

This Quarter's Highlights Include:

[DPDK 17.08 Released](#)
[DPDK Summit China 2017](#)
[DPDK Summit Userspace & DPDK Summit USA](#)
[DPDK & FD.io at OPNFV Summit](#)
[Terabit Software Router](#)

DPDK 17.08 Released

DPDK 17.08 was successfully released on Tuesday 10th August, and is now available for [download from dpdk.org](#). The release contains a number of significant enhancements, including:

- New PMDs:
 - Fail-Safe PMD: Added the new Fail-Safe PMD. This virtual device allows applications to support seamless hotplug of devices.
- PMD Enhancements:
 - BNXT: Support MTU modification, LRO, VLAN filter and strip, support for more dev_ops, added PMD-specific APIs to control VF from PF, updated HWRM version to 1.7.7.
 - CXGBE: The CXGBE PMD was updated to run Chelsio T6 family of adapters. The Tx and Rx path were reworked to improve performance and reduce latency for slow traffic.
 - DPAA2: Added support for MAC Filter configuration, segmented buffers, VLAN filter/strip, support for more dev_ops, optimized the packet receive path.
 - I40E: Added support for the latest firmware version (6), enhanced support for Dynamic Device Personalization (DDP) profiles, support for Adaptive Virtual Function (AVF) to enable an existing VF driver to run on future devices.
 - MLX5: Added vectorized Rx/Tx burst for x86, support for isolated mode from flow API, reworked the flow drop action to implement in hardware classifier, improved Rx interrupts management.
 - SZEDATA2: Added support for firmware with multiple Ethernet ports per physical port.
- Cryptodev Enhancements:
 - Multi-Core Packet Scheduler: The cryptodev scheduler was enhanced to use the packet distributor library to balance software crypto processing workload across multiple cores. This allows more efficient use of resources for software-accelerated crypto processing.
- Generic APIs:
 - Generic Traffic Management (QoS Scheduling) API: A generic API (rte_tm) has been created as part of ethdev for QoS scheduling. This currently includes support for I40E and IXGBE. Support for the existing QoS scheduling software implementation (librte_sched) will be added in 17.11.
 - Added support for the Generic Flow API (rte_flow) to the IGB and ENIC PMDs.
- Virtualization Enhancements:
 - vHost SCSI: The vhost-user protocol is common to many device types, including virtio_net, virtio_scsi and virtio_blk. This change introduces a simple memory-based block device that can be presented to a guest VM through the vhost-user-scsi-pci controller.

Open Source Project

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

2017 Events

The DPDK Summit Userspace event will be held in Dublin on Sept. 26th & 27th. Further details and a link for registration are available on the [DPDK Summit](#) website.

The DPDK Summit USA will be held in San Jose on November 14th and 15th. Further details, including registration and CFP, will be announced shortly on the [Linux Foundation Events](#) page.

Meet-Ups

Out of the Box Network Developers Meet-Ups are now taking place in the following locations:

- [Santa Clara](#) (live streamed and recordings available afterwards)
- [Portland](#)
- [Dublin](#)
- [Phoenix](#)

[Boston Software Networking](#)

[Intel® Developer Zone SDN/NFV Dev Lab](#), Tuesday September 19th, Santa Clara

Videos

[DPDK 101: Introduction to DPDK](#)
[DPDK 201: New Features Deep Dive](#)
[DPDK 16.04 New Features](#)
[DPDK 2.2 New Features](#)

This provides KVM guests with a fast, virtio-based connection to SCSI LUNs (Logical Unit Numbers).

- Misc:
 - GRO (Generic Receive Offload): Generic Receive Offload is a widely used SW-based offloading technique to reduce per-packet processing overhead. It improves performance by reassembling small packets into large ones. A new library (librte_gro) was added to DPDK which implements GRO.
 - Service Cores: This change provides the ability to define a set of services (components of DPDK that require CPU cycles to operate) and a set of service cores (DPDK lcores tasked with running services). Each service core then loops through the services that are enabled for that core, and invokes the function to run each service in turn. The benefit of this approach is that the mapping between service cores and services can be configured to abstract away the difference between platforms and environments.
 - 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.

The release contained 1023 patches from 125 different authors. The full feature list is available in the [Release Notes](#).

Thanks to everybody who contributed to making this another successful DPDK release!

DPDK Summit China 2017



The 2017 DPDK Summit China event took place in Shanghai on June 27th. There were over 260 attendees representing a broad range of networking companies including Huawei, ZTE, Nokia, China Mobile, Alibaba, Tencent, Ericsson and Cisco. Technical talks included presentations on DPDK in Containers, F-Stack (a DPDK-enabled userspace TCP/IP stack developed by Tencent), vHost Data Path Acceleration (vDPA), OVS-DPDK, Cryptodev, I40E enhancements and many more topics.

