

## **This Quarter's Highlights Include:**

[DPDK 17.02 Released](#)  
[DPDK-in-a-Box Available from Netgate](#)  
[Long Term Support](#)  
[FD.io 17.01 Released](#)  
[Moving DPDK to The Linux Foundation](#)  
[New DPDK Videos Available](#)

## **DPDK 17.02 Released**

DPDK 17.02 was successfully released on Tuesday February 14th and is now available for [download from dpdk.org](#). The release contains a number of significant enhancements, including:

### **Cryptodev Enhancements:**

- Cryptodev is a generic, multi-vendor API supporting hardware and software crypto via DPDK. In 16.11 we completed all the functionality required to support our IPsec efforts. In 17.02 we've added further performance optimizations and initial support for crypto requirements for the cable segment (which will be completed in 17.05). Details of the 17.02 enhancements include:
- AESNI\_GCM PMD Enhancements: The AESNI\_GCM PMD was migrated to a new software library which adds support for the following: GMAC algorithm; 256-bit cipher keys; session-less mode; out-of place processing; scatter-gather support for chained mbufs (only out-of place and destination mbuf must be contiguous).
- QAT PMD Enhancements: The QAT PMD was enhanced to support: scatter-gather lists (SGL); support for DES; cipher-only support for AES CBC and DES CBC.
- Support Single Operations in AESNI\_MB PMD: Support was added for cipher-only and hash-only operations in the AESNI\_MB PMD.
- Round-Robin Scheduler: A cryptodev scheduler PMD was created that acts as a software crypto PMD. The PMD supports attaching multiple

## **Open Source Project**

[Open source website](#)  
[Source code](#)  
[Documentation](#)  
[Mailing lists](#)  
[Public Roadmap](#)

## **Intel® Network Builders University (login req'd):**

[DPDK 101: Introduction to DPDK](#)

[DPDK 201: New Features Deep Dive](#)

[DPDK 2.2 New Features](#)

[Accelerating Your Cloud & Enterprise with DPDK](#)

[DPDK Setup and Configuration](#)

## **Other DPDK-Related Webinars (login req'd):**

[DPDK 16.04 New Features](#)

[DPDK 16.07 New Features](#)

[Enabling the Storage Transformation with SPDK](#)

[Open vSwitch with DPDK in](#)

crypto PMDs, software or hardware, as slaves, and distributes the crypto workload to them according to the specified mode. Currently the only supported mode is Round Robin, but other modes will be added in future.

- Performance Test App: A new performance test application allows measuring performance parameters of PMDs available in crypto tree.
- ARMv8 PMD which implements ARMv8 cryptographic extensions to process chained crypto operations in an optimized way.

#### Generic APIs:

- Generic Flow API: This API provides a generic means to configure hardware to match specific ingress or egress traffic, alter its fate and query related counters according to any number of user-defined rules. This is a generic API that will eventually be supported by all vendors with DPDK-enabled NICs, and is another step in establishing DPDK as the de facto data plane API for NFV.

#### Virtualization Enhancements:

- Interrupt Mode Support in Virtio PMD: Support for interrupt mode was added to the virtio PMD, which can be used to enable power savings for VNFs during periods with low throughput. This is only supported with VFIO no-IOMMU mode.
- Virtio-User as an Alternative Exception Path: Virtio-user, with vhost-user as the backend, was added in 16.07 as a way to improve IPC (Inter-Process Communication) and user space container networking. Now, virtio-user with vhost-kernel as the backend provides an alternative (to KNI) exception path which exchanges packets with the kernel networking stack.
- DPDK vHost Performance Optimization: Improved memory access efficiency for the common path which has mergeable Rx buffer support.

#### Miscellaneous:

- Support MACsec Offload for IXGBE: MACsec offload support was added for IXGBE. Ethdev APIs are available to enable/disable MACsec offload, configure Rx and Tx SCs, and enable Rx and Tx SAs.
- VF Daemon (SRIOV Policy Manager) for I40E: The VF Daemon (SRIOV Policy Manager) functionality that was added for IXGBE in 16.11 was extended to include I40E. This feature allows the PF to make policy decisions which determine which configuration changes a VF can/cannot make, and builds on AT&T's work in this area (see <https://github.com/att/vfd>). This feature is marked as Experimental because we need to determine whether in the longer term it belongs in DPDK or in the kernel.

The full feature list is available in the [Release Notes](#). A [DPDK 16.11 and 17.02 New Features webinar](#) has been scheduled for February 22nd which will give an overview of the new features included in the last two releases.

