



Contribution ID: 44

Type: **not specified**

Core scheduling

Monday, 9 September 2019 15:00 (45 minutes)

There have been two different approaches proposed on the LKML over the past year on core scheduling. One was the coscheduling approach by Jan Schön herr, originally posted at <https://lkml.org/lkml/2018/9/7/1521> and the next version posted at <https://lkml.org/lkml/2018/10/19/859>

Upstream chose a different route and decided to modify CFS, and only do “core-scheduling”. Vineeth picked up the patches from Peter Zijlstra. This is a discussion on how we can further that work, especially when there are security implications such as L1TF and MDS, which make important this work to go upstream.

Aubrey Li will talk about Core scheduling: Fixing when fast instructions go slow

Keeping system utilization high is important both to keep costs down and to keep energy efficiency up. That often means tightly packing compute jobs and using the latest processor features. However, these approaches can be at odds when a new processor feature like AVX512 is used. The performance of latency critical jobs can be reduced by 10% if co-located with deep learning training jobs. These jobs use AVX512 instructions to accelerate wide vector operations. Whenever a core executes AVX512 instructions, the core automatically reduces its frequency. This can lead to a significant overall performance loss for a non-AVX512 job on the same core. In this presentation, we will discuss how to preserve performance while still allowing AVX512-based acceleration.

AVX512 task detection

- From user space, PMU events can be used but it's expensive.
- In the kernel, I proposed to expose process AVX512 usage elapsed time as a heuristic hint.
- Discuss an interface for tasks in cgroup.

AVX512 task isolation

- Discuss kernel space solution, if the recent proposal core scheduling can be leveraged for isolation.
- Discuss user space solution, if user space job scheduler is better than kernel scheduler

I agree to abide by the anti-harassment policy

Yes

I confirm that I am already registered for LPC 2019

Primary authors: Mr LI, Aubrey; SCHÖNHERR, Jan; REIS, Hugo; REMANAN PILLAI, Vineeth

Presenters: Mr LI, Aubrey; SCHÖNHERR, Jan; REIS, Hugo; REMANAN PILLAI, Vineeth

Session Classification: Scheduler MC