



DPDK Summit

DPDK ARCHITECTURE AND ROADMAP DISCUSSION

KANNAN BABU RAMIA, INTEL
DEEPAK KUMAR JAIN, INTEL

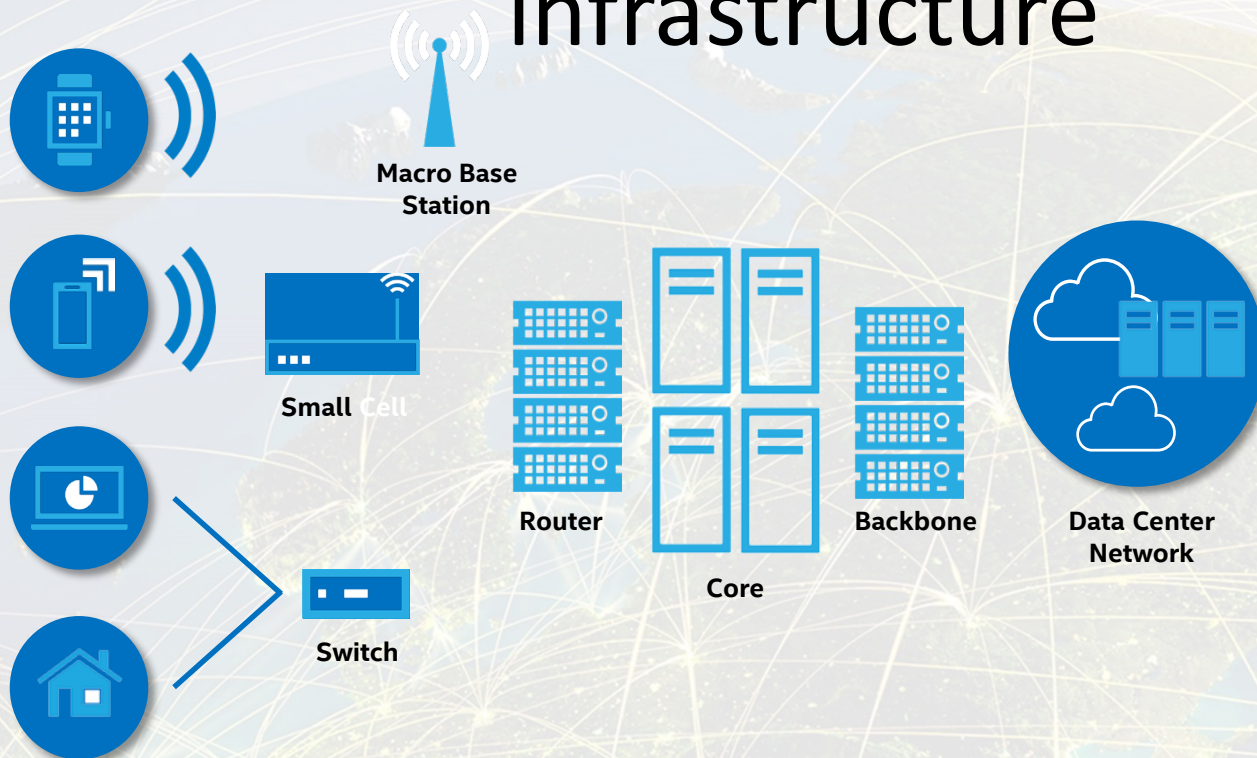
LEGAL DISCLAIMER

- 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.
- The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.
- Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm>
- Intel and the Intel 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.
- Copyright © 2017, Intel Corporation. All rights reserved.
- Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804
- Mileage may vary Disclaimer: Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks Test and System Configurations: Estimates are based on internal Intel analysis using at least Data Plane Development Kit IpSec sample application on Intel(R) Xeon(R) CPU E5-2658 v4@ 2.30GHz with atleast using Intel(R) Communications Chipset(s) 8955 with Intel(R) QuickAssist Technology.

Agenda

- Key trends in network transformation
- DPDK role
- DPDK Architecture
- Multi Architecture/ Multi vendor support
- Open source projects using DPDK
- DPDK Roadmap
- Open Questions

