Subject: Issue with stopping VE in Ubuntu 8.04 Posted by nixgeek on Tue, 19 Aug 2008 17:20:14 GMT

View Forum Message <> Reply to Message

Hi,

I have been very excited about OpenVZ and its similarities with Sun Zones. However, I am having an issue with stopping VE's after I have created them. After it fails I am unable to enter the VE. Nor can I reboot. I have to manually reboot it.

I am running the latest version.

Server: Ubuntu 8.04

Vzkernel: Linux vmserver 2.6.24-19-openvz #1 SMP Sat Jul 12 01:13:47 UTC 2008 x86_64

GNU/Linux

vzctl: 3.0.22

VE: ubuntu-8.04-amd64-minimal

Hardware:

CPU: Intel core 2 duo e6400

Mem: 4GB

Nic: Intel Pro 1000 ethernet adapter

I am using the cached template from:

http://wiki.openvz.org/Download/template/precreated

Attached is the conf file I am using to create the VE. It is based on vzsplit.

Any ideas?

This is stopping my progression in this adventure.

File Attachments

1) ve-64hosts.conf-sample, downloaded 347 times

Subject: Re: Issue with stopping VE in Ubuntu 8.04 Posted by piavlo on Wed, 20 Aug 2008 22:12:36 GMT

View Forum Message <> Reply to Message

First of all 2.6.24 is unstable branch, you can try more stable 2.6.18

Check if you have any processes which are in D state (uninterruptable sleep) inside of VE or kernel threads (like events) which are also in D state.

IMHO this the MAIN HUGE problem with OpenVZ - once you get one of those D state processes you can stop the VE

also you actually can't stop the HN also. This happens a lot with

processes accessing nfs shares but only inside VE, while there are no problems with remote nfs server.

Once OpenVZ releases kernel 2.6.25+

this will be partially resolved, as there are new TASK_KILLABLE

state which allows processes in this state to be killed with SIGKILL, and nfs client code is already utilizing this new state, which will allow those VE to be stopped (only if the stuck processes are using TASK_KILLABLE sleep).

Subject: Re: Issue with stopping VE in Ubuntu 8.04 Posted by nixgeek on Thu, 21 Aug 2008 02:36:35 GMT

View Forum Message <> Reply to Message

Ah interesting...

I am using NFS in the VE.

And it does appear after I mount the NFS share that this issue arises. However, I've tried to umount the shares prior to shutting it down. Still same issue. So it might be a /etc/init.d/portmap not stopping or some other NFS related app.

I noticed that there are 2-3 process still running even after the 'vzctl stop' command times out.

After that I am hosed... Cannot even reboot. I have to do a 'halt -fp' in the HW node.

I did not know this was the unstable release.

Thanks for the info...

Subject: Re: Issue with stopping VE in Ubuntu 8.04 Posted by nixgeek on Thu, 21 Aug 2008 02:51:56 GMT

View Forum Message <> Reply to Message

Thanks for the info...

So what is the best way to deal with this in the meantime?

What is the work-a-round? {if any}

Subject: Re: Issue with stopping VE in Ubuntu 8.04 Posted by piavlo on Thu, 21 Aug 2008 11:25:02 GMT

View Forum Message <> Reply to Message

Even with so called stable 2.6.18 release there are same problems.

When a process get stuck in D state inside VE (probably trying to access files on nfs mounts), there are usually corresponding one of the [events/cpunum] kernel thread in D state, meaning the only way

is to hard reset the node (or little better use MagicSysRq for cleaner reset).

One of the things i've noticed that makes these hangup D processes to appear much less frequently, is that you should run portmap on HN too (since i've noticed there are HN dmesg notifications smthing like portmap not responding, i don't know if it's regaring portmap not reponding inside VE of the HN itself, since there is one kernel and one dmesg, and i can't if a specific dmessage is related to some VE or the HN itself). Another thing i did ,out of paranoia, is just install nfs-utils on HN too. I'd love to hear from the developers what they think about this portmap issue?

In general the nfs support in OpenVZ is very BAD, i've bugged the developers several time about a year ago, but nothing changed since. Since i've got the feeling that nfs support is not important feature for them i've not bugged them since then.