Contribution ID: **243** Type: **not specified**

LTTng as a fast system call tracer

Wednesday, September 22, 2021 9:00 AM (25 minutes)

Upstreaming the LTTng kernel tracer [1] (originally created in 2005) into the Linux kernel has been a long-term goal of the LTTng project.

Today, various tracing technologies are available in the Linux kernel: instrumentation with tracepoints, kprobes, kretprobes, function tracing, performance counters through perf, as well as user-visible ABIs, namely Ftrace, Perf, and eBPF. There are however areas in which the LTTng kernel tracer has unique capabilities which other tracers lack.

Efficiently tracing system call entry/exit while fetching system call input/output parameters from user-space is a use-case the LTTng kernel tracer can cover, thanks to its ring buffer design which allows preemption.

Discuss the challenges and establish a roadmap towards upstreaming the pieces of the LTTng kernel tracer required to trace system calls into the Linux kernel.

[1] https://lttng.org

I agree to abide by the anti-harassment policy

I agree

Primary author: DESNOYERS, Mathieu (EfficiOS Inc.)

Presenter: DESNOYERS, Mathieu (EfficiOS Inc.)

Session Classification: Tracing MC

Track Classification: Tracing MC