
Subject: Re: [PATCH 5/12] Handle pid namespaces in cgroups code

Posted by [Paul Menage](#) on Tue, 29 Jan 2008 18:08:12 GMT

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

On Jan 29, 2008 5:52 AM, Pavel Emelyanov <xemul@openvz.org> wrote:

> There's one place that works with task pids - its the "tasks" file
> in cgroups. The read/write handlers assume, that the pid values
> go to/come from the user space and thus it is a virtual pid, i.e.
> the pid as it is seen from inside a namespace.
>
> Tune the code accordingly.
>
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Acked-by: Paul Menage <menage@google.com>

>
> ---
> kernel/cgroup.c | 4 ++--
> 1 files changed, 2 insertions(+), 2 deletions(-)
>
> diff --git a/kernel/cgroup.c b/kernel/cgroup.c
> index 2c5cccb..4766bbb 100644
> --- a/kernel/cgroup.c
> +++ b/kernel/cgroup.c
> @@ -1269,7 +1269,7 @@ static int attach_task_by_pid(struct cgroup *cgrp, char *pidbuf)
>
> if (pid) {
> rcu_read_lock();
> - tsk = find_task_by_pid(pid);
> + tsk = find_task_by_vpid(pid);
> if (!tsk || tsk->flags & PF_EXITING) {
> rcu_read_unlock();
> return -ESRCH;
> @@ -1955,7 +1955,7 @@ static int pid_array_load(pid_t *pidarray, int npids, struct cgroup
*cgrp)
> while ((tsk = cgroup_iter_next(cgrp, &it))) {
> if (unlikely(n == npids))
> break;
> - pidarray[n++] = task_pid_nr(tsk);
> + pidarray[n++] = task_pid_vnr(tsk);
> }
> cgroup_iter_end(cgrp, &it);
> return n;
> --
> 1.5.3.4
>
>
