



# DPDK Roadmap

Tim O'Driscoll

DPDK Summit USA - San Jose - 2016



# Agenda

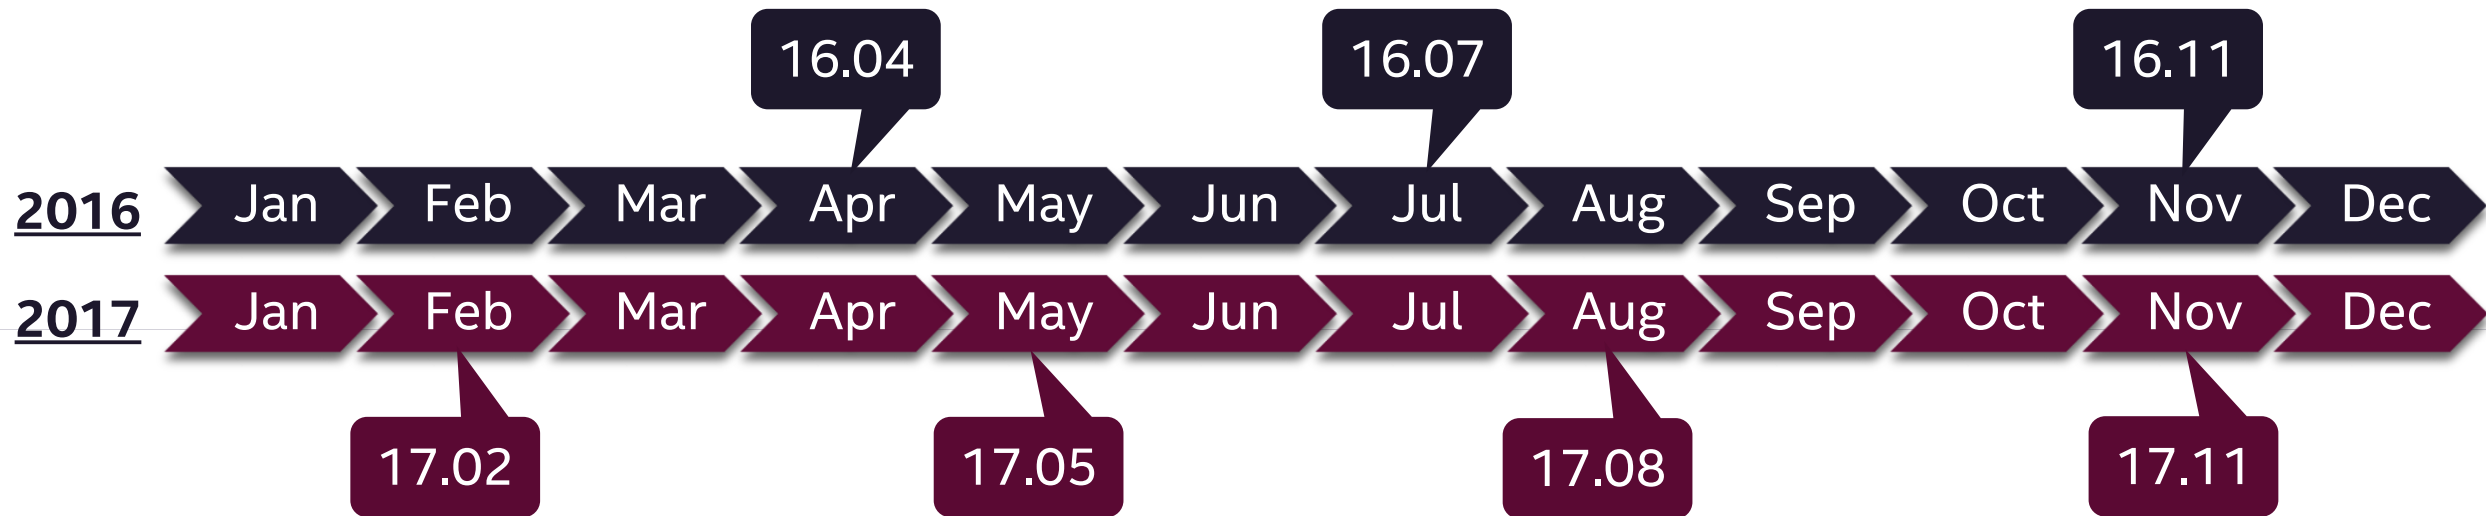


- ▶ 2016/2017 Releases
- ▶ Roadmap for 16.11 and 17.02
- ▶ Multi-Architecture Support
- ▶ Cryptodev
- ▶ Virtualization
- ▶ Community Survey Roadmap Feedback
- ▶ FD.io

# 2016 & 2017 Releases



- Since 16.04, releases use the Ubuntu\* numbering scheme of YY.MM.
- We're transitioning gradually from 3 major releases per year to 4.
- Frequency and dates of releases will be fixed from 2017 onwards.



