Subject: vzctl --wait problem

Posted by sodi on Thu, 10 May 2007 20:08:42 GMT

View Forum Message <> Reply to Message

Hi,

Sometimes

vzctl start 101 --wait 3

results in a kernel panic of the host, such that I have to reset the whole server.

This is not deterministic, sometimes it happens already on the first start and sometimes after some starts and stops.

It never happens without the wait option. As far as I can see there is some usleep() call which causes the crash. So this might also be a problem with the linux kernel of the host

Any ideas? Why? How to solve?

I compiled a linux 2.6.18 kernel with current patches. If the start works everything is fine.

The machine is a 8 core x86.

And yes, the --wait parameter is important since we want to execute some scripts after starting the ve, so I need a running linux.

Thanks for any help or hints.

Subject: Re: vzctl --wait problem

Posted by Vasily Tarasov on Fri, 11 May 2007 07:44:26 GMT

View Forum Message <> Reply to Message

Hello,

We need a kernel panic message. Please, post it here. Also specify the exact kernel version (we have several current in different series) and vzctl version.

Thank you, Vasily.

Subject: Re: vzctl --wait problem

Posted by sodi on Fri, 11 May 2007 08:07:38 GMT

View Forum Message <> Reply to Message

For compiling the kernel I followed the instruction of the wiki entry, so from kernel.org: linux-2.6.18.tar.bz2

with patch

patch-ovz028stab031.1-combined.gz

I also checked the .config settings against the entries mentioned in the wiki for kernel build.

Sorry for the kernel log, but it's a problem for me, since the machine is somewhere else and I use putty to login, so I can't really see the console, is there a file where the kernel messages could be found, I checked in /var/log, but can't see a file with kernel panic messages.

As far as I know there was something wrt "recursive" and something with "usleep(50000)" mentioned on the console. I'll ask for the console screen next time it'll happen.

Thanks for reply, I know kernel stack would be helpfull.

Subject: Re: vzctl --wait problem

Posted by kir on Fri, 11 May 2007 08:12:23 GMT

View Forum Message <> Reply to Message

See http://wiki.openvz.org/Remote_console_setup on the info how to obtain kernel panic messages.

Also, since the oops happening sporadically, you might want to test your hardware before: http://wiki.openvz.org/Hardware_testing

Subject: Re: vzctl --wait problem

Posted by sodi on Sat, 12 May 2007 18:37:30 GMT

View Forum Message <> Reply to Message

Thanks for the serial console/netconsole hints, I think this would be some last chance try (netconsole might work for me)

Also since the installation usual is done remotely and the KVM is not always on, it's only by chance to get a photo from the console.

Anyway, after changing the HPET support in kernel compile options to "no" and recompiling the kernel, I didn't get the kernel panic any more. May be this is a specific problem to a 2 Quad-Core Intel Xeon system. I don't think this is really an openVZ problem, but a general kernel problem. (It was the same .config except the HPET was disabled)

Since then I created several ve, stopped / started them w/o any problems.

Sorry, that I can't be more helpful with console information. I know that if there is a problem, it could generally improve openVZ, if the reason is known.

Subject: Re: vzctl --wait problem

View Forum Message <> Reply to Message

Hello,

Thank you for your information. I just compiled OpenVZ kernel with HPET turned on but unfortunately on my system (2 x Opteron 25) `vzctl start <veid> --wait` works smoothly.

Can you tell me the template you're using and attach here (as attachment!) the .config file you're using when `vzctl start <veid> --wait` fails.

Thank you for your help, Vasily

Subject: Re: vzctl --wait problem

Posted by sodi on Mon, 14 May 2007 16:21:21 GMT

View Forum Message <> Reply to Message

The configuration for the "working" kernel is appended (I didn't change most of the default options). I assume at least that it's working, since I didn't have any problems with it by now. With HPET enabled the crash happened about once in ten trials.

As template I used "opensuse-10-i386-default.tar.gz", since the machine is not allowed to talk to web, I used the precreated template. But still I don't think this is really a problem of the template, I didn't change it.

I was guessing for HPET, since it seemed to me that there had been several patches applied between .18 to .21 wrt SMP. I'm not a kernel expert, just looking for patches for timer / HPET.

File Attachments

1) working.config, downloaded 310 times