
Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by [Randy Dunlap](#) on Tue, 11 Sep 2007 15:22:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 11 Sep 2007 09:20:33 +0200 Cedric Le Goater wrote:

> Paul Menage wrote:
> > On 9/10/07, Dmitry Adamushko <dmitry.adamushko@gmail.com> wrote:
> >> On 10/09/2007, Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com> wrote:
> >>> On Mon, Sep 10, 2007 at 10:22:59AM -0700, Andrew Morton wrote:
> >>>> objection ;) "cpuctlr" isn't memorable. Kernel code is write-rarely,
> >>>> read-often. "cpu_controller", please. The extra typing is worth it ;)
> >>> Ok! Here's the modified patch (against 2.6.23-rc4-mm1).
> >> as everyone seems to be in a quest for a better name... I think, the
> >> obvious one would be just 'group_sched'.
> >>
> >
> > But "sched" on its own could refer to CPU scheduling, I/O scheduling,
> > network scheduling, ...
> >
> > And "group" is more or less implied by the fact that it's in the
> > containers/control groups filesystem.
>
> "control groups" is the name of your framework. right ?
>
> > So "group_sched" isn't really all that informative. The name should
> > definitely contain either "cpu" or "cfs".
>
> "cfs" control group subsystem.

That looks odd, like it's a filesystem.
What does cfs really mean?

> "cfs" looks good enough to identify the subsystem.

~Randy

*** Remember to use Documentation/SubmitChecklist when testing your code ***

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