



Contribution ID: 292

Type: **not specified**

Container Networking: The Play of BPF and Network NS with different Virtual Devices

Wednesday, 15 November 2023 15:00 (30 minutes)

In the container-centric ecosystem, achieving efficient network isolation without compromising on performance has become paramount. Not all containers require the stringent network isolation akin to VMs. Many can benefit from a more flexible approach, like using eBPF hooks, to mark and manage network traffic with QOS. This presentation delves into the application of cgroup-bpf based hooks (bind/connect/sendmsg) in crafting lightweight network isolation solutions.

Yet, there exist cases where the above cgroup-bpf solution falls short. Here, we'll explore the indispensable role of network namespaces in ensuring robust network isolation, separating container traffic from the host network. But with stringent isolation, come challenges. Latency-sensitive applications can experience performance bottlenecks in this setup.

To address these challenges, we have investigated various techniques such as veth, netkit, IPVLAN, and SR-IOV. Each has its merits and drawbacks in terms of performance, usability, configuration complexities, and kernel support requirements. Join us as we decode the intricacies of network isolation in a containerized world, offering insights for both novice and seasoned developers.

Primary authors: LAU, Martin (Meta); CHAHANDE, Takshak (Meta)

Presenters: LAU, Martin (Meta); CHAHANDE, Takshak (Meta)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track