
Subject: Re: [RFC][PATCH 0/6] Add group fairness to CFS - v1
Posted by [Srivatsa Vaddagiri](#) on Fri, 15 Jun 2007 14:06:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, Jun 15, 2007 at 04:46:13PM +0400, Kirill Korotaev wrote:

> Ingo Molnar wrote:

>

> > i'd still like to hear back from Kirill & co whether this framework is

> > flexible enough for their work (OpenVZ, etc.) too.

>

> My IMHO is that so far the proposed group scheduler doesn't look ready/suitable.

Hi Kirill,

Yes its work-in-progress and hence is not ready/fully-functional (yet). The patches I posted last gives an idea of the direction it is heading. For ex: <http://lkml.org/lkml/2007/6/11/162> and <http://lkml.org/lkml/2007/5/25/146> gives an idea of how SMP load balance will work.

IMHO the nice thing about this approach is it (re)uses lot of code in scheduler which is there already to achieve fairness between tasks and higher schedulable elements (users/containers etc).

Also with CFS engine's precise nanosecond accurate accounting and time-sorted list of tasks/entities, I feel we will get much tighter control over distribution of CPU between tasks/users/containers.

> We need to have a working SMP version before it will be clear
> whether the whole approach is good and works correct on variety of load patterns.

If you have any heads up thoughts on areas/workloads where this may pose problems for container/user scheduling, I would be glad to hear them. Otherwise I would greatly wellcome any help in developing/reviewing these patches which meets both our goals!

--

Regards,
vatsa

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