# Release 16.11 Roadmap



Feature	Description	Feature	Description
Cryptodev: Additional Hardware Algorithms	Add support for additional algorithms which are supported by Intel® QuickAssist Technology but are not currently supported in DPDK. This includes: 3DES_CBC_128/192, KASUMI, NULL, SHA224_HMAC, SHA384_HMAC and AES-GMAC.	Cryptodev: Performance Optimisation	Analyze the performance of the cryptodev API, identify bottlenecks, and optimize where required.
Cryptodev: Additional Software Algorithms	Add support for software implementations of additional crypto algorithms. This includes: ZUC (EEA3 and EIA3) and 3DES_CBC_128/192 with MD5_HMAC, SHA1/SHA224/SHA256_HMAC and AES-GMAC. ZUC will be supported via an optimized software library similar to KASUMI and SNOW3G. 3DES, and potentially other software implementations in future, will not be optimized.	IPsec Sample App Enhancements	Add support for AES-GCM, AES-CTR, config file support to remove hard-coding of SAs/SPs etc., use forward cipher function to generate IV on CBC mode.
		Consistent Filter API	When a Generic Filtering API is agreed (see RFC at: <a href="http://dpdk.org/ml/archives/dev/2016-July/043365.html">http://dpdk.org/ml/archives/dev/2016-July/043365.html</a> ), implement support for that API for IXGBE and I40E.
		Cuckoo Hash Enhancements	Optimize the Cuckoo Hash lookup stages by using intelligent prefetching for keys and using IA AVX instructions for vector processing of keys.

**Note: Contents not yet finalized and subject to change without prior notice**

# Release 16.11 Roadmap



Feature	Description
Add vHost PMD xStats	Update the vHost PMD to support the extended statistics API.
Delay Packet Copy in vHost-User Dequeue	It may be possible to increase vhost-user performance by delaying the packet copy on Tx until a point where we know for certain whether the copy is required or not. This would avoid copying the packet in cases where it is not definitely required. Further investigation is required to determine how much of a performance gain can be achieved.
vHost-User Optimisation	Previously, most vhost-user optimization focused on the simple Tx path. In 16.11, we'll optimize the full Tx path. This will improve performance for unmodified guests.

## Version 16.11 (2016 November)

- Device Objects Refactoring
- Flattened Device Tree Access
- Generic Rx Filtering API
- Extend Hardware and Software Crypto algorithms support
- Optimize Crypto Performance
- IPsec sample app: AES-GCM, AES-CTR and configuration files
- vHost PMD: XStats
- vHost-User: Delayed packet Copy in dequeue
- Cuckoo Hash: Optimize lookup for x86

**Note: Contents not yet finalized and subject to change without prior notice**

# Release 17.02 Roadmap



Feature	Description	Feature	Description
Consistent Filter API phase 2	Extend support for the Consistent Filter API that will be implemented in 16.11 to IGB and FM10K.	Cryptodev: Additional Software Algorithms	<ul style="list-style-type: none"><li>• Optimize the AES-GCM software PMD.</li><li>• Enhance the Intel® AES-NI MB PMD to add support for cipher-only and authentication-only operations. Add support for AES_CBC_MAC, AES_CMAC_128, AES_GMAC_128.</li><li>• Add software support for: 3DES_ECB_128/192, AES_ECB_128/192/256, AES_F8, AES_XTS, ARC4.</li></ul>
Elastic Flow Distributor	The Elastic Flow Distributor (EFD) is a flow-based load balancing library which scales linearly for both lookup and insert with the number of threads or cores. EFD lookup uses a “perfect hashing” scheme where only the information needed to compute a key’s value (and not the key itself) is stored in the lookup table, thus reducing CPU cache storage requirements.	Run-Time Configuration of Flow Type (PCTYPE) and Packet Type (PTYPE) for I40E	At the moment all flow types and packet types in DPDK are statically defined. This makes impossible to add new values without first defining them statically and then recompiling DPDK. The ability to configure flow types and packet types at run time will be added for I40E.
GCC Support for Microarchitecture Names	As of GCC 4.9 the -march option supports code names for processors (e.g. -march=broadwell). Support for this option will be added in DPDK.		

