



Contribution ID: 161

Type: **not specified**

Using eBPF as an Abstraction for Switching

Wednesday, 14 November 2018 16:35 (35 minutes)

eBPF (extended Berkeley Packet Filter) has been shown to be a flexible kernel construct used for a variety of use cases, such as load balancing, intrusion detection systems (IDS), tracing and many others. One such emerging use case revolves around the proposal made by William Tu for the use of eBPF as a data path for Open vSwitch. However, there are broader switching use cases developing around the use of eBPF capable hardware. This talk is designed to explore the bottlenecks that exist in generalising the application of eBPF further to both container switching as well as physical switching.

Topics that will be covered include proposals for container isolation through the use of features such as programmable RSS, the viability of physical switching using eBPF capable hardware as well as integrations with other subsystems or additional helper functions which may improve the possible functionality.

Presenter: VILJOEN, Nick (Netronome)

Session Classification: Networking Track