### DPDK-in-a-Box Available from Netgate

DPDK-in-a-Box started out as an initiative to develop a low-cost, DPDK-capable platform that could be used by universities, start-ups and other DPDK developers with limited budget. Following the level of interest shown in

[OVS 2.4.0](#)

[Open vSwitch with DPDK in OVS 2.5.0](#)

[Open vSwitch with DPDK in OVS 2.6.0](#)

[Accelerating Your Cloud with DPDK](#)

[Intel Software Defined Infrastructure: Tips, Tricks and Tools for Network Design and Optimization](#)

[IP Flow Analytics Enabled by Saisei and DPDK](#)

[Ubuntu 16.04 + Intel: Expanding the Possibilities of Data Center Networking](#)

#### Intel® Developer Zone Articles:

[Data Plane Development Kit \(DPDK\): Getting Started](#)

[Using Open vSwitch with DPDK on Ubuntu](#)

[QoS Configuration and usage for Open vSwitch with DPDK](#)

[vHost User Multiqueue for Open vSwitch with DPDK](#)

[Profiling DPDK Code with Intel® VTune™ Amplifier](#)

[DPDK Packet Capture Framework](#)

#### Meet-Ups and Training:

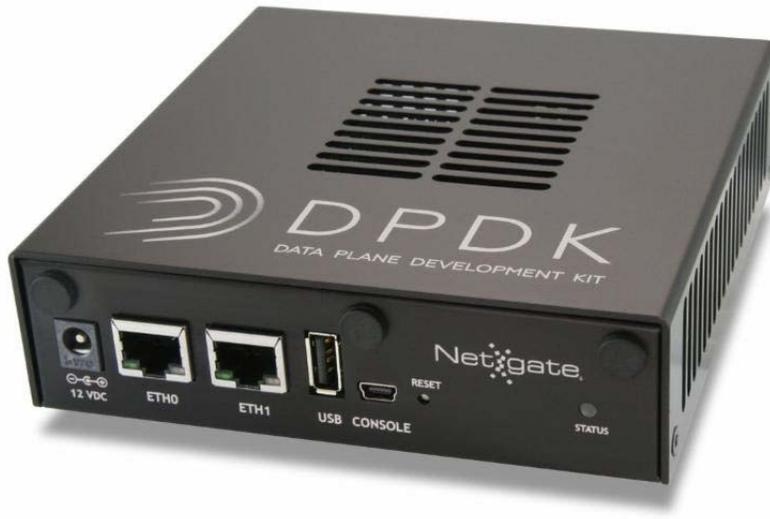
[Out of the Box Network Developers Meet-Up](#)

[Intel® Developer Zone NFV/DPDK Dev Lab](#), March 1st/2nd, Santa Clara

[NTT and Intel on Routing and Switching in Core and Edge in the SDN world](#), March 9th, Santa Clara

#### 2017 Events

the initial prototypes, the platform is now available from Netgate.



Dates and locations for DPDK Summits in 2017 have not yet been finalized, but we're currently planning the following:

- Bangalore in ~April
- PRC (Shanghai or Beijing or Shenzhen) in ~May
- San Jose in ~September
- Europe (possibly Dublin again) in ~November

Further details will be communicated in future newsletters and via the [DPDK Summit](#) website.

The DPDK-in-a-Box features a 2 core Intel® Atom™ C2338 CPU, 2GB memory, 2 GbE Ethernet Ports, nano-ITX form factor, and comes ready to use with DPDK pre-loaded. It's available [here](#).

## Long Term Support

We've recently added long-term support to DPDK. DPDK releases that are designated as LTS will be maintained for 2 years. During this interval, bug fixes will be backported from the latest releases to the LTS release. This means that users who base their applications on the LTS release will benefit from having a stable DPDK platform with no ABI changes for a 2 year period, but will still be able to avail of the latest bug fixes.

Our first LTS release is 16.11. At the moment, the plan is to maintain only one LTS release at a time, so the next one is planned to be 18.11.

## FD.io Release 17.01



The third open source release the [Fast Data \(FD.io\)](#) project was completed in January. FD.io builds on DPDK and supports the creation of high performance, flexible, and scalable software packet processing solutions. The 17.01 release includes updates to two FD.io projects:

### [VPP \(Vector Packet Processing\):](#)

- Integrated the DPDK 16.11 release and added support for DPDK's:
  - CryptODEV API. This allows VPP to use hardware and software accelerated crypto via DPDK.
  - Hierarchical QoS. This enhancement was required to support the vPE use case.
- Complete rework of Forwarding Information Base (FIB) and various performance improvements.
- New plugins for ACL and flow per packet.
- Further details, including the full feature list, are available in the [Release Notes](#).

### [NSH SFC \(Network Service Header Service Function Chaining\):](#)

- NSH Classifier support.
- Enable NSH Proxy feature to support integrating NSH-unaware Service Functions into a Service Function Chain.
- Further details are available in the jira entries ([NSH Classifier](#), [NSH Proxy](#)).

### **Moving DPDK to the Linux Foundation**

Work on moving DPDK to The Linux Foundation is continuing. Migrating an established open source project like DPDK is a complex task because it means gaining agreement from the community on all the proposed changes. Good progress is being made, with the new Linux Foundation project expected to officially launch in Q2. The major work items include:

- A draft [Project Charter](#) has been created which defines the way the project will operate after the transition. This has been through multiple reviews and is now almost complete. This includes a two-tier membership structure with Gold and Silver members.
- The technical governance of the project has been updated and documented in the [Contributor's Guidelines](#).
- The composition of the DPDK Technical Board (commonly referred to as a TSC in many other projects) has been updated.
- A draft [budget estimate](#) has been created. This will need to be revised once the actual project budget is known.
- A draft proposal for an open [DPDK Reference Lab](#) has been created. Discussions on this are continuing.

