



Contribution ID: 61

Type: **not specified**

## Limit the idle CPU search depth and use CPU filter during task wake up

*Tuesday, 13 September 2022 17:10 (25 minutes)*

When a task is woken up in the last level cache (LLC) domain, the scheduler tries to find an idle CPU for the task. But when the LLC domain is fully busy, the search for an idle CPU may be in vain, adding long latency to the task wakeup and yet does not lead to an idle CPU. The latency gets worse when the number of CPUs in the LLC increases, which will be the case for future platforms.

During LPC2021 there was a discussion on how to find the idle CPU effectively. This talk is an extended discussion on that. We will illustrate how we encountered the issue and how to debug this issue on a high core count system. We'll introduce the proposal that leverages the `util_avg` of the LLC to decide how much effort is spent to scan for idle CPUs. We'll also present a proposal to filter out the busy CPUs so as to speed up the scan. We'll share the current test data using the mechanism. We hope to get feedback on advice/feedback on tuning the scan policy and making it viable for upstreaming.

### I agree to abide by the anti-harassment policy

Yes

**Primary authors:** YU, Chen; WU, Yun (Bytedance)

**Presenters:** YU, Chen; WU, Yun (Bytedance)

**Session Classification:** Real-time and Scheduling MC

**Track Classification:** LPC Microconference: Real-time and Scheduling MC