

Compute Express Link + Linux + QEMU = Yes

Friday, 24 September 2021 07:00 (45 minutes)

Compute Express Link (CXL) is a cache coherent protocol designed to boost the performance of accelerators and memory expanders. By nature of being an open standard, it also allows vendors to design hardware which will work with a generic CXL driver in the same vein as AHCI, or NVMe (and others). With CXL 2.0 spec release late last year, we're starting to see usage models being actively pushed extend to client, enterprise, and cloud.

Our Linux driver team was tasked with writing the Linux driver without the luxury of having reference hardware to develop on. QEMU is a common solution for this problem. However, a couple of aspects set this journey aside from many others. First, this isn't just a device that requires emulation, but an entire bus, and second, we began development before the spec was even finalized and this allowed us to find spec bugs and gaps.

The talk will cover the plumbing in QEMU used to enable our driver development. It will begin with Compute Express Link 2.0 fundamentals, the challenges that posed, and how those were solved or deferred, in detail, inside QEMU. There will be some time spent on the Linux kernel development done to date, and why QEMU is such an ideal environment for a task like this.

The current work is very bare-bones with respect to the complex topologies and configurations that CXL 2.0 enables. The remainder will be spent on the current gaps in the QEMU emulation and a call for help on how we can fix those and what should be done [if anything] to get this work upstream.

I agree to abide by the anti-harassment policy

I agree

Primary author: WIDAWSKY, Ben

Presenter: WIDAWSKY, Ben

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track (Closed)