
Subject: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [wishd](#) on Wed, 10 Jan 2018 22:08:55 GMT

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

3.10.0-693.11.6.vz7.40.4 - with meltdown on most servers have been fine. However one server basically locks out (out of memory). It is in the older kernel for the moment. Whats wierd is its not happening on all. Dmesg shows:

```
ive_file:5512KB active_file:3400KB unevictable:0KB
```

```
[ 76.038277] Memory cgroup out of memory: Kill process 3346 (init) score 0 or sacrifice child
```

```
[ 76.045106] Killed process 10947 (sh) in VE "5408" total-vm:1936kB, anon-rss:68kB,  
file-rss:444kB, shmем-rss:0kB
```

```
[ 78.795738] SLUB: Unable to allocate memory on node -1 (gfp=0xd0)
```

```
[ 78.795749] cache: inode_cache(7:5185), object size: 592, buffer size: 600, default order: 3,  
min order: 0
```

```
[ 78.795756] node 0: slabs: 4, objs: 120, free: 0
```

```
[ 78.795762] node 1: slabs: 10, objs: 540, free: 0
```

```
[ 79.124919] lsb_release invoked oom-killer: gfp_mask=0x3084d0, order=0, oom_score_adj=0
```

```
[ 79.124930] lsb_release cpuset=38104 mems_allowed=0-1
```

```
[ 79.124940] CPU: 14 PID: 15221 Comm: lsb_release ve: 38104 Not tainted
```

```
3.10.0-693.11.6.vz7.40.4 #1 40.4
```

```
[ 79.124947] Hardware name: Supermicro SYS-6018R-TDW/X10DDW-i, BIOS 2.0a 08/17/2016
```

```
[ 79.124953] Call Trace:
```

```
[ 79.124969] [<ffffffff816a2398>] dump_stack+0x19/0x1b
```

```
[ 79.124981] [<ffffffff8169ebcb>] dump_header+0x90/0x229
```

```
[ 79.124996] [<ffffffff811a4a67>] ? release_pages+0x257/0x440
```

```
[ 79.125007] [<ffffffff811998c8>] oom_kill_process+0x5e8/0x640
```

```
[ 79.125019] [<ffffffff811c0fde>] ? get_task_oom_score_adj+0xee/0x100
```

```
[ 79.125034] [<ffffffff8120eb49>] mem_cgroup_oom_synchronize+0x4a9/0x4f0
```

```
[ 79.125046] [<ffffffff81199e73>] pagefault_out_of_memory+0x13/0x50
```

```
[ 79.125056] [<ffffffff8169ceae>] mm_fault_error+0x68/0x12b
```

```
[ 79.125068] [<ffffffff816afa91>] __do_page_fault+0x391/0x450
```

```
[ 79.125079] [<ffffffff816afb85>] do_page_fault+0x35/0x90
```

```
[ 79.125089] [<ffffffff816ab8f8>] page_fault+0x28/0x30
```

```
[ 79.125100] Task in /machine.slice/38104 killed as a result of limit of /machine.slice/38104
```

```
[ 79.125106] memory: usage 41928kB, limit 9007199254740988kB, failcnt 0
```

```
[ 79.125113] memory+swap: usage 41928kB, limit 9007199254740988kB, failcnt 0
```

```
[ 79.125119] kmem: usage 17488kB, limit 17488kB, failcnt 49
```

```
[ 79.125125] Memory cgroup stats for /machine.slice/38104: rss_huge:0KB mapped_file:4188KB
```

```
shmем:24KB slab_unreclaimable:11064KB swap:0KB cache:18308KB rss:6076KB
```

```
slab_reclaimable:3272KB inactive_anon:12KB active_anon:6088KB inactive_file:13680KB
```

```
active_file:4604KB unevictable:0KB
```

```
[ 79.125166] Memory cgroup out of memory: Kill process 15221 (lsb_release) score 0 or  
sacrifice child
```

```
[ 79.134141] Killed process 15221 (lsb_release) in VE "38104" total-vm:22968kB,  
anon-rss:1432kB, file-rss:1944kB, shmем-rss:0kB
```

```
[ 83.173254] SLUB: Unable to allocate memory on node -1 (gfp=0x1080d0)
```

```
[ 83.173265] cache: kmalloc-1024(7:5185), object size: 1024, buffer size: 1024, default order: 3,
min order: 0
[ 83.173272] node 0: slabs: 4, objs: 72, free: 0
[ 83.173279] node 1: slabs: 8, objs: 256, free: 0
[ 85.357464] SLUB: Unable to allocate memory on node -1 (gfp=0x1080d0)
[ 85.357475] cache: kmalloc-1024(7:5185), object size: 1024, buffer size: 1024, default order: 3,
min order: 0
[ 85.357482] node 0: slabs: 4, objs: 72, free: 0
[ 85.357487] node 1: slabs: 8, objs: 256, free: 0
```

