

Linux Plumbers Conference 2024



Contribution ID: 99

Type: **not specified**

Paravirt scheduling with eBPF

Friday, 20 September 2024 15:00 (45 minutes)

Double scheduling is a concern with virtualization hosts where the host schedules vcpus without knowing what's run by the vcpu and guest schedules tasks without knowing where the vcpu is physically running. This causes issues related to latencies, power consumption, resource utilization etc. An ideal solution would be to have a cooperative scheduling framework where the guest and host shares scheduling related information and makes an educated scheduling decision to optimally handle the workloads.

The initial approach was to have all the logic in KVM and this was discussed in LPC 2023 and upstream. KVM maintainers were not favourable with this idea as it puts a policies and decisions related to scheduling in KVM. And the resulting paravirt protocol is not generic. So the consensus is to have the policies and decisions as a separate entity which could be generic and implemented outside of kvm(A bpf program or kernel module).

We are working on the next revision where the guest and host implements the policies as a bpf program. Host exposes a virtual device through which initial handshake and negotiations happen and bpf program takes care of the guest/host communication, policies scheduling decision etc.

This talk is about our journey in designing a paravirt scheduling protocol, use of bpf while trying to maintain the generic nature of the protocol, advantages and challenges of bpf in this project etc. We would also be discussing the future of the project to include other use cases other than minimizing latencies..

Primary author: Mr REMANAN PILLAI, Vineeth (Google)

Co-author: FERNANDES, Joel

Presenters: FERNANDES, Joel; Mr REMANAN PILLAI, Vineeth (Google)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track