End-to-end network infrastructure



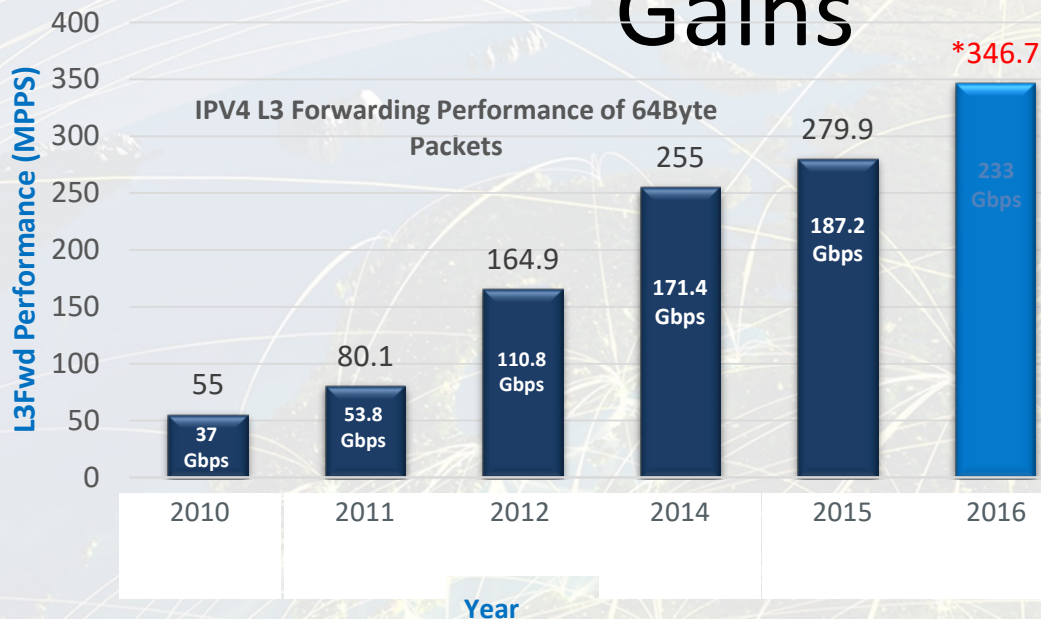
**WORKLOAD
CONVERGENCE**

**NETWORK
VIRTUALIZATION**

**END-TO-END
TRANSFORMATION**

DPDK Generational Performance

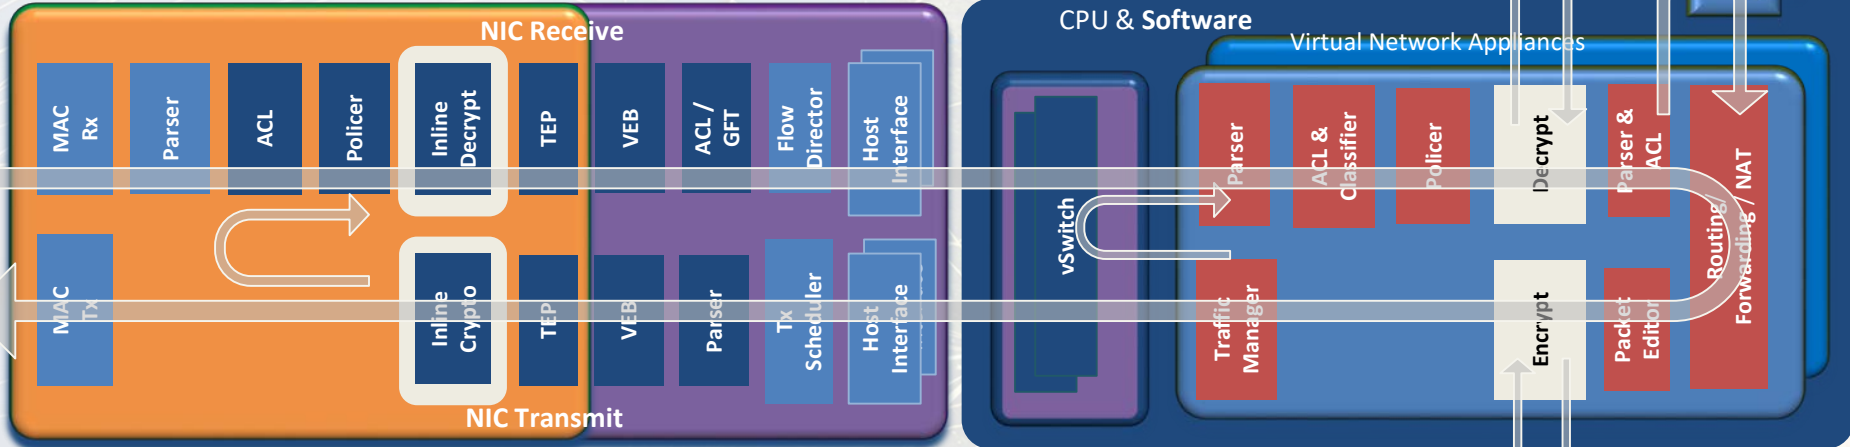
Gains



*System Configuration	
Hardware	
Platform	SuperMicro* - X10DRX
CPU	Intel® Xeon® E5-2658 v4 Processor
Chipset	Intel® C612 chipset
Sockets	2
Cores per Socket	14 (28 threads)
LL CACHE	30 MB
QPI/DMI	9.6GT/s
PCIe	Gen3x8
MEMORY	
DDR4 2400 MHz, 1Rx4 8GB (total 64GB), 4 Channel per Socket	
NIC	
10 x Intel® Ethernet CNA XL710-QDA2PCI-Express Gen3 x8 Dual Port 40 GbE Ethernet NIC (1x40G/card)	
NIC Mbps	40,000
BIOS	BIOS version: 1.0c (02/12/2015)
Software	
OS	Debian* 8.0
Kernel version	3.18.2
Other	DPDK 2.2.0

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/performance](#)

NFV – Life of a Packet



VEB – Virtual Embedded Bridge
 TEP – Tunnel End Point
 ACL – Access Control List
 GFT – Generic Flow Table

Outer Header (Underlay)


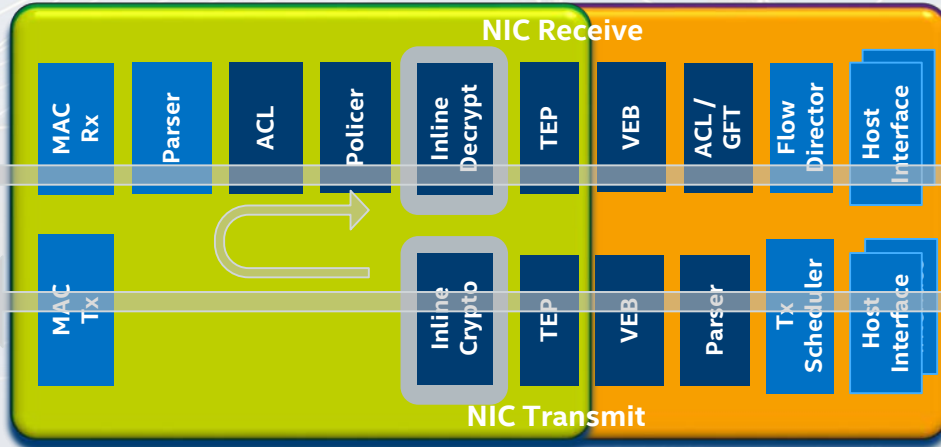
Inner Payload (Overlay)

Outer Ethernet Header	Outer IP Header	IPSec ESP transport	Outer UDP Header	Overlay Header (e.g. VXLAN)	Inner Ethernet Header	Inner IP Header	Inner L4 Header	Inner Data	IPSec ESP trailer	Outer CRC
-----------------------	-----------------	---------------------	------------------	-----------------------------	-----------------------	-----------------	-----------------	------------	-------------------	-----------

