



Linux plumbing of CXL error reporting

Handling of CXL memory device errors in the kernel
Compute Express Link MC

Robert Richter
Yazen Ghannam

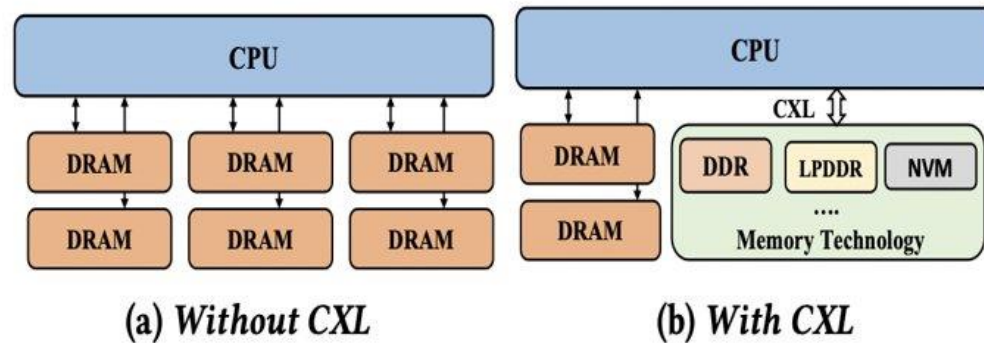
Linux Plumbers Conference 2022
Dublin, Sep 14, 2022

AMD 
together we advance_

Agenda


-
1. CXL RAS high level overview
 2. CXL RAS considerations
 3. CXL Error Handling
 4. CXL VH vs. RCD mode
 5. CXL kernel and userland
 6. Summary
 7. Discussion

CXL Overview - CXL Memory Devices (Type 3)



- Enables flexibility to add a variety of memory in addition to host-attached memory (DRAM).
- CXL Memory can have different characteristics compared to host-attached memory, e.g. considering system cost, capacity, power, bandwidth.
- Device can provide Host-managed Device Memory (HDM), that memory is located on a CXL device (Device-attached Memory).
- HDM is mapped to the system coherent address space and accessible through load/store semantics.

CXL Overview - Memory Devices and RAS

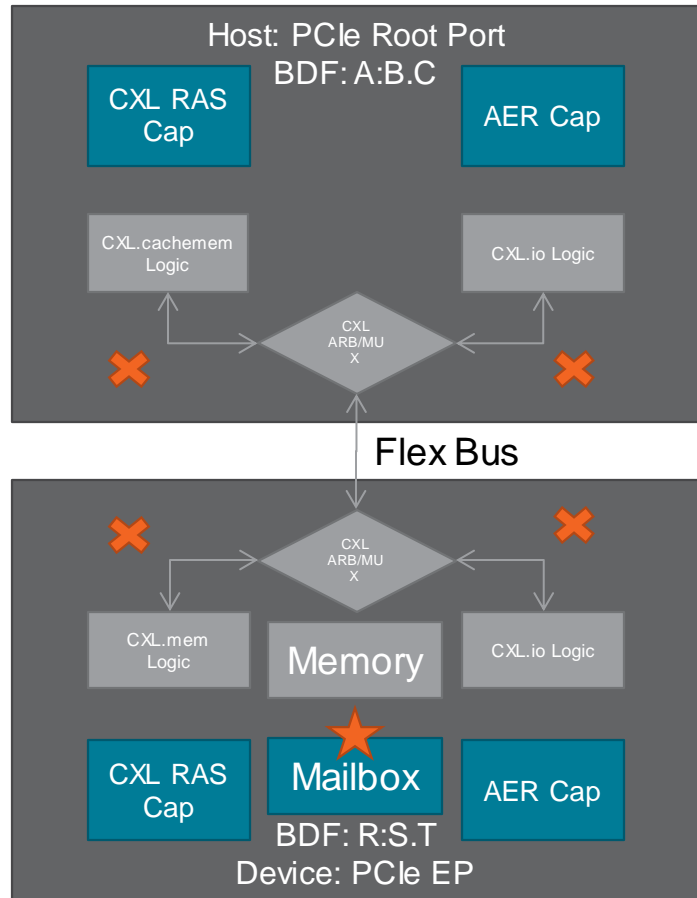


- Different to native memory controllers: Device and host have components to control CXL mem.
- There is no longer a unique memory controller in the system that controls all the memory.
- A device can report errors and error details.
- CXL protocol must be used to propagate errors from the device to the system.
- Many components are involved where errors can happen.
 - Errors are handled differently compared to host-attached memory.
 - More error locations with different error types and different error flows exist compared to non-CXL systems.