**Note: Contents not yet finalized and subject to change without prior notice**

# Release 17.02 Roadmap



Feature	Description
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.
Hardware QoS for IXGBE and I40E	Implement support for Hardware QoS for IXGBE and I40E. This involves using DCB so that a PF or VF can receive packets for each of the Traffic Classes (TCs) based on the User Priority field (in the VLAN header). Bandwidth on transmit for each TC is programmable.

## Version 17.02 (2017 February)

- Hotplug Notification
- Automatic Device Binding
- Xen Netfront Driver

**Note: Contents not yet finalized and subject to change without prior notice**

# Multi-Architecture Support



CPU:  
Power8 (IBM)

NIC:  
ENIC (Cisco)

CPU:  
Tile-Gx (EZchip)

NIC:  
BNX2X (Broadcom)  
CXGBE (Chelsio)

NIC:  
ENA (Semihalf)

Better interworking with  
SoCs.

Consistent APIs across  
NICs (filtering etc.)

1.8

2.0

2.1

2.2

16.04

16.07

16.11 +

NIC:  
FM10K (Intel)  
MLX4 (Mellanox)

CPU:  
ARM (Cavium,  
RehiveTech)

NIC:  
NFP (Netronome)  
SZEDATA2 (Liberouter)  
MLX5 (Mellanox)

CPU:  
ARM (NXP)  
Mempool Handler (Intel)

NIC:  
QEDE (QLogic)  
ThunderX (Cavium)  
BNXT (Broadcom)



# Cryptodev Roadmap



- Cryptodev API introduced.
- PMD for Intel® QuickAssist Technology.
- PMD for Intel® AES-NI Multi-Buffer software lib.

2.2

16.04

16.07

16.11

17.02

Future?

- SNOW 3G software PMD support for bit-level operations.
- Software implementation of the KASUMI algorithm.

