



## DPDK in 2015

2015 was another hugely successful year for DPDK, with many significant highlights:

- We've seen tremendous growth in the open source community, with both the number of patches and the number of contributors increasing from release to release.
- To help manage this growth we've added an Architecture Board (see below), and are working on other governance elements for the project.
- We've expanded our CPU and NIC support well beyond just x86 and Intel NICs. We now support multiple CPU architectures (x86, ia32, Power 8, TILE-Gx, ARMv7/v8) and multiple NIC vendors (Intel, Mellanox, Chelsio, Broadcom, Netronome etc.).
- We've created and implemented a new DPDK-AE architecture, which expands DPDK to include new device types (initially crypto hardware and software offloads).
- We've continued our global outreach and evangelism with DPDK Summit events in all three major geos (US, PRC, Europe) and also with a strong presence at key NFV/SDN shows.
- We completed 3 major software releases - 2.0, 2.1 and 2.2.



## DPDK 2.2 Released

The DPDK 2.2 release was successfully completed and can now be [downloaded from dpdk.org](#). Some of the key features that are included in the release are below. For the full list, see the 2.2 section of the [Release Notes](#).

- **Multi-Architecture Support:** Support for the ARMv7 and ARMv8 architectures, support for [Mellanox](#)

### Useful Links

- [Open source website](#)
- [Mailing lists](#)
- [Documentation](#)
- [Roadmap](#)
- [Latest stable release \(2.2\)](#)
- [DPDK Summit events](#)

### Upcoming Webinars

[DPDK 2.2 New Features](#) by  
Tim O'Driscoll  
7th Jan 2016, 17:00 GMT

### Previous Webinars

[DPDK 101: Introduction to Data Plane Development Kit](#) by Andrew Duignan  
[Data Plane Development Kit - Sample Applications and New Features Deep Dive](#) by M Jay

- [ConnectX-4 EN and Lx EN](#), [Liberouter COMBO-80G and COMBO-100G](#) and [Netronome NFP-6xxx NI Multi-Buffer PMDs](#)
- **DPDK-AE (Expansion of DPDK to Include New Device Types):** [Cryptodev API with QAT and AES-NI Multi-Buffer PMDs](#)
- **NIC Support:** FM10420 support, DCB for I40e and X550, Configuration Granularity of RSS for I40e, I40e 32-Nit GRE Keys, Enhanced TSO support for IGB and Fm10k
- **NFV/OPNFV Enabling Features:** [IEEE 1588 Ethdev API and sample application](#), [Keep Alive \(Fault Detection\)](#), Packet Framework (Edge Router use case), Extended Statistics, Vhost Multi-Queue support, Virtio enhancements, PCI Hotplug enhancements
- **Performance Improvements:** Vector PMD for I40e and Fm10k
- **Core DPDK Enhancements:** Interrupt Mode for I40e and E1000, [Performance Thread sample application](#), Keep Alive, Link Bonding Dynamic RSS Configuration
- **Usability Improvements:** [Userspace Ethtool sample application](#)

A [DPDK 2.2 New Features](#) webinar describing the new features included in the release will be held on 7th Jan 2016 at 17:00 GMT/09:00 PST/12:00 EST/18:00 CET.

## DPDK Architecture Board Established

The DPDK community has grown significantly since we decided to open source it in 2013. However, we have received some feedback expressing the concern that, as the project continues to grow, it needs a more formal governance structure and a clearer decision-making process. To that end, we've just put in place a DPDK Architecture Board to oversee the technical direction of the project and to be the final decision-making body on all technical issues.

Members of the board were chosen based on their technical expertise and history of contributions, and to represent a broad range of perspectives on DPDK:

- Bruce Richardson and Konstantin Ananyev (Intel)
- Thomas Monjalon and Olivier Matz (6WIND)
- Stephen Hemminger (Brocade)
- Panu Matilainen (Red Hat)
- Jerin Jacob (Cavium)



## New Release Numbering Scheme

We're changing the numbering scheme for DPDK releases. Until now, we've used a major.minor release number, with our latest example being 2.2. From now on, we'll adopt the approach used by Ubuntu, where releases are designated by the year and month in which they're completed. The first example of this will be 16.04 (formerly referred to as 2.3), which is planned for April next year. The advantage of the new approach is that it makes it much easier to track release dates - instead of having to remember when 2.3 is due, you'll know immediately when 16.04 will be released.

We're also moving from 3 major releases per year (2.0, 2.1 and 2.2 in 2015) to 4. This change is being implemented gradually so as to avoid impacting existing plans, and will be fully in place from 2017 onwards. Currently planned releases are:

- 2016: 16.04, 16.07, 16.11
- 2017: 17.02, 17.05, 17.08, 17.11

## Release 16.04 (formerly 2.3) Roadmap

Intel's planned contributions to the DPDK 16.04 release have been published. Some of the key features are included below. Further details are included in the [16.04 \(2.3\) Roadmap](#) thread on the dpdk.org mailing list.

