
Subject: Server crash

Posted by [dvazart](#) on Tue, 30 Jun 2009 12:13:17 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi again !

I'm running a 2.6.18-14-ovz-686-enterprise kernel under Debian Etch. I have about 150 VEs running in 2 quad core Intel Xeon processors with 32 Gb RAM.

My HN crashes every weekend at different times, this weekend it crashed 2 times... and my customers are not really happy

I think I have seen on the console an output like an "oops" (<http://wiki.openvz.org/Oops>) but I'm not sure that the logs do not see anything abnormal, except this:

`tail /var/log/messages`

```
Jun 23 18:31:59 sht2 kernel: oom-killer: gfp_mask=0xd0, order=0
Jun 23 18:31:59 sht2 kernel: [<c0159a39>] out_of_memory+0x109/0x150
Jun 23 18:31:59 sht2 kernel: [<c015b5e8>] __alloc_pages+0x328/0x3a0
Jun 23 18:31:59 sht2 kernel: [<c015b681>] __get_free_pages+0x21/0x50
Jun 23 18:31:59 sht2 kernel: [<c0190311>] __pollwait+0xb1/0x110
Jun 23 18:31:59 sht2 kernel: [<c04287cf>] tcp_poll+0x2f/0x220
Jun 23 18:31:59 sht2 kernel: [<c03f26c0>] sock_poll+0x20/0x30
Jun 23 18:31:59 sht2 kernel: [<c018f9d1>] do_select+0x291/0x4d0
Jun 23 18:31:59 sht2 kernel: [<c0190260>] __pollwait+0x0/0x110
Jun 23 18:31:59 sht2 kernel: [<c0119d20>] default_wake_function+0x0/0x20
Jun 23 18:31:59 sht2 last message repeated 19 times
Jun 23 18:31:59 sht2 kernel: [<c018fdf1>] core_sys_select+0x1e1/0x330
Jun 23 18:31:59 sht2 kernel: [<c01794f8>] do_sync_write+0xc8/0x110
Jun 23 18:31:59 sht2 kernel: [<c013a900>] autoremove_wake_function+0x0/0x60
Jun 23 18:31:59 sht2 kernel: [<c01b47b8>] dnotify_parent+0x38/0xd0
Jun 23 18:31:59 sht2 kernel: [<c019065d>] sys_select+0x4d/0x1c0
Jun 23 18:31:59 sht2 kernel: [<c017a891>] sys_write+0xb1/0xc0
Jun 23 18:31:59 sht2 kernel: [<c010322f>] syscall_call+0x7/0xb
```

Is this normal ??

I have 3 other questions :

- it is possible that a misconfiguration in VEs can crash the server?
- it may be a bug in OpenVZ kernel? (I use : 028stab056.1dso1)
- because my server has 32 GB of RAM, I don't want to run a memory test, that could take a long time...

Can you advise me ?

thanks !

Subject: Re: Server crash

Posted by [maratrus](#) on Tue, 30 Jun 2009 12:47:54 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello,

Quote:

- it is possible that a misconfiguration in VEs can crash the server?

A misconfiguration mustn't crash the server. But it may drastically decrease performance.

Quote:

- it may be a bug in OpenVZ kernel?

Don't hesitate to make a new bug report if any bug is found.

<http://bugzilla.openvz.org/>

Quote:

My HN crashes every weekend at different times, this weekend it crashed 2 times... and my customers are not really happy

We have to be provided with textual logs describing the crash in detail.

A serial console might be very helpful to get the full crash output.

http://wiki.openvz.org/Remote_console_setup#Serial_console

Quote:

I think I have seen on the console an output like an "oops" (<http://wiki.openvz.org/Oops>) but I'm not sure that the logs do not see anything abnormal, except this:

```
tail /var/log/messages
```

These messages indicate that out of memory was invoked

http://en.wikipedia.org/wiki/Out_of_memory

http://linux-mm.org/OOM_Killer

But I cannot see any crash related things.

If it is an out-of-memory situation that bothers your customers so much you may have to set up `user_beancounters` parameters more accurately. The best thing I can suggest you is to read an article describing UBC in general and in detail.

<http://wiki.openvz.org/UBC>

Subject: Re: Server crash

Posted by [dvazart](#) on Tue, 30 Jun 2009 13:31:16 GMT

[View Forum Message](#) <> [Reply to Message](#)

Thanks for your answer.

So, if my HN enters in a state of out of memory, with only half the memory used, this can lead to produce the "oops" in OpenVZ?

and this means that I have a problem in some of the memory bars...

am I right?

Subject: Re: Server crash

Posted by [maratrus](#) on Tue, 30 Jun 2009 15:39:49 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

Quote:

So, if my HN enters in a state of out of memory, with only half the memory used, this can lead to produce the "oops" in OpenVZ?

Please, examine log messages carefully. Examine what in particular oom-killer wrote in log files.

Quote:

this can lead to produce the "oops" in OpenVZ?

It is not an oops. It's just a dump of a stack i.e. those functions to which the addresses on a stack correspond.

Are you running x86_64 kernel?

Subject: Re: Server crash

Posted by [dvazart](#) on Wed, 01 Jul 2009 12:35:42 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi, thanks for your answer.

No, i'm using a 32 bits kernel.

It is possible that a misconfiguration of VMGUARPAGES, PRIVMPAGES and OOMGUARPAGES cause an OOM situation (as mentioned in the first post) in the HN?

This can also lead to an oops in OpenVZ?

Subject: Re: Server crash

Posted by [dvazart](#) on Thu, 02 Jul 2009 06:33:32 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi !

My server was crashed today too... I take an screenshot of the tty with the KVM and i have this output:

<http://picasaweb.google.com/lh/photo/ANuyVCUNiPFZktbclv6adA?feat=directlink>

it not seems to be an oops, you know it can be?

I really need help with this... thanks !

Subject: Re: Server crash

Posted by [maratrus](#) on Thu, 02 Jul 2009 14:38:39 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Daniel,

it's neither an oops not crash.

It's just an information from oom-killer.

Oom-killer is not an OpenVZ specific thing so you can read about it wherever you want. I provided you with two links.

To find out the real reason of oom-killer you have to examine logs carefully.

A possible reason why this can happen is that your system might have bumped into lack of the size of the "normal zone", i.e. that zone where kernel holds its objects. The size of "normal zone" is limited to ~800mb regardless of the total amount of RAM and it is a restriction of x86 architecture. As far as I understand your server has 32Gb RAM and it is heavily loaded. So, there is a nonzero probability that "normal zone" might be exhausted.

Subject: Re: Server crash

Posted by [dvazart](#) on Fri, 03 Jul 2009 07:21:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

Thanks for your explanation.

do you think its a misconfiguration of KMEMSIZE in the VE's ??

thanks.

Subject: Re: Server crash

Posted by [maratrus](#) on Fri, 03 Jul 2009 09:11:38 GMT

[View Forum Message](#) <> [Reply to Message](#)

It may be a lot of VEs running on the HN. Try to reduce the number of them.
