

---

Subject: Re: [RFD][PATCH] memcg: Move Usage at Task Move  
Posted by [Paul Menage](#) on Wed, 11 Jun 2008 08:48:20 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

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.

>> 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?

Paul

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---