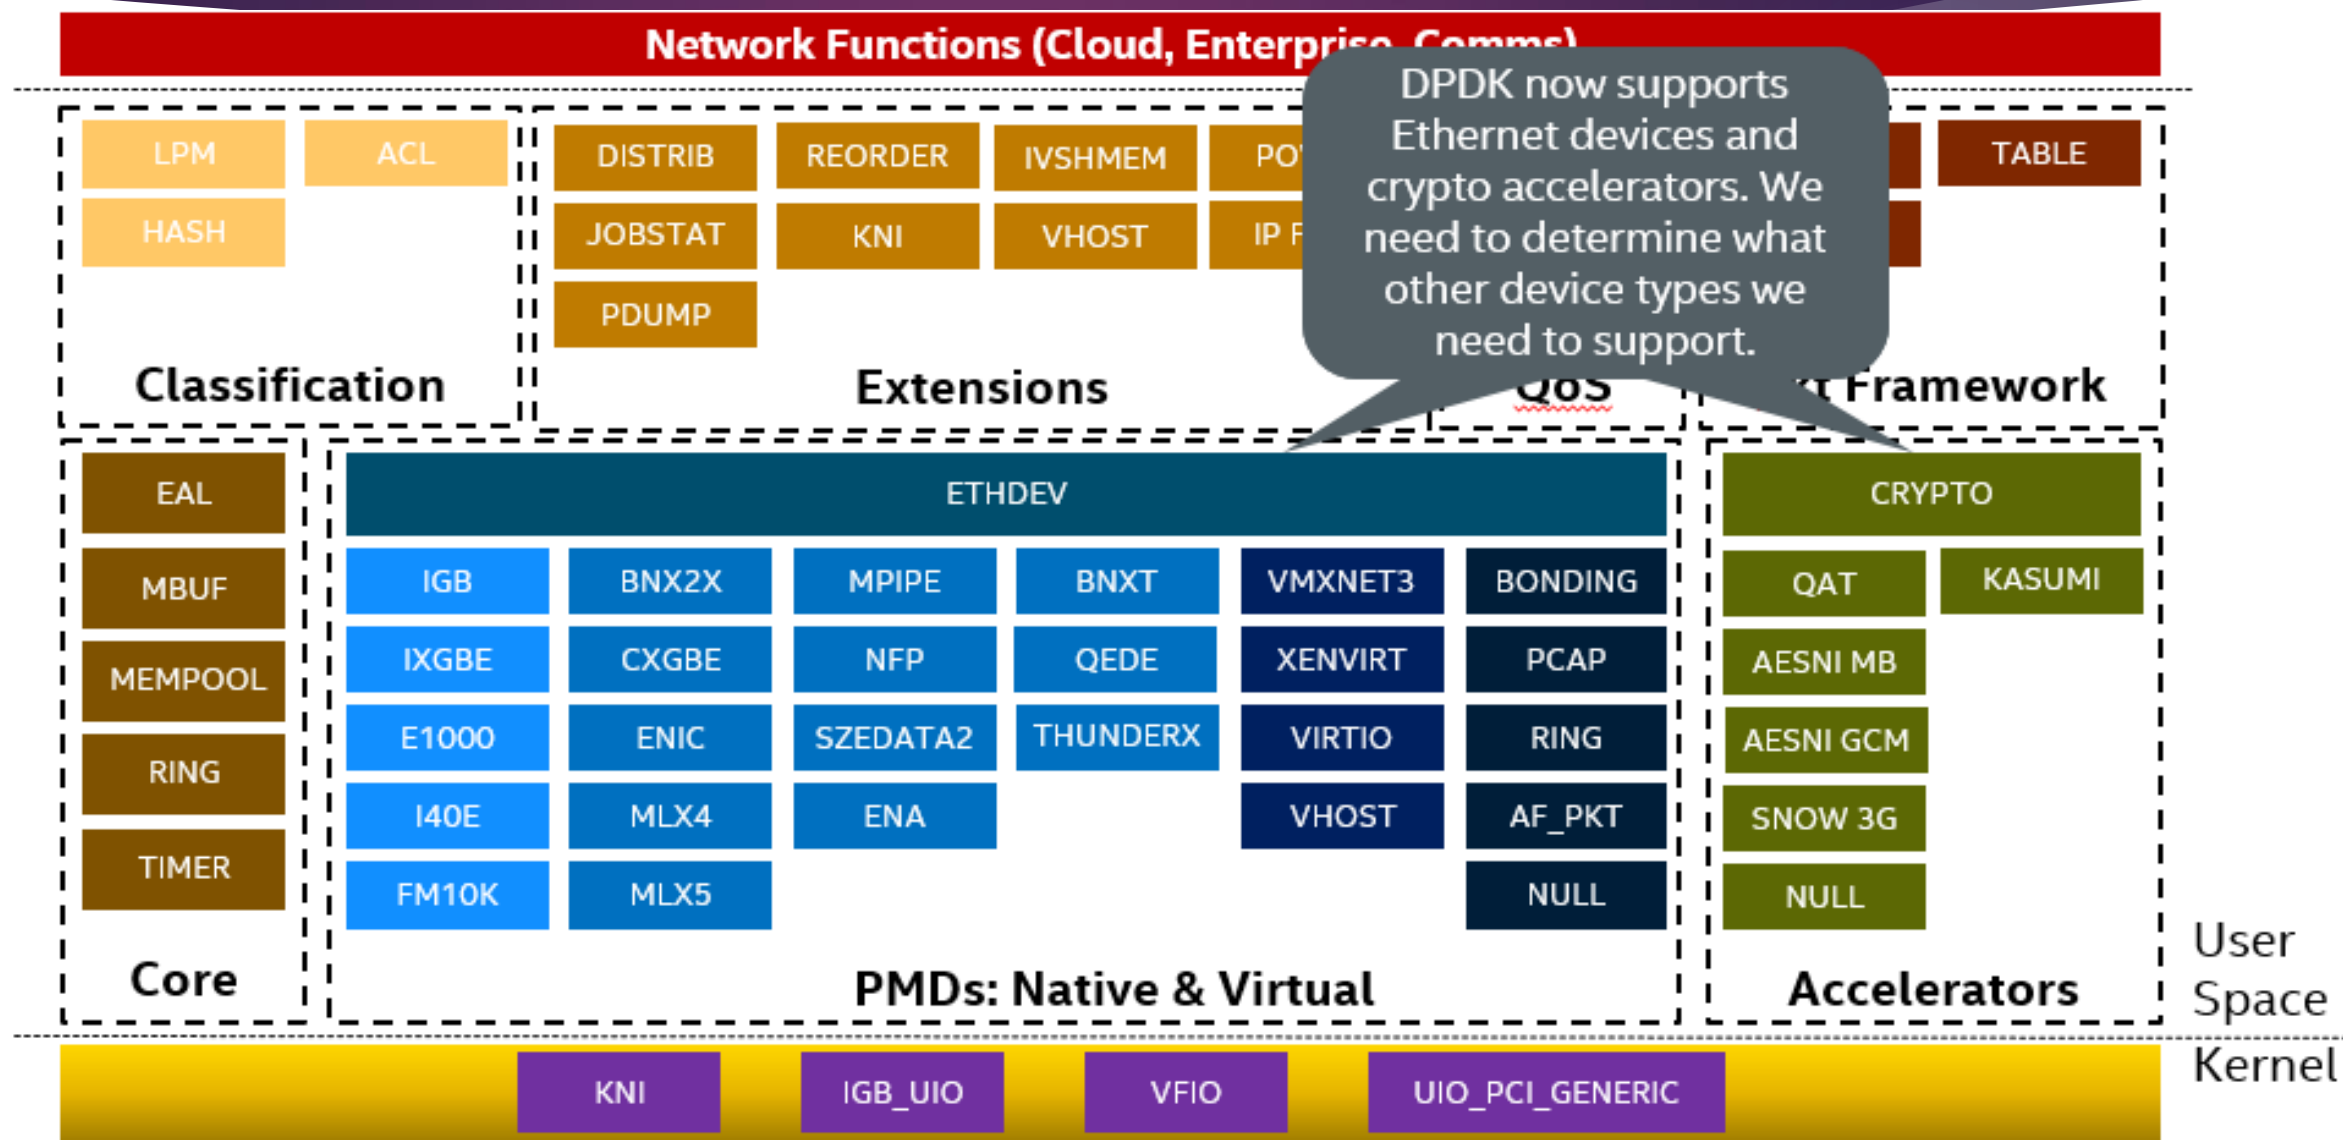
- Optimize AES-GCM SW PMD.
- Software support for: 3DES\_ECB, AES\_ECB, AES\_F8, AES\_XTS, ARC4.

