

## 2.6.10 devfs oops without devfs mounted at all

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-02/2284.html>

---

**From:** Sergey S. Kostyliov (*rathamahata\_at\_ehouse.ru*)

**Date:** 02/08/05

To: linux-kernel@vger.kernel.org  
Date: Tue, 8 Feb 2005 20:13:50 +0300

Hello all,

Here is an oops I've just get on my smp system:

```
Unable to handle kernel NULL pointer dereference at virtual address 0000001c
printing eip:
c01afe5b
*pde = 00000000
Oops: 0000 [#1]
PREEMPT SMP
Modules linked in: ipt_REJECT ipt_state ip_contrack iptable_filter
CPU: 2
EIP: 0060:[<c01afe5b>] Not tainted VLI
EFLAGS: 00010286 (2.6.10)
EIP is at devfsd_close+0x1b/0xc8
eax: f7440a00 ebx: 00000000 ecx: c01afe40 edx: ed395280
esi: 00000000 edi: f7f17800 ebp: f74f96c8 esp: cdc70f84
ds: 007b es: 007b ss: 0068
Process megamgr.bin (pid: 12844, threadinfo=cdc70000 task=dd81e520)
Stack: ed395280 ed395280 00000000 f7f17800 c0150c76 ee9e87f8 ed395280 00000000
       f1985c80 cdc70000 c014f50f 00000003 00000003 080caa60 00000000 c01024df
       00000003 080cc700 bfffe4f8 080caa60 00000000 bfffe4fc 00000006 0000007b
Call Trace:
 [<c0150c76>] __fput+0x106/0x120
 [<c014f50f>] filp_close+0x4f/0x80
 [<c01024df>] syscall_call+0x7/0xb
Code: fe ff ff ba ff ff ff ff e9 e5 fe ff ff 8d 76 00 83 ec 10 89 5c 24 04 89 74 24 08 89 7c 24 0c 8b 80 98 00 00
00 8b 98 68 01 00 00 <39> 53 1c 74 12 8b 5c 24 04 31 c0 8b 74 24 08 8b 7c 24 0c 83 c4
```

```
rathamahata@terror rathamahata $ dmesg | ksymoops -o /lib/modules/2.6.10/ -m
/usr/local/src/linux-2.6.10/System.map
ksymoops 2.4.9 on i686 2.6.10. Options used
-V (default)
-k /proc/ksyms (default)
-l /proc/modules (default)
-o /lib/modules/2.6.10/ (specified)
-m /usr/local/src/linux-2.6.10/System.map (specified)
```

## Linux-Kernel: 2.6.10 devfs oops without devfs mounted at all

```
Error (regular_file): read_ksyms stat /proc/ksyms failed
ksymoops: No such file or directory
No modules in ksyms, skipping objects
No ksyms, skipping lsmod
Unable to handle kernel NULL pointer dereference at virtual address 0000001c
c01afe5b
*pde = 00000000
Oops: 0000 [#1]
CPU: 2
EIP: 0060:[<c01afe5b>] Not tainted VLI
Using defaults from ksymoops -t elf32-i386 -a i386
EFLAGS: 00010286 (2.6.10)
eax: f7440a00 ebx: 00000000 ecx: c01afe40 edx: ed395280
esi: 00000000 edi: f7f17800 ebp: f74f96c8 esp: cdc70f84
ds: 007b es: 007b ss: 0068
Stack: ed395280 ed395280 00000000 f7f17800 c0150c76 ee9e87f8 ed395280 00000000
       f1985c80 cdc70000 c014f50f 00000003 00000003 080caa60 00000000 c01024df
       00000003 080cc700 bfffe4f8 080caa60 00000000 bfffe4fc 00000006 0000007b
Call Trace:
[<c0150c76>] __fput+0x106/0x120
[<c014f50f>] filp_close+0x4f/0x80
[<c01024df>] syscall_call+0x7/0xb
Code: fe ff ff ba ff ff ff ff e9 e5 fe ff ff 8d 76 00 83 ec 10 89 5c 24 04 89 74 24 08 89 7c 24 0c 8b 80 98 00 00
00 8b 98 68 01 00 00 <39> 53 1c 74 12 8b 5c 24 04 31 c0 8b 74 24 08 8b 7c 24 0c 83 c4
```

```
>>EIP; c01afe5b <devfsd_close+1b/c8> <=====
```

```
>>eax; f7440a00 <pg0+3708ea00/3fc4c400>
>>ecx; c01afe40 <devfsd_close+0/c8>
>>edx; ed395280 <pg0+2cfe3280/3fc4c400>
>>edi; f7f17800 <pg0+37b65800/3fc4c400>
>>ebp; f74f96c8 <pg0+371476c8/3fc4c400>
>>esp; cdc70f84 <pg0+d8bef84/3fc4c400>
```

```
Trace; c0150c76 <__fput+106/120>
Trace; c014f50f <filp_close+4f/80>
Trace; c01024df <syscall_call+7/b>
```

