

Linux Plumbers Conference 2024



Contribution ID: 101

Type: **not specified**

What is missing to use `fexecve` (fd-based `execve`) to launch services?

Thursday 19 September 2024 11:00 (25 minutes)

Systemd does various checks and extensive preparation of the environment in which it'll spawn an executable. Currently, this is subject to a TOCTOU race, because we access the binary by path. We have code ready to use an fd for everything, but unfortunately the process that is spawned has a bogus `COMM` value (the fd number), which breaks `ps -C ...`. To make `fexecve` / `execveat` fully usable for userspace, we need to have a way to override `/proc/self/comm` for the executed process.

In the talk, I'll provide a short motivation why this feature is useful, what the current shortcomings are, and open the discussion to hopefully come up with an (simple) addition to the kernel API to fill in this missing bit.

Primary author: JĘDRZEJEWSKI-SZMEK, Zbigniew (Red Hat)

Presenter: JĘDRZEJEWSKI-SZMEK, Zbigniew (Red Hat)

Session Classification: Kernel <-> Userspace/Init/System Management boundaries and APIs MC

Track Classification: Kernel <-> Userspace/Init/System Management boundaries and APIs MC