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server
Posted by [ikbenut](#) on Thu, 11 Jan 2018 15:50:59 GMT

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

Hello,

I have a similar problem with the latest kernel.

Restarting openvz vps server hangs and only a mount is active, after that nothing works to stop or unmount that vps. Only reboot node fixes the problem.

When i boot with previous kernel restart or stop/start works great.

I discovered this problem when i was running a script to find certain files and change ownership and saw that i had memory allocation errors.

```
SLUB: Unable to allocate memory on node -1
cache: ext4_inode_cache, object size: 1056, buffer size: 1064, default order: 3, min order: 0
node 0: slabs: 4466, objs: 132900, free: 0
node 1: slabs: 1760, objs: 51477, free: 0
```

With both the latest kernel and the previous i get those errors.
Memory itself is ok.

CT-965 /# free -m

	total	used	free	shared	buff/cache	available
Mem:	6144	368	4680	29	1095	5388
Swap:	1024	0	1024			

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server
Posted by [wishd](#) on Thu, 11 Jan 2018 16:15:34 GMT

I don't have more info to provide here unfortunately. I did migrate all vm's off the server as a resolution to a system with out issues. It doesn't appear to be hardware relate - all tests pass. I may test further with a BIOS upgrade in the future. All software really matched the other openvz7 servers.

This is surprisingly the same behavior in <https://forum.openvz.org/index.php?t=tree&th=13348&start=0> which was fixed in a kernel update - except it occurred with in 15 minutes of bootup.

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [khorenko](#) on Fri, 12 Jan 2018 07:22:15 GMT

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

> kmem: usage 17488kB, limit 17488kB, failcnt 49

Somehow you've got a kernel memory limit (and quite low limit) and it's just not enough. Get rid of this limit and that's it.

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [wishd](#) on Fri, 12 Jan 2018 17:06:56 GMT

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

I may not have been clear. The entire server was crashing similar to what 'ikbenut' said not a single vm only. I saw the same vm's running, but could not enter with mounts active. Had to do a full restart. The common error seems to be 'SLUB: Unable to allocate memory on node -1'. There was plenty of free memory as well, on boot also no swap space was used in any way before it locked up again.

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [khorenko](#) on Fri, 12 Jan 2018 20:28:55 GMT

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

Sorry, but you've provided only part of messages and i can tell that the process was killed in /machine.slice/38104 cgroup (CT 38104) because of low kmem limit in that cgroup:

```
[ 79.125100] Task in /machine.slice/38104 killed as a result of limit of /machine.slice/38104
...
[ 79.125119] kmem: usage 17488kB, limit 17488kB, failcnt 49
```

Other issues may be caused by similar limits or it might be something else.

Anyway i'd start with reviewing logs, checking failcounters in messages like

[79.125119] kmem: usage 17488kB, limit 17488kB, failcnt 49

and getting rid of kmem limits at all for Containers.

kmem is accounted into total ram+swap, so no need to limit it independently.

