

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-03/msg06401.html>

- *From:* Leroy van Logchem <leroy.vanlogchem@xxxxxxxxxx>
 - *Date:* Thu, 15 Mar 2007 15:12:05 +0100
-

Andi Kleen wrote:

Where does it hang exactly? Do you have a boot log?

Linux version 2.6.20 (root@linus) (gcc version 3.4.6 20060404 (Red Hat 3.4.6-3)) #1 SMP Thu Mar 15 11:06:29 CET 2007

BIOS-provided physical RAM map:

sanitize start

sanitize end

copy_e820_map() start: 0000000000000000 size: 000000000009b800 end: 000000000009b800 type: 1

copy_e820_map() type is E820_RAM

copy_e820_map() start: 000000000009b800 size: 000000000004800 end: 00000000000a0000 type: 2

copy_e820_map() start: 00000000000ca000 size: 000000000002000 end: 00000000000cc000 type: 2

copy_e820_map() start: 00000000000e4000 size: 000000000001c000 end: 00000000000100000 type: 2

copy_e820_map() start: 00000000000100000 size: 00000000cfe70000 end: 00000000cff70000 type: 1

copy_e820_map() type is E820_RAM

copy_e820_map() start: 00000000cff70000 size: 000000000008000 end: 00000000cff78000 type: 3

copy_e820_map() start: 00000000cff78000 size: 000000000008000 end: 00000000cff80000 type: 4

copy_e820_map() start: 00000000cff80000 size: 000000000008000 end: 00000000d0000000 type: 2

copy_e820_map() start: 00000000e0000000 size: 0000000010000000 end: 00000000f0000000 type: 2

copy_e820_map() start: 00000000fec00000 size: 000000000010000 end: 00000000fec10000 type: 2

copy_e820_map() start: 00000000fee00000 size: 000000000001000 end: 00000000fee01000 type: 2

copy_e820_map() start: 00000000ff800000 size: 000000000400000 end: 00000000ffc00000 type: 2

copy_e820_map() start: 00000000ffffc00 size: 000000000000400 end: 0000000100000000 type: 2

copy_e820_map() start: 0000000100000000 size: 0000000030000000 end: 0000000130000000 type: 1

copy_e820_map() type is E820_RAM

BIOS-e820: 0000000000000000 - 000000000009b800 (usable)

BIOS-e820: 000000000009b800 - 00000000000a0000 (reserved)

BIOS-e820: 00000000000ca000 - 00000000000cc000 (reserved)

BIOS-e820: 00000000000e4000 - 00000000000100000 (reserved)

BIOS-e820: 00000000000100000 - 00000000cff70000 (usable)

BIOS-e820: 00000000cff70000 - 00000000cff78000 (ACPI data)

BIOS-e820: 00000000cff78000 - 00000000cff80000 (ACPI NVS)

BIOS-e820: 00000000cff80000 - 00000000d0000000 (reserved)

BIOS-e820: 00000000e0000000 - 00000000f0000000 (reserved)

BIOS-e820: 00000000fec00000 - 00000000fec10000 (reserved)

