
Subject: Re: [ckrm-tech] [RFC][PATCH 1/6] Introduce struct sched_entity and struct
IRQ

Posted by [Srivatsa Vaddagiri](#) on Tue, 12 Jun 2007 03:52:45 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, Jun 12, 2007 at 07:45:59AM +0530, Balbir Singh wrote:

```
> > /* CFS-related fields in a runqueue */
> > +struct IRQ {
> > + unsigned long raw_weighted_load;
> > + #define CPU_LOAD_IDX_MAX 5
> > + unsigned long cpu_load[CPU_LOAD_IDX_MAX];
> > + unsigned long nr_load_updates;
> > +
> > + u64 fair_clock, delta_fair_clock;
> > + u64 exec_clock, delta_exec_clock;
> > + s64 wait_runtime;
> > + unsigned long wait_runtime_overruns, wait_runtime_underruns;
> > +
> > + struct rb_root tasks_timeline;
> > + struct rb_node *rb_leftmost;
> > + struct rb_node *rb_load_balance_curr;
> > +};
> > +
>
> Shouldn't the rq->lock move into IRQ?
```

Right now, the per-cpu rq lock protects all (local) runqueues attached with the
cpu. At some point, for scalability reasons, we may want to split that to
be per-cpu per-local runqueue (as you point out). I will put that in my todo
list of things to consider. Thanks for the review!

--

Regards,
vatsa

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