Subject: ovzkernels not booting - where to look first Posted by ugob on Wed, 01 Aug 2007 03:02:24 GMT

View Forum Message <> Reply to Message

Hi,

This morning, I upgraded the OpenVZ kernel on a server that is in a colocation facility.

Old kernel: ovzkernel-smp-2.6.9-023stab040.1

New kernel: ovzkernel-PAE-2.6.18-8.1.8.el5.028stab039.1

I rebooted using the new kernel, the machine didn't come up. The sysadmin at the colocation facility told me that none of the ovzkernel would boot, they would hang either at the partition check or before.

I'm going to the colocation facility tomorrow, but I currently have access to the server, which is booted with 2.6.9-55.0.2.ELsmp (stock centos4 kernel).

Any ideas of what I could check now? And what I should check first when I'm there?

Thanks, Ugo

Subject: Re: ovzkernels not booting - where to look first Posted by vaverin on Wed, 01 Aug 2007 04:49:24 GMT View Forum Message <> Reply to Message

Hi Ugo,

- 1) is it x86 or x86 64 node? new kernel is 32-bit, but old one can be 64-bit.
- 2) it may be initrd related issue. are you sure that initrd image for new kernel has been created correctly? Could you please try to install RHEL5/CentOs5 kernel on the node?
- 3) It is important to understand on where node hangs. Could you please check /var/log/messages file on your node? If it does not have any messages from new kernel -- have you possibility to attach KVM or serial console to the node?

thank you, Vasily Averin

Subject: Re: ovzkernels not booting - where to look first Posted by ugob on Wed, 01 Aug 2007 11:58:43 GMT

View Forum Message <> Reply to Message

vaverin wrote on Wed, 01 August 2007 00:49Hi Ugo,

1) is it x86 or x86_64 node? new kernel is 32-bit, but old one can be 64-bit.

All the ovzkernels are i686.

vaverin wrote on Wed, 01 August 2007 00:49

2) it may be initrd related issue. are you sure that initrd image for new kernel has been created correctly? Could you please try to install RHEL5/CentOs5 kernel on the node?

I don't really understand what you mean here.

vaverin wrote on Wed, 01 August 2007 00:49

3) It is important to understand on where node hangs. Could you please check /var/log/messages file on your node? If it does not have any messages from new kernel -- have you possibility to attach KVM or serial console to the node?

/var/log/messages didn't show anything about the new kernel. I'll be there tonight so I'll be able to see the output. I was just asking for advice in advance.

Thanks,

Subject: Re: ovzkernels not booting - where to look first Posted by khorenko on Wed, 01 Aug 2007 13:33:45 GMT View Forum Message <> Reply to Message

Hello Ugo,

1) can you please check that the processor supports the PAE? # cat /proc/cpuinfo

. . .

flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse sse2 syscall nx mmxext lm 3dnowext 3dnow pni

. . .

e.g. 'pae' flag should present.

2)

Quote:it may be initrd related issue. are you sure that initrd image for new kernel has been created correctly?

Could you please check the /etc/grub.conf (or lilo.conf if you are using lilo). There should be an entry for the newly installed ovzkernel-PAE-2.6.18-8.1.8.el5.028stab039.1 Do those files mentioned in 'kernel' and 'initrd' sections exist in /boot? Could you please try to recreate initrd file using the following command? Are there any errors reported? # mkinitrd -v -f /boot/initrd-2.6.18-8.1.8.el5.028stab039.1PAE.img

2.6.18-8.1.8.el5.028stab039.1PAE

Quote: Could you please try to install RHEL5/CentOs5 kernel on the node?

Well, could you try to install a stock kernel from CentOS 5 and boot in it? For example that one: http://isoredirect.centos.org/centos/5/updates/i386/RPMS/kernel-PAE-2.6.18-8.1.8.el5.i686.rpm

3) If all of this won't work, please, try to attach a serial console to the node or at least KVM. This will allow us to collect the boot logs or at least to see the last messages. If you are going to the colocation facility, please, take a photo with you and take a photo of a screen of a hanged node. If we are lucky the last messages can contain useful information. http://wiki.openvz.org/Remote_console_setup

Hope this helps.

Thank you, Konstantin.

Subject: Re: ovzkernels not booting - where to look first Posted by ugob on Wed, 01 Aug 2007 19:38:24 GMT View Forum Message <> Reply to Message

finist wrote on Wed, 01 August 2007 09:33Hello Ugo,

1) can you please check that the processor supports the PAE? # cat /proc/cpuinfo

...

flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse sse2 syscall nx mmxext lm 3dnowext 3dnow pni

e.g. 'pae' flag should present.

flags : fpu vme de pse tsc msr pae mce cx8 apic mtrr pge mca cmov pat pse36 mmx fxsr sse pni syscall mp mmxext 3dnowext 3dnow

So it looks OK.

finist wrote on Wed, 01 August 2007 09:33

2) it may be initrd related issue. are you sure that initrd image for new kernel has been created correctly?

