

Crash-on-boot in init_i440gx SMP

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-10/6307.html>

From: Bruce Perens (bruce_at_perens.com)

Date: 10/30/03

To: linux-kernel@vger.kernel.org

Date: Thu, 30 Oct 2003 12:02:03 -0800 (PST)

Hi,

Using -test9, I'm hitting BUG_ON(phys_addr + size < phys_addr) in ioremap.c, called from init_i440gx with a dual Pentium 3 SMP motherboard. Boot log with oops, and ksymoops output are attached. This system is main server for perens.com, and thus hasn't exercised 2.6 much. Please email me if I can be of any further assistance.

Thanks

Bruce Perens

Linux version 2.6.0-test9 (root@server) (gcc version 3.3.2 (Debian)) #5 SMP Wed Oct 29 15:21:09 PST 2003

BIOS-provided physical RAM map:

BIOS-e820: 0000000000000000 - 000000000009fc00 (usable)
BIOS-e820: 000000000009fc00 - 00000000000a0000 (reserved)
BIOS-e820: 00000000000e0000 - 0000000000100000 (reserved)
BIOS-e820: 0000000000100000 - 0000000000c00000 (usable)
BIOS-e820: 00000000fec00000 - 00000000fec01000 (reserved)
BIOS-e820: 00000000fee00000 - 00000000fee01000 (reserved)
BIOS-e820: 00000000ffc00000 - 0000000100000000 (reserved)

192MB LOWMEM available.

found SMP MP-table at 000fb4f0

hm, page 000fb000 reserved twice.

hm, page 000fc000 reserved twice.

hm, page 000f2000 reserved twice.

hm, page 000f3000 reserved twice.

On node 0 totalpages: 49152

DMA zone: 4096 pages, LIFO batch:1

Normal zone: 45056 pages, LIFO batch:11

HighMem zone: 0 pages, LIFO batch:1

DMI 2.1 present.

ACPI disabled because your bios is from 99 and too old

You can enable it with acpi=force

ACPI: Unable to locate RSDP

Intel MultiProcessor Specification v1.4

Linux-Kernel: Crash-on-boot in init_i440gx SMP

Virtual Wire compatibility mode.
OEM ID: INTEL Product ID: 440BX APIC at: 0xFEE00000
Processor #0 6:7 APIC version 17
Processor #1 6:7 APIC version 17
I/O APIC #2 Version 17 at 0xFEC00000.
Enabling APIC mode: Flat. Using 1 I/O APICs
Processors: 2
Building zonelist for node : 0
Kernel command line: root=/dev/hda3 ro console=tty1 console=ttyS0,38400n8
Initializing CPU#0
PID hash table entries: 1024 (order 10: 8192 bytes)
Detected 451.139 MHz processor.
Console: colour VGA+ 80x25
Memory: 190528k/196608k available (2215k kernel code, 5420k reserved, 784k data, 164k init, 0k highmem)
Calibrating delay loop... 890.88 BogoMIPS
Dentry cache hash table entries: 32768 (order: 5, 131072 bytes)
Inode-cache hash table entries: 16384 (order: 4, 65536 bytes)
Mount-cache hash table entries: 512 (order: 0, 4096 bytes)
CPU: L1 I cache: 16K, L1 D cache: 16K
CPU: L2 cache: 512K
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#0.
Enabling fast FPU save and restore... done.
Enabling unmasked SIMD FPU exception support... done.
Checking 'hlt' instruction... OK.
POSIX conformance testing by UNIFIX
CPU0: Intel Pentium III (Katmai) stepping 03
per-CPU timeslice cutoff: 1461.42 usecs.
task migration cache decay timeout: 2 msecs.
enabled ExtINT on CPU#0
ESR value before enabling vector: 00000004
ESR value after enabling vector: 00000000
Booting processor 1/1 eip 3000
Initializing CPU#1
masked ExtINT on CPU#1
ESR value before enabling vector: 00000000
ESR value after enabling vector: 00000000
Calibrating delay loop... 901.12 BogoMIPS
CPU: L1 I cache: 16K, L1 D cache: 16K
CPU: L2 cache: 512K
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#1.
CPU1: Intel Pentium III (Katmai) stepping 03
Total of 2 processors activated (1792.00 BogoMIPS).
ENABLING IO-APIC IRQs
Setting 2 in the phys_id_present_map
...changing IO-APIC physical APIC ID to 2 ... ok.
..TIMER: vector=0x31 pin1=2 pin2=0
testing the IO APIC.....
..... done.
Using local APIC timer interrupts.