CXL RAS - More considerations

- Errors can be handled either by firmware (FW FIRST) or the OS (OS FIRST).
→ Even more ways of reporting errors.
- CXL VH mode vs. RCD mode
→ Different topology causes a different error flow.
- Same look and feel as for native memory?
→ Existing subsystems (e.g. EDAC) might not fit well.
→ Is there a need to implement new or extend other existing tools (e.g. ndctl)?
- Error injection
→ Various components and tools are involved.

CXL RAS - Where errors can happen in a CXL memory device



CXL VH mode

- The three CXL protocols (io, mem, and cache) form two groups: CXL.io and CXL.cachemem (or CXL.cache + CXL.mem). These two groups of protocols are multiplexed across the same physical “Flex Bus” link. See CXL spec section “Flex Bus Layering Overview”.
- There are four points where protocol/link errors can be detected marked with an ✖.
- Component Events is a collection of other device-related errors, e.g. memory ECC, thermal, etc. These are reported through the Mailbox (Device Command Interface and Event Records) marked with a ★.

CXL Error Handling

Error Type	Error Reporting	Error Handling
Poison	Reported synchronously with the corrupted data and traverses to the requestor.	Host processor-specific. Mailbox command used for Poison Management.
Viral	Reported via CXL Link	Host processor-specific.
CXL Protocol Errors	Reported via RP's PCIe AER. AER "Internal Error" (UIE/CIE) indicates information is also logged the RP's CXL RAS caps (CXL RP/DSP/USP).	PCIe Root Port collects the error information and signals the FW/OS. Handling is analogous to PCIe RP handling but also reading CXL RAS caps if there are internal errors.
CXL component events	Reported via CXL Device/Component Command Interface (mailbox, etc.) using CXL-defined record formats.	Device may signal FW using Vendor Defined Messages (VDM) or OS using interrupts (MSI/MSI-X). FW/OS issues mailbox commands to device to gather records.

CXL Protocol Errors

FW FIRST

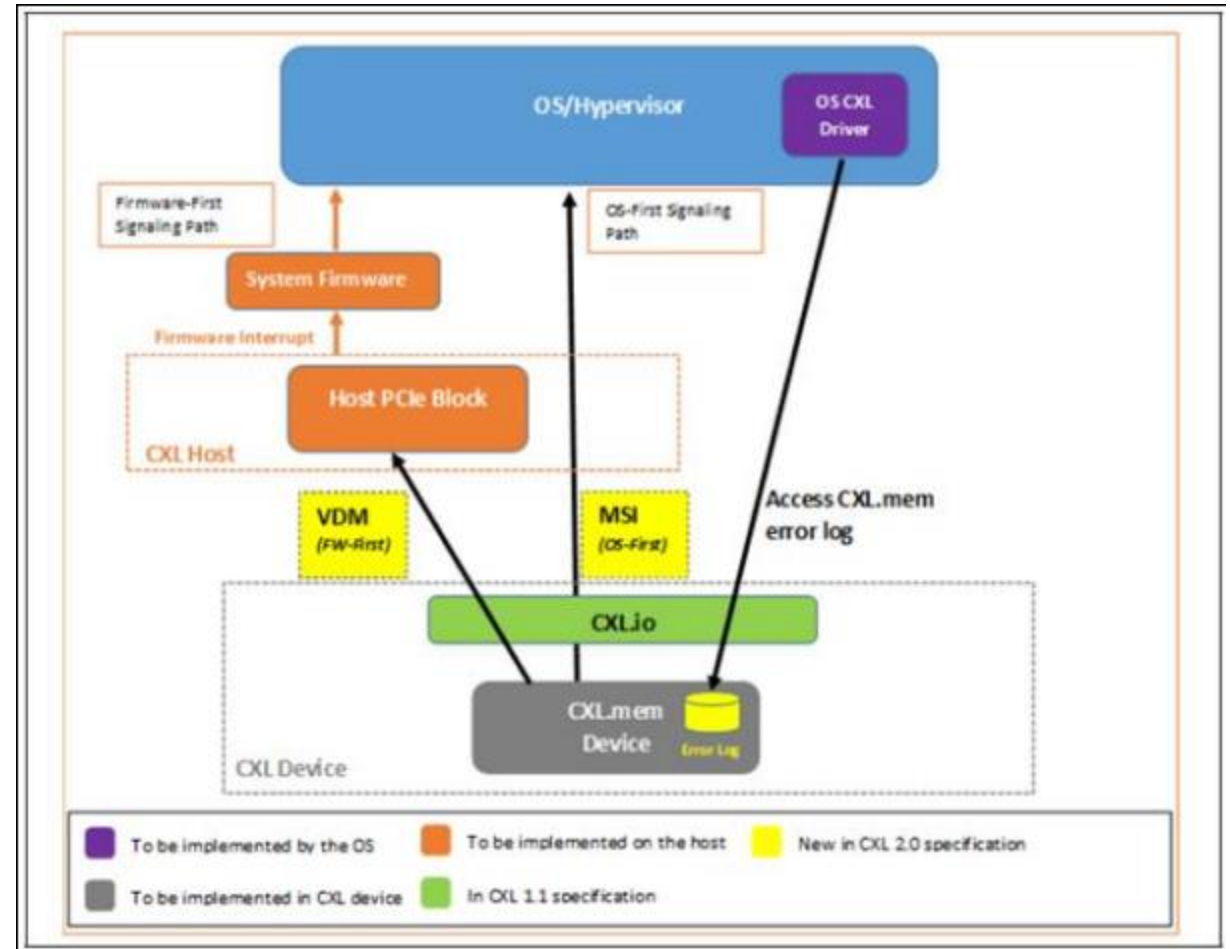
- Platform firmware consumes the error.
- Platform firmware may signal OS through existing ACPI methods.
- Platform firmware uses CXL Protocol Error record format (CPER) to represent the hardware error (see UEFI 2.9 spec, N.2.13).