Could you please check the /etc/grub.conf (or lilo.conf if you are using lilo). There should be an entry for the newly installed ovzkernel-PAE-2.6.18-8.1.8.el5.028stab039.1

Do those files mentioned in 'kernel' and 'initrd' sections exist in /boot? Could you please try to recreate initrd file using the following command? Are there any errors reported? # mkinitrd -v -f /boot/initrd-2.6.18-8.1.8.el5.028stab039.1PAE.img 2.6.18-8.1.8.el5.028stab039.1PAE

All files are there. I tried recreating the file like you suggested, no error.

Creating initramfs

Looking for deps of module scsi mod

Looking for deps of module sd mod scsi mod

Looking for deps of module scsi mod

Looking for deps of module unknown

Looking for deps of module 3w-xxxx scsi_mod

Looking for deps of module scsi_mod

Looking for deps of module ide-disk

Looking for deps of module ext3 jbd

Looking for deps of module jbd

Using modules: ./kernel/drivers/scsi/scsi_mod.ko ./kernel/drivers/scsi/sd_mod.ko

./kernel/drivers/scsi/3w-xxxx.ko ./kernel/fs/jbd/jbd.ko ./kernel/fs/ext3/ext3.ko

/sbin/nash -> /tmp/initrd.n19895/bin/nash

/sbin/insmod.static -> /tmp/initrd.n19895/bin/insmod

/sbin/udev.static -> /tmp/initrd.n19895/sbin/udev

/etc/udev/udev.conf -> /tmp/initrd.n19895/etc/udev/udev.conf

copy from

/lib/modules/2.6.18-8.1.8.el5.028stab039.1PAE/./kernel/drivers/scsi/scsi_mod.ko(elf32-i386) to /tmp/initrd.n19895/lib/scsi_mod.ko(elf32-i386)

copy from

/lib/modules/2.6.18-8.1.8.el5.028stab039.1PAE/./kernel/drivers/scsi/sd_mod.ko(elf32-i386) to /tmp/initrd.n19895/lib/sd_mod.ko(elf32-i386)

copy from

/lib/modules/2.6.18-8.1.8.el5.028stab039.1PAE/./kernel/drivers/scsi/3w-xxxx.ko(elf32-i386) to /tmp/initrd.n19895/lib/3w-xxxx.ko(elf32-i386)

copy from /lib/modules/2.6.18-8.1.8.el5.028stab039.1PAE/./kernel/fs/jbd/jbd.ko(elf32-i386) to /tmp/initrd.n19895/lib/jbd.ko(elf32-i386)

copy from /lib/modules/2.6.18-8.1.8.el5.028stab039.1PAE/./kernel/fs/ext3/ext3.ko(elf32-i386) to /tmp/initrd.n19895/lib/ext3.ko(elf32-i386)

Loading module scsi mod

Loading module sd mod

Loading module 3w-xxxx

Loading module jbd

Loading module ext3

finist wrote on Wed, 01 August 2007 09:33

Quote: Could you please try to install RHEL5/CentOs5 kernel on the node?

Well, could you try to install a stock kernel from CentOS 5 and boot in it?

For example that one: http://isoredirect.centos.org/centos/5/updates/i386/RPMS/kernel-PAE-2.6.18-8.1.8.el5.i686.rpm

Can't install:[root@bibitte ~]# rpm -ivhf kernel-PAE-2.6.18-8.1.8.el5.i686.rpm warning: kernel-PAE-2.6.18-8.1.8.el5.i686.rpm: V3 DSA signature: NOKEY, key ID e8562897 error: Failed dependencies:

initscripts >= 8.11.1-1 is needed by kernel-PAE-2.6.18-8.1.8.el5.i686 mkinitrd >= 4.2.21-1 is needed by kernel-PAE-2.6.18-8.1.8.el5.i686 ppp < 2.4.3-3 conflicts with kernel-PAE-2.6.18-8.1.8.el5.i686 e2fsprogs < 1.37-4 conflicts with kernel-PAE-2.6.18-8.1.8.el5.i686 procps < 3.2.5-6.3 conflicts with kernel-PAE-2.6.18-8.1.8.el5.i686 udev < 063-6 conflicts with kernel-PAE-2.6.18-8.1.8.el5.i686 iptables < 1.3.2-1 conflicts with kernel-PAE-2.6.18-8.1.8.el5.i686

It is quite troubling, that, according to the sysadmin at the colocation facility, even the -smp kernel, which was running fine before the upgrade, doesn't boot anymore...

finist wrote on Wed, 01 August 2007 09:33

3) If all of this won't work, please, try to attach a serial console to the node or at least KVM. This will allow us to collect the boot logs or at least to see the last messages. If you are going to the colocation facility, please, take a photo with you and take a photo of a screen of a hanged node. If we are lucky the last messages can contain useful information. http://wiki.openvz.org/Remote_console_setup

Hmmm, I don't have my digital camera with me, only my cell. Not sure if I have a null modem cable. I'll see. I will have, for sure, a paper notepad, and a pen.

