## Subject: Re: problem with ZONE\_MOVABLE. Posted by Balbir Singh on Thu, 13 Sep 2007 15:53:30 GMT

View Forum Message <> Reply to Message

```
Mel Gorman wrote:
> On (13/09/07 19:07), KAMEZAWA Hiroyuki didst pronounce:
>> Hi,
>>
>> While I'm playing with memory controller of 2.6.23-rc4-mm1, I met following.
>>
>> ==
>> [root@drpg test-2.6.23-rc4-mm1]# echo $$ > /opt/mem_control/group_1/tasks
>> [root@drpq test-2.6.23-rc4-mm1]# cat /opt/mem_control/group_1/memory.limit
>> 32768
>> [root@drpq test-2.6.23-rc4-mm1]# cat /opt/mem_control/group_1/memory.usage
>> 286
>> // Memory is limited to 512 GiB. try "dd" 1GiB (page size is 16KB)
>> [root@drpg test-2.6.23-rc4-mm1]# dd if=/dev/zero of=/tmp/tmpfile bs=1024 count=1048576
>> Killed
>> [root@drpg test-2.6.23-rc4-mm1]# ls
>> Killed
>> //above are caused by OOM.
>> [root@drpq test-2.6.23-rc4-mm1]# cat /opt/mem_control/group_1/memory.usage
>> 32763
>> [root@drpq test-2.6.23-rc4-mm1]# cat /opt/mem_control/group_1/memory.limit
>> 32768
>> // fully filled by page cache. no reclaim run.
>> ==
>>
>> The reason this happens is because I used kernelcore= boot option, i.e
>> ZONE_MOVABLE. Seems try_to_free_mem_container_pages() ignores ZONE_MOVABLE.
>> Quick fix is attached, but Mel's one-zonelist-pernode patch may change this.
>> I'll continue to watch.
>>
> You are right on both counts. This is a valid fix but
> one-zonelist-pernode overwrites it. Specifically the code in question
> with one-zonelist will look like:
>
> for_each_online_node(node) {
> zonelist = &NODE DATA(node)->node zonelist;
   if (do_try_to_free_pages(zonelist, sc.gfp_mask, &sc))
   return 1;
>
> }
> We should be careful that this problem does not get forgotten about if
```

- > one-zonelist gets delayed for a long period of time. Have the fix at the
- > end of the container patchset where it can be easily dropped if
- > one-zonelist is merged.

>

> Thanks

Yes, I second that. So, we should get KAMEZAWA's fix in.

--

Warm Regards, Balbir Singh Linux Technology Center IBM, ISTL

Containers mailing list Containers@lists.linux-foundation.org https://lists.linux-foundation.org/mailman/listinfo/containers