



Contribution ID: 139

Type: **not specified**

How do we make a Steamdeck scheduler work on large servers

With the proliferations of many sched_ext schedulers, including ones that caters for very specific workloads within Meta. There exists a need for a “default” fleet scheduler that “just works” for a wide range of hardware and use cases. SCX_LAVD is one such candidate as one of the more mature sched_ext schedulers out there with various heuristics to favor latency critical threads.

The talk will focus on various challenges and strategies in bringing in SCX_LAVD and trying to run it on large production workloads and large topologies:

1. How do we handle large and varied topologies and cache hierarchies that exists in the fleet to take optimal advantage of the hardware?
2. How do we tune LAVD such that it performs well in throughput bound use cases without sacrificing its latency advantages?

Primary author: DAI, David

Presenter: DAI, David

Session Classification: sched_ext: The BPF extensible scheduler class MC

Track Classification: sched_ext: The BPF extensible scheduler class MC