

Linux Plumbers Conference 2024



Contribution ID: 230

Type: not specified

Policy zones: memory partitioning for fun and profit

Friday 20 September 2024 18:15 (15 minutes)

There are three types of zones:

1. The first four zones partition the physical address space of CPU memory.
2. The device zone provides interoperability between CPU and device memory.
3. The movable zone commonly represents a memory allocation policy.

Though originally designed for memory hot removal, the movable zone is instead widely used for other purposes, e.g., CMA and kdump kernel, on platforms that do not support hot removal. Nowadays, it is legitimately a zone independent of any physical characteristics. In spite of being somewhat regarded as a hack, largely due to the lack of a generic design concept for its true major use cases, the movable zone naturally resembles a policy (virtual) zone overlayed on the first four (physical) zones.

This proposal formally generalizes this concept as policy zones so that additional policies can be implemented and enforced, to for example optimize page allocations based on properties known at allocation time, or even runtime. Those properties include memory object size and mobility, or hotness and lifetime, respectively.

Jon at LWN kindly wrote a nice article on this topic: <https://lwn.net/Articles/964239/>

Primary author: ZHAO, Yu (Google)

Presenter: ZHAO, Yu (Google)

Session Classification: Kernel Memory Management MC

Track Classification: Kernel Memory Management MC