

This Quarter's Highlights Include:

[DPDK 16.11 Released](#)
[FD.io 16.09 Released](#)
[Successful DPDK Summit Userspace Event](#)
[Moving DPDK to the Linux Foundation](#)
[Large Hadron Collider to be Powered by DPDK?](#)
[Ecosystem Updates](#)
[Successful DPDK Summit USA Event](#)

DPDK 16.11 Released

DPDK 16.11 was successfully released on Sunday 13th November, and is now available for [download from dpdk.org](#). The release contains a number of significant enhancements, including:

- **Cryptodev Support for Additional Algorithms:**
 - The scope of the cryptodev API has been increased with hardware and software support for additional crypto algorithms. This makes the cryptodev API applicable to a wider variety of customers and use cases.
 - Support has been added for hardware acceleration (via Intel® QuickAssist Technology) of the following algorithms:
 - Cipher algorithms: 3DES_CBC, 3DES_CTR, KASUMI (F8), NULL
 - Authentication algorithms: KASUMI (F9), NULL, SHA224/384_HMAC, MD5_HMAC, AES_GMAC
 - A software implementation of the ZUC algorithm was also added, which supports both EEA3 (cipher) and EIA3 (authentication).
 - In addition, an OpenSSL PMD was created which uses the openssl library to provide software support for multiple crypto algorithms:
 - Cipher algorithms: 3DES_CBC, 3DES_CTR, AES_CBC, AES_CTR, AES_GCM
 - Authentication algorithms: AES_GMAC, MD5, SHA1/224/256/384/512, MD5_HMAC, SHA1/224/256/384/512_HMAC
- **Cryptodev Performance Optimization:** The performance of the QAT PMD was optimized. This is expected to result in a performance increase of up to 13% for small packets compared to the 16.07 release.
- **IPsec Sample Application Enhancements:** The IPsec security Gateway sample application that was originally introduced in the 16.04 release has been enhanced to add support for the AES_GCM and AES_CTR algorithms. Support for a config file has also been added to remove the need to hard-code security policies and security associations in the application itself.
- **vHost-User Performance Optimization:**
 - **vHost-User Optimization:** The common path, which supports mergeable Rx buffers and is used for unmodified guests, was optimized to improve performance. The performance gain in vhost-user scalability is expected to be ~20-30%. Previous optimization efforts had focused on the alternative, simple path.

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 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](#)

- **vHost Multi-Queue Enhancement:** Allows dynamic queue number configuration.
- **Delay Packet Copy in Vhost-User Dequeue:** This is a performance enhancement that delays and avoid one packet copy from vhost-user dequeue to vhost-user enqueue or to NIC PMD Transmit Queue. In some scenarios the copy can then be avoided, resulting in a performance increase for large packet sizes.
- **Vhost PMD Xstats:** Over the last few releases, we've made lots of improvements to statistics in DPDK. In 16.11, the extended stats (Xstats) feature was implemented for the vhost PMD. The statistics include packet type (unicast, multicast, broadcast) and length counters.
- **VF Daemon (SR-IOV Policy Manager):** The VF Daemon (VFD) runs on the host and makes policy decisions which control the changes that a VF can and cannot make. In a multi-tenant environment, this allows the PF to implement policies that prevent a VM from making configuration changes that may adversely impact other VMs. This feature is currently implemented for IXGBE, but support will be extended to I40E in the 17.02 release.
- **Cuckoo Hash Enhancements:** The cuckoo hash, which was originally introduced in the 2.1 release, has been enhanced to improve performance. Lookup operations have been changed to use intelligent prefetching of keys and AVX instructions for vector processing to improve performance by up to 1.4x. For insert operations, support has been added for multiple concurrent writers so that performance scales linearly with the number of threads.

The full feature list is available in the [Release Notes](#). A DPDK 16.11 New Features webinar will be scheduled soon via [BrightTalk](#) which will give an overview of the new features included in the release.

FD.io Release 16.09



The second open source release the new [Fast Data \(FD.io\)](#) project was completed in September. FD.io builds on DPDK and supports the creation of high performance, flexible, and scalable software packet processing solutions. The 16.09 release includes updates to three FD.io projects:

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

- New in-tree plugins for IPv6 ILA (Identifier locator addressing), iOAM, and Secure Network Address Translation (SNAT).
- LISP functionality has been expanded to include Map-resolver failover, L2 overlays, multi-tenancy, multi-homing, and RTRs (Re-encapsulation tunnel routers).
- Other new features include: high-performance (line-rate) "neutron like" L4 port-filtering, dynamically ordered subgraphs, 64-bit vector lengths, huge shared-memory segments, and support for DPDK 16.07.

[Honeycomb Agent:](#)

- Provides netconf/restconf for the following yang models on top of VPP: ietf-interface, ietf-ip, vhost-user interface management, tap interface management, encap management for vxlan, vxlan-gpe, gre, vlan, and nsh, management of bridge domains, I2 and I3 acls via ietf-acl model, and configuration of LISP mapping server.

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

- Support for an NSH Service Function Forwarder (SFF).

[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](#)

The [Intel Developer Zone NFV/DPDK Devlab](#) is being held on Thursday 8th and Friday 9th December. It includes a DPDK deep-dive as well as sessions on FD.io, IPsec, Containers and ONOS.

