Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by Paul Menage on Mon, 10 Sep 2007 22:45:15 GMT
View Forum Message <> Reply to Message

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.

So "group_sched" isn't really all that informative. The name should definitely contain either "cpu" or "cfs".

Paul

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

Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by Cedric Le Goater on Tue, 11 Sep 2007 07:20:33 GMT
View Forum Message <> Reply to Message

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.

"cfs" looks good enough to identify the subsystem.

C.

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

---

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

---

Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by Jan Engelhardt on Tue, 11 Sep 2007 15:37:49 GMT
View Forum Message <> Reply to Message

On Sep 11 2007 08:22, Randy Dunlap wrote:
>> "cfs" control group subsystem.
>
>That looks odd, like it's a filesystem.
>What does cfs really mean?

http://en.wikipedia.org/wiki/CFS
(scnr)

It misses the C...something Filesystem tho.

Jan
--

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

---

Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by Srivatsa Vaddagiri on Tue, 11 Sep 2007 15:43:02 GMT
View Forum Message <> Reply to Message

On Tue, Sep 11, 2007 at 08:22:43AM -0700, Randy Dunlap wrote:
> That looks odd, like it's a filesystem.
> What does cfs really mean?

cfs = completely fair scheduler :)

In this thread, we are talking of hooking the cfs cpu scheduler with the task-container framework in -mm tree, so that the scheduler can deal with groups of tasks rather than just tasks, while handling fairness of cpu allocation.

I agree "cfs" control subsystem does look odd a bit here. "cpu" control subsystem seems better.

--
Regards,
vatsa

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

---

Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by Paul Menage on Tue, 11 Sep 2007 15:44:47 GMT
View Forum Message <> Reply to Message

On 9/11/07, Cedric Le Goater <clg@fr.ibm.com> wrote:
> >
> > 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 ?

That's the main contender for the new name, to replace "task containers" since there were enough objections. (Renamed patches coming out today, I hope).

Paul

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

---

Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by Randy Dunlap on Tue, 11 Sep 2007 15:53:01 GMT

On Tue, 11 Sep 2007 21:21:19 +0530 Srivatsa Vaddagiri wrote:

> On Tue, Sep 11, 2007 at 08:22:43AM -0700, Randy Dunlap wrote:
> > That looks odd, like it's a filesystem.
> > What does cfs really mean?
>
> cfs = completely fair scheduler :)
>
> In this thread, we are talking of hooking the cfs cpu scheduler with the
> task-container framework in -mm tree, so that the scheduler can deal
> with groups of tasks rather than just tasks, while handling fairness of
> cpu allocation.
>
> I agree "cfs" control subsystem does look odd a bit here. "cpu" control
> subsystem seems better.

Thanks. I agree that using "cpu" is better. I.e., don't tie it
to a particular scheduler name. It would just need to change the
next time we have a new scheduler. ;)

~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