® Xeon® CPU E5-2680 0 @ 2.70GHz
 - * Intel® Xeon® CPU E5-2680 v2 @ 2.80GHz
 - * Intel® Xeon® CPU E5-2699 v3 @ 2.30GHz
 - * Intel® Xeon® CPU E5-2699 v4 @ 2.20GHz
 - * Intel® Xeon® Gold 6139 CPU @ 2.30GHz
 - * Intel® Xeon® Gold 6252N CPU @ 2.30GHz
 - * Intel® Xeon® Platinum 8180 CPU @ 2.50GHz
 - * Intel® Xeon® Platinum 8280M CPU @ 2.70GHz
 - OS:
 - * CentOS 7.6
 - * Fedora 30
 - * FreeBSD 12.0
 - * Red Hat Enterprise Linux Server release 8.0
 - * Red Hat Enterprise Linux Server release 7.6

- * Suse12SP3
 - * Ubuntu 14.04
 - * Ubuntu 16.04
 - * Ubuntu 16.10
 - * Ubuntu 18.04
 - * Ubuntu 19.04
- NICs:
- * Intel® Corporation Ethernet Controller E810-C for SFP (2x25G)
 - Firmware version: 1.02 0x80002084 1.2538.0/1.02 0x80002082 1.2538.0
 - Device id (pf): 8086:1593
 - Driver version: 0.12.25 (ice)
 - * Intel® Corporation Ethernet Controller E810-C for SFP (2x100G)
 - Firmware version: 1.02 0x80002081 1.2538.0
 - Device id (pf): 8086:1592
 - Driver version: 0.12.25 (ice)
 - * Intel® 82599ES 10 Gigabit Ethernet Controller
 - Firmware version: 0x61bf0001
 - Device id (pf/vf): 8086:10fb / 8086:10ed
 - Driver version: 5.6.1 (ixgbe)
 - * Intel® Corporation Ethernet Connection X552/X557-AT 10GBASE-T
 - Firmware version: 0x800003e7
 - Device id (pf/vf): 8086:15ad / 8086:15a8
 - Driver version: 5.1.0 (ixgbe)
 - * Intel® Corporation Ethernet Controller 10G X550T
 - Firmware version: 0x80000482
 - Device id (pf): 8086:1563
 - Driver version: 5.6.1 (ixgbe)
 - * Intel® Ethernet Converged Network Adapter X710-DA4 (4x10G)
 - Firmware version: 7.00 0x80004cdb
 - Device id (pf/vf): 8086:1572 / 8086:154c
 - Driver version: 2.9.21 (i40e)
 - * Intel® Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 4.10 0x80001a3c
 - Device id (pf/vf): 8086:37d0 / 8086:37cd

- Driver version: 2.9.21 (i40e)
- * Intel® Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 7.00 0x80004cf8
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.9.21 (i40e)
- * Intel® Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
 - Firmware version: 7.00 0x80004c97
 - Device id (pf/vf): 8086:1583 / 8086:154c
 - Driver version: 2.9.21 (i40e)
- * Intel® Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x80000cbc
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- * Intel® Corporation I210 Gigabit Network Connection
 - Firmware version: 3.25, 0x800006eb
 - Device id (pf): 8086:1533
 - Driver version: 5.4.0-k(igb)
- ARMv8 SoC combinations from Marvell (with integrated NICs)
 - SoC:
 - * CN83xx, CN96xx, CN93xx
 - OS (Based on Marvell OCTEON TX SDK-10.1.2.0):
 - * Arch Linux
 - * Buildroot 2018.11
 - * Ubuntu 16.04.1 LTS
 - * Ubuntu 16.10
 - * Ubuntu 18.04.1
 - * Ubuntu 19.04
- Intel® platforms with Mellanox® NICs combinations
 - CPU:
 - * Intel® Xeon® Gold 6154 CPU @ 3.00GHz
 - * Intel® Xeon® CPU E5-2697A v4 @ 2.60GHz
 - * Intel® Xeon® CPU E5-2697 v3 @ 2.60GHz
 - * Intel® Xeon® CPU E5-2680 v2 @ 2.80GHz
 - * Intel® Xeon® CPU E5-2650 v4 @ 2.20GHz

- * Intel® Xeon® CPU E5-2640 @ 2.50GHz
- * Intel® Xeon® CPU E5-2620 v4 @ 2.10GHz
- OS:
 - * Red Hat Enterprise Linux Server release 8.0 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.7 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.6 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.3 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.2 (Maipo)
 - * Ubuntu 19.04
 - * Ubuntu 18.10
 - * Ubuntu 18.04
 - * Ubuntu 16.04
 - * SUSE Linux Enterprise Server 15
- OFED:
 - * MLNX_OFED 4.6-1.0.1.1
 - * MLNX_OFED 4.7-1.0.0.1
 - * MLNX_OFED 4.7-3.1.9.0 and above
- upstream kernel:
 - * Linux 5.3 and above
- rdma-core:
 - * rdma-core-24.1-1 and above
- NICs:
 - * Mellanox® ConnectX®-3 Pro 40G MCX354A-FCC_Ax (2x40G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1007
 - Firmware version: 2.42.5000
 - * Mellanox® ConnectX®-4 10G MCX4111A-XCAT (1x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
 - * Mellanox® ConnectX®-4 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8

- Device ID: 15b3:1013
- Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 25G MCX4111A-ACAT (1x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 40G MCX415A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 50G MCX414A-BCAT (2x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 50G MCX415A-GCAT/MCX416A-BCAT/MCX416A-GCAT (2x50G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 50G MCX415A-CCAT (1x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013

- Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 100G MCX416A-CCAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.26.2032 and above
- * Mellanox® ConnectX®-4 Lx 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.26.2032 and above
- * Mellanox® ConnectX®-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.26.2032 and above
- * Mellanox® ConnectX®-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.26.2032 and above
- * Mellanox® ConnectX®-5 Ex EN 100G MCX516A-CDAT (2x100G)
 - Host interface: PCI Express 4.0 x16
 - Device ID: 15b3:1019
 - Firmware version: 16.26.2032 and above
- IBM Power 9 platforms with Mellanox® NICs combinations
 - CPU:
 - * POWER9 2.2 (pvr 004e 1202) 2300MHz
 - OS:
 - * Ubuntu 18.04.1 LTS (Bionic Beaver)
 - NICs:
 - * Mellanox® ConnectX®-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.26.1040
 - OFED:
 - * MLNX_OFED 4.7-1.0.0.2

1.8 19.11.1 Release Notes

1.8.1 19.11.1 Fixes

- acl: fix 32-bit match for range field
- app/eventdev: fix pipeline test with meson build
- app/pdump: fix build with clang
- app/testpmd: add port check before manual detach
- app/testpmd: call cleanup on exit
- app/testpmd: fix device mcast list error handling
- app/testpmd: fix GENEVE flow item
- app/testpmd: fix hot-unplug detaching
- app/testpmd: fix identifier size for port attach
- app/testpmd: fix initial value when setting PFC
- app/testpmd: fix RFC addresses for Tx only
- app/testpmd: fix txonly flow generation entropy
- app/testpmd: fix uninitialized members of MPLS
- app/testpmd: fix uninitialized members when setting PFC
- app/testpmd: rename function for detaching by devargs
- app/testpmd: update Rx offload after setting MTU
- app/test: remove meson dependency on file in /sys
- bpf: fix headers install with meson
- build: explicitly enable sse4 for meson
- build: fix libm detection in meson
- build: remove unneeded function versioning
- bus/fslmc: remove conflicting memory barrier macro
- cfgfile: fix symbols map
- ci: fix Travis config warnings
- ci: use meson 0.47.1
- common/cpt: check cipher and auth keys are set
- common/cpt: fix component for empty IOV buffer
- crypto/armv8: fix clang build
- crypto/ccp: fix queue alignment
- crypto/dpaa_sec: fix IOVA conversions
- crypto/octeontx2: add kmod dependency info
- devtools: add fixes flag to commit listing

- devtools: fix debug build test
- doc: add module EEPROM dump to mlx5 features
- doc: clarify memory write combining in mlx5 guide
- doc: fix build with python 3.8
- doc: fix devargs in OCTEON TX2 event device guide
- doc: fix igb_uio parameter in ntb guide
- doc: fix multi-producer enqueue figure in ring guide
- doc: fix naming of Mellanox devices
- doc: fix quiescent state description in RCU guide
- doc: fix typos in 19.11 release notes
- doc: fix warning with meson
- doc: reduce indentation in meson build file
- doc: reduce whitespace in meson build file
- doc: update recommended versions for i40e
- drivers/crypto: fix session-less mode
- eal/linux: fix build error on RHEL 7.6
- eal/linux: fix build when VFIO is disabled
- eal/linux: fix uninitialized data valgrind warning
- eal/windows: fix cpuset macro name
- ethdev: fix callback unregister with wildcard argument list
- ethdev: fix flow API doxygen comment
- ethdev: fix secondary process memory overwrite
- ethdev: fix switching domain allocation
- ethdev: fix VLAN offloads set if no driver callback
- event/dpaa2: set number of order sequences
- event/dsw: avoid credit leak on oversized enqueue bursts
- event/dsw: flush buffers immediately on zero-sized enqueue
- event/octeontx2: fix device name in device info
- examples/ethtool: fix unchecked return value
- examples/fips_validation: fix AES-GCM cipher length parsing
- examples/fips_validation: fix cipher length for AES-GCM
- examples/fips_validation: fix string token for CT length
- examples/ioat: fix failure check for ioat dequeue
- examples/ioat: fix invalid link status check

- examples/ioat: fix unchecked return value
- examples/ipsec-secgw: extend inline session to non AES-GCM
- examples/ipsec-secgw: fix crash on unsupported algo
- examples/l2fwd-event: fix core allocation in poll mode
- examples/l2fwd-event: fix error checking
- examples/l2fwd-event: fix ethdev RSS setup
- examples/l2fwd-event: fix event device config
- examples/l3fwd-power: fix a typo
- examples/l3fwd-power: fix interrupt disable
- examples/ntb: fix mempool ops setting
- examples/power: fix ack for enable/disable turbo
- examples/tep_term: remove redundant info get
- examples/vhost_blk: check unused value on init
- examples/vhost_blk: fix check of device path
- fib: fix possible integer overflow
- fix Mellanox copyright and SPDX tag
- hash: fix lock-free flag doxygen
- hash: fix meson headers packaging
- kni: fix build with Linux 5.6
- kni: fix meson warning about console keyword
- kni: fix not contiguous FIFO
- kni: rename variable with namespace prefix
- latency: fix calculation for multi-thread
- lib: fix unnecessary double negation
- maintainers: resign from flow API maintenance
- maintainers: update for failsafe and PCI library
- mem: fix munmap in error unwind
- mempool: fix anonymous populate
- mempool: fix populate with small virtual chunks
- mempool: fix slow allocation of large pools
- mempool/octeonx: fix error handling in initialization
- mk: avoid combining -r and -export-dynamic linker options
- net/af_xdp: fix fill queue addresses
- net/af_xdp: fix maximum MTU

- net/af_xdp: fix redundant check for wakeup need
- net/af_xdp: fix umem frame size and headroom
- net/bnx2x: fix reset of scan FP flag
- net/bnx2x: fix to sync fastpath Rx queue access
- net/bnx2x: fix VLAN stripped flag
- net/bnx2x: support secondary process
- net/bnxt: add a field for FW capabilities
- net/bnxt: allow group ID 0 for RSS action
- net/bnxt: do not log error if stats queried before start
- net/bnxt: fix alloc filter to use a common routine
- net/bnxt: fix buffer allocation reattempt
- net/bnxt: fix bumping of L2 filter reference count
- net/bnxt: fix crash in port stop while handling events
- net/bnxt: fix default timeout for getting FW version
- net/bnxt: fix enable/disable VLAN filtering
- net/bnxt: fix flow creation
- net/bnxt: fix flow flush to sync with flow destroy
- net/bnxt: fix IOVA mapping
- net/bnxt: fix link during port toggle
- net/bnxt: fix MAC address setting when port is stopped
- net/bnxt: fix max rings calculation
- net/bnxt: fix non matching flow hitting filter rule
- net/bnxt: fix overwriting error message
- net/bnxt: fix port stop on error recovery failure
- net/bnxt: fix probe in FreeBSD
- net/bnxt: fix race condition when port is stopped
- net/bnxt: fix recovery alarm race condition in port close
- net/bnxt: fix request for hot reset support
- net/bnxt: fix return code handling in VLAN config
- net/bnxt: fix reusing L2 filter
- net/bnxt: fix Tx queue profile selection
- net/bnxt: fix unnecessary delay in port stop
- net/bnxt: fix VLAN strip
- net/bnxt: fix VLAN strip flags in SSE Rx

- net/bnxt: handle HW filter setting when port is stopped
- net/bnxt: remove a redundant variable
- net/bnxt: remove redundant if statement
- net/bnxt: remove redundant macro
- net/bnxt: remove unnecessary memset
- net/bnxt: remove unnecessary structure variable
- net/bnxt: restore MAC filters during reset recovery
- net/bnxt: restore VLAN filters during reset recovery
- net/bnxt: use macro for PCI log format
- net/cxgbe: announce Tx multi-segments offload
- net/dpaa: fix Rx offload flags on jumbo MTU set
- net/failsafe: fix reported hash key size in device info
- net/fm10k: fix descriptor VLAN field filling in Tx
- net/fm10k: fix non-x86 build
- net/hns3: fix crash when closing port
- net/hns3: fix dumping VF register information
- net/hns3: fix link status on failed query
- net/hns3: fix ring vector related mailbox command format
- net/hns3: fix Rx queue search with broadcast packet
- net/hns3: fix triggering reset procedure in slave process
- net/i40e/base: add new link speed constants
- net/i40e/base: fix buffer address
- net/i40e/base: fix display of FEC settings
- net/i40e/base: fix error message
- net/i40e/base: fix missing link modes
- net/i40e/base: fix retrying logic
- net/i40e/base: fix Tx descriptors number
- net/i40e: fix port close in FreeBSD
- net/i40e: fix Tx when TSO is enabled
- net/i40e: fix unchecked Tx cleanup error
- net/i40e: set fixed flag for exact link speed
- net/iavf: add TSO offload use basic path
- net/iavf/base: fix adminq return
- net/iavf/base: fix command buffer memory leak

- net/iavf: fix Rx total stats
- net/iavf: fix virtual channel return
- net/ice: add outer IPv4 matching for GTP-U flow
- net/ice/base: fix loop limit
- net/ice/base: increase PF reset wait timeout
- net/ice: disable TSO offload in vector path
- net/ice: fix flow director flag
- net/ice: fix flow director GTP-U pattern
- net/ice: fix flow director passthru
- net/ice: fix flow FDIR/switch memory leak
- net/ice: fix GTP-U rule conflict
- net/ice: fix packet type table
- net/ice: fix queue MSI-X interrupt binding
- net/ice: fix Tx when TSO is enabled
- net/ice: fix unchecked Tx cleanup error
- net/ice: fix VSI context
- net/ice: use ethernet copy API to do MAC assignment
- net/ipn3ke: fix line side statistics register read
- net/ipn3ke: fix meson build
- net/ixgbe: check for illegal Tx packets
- net/ixgbe: fix blocking system events
- net/ixgbe: fix flow control mode setting
- net/ixgbe: fix link status
- net/ixgbe: fix link up in FreeBSD
- net/ixgbe: remove dead code
- net/ixgbe: remove duplicate function declaration
- net/ixgbe: set fixed flag for exact link speed
- net/mlx5: add free on completion queue
- net/mlx5: allow push VLAN without VID
- net/mlx5: block pop VLAN action on Tx
- net/mlx5: block push VLAN action on Rx
- net/mlx5: clean up redundant assignment
- net/mlx5: engage free on completion queue
- net/mlx5: fix bit mask to validate push VLAN

- net/mlx5: fix blocker for push VLAN on Rx
- net/mlx5: fix build with clang 3.4.2
- net/mlx5: fix check for VLAN actions
- net/mlx5: fix crash when meter action conf is null
- net/mlx5: fix crash when setting hairpin queues
- net/mlx5: fix dirty array of actions
- net/mlx5: fix doorbell register offset type
- net/mlx5: fix encap/decap validation
- net/mlx5: fix flow match on GRE key
- net/mlx5: fix GENEVE tunnel flow validation
- net/mlx5: fix hairpin queue capacity
- net/mlx5: fix ICMPv6 header rewrite actions
- net/mlx5: fix ICMPv6 header rewrite action validation
- net/mlx5: fix inline packet size for ConnectX-4 Lx
- net/mlx5: fix item flag on GENEVE item validation
- net/mlx5: fix L3 VXLAN RSS expansion
- net/mlx5: fix layer flags missing in metadata
- net/mlx5: fix layer type in header modify action
- net/mlx5: fix layer validation with decapsulation
- net/mlx5: fix legacy multi-packet write session
- net/mlx5: fix masks of encap and decap actions
- net/mlx5: fix matcher field usage for metadata entities
- net/mlx5: fix match information in meter
- net/mlx5: fix matching for ICMP fragments
- net/mlx5: fix match on ethertype and CVLAN tag
- net/mlx5: fix memory regions release deadlock
- net/mlx5: fix metadata item endianness conversion
- net/mlx5: fix metadata split with encap action
- net/mlx5: fix meter header modify before decap
- net/mlx5: fix meter suffix flow
- net/mlx5: fix modify actions support limitation
- net/mlx5: fix multiple flow table hash list
- net/mlx5: fix pop VLAN action validation
- net/mlx5: fix register usage in meter

- net/mlx5: fix running without Rx queue
- net/mlx5: fix setting of port ID for egress rules
- net/mlx5: fix setting of Rx hash fields
- net/mlx5: fix shared metadata matcher field setup
- net/mlx5: fix tunnel flow priority
- net/mlx5: fix Tx burst routines set
- net/mlx5: fix VLAN actions in meter
- net/mlx5: fix VLAN ID action offset
- net/mlx5: fix VLAN match for DV mode
- net/mlx5: fix VLAN VID action validation
- net/mlx5: fix VXLAN-GPE item translation
- net/mlx5: fix zero out UDP checksum in encaps data
- net/mlx5: make FDB default rule optional
- net/mlx5: move Tx complete request routine
- net/mlx5: optimize Rx hash fields conversion
- net/mlx5: support maximum flow id allocation
- net/mlx5: unify validation of drop action
- net/mlx5: update description of validation functions
- net/mlx5: update Tx error handling routine
- net/mlx: add static ibverbs linkage with meson
- net/mlx: fix build with clang 9
- net/mlx: fix overlinking with meson and glue dlopen
- net/mlx: rename meson variable for dlopen option
- net/mlx: workaround static linkage with meson
- net/netvsc: disable before changing RSS parameters
- net/netvsc: fix crash in secondary process
- net/netvsc: fix RSS offload flag
- net/netvsc: initialize link state
- net/octeontx2: fix flow control initial state
- net/octeontx2: fix getting supported packet types
- net/octeontx2: fix PTP
- net/octeontx2: fix PTP and HIGIG2 coexistence
- net/octeontx2: fix Tx flow control for HIGIG
- net/octeontx2: fix VF configuration

- net/octeontx: fix memory leak of MAC address table
- net/qede/base: fix number of ports per engine
- net/qede: do not stop vport if not started
- net/qede: fix VF reload
- net/sfc: fix log format specifiers
- net/tap: fix memory leak when unregister intr handler
- net/vhost: allocate interface name from heap
- net/vhost: check creation failure
- net/vhost: delay driver setup
- net/vhost: fix probing in secondary process
- net/vhost: fix setup error path
- net/vhost: prevent multiple setups on reconfiguration
- net/virtio-user: check file descriptor before closing
- net/virtio-user: check tap offload setting failure
- net/virtio-user: do not close tap when disabling queue pairs
- net/virtio-user: do not reset virtqueues for split ring
- net/virtio-user: fix packed ring server mode
- raw/ntb: fix write memory barrier
- service: don't walk out of bounds when checking services
- test/common: fix log2 check
- test/compress: replace test vector
- test/crypto: fix missing operation status check
- test/event: fix OCTEON TX2 event device name
- test/event: fix unintended vdev creation
- test: fix build without ring PMD
- test/ipsec: fix a typo in function name
- usertools: fix syntax warning in python 3.8
- usertools: fix telemetry client with python 3
- vfiio: fix mapping failures in ppc64le
- vhost: catch overflow causing mmap of size 0
- vhost: check message header size read
- vhost/crypto: fix fetch size
- vhost: do not treat empty socket message as error
- vhost: fix crash on port deletion

- vhost: fix deadlock on port deletion
- vhost: fix inflight resubmit check
- vhost: fix packed virtqueue ready condition
- vhost: fix socket initial value
- vhost: flush shadow Tx if no more packets
- vhost: protect log address translation in IOTLB update

1.8.2 19.11.1 Validation

- Red Hat(R) Testing
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 4.2
 - * X540-AT2 NIC(ixgbe, 10G)
 - Functionality
 - * Guest with device assignment(PF) throughput testing(1G hugepage size)
 - * Guest with device assignment(PF) throughput testing(2M hugepage size)
 - * Guest with device assignment(VF) throughput testing
 - * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing
 - * PVP vhost-user 2Q throughput testing
 - * PVP vhost-user 1Q - cross numa node throughput testing
 - * Guest with vhost-user 2 queues throughput testing
 - * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect
 - * PVP 1Q live migration testing
 - * PVP 1Q cross numa node live migration testing
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M)
 - * Guest with ovs+dpdk+vhost-user 2Q live migration testing
- Intel(R) Testing
 - **Basic Intel(R) NIC(ixgbe, i40e and ice) testing**
 - * PF (i40e)
 - * PF (ixgbe)
 - * PF (ice)
 - * 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
 - * l3fwd-power example application
 - * Multi-process example applications
 - ConnectX-5
 - * RHEL 7.4
 - * Kernel 3.10.0-693.el7.x86_64
 - * Driver MLNX_OFED_LINUX-5.0-1.0.0.0
 - * fw 16.27.1016
 - ConnectX-4 Lx
 - * RHEL 7.4
 - * Kernel 3.10.0-693.el7.x86_64
 - * Driver MLNX_OFED_LINUX-5.0-1.0.0.0
 - * fw 14.27.1016
- Broadcom(R) Testing
 - Functionality
 - * Tx/Rx
 - * Link status
 - * RSS
 - * Checksum/TSO
 - * VLAN filtering

- * statistics
- * MTU
- Platform
 - * BCM57414 NetXtreme-E 10Gb/25Gb Ethernet Controller, Firmware: 216.1.169.0
 - * BCM57508 NetXtreme-E 10Gb/25Gb/40Gb/50Gb/100Gb/200Gb Ethernet, Firmware : 216.0.314.0
- IBM(R) Testing
 - Functionality
 - * Basic PF on Mellanox
 - * Single port stability test using 13fwd (16 cpus) and TRex, tested 64 and 1500 byte packets at a 0.0% drop rate for 4 hours each
 - * Performance: no degradation compared to 19.11.0
 - Platform
 - * Ubuntu 18.04.4 LTS
 - * Kernel 4.15.0-88-generic
 - * IBM Power9 Model 8335-101 CPU: 2.3 (pvr 004e 1203)
 - * Mellanox Technologies MT28800 Family [ConnectX-5 Ex], firmware version: 16.26.4012, MLNX_OFED_LINUX-4.7-3.2.9.1

1.9 19.11.2 Release Notes

1.9.1 19.11.2 Fixes

- 2cf9c470eb vhost: check log mmap offset and size overflow (CVE-2020-10722)
- 8e9652b0b6 vhost: fix translated address not checked (CVE-2020-10723)
- 95e1f29c26 vhost/crypto: validate keys lengths (CVE-2020-10724)
- 963b6eea05 vhost: fix potential memory space leak (CVE-2020-10725)
- c9c630a117 vhost: fix potential fd leak (CVE-2020-10726)
- cd0ea71bb6 vhost: fix vring index check (CVE-2020-10726)

1.9.2 19.11.2 Validation

- Red Hat(R) Testing
 - Platform
 - * RHEL 8.3
 - * Kernel 4.18
 - * Qemu 4.2
 - * X540-AT2 NIC(ixgbe, 10G)
 - Functionality

- * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing
- * PVP vhost-user 2Q throughput testing
- * PVP vhost-user 1Q - cross numa node throughput testing
- * PVP 1Q live migration testing
- * PVP 1Q cross numa node live migration testing
- Intel(R) Testing
 - Virtio features
 - * vhost/virtio loopback test with virtio user as server mode
 - * loopback multi queues
 - * loopback multi paths port restart
 - * vhost/virtio pvp multi-paths performance
 - * pvp multi-queues and port restart
 - * vhost dequeue zero copy
 - * pvp share lib
 - * pvp vhost user reconnect
 - * pvp test with 4k pages
 - * pvp test with 2M hugepages
 - * pvp virtio bonding
 - * pvp test with diff qemu version
 - * vhost enqueue interrupt
 - * vhost event idx interrupt
 - * vhost virtio pmd interrupt
 - * vhost virtio user interrupt
 - * virtio event idx interrupt
 - * virtio user for container networking
 - * virtio user as exceptional path
 - * vhost xstats
 - * virtio-pmd multi-process
 - * vm2vm virtio pmd
 - * vm2vm virtio-net iperf
 - * vm2vm virtio-user
 - * vhost user live migration

1.10 19.11.3 Release Notes

1.10.1 19.11.3 Fixes

- app/crypto-perf: fix display of sample test vector
- app/eventdev: check Tx adapter service ID
- app: fix usage help of options separated by dashes
- app/pipeline: fix build with gcc 10
- app: remove extra new line after link duplex
- app/testpmd: add parsing for QinQ VLAN headers
- app/testpmd: fix DCB set
- app/testpmd: fix memory failure handling for i40e DDP
- app/testpmd: fix PPPoE flow command
- app/testpmd: fix statistics after reset
- baseband/turbo_sw: fix exposed LLR decimals assumption
- bbdev: fix doxygen comments
- build: disable gcc 10 zero-length-bounds warning
- build: fix linker warnings with clang on Windows
- build: support MinGW-w64 with Meson
- buildtools: get static mlx dependencies for meson
- bus/fslmc: fix dereferencing null pointer
- bus/fslmc: fix size of qman fq descriptor
- bus/pci: fix devargs on probing again
- bus/pci: fix UIO resource access from secondary process
- bus/vmbus: fix comment spelling
- ci: fix telemetry dependency in Travis
- common/iavf: update copyright
- common/mlx5: fix build with -fno-common
- common/mlx5: fix build with rdma-core 21
- common/mlx5: fix netlink buffer allocation from stack
- common/mlx5: fix umem buffer alignment
- common/octeontx: fix gcc 9.1 ABI break
- common/qat: fix GEN3 marketing name
- contigmem: cleanup properly when load fails
- crypto/caam_jr: fix check of file descriptors
- crypto/caam_jr: fix IRQ functions return type

- crypto/ccp: fix fd leak on probe failure
- cryptodev: add asymmetric session-less feature name
- cryptodev: fix missing device id range checking
- cryptodev: fix SHA-1 digest enum comment
- crypto/kasumi: fix extern declaration
- crypto/nitrox: fix CSR register address generation
- crypto/nitrox: fix oversized device name
- crypto/octeontx2: fix build with gcc 10
- crypto/openssl: fix out-of-place encryption
- crypto/qat: fix cipher descriptor for ZUC and SNOW
- crypto/qat: support plain SHA1..SHA512 hashes
- devtools: fix symbol map change check
- doc: add i40e limitation for flow director
- doc: add NASM installation steps
- doc: fix API index
- doc: fix build issue in ABI guide
- doc: fix build with doxygen 1.8.18
- doc: fix default symbol binding in ABI guide
- doc: fix log level example in Linux guide
- doc: fix LTO config option
- doc: fix matrix CSS for recent sphinx
- doc: fix multicast filter feature announcement
- doc: fix number of failsafe sub-devices
- doc: fix reference in ABI guide
- doc: fix sphinx compatibility
- doc: fix typo in contributors guide
- doc: fix typo in contributors guide
- doc: fix typos in ABI policy
- doc: prefer https when pointing to dpdk.org
- drivers: add crypto as dependency for event drivers
- drivers/crypto: disable gcc 10 no-common errors
- drivers/crypto: fix build with make 4.3
- drivers/crypto: fix log type variables for -fno-common
- drivers: fix log type variables for -fno-common

- eal/arm64: fix precise TSC
- eal: fix C++17 compilation
- eal: fix comments spelling
- eal: fix log message print for regex
- eal: fix PRNG init with HPET enabled
- eal: fix typo in endian conversion macros
- eal/freebsd: fix queuing duplicate alarm callbacks
- eal/ppc: fix bool type after altivec include
- eal/ppc: fix build with gcc 9.3
- eal/x86: ignore gcc 10 stringop-overflow warnings
- ethdev: fix build when vtune profiling is on
- ethdev: fix spelling
- eventdev: fix probe and remove for secondary process
- event/dsw: avoid reusing previously recorded events
- event/dsw: fix enqueue burst return value
- event/dsw: remove redundant control ring poll
- event/dsw: remove unnecessary read barrier
- event/octeonx2: fix build for O1 optimization
- event/octeonx2: fix queue removal from Rx adapter
- examples/eventdev: fix build with gcc 10
- examples/eventdev: fix crash on exit
- examples/fips_validation: fix parsing of algorithms
- examples/ip_pipeline: remove check of null response
- examples/ipsec-gw: fix gcc 10 maybe-uninitialized warning
- examples/kni: fix crash during MTU set
- examples/kni: fix MTU change to setup Tx queue
- examples/l2fwd-keepalive: fix mbuf pool size
- examples/qos_sched: fix build with gcc 10
- examples: remove extra new line after link duplex
- examples/vhost_blk: fix build with gcc 10
- examples/vmdq: fix output of pools/queues
- examples/vmdq: fix RSS configuration
- examples/vm_power: drop Unix path limit redefinition
- examples/vm_power: fix build with -fno-common

- fib: fix headers for C++ support
- fix same typo in multiple places
- fix various typos found by Lintian
- ipsec: check SAD lookup error
- ipsec: fix build dependency on hash lib
- kvargs: fix buffer overflow when parsing list
- kvargs: fix invalid token parsing on FreeBSD
- kvargs: fix strcmp helper documentation
- log: fix level picked with globbing on type register
- lpm6: fix comments spelling
- lpm6: fix size of tbl8 group
- mem: fix overflow on allocation
- mem: mark pages as not accessed when freeing memory
- mem: mark pages as not accessed when reserving VA
- mempool/dpaa2: install missing header with meson
- mempool/octeonx2: fix build for gcc O1 optimization
- mempool: remove inline functions from export list
- mem: preallocate VA space in no-huge mode
- mk: fix static linkage of mlx dependency
- net/avp: fix gcc 10 maybe-uninitialized warning
- net/bnxt: do not use PMD log type
- net/bnxt: fix error log for command timeout
- net/bnxt: fix FW version query
- net/bnxt: fix HWRM command during FW reset
- net/bnxt: fix max ring count
- net/bnxt: fix memory leak during queue restart
- net/bnxt: fix number of TQM ring
- net/bnxt: fix port start failure handling
- net/bnxt: fix possible stack smashing
- net/bnxt: fix Rx ring producer index
- net/bnxt: fix storing MAC address twice
- net/bnxt: fix TQM ring context memory size
- net/bnxt: fix using RSS config struct
- net/bnxt: fix VLAN add when port is stopped