NFV - Service Function Chaining

Intel® QuickAssist
Technology
Crypto
Accelerator

VM to VM communication
between Virtual Network
Functions, needs multi-
100s of Gbps throughput

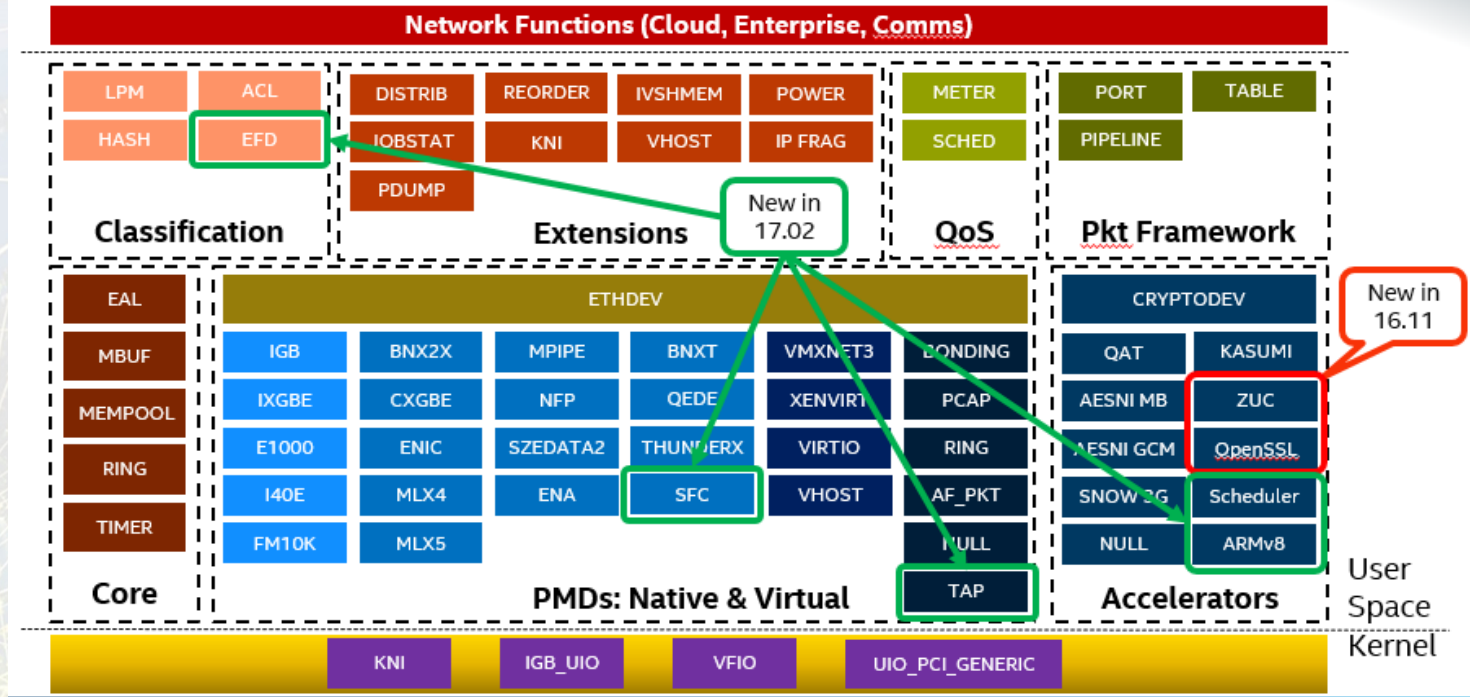


VEB – Virtual Embedded Bridge
TEP – Tunnel End Point
ACL – Access Control List
GFT – Generic Flow Table
VNF – Virtual Network Functions

DPDK Architecture

DPDK Fundamentals

- Implements run-to-completion and pipeline models
- No scheduler - all devices accessed by polling
- Supports 32-bit and 64-bit OSs with and without NUMA
- Scales from Intel® Atom™ to Intel® Xeon® processors
- Number of cores and processors is not limited
- Optimal packet allocation across DRAM channels
- Use of 2M & 1G hugepages and cache aligned structures
- Uses bulk concepts - processing 'n' packets simultaneously



Multi-Architecture/ Multi-Vendor Support

CPU Architectures

POWER 8
IBM

TILE-Gx



ARM v7/v8



Enhanced ARM Support



Event API



SoC Enhancements



2014

First non-IA
contributions.

2015

Non-Intel NIC support.

2016

Significant ARM vendor
engagement.

2017

SoC enhancements.
Non-Intel crypto.

ENIC
CISCO

BNX2X



MLX4/MLX5



ThunderX PMD



BNXT



ARMv8 Crypto

OcteonTX

LiquidIO



DPAA2



AVP



SZEDATA2



CXGBE



QEDE



ENA



SFC



ARK




NFP



Poll Mode Drivers

DPDK CONSUMPTION

vSwitches



Lagopus

BESS

DPDK in OS Distros



CentOS
Version 7.1 +



vRouters



Packet Generators



Pktgen



MoonGen

Ostinato

Storage

SPDK

+ Many more

TCP/IP Stacks

mTCP



Seastar

LWIP DPDK

DPDK Roadmap

Q1'17 (v 17.02)

- Added Elastic Flow Distributor library (`rte_efd`).L
- Added generic flow API (`rte_flow`).
- Added support for representing buses in EAL.
- Added APIs for MACsec offload support to the ixgbe PMD.
- Added VF Daemon (VFD) for i40e. – EXPERIMENTAL.
- virtio-user with vhost-kernel as another exceptional path.
- Added ARMv8 crypto PMD and updates to QAT, AESNI-MB PMDs.

