

Use of eBPF in cpu scheduler

Monday, September 20, 2021 7:50 AM (40 minutes)

eBPF has been used extensively in performance profiling and monitoring. In this talk, I will describe a set of eBPF applications that help monitor and enhance cpu scheduling performances. These applications include:

- Profiling scheduling latencies. I will talk about an application of eBPF to collect scheduling latency stats.
- Profiling resource efficiency. For background, I will first introduce the scheduler feature core scheduling which is developed for mitigating L1TF cpu vulnerability. Then I will introduce the eBPF feature ksym which enables this application and describe how eBPF can help report the forced idle time, a type of cpu usage inefficiency caused by core scheduling.
- The third application of eBPF is to assist userspace scheduling. ghOSt is a framework open sourced by Google to enable general-purpose delegation of scheduling policy to userspace processes in a Linux environment. ghOSt uses BPF acceleration for policy actions that need to happen closer to scheduling edges. We use this to maximize CPU utilization (pick_next_task), minimize jitter (task_tick elision) and control tail latency (select_task_rq on wakeup). We are also experimenting with BPF to implement a scaled-down variant of the scheduling policy while upgrading the main userspace ghOSt agent.

I agree to abide by the anti-harassment policy

I agree

Primary authors: LUO, Hao (Google); RHODEN, Barret (Google)

Presenters: LUO, Hao (Google); RHODEN, Barret (Google)

Session Classification: BPF & Networking Summit

Track Classification: Networking & BPF Summit (Closed)