Thanks!

Subject: Re: ovzkernels not booting - where to look first Posted by ugob on Thu, 02 Aug 2007 01:14:33 GMT View Forum Message <> Reply to Message

Coming back from the datacenter.

The ovzkernel-smp-2.6.9-023stab040.1 boots fine, like before.

It is the newer ones that don't boot:

ovzkernel-PAE-2.6.18-8.1.3.el5.028stab033.1 ovzkernel-PAE-2.6.18-8.1.4.el5.028stab035.1 ovzkernel-PAE-2.6.18-8.1.8.el5.028stab039.1 So the problem is likely to be either the version 2.6.18 of the kernel, or the PAE version of the kernel. Grrr, I just realized I should have tried booting in the 2.6.18 non pae kernel... I forgot.

The good news is that my server is up and running. The bad news are:

- I can't help much to solve the problem as the PAE kernels hang randomly (one hung at the startup of VE 109, another at init startup, another at the filesystem check, no logs of course.
- How can I upgrade my kernel if only non-pae or 2.6.9 kernels boot?

Thanks,

Ugo

Subject: Re: ovzkernels not booting - where to look first Posted by khorenko on Thu, 02 Aug 2007 06:49:37 GMT

View Forum Message <> Reply to Message

1) may be the question is a bit late, but still: and what do you want to do globally? Why do you want to upgrade the kernel? Do you want just a kernel which can support more than 4 Gb RAM? If yes, you can safely use the enterprise kernel from 2.6.9 branch (e.g. ovzkernel-enterprise-2.6.9-023stab044.4.i686.rpm). It should work fine. Or you just want to use the most modern kernel?

- 2) i still think that the problem with 2.6.18 kernels is in initrd.
- i just tried to reproduce: took a RHEL4.3 and tried to install 2.6.18 OVZ kernel. The boot failed unable to find root.
- i haven't had a time to get the reason and don't know the exact solution how to fix this i'll certainly do this but later.

At the moment i've done a workaround - took a CentOS5 node, install there the same OVZ kernel and just copy an initrd to the RHEL4.3 node. It works. You can try the same while i'm looking for the correct solution.

Note: to make sure the initrd created on the CentOS5 node contains all the modules required on your RHEL4 node it's better to recreate initrd manually (on the CentOS5 node): # mkinitrd -v -f /boot/initrd-2.6.18-8.1.8.el5.028stab039.1PAE.img 2.6.18-8.1.8.el5.028stab039.1PAE --preload=scsi_mod --preload=sd_mod --preload=3w-xxxx

3) if nothing helps please provide me access to the node, i'll try to boot the kernel. Of course the permission to reboot is required. :\ And if my described workaround won't workout, we have to get the ability to collect the logs somehow. Just find a COM-to-COM cable and connect this node with any other (preferably Linux, but not required). We can help you to configure it later.

You can safely send the access through the private messaging. Just one more thing: i'll be unavailable in a few days so please, if you'll send something private, send the copy to Vasily (vaverin), he can help you too.

Subject: Re: ovzkernels not booting - where to look first Posted by ugob on Thu, 02 Aug 2007 11:50:14 GMT

View Forum Message <> Reply to Message

finist wrote on Thu, 02 August 2007 02:491) may be the question is a bit late, but still: and what do you want to do globally? Why do you want to upgrade the kernel? Do you want just a kernel which can support more than 4 Gb RAM? If yes, you can safely use the enterprise kernel from 2.6.9 branch (e.g. ovzkernel-enterprise-2.6.9-023stab044.4.i686.rpm). It should work fine. Or you just want to use the most modern kernel?

I just want to use the most modern kernel. I don't need the PAE, I just need an SMP kernel. finist wrote on Thu, 02 August 2007 02:49

2) i still think that the problem with 2.6.18 kernels is in initrd.

i just tried to reproduce: took a RHEL4.3 and tried to install 2.6.18 OVZ kernel. The boot failed - unable to find root.

i haven't had a time to get the reason and don't know the exact solution how to fix this - i'll certainly do this but later.

At the moment i've done a workaround - took a CentOS5 node, install there the same OVZ kernel and just copy an initrd to the RHEL4.3 node. It works. You can try the same while i'm looking for the correct solution.