- net/bnxt: fix VNIC Rx queue count on VNIC free
- net/bnxt: use true/false for bool types
- net/dpaa2: fix 10G port negotiation
- net/dpaa2: fix congestion ID for multiple traffic classes
- net/dpaa: use dynamic log type
- net/e1000: fix port hotplug for multi-process
- net/ena/base: fix documentation of functions
- net/ena/base: fix indentation in CQ polling
- net/ena/base: fix indentation of multiple defines
- net/ena/base: fix testing for supported hash function
- net/ena/base: make allocation macros thread-safe
- net/ena/base: prevent allocation of zero sized memory
- net/ena: fix build for O1 optimization
- net/ena: set IO ring size to valid value
- net/enetc: fix Rx lock-up
- net/enic: fix flow action reordering
- net/failsafe: fix fd leak
- net/hinic: allocate IO memory with socket id
- net/hinic/base: fix PF firmware hot-active problem
- net/hinic/base: fix port start during FW hot update
- net/hinic: fix LRO
- net/hinic: fix queues resource free
- net/hinic: fix repeating cable log and length check
- net/hinic: fix snprintf length of cable info
- net/hinic: fix TSO
- net/hinic: fix Tx mbuf length while copying
- net/hns3: add free threshold in Rx
- net/hns3: add RSS hash offload to capabilities
- net/hns3: clear residual flow rules on init
- net/hns3: fix configuring illegal VLAN PVID
- net/hns3: fix configuring RSS hash when rules are flushed
- net/hns3: fix crash when flushing RSS flow rules with FLR
- net/hns3: fix default error code of command interface
- net/hns3: fix default VLAN filter configuration for PF

- net/hns3: fix mailbox opcode data type
- net/hns3: fix MSI-X interrupt during initialization
- net/hns3: fix packets offload features flags in Rx
- net/hns3: fix promiscuous mode for PF
- net/hns3: fix return value of setting VLAN offload
- net/hns3: fix return value when clearing statistics
- net/hns3: fix RSS indirection table configuration
- net/hns3: fix RSS key length
- net/hns3: fix Rx interrupt after reset
- net/hns3: fix status after repeated resets
- net/hns3: fix Tx interrupt when enabling Rx interrupt
- net/hns3: fix VLAN filter when setting promiscuous mode
- net/hns3: fix VLAN PVID when configuring device
- net/hns3: reduce judgements of free Tx ring space
- net/hns3: remove one IO barrier in Rx
- net/hns3: remove unnecessary assignments in Tx
- net/hns3: replace memory barrier with data dependency order
- net/hns3: support different numbers of Rx and Tx queues
- net/hns3: support Rx interrupt
- net/i40e/base: update copyright
- net/i40e: fix flow director enabling
- net/i40e: fix flow director for ARP packets
- net/i40e: fix flow director initialisation
- net/i40e: fix flush of flow director filter
- net/i40e: fix queue region in RSS flow
- net/i40e: fix queue related exception handling
- net/i40e: fix setting L2TAG
- net/i40e: fix wild pointer
- net/i40e: fix X722 performance
- net/i40e: relax barrier in Tx
- net/i40e: relax barrier in Tx for NEON
- net/iavf: fix link speed
- net/iavf: fix setting L2TAG
- net/iavf: fix stats query error code

- net/ice: add action number check for switch
- net/ice/base: check memory pointer before copying
- net/ice/base: fix binary order for GTPU filter
- net/ice/base: fix MAC write command
- net/ice/base: fix uninitialized stack variables
- net/ice/base: minor fixes
- net/ice/base: read PSM clock frequency from register
- net/ice/base: remove unused code in switch rule
- net/ice/base: update copyright
- net/ice: change default tunnel type
- net/ice: fix crash in switch filter
- net/ice: fix hash flow crash
- net/ice: fix input set of VLAN item
- net/ice: fix RSS advanced rule
- net/ice: fix RSS for GTPU
- net/ice: fix setting L2TAG
- net/ice: fix variable initialization
- net/ice: remove bulk alloc option
- net/ice: remove unnecessary variable
- net/ice: support mark only action for flow director
- net/ipn3ke: use control thread to check link status
- net/ixgbe/base: update copyright
- net/ixgbe: check driver type in MACsec API
- net/ixgbe: fix link state timing on fiber ports
- net/ixgbe: fix link status after port reset
- net/ixgbe: fix link status inconsistencies
- net/ixgbe: fix link status synchronization on BSD
- net/ixgbe: fix resource leak after thread exits normally
- net/ixgbe: fix statistics in flow control mode
- net/memif: fix init when already connected
- net/memif: fix resource leak
- net/mlx4: fix build with -fno-common
- net/mlx4: fix drop queue error handling
- net/mlx5: add device parameter for MPRQ stride size

- net/mlx5: add multi-segment packets in MPRQ mode
- net/mlx5: enable MPRQ multi-stride operations
- net/mlx5: fix actions validation on root table
- net/mlx5: fix assert in doorbell lookup
- net/mlx5: fix assert in dynamic metadata handling
- net/mlx5: fix assert in modify converting
- net/mlx5: fix build with separate glue lib for dlopen
- net/mlx5: fix call to modify action without init item
- net/mlx5: fix counter container usage
- net/mlx5: fix crash when releasing meter table
- net/mlx5: fix CVLAN tag set in IP item translation
- net/mlx5: fix doorbell bitmap management offsets
- net/mlx5: fix gcc 10 enum-conversion warning
- net/mlx5: fix header modify action validation
- net/mlx5: fix imissed counter overflow
- net/mlx5: fix jump table leak
- net/mlx5: fix mask used for IPv6 item validation
- net/mlx5: fix matching for UDP tunnels with Verbs
- net/mlx5: fix match on empty VLAN item in DV mode
- net/mlx5: fix metadata for compressed Rx CQEs
- net/mlx5: fix meter color register consideration
- net/mlx5: fix meter suffix table leak
- net/mlx5: fix packet length assert in MPRQ
- net/mlx5: fix push VLAN action to use item info
- net/mlx5: fix RSS enablement
- net/mlx5: fix RSS key copy to TIR context
- net/mlx5: fix Tx queue release debug log timing
- net/mlx5: fix validation of push VLAN without full mask
- net/mlx5: fix validation of VXLAN/VXLAN-GPE specs
- net/mlx5: fix VLAN flow action with wildcard VLAN item
- net/mlx5: fix VLAN ID check
- net/mlx5: fix VLAN PCP item calculation
- net/mlx5: fix zero metadata action
- net/mlx5: fix zero value validation for metadata

- net/mlx5: improve logging of MPRQ selection
- net/mlx5: reduce Tx completion index memory loads
- net/mlx5: set dynamic flow metadata in Rx queues
- net/mlx5: update VLAN and encap actions validation
- net/mlx5: use open/read/close for ib stats query
- net/mvneta: do not use PMD log type
- net/mvpp2: fix build with gcc 10
- net/netvsc: avoid possible live lock
- net/netvsc: do not configure RSS if disabled
- net/netvsc: do RSS across Rx queue only
- net/netvsc: fix comment spelling
- net/netvsc: fix memory free on device close
- net/netvsc: handle Rx packets during multi-channel setup
- net/netvsc: handle Tx completions based on burst size
- net/netvsc: propagate descriptor limits from VF
- net/netvsc: remove process event optimization
- net/netvsc: split send buffers from Tx descriptors
- net/nfp: fix dangling pointer on probe failure
- net/nfp: fix log format specifiers
- net/null: fix secondary burst function selection
- net/null: remove redundant check
- net/octeontx2: disable unnecessary error interrupts
- net/octeontx2: enable error and RAS interrupt in configure
- net/octeontx2: fix buffer size assignment
- net/octeontx2: fix device configuration sequence
- net/octeontx2: fix link information for loopback port
- net/octeontx: fix dangling pointer on init failure
- net/octeontx: fix meson build for disabled drivers
- net/pfe: do not use PMD log type
- net/pfe: fix double free of MAC address
- net/qede: fix link state configuration
- net/qede: fix port reconfiguration
- net/ring: fix device pointer on allocation
- net/sfc/base: fix build when EVB is enabled

- net/sfc/base: fix manual filter delete in EF10
- net/sfc/base: handle manual and auto filter clashes in EF10
- net/sfc/base: reduce filter priorities to implemented only
- net/sfc/base: refactor filter lookup loop in EF10
- net/sfc/base: reject automatic filter creation by users
- net/sfc/base: use simpler EF10 family conditional check
- net/sfc/base: use simpler EF10 family run-time checks
- net/sfc: fix initialization error path
- net/sfc: fix promiscuous and allmulticast toggles errors
- net/sfc: fix reported promiscuous/multicast mode
- net/sfc: fix Rx queue start failure path
- net/sfc: set priority of created filters to manual
- net/softnic: fix memory leak for thread
- net/softnic: fix resource leak for pipeline
- net/tap: do not use PMD log type
- net/tap: fix check for mbuf number of segment
- net/tap: fix crash in flow destroy
- net/tap: fix fd leak on creation failure
- net/tap: fix file close on remove
- net/tap: fix mbuf and mem leak during queue release
- net/tap: fix mbuf double free when writev fails
- net/tap: fix queues fd check before close
- net/tap: fix unexpected link handler
- net/tap: remove unused assert
- net/thunderx: use dynamic log type
- net/vhost: fix potential memory leak on close
- net/virtio: do not use PMD log type
- net/virtio: fix crash when device reconnecting
- net/virtio: fix outdated comment
- net/virtio: fix unexpected event after reconnect
- net/virtio-user: fix devargs parsing
- net/vmxnet3: fix RSS setting on v4
- net/vmxnet3: handle bad host framing
- pci: accept 32-bit domain numbers

- pci: fix build on FreeBSD
- pci: fix build on ppc
- pci: reject negative values in PCI id
- pci: remove unneeded includes in public header file
- remove references to private PCI probe function
- Revert “common/qat: fix GEN3 marketing name”
- Revert “net/bnxt: fix number of TQM ring”
- Revert “net/bnxt: fix TQM ring context memory size”
- security: fix crash at accessing non-implemented ops
- security: fix return types in documentation
- security: fix session counter
- security: fix verification of parameters
- service: fix crash on exit
- service: fix identification of service running on other lcore
- service: fix race condition for MT unsafe service
- service: remove rte prefix from static functions
- telemetry: fix port stats retrieval
- test/crypto: fix flag check
- test/crypto: fix statistics case
- test: fix build with gcc 10
- test/flow_classify: enable multi-sockets system
- test/ipsec: fix crash in session destroy
- test/kvargs: fix invalid cases check
- test/kvargs: fix to consider empty elements as valid
- test: load drivers when required
- test: remove redundant macro
- test: skip some subtests in no-huge mode
- timer: protect initialization with lock
- usertools: check for pci.ids in /usr/share/misc
- vfio: fix race condition with sysfs
- vfio: fix use after free with multiprocess
- vhost/crypto: add missing user protocol flag
- vhost: fix packed ring zero-copy
- vhost: fix peer close check

- vhost: fix shadowed descriptors not flushed
- vhost: fix shadow update
- vhost: fix zero-copy server mode
- vhost: handle mbuf allocation failure
- vhost: make IOTLB cache name unique among processes
- vhost: prevent zero-copy with incompatible client mode
- vhost: remove unused variable

1.10.2 19.11.3 Validation

- Intel(R) Testing
 - Basic Intel(R) NIC(ixgbe, i40e and ice) testing
 - * PF (i40e)
 - * PF (ixgbe)
 - * PF (ice)
 - * VF (i40e)
 - * VF (ixgbe)
 - * VF (ice)
 - * Compile Testing
 - * Intel NIC single core/NIC performance
 - Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance
- 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
 - * l3fwd-power example application
 - * Multi-process example applications

- ConnectX-5
 - * RHEL 7.4
 - * Driver MLNX_OFED_LINUX-5.0-2.1.8.0
 - * fw 16.27.2008
- ConnectX-4 Lx
 - * RHEL 7.4
 - * Driver MLNX_OFED_LINUX-5.0-2.1.8.0
 - * fw 14.27.1016
- Red Hat(R) Testing
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 4.2
 - * X540-AT2 NIC(ixgbe, 10G)
 - Functionality
 - * Guest with device assignment(PF) throughput testing(1G hugepage size)
 - * Guest with device assignment(PF) throughput testing(2M hugepage size)
 - * Guest with device assignment(VF) throughput testing
 - * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing
 - * PVP vhost-user 2Q throughput testing
 - * PVP vhost-user 1Q - cross numa node throughput testing
 - * Guest with vhost-user 2 queues throughput testing
 - * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect
 - * PVP 1Q live migration testing
 - * PVP 1Q cross numa node live migration testing
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M)
 - * Guest with ovs+dpdk+vhost-user 2Q live migration testing
 - * Allocate memory from the NUMA node which Virtio device locates
- Intel(R) Testing with Open vSwitch
 - OVS testing with OVS branches master and 2.13 with VSPERF
 - Tested NICs
 - * i40e (X710)
 - * ixgbe (82599ES)

- * ice
- * vhost
- Functionality
 - * P2P
 - * PVP
 - * PVPV
 - * PVVP
 - * Multiqueue RSS
 - * Vhost reconnect
 - * Jumbo frames 1500, 6000, 9702
- Microsoft(R) Testing
 - Platform
 - * **Azure**
 - Ubuntu 16.04-LTS
 - Ubuntu 18.04-DAILY-LTS
 - RHEL 7-RAW
 - RHEL 7.5
 - CentOS 7.5
 - SLES-15-sp1 gen1
 - * Mellanox(R) ConnectX-4
 - * LISAv2 test framework
 - Functionality
 - * VERIFY-DPDK-COMPLIANCE - verifies kernel is supported and that the build is successful
 - * VERIFY-DPDK-BUILD-AND-TESTPMD-TEST - verifies using testpmd that packets can be sent from a VM to another VM
 - * VERIFY-SRIOV-FAILSAFE-FOR-DPDK - disables/enables Accelerated Networking for the NICs under test and makes sure DPDK works in both scenarios
 - * VERIFY-DPDK-FAILSAFE-DURING-TRAFFIC - disables/enables Accelerated Networking for the NICs while generating traffic using testpmd
 - * PERF-DPDK-FWD-PPS-DS15 - verifies DPDK forwarding performance using testpmd on 2, 4, 8 cores, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-SINGLE-CORE-PPS-DS4 - verifies DPDK performance using testpmd on 1 core, rx and io mode on size Standard_DS4_v2
 - * PERF-DPDK-SINGLE-CORE-PPS-DS15 - verifies DPDK performance using testpmd on 1 core, rx and io mode on size Standard_DS15_v2

- * PERF-DPDK-MULTICORE-PPS-DS15 - verifies DPDK performance using testpmd on 2, 4, 8 cores, rx and io mode on size Standard_DS15_v2
- * PERF-DPDK-MULTICORE-PPS-F32 - verifies DPDK performance using testpmd on 2, 4, 8, 16 cores, rx and io mode on size Standard_F32s_v2
- * DPDK-RING-LATENCY - verifies DPDK CPU latency using dpdk-ring-ping
- * VERIFY-DPDK-PRIMARY-SECONDARY-PROCESSES - verifies primary / secondary processes support for DPDK. Runs only on RHEL and Ubuntu distros with Linux kernel ≥ 4.20
- * VERIFY-DPDK-OVS - builds OVS with DPDK support and tests if the OVS DPDK ports can be created. Runs only on Ubuntu distro.

1.11 19.11.4 Release Notes

1.11.1 19.11.4 Fixes

- app/eventdev: fix capability check in pipeline ATQ test
- app/testpmd: fix burst percentage calculation
- app/testpmd: fix CPU cycles per packet stats on Tx modes
- app/testpmd: fix error detection in MTU command
- app/testpmd: fix memory leak on error path
- app/testpmd: fix stats error message
- app/testpmd: remove hardcoded descriptors limit
- app/testpmd: use clock time in throughput calculation
- avoid libfdt checks adding full paths to pkg-config
- bpf: fix add/sub min/max estimations
- build: fix drivers library path on Windows
- bus/dpaa: fix iterating on a class type
- bus/fslmc: fix getting FD error
- bus/fslmc: fix iterating on a class type
- bus/fslmc: fix memory leak in secondary process
- bus/pci: fix VF memory access
- bus/vdev: fix a typo in doxygen comment
- bus/vmbus: fix ring buffer mapping
- cfgfile: fix stack buffer underflow
- common/cpt: fix encryption offset
- common/dpaax: fix 12-bit null auth case
- common/mlx5: fix code arrangement in tag allocation
- common/mlx5: fix queue doorbell record size

- common/mlx5: fix void parameters in glue wrappers
- common/octeontx2: fix crash on running procinfo
- common/qat: fix uninitialized variable
- common/qat: get firmware version
- common/qat: move max inflights param into qp
- common/qat: remove tail write coalescing
- common/qat: support dual threads for enqueue/dequeue
- crypto/armv8: remove debug option
- crypto/armv8: use dedicated log type
- crypto/dpaa2_sec: fix HFN override
- crypto/dpaa2_sec: fix 18-bit PDCP cases with HFN override
- crypto/dpaa2_sec: fix inline query for descriptors
- crypto/qat: add minimum enq threshold
- crypto/qat: fix AES-XTS capabilities
- crypto/qat: handle mixed hash-cipher on GEN2
- crypto/qat: handle mixed hash-cipher requests on GEN3
- devtools: fix path in forbidden token check
- doc: add RIB and FIB into the API index
- doc: fix a typo in mlx5 guide
- doc: fix doc build after qat threshold patch
- doc: fix ethtool app path
- doc: fix reference to master process
- doc: fix some typos in Linux guide
- doc: fix typo in bbdev test guide
- doc: rebuild with meson whenever a file changes
- doc: update build instructions in the Linux guide
- drivers/crypto: add missing OOP feature flag
- drivers/net: fix exposing internal headers
- drivers/qat: add handling of capabilities in multi process
- drivers/qat: add multi process handling of driver id
- drivers/qat: improve multi process on qat
- eal/arm: add vcopyq intrinsic for aarch32
- eal/armv8: fix timer frequency calibration with PMU
- eal: fix lcore accessors for non-EAL threads

- eal: fix parentheses in alignment macros
- eal: fix uuid header dependencies
- eal/linux: fix epoll fd list rebuild for interrupts
- eal: remove redundant newline in alert message
- eal/windows: fix symbol export
- ethdev: fix data room size verification in Rx queue setup
- ethdev: fix log type for some error messages
- ethdev: fix VLAN offloads set if no relative capabilities
- eventdev: fix race condition on timer list counter
- eventdev: relax SMP barriers with C11 atomics
- eventdev: remove redundant reset on timer cancel
- eventdev: use C11 atomics for lcore timer armed flag
- event/dpaa2: add all-types queue capability flag
- event/dpaa: remove dead code
- event/octeontx2: fix device reconfigure
- event/octeontx2: fix sub event type
- examples: add flush after stats printing
- examples/eventdev: fix 32-bit coremask
- examples/fips_validation: fix count overwrite for TDES
- examples/fips_validation: fix parsing of TDES vectors
- examples/fips_validation: fix TDES interim callback
- examples/packet_ordering: use proper exit method
- hash: fix out-of-memory handling in hash creation
- kni: fix reference to master/slave process
- lib: remind experimental status in headers
- mbuf: fix boundary check at dynamic field registration
- mbuf: fix dynamic field dump log
- mbuf: fix error code in dynamic field/flag registration
- mbuf: fix free space update for dynamic field
- mbuf: remove unused next member in dynamic flag/field
- mem: fix 32-bit init config with meson
- mempool: fix allocation in memzone during retry
- meter: remove inline functions from export list
- net/af_packet: fix check of file descriptors

- net/af_packet: fix memory leak on init failure
- net/af_packet: fix munmap on init failure
- net/af_xdp: remove mempool freeing on umem destruction
- net/bnxt: fix flow error on filter creation
- net/bnxt: fix freeing filters on flow creation failure
- net/bnxt: fix logical AND in if condition
- net/bnxt: fix performance for Arm
- net/bnxt: fix unnecessary HWRM command
- net/bnxt: remove unused enum declaration
- net/bonding: change state machine to defaulted
- net/bonding: delete redundant code
- net/bonding: fix dead loop on RSS RETA update
- net/bonding: fix error code on device creation
- net/bonding: fix LACP negotiation
- net/bonding: fix MAC address when one port resets
- net/bonding: fix MAC address when switching active port
- net/bonding: fix socket ID check
- net/cxgbe: fix CLIP leak in filter error path
- net/cxgbe: fix double MPS alloc by flow validate and create
- net/cxgbe: fix L2T leak in filter error and free path
- net/dpaa: fix FD offset data type
- net/e1000: fix crash on Tx done clean up
- net/e1000: report VLAN extend capability
- net/failsafe: fix RSS RETA size info
- net: fix checksum on big endian CPUs
- net: fix IPv4 checksum
- net: fix pedantic build
- net: fix unneeded replacement of TCP checksum 0
- net/hinic/base: avoid system time jump
- net/hinic/base: check output of management sync channel
- net/hinic/base: remove unused function parameters
- net/hinic: check memory allocations in flow creation
- net/hinic: fix setting promiscuous mode
- net/hinic: optimize Rx performance for x86

- net/hns3: add RSS hash offload to Rx configuration
- net/hns3: check multi-process action register result
- net/hns3: clear promiscuous on PF uninit
- net/hns3: clear residual hardware configurations on init
- net/hns3: fix adding multicast MAC address
- net/hns3: fix flow director error message
- net/hns3: fix key length when configuring RSS
- net/hns3: fix RSS configuration on empty RSS type
- net/hns3: fix Rx buffer size
- net/hns3: fix Tx less than 60 bytes
- net/hns3: fix unintended sign extension in dump operation
- net/hns3: fix unintended sign extension in fd operation
- net/hns3: fix VLAN strip configuration when setting PVID
- net/hns3: fix VLAN tags reported in Rx
- net/hns3: get link status change through mailbox
- net/hns3: ignore function return on reset error path
- net/hns3: optimize default RSS algorithm
- net/hns3: remove restriction on setting VF MTU
- net/hns3: remove unnecessary branch
- net/hns3: remove unsupported VLAN capabilities
- net/i40e: enable NEON Rx/Tx in meson
- net/i40e: enable QinQ stripping
- net/i40e: fix binding interrupt without MSI-X vector
- net/i40e: fix filter pctype
- net/i40e: fix flow director MSI-X resource allocation
- net/i40e: fix flow director Rx writeback packet
- net/i40e: fix getting EEPROM information
- net/i40e: fix queue pairs configuration in VF
- net/i40e: remove duplicate tunnel type check
- net/i40e: report VLAN filter capability
- net/i40e: support aarch32
- net/iavf: fix RSS RETA after restart
- net/iavf: fix uninitialized variable
- net/ice: add input set byte number check

- net/ice: add memory allocation check in RSS init
- net/ice/base: fix GTP-U inner RSS IPv4 IPv6 co-exist
- net/ice/base: fix initializing resource for field vector
- net/ice/base: fix memory leak on error path
- net/ice/base: fix memory leak on GTPU RSS
- net/ice/base: fix reference count on VSI list update
- net/ice/base: fix return value
- net/ice/base: fix RSS interference
- net/ice/base: fix RSS removal for GTP-U
- net/ice/base: fix VSI ID mask to 10 bits
- net/ice: calculate TCP header size for offload
- net/ice: fix bytes statistics
- net/ice: fix error log in generic flow
- net/ice: fix memory leak when releasing VSI
- net/ice: fix switch action number check
- net/ice: fix TCP checksum offload
- net/ice: fix Tx hang with TSO
- net/ice: revert fake TSO fixes
- net/ixgbe/base: fix host interface shadow RAM read
- net/ixgbe/base: fix infinite recursion on PCIe link down
- net/ixgbe/base: fix x550em 10G NIC link status
- net/ixgbe/base: remove dead code
- net/ixgbe: fix flow control status
- net/ixgbe: fix include of vector header file
- net/ixgbe: fix MAC control frame forward
- net/ixgbe: report 10Mbps link speed for x553
- net/kni: set packet input port in Rx
- net/mlx4: optimize stack memory size in probe
- net/mlx5: do not select legacy MPW implicitly
- net/mlx5: fix counter query
- net/mlx5: fix crash in NVGRE item translation
- net/mlx5: fix descriptors number adjustment
- net/mlx5: fix flow items size calculation
- net/mlx5: fix flow META item validation

- net/mlx5: fix hairpin Rx queue creation error flow
- net/mlx5: fix hairpin Tx queue creation error flow
- net/mlx5: fix HW counters path in switchdev mode
- net/mlx5: fix initialization of steering registers
- net/mlx5: fix interrupt installation timing
- net/mlx5: fix iterator type in Rx queue management
- net/mlx5: fix LRO checksum
- net/mlx5: fix metadata storing for NEON Rx
- net/mlx5: fix secondary process resources release
- net/mlx5: fix tunnel flow priority
- net/mlx5: fix typos in meter error messages
- net/mlx5: fix UAR lock sharing for multiport devices
- net/mlx5: fix unnecessary init in mark conversion
- net/mlx5: fix unreachable MPLS error path
- net/mlx5: fix vectorized Rx burst termination
- net/mlx5: fix VF MAC address set over BlueField
- net/mlx5: fix VLAN pop with decap action validation
- net/mlx5: fix VLAN push action on hairpin queue
- net/mlx5: remove ineffective increment in hairpin split
- net/mlx5: remove needless Tx queue initialization check
- net/mlx5: remove redundant newline from logs
- net/mvpp2: fix non-EAL thread support
- net/netvsc: do not query VF link state
- net/netvsc: do not spin forever waiting for reply
- net/netvsc: fix chimney index
- net/netvsc: fix crash during Tx
- net/netvsc: fix underflow when Rx external mbuf
- net/netvsc: fix warning when VF is removed
- net/nfp: fix RSS hash configuration reporting
- net/octeontx2: fix DMAC filtering
- net/qede: fix multicast drop in promiscuous mode
- net/qede: remove dead code
- net/sfc: do not enforce hash offload in RSS multi-queue
- net/virtio-user: check tap system call setting

- net/virtio-user: fix status management
- pci: fix address domain format size
- rawdev: allow getting info for unknown device
- rawdev: export dump function in map file
- rawdev: fill NUMA socket ID in info
- rawdev: remove remaining experimental tags
- raw/ufpga/base: fix NIOS SPI init
- raw/ufpga/base: fix SPI transaction
- rib: add C++ include guard
- sched: fix 64-bit rate
- sched: fix port time rounding
- sched: fix subport freeing
- service: fix C++ linkage
- service: fix core mapping reset
- service: fix lcore iteration
- test: allow no-huge mode for fast-tests
- test/bpf: fix few small issues
- test/crypto: add mixed encrypted-digest
- test/crypto: change cipher offset for ESN vector
- test/crypto: fix asymmetric session mempool creation
- test/cycles: restore default delay callback
- test: fix build with ring PMD but no bond PMD
- test: fix rpath for drivers with meson
- test/hash: move lock-free tests to perf tests
- test/mbuf: fix a dynamic flag log
- test/ring: fix statistics in bulk enq/dequeue
- version: 19.11.4-rc1
- vfiio: map contiguous areas in one go
- vfiio: remove unused variable
- vhost: fix double-free with zero-copy
- vhost: fix features definition location
- vhost: fix virtio ready flag check
- vhost: remove zero-copy and client mode restriction

1.11.2 19.11.4 Validation

- Canonical(R) Testing
 - Build tests on all Ubuntu architectures
 - OVS-DPDK tests on x86_64
- Red Hat(R) Testing
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 5.1
 - * X540-AT2 NIC(ixgbe, 10G)
 - Functionality
 - * Guest with device assignment(PF) throughput testing(1G hugepage size)
 - * Guest with device assignment(PF) throughput testing(2M hugepage size)
 - * Guest with device assignment(VF) throughput testing
 - * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing
 - * PVP vhost-user 2Q throughput testing
 - * PVP vhost-user 1Q - cross numa node throughput testing
 - * Guest with vhost-user 2 queues throughput testing
 - * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect
 - * PVP 1Q live migration testing
 - * PVP 1Q cross numa node live migration testing
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M)
 - * Guest with ovs+dpdk+vhost-user 2Q live migration testing
 - * Allocate memory from the NUMA node which Virtio device locates
 - * Host PF + DPDK testing
 - * Host VF + DPDK testing
- Intel(R) Testing
 - **Basic Intel(R) NIC(ixgbe, i40e and ice) testing**
 - * PF (i40e)
 - * PF (ixgbe)
 - * PF (ice)
 - * VF (i40e)
 - * VF (ixgbe)

- * VF (ice)
- * Compile Testing
- * Intel NIC single core/NIC performance
- Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance
- Intel(R) Testing with Open vSwitch
 - OVS testing with OVS branches master and 2.13 with VSPERF
 - Tested NICs
 - * i40e (X710)
 - * ixgbe (82599ES)
 - * ice
 - * vhost user client
 - Functionality
 - * Performance tests
 - * vHost zero-copy
 - * Flow control
 - * RSS
 - * Partial HW offloading
- 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
 - * l3fwd-power example application
 - * Multi-process example applications
 - ConnectX-5
 - * RHEL 7.4

- * Driver MLNX_OFED_LINUX-5.1-0.6.6.0
- * fw 16.28.1002
- ConnectX-4 Lx
 - * RHEL 7.4
 - * Driver MLNX_OFED_LINUX-5.1-0.6.6.0
 - * fw 14.28.1002
- Microsoft(R) Testing
 - Platform
 - * **Azure**
 - Ubuntu 16.04-LTS
 - Ubuntu 18.04-DAILY-LTS
 - RHEL 7-RAW
 - RHEL 7.5
 - CentOS 7.5
 - SLES-15-sp1 gen1
 - * Mellanox(R) ConnectX-4
 - * LISAv2 test framework
 - Functionality
 - * VERIFY-DPDK-COMPLIANCE - verifies kernel is supported and that the build is successful
 - * VERIFY-DPDK-BUILD-AND-TESTPMD-TEST - verifies using testpmd that packets can be sent from a VM to another VM
 - * VERIFY-SRIOV-FAILSAFE-FOR-DPDK - disables/enables Accelerated Networking for the NICs under test and makes sure DPDK works in both scenarios
 - * VERIFY-DPDK-FAILSAFE-DURING-TRAFFIC - disables/enables Accelerated Networking for the NICs while generating traffic using testpmd
 - * PERF-DPDK-FWD-PPS-DS15 - verifies DPDK forwarding performance using testpmd on 2, 4, 8 cores, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-SINGLE-CORE-PPS-DS4 - verifies DPDK performance using testpmd on 1 core, rx and io mode on size Standard_DS4_v2
 - * PERF-DPDK-SINGLE-CORE-PPS-DS15 - verifies DPDK performance using testpmd on 1 core, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-MULTICORE-PPS-DS15 - verifies DPDK performance using testpmd on 2, 4, 8 cores, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-MULTICORE-PPS-F32 - verifies DPDK performance using testpmd on 2, 4, 8, 16 cores, rx and io mode on size Standard_F32s_v2
 - * DPDK-RING-LATENCY - verifies DPDK CPU latency using dpdk-ring-ping

- * VERIFY-DPDK-PRIMARY-SECONDARY-PROCESSES - verifies primary / secondary processes support for DPDK. Runs only on RHEL and Ubuntu distros with Linux kernel >= 4.20
- * VERIFY-DPDK-OVS - builds OVS with DPDK support and tests if the OVS DPDK ports can be created. Runs only on Ubuntu distro.

