



Contribution ID: 72

Type: **not specified**

The hard work behind large physical allocations in the kernel

Thursday, November 15, 2018 9:00 AM (45 minutes)

The physical memory management in the Linux kernel is mostly based on single page allocations, but there are many situations where a larger physically continuous memory needs to be allocated. Some are for the benefit of userspace (e.g. huge pages), others for better performance in the kernel (SLAB/SLUB, networking, and others).

Making sure that contiguous physical memory is available for allocation is far from trivial, as pages are reclaimed for reuse roughly in last-recently-used (LRU) order, which is typically different from their physical placement. The freed memory is thus fragmented. The kernel has two complementary mechanisms to defragment the free memory. One is memory compaction, which migrates used pages to make the free pages contiguous. The other is page grouping by mobility, which tries to make sure that pages that cannot be migrated are grouped together, so the rest of pages can be effectively compacted. Both mechanisms employ various heuristics to balance the success of large allocations, and their overhead in terms of latencies due to processor and lock usage.

The talk will discuss the two mechanisms, focusing on the known problems and their possible solutions, that have been proposed by several memory management developers.

I agree to abide by the anti-harassment policy

Yes

Primary author: BABKA, Vlastimil (SUSE)

Presenter: BABKA, Vlastimil (SUSE)

Session Classification: LPC Main Track

Track Classification: Refereed talk