- Added Intel® AES-NI GCM PMD
- Support for SNOW 3G algorithm.
- Discovery mechanism.
- NULL PMD.

- HW support for 3DES\_CBC & KASUMI.
- SW support for ZUC and 3DES\_CBC.
- Goal is to support all algorithms required for IPsec data plane.

- Support for non-Intel crypto accelerators?
- Cryptodev scheduler?
- Additional algorithm support.

# Future Device Types?



# Virtualization



- Vhost multi-queue support.
- Virtio optimization.

- Virtio optimization.
- Live migration for SR-IOV.
- QEMU vhost back-end reconnect.
- Notify VF of PF reset (IGB & IXGBE).

2.2

16.04

16.07

16.11

17.02

Future?

- Virtio 1.0 support.
- Vhost-user live migration.
- Notify VF of PF reset (I40E).

- Vhost-user optimization for unmodified guest.
- Extended Stats in vhost PMD.

- Vhost PCI Device – faster VM-VM communications.

# Community Survey Roadmap Feedback



## ► Inter-working with Kernel (KNI, KCP/KDP, bifurcated driver etc.)

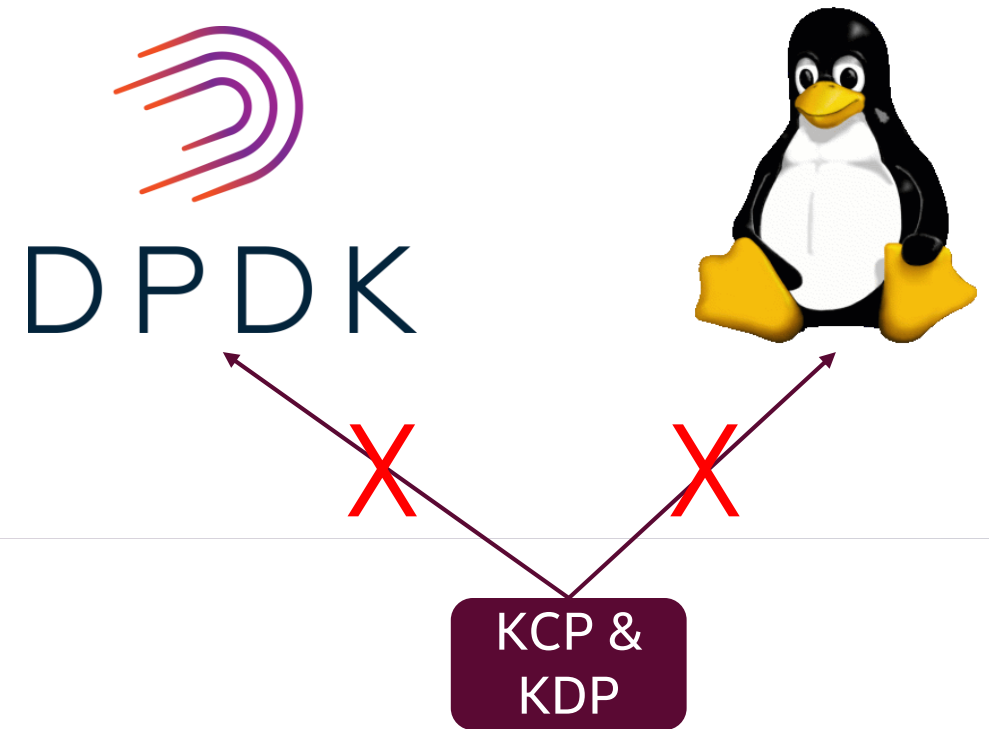
“Real Linux KCP Driver”

“KNI stability or a good KNI alternative”

“KCP/KDP as a replacement of KNI. We currently integrate those patches into the DPDK tree since they are of essential use to us.”

“Bifurcated driver”

“eBPF support”



# Community Survey Roadmap Feedback



## ► Better documentation:

“Improved documentation. Examples are very useful - I have no complaints there - but the major components of the framework could be described more fully.”

“Expecting very good documentation for all the services which is already provided by dpdk. Particularly Mbuf, mempool, rings init ,dealing with NUMA etc.”

“Documentation and examples. The mailing list seems to get similar questions over and over, so perhaps a little more detail would help the community.”

## ► Hardware offload abstraction:

“Better abstraction of offloading features. Example: tcp/ip checksum should be calculated by intermediate software layer if not supported by driver. The application should stay the same.”

“Ethdev api needs a cleanup to get rid of a lot of intel nics specifics.”

“better (unified) abstraction of PMDs. Application should ideally not have to be aware of the underlying driver type.”

# Community Survey Roadmap Feedback



## ▶ TCP/IP stack:

“production quality user-space TCP/IP stack that runs on top of DPDK”

“TCP/IP protocol stack”

## ▶ Usability:

“I fully support the current move to more of a library and less of a framework and would like to see that continued”

