Subject: How can we vzctl 101 set --cpuunits > 500 000
Posted by XavM on Tue, 09 Apr 2013 09:42:24 GMT
View Forum Message <> Reply to Message

The context :

Our server has 16 physical Cores (32 reported by /proc/cpuinfo because of HyperThreading)

On this server, vzcpucheck reports : Power of the node: 3 359 610

We want to configure several small CTs and a big one on this server and make sure that CPU performances will be predictable.

To achieve this predictability, we want to set the values for cpuunits and cpulimit to the same relative amount on the "big" container;

This would guarantee we will always have the same amout of available CPU in this big CT, wether the others small CTs are using CPU or not.

This would also guarantee that n% of used CPU reported from within the CT always corresponds to the same time of physical CPU on the hardware node.

The Big CT should have 50% of the available CPU and we assumed the following command would be appropriate :

vzctl set 101 --cpuunits 1679805 --cpulimit 50 --save

(with 1 679 805 = Power of the node/2 = 3 359 610/2)

But we can't set anything bigger than 500 000 for cpuunits.

We have seen bug 861, but still don't understand why setting a value higher than 500 000 to cpuunits is not allowed.

The message 28790 in the forum, says cpuunits are "relative values" <- but relative to what ? Isn't it relative to "vzcpucheck Power of the node"

The Question :

How can we configure a CT to have a Min/Max (cpuunits / cpulimit) CPU of 50% of all the available CPU , when vzcpucheck reports a Power of the node to be 3 359 610

Regards,

Subject: Re: How can we vzctl 101 set --cpuunits > 500 000
Posted by Paparaciz on Thu, 11 Apr 2013 13:20:06 GMT

cpuunits is relative value.

if ct1 has 1000 cpuunits and ct2 had 3000 cpuunits it means only that ct3 gets 3times more cpu.

better use cpulimit and give number of cpu cores (--cpus) or give exactly cpu cores you want (--cpumask)

from man vzctl:

--cpuunits num
CPU weight for a container. Argument is positive non-zero number, passed to and used in the kernel fair scheduler. The larger the number is, the more CPU time this container gets. Maximum value is 500000, minimal is 8. Number is relative to weights of all the other running containers. If cpuunits are not specified, default value of 1000 is used.
You can set CPU weight for CT0 (host system itself) as well (use vzctl set 0 --cpuunits num). Usually, OpenVZ initscript (/etc/init.d/vz) takes care of setting this.

--cpus num
sets number of CPUs available in the container.

--cpumask cpus | all
sets list of allowed CPUs for the container. Input format is a comma-separated list of decimal numbers and ranges. Consecutively set bits are shown as two hyphen- separated decimal numbers, the smallest and largest bit numbers set in the range. For example, if you want the container to execute on CPUs 0, 1, 2, 7, you should pass 0-2,7. Default value is all (the container can execute on any CPU).

Subject: Re: How can we vzctl 101 set --cpuunits > 500 000
Posted by XavM on Fri, 12 Apr 2013 05:59:35 GMT

Thanks a lot.

It is now crystal clear

(
The "man vzctl" is much better explained than the wiki on this sublect :
//openvz.org/User_Guide/Managing_Resources#Managing_Containe r_CPU_resources
)