Released

DPDK Roadmap

Q2'17 (v 17.05)

- **Added Eventdev PMD.**
- **Added event driven programming model library (rte_eventdev).**
- **Added bit-rate calculation, latency stats and information metric library.**
- **Kept consistent PMD batching behaviour.**
- **Added VFIO hotplug and vmxnet3 version 3 support.**
- **Added MTU feature support to Virtio and Vhost.**
- **Added interrupt mode support for virtio-user.**

DPDK Roadmap

Q2'17 (v 17.08)

- Generic QoS API
- Cryptodev Multi-Core SW Scheduler
- Generic Receive Offload
- Generic Flow Enhancements
- VF Port Reset for IXGBE
- API to Configure Queue Regions for RSS
- Support for IPFIX

OPEN QUESTIONS?

- ▶ What is missing from DPDK?
- ▶ What are the major pain-points in using DPDK?
- ▶ What can be improved in DPDK? Build process? Logging?
- ▶ What are the big performance bottlenecks?
- ▶ Working with Kernel?



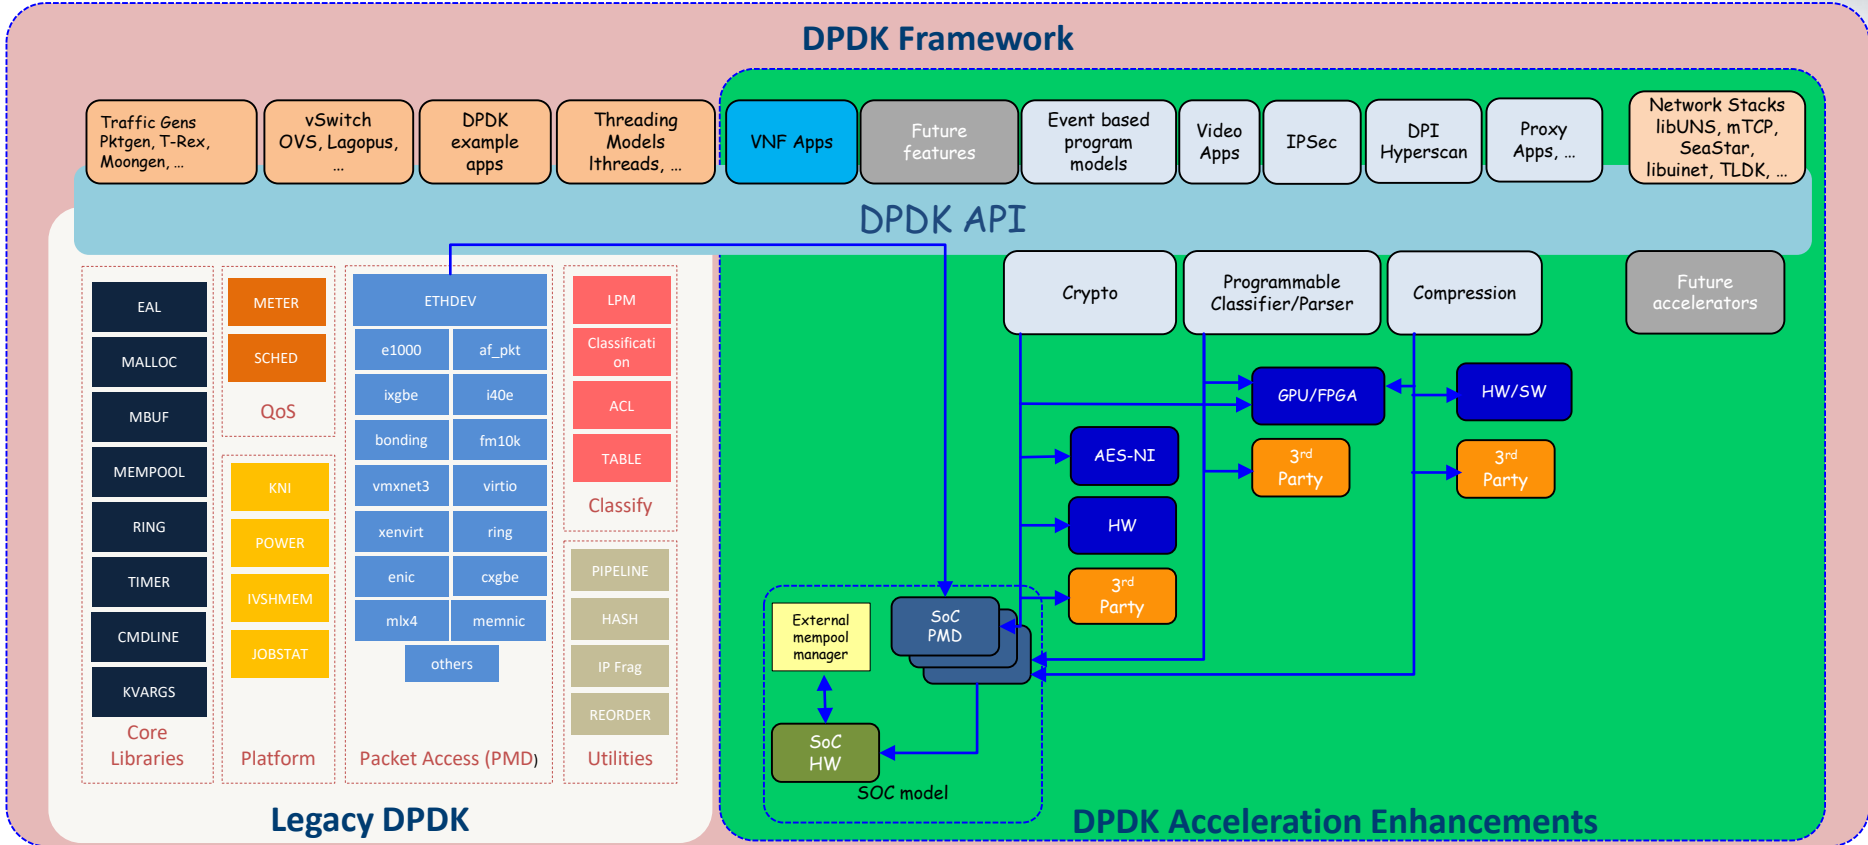
THANK YOU

DPDK Sample Apps



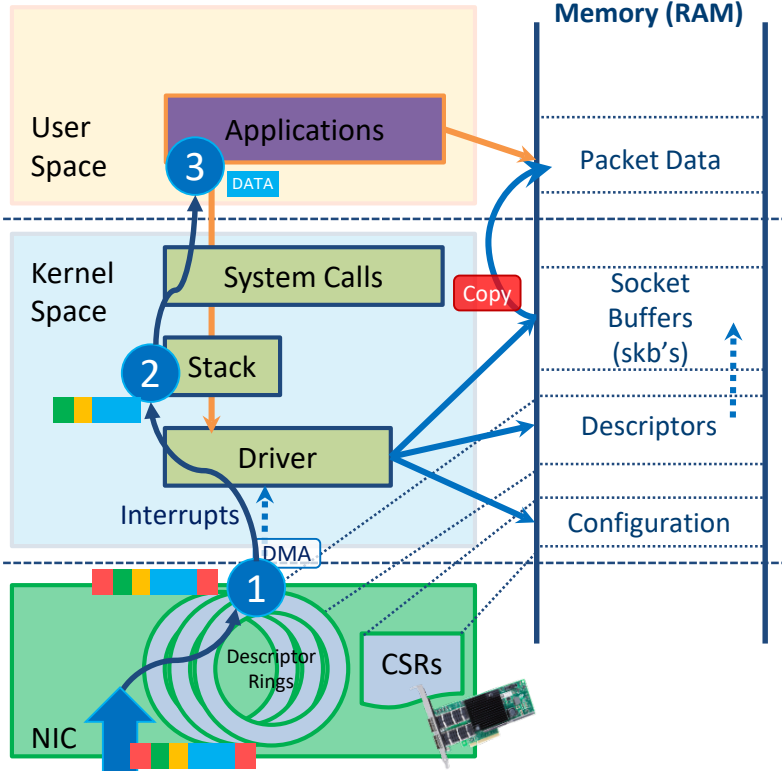
DPDK

DPDK Acceleration Enhancements

