Subject: Re: [RFC][only for review] memory controller bacground reclaim [1/5] spinlock fix in res_counter m
Posted by Pavel Emelianov on Wed, 28 Nov 2007 11:08:31 GMT
View Forum Message <> Reply to Message

KAMEZAWA Hiroyuki wrote:
> spinlock is necessary when someone changes res_counter value.
> splited out from YAMAMOTO's background page reclaim for memory cgroup set.
>
> Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
> From: YAMAMOTO Takashi <yamamoto@valinux.co.jp>
>
>
> kernel/res_counter.c | 5 +++--
> 1 file changed, 3 insertions(+), 2 deletions(-)
>
> Index: linux-2.6.24-rc3-mm1/kernel/res_counter.c
> ===
> --- linux-2.6.24-rc3-mm1.orig/kernel/res_counter.c 2007-11-27 14:07:44.000000000 +0900
> +++ linux-2.6.24-rc3-mm1/kernel/res_counter.c 2007-11-27 14:09:40.000000000 +0900
> @@ -98,7 +98,7 @@
> {
>   int ret;
>   char *buf, *end;
> - unsigned long long tmp, *val;
> + unsigned long long flags, tmp, *val;

Flags for irqsave should be unsigned long. No?

>   buf = kmalloc(nbytes + 1, GFP_KERNEL);
>   ret = -ENOMEM;
> @@ -121,9 +121,10 @@
>   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