CXL Cache (safe) support

Alejandro Lucero alucerop@amd.com
December, 2025
Linux Plumbers Conference

CXL Cache (safe) support

Goals:

1) Heads up about CXL.cache support
2) Discuss potential needs
3) Identify work to do

CXL Cache (safe) support

- CXL Type2: CXL.mem and CXL.cache
- CXL Type1: CXL.cache
- CXL.cache: device caching Host memory
- CXL.cache: device handling its HDM coherency (Host caching HDM device memory) → CXL.mem + HDM-DB

CXL Cache (safe) support

- CXL Type2: CXL.mem and CXL.cache
- CXL Type1: CXL.cache
- CXL.cache: device caching Host memory
- CXL.cache: device handling its HDM coherency (Host caching HDM device memory) → CXL.mem + HDM-DB

CXL Cache (safe) support

- Per CXL Host Bridge Snoop Cache
- A device advertises a cache size
- Host checking cache size <= snoop cache size
- Host enables Type2/1 cache functionality
- Per device tracking? HW? SW managing cache size available

CXL Cache (safe) support

- Per device HW tracking?
  - A device reading more than its cache size? Rogue or maquiavellian device … should we care?
  - A polite device suffering? Snoop cache asking to this device to release a cache line due to the other one (previous) activities?
- If HW does not control it, the only way SW can is to restrict one CXL.cache device per Snoop cache.

CXL Cache (safe) support

CXL.cache and Host addresses

- DMA-like Host control without DMAbility or coherency concerns
- Physical address or input-output virtual address (IOMMU later)
- From specs: CXL.cache addresses restricted (in number? Cache size?)
    - With IOMMU for sure (as DMA addresses)
    - Without IOMMU? Not all DMA addresses are suitable for CXL.cache (specs)

CXL Cache (safe) support

CXL.cache and Host addresses: **Do we need to change DMA API?**

- CXL wrapper or CXL API tracking per-device requests?
- DMA API being CXL-aware?

- DMA linked to IOMMU ...

CXL Cache (safe) support

CXL.cache and Host addresses (Virtualization)

- A passthrough device using CXL.cache? It makes sense but ...
- Hard to get answers from … feeling security is a second concern (hihertho?)
- IOMMU or ATS? Former is safe but counterproductive and counterintuitive … in CXL terms … Latter not enough (passthrough device not trusted)
- What should the kernel do?

CXL Cache (safe) support

CXL.cache and Host addresses (Virtualization)

- Create a working group with IOMMU kernel
- Get IOMMU/CXL HW (Intel, AMD, ARM) involved

  …
- Do not allow CXL.cache for guests/user space by now

CXL Cache (safe) support

Take aways:

- SW snoop cache control restricting devices (or having an option to)?
- CXL.cache addresses allocation/mapping API or DMA API changes?
- CXL.cache and virtualization: security with IOMMU?

CXL Cache (safe) support

Q & A