Subject: Re: [RFC] [PATCH 0/3] Add group fairness to CFS
Posted by Ingo Molnar on Wed, 23 May 2007 18:38:24 GMT
View Forum Message <> Reply to Message

* Guillaume Chazarain <guichaz@yahoo.fr> wrote:

> Some observations:
>
> 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.

hm, this shouldnt be the case. Can you see this with -v14?

> 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.

for containers it's exactly the right behavior: group scheduling is really a 'super' container concept that allows the allocation of CPU time regardless of how a group uses it. The only additional control we might want is to allocate different amount of CPU time to different groups. (i.e. a concept vaguely similar to "nice levels", but at the group level - using a different and saner API than nice levels.) Nice levels are really only meaningful at the lowest level.

for 'friendly users' it's perhaps not what we want - but those do not need to isolate themselves from each other anyway.

> Thanks for working in this very interesting direction.

seconded :)

Ingo

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