Contribution ID: **162** Type: **not specified**

BPF Host Network Resource Management

Wednesday, November 14, 2018 5:10 PM (35 minutes)

Linux currently provides mechanisms for managing and allocating many of the system resources such as CPU, Memory, etc. Network resource management is more complicated since networking deals not only with a local resource, such as CPU management does, but can also deal with a global resource. The goal is not only to provide a mechanism for allocating the local network resource (NIC bandwidth), but also to support management of network resources external to the host, such as link and switch bandwidths.

For networking, the primary mechanism for allocating and managing bandwidth has been the traffic control (tc) subsystem. While tc allows for shaping of outgoing traffic and policing of incoming traffic, it suffers from some drawbacks. The first drawback is a history of performance issues when using the Hierarchical Queuing Discipline (htb) which is usually required for anything other than simple shaping needs. A second drawback is the lack of flexibility usually provided by general programming constructs.

We are in the process of designing and implementing a BPF based framework for efficiently supporting shaping of both egress and ingress traffic based on both local and global network allocations.

Presenters: BRAKMO, Lawrence (Facebook); STAROVOITOV, Alexei (Facebook)

Session Classification: Networking Track