
Subject: Re: [RFC][only for review] memory controller bacground reclaim [1/5]
spinlock fix in res_counter m

Posted by [KAMEZAWA Hiroyuki](#) on Thu, 29 Nov 2007 01:14:39 GMT

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

On Wed, 28 Nov 2007 14:08:31 +0300

Pavel Emelyanov <xemul@openvz.org> wrote:

> 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 irqsav should be unsigned long. No?

>

yes. It must be.

Thanks, I'll fix.

-Kame

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
