



Contribution ID: 142

Type: **not specified**

Leveraging Kernel Tables with XDP

Tuesday, November 13, 2018 9:55 AM (35 minutes)

XDP is a framework for running BPF programs in the NIC driver to allow decisions about the fate of a received packet at the earliest point in the Linux networking stack. For the most part the BPF programs rely on maps to drive packet decisions, maps that are managed for example by a userspace agent. This architecture has implications on how the system is configured, monitored and debugged.

An alternative approach is to make the kernel networking tables accessible by BPF programs. This approach allows the use of standard Linux APIs and tools to manage networking configuration and state while still achieving the higher performance provided by XDP. An example of providing access to kernel tables is the recently added helper to allow IPv4 and IPv6 FIB (and nexthop) lookups in XDP programs. Routing suites such as FRR manage the FIB tables, and the XDP packet path benefits by automatically adapting to the FIB updates in real time. While a huge first step, a FIB lookup alone is not sufficient for general networking deployments.

This talk discusses the advantages of making kernel tables available to XDP programs to create a programmable packet pipeline, what features have been implemented as of October 2018, key missing features, and current challenges.

Presenter: AHERN, David (Cumulus Networks)

Session Classification: Networking Track