

Criteria of using VFIO mdev (vs. userspace DMA)

Wednesday 26 August 2020 07:00 (20 minutes)

VFIO mdev provides a framework for subdevice assignment and reuses existing VFIO uAPI to handle common passthrough-related requirements. However, subdevice (e.g. ADI defined in Intel Scalable IOV) might not be a PCI endpoint (e.g. just a work queue), thus requires some degree of emulation/mediation in kernel to fit into VFIO device API. Then there is a concern on putting emulation in kernel and how to judge abuse of mdev framework by simply using it as an easy path to hook into virtualization stack. An associated open is about differentiating mdev from userspace DMA framework (such as uacce), and whether building passthrough features on top of userspace DMA framework is a better choice than using mdev.

I agree to abide by the anti-harassment policy

I agree

Primary author: TIAN, Kevin (Intel)

Presenter: RAJ, Ashok

Session Classification: VFIO/IOMMU/PCI MC

Track Classification: VFIO/IOMMU/PCI MC