



Contribution ID: 296

Type: **not specified**

Cilium's BPF kernel datapath revamped

Tuesday, 13 September 2022 10:30 (30 minutes)

Since the early days of eBPF, Cilium's core building block for its datapath is tc BPF. With more adopters of eBPF in the Kubernetes landscape, there is growing risk from a user perspective that Pods orchestrating tc BPF programs might step on each other, leading to hard to debug problems.

We dive into a recently experienced incident, followed by our proposal of a revamped tc ingress/egress BPF datapath for the kernel which incorporates lessons learned from production use, lower overhead as a framework, and supporting BPF links for tc BPF programs in a native, seamless manner (that is, not conflicting with tc's object model). In particular the latter solve the program ownership and allow for better debugability through a common interface for BPF. We also discuss our integration approach into libbpf and bpftool, dive into the uapi extensions and next steps.

I agree to abide by the anti-harassment policy

Yes

Primary authors: BORKMANN, Daniel (Isovalent); ALEKSANDROV, Nikolay (Isovalent)

Presenter: BORKMANN, Daniel (Isovalent)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track