
Subject: Re: [RFD][PATCH] memcg: Move Usage at Task Move
Posted by [serue](#) on Thu, 12 Jun 2008 13:17:48 GMT

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

Quoting KAMEZAWA Hiroyuki (kamezawa.hiroyu@jp.fujitsu.com):

> On Wed, 11 Jun 2008 01:48:20 -0700

> "Paul Menage" <menage@google.com> wrote:

>

> > On Wed, Jun 11, 2008 at 1:27 AM, KAMEZAWA Hiroyuki

> > <kamezawa.hiroyu@jp.fujitsu.com> wrote:

> > > Sorry. try another sentence..

> > >

> > > I think cgroup itself is designed to be able to be used without middleware.

> >

> > True, but it shouldn't be hostile to middleware, since I think that

> > automated use will be much more common. (And certainly if you count

> > the number of servers :-)

> >

> > > IOW, whether using middleware or not is the matter of users not of developpers.

> > > There will be a system that system admin controlles all and move tasks by hand.

> > > ex)...personal notebooks etc..

> > >

> >

> > You think so? I think that at the very least users will be using tools

> > based around config scripts, rule engines and libcgroup, if not a

> > persistent daemon.

> >

> > I believe some users will never use middlewares because of their special

> > usage of linux.

>

>

>

> > >> If the common mode for middleware starting a new cgroup is fork() /

> > >> move / exec() then after the fork(), the child will be sharing pages

> > >> with the main daemon process. So the move will pull all the daemon's

> > >> memory into the new cgroup

> > >>

> > > My patch (this patch) just moves Private Anon page to new cgroup. (of mapcount=1)

> >

> > OK, well that makes it more reasonable regarding the above problem.

> > But I can still see problems if, say, a single thread moves into a new

> > cgroup, you move the entire memory. Perhaps you should only do so if

> > the mm->owner changes task?

> >

>

> Thank you for pointing out. I'll add mm->owner check.

>

> BTW, should we have a cgroup for SYSVIPC resource controller and devide it

> from memory resource controller ? I think that per-task on-demand usage
> accounting is not suitable for shmem (and hugepage).
> per-creator (caller of shmget()) accounting seems to be better for me.
>
> Just a question:
> What happens when a thread (not thread-group-leader) changes its ns by
> ns-cgroup ? not-allowed ?

I don't quite understand the question. I assume you're asking whether
your cgroup, when composed with ns, will refuse a task in cgroup /cg/1/2
from being able to

```
mkdir /cg/1/2/3  
echo $$ > /cg/1/2/3/tasks
```

or

```
unshare(CLONE_NEWNS)
```

which the ns cgroup would allow, and what your cgroup would do in that
case. If your question ("not-allowed ?") is about ns cgroup behavior
then please rephrase.

thanks,
-serge

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