OS FIRST

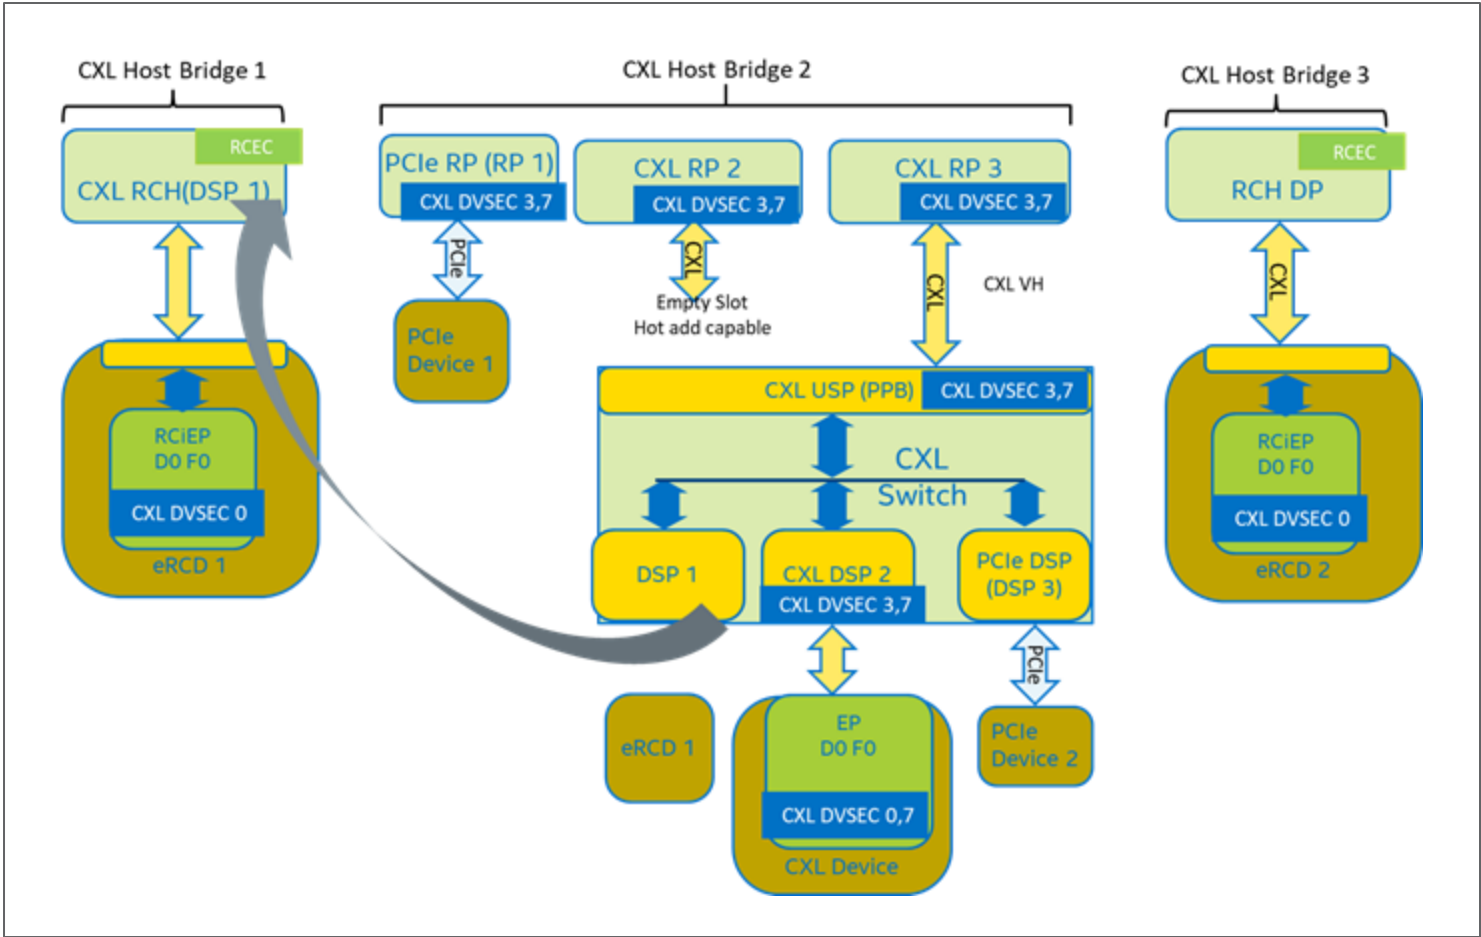
- OS must handle AER events through existing methods.
- If AER status indicates an “Internal Error” (UIE/CIE), and the source is a RCEC, then the OS must inspect the AER and CXL RAS structures of all associated DPs.
- If AER status indicates an “Internal Error” (UIE/CIE), and the source is an Endpoint, then the OS must inspect the CXL RAS structure of the associated UP.