Further discussion on this is taking place on the new [moving@dppk.org](mailto:moving@dppk.org) mailing list and during [weekly public meetings](#).

### **New Intel® Network Builders University Videos**

The DPDK module in the Intel® Network Builders University program has been enhanced with a number of new videos:

- [Testing VNF Performance Using Data Plane Performance Demonstrator](#)
- [Maximising NFV Performance on IA](#)
- [Setting up DPDK on Different Operating Systems](#)
- [DPDK Sample Applications](#)
- [Writing a Simple DPDK Forwarding Application](#)
- [DPDK Packet Framework](#)

### **DPDK at FOSDEM**

# FOSDEM'17



There was a strong DPDK presence at the FOSDEM event again this year. As part of the [SDN/NFV DevRoom](#), Ferruh Yigit presented on [Writing a Functional DPDK Application from Scratch](#), and Ray Kinsella presented on [Accelerating TCP with TLDK](#) which described the FD.io [Transport Layer Development Kit](#) project designed to support L4 protocols on top of DPDK.

## Intel 17.05 Roadmap

The Intel roadmap for DPDK 17.05 was published on the [dev@dptk.org](mailto:dev@dptk.org) mailing list. The key features targeted for the release include:

I40E Enhancements, including:

- I40E Hardware QoS: Hardware QoS will be supported on the I40E. This will include Tx bandwidth control (min and max), and Rx Traffic Class assignment.
- Configurable Tunnel Filters for I40E: DPDK support will be added for a new I40E admin queue which allows configuration of filter types for cloud filters.
- Enable MPLS on I40E: MPLSoUDP and MPLSoGRE will be supported for I40E, including the new protocols and filtering (RSS and Flow Director), checksum offload and packet type recognition.

Cryptodev Enhancements:

- DOCSIS BPI+: The cryptodev API will be enhanced to enable processing of packets according to the Baseline Privacy Interface Plus (BPI+) specification described in the Data-over-Cable Service Interface Specification (DOCSIS) Security Specification. See the RFC (<http://dptk.org/ml/archives/dev/2016-December/052433.html>) for further details. Support will be added to the QAT PMD for AES and DES, and to the existing AESNI\_MB PMD.
- Packet-Based Scheduler: This allows packets to be encrypted/decrypted in either SW or HW depending on packet size and HW utilization. Reordering will be supported so that packet order can be preserved.

Generic APIs:

- API to Configure Programmable Devices: More devices are now supporting programmable components, for example the Pipeline Personalization Profiles in I40E. An API will be added to allow any programmable device to be configured.
- Software Eventdev Implementation: The libeventdev API (<http://dptk.org/ml/archives/dev/2016-December/052877.html>) has been added to the dpdk-next-eventdev tree. A software implementation of this API will be added. This will be the first implementation of the eventdev API.
- Abstraction Layer for QoS: An abstraction layer for Quality of Service (QoS) hierarchical schedulers

will be implemented. The goal of the abstraction layer is to provide a simple generic API that is agnostic of the underlying HW, SW or mixed HW-SW implementation. See the RFC (<http://dpdk.org/ml/archives/dev/2016-November/050956.html>) for details.

#### Virtualization Enhancements:

- New vHost Device Type: The vhost-user framework will be expanded so that it can support additional device types. Support for SCSI will be added initially, but block devices and other device types may be added in future.
- Interrupt Mode for Virtio-User: Interrupt mode support will be added for virtio-user, which is a virtual device for high performance container networking added in 16.07.
- Automated VF Processing of PF Reset for I40E: In 16.07, changes were made for both IXGBE and I40E to notify a VF when a PF reset occurs. This will be further enhanced for I40E so that the driver handles as much of the processing as possible, including things like resetting VF ports.

#### Service Assurance:

- Extended Stats: Latency and Bit Rate Statistics: Enhance the Extended NIC Stats (Xstats) implementation to support the collection and reporting of latency and bit rate measurements. Latency statistics will include min, max and average latency, and jitter. Bit rate statistics will include peak and average bit rate aggregated over a user-defined time period. This will be implemented for IXGBE and I40E.

#### Performance Optimization:

- Packet Distributor Enhancements: Enhancements will be made to the Packet Distributor library to improve performance: 1. Introduce burst functionality to allow batches of packets to be sent to workers. 2. Improve the performance of the flow/core affinity through the use of vector instructions.

#### Miscellaneous:

- Ethernet 32-bit CRC Generation: An optimized x86 library for CRC-32 will be added. A CRC may need to be generated for a downstream frame by the DOCSIS MAC layer because the CRC may be removed by the NIC or the frame may be modified by the MAC for PHS (packet eader suppression).

Copyright © 2017 Intel Corporation.

[Legal Notices](#). Other names and brands may be claimed as the property of others.