- **Multi-Architecture Support:** External Mempool Manager
- **DPDK-AE (Expansion of DPDK to Include New Device Types):** Cryptodev Support for SNOW 3G

- **NIC Support:** Fm10k enhancements (FTAG-based forwarding, performance tuning), I40e enhancements (Flow Director enhancements, VEB Switching, Floating VEB, IPGRE support, set VF MAC address etc.)
- **NFV/OPNFV Enabling Features:** Packet Framework (Edge Router use case), Live Migration (KVM Vhost-User migration), Virtio/Vhost enhancements (virtio 1.0 support, Vhost software TSO, performance tuning), Container enhancements (Virtio for containers, Hugetlbfs mount point size, Cgroup resource awareness, Enable short-lived DPDK applications)
- **Performance Improvements:** Support Intel Resource Director Technology
- **Core DPDK Enhancements:** IPsec sample application, Ethdev enhancements, Increase Next Hops for LPM, Generic Tunneling API
- **Usability Improvements:** Tcpcdump support

Hopefully others will also provide details of enhancements that they plan to submit so that we can build up a complete picture of what will be in the release and make sure we avoid duplication/overlaps.



## Ecosystem Updates



In Red Hat's blog on "[Getting the Best of Both Worlds with Queue Splitting \(Bifurcated Driver\)](#)", they described how traffic can be split between DPDK and the Linux kernel using the Queue Splitting (Bifurcated Driver) approach:

*"What if there was a way to have ultra low latency and high throughput for some traffic, and full feature-set from Linux networking, all at the same time? This 'utopia' is now possible with Queue Splitting (Bifurcated Driver)!"*



In Sandvine's press release on "[Sandvine Virtual Series Achieves 1.1 Tbps of NFV Performance](#)", they describe how DPDK's "exceptional packet processing performance" helped them to achieve an unprecedented level of performance:

*"The 1.1 Tbps benchmark Sandvine has achieved is more bandwidth than any current single location needs. We believe this indicates that the race for NFV performance is over and that the industry can now focus on continuing to build out NFV functionality."*



**Hewlett Packard  
Enterprise**

HP's ConteXtream BU announced that they've achieved [1 Tbps of throughput with one billion OpenFlow rules on a six-node virtual switch cluster](#). The setup uses OpenStack and ContextSwitch, HP ConteXtream's OpenFlow software based-switching component, which uses DPDK to achieve high performance:

*"We've heard doubts from many CSPs time and again: 'Sure, NFV and SDN are all cool technologies and we understand the benefits they bring. But what about performance? Can they be used where we need both scale AND performance?' This is why it was important for us to demonstrate that SDN/NFV can be used in high performance demanding use cases."*



Based on DPDK, Silicom's [SPDK-1.1 Traffic Recorder](#) provides an efficient and scalable packet processing

infrastructure for monitoring and network visibility implementations at 10GbE wire speed and beyond:  
*"Integrating SPDK-1.1 with commercial off the shelf hardware enables quick and easy assembly of a capable packet recording device. DPDK kernel bypass approach proves once again the superiority of efficient polling mode, in order to reach line rate processing, at 10Gbps and beyond."*



DPDK is now packaged in Ubuntu 15.10, which is available from Canonical. Their post on "[What's new in Ubuntu 15.10: cloud and server?](#)" says:

*"High performance networking: New in 15.10 for telcos and enterprises with heavy networking requirements, DPDK (Data Plane Development Kit) previews. DPDK enables virtual network functions to deliver the high performance network throughput required in core network services."*



In their announcement on their new 7.2 release, "[Red Hat Drives Networking, Linux Container Innovation in Latest Version of Red Hat Enterprise Linux 7](#)", Red Hat stressed the benefits that DPDK gives for NFV applications:

*"Inclusion of the Data Plane Development Kit (DPDK), which makes it possible to rapidly develop low-latency and high throughput custom applications capable of direct packet processing in user space for NFV and other use cases."*



China Telecom released a [DPDK Technology White Paper](#). The published version is in Mandarin, but there is an internal English version. The paper describes DPDK usage for NFV:

*"In various kinds of NFV infrastructure performance optimization technical solutions, DPDK (Data Plane Development Kit) software acceleration solution has become a widely used basic method. ... the method has formed a complete performance acceleration overall architecture, with user-mode APIs that could be used by high speed forwarding applications."*



## Userspace 2015

Our first ever DPDK developers' conference, Userspace 2015, took place in Dublin on October 8th and 9th. This two day event was sponsored by Intel, Red Hat and Cisco, and was attended by all of the major contributors to the DPDK open source project.

The agenda included presentations on a wide range of topics, as well as a number of shorter lightning talks, and several interactive discussion sessions.

Videos and presentations from the event are accessible from the [Past Events page on the DPDK Summit website](#). Select "2015 Dublin DPDK Userspace" from the drop-down menu.



## DPDK in Containers

DPDK already runs in Linux containers - see Red Hat's blog post "[Can you run Intel's Data-plane Development Kit \(DPDK\) in a Docker container? Yep](#)". We're continually looking for ways to make it run faster though.

One promising approach is to use virtio within containers. We don't yet have any definite figures, but we're hoping to achieve a significant performance boost. We've provided [an early version of this](#) to the open source community, and it's generated a lot of interest.

We're also planning other container-related enhancements for our 16.04 release, which are outlined in the Release 16.04 Roadmap section above.



\* 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](mailto:dpdk.summit@intel.com) in collaboration with

**Constant Contact** 

Try it free today