
Subject: Re: [-mm PATCH 5/10] Memory controller task migration (v7)

Posted by [yamamoto](#) on Tue, 28 Aug 2007 08:32:52 GMT

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

> YAMAMOTO Takashi wrote:

> >> Allow tasks to migrate from one container to the other. We migrate

> >> mm_struct's mem_container only when the thread group id migrates.

> >

> >> + /*

> >> + * Only thread group leaders are allowed to migrate, the mm_struct is

> >> + * in effect owned by the leader

> >> + */

> >> + if (p->tgid != p->pid)

> >> + goto out;

> >

> > does it mean that you can't move a process between containers

> > once its thread group leader exited?

> >

> > YAMAMOTO Takashi

>

>

> Hi,

>

> Good catch! Currently, we treat the mm as owned by the thread group leader.

> But this policy can be easily adapted to any other desired policy.

> Would you like to see it change to something else?

>

> --

> Warm Regards,

> Balbir Singh

> Linux Technology Center

> IBM, ISTL

although i have no good idea right now, something which allows
to move a process with its thread group leader dead would be better.

YAMAMOTO Takashi

Containers mailing list

Containers@lists.linux-foundation.org

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