
Subject: OpenVZ Hangs on PAE Kernel
Posted by [Echelon](#) on Sun, 12 Aug 2007 15:29:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

I am trying to run the OpenVZ PAE kernel with 8 gigs of ram to back the server, however the server keeps suffering from hangings.

It starts 2 or 3 vpses fine, but once it hits 4 or so, it hangs without a kernel panic or any other proper error message.

The kernel version is 2.6.18-8.1.8.el5.028stab039.1PAE but it also does it on 2.6.18-ovz028stab039.1-enterprise

Anybody have the same problem and come up with a solution? Node is running CentOS 5. Thanks.

BTW: The server is running a 32 bit os.

Subject: Re: OpenVZ Hangs on PAE Kernel
Posted by [WireSix](#) on Sun, 12 Aug 2007 17:44:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

a little more background info

Motherboard: SuperMicro PDSMA
CPU: Intel q6600
Ram: 4 x 2GB DDR2-667
OS: CentOS 5.0 x32

From IP console the system shows absolutely no errors, kernel does not panic. The actual console sits at the login prompt with a cursor that blinks slightly faster than normal.

I can not attribute this to a hardware problem as this has happend with 3 nodes that I've tested with. There appear to be minimal DIFF's between the PAE, Enterprise and non-pae kernels I don't think it's a kernel issue but possibly something somewhere else.

Anyone have any ideas?

Subject: Re: OpenVZ Hangs on PAE Kernel
Posted by [vaverin](#) on Mon, 13 Aug 2007 05:45:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Ryan,

Could you please describe the node state in more details:

- is the node accessible via ssh?
- is the node pingable?
- what happened on active shells? do you have any reaction on keyboard?
- do you have local console? could you try to use Magic Sysrq keys?
http://wiki.openvz.org/Magic_SysRq_Key
- could you please attach remote console (serial- or netconsole) to node log collection?
http://wiki.openvz.org/Remote_console_setup

Is the problem reproduceable?

In general in case of node hangs the troubleshooting procedure is following:

- attach remote console and tune the node logs collection
- reproduce the hang
- describe the node state by using Magic SysRq keys:

We need to know:

"Show Pc" (alt+sysrq+p) -- several times

"Show CPUs" (alt+sysrq+w) -- several times too

This debug describes the state of CPUs, what the tasks they executed.

Then please press "Show Tasks" (alt+sysrq+t). this debug outputs information about all the tasks on the node and it may take a lot of time, up to several minutes or even more, depending on number of processes running on your node. It is most important information, and it should be collected without the loss. In case of local logs some part of this information is lost, that's why it's important to attach the remote console to the node.

Then please press again alt+sysrq+p and alt+sysrq+w again,
then "Show Mem" (alt+sysrq+m) and "Show Vsched" (alt+sysrq+V).

Now you can reboot the node and send the collected logs to us, via attachment or via bugzilla:
<http://bugzilla.openvz.org/>

thank you,
Vasily Averin

Subject: Re: OpenVZ Hangs on PAE Kernel
Posted by [WireSix](#) on Tue, 14 Aug 2007 03:18:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

I'll have to get that setup on the new node, the 3 existing nodes are in production running on standard kernels until we can get this figured out, I have another node arriving in the next two days we can use for testing.

Subject: Re: OpenVZ Hangs on PAE Kernel
Posted by [WireSix](#) on Tue, 28 Aug 2007 03:58:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

Q: is the node accessible via ssh?

A: no, entirely offline, unpingable

Q: is the node pingable?

A: see above

Q: what happened on active shells? do you have any reaction on keyboard?

A: nothing, totally dead

Q: do you have local console? could you try to use Magic Sysrq keys?

A: KVM/IP is available, absolutely no response, also tried on local console, no response

Q: could you please attach remote console (serial- or netconsole) to node log collection?

A: installed and netconsole had no output different from true console.

vaverin wrote on Mon, 13 August 2007 01:45Hi Ryan,

Could you please describe the node state in more details:

- is the node accessible via ssh?
 - is the node pingable?
 - what happened on active shells? do you have any reaction on keyboard?
 - do you have local console? could you try to use Magic Sysrq keys?
- http://wiki.openvz.org/Magic_SysRq_Key
- could you please attach remote console (serial- or netconsole) to node log collection?
- http://wiki.openvz.org/Remote_console_setup

Is the problem reproduceable?

In general in case of node hangs the troubleshooting procedure is following:

- attach remote console and tune the node logs collection
- reproduce the hang
- describe the node state by using Magic SysRq keys:

We need to know:

"Show Pc" (alt+sysrq+p) -- several times

"Show CPUs" (alt+sysrq+w) -- several times too

This debug describes the state of CPUs, what the tasks they executed.

Then please press "Show Tasks" (alt+sysrq+t). this debug outputs information about all the tasks on the node and it may take a lot of time, up to several minutes or even more, depending on number of processes running on your node. It is most important information, and it should be collected without the loss. In case of local logs some part of this information is lost, that's why it's

important to attach the remote console to the node.

Then please press again alt+sysrq+p and alt+sysrq+w again,
then "Show Mem" (alt+sysrq+m) and "Show Vsched" (alt+sysrq+V).

Now you can reboot the node and send the collected logs to us, via [attachement](#) or via bugzilla:
<http://bugzilla.openvz.org/>

thank you,
Vasily Averin

Subject: Re: OpenVZ Hangs on PAE Kernel
Posted by [vaverin](#) on Tue, 28 Aug 2007 06:56:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

It's real nightmare.

Do you probably know how to reproduce this situation?

This issue can be software or hardware-related.

If this issue is software-related, then blank screen and non-working Magic sysrq can indicate some lockup in interrupt handlers: all the CPUs are busy and interrupts are disabled.

However this sort of problems should be detected by NMI watchdog. Could you please check /proc/interrupts file: is NMI watchdog enabled on your node? is it works properly (number of NMI interrupts should be increased evenly in real time).

If NMI watchdog is disabled on your node -- you can enable it (or switch its mode) by using "nmi_watchdog=" option into kernel commandline. You can try to set NMI_IO_APIC mode (nmi_watchdog=1) or NMI_LOCAL_APIC mode (nmi_watchdog=2).

However if NMI watchdog is enabled and works properly on your node -- you can try to re-install 64-bit OS on HW-node (you will be able to use 32-bit VE in this case).

In the ways described above will not help and the node will hangs again -- then situation looks like hardware fault. In this case I would like to recommend you to check and replace some parts of your hardware.

thank you,
Vasily Averin