1.11.3 19.11.4 Known Issues

- ICE
 - Exception on VF port reset
 - MD5 is not same between kernel ethtool and dpdk ethtool when testing userspace_ethtool/retrieve_eeprom
- vhost/virtio
 - udp-fragmentation-offload cannot be setup on Ubuntu 19.10 VMs. https://bugzilla.kernel.org/show_bug.cgi?id=207075
 - l3fwd-power can wake up lcore, but then cannot sleep again
- cryptodev
 - fips_cryptodev test fails for TDES
- vdev_netvsc
 - hot-removal of VF driver can fail

1.12 19.11.5 Release Notes

1.12.1 19.11.5 Fixes

- vhost/crypto: fix data length check (CVE-2020-14374)
- vhost/crypto: fix incorrect descriptor deduction (CVE-2020-14378)
- vhost/crypto: fix incorrect write back source
- vhost/crypto: fix missed request check for copy mode (CVE-2020-14376 CVE-2020-14377)
- vhost/crypto: fix pool allocation
- vhost/crypto: fix possible TOCTOU attack (CVE-2020-14375)

1.12.2 19.11.5 Validation

- Intel(R) Testing
 - Basic cryptodev testing
 - * vhost_crypto Unit test and Function/Performance test

1.13 19.11.6 Release Notes

1.13.1 19.11.6 Fixes

- acl: fix x86 build for compiler without AVX2

- app/bbdev: fix test vector symlink
- app/eventdev: check timer adapters number
- app: fix ethdev port id size
- app: fix missing dependencies
- app/testpmd: do not allow dynamic change of core number
- app/testpmd: fix bonding xmit balance policy command
- app/testpmd: fix build with gcc 11
- app/testpmd: fix descriptor id check
- app/testpmd: fix displaying Rx/Tx queues information
- app/testpmd: fix max Rx packet length for VLAN packet
- app/testpmd: fix MTU after device configure
- app/testpmd: fix name of bitrate library in meson build
- app/testpmd: fix packet header in txonly mode
- app/testpmd: fix port id check in Tx VLAN command
- app/testpmd: fix RSS key for flow API RSS rule
- app/testpmd: fix VLAN configuration on failure
- app/testpmd: remove restriction on Tx segments set
- app/testpmd: revert max Rx packet length adjustment
- app/testpmd: revert setting MTU explicitly after configure
- app/test-sad: fix uninitialized variable
- baseband/fpga_lte_fec: fix crash with debug
- baseband/turbo_sw: fix memory leak in error path
- build: fix gcc warning requiring Wformat
- build: fix install on Windows
- build: fix MS linker flag with meson 0.54
- build: skip detecting libpcap via pcap-config
- bus/dpaa: fix fd check before close
- bus/dpaa: remove logically dead code
- bus/fslmc: fix atomic queues on NXP LX2 platform
- bus/fslmc: fix dpio close
- bus/fslmc: fix VFIO group descriptor check
- bus/pci: fix leak on VFIO mapping error
- bus/pci: fix memory leak when unmapping VFIO resource
- bus/pci: remove duplicate declaration

- bus/pci: remove unused scan by address
- common/mlx5: fix DevX SQ object creation
- common/mlx5: fix name for ConnectX VF device ID
- common/mlx5: fix PCI address lookup
- common/qat: add missing kmod dependency info
- compress/isal: check allocation in queue setup
- config: add Graviton2(arm64) defconfig
- config: enable packet prefetching with Meson
- crypto/aesni_mb: fix CCM digest size check
- crypto/aesni_mb: fix GCM digest size check
- crypto/armv8: fix mempool object returning
- crypto/caam_jr: fix device tree parsing for SEC_ERA
- cryptodev: fix parameter parsing
- crypto/dpaa2_sec: fix stats query without queue pair
- crypto/dpaa2_sec: remove dead code
- crypto/dpaa_sec: fix a null pointer dereference
- crypto/octeontx2: fix multi-process
- crypto/octeontx2: fix out-of-place support
- crypto/octeontx2: fix session-less mode
- crypto/octeontx: fix out-of-place support
- crypto/scheduler: fix header install with meson
- crypto/scheduler: remove unused internal seqn
- devtools: fix build test config inheritance from env
- devtools: fix directory filter in forbidden token check
- devtools: fix x86-default build test install env
- distributor: fix API documentation
- distributor: fix buffer use after free
- distributor: fix clearing returns buffer
- distributor: fix flushing in flight packets
- distributor: fix handshake deadlock
- distributor: fix handshake synchronization
- distributor: fix return pkt calls in single mode
- distributor: fix scalar matching
- distributor: handle worker shutdown in burst mode

- doc: add SPDX license tag header to Intel performance guide
- doc: add SPDX license tag header to meson guide
- doc: clarify instructions on running as non-root
- doc: fix diagram in dpaa2 guide
- doc: fix EF10 Rx mode name in sfc guide
- doc: fix ethdev port id size
- doc: fix formatting of notes in meson guide
- doc: fix grammar
- doc: fix missing classify methods in ACL guide
- doc: fix rule file parameters in l3fwd-acl guide
- doc: fix typo in ipsec-secgw guide
- doc: fix typo in KNI guide
- doc: fix typo in pcap guide
- doc: improve multiport PF in nfp guide
- doc: remove notice about AES-GCM IV and J0
- doc: remove obsolete deprecation notice for power library
- doc: update information on using hugepages
- drivers/net: fix port id size
- eal/arm: fix build with gcc optimization level 0
- eal/arm: fix clang build of native target
- eal: fix doxygen for EAL cleanup
- eal: fix leak on device event callback unregister
- eal: fix MCS lock and ticketlock headers install
- eal/linux: change udev debug message
- eal/linux: fix memory leak in uevent handling
- eal/x86: fix memcpy AVX-512 enablement
- efd: fix tailq entry leak in error path
- ethdev: fix data type for port id
- ethdev: fix memory ordering for callback functions
- ethdev: fix RSS flow expansion in case of mismatch
- ethdev: move non-offload capabilities
- ethdev: remove redundant license text
- eventdev: check allocation in Tx adapter
- eventdev: fix adapter leak in error path

- event/dpaa2: fix dereference before null check
- event/dpaa2: remove dead code from self test
- event/octeontx2: unlink queues during port release
- examples/fips_validation: fix buffer overflow
- examples/fips_validation: fix build with pkg-config
- examples/fips_validation: fix missed version line
- examples/fips_validation: fix version compatibility
- examples: fix flattening directory layout on install
- examples/ioat: fix stats print
- examples/ip_pipeline: fix external build
- examples/ip_pipeline: use POSIX network address conversion
- examples/ipsec-secgw: use POSIX network address conversion
- examples/kni: fix build with pkg-config
- examples/l2fwd-crypto: fix build with pkg-config
- examples/l2fwd-crypto: fix missing dependency
- examples/l2fwd-keepalive: skip meson build if no librt
- examples/l3fwd-power: check packet types after start
- examples/multi_process: fix build on Ubuntu 20.04
- examples/ntb: fix clean target
- examples/performance-thread: fix build with low core count
- examples/performance-thread: fix build with pkg-config
- examples/qos_sched: fix usage string
- examples/rxtx_callbacks: fix build with pkg-config
- examples/vhost_blk: check driver start failure
- examples/vhost_blk: fix build with pkg-config
- examples/vhost_crypto: add new line character in usage
- examples/vm_power: fix 32-bit build
- examples/vm_power: fix build on Ubuntu 20.04
- fix spellings that Lintian complains about
- gro: fix packet type detection with IPv6 tunnel
- gso: fix payload unit size for UDP
- ipc: fix spelling in log and comment
- kni: fix build on RHEL 8.3
- kni: fix build with Linux 5.9

- license: add licenses for exception cases
- maintainers: update Mellanox emails
- malloc: fix style in free list index computation
- mbuf: fix dynamic fields and flags with multiprocess
- mbuf: fix typo in dynamic field convention note
- mcslock: fix hang in weak memory model
- mem: fix allocation failure on non-NUMA kernel
- mem: fix allocation in container with SELinux
- mem: fix config name in error logs
- mempool/octeontx: fix aura to pool mapping
- net/af_xdp: avoid deadlock due to empty fill queue
- net/af_xdp: change return value from Rx to unsigned
- net/af_xdp: fix pointer storage size
- net/af_xdp: fix umem size
- net/af_xdp: use strlcpy instead of strncpy
- net/bnx2x: add QLogic vendor id for BCM57840
- net/bnxt: add memory allocation check in VF info init
- net/bnxt: add separate mutex for FW health check
- net/bnxt: fix boolean operator usage
- net/bnxt: fix checking VNIC in shutdown path
- net/bnxt: fix crash in vector mode Tx
- net/bnxt: fix doorbell barrier location
- net/bnxt: fix drop enable in get Rx queue info
- net/bnxt: fix endianness while setting L4 destination port
- net/bnxt: fix L2 filter allocation
- net/bnxt: fix link status during device recovery
- net/bnxt: fix link update
- net/bnxt: fix LRO configuration
- net/bnxt: fix memory leak when freeing VF info
- net/bnxt: fix queue get info
- net/bnxt: fix queue release
- net/bnxt: fix resetting mbuf data offset
- net/bnxt: fix Rx performance by removing spinlock
- net/bnxt: fix shift operation

- net/bnxt: fix structure variable initialization
- net/bnxt: fix vnic Rx queue cnt updation
- net/bnxt: fix xstats by id
- net/bnxt: increase size of Rx CQ
- net/bnxt: remove useless prefetches
- net/bonding: fix possible unbalanced packet receiving
- net/bonding: fix Rx queue conversion
- net: check segment pointer in raw checksum processing
- net/cxgbe: fix crash when accessing empty Tx mbuf list
- net/cxgbe: fix duplicate MAC addresses in MPS TCAM
- net/cxgbe: fix queue DMA ring leaks during port close
- net/dpaa2: fix build with timesync functions
- net/dpaa2: fix misuse of interface index
- net/dpaa: fix port ID type in API
- net/ena/base: align IO CQ allocation to 4K
- net/ena/base: fix release of wait event
- net/ena/base: specify delay operations
- net/ena/base: use min/max macros with type conversion
- net/ena: fix getting xstats global stats offset
- net/ena: fix setting Rx checksum flags in mbuf
- net/ena: remove unused macro
- net/enic: fix header sizes when copying flow patterns
- net/enic: generate VXLAN src port if it is zero in template
- net/enic: ignore VLAN inner type when it is zero
- net/failsafe: fix double space in warning log
- net/failsafe: fix state synchro cleanup
- net/fm10k: fix memory leak when thresh check fails
- net/fm10k: fix memory leak when Tx thresh check fails
- net/fm10k: fix vector Rx
- net/hinic/base: add message check for command channel
- net/hinic/base: fix clock definition with glibc version
- net/hinic/base: fix log info for PF command channel
- net/hinic/base: get default cos from chip
- net/hinic/base: remove queue number limitation

- net/hinic/base: support two or more AEQS for chip
- net/hinic: fix filters on memory allocation failure
- net/hinic: fix negative array index read
- net/hinic: fix Rx nombuf stats
- net/hinic: remove optical module operation
- net/hns3: check PCI config space reads
- net/hns3: check PCI config space write
- net/hns3: check setting VF PCI bus return value
- net/hns3: decrease non-nearby memory access in Rx
- net/hns3: fix configurations of port-level scheduling rate
- net/hns3: fix configuring device with RSS enabled
- net/hns3: fix config when creating RSS rule after flush
- net/hns3: fix crash with multi-TC
- net/hns3: fix data type to store queue number
- net/hns3: fix default MAC address from firmware
- net/hns3: fix deleting default VLAN from PF
- net/hns3: fix error type when validating RSS flow action
- net/hns3: fix flow error type
- net/hns3: fix flow RSS queue number 0
- net/hns3: fix flushing RSS rule
- net/hns3: fix out of bounds access
- net/hns3: fix queue offload capability
- net/hns3: fix reassembling multiple segment packets in Tx
- net/hns3: fix RSS max queue id allowed in multi-TC
- net/hns3: fix some incomplete command structures
- net/hns3: fix storing RSS info when creating flow action
- net/hns3: fix TX checksum with fix header length
- net/hns3: reduce address calculation in Rx
- net/hns3: report Rx drop packets enable configuration
- net/hns3: report Rx free threshold
- net/hns3: skip VF register access when PF in FLR
- net/i40e: add C++ include guard
- net/i40e/base: fix function header arguments
- net/i40e/base: fix Rx only for unicast promisc on VLAN

- net/i40e: fix build for log format specifier
- net/i40e: fix byte counters
- net/i40e: fix flow director for eth + VLAN pattern
- net/i40e: fix incorrect FDIR flex configuration
- net/i40e: fix link status
- net/i40e: fix QinQ flow pattern to allow non full mask
- net/i40e: fix recreating flexible flow director rule
- net/i40e: fix vector Rx
- net/i40e: fix virtual channel conflict
- net/iavf: downgrade error log
- net/iavf: enable port reset
- net/iavf: fix command after PF reset
- net/iavf: fix flow flush after PF reset
- net/iavf: fix iterator for RSS LUT
- net/iavf: fix performance drop after port reset
- net/iavf: fix port start during configuration restore
- net/iavf: fix releasing mbufs
- net/iavf: fix scattered Rx enabling
- net/iavf: fix setting of MAC address
- net/iavf: fix unchecked Tx cleanup error
- net/iavf: fix vector Rx
- net/iavf: support multicast configuration
- net/ice/base: fix issues around move nodes
- net/ice/base: fix parameter name in comment
- net/ice: fix flow validation for unsupported patterns
- net/ice: fix ptype parsing
- net/ice: fix Rx offload flags in SSE path
- net/ice: fix vector Rx
- net/ice: update writeback policy to reduce latency
- net/ixgbe: check switch domain allocation result
- net/ixgbe: fix vector Rx
- net/ixgbe: fix VF reset HW error handling
- net/ixgbe: remove redundant MAC flag check
- net/memif: do not update local copy of tail in Tx

- net/memif: relax load of ring head for M2S ring
- net/memif: relax load of ring head for S2M ring
- net/memif: relax load of ring tail for M2S ring
- net/mlx5: fix debug configuration build issue
- net/mlx5: fix hairpin dependency on destination DevX TIR
- net/mlx5: fix meter table definitions
- net/mlx5: fix missing meter packet
- net/mlx5: fix port shared data reference count
- net/mlx5: fix raw encap/decap limit
- net/mlx5: fix representor interrupts handler
- net/mlx5: fix RSS queue type validation
- net/mlx5: fix RSS RETA reset on start
- net/mlx5: fix Rx descriptor status
- net/mlx5: fix Rx packet padding config via DevX
- net/mlx5: fix Rx queue completion index consistency
- net/mlx5: fix Rx queue count calculation
- net/mlx5: fix Rx queue count calculation
- net/mlx5: fix switch port id when representor in bonding
- net/mlx5: fix xstats reset reinitialization
- net/mlx5: free MR resource on device DMA unmap
- net/mlx5: remove unused includes
- net/mlx5: remove unused log macros
- net/mlx5: remove unused variable in Tx queue creation
- net/mlx5: validate MPLSoGRE with GRE key
- net/mlx: do not enforce RSS hash offload
- net/mvpp2: fix memory leak in error path
- net/netvsc: allocate contiguous physical memory for RNDIS
- net/netvsc: check for overflow on packet info from host
- net/netvsc: disable external mbuf on Rx by default
- net/netvsc: fix multiple channel Rx
- net/netvsc: fix rndis packet addresses
- net/netvsc: fix stale value after free
- net/netvsc: fix Tx queue leak in error path
- net/netvsc: manage VF port under read/write lock

- net/netvsc: replace compiler builtin overflow check
- net/nfp: expand device info get
- net/octeontx2: fix multi segment mode for jumbo packets
- net/octeontx2: fix RSS flow create
- net/octeontx2: remove useless check before free
- net/pcap: fix crash on exit for infinite Rx
- net/pcap: fix input only Rx
- net/pfe: fix misuse of interface index
- net/qede: fix dereference before null check
- net/qede: fix getting link details
- net/qede: fix milliseconds sleep macro
- net/ring: check internal arguments
- net/ring: fix typo in log message
- net/sfc/base: fix tunnel configuration
- net/sfc: fix RSS hash flag when offload is disabled
- net/sfc: fix RSS hash offload if queue action is used
- net/softnic: use POSIX network address conversion
- net/tap: free mempool when closing
- net/thunderx: fix memory leak on rbdr desc ring failure
- net/vdev_netvsc: fix device probing error flow
- net/vhost: fix xstats after clearing stats
- net/virtio: check raw checksum failure
- net/virtio: fix packed ring indirect descriptors setup
- pmdinfogen: fix build with gcc 11
- port: remove useless assignment
- power: fix current frequency index
- raw/dpaa2_qdma: fix reset
- raw/ifpga/base: fix interrupt handler instance usage
- raw/ifpga/base: fix return of IRQ unregister
- raw/ifpga/base: handle unsupported interrupt type
- raw/ifpga: terminate string filled by readlink with null
- raw/ifpga: use trusted buffer to free
- raw/ioat: fix missing close function
- raw/skeleton: allow closing already closed device

- raw/skeleton: reset test statistics
- rcu: avoid literal suffix warning in C++ mode
- Revert “app/testpmd: fix name of bitrate library in meson build”
- Revert “Revert “build: always link whole DPDK static libraries””
- Revert “Revert “build/pkg-config: improve static linking flags””
- Revert “Revert “build/pkg-config: move pkg-config file creation””
- Revert “Revert “build/pkg-config: output drivers first for static build””
- Revert “Revert “build/pkg-config: prevent overlinking””
- Revert “Revert “devtools: test static linkage with pkg-config””
- stack: fix uninitialized variable
- stack: reload head when pop fails
- table: fix hash for 32-bit
- test/crypto: fix device number
- test/crypto: fix stats test
- test/distributor: collect return mbufs
- test/distributor: ensure all packets are delivered
- test/distributor: fix freeing mbufs
- test/distributor: fix lcores statistics
- test/distributor: fix mbuf leak on failure
- test/distributor: fix quitting workers in burst mode
- test/distributor: fix race conditions on shutdown
- test/distributor: fix shutdown of busy worker
- test/event_crypto_adapter: fix configuration
- test/event: fix function arguments for crypto adapter
- test/mbuf: skip field registration at busy offset
- test/rcu: fix build with low core count
- test/ring: fix number of single element enqueue/dequeue
- timer: add limitation note for sync stop and reset
- usertools: fix CPU layout script to be PEP8 compliant
- usertools: fix pmdinfo parsing
- vdp/afc: fix build with recent kernels
- version: 19.11.6-rc1
- vfio: fix group descriptor check
- vhost: fix error path when setting memory tables

- vhost: fix external mbuf creation
- vhost: fix fd leak in dirty logging setup
- vhost: fix fd leak in kick setup
- vhost: fix IOTLB mempool single-consumer flag
- vhost: fix virtio-net header length with packed ring
- vhost: fix virtqueue initialization
- vhost: fix virtqueues metadata allocation
- vhost: validate index in available entries API
- vhost: validate index in guest notification API
- vhost: validate index in inflight API
- vhost: validate index in live-migration API

1.13.2 19.11.6 Validation

- Intel(R) Testing
 - **Basic Intel(R) NIC(ixgbe, i40e and ice) testing**
 - * PF (i40e)
 - * PF (ixgbe)
 - * PF (ice)
 - * VF (i40e)
 - * VF (ixgbe)
 - * VF (ice)
 - * Compile Testing
 - * Intel NIC single core/NIC performance
 - Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance
- Microsoft(R) Testing
 - Platform
 - * **Azure**
 - Ubuntu 16.04-LTS
 - Ubuntu 18.04-DAILY-LTS
 - RHEL 7.5
 - Openlogic CentOS 7.5
 - SLES-15-sp1 gen1
 - * Mellanox(R) ConnectX-4

- * LISAv2 test framework
- Functionality
 - * VERIFY-DPDK-COMPLIANCE * verifies kernel is supported and that the build is successful
 - * VERIFY-DPDK-BUILD-AND-TESTPMD-TEST * verifies using testpmd that packets can be sent from a VM to another VM
 - * VERIFY-SRIOV-FAILSAFE-FOR-DPDK * disables/enables Accelerated Networking for the NICs under test and makes sure DPDK works in both scenarios
 - * VERIFY-DPDK-FAILSAFE-DURING-TRAFFIC * disables/enables Accelerated Networking for the NICs while generating traffic using testpmd
 - * PERF-DPDK-FWD-PPS-DS15 * verifies DPDK forwarding performance using testpmd on 2, 4, 8 cores, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-SINGLE-CORE-PPS-DS4 * verifies DPDK performance using testpmd on 1 core, rx and io mode on size Standard_DS4_v2
 - * PERF-DPDK-SINGLE-CORE-PPS-DS15 * verifies DPDK performance using testpmd on 1 core, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-MULTICORE-PPS-DS15 * verifies DPDK performance using testpmd on 2, 4, 8 cores, rx and io mode on size Standard_DS15_v2
 - * PERF-DPDK-MULTICORE-PPS-F32 * verifies DPDK performance using testpmd on 2, 4, 8, 16 cores, rx and io mode on size Standard_F32s_v2
 - * DPDK-RING-LATENCY * verifies DPDK CPU latency using dpdk-ring-ping
 - * VERIFY-DPDK-OVS * builds OVS with DPDK support and tests if the OVS DPDK ports can be created. Runs only on Ubuntu distro.
 - * VERIFY-DPDK-BUILD-AND-NETVSCPMD-TEST * verifies using testpmd with netvsc pmd that packets can be sent from a VM to another VM.
 - * VERIFY-SRIOV-FAILSAFE-FOR-DPDK-NETVSCPMD * disables/enables Accelerated Networking for the NICs under test and makes sure DPDK with netvsc pmd works in both scenarios.
 - * VERIFY-DPDK-FAILSAFE-DURING-TRAFFIC-NETVSCPMD * Verify Accelerated Networking (VF) removed and readded for the NICs while generating traffic using testpmd with netvsc pmd.
- Red Hat(R) Testing
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 5.2
 - * X540-AT2 NIC(ixgbe, 10G)
 - Functionality
 - * Guest with device assignment(PF) throughput testing(1G hugepage size)

- * Guest with device assignment(PF) throughput testing(2M hugepage size)
- * Guest with device assignment(VF) throughput testing
- * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing
- * PVP vhost-user 2Q throughput testing
- * PVP vhost-user 1Q * cross numa node throughput testing
- * Guest with vhost-user 2 queues throughput testing
- * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect
- * vhost-user reconnect with dpdk-client, qemu-server: ovs reconnect: PASS
- * PVP 1Q live migration testing
- * PVP 1Q cross numa node live migration testing
- * Guest with ovs+dpdk+vhost-user 1Q live migration testing
- * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M)
- * Guest with ovs+dpdk+vhost-user 2Q live migration testing
- * Allocate memory from the NUMA node which Virtio device locates
- * Host PF + DPDK testing
- * Host VF + DPDK testing
- Intel(R) Testing with Open vSwitch
 - OVS testing with OVS 2.14.1
 - Performance
 - * ICE Device
 - Basic performance tests (RFC2544 P2P, PVP_CONT, RFC2544 PVP_TPUT, RFC2544 PVVP_TPUT, PVPV) Jumbo frames RSS
 - * i40e Device
 - Basic performance (RFC2544 P2P, PVP_CONT, RFC2544 PVP_TPUT, RFC2544 PVVP_TPUT, PVPV) Jumbo frames RSS Flow control
 - * ixgbe Device
 - Basic performance tests (RFC2544 P2P, PVP_CONT, RFC2544 PVP_TPUT, RFC2544 PVVP_TPUT, PVPV) Jumbo frames RSS
 - Functionality
 - * vhostuserclient device
 - * jumbo frames
 - * dpdkvhostuserclient re-connect
 - * dpdkvhostuserclient NUMA node
- Nvidia(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
- * l3fwd-power example application
- * Multi-process example applications
- Build tests
 - * Ubuntu 20.04 with MLNX_OFED_LINUX-5.1-2.5.8.0.
 - * Ubuntu 20.04 with rdma-core master (6a5c1b7).
 - * Ubuntu 20.04 with rdma-core v28.0.
 - * Ubuntu 18.04 with rdma-core v17.1.
 - * Ubuntu 18.04 with rdma-core master (6a5c1b7) (i386).
 - * Ubuntu 16.04 with rdma-core v22.7.
 - * Fedora 32 with rdma-core v32.0.
 - * CentOS 7 7.9.2009 with rdma-core master (6a5c1b7).
 - * CentOS 7 7.9.2009 with MLNX_OFED_LINUX-5.1-2.5.8.0.
 - * CentOS 8 8.3.2011 with rdma-core master (6a5c1b7).
 - * openSUSE Leap 15.2 with rdma-core v27.1.
- ConnectX-5
 - * RHEL 7.4
 - * Driver MLNX_OFED_LINUX-5.1-2.5.8.0
 - * fw 14.28.2006
- ConnectX-4 Lx
 - * RHEL 7.4
 - * Driver MLNX_OFED_LINUX-5.1-2.5.8.0
 - * fw 16.28.2006

1.13.3 19.11.6 Known Issues

- i40e
 - rss_to_rte_flow/set_key_keylen: create rule can fail. https://bugs.dpdk.org/show_bug.cgi?id=573
 - inconsistency with expected queue after creating a flow rule - firmware issue.
- vhost/virtio
 - udp-fragmentation-offload cannot be setup on Ubuntu 19.10 VMs. https://bugzilla.kernel.org/show_bug.cgi?id=207075
- vdev_netvsc
 - hot-removal of VF driver can fail

1.14 19.11.7 Release Notes

1.14.1 19.11.7 Fixes

- app/crypto-perf: fix CSV output format
- app/crypto-perf: fix latency CSV output
- app/crypto-perf: fix spelling in output
- app/crypto-perf: remove always true condition
- app/eventdev: adjust event count order for pipeline test
- app/eventdev: fix SMP barrier in performance test
- app/eventdev: remove redundant enqueue in burst Tx
- app: fix build with extra include paths
- app/procinfo: fix check on xstats-ids
- app/procinfo: fix _filters stats reporting
- app/procinfo: remove useless memset
- app/testpmd: avoid exit without terminal restore
- app/testpmd: fix help of metering commands
- app/testpmd: fix IP checksum calculation
- app/testpmd: fix key for RSS flow rule
- app/testpmd: fix max Rx packet length for VLAN packets
- app/testpmd: fix packets dump overlapping
- app/testpmd: release flows left before port stop
- build: fix linker flags on Windows
- build: fix plugin load on static build
- build: fix scheduler macro definition for meson
- build: force pkg-config for dependency detection

- build: provide suitable error for “both” libraries option
- common/mlx5: fix pointer cast on Windows
- common/octeontx2: fix build with SVE
- crypto/dpaa2_sec: fix memory allocation check
- doc: fix mark action zero value in mlx5 guide
- doc: fix product link in hns3 guide
- doc: fix QinQ flow rules in testpmd guide
- doc: update flow mark action in mlx5 guide
- eal/arm: fix debug build with gcc for 128-bit atomics
- eal: fix MCS lock header include
- eal: fix reciprocal header include
- eal/linux: fix handling of error events from epoll
- ethdev: fix max Rx packet length check
- ethdev: fix missing header include
- eventdev: fix a return value comment
- examples/eventdev: add info output for main core
- examples/eventdev: check CPU core enabling
- examples/eventdev: move ethdev stop to the end
- examples/l3fwd: remove limitation on Tx queue count
- fbarray: fix overlap check
- fib: fix missing header includes
- ip_frag: remove padding length of fragment
- ipsec: fix missing header include
- lib/power: fix make build error
- lpm: fix vector IPv4 lookup
- mbuf: add C++ include guard for dynamic fields header
- mbuf: fix missing header include
- mem: exclude unused memory from core dump
- net/af_xdp: remove useless assignment
- net/avp: remove always true condition
- net/bnxt: disable end of packet padding for Rx
- net/bnxt: fix cleanup on mutex init failure
- net/bnxt: fix doorbell write ordering
- net/bnxt: fix fallback mbuf allocation logic

- net/bnxt: fix FW version log
- net/bnxt: fix max rings computation
- net/bnxt: fix memory leak when mapping fails
- net/bnxt: fix null termination of Rx mbuf chain
- net/bnxt: fix outer UDP checksum Rx offload capability
- net/bnxt: fix Rx rings in RSS redirection table
- net/bnxt: fix VNIC config on Rx queue stop
- net/bnxt: fix VNIC RSS configure function
- net/bnxt: propagate FW command failure to application
- net/bnxt: release HWRM lock in error
- net/bonding: fix PCI address comparison on non-PCI ports
- net/bonding: fix port id validity check on parsing
- net/bonding: remove local variable shadowing outer one
- net/cxgbe: fix jumbo frame flag condition
- net/dpaa2: fix jumbo frame flag condition for MTU set
- net/dpaa: fix jumbo frame flag condition for MTU set
- net/e1000: fix flow control mode setting
- net/e1000: fix jumbo frame flag condition for MTU set
- net/ena: flush Rx buffers memory pool cache
- net/enetc: fix jumbo frame flag condition for MTU set
- net/enic: fix filter log message
- net/enic: fix filter type used for flow API
- net/hns3: adjust format specifier for enum
- net/hns3: fix data overwriting during register dump
- net/hns3: fix dump register out of range
- net/hns3: fix error code in xstats
- net/hns3: fix firmware exceptions by concurrent commands
- net/hns3: fix flow director rule residue on malloc failure
- net/hns3: fix interception with flow director
- net/hns3: fix jumbo frame flag condition for MTU set
- net/hns3: fix memory leak on secondary process exit
- net/hns3: fix register length when dumping registers
- net/hns3: fix Rx/Tx errors stats
- net/hns3: fix VF query link status in dev init

- net/hns3: fix VF reset on mailbox failure
- net/hns3: fix xstats with id and names
- net/hns3: remove MPLS from supported flow items
- net/hns3: use new opcode for clearing hardware resource
- net/hns3: validate requested maximum Rx frame length
- net/i40e: add null input checks
- net/i40e: fix global register recovery
- net/i40e: fix jumbo frame flag condition
- net/i40e: fix L4 checksum flag
- net/i40e: fix returned code for RSS hardware failure
- net/i40e: fix Rx bytes statistics
- net/i40e: fix stats counters
- net/i40e: fix VLAN stripping in VF
- net/i40e: fix X722 for 802.1ad frames ability
- net/iavf: fix jumbo frame flag condition
- net/iavf: fix vector mapping with queue
- net/ice/base: fix memory handling
- net/ice/base: fix null pointer dereference
- net/ice/base: fix tunnel destroy
- net/ice: check Rx queue number on RSS init
- net/ice: disable IPv4 checksum offload in vector Tx
- net/ice: enlarge Rx queue rearm threshold to 64
- net/ice: fix jumbo frame flag condition
- net/ice: fix outer checksum flags
- net/ice: fix outer UDP Tx checksum offload
- net/ice: fix RSS lookup table initialization
- net/ipn3ke: fix jumbo frame flag condition for MTU set
- net/ixgbe: add new flag of stripped VLAN for NEON
- net/ixgbe: detect failed VF MTU set
- net/ixgbe: disable NFS filtering
- net/ixgbe: enable IXGBE NEON vector PMD when CHECKSUM need to checksum
- net/ixgbe: fix configuration of max frame size
- net/ixgbe: fix flex bytes flow director rule
- net/ixgbe: fix jumbo frame flag condition

- net/ixgbe: fix UDP zero checksum on x86
- net/ixgbe: support bad checksum flag for NEON
- net/ixgbe: support good checksum flag for NEON
- net/liquidio: fix jumbo frame flag condition for MTU set
- net/mlx4: fix handling of probing failure
- net/mlx4: fix port attach in secondary process
- net/mlx5: fix comparison sign in flow engine
- net/mlx5: fix crash on secondary process port close
- net/mlx5: fix leak on Rx queue creation failure
- net/mlx5: fix leak on Tx queue creation failure
- net/mlx5: refuse empty VLAN in flow pattern
- net/mvnet: check allocation in Rx queue flush
- net/mvpp2: fix frame size checking
- net/mvpp2: fix stack corruption
- net/mvpp2: remove CRC length from MRU validation
- net/mvpp2: remove debug log on fast-path
- net/mvpp2: remove VLAN flush
- net/netvsc: ignore unsupported packet on sync command
- net/nfp: fix jumbo frame flag condition for MTU set
- net/nfp: read chip model from PluDevice register
- net/octeontx2: fix corruption in segments list
- net/octeontx2: fix jumbo frame flag condition for MTU
- net/octeontx2: fix PF flow action for Tx
- net/octeontx: fix build with SVE
- net/pcap: fix byte stats for drop Tx
- net/pcap: fix infinite Rx with large files
- net/pcap: remove local variable shadowing outer one
- net/qede: fix jumbo frame flag condition for MTU set
- net/qede: fix promiscuous enable
- net/sfc: fix generic byte statistics to exclude FCS bytes
- net/sfc: fix jumbo frame flag condition for MTU set
- net/sfc: fix TSO and checksum offloads for EF10
- net/thunderx: fix jumbo frame flag condition for MTU set
- net/virtio-user: fix run closing stdin and close callfd

- power: clean up includes
- power: create guest channel public header file
- power: export guest channel header file
- power: fix missing header includes
- power: make channel message functions public
- power: rename constants
- power: rename public structs
- Revert “app/testpmd: release flows left before port stop”
- rib: fix insertion in some cases
- rib: fix missing header include
- rib: fix missing header includes
- app/testpmd: fix setting maximum packet length
- net/mlx5: fix port attach in secondary process
- net/mlx5: fix VXLAN decap on non-VXLAN flow
- mempool: fix panic on dump or audit
- mem: fix build
- usertools: fix binding built-in kernel driver
- test/distributor: fix return buffer queue overload
- test/event_crypto: set cipher operation in transform
- test: fix buffer overflow in Tx burst
- test: fix terminal settings on exit
- test/ipsec: fix result code for not supported
- test/mcslock: remove unneeded per lcore copy
- usertools: fix Python compatibility issue
- vhost: fix packed ring dequeue offloading
- vhost: fix vid allocation race

1.14.2 19.11.7 Validation

- Red Hat(R) Testing
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 5.2
 - * X540-AT2 NIC(ixgbe, 10G)
 - * Tested on 19.11.7-RC1

– Tests

- * Guest with device assignment(PF) throughput testing(1G hugepage size): PASS
- * Guest with device assignment(PF) throughput testing(2M hugepage size) : PASS
- * Guest with device assignment(VF) throughput testing: PASS
- * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing: PASS
- * PVP vhost-user 2Q throughput testing: PASS
- * PVP vhost-user 1Q - cross numa node throughput testing: PASS
- * Guest with vhost-user 2 queues throughput testing: PASS
- * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect: PASS
- * vhost-user reconnect with dpdk-client, qemu-server: ovs reconnect: PASS
- * PVP 1Q live migration testing: PASS
- * PVP 1Q cross numa node live migration testing: PASS
- * Guest with ovs+dpdk+vhost-user 1Q live migration testing: PASS
- * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M): PASS
- * Guest with ovs+dpdk+vhost-user 2Q live migration testing: PASS
- * Host PF + DPDK testing: PASS
- * Host VF + DPDK testing: PASS

- Nvidia (R) Testing

- functional tests on Mellanox hardware

- * NIC: ConnectX-4 Lx / OS: RHEL7.4 / Driver: MLNX_OFED_LINUX-5.2-2.2.0.0
Firmware: 14.29.2002
- * NIC: ConnectX-5 / OS: RHEL7.4 / Driver: MLNX_OFED_LINUX-5.2-2.2.0.0
Firmware: 16.29.2002

- Basic functionality:

- * Send and receive multiple types of traffic.
- * testpmd xstats counter test.
- * testpmd timestamp test.
- * Changing/checking link status through testpmd.
- * RTE flow tests: Items: eth / vlan / ipv4 / ipv6 / tcp / udp / icmp / gre / nvgre / vxlan ip
in ip / mplsoudp / mplsogre
- * Actions: drop / queue / rss / mark / flag / jump / count / raw_encap / raw_decap /
vxlan_encap / vxlan_decap / NAT / dec_ttl
- * Some RSS tests.
- * VLAN filtering, stripping and insertion tests.
- * Checksum and TSO tests.
- * ptype tests.

