
Subject: Re: [RFD][PATCH] memcg: Move Usage at Task Move
Posted by [KAMEZAWA Hiroyuki](#) on Wed, 11 Jun 2008 04:37:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 11 Jun 2008 13:29:09 +0900

Daisuke Nishimura <nishimura@mxp.nes.nec.co.jp> wrote:

```
> On Wed, 11 Jun 2008 13:14:37 +0900, KAMEZAWA Hiroyuki
<kamezawa.hiroyu@jp.fujitsu.com> wrote:
> > On Wed, 11 Jun 2008 12:44:46 +0900 (JST)
> > yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:
> >
> > > I'm now considering following logic. How do you think ?
> > >
> > > Assume: move TASK from group:CURR to group:DEST.
> > >
> > > == move_task(TASK, CURR, DEST)
> > >
> > > if (DEST's limit is unlimited)
> > >   moving TASK
> > >   return success.
> > >
> > > usage = check_usage_of_task(TASK).
> > >
> > > /* try to reserve enough room in destination */
> > > if (try_to_reserve_enough_room(DEST, usage)) {
> > >   move TASK to DEST and move pages AMAP.
> > >   /* usage_of_task(TASK) can be changed while we do this.
> > >     Then, we move AMAP. */
> > >   return success;
> > > }
> > > return failure.
> > > ==
> >
> > > AMAP means that you might leave some random charges in CURR?
> >
> > yes. but we can reduce bad case by some way
> > - reserve more than necessary.
> > or
> > - read_lock mm->sem while move.
> >
> > I preffer the latter.
> > Though it's expensive, I think moving a task would not happen
> > so often.
>
> Sure.
```

I'd like to write one and post as RFC. (hopefully in this week)

Thanks,
-Kame

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