Subject: Re: Question : memrlimit cgroup's task_move (2.6.26-rc5-mm3)
Posted by KAMEZAWA Hiroyuki on Fri, 20 Jun 2008 00:09:48 GMT
View Forum Message <> Reply to Message

On Thu, 19 Jun 2008 23:55:56 +0530
Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

> * KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> [2008-06-19 12:14:35]:
>
> > I used memrlimit cgroup at the first time.
> >
> > May I ask a question about memrlimit cgroup ?
> >
>
> Hi, Kamezawa-San,
>
> Could you please review/test the patch below to see if it solves your
> problem? If it does, I'll push it up to Andrew
>

At quick glance,
> + /*
> + * NOTE: Even though we do the necessary checks in can_attach(),
> + * by the time we come here, there is a chance that we still
> + * fail (the memrlimit cgroup has grown its usage, and the
> + * addition of total_vm will no longer fit into its limit)
> + */
I don't like this kind of holes. Considering tests which are usually done
by developpers, the problem seems not to be mentioned as "rare"..
It seems we can easily cause Warning. right ?

Even if you don't want to handle this case now, please mention as "TBD"
rather than as "NOTE".

> +
> +/*
> + * Add the value val to the resource counter and check if we are
> + * still under the limit.
> + */
> +static inline bool res_counter_add_check(struct res_counter *cnt,
> +   unsigned long val)
> +{
> + bool ret = false;
> + unsigned long flags;
> +
> + spin_lock_irqsave(&cnt->lock, flags);
> + if (cnt->usage + val < cnt->limit)

> + ret = true;
cnt->usage + val <= cnt->limit.

Thanks,
-Kame

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