



Contribution ID: 301

Type: **not specified**

Handling User Page Faults from Kernel Tracers

Friday, 20 September 2024 12:44 (22 minutes)

There are scenarios where tracer inputs come from user-space and are not paged in. Tracepoints invoked at system call entry immediately after an `exec(2)` system call are very likely to require page faults to access arguments located in the ELF data section.

This issue is not limited to system call instrumentation, it also affects instrumentation ABIs such as User Events.

Discuss our “Faultable tracepoints” proposal [1], which lays the groundwork required to allow kernel tracers to handle page faults, and the approaches which can be taken by Ftrace, Perf, eBPF and LTTng to handle those faults.

[1] “Faultable tracepoints” <https://lore.kernel.org/lkml/20240626185941.68420-1-mathieu.desnoyers@efficios.com/>

Primary author: DESNOYERS, Mathieu (EfficiOS Inc.)

Co-author: JEANSON, Michael (EfficiOS)

Presenters: DESNOYERS, Mathieu (EfficiOS Inc.); JEANSON, Michael (EfficiOS)

Session Classification: Tracing / Perf events MC

Track Classification: Tracing / Perf events MC