Linux–Kernel: Crash–on–boot in init_l440gx SMP

```
calibrating APIC timer ...
..... CPU clock speed is 450.0965 MHz.
..... host bus clock speed is 100.0214 MHz.
checking TSC synchronization across 2 CPUs: passed.
Starting migration thread for cpu 0
Bringing up 1
CPU 1 IS NOW UP!
Starting migration thread for cpu 1
CPUS done 2
NET: Registered protocol family 16
PCI: PCI BIOS revision 2.10 entry at 0xfdb81, last bus=1
PCI: Using configuration type 1
mtrr: v2.0 (20020519)
ACPI: Subsystem revision 20031002
ACPI: Interpreter disabled.
ACPI: ACPI tables contain no PCI IRQ routing entries
PCI: Invalid ACPI–PCI IRQ routing table
PCI: Probing PCI hardware
PCI: Probing PCI hardware (bus 00)
Machine check exception polling timer started.
Starting balanced_irq
Total HugeTLB memory allocated, 0
Initializing Cryptographic API
Limiting direct PCI/PCI transfers.
pty: 256 Unix98 ptys configured
Linux agpgart interface v0.100 (c) Dave Jones
agpgart: Detected an Intel 440BX Chipset.
agpgart: Maximum main memory to use for agp memory: 150M
agpgart: AGP aperture is 4M @ 0xfa400000
[drm] Initialized tdfx 1.0.0 20010216 on minor 0
Serial: 8250/16550 driver $Revision: 1.90 $ 8 ports, IRQ sharing enabled
ÿttyS0 at I/O 0x3f8 (irq = 4) is a 16550A
ttyS1 at I/O 0x2f8 (irq = 3) is a 16550A
RAMDISK driver initialized: 16 RAM disks of 4096K size 1024 blocksize
Uniform Multi–Platform E–IDE driver Revision: 7.00alpha2
ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx
PIIX4: IDE controller at PCI slot 0000:00:07.1
PIIX4: chipset revision 1
PIIX4: not 100% native mode: will probe irqs later
   ide0: BM–DMA at 0xffa0–0xffa7, BIOS settings: hda:DMA, hdb:DMA
   ide1: BM–DMA at 0xffa8–0xffaf, BIOS settings: hdc:DMA, hdd:pio
hda: Maxtor 96147H8, ATA DISK drive
hdb: SAF CD–RW6424A, ATAPI CD/DVD–ROM drive
Using anticipatory io scheduler
ide0 at 0x1f0–0x1f7,0x3f6 on irq 14
hdc: Maxtor 93652U8, ATA DISK drive
ide1 at 0x170–0x177,0x376 on irq 15
hda: max request size: 128KiB
hda: 120060864 sectors (61471 MB) w/2048KiB Cache, CHS=65535/16/63, UDMA(33)
   hda: hda1 hda2 hda3
hdc: max request size: 128KiB
```

