
Subject: Re: containers (was Re: -mm merge plans for 2.6.23)

Posted by [akpm](#) on Wed, 11 Jul 2007 05:29:42 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, 11 Jul 2007 10:25:16 +0530 Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com> wrote:

> On Tue, Jul 10, 2007 at 11:53:19AM -0700, Andrew Morton wrote:

> > On Tue, 10 Jul 2007 11:34:38 -0700

> > "Paul Menage" <menage@google.com> wrote:

> >

> > > Andrew, how about we merge enough of the container framework to

> > > support CFS? Bits we could leave out for now include container_clone()

> > > support and the nsproxy subsystem, fork/exit callback hooks, and

> > > possibly leave cpusets alone for now (which would also mean we could

> > > skip the automatic release-agent stuff). I'm in Tokyo for the Linux

> > > Foundation Japan symposium right now, but I should be able to get the

> > > new patchset to you for Friday afternoon.

> >

> > mm.. Given that you propose leaving bits out for the 2.6.23 merge, and

> > that changes are still pending and that nothing will _use_ the framework in

> > 2.6.23 [...]

>

> Andrew,

> The cpu group scheduler is ready and waiting for the container patches

> in 2.6.23 :)

>

> Here are some options with us:

>

> a. (As Paul says) merge enough of container patches to enable

> its use with cfs group scheduler (and possibly cpusets?)

>

> b. Enable group scheduling bits in 2.6.23 using the user-id grouping

> mechanism (aka fair user scheduler). For 2.6.24, we could remove

> this interface and use Paul's container patches instead. Since this

> means change of API interface between 2.6.23 and 2.6.24, I don't

> prefer this option.

>

> c. Enable group scheduling bits only in -mm for now (2.6.23-mmX), using

> Paul's container patches. I can send you a short patch that hooks up

> cfs group scheduler with Paul's container infrastructure.

>

> If a. is not possible, I would prefer c.

>

> Let me know your thoughts ..

I'm inclined to take the cautious route here - I don't think people will be dying for the CFS thingy (which I didn't even know about?) in .23, and it's rather a lot of infrastructure to add for a CPU scheduler configurator

gadget (what does it do, anyway?)

We have plenty of stuff for 2.6.23 already ;)

Is this liveable with??

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
