



Contribution ID: 219

Type: **not specified**

Towards Better Memory Allocation for Device Drivers

Friday, 20 September 2024 15:15 (15 minutes)

For decades, Linux memory management has been mostly focused on the needs of user space and generic kernel-space users (memory control groups, transparent huge pages, compression). Other big changes are good for maintenance and/or debugging (removal of DISCONTIGMEM, compaction, kmemleak, folios, removal of redundant slab-style allocators and many other). Little has been done for device drivers (only CMA comes to mind).

It appears that Linux memory management does not match the needs of device drivers very well. At some point, the ugly, quirky and subtly or grossly broken hardware kicks in and spoils the party. It's usually the job of device drivers to bridge the gap. They often have to care about the placement of memory buffers in physical address space. Surprisingly or not, these constraints rarely match the constraints of DMA and DMA32 zones (if they even exist). As a result, these zones add complexity to the buddy allocator but they do not bring much benefit. CMA may help sometimes but not always.

I would like to discuss the alternatives and possible ways to remove DMA and DMA32 from the kernel.

Primary author: TESÁŘÍK, Petr

Presenter: TESÁŘÍK, Petr

Session Classification: Kernel Memory Management MC

Track Classification: Kernel Memory Management MC