

2.6.9rc2-mm4 oops

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/9143.html>

From: Fernando Pablo Lopez-Lezcano (*nando_at_ccrma.Stanford.EDU*)

Date: 09/30/04

To: Andrew Morton <akpm@osdl.org>, linux-kernel <linux-kernel@vger.kernel.org>

Date: 30 Sep 2004 12:14:13 -0700

I ran into this while testing the voluntary preemption S7 patch. It still happens in vanilla 2.6.9rc2-mm4 (with the small patch below that Ingo sent me, otherwise plain rc2-mm4 would not boot). The machine is an athlon64 based workstation and has no floppy drive:

inserting floppy driver for 2.6.8.1-1.520.1nov.rhfc2.ccrma

Unable to handle kernel paging request at virtual address f8881920

printing eip:

c0251d3d

*pde = 37f5f067

Oops: 0002 [#1]

PREEMPT

Modules linked in: floppy(U) sg(U) dm_mod(U) uhci_hcd(U) ehci_hcd(U)

button(U) battery(U) asus_acpi(U) ac(U) ext3(U) jbd(U) raid5(U) xor(U)

sata_via(U) sata Promise(U) libata(U) sd_mod(U) scsi_mod(U)

CPU: 0

EIP: 0060:[<c0251d3d>] Not tainted VLI

EFLAGS: 00010246 (2.6.8.1-1.520.1nov.rhfc2.ccrma)

EIP is at acpi_bus_register_driver+0xd2/0x165

eax: f8881920 ebx: f88eefe0 ecx: c03d6b40 edx: f88ebd30

esi: ffffffff edi: f6ed4000 ebp: c03d9460 esp: f6ed4f5c

ds: 007b es: 007b ss: 0068

Process modprobe (pid: 2119, threadinfo=f6ed4000 task=f6e71870)

Stack: c03d94a0 f88e9126 00000015 00000014 f8870bc1 c03d94a0 f88ef280

f6ed4000

c0129ef7 c03d94a0 f88ef280 f6ed4000 c03d9460 c014dd52 00000246

f62f29ac

f6e1fc40 f6d8b8ac f6de4380 f6de43ac 00000000 b7fde008 0807a1a0

006a809d

Call Trace:

[<f88e9126>] acpi_floppy_init+0x16/0x50 [floppy]

[<f8870bc1>] floppy_init+0x11/0x600 [floppy]

[<c0129ef7>] printk+0x17/0x20

[<c014dd52>] sys_init_module+0x252/0x3b0

[<c0106afd>] sysenter_past_esp+0x52/0x71

Code: 00 00 00 a1 ec 67 3e c0 c7 05 78 67 3e c0 a0 7b 3a c0 c7 05 7c 67

3e c0 bd 01 00 00 89 1d ec 67 3e c0 c7 03 e8 67 3e c0 89 43 04 <89> 18

Linux-Kernel: 2.6.9rc2-mm4 oops

```
81 3d 68 67 3e c0 3c 4b 24 1d 74 1c 68 68 67 3e c0 68 bf
<6>note: modprobe[2119] exited with preempt_count 1
Debug: sleeping function called from invalid context at
include/linux/rwsem.h:43
in_atomic():1, irqs_disabled():0
[<c0125652>] __might_sleep+0xa2/0xb0
[<c012d222>] do_exit+0xa2/0x980
[<c010775f>] die+0x2bf/0x2c0
[<c012a0b6>] vprintk+0x1b6/0x340
[<c011e1e4>] do_page_fault+0x314/0x56c
[<c01d8155>] sysfs_new_dirent+0x25/0x80
[<c01d81cd>] sysfs_make_dirent+0x1d/0x90
[<c0172a29>] unmap_area_pmd+0x49/0x60
[<c01d7d84>] sysfs_add_file+0x74/0xa0
[<c0172bb0>] unmap_vm_area+0x30/0x80
[<c0173136>] __vunmap+0xb6/0xf0
[<c0129cf0>] call_console_drivers+0x80/0x110
[<c011ded0>] do_page_fault+0x0/0x56c
[<c0106cf9>] error_code+0x2d/0x38
[<c0251d3d>] acpi_bus_register_driver+0xd2/0x165
[<f88e9126>] acpi_floppy_init+0x16/0x50 [floppy]
[<f8870bc1>] floppy_init+0x11/0x600 [floppy]
[<c0129ef7>] printk+0x17/0x20
[<c014dd52>] sys_init_module+0x252/0x3b0
[<c0106afd>] sysenter_past_esp+0x52/0x71
ohci1394: $Rev: 1226 $ Ben Collins <bcollins@debian.org>
```

-- Fernando

```
--- linux/init/main.c.orig
+++ linux/init/main.c
@@ -435,6 +435,12 @@ static void noinline rest_init(void)
{
    kernel_thread(init, NULL, CLONE_FS | CLONE_SIGHAND);
    numa_default_policy();
+ /*
+ * Re-enable preemption but disable interrupts to make sure
+ * we dont get preempted until we schedule() in cpu_idle().
+ */
+ local_irq_disable();
+ preempt_enable_no_resched();
    unlock_kernel();
    cpu_idle();
}
@@ -501,6 +507,7 @@ asmlinkage void __init start_kernel(void
    * time - but meanwhile we still have a functioning scheduler.
    */
    sched_init();
+ preempt_disable();
    build_all_zonelists();
    page_alloc_init();
```

Linux-Kernel: 2.6.9rc2-mm4 oops

```
trap_init();
--- linux/arch/i386/kernel/entry.S.orig
+++ linux/arch/i386/kernel/entry.S
@@ -197,10 +197,8 @@ need_resched:
    jz restore_all
    testl $IF_MASK,EFLAGS(%esp) # interrupts off (exception
path) ?
    jz restore_all
- movl $PREEMPT_ACTIVE,TI_preempt_count(%ebp)
  sti
- call schedule
- movl $0,TI_preempt_count(%ebp)
+ call preempt_schedule
  cli
  jmp need_resched
#endif
```

—

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/>