
Subject: *SOLVED* 2.6.18 crash w/ kernel log :: oom-killer

Posted by [rickb](#) on Sun, 20 May 2007 20:31:39 GMT

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

Hello. I have an openvz server with 4GB of memory which continually hits oom-killer after 15-60 minutes of activity- even though I think it has free memory left. I loaded 8GB of swap to try and guard against oom-killer but that did not help.

I have this crash report:

<http://66.97.172.77/cr.txt>

If this is a bug I will create a bugzilla incident but I want the devs opinion on this first since I may be interpreting it incorrectly.

2.6.18 kernel

Quote:HighMem free:373380kB min:512kB low:4064kB high:7620kB active:2106760kB
inactive:747916kB present:3407872kB pages_scanned:0 all_unreclaimable? no

Free swap = 8131196kB

Total swap = 8457336kB

my limited knowledge tells me this server has approx 4GB physical memory with 748MB unused, 8.4GB swap with 8.1GB unused. If this is the case, should/can oom-killer be triggered?

I loaded an older kernel for right now, to try and isolate the problem.

Thanks!

Rick

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer

Posted by [Vasily Tarasov](#) on Mon, 21 May 2007 06:41:10 GMT

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

Hello,

Rick, can you tell us please specific OpenVz kernel version you're using?

We introduced some changes to OOM, so, we want to ensure that they cause your problem...

Thank you,

Vasily.

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer
Posted by [rickb](#) on Mon, 21 May 2007 06:43:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

I backdated the kernel way back to 2.6.8 and had the same problem. I believe its a problem with my physical server and not the kernel, as the server was stable for over 1 year on 2.6.8

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer
Posted by [Vasily Tarasov](#) on Mon, 21 May 2007 06:51:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

I don't think that any hardware problem can cause OOM and after your answer about situation on 2.6.8 OVZ kernel I'm sure that the problem is not in OVZ kernel also. So my assumption is that you have installed some memory eating monster on your node.

Thank you,
Vasily.

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer
Posted by [rickb](#) on Mon, 21 May 2007 06:53:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks Vasily,

The node has 4GB mem, 8GB swap, problem does occur with 2.6.8 and 2.6.18. To my understanding, oom killer should not activate when there is swap or memory left. is this correct?

does the crash report I linked to, say that there is memory in the system? further, all of my ve privmpages are set to 3GB which prevents any one ve from causing such a problem. Right before the system goes nutty with oomkiller, the mem free is around 1GB. In the crash report, I would expect the memory to be completely exhausted.

I am more interested in knowing if I am reading the oomkiller report properly. I feel if there is truly no memory left in the system, this is not an openvz problem, and an application being too hungry like you said. But, if there is memory left, and plenty of swap, I do not feel oomkiller should be woken up to do its dirty work.

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer

Posted by [xemul](#) on Mon, 21 May 2007 13:21:11 GMT

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

The oom killer prints the memory info on the first kill and some subsequent ones. Show it to us please.

One of the reasons why this can happen is that the kernel has run out of the so called "normal zone", i.e. the memory that kernel can place its objects to. This area is limited with ~800Mb of RAM regardless of the total size of the memory on the node.

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer

Posted by [rickb](#) on Mon, 21 May 2007 13:29:34 GMT

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

Thanks Xemul, that makes sense. Here is the klog from crash start to finish,

<http://66.97.172.77/oom.txt>

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer

Posted by [rickb](#) on Tue, 22 May 2007 22:53:19 GMT

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

Pavel, thanks for your help so far. I attached the full kernel log of the "crash" in my previous message. Can you tell me if this oom killer action is normal/expected?

If the server is dangerously low on physical memory for kernel data (right on that ~800mb threshold), I would think the kernel would temporarily fail mallocs and start swapping user memory to the disk, rather than honoring the malloc and then invoking oomkiller. Or, as a last resort, honor mallocs using the swap.

Thanks!

Rick

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer

Posted by [xemul](#) on Fri, 25 May 2007 09:39:18 GMT

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

Sorry for the late answer.

I've found in your log that you have

Normal free:3696kB ...

somewhat about 3MiB of normal zone. This is a good reason for oom killing.

Moreover, the info on top slabs is

```
dentry_cache      : size 18817024 objsize 144
page_beancounter  : size 23429120 objsize 32
size-4096(UBC)    : size 7970560 objsize 4096
filp              : size 4075520 objsize 192
size-64           : size 710627328 objsize 64
radix_tree_node   : size 10047488 objsize 276
vm_area_struct    : size 8720384 objsize 88
pmd               : size 16024320 objsize 4096
buffer_head       : size 4988928 objsize 52
ext3_inode_cache  : size 20996096 objsize 524
```

This shows that ~700MiB is busy with size-64 slab. This is a generic slab and many different objects can be stored in it. What exactly - this requires deeper investigation.

Subject: Re: 2.6.18 crash w/ kernel log :: oom-killer
Posted by [rickb](#) on Fri, 25 May 2007 10:41:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks Pavel, I'll do more research on this.