CXL Component Errors

- Boot-time tasks
 - OS must request and be granted “Memory Error reporting” control through _OSC.
 - OS may load a “CXL Memory Device” driver to manage the Device/Component Command Interface.
 - OS may load a “CXL Event” service driver.
 - Leverage “CXL Memory Device” driver to interact with mailbox, etc.
 - Enable polling timers and MSI/MSI-X interrupts for Events
- Run-time tasks
 - Event driver handlers will execute mailbox commands to get and clear records.
 - Event driver will parse records and report to user.
 - Event driver may call other OS facilities to act on error information

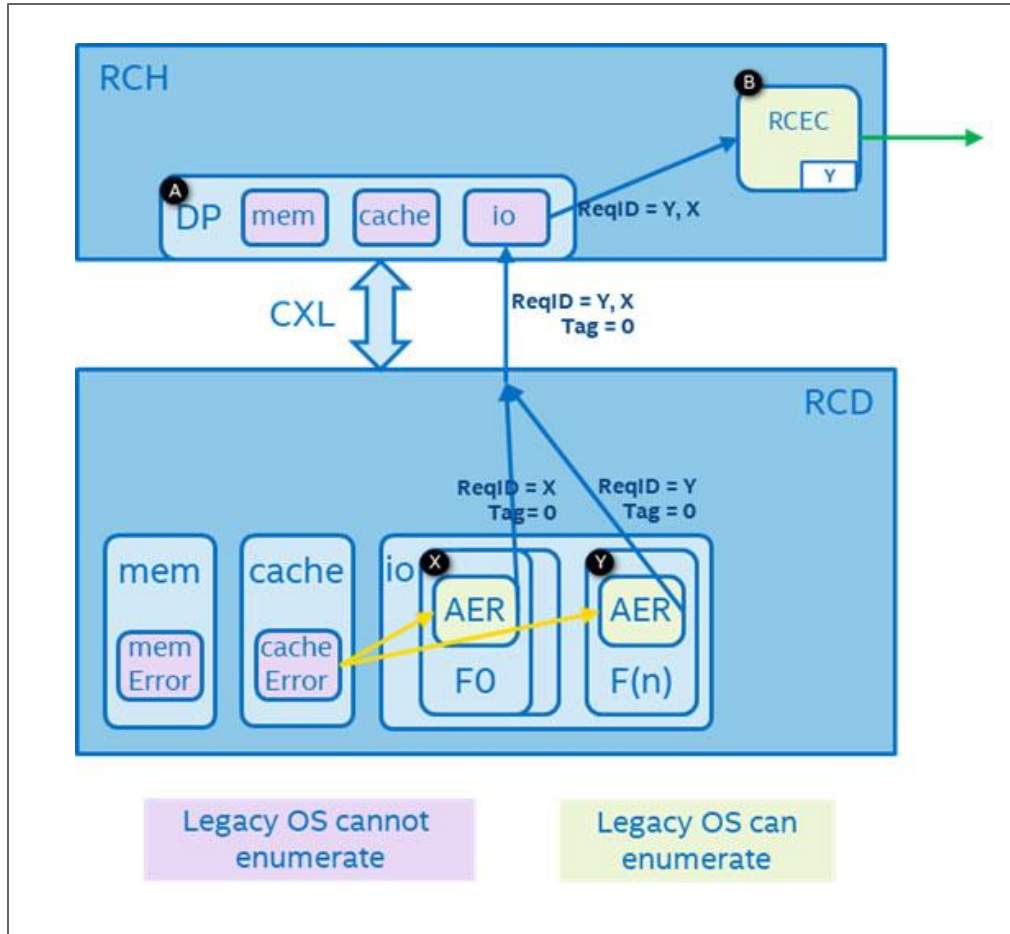


CXL VH vs. RCD mode - VH software view



Virtual Hierarchy (VH), Source: CXL-3.0, Figure 9-13

CXL VH vs. RCD mode - RCD software view



Source: CXL-3.0, Figure 12-4

Restricted CXL Device (RCD) mode (former CXL 1.1):

- RCH Downstream Port and RCD Upstream Port is not in the PCIe hierarchy.
- Device is an RCiEP.
- RCH contains an RCEC.
- A Root Complex Register Block (RCRB) exists for the DP and UP (containing PCI type 1 and 0 header) with MEMBAR0 pointing to the component registers.
- Component registers are the same as in VH mode.

CXL RAS - Missing in the kernel

- CXL RCD mode
- CPER extension for CXL Protocol Errors
- AER handling including support for CXL RAS caps and RCD (RCEC, CXL DP/UP) support (extend PCIe infrastructure)
- CXL mailbox interrupt support
- CXL event driver
- Notify userland, use existing subsystems (e.g. kernel log, tracepoints, EDAC)

CXL RAS - Plan for userspace

What could be in userspace?

- Mailbox interaction
