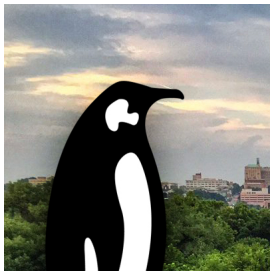

Contribution ID: **269** Type: **not specified**

RISC-V irqbypass with KVM

Monday, 13 November 2023 12:45 (20 minutes)

KVM and VFIO provide an architecture-neutral irqbypass framework, but its enablement requires an implementation of an architecture-specific function, kvm_arch_irq_bypass_add_producer(). The RISC-V AIA and IOMMU specifications provide novel support for guest interrupt delivery (most notably MRIFs), which must be considered for RISC-V KVM's irqbypass implementation. We have an initial proposal which includes the RISC-V IOMMU driver implementing an IRQ domain in order to provide irq_set_vcpu_affinity(). This discussion is seeking feedback on that approach. Additionally, the RISC-V IOMMU will send notice MSIs when guest vIMSICs are backed by MRIFs, requiring a policy to select where the notice MSIs are delivered. This means we need to define new uAPI in order to involve the user. Feedback on the uAPI proposals would also be welcome. Also, we acknowledge that irqbypass performance will differ for guests with assigned interrupt files vs. those with MRIFs and we would like to discuss how best to modify or extend accounting in order to improve accuracy of measurements. Finally, the PoC is just getting started, by the time Plumbers meets, there should be enough done to have made other design decisions which could be discussed.

Primary author: JONES, Andrew (Ventana Micro Systems)

Presenter: JONES, Andrew (Ventana Micro Systems)

Session Classification: RISC-V MC

Track Classification: LPC Microconference: RISC-V MC