



Contribution ID: 203

Type: **not specified**

Scheduler task accounting for cgroups

Wednesday, 14 November 2018 09:00 (15 minutes)

Cgroup accounting has significant overhead due to the need to constantly loop over all cpus to update statistics of cpu usages and blocked averages. We have seen that on 4 socket Haswell, database benchmarks like TPCC have 8% performance regression at the time of Haswell and 4.4 kernel when it is run under cgroup. On recent Cannon Lake platform using latest PCIE SSDs and 4.18 kernel, the regression in the scheduler has gotten worse to 12%. We will highlight the bottlenecks in the scheduler with detailed profiles of the hot path. We'll like to explore possible avenues to improve cgroup accounting.

I agree to abide by the anti-harassment policy

Presenter: CHEN, Tim

Session Classification: Performance and Scalability MC