

Linux Plumbers Conference 2024



Contribution ID: 59

Type: **not specified**

CXL benchmarking

Thursday, 19 September 2024 13:05 (20 minutes)

Benchmarking and efficiency estimation of CXL infrastructure is a crucial task for the whole CXL ecosystem. Which tool(s) can be used and how can we execute such benchmarking? Potentially, a benchmarking tool could simulate the target use-case (for example, huge relational database, in-memory database, huge social network, ML model training, Virtual Machine use-case, HPC use-case, and so on). But, technically speaking, we need a tool that is capable of generating workloads with the opportunity of tuning various parameters (allocation/deallocation size, total allocation size, allocation/deallocation pattern, read/write pattern, memory type selection policy, migration policy, threads number and so on). I would like to suggest the discussion which tuning parameters are really crucial and how we can implement and deliver the CXL benchmarking infrastructure.

Primary authors: MANZANARES, Adam; DUBEYKO, Viacheslav

Presenters: MANZANARES, Adam; DUBEYKO, Viacheslav

Session Classification: Compute Express Link MC

Track Classification: Compute Express Link MC