
Subject: Re: [RFC] [PATCH] cgroup: add "procs" control file
Posted by [Paul Menage](#) on Fri, 20 Jun 2008 05:37:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Jun 18, 2008 at 1:02 AM, Li Zefan <lizf@cn.fujitsu.com> wrote:
> - What to do if the attaching of a thread failed? continue to attach
> other threads, or stop and return error?

I think this is something that will have to be handled in the design
of transactional cgroup attach.

> - When a sub-thread of a process is in the cgroup, but not its thread
> cgroup leader, what to do when 'cat procs'? just skip those threads?

Sounds reasonable. I think that in general the procs file is more
useful as a write API than a read API anyway, for the reasons you
indicate there.

```
> +   tsk = attach_get_task(cgrp, pidbuf);  
> +   if (IS_ERR(tsk))  
> +       return PTR_ERR(tsk);  
> +  
> +   /* attach thread group leader */
```

Should we check that this is in fact a thread group leader?

```
> +  
> +   /* attach all sub-threads */  
> +   rcu_read_lock();
```

cgroup_attach_task() calls synchronize_rcu(), so it doesn't seem
likely that rcu_read_lock() is useful here, and might even deadlock?

What are you trying to protect against with the RCU lock?

```
>   {  
> +       .name = "procs",
```

Maybe call it "cgroup.procs" to avoid name clashes in future? We had a
debate a while back where I tried to get the cgroup files like "tasks"
and "notify_on_release" prefixed with "cgroup." , which were argued
against on grounds of backwards compatibility. But there's no
compatibility issue here. The only question is whether it's too ugly
to have the legacy filenames without a prefix and the new ones with a
prefix.

Paul

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