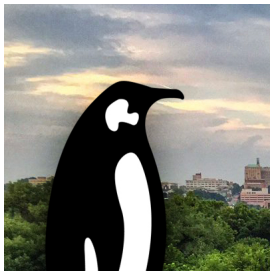

Contribution ID: **183** Type: **not specified**

Improve Xeon IRQ throughput with posted interrupt

Wednesday, 15 November 2023 09:30 (30 minutes)

Server SoCs today offer more PCIe lanes as well as the ability to stack more IO devices on a single port. Out of the box, devices such as high-speed NVMe drives can generate a significant number of interrupts at high frequencies. Due to microarchitecture choice and PCIe strong ordering requirements, limited IRQ throughput on Intel Xeon has also become a limiting factor for DMA throughput. IOPS can drop more than 50% as evidenced by the FIO test on multiple NVMe disks. This talk describes a scheme and RFC patch that optimize IRQ throughput by enabling posted interrupts on bare metal (beyond the VM usage today). The result is that much of the performance loss can be recovered without any new hardware or driver changes.

Primary author: PAN, Jacob

Presenter: PAN, Jacob

Session Classification: VFIO/IOMMU/PCI MC

Track Classification: LPC Microconference: VFIO/IOMMU/PCI MC