

Linux Plumbers Conference

Dublin, Ireland September 12-14, 2022

A decorative graphic of a green pipe network with various fittings, valves, and elbows, set against a white background with soft shadows.

CXL Hotplug: spec to reality


PJ Waskiewicz

Jump Trading



Linux

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

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

Agenda

- CXL overview
- CXL roadmaps
- FPGA use cases
- Implementation vs. Spec
- Hotplug proposals, next steps

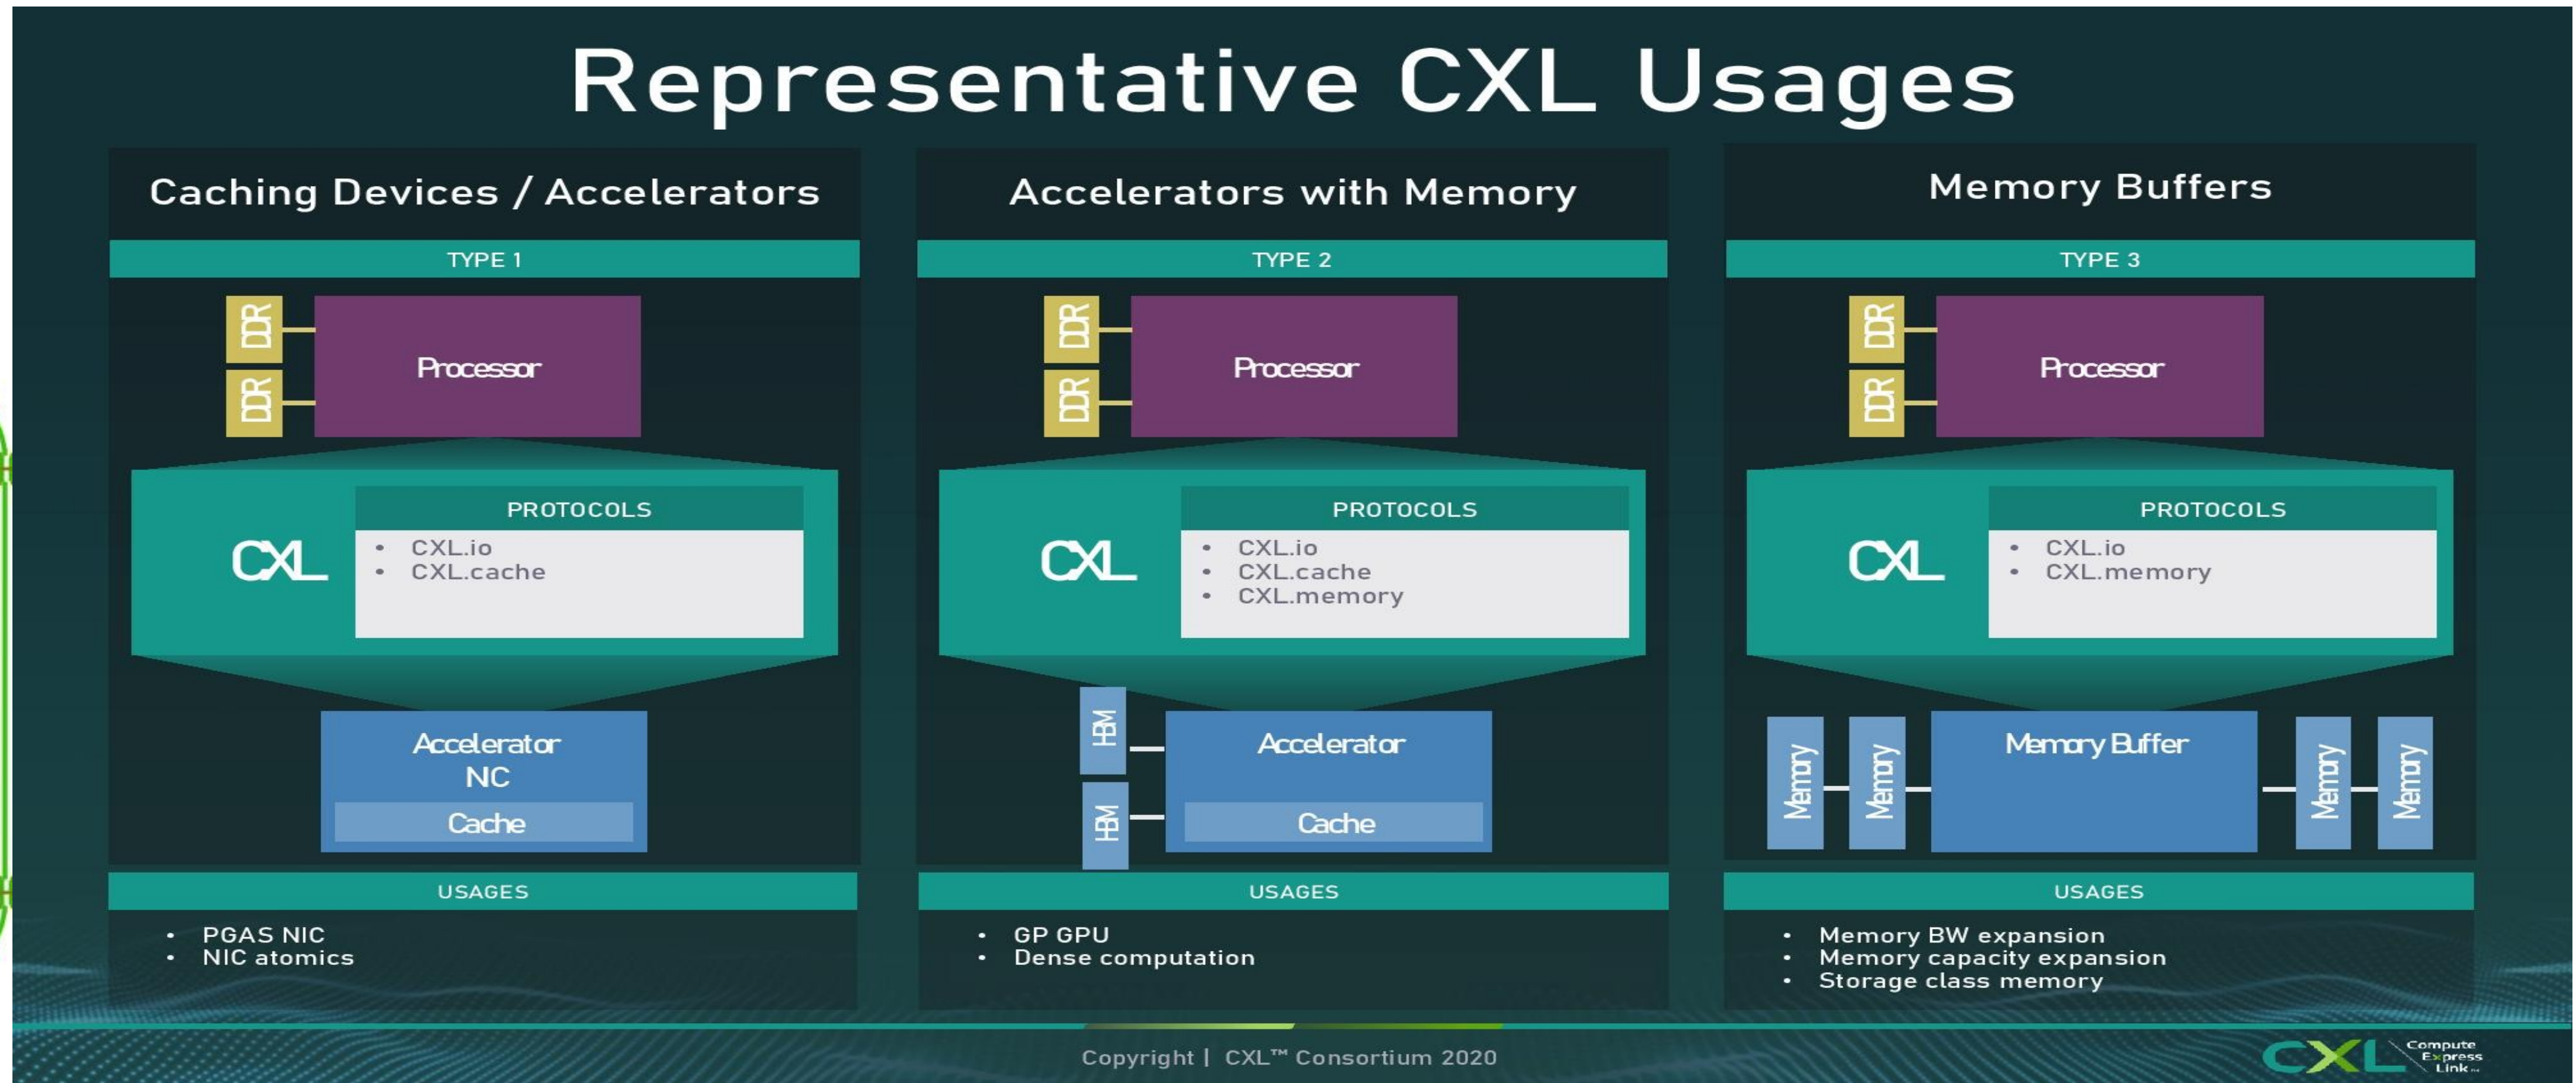


Linux

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

CXL Overview

Representative CXL Usages



CXL roadmaps

	Intel	AMD	Nvidia
CXL 1.1	Sapphire Rapids	Genoa	Grace A100
CXL 2.0	Granite Rapids ??	Turin ??	Grace A100
CXL 3.0	???	???	???




FPGA usage

- Endpoint devices reprogram on-demand
 - e.g. - pick up new network filters
- Most write to flash, initiate FPGA fabric rewrite
- PCIe hotplug used to “reboot” card
- Rescan bus to rediscover card, start driver load process from probe()
- Avoids host reboot (potentially minutes)



Linux

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



CXL 1.1 vs Hotplug

- Per spec, CXL 1.1 does not support hotplug
- Hotplug arrives “officially” in CXL 2.0

- Issue brought up at SWWG
 - Was told this is implementation problem, not spec problem
 - Spawned more questions from other members around hotplug



Linux

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



CXL 1.1 Implementation

- According to some vendors, initial host-side support might be “1.1+”
 - Still unsure what the “+” means
 - Hinted that limited hotplug support might be included
- 1.1 implementation mess will be difficult to manage until 2.0 arrives



