

A decorative graphic of a green pipe network with various fittings, valves, and elbows, set against a white background. The pipes are arranged in a complex, interconnected pattern, with some sections being thicker than others. The overall style is clean and modern, with a focus on the color green.

# CXL QEMU Update / Plans

Jonathan Cameron - Huawei



**Linux  
Plumbers Conference** | Dublin, Ireland **Sept. 12-14, 2022**

A decorative graphic of a green pipe network with various fittings, elbows, and valves, framing the central text and list. The pipes are a vibrant green color and are set against a white background with soft shadows.

# What I'm looking for today...

- Input on priorities.
- Path to meeting those priorities.



**Linux  
Plumbers Conference** | Dublin, Ireland **Sept. 12-14, 2022**



# Why QEMU emulation matters?

- See Ben Widawsky's [talk](#) from last year.
- Allows flexible emulation of complex configurations.
- Poke all the paths, not just the ones that are easy to Mock.
- Different purpose but some overlap with in-kernel Mocking Interface.
- Not done yet, but CI testing...



# The Dream!

- CXL QEMU emulation sufficient to exercise ALL operating system support.
- I'm dreaming, so Fabric management as well.
  - Maybe RTOS on management controller?
- Platform for specification prove out.
  - Not public until spec release: need a baseline to do this fast.
  - CXL PMU definition in CXL 3.0 benefited from this.



A decorative graphic of a green pipe network with various fittings, elbows, and valves, framing the slide content.

# The BIG Trade-off

- PoC code is easy, upstreaming takes much longer.
- Often a PoC is enough to be *fairly* sure the kernel code is good...
- Doesn't provide a good long term platform to test or build on.



# What is upstream?

- CFMWS (x86), Host Bridge, Root Port, Switches, Type 3 Devices.
- Everything\* needed to bring up interleaved persistent memory Type 3 devices.



Linux

Plumbers Conference | Dublin, Ireland Sept. 12-14, 2022

\*If you want to bring it up twice you need serial number patch!

# What else is 'done'

- [ARM support](#) (blocked on lack of DT)
- [DOE / CDAT](#) (runtime autogen of content needed)
- [DOE / Compliance](#) (not of interest to me!)



A decorative graphic of a green pipe network with various fittings, elbows, and valves, framing the slide content.

# Hacks / in progress

- [CXL PMU](#)
- [Fabric manger over I2C / MCTP](#)
- Switch CCI mailbox (stand alone PCI function) – not posted yet.
- [SPDM](#)
- [Poison Injection](#)





A decorative graphic of green pipes with valves and fittings, running vertically on the left side of the slide and curving at the top and bottom.

# In Kernel, not QEMU

- [Event logs.](#)
- Security commands
- CXL 1.1 support (Restricted CXL Device – RCD)



# Fun for the future?

- Volatile capacity.
- Dynamic Capacity
- FM-API via mailbox end point – MLDs!
  - Connect them to a switch with VCS roots in different RPs / PCI Expander Bridges.
  - Exercise all the ‘fun’ corners.
- Fabric (small) – once the hardware interfaces defined.



# What we need!

- More review.
  - QEMU community are great but we can't expect them to review the CXL specific parts.
- More features.
  - All contributions welcome!



# Thanks

- Ben Widawsky who started it all.
- Chris Browy and Huai-Cheng Kuo (DOE / CDAT etc)
- Michael Tsirkin - Upstream maintainer.
- QEMU reviewers (particularly those I kept pestering ;)