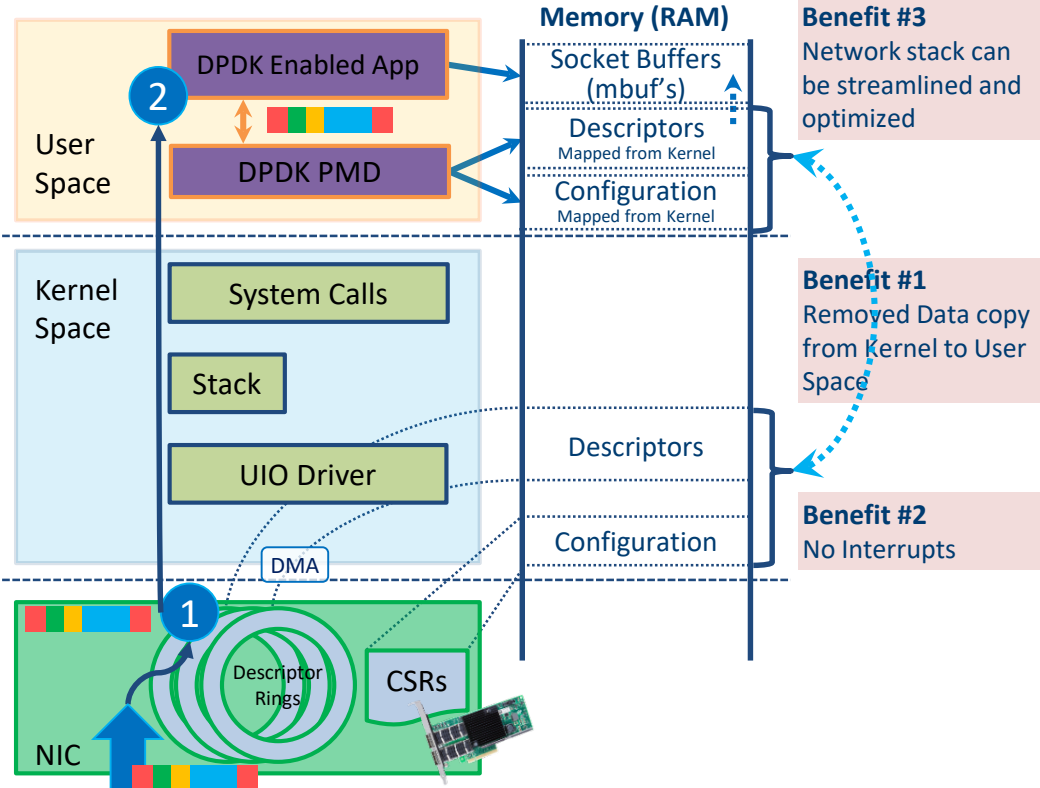


Packet Processing Kernel vs. User Space

Kernel Space Driver



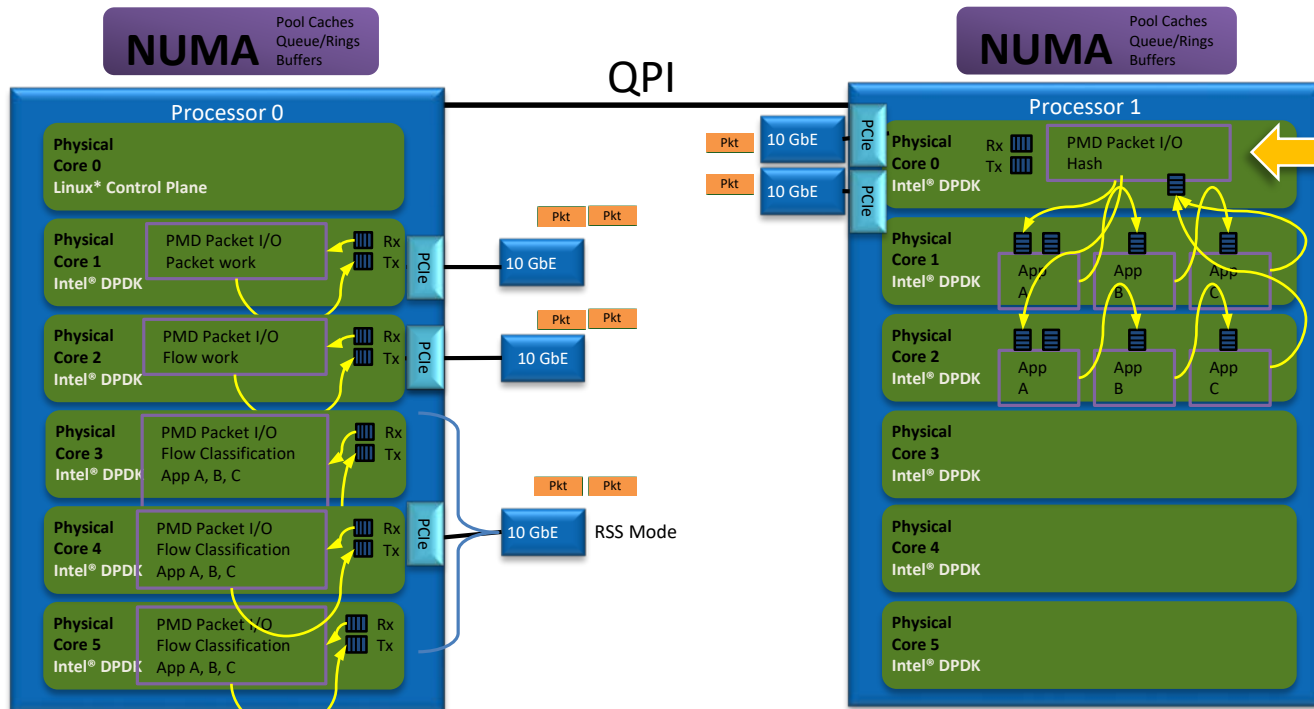
User Space Driver with Zero Copy



DPDK IN-DEPTH

PCIe* Connectivity and Core Usage

Using run-to-completion or pipeline software models



Can handle more I/O on fewer cores with vectorization

Run to Completion Model

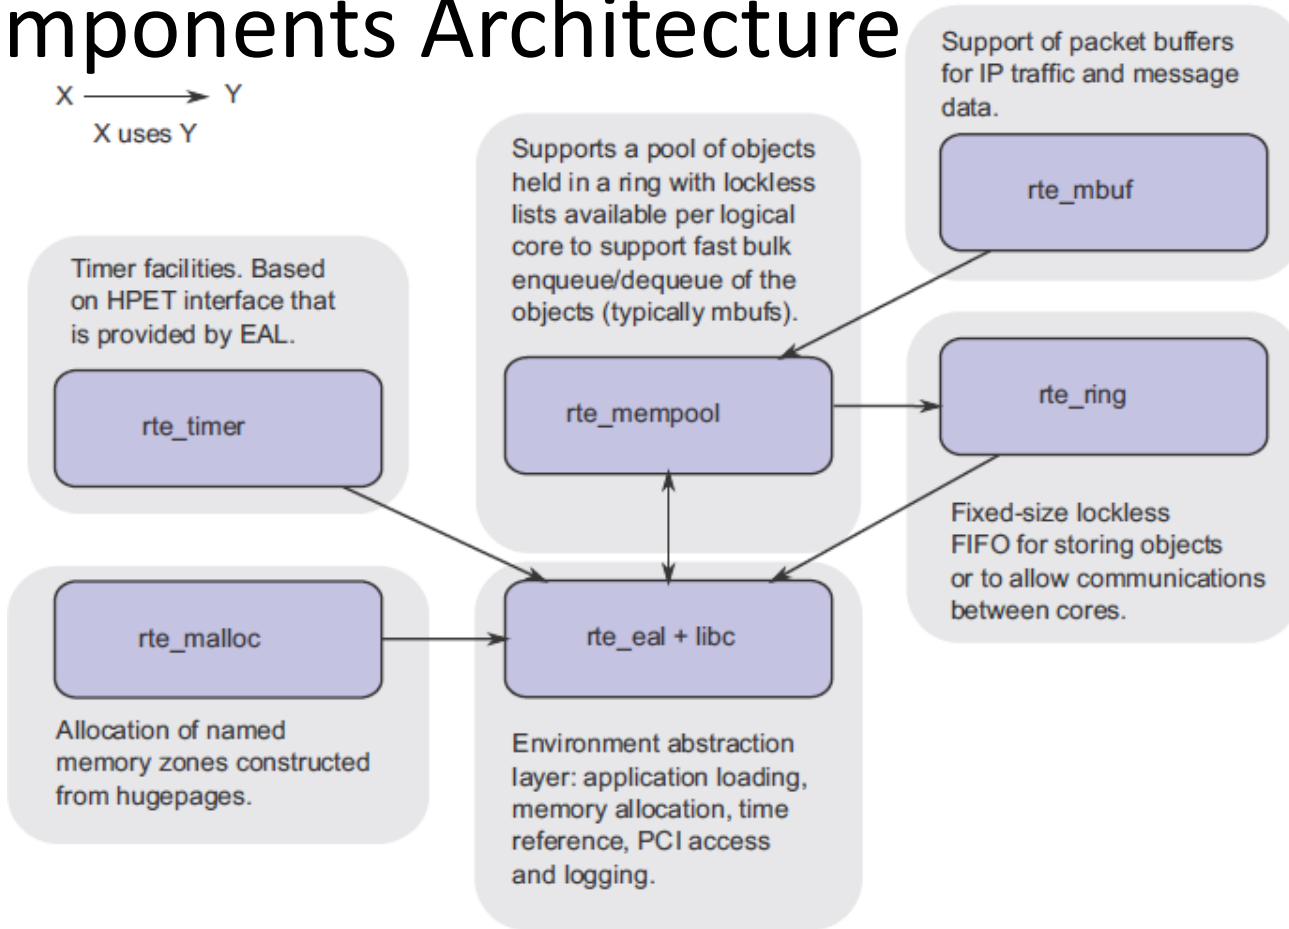
- I/O and Application workload can be handled on a single core
- I/O can be scaled over multiple cores