- * link_status_interrupt example application tests.
- * l3fwd-power example application tests.
- * Multi-process example applications tests.
- Compilation tests with multiple configurations in the following OS/driver combinations are also passing:
 - * Ubuntu 20.04.2 with MLNX_OFED_LINUX-5.2-2.2.0.0.
 - * Ubuntu 20.04.2 with rdma-core master (a1a9ffb).
 - * Ubuntu 20.04.2 with rdma-core v28.0.
 - * Ubuntu 18.04.5 with rdma-core v17.1.
 - * Ubuntu 18.04.5 with rdma-core master (a1a9ffb) (i386).
 - * Ubuntu 16.04.7 with rdma-core v22.7.
 - * Fedora 32 with rdma-core v33.0.
 - * CentOS 7 7.9.2009 with rdma-core master (a1a9ffb).
 - * CentOS 7 7.9.2009 with MLNX_OFED_LINUX-5.2-2.2.0.0.
 - * CentOS 8 8.3.2011 with rdma-core master (a1a9ffb).
 - * OpenSUSE Leap 15.2 with rdma-core v27.1.
- Intel(R) Testing
 - Basic Intel NIC (ixgbe, i40e) testing
 - * PF (i40e)
 - * PF (ixgbe)
 - * VF (i40e)
 - * VF (ixgbe)
 - * Compile Testing
 - * Intel NIC single core/NIC performance
 - Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance
- Intel(R) Testing with Open vSwitch
 - OVS testing with OVS 2.14.3
 - * I40e: Performance Tests Jumbo frames RSS
 - * Niantic: Performance Tests Jumbo frames RSS
 - * Ice: Performance Tests Jumbo frames RSS
 - * Vhost: Port addition, deletion, jumbo frames and RSS multi-queue tests.
- Canonical(R) Testing
 - Build tests on Ubuntu 21.04

- OVS-DPDK tests on x86_64
 - * performance tests
 - test guest-openvswitch for OVS-5CPU => Pass
 - test guest-dpdk-vhost-user-client-multiq for OVSDPDK-VUC => Pass
 - * VUC endurance tests
 - start stop guests (client) => Pass
 - add/remove ports (client) => Pass

1.14.3 19.11.7 Known Issues

- The UDP fragmentation offload feature of Virtio-net device can not be turned on in the VM. Bug: https://bugzilla.kernel.org/show_bug.cgi?id=207075
- mlx5 VLAN packets will not do RSS. Bug: https://bugs.dpdk.org/show_bug.cgi?id=661

1.15 19.11.8 Release Notes

1.15.1 19.11.8 Fixes

- Regenerate meson.build changes required due to reverts
- Revert “Revert “Revert “build: always link whole DPDK static libraries”””
- Revert “Revert “Revert “build/pkg-config: improve static linking flags”””
- Revert “Revert “Revert “build/pkg-config: move pkg-config file creation”””
- Revert “Revert “Revert “build/pkg-config: output drivers first for static build”””
- Revert “Revert “Revert “build/pkg-config: prevent overlinking”””
- Revert “Revert “Revert “devtools: test static linkage with pkg-config”””

1.15.2 19.11.8 Validation

- Intel(R) Testing
 - Builds
 - Basic Intel NIC (ixgbe, i40e) testing
 - * PF (i40e)
 - * PF (ixgbe)
 - * VF (i40e)
 - * VF (ixgbe)
 - * Compile Testing
 - * Intel NIC single core/NIC performance
 - Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance

- Intel(R) Testing with Open vSwitch
 - OVS build and link testing with OVS 2.13 (make based)
 - OVS build and link testing with OVS 2.14 (make based)
- Canonical(R) Testing
 - Build tests of DPDK & OVS 2.13.3 on Ubuntu 20.04 (meson based)
 - Build tests of DPDK & OVS 2.13.3 on Ubuntu 20.10 (meson based)
 - OVS-DPDK tests on x86_64

1.16 19.11.9 Release Notes

1.16.1 19.11.9 Fixes

- app/bbdev: fix HARQ error messages
- app/crypto-perf: check memory allocation
- app/eventdev: fix lcore parsing skipping last core
- app/eventdev: fix overflow in lcore list parsing
- app: fix exit messages
- app/testpmd: fix bitmap of link speeds when force speed
- app/testpmd: fix build with musl
- app/testpmd: fix DCB forwarding configuration
- app/testpmd: fix DCB re-configuration
- app/testpmd: fix forward lcores number for DCB
- app/testpmd: fix max queue number for Tx offloads
- app/testpmd: fix NVGRE encap configuration
- app/testpmd: fix segment number check
- app/testpmd: fix Tx/Rx descriptor query error log
- app/testpmd: remove unnecessary UDP tunnel check
- bpf: fix JSLT validation
- build: detect execinfo library on Linux
- build: exclude meson files from examples installation
- buildtools: fix build with busybox
- bus/dpaa: fix 64-bit arch detection
- bus/dpaa: fix build with musl
- bus/dpaa: fix statistics reading
- bus/fslmc: fix random portal hangs with qbman 5.0
- bus/fslmc: remove unused debug macro

- common/dpaax/caamflib: fix build with musl
- common/dpaax: fix possible null pointer access
- common/iavf: use pad byte to specify MAC type
- common/qat: increase IM buffer size for GEN3
- common/sfc_efx/base: limit reported MCDI response length
- compress/qat: enable compression on GEN3
- config/ppc: reduce number of cores and NUMA nodes
- crypto/dpaa2_sec: fix close and uninit functions
- crypto/dpaa2_sec: affine the thread portal affinity
- crypto/qat: fix null authentication request
- crypto/qat: fix offset for out-of-place scatter-gather
- doc: fix build with Sphinx 4
- doc: fix HiSilicon copyright syntax
- doc: fix multiport syntax in nfp guide
- doc: fix names of UIO drivers
- doc: fix sphinx rtd theme import in GHA
- drivers/crypto: fix build with -fno-common
- drivers/net: check process type in close operation
- eal: add C++ include guard for reciprocal header
- eal/arm64: fix platform register bit
- eal: fix comment of OS-specific header files
- eal: fix hang in control thread creation
- eal: fix race in control thread creation
- eal: fix service core list parsing
- eal/windows: add cleanup function stub
- eal/windows: fix default thread priority
- ethdev: add missing buses in device iterator
- ethdev: fix redundant flow after RSS expansion
- ethdev: update flow item GTP QFI definition
- ethdev: validate input in EEPROM info
- ethdev: validate input in module EEPROM dump
- ethdev: validate input in register info
- eventdev: fix case to initiate crypto adapter service
- eventdev: fix memory leakage on thread creation failure

- eventdev: remove redundant thread name setting
- event/dpaa2: remove unused macros
- event/octeontx2: fix device reconfigure for single slot
- examples: add eal cleanup to examples
- examples/bbdev: fix header include for musl
- examples/ethtool: remove unused parsing
- examples/flow_classify: fix NUMA check of port and core
- examples/l2fwd-cat: fix NUMA check of port and core
- examples/l2fwd-crypto: fix packet length while decryption
- examples/l2fwd-crypto: skip masked devices
- examples/l3fwd: fix LPM IPv6 subnets
- examples/l3fwd-power: fix empty poll thresholds
- examples/packet_ordering: fix port configuration
- examples/ptpclient: remove wrong comment
- examples/rxtx_callbacks: fix port ID format specifier
- examples/skeleton: fix NUMA check of port and core
- examples/timer: fix time interval
- examples/vhost: check memory table query
- examples/vhost_crypto: remove unused short option
- fbararray: fix log message on truncation error
- net/mlx5: fix flow split combined with counter
- ipc: check malloc sync reply result
- ipc: use monotonic clock
- kni: fix compilation on SLES15-SP3
- license: fix typos
- log/linux: make default output stderr
- mbuf: check shared memory before dumping dynamic space
- mem: fix freeing segments in -huge-unlink mode
- net/af_xdp: fix error handling during Rx queue setup
- net/ark: refactor Rx buffer recovery
- net/ark: update packet director initial state
- net/bnx2x: fix build with GCC 11
- net/bnx2x: fix build with GCC 11
- net/bnxt: check PCI config read

- net/bnxt: drop unused attribute
- net/bnxt: fix build failures after merging patches
- net/bnxt: fix configuring LRO
- net/bnxt: fix device readiness check
- net/bnxt: fix double free in port start failure
- net/bnxt: fix dynamic VNIC count
- net/bnxt: fix firmware fatal error handling
- net/bnxt: fix FW readiness check during recovery
- net/bnxt: fix handling of null flow mask
- net/bnxt: fix health check alarm cancellation
- net/bnxt: fix HWRM and FW incompatibility handling
- net/bnxt: fix link state operations
- net/bnxt: fix memory allocation for command response
- net/bnxt: fix mismatched type comparison in MAC restore
- net/bnxt: fix PCI write check
- net/bnxt: fix PTP support for Thor
- net/bnxt: fix queues per VNIC
- net/bnxt: fix ring count calculation for Thor
- net/bnxt: fix RSS context cleanup
- net/bnxt: fix Rx and Tx timestamps
- net/bnxt: fix Rx buffer posting
- net/bnxt: fix Rx descriptor status
- net/bnxt: fix Rx queue count
- net/bnxt: fix Rx timestamp when FIFO pending bit is set
- net/bnxt: fix single PF per port check
- net/bnxt: fix timesync when PTP is not supported
- net/bnxt: fix Tx length hint threshold
- net/bnxt: fix Tx timestamp init
- net/bnxt: fix VF info allocation
- net/bnxt: fix VNIC configuration
- net/bnxt: fix xstats get
- net/bnxt: prevent device access in error state
- net/bnxt: refactor multi-queue Rx configuration
- net/bnxt: remove unnecessary forward declarations

- net/bnxt: remove unused function parameters
- net/bnxt: remove unused macro
- net/bnxt: reset filter indices on free
- net/bnxt: use prefix on global function
- net/bonding: fix adding itself as its slave
- net/bonding: fix LACP system address check
- net/bonding: fix leak on remove
- net/bonding: fix socket ID check
- net/cxgbe: remove use of uint type
- net/dpaa2: fix getting link status
- net/e1000/base: fix timeout for shadow RAM write
- net/e1000: fix flow error message object
- net/e1000: fix max Rx packet size
- net/e1000: fix Rx error counter for bad length
- net/e1000: remove MTU setting limitation
- net/ena/base: fix type conversions by explicit casting
- net/ena: fix releasing Tx ring mbufs
- net/ena: remove endian swap functions
- net/ena: switch memcpy to optimized version
- net/failsafe: fix RSS hash offload reporting
- net/failsafe: report minimum and maximum MTU
- net: fix comment in IPv6 header
- net/hinic: fix crash in secondary process
- net/hns3: clear hash map on flow director clear
- net/hns3: delete redundant blank line
- net/hns3: fix DCB mode check
- net/hns3: fix DCB reconfiguration
- net/hns3: fix flow control exception
- net/hns3: fix flow control mode
- net/hns3: fix flow counter value
- net/hns3: fix FLR miss detection
- net/hns3: fix handling link update
- net/hns3: fix HW buffer size on MTU update
- net/hns3: fix mailbox error message

- net/hns3: fix MTU config complexity
- net/hns3: fix ordering in secondary process initialization
- net/hns3: fix processing Tx offload flags
- net/hns3: fix querying flow director counter for out param
- net/hns3: fix requested FC mode rollback
- net/hns3: fix rollback after setting PVID failure
- net/hns3: fix secondary process request start/stop Rx/Tx
- net/hns3: fix typos on comments
- net/hns3: fix VF mailbox head field
- net/hns3: fix VMDq mode check
- net/hns3: log time delta in decimal format
- net/hns3: remove meaningless packet buffer rollback
- net/hns3: remove unused macro
- net/hns3: remove unused macros
- net/hns3: remove unused macros
- net/hns3: remove unused mailbox macro and struct
- net/hns3: remove unused VMDq code
- net/hns3: return error on PCI config write failure
- net/hns3: support get device version when dump register
- net/hns3: update HiSilicon copyright syntax
- net/i40e: announce request queue capability in PF
- net/i40e: fix input set field mask
- net/i40e: fix IPv4 fragment offload
- net/i40e: fix lack of MAC type when set MAC address
- net/i40e: fix negative VEB index
- net/i40e: fix parsing packet type for NEON
- net/i40e: fix primary MAC type when starting port
- net/i40e: fix VF RSS configuration
- net/i40e: remove redundant VSI check in Tx queue setup
- net/i40evf: fix packet loss for X722
- net/iavf: fix lack of MAC type when set MAC address
- net/iavf: fix primary MAC type when starting port
- net/iavf: fix TSO max segment size
- net/iavf: fix Tx context descriptor

- net/iavf: fix VF to PF command failure handling
- net/ice/base: cleanup filter list on error
- net/ice/base: fix build with GCC 11
- net/ice/base: fix memory allocation for MAC addresses
- net/ice/base: fix memory allocation wrapper
- net/ice/base: fix payload indicator on ptype
- net/ice: check some functions return
- net/ice: fix disabling promiscuous mode
- net/ice: fix fast mbuf freeing
- net/ice: fix illegal access when removing MAC filter
- net/ice: fix VLAN filter with PF
- net/ice: fix VSI array out of bounds access
- net/ixgbe: fix RSS RETA being reset after port start
- net/ixgbe: fix Rx errors statistics for UDP checksum
- net/kni: check init result
- net/memif: fix Tx bps statistics for zero-copy
- net/mlx4: fix buffer leakage on device close
- net/mlx4: fix leak when configured repeatedly
- net/mlx4: fix RSS action with null hash key
- net/mlx4: fix secondary process initialization ordering
- net/mlx5: fix drop action for Direct Rules/Verbs
- net/mlx5: fix leak when configured repeatedly
- net/mlx5: fix metadata item validation for ingress flows
- net/mlx5: fix probing device in legacy bonding mode
- net/mlx5: fix Rx metadata leftovers
- net/mlx5: fix Rx segmented packets on mbuf starvation
- net/mlx5: fix secondary process initialization ordering
- net/mlx5: remove drop queue function prototypes
- net/mlx5: support RSS expansion for IPv6 GRE
- net/nfp: fix reporting of RSS capabilities
- net/octeontx2: fix VLAN filter
- net/pcap: fix format string
- net/qede: accept bigger RSS table
- net/qede: reduce log verbosity

- net/tap: check ioctl on restore
- net/tap: fix build with GCC 11
- net/tap: fix interrupt vector array size
- net/virtio: fix interrupt unregistering for listening socket
- net/virtio: fix vectorized Rx queue rearm
- power: do not skip saving original P-state governor
- power: fix sanity checks for guest channel read
- power: remove duplicated symbols from map file
- power: save original ACPI governor always
- raw/ntb: check memory allocations
- raw/ntb: check SPAD user index
- raw/skeleton: add missing check after setting attribute
- Revert “kni: fix compilation on SLES15-SP3”
- service: clean references to removed symbol
- stack: allow lock-free only on relevant architectures
- test/bpf: fix error message
- test: check flow classifier creation
- test: check thread creation
- test/cmdline: fix inputs array
- test/cmdline: silence clang 12 warning
- test/crypto: fix auth-cipher compare length in OOP
- test/crypto: fix build with GCC 11
- test/distributor: fix burst flush on worker quit
- test/distributor: fix worker notification in burst mode
- test/event: fix timeout accuracy
- test: fix autotest handling of skipped tests
- test: fix build with GCC 11
- test: fix division by zero
- test: fix TCP header initialization
- test/kni: check init result
- test/kni: fix a comment
- test/mempool: fix object initializer
- test/power: add delay before checking CPU frequency
- test/power: add turbo mode to frequency check

- test/power: fix CPU frequency check
- test/power: fix low frequency test when turbo enabled
- test/power: fix turbo test
- test/power: round CPU frequency to check
- test: proceed if timer subsystem already initialized
- test/table: fix build with GCC 11
- test/timer: check memzone allocation
- vdpa/ifc: check PCI config read
- vfio: do not merge contiguous areas
- vfio: fix API description
- vfio: fix DMA mapping granularity for IOVA as VA
- vfio: fix duplicated user mem map
- vfio: fix stdbool usage without include
- vhost: fix batch dequeue potential buffer overflow
- vhost: fix initialization of temporary header
- vhost: fix packed ring potential buffer overflow
- vhost: fix queue initialization
- vhost: fix split ring potential buffer overflow

1.16.2 19.11.9 Validation

- [Intel(R) Testing](<http://mails.dpdk.org/archives/stable/2021-July/032349.html>)
 - Builds
 - Basic Intel NIC (ixgbe, i40e) testing
 - * PF (i40e)
 - * PF (ixgbe)
 - * VF (i40e)
 - * VF (ixgbe)
 - * Compile Testing
 - * Intel NIC single core/NIC performance
 - Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance
- [Intel(R) Testing with Open vSwitch](<http://mails.dpdk.org/archives/stable/2021-June/031702.html>)
 - OVS build and link testing with OVS 2.13 (make based)

- OVS build and link testing with OVS 2.14 (make based)
- VSPERF tests
 - * phy2phy_tput
 - * pvp_tput
 - * pvp_cont
 - * pvvp_tput
 - * pvpv_cont
 - * ovspdk_jumbo_p2p
 - * ovspdk_jumbo_pvp
 - * ovspdk_jumbo_p2p_upper_bound
 - * ovspdk_mq_pvp_rxqs_testpmd
 - * ovspdk_mq_pvp_rxqs_linux_bridge
- [Canonical(R) Testing](<http://mails.dpdk.org/archives/stable/2021-July/031997.html>)
 - Build tests of DPDK & OVS 2.13.3 on Ubuntu 20.04 (meson based)
 - Build tests of DPDK & OVS 2.13.3 on Ubuntu 20.10 (meson based)
 - Functional and performance tests based on OVS-DPDK on x86_64
 - Autopkgtests for DPDK and OpenvSwitch
- [Red Hat(R) Testing](<http://mails.dpdk.org/archives/stable/2021-June/031619.html>)
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 6.0
 - * X540-AT2 NIC(ixgbe, 10G)
 - * Tested on 19.11.9-RC3
 - Tests
 - * Guest with device assignment(PF) throughput testing(1G hugepage size): PASS
 - * Guest with device assignment(PF) throughput testing(2M hugepage size) : PASS
 - * Guest with device assignment(VF) throughput testing: PASS
 - * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing: PASS
 - * PVP vhost-user 2Q throughput testing: PASS
 - * PVP vhost-user 1Q - cross numa node throughput testing: PASS
 - * Guest with vhost-user 2 queues throughput testing: PASS
 - * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect: PASS
 - * vhost-user reconnect with dpdk-client, qemu-server: ovs reconnect: PASS

- * PVP 1Q live migration testing: PASS
- * PVP 1Q cross numa node live migration testing: PASS
- * Guest with ovs+dpdk+vhost-user 1Q live migration testing: PASS
- * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M): PASS
- * Guest with ovs+dpdk+vhost-user 2Q live migration testing: PASS
- * Host PF + DPDK testing: PASS
- * Host VF + DPDK testing: PASS
- [Nvidia (R) Testing](<http://mails.dpdk.org/archives/stable/2021-July/032381.html>)
 - functional tests on Mellanox hardware
 - * NIC: ConnectX-4 Lx / OS: Ubuntu 20.04 LTS / Driver: MLNX_OFED_LINUX-5.3-1.0.0.1 / Firmware: 14.30.1004
 - * NIC: ConnectX-5 / OS: Ubuntu 20.04 LTS / Driver: MLNX_OFED_LINUX-5.3-1.0.0.1 / Firmware: 16.30.1004
 - Basic functionality:
 - * Send and receive multiple types of traffic.
 - * testpmd xstats counter test.
 - * testpmd timestamp test.
 - * Changing/checking link status through testpmd.
 - * RTE flow tests: Items: eth / vlan / ipv4 / ipv6 / tcp / udp / icmp / gre / nvgre / vxlan ip in ip / mplsoudp / mplsogre
 - * Actions: drop / queue / rss / mark / flag / jump / count / raw_encap / raw_decap / vxlan_encap / vxlan_decap / NAT / dec_ttl
 - * Some RSS tests.
 - * VLAN filtering, stripping and insertion tests.
 - * Checksum and TSO tests.
 - * ptype tests.
 - * link_status_interrupt example application tests.
 - * l3fwd-power example application tests.
 - * Multi-process example applications tests.
 - Compilation tests with multiple configurations in the following OS/driver combinations are also passing:
 - * Ubuntu 20.04.2 with MLNX_OFED_LINUX-5.3-1.0.0.1.
 - * Ubuntu 20.04.2 with rdma-core master (a66e2a5).
 - * Ubuntu 20.04.2 with rdma-core v28.0.
 - * Ubuntu 18.04.5 with rdma-core v17.1.
 - * Ubuntu 18.04.5 with rdma-core master (a66e2a5) (i386).

- * Ubuntu 16.04.7 with rdma-core v22.7.
- * Fedora 34 with rdma-core v35.0.
- * Fedora 35 (Rawhide) with rdma-core v35.0 (only with gcc).
- * CentOS 7 7.9.2009 with rdma-core master (a66e2a5).
- * CentOS 7 7.9.2009 with MLNX_OFED_LINUX-5.3-1.0.0.1.
- * CentOS 8 8.3.2011 with rdma-core master (7f2d460).
- * OpenSUSE Leap 15.3 with rdma-core v31.0.

1.16.3 19.11.9 Known Issues

This time there are a few issues which were identified in the tests. But none of them is an in-place regression. So far they are all issues with newer compilers, toolchains and kernels that happen to be incompatible with the 19.11.x code.

- Due to a kernel patch backport in SUSE Linux Enterprise Server 15 SP3 [compilation of kni fails there](https://bugs.dpdk.org/show_bug.cgi?id=728). A fix would break existing SP2 installation and [further options are discussed](<http://mails.dpdk.org/archives/stable/2021-June/031210.html>)
- A [Build issue with clang 12.0.1](https://bugs.dpdk.org/show_bug.cgi?id=745)
- A make [build issue with Fedora 34 with GCC11](https://bugs.dpdk.org/show_bug.cgi?id=744)
- Meson based [build issues with ICC-19.1.1](https://bugs.dpdk.org/show_bug.cgi?id=747)

19.11.10 Release Notes

1.16.4 19.11.10 Fixes

- app/crypto-perf: fix out-of-place mempool allocation
- app/test: fix IPv6 header initialization
- app/testpmd: change port link speed without stopping all
- app/testpmd: fix help string for port reset
- app/testpmd: fix IPv4 checksum
- app/testpmd: fix MAC address after port reset
- app/testpmd: fix offloads for newly attached port
- app/testpmd: fix Tx checksum calculation for tunnel
- bitmap: fix buffer overrun in bitmap init
- bus: clarify log for non-NUMA-aware devices
- bus/dpaa: fix freeing in FMAN interface destructor
- common/mlx5: fix compatibility with OFED port query API
- common/mlx5: fix Netlink port name padding in probing
- common/mlx5: use new port query API if available
- cryptodev: fix freeing after device release
- crypto/mvsam: fix AES-GCM session parameters

- crypto/mvsam: fix capabilities
- crypto/mvsam: fix options parsing
- crypto/mvsam: fix session data reset
- crypto/octeontx: fix freeing after device release
- crypto/qat: fix Arm build with special memcopy
- devtools: fix file listing in maintainers check
- distributor: fix 128-bit write alignment
- doc: add limitation for ConnectX-4 with L2 in mlx5 guide
- doc: announce common prefix for ethdev
- doc: fix default burst size in testpmd
- doc: fix spelling
- drivers/net: fix memzone allocations for DMA memory
- ethdev: fix doc of flow action
- eventdev: fix event port setup in tx adapter
- flow_classify: fix leaking rules on delete
- (github-cpaelzer-stable/19.11) telemetry: fix race in telemetry control thread creation
- (HEAD -> 19.11, tag: v19.11.10-rc2, origin-stable/19.11) version: 19.11.10-rc2
- ipc: stop mp control thread on cleanup
- kni: fix crash on userspace VA for segmented packets
- kni: fix mbuf allocation for kernel side use
- mempool/octeontx2: fix shift calculation
- net/bnxt: check access to possible null pointer
- net/bnxt: cleanup code
- net/bnxt: clear cached statistics
- net/bnxt: detect bad opaque in Rx completion
- net/bnxt: fix auto-negotiation on Whitney+
- net/bnxt: fix check for PTP support in FW
- net/bnxt: fix error handling in VNIC prepare
- net/bnxt: fix error messages in VNIC prepare
- net/bnxt: fix missing barriers in completion handling
- net/bnxt: fix nested lock during bonding
- net/bnxt: fix ring and context memory allocation
- net/bnxt: fix Rx burst size constraint
- net/bnxt: fix Rx interrupt setting

- net/bnxt: fix scalar Tx completion handling
- net/bnxt: fix Tx descriptor status implementation
- net/bnxt: fix typo in log message
- net/bnxt: improve probing log message
- net/bnxt: invoke device removal event on recovery failure
- net/bnxt: remove unnecessary code
- net/bnxt: remove unnecessary comment
- net/bnxt: set flow error after tunnel redirection free
- net/bnxt: set flow error when free filter not available
- net/bnxt: use common function to free VNIC resource
- net/bnxt: workaround spurious zero stats in Thor
- net/bonding: check flow setting
- net/bonding: fix error message on flow verify
- net/ena: enable multi-segment in Tx offload flags
- net/ena: trigger reset on Tx prepare failure
- net/hinic/base: fix LRO
- net/hinic: increase protection of the VLAN
- net/hns3: fix delay for waiting to stop Rx/Tx
- net/hns3: fix filter parsing comment
- net/hns3: fix Tx prepare after stop
- net/hns3: fix VLAN strip log
- net/hns3: increase VF reset retry maximum
- net/i40e: fix descriptor scan on Arm
- net/i40e: fix multi-process shared data
- net/iavf: fix RSS key access out of bound
- net/iavf: fix Tx threshold check
- net/ice/base: fix first profile mask
- net/ice/base: revert change of first profile mask
- net/ice: fix default RSS key generation
- net/ice: fix memzone leak when firmware is missing
- net/ixgbe: fix flow entry access after freeing
- net/mlx5: fix incorrect r/w lock usage in DMA unmap
- net/mlx5: fix IPIP multi-tunnel validation
- net/mlx5: fix match MPLS over GRE with key

- net/mlx5: fix MPLS RSS expansion
- net/mlx5: fix overflow in mempool argument
- net/mlx5: fix representor interrupt handler
- net/mlx5: fix RoCE LAG bond device probing
- net/mlx5: fix RSS flow item expansion for GRE key
- net/mlx5: fix Rx/Tx queue checks
- net/mlx5: fix switchdev mode recognition
- net/mlx5: fix typo in vectorized Rx comments
- net/mlx5: limit implicit MPLS RSS expansion over GRE
- net/mlx5: limit inner RSS expansion for MPLS
- net/mlx5: remove redundant operations in NEON Rx
- net/mlx5: remove unsupported flow item MPLS over IP
- net/mlx5: workaround drop action with old kernel
- net/mvpp2: fix configured state dependency
- net/mvpp2: fix port speed overflow
- net/octeontx2: fix default MCAM allocation size
- net/octeontx2: fix flow creation limit on CN98xx
- net/octeontx2: use runtime LSO format indices
- net/octeontx/base: fix debug build with clang
- net/pfe: remove unnecessary null check
- net/sfc: fix MAC stats lock in xstats query by ID
- net/sfc: fix MAC stats update for stopped device
- net/sfc: fix reading adapter state without locking
- net/sfc: fix xstats query by ID according to ethdev
- net/sfc: fix xstats query by unsorted list of IDs
- net/softnic: fix connection memory leak
- net/softnic: fix memory leak as profile is freed
- net/virtio: fix aarch32 build
- net/virtio: fix refill order in packed ring datapath
- net/virtio: report maximum MTU in device info
- power: fix namespace for internal struct
- rib: fix max depth IPv6 lookup
- (tag: v19.11.10-rc1) version: 19.11.10-rc1
- test/crypto: fix mbuf reset after null check

- test/crypto: fix mempool size for session-less
- test/crypto: fix typo in AES case
- test/crypto: fix typo in ESN case
- test/mbuf: fix virtual address conversion
- tests/eal: fix memory leak
- vhost/crypto: check request pointer before dereference
- vhost: fix crash on reconnect
- vhost: fix missing guest pages table NUMA realloc
- vhost: fix missing memory table NUMA realloc
- vhost: utilize dynamic memory allocator

