
Subject: OpenVZ,GFS2 and locks_remove_flock kernel BUG

Posted by [juriskrumins](#) on Mon, 07 Jan 2013 09:17:31 GMT

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

We've experienced kernel panic using 2.6.32-042stab065.3 openvz kernel version.
Running VZ CT ot top of GFS2 fs volume. Here is full kernel backtrace:

```
[20319.037070] GFS2: fsid=tengu2legion:vz22091.0: fatal: filesystem consistency error
[20319.037074] GFS2: fsid=tengu2legion:vz22091.0:  inode = 1784866 244410
[20319.037078] GFS2: fsid=tengu2legion:vz22091.0:  function = gfs2_dinode_dealloc, file =
fs/gfs2/super.c, line = 1384
[20319.037086] GFS2: fsid=tengu2legion:vz22091.0: about to withdraw this file system
[20319.037123] GFS2: fsid=tengu2legion:vz22091.0: telling LM to unmount
[22385.471217] -----[ cut here ]-----
[22385.471240] kernel BUG at fs/locks.c:2079!
[22385.471255] invalid opcode: 0000 [#1] SMP
[22385.471277] last sysfs file: /sys/devices/virtual/block/dm-3/dev
[22385.471296] CPU 19
[22385.471306] Modules linked in: vzethdev vznetdev pio_nfs pio_direct pfmt_raw pfmt_ploop1
ploop simfs vzrst vzcpt nfs lockd fscache nfs_acl auth_rpcgss vzmon ip6t_REJECT
ip6table_mangle ip6table_filter ip6_tables xt_owner xt_recent xt_mac ipt_REDIRECT nf_nat_irc
nf_nat_ftp xt_helper xt_conntrack nf_conntrack_irc nf_conntrack_ftp xt_length ipt_LOG xt_hl
xt_tcpmss xt_TCPMSS xt_DSCP xt_dscp xt_multiport xt_limit gfs2 ebtable_nat ebtables
ipt_MASQUERADE xt_CHECKSUM vxdquota xt_state ipt_REJECT iptable_nat nf_nat
nf_conntrack_ipv4 nf_conntrack nf_defrag_ipv4 iptable_mangle iptable_filter ip_tables vzevent
drbd dlm configfs sunrpc bridge vzdev bonding 8021q garp stp llc ipv6 vhost_net macvtap
macvlan tun kvm_intel kvm sg ses enclosure igb cdc_ether usbnet mii serio_raw i2c_i801
i2c_core iTCO_wdt iTCO_vendor_support ioatdma dca i7core_edac edac_core shpchp ext4
mbcache jbd2 sd_mod crc_t10dif pata_acpi ata_generic ata_piix megaraid_sas dm_mirror
dm_region_hash dm_log dm_mod [last unloaded: ploop]
[22385.472016]
[22385.472101] Pid: 161511, comm: atd veid: 22091 Not tainted 2.6.32-042stab065.3 #1
042stab065_3 IBM System x3630 M3 -[7377F2G]-/69Y1101
[22385.472293] RIP: 0010:[<ffffffff811e6bf3>] [<ffffffff811e6bf3>]
locks_remove_flock+0x103/0x130
[22385.472475] RSP: 0018:ffff8817bbc6dcd8 EFLAGS: 00010246
[22385.472567] RAX: 0000000000000001 RBX: ffff88307288b8c0 RCX: 00000000000007baf
[22385.472663] RDX: ffff8817be2682c0 RSI: 0000000000000002 RDI: ffff88303a67f1c0
[22385.472759] RBP: ffff8817bbc6dda8 R08: 0000000000000001 R09: 0000000000000002
[22385.472854] R10: 0000000000000000 R11: 0000000000000002 R12: ffff882fdc01d150
[22385.472948] R13: ffff8817bbc6dcd8 R14: ffff8817d829cd20 R15: ffff88180440b500
[22385.473044] FS: 0000000000000000(0000) GS:ffff8818e0bc0000(0000)
knIGS:0000000000000000
[22385.473215] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[22385.473306] CR2: 00007fc36eca2020 CR3: 0000000001a85000 CR4: 000000000000006e0
[22385.473402] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
[22385.473498] DR3: 0000000000000000 DR6: 00000000ffff0ff0 DR7: 00000000000000400
[22385.473594] Process atd (pid: 161511, veid: 22091, threadinfo ffff8817bbc6c000, task
```

```
ffff8817be2682c0)
[22385.473773] Stack:
[22385.473859] 0000000000000000 0000000000000000 0000000000000000
0000000000000000
[22385.473970] <d> 0000000000000000 0000000000000000 000276e700020002
0000000000000000
[22385.474160] <d> 0000000000000000 0000000000000000 0000000000000000
ffff88307288b8c0
[22385.474430] Call Trace:
[22385.474523] [<ffffffff81199f30>] __fput+0xd0/0x280
[22385.474617] [<ffffffff8119a105>] fput+0x25/0x30
[22385.474711] [<ffffffff8119525d>] filp_close+0x5d/0x90
[22385.474809] [<ffffffff8107189f>] put_files_struct+0x7f/0xf0
[22385.474902] [<ffffffff81071963>] exit_files+0x53/0x70
[22385.474995] [<ffffffff810735f5>] do_exit+0x1b5/0x920
[22385.475091] [<ffffffff810e814a>] ? audit_syscall_entry+0x26a/0x290
[22385.475188] [<ffffffff81073db8>] do_group_exit+0x58/0xd0
[22385.475283] [<ffffffff81073e47>] sys_exit_group+0x17/0x20
[22385.475380] [<ffffffff8100b182>] system_call_fastpath+0x16/0x1b
[22385.475472] Code: 49 89 c4 49 8b 04 24 48 85 c0 75 ee e8 d7 7f 31 00 48 81 c4 b8 00 00 00
5b 41 5c 41 5d c9 c3 0f b7 40 30 a8 02 75 08 a8 20 75 14 <0f> 0b eb fe 4c 89 e7 66 0f 1f 44 00
00 e8 db fc ff ff eb b2 be
[22385.475993] RIP [<ffffffff811e6bf3>] locks_remove_flock+0x103/0x130
[22385.476092] RSP <ffff8817bbc6dcd8>
```

We've also found some similar bug in RedHat bugzilla for the version 4 and 5 of RHEL: BUG ID 456282

So maybe anybody from OpenVZ project can take a look at this and maybe comment since it's community OpenVZ kernel and you have secret knowledge :).

Thanks in advance.
Juris Krumins.