Linux-Kernel: Crash-on-boot in init_l440gx SMP

```
hdc: 71346240 sectors (36529 MB) w/2048KiB Cache, CHS=65535/16/63, UDMA(33)
hdc: hdc1 hdc2 hdc3
ieee1394: Loaded CMP driver
-----[ cut here ]-----
kernel BUG at arch/i386/mm/ioremap.c:202!
invalid operand: 0000 [#1]
CPU: 0
EIP: 0060:[<c011e778>] Not tainted
EFLAGS: 00010207
EIP is at ioremap_nocache+0xb8/0xd0
eax: c37fd800 ebx: cc000000 ecx: 00000000 edx: 00000100
esi: cc814000 edi: 00100000 ebp: fff00000 esp: cbfc1f84
ds: 007b es: 007b ss: 0068
Process swapper (pid: 1, threadinfo=cbfc0000 task=c11ff900)
Stack: fff00000 00100000 00000010 c37fd800 00007113 cbf23000 cbf23c00 00000000
       c040b465 fff00000 00100000 00000000 00000000 c04145f0 c04145f4 00000001
       00000000 c03f29ab 00000000 00000000 c013739f 00000000 cbfc0000 c01050f6
Call Trace:
 [<c040b465>] init_l440gx+0x65/0x290
 [<c03f29ab>] do_initcalls+0x2b/0xa0
 [<c013739f>] init_workqueues+0xf/0x24
 [<c01050f6>] init+0x56/0x180
 [<c01050a0>] init+0x0/0x180
 [<c01092b9>] kernel_thread_helper+0x5/0xc

Code: 0f 0b ca 00 f8 db 33 c0 eb af 0f 0b c9 00 f8 db 33 c0 eb a1
<0>Kernel panic: Attempted to kill init!

-----
Ksymoops output:
ksymoops 2.4.9 on i686 2.4.22-1-686-smp. Options used
 -v /usr/src/linux/vmlinux (specified)
 -K (specified)
 -L (specified)
 -o /lib/modules/2.6.0-test9 (specified)
 -m /boot/System.map-2.6.0-test9 (specified)

No modules in ksyms, skipping objects
CPU 1 IS NOW UP!
Machine check exception polling timer started.
kernel BUG at arch/i386/mm/ioremap.c:202!
invalid operand: 0000 [#1]
CPU: 0
EIP: 0060:[<c011e778>] Not tainted
Using defaults from ksymoops -t elf32-i386 -a i386
EFLAGS: 00010207
eax: c37fd800 ebx: cc000000 ecx: 00000000 edx: 00000100
esi: cc814000 edi: 00100000 ebp: fff00000 esp: cbfc1f84
ds: 007b es: 007b ss: 0068
Stack: fff00000 00100000 00000010 c37fd800 00007113 cbf23000 cbf23c00 00000000
       c040b465 fff00000 00100000 00000000 00000000 c04145f0 c04145f4 00000001
```

Linux-Kernel: Crash-on-boot in init_l440gx SMP

00000000 c03f29ab 00000000 00000000 c013739f 00000000 cbfc0000 c01050f6

Call Trace:

```
[<c040b465>] init_l440gx+0x65/0x290
[<c03f29ab>] do_initcalls+0x2b/0xa0
[<c013739f>] init_workqueues+0xf/0x24
[<c01050f6>] init+0x56/0x180
[<c01050a0>] init+0x0/0x180
[<c01092b9>] kernel_thread_helper+0x5/0xc
Code: 0f 0b ca 00 f8 db 33 c0 eb af 0f 0b c9 00 f8 db 33 c0 eb a1
```

>>EIP; c011e778 <ioremap_nocache+b8/d0> <=====

```
>>eax; c37fd800 <__crc_nobh_prepare_write+24e27/745e2>
>>ebx; cc000000 <__crc__down_failed_trylock+6d475/2b619b>
>>esi; cc814000 <__crc_bio_unmap_user+14ce38/3f930c>
>>ebp; fff00000 <__crc_radix_tree_delete+31421e/396f3d>
>>esp; cbfc1f84 <__crc__down_failed_trylock+2f3f9/2b619b>
```

```
Trace; c040b465 <init_l440gx+65/290>
Trace; c03f29ab <do_initcalls+2b/a0>
Trace; c013739f <init_workqueues+f/24>
Trace; c01050f6 <init+56/180>
Trace; c01050a0 <init+0/180>
Trace; c01092b9 <kernel_thread_helper+5/c>
```

```
Code; c011e778 <ioremap_nocache+b8/d0>
00000000 <_EIP>;
Code; c011e778 <ioremap_nocache+b8/d0> <=====
0: 0f 0b ud2a <=====
Code; c011e77a <ioremap_nocache+ba/d0>
2: ca 00 f8 lret $0xf800
Code; c011e77d <ioremap_nocache+bd/d0>
5: db 33 (bad) (%ebx)
Code; c011e77f <ioremap_nocache+bf/d0>
7: c0 eb af shr $0xaf,%bl
Code; c011e782 <ioremap_nocache+c2/d0>
a: 0f 0b ud2a
Code; c011e784 <ioremap_nocache+c4/d0>
c: c9 leave
Code; c011e785 <ioremap_nocache+c5/d0>
d: 00 f8 add %bh,%al
Code; c011e787 <ioremap_nocache+c7/d0>
f: db 33 (bad) (%ebx)
Code; c011e789 <ioremap_nocache+c9/d0>
11: c0 eb a1 shr $0xa1,%bl
```

<0>Kernel panic: Attempted to kill init!

—

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

Linux-Kernel: Crash-on-boot in init_I440gx SMP

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

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