

DPDK Summit North America 2017 - November 14 - 15, 2017

Title	Description	Presenter	Company
Opening Remarks & Governing Board	Introduction to the event, including a review of the agenda, logistics and expectations. An update from the Governing Board on who the Governing Board are, what their responsibilities are, progress to date, future priorities/challenges for the project.	Jim St. Leger	Intel
Community Survey Feedback	We conducted a survey of the DPDK community, soliciting input on a variety of topics including DPDK usage, roadmap, performance, patch submission process, documentation and tools. This session will present the results of the survey, which will help to guide the future direction of the project.	John McNamara	Intel
Reducing Barriers to Adoption - Making DPDK Easier to Integrate into Your Application	While DPDK is a widely-adopted software package for high-performance networking applications, there are a number of ways in which it is harder to use than it otherwise needs to be. This is especially true when it comes to integrating DPDK with an existing legacy codebase. This presentation will look at some of the issues and provide an update on current development and prototyping work to simplify DPDK integration with existing code.	Bruce Richardson	Intel
New Command Line Interface for DPDK	The current command line interface for DPDK called cmdline has a number of limitation and a complex user design. The next command line for DPDK called CLI is more dynamic with a simple directory style design. The directory style design allows for commands to be placed in a hierarchy for easy integration, plus supporting a simple argc/argv function interface. Using these features reduced the LOC in test-pmd cmdline file from 12K to ~4K. The presentation includes an example usage.	Keith Wiles	Intel
Event Adapters - Connecting Devices to Eventdev	Recently, the DPDK has enabled applications to use dynamically load balanced pipelines with the introduction of libeventdev. In addition to using eventdev for CPU to CPU pipelines, devices such as ethdev, cryptodev and timers need to be able to inject events into eventdev. Currently, we are in the process of upstreaming extensions to eventdev called eventdev adapters for each of these devices that would allow applications to configure event input from these devices to the event device. We will discuss each of the adapter APIs and show example code that allow event based applications to be written in a platform independent manner.	Nikhil Rao	Intel

Title	Description	Presenter	Company
GRO/GSO Libraries: Bring Significant Performance Gains to DPDK-based Applications	A major part of packet processing has to be done on a per-packet basis, such as switching and TCP/IP header processing. The overhead of the per-packet routines, however, exerts a significant impact on the performance of network processing. Generic Receive Offload (GRO) and Generic Segmentation Offload (GSO) are two effective techniques for mitigating the per-packet processing overhead by reducing the number of packets to be processed. Specifically, GRO merges the receiving packets of the same flow in RX, while GSO delays packet segmentation in TX.	Jiayu Hu	Intel
Power Aware Packet Processing	A drive to deliver OPEX saving and performance where and when it's needed. Enter a new era of power optimized packet processing. This talk reviews new & existing DPDK extensions for policy based power control proposed in August and the associated performance benefits.	Chris MacNamara	Intel
Enhanced Memory Management	In this presentation we will be reviewing Enhanced Memory Management techniques and multi-process enhancements as a possible way to seamlessly	Laszlo Vadkerti	Ericsson
		Jiangtao Zhang	Ericsson
Making networking apps scream on Windows with DPDK	Network bandwidth is precious and milliseconds matter for many user-mode applications and virtual appliances running on both Linux and Windows. In order	Jason Messer	Microsoft
Mediated Devices: Better Userland IO	Unbinding Linux kernel drivers to allow userland IO through VFIO has a number of disadvantages such as another large touchy code base to deal with the hardware, loss of standard Linux tools (ifconfig, ethtool, tcpdump, SNMPd...) and impossibility to accelerate container networking. Mediated device introduced in Linux kernel 4.10 for GPUs and provisions for additional devices hold the promise of collaboration between kernel drivers and userland application in need of direct datapath steering.	Manasi Deval	Intel
		François-Frédéric Ozog	Linaro
Mellanox bifurcated driver model	Mellanox PMD uses verbs instead of taking full control over the device (PCI). That allows the kernel (netdev) and more than a single PMD to run on a single PCI function. If the DPDK app is not steering by <code>rte_flow</code> , all the traffic the packets be processed by the kernel net device.	Rony Efraim	Mellanox

