
Subject: [PATCH][BUGFIX][for -mm] Misc fix for memory cgroup [1/5] -- fix refcnt check in mem_cgroup_charge

Posted by [KAMEZAWA Hiroyuki](#) on Thu, 11 Oct 2007 04:55:49 GMT

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

The logic of uncharge is

- decrement refcnt -> lock page cgroup -> remove page cgroup.

But the logic of charging is

- lock page cgroup -> increment refcnt -> return.

Then, one charge will be added to a page_cgroup under being removed.

This makes no big trouble (like panic) but one charge is lost.

This patch add a test at charging to verify page_cgroup's refcnt is greater than 0. If not, unlock and retry.

Changelog v1->v2:

- * added cpu_relax() before retry.

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

mm/memcontrol.c | 10 ++++++--

1 file changed, 8 insertions(+), 2 deletions(-)

Index: devel-2.6.23-rc8-mm2/mm/memcontrol.c

=====

--- devel-2.6.23-rc8-mm2.orig/mm/memcontrol.c

+++ devel-2.6.23-rc8-mm2/mm/memcontrol.c

@ @ -271,14 +271,20 @ @ int mem_cgroup_charge(struct page *page,

 * to see if the cgroup page already has a page_cgroup associated

 * with it

 */

+retry:

 lock_page_cgroup(page);

 pc = page_get_page_cgroup(page);

 /*

 * The page_cgroup exists and the page has already been accounted

 */

 if (pc) {

- atomic_inc(&pc->ref_cnt);

- goto done;

+ if (unlikely(!atomic_inc_not_zero(&pc->ref_cnt))) {

+ /* this page is under being uncharged ? */

+ unlock_page_cgroup(page);

+ cpu_relax();

+ goto retry;

+ } else

```
+ goto done;  
}  
  
unlock_page_cgroup(page);
```

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