Linux

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

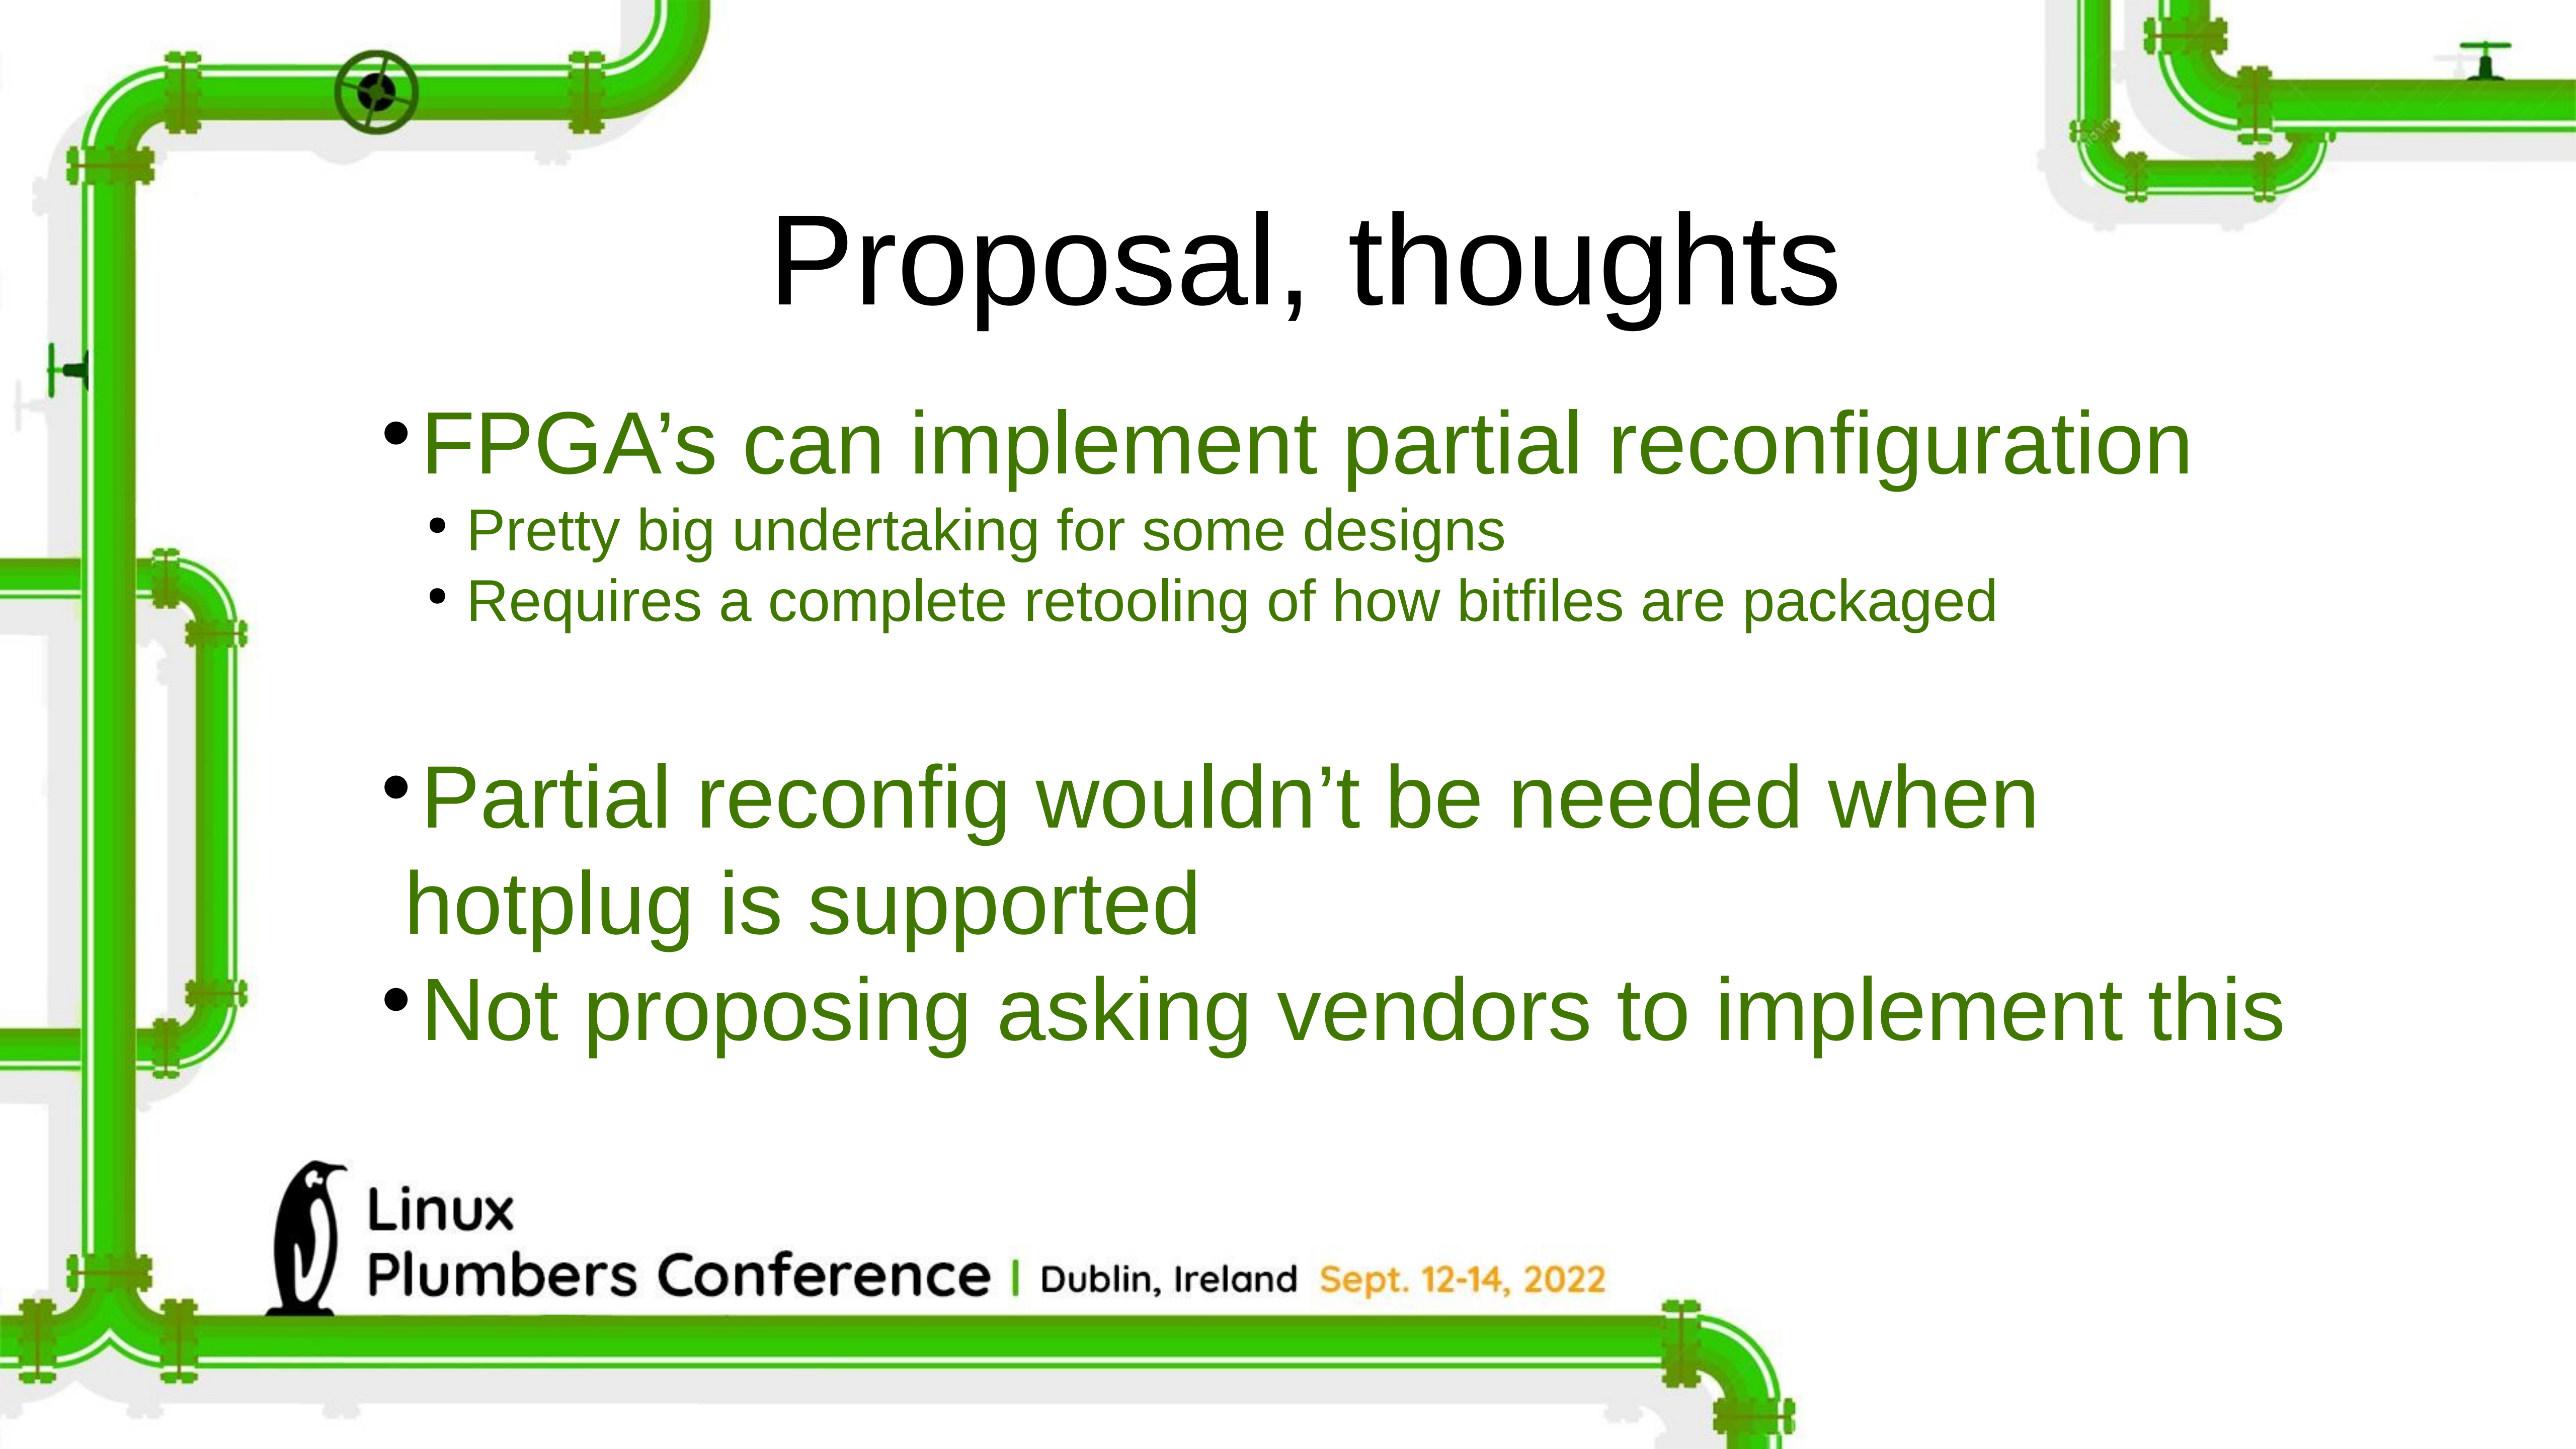
Proposal, thoughts

- Expose a “best-effort” hotplug hook in sysfs?
 - Keep the same remove/rescan from PCIe today
 - Marked as “use at your own risk” hook, could still cause panics
- CXL.cache could be tricky, effectively CPU hotplug
- CXL.mem could also be tricky
 - Need to quiesce memory transactions before hotplugging memory



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.


Proposal, thoughts

- FPGA's can implement partial reconfiguration
 - Pretty big undertaking for some designs
 - Requires a complete retooling of how bitfiles are packaged
- Partial reconfig wouldn't be needed when hotplug is supported
- Not proposing asking vendors to implement this



Linux

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



Linux Plumbers Conference

Dublin, Ireland **September 12-14, 2022**