

---

Subject: Re: Re: [RFC] Virtualization steps

Posted by [Herbert Poetzl](#) on Wed, 29 Mar 2006 13:45:24 GMT

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

---

On Wed, Mar 29, 2006 at 01:13:14PM +0400, Kirill Korotaev wrote:

> Sam,

>

> >>Why do you think it can not be measured? It either can be, or it is too

> >>low to be measured reliably (a fraction of a per cent or so).

> >

> >Well, for instance the fair CPU scheduling overhead is so tiny it may as

> >well not be there in the VServer patch. It's just a per-vserver TBF

> >that feeds back into the priority (and hence timeslice length) of the

> >process. ie, you get "CPU tokens" which deplete as processes in your

> >vserver run and you either get a boost or a penalty depending on the

> >level of the tokens in the bucket. This doesn't provide guarantees, but

> >works well for many typical workloads.

> I wonder what is the value of it if it doesn't do guarantees or QoS?

> In our experiments with it we failed to observe any fairness.

probably a misconfiguration on your side ...

> So I suppose the only goal of this is too make sure that maliscuios

> user want consume all the CPU power, right?

the currently used scheduler extensions do much

more than that, basically all kinds of scenarios

can be satisfied with it, at almost no overhead

> >How does your fair scheduler work?

> >Do you just keep a runqueue for each vps?

> we keep num\_online\_cpus runqueues per VPS.

> Fairs scheduler is some kind of SFQ like algorithm which selects VPS

> to be scheduled, than standart linux scheduler selects a process in a

> VPS runqueues to run.

>

> >To be honest, I've never needed to determine whether its overhead is 1%

> >or 0.01%, it would just be a meaningless benchmark anyway :-). I know

> >it's "good enough for me".

> Sure! We feel the same, but people like numbers :)

well, do you have numbers?

best,

Herbert

> Thanks,  
> Kirill

---