



Contribution ID: 178

Type: **not specified**

Expressing per-task/per-cgroup performance hints

Thursday, November 15, 2018 9:25 AM (25 minutes)

The Linux scheduler is able to drive frequency selection, when the schedutil cpufreq's governor is in use, based on task utilization aggregated at CPU level. The CPU utilization is then used to select the frequency which better fits the task's generated workload. The current translation of utilization values into a frequency selection is pretty simple: we just go to max for RT tasks or to the minimum frequency which can accommodate the utilization of DL+FAIR tasks.

While this simple mechanism is good enough for DL tasks, for RT and FAIR tasks we can aim at some better frequency driving which can take into consideration hints coming from user-space.

Utilization clamping is a mechanism which allows to filter the utilization generated by RT and FAIR tasks within a range defined from user-space, either for a task or for task groups. The clamped utilization requirements of RUNNABLE tasks are aggregated at CPU level and used to enforce its minimum and/or maximum frequency.

This session is meant to give an update on the most recent LKML posting of the utilization clamping patchset and to open a discussion on how to better progress this proposal.

I agree to abide by the anti-harassment policy

Primary authors: RASMUSSEN, Morten (Arm); BELLASI, Patrick (Arm Ltd.)

Presenters: RASMUSSEN, Morten (Arm); BELLASI, Patrick (Arm Ltd.)

Session Classification: Power Management and Energy-awareness MC