Enabling PCI Features via Device Driver vs. IOMMU implicitly

- Ashok Raj & Lu Baolu

MANY THANKS TO OUR SPONSORS

Diamond

facebook

Platinum

NetApp

IBM

Western Digital.

Gold

arm

Microsoft Azure

Catchboxes

COLLABORA

Silver

AMD

Red Hat

Event Services

THE LINUX FOUNDATION

ANTI-HARASSMENT POLICY

The Linux Plumbers conference is a working conference intended for professional networking and collaboration in the Linux community. Participants are expected to behave according to professional standards and in accordance with their employer's policies on appropriate workplace behaviour.

While at the Linux Plumbers Conference or related social networking opportunities, attendees should not engage in discriminatory or offensive speech or actions regarding gender, sexuality, race, or religion. Speakers should be especially aware of these concerns.

We do not condone any statements by speakers contrary to these standards, and we reserve the right to deny entrance to any individual.

Please bring any concerns to to the immediate attention of a member of the Linux Plumbers Conference planning committee. We thank our attendees for their help in keeping the Linux Plumbers Conference professional, welcoming, and friendly.

This policy is derived from the Linux Foundation's Code of Conduct which applies to all Linux Foundation events.

HOUSEKEEPING

Thank you for joining the LPC Sessions Leads Training

Please keep your microphone and camera muted when not actively participating

Ask your questions at any time. Please chime in

Enable your webcam before asking questions

Device Driver vs. IOMMU Enabling PCI capabilities

- Some PCI capabilities are automatically enabled by IOMMU sub-system during attach_device()
  - PASID is enabled before ATS
  - Driver responsible for pci_restore_state() after an flr()
- Moving to Device Driver enable preferred for the following reasons.
  - Some devices have issues with certain features. Requires some quirks which otherwise could simply be driver managed.
  - Some devices like SRIOV requires knowledge about IOMMU.
    - Features like PASID/PRI/ATS require IOMMU. We have no such discovery for device drivers.
    - Devices might require ATS for functionality vs optional for performance only.

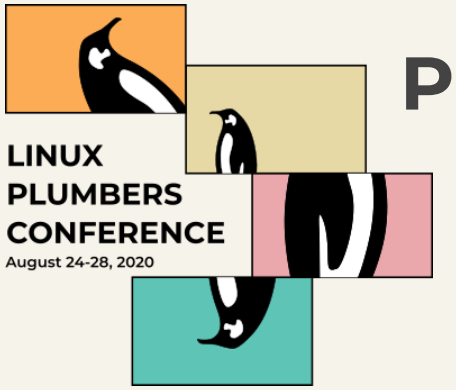

Proposed Extensions

- Check if IOMMU is present and enabled
  - `iommu_cap_supported(dev, features);`
  - Features – PASID, PRI, ATS
- Enable features through iommu
  - `iommu_enable_caps(dev, features);`
  - Will help some shared code for PF/VF checks for instance.
- Drivers continue to use `pci_restore_state()` during resume or flr flows.

WE ARE ON A BREAK

Let us know you are back by clicking the “Ready” button on the lower-right corner