
Subject: Re: Re: [RFC] Virtualization steps
Posted by [dev](#) on Wed, 12 Apr 2006 08:22:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sam,

> 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)
>
> 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, what we observe is the situation, when Linux cpu scheduler spreads 2 tasks on 1st CPU and 1 task on the 2nd CPU. Std linux scheduler doesn't do any rebalancing after that, so no plays with tokens make the spread to be 3:2:1, since the lowest priority process gets a full 2nd CPU (100% instead of 33% of CPU).

Where is my mistake? Can you provide a configuration where we could test or the instructions on how to avoid this?

Thanks,
Kirill
