

---

Subject: Re: Re: [RFC] Virtualization steps  
Posted by [Sam Vilain](#) on Wed, 29 Mar 2006 21:37:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2006-03-29 at 18:47 +0400, Kirill Korotaev wrote:

> >> 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 ...  
> maybe you can provide some instructions on which kernel version to use  
> and how to setup the following scenario:  
> 2CPU box. 3 VPSs which should run with 1:2:3 ratio of CPU usage.

Ok, I'll call those three VPSes fast, faster and fastest.

"fast" : fill rate 1, interval 3  
"faster" : fill rate 2, interval 3  
"fastest" : fill rate 3, interval 3

That all adds up to a fill rate of 6 with an interval of 3, but that is right because with two processors you have 2 tokens to allocate per jiffie. Also set the bucket size to something of the order of HZ.

You can watch the processes within each vserver's priority jump up and down with `vtop' during testing. Also you should be able to watch the vserver's bucket fill and empty in /proc/virtual/XXX/sched (IIRC)

> > well, do you have numbers?  
> just run the above scenario with one busy loop inside each VPS. I was  
> not able to observe 1:2:3 cpu distribution. Other scenarios also didn't  
> showed my any fairness. The results were different. Sometimes 1:1:2,  
> sometimes others.

I mentioned this earlier, but for the sake of the archives I'll repeat - if you are running with any of the buckets on empty, the scheduler is imbalanced and therefore not going to provide the exact distribution you asked for.

However with a single busy loop in each vserver I'd expect the above to yield roughly 100% for fastest, 66% for faster and 33% for fast, within 5 seconds or so of starting those processes (assuming you set a bucket size of HZ).

Sam.

---