
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 02:24:06 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Thu, 29 Nov 2007 10:37:02 +0900

KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:

> 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!

>
Could you try this ?

Thanks,
-Kame
==

Don't call kmalloc() against possible but offline node.

Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

mm/memcontrol.c | 10 ++++++---
1 file changed, 8 insertions(+), 2 deletions(-)

Index: test-2.6.24-rc3-mm1/mm/memcontrol.c

```
=====
--- test-2.6.24-rc3-mm1.orig/mm/memcontrol.c
+++ test-2.6.24-rc3-mm1/mm/memcontrol.c
@@ -1117,8 +1117,14 @@ static int alloc_mem_cgroup_per_zone_inf
 struct mem_cgroup_per_node *pn;
 struct mem_cgroup_per_zone *mz;
 int zone;
-
- pn = kmalloc_node(sizeof(*pn), GFP_KERNEL, node);
+ /*
+  * This routine is called against possible nodes.
+  * But it's BUG to call kmalloc() against offline node.
+  */
+ if (node_state(N_ONLINE, node))
+ pn = kmalloc_node(sizeof(*pn), GFP_KERNEL, node);
+ else
+ pn = kmalloc(sizeof(*pn), GFP_KERNEL);
+ if (!pn)
+ return 1;
```

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