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO]1<- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

```
BIOS-e820: 00000000fee00000 – 00000000fee01000 (reserved)
BIOS-e820: 00000000ff800000 – 00000000ffc00000 (reserved)
BIOS-e820: 00000000ffffc000 – 0000000100000000 (reserved)
BIOS-e820: 0000000100000000 – 0000000130000000 (usable)
3968MB HIGHMEM available.
896MB LOWMEM available.
found SMP MP-table at 000f5d00
NX (Execute Disable) protection: active
Zone PFN ranges:
DMA 0 -> 4096
Normal 4096 -> 229376
HighMem 229376 -> 1245184
early_node_map[1] active PFN ranges
0: 0 -> 1245184
DMI present.
Using APIC driver default
ACPI: PM-Timer IO Port: 0x1008
ACPI: LAPIC (acpi_id[0x00] lapic_id[0x00] enabled)
Processor #0 15:4 APIC version 20
ACPI: LAPIC (acpi_id[0x01] lapic_id[0x06] enabled)
Processor #6 15:4 APIC version 20
ACPI: LAPIC (acpi_id[0x02] lapic_id[0x01] enabled)
Processor #1 15:4 APIC version 20
ACPI: LAPIC (acpi_id[0x03] lapic_id[0x07] enabled)
Processor #7 15:4 APIC version 20
ACPI: LAPIC_NMI (acpi_id[0x00] high edge lint[0x1])
ACPI: LAPIC_NMI (acpi_id[0x01] high edge lint[0x1])
ACPI: LAPIC_NMI (acpi_id[0x02] high edge lint[0x1])
ACPI: LAPIC_NMI (acpi_id[0x03] high edge lint[0x1])
ACPI: IOAPIC (id[0x02] address[0xfec00000] gsi_base[0])
IOAPIC[0]: apic_id 2, version 32, address 0xfec00000, GSI 0–23
ACPI: IOAPIC (id[0x03] address[0xfec10000] gsi_base[24])
IOAPIC[1]: apic_id 3, version 32, address 0xfec10000, GSI 24–47
ACPI: IOAPIC (id[0x04] address[0xfec80000] gsi_base[48])
IOAPIC[2]: apic_id 4, version 32, address 0xfec80000, GSI 48–71
ACPI: IOAPIC (id[0x05] address[0xfec80400] gsi_base[72])
IOAPIC[3]: apic_id 5, version 32, address 0xfec80400, GSI 72–95
ACPI: INT_SRC_OVR (bus 0 bus_irq 0 global_irq 2 high edge)
ACPI: INT_SRC_OVR (bus 0 bus_irq 9 global_irq 9 high level)
Enabling APIC mode: Flat. Using 4 I/O APICs
Using ACPI (MADT) for SMP configuration information
Allocating PCI resources starting at d1000000 (gap: d0000000:10000000)
Detected 3200.406 MHz processor.
Built 1 zonelists. Total pages: 1233024
Kernel command line: ro root=/dev/VolGroup00/LogVol100 rhgb console=tty0 console=ttyS0,9600n8
Enabling fast FPU save and restore... done.
Enabling unmasked SIMD FPU exception support... done.
Initializing CPU#0
CPU 0 irqstacks, hard=c0485000 soft=c0465000
PID hash table entries: 4096 (order: 12, 16384 bytes)
Console: colour VGA+ 80x25
```

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Dentry cache hash table entries: 131072 (order: 7, 524288 bytes)
Inode-cache hash table entries: 65536 (order: 6, 262144 bytes)
Memory: 4137600k/4980736k available (2344k kernel code, 54916k reserved, 855k data, 236k init, 3276224k highmem)
virtual kernel memory layout:
fixmap : 0xffc56000 – 0xfffff000 (3748 kB)
pkmap : 0xffa00000 – 0xffc00000 (2048 kB)
vmalloc : 0xf8800000 – 0xff9fe000 (113 MB)
lowmem : 0xc0000000 – 0xf8000000 (896 MB)
.init : 0xc0425000 – 0xc0460000 (236 kB)
.data : 0xc034a2ea – 0xc04201d4 (855 kB)
.text : 0xc0100000 – 0xc034a2ea (2344 kB)
Checking if this processor honours the WP bit even in supervisor mode... Ok.
Calibrating delay using timer specific routine.. 6405.55 BogoMIPS (lpj=12811105)
Security Framework v1.0.0 initialized
SELinux: Initializing.
SELinux: Starting in permissive mode
selinux_register_security: Registering secondary module capability
Capability LSM initialized as secondary
Mount-cache hash table entries: 512
monitor/mwait feature present.
using mwait in idle threads.
CPU: Trace cache: 12K uops, L1 D cache: 16K
CPU: L2 cache: 2048K
CPU: Physical Processor ID: 0
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#0.
CPU0: Intel P4/Xeon Extended MCE MSR (24) available
CPU0: Thermal monitoring enabled
Compat vDSO mapped to fffff000.
Checking 'hlt' instruction... OK.
Freeing SMP alternatives: 11k freed
ACPI: Core revision 20060707
CPU0: Intel(R) Xeon(TM) CPU 3.20GHz stepping 03
Booting processor 1/1 eip 2000
CPU 1 irqstacks, hard=c0486000 soft=c0466000
Initializing CPU#1
Calibrating delay using timer specific routine.. 6433.00 BogoMIPS (lpj=12866005)
monitor/mwait feature present.
CPU: Trace cache: 12K uops, L1 D cache: 16K
CPU: L2 cache: 2048K
CPU: Physical Processor ID: 0
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#1.
CPU1: Intel P4/Xeon Extended MCE MSR (24) available
CPU1: Thermal monitoring enabled
CPU1: Intel(R) Xeon(TM) CPU 3.20GHz stepping 03
Booting processor 2/6 eip 2000
CPU 2 irqstacks, hard=c0487000 soft=c0467000
Initializing CPU#2
Calibrating delay using timer specific routine.. 6400.78 BogoMIPS (lpj=12801564)

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

monitor/mwait feature present.

CPU: Trace cache: 12K uops, L1 D cache: 16K

CPU: L2 cache: 2048K

CPU: Physical Processor ID: 3

Intel machine check architecture supported.

Intel machine check reporting enabled on CPU#2.

