

---

## Subject: Problems with NFS?

Posted by [arpad](#) on Mon, 26 Mar 2007 14:16:39 GMT

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

---

Hi!

I have a VE running UNFSd and another two VE which are using the first ones as NFS servers. Today one of the VE's lost the NFS servers connection and a process stucked in the D state (I opened a topic from it because I can't shutdown that VE). After several tries I get in to the dmesg and I have to reboot the server:

```
nfs: server nfs not responding, still trying
nfs: server nfs not responding, still trying
nfs: server nfs not responding, still trying
Unable to handle kernel NULL pointer dereference at 0000000000000068 RIP:
[<ffffff8025d438>] __ip_route_output_key+0x108/0x910
PGD 2c60a067 PUD 2b1f1067 PMD 0
Oops: 0000 [1] SMP
CPU: 1
Modules linked in: simfs vznetdev vzrst vzcpt vzdquota nfsd
exportfs it87 hwmon_vid hwmon i2c_isa ip_vs_rr ip_vs bridge llc
iptable_filter ipt_MASQUERADE iptable_nat ip_nat ip_conntrack
nfnetlink ip_tables xt_tcpudp ip6table_filter ip6_tables
x_tables vzethdev vzmon ipv6 vzdev tun crc32 i2c_nforce2
i2c_core forcedeth sg dm_snapshot e1000 nfs lockd sunrpc
dm_mirror usbhid ohci_hcd usb_storage ehci_hcd usbcore
Pid: 2960, comm: rpciod/1 Not tainted 2.6.18-028test018-syrius #1
RIP: 0060:[<ffffff8025d438>] [<ffffff8025d438>] __ip_route_output_key+0x108/0x910
RSP: 0068:ffff81003e76b980 EFLAGS: 00010246
RAX: 0000000000000000 RBX: 0000000000000000 RCX: 0000000000000000
RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000000
RBP: ffff81003e76bab0 R08: 0000000000000000 R09: ffff81003e76b9c0
R10: ffff81002c989040 R11: 0000000000000000 R12: 0000000000000000
R13: ffff81003e76b9c0 R14: 000000000000000a R15: 0000000000000000
FS: 00002b0cb79f73e0(0000) GS:ffff81003ee52400(0000) knlGS:00000000b7ba56c0
CS: 0060 DS: 0068 ES: 0068 CR0: 000000008005003b
CR2: 0000000000000068 CR3: 0000000027703000 CR4: 000000000000006e
Process rpciod/1 (pid: 2960, veid=0, threadinfo ffff81003e76a000, task ffff81003e49e000)
Stack: 0000000000000000 0000000000000000 ffff81003d104024 ffff81003e76bb18
0000000000000000 ffffffff80412fa9 000000000000005dc 0000000000000000
0000000000000000 0000000000000000 0000000000000000 0000000000000000
Call Trace:
[<ffffff80412fa9>] ip_generic_getfrag+0x49/0xa0
[<ffffff8040d109>] ip_route_output_flow+0x19/0x70
[<ffffff80256bf6>] udp_sendmsg+0x336/0x660
[<ffffff80259524>] sock_sendmsg+0xd4/0xf0
[<ffffff8029d2d0>] autoremove_wake_function+0x0/0x30
```

```

[<ffffff8022fb73>] __wake_up+0x43/0x70
[<ffffff88059c90>] :sunrpc:rpc_async_schedule+0x0/0x10
[<ffffff80216488>] release_console_sem+0x58/0x70
[<ffffff8042be80>] inet_sendpage+0x70/0xe0
[<ffffff803ed545>] kernel_sendmsg+0x35/0x50
[<ffffff880587d1>] :sunrpc:xs_udp_send_request+0xf1/0x2e0
[<ffffff8805901a>] :sunrpc:rpc_sleep_on+0x3a/0x60
[<ffffff88059c90>] :sunrpc:rpc_async_schedule+0x0/0x10
[<ffffff88055d6b>] :sunrpc:xprt_transmit+0x12b/0x260
[<ffffff880a3260>] :nfs:nfs3_xdr_writeargs+0x0/0xa0
[<ffffff88053577>] :sunrpc:call_transmit+0x207/0x270
[<ffffff88059b28>] :sunrpc:__rpc_execute+0xc8/0x230
[<ffffff80251522>] run_workqueue+0xb2/0x110
[<ffffff8024d560>] worker_thread+0x0/0x170
[<ffffff8029d090>] keventd_create_kthread+0x0/0x80
[<ffffff8024d689>] worker_thread+0x129/0x170
[<ffffff802845b0>] default_wake_function+0x0/0x10
[<ffffff8024d560>] worker_thread+0x0/0x170
[<ffffff802343e9>] kthread+0xd9/0x120
[<ffffff80262fe0>] child_rip+0xa/0x12
[<ffffff8029d090>] keventd_create_kthread+0x0/0x80
[<ffffff802758a0>] flat_send_IPI_mask+0x0/0x50
[<ffffff80234310>] kthread+0x0/0x120
[<ffffff80262fd6>] child_rip+0x0/0x12

```

Code: 8b 40 68 89 7c 24 4c 44 88 64 24 54 88 5c 24 55 89 44 24 44

RIP [<ffffff8025d438>] \_\_ip\_route\_output\_key+0x108/0x910

RSP <ffff81003e76b980>

CR2: 0000000000000068

Unable to handle kernel NULL pointer dereference at 0000000000000020 RIP:

[<ffffff802f2a24>] proc\_flush\_task+0x54/0x70

PGD 2c60a067 PUD 2b1f1067 PMD 0

Oops: 0000 [2] SMP

CPU: 1

Modules linked in: simfs vznetdev vzrst vzcpt vzdquota nfsd  
 exportfs it87 hwmon\_vid hwmon i2c\_isa ip\_vs\_rr ip\_vs bridge llc  
 iptable\_filter ipt\_MASQUERADE iptable\_nat ip\_nat ip\_contrack  
 nfnetlink ip\_tables xt\_tcpudp ip6table\_filter ip6\_tables  
 x\_tables vzethdev vzmon ipv6 vzdev tun crc32 i2c\_nforce2  
 i2c\_core forcedeth sg dm\_snapshot e1000 nfs lockd sunrpc  
 dm\_mirror usbhid ohci\_hcd usb\_storage ehci\_hcd usbcore

Pid: 2960, comm: rpciod/1 Not tainted 2.6.18-028test018-syrius #1

RIP: 0060:[<ffffff802f2a24>] [<ffffff802f2a24>] proc\_flush\_task+0x54/0x70

RSP: 0068:ffff81003e76b6e0 EFLAGS: 00010287

RAX: 0000000000000000 RBX: ffff81003e49e000 RCX: 0000000000000102

RDY: 00000000000000b90 RSI: 00000000000000b90 RDI: ffff81003e49e000

RBP: ffff81003e49e000 R08: 000000000000039c R09: ffff81000688c988

R10: ffff81000688c9e8 R11: ffffffff802f1e50 R12: 0000000000000000  
R13: ffff81003e49e0c8 R14: ffffffff804ae500 R15: ffff81003e49e148  
FS: 00002b0cb79f73e0(0000) GS:ffff81003ee52400(0000) knlGS:00000000b7ba56c0  
CS: 0060 DS: 0068 ES: 0068 CR0: 000000008005003b  
CR2: 0000000000000020 CR3: 0000000027703000 CR4: 000000000000006e0  
Process rpciod/1 (pid: 2960, veid=0, threadinfo ffff81003e76a000, task ffff81003e49e000)  
Stack: ffffffff80216eff ffff81003e76b728 0000000000000020 ffff81003e49e000  
fffffff80215211 0000003000000010 ffff81003e76b728 ffff81003e49e1b8  
0000000000000009 ffff81003e76b728 ffff81003e76b728 ffffffff804a9c08

