Subject: Re: Re: [RFC] Virtualization steps
Posted by Kirill Korotaev on Wed, 29 Mar 2006 09:13:14 GMT
View Forum Message <> Reply to Message

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. So I
suppose the only goal of this is too make sure that maliscuios user want
consume all the CPU power, right?

> 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 :)

Thanks,
Kirill