Title	Description	Presenter	Company
DPDK with KNI – Pushing the Performance of an SDWAN Gateway to Highway Limits!	An SDWAN gateway is usually built with an x86 commercial off-the-shelf (COTS) hardware that often runs a variant of Linux Operating System and requires high throughput for connecting a corporate’s branch network with its Data Centers. However owing to the inherent limitations of standard 4K sized pages without dedicated resource allocations in a general-purpose Linux kernel, it has been seen that even a high-end SDWAN gateway hardware cannot forward traffic to its full potential.	Sabyasachi Sengupta	Nuage Networks
DPDK as microservices in ZTE Paas	To provide high performance for ICT (Information Communications Technology) area, we use DPDK as a micro service in container networking. We used	Yong Wang	ZTE
Accelerate Clear Container Network performance	Clear Container is a great technology to secure a container with a fast and lightweight hypervisor, and there might be very different type of workloads running inside Clear Containers, e.g. some workloads require high packet processing rate (PPS) and some workloads require massive data transfer (BPS), given Clear Container’s much higher density than Virtual Machine, a high performance virtual switch is very critical and demands is highly emerged, but current available virtual switches is still far behind those demands.	Songming Yan	ZTE
The Path to Data Plane Microservices	Clear Container is a great technology to secure a container with a fast and lightweight hypervisor, and there might be very different type of workloads running inside Clear Containers, e.g. some workloads require high packet processing rate (PPS) and some workloads require massive data transfer (BPS), given Clear Container’s much higher density than Virtual Machine, a high performance virtual switch is very critical and demands is highly emerged, but current available virtual switches is still far behind those demands.	Jun Xiao	CloudNetEngine
The Path to Data Plane Microservices	DPDK revolutionized software packet processing initially for discrete appliances and then for Virtual Network Functions. Containers and μ Services technology are extensively used as a means to scale up and out in the Cloud. These technologies now include Comms Service Providers among their advocates, and embracing these technologies with their scaling model and resiliency is the new frontier in software packet processing.	Ray Kinsella	Intel
Container Panel Discussion	A panel discussion with Yong Wang, Songming Yan, Jun Xiao and Ray Kinsella to discuss DPDK enablement of containers and micro-services.	N/A	N/A
Accelerate storage service via SPDK	SPDK (storage performance development kit, http://spdk.io) is an open source library used to accelerate the storage service (e.g., file, block) especially for PCIe SSDs (e.g., 3D Xpoint SSDs). The foundation of SPDK is the user space, asynchronous and polled mode drivers (e.g., IOAT and NVMe), and the idea of which is similar to DPDK.	Jim Harris	Intel

