



Contribution ID: 202

Type: **not specified**

## Adding RAS Support for CXL Port Devices

*Monday, 13 November 2023 16:20 (25 minutes)*

Title: Adding RAS Support for CXL Port Devices

Authors: Robert Richter, Terry Bowman

Abstract:

CXL RAS error registers are present in all CXL HW entities[1]. The CXL driver currently only implements RAS OS support for endpoints. Support for the following CXL entities must be added as well: CXL RCH down ports, root ports, USP/DSP switch ports, and CXL host bridges (referred hereafter as ports). The existing endpoint driver's handlers do not apply well for reuse by ports. CXL ports are presented as separate PCIe devices, implement the AER extended capabilities, and as a result bind to the AER service port driver. It should be noted CXL ports are not associated with a driver aware of CXL or RAS. There are also cases for port RAS error handling when an endpoint device may not be present. For these reasons, the current CXL error handler must be extended and or refactored for port RAS handling.

The challenge is providing common CXL RAS handling without overcommitting to a specific PCIe device type. Adding a CXL port service driver is one potential solution. A CXL port service driver could include the RAS handling and logging used by all the CXL HW entities. The CXL port driver would interact with the AER port service driver. This presentation will include the constraints and requirements related to such a CXL port service driver.

This includes binding the service driver to a CXL port device, registration of error handlers, common and device specific required handling. A challenge to consider is refactoring existing CXL endpoint RAS handling to make it common.

Will provide slides after abstract acceptance. Plan to define current CXL RAS state, issues overcome, and moving forward with CXL2.0 and port error handling.

[1] CXL 3.0, 8.2.4 - CXL.cache and CXL.mem Registers

**Primary author:** BOWMAN, Terry (AMD)

**Co-author:** RICHTER, Robert (AMD Inc.)

**Presenter:** BOWMAN, Terry (AMD)

**Session Classification:** Compute Express Link MC

**Track Classification:** LPC Microconference: Compute Express Link MC