



Contribution ID: 302

Type: **not specified**

Networking resource control with per-cgroup LSM

Tuesday, September 13, 2022 12:30 PM (30 minutes)

Google's container management system runs different workloads on the same host. To effectively manage networking resources, the kernel has to apply different networking policies to different containers.

Historically, most of the networking resource control happened inside proprietary Google networking cgroup. That cgroup is an interesting cross between upstream `net_cls` and `net_prio`, has a lot of Google-specific business logic and has no chance of being accepted upstream.

In this talk I'm going to talk about what we'd like to manage on the networking resource side and which BPF mechanisms were added to achieve this (`lsm_cgroup`).

I agree to abide by the anti-harassment policy

Yes

Primary author: FOMICHEV, Stanislav (Google)

Co-author: BANDEWAR, Mahesh

Presenter: FOMICHEV, Stanislav (Google)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track