



Contribution ID: 3

Type: **not specified**

Sched-Ext: The BPF extensible scheduler class

Overview

`sched_ext` is a Linux kernel feature which enables implementing host-wide, safe kernel thread schedulers in BPF, and dynamically loading them at runtime. `sched_ext` enables safe and rapid iterations of scheduler implementations, thus radically widening the scope of scheduling strategies that can be experimented with and deployed, even in massive and complex production environments.

`sched_ext` was first sent to the upstream list as an RFC patch set back in November 2022. Since then, the project has evolved a great deal, both technically, as well as in the significant growth of the community of `sched_ext` users and contributors.

Discussions

The following are the main topics that will be discussed during the MC:

- An update on the current status of `sched_ext`, initiating discussions about potential future directions, and the growth of the community.
- The componentization of the scheduler using a library-oriented approach to minimize code duplication among the scx schedulers.
- Defining a set of APIs for user-space scheduling, including Rust schedulers and Rust/C hybrid schedulers.
- Distro integration: addressing challenges and open issues related to shipping `sched_ext` and scx-based schedulers in Linux distributions.
- Practical applications of `sched_ext`: exploring how `sched_ext` can be utilized and its benefits in real production environments (e.g., SteamDeck, Meta).
- The use of `sched_ext` in virtualized and para-virtualized environments.

Logistics notes

- This is the first time we've proposed a `sched_ext` MC, so we have no results and accomplishments to discuss from prior meetings
- We've discussed the possibility of combining with either the `sched` or `sched-RT` microconferences with the authors of those MCs, and we all agree that there are already too many topics to combine the MCs

Primary authors: RIGHI, Andrea (NVIDIA); VERNET, David (Meta)

Co-authors: MIN, Changwoo (Igalia); Mr GHERDOVICH, Giovanni (SUSE); Mrs CHHAYA-SHAILESH, Himadri (Inria-Paris); Mr JUNG, Peter (CachyOS); Mr GÓRSKI, Piotr (CachyOS)

Presenters: RIGHI, Andrea (NVIDIA); MIN, Changwoo (Igalia); VERNET, David (Meta); Mr GHERDOVICH, Giovanni (SUSE); Mrs CHHAYA-SHAILESH, Himadri (Inria-Paris); Mr JUNG, Peter (CachyOS); Mr GÓRSKI, Piotr (CachyOS)

Track Classification: LPC Microconference Proposals