Subject: Question : memrlimit cgroup's task_move (2.6.26-rc5-mm3)
Posted by KAMEZAWA Hiroyuki on Thu, 19 Jun 2008 03:14:35 GMT
View Forum Message <> Reply to Message

I used memrlimit cgroup at the first time.

May I ask a question about memrlimit cgroup ?

In following
==
```
static void memrlimit_cgroup_move_task(struct cgroup_subsys *ss,
                                struct cgroup *cgrp,
                                struct cgroup *old_cgrp,
                                struct task_struct *p)
{
        struct mm_struct *mm;
        struct memrlimit_cgroup *memrcg, *old_memrcg;

<snip>
        if (res_counter_charge(&memrcg->as_res, (mm->total_vm << PAGE_SHIFT)))
                goto out;
        res_counter_uncharge(&old_memrcg->as_res, (mm->total_vm << PAGE_SHIFT));
```
==
This is a callback for task_attach(). and this never fails.

What happens when the moved task, which move-of-charge fails, exits ?

==
```
% mkdir /dev/cgroup/memrlimit/group_01
% mkdir /dev/cgroup/memrlimit/group_02
% echo 1G > /dev/cgroup/memrlimit/group_01/memrlimit.limit_in_bytes
% echo 0 >  /dev/cgroup/memrlimit/group_02/memrlimit.limit_in_bytes
% echo $$ > /dev/cgroup/memrlimit/group_01/tasks
% echo $$ > /dev/cgroup/memrlimit/group_02/tasks
% exit
```
== you'll see WARNING ==

I think the charge of the new group goes to minus. right ?
(and old group's charge never goes down.)
I don't think this is "no problem".

What kind of patch is necessary to fix this ?
task_attach() should be able to fail in future ?

I'm sorry if I misunderstand something or this is already in TODO list.

Thanks,
-Kame

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