Subject: Re: cgroup tasks file error
Posted by ccmail111 on Tue, 14 Dec 2010 07:28:19 GMT
View Forum Message <> Reply to Message

Thanks Matt.
After doing what Serge suggests, I was able to create a dummy cgroup and
move task (#580) below to the group. But now..
I see another issue: cannot move the task back to its parent (root), from the group (hello) as in:

I am trying to move back task #580..

[host:/dev/cgroup]$ id
uid=0(root) gid=0(root) groups=0(root)

[host:/dev/cgroup]$ echo 580 > tasks
-bash: echo: write error: Operation not permitted

[host:/dev/cgroup]$ cat hello/tasks
580
610
2104
[host:/dev/cgroup]$

--- On Mon, 12/13/10, Matt Helsley <matthltc@us.ibm.com> wrote:

> From: Matt Helsley <matthltc@us.ibm.com>
> Subject: Re: cgroup tasks file error
> To: "Serge E. Hallyn" <serge.hallyn@canonical.com>
> Cc: "ccmail111" <ccmail111@yahoo.com>, containers@lists.linux-foundation.org
> Date: Monday, December 13, 2010, 8:02 PM
> On Mon, Dec 13, 2010 at 05:16:28PM
> -0600, Serge E. Hallyn wrote:
> > Quoting ccmail111 (ccmail111@yahoo.com):
> > >
> > > I see error:[host:/dev/cgroup]$ echo 693 >
> hello-test/tasks
> > > -bash: echo: write error: No space left on
> device
>
> This does seem quite odd so I spent a little time looking
> at this and I agree with Serge.
>
> > > [host:/dev/cgroup]$ pwd/dev/cgroup
> > >
> > > But the user process is up and running..
> > >

> > > [host:/dev/cgroup]$ ps aux | grep procroot
>
> > >
> > > 693 0.0 0.4 34720 1112 ttyS0   Sl
> 19:11   0:00 /opt/bin/myproc -ext
> > >
> > > Also the cgroup exists and valid..
> > >
> > > [host:/dev/cgroup]$ ls | grep hello-test
> > > hello-test
> > >
> > > What above error mean and any suggestions ?
> > > Please email.
> >
> > Which cgroups do you have composed on that
> mount? I'm guess you
> > have cpuset, and you need to set the cpuset.mems and
> cpuset.cpus.
> > Until you do that, no tasks can be assigned to it.
>
> I looked a a few places in kernel/cgroup.c which return
> ENOSPC
> or could potentially forward such an error. The only place
> that
> fits is in the attach path and is consistent with the
> notion that
> it's a cpuset issue:
>
> echo <pid> > tasks =>
> cgroup_tasks_write() =>
> attach_task_by_pid() =>
> cgroup_attach_task() => (via ss->can_attach() where
> ss is the cpuset subsystem)
> cpuset_can_attach():
>   if
> (cpumask_empty(cs->cpus_allowed) ||
> nodes_empty(cs->mems_allowed))
>
> return -ENOSPC;
>
> No other cgroup subsystem that I looked at (freezer,
> memcontrol, ns,
>   blkio, devcgroup) returns ENOSPC when
> attaching a task.
>
> So not only do you need to set those masks but each mask
> must have at
> least one cpu and "mem" respectively.

>
> Cheers,
> -Matt Helsley
>

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containe rs