2017 Events

We've now concluded our DPDK Summit event series for this year, having held successful events in Shanghai, San Jose and Dublin, which help to further expand the reach of DPDK, communicate the benefits to new users, and build a stronger open source community.

Planning for next year's events is starting now. In addition to repeating our PRC, USA and European events in 2017, we're also investigating the possibility of hosting our first DPDK Summit in Bangalore, and are considering a Summit in the APAC region, possibly in Tokyo.

- A Network Service Header proxy was added to support non-NSH aware VNFs.
- VXLAN-GPE/GRE encapsulation support was added.

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

Successful DPDK Summit Userspace Event



A very successful DPDK Summit Userspace event was held in The Clayton Hotel in Dublin on October 20th and 21st. Our Userspace event is an annual gathering of the major contributors to the DPDK open source community. It differs from our other DPDK Summits in that it's targeted specifically at software developers contributing to DPDK, and that we keep it small (<100 people) in order to encourage interactive discussion and debate.

This year, all the main contributors to DPDK were represented. In addition to a number of technical presentations we also held interactive discussions on the DPDK roadmap, mbuf structure, interworking with the Linux kernel, usability improvements and project governance. Presentations and videos from the event are available on the [DPDK Summit](#) website.

This was also the first DPDK event that was live-streamed on [Periscope](#). Following the success of this trial we'll be doing this again for future DPDK Summit events.


Moving DPDK to the Linux Foundation

During our DPDK Summit Userspace event in Dublin we held an open discussion on project governance. During that, we agreed to move DPDK into the Linux Foundation. This is a big step for the project, but it's important to emphasize that there's still a lot of work to be done. We have agreement in principle at a high level, but the DPDK community contains a very large and diverse number of stakeholders and reaching consensus on the details is going to take a lot more work.

A [summary of the discussion in Dublin](#) is available. Further discussion on this is taking place on the new moving@dpdk.org mailing list and during [weekly public meetings](#). Anybody who wishes to be involved in this process should register for the mailing list and participate in the meetings.

Large Hadron Collider to be Powered by DPDK?

In a recent [blog post](#) and [white paper](#), CERN described a prototype they've built which uses DPDK to accelerate processing of the huge volume of data (30 petabytes annually), generated by the Large Hadron Collider. It's just a prototype so far, but is being considered for the next upgrade:



"The bandwidth results obtained so far (i.e. 32Gb/s approximately) signify that DPDK can be an interesting candidate for the LHCb upgrade in 2020."

The [original image](#) was made available by [CERN](#) under a Creative Commons license (CC-BY-SA).

Ecosystem Updates

In a blog post titled "[How to Enable BreakingPoint VE Performance Acceleration Using DPDK](#)", Ixia describe the benefits that DPDK gives to their BreakingPoint Virtual Edition (VE):

"One key technology that is making its way into a wide variety of VNFs is the data plane development kit (DPDK)."

"Packet processing and throughput is significantly enhanced to the tune of up to 10X improvement compared with systems that have not used the DPDK. Software engineers can compile their VNFs with the DPDK code base to take advantage of the kit's libraries for executing code in the Linux user space"

In a blog titled "How [the open source Data Plane Development Kit enables high-performance Linux networking](#)", TechRepublic describe the limitations of Linux for networking and the benefits of DPDK:

"As a general purpose OS, Linux is limited in its effective network throughput, but Intel's latest open source project could help better enable impressive network performance from it."

"DPDK is an enabler for software, driving the replacement of dedicated hardware in the data center."

In a recent press release titled "[Creanord Supercharges RFC 6349 TCP Testing with DPDK. Launches 100G vProbe VNF Supporting Millions of TCP Connections](#)", Creanord describe how DPDK has helped to

"supercharge" their TCP testing solutions:

"Powerful TCP testing has become a must. You simply need to know how many users your network can sustain. RFC 6349 is a good starting point, but does not go nearly far enough. So we took it and supercharged it with the help of DPDK and run it on our vProbe - enabling automated pre-testing of SDN-driven connections with NFV middleboxes."

Successful DPDK Summit USA Event

A very successful DPDK Summit USA 2016 was held in The Tech Museum of Innovation in San Jose on August 10th and 11th. Over 280 people attended representing all the major DPDK contributors and users in the Bay Area. Speakers included representatives from Intel, Oracle, Silicom, Microsoft, University of Massachusetts Lowell, Red Hat, Stanford University, UC Berkeley, University of British Columbia, NXP, Cavium and others. Topics that were covered included the DPDK cryptodev API, accelerating SSL and OVS via DPDK, userspace networking in containers, various vswitch presentations (OVS, VPP, PISCES, BESS, Lagopus), use of DPDK for storage, and DPDK-accelerated TCP/IP stacks.

Presentations and videos from the event are available on the [DPDK Summit](#) website.

Copyright © 2016 Intel Corporation.

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

Intel Shannon Ltd, Dromore House, East Park, Shannon, Co. Clare V14 AN23 Ireland

[SafeUnsubscribe™ {recipient's email}](#)

[Forward this email](#) | [Update Profile](#) | [About our service provider](#)

Sent by dpdk.summit@intel.com in collaboration with



Try it free today