
Subject: Re: [PATCH 1/1, v9] cgroup/freezer: add per freezer duty ratio control
Posted by [Matt Helsley](#) on Wed, 16 Feb 2011 03:23:21 GMT

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

On Tue, Feb 15, 2011 at 04:38:12PM -0800, Jacob Pan wrote:

> On Wed, 16 Feb 2011 01:00:15 +0100

> "Rafael J. Wysocki" <rjw@sisk.pl> wrote:

<snip>

> > Well, quite frankly, I'm not going to take this patch unless it gets

> > an ACK from the scheduler people (which I'm guessing is not going to

> > happen before hell freezes over).

> >

> > IOW, please find a better way to address the issue at hand.

> >

> We do have a real need that there is no exist feature in the kernel can

> provide solution for. You want ACK from scheduler people because it has

> impact on disabling irq? or you think scheduler should be the one that

> provide the solution. I did try cpu subsystem, but it seems to be

> limited to RT and certain scheduling policy e.g. RR and FIFO.

I agree with Rafael. I think the scheduler should provide the solution

and it can be done via modifications to the cpu cgroup subsystem.

Yes, it only has the shares and rt-related files *right now*. However,

Kame replied earlier with a link to some patches for extending it

that introduce files with similar (granted: not the same) semantics:

<http://linux.derkeiler.com/Mailing-Lists/Kernel/2010-10/msg04335.html>

I think you could build on that and help push those patches towards
mainline along with your enhancements for reducing wakeups rather
than modify the freezer cgroup subsystem.

Cheers,

-Matt Helsley

Containers mailing list

Containers@lists.linux-foundation.org

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

Subject: Re: [PATCH 1/1, v9] cgroup/freezer: add per freezer duty ratio control
Posted by [jacob.jun.pan](#) on Wed, 16 Feb 2011 18:18:05 GMT

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

On Tue, 15 Feb 2011 19:23:21 -0800

Matt Helsley <matthltc@us.ibm.com> wrote:

> On Tue, Feb 15, 2011 at 04:38:12PM -0800, Jacob Pan wrote:
> > On Wed, 16 Feb 2011 01:00:15 +0100
> > "Rafael J. Wysocki" <rjw@sisk.pl> wrote:
>
> <snip>
>
> > > Well, quite frankly, I'm not going to take this patch unless it
> > > gets an ACK from the scheduler people (which I'm guessing is not
> > > going to happen before hell freezes over).
> > >
> > > IOW, please find a better way to address the issue at hand.
> > >
> > We do have a real need that there is no exist feature in the kernel
> > can provide solution for. You want ACK from scheduler people
> > because it has impact on disabling irq? or you think scheduler
> > should be the one that provide the solution. I did try cpu
> > subsystem, but it seems to be limited to RT and certain scheduling
> > policy e.g. RR and FIFO.
>
> I agree with Rafael. I think the scheduler should provide the solution
> and it can be done via modifications to the cpu cgroup subsystem.
> Yes, it only has the shares and rt-related files *right now*. However,
> Kame replied earlier with a link to some patches for extending it
> that introduce files with similar (granted: not the same) semantics:
>
> <http://linux.derkeiler.com/Mailing-Lists/Kernel/2010-10/msg04335.html>
>
> I think you could build on that and help push those patches towards
> mainline along with your enhancements for reducing wakeups rather
> than modify the freezer cgroup subsystem.
>
I will give CFS bandwidth control patches a try. See if I can be any
help. Thanks.

Jacob

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