# DPDK Summit 2014

**Date:** September 8, 2014
**Time:** 8am - 5:30pm
**Location:** San Francisco Marriott Marquis
780 Mission Street
San Francisco - CA 94103

**A one-day event that brings together the DPDK open source community for technical presentations and dialogue**

# PROGRAM AT A GLANCE

| Time | Speaker | Session Title & Abstract |
|---|---|---|
| 8:00am | Registration & Refreshments | |
| 8:30am | **Jim St Leger**<br>*Software Product Line Manager, Intel* | **DPDK Summit Kick-off** |
| 9:00am | **Jun Xu**<br>*Principal Engineer, Futurewei Technologies, Inc.* | **Is It Time to Revisit the IP Stack in the Linux Kernel and KVM?**<br>We might take too many things for granted, like the Linux kernel providing a TCP/IP stack since its inception, whereas UNIX did not. Fast forward to today's world with virtualization, where most hypervisors derived from modern OSs also supply an IP stack. Should we take out the IP stack to let the OS and hypervisor focus on their main tasks, including process scheduling, resource management, and virtualization, or let them be the monolithic piece for all these elements? The Data Plane Development Kit (DPDK) provides an interesting starting point to explore other options. |
| 10:00am | **László Vadkerti**<br>*Ericsson Lead Software Developer*<br>**András Kovács**<br>*Ericsson Lead Software Developer* | **Multi-Socket Ferrari for NFV**<br>This presentation describes an approach to support latency-sensitive applications by diving into best practices in augmenting DPDK to deliver lower jitter and high availability in addition to higher packet throughputs. We will review Enhanced Memory Management techniques and multi-process enhancements to the DPDK library foundation.<br>We also will describe our experience and optimizations in using DPDK with a XEN Hypervisor including the addition of NUMA awareness. |
| 11:00 - 11:15pm | Coffee Break | |
| 11:15am | **Gabriel Silva**<br>*Microsoft Program Manager Windows Server Networking* | **Lightning Fast I/O for Windows Server v.Next with PacketDirect**<br>Microsoft operates some of the world's largest data centers, such as Bing, Office365, Xbox Live, and Azure, to name a few. One of the key fundamentals enabling Microsoft to operate efficiently at such hyper scale is NIC performance. In fact, many cloud operators are trying to address the same challenge – how to drive up packet processing performance for the network functions running in their data centers. With Windows Server v.Next, Microsoft is exploring a new I/O path for accelerated packet processing. Early experiments have shown nearly an order of magnitude improvement using PacketDirect on a DPDK-capable, Intel® 82599 10 Gigabit Ethernet Controller. In this talk, we describe some of the scenarios and an early architecture, and show a demonstration of this capability in action. |
| 12:15pm | Lunch Break | |
| 1:00pm | **Yoshihiro Nakajima**<br>*Researcher, Media Innovation Laboratory, NTT Network Innovation Laboratories* | **A High-Performance vSwitch of the User, by the User, for the User**<br>A high-performance software switch is a key component for next-generation telecom infrastructure, especially for NFV and SDN. NTT Laboratories developed a high-performance and highly-scalable SDN/OpenFlow software switch, called Lagopus, that leverages state-of-the-art server and software technologies, including Intel® processors, Intel® Ethernet Controllers, and the DPDK. We present implementation details about Lagopus and demonstrate a carrier use case using Ryu, an SDN/OpenFlow controller. |
| 2:00pm | **Venky Venkatesan**<br>*DPDK architect, Intel* | **Application Performance Tuning and Future Optimizations in DPDK**<br>In this session, one of the original authors of the DPDK library will share insight into how to best use the available tools and library hooks when looking to optimize system packet performance. The session will also provide insight into concepts under consideration to facilitate discussion and prioritization feedback into the future planning process. |
| 3:00 - 3:15pm | Coffee Break | |
| 3:15pm | **Bhavesh Davda**<br>*VMware Senior Staff Engineer, Office of CTO*<br>**Rashmin Patel**<br>*DPDK Virtualization Engineer, Intel* | **DPDK in a Virtual World**<br>The usage of virtualized DPDK applications has increased tremendously. This session will review how the DPDK APIs take advantage of platform technologies like SR-IOV, direct device assignment (VT-d) and para-virtual as well as emulated devices offered by the underlying platform to achieve higher packet throughput at predictable latency. The session will primarily focus on the virtualization options offered by DPDK for the VMware ESXi Hypervisor environment. The session will conclude by sharing the future vision and commitment to enhance the API even further to enable community developers and end users to get most out of underlying Intel Architecture and Hypervisor target. |
| 4:15pm | **Thomas Monjalon**<br>*6WIND Packet Processing Engineer and DPDK.org Maintainer* | **High-Performance Networking Leveraging the DPDK and the Growing Community**<br>High-performance networking stacks can be designed using the DPDK and packet processing software. This presentation covers the development of high-performance applications with examples for IPsec, TCP, virtual switching, and virtual networking functions for NFV. We also describe how the DPDK.org community is growing for the benefit of high-performance networking enthusiasts. |
| 5:15 - 5:30pm | **Tim O'Driscoll**<br>*Software Engineering Manager for DPDK, Intel* | **Closing Remarks** |