#### Call Trace:

[<fffffff80216eff>] release\_task+0x31f/0x370  
[<fffffff80215211>] do\_exit+0xc1/0xda0  
[<fffffff80370af0>] vgacon\_cursor+0x0/0x1c7  
[<fffffff80267c38>] \_spin\_unlock\_irqrestore+0x8/0x10  
[<fffffff8020a9ed>] do\_page\_fault+0x71d/0x7e0  
[<fffffff8022d03b>] local\_bh\_enable+0x8b/0xb0  
[<fffffff80231729>] dev\_queue\_xmit+0x259/0x280  
[<fffffff80412683>] ip\_fragment+0x363/0x830  
[<fffffff80412d30>] ip\_finish\_output2+0x0/0x190  
[<fffffff80233799>] ip\_output+0x109/0x270  
[<fffffff80262e25>] error\_exit+0x0/0x84  
[<fffffff8025d438>] \_\_ip\_route\_output\_key+0x108/0x910  
[<fffffff80412fa9>] ip\_generic\_getfrag+0x49/0xa0  
[<fffffff8040d109>] ip\_route\_output\_flow+0x19/0x70  
[<fffffff80256bf6>] udp\_sendmsg+0x336/0x660  
[<fffffff80259524>] sock\_sendmsg+0xd4/0xf0  
[<fffffff8029d2d0>] autoremove\_wake\_function+0x0/0x30  
[<fffffff8022fb73>] \_\_wake\_up+0x43/0x70  
[<fffffff88059c90>] :sunrpc:rpc\_async\_schedule+0x0/0x10  
[<fffffff80216488>] release\_console\_sem+0x58/0x70  
[<fffffff8042be80>] inet\_sendpage+0x70/0xe0  
[<fffffff803ed545>] kernel\_sendmsg+0x35/0x50  
[<fffffff880587d1>] :sunrpc:xs\_udp\_send\_request+0xf1/0x2e0  
[<fffffff8805901a>] :sunrpc:rpc\_sleep\_on+0x3a/0x60  
[<fffffff88059c90>] :sunrpc:rpc\_async\_schedule+0x0/0x10  
[<fffffff88055d6b>] :sunrpc:xprt\_transmit+0x12b/0x260  
[<fffffff880a3260>] :nfs:nfs3\_xdr\_writeargs+0x0/0xa0  
[<fffffff88053577>] :sunrpc:call\_transmit+0x207/0x270  
[<fffffff88059b28>] :sunrpc:\_\_rpc\_execute+0xc8/0x230  
[<fffffff80251522>] run\_workqueue+0xb2/0x110  
[<fffffff8024d560>] worker\_thread+0x0/0x170  
[<fffffff8029d090>] keventd\_create\_kthread+0x0/0x80  
[<fffffff8024d689>] worker\_thread+0x129/0x170  
[<fffffff802845b0>] default\_wake\_function+0x0/0x10  
[<fffffff8024d560>] worker\_thread+0x0/0x170  
[<fffffff802343e9>] kthread+0xd9/0x120  
[<fffffff80262fe0>] child\_rip+0xa/0x12  
[<fffffff8029d090>] keventd\_create\_kthread+0x0/0x80

```
[<ffffff802758a0>] flat_send_IPI_mask+0x0/0x50  
[<ffffff80234310>] kthread+0x0/0x120  
[<ffffff80262fd6>] child_rip+0x0/0x12
```

```
Code: 48 8b 48 20 e9 13 fd ff ff 66 66 90 5b c3 00 00 00 00 00  
RIP [<ffffff802f2a24>] proc_flush_task+0x54/0x70  
RSP <ffff81003e76b6e0>  
CR2: 0000000000000020  
Fixing recursive fault but reboot is needed!
```

After it I had random freezes during disk I/O (thats why I decided to reboot the machine). I think from this dump there is something bad with the NFS under OpenVZ. I'm currently running 2.6.18-028test018. The host and the guests are all Gentoo Linuxes. I deconfigured the NFS after reboot and now I'm using bind mount. I hope it helps you.

Arpad

---

---

Subject: Re: Problems with NFS?  
Posted by [Vasily Tarasov](#) on Tue, 27 Mar 2007 08:22:30 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hello,

thank you for information, I filled the bug [http://bugzilla.openvz.org/show\\_bug.cgi?id=513](http://bugzilla.openvz.org/show_bug.cgi?id=513). Seems, that you use a self-compiled kernel. Can you, please, attach the .config file you used to bug report.

Thank you,  
Vasily.

---