

Linux Plumbers Conference 2024



Contribution ID: 354

Type: **not specified**

IPI deferral

CPU isolation allows us to shield a subset of CPUs from a lot of kernel interference, but not all of it. Activity on the housekeeping CPUs can and does trigger IPIs which can still end up targeting isolated CPUs. The main culprits here are static key updates and `vunmap()` + the resulting `flush_tlb_kernel_range()`.

As discussed in previous editions, since these IPIs are only relevant to the kernel (and not to userspace), an approach to remove the interference is to defer these IPIs until the targeted (isolated) CPUs next enter the kernel.

This talk will present the changes that have happened since this was last presented, and open a discussion concerning the remaining challenges.

Primary author: SCHNEIDER, Valentin (Red Hat)

Session Classification: Real-time MC

Track Classification: Real-time MC