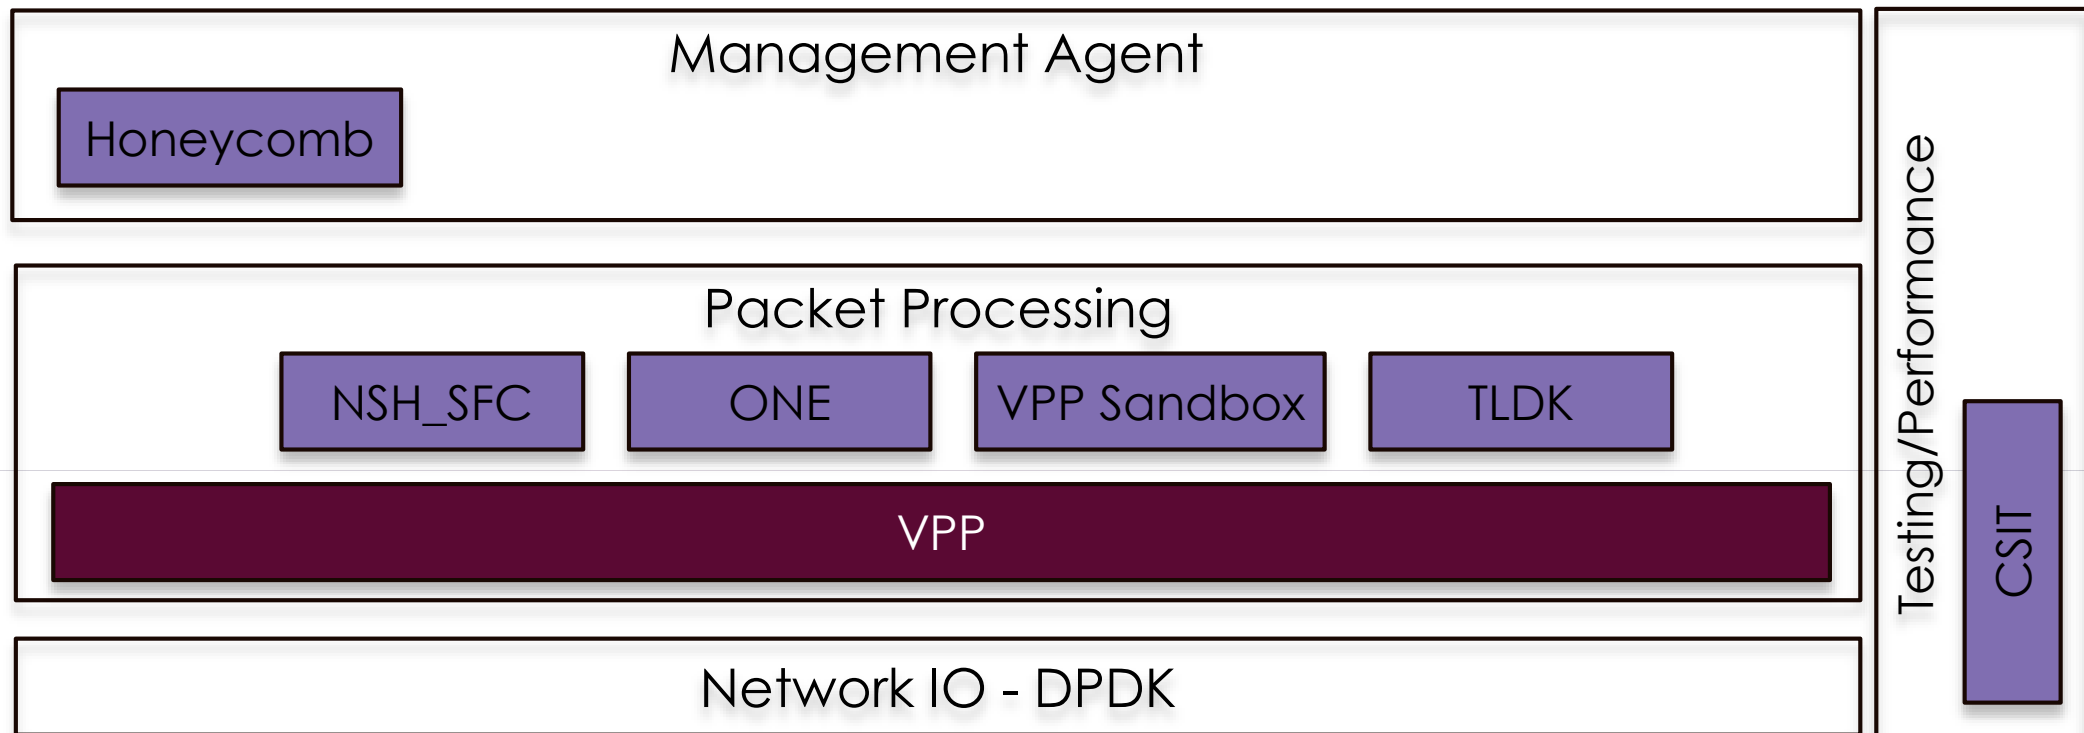
“A better way to use DPDK as an independent library, and not the complex extapp.mk system”

“ability to use PMDs without EAL”

## ▶ API/ABI stability, long-term support:

“greater compatibility from release to release; long-term support”

“stabilization of mbuf interface between releases (more access macros, etc.)”





1. Service Function Forwarder (SFF)
2. NSH Proxy for non NSH-aware VNFs

Honeycomb

Packet Processing

NSH\_SFC

ONE

VPP Sandbox

TLDK

VPP

Network IO - DPDK

Testing/Performance

CSIT





The TLDK project scope is:

- The project scope includes implementing a set of libraries for L4 protocol processing (UDP, TCP etc.) for both IPv4 and IPv6.
- The project scope includes creating VPP graph nodes, plugins etc. using those libraries to implement a host stack.

