Linux Plumbers Conference 2023



Contribution ID: 132

Type: not specified

Dynamic vCPU priority boosting in KVM for latency sensitive workloads

Tuesday, 14 November 2023 17:15 (45 minutes)

KVM, a virtualization technology in Linux, delegates memory and virtual CPU execution management of virtual machines to the Linux kernel. This has both advantages and disadvantages. One disadvantage is that it can lead to latency issues in time-sensitive workloads in the VM (such as audio and video). This is because KVM creates one task per vCPU for the VM and then delegates the scheduling of these tasks to the kernel. The kernel does not differentiate between vCPU tasks and other tasks, and it does not have insight into the workloads executed by the vCPUs. As a result, scheduling latencies for the vCPUs in a busy system can translate to apparent latency issues within the VM.

This talk discusses an effort to minimize the latencies in VMs by having a framework of communication between the host and guest. The basic idea is to have a communication mechanism between the host and guest such that the host can make an educated decision on the priority of vCPU threads. The guest communicates its scheduling needs to the host, and the host can boost the priority of vCPU threads as needed. The host does have some insight about the guest (eg: interrupt injection) and can boost the vCPU explicitly. This is communicated to the guest, and the guest can determine when to boost/unboost.

The talk details the design, implementation, and performance details of a x86_64 prototype that implements this idea using shared memory and hypercalls. We would also be discussing technical issues/challenges during the design and implementation of the prototype.

Primary authors: FERNANDES, Joel (Google); Mr REMANAN PILLAI, Vineeth (Google)
Presenters: FERNANDES, Joel (Google); Mr REMANAN PILLAI, Vineeth (Google)
Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track