

Linux Plumbers Conference 2024



Contribution ID: 140

Type: **not specified**

PPC64 - Bridging the pSeries and PowerNV islands for VFIO and IOMMUFD

Wednesday, 18 September 2024 12:00 (10 minutes)

PPC64 implementation of VFIO is spread across two vastly different machine types (pSeries, PowerNV) trying to share a lot of common code driven by PPC specific SPAPR IOMMU API.

The support of the PCI device assignment on these sub arch's have gone through many cycles of breakages and fixes with ongoing efforts to add support for IOMMUFD, which PPC64 is yet to catch up to. Enhancements[1] to bring consistency to the SPAPR IOMMU management for both the machine types exposed quite a few challenges due to underlying arch differences and the way VFIO and IOMMUFD models containers. One of the key departure from earlier attempts at VFIO support was to move away from "borrowing" the DMA window which limited the customization of the window size.

The talk aims the below,

- Overview of PPC64 IOMMU architecture
- The SPAPR TCE v1, v2 approaches
- Challenges and possible approaches to support IOMMUFD for KVM guests on pSeries and PowerNV.
- Taking community feedback for upcoming patches for IOMMUFD.

References:

[1] - <https://lore.kernel.org/linuxppc-dev/171810893836.1721.2640631616827396553.stgit@linux.ibm.com/>

Primary author: BHAT, Shivaprasad G (IBM)

Co-author: Mr N, Narayana Murty (IBM)

Presenter: BHAT, Shivaprasad G (IBM)

Session Classification: VFIO/IOMMU/PCI MC

Track Classification: VFIO/IOMMU/PCI MC