



Contribution ID: 17

Type: **not specified**

## **sched\_ext: The BPF extensible scheduler class MC**

### **CFP ends on September 30th (CLOSED)**

sched\_ext[1] is a Linux kernel feature which enables implementing safe task 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.

This MC is the space for the community to discuss the developments of sched\_ext, its impact on the community, and to outline future strategies aimed at improving the integration with the other Linux kernel subsystems.

Last year the sched\_ext MC proved highly productive in facilitating coordination with distribution maintainers, allowing us to clarify their requirements and ease potential maintenance burdens. This collaboration directly contributed to upstream changes, including patches such as [2].

Ideas of topics to be discussed include (but are not limited to):

- Use of BPF arenas for task/CPU context sharing between kernel, BPF, and user space
- Composable schedulers/scheduler libraries with BPF arenas
- Deadline server(s) for the SCHED\_EXT class
- Integration with other scheduling-related features (RCUs, proxy execution, PREEMPT\_RT, etc.)
- Potential integration with other Linux subsystems (e.g., Rust-for-Linux)
- Scheduling for gaming and latency-sensitive workloads
- User-space scheduling
- Tickless scheduling
- Tools and benchmarks to analyze and understand scheduler activities

While we already have a tentative schedule with existing talk proposals to cover the topics mentioned above, we are also planning to open a public CFP to accept additional topics to discuss. Time permitting, we are open to readjust the schedule to accommodate further discussions that are relevant to the Linux community.

[1] <https://github.com/sched-ext/scx>

[2] <https://lore.kernel.org/all/20240921193921.75594-1-andrea.righi@linux.dev/>

**Primary author:** RIGHI, Andrea (NVIDIA)

**Co-authors:** Mr FERNANDES, Joel (NVIDIA); Mr HODGES, Daniel (Meta); MIN, Changwoo (Igalia)

**Presenters:** Mr FERNANDES, Joel (NVIDIA); Mr HODGES, Daniel (Meta); MIN, Changwoo (Igalia); RIGHI, Andrea (NVIDIA)