Subject: Re: [RFC] [PATCH 0/3] Add group fairness to CFS
Posted by William Lee Irwin III on Wed, 23 May 2007 18:23:27 GMT
View Forum Message <> Reply to Message

On Wed, May 23, 2007 at 08:12:12PM +0200, Guillaume Chazarain wrote:
> o Doing an infinite loop as root seems to badly affect interactivity
> much more than with a normal user. Note that this is subjective, so
> maybe I'm smocking crack here.

CPU hogs as a distinct user will end up entitled to a whole user's
share of CPU bandwidth. That's just what it implements.

On Wed, May 23, 2007 at 08:12:12PM +0200, Guillaume Chazarain wrote:
> o Nice values are not reflected across users. From my test, if user1
> has a single busy loop at nice 19, and user2 a single busy loop at
> nice 0, both process will have a 50% CPU share, this looks wrong. Note
> that I have no idea how to solve this one.

This depends on how nice levels nest with task group weighting. If
they nest within task group weightings, the shares of CPU bandwidth
for single tasks having their own dedicated task groups will not
vary with nice number. It's harder to do the other way around with
hierarchical scheduling, but easy if the task groups merely determine
task weights.

-- wli

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