1.16.5 19.11.10 Validation

- [Intel(R) Testing](<http://mails.dpdk.org/archives/stable/2021-August/033130.html>)
 - Builds
 - Basic Intel NIC (ixgbe, i40e) testing
 - * PF (i40e)
 - * PF (ixgbe)
 - * VF (i40e)
 - * VF (ixgbe)
 - * Compile Testing
 - * Intel NIC single core/NIC performance
 - Basic cryptodev and virtio testing
 - * vhost/virtio basic loopback, PVP and performance test
 - * cryptodev Function/Performance
- [Intel(R) Testing with Open vSwitch]
 - OVS build and link testing with OVS 2.13 (make based)
 - OVS build and link testing with OVS 2.14 (make based)
 - VSPERF tests
 - * phy2phy_tput
 - * pvp_tput
 - * pvp_cont
 - * pvvp_tput
 - * pvvp_cont
 - * ovspdk_jumbo_p2p
 - * ovspdk_jumbo_pvp

- * ovswdpdk_jumbo_p2p_upper_bound
- * ovswdpdk_mq_pvp_rxqs_testpmd
- * ovswdpdk_mq_pvp_rxqs_linux_bridge
- [Canonical(R) Testing](<http://mails.dpdk.org/archives/stable/2021-August/033126.html>)
 - Build tests of DPDK & OVS 2.13.3 on Ubuntu 20.04 (meson based)
 - Build tests of DPDK & OVS 2.13.3 on Ubuntu 20.10 (meson based)
 - Functional and performance tests based on OVS-DPDK on x86_64
 - Autopkgtests for DPDK and OpenvSwitch
- [Red Hat(R) Testing](<http://mails.dpdk.org/archives/stable/2021-September/033218.html>)
 - Platform
 - * RHEL 8
 - * Kernel 4.18
 - * Qemu 6.0
 - * X540-AT2 NIC(ixgbe, 10G)
 - * Tested on 19.11.9-RC3
 - Tests
 - * Guest with device assignment(PF) throughput testing(1G hugepage size): PASS
 - * Guest with device assignment(PF) throughput testing(2M hugepage size) : PASS
 - * Guest with device assignment(VF) throughput testing: PASS
 - * PVP (host dpdk testpmd as vswitch) 1Q: throughput testing: PASS
 - * PVP vhost-user 2Q throughput testing: PASS
 - * PVP vhost-user 1Q - cross numa node throughput testing: PASS
 - * Guest with vhost-user 2 queues throughput testing: PASS
 - * vhost-user reconnect with dpdk-client, qemu-server: qemu reconnect: PASS
 - * vhost-user reconnect with dpdk-client, qemu-server: ovs reconnect: PASS
 - * PVP 1Q live migration testing: PASS
 - * PVP 1Q cross numa node live migration testing: PASS
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing: PASS
 - * Guest with ovs+dpdk+vhost-user 1Q live migration testing (2M): PASS
 - * Guest with ovs+dpdk+vhost-user 2Q live migration testing: PASS
 - * Host PF + DPDK testing: PASS
 - * Host VF + DPDK testing: PASS
- [Nvidia (R) Testing](<http://mails.dpdk.org/archives/stable/2021-August/033131.html>)
 - functional tests on Mellanox hardware

- * NIC: ConnectX-4 Lx / OS: Ubuntu 20.04 LTS / Driver: MLNX_OFED_LINUX-5.4-1.0.3.0 / Firmware: 14.31.1014
 - * NIC: ConnectX-4 Lx / OS: Ubuntu 20.04 LTS / kernel: 5.14.0-rc7 / Driver: rdma-core v36.0 / Firmware: 14.31.1014
 - * NIC: ConnectX-5 / OS: Ubuntu 20.04 LTS / Driver: MLNX_OFED_LINUX-5.4-1.0.3.0 / Firmware: 16.31.1014
 - * NIC: ConnectX-5 / OS: Ubuntu 20.04 LTS / kernel: 5.14.0-rc7 / Driver: v36.0 / Firmware: 16.31.1014
- Basic functionality:
- * Send and receive multiple types of traffic.
 - * testpmd xstats counter test.
 - * testpmd timestamp test.
 - * Changing/checking link status through testpmd.
 - * RTE flow tests: Items: eth / vlan / ipv4 / ipv6 / tcp / udp / icmp / gre / nvgre / vxlan ip in ip / mplsoudp / mplsogre
 - * Actions: drop / queue / rss / mark / flag / jump / count / raw_encap / raw_decap / vxlan_encap / vxlan_decap / NAT / dec_ttl
 - * Some RSS tests.
 - * VLAN filtering, stripping and insertion tests.
 - * Checksum and TSO tests.
 - * ptype tests.
 - * link_status_interrupt example application tests.
 - * l3fwd-power example application tests.
 - * Multi-process example applications tests.
- Compilation tests with multiple configurations in the following OS/driver combinations are also passing:
- * Ubuntu 20.04.2 with MLNX_OFED_LINUX-5.4-1.0.3.0.
 - * Ubuntu 20.04.2 with rdma-core master (64d1ae5).
 - * Ubuntu 20.04.2 with rdma-core v28.0.
 - * Ubuntu 18.04.5 with rdma-core v17.1.
 - * Ubuntu 18.04.5 with rdma-core master (5b0f5b2) (i386).
 - * Ubuntu 16.04.7 with rdma-core v22.7.
 - * Fedora 34 with rdma-core v36.0.
 - * Fedora 35 (Rawhide) with rdma-core v36.0 (only with gcc).
 - * CentOS 7 7.9.2009 with rdma-core master (64d1ae5).
 - * CentOS 7 7.9.2009 with MLNX_OFED_LINUX-5.4-1.0.3.0.
 - * CentOS 8 8.3.2011 with rdma-core master (64d1ae5).

* OpenSUSE Leap 15.3 with rdma-core v31.0.

1.16.6 19.11.10 Known Issues

A few issues are still present, but none of them is an in-place regression. So far they are all issues with newer compilers, toolchains and kernels that happen to be incompatible with the 19.11.x code.

- Due to a kernel patch backport in SUSE Linux Enterprise Server 15 SP3 [compilation of kni fails there](https://bugs.dpdk.org/show_bug.cgi?id=728). A fix would break existing SP2 installation and [further options are discussed](<http://mails.dpdk.org/archives/stable/2021-June/031210.html>)
- A [Build issue with clang 12.0.1](https://bugs.dpdk.org/show_bug.cgi?id=745)
- A make [build issue with Fedora 34 with GCC11](https://bugs.dpdk.org/show_bug.cgi?id=744)
- Meson based [build issues with ICC-19.1.1](https://bugs.dpdk.org/show_bug.cgi?id=747)

2.1 New Features

- **EAL will now pick IOVA as VA mode as the default in most cases.**

Previously, the preferred default IOVA mode was selected to be IOVA as PA. The behavior has now been changed to handle IOVA mode detection in a more complex manner, and will default to IOVA as VA in most cases.

- **Added MCS lock.**

MCS lock provides scalability by spinning on a CPU/thread local variable which avoids expensive cache bouncing. It provides fairness by maintaining a list of acquirers and passing the lock to each CPU/thread in the order they acquired the lock.

- **Updated the EAL Pseudo-random Number Generator.**

The `lrand48()` based `rte_rand()` function is replaced with a DPDK-native combined Linear Feedback Shift Register (LFSR) pseudo-random number generator (PRNG).

This new PRNG implementation is multi-thread safe, provides higher-quality pseudo-random numbers (including full 64 bit support) and improved performance.

In addition, `<rte_random.h>` is extended with a new function `rte_rand_max()` which supplies unbiased, bounded pseudo-random numbers.

- **Updated the Broadcom bnxt PMD.**

Updated the Broadcom bnxt PMD. The major enhancements include:

- Performance optimizations in non-vector Tx path.
- Added support for SSE vector mode.
- Updated HWRM API to version 1.10.0.91.

- **Added support for Broadcom NetXtreme-E BCM57500 Ethernet controllers.**

Added support to the Broadcom bnxt PMD for the BCM57500 (a.k.a. “Thor”) family of Ethernet controllers. These controllers support link speeds up to 200Gbps, 50G PAM-4, and PCIe 4.0.

- **Added Huawei hinic PMD.**

Added the new `hinic` net driver for Huawei Intelligent PCIE Network Adapters based on the Huawei Ethernet Controller Hi1822. See the `../nics/hinic` guide for more details on this new driver.

- **Updated the Intel ice driver.**

Updated the Intel ice driver with new features and improvements, including:

- Enabled Tx outer/inner L3/L4 checksum offload.
- Enabled generic filter framework and supported switch filter.
- Supported UDP tunnel port add.

• **Updated the Intel i40e driver.**

Updated the Intel i40e driver with new features and improvements, including:

- Added support for MARK + RSS action in `rte_flow` (non-vector RX path only)

• **Updated Mellanox mlx5 driver.**

Updated Mellanox mlx5 driver with new features and improvements, including:

- Updated the packet header modification feature. Added support of TCP header sequence number and acknowledgment number modification.
- Added support for match on ICMP/ICMP6 code and type.
- Added support for matching on GRE's key and C,K,S present bits.
- Added support for IP-in-IP tunnel.
- Accelerated flows with count action creation and destroy.
- Accelerated flows counter query.
- Improved Tx datapath performance with enabled HW offloads.
- Added support for LRO.

• **Updated Solarflare network PMD.**

Updated the Solarflare `sf_c_efx` driver with changes including:

- Added support for Rx interrupts.

• **Added memif PMD.**

Added a new Shared Memory Packet Interface (`memif`) PMD. See the `./nics/memif` guide for more details on this new driver.

• **Updated the AF_XDP PMD.**

Updated the AF_XDP PMD. The new features include:

- Enabled zero copy through `mbuf`'s external memory mechanism to achieve higher performance.
- Added multi-queue support to allow one `af_xdp` vdev with multiple netdev queues.
- Enabled “`need_wakeup`” feature which can provide efficient support for the use case where the application and driver executing on the same core.

• **Enabled infinite Rx in the PCAP PMD.**

Added an infinite Rx feature to the PCAP PMD which allows packets in the Rx PCAP to be received repeatedly at a high rate. This can be useful for quick performance testing of DPDK apps.

- **Enabled receiving no packet in the PCAP PMD.**

Added function to allow users to run the PCAP PMD without receiving any packets on PCAP Rx. When the function is called, a dummy queue is created for each Tx queue argument passed.

- **Added a FPGA_LTE_FEC bbdev PMD.**

Added a new `fpga_lte_fec` bbdev driver for the Intel® FPGA PAC (Programmable Acceleration Card) N3000. See the `../bbdevs/fpga_lte_fec` BBDEV guide for more details on this new driver.

- **Updated the TURBO_SW bbdev PMD.**

Updated the `turbo_sw` bbdev driver with changes including:

- Added option to build the driver with or without dependency of external SDK libraries.
- Added support for 5G NR encode/decode operations.

- **Updated the Intel QuickAssist Technology (QAT) symmetric crypto PMD.**

Added support for digest-encrypted cases where digest is appended to the data.

- **Added the Intel QuickData Technology PMD.**

Added a PMD for the Intel® QuickData Technology, part of Intel® I/O Acceleration Technology (Intel I/OAT), which allows data copies to be done by hardware instead of via software, reducing cycles spent copying large blocks of data in applications.

- **Added Marvell OCTEON TX2 drivers.**

Added the new `ethdev`, `eventdev`, `mempool`, `eventdev` Rx adapter, `eventdev` Tx adapter, `eventdev` Timer adapter and `rawdev` DMA drivers for various HW co-processors available in OCTEON TX2 SoC.

See `../platform/octeontx2` and driver information:

- `../nics/octeontx2`
- `../mempool/octeontx2`
- `../eventdevs/octeontx2`
- `../rawdevs/octeontx2_dma`

- **Introduced the Intel NTB PMD.**

Added a PMD for Intel NTB (Non-transparent Bridge). This PMD implements a handshake between two separate hosts and can share local memory for peer host to directly access.

- **Updated the IPsec library and IPsec Security Gateway application.**

Added the following features to `librte_ipsec`. Corresponding changes are also added in the `ipsec-secgw` sample application.

- ECN and DSCP field header reconstruction as per RFC4301.
- Transport mode with IPv6 extension headers.
- Support packets with multiple segments.

- **Updated telemetry library for global metrics support.**

Updated `librte_telemetry` to fetch the global metrics from the `librte_metrics` library.

- **Added new telemetry mode for l3fwd-power application.**

Added a telemetry mode to the `l3fwd-power` application to report application level busyness, empty and full polls of `rte_eth_rx_burst()`.

- **Updated the pdump application.**

Add support for pdump to exit with primary process.

- **Updated test-compress-perf tool application.**

Added a multiple cores feature to the compression perf tool application.

2.2 Removed Items

- Removed KNI ethtool, `CONFIG_RTE_KNI_KMOD_ETHTOOL`, support.
- build: armv8 crypto extension is disabled.

2.3 API Changes

- The `rte_mem_config` structure has been made private. New accessor `rte_mcfg_*` functions were introduced to provide replacement for direct access to the shared mem config.
- The network structures, definitions and functions have been prefixed by `rte_` to resolve conflicts with libc headers.
- malloc: The function `rte_malloc_set_limit()` was never implemented. It is deprecated and will be removed in a future release.
- cryptodev: the `uint8_t *data` member of the `key` structure in the `xforms` structure (`rte_crypto_cipher_xform`, `rte_crypto_auth_xform`, and `rte_crypto_aead_xform`) have been changed to `const uint8_t *data`.
- eventdev: No longer marked as experimental.

The eventdev functions are no longer marked as experimental, and have become part of the normal DPDK API and ABI. Any future ABI changes will be announced at least one release before the ABI change is made. There are no ABI breaking changes planned.

- ip_frag: The IP fragmentation library converts input mbuf into fragments using input MTU size via the `rte_ipv4_fragment_packet()` interface. Once fragmentation is done, each `mbuf->ol_flags` are set to enable IP checksum H/W offload irrespective of the platform capability. Cleared IP checksum H/W offload flag from the library. The application must set this flag if it is supported by the platform and application wishes to use it.
- ip_frag: IP reassembly library converts the list of fragments into a reassembled packet via `rte_ipv4_frag_reassemble_packet()` interface. Once reassembly is done, `mbuf->ol_flags` are set to enable IP checksum H/W offload irrespective of the platform capability. Cleared IP checksum H/W offload flag from the library. The application must set this flag if it is supported by the platform and application wishes to use it.
- sched: Macros `RTE_SCHED_QUEUES_PER_TRAFFIC_CLASS` and `RTE_SCHED_PIPE_PROFILES_PER_PORT` are removed for flexible configuration of pipe traffic classes and their queues size, and for runtime configuration of the maximum number of pipe profiles, respectively. In addition, the `wrr_weights` field of struct `rte_sched_pipe_params` is modified to be used only for best-effort tc, and the `qsize` field of struct `rte_sched_port_params` is changed to allow different sizes for each queue.

2.4 ABI Changes

- eventdev: Event based Rx adapter callback

The mbuf pointer array in the event eth Rx adapter callback has been replaced with an event array. Using an event array allows the application to change attributes of the events enqueued by the SW adapter.

The callback can drop packets and populate a callback argument with the number of dropped packets. Add a Rx adapter stats field to keep track of the total number of dropped packets.

- cryptodev: New member in `rte_cryptodev_config` to allow applications to disable features supported by the crypto device. Only the following features would be allowed to be disabled this way,

- `RTE_CRYPTODEV_FF_SYMMETRIC_CRYPTO`.
- `RTE_CRYPTODEV_FF_ASYMMETRIC_CRYPTO`.
- `RTE_CRYPTODEV_FF_SECURITY`.

Disabling unused features would facilitate efficient usage of HW/SW offload.

- bbdev: New operations and parameters have been added to support new 5G NR operations. The bbdev ABI is still kept experimental.
- rawdev: The driver names have been changed to `librte_rawdev_*`. Now they all have the same prefix, and same name with make and meson builds.

2.5 Shared Library Versions

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

```

librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
librte_bpf.so.1
librte_bus_dpaa.so.2
librte_bus_fslmc.so.2
librte_bus_ifpga.so.2
librte_bus_pci.so.2
librte_bus_vdev.so.2
librte_bus_vmbus.so.2
librte_cfgfile.so.2
librte_cmdline.so.2
librte_compressdev.so.1
+ librte_cryptodev.so.8
librte_distributor.so.1
+ librte_eal.so.11
librte_efd.so.1
librte_ethdev.so.12
+ librte_eventdev.so.7
librte_flow_classify.so.1
librte_gro.so.1
librte_gso.so.1
librte_hash.so.2
librte_ip_frag.so.1
librte_ipsec.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.5
librte_member.so.1
librte_mempool.so.5
librte_meter.so.3
librte_metrics.so.1
librte_net.so.1
librte_pci.so.1
librte_pdump.so.3
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_dpaa2_qdma.so.1
librte_pmd_ring.so.2
librte_pmd_softnic.so.1
librte_pmd_vhost.so.2
librte_port.so.3
librte_power.so.1
librte_rawdev.so.1
librte_rcu.so.1
librte_reorder.so.1
librte_ring.so.2
+ librte_sched.so.3
librte_security.so.2
librte_stack.so.1
librte_table.so.3
librte_timer.so.1
librte_vhost.so.4
```

2.6 Known Issues

- **Unsuitable IOVA mode may be picked as the default.**

Not all kernel drivers and not all devices support all IOVA modes. EAL will attempt to pick a reasonable default based on a number of factors, but there may be cases where the default is unsuitable.

It is recommended to use the `-iova-mode` command-line parameter if the default is not suitable.

2.7 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
 - CPU
 - * Intel(R) Atom(TM) CPU C3758 @ 2.20GHz
 - * Intel(R) Xeon(R) CPU D-1541 @ 2.10GHz
 - * Intel(R) Xeon(R) CPU D-1553N @ 2.30GHz
 - * 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 E5-2699 v4 @ 2.20GHz
 - * Intel(R) Xeon(R) Gold 6139 CPU @ 2.30GHz
 - * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz

- * Intel(R) Xeon(R) Platinum 8280M CPU @ 2.70GHz
- OS:
 - * CentOS 7.6
 - * Fedora 30
 - * FreeBSD 12.0
 - * Red Hat Enterprise Linux Server release 8.0
 - * Red Hat Enterprise Linux Server release 7.6
 - * Suse12SP3
 - * Ubuntu 16.04
 - * Ubuntu 16.10
 - * Ubuntu 18.04
 - * Ubuntu 19.04
- NICs:
 - * Intel(R) 82599ES 10 Gigabit Ethernet Controller
 - Firmware version: 0x61bf0001
 - Device id (pf/vf): 8086:10fb / 8086:10ed
 - Driver version: 5.6.1 (ixgbe)
 - * Intel(R) Corporation Ethernet Connection X552/X557-AT 10GBASE-T
 - Firmware version: 0x800003e7
 - Device id (pf/vf): 8086:15ad / 8086:15a8
 - Driver version: 5.1.0 (ixgbe)
 - * Intel Corporation Ethernet Controller 10G X550T
 - Firmware version: 0x80000482
 - Device id (pf): 8086:1563
 - Driver version: 5.6.1 (ixgbe)
 - * Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
 - Firmware version: 7.00 0x80004cdb
 - Device id (pf/vf): 8086:1572 / 8086:154c
 - Driver version: 2.9.21 (i40e)
 - * Intel(R) Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 4.10 0x80001a3c
 - Device id (pf/vf): 8086:37d0 / 8086:37cd
 - Driver version: 2.9.21 (i40e)
 - * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)

- Firmware version: 7.00 0x80004cf8
- Device id (pf/vf): 8086:158b / 8086:154c
- Driver version: 2.9.21 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
 - Firmware version: 7.00 0x80004c97
 - Device id (pf/vf): 8086:1583 / 8086:154c
 - Driver version: 2.9.21 (i40e)
- * Intel(R) Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x80000cbc
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- * Intel Corporation I210 Gigabit Network Connection
 - Firmware version: 3.25, 0x800006eb
 - Device id (pf): 8086:1533
 - Driver version: 5.4.0-k(igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
 - * 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 7.6 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.3 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.2 (Maipo)
 - * Ubuntu 19.04
 - * Ubuntu 18.10
 - * Ubuntu 18.04
 - * Ubuntu 16.04

- * SUSE Linux Enterprise Server 15
- OFED:
 - * MLNX_OFED 4.6-1.0.1.1
 - * MLNX_OFED 4.6-4.1.2.0
- 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: 2.42.5000
 - * Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
 - * Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
 - * Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
 - * Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
 - * Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)

- Host interface: PCI Express 3.0 x8
- Device ID: 15b3:1013
- Firmware version: 12.25.6406 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
- * 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: 12.25.6406 and above
 - Firmware version: 12.25.6406 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
- * Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.25.6406 and above
- * Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.25.6406 and above
- * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.25.6406 and above
- * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.25.6406 and above
- * Mellanox(R) ConnectX(R)-5 Ex EN 100G MCX516A-CDAT (2x100G)
 - Host interface: PCI Express 4.0 x16

- Device ID: 15b3:1019
- Firmware version: 16.25.6406 and above
- Mellanox(R) BlueField SmartNIC
 - Mellanox(R) BlueField SmartNIC MT416842 (2x25G)
 - * Host interface: PCI Express 3.0 x16
 - * Device ID: 15b3:a2d2
 - * Firmware version: 18.25.6600
 - SoC Arm cores running OS:
 - * CentOS Linux release 7.5.1804 (AltArch)
 - * MLNX_OFED 4.6-3.5.8.0
 - DPDK application running on Arm cores inside SmartNIC
- IBM Power 9 platforms with Mellanox(R) NICs combinations
 - CPU:
 - * POWER9 2.2 (pvr 004e 1202) 2300MHz
 - OS:
 - * Ubuntu 18.04.1 LTS (Bionic Beaver)
 - NICs:
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.25.6406
 - OFED:
 - * MLNX_OFED 4.6-4.1.2.0
- ARMv8 SoC combinations from Marvell (with integrated NICs)
 - SoC:
 - * CN83xx, CN96xx, CNF95xx, CN93xx
 - OS (Based on Marvell OCTEON TX SDK 10.0):
 - * Arch Linux
 - * Buildroot 2018.11
 - * Ubuntu 16.04.1 LTS
 - * Ubuntu 16.10
 - * Ubuntu 18.04.1
 - * Ubuntu 19.04

3.1 New Features

- **Added new armv8 machine targets.**

Added new armv8 machine targets:

- BlueField (Mellanox)
- OcteonTX2 (Marvell)
- ThunderX2 (Marvell)

- **Added Windows Support.**

Added Windows support to build Hello World sample application.

- **Added Stack Library.**

Added a new stack library and APIs for configuration and use of a bounded stack of pointers. The API provides multi-thread safe push and pop operations that can operate on one or more pointers per operation.

The library supports two stack implementations: standard (lock-based) and lock-free. The lock-free implementation is currently limited to x86-64 platforms.

- **Added Lock-Free Stack Mempool Handler.**

Added a new lock-free stack handler, which uses the newly added stack library.

- **Added RCU library.**

Added RCU library supporting a quiescent state based memory reclamation method. This library helps identify the quiescent state of the reader threads so that the writers can free the memory associated with the lock free data structures.

- **Updated KNI module and PMD.**

Updated the KNI kernel module to set the `max_mtu` according to the given initial MTU size. Without it, the maximum MTU was 1500.

Updated the KNI PMD driver to set the `mbuf_size` and MTU based on the given mb-pool. This provide the ability to pass jumbo frames if the mb-pool contains a suitable buffer size.

- **Added the AF_XDP PMD.**

Added a Linux-specific PMD driver for AF_XDP. This PMD can create an AF_XDP socket and bind it to a specific netdev queue. It allows a DPDK application to send and receive raw packets

through the socket which would bypass the kernel network stack to achieve high performance packet processing.

- **Added a net PMD NFB.**

Added the new `nfb` net driver for Netcope NFB cards. See the `../nics/nfb` NIC guide for more details on this new driver.

- **Added IPN3KE net PMD.**

Added the new `ipn3ke` net driver for the Intel® FPGA PAC (Programmable Acceleration Card) N3000. See the `../nics/ipn3ke` NIC guide for more details on this new driver.

In addition `ifpga_rawdev` was also updated to support Intel® FPGA PAC N3000 with SPI interface access, I2C Read/Write, and Ethernet PHY configuration.

- **Updated Solarflare network PMD.**

Updated the Solarflare `sfc_efx` driver with changes including:

- Added support for Rx descriptor status and related API in a secondary process.
- Added support for Tx descriptor status API in a secondary process.
- Added support for RSS RETA and hash configuration reading API in a secondary process.
- Added support for Rx packet types list in a secondary process.
- Added Tx prepare to do Tx offloads checks.
- Added support for VXLAN and GENEVE encapsulated TSO.

- **Updated Mellanox mlx4 driver.**

Updated Mellanox `mlx4` driver with new features and improvements, including:

- Added firmware version reading.
- Added support for secondary processes.
- Added support of per-process device registers. Reserving identical VA space is not needed anymore.
- Added support for multicast address list interfaces.

- **Updated Mellanox mlx5 driver.**

Updated Mellanox `mlx5` driver with new features and improvements, including:

- Added firmware version reading.
- Added support for new naming scheme of representor.
- Added support for new PCI device DMA map/unmap API.
- Added support for multiport InfiniBand device.
- Added control of excessive memory pinning by kernel.
- Added support of DMA memory registration by secondary process.
- Added support of per-process device registers. Reserving identical VA space is not required anymore.
- Added support for jump action for both E-Switch and NIC.

- Added Support for multiple `rte_flow` groups in NIC steering.
- **Flow engine re-designed to support large scale deployments. this includes:**
 - * Support millions of offloaded flow rules.
 - * Fast flow insertion and deletion up to 1M flow update per second.

- **Renamed `avf` to `iavf`.**

Renamed Intel Ethernet Adaptive Virtual Function driver `avf` to `iavf`, which includes the directory name, lib name, filenames, makefile, docs, macros, functions, structs and any other strings in the code.

- **Updated the `enic` driver.**

Updated `enic` driver with new features and improvements, including:

- Fixed several flow (director) bugs related to MARK, SCTP, VLAN, VXLAN, and inner packet matching.
- Added limited support for RAW.
- Added limited support for RSS.
- Added limited support for PASSTHRU.

- **Updated the `ixgbe` driver.**

Updated the `ixgbe` driver to add promiscuous mode support for the VF.

- **Updated the `ice` driver.**

Updated `ice` driver with new features and improvements, including:

- Added support of SSE and AVX2 instructions in Rx and Tx paths.
- Added package download support.
- Added Safe Mode support.
- Supported RSS for UPD/TCP/SCTP+IPV4/IPV6 packets.

- **Updated the `i40e` driver.**

New features for PF in the `i40e` driver:

- Added support for VXLAN-GPE packet.
- Added support for VXLAN-GPE classification.

- **Updated the `ENETC` driver.**

Updated `ENETC` driver with new features and improvements, including:

- Added physical addressing mode support.
- Added SXGMII interface support.
- Added basic statistics support.
- Added promiscuous and allmulticast mode support.
- Added MTU update support.
- Added jumbo frame support.

- Added queue start/stop.
- Added CRC offload support.
- Added Rx checksum offload validation support.
- **Updated the atlantic PMD.**
Added MACSEC hardware offload experimental API.
- **Updated the Intel QuickAssist Technology (QAT) compression PMD.**
Updated the Intel QuickAssist Technology (QAT) compression PMD to simplify, and make more robust, the handling of Scatter Gather Lists (SGLs) with more than 16 segments.
- **Updated the QuickAssist Technology (QAT) symmetric crypto PMD.**
Added support for AES-XTS with 128 and 256 bit AES keys.
- **Added Intel QuickAssist Technology PMD for asymmetric crypto.**
Added a new QAT Crypto PMD which provides asymmetric cryptography algorithms. Modular exponentiation and modular multiplicative inverse algorithms were added in this release.
- **Updated AESNI-MB PMD.**
Added support for out-of-place operations.
- **Updated the IPsec library.**
The IPsec library has been updated with AES-CTR and 3DES-CBC cipher algorithms support. The related `ipsec-secgw` test scripts have been added.
- **Updated the testpmd application.**
Improved the `testpmd` application performance on ARM platform. For `macswap` forwarding mode, NEON intrinsics are now used to do swap to save CPU cycles.
- **Updated power management library.**
Added support for Intel Speed Select Technology - Base Frequency (SST-BF). The `rte_power_get_capabilities` struct now has a bit in its returned mask indicating if it is a high frequency core.
- **Updated distributor sample application.**
Added support for the Intel SST-BF feature so that the distributor core is pinned to a high frequency core if available.

3.2 API Changes

- `eal`: the type of the `attr_value` parameter of the function `rte_service_attr_get()` has been changed from `uint32_t *` to `uint64_t *`.
- `meter`: replace enum `rte_meter_color` in the `meter` library with new `rte_color` definition added in 19.02. Replacements with `rte_color` values has been performed in many places such as `rte_mtr.h` and `rte_tm.h` to consolidate multiple color definitions.
- `vfio`: Functions `rte_vfio_container_dma_map` and `rte_vfio_container_dma_unmap` have been extended with an option to request mapping or un-mapping to the default `vfio` container `fd`.

- **power:** `rte_power_set_env` and `rte_power_unset_env` functions have been modified to be thread safe.
- **timer:** Functions have been introduced that allow multiple instances of the timer lists to be created. In addition they are now allocated in shared memory. New functions allow particular timer lists to be selected when timers are being started, stopped, and managed.

3.3 ABI Changes

- **ethdev:** Additional fields in `rte_eth_dev_info`.

The `rte_eth_dev_info` structure has had two extra fields added: `min_mtu` and `max_mtu`. Each of these are of type `uint16_t`. The values of these fields can be set specifically by the PMD drivers as supported values can vary from device to device.

- **cryptodev:** in 18.08 a new structure `rte_crypto_asym_op` was introduced and included into `rte_crypto_op`. As the `rte_crypto_asym_op` structure was defined as cache-line aligned that caused unintended changes in `rte_crypto_op` structure layout and alignment. Remove cache-line alignment for `rte_crypto_asym_op` to restore expected `rte_crypto_op` layout and alignment.
- **timer:** `rte_timer_subsystem_init` now returns success or failure to reflect whether it was able to allocate memory.

3.4 Shared Library Versions

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

```

librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
librte_bpf.so.1
librte_bus_dpaa.so.2
librte_bus_fslmc.so.2
librte_bus_ifpga.so.2
librte_bus_pci.so.2
librte_bus_vdev.so.2
librte_bus_vmbus.so.2
librte_cfgfile.so.2
librte_cmdline.so.2
librte_compressdev.so.1
+ librte_cryptodev.so.7
librte_distributor.so.1
+ librte_eal.so.10
librte_efd.so.1
+ librte_ethdev.so.12
librte_eventdev.so.6
librte_flow_classify.so.1
librte_gro.so.1
librte_gso.so.1
librte_hash.so.2
librte_ip_frag.so.1
librte_ipsec.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.5

```

```

librte_member.so.1
librte_mempool.so.5
librte_meter.so.3
librte_metrics.so.1
librte_net.so.1
librte_pci.so.1
librte_pdump.so.3
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_dpaa2_qdma.so.1
librte_pmd_ring.so.2
librte_pmd_softnic.so.1
librte_pmd_vhost.so.2
librte_port.so.3
librte_power.so.1
librte_rawdev.so.1
+ librte_rcu.so.1
librte_reorder.so.1
librte_ring.so.2
librte_sched.so.2
librte_security.so.2
+ librte_stack.so.1
librte_table.so.3
librte_timer.so.1
librte_vhost.so.4

```

3.5 Known Issues

- **On x86 platforms, AVX512 support is disabled with binutils 2.31.**

Due to a defect in binutils 2.31 AVX512 support is disabled. DPDK defect: https://bugs.dpdk.org/show_bug.cgi?id=249 GCC defect: https://gcc.gnu.org/bugzilla/show_bug.cgi?id=90028

- **No software AES-XTS implementation.**

There are currently no cryptodev software PMDs available which implement support for the AES-XTS algorithm, so this feature can only be used if compatible hardware and an associated PMD is available.

