



Contribution ID: 112

Type: **not specified**

CPU controller on a single runqueue

Tuesday, 10 September 2019 17:00 (45 minutes)

The cgroups CPU controller in the Linux scheduler is implemented using hierarchical runqueues, which introduces a lot of complexity, and incurs a large overhead with frequently scheduling workloads. This presentation is about a new design for the cgroups CPU controller, which uses just one runqueue, and instead scales the vruntime by the inverse of the task priority. The goal is to make people familiar with the new design, so they know what is going on, and do not need to spend a month examining kernel/sched/fair.c to figure things out.

I agree to abide by the anti-harassment policy

Yes

Primary author: VAN RIEL, Rik (Facebook)

Presenter: VAN RIEL, Rik (Facebook)

Session Classification: LPC Refereed Track