Slides from the event are available on the [DPDK Summit Past Events](#) page.

DPDK Summit Userspace & DPDK Summit USA



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

[DPDK Setup and Configuration](#)

[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](#)

[Testing DPDK performance and features with TestPMD](#)

[Building and Installing Vector Packet Processing \(VPP\) with Vagrant](#)

[DPDK-in-a-Box - The DPDK Starter Kit](#)

[DPDK 16.11 & 17.02 New Features](#)

[DPDK 16.04 New Features](#)

[DPDK 16.07 New Features](#)

[Enabling the Storage Transformation with SPDK](#)

[Building Blocks for Scalable, High Performance Storage](#)

[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](#)

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

[Intel® Multi-buffer Crypto for IPsec on DPDK - Get Started](#)

[DPDK Crypto - Get Started with Intel® QuickAssist Technology](#)

Blogs/Articles

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

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

[QoS Configuration and usage](#)

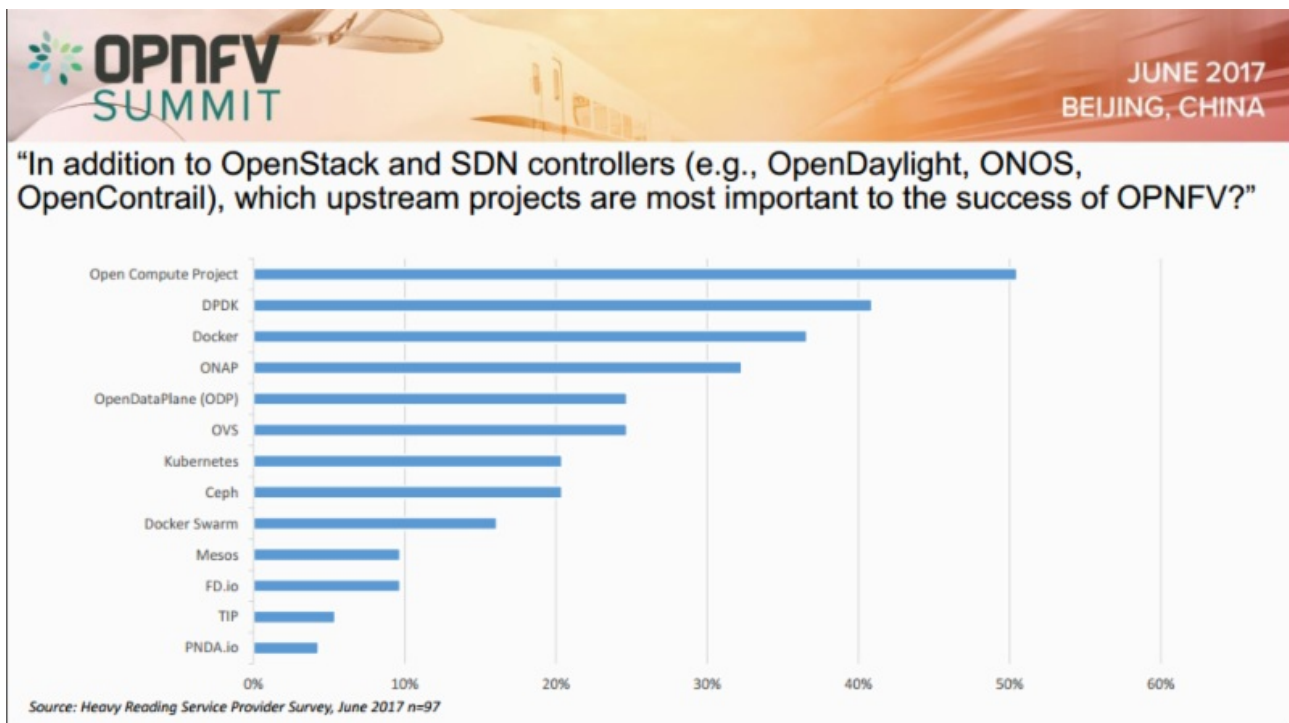
Plans for our USA and European DPDK Summits have been finalised. For Europe, our DPDK Summit Userspace event will take place at [The Clayton Hotel](#), Dublin on September 26th and 27th. Registration is open now via the [DPDK Summit](#) website. DPDK Summit USA will take place at [Club Auto Sport](#), San Jose on November 14th and 15th. Registration and CFP will be open soon via the [Linux Foundation Events](#) page.

