

Linux Plumbers Conference 2024



Contribution ID: 311

Type: **not specified**

Improving the Perf event subsystem after 15 years

Friday, 20 September 2024 10:00 (22 minutes)

15 years ago the perf events subsystem evolved from prior subsystems like OProfile, adding features like kernel profiling, tracing and system wide profiling. The subsystem continues to evolve, often driven by hardware, tracing and BPF developments. However, the purpose of this talk isn't to look back at the perf event subsystem, the perf tool and ongoing improvements, it is to take advantage of a roomful of perf subsystem contributors and think about what is coming next and how to keep the subsystem fit for purpose.

Improving a subsystem would be beyond the time limit of any microconference talk. We aim to introduce broad topics, and have topics introduced to us, that can then serve as the basis for breakouts.

Areas for discussion include:

- Uncore events and sampling, modern and future topologies, the growth of accelerators.
- Context switch performance in the era of 100s of performance monitoring units.
- Scalability of perf events on heavily consolidated server platforms.
- BPF and perf event implementation convergence.
- The kernel PMU abstraction, challenges in its specification, use and implementation.
- Getting richer information at low cost, such as owner of locks during contention, extra details in stack traces.
- Overhauling points of pain, such as the mmap and lack of munmap, mremap events.
- PMU precise event abstractions on non-Intel and simplifying tool support.
- Kernel vs tool separation and licensing challenges.
- Perf tool evolution, performance, dependency and distribution challenges.

Primary author: ROGERS, Ian (Google)

Presenter: ROGERS, Ian (Google)

Session Classification: Tracing / Perf events MC

Track Classification: Tracing / Perf events MC