



Contribution ID: 214

Type: **not specified**

Intel Low Power Mode Daemon on Hybrid CPUs

Tuesday, November 14, 2023 5:15 PM (15 minutes)

Hardware vendors increasingly employ hybrid-CPU processors. Hybrid-CPU processors mix multiple CPU types within the same processor, to satisfy previously conflicting tradeoffs in power, performance, and cost.

Some hybrid-CPU processors benefit from a “low-power mode”, a non-linear power savings resulting from all work being restricted to a “low-power CPUs”, while no work runs on the “high-performance CPUs”. The power benefit of low-power mode depends on how long it can be sustained, while the performance impact depends on how quickly low-power mode can exit.

Introducing, `intel_lpm` (Intel Low-Power Mode Daemon), which monitors system activity, hardware hints and user preferences to decide when to enter and exit low-power mode. `intel_lpm` is successful when it recognizes near idle workloads that can fit on the low-power CPU. Conversely, `intel_lpm` must recognize when to exit low-power mode to avoid impacting performance.

`intel_lpm`'s power goal is to get low-power mode savings to approach that of using CPU hotplug offline. `intel_lpm` works best with the cgroup v2 CPU controller and its “isolated”partition feature. However, there are multiple sources of noise on the quiesced CPUs that we must handle, including RCU, idle load balancing, and recurring timers –all of which can defeat the power saving impact of low-power mode.

The goal of this presentation is to describe challenges, workable solutions and get community feedback on user space intelligence.

Note that this topic has been submitted to the main track already, but in case it is not accepted there, I'd still like to present it in the Power Management and Thermal Control MC.

Primary authors: ZHANG, Rui; PANDRUVADA, Srinivas

Presenters: ZHANG, Rui; PANDRUVADA, Srinivas

Session Classification: Power Management and Thermal Control MC

Track Classification: LPC Microconference: Power Management and Thermal Control MC