[for Open vSwitch with DPDK](#)
[vHost User Multiqueue for Open vSwitch with DPDK](#)
[Profiling DPDK Code with Intel® VTune™ Amplifier](#)
[DPDK Packet Capture Framework](#)
[Build Your Own Packet Generator - DPDK-in-a-Box](#)

DPDK & FD.io at OPNFV Summit

The DPDK and FD.io projects held a joint mini summit at the [OPNFV Summit](#) in Beijing in June. Presentations included sessions on the DPDK roadmap, data plane virtualization enhancements, container networking, cryptodev and more. A youtube playlist containing [videos of all of the presentations at the OPNFV Summit](#) is available.

The importance of DPDK to NFV and OPNFV was shown in the results of a Heavy Reading survey on [Telco Perceptions of OPNFV](#):



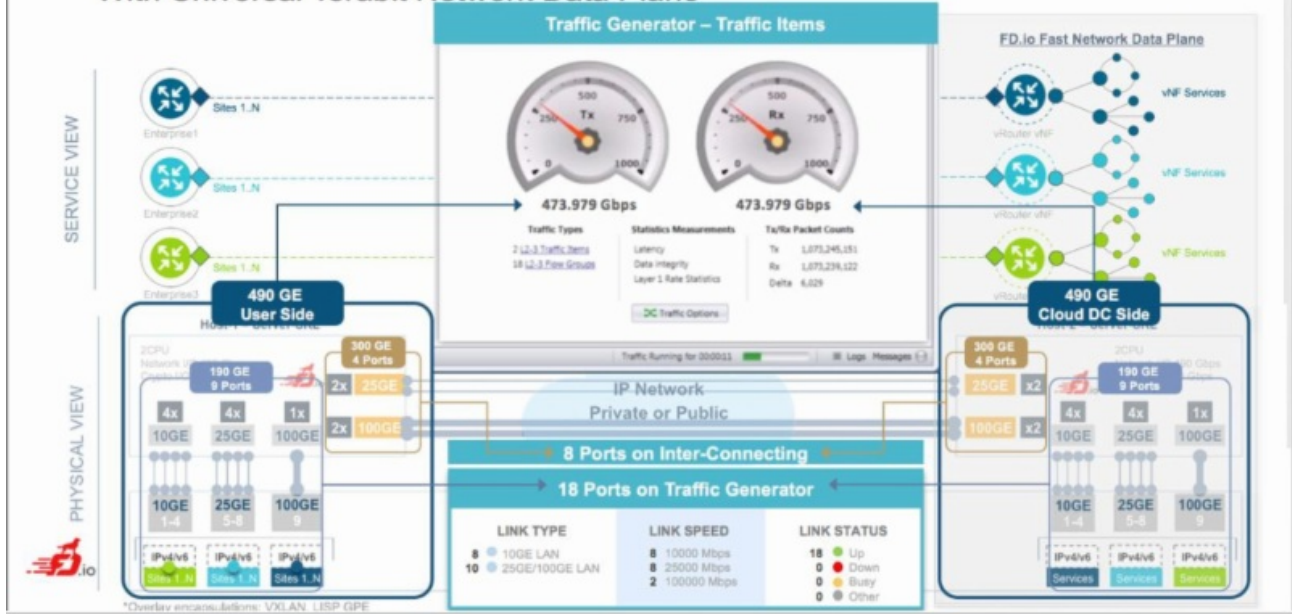
Terabit Software Router

In a blog, white paper and video, the FD.io project described how DPDK, FD.io and the Intel® Xeon® Processor Scalable family can be combined to build a software router capable of processing a terabit of data. The following extract from the white paper provides more details on the configuration and results:

"Recent testing of FD.io release 17.04 shows impressive gains in performance on Intel's newest platform when switching and routing layer 2/3 traffic. With the prior generation Intel® Xeon® Processor E7-8890v3, FD.io testing showed aggregate forwarding rate of 480 Gbps (200 Mpps) for 4-Socket machine (using 4 of E7-8890v3 CPU configuration); however, the same FD.io tests run on two 2-Socket blades (e.g. a modern 2RU server) with the new Intel® Xeon® Platinum 8168 CPUs (using four of 8168 CPUs in two by two-socket configuration), within the same power budget, show increase of forwarding rate to 948 Gbps (400 Mpps) benefiting from the PCIe bandwidth increase of the new CPUs, and the overall decrease in cycles-per-packet due to CPU micro-architecture improvements."

Fast Cloud Network Services

With Universal Terabit Network Data Plane



Copyright © 2017 Intel Corporation.

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