
Subject: Re: [RFC] Prefixing cgroup generic control filenames with "cgroup."
Posted by [Paul Menage](#) on Thu, 28 Feb 2008 22:06:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Feb 28, 2008 at 1:33 PM, <serge@hallyn.com> wrote:

- >
- > You said the set of files belong to cgroup itself is likely to increase
- > - do you have some candidates in mind?

Nothing concrete right now. One example that I already proposed was the "cgroup.api" file but that's shelved for now, until such time as I actually propose the binary API that it was intended to help support.

- > Perhaps ones which are likely
- > to conflict with choice taskgroup names?

It's hard to determine what likely taskgroup names would be. For my own use, pretty much every group has a unique 64-bit ID in the name so this isn't something I worry about directly affecting our systems. But I don't know what other users might want to do.

- >
- > If anything I'd say add a 'prefix_cgroup' mount option and use it to
- > decide whether to prefix or not (rather than use the noprefix option).

The existing "noprefix" option controls whether the *subsystems* get prefixes. It's mainly there to provide backwards compatibility for cpusets. Existing cpusets clients would be expecting to find files with names like "mems_allowed" rather than "cpuset.mems_allowed". So mounting with the "noprefix" option (which happens automatically when you mount the "cpuset" filesystem wrapper) gives the same result as before.

Currently "noprefix" has no effect on cgroup files, since they never have a prefix anyway.

Yes, we could add a new mount option "prefixcgroup", and let people decide which they want. But I still don't see any argument *against* doing the prefixing automatically (rather than an argument that it's no better or worse) other than wanting 2.6.24 compatibility.

Paul

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