

Linux Plumbers Conference 2025



Contribution ID: 35

Type: **not specified**

VFIO/IOMMU/PCI MC

CFP ends on September 30th (CLOSED)

The PCI interconnect specification, the devices that implement it, and the system IOMMUs that provide memory and access control to them are nowadays a de-facto standard for connecting high-speed components, incorporating more and more features such as:

- Address Translation Service (ATS)/Page Request Interface (PRI)
- Single-root I/O Virtualization (SR-IOV)/Process Address Space ID (PASID)
- Shared Virtual Addressing (SVA)
- Remote Direct Memory Access (RDMA)
- Peer-to-Peer DMA (P2PDMA)
- Cache Coherent Interconnect for Accelerators (CCIX)
- Compute Express Link (CXL)/Data Object Exchange (DOE)
- Component Measurement and Authentication (CMA)
- Integrity and Data Encryption (IDE)
- Security Protocol and Data Model (SPDM)

These features are aimed at high-performance systems, server and desktop computing, embedded and SoC platforms, virtualisation, and ubiquitous IoT devices.

The kernel code that enables these new system features focuses on coordination between the PCI devices, the IOMMUs they are connected to, and the VFIO layer used to manage them (for userspace access and device passthrough) with related kernel interfaces and userspace APIs to be designed in-sync and in a clean way for all three sub-systems.

The VFIO/IOMMU/PCI MC focuses on the kernel code that enables these new system features, often requiring coordination between the VFIO, IOMMU and PCI subsystems.

Following the success of LPC 2017, 2019, 2020, 2021, 2022, 2023 and 2024 VFIO/IOMMU/PCI MC, the Linux Plumbers Conference 2024 VFIO/IOMMU/PCI track will focus on promoting discussions on the PCI core and current kernel patches aimed at VFIO/IOMMU/PCI subsystems. Specific sessions will focus on discussions that require coordination between the three subsystems.

See the following video recordings from 2024: [LPC 2024 - VFIO/IOMMU/PCI MC](#).

Older recordings are available through the official YouTube channel of the Linux Plumbers Conference and the archived LPC 2017 VFIO/IOMMU/PCI MC web page at Linux Plumbers Conference 2017 the audio recordings from the MC track and links to presentation materials are available.

The tentative schedule will provide an update on the current state of VFIO/IOMMU/PCI kernel subsystems, followed by a discussion of current issues related to the proposed topics.

The following was a result of last year's successful Linux Plumbers MC:

- The first version of work on solving the complex and pressing issue of secure device assignment that spans across the PCI, IOMMU, and CXL subsystems has been completed, and a series of patches has been sent to view and spark more discussion and debate about how to solve this challenging problem.

Tentative topics that are under consideration for this year include (but are not limited to):

- PCI
 - Cache Coherent Interconnect for Accelerators (CCIX)/Compute Express Link (CXL) expansion memory and accelerators management
 - Data Object Exchange (DOE)
 - Integrity and Data Encryption (IDE)
 - Component Measurement and Authentication (CMA)
 - Security Protocol and Data Model (SPDM)
 - I/O Address Space ID Allocator (IOASID)
 - INTX/MSI IRQ domain consolidation
 - Gen-Z interconnect fabric
 - PCI error handling and management, e.g., Advanced Error Reporting (AER), Downstream Port Containment (DPC), ACPI Platform Error Interface (APEI) and Error Disconnect Recovery (EDR)
 - Power management and devices supporting Active-state Power Management (ASPM)
 - Peer-to-Peer DMA (P2PDMA)
 - Resources claiming/assignment consolidation
 - DMA ownership models
 - Thunderbolt, DMA, RDMA and USB4 security
- VFIO
 - I/O Page Fault (IOPF) for passthrough devices
 - Shared Virtual Addressing (SVA) interface
 - Single-root I/O Virtualization(SRIOV)/Process Address Space ID (PASID) integration
 - PASID in SRIOV virtual functions
 - TDISP/TSM Device assignment/sub-assignment
- IOMMU
 - /dev/iommufd development
 - IOMMU virtualisation
 - IOMMU drivers SVA interface
 - DMA-API layer interactions and the move towards generic dma-ops for IOMMU drivers
 - Possible IOMMU core changes (e.g., better integration with the device-driver core, etc.)

If you are interested in participating in this MC and have topics to propose, please use the Call for Proposals (CfP) process. More topics might be added based on CfP for this MC.

Otherwise, join us in discussing how to help Linux keep up with the new features added to the PCI interconnect specification. We hope to see you there!

Key Attendees:

- Alex Williamson
- Benjamin Herrenschmidt
- Bjorn Helgaas
- Dan Williams
- Ilpo Järvinen
- Jacob Pan
- James Gowans
- Jason Gunthorpe
- Jonathan Cameron
- Jörg Rödel
- Kevin Tian
- Lorenzo Pieralisi
- Lu Baolu
- Manivannan Sadhasivam

Contacts:

- Alex Williamson (alex.williamson@redhat.com)
- Bjorn Helgaas (helgaas@kernel.org)
- Jörg Roedel (joro@8bytes.org)
- Lorenzo Pieralisi (lpieralisi@kernel.org)
- Krzysztof Wilczyński (kwilczynski@kernel.org)

Primary author: WILCZYŃSKI, Krzysztof

Co-authors: WILLIAMSON, Alex; HELGAAS, Bjorn (Google); ROEDEL, Joerg (AMD); PIERALISI, Lorenzo

Presenters: WILLIAMSON, Alex; HELGAAS, Bjorn (Google); ROEDEL, Joerg (AMD); WILCZYŃSKI, Krzysztof; PIERALISI, Lorenzo