

Passthrough of VMD subdevices

Wednesday, 26 August 2020 10:25 (15 minutes)

The Intel Volume Management Device (VMD) behaves similar to a PCI-to-PCI bridge that changes the subdevice's requester ID to VMD's. VMD also remaps subdevice MSI/X into its own MSI/X mapping table. Because of the requester ID factor, the VMD device and subdevice domain fall under a single IOMMU group.

VMD is being integrated more and more into Intel chipsets and the desire to assign individual subdevices is only going to become more of an outstanding problem as time goes on. The existing model of assignment of the whole IOMMU group to a VM is problematic to VMD subdevices, as well as any expectation surrounding interrupt remapping.

VFIO/IOMMU may need an (unsafe) DMA remapping provider-consumer relationship to assigning individual subdevice DMA contexts. To handle MSI/X, the guest may need to avoid using it in the first place, or have VMD in the host deliver the interrupts.

I agree to abide by the anti-harassment policy

I agree

Primary author: DERRICK, Jonathan

Presenter: DERRICK, Jonathan

Session Classification: VFIO/IOMMU/PCI MC

Track Classification: VFIO/IOMMU/PCI MC