

Linux Plumbers Conference 2024



Contribution ID: 20

Type: **not specified**

Compute Express Link MC

CFP closes on July 5th.

Compute Express Link is a cache coherent fabric that has been gaining momentum in the industry. Whilst the ecosystem is still catching up with CXL 3.0 and earlier features, CXL 3.1 launched just after the 2023 CXL uconf, bringing yet more challenges for the community (temporal sharing, advanced RAS features). There also has been controversy and confusion in the Linux kernel community about the state and future of CXL, regarding its usage and integration into, for example, the core memory management subsystem. Many concerns have been put to rest through proper clarification and setting of expectations.

The Compute Express Link microconference focuses on how to evolve the Linux CXL kernel driver and userspace components for support of the CXL specifications. The microconference provides a place to open the discussion, incorporate more perspectives, and grow the CXL community with a goal that the CXL Linux plumbing serves the needs of the CXL ecosystem while balancing the needs of the Linux project. Specifically, this microconference welcomes submissions detailing industry and academia use cases in order to develop usage model scenarios. Finally, it will be a good opportunity to have existing upstream CXL developers available in a forum to discuss current CXL support and to communicate areas that need additional involvement.

The earlier editions of the microconference resolved a number of open questions (CXL 1.1 RAS now upstream), and introduced new topics we expect to revisit this year (e.g. dynamic capacity / shared memory and error handling)

Suggested topics:

- Ecosystem & Architectural review
- Dynamic Capacity Devices - Status and next steps
- Inter host shared capacity
- Fabric Management - What should Linux enable (blast radius concerns)? Open source solutions?
- Error handling and RAS (including OCP RAS API)
- Testing and emulation
- Security (ie: IDE/SPDM)
- Managing vendor specificity
- Virtualization of dynamic capacity.
- Type 2 accelerator support - CXL 3.0+ approaches.
- Coherence management of type2/3 memory (back-invalidation)
- Peer2Peer (ie: Unordered IO)
- Reliability, availability and serviceability (ie: Advanced Error Reporting, Isolation, Maintenance).
- Hotplug (QoS throttling, policies, daxctl)
- Hot remove
- Documentation
- Memory tiering topics that can relate to cxl (out of scope of MM/performance MCs)
- Industry and academia use cases

Primary authors: MANZANARES, Adam (Samsung Electronics); WILLIAMS, Dan (Intel Open Source Technology Center); BUESO, Davidlohr (Samsung Semiconductor); CAMERON, Jonathan (Huawei Technologies R&D (UK))

Presenters: MANZANARES, Adam (Samsung Electronics); WILLIAMS, Dan (Intel Open Source Technology Center); BUESO, Davidlohr (Samsung Semiconductor); CAMERON, Jonathan (Huawei Technologies R&D (UK))

Track Classification: LPC Microconference Proposals