
Subject: A question about group CFS scheduling
Posted by [Zhao Forrest](#) on Thu, 26 Jun 2008 07:19:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi experts,

In Documentation/sched-design-CFS.txt it reads:
Group scheduler tunables:

When CONFIG_FAIR_USER_SCHED is defined, a directory is created in sysfs for each new user and a "cpu_share" file is added in that directory.

```
# cd /sys/kernel/uids
# cat 512/cpu_share      # Display user 512's CPU share
1024
# echo 2048 > 512/cpu_share  # Modify user 512's CPU share
# cat 512/cpu_share      # Display user 512's CPU share
2048
#
```

CPU bandwidth between two users are divided in the ratio of their CPU shares.
For ex: if you would like user "root" to get twice the bandwidth of user "guest", then set the cpu_share for both the users such that "root"'s cpu_share is twice "guest"'s cpu_share.

My question is: how is CPU bandwidth divided between cgroup and regular processes?

For example,

- 1 the cpu_share of user "root" is set to 2048
- 2 the cpu_share of user "guest" is set to 1024
- 3 there're many processes owned by other users, which don't belong to any cgroup

if the relative CPU bandwidth allocated to cgroup of "root" is 2,
allocated to cgroup of "guest" is 1, then what's the relative CPU
bandwidth allocated to other regular processes? 2 or 1?

Thanks,
Forrest

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