



Contribution ID: 165

Type: **not specified**

VM-CPUFreq for x86: Scaling the guest frequency for performance and power savings

Tuesday, November 14, 2023 4:00 PM (15 minutes)

With the increasing size of virtual machines (VMs) and the growing deployment of workloads within VMs, power management has become a crucial aspect to optimize performance and energy efficiency. However, the absence of frequency scaling within VMs and the lack of workload utilization visibility for the hypervisor can result in suboptimal performance and power management. To address this issue, Google proposed patches to improve dynamic voltage and frequency scaling (DVFS) for VMs on ARM systems [1]. By providing the hypervisor with VMs' vCPU utilization data, it can make informed decisions regarding frequency scaling, leading to improved performance and power efficiency.

In this talk, we discuss what is needed to make this work on x86 servers with AMD EPYC as an example.

We discuss

- The metrics that need to be communicated from the guest to the hypervisor and vice-versa.
- The potential interfaces through which such guest-host communication can be achieved with low-overhead.
- Experimental results across VMs of different sizes running different results in order to demonstrate the effectiveness of the vCPU utilization information being communicated to the hypervisor and the impact that it has on the performance and power-consumption of the workloads.

[1]: [RFC PATCH 0/6] Improve VM DVFS and task placement behavior: <https://lore.kernel.org/lkml/20230330224348.1006691-1-davidai@google.com/>

Primary author: Mr KARNY, Wyes (AMD)

Co-author: SHENOY, Gautham R (AMD Inc.)

Presenter: Mr KARNY, Wyes (AMD)

Session Classification: Power Management and Thermal Control MC

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