

Linux Plumbers Conference 2024



Contribution ID: 327

Type: **not specified**

”Hey, psst, try this.” The underground culture around custom CPU schedulers.

Wednesday, 18 September 2024 12:10 (20 minutes)

Unbeknownst to the upstream Linux community, there exists a variety of alternative Linux CPU schedulers, circulating downstream for years.

The CachyOS Linux distribution, an Arch Linux derivative founded in May 2021, has painstakingly collected these patches into a nicely organized repository, bringing them to the fruition of a larger audience. The project itself derived its name from one of these schedulers, Cachy by Hamad Salim Al Marri; the Cachy scheduler has later been renamed to CacULE. Some of the most popular schedulers currently in circulation are:

- The **BitMap Queue CPU scheduler**, or BMQ, by Alfred Chen
- The **Priority and Deadline based Skiplist multiple queue CPU scheduler**, or PDS, by Alfred Chen
- The **Burst-Oriented Response Enhancer CPU scheduler**, or BORE, by Masahito Suzuki
- The **CacULE CPU scheduler**, inspired by the ULE scheduler from FreeBSD and formerly known as Cachy, by Hamad Salim Al Marri
- The **Enhanced CPU Handling Orchestrator scheduler**, or ECHO, by Hamad Salim Al Marri

The goals for this session are to **illustrate the algorithms employed** by these schedulers, bring them **to the attention of the wider upstream** kernel community, and **explore the opportunity to implement them using the sched-ext** framework to allow for an even wider circulation of their underpinning ideas.

Primary authors: GHERDOVICH, Giovanni (SUSE); JUNG, Peter; GÓRSKI, Piotr (CachyOS)

Presenters: GHERDOVICH, Giovanni (SUSE); JUNG, Peter; GÓRSKI, Piotr (CachyOS)

Session Classification: Sched-Ext: The BPF extensible scheduler class MC

Track Classification: Sched-Ext: The BPF extensible scheduler class MC