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [martv](#) on Sat, 13 Jan 2018 14:56:31 GMT

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

I have similar problem. After upgrade 2 nodes to latest kernel (3.10.0-693.11.6.vz7.40.4), on both occurred problems with memory. After reboot the server runs fine for 30 - 60 minutes. After that, one or more containers crashes and after while is the whole node unaccessible (I cannot log in, only hard reboot is possible). This is from /var/log/messages:

Jan 13 15:02:13 vs15 kernel: INFO: task monit:10284 blocked for more than 120 seconds.

Jan 13 15:02:13 vs15 kernel: "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.

Jan 13 15:02:13 vs15 kernel: monit D ffff882012bd9010 0 10284 3910 939
0x00000000

Jan 13 15:02:13 vs15 kernel: Call Trace:

Jan 13 15:02:13 vs15 kernel: [<ffffffff816a7a79>] schedule+0x29/0x70

Jan 13 15:02:13 vs15 kernel: [<ffffffff816a90fd>] rwsem_down_read_failed+0x10d/0x1a0

Jan 13 15:02:13 vs15 kernel: [<ffffffff8132f618>] call_rwsem_down_read_failed+0x18/0x30

Jan 13 15:02:13 vs15 kernel: [<ffffffff816a6d60>] down_read+0x20/0x40

Jan 13 15:02:13 vs15 kernel: [<ffffffff812923d5>] proc_pid_cmdline_read+0xb5/0x560

Jan 13 15:02:13 vs15 kernel: [<ffffffff8121cefc>] vfs_read+0x9c/0x170

Jan 13 15:02:13 vs15 kernel: [<ffffffff8121ddbf>] SyS_read+0x7f/0xe0

Jan 13 15:02:13 vs15 kernel: [<ffffffff816b4a7d>] system_call_fastpath+0x16/0x1b

On the second node I can see errors like these: <http://prntscr.com/hzwnw6>

Funny thing is that after reboot to older kernel, these errors are still there! But before the update to latest kernel (and utilities), none of this happened. So I believe that it has to be linked to this (it is happening on two physical nodes right after the upgrade).

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [ikbenut](#) on Wed, 17 Jan 2018 09:25:51 GMT

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

I did some more testing.

After reboot in older kernel the slub error seems to reappear, but less than with the latest kernel. So it does seem to be related with the update, I think that some other kernel settings are changed during the latest kernel install and are still active after booting an older kernel.

Before I did the kernel update everything was fine and no errors whatsoever. The scripts that are causing the errors are a quota tally script from DirectAdmin and an own made script to set ownership of files in the Home folders of users.

Also when I boot in the latest kernel I also get the same issue as martv describes. I can restart vps server but after about 15 minutes this does not work anymore and the vps is mounted as ploop device, but no vps start.

I thought I let the server do his thing, but after 1 hour still nothing happened and the whole server hangs. Only a reboot helps!

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [martv](#) on Wed, 17 Jan 2018 10:21:44 GMT

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

After several days without problem, I must say that the solution proposed by khorenko worked. I disabled KMEM limits for all containers and now is everything stable again. I used this simple shell script:

```
for CTID in $(/usr/sbin/vzlist -H -o ctid|awk '{print $1;}');
```

```
do
```

```
    vzctl set $CTID --kmemsize unlimited --save
```

```
done
```

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [ikbenut](#) on Wed, 17 Jan 2018 14:11:52 GMT

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

Changing those settings helped, no more errors and after testing with some script also the restart of vps servers is working ok.

But does this change have impact on other resources of the node? Need some other things be

limited?

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [wishd](#) on Wed, 17 Jan 2018 15:21:50 GMT

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

Make sure physpages and swappages are set and it will be fine.

Subject: Re: 3.10.0-693.11.6.vz7.40.4 Out of memory on one Server

Posted by [khorenko](#) on Thu, 18 Jan 2018 07:21:03 GMT

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

wishd wrote on Wed, 17 January 2018 18:21 Make sure physpages and swappages are set and it will be fine.

+1