3.6 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
 - CPU
 - * Intel(R) Atom(TM) CPU C3758 @ 2.20GHz
 - * Intel(R) Xeon(R) CPU D-1541 @ 2.10GHz
 - * 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 E5-2699 v4 @ 2.20GHz
 - * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
 - * Intel(R) Xeon(R) Gold 6139 CPU @ 2.30GHz

– OS:

- * CentOS 7.4
- * CentOS 7.5
- * Fedora 25
- * Fedora 28
- * Fedora 29
- * FreeBSD 12.0
- * Red Hat Enterprise Linux Server release 7.4
- * Red Hat Enterprise Linux Server release 7.5
- * Red Hat Enterprise Linux Server release 7.6
- * SUSE12SP3
- * Open SUSE 15
- * Wind River Linux 8
- * Ubuntu 14.04
- * Ubuntu 16.04
- * Ubuntu 16.10
- * Ubuntu 18.04
- * Ubuntu 18.10

– 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 Corporation Ethernet Controller 10G X550T
 - Firmware version: 0x80000482
 - Device id (pf): 8086:1563
 - Driver version: 5.1.0-k(ixgbe)
- * Intel(R) Ethernet Converged Network Adapter X710-DA4 (4x10G)
 - Firmware version: 6.80 0x80003cc1
 - Device id (pf/vf): 8086:1572 / 8086:154c

- Driver version: 2.7.29 (i40e)
- * Intel(R) Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 3.33 0x8000fd5 0.0.0
 - Device id (pf/vf): 8086:37d0 / 8086:37cd
 - Driver version: 2.7.29 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 6.80 0x80003d05
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.7.29 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
 - Firmware version: 6.80 0x80003cfb
 - Device id (pf/vf): 8086:1583 / 8086:154c
 - Driver version: 2.7.29 (i40e)
- * Intel(R) Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x8000dda
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- * Intel Corporation I210 Gigabit Network Connection
 - Firmware version: 3.25, 0x800006eb, 1.1824.0
 - Device id (pf): 8086:1533
 - Driver version: 5.4.0-k(igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
 - * 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 7.6 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)

- * Red Hat Enterprise Linux Server release 7.3 (Maipo)
- * Red Hat Enterprise Linux Server release 7.2 (Maipo)
- * Ubuntu 19.04
- * Ubuntu 18.10
- * Ubuntu 18.04
- * Ubuntu 16.04
- * SUSE Linux Enterprise Server 15
- MLNX_OFED: 4.5-1.0.1.0
- MLNX_OFED: 4.6-1.0.1.1
- 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: 12.25.1020 and above
 - * Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
 - * Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above

- * Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
- * 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: 12.25.1020 and above
 - Firmware version: 12.25.1020 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
- * Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.25.1020 and above
- * Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.25.1020 and above
- * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.25.1020 and above
- * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)

- Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.25.1020 and above
 - * Mellanox(R) ConnectX(R)-5 Ex EN 100G MCX516A-CDAT (2x100G)
 - Host interface: PCI Express 4.0 x16
 - Device ID: 15b3:1019
 - Firmware version: 16.25.1020 and above
- Arm platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Qualcomm Arm 1.1 2500MHz
 - OS:
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - NICs:
 - * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.24.0220
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.24.0220
- Mellanox(R) BlueField SmartNIC
 - Mellanox(R) BlueField SmartNIC MT416842 (2x25G)
 - * Host interface: PCI Express 3.0 x16
 - * Device ID: 15b3:a2d2
 - * Firmware version: 18.25.1010
 - SoC Arm cores running OS:
 - * CentOS Linux release 7.4.1708 (AltArch)
 - * MLNX_OFED 4.6-1.0.0.0
 - DPDK application running on Arm cores inside SmartNIC
- IBM Power 9 platforms with Mellanox(R) NICs combinations
 - CPU:
 - * POWER9 2.2 (pvr 004e 1202) 2300MHz
 - OS:

- * Ubuntu 18.04.1 LTS (Bionic Beaver)
- NICs:
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.24.1000
- OFED:
 - * MLNX_OFED_LINUX-4.6-1.0.1.0

4.1 New Features

- **Added support for freeing hugepages exactly as originally allocated.**

Some applications using memory event callbacks (especially for managing RDMA memory regions) require that memory be freed back to the system exactly as it was originally allocated. These applications typically also require that a malloc allocation not span across two separate hugepage allocations. A new `--match-allocations` EAL init flag has been added to fulfill both of these requirements.

- **Added API to register external memory in DPDK.**

A new `rte_extmem_register/rte_extmem_unregister` API was added to allow chunks of external memory to be registered with DPDK without adding them to the malloc heap.

- **Added support for using virtio-user without hugepages.**

The `--no-huge` mode was augmented to use memfd-backed memory (on systems that support memfd), to allow using virtio-user-based NICs without hugepages.

- **Release of the ENA PMD v2.0.0.**

Version 2.0.0 of the ENA PMD was added with the following additions:

- Added Low Latency Queue v2 (LLQv2). This feature reduces the latency of the packets by pushing the header directly through the PCI to the device. This allows the NIC to start handle packets right after the doorbell without waiting for DMA.
- Added independent configuration of HW Tx and Rx ring depths.
- Added support for up to 8k Rx descriptors per ring.
- Added additional doorbell check on Tx, to handle Tx more efficiently for big bursts of packets.
- Added per queue statistics.
- Added extended statistics using `xstats` DPDK API.
- The reset routine was aligned with the DPDK API, so now it can be handled as in other PMDs.
- Fixed out of order (OOO) completion.
- Fixed memory leaks due to port stops and starts in the middle of traffic.
- Updated documentation and features list of the PMD.

- **Updated mlx5 driver.**

Updated the mlx5 driver including the following changes:

- Fixed `imissed` counter to be reported through `rte_eth_stats` instead of `rte_eth_xstats`.
- Added packet header modification through Direct Verbs flow driver.
- Added ConnectX-6 PCI device ID to be probed by `mlx5` driver.
- Added flow counter support to Direct Verbs flow driver through DevX.
- Renamed build options for the glue layer to `CONFIG_RTE_IBVERBS_LINK_DLOPEN` for `make` and `ibverbs_link` for `meson`.
- Added static linkage of `mlx` dependency.
- Improved stability of E-Switch flow driver.
- Added new `make` build configuration to set the cacheline size for BlueField correctly - `arm64-bluefield-linux-gcc`.

- **Updated the enic driver.**

- Added support for the `RTE_ETH_DEV_CLOSE_REMOVE` flag.
- Added a handler to get the firmware version string.
- Added support for multicast filtering.

- **Added dynamic queues allocation support for i40e VF.**

Previously, the available VF queues were reserved by PF at initialization stage. Now both DPDK PF and Kernel PF ($\geq 2.1.14$) will support dynamic queue allocation. At runtime, when VF requests for more queue exceed the initial reserved amount, the PF can allocate up to 16 queues as the request after a VF reset.

- **Added ICE net PMD.**

Added the new `ice` net driver for Intel(R) Ethernet Network Adapters E810. See the `../nics/ice` NIC guide for more details on this new driver.

- **Added support for SW-assisted VDPA live migration.**

This SW-assisted VDPA live migration facility helps VDPA devices without logging capability to perform live migration, a mediated SW relay can help devices to track dirty pages caused by DMA. the IFC driver has enabled this SW-assisted live migration mode.

- **Added security checks to the cryptodev symmetric session operations.**

Added a set of security checks to the access cryptodev symmetric session. The checks include the session's user data read/write check and the session private data referencing status check while freeing a session.

- **Updated the AESNI-MB PMD.**

- Added support for `intel-ipsec-mb` version 0.52.
- Added AES-GMAC algorithm support.
- Added Plain SHA1, SHA224, SHA256, SHA384, and SHA512 algorithms support.

- **Added IPsec Library.**

Added an experimental library `librte_ipsec` to provide ESP tunnel and transport support for IPv4 and IPv6 packets.

The library provides support for AES-CBC ciphering and AES-CBC with HMAC-SHA1 algorithm-chaining, and AES-GCM and NULL algorithms only at present. It is planned to add more algorithms in future releases.

See `../prog_guide/ipsec_lib` for more information.

- **Updated the ipsec-secgw sample application.**

The `ipsec-secgw` sample application has been updated to use the new `librte_ipsec` library, which has also been added in this release. The original functionality of `ipsec-secgw` is retained, a new command line parameter `-l` has been added to `ipsec-secgw` to use the IPsec library, instead of the existing IPsec code in the application.

The IPsec library does not support all the functionality of the existing `ipsec-secgw` application. It is planned to add the outstanding functionality in future releases.

See `../sample_app Ug/ipsec_secgw` for more information.

- **Enabled checksum support in the ISA-L compressdev driver.**

Added support for both `adler` and `crc32` checksums in the ISA-L PMD. This aids data integrity across both compression and decompression.

- **Added a compression performance test tool.**

Added a new performance test tool to test the `compressdev` PMD. The tool tests compression ratio and compression throughput.

- **Added intel_pstate support to Power Management library.**

Previously, using the power management library required the disabling of the `intel_pstate` kernel driver, and the enabling of the `acpi_cpufreq` kernel driver. This is no longer the case, as the use of the `intel_pstate` kernel driver is now supported, and automatically detected by the library.

4.2 API Changes

- `eal`: Function `rte_bsf64` in `rte_bitmap.h` has been renamed to `rte_bsf64_safe` and moved to `rte_common.h`. A new `rte_bsf64` function has been added in `rte_common.h` that follows the convention set by the existing `rte_bsf32` function.
- `eal`: Segment `fd` API on Linux now sets error code to `ENOTSUP` in more cases where segment the `fd` API is not expected to be supported:
 - On attempt to get a segment `fd` for an externally allocated memory segment
 - In cases where `memfd` support would have been required to provide segment `fds` (such as `in-memory` or `no-huge` mode)
- `eal`: Functions `rte_malloc_dump_stats()`, `rte_malloc_dump_heaps()` and `rte_malloc_get_socket_stats()` are no longer safe to call concurrently with `rte_malloc_heap_create()` or `rte_malloc_heap_destroy()` function calls.
- `mbuf`: `RTE_MBUF_INDIRECT()`, which was deprecated in 18.05, was replaced with `RTE_MBUF_CLONED()` and removed in 19.02.

- **sched:** As result of the new format of the mbuf sched field, the functions `rte_sched_port_pkt_write()` and `rte_sched_port_pkt_read_tree_path()` got an additional parameter of type `struct rte_sched_port`.
- **pdump:** The `rte_pdump_set_socket_dir()`, the parameter `path` of `rte_pdump_init()` and enum `rte_pdump_socktype` were deprecated since 18.05 and are removed in this release.
- **cryptodev:** The parameter `session_pool` in the function `rte_cryptodev_queue_pair_setup()` is removed.
- **cryptodev:** a new function `rte_cryptodev_sym_session_pool_create()` has been introduced. This function is now mandatory when creating symmetric session header mempool. Please note all crypto applications are required to use this function from now on. Failed to do so will cause a `rte_cryptodev_sym_session_create()` function call return error.

4.3 ABI Changes

- **mbuf:** The format of the sched field of `rte_mbuf` has been changed to include the following fields: `queue ID`, `traffic class`, `color`.
- **cryptodev:** as shown in the 18.11 deprecation notice, the structure `rte_cryptodev_qp_conf` has added two parameters for symmetric session mempool and symmetric session private data mempool.
- **cryptodev:** as shown in the 18.11 deprecation notice, the structure `rte_cryptodev_sym_session` has been updated to contain more information to ensure safely accessing the session and session private data.
- **security:** A new field `uint64_t opaque_data` has been added to `rte_security_session` structure. That would allow upper layer to easily associate/de-associate some user defined data with the security session.

4.4 Shared Library Versions

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

```

librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
librte_bpf.so.1
librte_bus_dpaa.so.2
librte_bus_fslmc.so.2
librte_bus_ifpga.so.2
librte_bus_pci.so.2
librte_bus_vdev.so.2
librte_bus_vmbus.so.2
librte_cfgfile.so.2
librte_cmdline.so.2
librte_compressdev.so.1
+ librte_cryptodev.so.6
librte_distributor.so.1
librte_eal.so.9
librte_efd.so.1
librte_ethdev.so.11
librte_eventdev.so.6
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.5
librte_member.so.1
librte_mempool.so.5
librte_meter.so.2
librte_metrics.so.1
librte_net.so.1
librte_pci.so.1
+ librte_pdump.so.3
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_dpaa2_qdma.so.1
librte_pmd_ring.so.2
librte_pmd_softnic.so.1
librte_pmd_vhost.so.2
librte_port.so.3
librte_power.so.1
librte_rawdev.so.1
librte_reorder.so.1
librte_ring.so.2
+ librte_sched.so.2
+ librte_security.so.2
librte_table.so.3
librte_timer.so.1
librte_vhost.so.4

```

4.5 Known Issues

- AVX-512 support has been disabled for GCC builds when `binutils 2.30` is detected [1] because of a crash [2]. This can affect native machine type build targets on the platforms that support AVX512F like Intel Skylake processors, and can cause a possible performance drop. The immediate workaround is to use `clang` compiler on these platforms. Initial workaround in DPDK v18.11 was to disable AVX-512 support for GCC completely, but based on information on defect submitted to GCC community [3], issue has been identified as `binutils 2.30` issue. Since currently only GCC generates AVX-512 instructions, the scope is limited to GCC and `binutils 2.30`
 - [1]: Commit (“mk: fix scope of disabling AVX512F support”)
 - [2]: https://bugs.dpdk.org/show_bug.cgi?id=97
 - [3]: https://gcc.gnu.org/bugzilla/show_bug.cgi?id=88096

4.6 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
 - CPU
 - * Intel(R) Atom(TM) CPU C3758 @ 2.20GHz

- * Intel(R) Xeon(R) CPU D-1541 @ 2.10GHz
- * 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 E5-2699 v4 @ 2.20GHz
- * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
- * Intel(R) Xeon(R) Gold 6139 CPU @ 2.30GHz

– OS:

- * CentOS 7.4
- * CentOS 7.5
- * Fedora 25
- * Fedora 28
- * FreeBSD 11.2
- * FreeBSD 12.0
- * Red Hat Enterprise Linux Server release 7.4
- * Red Hat Enterprise Linux Server release 7.5
- * Open SUSE 15
- * Wind River Linux 8
- * Ubuntu 14.04
- * Ubuntu 16.04
- * Ubuntu 16.10
- * Ubuntu 18.04
- * Ubuntu 18.10

– 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: 6.80 0x80003cc1
 - Device id (pf/vf): 8086:1572 / 8086:154c

- Driver version: 2.7.26 (i40e)
- * Intel(R) Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 3.33 0x80000fd5 0.0.0
 - Device id (pf/vf): 8086:37d0 / 8086:37cd
 - Driver version: 2.7.26 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 6.80 0x80003d05
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.7.26 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
 - Firmware version: 6.80 0x80003cfb
 - Device id (pf/vf): 8086:1583 / 8086:154c
 - Driver version: 2.7.26 (i40e)
- * Intel(R) Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x80000dda
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
 - * 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 7.6 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.3 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.2 (Maipo)
 - * Ubuntu 18.10
 - * Ubuntu 18.04

- * Ubuntu 17.10
- * Ubuntu 16.04
- * SUSE Linux Enterprise Server 15
- MLNX_OFED: 4.4-2.0.1.0
- MLNX_OFED: 4.5-1.0.1.0
- 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: 2.42.5000
 - * Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
 - * Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
 - * Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
 - * Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above

- * Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
- * 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: 12.24.1000 and above
 - Firmware version: 12.24.1000 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
- * Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.24.1000 and above
- * Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.24.1000 and above
- * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.24.1000 and above
- * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.24.1000 and above
- * Mellanox(R) ConnectX(R)-5 Ex EN 100G MCX516A-CDAT (2x100G)

- Host interface: PCI Express 4.0 x16
 - Device ID: 15b3:1019
 - Firmware version: 16.24.1000 and above
- ARM platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Qualcomm ARM 1.1 2500MHz
 - OS:
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - NICs:
 - * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.24.0220
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.24.0220
- Mellanox(R) BlueField SmartNIC
 - Mellanox(R) BlueField SmartNIC MT416842 (2x25G)
 - * Host interface: PCI Express 3.0 x16
 - * Device ID: 15b3:a2d2
 - * Firmware version: 18.24.0328
 - SoC ARM cores running OS:
 - * CentOS Linux release 7.4.1708 (AltArch)
 - * MLNX_OFED 4.4-2.5.9.0
 - DPDK application running on ARM cores inside SmartNIC
- Power 9 platforms with Mellanox(R) NICs combinations
 - CPU:
 - * POWER9 2.2 (pvr 004e 1202) 2300MHz
 - OS:
 - * Ubuntu 18.04.1 LTS (Bionic Beaver)
 - NICs:
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16

- Device ID: 15b3:1017
- Firmware version: 16.23.1020
- OFED:
 - * MLNX_OFED_LINUX-4.5-1.0.1.0

DPDK RELEASE 18.11

5.1 New Features

- **Added support for using externally allocated memory in DPDK.**

DPDK has added support for creating new `rte_malloc` heaps referencing memory that was created outside of DPDK's own page allocator, and using that memory natively with any other DPDK library or data structure.

- **Added check for ensuring allocated memory is addressable by devices.**

Some devices can have addressing limitations so a new function, `rte_mem_check_dma_mask()`, has been added for checking that allocated memory is not out of the device range. Since memory can now be allocated dynamically after initialization, a DMA mask is stored and any new allocated memory will be checked against it and rejected if it is out of range. If more than one device has addressing limitations, the DMA mask is the more restrictive one.

- **Updated the C11 memory model version of the ring library.**

Added changes to decrease latency for architectures using the C11 memory model version of the ring library.

On Cavium ThunderX2 platform, the changes decreased latency by 27-29% and 3-15% for MPMC and SPSC cases respectively (with 2 lcores). The real improvements may vary with the number of contending lcores and the size of the ring.

- **Added hot-unplug handle mechanism.**

Added `rte_dev_hotplug_handle_enable()` and `rte_dev_hotplug_handle_disable()` for enabling or disabling the hotplug handle mechanism.

- **Added support for device multi-process hotplug.**

Added support for hotplug and hot-unplug in a multiprocessing scenario. Any `ethdev` devices created in the primary process will be regarded as shared and will be available for all DPDK processes. Synchronization between processes will be done using DPDK IPC.

- **Added new Flow API actions to rewrite fields in packet headers.**

Added new Flow API actions to:

- Modify source and destination IP addresses in the outermost IPv4/IPv6 headers.
- Modify source and destination port numbers in the outermost TCP/UDP headers.

- **Added new Flow API action to swap MAC addresses in Ethernet header.**

Added new Flow API action to swap the source and destination MAC addresses in the outermost Ethernet header.

- **Add support to offload more flow match and actions for CXGBE PMD.**

The Flow API support has been enhanced for the CXGBE Poll Mode Driver to offload:

- Match items: destination MAC address.
- Action items: push/pop/rewrite vlan header, rewrite IP addresses in outermost IPv4/IPv6 header, rewrite port numbers in outermost TCP/UDP header, swap MAC addresses in outermost Ethernet header.

- **Added a devarg to use the latest supported vector path in i40e.**

A new devarg `use-latest-supported-vec` was introduced to allow users to choose the latest vector path that the platform supported. For example, users can use AVX2 vector path on BDW/HSW to get better performance.

- **Added support for SR-IOV in netvsc PMD.**

The `netvsc` poll mode driver now supports the Accelerated Networking SR-IOV option in Hyper-V and Azure. This is an alternative to the previous `vdev_netvsc`, `tap`, and `failsafe` drivers combination.

- **Added a new net driver for Marvell Armada 3k device.**

Added the new `mvneta` net driver for Marvell Armada 3k device. See the `../nics/mvneta` NIC guide for more details on this new driver.

- **Added NXP ENETC PMD.**

Added the new `enetc` driver for the NXP `enetc` platform. See the `../nics/enetc` NIC driver guide for more details on this new driver.

- **Added Ethernet poll mode driver for Aquantia aQtion family of 10G devices.**

Added the new `atlantic` ethernet poll mode driver for Aquantia XGBE devices. See the `../nics/atlantic` NIC driver guide for more details on this driver.

- **Updated mlx5 driver.**

Updated the `mlx5` driver including the following changes:

- Improved security of PMD to prevent the NIC from getting stuck when the application misbehaves.
- Reworked flow engine to supported e-switch flow rules (transfer attribute).
- Added support for header re-write(L2-L4), VXLAN encap/decap, count, match on TCP flags and multiple flow groups with e-switch flow rules.
- Added support for match on metadata, VXLAN and MPLS encap/decap with flow rules.
- Added support for `RTE_ETH_DEV_CLOSE_REMOVE` flag to provide better support for representors.
- Added support for meson build.
- Fixed build issue with PPC.
- Added support for BlueField VF.
- Added support for externally allocated static memory for DMA.

- **Updated Solarflare network PMD.**

Updated the `sfc_efx` driver including the following changes:

- Added support for Rx scatter in EF10 datapath implementation.
- Added support for Rx descriptor status API in EF10 datapath implementation.
- Added support for TSO in EF10 datapath implementation.
- Added support for Tx descriptor status API in EF10 (`ef10` and `ef10_simple`) datapaths implementation.

- **Updated the enic driver.**

- Added AVX2-based vectorized Rx handler.
- Added VLAN and checksum offloads to the simple Tx handler.
- Added the “count” flow action.
- Enabled the virtual address IOVA mode.

- **Updated the failsafe driver.**

Updated the failsafe driver including the following changes:

- Added support for Rx and Tx queues start and stop.
- Added support for Rx and Tx queues deferred start.
- Added support for runtime Rx and Tx queues setup.
- Added support multicast MAC address set.

- **Added a devarg to use a PCAP interface physical MAC address.**

A new devarg `phy_mac` was introduced to allow users to use the physical MAC address of the selected PCAP interface.

- **Added TAP Rx/Tx queues sharing with a secondary process.**

Added support to allow a secondary process to attach a TAP device created in the primary process, probe the queues, and process Rx/Tx in a secondary process.

- **Added classification and metering support to SoftNIC PMD.**

Added support for flow classification (`rte_flow` API), and metering and policing (`rte_mtr` API) to the SoftNIC PMD.

- **Added Crypto support to the Softnic PMD.**

The Softnic is now capable of processing symmetric crypto workloads such as cipher, cipher-authentication chaining, and AEAD encryption and decryption. This is achieved by calling DPDK Cryptodev APIs.

- **Added cryptodev port to port library.**

Cryptodev port is a shim layer in the port library that interacts with DPDK Cryptodev PMDs including burst enqueueing and dequeueing crypto operations.

- **Added symmetric cryptographic actions to the pipeline library.**

In the pipeline library support was added for symmetric crypto action parsing and an action handler was implemented. The action allows automatic preparation of the crypto operation with the rules specified such as algorithm, key, and IV, etc. for the cryptodev port to process.

- **Updated the AESNI MB PMD.**

The AESNI MB PMD has been updated with additional support for the AES-GCM algorithm.

- **Added NXP CAAM JR PMD.**

Added the new caam job ring driver for NXP platforms. See the `../cryptodevs/caam_jr` guide for more details on this new driver.

- **Added support for GEN3 devices to Intel QAT driver.**

Added support for the third generation of Intel QuickAssist devices.

- **Updated the QAT PMD.**

The QAT PMD was updated with additional support for:

- The AES-CMAC algorithm.

- **Added support for Dynamic Huffman Encoding to Intel QAT comp PMD.**

The Intel QuickAssist (QAT) compression PMD has been updated with support for Dynamic Huffman Encoding for the Deflate algorithm.

- **Added Event Ethernet Tx Adapter.**

Added event ethernet Tx adapter library that provides configuration and data path APIs for the ethernet transmit stage of an event driven packet processing application. These APIs abstract the implementation of the transmit stage and allow the application to use eventdev PMD support or a common implementation.

- **Added Distributed Software Eventdev PMD.**

Added the new Distributed Software Event Device (DSW), which is a pure-software eventdev driver distributing the work of scheduling among all eventdev ports and the lcores using them. DSW, compared to the SW eventdev PMD, sacrifices load balancing performance to gain better event scheduling throughput and scalability.

- **Added extendable bucket feature to hash library (`rte_hash`).**

This new “extendable bucket” feature provides 100% insertion guarantee to the capacity specified by the user by extending hash table with extra buckets when needed to accommodate the unlikely event of intensive hash collisions. In addition, the internal hashing algorithm was changed to use partial-key hashing to improve memory efficiency and lookup performance.

- **Added lock free reader/writer concurrency to hash library (`rte_hash`).**

Lock free reader/writer concurrency prevents the readers from getting blocked due to a preempted writer thread. This allows the hash library to be used in scenarios where the writer thread runs on the control plane.

- **Added Traffic Pattern Aware Power Control Library.**

Added an experimental library that extends the Power Library and provides `empty_poll` APIs. This feature measures how many times `empty_polls` are executed per core and uses the number of empty polls as a hint for system power management.

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

- **Added JSON power policy interface for containers.**

Extended the Power Library and `vm_power_manager` sample app to allow power policies to be submitted via a FIFO using JSON formatted strings. Previously limited to Virtual Machines, this feature extends power policy functionality to containers and host applications that need to have their cores frequency controlled based on the rules contained in the policy.

- **Added Telemetry API.**

Added a new telemetry API which allows applications to transparently expose their telemetry in JSON via a UNIX socket. The JSON can be consumed by any Service Assurance agent, such as CollectD.

- **Updated KNI kernel module, `rte_kni` library, and KNI sample application.**

Updated the KNI kernel module with a new kernel module parameter, `carrier=[on|off]` to allow the user to control the default carrier state of the KNI kernel network interfaces. The default carrier state is now set to `off`, so the interfaces cannot be used until the carrier state is set to `on` via `rte_kni_update_link` or by writing `1` to `/sys/devices/virtual/net/<iface>/carrier`. In previous versions the default carrier state was left undefined. See `./prog_guide/kernel_nic_interface` for more information.

Also added the new API function `rte_kni_update_link()` to allow the user to set the carrier state of the KNI kernel network interface.

Also added a new command line flag `-m` to the KNI sample application to monitor and automatically reflect the physical NIC carrier state to the KNI kernel network interface with the new `rte_kni_update_link()` API. See `./sample_app Ug/kernel_nic_interface` for more information.

- **Added ability to switch queue deferred start flag on `testpmd` app.**

Added a console command to `testpmd` app, giving ability to switch `rx_deferred_start` or `tx_deferred_start` flag of the specified queue of the specified port. The port must be stopped before the command call in order to reconfigure queues.

- **Add a new sample application for vDPA.**

The `vdpa` sample application creates vhost-user sockets by using the vDPA backend. vDPA stands for vhost Data Path Acceleration which utilizes virtio ring compatible devices to serve virtio driver directly to enable datapath acceleration. As vDPA driver can help to set up vhost datapath, this application doesn't need to launch dedicated worker threads for vhost enqueue/dequeue operations.

- **Added cryptodev FIPS validation example application.**

Added an example application to parse and perform symmetric cryptography computation to the NIST Cryptographic Algorithm Validation Program (CAVP) test vectors.

- **Allow unit test binary to take parameters from the environment.**

The unit test "test", or "dpdk-test", binary is often called from scripts, which can make passing additional parameters, such as a `coremask`, difficult. Support has been added to the application to allow it to take additional command-line parameter values from the `DPDK_TEST_PARAMS` environment variable to make this application easier to use.

5.2 API Changes

- eal: `rte_memseg_list` structure now has an additional flag indicating whether the memseg list is externally allocated. This will have implications for any users of memseg-walk-related functions, as they will now have to skip externally allocated segments in most cases if the intent is to only iterate over internal DPDK memory.

In addition the `socket_id` parameter across the entire DPDK has gained additional meaning, as some socket ID's will now be representing externally allocated memory. No changes will be required for existing code as backwards compatibility will be kept, and those who do not use this feature will not see these extra socket ID's. Any new API's must not check socket ID parameters themselves, and must instead leave it to the memory subsystem to decide whether socket ID is a valid one.

- eal: The following devargs functions, which were deprecated in 18.05, were removed in 18.11: `rte_eal_parse_devargs_str()`, `rte_eal_devargs_add()`, `rte_eal_devargs_type_count()`, and `rte_eal_devargs_dump()`.
- eal: The parameters of the function `rte_devargs_remove()` have changed from bus and device names to `struct rte_devargs`.
- eal: The deprecated functions `attach/detach` were removed in 18.11. `rte_eal_dev_attach` can be replaced by `rte_dev_probe` or `rte_eal_hotplug_add`. `rte_eal_dev_detach` can be replaced by `rte_dev_remove` or `rte_eal_hotplug_remove`.
- eal: The scope of `rte_eal_hotplug_add()/rte_dev_probe()` and `rte_eal_hotplug_remove()/rte_dev_remove()` has been extended. In the multi-process model, they will guarantee that the device is attached or detached on all processes.
- mbuf: The `__rte_mbuf_raw_free()` and `__rte_pktmbuf_prefree_seg()` functions were deprecated since 17.05 and are replaced by `rte_mbuf_raw_free()` and `rte_pktmbuf_prefree_seg()`.
- ethdev: The deprecated functions `attach/detach` were removed in 18.11. `rte_eth_dev_attach()` can be replaced by `RTE_ETH_FOREACH_MATCHING_DEV` and `rte_dev_probe()` or `rte_eal_hotplug_add()`. `rte_eth_dev_detach()` can be replaced by `rte_dev_remove()` or `rte_eal_hotplug_remove()`.
- ethdev: A call to `rte_eth_dev_release_port()` has been added in `rte_eth_dev_close()`. As a consequence, a closed port is freed and seen as invalid because of its state `RTE_ETH_DEV_UNUSED`. This new behavior is enabled per driver for a migration period.
- A new device flag, `RTE_ETH_DEV_NOLIVE_MAC_ADDR`, changes the order of actions inside `rte_eth_dev_start()` regarding MAC set. Some NICs do not support MAC changes once the port has started and with this new device flag the MAC can be properly configured in any case. This is particularly important for bonding.
- The default behavior of CRC strip offload has changed in this release. Without any specific Rx offload flag, default behavior by a PMD is now to strip CRC. `DEV_RX_OFFLOAD_CRC_STRIP` offload flag has been removed. To request keeping CRC, application should set `DEV_RX_OFFLOAD_KEEP_CRC` Rx offload.
- eventdev: The type of the second parameter to `rte_event_eth_rx_adapter_caps_get()` has been changed from `uint8_t` to `uint16_t`.

- **kni:** By default, interface carrier status is `off` which means there won't be any traffic. It can be set to `on` via `rte_kni_update_link()` API or via `sysfs` interface: `echo 1 > /sys/class/net/vEth0/carrier`.

Note interface should be up to be able to read/write `sysfs` interface. When KNI sample application is used, `-m` parameter can be used to automatically update the carrier status for the interface.

- **kni:** When `ethtool` support is enabled (`CONFIG_RTE_KNI_KMOD_ETHTOOL=y`) `ethtool` commands `ETHTOOL_GSET` & `ETHTOOL_SSET` are no longer supported for kernels that have `ETHTOOL_GLINKSETTINGS` & `ETHTOOL_SLINKSETTINGS` support. This means `ethtool "-a|--show-pause"`, `ethtool "-s|--change"` won't work, and `ethtool <iface>` output will have less information.

5.3 ABI Changes

- **eal:** added `legacy_mem` and `single_file_segments` values to `rte_config` structure on account of improving DPDK usability when using either `--legacy-mem` or `--single-file-segments` flags.
- **eal:** EAL library ABI version was changed due to previously announced work on supporting external memory in DPDK:
 - Structure `rte_memseg_list` now has a new field indicating length of memory addressed by the segment list
 - Structure `rte_memseg_list` now has a new flag indicating whether the memseg list refers to external memory
 - Structure `rte_malloc_heap` now has a new field indicating socket ID the malloc heap belongs to
 - Structure `rte_mem_config` has had its `malloc_heaps` array resized from `RTE_MAX_NUMA_NODES` to `RTE_MAX_HEAPS` value
 - Structure `rte_malloc_heap` now has a `heap_name` member
 - Structure `rte_eal_memconfig` has been extended to contain next socket ID for externally allocated segments
- **eal:** Added `dma_maskbits` to `rte_mem_config` for keeping the most restrictive DMA mask based on the devices addressing limitations.
- **eal:** The structure `rte_device` has a new field to reference a `rte_bus`. It thus changes the size of the `struct rte_device` and the inherited device structures of all buses.

