
Subject: Kernel Panic after booting to openvz2.6 kernel
Posted by [arunksasi](#) on Wed, 13 Jun 2018 09:55:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

Below are the logs for kernel panic

```
root@cpu-5110:~# tail -f /var/log/kern.log
Jun 13 07:40:26 cpu-5110 kernel: [ 14.269578] FS-Cache: Loaded
Jun 13 07:40:26 cpu-5110 kernel: [ 14.318379] NFS: Registering the id_resolver key type
Jun 13 07:40:26 cpu-5110 kernel: [ 14.318786] Key type id_resolver registered
Jun 13 07:40:26 cpu-5110 kernel: [ 14.319175] FS-Cache: Netfs 'nfs' registered for caching
Jun 13 07:40:26 cpu-5110 kernel: [ 14.412721] ploop_dev: module loaded
Jun 13 07:40:26 cpu-5110 kernel: [ 14.588418] ADDRCONF(NETDEV_UP): enp1s0: link is not ready
Jun 13 07:40:29 cpu-5110 kernel: [ 17.479470] igb 0000:01:00.0: enp1s0: igb: enp1s0 NIC Link is Up 1000 Mbps Full Duplex, Flow Control: RX
Jun 13 07:40:29 cpu-5110 kernel: [ 17.480627] ADDRCONF(NETDEV_CHANGE): enp1s0: link becomes ready
Jun 13 07:40:35 cpu-5110 kernel: [ 23.509432] venet0: no IPv6 routers present
Jun 13 07:40:40 cpu-5110 kernel: [ 28.043020] enp1s0: no IPv6 routers present
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061149] dmidecode: Corrupted page table at address 7f6b1bcc000
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061174] Kernel PGD 80000010641e5067 PUD 10641aa067 PMD 1065da1067 PTE ffffffff8b235225
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061209] User PGD 10641e5067 PUD 10641aa067 PMD 1065da1067 PTE ffffffff8b235225
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061243] Bad pagetable: 000d [#1] SMP
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061261] last sysfs file: /sys/firmware/efi/systab
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061280] CPU 3
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061288] Modules linked in: vzethdev vznetdev pio_kaio pio_nfs pio_direct pfmt_raw pfmt_ploop1 ploop simfs vzrst nf_nat nf_contrack_ipv4 nf_defrag_ipv4 vzcpt nfs nfs_acl auth_rpcgss fscache lockd sunrpc nf_contrack vziolimit vzdquota vzmon vzdev ip6t_REJECT ip6table_mangle ip6table_filter ip6_tables xt_length xt_hi xt_tcpmss xt_TCPMSS iptable_mangle iptable_filter xt_multiport xt_limit xt_dscp ipt_REJECT ip_tables vzevent ipv6 vfat fat snd_hda_codec_hdmi snd_hda_codec_realtek snd_hda_codec_generic snd_hda_intel snd_hda_codec snd_pcsp snd_hwdep snd_pcm snd_page_alloc snd_timer snd soundcore shpchp serio_raw video wmi output ext4 jbd2 mbcache btrfs(T) lzo_compress lzo_decompress zlib_deflate raid10 linear pata_via netxen_nic 3w_9xxx qlge ixgbe mdio sata_nv forcedeth via686a mptctl mptsas scsi_transport_sas mptspi mptscsih mptbase dm_crypt raid456 async_raid6_recov async_memcpy async_pq async_xor xor async_tx raid6_pq raid0 raid1 dm_mirror dm_region_hash dm_log dm_mod sata_via ata_piix sata_sis pata_sis sym53c8xx megaraid_sas aic7xxx scsi_transport_spi sd_mod crc_t10dif 3w_xxxx sky2 r8169 skge e1000e e1000 via_rhine sis900 8139too e100 mii igb i2c_algo_bit dca ptp pps_core ahci nvme xhci_hcd i2c_core
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061840]
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061849] Pid: 2563, comm: dmidecode veid: 0 Tainted: G
```

```

-- ----- T 2.6.32-openvz-042stab130.1-amd64 #1 042stab130 System manufacturer
System Product Name/STRIX Z270F GAMING
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061917] RIP: 0033:[<00007f6b1bac8b51>]
[<00007f6b1bac8b51>] 0x7f6b1bac8b51
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061958] RSP: 002b:00007fff719095f0 EFLAGS:
00010213
Jun 13 07:41:01 cpu-5110 kernel: [ 49.061980] RAX: 00007f6b1bccc000 RBX:
00007f6b1cda8020 RCX: 0000000000000020
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062009] RDX: 0000000000000001 RSI:
00007f6b1bccc000 RDI: 00007f6b1cda8020
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062038] RBP: 0000000000000020 R08:
0000000000000003 R09: ffffffff8b235000
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062067] R10: 0000000000000001 R11:
0000000000000246 R12: 00007f6b1bacaef8
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062094] R13: 0000000000000003 R14:
0000000000000000 R15: ffffffff8b235000
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062123] FS: 00007f6b1bcc5700(0000)
GS:ffff8800282c0000(0000) knlGS:0000000000000000
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062154] CS: 0010 DS: 0000 ES: 0000 CR0:
0000000080050033
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062177] CR2: 00007f6b1bccc000 CR3:
000000106694c000 CR4: 0000000001607e0
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062205] DR0: 0000000000000000 DR1:
0000000000000000 DR2: 0000000000000000
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062233] DR3: 0000000000000000 DR6: 00000000fffe0ff0
DR7: 0000000000000400
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062262] Process dmidecode (pid: 2563, veid: 0,
threadinfo ffff881063bac000, task ffff881067d6a680)
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062297]
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062304] RIP [<00007f6b1bac8b51>] 0x7f6b1bac8b51
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062330] RSP <00007fff719095f0>
Jun 13 07:41:01 cpu-5110 kernel: [ 49.062344] Tainting kernel with flag 0x7
Jun 13 07:41:01 cpu-5110 kernel: [ 49.063085] Pid: 2563, comm: dmidecode veid: 0 Tainted: G
-- ----- T 2.6.32-openvz-042stab130.1-amd64 #1
Jun 13 07:41:01 cpu-5110 kernel: [ 49.063971] Call Trace:
Jun 13 07:41:01 cpu-5110 kernel: [ 49.064951] [<ffffffff810855e1>] ? add_taint+0x71/0x80
Jun 13 07:41:01 cpu-5110 kernel: [ 49.065933] [<ffffffff81561314>] ? oops_end+0x54/0x100
Jun 13 07:41:01 cpu-5110 kernel: [ 49.066870] [<ffffffff81053809>] ? pgtable_bad+0x99/0xb0
Jun 13 07:41:01 cpu-5110 kernel: [ 49.067785] [<ffffffff8105472a>] ?
__do_page_fault+0x3fa/0x500
Jun 13 07:41:01 cpu-5110 kernel: [ 49.068540] [<ffffffff81188aed>] ?
do_mmap_pgoff+0x33d/0x3a0
Jun 13 07:41:01 cpu-5110 kernel: [ 49.069336] [<ffffffff81171e29>] ?
sys_mmap_pgoff+0x1c9/0x380
Jun 13 07:41:01 cpu-5110 kernel: [ 49.070137] [<ffffffff8156330e>] ? do_page_fault+0x3e/0xa0
Jun 13 07:41:01 cpu-5110 kernel: [ 49.070869] [<ffffffff81560265>] ? page_fault+0x25/0x30
Jun 13 07:41:01 cpu-5110 kernel: [ 49.071431] ---[ end trace ed120bd972933c2c ]---