CPU2: Intel P4/Xeon Extended MCE MSR (24) available

CPU2: Thermal monitoring enabled

CPU2: Intel(R) Xeon(TM) CPU 3.20GHz stepping 03

Booting processor 3/7 eip 2000

CPU 3 irqstacks, hard=c0488000 soft=c0468000

Initializing CPU#3

Calibrating delay using timer specific routine.. 6400.79 BogoMIPS (lpj=12801592)

monitor/mwait feature present.

CPU: Trace cache: 12K uops, L1 D cache: 16K

CPU: L2 cache: 2048K

CPU: Physical Processor ID: 3

Intel machine check architecture supported.

Intel machine check reporting enabled on CPU#3.

CPU3: Intel P4/Xeon Extended MCE MSR (24) available

CPU3: Thermal monitoring enabled

CPU3: Intel(R) Xeon(TM) CPU 3.20GHz stepping 03

Total of 4 processors activated (25640.13 BogoMIPS).

ENABLING IO-APIC IRQs

..TIMER: vector=0x31 apic1=0 pin1=2 apic2=-1 pin2=-1

checking TSC synchronization across 4 CPUs: passed.

Brought up 4 CPUs

migration_cost=14,1477

NET: Registered protocol family 16

ACPI: bus type pci registered

PCI: Using MMCONFIG

PCI: No mmconfig possible on 9:1

Setting up standard PCI resources

ACPI: Interpreter enabled

ACPI: Using IOAPIC for interrupt routing

ACPI: PCI Root Bridge [PCI0] (0000:00)

PCI quirk: region 1000-107f claimed by ICH4 ACPI/GPIO/TCO

PCI quirk: region 1180-11bf claimed by ICH4 GPIO

PCI: PXH quirk detected, disabling MSI for SHPC device

PCI: PXH quirk detected, disabling MSI for SHPC device

PCI: Transparent bridge - 0000:00:1e.0

ACPI: PCI Interrupt Link [LNKA] (IRQs 3 4 *5 6 7 10 11 14 15)

ACPI: PCI Interrupt Link [LNKB] (IRQs 3 4 5 6 7 *10 11 14 15)

ACPI: PCI Interrupt Link [LNKC] (IRQs 3 4 5 6 7 *10 11 14 15)

ACPI: PCI Interrupt Link [LNKD] (IRQs 3 4 5 6 7 10 *11 14 15)

ACPI: PCI Interrupt Link [LNKE] (IRQs 3 4 5 6 7 10 11 14 15) *0, disabled.

ACPI: PCI Interrupt Link [LNKF] (IRQs 4 5 6 7 10 11 14 15) *0, disabled.

ACPI: PCI Interrupt Link [LNKG] (IRQs 3 4 5 6 7 10 *11 14 15)

ACPI: PCI Interrupt Link [LNKH] (IRQs 4 5 6 7 *10 11 14 15)

Linux Plug and Play Support v0.97 (c) Adam Belay

PCI: Using ACPI for IRQ routing

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

PCI: If a device doesn't work, try "pci=routeirq". If it helps, post a report
PCI: Bridge: 0000:01:00.0
IO window: disabled.
MEM window: disabled.
PREFETCH window: disabled.
PCI: Bridge: 0000:01:00.2
IO window: disabled.
MEM window: disabled.
PREFETCH window: disabled.
PCI: Bridge: 0000:00:02.0
IO window: disabled.
MEM window: dd100000-dd1fffff
PREFETCH window: disabled.
PCI: Bridge: 0000:04:00.0
IO window: disabled.
MEM window: dd200000-dd2fffff
PREFETCH window: df400000-df7fffff
PCI: Bridge: 0000:04:00.2
IO window: disabled.
MEM window: disabled.
PREFETCH window: disabled.
PCI: Bridge: 0000:00:04.0
IO window: disabled.
MEM window: dd200000-dd2fffff
PREFETCH window: df400000-df7fffff
PCI: Bridge: 0000:00:06.0
IO window: disabled.
MEM window: disabled.
PREFETCH window: disabled.
PCI: Bridge: 0000:00:1c.0
IO window: 2000-2fff
MEM window: dd300000-dd3fffff
PREFETCH window: disabled.
PCI: Bridge: 0000:00:1e.0
IO window: 3000-3fff
MEM window: dd400000-deffffff
PREFETCH window: d1000000-d10fffff
ACPI: PCI Interrupt 0000:00:02.0[A] -> GSI 16 (level, low) -> IRQ 16
ACPI: PCI Interrupt 0000:00:04.0[A] -> GSI 16 (level, low) -> IRQ 16
ACPI: PCI Interrupt 0000:00:06.0[A] -> GSI 16 (level, low) -> IRQ 16
NET: Registered protocol family 2
IP route cache hash table entries: 32768 (order: 5, 131072 bytes)
TCP established hash table entries: 131072 (order: 9, 2621440 bytes)
TCP bind hash table entries: 65536 (order: 8, 1310720 bytes)
TCP: Hash tables configured (established 131072 bind 65536)
TCP reno registered
checking if image is initramfs... it is
Freeing initrd memory: 1087k freed
Simple Boot Flag at 0x39 set to 0x1
apm: BIOS not found.
audit: initializing netlink socket (disabled)

