



Contribution ID: 121

Type: **not specified**

Linux plumbing of CXL error reporting

Wednesday, September 14, 2022 5:35 PM (25 minutes)

With the introduction of CXL Type 3 Memory Devices a system may contain multiple different memory controllers to support and provide volatile memory. To add support of all those, generic and architectural specific implementations across different subsystems (CXL, PCI, ACPI, MCA, EDAC, etc.) are involved. CXL introduces following errors:

- CXL link and protocol errors and
- CXL type 3 memory device errors.

So now a broad variety of error types and sources must be handled additionally compared to what typically exists for mem controllers or pci devices.

The CXL kernel interface also provides an ioctl user interface to retrieve error events. And, there are a couple of new tools to control all that.

All this raises new questions on how to share, rework and plumb existing subsystems to make CXL work. And, should new APIs and tools for collecting errors be introduced what would be more suitable?

Topics of the discussion include:

- Should there be the same look and feel as with dram mem controllers when reporting cxl memory errors?
- Should errors be handled in userspace? How can access be serialized, how to support multiple users? Which tools should be the focus on?
- How can we reuse PCIe AER and RCEC implementations of the pci stack for cxl? Should we join pci and cxl, in particular maintaining a struct pci_dev for each cxl dev?
- What are the challenges in supporting multiple mem controllers in the system?
- Is there a sufficient need to integrate cxl error reporting into the EDAC subsystem?

A very brief overview of CXL error reporting and involved Linux subsystems will be presented to further discuss and find answers for above questions.

I agree to abide by the anti-harassment policy

Yes

Primary authors: Mr RICHTER, Robert (Advanced Micro Devices); Mr GHANNAM, Yazen (Advanced Micro Devices)

Presenters: Mr RICHTER, Robert (Advanced Micro Devices); Mr GHANNAM, Yazen (Advanced Micro Devices)

Session Classification: Compute Express Link MC

Track Classification: LPC Microconference: Compute Express Link MC