

A programmable Qdisc with eBPF

Wednesday, 26 August 2020 09:00 (45 minutes)

Today we have a few dozens of Qdisc's available in Linux kernel, offering various algorithms to schedule network packets. You can change the parameters of each Qdisc, but you can not change the core algorithm of a given Qdisc. A programmable Qdisc offers a way to customize your own scheduling algorithms without writing a Qdisc kernel module from scratch. With eBPF emerges across the Linux network stack, it is time to explore how to integrate eBPF with Qdisc's.

Unlike the existing eBPF TC filter and action, a programmable Qdisc is much more complicated, because we have to think about how to store skb's and what we can offer for users to program. More importantly, a hierarchical Qdisc is even harder while could offer more flexibility.

We will examine the latest eBPF functionalities and packet scheduler architecture, discuss those challenges with possible solutions for a programmable Qdisc with eBPF.

I agree to abide by the anti-harassment policy

I agree

Primary author: WANG, Cong (Bytedance)

Presenter: WANG, Cong (Bytedance)

Session Classification: Networking and BPF Summit

Track Classification: Networking & BPF Summit