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

```
audit(1173966820.156:1): initialized
highmem bounce pool size: 64 pages
Total HugeTLB memory allocated, 0
VFS: Disk quotas dquot_6.5.1
Dquot-cache hash table entries: 1024 (order 0, 4096 bytes)
SELinux: Registering netfilter hooks
io scheduler noop registered
io scheduler anticipatory registered
io scheduler deadline registered
io scheduler cfq registered (default)
Real Time Clock Driver v1.12ac
Linux agpgart interface v0.101 (c) Dave Jones
[drm] Initialized drm 1.1.0 20060810
Serial: 8250/16550 driver $Revision: 1.90 $ 4 ports, IRQ sharing enabled
ÿserial8250: ttyS0 at I/O 0x3f8 (irq = 4) is a 16550A
serial8250: ttyS1 at I/O 0x2f8 (irq = 3) is a 16550A
RAMDISK driver initialized: 16 RAM disks of 16384K size 1024 blocksize
Intel(R) PRO/1000 Network Driver - version 7.3.15-k2-NAPI
Copyright (c) 1999-2006 Intel Corporation.
ACPI: PCI Interrupt 0000:08:01.0[A] -> GSI 24 (level, low) -> IRQ 17
e1000: 0000:08:01.0: e1000_probe: (PCI:66MHz:32-bit) 00:30:48:76:d7:70
e1000: eth0: e1000_probe: Intel(R) PRO/1000 Network Connection
ACPI: PCI Interrupt 0000:08:02.0[A] -> GSI 25 (level, low) -> IRQ 18
e1000: 0000:08:02.0: e1000_probe: (PCI:66MHz:32-bit) 00:30:48:76:d7:71
e1000: eth1: e1000_probe: Intel(R) PRO/1000 Network Connection
Uniform Multi-Platform E-IDE driver Revision: 7.00alpha2
ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx
ICH5: IDE controller at PCI slot 0000:00:1f.1
ACPI: PCI Interrupt 0000:00:1f.1[A] -> GSI 18 (level, low) -> IRQ 19
ICH5: chipset revision 2
ICH5: not 100% native mode: will probe irqs later
ide0: BM-DMA at 0x1460-0x1467, BIOS settings: hda:pio, hdb:pio
ide1: BM-DMA at 0x1468-0x146f, BIOS settings: hdc:DMA, hdd:pio
hdc: _NEC CD-ROM CD-3002C, ATAPI CD/DVD-ROM drive
ide1 at 0x170-0x177,0x376 on irq 15
hdc: ATAPI 52X CD-ROM drive, 96kB Cache, UDMA(33)
Uniform CD-ROM driver Revision: 3.20
ide-floppy driver 0.99.newide
PNP: No PS/2 controller found. Probing ports directly.
serio: i8042 KBD port at 0x60,0x64 irq 1
serio: i8042 AUX port at 0x60,0x64 irq 12
mice: PS/2 mouse device common for all mice
input: AT Translated Set 2 keyboard as /class/input/input0
TCP cubic registered
Initializing XFRM netlink socket
NET: Registered protocol family 1
NET: Registered protocol family 17
Using IPI Shortcut mode
Freeing unused kernel memory: 236k freed
Time: tsc clocksource has been installed.
```

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO] <- Bad?

At this point SysRq and ctrl+alt+del still work but it doesnt continue.

2.6.20.2 and 3 would continue at this point with:

SCSI subsystem initialized

ACPI: PCI Interrupt 0000:05:0e.0[A] -> GSI 18 (level, low) -> IRQ 17

etcetera etcetera

Can you please double check this by trying with/without again -- sometimes bisepts go bad.

As requested I started to redo the test but now without git using kernel.org tars.

The results now are, still using the same .config:

linux-2.6.20.tar.gz : bad

linux-2.6.20.1.tar.gz: bad (boot log equal)

linux-2.6.20.2.tar.gz: good

linux-2.6.20.3.tar.gz: good

(triple checked)

Really strange. Nilshar, please try these kernels too with:

COMPAT_VDSO=y

CONFIG_HIGHMEM64G=y

-

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

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

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

Re: [Bug 8040] Hang before INIT when CONFIG_HIGHMEM4G=y [Fix CONFIG_COMPAT_VDSO]7<- Bad?