Honeycomb

Packet Processing

NSH\_SFC

ONE

VPP Sandbox

TLDK

VPP

Network IO - DPDK

Testing/Performance

CSIT



1. Remove DPDK patches from VPP so it can use DPDK unmodified.
2. Optimise DPDK vhost-user and replace vhost-vpp with DPDK vhost PMD.
3. Use DPDK cryptodev API to accelerate VPP IPsec.
4. Integrate DPDK QoS with VPP to support vPE use case.

Honeycomb

Pack Processing

NSH\_SFC

ONE

VPP Sandbox

TLDK

VPP

Network IO - DPDK

Testing/Performance

CSIT

# How You Can Help



**If you've made enhancements to DPDK, upstream them.**

See the Contributor's Guidelines at: <http://dpdk.org/doc/guides/contributing/index.html>

**If you are planning to submit enhancements to DPDK, communicate them as early as you can on the dev mailing list, and add them to the public roadmap.**

**Helps to avoid duplication.**

**Gives us a complete picture of what will be in the next release.**

**Allows others to comment earlier, which can save you time and effort later.**

# Further Information



**Data Plane Development Kit (DPDK) Open Source Website:** <http://dpdk.org/>

Source code: <http://dpdk.org/browse/dpdk/>

Documentation: <http://dpdk.org/doc>. Includes Getting Started Guides, Release Notes, Programmer's Guide, NIC Guides, Sample App Guide, API Guide etc.

Mailing Lists: <http://dpdk.org/ml>. The most commonly used are “dev” for development discussions and patches, and “users” for usage discussions.

Roadmap: <http://dpdk.org/dev/roadmap>

**Intel® Network Builders University:** <https://networkbuilders.intel.com/university/coursecategory/dpdk>

Includes “DPDK 101: Introduction to Data Plane Development Kit”, “DPDK 201: Sample Applications and New Features Deep Dive” and other DPDK-related training content.

**Other DPDK-Related Webinars:** <https://www.brighttalk.com/>, search for “DPDK”

Examples include “DPDK 16.04 New Features”, “Open vSwitch with DPDK in OVS 2.5.0”, “Accelerating Your Cloud With DPDK”, “IP Flow Analytics Enabled by Saisei and DPDK” etc.

**DPDK Events:** <https://dpdksummit.com/>

Provides information on past and future DPDK events. Videos of presentations at previous events are available on the Past Events page.

# Legal Disclaimers



No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. **No computer system can be absolutely secure.** Check with your system manufacturer or retailer or learn more at [intel.com](http://intel.com).

© 2016 Intel Corporation. Intel, the Intel logo, Intel. Experience What's Inside, and the Intel. Experience What's Inside logo are trademarks of Intel. Corporation in the U.S. and/or other countries.

\*Other names and brands may be claimed as the property of others.

Questions?

Tim O'Driscoll

[tim.odriscoll@intel.com](mailto:tim.odriscoll@intel.com)

# Service Assurance – OPNFV SFQM Project



Spaces ▾

People



Log in

Community

Developer

Documentation

Events

Infrastructure

Meetings

Others/Sketchpad

Projects

ware/Releases

No Permission

Pages

## SW Fastpath Service Quality Metrics Home

Created by Bryan Sullivan, last modified by Maryam Tahhan on Aug 05, 2016

### Project: Software Fastpath Service Quality Metrics

Project "Software Fastpath Service Quality Metrics" is focused on the development of utilities and libraries which supports low latency, high performance packet processing paths (fast path) through the NFVI such that VNFs can:

- Measure Telco Traffic and Performance KPIs through that software fast path
- Detect and report violations that can be consumed by VNFs and higher level EMS/OSS systems

An ability to measure and enforce Traffic Quality KPI's (Service Assurance) in the data-plane will be mandatory for any Telco-grade NFVI Implementation, regardless of whether the data plane is implemented in hardware or software.

# Service Assurance Roadmap



- Define Extended Stats naming scheme to store metadata for each statistic.
- Added Keep Alive, to detect core failures.

- Improved the Extended Stats API to use ID-value pairs.
- Modified the keep-alive sample app to use the alarm/interrupt API.

- Add collection and reporting of latency measurements (min, max, average, jitter) and bit rate (average and peak) for IXGBE and I40E.

2.2

16.04

16.07

16.11

17.02

Future?

- Update the vHost PMD to support the extended statistics API.

- Extend Extended Stats to additional PMDs?
- We're working closely with the OPNFV SFQM project to determine future requirements.