Title	Description	Presenter	Company
Accelerating P4-based Dataplane with DPDK	The high-level P4 programming language promises protocol and hardware-agnostic design of network functions. As the low-level functional implementation, the P4 Behavior Model (BMv2) provides the necessary constructional blocks (parser, deparser, lookup tables, and action primitives, etc.) into which any P4 dataplane programs can be compiled.	Peilong Li	University of Massachusetts Lowell
Implementation and Testing of Soft Patch Panel	SPP is a framework to easily interconnect DPDK applications on host and guest virtual machines together, and assign resources dynamically to these	Tetsuro Nakamura	NTT
Reflections on Mirroring With DPDK	Debugging network problems is often hard, and further complicated when a guest O/S is provided with an SR-IOV VF bound to a DPDK driver because tools running on the physical host (e.g. tcpdump) lose visibility to the interface. Hardware mirroring of traffic to another VF provides the ability to regain visibility and to help facilitate the troubleshooting process.	E. Scott Daniels	AT&T Labs
A network application API on top of device APIs	NFV promise is to be able to instantiate or even live migrate VMs on different platforms and have applications benefit from whatever acceleration is available. As a result, the application developer shall not make compilation or define application architecture based on what he/she expects from the runtime environment. ODP and DPDK have in common the concept of "device" APIs (Ethernet, crypto, events, IPsec, compression....) with distinct approaches.	François-Frédéric Ozog	Linaro
SafetyOrange - a tiny server class multi-purpose box with DPDK	SafetyOrange is a portable (4.3 liter) and silent Xeon computer. Well, it is larger than 'DPDK in a box' but it supports two NICs (as of now sporting 2 XL710 cards),	Andras Kovacs	Ericsson
Technical Roadmap	An update from the Technical Board covering the future roadmap and technical challenges for the project.	Laszlo Vadkerti	Ericsson
rte_raw_device: implementing programmable accelerators using generic offload	An update from the Technical Board covering the future roadmap and technical challenges for the project.	Technical Board	N/A
rte_raw_device: implementing programmable accelerators using generic offload	There are various kinds of HW accelerators available with SoCs. Each of the accelerators may support different capabilities and interfaces. Many of these	Hemant Agrawal	NXP
rte_raw_device: implementing programmable accelerators using generic offload	There are various kinds of HW accelerators available with SoCs. Each of the accelerators may support different capabilities and interfaces. Many of these	Shreyansh Jain	NXP
DPDK support for new hardware offloads	Fully programmable SmartNICs allow new offloads like OVS, eBPF, P4 or vRouter, and the Linux kernel is changing for supporting them. Having these same offloads when using DPDK is a possibility although the implications are not clear yet. We present Netronome's perspective for adding such a support to DPDK mainly for OVS and eBPF.	Alejandro Lucero	Netronome
Flexible and Extensible support for new protocol processing with DPDK using Dynamic Device Personalization	Dynamic Device Personalization allows a DPDK application to enable identification of new protocols, for example, GTP, PPPoE, QUIC, without changing the hardware. The demo showcases a DPDK application identifying and spreading traffic on GTP and QUIC. Dynamic Device Personalization can be	Andrey Chilikin	Intel
Flexible and Extensible support for new protocol processing with DPDK using Dynamic Device Personalization	Dynamic Device Personalization allows a DPDK application to enable identification of new protocols, for example, GTP, PPPoE, QUIC, without changing the hardware. The demo showcases a DPDK application identifying and spreading traffic on GTP and QUIC. Dynamic Device Personalization can be	Brian Johnson	Intel

Title	Description	Presenter	Company
Serverless DPDK - How SmartNIC resident DPDK Accelerates Packet Processing	Cloud architectures and business models are driving the need to ensure that all server compute resources have a revenue tie-in, heralding the march towards the serverless dataplane. This session presents a unique way to harness the power of DPDK to accelerate packet processing by pushing the data plane into a SmartNIC. We will discuss the motivation, benefits and challenges of implementing a DPDK based data plane running on the compute resources embedded in a SmartNIC.	Nishant Lodha	Cavium
Enabling hardware acceleration in DPDK data plane applications	This presentation will look at the challenges faced in leveraging hardware acceleration in DPDK enabled applications, addressing some of the problems posed in creating consistent hardware agnostic APIs to support multiple accelerators with non-aligned features, and the knock implications this can have to application designs.	Declan Doherty	Intel
rte_security: enhancing IPSEC offload	In this talk we present a joint work of NXP, Intel and Mellanox on offloading security protocol processing to hardware providing better utilization of host CPU for packet processing. This talk provides the overview of new enhancement	Hemant Agrawal	NXP
		Declan Doherty	Intel
		Boris Pismenny	Mellanox
Mellanox FPGA	The FPGA allows a wide variety of features to be supported in DPDK. We observe that programmable HW is useful for packet-processing pipelines. For example, consider a pipeline of multiple match-action operations, in which actions may also specify generic packet modifications that are carried out by accelerators. In this case, the CPU is only involved at the beginning (transmission) or end (reception) of the pipeline, while the accelerator invocations are initiated by NIC matching operations.	Boris Pismenny	Mellanox
SMARTNIC, FPGA, IPSEC Panel discussion	A panel discussion with Hemant Agrawal, Alejandro Lucero, Andrey Chilikin, Brian Johnson, Nishant Lodha, Declan Doherty and Boris Pismenny to discuss DPDK enablement for smart NICs, FPGA and IPsec.	N/A	N/A
VPP Host Stack	Although packet forwarding with VPP and DPDK can now scale to tens of millions of packets per second per core, lack of alternatives to kernel-based sockets means that containers and host applications cannot take full advantage of this speed. To fill this gap, VPP was recently added functionality specifically designed to allow containerized or host applications to communicate via shared-memory if co-located, or via a high-performance TCP stack inter-host.	Florin Coras	Cisco

Title	Description	Presenter	Company
DPDK's best kept secret – Micro-benchmark performance tests	To have apple to apple comparisons, developers need a common ground of base level metrics. That common ground is ability to identify the basic DPDK building block of importance (as well as relevance to the work load) e.g., producer/consumer rings and measure the cycle cost associated with basic operation like enqueue/dequeueing – bulk versus single.	Muthurajan Jayakumar	Intel
DPDK on Microsoft Azure	SDN is at the foundation of all large scale networks in the public cloud, such as Microsoft Azure. But how do we make a software network scale to an era of	Daniel Firestone	Microsoft
OpenNetVM: A high-performance NFV platforms to meet future communication challenges	To truly achieve the vision of a high-performance software-based network that is flexible, lower-cost, and agile, a fast and carefully designed NFV platform along with a comprehensive SDN control plane is needed. Our high-performance NFV platform, OpenNetVM, exploits DPDK and enables high bandwidth network functions to operate at near line speed, while taking advantage of the flexibility and customization of low cost commodity servers.	K. K. Ramakrishnan	Univ. of California, Riverside
Make DPDK's software traffic manager a deployable solution for vBNG	Achieving network functions parity across purpose-built ASIC implementation and virtual implementation is not straightforward. Irrespective of differences in performance capability between purpose-built and virtual environments. Functional disfiguration represents a significant obstacle in operators' adoption of virtualization as it implies a dependency on access/aggregation network topology and configuration.	Csaba Keszei	Ericsson
testpmd: swissknife for NFV	testpmd supports all DPDK drivers, exposes all of their knobs in its command line: this has been a great tool for developers from dpdk.org origin. But beyond dpdk.org community, I would like to present how testpmd is used for NFV developments, performances characterization, and also for post deployment validation: how it's done today, what is under addition, and what we need next.	Franck Baudin	Red Hat
OpenVswitch hardware offload over DPDK	Telcos and Cloud providers are looking for higher performance and scalability when building nextgen datacenters for NFV & SDN deployments. While running OVS over DPDK reduces the CPU overload of interrupt driven packet processing, CPU cores are still not completely freed up from polling of packet queues.	Rony Efraim	Mellanox

Title	Description	Presenter	Company
Accelerating NFV with VMware's Enhanced Network Stack (ENS) and Intel's Poll Mode Drivers (PMD)	Network Functions Virtualization (NFV) deployments are happening at a rapid pace. This is driving the need to more efficiently consolidate compute, storage	Jin Heo	VMware
DPDK Membership Library	In this talk we will present the new DPDK Membership Library, this library is used to create what we call a "set-summary" which is a new data structure that is used to summarize large set of elements. It is the generalization and extension to the traditional filter structure, e.g. bloom filter, cuckoo filter, etc to efficiently test if a key belongs to a large set.	Rahul Shah Sameh Gobriel	Intel
Integrating and using DPDK with Open vSwitch	Some applications are written from the ground up with DPDK in mind, but Open vSwitch is not one of them. This talk will look at how Open vSwitch integrated and uses DPDK. It will look at various aspects such as DPDK initialization, threading, and the usage of DPDK PMD's and libraries. It will also talk about	Aaron Conole	Red Hat
Lagopus Router	In this talk, we introduce a new open source router implementation called Lagopus Router. It is an extensible microservice architecture router that consists	Kevin Traynor	Red Hat
vSwitch Panel Discussion	In this talk, we introduce a new open source router implementation called Lagopus Router. It is an extensible microservice architecture router that consists	Tomoya Hibi	NTT
Closing Remarks	A panel discussion with Rony Efraim, Jin Heo, Rahul Shah, Sameh Gobriel, Charlie Tai, Aaron Conole, Kevin Traynor, Tomoya Hibi and Hirokazu Takahashi to discuss DPDK acceleration of vswitches.	Hirokazu Takahashi	NTT
		N/A	N/A
		N/A	N/A

