
Subject: [RFC][for -mm] memory controller enhancements for reclaiming take2 [4/8]
possible race fix in res_co

Posted by [KAMEZAWA Hiroyuki](#) on Mon, 03 Dec 2007 09:38:09 GMT

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

spinlock is necessary when someone changes res->counter value.
splitted out from YAMAMOTO's background page reclaim for memory cgroup set.

Changelog v1 -> v2:
- fixed type of "flags".

Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
From: YAMAMOTO Takashi <yamamoto@valinux.co.jp>

kernel/res_counter.c | 4 +++-
1 file changed, 3 insertions(+), 1 deletion(-)

Index: linux-2.6.24-rc3-mm2/kernel/res_counter.c

```
=====
--- linux-2.6.24-rc3-mm2.orig/kernel/res_counter.c
+++ linux-2.6.24-rc3-mm2/kernel/res_counter.c
@@ -98,6 +98,7 @@ ssize_t res_counter_write(struct res_cou
{
    int ret;
    char *buf, *end;
+ unsigned long flags;
    unsigned long long tmp, *val;

    buf = kmalloc(nbytes + 1, GFP_KERNEL);
@@ -121,9 +122,10 @@ ssize_t res_counter_write(struct res_cou
    if (*end != '\0')
        goto out_free;
}
-
+ spin_lock_irqsave(&counter->lock, flags);
    val = res_counter_member(counter, member);
    *val = tmp;
+ spin_unlock_irqrestore(&counter->lock, flags);
    ret = nbytes;
out_free:
    kfree(buf);
```

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