



Contribution ID: 157

Type: **not specified**

What's Happened to the World of Networking Hardware Offloads?

Wednesday, 14 November 2018 14:00 (35 minutes)

Over the last 10 years the world has seen NICs go from single port, single netdev devices, to multi-port, hardware switching, CPU/NFP having, FPGA carrying, hundreds of attached netdev providing, behemoths. This presentation will begin with an overview of the current state of filtering and scheduling, and the evolution of the kernel and networking hardware interfaces. (HINT: it's a bit of a jungle we've helped grow!) We'll summarize the different kinds of networking products available from different vendors, and show the workflows of how a user can use the network hardware offloads/accelerations available and where there are still gaps. Of particular interest to us is how to have a useful, generic hardware offload supporting infrastructure (with seamless software fallback!) within the kernel, and we'll explain the differences between deploying an eBPF program that can run in software, and one that can be offloaded by a programmable ASIC based NIC. We will discuss our analysis of the cost of an offload, and when it may not be a great idea to do so, as hardware offload is most useful when it achieves the desired speed and requires no special software (kernel changes). Some other topics we will touch on: the programmability exposed by smart NICs is more than that of a data plane packet processing engine and hence any packet processing programming language such as eBPF or P4 will require certain extensions to take advantage of the device capabilities in a holistic way. We'll provide a look into the future and how we think our customers will use the interfaces we want to provide both from our hardware, and from the kernel. We will also go over the matrix of most important parameters that are shaping our HW designs and why.

Presenters: BRANDEBURG, Jesse (Intel); JAIN, Anjali Singhai (Intel)

Session Classification: Networking Track