- Monitoring tool (Event interrupts, etc.)
- Event handlers
- Address translation (hw to hw: device phys address to system phys address).
Memory interleaving, device's base address, etc.
- Fault analysis
- Interleaving and region setup

Userspace or kernel drivers?

- Keep the kernel simple with reference implementation in userspace?
- Shouldn't (must) the kernel handle the errors?

Summary

- CXL is a common, standard interface for many types of memory.
- CXL enables System design flexibility and optimizations for a variety of memory.
- Needs RAS to use CXL mem interchangeable with native memory.
- Many components are involved, and a variety of errors need to be handled for CXL RAS.
- MCA and PCIe error reporting can be reused partly.
- Support of CXL protocol and component errors handling is new.

→ Fresh patches, fresh kernels welcome.

References

- [1] Compute Express Link (CXL) Specification
August 1, 2022
Revision 3.0, Version 1.0
Available: <https://www.computeexpresslink.org/spec-landing>
- [2] PCI Express® Base Specification Revision 6.0
16 December 2021
PCI-SIG
Available: <https://pcisig.com/specifications>
- [3] Advanced Configuration and Power Interface (ACPI) Specification
UEFI Forum, Inc.
Version 6.4
January 2021
Available: <https://uefi.org/specifications>
- [4] Unified Extensible Firmware Interface (UEFI) Specification
Version 2.9
March 2021
Available: <https://uefi.org/specifications>

CXL RAS

Thank You!

CXL RAS

Discussion

COPYRIGHT AND DISCLAIMER

©2022 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate releases, for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD 