
Subject: Re: Is there a stable OpenVZ kernel, and which should be fit for production

Posted by [MailingListe](#) on Wed, 23 Nov 2011 16:59:29 GMT

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

Zitat von Kir Kolyshkin <kir@openvz.org>:

> On 11/23/2011 04:31 PM, Dariush Pietrzak wrote:

>>> I am very sad to hear this. Could you please file a bug to

>>> bugzilla.openvz.org so our kernel guys will start working on that?

>> Looking at bugzilla there are many other similiar reports, one of mine has

>> been closed as fixed, but then returned in exactly the same function after

>> just 6 minutes of stress-testing new kernel.

>> It's easy to reproduce, just put enough load on the system.

>

> Have you reopened it already? Can you provide bug number?

>

>>

>> It looks really troubling, both vSwap and 042.x branches look very nice

>> feature-wise, even vzmigrate seems to work fine, which is no small feat,

>> but it kinda feels like stability has been sacrificed to get there.

>>

>> best regards, Eyck

>

> Guys,

>

> I do understand reasons for your frustration, but so far I have only

> seen one specific bug mentioned in this thread, namely

> <http://bugzilla.openvz.org/2095> it was filed yesterday and there is

> a patch already available for testing. Any other statements like

> "there are many bugs", "this kernel is unstable" are just not

> specific enough for me to deal with.

>

> If there are bugs, they need to be reported and fixed, and we,

> OpenVZ team, partly rely on you, our users. We do have internal QA

> but can't possibly test all the use cases and scenarios.

>

> Specifically, we rely on having bug reports from you, with full

> kernel logs (see http://wiki.openvz.org/Remote_console_setup), test

> cases (as specific and reproducible as possible), and ideally your

> ability to test patches that developers provide and report your

> results back to bugzilla.

Okay, can someone with a bugzilla account please confirm and create a bug with this one:

Kernel (uname -a): 2.6.32-042stab039.11 as x86_64 installed on CentOS 6 HN

steps to reproduce:

- use a Ubuntu 10.04 i386 (32 bit) template and create a VE
- load ipt_recent in vz.conf
- start the VE and use something like this "iptables -A INPUT -p tcp --dport 25 -m conntrack --ctstate NEW -m recent --name SMTP --set" inside the VE
- exit the VE and execute vzctl stop <VE-ID> at the HN

The result is the following kernel panic

```
[ 202.003789] libfcoe_device_notification: NETDEV_UNREGISTER venet0
[ 202.050580] libfcoe_device_notification: NETDEV_UNREGISTER lo
[ 202.089227] BUG: unable to handle kernel NULL pointer dereference
at 00000000000000038
[ 202.089250] IP: [<fffffffa05690ff>] fini_ipt_recent+0x1f/0x50 [xt_recent]
[ 202.089271] PGD 6bcb2067 PUD 64d37067 PMD 0
[ 202.089290] Oops: 0000 [#1] SMP
[ 202.089309] last sysfs file:
/sys/devices/pci0000:00/0000:00:07.0/net/eth0/type
[ 202.089322] CPU 0
[ 202.089329] Modules linked in: netconsole configfs vzevent simfs
vzrst nf_nat vzcpt nfs lockd fscache nfs_acl auth_rpcgss vzdquota
xt_conntrack ip6t_REJECT ip6table_mangle ip6table_filter ip6_tables
nf_conntrack_ftp nf_conntrack_ipv4 nf_conntrack nf_defrag_ipv4
xt_recent xt_length xt_hl xt_tcpmss xt_TCPMSS iptable_mangle
iptable_filter xt_multiport xt_limit xt_dscp ipt_REJECT ip_tables
vzevent fcoe libfcoe libfc scsi_transport_fc scsi_tgt 8021q garp stp
llc sunrpc vznetdev vzmon vzdev ipv6 ppdev parport_pc parport k10temp
hwmon edac_core edac_mce_amd shpchp sg snd_hda_codec_via snd_hda_intel
snd_hda_codec snd_hwdep snd_seq snd_seq_device snd_pcm snd_timer snd
soundcore snd_page_alloc i2c_nforce2 ext4 mbcache jbd2 sr_mod cdrom
sd_mod crc_t10dif pata_amd ata_generic pata_acpi sata_nv forcedeth
nouveau ttm drm_kms_helper drm i2c_algo_bit i2c_core video output
dm_mod [last unloaded: scsi_wait_scan]
[ 202.089691]
[ 202.089696] Modules linked in: netconsole configfs vzevent simfs
vzrst nf_nat vzcpt nfs lockd fscache nfs_acl auth_rpcgss vzdquota
xt_conntrack ip6t_REJECT ip6table_mangle ip6table_filter ip6_tables
nf_conntrack_ftp nf_conntrack_ipv4 nf_conntrack nf_defrag_ipv4
xt_recent xt_length xt_hl xt_tcpmss xt_TCPMSS iptable_mangle
iptable_filter xt_multiport xt_limit xt_dscp ipt_REJECT ip_tables
vzevent fcoe libfcoe libfc scsi_transport_fc scsi_tgt 8021q garp stp
llc sunrpc vznetdev vzmon vzdev ipv6 ppdev parport_pc parport k10temp
hwmon edac_core edac_mce_amd shpchp sg snd_hda_codec_via snd_hda_intel
snd_hda_codec snd_hwdep snd_seq snd_seq_device snd_pcm snd_timer snd
soundcore snd_page_alloc i2c_nforce2 ext4 mbcache jbd2 sr_mod cdrom
sd_mod crc_t10dif pata_amd ata_generic pata_acpi sata_nv forcedeth
nouveau ttm drm_kms_helper drm i2c_algo_bit i2c_core video output
```

```
dm_mod [last unloaded: scsi_wait_scan]
[ 202.090036] Pid: 25, comm: netns Not tainted 2.6.32-042stab039.11
#1 042stab039_11 To Be Filled By O.E.M.
[ 202.090047] RIP: 0010:[<ffffffffffa05690ff>] [<ffffffffffa05690ff>]
fini_ipt_recent+0x1f/0x50 [xt_recent]
[ 202.090065] RSP: 0018:ffff88006f115ce0 EFLAGS: 00010282
[ 202.090073] RAX: 0000000000000000 RBX: ffff8800378b3800 RCX:
0000000000000000
[ 202.090083] RDX: 0000000000000003 RSI: 0000000000000000 RDI:
ffffffffffa056a3ac
[ 202.090094] RBP: ffff88006f115cf0 R08: 0000000000000000 R09:
000000000000001f8
[ 202.090103] R10: ffff8800378a6000 R11: 0000000000000000 R12:
ffff88006bb058e0
[ 202.090112] R13: ffff88006bb05000 R14: ffff88006bb058e0 R15:
0000000000000080
[ 202.090123] FS: 00007f22b35fe700(0000) GS:ffff880002600000(0000)
knlGS:00000000b772b8d0
[ 202.090193] CS: 0010 DS: 0018 ES: 0018 CR0: 000000008005003b
[ 202.090193] CR2: 0000000000000038 CR3: 00000000372e1000 CR4:
000000000000006f0
[ 202.090193] DR0: 0000000000000000 DR1: 0000000000000000 DR2:
0000000000000000
[ 202.090193] DR3: 0000000000000000 DR6: 00000000ffff0ff0 DR7:
00000000000000400
[ 202.090193] Process netns (pid: 25, veid=0, threadinfo
ffff88006f114000, task ffff88006f112ec0)
[ 202.090193] Stack:
[ 202.090193] ffff88006bb058e0 ffff8800378b3800 ffff88006f115d30
ffffffffffa05692e9
[ 202.090193] <0> ffff88006b099800 00000000000000158 ffff88006b099958
ffff88006b099800
[ 202.090193] <0> ffff88006b099800 fffffffffffa0545660 ffff88006f115d60
ffffffffffa05222c5
[ 202.090193] Call Trace:
[ 202.090193] [<ffffffffffa05692e9>] recent_mt_destroy+0x149/0x150 [xt_recent]
[ 202.090193] [<ffffffffffa05222c5>] cleanup_match+0x45/0x60 [ip_tables]
[ 202.090193] [<ffffffffff810a3f77>] ? uncharge_beancounter+0x57/0x70
[ 202.090193] [<ffffffffffa05223d5>] cleanup_entry+0x65/0xc0 [ip_tables]
[ 202.090193] [<ffffffffffa0524def>] ipt_unregister_table+0x5f/0x90
[ip_tables]
[ 202.090193] [<ffffffffffa054502c>] iptable_filter_net_exit+0x2c/0x30
[iptable_filter]
[ 202.090193] [<ffffffffff8140a94e>] cleanup_net+0x8e/0xe0
[ 202.090193] [<ffffffffff8140a8c0>] ? cleanup_net+0x0/0xe0
[ 202.090193] [<ffffffffff8108c220>] worker_thread+0x190/0x2d0
[ 202.090193] [<ffffffffff81092680>] ? autoremove_wake_function+0x0/0x40
[ 202.090193] [<ffffffffff8108c090>] ? worker_thread+0x0/0x2d0
```

```
[ 202.090193] [<ffffffff810920a6>] kthread+0x96/0xa0
[ 202.090193] [<ffffffff8100c2ca>] child_rip+0xa/0x20
[ 202.090193] [<ffffffff81092010>] ? kthread+0x0/0xa0
[ 202.090193] [<ffffffff8100c2c0>] ? child_rip+0x0/0x20
[ 202.090193] Code: 1c 24 c9 c3 0f 1f 84 00 00 00 00 00 55 48 89 e5
53 48 83 ec 08 0f 1f 44 00 00 48 8b 87 70 03 00 00 48 89 fb 48 c7 c7
ac a3 56 a0 <48> 8b 70 38 e8 68 68 c8 e0 48 8b bb 20 02 00 00 e8 ec 9f
c0 e0
[ 202.090193] RIP [<ffffffffffa05690ff>] fini_ipt_recent+0x1f/0x50 [xt_recent]
[ 202.090193] RSP <ffff88006f115ce0>
[ 202.090193] CR2: 0000000000000038
[ 202.120364] ---[ end trace 6b08bce91c3d45d5 ]---
[ 202.121379] Kernel panic - not syncing: Fatal exception
[ 202.122386] Pid: 25, comm: netns Tainted: G    D
----- 2.6.32-042stab039.11 #1
[ 202.122389] Call Trace:
[ 202.122398] [<ffffffff814c3a31>] ? panic+0x78/0x143
[ 202.122402] [<ffffffff814c7d14>] ? oops_end+0xe4/0x100
[ 202.122407] [<ffffffff81040c5b>] ? no_context+0xfb/0x260
[ 202.122410] [<ffffffff81040ed5>] ? __bad_area_nosemaphore+0x115/0x1e0
[ 202.122413] [<ffffffff81040fb3>] ? bad_area_nosemaphore+0x13/0x20
[ 202.122417] [<ffffffff8104168d>] ? __do_page_fault+0x31d/0x480
[ 202.122420] [<ffffffff814c427a>] ? thread_return+0x4e/0x854
```

Thanks

Andreas

File Attachments

1) [smime.p7s](#), downloaded 562 times
