
Subject: Running procs in VPS context :: possible to limit/throttle?

Posted by [HaroldB](#) on Wed, 16 Aug 2006 10:15:12 GMT

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

Hello, the cpulimit resource does a fine job, but lets say a vps has 20 programs trying to run at once. They will use all of the cpulimit and beause they take longer to complete, the number of average running procs will be much higher than normal, which causes the hardware server's load to be higher.

So, it seems to me that cpulimit can hurt overal system performance under certain circumstances (more than a few concurrent, long-running processes per ve).

Is it possible to somehow limit the number of running procs in a ve? Am I missing something, like a simple workaround the problem stated above.

Thank you
Harold

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [dim](#) on Wed, 16 Aug 2006 15:07:06 GMT

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

```
# man vzctl  
# vzctl set VEID --numproc 80:80 --save
```

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [scythe](#) on Wed, 16 Aug 2006 15:09:15 GMT

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

You mean --numproc parameter of vzctl?

EDIT: hehe, 2 seconds late, dim won

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [HaroldB](#) on Wed, 16 Aug 2006 15:14:51 GMT

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

I am talking about running procs, not total procs.

total procs is useful but running procs can take down a system. Is it possible to limit running procs?

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [dim](#) on Wed, 16 Aug 2006 15:21:16 GMT

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

I don't understand. There is only one task per CPU in the each moment, what do you mean?
Number of tasks, which are ready to be executed for each VE? But how does it help?

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [HaroldB](#) on Wed, 16 Aug 2006 15:23:40 GMT

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

for example:

```
[root@server~]# vmstat
procs -----memory----- ---swap-- -----io----- --system-- ----cpu----
 r b swpd free buff cache si so bi bo in cs us sy id wa
 5 2 6654616 38736 137608 1754916 0 2 1 0 1 0 16 10 59 15
```

5 running procs in this server. it has less than 5 cpus.

make sense now? how can we limit the number of running procs inside a vps.

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [dim](#) on Wed, 16 Aug 2006 15:26:32 GMT

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

5 tasks with state RUNNING in scheduler queue.
It doesn't make sense to limit this queue.

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [HaroldB](#) on Wed, 16 Aug 2006 15:28:36 GMT

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

say a vps tries to run 100 concurrent processes... can openvz limit/throttle this?

like I said earlier, setting the cpulimit lower will only increase the number of running procs because they take longer to complete.

Do you agree on this point?

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [dim](#) on Wed, 16 Aug 2006 15:35:48 GMT

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

HaroldB wrote on Wed, 16 August 2006 11:28say a vps tries to run 100 concurrent processes... can openvz handle/limit/throttle this?

like I said earlier, setting the cpulimit lower will only increase the number of running procs because they take longer to complete.

Do you agree on this point?

Agree.

There is cpuunits parameter also in OpenVZ. This parameter allow to set 'weight' of VE in respect of other VEs. If all VEs have equal weights, then one VE will have 1/number_of_ves CPU time no matter how much concurrent processes it runs. At the same time, if all VE's processes are in idle state, it's CPU time will be used by processes of other VEs.

Our scheduler operates in term of VEs, not processes.

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [HaroldB](#) on Wed, 16 Aug 2006 16:19:22 GMT

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

Thanks for the info, that makes sense. Here is my question another way.

Lets say a ve runs a program like "cpu-burn" (program which uses all available cpu cycles) once. If that process stays running, it would add a load of 1 to the hardware node because its exactly one running process. However, if that program is run 50 times concurrently, a load of 50 is added. All the while, the same amount of cpu time is used by the vps. maybe this is an inherent part of an operating system, I am not sure.

it seems like an easy way to DOS a openvz/vz hn, by running tons of procs in the vps, even though your cpu usage is the same. I have tested this on vmware and the hardware node feels no effect, because of the separate schedulers I believe.

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [aistis](#) on Wed, 16 Aug 2006 20:01:54 GMT

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

Harold, think for a second

cpulimit allows you to take care of CPU time the VE processes will get

cpuunits allows you to assign a priority for that CPU time the VE processes will get compared to other VEs (or in a words of dim)

now how you gonna DOS OVZ/VZ box? Are you afraid of tasks in the scheduler queue or what?

Subject: Re: Running procs in VPS context :: possible to limit/throttle?

Posted by [dev](#) on Wed, 16 Aug 2006 20:40:21 GMT

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

it is exactly the same for OVZ - we have 2-level CPU scheduler. On the first level VE to run is selected, on the second level a processes inside VE is selected.

So the first level CPU scheduler `_limits_` and `_makes_` sure that no VE can consume more CPU then it is allowed to. And VE is limited in CPU both due to `cpulimits` and its weight.

So what happens if you run 50 processes inside VE and 1 process on the node itself? Your `loadavg` on the node is 51 (`/proc/loadavg` or `uptime` output). But it has nothing to do with the `_real_` load. It has nothing to do with the real load even on std Linux. It is just one of the parameters. And if your VE is limited to 10% of CPU these 50 processes won't be able to get more then 10% of CPU altogether. And the ONLY! process in host system will be able to consume 90% of CPU regardless of those 50.