
Subject: Re: [RFC] Prefixing cgroup generic control filenames with "cgroup."
Posted by [Paul Menage](#) on Thu, 28 Feb 2008 22:26:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Feb 28, 2008 at 2:21 PM, Andrew Morton

<akpm@linux-foundation.org> wrote:

> On Thu, 28 Feb 2008 14:06:30 -0800

>

> "Paul Menage" <menage@google.com> wrote:

>

>

>

> > On Thu, Feb 28, 2008 at 1:40 PM, Andrew Morton

> > <akpm@linux-foundation.org> wrote:

> > >

> > > Maybe cgroups shouldn't be putting kernel-generated files in places where

> > > user-specified files appear?

> > >

> >

> > Well, that API (mixing control files and group directories in the same
> > directory namespace) was inherited directly from cpusets.

> >

> > It wouldn't be hard to throw that away and move all the user-created

> > group directories into their own subdirectory, i.e. change the

> > existing directory layout from something like:

> >

> > /mnt/cgroup/

> > tasks

> > cpu.shares

> > memory.limit_in_bytes

> > memory.usage_in_bytes

> > user_created_groupname1/

> > tasks

> > cpu.shares

> > memory.limit_in_bytes

> > memory.usage_in_bytes

> > user_created_groupname2/

> > tasks

> > cpu.shares

> > memory.limit_in_bytes

> > memory.usage_in_bytes

> >

> > to something like:

> >

> > /mnt/cgroup/

> > tasks

> > cpu.shares

> > memory.limit_in_bytes

```
> > memory.usage_in_bytes
> > groups/
> >   user_created_groupname1/
> >     tasks
> >     cpu.shares
> >     memory.limit_in_bytes
> >     memory.usage_in_bytes
> >     groups/
> >   user_created_groupname2/
> >     tasks
> >     cpu.shares
> >     memory.limit_in_bytes
> >     memory.usage_in_bytes
> >     groups/
>
> That looks nice.
>
```

I'll put a patch together for consideration.

```
> Could we do something like auto-prefixing user-created directories with a
> fixed string so that there is no way in which the user can cause a
> collision with kernel-created files?
```

That's something like putting them all in their own sub-directory, but sounds less clean to me.

```
>
> I suppose that would break cpusets back-compatibility as well? If so, we
> could do the prefixing only for non-cpusets directories, but that's getting
> a bit weird.
```

We already have something like that in place, actually.

When you mount the legacy "cpuset" filesystem, it just passes through to the cgroup filesystem with the mount options "cpuset,noprefix", i.e. mount a cgroups hierarchy with just cpusets bound to it, and *don't* prefix subsystem control files with the subsystem name. It wouldn't be hard to also have a "nosubdir" mount option that keeps the existing single-level style, have the cpuset filesystem pass that option.

Paul

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