

Hi,

```
>          uid "vatsa"          uid "guest"
>          (make -s -j4 bzImage) (make -s -j20 bzImage)
>
> 2.6.22-rc1          772.02 sec      497.42 sec (real)
> 2.6.22-rc1+cfs-v13      780.62 sec      478.35 sec (real)
> 2.6.22-rc1+cfs-v13+this patch  776.36 sec      776.68 sec (real)
```

Impressive numbers :-)

Testing this in qemu/UP/i386, I had to do this:

```
--- linux/kernel/sched_fair.c
+++ linux/kernel/sched_fair.c
@@ -350,9 +350,10 @@ @@

    if (p->wait_start_fair) {
        delta_fair = lrq->fair_clock - p->wait_start_fair;
-   if (unlikely(p->load_weight != lrq->nice_0_load))
-       delta_fair = (delta_fair * p->load_weight) /
-       lrq->nice_0_load;
+   if (unlikely(p->load_weight != lrq->nice_0_load)) {
+       s64 m = delta_fair * p->load_weight;
+       delta_fair = do_div(m, lrq->nice_0_load);
+   }
    add_wait_runtime(lrq, p, delta_fair);
}
```

to make it compile, otherwise it ends with:

```
kernel/built-in.o: In function `update_stats_wait_end':
/home/g/linux-group-fair/linux-2.6.21-rc1-cfs-v13-fair/kernel/sched_fair.c:354:
undefined reference to `__divdi3'
/home/g/linux-group-fair/linux-2.6.21-rc1-cfs-v13-fair/kernel/sched_fair.c:354:
undefined reference to `__divdi3'
```

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.

- 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.

Thanks for working in this very interesting direction.

--

Guillaume

Containers mailing list

Containers@lists.linux-foundation.org

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