Pipeline Model

- I/O application disperses packets to other cores
- Application work performed on other cores


Core Components Architecture

X \longrightarrow Y
 X uses Y



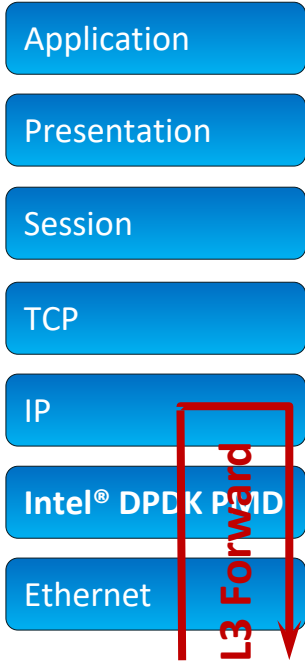
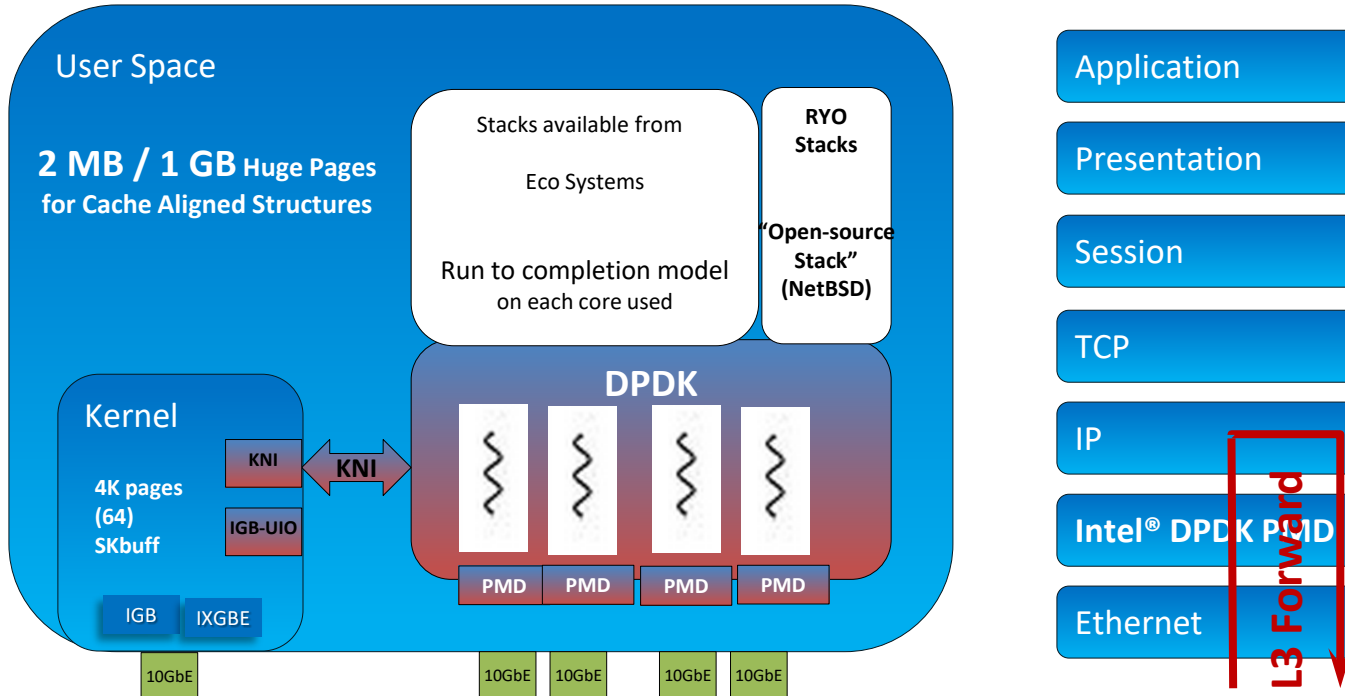
DPDK model

Intel® DPDK allocates packet memory equally across 2, 3, 4 channels.
Aligned to have equal load over channels