This architecture has variable length instructions, decoding before eip is unreliable, take these instructions with a pinch of salt.

```
Code; c01afe30 <devfsd_ioctl+150/160>
00000000 <_EIP>;
Code; c01afe30 <devfsd_ioctl+150/160>
0: fe (bad)
Code; c01afe31 <devfsd_ioctl+151/160>
1: ff (bad)
Code; c01afe32 <devfsd_ioctl+152/160>
2: ff (bad)
Code; c01afe33 <devfsd_ioctl+153/160>
3: ba ff ff ff ff mov $0xffffffff,%edx
```

## Linux-Kernel: 2.6.10 devfs oops without devfs mounted at all

```
Code; c01afe38 <devfsd_ioctl+158/160>
  8: e9 e5 fe ff ff jmp ffffffff2 <_EIP+0xffffffff2>
Code; c01afe3d <devfsd_ioctl+15d/160>
  d: 8d 76 00 lea 0x0(%esi),%esi
Code; c01afe40 <devfsd_close+0/c8>
 10: 83 ec 10 sub $0x10,%esp
Code; c01afe43 <devfsd_close+3/c8>
 13: 89 5c 24 04 mov %ebx,0x4(%esp)
Code; c01afe47 <devfsd_close+7/c8>
 17: 89 74 24 08 mov %esi,0x8(%esp)
Code; c01afe4b <devfsd_close+b/c8>
 1b: 89 7c 24 0c mov %edi,0xc(%esp)
Code; c01afe4f <devfsd_close+f/c8>
 1f: 8b 80 98 00 00 00 mov 0x98(%eax),%eax
Code; c01afe55 <devfsd_close+15/c8>
 25: 8b 98 68 01 00 00 mov 0x168(%eax),%ebx
```

This decode from eip onwards should be reliable

```
Code; c01afe5b <devfsd_close+1b/c8>
00000000 <_EIP>:
Code; c01afe5b <devfsd_close+1b/c8> <=====
  0: 39 53 1c cmp %edx,0x1c(%ebx) <=====
Code; c01afe5e <devfsd_close+1e/c8>
  3: 74 12 je 17 <_EIP+0x17>
Code; c01afe60 <devfsd_close+20/c8>
  5: 8b 5c 24 04 mov 0x4(%esp),%ebx
Code; c01afe64 <devfsd_close+24/c8>
  9: 31 c0 xor %eax,%eax
Code; c01afe66 <devfsd_close+26/c8>
  b: 8b 74 24 08 mov 0x8(%esp),%esi
Code; c01afe6a <devfsd_close+2a/c8>
  f: 8b 7c 24 0c mov 0xc(%esp),%edi
Code; c01afe6e <devfsd_close+2e/c8>
 13: 83 .byte 0x83
Code; c01afe6f <devfsd_close+2f/c8>
 14: c4 .byte 0xc4
```

1 error issued. Results may not be reliable.  
rathamahata@terror rathamahata \$

What is interesting is that devfs doesn't seem to be mounted at all ...

```
rathamahata@terror rathamahata $ cat /proc/mounts
rootfs / rootfs rw 0 0
/dev/root / reiserfs rw,noatime,nodiratime 0 0
none /proc proc rw,nodiratime 0 0
none /sys sysfs rw 0 0
none /dev ramfs rw 0 0
none /dev/pts devpts rw 0 0
/dev/sda5 /var/www reiserfs rw,noatime,nodiratime 0 0
```

## Linux-Kernel: 2.6.10 devfs oops without devfs mounted at all

```
/dev/sda6 /var/lib reiserfs rw,noatime,nodiratime 0 0  
/dev/sda7 /var/lib/mysql-innodb reiserfs rw,noatime,nodiratime 0 0  
/dev/sda8 /var/log reiserfs rw,noatime,nodiratime 0 0  
/dev/sda9 /var/log/innodb reiserfs rw,noatime,nodiratime 0 0
```

--

Sergey S. Kostyliov <rathamahata@ehouse.ru>

Jabber ID: rathamahata@jabber.org

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>