

---

Subject: Re: cgroup tasks file error  
Posted by [ccmail111](#) on Tue, 14 Dec 2010 17:33:18 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Isn't ns mounted by default ?  
I rebooted machine,  
based on 2.6.32 kernel.

Then,

```
[host:~]$ mkdir /dev/cgroup
[host:~]$ mount -t cgroup cpuset -ocpuset,ns /dev/cgroup
[host:~]$ ps aux | grep libvirt
root    575  0.6  0.9 88816 2216 ?      Sl  17:28   0:00 myproc --daemon
root    654  0.0  0.2  4044  664 ttyS0   S+  17:29   0:00 grep libvirt
[host:~]$ mkdir /dev/cgroup/hello
[host:~]$ cat /dev/cgroup/hello/tasks
[host:~]$ echo 0 > /dev/cgroup/hello/cpuset.cpus
[host:~]$ echo 0 > /dev/cgroup/hello/cpuset.mems
[host:~]$ echo 575 > /dev/cgroup/hello/tasks
[host:~]$ cat /dev/cgroup/hello/tasks
575
[host:~]$ id
uid=0(root) gid=0(root) groups=0(root)
[host:~]$
```

Now, I try to move the task out - the parent/root:

```
[host:~]$ echo 575 > /dev/cgroup/tasks
-bash: echo: write error: Operation not permitted
```

Any thoughts !?

--- On Tue, 12/14/10, Jue Hong <[honshj@gmail.com](mailto:honshj@gmail.com)> wrote:

```
> From: Jue Hong <honshj@gmail.com>
> Subject: Re: cgroup tasks file error
> To: "ccmail111" <ccmail111@yahoo.com>
> Cc: "Serge E. Hallyn" <serge.hallyn@canonical.com>, "Matt Helsley" <matthlhc@us.ibm.com>,
containers@lists.linux-foundation.org
> Date: Tuesday, December 14, 2010, 4:39 AM
> Have you mounted the ns subsys?
> The ns subsys will place some restrictions on the task
> movement.
```

```

>
> On Tue, Dec 14, 2010 at 3:28 PM, ccm111 <ccmail111@yahoo.com>
> wrote:
> >
> > 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 <matthlhc@us.ibm.com>
> > wrote:
> >
> >
> >> From: Matt Helsley <matthlhc@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 ccm111 (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/containers
> >
>
>
>
> --
>
> Department of Network Infrastructure
> Tencent Company, Shenzhen, China
>

```

---

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

---