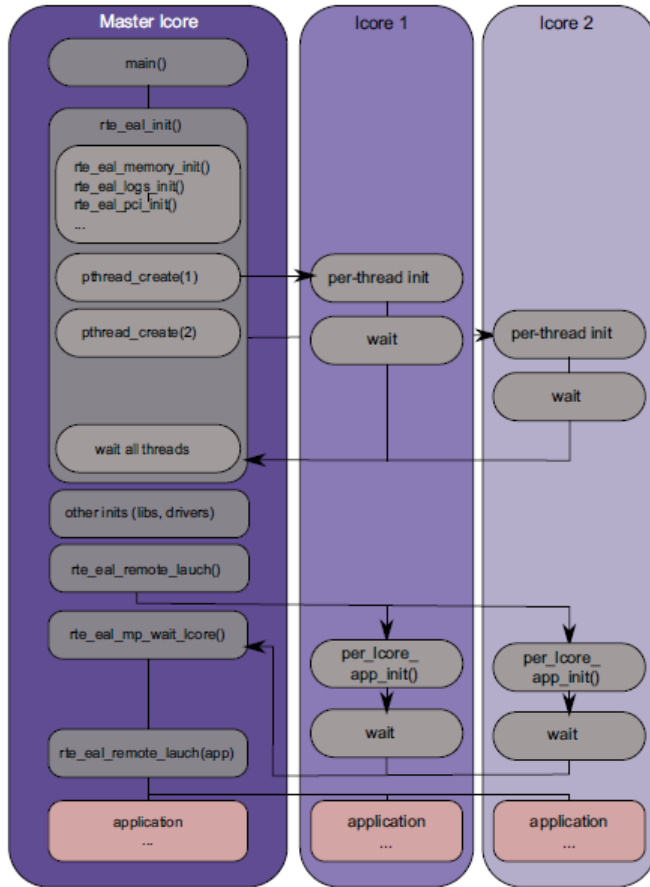
Rings for cached buffers

Per core lists, unique per lcore. Allows packet movement without locks



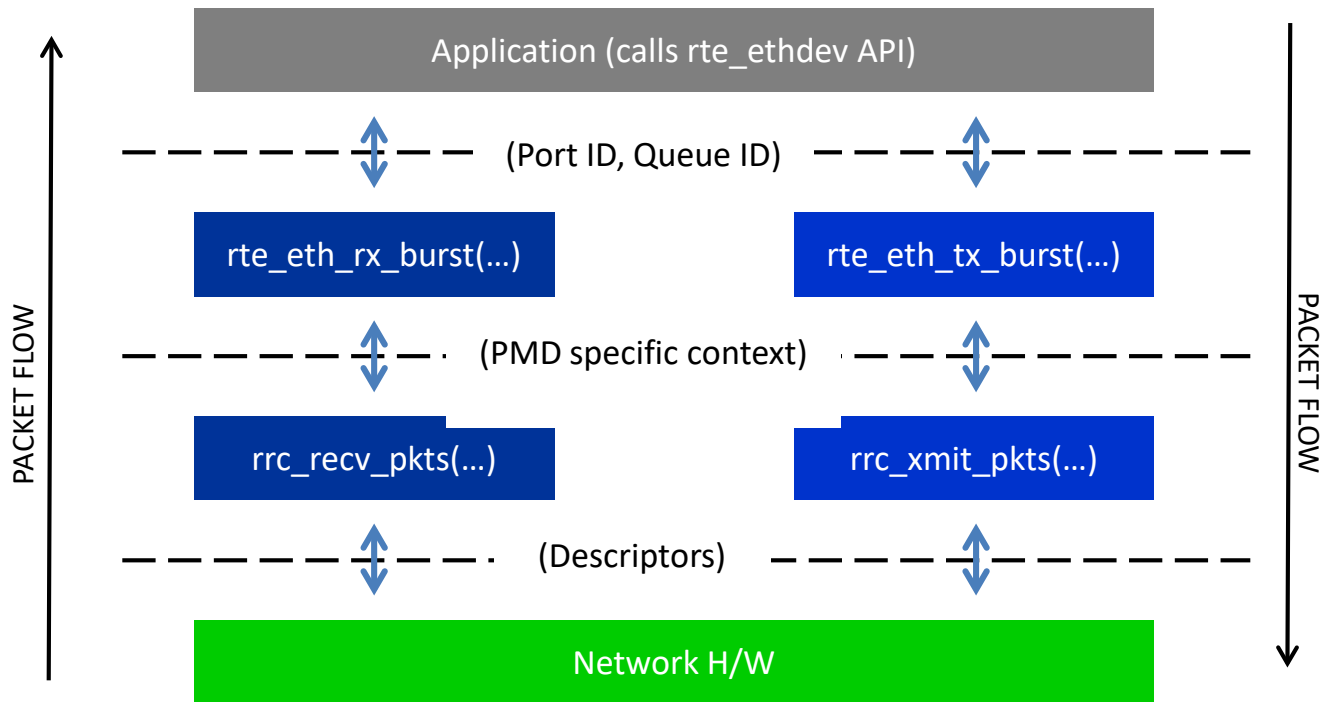
High Performance Components of DPDK

- Environment Abstraction Layer
 - Abstracts huge-page file system, provides multi-thread and multi-process support, etc.
- Memory Manager
 - Responsible for allocating pools of objects in memory. A pool is created in huge page memory space and uses a ring to store free objects. It also provides an alignment helper to ensure that objects are padded to spread them equally on all DRAM channels.
- Buffer Manager
 - Reduces by a significant amount the time the operating system spends allocating and de-allocating buffers. The Intel® DPDK pre-allocates fixed size buffers which are stored in memory pools.
- Queue Manager
 - Implements safe lockless queues, instead of using spinlocks, that allow different software components to process packets, while avoiding unnecessary wait times.
- Flow Classification
 - Provides an efficient mechanism which incorporates Intel® Streaming SIMD Extensions (Intel® SSE) to produce a hash based on tuple information so that packets may be placed into flows quickly for processing, thus greatly improving throughput.

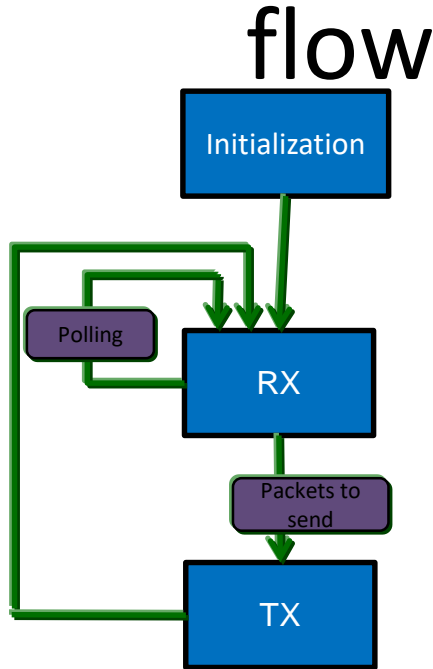


EAL Initialization in a Linux Environment

Ethernet Device Framework



30,000 ft overview of packet



1. Initialization

- Init Memory Zones and Pools
- Init Devices and Device Queues
- Start Packet Forwarding Application

2. Packet Reception (RX)

- Poll Devices' RX queues and receive packets in bursts
- Allocate new RX buffers from per queue memory pools to stuff into descriptors

3. Packet Transmission (TX)

- Transmit the received packets from RX
- Free the buffers that we used to store the packets