5.4 Shared Library Versions

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

```
librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
librte_bpf.so.1
+ librte_bus_dpaa.so.2
+ librte_bus_fslmc.so.2
+ librte_bus_ifpga.so.2
+ librte_bus_pci.so.2
+ librte_bus_vdev.so.2
```

```

+ librte_bus_vmbus.so.2
  librte_cfgfile.so.2
  librte_cmdline.so.2
  librte_compressdev.so.1
  librte_cryptodev.so.5
  librte_distributor.so.1
+ librte_eal.so.9
  librte_efd.so.1
+ librte_ethdev.so.11
+ librte_eventdev.so.6
  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.4
  librte_member.so.1
  librte_mempool.so.5
  librte_meter.so.2
  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_dpaa2_qdma.so.1
  librte_pmd_ring.so.2
  librte_pmd_softnic.so.1
  librte_pmd_vhost.so.2
  librte_port.so.3
  librte_power.so.1
  librte_rawdev.so.1
  librte_reorder.so.1
  librte_ring.so.2
  librte_sched.so.1
  librte_security.so.1
  librte_table.so.3
  librte_timer.so.1
  librte_vhost.so.4

```

5.5 Known Issues

- When using SR-IOV (VF) support with netvsc PMD and the Mellanox mlx5 bifurcated driver the Linux netvsc device must be brought up before the netvsc device is unbound and passed to the DPDK.
- IBM Power8 is not supported in this release of DPDK. IBM Power9 is supported.
- AVX-512 support has been disabled for GCC builds [1] because of a crash [2]. This can affect native machine type build targets on the platforms that support AVX512F like Intel Skylake processors, and can cause a possible performance drop. The immediate workaround is to use clang compiler on these platforms. The issue has been identified as a GCC defect and reported to the GCC community [3]. Further actions will be taken based on the GCC defect result.

- [1]: Commit 8d07c82b239f (“mk: disable gcc AVX512F support”)
- [2]: https://bugs.dpdk.org/show_bug.cgi?id=97
- [3]: https://gcc.gnu.org/bugzilla/show_bug.cgi?id=88096

5.6 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
 - CPU
 - * Intel(R) Atom(TM) CPU C3758 @ 2.20GHz
 - * Intel(R) Xeon(R) CPU D-1541 @ 2.10GHz
 - * 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 E5-2699 v4 @ 2.20GHz
 - * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
 - OS:
 - * CentOS 7.5
 - * Fedora 25
 - * Fedora 28
 - * FreeBSD 11.2
 - * Red Hat Enterprise Linux Server release 7.5
 - * Open SUSE 15
 - * Wind River Linux 8
 - * Ubuntu 14.04
 - * Ubuntu 16.04
 - * Ubuntu 16.10
 - * Ubuntu 17.10
 - * Ubuntu 18.04
 - 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 0x80003221
 - Device id (pf/vf): 8086:1572 / 8086:154c
 - Driver version: 2.4.6 (i40e)
- * Intel(R) Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 3.33 0x80000fd5 0.0.0
 - Device id (pf/vf): 8086:37d0 / 8086:37cd
 - Driver version: 2.4.6 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 6.01 0x80003221
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.4.6 (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: 2.4.6 (i40e)
- * Intel(R) Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x80000dda
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
 - * 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 7.6 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)

- * Red Hat Enterprise Linux Server release 7.3 (Maipo)
- * Red Hat Enterprise Linux Server release 7.2 (Maipo)
- * Ubuntu 18.10
- * Ubuntu 18.04
- * Ubuntu 17.10
- * Ubuntu 16.04
- * SUSE Linux Enterprise Server 15
- MLNX_OFED: 4.4-2.0.1.0
- MLNX_OFED: 4.5-0.3.1.0
- 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: 2.42.5000
 - * Mellanox(R) ConnectX(R)-4 10G MCX4111A-XCAT (1x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
 - * Mellanox(R) ConnectX(R)-4 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4111A-ACAT (1x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
 - * Mellanox(R) ConnectX(R)-4 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
 - * Mellanox(R) ConnectX(R)-4 40G MCX4131A-BCAT/MCX413A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above

- * Mellanox(R) ConnectX(R)-4 40G MCX415A-BCAT (1x40G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX4131A-GCAT/MCX413A-GCAT (1x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX414A-BCAT (2x50G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
- * 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: 12.23.8022 and above
 - Firmware version: 12.23.8022 and above
- * Mellanox(R) ConnectX(R)-4 50G MCX415A-CCAT (1x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
- * Mellanox(R) ConnectX(R)-4 100G MCX416A-CCAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1013
 - Firmware version: 12.23.8022 and above
- * Mellanox(R) ConnectX(R)-4 Lx 10G MCX4121A-XCAT (2x10G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.23.8022 and above
- * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.23.8022 and above
- * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)

- Host interface: PCI Express 3.0 x16
- Device ID: 15b3:1017
- Firmware version: 16.23.8022 and above
- * Mellanox(R) ConnectX(R)-5 Ex EN 100G MCX516A-CDAT (2x100G)
 - Host interface: PCI Express 4.0 x16
 - Device ID: 15b3:1019
 - Firmware version: 16.23.8022 and above
- ARM platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Qualcomm ARM 1.1 2500MHz
 - OS:
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - NICs:
 - * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.24.0220
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.24.0220
- Mellanox(R) BlueField SmartNIC
 - Mellanox(R) BlueField SmartNIC MT416842 (2x25G)
 - * Host interface: PCI Express 3.0 x16
 - * Device ID: 15b3:a2d2
 - * Firmware version: 18.24.0246
 - SoC ARM cores running OS:
 - * CentOS Linux release 7.4.1708 (AltArch)
 - * MLNX_OFED 4.4-2.5.3.0
 - DPDK application running on ARM cores inside SmartNIC
- ARM SoC combinations from NXP (with integrated NICs)
 - SoC:
 - * NXP/Freescale QorIQ LS1046A with ARM Cortex A72
 - * NXP/Freescale QorIQ LS2088A with ARM Cortex A72

– OS:

- * Ubuntu 18.04.1 LTS with NXP QorIQ LSDK 1809 support packages
- * Ubuntu 16.04.3 LTS with NXP QorIQ LSDK 1803 support packages

6.1 New Features

- **Added support for Hyper-V netvsc PMD.**

The new `netvsc` poll mode driver provides native support for networking on Hyper-V. See the `./nics/netvsc` NIC driver guide for more details on this new driver.

- **Added Flow API support for CXGBE PMD.**

Flow API support has been added to CXGBE Poll Mode Driver to offload flows to Chelsio T5/T6 NICs. Support added for:

- Wildcard (LE-TCAM) and Exact (HASH) match filters.
- Match items: physical ingress port, IPv4, IPv6, TCP and UDP.
- Action items: queue, drop, count, and physical egress port redirect.

- **Added ixgbe preferred Rx/Tx parameters.**

Rather than applications providing explicit Rx and Tx parameters such as queue and burst sizes, they can request that the EAL instead uses preferred values provided by the PMD, falling back to defaults within the EAL if the PMD does not provide any. The provision of such tuned values now includes the ixgbe PMD.

- **Added descriptor status check support for fm10k.**

The `rte_eth_rx_descriptor_status` and `rte_eth_tx_descriptor_status` APIs are now supported by fm10K.

- **Updated the enic driver.**

- Add low cycle count Tx handler for no-offload Tx.
- Add low cycle count Rx handler for non-scattered Rx.
- Minor performance improvements to scattered Rx handler.
- Add handlers to add/delete VxLAN port number.
- Add `devarg` to specify ingress VLAN rewrite mode.

- **Updated mlx5 driver.**

Updated the mlx5 driver including the following changes:

- Added port representors support.

- Added Flow API support for e-switch rules. Added support for ACTION_PORT_ID, ACTION_DROP, ACTION_OF_POP_VLAN, ACTION_OF_PUSH_VLAN, ACTION_OF_SET_VLAN_VID, ACTION_OF_SET_VLAN_PCP and ITEM_PORT_ID.
- Added support for 32-bit compilation.
- **Added TSO support for the mlx4 driver.**
Added TSO support for the mlx4 drivers from MLNX_OFED_4.4 and above.
- **SoftNIC PMD rework.**
The SoftNIC PMD infrastructure has been restructured to use the Packet Framework, which makes it more flexible, modular and easier to add new functionality in the future.
- **Updated the AESNI MB PMD.**
The AESNI MB PMD has been updated with additional support for:
 - 3DES for 8, 16 and 24 byte keys.
- **Added a new compression PMD using Intel’s QuickAssist (QAT) device family.**
Added the new QAT compression driver, for compression and decompression operations in software. See the ../compressdevs/qat_comp compression driver guide for details on this new driver.
- **Updated the ISA-L PMD.**
Added support for chained mbufs (input and output).

6.2 API Changes

- The path to the runtime config file has changed. The new path is determined as follows:
 - If DPDK is running as root, /var/run/dpdk/<prefix>/config
 - If DPDK is not running as root:
 - * If \$XDG_RUNTIME_DIR is set, \${XDG_RUNTIME_DIR}/dpdk/<prefix>/config
 - * Otherwise, /tmp/dpdk/<prefix>/config
- eal: The function `rte_eal_mbuf_default_mempool_ops` was deprecated and is removed in 18.08. It shall be replaced by `rte_mbuf_best_mempool_ops`.
- mempool: Following functions were deprecated and are removed in 18.08:
 - `rte_mempool_populate_iova_tab`
 - `rte_mempool_populate_phys_tab`
 - `rte_mempool_populate_phys` (`rte_mempool_populate_iova` should be used)
 - `rte_mempool_virt2phy` (`rte_mempool_virt2iova` should be used)
 - `rte_mempool_xmem_create`
 - `rte_mempool_xmem_size`
 - `rte_mempool_xmem_usage`
- ethdev: The old offload API is removed:

- Rx per-port `rte_eth_conf.rxmode`. [bit-fields]
- Tx per-queue `rte_eth_txconf.txq_flags`
- `ETH_TXQ_FLAGS_NO*`

The transition bits are removed:

- `rte_eth_conf.rxmode.ignore_offload_bitfield`
- `ETH_TXQ_FLAGS_IGNORE`

- cryptodev: The following API changes have been made in 18.08:

- In struct `struct rte_cryptodev_info`, field `rte_pci_device *pci_dev` has been replaced with field `struct rte_device *device`.
- Value 0 is accepted in `sym.max_nb_sessions`, meaning that a device supports an unlimited number of sessions.
- Two new fields of type `uint16_t` have been added: `min_mbuf_headroom_req` and `min_mbuf_tailroom_req`. These parameters specify the recommended headroom and tailroom for mbufs to be processed by the PMD.

- cryptodev: The following functions were deprecated and are removed in 18.08:

- `rte_cryptodev_queue_pair_start`
- `rte_cryptodev_queue_pair_stop`
- `rte_cryptodev_queue_pair_attach_sym_session`
- `rte_cryptodev_queue_pair_detach_sym_session`

- cryptodev: The following functions were deprecated and are replaced by other functions in 18.08:

- `rte_cryptodev_get_header_session_size` is replaced with `rte_cryptodev_sym_get_header_session_size`
- `rte_cryptodev_get_private_session_size` is replaced with `rte_cryptodev_sym_get_private_session_size`

- cryptodev: Feature flag `RTE_CRYPTODEV_FF_MBUF_SCATTER_GATHER` is replaced with the following more explicit flags:

- `RTE_CRYPTODEV_FF_IN_PLACE_SGL`
- `RTE_CRYPTODEV_FF_OOP_SGL_IN_SGL_OUT`
- `RTE_CRYPTODEV_FF_OOP_SGL_IN_LB_OUT`
- `RTE_CRYPTODEV_FF_OOP_LB_IN_SGL_OUT`
- `RTE_CRYPTODEV_FF_OOP_LB_IN_LB_OUT`

- cryptodev: Renamed cryptodev experimental APIs:

Used `user_data` instead of `private_data` in following APIs to avoid confusion with the existing session parameter `sess_private_data[]` and related APIs.

- `rte_cryptodev_sym_session_set_private_data()` changed to `rte_cryptodev_sym_session_set_user_data()`
- `rte_cryptodev_sym_session_get_private_data()` changed to `rte_cryptodev_sym_session_get_user_data()`

- **compressdev:** Feature flag `RTE_COMP_FF_MBUF_SCATTER_GATHER` is replaced with the following more explicit flags:
 - `RTE_COMP_FF_OOP_SGL_IN_SGL_OUT`
 - `RTE_COMP_FF_OOP_SGL_IN_LB_OUT`
 - `RTE_COMP_FF_OOP_LB_IN_SGL_OUT`

6.3 Shared Library Versions

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

```

librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
librte_bpf.so.1
librte_bus_dpaa.so.1
librte_bus_fslmc.so.1
librte_bus_pci.so.1
librte_bus_vdev.so.1
+ librte_bus_vmbus.so.1
librte_cfgfile.so.2
librte_cmdline.so.2
librte_common_octeontx.so.1
librte_compressdev.so.1
+ librte_cryptodev.so.5
librte_distributor.so.1
+ librte_eal.so.8
+ librte_ethdev.so.10
+ librte_eventdev.so.5
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.4
+ librte_mempool.so.5
librte_meter.so.2
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_dpaa2_cmdif.so.1
librte_pmd_dpaa2_qdma.so.1
librte_pmd_ring.so.2
librte_pmd_softnic.so.1
librte_pmd_vhost.so.2
librte_port.so.3
librte_power.so.1
librte_rawdev.so.1
librte_reorder.so.1

```

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

6.4 Tested Platforms

- Intel(R) platforms with Intel(R) NICs combinations
 - CPU
 - * Intel(R) Atom(TM) CPU C3858 @ 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 v4 @ 2.20GHz
 - * 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
 - * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
 - OS:
 - * CentOS 7.4
 - * Fedora 25
 - * Fedora 27
 - * Fedora 28
 - * FreeBSD 11.1
 - * Red Hat Enterprise Linux Server release 7.5
 - * SUSE Enterprise Linux 12
 - * Wind River Linux 8
 - * Ubuntu 14.04
 - * Ubuntu 16.04
 - * Ubuntu 16.10
 - * Ubuntu 17.10
 - * Ubuntu 18.04
 - 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 0x80003221
 - Device id (pf/vf): 8086:1572 / 8086:154c
 - Driver version: 2.4.6 (i40e)
- * Intel Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 3.33 0x80000fd5 0.0.0
 - Device id (pf/vf): 8086:37d0 / 8086:37cd
 - Driver version: 2.4.3 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 6.01 0x80003221
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.4.6 (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: 2.4.6 (i40e)
- * Intel(R) Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x80000dda
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
 - * 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 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.3 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.2 (Maipo)
 - * Ubuntu 18.04
 - * Ubuntu 17.10
 - * Ubuntu 16.04
 - * SUSE Linux Enterprise Server 15
- MLNX_OFED: 4.3-2.0.2.0
- MLNX_OFED: 4.4-2.0.1.0
- 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: 2.42.5000
 - * 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 and above
 - * 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 and above
 - * 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 and above
 - * 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 and above
 - * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
 - Firmware version: 12.21.1000 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- ARM platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Qualcomm ARM 1.1 2500MHz
 - OS:
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - NICs:
 - * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8
 - Device ID: 15b3:1015
 - Firmware version: 14.23.1000
 - * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.23.1000
- Mellanox BlueField SmartNIC
 - **Mellanox(R) BlueField SmartNIC MT416842 (2x25G)**
 - * Host interface: PCI Express 3.0 x16
 - * Device ID: 15b3:a2d2
 - * Firmware version: 18.99.3950
 - SoC ARM cores running OS: * CentOS Linux release 7.4.1708 (AltArch) * Mellanox MLNX_OFED 4.2-1.4.21.0
 - DPDK application running on ARM cores inside SmartNIC
 - BlueField representors support planned for next release.

7.1 New Features

- **Reworked memory subsystem.**

Memory subsystem has been reworked to support new functionality.

On Linux, support for reserving/unreserving hugepage memory at runtime has been added, so applications no longer need to pre-reserve memory at startup. Due to reorganized internal workings of memory subsystem, any memory allocated through `rte_malloc()` or `rte_memzone_reserve()` is no longer guaranteed to be IOVA-contiguous.

This functionality has introduced the following changes:

- `rte_eal_get_physmem_layout()` was removed.
- A new flag for memzone reservation (`RTE_MEMZONE_IOVA_CONTIG`) was added to ensure reserved memory will be IOVA-contiguous, for use with device drivers and other cases requiring such memory.
- New callbacks for memory allocation/deallocation events, allowing users (or drivers) to be notified of new memory being allocated or deallocated
- New callbacks for validating memory allocations above a specified limit, allowing user to permit or deny memory allocations.
- A new command-line switch `--legacy-mem` to enable EAL behavior similar to how older versions of DPDK worked (memory segments that are IOVA-contiguous, but hugepages are reserved at startup only, and can never be released).
- A new command-line switch `--single-file-segments` to put all memory segments within a segment list in a single file.
- A set of convenience function calls to look up and iterate over allocated memory segments.
- `-m` and `--socket-mem` command-line arguments now carry an additional meaning and mark pre-reserved hugepages as “unfree-able”, thereby acting as a mechanism guaranteeing minimum availability of hugepage memory to the application.

Reserving/unreserving memory at runtime is not currently supported on FreeBSD.

- **Added bucket mempool driver.**

Added a bucket mempool driver which provides a way to allocate contiguous block of objects. The number of objects in the block depends on how many objects fit in the `RTE_DRIVER_MEMPOOL_BUCKET_SIZE_KB` memory chunk which is a build time option.

The number may be obtained using `rte_mempool_ops_get_info()` API. Contiguous blocks may be allocated using `rte_mempool_get_contig_blocks()` API.

- **Added support for port representors.**

Added DPDK port representors (also known as “VF representors” in the specific context of VFs), which are to DPDK what the Ethernet switch device driver model (**switchdev**) is to Linux, and which can be thought as a software “patch panel” front-end for applications. DPDK port representors are implemented as additional virtual Ethernet device (**ethdev**) instances, spawned on an as-needed basis through configuration parameters passed to the driver of the underlying device using `devargs`.

- **Added support for VXLAN and NVGRE tunnel endpoint.**

New actions types have been added to support encapsulation and decapsulation operations for a tunnel endpoint. The new action types are `RTE_FLOW_ACTION_TYPE_[VXLAN/NVGRE]_ENCAP`, `RTE_FLOW_ACTION_TYPE_[VXLAN/NVGRE]_DECAP`, `RTE_FLOW_ACTION_TYPE_JUMP`. A new item type `RTE_FLOW_ACTION_TYPE_MARK` has been added to match a flow against a previously marked flow. A shared counter has also been introduced to the flow API to count a group of flows.

- **Added PMD-recommended Tx and Rx parameters.**

Applications can now query drivers for device-tuned values of ring sizes, burst sizes, and number of queues.

- **Added RSS hash and key update to CXGBE PMD.**

Added support for updating the RSS hash and key to the CXGBE PMD.

- **Added CXGBE VF PMD.**

CXGBE VF Poll Mode Driver has been added to run DPDK over Chelsio T5/T6 NIC VF instances.

- **Updated mlx5 driver.**

Updated the mlx5 driver including the following changes:

- Introduced Multi-packet Rx to enable 100Gb/sec with 64B frames.
- Support for being run by non-root users given a reduced set of capabilities `CAP_NET_ADMIN`, `CAP_NET_RAW` and `CAP_IPC_LOCK`.
- Support for TSO and checksum for generic UDP and IP tunnels.
- Support for inner checksum and RSS for GRE, VXLAN-GPE, MPLSoGRE and MPLSoUDP tunnels.
- Accommodate the new memory hotplug model.
- Support for non virtually contiguous mempools.
- Support for MAC adding along with allmulti and promiscuous modes from VF.
- Support for Mellanox BlueField SoC device.
- Support for PMD defaults for queue number and depth to improve the out of the box performance.

- **Updated mlx4 driver.**

Updated the mlx4 driver including the following changes:

- Support for to being run by non-root users given a reduced set of capabilities CAP_NET_ADMIN, CAP_NET_RAW and CAP_IPC_LOCK.
- Supported CRC strip toggling.
- Accommodate the new memory hotplug model.
- Support non virtually contiguous mempools.
- Dropped support for Mellanox OFED 4.2.

- **Updated Solarflare network PMD.**

Updated the sfc_efx driver including the following changes:

- Added support for Solarflare XtremeScale X2xxx family adapters.
- Added support for NVGRE, VXLAN and GENEVE filters in flow API.
- Added support for DROP action in flow API.
- Added support for equal stride super-buffer Rx mode (X2xxx only).
- Added support for MARK and FLAG actions in flow API (X2xxx only).

- **Added Ethernet poll mode driver for AMD XGBE devices.**

Added the new `axgbe` ethernet poll mode driver for AMD XGBE devices. See the `../nics/axgbe` nic driver guide for more details on this new driver.

- **Updated szedata2 PMD.**

Added support for new NFB-200G2QL card. A new API was introduced in the `libsze2` library which the `szedata2` PMD depends on, thus the new version of the library was needed. New versions of the packages are available and the minimum required version is 4.4.1.

- **Added support for Broadcom NetXtreme-S (BCM58800) family of controllers (aka Stingray).**

Added support for the Broadcom NetXtreme-S (BCM58800) family of controllers (aka Stingray). The BCM58800 devices feature a NetXtreme E-Series advanced network controller, a high-performance ARM CPU block, PCI Express (PCIe) Gen3 interfaces, key accelerators for compute offload and a high-speed memory subsystem including L3 cache and DDR4 interfaces, all interconnected by a coherent Network-on-chip (NOC) fabric.

The ARM CPU subsystem features eight ARMv8 Cortex-A72 CPUs at 3.0 GHz, arranged in a multi-cluster configuration.

- **Added vDPA in vhost-user lib.**

Added support for selective datapath in the `vhost-user lib`. vDPA stands for vhost Data Path Acceleration. It supports virtio ring compatible devices to serve the virtio driver directly to enable datapath acceleration.

- **Added IFCVF vDPA driver.**

Added IFCVF vDPA driver to support Intel FPGA 100G VF devices. IFCVF works as a HW vhost data path accelerator, it supports live migration and is compatible with virtio 0.95 and 1.0.

This driver registers the ifcvf vDPA driver to vhost lib, when virtio connects. With the help of the registered vDPA driver the assigned VF gets configured to Rx/Tx directly to VM's virtio vrings.

- **Added support for vhost dequeue interrupt mode.**

Added support for vhost dequeue interrupt mode to release CPUs to others when there is no data to transmit. Applications can register an epoll event file descriptor to associate Rx queues with interrupt vectors.

- **Added support for virtio-user server mode.**

In a container environment if the vhost-user backend restarts, there's no way for it to reconnect to virtio-user. To address this, support for server mode has been added. In this mode the socket file is created by virtio-user, which the backend connects to. This means that if the backend restarts, it can reconnect to virtio-user and continue communications.

- **Added crypto workload support to vhost library.**

New APIs have been introduced in the vhost library to enable virtio crypto support including session creation/deletion handling and translating virtio-crypto requests into DPDK crypto operations. A sample application has also been introduced.

- **Added virtio crypto PMD.**

Added a new Poll Mode Driver for virtio crypto devices, which provides AES-CBC ciphering and AES-CBC with HMAC-SHA1 algorithm-chaining. See the `../cryptodevs/virtio` crypto driver guide for more details on this new driver.

- **Added AMD CCP Crypto PMD.**

Added the new `ccp` crypto driver for AMD CCP devices. See the `../cryptodevs/ccp` crypto driver guide for more details on this new driver.

- **Updated AESNI MB PMD.**

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

- AES-CMAC (128-bit key).

- **Added the Compressdev Library, a generic compression service library.**

Added the Compressdev library which provides an API for offload of compression and decompression operations to hardware or software accelerator devices.

- **Added a new compression poll mode driver using Intels ISA-L.**

Added the new `ISA-L` compression driver, for compression and decompression operations in software. See the `../compressdevs/isal` compression driver guide for details on this new driver.

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

The Event Timer Adapter Library extends the event-based model by introducing APIs that allow applications to arm/cancel event timers that generate timer expiry events. This new type of event is scheduled by an event device along with existing types of events.

- **Added OcteonTx TIM Driver (Event timer adapter).**

The OcteonTx Timer block enables software to schedule events for a future time, it is exposed to an application via the Event timer adapter library.

See the `../eventdevs/octeontx` guide for more details

- **Added Event Crypto Adapter Library.**

Added the Event Crypto Adapter Library. This library extends the event-based model by introducing APIs that allow applications to enqueue/dequeue crypto operations to/from cryptodev as events scheduled by an event device.

- **Added Ifpga Bus, a generic Intel FPGA Bus library.**

Added the Ifpga Bus library which provides support for integrating any Intel FPGA device with the DPDK framework. It provides Intel FPGA Partial Bit Stream AFU (Accelerated Function Unit) scan and drivers probe.

- **Added IFPGA (Intel FPGA) Rawdev Driver.**

Added a new Rawdev driver called IFPGA (Intel FPGA) Rawdev Driver, which cooperates with OPAE (Open Programmable Acceleration Engine) shared code to provide common FPGA management ops for FPGA operation.

See the `../rawdevs/ifpga` programmer's guide for more details.

- **Added DPAA2 QDMA Driver (in rawdev).**

The DPAA2 QDMA is an implementation of the rawdev API, that provide a means of initiating a DMA transaction from CPU. The initiated DMA is performed without the CPU being involved in the actual DMA transaction.

See the `../rawdevs/dpaa2_qdma` guide for more details.

- **Added DPAA2 Command Interface Driver (in rawdev).**

The DPAA2 CMDIF is an implementation of the rawdev API, that provides communication between the GPP and NXP's QorIQ based AIOP Block (Firmware). Advanced IO Processor i.e. AIOP are clusters of programmable RISC engines optimized for flexible networking and I/O operations. The communication between GPP and AIOP is achieved via using DPCI devices exposed by MC for GPP <-> AIOP interaction.

See the `../rawdevs/dpaa2_cmdif` guide for more details.

- **Added device event monitor framework.**

Added a general device event monitor framework to EAL, for device dynamic management to facilitate device hotplug awareness and associated actions. The list of new APIs is:

- `rte_dev_event_monitor_start` and `rte_dev_event_monitor_stop` for the event monitor enabling and disabling.
- `rte_dev_event_callback_register` and `rte_dev_event_callback_unregister` for registering and un-registering user callbacks.

Linux uevent is supported as a backend of this device event notification framework.

- **Added support for procinfo and pdump on eth vdev.**

For ethernet virtual devices (like TAP, PCAP, etc.), with this feature, we can get stats/xstats on shared memory from a secondary process, and also pdump packets on those virtual devices.

- **Enhancements to the Packet Framework Library.**

Design and development of new API functions for Packet Framework library that implement a common set of actions such as traffic metering, packet encapsulation, network address translation, TTL update, etc., for pipeline table and input ports to speed up application development. The API

functions includes creating action profiles, registering actions to the profiles, instantiating action profiles for pipeline table and input ports, etc.

- **Added the BPF Library.**

The BPF Library provides the ability to load and execute Enhanced Berkeley Packet Filters (eBPF) within user-space DPDK applications. It also introduces a basic framework to load/unload BPF-based filters on Eth devices (right now only via SW RX/TX callbacks). It also adds a dependency on libelf.

7.2 API Changes

- service cores: No longer marked as experimental.

The service cores functions are no longer marked as experimental, and have become part of the normal DPDK API and ABI. Any future ABI changes will be announced at least one release before the ABI change is made. There are no ABI breaking changes planned.

- eal: The `rte_lcore_has_role()` return value changed.

This function now returns true or false, respectively, rather than 0 or < 0 for success or failure. It makes use of the function more intuitive.

- mempool: The capability flags and related functions have been removed.

Flags `MEMPOOL_F_CAPA_PHYS_CONTIG` and `MEMPOOL_F_CAPA_BLK_ALIGNED_OBJECTS` were used by ooctontx mempool driver to customize generic mempool library behavior. Now the new driver callbacks `calc_mem_size` and `populate` may be used to achieve it without specific knowledge in the generic code.

- mempool: The following xmem functions have been deprecated:

- `rte_mempool_xmem_create`
- `rte_mempool_xmem_size`
- `rte_mempool_xmem_usage`
- `rte_mempool_populate_iova_tab`

- mbuf: The control mbuf API has been removed in v18.05. The impacted functions and macros are:

- `rte_ctrlmbuf_init()`
- `rte_ctrlmbuf_alloc()`
- `rte_ctrlmbuf_free()`
- `rte_ctrlmbuf_data()`
- `rte_ctrlmbuf_len()`
- `rte_is_ctrlmbuf()`
- `CTRL_MBUF_FLAG`

The packet mbuf API should be used as a replacement.

- meter: API updated to accommodate configuration profiles.

The meter API has been changed to support meter configuration profiles. The configuration profile represents the set of configuration parameters for a given meter object, such as the rates and sizes for the token buckets. These configuration parameters were previously part of the meter object internal data structure. The separation of the configuration parameters from the meter object data structure results in reducing its memory footprint which helps in better cache utilization when a large number of meter objects are used.

- **ethdev:** The function `rte_eth_dev_count()`, often mis-used to iterate over ports, is deprecated and replaced by `rte_eth_dev_count_avail()`. There is also a new function `rte_eth_dev_count_total()` to get the total number of allocated ports, available or not. The hotplug-proof applications should use `RTE_ETH_FOREACH_DEV` or `RTE_ETH_FOREACH_DEV_OWNED_BY` as port iterators.
- **ethdev:** In struct `struct rte_eth_dev_info`, field `rte_pci_device *pci_dev` has been replaced with field `struct rte_device *device`.
- **ethdev:** Changes to the semantics of `rte_eth_dev_configure()` parameters.

If both the `nb_rx_q` and `nb_tx_q` parameters are zero, `rte_eth_dev_configure()` will now use PMD-recommended queue sizes, or if recommendations are not provided by the PMD the function will use ethdev fall-back values. Previously setting both of the parameters to zero would have resulted in `-EINVAL` being returned.

- **ethdev:** Changes to the semantics of `rte_eth_rx_queue_setup()` parameters.

If the `nb_rx_desc` parameter is zero, `rte_eth_rx_queue_setup` will now use the PMD-recommended Rx ring size, or in the case where the PMD does not provide a recommendation, will use an ethdev-provided fall-back value. Previously, setting `nb_rx_desc` to zero would have resulted in an error.

- **ethdev:** Changes to the semantics of `rte_eth_tx_queue_setup()` parameters.

If the `nb_tx_desc` parameter is zero, `rte_eth_tx_queue_setup` will now use the PMD-recommended Tx ring size, or in the case where the PMD does not provide a recommendation, will use an ethdev-provided fall-back value. Previously, setting `nb_tx_desc` to zero would have resulted in an error.

