Subject: Re: [RFC] Prefixing cgroup generic control filenames with "cgroup."
Posted by Paul Menage on Mon, 03 Mar 2008 09:11:51 GMT
View Forum Message <> Reply to Message

So, there have been various options suggested over the course of this thread:

--

1) no code changes, just stake out all names matching a certain regexp
(e.g. "[a-z].*") as being potentially used by the kernel in the
future; document this, and let users who are worried about name
clashes avoid these names

pros: no work involved, avoids potentially complex changes to solve a
possibly non-problem.

cons: leaves an intermingled namespace; since this would be a
convention rather than an enforced rule, users might be unaware that
they're setting themselves up for a fall

--

2) separate out the kernel-generated names and user-generated names by
putting the user-generated names in a "groups" sub-directory (can be a
mount option that's automatically disabled for cpusets).

pros: completely solves problem of intermingled namespaces; makes it
easier to see sub-groups at a glance

cons: extra code, slightly more awkward to deal with in the general
case, is incompatible with the code that was in mainline in the brief
period of time since 2.6.24 was finalized.

--

3) prefix all cgroup-provided files with "cgroup."

pros: hardly any extra code; mostly solves the namespace problem since
user names are much less likely to begin with "cgroup."

cons: changes name of "tasks" file visible in 2.6.24; doesn't help if
a future new subsystem is introduced with a commonly-used name that
might clash with existing user-generated names.

--

4) prefix all future cgroup files with "cgroup." (possibly including
existing notify_on_release and release_agent files?)

pros: no extra code, just involves slightly longer strings for future
new files; no incompatibility issues

cons: ugly inconsistency between new cgroup files and grandfathered
old ones, plus same clash problems as option 3

--

Paul

_____

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