
Subject: Re: [RFC] Prefixing cgroup generic control filenames with "cgroup."
Posted by [Paul Menage](#) on Fri, 29 Feb 2008 01:03:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Feb 28, 2008 at 3:36 PM, Paul Jackson <pj@sgi.com> wrote:

- >
- > We could accomplish that much by decreeing that future new kernel
- > generated names that we might add follow some stronger convention,
- > such as the cgroup_ or appropriate subsystem prefix.

Subsystem-created files already have an appropriate prefix.

- > No need to
- > change the existing well known names for this reason.

But that's part of my point - is it reasonable to describe a system that was only introduced in 2.6.24 as "well-known"?

- >
- > Actually, in terms of 'common names used
- > by humans' some of these names, "tasks" and "notify_on_release", date
- > back much earlier than that. Please don't rename these two files in
- > cgroups; and of course absolutely don't rename them in cpusets.

No, I wasn't planning to make any changes to cpusets.

- >
- > Please don't end up with different names of these files, depending on
- > whether you're in cgroups or cpusets, either.

That already happens - when mounted as the "cpuset" filesystem, we have names like "mems_allowed". When mounted as cgroups, we have names like cpuset.mems_allowed.

- >
- > > 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?
- >
- > Lordy lordy -- a bunch of intrusive, complicating crap to solve a
- > non-existent problem (sorry for the indelicate choice of words ;).

No, I don't like that idea either.

Paul

Containers mailing list
Containers@lists.linux-foundation.org