- **ethdev:** Several changes were made to the flow API.
 - The unused DUP action was removed.
 - Actions semantics in flow rules: list order now matters (“first to last” instead of “all simultaneously”), repeated actions are now all performed, and they do not individually have (non-)terminating properties anymore.
 - Flow rules are now always terminating unless a PASSTHRU action is present.
 - C99-style flexible arrays were replaced with standard pointers in RSS action and in RAW pattern item structures due to compatibility issues.
 - The RSS action was modified to not rely on external `struct rte_eth_rss_conf` anymore to instead expose its own and more appropriately named configuration fields directly (`rss_conf->rss_key => key`, `rss_conf->rss_key_len => key_len`, `rss_conf->rss_hf => types`, `num => queue_num`), and the addition of missing RSS parameters (`func` for RSS hash function to apply and `level` for the encapsulation level).
 - The VLAN pattern item (`struct rte_flow_item_vlan`) was modified to include inner EtherType instead of outer TPID. Its default mask was also modified to cover the VID

- part (lower 12 bits) of TCI only.
- A new transfer attribute was added to `struct rte_flow_attr` in order to clarify the behavior of some pattern items.
 - PF and VF pattern items are now only accepted by PMDs that implement them (bnxt and i40e) when the transfer attribute is also present, for consistency.
 - Pattern item `PORT` was renamed `PHY_PORT` to avoid confusion with DPDK port IDs.
 - An action counterpart to the `PHY_PORT` pattern item was added in order to redirect matching traffic to a specific physical port.
 - `PORT_ID` pattern item and actions were added to match and target DPDK port IDs at a higher level than `PHY_PORT`.
 - `RTE_FLOW_ACTION_TYPE_[VXLAN/NVGRE]_ENCAP` action items were added to support tunnel encapsulation operation for VXLAN and NVGRE type tunnel endpoint.
 - `RTE_FLOW_ACTION_TYPE_[VXLAN/NVGRE]_DECAP` action items were added to support tunnel decapsulation operation for VXLAN and NVGRE type tunnel endpoint.
 - `RTE_FLOW_ACTION_TYPE_JUMP` action item was added to support a matched flow to be redirected to the specific group.
 - `RTE_FLOW_ACTION_TYPE_MARK` item type has been added to match a flow against a previously marked flow.
- ethdev: Change flow APIs regarding count action:
 - `rte_flow_create()` API count action now requires the `struct rte_flow_action_count`.
 - `rte_flow_query()` API parameter changed from action type to action structure.
 - ethdev: Changes to offload API

A pure per-port offloading isn't requested to be repeated in `[rt]x_conf->offloads` to `rte_eth_[rt]x_queue_setup()`. Now any offloading enabled in `rte_eth_dev_configure()` can't be disabled by `rte_eth_[rt]x_queue_setup()`. Any new added offloading which has not been enabled in `rte_eth_dev_configure()` and is requested to be enabled in `rte_eth_[rt]x_queue_setup()` must be per-queue type, or otherwise trigger an error log.
 - ethdev: Runtime queue setup

`rte_eth_rx_queue_setup` and `rte_eth_tx_queue_setup` can be called after `rte_eth_dev_start` if the device supports runtime queue setup. The device driver can expose this capability through `rte_eth_dev_info_get`. A Rx or Tx queue set up at runtime need to be started explicitly by `rte_eth_dev_rx_queue_start` or `rte_eth_dev_tx_queue_start`.

7.3 ABI Changes

- ring: The alignment constraints on the ring structure has been relaxed to one cache line instead of two, and an empty cache line padding is added between the producer and consumer structures. The size of the structure and the offset of the fields remains the same on platforms with 64B cache line, but changes on other platforms.

- mempool: Some ops have changed.

A new callback `calc_mem_size` has been added to `rte_mempool_ops` to allow customization of the required memory size calculation. A new callback `populate` has been added to `rte_mempool_ops` to allow customized object population. Callback `get_capabilities` has been removed from `rte_mempool_ops` since its features are covered by `calc_mem_size` and `populate` callbacks. Callback `register_memory_area` has been removed from `rte_mempool_ops` since the new callback `populate` may be used instead of it.

- ethdev: Additional fields in `rte_eth_dev_info`.

The `rte_eth_dev_info` structure has had two extra entries appended to the end of it: `default_rxportconf` and `default_txportconf`. Each of these in turn are `rte_eth_dev_portconf` structures containing three fields of type `uint16_t`: `burst_size`, `ring_size`, and `nb_queues`. These are parameter values recommended for use by the PMD.

- ethdev: ABI for all flow API functions was updated.

This includes functions `rte_flow_copy`, `rte_flow_create`, `rte_flow_destroy`, `rte_flow_error_set`, `rte_flow_flush`, `rte_flow_isolate`, `rte_flow_query` and `rte_flow_validate`, due to changes in error type definitions (enum `rte_flow_error_type`), removal of the unused DUP action (enum `rte_flow_action_type`), modified behavior for flow rule actions (see API changes), removal of C99 flexible array from RAW pattern item (struct `rte_flow_item_raw`), complete rework of the RSS action definition (struct `rte_flow_action_rss`), sanity fix in the VLAN pattern item (struct `rte_flow_item_vlan`) and new transfer attribute (struct `rte_flow_attr`).

- bbdev: New parameter added to `rte_bbdev_op_cap_turbo_dec`.

A new parameter `max_llr_modulus` has been added to `rte_bbdev_op_cap_turbo_dec` structure to specify maximal LLR (likelihood ratio) absolute value.

- bbdev: Queue Groups split into UL/DL Groups.

Queue Groups have been split into UL/DL Groups in the Turbo Software Driver. They are independent for Decode/Encode. `rte_bbdev_driver_info` reflects introduced changes.

7.4 Known Issues

- **Secondary process launch is not reliable.**

Recent memory hotplug patches have made multiprocess startup less reliable than it was in past releases. A number of workarounds are known to work depending on the circumstances. As such it isn't recommended to use the secondary process mechanism for critical systems. The underlying issues will be addressed in upcoming releases.

The issue is explained in more detail, including potential workarounds, in the Bugzilla entry referenced below.

Bugzilla entry: https://bugs.dpdk.org/show_bug.cgi?id=50

- **pdump is not compatible with old applications.**

As we changed to use generic multi-process communication for pdump negotiation instead of previous dedicated unix socket way, pdump applications, including the dpdk-pdump example and any other applications using `librte_pdump`, will not work with older version DPDK primary applications.

- **rte_abort takes a long time on FreeBSD.**

DPDK processes now allocates a large area of virtual memory address space. As a result `rte_abort` on FreeBSD now dumps the contents of the whole reserved memory range, not just the used portion, to a core dump file. Writing this large core file can take a significant amount of time, causing processes to appear to hang on the system.

The work around for the issue is to set the system resource limits for core dumps before running any tests, e.g. `limit coredumpsize 0`. This will effectively disable core dumps on FreeBSD. If they are not to be completely disabled, a suitable limit, e.g. 1G might be specified instead of 0. This needs to be run per-shell session, or before every test run. This change can also be made persistent by adding `kern.coredump=0` to `/etc/sysctl.conf`.

Bugzilla entry: https://bugs.dpdk.org/show_bug.cgi?id=53

- **ixgbe PMD crash on hotplug detach when no VF created.**

ixgbe PMD uninit path cause null pointer dereference because of port representor cleanup when number of VF is zero.

Bugzilla entry: https://bugs.dpdk.org/show_bug.cgi?id=57

- **Bonding PMD may fail to accept new slave ports in certain conditions.**

In certain conditions when using `testpmd`, bonding may fail to register new slave ports.

Bugzilla entry: https://bugs.dpdk.org/show_bug.cgi?id=52.

- **Unexpected performance regression in Vhost library.**

Patches fixing CVE-2018-1059 were expected to introduce a small performance drop. However, in some setups, bigger performance drops have been measured when running micro-benchmarks.

Bugzilla entry: https://bugs.dpdk.org/show_bug.cgi?id=48

7.5 Shared Library Versions

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

```
librte_acl.so.2
librte_bbdev.so.1
librte_bitratestats.so.2
+ librte_bpf.so.1
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_common_octeontx.so.1
+ librte_compressdev.so.1
librte_cryptodev.so.4
librte_distributor.so.1
+ librte_eal.so.7
+ librte_ethdev.so.9
+ librte_eventdev.so.4
```

```

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.4
+ librte_mempool.so.4
+ librte_meter.so.2
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_dpaa2_cmdif.so.1
+ librte_pmd_dpaa2_qdma.so.1
librte_pmd_ring.so.2
librte_pmd_softnic.so.1
librte_pmd_vhost.so.2
librte_port.so.3
librte_power.so.1
librte_rawdev.so.1
librte_reorder.so.1
+ librte_ring.so.2
librte_sched.so.1
librte_security.so.1
librte_table.so.3
librte_timer.so.1
librte_vhost.so.3

```

7.6 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-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 v4 @ 2.20GHz
 - * 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
 - * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
 - OS:

- * CentOS 7.4
- * Fedora 25
- * Fedora 27
- * Fedora 28
- * FreeBSD 11.1
- * Red Hat Enterprise Linux Server release 7.3
- * SUSE Enterprise Linux 12
- * Wind River Linux 8
- * Ubuntu 14.04
- * Ubuntu 16.04
- * Ubuntu 16.10
- * Ubuntu 17.10

– 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 0x80003221
 - Device id (pf/vf): 8086:1572 / 8086:154c
 - Driver version: 2.4.6 (i40e)
- * Intel Corporation Ethernet Connection X722 for 10GbE SFP+ (4x10G)
 - Firmware version: 3.33 0x8000fd5 0.0.0
 - Device id (pf/vf): 8086:37d0 / 8086:37cd
 - Driver version: 2.4.3 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 6.01 0x80003221
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.4.6 (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: 2.4.6 (i40e)
- * Intel(R) Corporation I350 Gigabit Network Connection
 - Firmware version: 1.63, 0x80000dda
 - Device id (pf/vf): 8086:1521 / 8086:1520
 - Driver version: 5.4.0-k (igb)
- Intel(R) platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
 - * 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 7.5 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.3 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.2 (Maipo)
 - * Ubuntu 18.04
 - * Ubuntu 17.10
 - * Ubuntu 16.10
 - * Ubuntu 16.04
 - * SUSE Linux Enterprise Server 15
 - MLNX_OFED: 4.2-1.0.0.0
 - MLNX_OFED: 4.3-2.0.2.0
 - 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: 2.42.5000
 - * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- Firmware version: 12.21.1000 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- ARM platforms with Mellanox(R) NICs combinations
 - CPU:
 - * Qualcomm ARM 1.1 2500MHz
 - OS:
 - * Red Hat Enterprise Linux Server release 7.5 (Maipo)
 - NICs:
 - * Mellanox(R) ConnectX(R)-4 Lx 25G MCX4121A-ACAT (2x25G)
 - Host interface: PCI Express 3.0 x8

- Device ID: 15b3:1015
- Firmware version: 14.22.0428
- * Mellanox(R) ConnectX(R)-5 100G MCX556A-ECAT (2x100G)
 - Host interface: PCI Express 3.0 x16
 - Device ID: 15b3:1017
 - Firmware version: 16.22.0428
- ARM SoC combinations from Cavium (with integrated NICs)
 - SoC:
 - * Cavium CN81xx
 - * Cavium CN83xx
 - OS:
 - * Ubuntu 16.04.2 LTS with Cavium SDK-6.2.0-Patch2 release support package.
- ARM SoC combinations from NXP (with integrated NICs)
 - SoC:
 - * NXP/Freescale QorIQ LS1046A with ARM Cortex A72
 - * NXP/Freescale QorIQ LS2088A with ARM Cortex A72
 - OS:
 - * Ubuntu 16.04.3 LTS with NXP QorIQ LSDK 1803 support packages

8.1 New Features

- **Added function to allow releasing internal EAL resources on exit.**

During `rte_eal_init()` EAL allocates memory from hugepages to enable its core libraries to perform their tasks. The `rte_eal_cleanup()` function releases these resources, ensuring that no hugepage memory is leaked. It is expected that all DPDK applications call `rte_eal_cleanup()` before exiting. Not calling this function could result in leaking hugepages, leading to failure during initialization of secondary processes.

- **Added igb, ixgbe and i40e ethernet driver to support RSS with flow API.**

Added support for igb, ixgbe and i40e NICs with existing RSS configuration using the `rte_flow` API.

Also enabled queue region configuration using the `rte_flow` API for i40e.

- **Updated i40e driver to support PPPoE/PPPoL2TP.**

Updated i40e PMD to support PPPoE/PPPoL2TP with PPPoE/PPPoL2TP supporting profiles which can be programmed by dynamic device personalization (DDP) process.

- **Added MAC loopback support for i40e.**

Added MAC loopback support for i40e in order to support test tasks requested by users. It will setup Tx -> Rx loopback link according to the device configuration.

- **Added support of run time determination of number of queues per i40e VF.**

The number of queue per VF is determined by its host PF. If the PCI address of an i40e PF is `aaaa:bb.cc`, the number of queues per VF can be configured with EAL parameter like `-w aaaa:bb.cc,queue-num-per-vf=n`. The value `n` can be 1, 2, 4, 8 or 16. If no such parameter is configured, the number of queues per VF is 4 by default.

- **Updated mlx5 driver.**

Updated the mlx5 driver including the following changes:

- Enabled compilation as a plugin, thus removed the mandatory dependency with `rdma-core`. With the special compilation, the `rdma-core` libraries will be loaded only in case Mellanox device is being used. For binaries creation the PMD can be enabled, still not requiring from every end user to install `rdma-core`.
- Improved multi-segment packet performance.
- Changed driver name to use the PCI address to be compatible with OVS-DPDK APIs.

- Extended statistics for physical port packet/byte counters.
- Converted to the new offloads API.
- Supported device removal check operation.

- **Updated mlx4 driver.**

Updated the mlx4 driver including the following changes:

- Enabled compilation as a plugin, thus removed the mandatory dependency with rdma-core. With the special compilation, the rdma-core libraries will be loaded only in case Mellanox device is being used. For binaries creation the PMD can be enabled, still not requiring from every end user to install rdma-core.
- Improved data path performance.
- Converted to the new offloads API.
- Supported device removal check operation.

- **Added NVGRE and UDP tunnels support in Solarflare network PMD.**

Added support for NVGRE, VXLAN and GENEVE tunnels.

- Added support for UDP tunnel ports configuration.
- Added tunneled packets classification.
- Added inner checksum offload.

- **Added AVF (Adaptive Virtual Function) net PMD.**

Added a new net PMD called AVF (Adaptive Virtual Function), which supports Intel® Ethernet Adaptive Virtual Function (AVF) with features such as:

- Basic Rx/Tx burst
- SSE vectorized Rx/Tx burst
- Promiscuous mode
- MAC/VLAN offload
- Checksum offload
- TSO offload
- Jumbo frame and MTU setting
- RSS configuration
- stats
- Rx/Tx descriptor status
- Link status update/event

- **Added feature supports for live migration from vhost-net to vhost-user.**

Added feature supports for vhost-user to make live migration from vhost-net to vhost-user possible. The features include:

- VIRTIO_F_ANY_LAYOUT
- VIRTIO_F_EVENT_IDX

- `VIRTIO_NET_F_GUEST_ECN`, `VIRTIO_NET_F_HOST_ECN`
- `VIRTIO_NET_F_GUEST_UFO`, `VIRTIO_NET_F_HOST_UFO`
- `VIRTIO_NET_F_GSO`

Also added `VIRTIO_NET_F_GUEST_ANNOUNCE` feature support in virtio pmd. In a scenario where the vhost backend doesn't have the ability to generate RARP packets, the VM running virtio pmd can still be live migrated if `VIRTIO_NET_F_GUEST_ANNOUNCE` feature is negotiated.

- **Updated the AESNI-MB PMD.**

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

- AES-CCM algorithm.

- **Updated the DPAA_SEC crypto driver to support `rte_security`.**

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

- **Added Wireless Base Band Device (bbdev) abstraction.**

The Wireless Baseband Device library is an acceleration abstraction framework for 3gpp Layer 1 processing functions that provides a common programming interface for seamless operation on integrated or discrete hardware accelerators or using optimized software libraries for signal processing.

The current release only supports 3GPP CRC, Turbo Coding and Rate Matching operations, as specified in 3GPP TS 36.212.

See the `../prog_guide/bbdev` programmer's guide for more details.

- **Added New eventdev Ordered Packet Distribution Library (OPDL) PMD.**

The OPDL (Ordered Packet Distribution Library) eventdev is a specific implementation of the eventdev API. It is particularly suited to packet processing workloads that have high throughput and low latency requirements. All packets follow the same path through the device. The order in which packets follow is determined by the order in which queues are set up. Events are left on the ring until they are transmitted. As a result packets do not go out of order.

With this change, applications can use the OPDL PMD via the eventdev api.

- **Added new pipeline use case for `dpdk-test-eventdev` application.**

Added a new "pipeline" use case for the `dpdk-test-eventdev` application. The pipeline case can be used to simulate various stages in a real world application from packet receive to transmit while maintaining the packet ordering. It can also be used to measure the performance of the event device across the stages of the pipeline.

The pipeline use case has been made generic to work with all the event devices based on the capabilities.

- **Updated Eventdev sample application to support event devices based on capability.**

Updated the Eventdev pipeline sample application to support various types of pipelines based on the capabilities of the attached event and ethernet devices. Also, renamed the application from software PMD specific `eventdev_pipeline_sw_pmd` to the more generic `eventdev_pipeline`.

- **Added Rawdev, a generic device support library.**

The Rawdev library provides support for integrating any generic device type with the DPDK framework. Generic devices are those which do not have a pre-defined type within DPDK, for example, ethernet, crypto, event etc.

A set of northbound APIs have been defined which encompass a generic set of operations by allowing applications to interact with device using opaque structures/buffers. Also, southbound APIs provide a means of integrating devices either as part of a physical bus (PCI, FSLMC etc) or through `vdev`.

See the `../prog_guide/rawdev` programmer's guide for more details.

- **Added new multi-process communication channel.**

Added a generic channel in EAL for multi-process (primary/secondary) communication. Consumers of this channel need to register an action with an action name to response a message received; the actions will be identified by the action name and executed in the context of a new dedicated thread for this channel. The list of new APIs:

- `rte_mp_register` and `rte_mp_unregister` are for action (un)registration.
- `rte_mp_sendmsg` is for sending a message without blocking for a response.
- `rte_mp_request` is for sending a request message and will block until it gets a reply message which is sent from the peer by `rte_mp_reply`.

- **Added GRO support for VxLAN-tunneled packets.**

Added GRO support for VxLAN-tunneled packets. Supported VxLAN packets must contain an outer IPv4 header and inner TCP/IPv4 headers. VxLAN GRO doesn't check if input packets have correct checksums and doesn't update checksums for output packets. Additionally, it assumes the packets are complete (i.e., `MF==0 && frag_off==0`), when IP fragmentation is possible (i.e., `DF==0`).

- **Increased default Rx and Tx ring size in sample applications.**

Increased the default `RX_RING_SIZE` and `TX_RING_SIZE` to 1024 entries in `testpmd` and the sample applications to give better performance in the general case. The user should experiment with various Rx and Tx ring sizes for their specific application to get best performance.

- **Added new DPDK build system using the tools “meson” and “ninja” [EXPERIMENTAL].**

Added support for building DPDK using `meson` and `ninja`, which gives additional features, such as automatic build-time configuration, over the current build system using `make`. For instructions on how to do a DPDK build using the new system, see the instructions in `doc/build-sdk-meson.txt`.

Note: This new build system support is incomplete at this point and is added as experimental in this release. The existing build system using `make` is unaffected by these changes, and can continue to be used for this and subsequent releases until such time as it's deprecation is announced.

8.2 Shared Library Versions

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

```
librte_acl.so.2
+ librte_bbdev.so.1
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_rawdev.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
```

8.3 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 v4 @ 2.20GHz
 - * 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
 - * Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz
- OS:
- * CentOS 7.2
 - * Fedora 25
 - * Fedora 26
 - * Fedora 27
 - * FreeBSD 11
 - * Red Hat Enterprise Linux Server release 7.3
 - * SUSE Enterprise Linux 12
 - * Wind River Linux 8
 - * Ubuntu 14.04
 - * Ubuntu 16.04
 - * Ubuntu 16.10
 - * Ubuntu 17.10
- 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 0x80003221
 - Device id (pf/vf): 8086:1572 / 8086:154c
 - Driver version: 2.4.3 (i40e)
 - * Intel Corporation Ethernet Connection X722 for 10GBASE-T
 - firmware-version: 6.01 0x80003221
 - Device id: 8086:37d2 / 8086:154c

- Driver version: 2.4.3 (i40e)
- * Intel(R) Ethernet Converged Network Adapter XXV710-DA2 (2x25G)
 - Firmware version: 6.01 0x80003221
 - Device id (pf/vf): 8086:158b / 8086:154c
 - Driver version: 2.4.3 (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: 2.4.3 (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
 - CPU:
 - * 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 7.5 Beta (Maipo)
 - * Red Hat Enterprise Linux Server release 7.4 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.3 (Maipo)
 - * Red Hat Enterprise Linux Server release 7.2 (Maipo)
 - * Ubuntu 17.10
 - * Ubuntu 16.10
 - * Ubuntu 16.04
 - MLNX_OFED: 4.2-1.0.0.0
 - MLNX_OFED: 4.3-0.1.6.0
 - 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above

- * 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 and above
 - Firmware version: 12.21.1000 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- * 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 and above
- ARM platforms with Mellanox(R) NICs combinations
 - CPU:
 - * 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

DPDK RELEASE 17.11

9.1 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 `mvnl` net driver for Marvell Armada 7k/8k devices. See the `./nics/mvpp2` 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/mvsam` 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 i40e driver to support GTP-C/GTP-U.**

Updated `i40e` PMD to support GTP-C/GTP-U with GTP-C/GTP-U supporting profiles which can be programmed by dynamic device personalization (DDP) process.

- **Added the i40e ethernet driver to support queue region feature.**

This feature enable queue regions configuration for RSS in PF, so that different traffic classes or different packet classification types can be separated into different queues in different queue regions.

- **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.

9.2 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.

9.3 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.

9.4 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.

9.5 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.

9.6 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

```

9.7 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

10.1 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 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.

10.2 Known Issues

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

10.3 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`.

10.4 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`.

10.5 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

```

10.6 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)

11.1 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 2²¹.**

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 i40.

- **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.

- **Updated i40e driver to support MPLSoUDP/MPLSoGRE.**

Updated i40e PMD to support MPLSoUDP/MPLSoGRE with MPLSoUDP/MPLSoGRE supporting profiles which can be programmed by dynamic device personalization (DDP) process.

- **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.

11.2 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.

11.3 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.

11.4 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`

11.5 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.

11.6 Removed Items

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

11.7 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
```

11.8 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

12.1 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 `../prog_guide/rte_flow` 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.

12.2 Resolved Issues

12.2.1 Drivers

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

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

12.2.2 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.

12.3 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`.

12.4 ABI Changes

12.5 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
```

12.6 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)

13.1 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.

13.2 Resolved Issues

13.2.1 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.**

13.3 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 work around this by always supplying the digest virtual and physical address fields in the `rte_crypto_sym_op` for an auth operation.

13.4 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`.

13.5 ABI Changes

13.6 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
```

13.7 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

13.8 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

13.9 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

14.1 New Features

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

The size of the mempool structure is reduced if the per-core 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.

14.2 Resolved Issues

14.2.1 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.

14.2.2 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.**

14.2.3 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.

14.2.4 Examples

14.2.5 Other

14.3 Known Issues

14.4 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`.

14.5 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.

14.6 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
```

14.7 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

14.8 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)

14.9 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

15.1 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 ≥ 3.2 .
- **Added mlx5 RX VLAN stripping support.**

Added support for RX VLAN stripping.

Only available with Mellanox OFED ≥ 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 ≥ 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 ≥ 3.2 .
- **Added mlx5 TX VLAN insertion support.**

Added support for TX VLAN insertion.

Only available with Mellanox OFED ≥ 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.

15.2 Resolved Issues

15.2.1 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.

15.2.2 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 `dropless` is set to `yes`.

15.2.3 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.

15.3 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 renamed 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.

15.4 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.

15.5 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
```

15.6 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

15.7 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)

DPDK RELEASE 2.2

16.1 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.

16.2 Resolved Issues

16.2.1 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.

16.2.2 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 interoperability 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 (13 / 14 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 exceed 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.

16.2.3 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.

16.2.4 Examples

- **l3fwd: Fixed crash with IPv6.**
- **vhost_xen: Fixed compile error.**

16.2.5 Other

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

16.3 Known Issues

- Some drivers do not fill in the packet type when receiving. As the l3fwd 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.

16.4 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.

16.5 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.

16.6 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
```

17.1 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 mempools, 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 linux 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.

17.2 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 vfi MSI.**

- **eal/linux: Fix irq handling with igb_uio.**

Fixed an issue where 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 andbsd.**

- **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-lcore pending-lists (implemented as skip-lists).

- **vfiio: 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.**

17.3 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.

17.4 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.

17.5 ABI Changes

- The `rte_hash` structure has been changed to internal use only.

18.1 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.

19.1 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.

20.1 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.

20.2 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).

20.3 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).

20.4 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`.

20.5 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 lcore and is a cycle reference, not a time reference.

Affected Environment/Platform: All.

Driver/Module: Environment Abstraction Layer (EAL).

20.6 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`

20.7 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

20.8 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

20.9 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).

20.10 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).

20.11 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).

20.12 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).

20.13 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).

20.14 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).

20.15 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).

20.16 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).

20.17 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.

20.18 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).

20.19 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

20.20 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).

20.21 Ethernet filter could receive other packets (non-assigned) in Niantic

Description: On Intel® Ethernet Controller 82599EB When Ethernet 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 ethernet 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 Ethernet 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).

20.22 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).

20.23 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 black-listing 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.

20.24 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.

20.25 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).

20.26 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).

20.27 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).

20.28 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).

20.29 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).

20.30 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).

20.31 igb_uio can not be used when running l3fwd-power

Description: Link Status Change(LSC) interrupt and packet receiving interrupt are all enabled in l3fwd-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 l3fwd-power APP, link status getting doesn't work properly.

Resolution/Workaround: Use `vfio-pci` when LSC and packet receiving interrupt enabled.

Affected Environment/Platform: ALL.

Driver/Module: `igb_uio` module.

20.32 Linux kernel 4.10.0 iommu attribute read error

Description: When VT-d is enabled (`iommu=pt intel_iommu=on`), reading IOMMU attributes from `/sys/devices/virtual/iommu/dmarXXX/intel-iommu/cap` on Linux kernel 4.10.0 error. This bug is fixed in [Linux commit a7fdb6e648fb](#), This bug is introduced in [Linux commit 39ab9555c241](#),

Implication: When binding devices to VFIO and attempting to run `testpmd` application, `testpmd` (and other DPDK applications) will not initialize.

Resolution/Workaround: Use other linux kernel version. It only happens in linux kernel 4.10.0.

Affected Environment/Platform: ALL OS of linux kernel 4.10.0.

Driver/Module: `vfio-pci` module.

20.33 Netvsc driver and application restart

Description: The Linux kernel `uio_hv_generic` driver does not completely shutdown and clean up resources properly if application using Netvsc PMD exits.

Implication: When application using Netvsc PMD is restarted it can not complete initialization handshake sequence with the host.

Resolution/Workaround: Either reboot the guest or remove and reinsert the `uio_hv_generic` module.

Affected Environment/Platform: Linux Hyper-V.

Driver/Module: `uio_hv_generic` module.

20.34 PHY link up fails when rebinding i40e NICs to kernel driver

Description: Some kernel drivers are not able to handle the link status correctly after DPDK application sets the PHY to link down.

Implication: The link status can't be set to "up" after the NIC is rebound to the kernel driver. Before a DPDK application quits it will invoke the function `i40e_dev_stop()` which will sets the PHY to link down. Some kernel drivers may not be able to handle the link status correctly after it retakes control of the device. This is a known PHY link configuration issue in the `i40e` kernel driver. The fix has been addressed in the 2.7.4 rc version. So if the `i40e` kernel driver is < 2.7.4 and doesn't have the fix backported it will encounter this issue.

Resolution/Workaround: First try to remove and reinsert the `i40e` kernel driver. If that fails reboot the system.

Affected Environment/Platform: All.

Driver/Module: Poll Mode Driver (PMD).

20.35 Restricted vdev ethdev operations supported in secondary process

Description In current virtual device sharing model, Ethernet device data structure will be shared between primary and secondary process. Only those Ethernet device operations which based on it are workable in secondary process.

Implication Some Ethernet device operations like device start/stop will be failed on virtual device in secondary process.

Affected Environment/Platform: ALL.

Driver/Module: Virtual Device Poll Mode Driver (PMD).

20.36 Kernel crash when hot-unplug igb_uio device while DPDK application is running

Description: When device has been bound to `igb_uio` driver and application is running, hot-unplugging the device may cause kernel crash.

Reason: When device is hot-unplugged, `igb_uio` driver will be removed which will destroy UIO resources. Later trying to access any UIO resource will cause kernel crash.

Resolution/Workaround: If using DPDK for PCI HW hot-unplug, prefer to bind device with VFIO instead of IGB_UIO.

Affected Environment/Platform: ALL.

Driver/Module: `igb_uio` module.

20.37 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 accrued from the GCC community defect.

Reason: Generated AVX-512 code cause crash: https://bugs.dpdk.org/show_bug.cgi?id=97 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.

20.38 Unsuitable IOVA mode may be picked as the default

Description Not all kernel drivers and not all devices support all IOVA modes. EAL will attempt to pick a reasonable default based on a number of factors, but there may be cases where the default may be unsuitable (for example, hotplugging devices using `igb_uio` driver while having picked IOVA as VA mode on EAL initialization).

Implication Some devices (hotplugged or otherwise) may not work due to incompatible IOVA mode being automatically picked by EAL.

Resolution/Workaround: It is possible to force EAL to pick a particular IOVA mode by using the `-iova-mode` command-line parameter. If conflicting requirements are present (such as one device requiring IOVA as PA and one requiring IOVA as VA mode), there is no workaround.

Affected Environment/Platform: Linux.

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.

21.1 Deprecation Notices

- **meson:** The minimum supported version of meson for configuring and building DPDK will be increased to v0.47.1 (from 0.41) from DPDK 19.05 onwards. For those users with a version earlier than 0.47.1, an updated copy of meson can be got using the `pip`, or `pip3`, tool for downloading python packages.
- **kvargs:** The function `rte_kvargs_process` will get a new parameter for returning key match count. It will ease handling of no-match case.
- **eal:** The function `rte_eal_remote_launch` will return new error codes after read or write error on the pipe, instead of calling `rte_panic`.
- **eal:** both declaring and identifying devices will be streamlined in v18.11. New functions will appear to query a specific port from buses, classes of device and device drivers. Device declaration will be made coherent with the new scheme of device identification. As such, `rte_devargs` device representation will change.
 - The enum `rte_devtype` was used to identify a bus and will disappear.
 - Functions previously deprecated will change or disappear:
 - * `rte_eal_devargs_type_count`
- **eal:** The `rte_logs` struct and global symbol will be made private to remove it from the externally visible ABI and allow it to be updated in the future.
- **dpaa2:** removal of `rte_dpaa2_memsegs` structure which has been replaced by a pa-va search library. This structure was earlier being used for holding memory segments used by dpaa2 driver for faster pa->va translation. This structure would be made internal (or removed if all dependencies are cleared) in future releases.
- **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`.
- **ethdev:** Update API functions returning `void` to return `int` with negative `errno` values to indicate various error conditions (e.g. invalid port ID, unsupported operation, failed operation):

- `rte_eth_dev_stop`
- `rte_eth_dev_close`
- **ethdev:** Will add `RTE_ETH_` prefix to all ethdev macros/enums in v21.11. Macros will be added for backward compatibility. Backward compatibility macros will be removed on v22.11. A few old backward compatibility macros from 2013 that does not have proper prefix will be removed on v21.11.
- **ethdev:** New offload flags `DEV_RX_OFFLOAD_FLOW_MARK` will be added in 19.11. This will allow application to enable or disable PMDs from updating `rte_mbuf::hash::fdir`. This scheme will allow PMDs to avoid writes to `rte_mbuf` fields on Rx and thereby improve Rx performance if application wishes do so. In 19.11 PMDs will still update the field even when the offload is not enabled.
- **sched:** To allow more traffic classes, flexible mapping of pipe queues to traffic classes, and subport level configuration of pipes and queues changes will be made to macros, data structures and API functions defined in “`rte_sched.h`”. These changes are aligned to improvements suggested in the RFC <https://mails.dpdk.org/archives/dev/2018-November/120035.html>.
- **metrics:** The function `rte_metrics_init` will have a non-void return in order to notify errors instead of calling `rte_exit`.