```

a reply to be appreciated .

Subject: Re: Kernel Panic after booting to openvz2.6 kernel
Posted by [vaverin](#) on Wed, 13 Jun 2018 12:23:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Bad pagetable: 000d"
0xd -- so called error code, it describes the reason of the page fault.

```
/*  
 * Page fault error code bits:  
 *  
 * bit 0 == 0: no page found      1: protection fault  
 * bit 1 == 0: read access        1: write access  
 * bit 2 == 0: kernel-mode access 1: user-mode access  
 * bit 3 ==                      1: use of reserved bit detected  
 * bit 4 ==                      1: fault was an instruction fetch  
*/
```

In this case kernel was not expected that bit3 us set.

I never saw such kind of issues before.

Error code is generated by CPU, so reason of this incident should be explained in CPU documentation.

Could you please clarify exact version of CPU used on affected node?

In general I can advise to update firmware (i.e motherboard BIOS) on affected node.

Thank you,
Vasily Averin

Subject: Re: Kernel Panic after booting to openvz2.6 kernel
Posted by [vaverin](#) on Wed, 13 Jun 2018 12:51:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

I've found related article in Red Hat knowledebase
<https://access.redhat.com/solutions/1379213>
So I'm going to follow these instructions to clarify the situation.

Also, just for record: Arun answered he uses overclocked Intel(R) Core(TM) i7-7700K CPU @ 4.20GHz

Subject: Re: Kernel Panic after booting to openvz2.6 kernel

Posted by [vaverin](#) on Wed, 13 Jun 2018 15:03:57 GMT

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

In our case userspace application dmidecde tried to access virtual memory address 7f6b1bccc000

CPU translated virtual memory address to physical by using few levels of page table directories, It finds address of top-level directory, reads its content, analyses it, found address of next-level directory, and so on. Finally it found page table entry that points to physical memory address.

Message below shows content of these directories.

```
"User PGD 10641e5067 PUD 10641aa067 PMD 1065da1067 PTE ffffffff8b235225"
```

I've checked what means these numbers, and found that 3 first entries (PGD, PUD and PMD) looks correct, but last one (PTE) looks incorrect. Some bits in this entry are reserved and should be set to 0, but in our case all of them are set into 1.

So as far as I understand CPU worked correctly and incident was caused by memory corruption.

Unfortunately I can say nothing about reason of this memory corruption, it can be both software and hardware-related.

We do not have such bugreports from other nodes, so if it is software-related issue -- it can be caused by some rarely used driver, for example self-compiled btrfs driver is guilty here.

On the other hand I cannot exclude hardware-related issue, over-clocked node can corrupt the memory.

Thank you,
Vasily Averin