Note: to make sure the initrd created on the CentOS5 node contains all the modules required on your RHEL4 node it's better to recreate initrd manually (on the CentOS5 node): # mkinitrd -v -f /boot/initrd-2.6.18-8.1.8.el5.028stab039.1PAE.img 2.6.18-8.1.8.el5.028stab039.1PAE --preload=scsi_mod --preload=sd_mod --preload=3w-xxxx

I doubt the problem is initrd, since some kernels hung while trying to start the 9th VE.

finist wrote on Thu, 02 August 2007 02:49

3) if nothing helps please provide me access to the node, i'll try to boot the kernel. Of course the permission to reboot is required. :\ And if my described workaround won't workout, we have to get the ability to collect the logs somehow. Just find a COM-to-COM cable and connect this node with any other (preferably Linux, but not required). We can help you to configure it later.

You can safely send the access through the private messaging. Just one more thing: i'll be unavailable in a few days so please, if you'll send something private, send the copy to Vasily (vaverin), he can help you too.

I can give you access as long as the people at the datacenter are available and I have admin time free (I have 15min per month free), or if I'm at the datacenter.

I'm leaving for a 3-week vacation tomorrow, so I think I'll just stick to the kernel that works (I've configured grub.conf accordingly) and wait until I come back. Now that the machine is running fine, I'm happy. I'll try to help as much as possible when I come back from vacation.

Subject: Re: ovzkernels not booting - where to look first Posted by vaverin on Fri, 03 Aug 2007 01:08:09 GMT

View Forum Message <> Reply to Message

Hi Ugo

ugob wrote on Thu, 02 August 2007 15:50finist wrote on Thu, 02 August 2007 02:49 2) i still think that the problem with 2.6.18 kernels is in initrd.

I doubt the problem is initrd, since some kernels hung while trying to start the 9th VE.

There is the following collision:

If 2.6.18 kernel really tried to start any VEs -- you should have kernel booting messages in /var/log/messages file. However you said it didn't show anything about the new kernel.

I agree with Konstantin, IMHO the situation on your node looks like initrd troubles.

Thank you, Vasily Averin

Subject: Re: ovzkernels not booting - where to look first Posted by ugob on Fri, 03 Aug 2007 01:52:22 GMT View Forum Message <> Reply to Message

Tion i ciam moodage as reopty to moodage

You are right, I should have looked at the log this time. I got syslog messages only once, although I tried all of the PAE kernels.

Not much luck. Here is how it ends:

Aug 1 17:44:45 bibitte kernel: VE: 101: started Aug 1 17:44:49 bibitte kernel: VE: 102: started Aug 1 17:44:54 bibitte kernel: VE: 103: started Aug 1 17:45:00 bibitte kernel: VE: 104: started Aug 1 17:45:06 bibitte kernel: VE: 109: started

Here is how it looks with the working kernel:

Aug 1 17:40:53 bibitte kernel: ip_conntrack version 2.1 (8188 buckets, 65504 max) - 312 bytes per conntrack

Aug 1 17:40:53 bibitte kernel: NET: Registered protocol family 17

Aug 1 17:40:53 bibitte kernel: e1000: eth0: e1000 watchdog: NIC Link is Up 10 Mbps Half

Duplex

Aug 1 17:40:55 bibitte kernel: device eth0 entered promiscuous mode

Aug 1 17:40:57 bibitte ntpd[8538]: kernel time sync status 0040

Aug 1 17:41:05 bibitte kernel: VPS: 101: started

Aug 1 17:41:08 bibitte kernel: VPS: 102: started

Aug 1 17:41:12 bibitte kernel: VPS: 103: started

Aug 1 17:41:17 bibitte kernel: VPS: 104: started

Aug 1 17:41:20 bibitte kernel: VPS: 109: started

Aug 1 17:41:23 bibitte kernel: VPS: 110: started

Aug 1 17:41:29 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Up 10 Mbps Full Duplex

One weird thing, I've got many, many, many of these on the working kernel:

Aug 1 17:56:36 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Down

Aug 1 17:56:37 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Up 10 Mbps Full Duplex

Aug 1 17:56:37 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Down

Aug 1 17:56:37 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Up 10 Mbps Full Duplex

Aug 1 17:56:39 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Down

Aug 1 17:56:39 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Up 10 Mbps Full Duplex

Aug 1 17:56:42 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Down

Aug 1 17:56:42 bibitte kernel: e1000: eth0: e1000_watchdog: NIC Link is Up 10 Mbps Full Duplex

I can give root access if you want, but you can't reboot. Would it help?

Subject: Re: ovzkernels not booting - where to look first Posted by vaverin on Fri, 03 Aug 2007 02:12:47 GMT

View Forum Message <> Reply to Message

ugob wrote on Fri, 03 August 2007 05:52I can give root access if you want, but you can't reboot. Would it help?

I'm ready to look on your node, probably I'll be able to find something.

Please give me access via PM.

Thank you, Vasily Averin