Subject: Re: [PATCH][for -mm] per-zone and reclaim enhancements for memory controller take 3 [3/10] per-zone
Posted by KAMEZAWA Hiroyuki on Thu, 29 Nov 2007 01:37:02 GMT
View Forum Message <> Reply to Message

On Wed, 28 Nov 2007 16:19:59 -0500
Lee Schermerhorn <Lee.Schermerhorn@hp.com> wrote:

> As soon as this loop hits the first non-existent node on my platform, I
> get a NULL pointer deref down in __alloc_pages. Stack trace below.
>
> Perhaps N_POSSIBLE should be N_HIGH_MEMORY? That would require handling
> of memory/node hotplug for each memory control group, right? But, I'm
> going to try N_HIGH_MEMORY as a work around.
>
Hmm, ok. (>_<

> Call Trace:
> [<a000000100014de0>] show_stack+0x80/0xa0
> sp=a0000001008e39c0 bsp=a0000001008dd1b0
> [<a000000100015a70>] show_regs+0x870/0x8a0
> sp=a0000001008e3b90 bsp=a0000001008dd158
> [<a00000010003d130>] die+0x190/0x300
> sp=a0000001008e3b90 bsp=a0000001008dd110
> [<a000000100071b80>] ia64_do_page_fault+0x8e0/0xa20
> sp=a0000001008e3b90 bsp=a0000001008dd0b8
> [<a00000010000b5c0>] ia64_leave_kernel+0x0/0x270
> sp=a0000001008e3c20 bsp=a0000001008dd0b8
> [<a000000100132e10>] __alloc_pages+0x30/0x6e0
> sp=a0000001008e3df0 bsp=a0000001008dcfe0
> [<a000000100187370>] new_slab+0x610/0x6c0
> sp=a0000001008e3e00 bsp=a0000001008dcf80
> [<a000000100187470>] get_new_slab+0x50/0x200
> sp=a0000001008e3e00 bsp=a0000001008dcf48
> [<a000000100187900>] __slab_alloc+0x2e0/0x4e0
> sp=a0000001008e3e00 bsp=a0000001008dcf00
> [<a000000100187c80>] kmem_cache_alloc_node+0x180/0x200
> sp=a0000001008e3e10 bsp=a0000001008dcec0
> [<a0000001001945a0>] mem_cgroup_create+0x160/0x400
> sp=a0000001008e3e10 bsp=a0000001008dce78
> [<a0000001000f0940>] cgroup_init_subsys+0xa0/0x400
> sp=a0000001008e3e20 bsp=a0000001008dce28
> [<a0000001008521f0>] cgroup_init+0x90/0x160
> sp=a0000001008e3e20 bsp=a0000001008dce00
> [<a000000100831960>] start_kernel+0x700/0x820
> sp=a0000001008e3e20 bsp=a0000001008dcd80
>

Maybe zonelists of NODE_DATA() is not initialized. you are right.
I think N_HIGH_MEMORY will be suitable here...(I'll consider node-hotplug case later.)

Thank you for test!

Regards,
-Kame